RUNNING HEAD: Understanding person acquisition

Understanding Person Acquisition using an Interactive Activation and Competition Network.

Sarah V Stevenage and Hugh G Lewis

University of Southampton

AFFILIATIONS: Dr Sarah Stevenage is a member of the Department of Psychology, University of Southampton. Dr Hugh Lewis is a member of the Department of Aeronautics and Astronautics, School of Engineering, University of Southampton.

ACKNOWLEDGEMENTS: The authors would like to thank Professor Mike Burton, and Professor Vicki Bruce for their help in the preparation of this manuscript.

CORRESPONDENCE may be sent to: Dr Sarah Stevenage, Department of Psychology, University of Southampton, Highfield, Southampton, SO17 1BJ, UK. Electronic mail may be sent to svs1@soton.ac.uk

Understanding Person Acquisition using an Interactive Activation and Competition Network.

Abstract

Face perception is one of the most developed visual skills that humans display, and recent work has attempted to examine the mechanisms involved in face perception through noting how neural networks achieve the same performance. The purpose of the present paper is to extend this approach to look not just at human face recognition, but also at human face acquisition. Experiment 1 presents empirical data to describe the acquisition over time of appropriate representations for newly encountered faces. These results are compared with those of Simulation 1, in which a modified IAC network capable of modelling the acquisition process is generated. Experiment 2 and Simulation 2 explore the mechanisms of learning further, and it is demonstrated that the acquisition of a set of associated new facts is easier than the acquisition of individual facts in isolation of one another. This is explained in terms of the advantage gained from additional inputs and mutual reinforcement of developing links within an interactive neural network system.

Understanding Person Acquisition using an Interactive Activation and Competition Network.

Introduction

Imagine a situation in which you see a person approaching you. From a distance, you think you recognise their jacket and the manner of their walk. You think it may be a friend of yours. As they get closer, you are able to gather more information: they do not seem to be avoiding you, in fact, they are smiling at you. Closer still and you can see that their face is familiar. You start to remember things about them - where you know them from, and their name - and you can see that they are pleased to see you. You have progressed from a situation of not knowing who was approaching you, to a situation of confident identification. The face, while not the only source of information, is an important part of that recognition process.

Research over the past decade has taken great strides forward in terms of our understanding of the processes underlying face recognition. Importantly, recent developments in connectionist modelling have led to the development of neural network models capable of simulating a host of phenomena associated with face recognition (see Brédart, Valentine, Calder & Gassi, 1995; Burton & Bruce, 1992; 1993; Burton, Bruce & Johnston, 1990; Burton, Young, Bruce, Johnston & Ellis, 1991) and the possibility has been explored that the mechanisms involved in such computer simulations might be instructive in terms of our understanding of the processes involved in human face recognition. However, recognition of familiar faces is far from being the limit of our human face processing capacity. Humans also have the ability to <u>learn</u> the faces of new people. This acquisition of new faces presents a challenge to any connectionist model of face processing, and the development of a network capable of modelling human face learning is the focus of the present paper.

The Interactive Activation and Competition Model (IAC) of Face Processing

Our ability to learn new faces allows us to hold faces in our short term memory for as long as a casual social situation requires. Furthermore, we can learn these faces in the sense that they gradually become part of the corpus of faces that we call familiar. A new connectionist architecture developed from the well established IAC model of familiar face recognition is presented here with the aim of modelling these learning processes. Before this can be described, a brief overview of IAC and its major mechanisms is provided (see Burton, Bruce & Johnston, 1990 for a fuller discussion).

IAC is based on the 'Jets and Sharks' framework presented by McClelland and Rumelhart (1988). It consists of pools of Face Recognition Units (FRUs), Person Identity Nodes (PINs) and Name Output Units (NOUs), and a pattern of interconnectivity exists in which units associated with the same person are connected by excitatory links across pools, and units associated with different people are connected by inhibitory links within pools. All links are bi-directional and in this sense, the IAC system is truly interactive (see Figure 1).

(Please insert Figure 1 about here)

IAC is successful in modelling a range of empirical and neuropsychological phenomena associated with face processing. For instance, the model can successfully activate the correct set of units on presentation of a known face, and so simulate human face recognition. In addition, it can account for the differential timescale in making familiarity, occupation and name decisions (Young, McWeeny, Ellis & Hay, 1986), the difficulty associated with naming and tip-of-the-tongue states (Burton & Bruce, 1992; Burton & Bruce, 1993 but see Brédart, Valentine, Calder & Gassi, 1995), and the curious phenomenon of covert recognition (Burton, Young, Bruce, Johnston & Ellis, 1991). Perhaps most importantly for the present paper, IAC provides a parsimonious account of the phenomena of associative (semantic) and repetition (identity) priming, and the mechanisms underlying these phenomena are outlined below.

Associative and Repetition Priming within IAC

Associative priming occurs when seeing one face (for example, Prince Charles) primes, or speeds up, the recognition of a second associated face (for example, Princess Diana). IAC can explain this with reference to the influence that the shared semantic units have. From Figure 1, seeing the face of Prince Charles causes activation of the FRU, and subsequently, the PIN for Prince Charles. Activation of Charles' PIN spreads to the semantic units, and so allows retrieval of information specific to Prince Charles. This includes some facts that are shared between Charles and Diana (i.e., 'royals'), and because all units are connected by bi-directional links, activation is passed back to Diana's PIN, and subsequently to her FRU. The units associated with Diana have thus been raised from their resting level. On subsequent presentation of Diana's face the relevant units will require less activation to reach a critical threshold level. The result is that Diana is recognised faster than she would have been had Charles not just been seen.

Repetition priming effects can also be explained by the IAC architecture, though a slightly different mechanism is implicated. Repetition priming occurs when the recognition of a face is primed, or speeded up, by prior exposure to that face. These effects are long lasting and domain specific (a name will prime later recognition of that name, and a face will prime later recognition of that face, but a name will not prime later recognition of the face or vice versa (Bruce & Valentine, 1985). Burton et al. (1990) suggest that repetition priming arises as a consequence of the way that links in the network are created and updated. Each link has a strength associated with it that dictates how readily activation can be passed along it. This strength is increased when the units at either end of the link are simultaneously active. In the case of repetition priming, the initial sight of a face (Prince Charles) will cause activation of both the FRU and the PIN for Prince Charles. With both active at the same time, the FRU-PIN link will become strengthened. On a second presentation, activation will be able to flow down this link faster, and thus recognition will be achieved more quickly.

Recent evidence has demonstrated that repetition priming may actually be located at two levels of an IAC architecture: between the feature input units and the FRUs, and between the FRUs and the associated PINs (Ellis, Burton, Young & Flude, 1997). Strengthening of links at both levels is responsible for entire face repetition priming (as above) and for the fact that presentation of part of a face (i.e., internal features only) can successfully prime subsequent recognition of the whole face. However, presentation of one part of a recognisable face (i.e., internal features only) can also prime the subsequent recognition of another non-overlapping part of the face (i.e., external features). IAC is able to explain this by implicating the only links common to the two situations - those between the FRU and the PIN.

The Importance of Hebbian Update as a Learning Mechanism

The notion of link update, using a traditional Hebbian Learning Mechanism (Hebb, 1949) can be extended to account not only for link strengthening, but also for link creation. In this sense, it is possible to develop an IAC system capable of learning by utilising mechanisms already implicit and accepted within IAC. This was the basis of a simple two-layer network (IAC-L) presented by Burton (1994), which modelled the recruitment of a new FRU for a new face input. The work here represents an extension of Burton's work inasmuch as learning of the features, face and name are involved rather than just the learning of the features themselves. As with Burton's work, the use of a Hebbian learning mechanism here has considerable advantages over alternative learning algorithms within connectionist architectures. The most common alternative for pattern recognition systems is the back propagation of error (Rumelhart, Hinton & Williams, 1986), which involves the global modification of all weights in the network so as to minimise recognition error. Back propagation presents a very popular and well-understood mechanism for supervised learning, but the notion of change to all weights throughout a network limits its appropriateness when what is called for here is the adjustment of weights in a more selective and unsupervised fashion.

Hebbian learning presents an alternative learning mechanism. It operates by strengthening links connecting two simultaneously active units, whilst weakening links connecting an active and an inactive unit (Hebb, 1949). Consequently, it provides an elegant, parsimonious and biologically plausible way of modelling both priming and learning. Given that the purpose of this paper is to explore these learning mechanisms more fully, the aim of Experiment 1 is to determine the pattern of human learning of new faces while the aim of Simulation 1 is to develop an IAC-like system capable of modelling this pattern.

Experiment 1

Design

The purpose of the present study was to examine the speed and manner of new face acquisition by human subjects. In order to slow the learning process down, and so examine its form, the visually difficult task of learning twins' faces was used.

<u>Participants</u>

Fifteen undergraduate participants (8 male, 7 female, age range 19-48 years (mean age = 23 years) acted as volunteers in the present study. All had normal or corrected-to-normal vision, and were unfamiliar with the stimulus faces.

<u>Materials</u>

The stimuli consisted of 80 head and shoulders photographs of a pair of identical twin girls (Lizzie and Rosie). Care was taken to ensure consistency in clothing and grooming, and all photographs were manipulated within Corel PhotoPaint to be matched on inter-ocular distance and to remove all background details. Photographs that varied in pose and expression were used to ensure that any emergent effects were in response to differences in the faces rather than differences in particular viewpoints. All stimuli were scanned into a Power Macintosh 7200/90 computer and were presented via a 17” colour monitor, at a viewing distance of 60-70 cm.

<u>Procedure</u>

Participants completed a training phase in which their task was to learn to identify the two twin girls by name. This training phase consisted of consecutive blocks of trials, with 10 trials (5 of each twin) in each block. Participants viewed one photograph at a time, for a fixed period of 2 seconds, and their task was to say whether it was Lizzie or Rosie. They indicated their response by pressing one of two keys labelled with the twins’ names. Following each response, the computer provided appropriate verbal corrective feedback (i.e., ‘Yes, Lizzie’ or ‘No, Rosie’). After a 500ms inter-stimulus interval, the next face was presented, and training continued until the participant reached a pre-set criterion of 100% accuracy over two successive blocks of 10 trials. Identification accuracy was recorded on each trial, as was the number of trials taken to reach the preset learning criterion.

Results and Discussion

The results were examined in terms of the number of trials taken to reach the preset learning criterion, and the shape of the learning curve, or the pattern of acquisition shown. Examination of the data confirmed that all fifteen participants were successful in learning the identities of the two twins. This was achieved in an average of 21.1 presentations (sd = 9.45). Consequently, although faced with a difficult perceptual task, human perceivers were able to learn the features that reliably distinguished one face from another.

Examination of the learning curve across participants revealed that this learning progressed rather slowly and erratically at first, but was followed by a burst of acquisition and then a levelling off of performance before the steady state of perfect performance was achieved (see Figure 2). This pattern of acquisition is characteristic of complex human learning. Consequently, any simulation of human learning should reflect both acquisition of the task, and a pattern of learning which shows the characteristic sigmoidal, or ‘s’-shaped, learning curve.

(Please insert Figure 2 about here)

Simulation 1: Learning a new person.

Structure of the IACAPA Architecture

IACAPA (Interactive Activation and Competition Account of Person Acquisition) is shown in Figure 3 (below). As with previous IAC models, IACAPA consisted of a large number of units, arranged into pools. Forty feature units (FTUs) provided the facial input and these were further arranged in pools according to the features that they described (i.e., hair, eyes, nose, mouth, etc.). IACAPA thus had 4 types of each of 10 features to describe the face. Photofit-type terms are used here for the sake of convenience and consistency with the IAC-L model (Burton, 1994) but, as

Burton noted, the face may be represented by more wholistic parameters such as eigenfaces. These represent a statistical way of describing a set of faces, with each eigenface representing a different principle component across the face set. Being based on the variation in pixel intensities across the entire face, eigenfaces capture global properties of the face such as hairstyle, face shape, and gender, although more individuating qualities can be revealed through the higher order eigenfaces (see Bruce & Humphreys, 1994 for a discussion). The nature of the representation at this feature unit level does not, however, affect the functioning of the network as a whole.

The FTUs fed a pool of 51 face recognition units (FRUs) which, in turn, fed a pool of 51 person identity nodes (PINs). Following from Brédart et al.'s (1995) reformulation of IAC, the present architecture echoed the notion that names and semantic information are likely to be stored in separate pools. The PINs thus fed two pools of units corresponding to the semantic information that we have about known people (SIUs - not included here) and their name output units (NOUs).

(Please insert Figure 3 about here)

IACAPA was initialised to have one known face. Units corresponding to this known face were connected via maximal excitatory links, while units that did not correspond to this face were connected to the 'known unit' via maximal inhibitory links. In addition, within-pool inhibition operated on all units within a pool so that the activation of one unit resulted in the inhibition of all other units within that pool. The architecture thus possessed strong links for known people, and pilot work established that the manual activation of the FTUs corresponding to the known face resulted in the successful activation of the FRU, PIN, and finally the NOU for the known person. In this sense, IACAPA remained able to model familiar face recognition.

The remaining 50 units in each of the FRU, PIN and NOU pools represented ‘free units’ which had not previously been recruited to represent an identity and were thus available to a learning situation. As with all units, within-pool inhibition was maximal between these free units. However, initially, all free units were connected to units from previous and subsequent network layers by links carrying small and random weights varying between –0.1 and +0.1. Activation of a novel set of feature inputs would be propagated forward most readily to free FRUs possessing a relatively large net random weight. Within-pool inhibition would then ensure that only one free FRU remained active. This free FRU could then be said to represent the new face. Similarly, recruitment of free units at the PIN and NOU layers of the network would model person and name acquisition respectively.

Mechanisms within the IACAPA Architecture

1. Activation. To preserve biological plausibility, the activation of all units was regulated by two parameters. First, activation was propagated forwards as a proportion (*alpha*) of the net input from the previous layer rather than by means of a perfect feed-forward mechanism. The activation received by units in the subsequent layer was thus subject to propagation loss. Second, all units were subject to a ceiling level of activation rather than responding in a boundless fashion. Unit activation was kept within its maximum and minimum levels by the influence of a *decay* function. This operated to reduce activity by an amount proportional to the current activation state: a unit with low activation received a minimal modification by the decay function while a highly active unit required a larger modification to prevent exceeding its maximum activation level. Activation of units was thus a balance between a proportion of the net input, and a decay term to prevent continuous escalation of activation.

2. Learning. The process of recruiting a set of free units will not, on its own, lead to lasting learning. For lasting learning to occur, the links between the novel input units and the newly recruited units must be strengthened. As utilised by Burton (1994) and endorsed by Medler (1998), the application of a Hebbian update mechanism will modify the network links to simulate learning (Hebb, 1949). Hebbian update serves to increase the weight between two units when, and only when, the units at either end of the link are simultaneously active. This reinforces the development of an association between units, and the link is strengthened by an amount proportional to the activity of the units themselves. Thus, while activity of the units at either end of the link is low, the link between the units is strengthened by only a small amount, but as the activation of the units increases, so the strengthening of the link accelerates until its weight reaches a maximum excitatory level (+1). The result is the modelling of gradual learning.

The Hebbian update function also serves to create inhibitory links between two non-associated units, when, and only when, the unit at one end of the link is active. This serves to reinforce the development of a negative association so that one unit will not mistakenly lead to activation of a unit with which it is not associated. Again, the degree to which the inhibitory link is made more inhibitory is proportional to the activation level of the units at either end of the link. When activation is low, the inhibitory link is modified only a little, while when activation is higher, the inhibitory link is modified to a greater extent, until its weight reaches a maximum inhibitory level (-1).

The Simulation Environment

All simulations were implemented in C using the Rochester Connectionist Simulator (Goddard, Lynne, Mintz & Bukys, 1989) and were run in a Unix environment on a Sun Workstation. The network parameters for all simulations were based on those used in previous research (see Burton et al., 1990). Unit activation varied between +1 and -0.2, with a slightly negative resting level of -0.1. In addition, established link strengths varied between a maximum excitatory weight of +1 and a maximum inhibitory weight of -1. The constant parameters governing the propagation of activation (*alpha*) and *decay* were selected on the basis of pilot work to ensure that activation was propagated successfully throughout the three layers of the IACAPA network. Details of all parameters and functions adopted in these simulations are provided in the Appendix.

<u>Procedure</u>

The network was presented with a novel face described by a set of FTUs which did not correspond to the known face. An alternate excitation and link-update process was used to achieve learning as follows. First, the FTUs for the novel face were maximally activated and the network was allowed to step through a series of cycles. The activation of all FRUs was monitored until they had reached a stable state. For the purposes of this and all subsequent simulations, stability was defined as having a constant activation level across fifteen consecutive cycles. The number of cycles taken to reach stability varied across the learning sequence (25 - 110 cycles), with longer learning periods being required earlier in the learning sequence. Once stability had been reached, the Hebbian update function was applied. The network was then stepped through a further set of cycles until stability was reached, and the Hebbian update function was again applied. This process was repeated until a free FRU, PIN and NOU had been recruited and their activation showed no further increase, at which time the network could be said to have created a representation for the face, the person, and the name of the newly acquired pattern.

Results of Simulation One

Figure 4 presents the results from this simulation, with activation levels being allowed to rise continuously across the learning sequence rather than being reset to resting level after each update phase. The figure thus represents the cumulative activation of the winning FRU, PIN and NOU across the learning phase.

(Please insert Figure 4 about here)

From these results, the most important factor was the affirmation that IACAPA was able to recruit a free FRU, PIN and NOU to represent the novel face. This can be interpreted as the successful acquisition of the representations necessary for complete person acquisition and represents a step on from Burton's (1994) IAC-L where names were not involved. As such, IACAPA represents the first neural network capable of learning the features, face and name associated with an unfamiliar person.

To be sure that these results were not limited in some way to the particular novel face presented, these results were replicated across a further four novel face presentations. These four new faces were created by the unique combinations of ten feature units such that the new faces bore limited resemblance to one another and to the 'known' face. Using an identical method of alternate excitation and link-update as above, the results presented a consistent pattern. For all new faces, a different set of free FRUs, PINs and NOUs were successfully recruited (see Table 1), and for all new faces, the rate of increase of activation associated with recruited units conformed to Figure 4 presented above. This confirmed that learning within IACAPA represented a generalisable learning mechanism rather than being the result of a particular starting set of links and a particular input.

(Please insert Table 1 about here)

Taking the results of all learning occasions together, several factors are worthy of note. First, the pace of learning increased as the learning sequence proceeded. This was shown by the decreasing number of cycles taken to reach stability after each Hebbian update. The result is a gradual learning curve which concords with the typical pattern of human learning from Experiment 1. Gradual learning was achieved in IACAPA through the fact that links were updated by an amount proportional to their level of activation. Low activation resulted in small increases in the strength of the associated links, while greater activation resulted in an accelerated increase in the strength of the links.

Second, it was notable that the final activation levels of the recruited units in Figure 4 were very similar to those represented by the known FRU, PIN and NOU during pilot work. Interrogation of the network also revealed that the strength of the links connecting these newly recruited units approached the maximum level manually coded for the known face. The implication is that, given sufficient time and/or exposure to a face, the network, like the human, can acquire a new set of representations and can come to recognise a once-unfamiliar person with the ease associated with the recognition of a very well known person.

Finally, some comment must be given to the performance of the architecture at the NOU level. If the face were familiar to us, activation of a NOU would indicate that the familiar person’s name had successfully been retrieved. In a learning situation, however, it is quite possible that one does not actually know the name of the person whose face has been learned. In this case, how should successful activation of a single NOU be interpreted? It is suggested here that activation of a single NOU corresponds to the potential to retrieve a piece of information unique to a person - a unique identifier. In many cases this will be a proper name but in cases where the proper name is not known we can think of the NOU as representing a label or nickname which may be assigned as a unique identifier in the absence of a proper name. Activation of the NOU would thus correspond to the retrieval of this unique identifier. In time, it may be that the proper name for this newly learned face is provided and acquired. At that point, the information represented by the NOU could be updated to represent the proper name provided.

Interim Discussion of Mechanisms of learning

In evaluating IACAPA, it is first important to note that the network has acquired units to represent a novel face and can thus be considered to have learned the new face. Excitatory links connecting the newly recruited units for a given face were created through Hebbian weight update. Similarly, inhibitory links connecting a recruited unit with other non-recruited units in previous and subsequent pools, were created. Consequently, the pattern of connectivity ensured an economic learning strategy such that only one unit at each layer of the network would come to represent the new person, and the pattern of connectivity came to resemble that for already familiar people. Perhaps more importantly, this learning was achieved using a Hebbian learning strategy which has rigour and biological plausibility, both of which were considered important in the attempt to model human capabilities.

The development of IACAPA has, however, revealed three important learning factors that make IACAPA different from Burton's (1994) IAC-L. First, while the use of a Hebbian update rule was effective in modelling gradual learning, it was necessary to apply it in a local fashion rather than in a global fashion. In this sense, a stepped Hebbian update function was implemented here in which link strengths were only updated when a change had occurred at that layer of the network. In other words, on recruitment of a free FRU to represent a novel face, the links between

FTUs and the recruited FRU were updated, but links between the FRU and all possible PINs, and indeed between all possible PINs and NOUs, were left alone until such a time as a unit became recruited at these levels. The application of this learning mechanism provided an *economical learning strategy*. However, such a mechanism was also essential to prevent the inadvertent creation of inhibitory links at deeper layers of the network early on in the learning cycle when, for example, the propagation of activation had created a winning FRU but had not as yet reached the level of the PINs or NOUs. Burton's earlier work did not need to address this issue because IAC-L possessed only two network layers.

Second, the learning rule was applied throughout IACAPA in an exclusive fashion such that link weights between units were only updated if the units at one or other end of the link were involved. In other words, links between FRUs and PINs were only updated when there was a winner at the PIN level (stepped update rule above) and on links (a) from the recruited FRU to the recruited PIN (excitatory link strengthened) (b) from the recruited FRU to all other PINs (inhibitory link strengthened) and (c) from the recruited PIN to all other FRUs (inhibitory link strengthened). This procedure prevented the inadvertent reinforcement of links between non-recruited units. Instead, these non-recruited links were maintained at a random level and were thus available for future learning. The application of this learning mechanism provided a *parsimonious learning strategy*. Again, Burton's IAC-L did not need to apply its learning rule in such an exclusive fashion because IAC-L was developed for the purpose of demonstrating learning per se, rather than for the demonstration of repeated learning.

Finally, IACAPA modelled learning through the use of a flexible Hebbian learning rule in which the weight updates from an input layer to the next layer (i.e., from FTUs to FRUs) were updated in a selective fashion. In other words, when a free FRU became recruited to represent a new face, the links between the FTUs describing this face, and the recruited FRU were strengthened to create excitatory links. In addition, the links between the recruited FRU and the FTUs describing features not held by the novel face were strengthened to create inhibitory links. However, the links between active FTUs and non-active FRUs which, at any other layer would be strengthened to create inhibitory links, remained untouched. This modification allowed IACAPA the flexibility to learn future similar inputs. As such, the application of this learning mechanism provided a *realistic learning strategy*. Again, IAC-L did not need such a modification to be able to learn a single instance of a new face.

In combination, these learning rules allowed IACAPA to model learning in an economical, parsimonious and realistic fashion and in a manner in sympathy with the workings of neurones within the human brain. These results consequently add to the developments outlined by Burton (1994) who demonstrated the recruitment of a new FRU upon presentation of a new combination of features. However, the use here of a network with more than two layers, and the aim here of modelling the learning of several new instances, has highlighted the need to modify the application of the Hebbian learning rule in the pursuit of a neural network capable of mimicking human learning (see Appendix for summary of the three learning rules).

Nevertheless, humans do not just learn isolated facts; they learn complex associations of new information. Within the present sphere, humans have the capacity to learn not just the face but also the name of someone new, and then gradually acquire person-specific information about their new acquaintance. In some sense, we haven't modelled the entirety of human learning unless we model its complexity. One question that arises is whether the learning of several pieces of information concurrently makes the overall learning task harder or easier. In other words, does the possibility of associations between a set of new facts offset the escalated number of facts and ultimately facilitate learning? Experiment 2 in conjunction with Simulation 2 will examine this issue by considering learning of faces and names, as compared to faces alone.

Experiment 2

Design

A between-subjects design was used to assess the relative difficulty of a learning task when either faces alone, or faces and names were provided. Three different experimental conditions existed consisting of (i) a control condition (identical to Experiment 1) in which participants viewed faces alone; (ii) an intermediate learning condition in which participants learned faces but had access to standard photographs which provided a visual cue, and (iii) a condition in which participants learned faces but had access to named standard photographs which thus provided a visual and a naming cue. All participants completed a sequence of training trials, with corrective feedback, and the dependent variable was the number of trials taken to reach a preset criterion of identification accuracy.

Participants

Forty-five undergraduate participants (22 male, 23 female; age range 18-49 years, mean age = 23 years) acted as volunteers in the present experiment. All participants had normal, or corrected-to-normal vision and were unfamiliar with the stimuli. Participants were randomly assigned to one of three conditions, with age and gender matched as far as possible across all conditions.

Materials

The 80 photographs of twin girls Lizzie and Rosie from Experiment 1 were used again in Experiment 2. In addition, two further photographs (one of each twin) were prepared and served the function of visual cues to assist learning.

All 80 training photographs were viewed via a 17” colour monitor at a viewing distance of 60-70 cm. In contrast, the 2 ‘help’ photographs were reproduced as black and white hard copy photographs and were available for visual inspection throughout the experiment.

Procedure

Participants in all three experimental conditions completed a learning phase, identical to that in Experiment 1. The faces of a pair of identical twin girls were presented one at a time for 2 seconds, and the participants was required to respond by means of a button press to indicate which girl (Lizzie or Rosie) they were looking at. Corrective feedback was provided and was followed by an inter-stimulus interval of 500msecs and then the next presentation. The learning phase continued until the participant reached 100% accuracy over two successive blocks of 10 trials each.

The difference between the three groups was in the nature of assistance they received prior to the learning phase. Group One represented a control condition and received no prior information except that they were about to view the faces of identical twin girls Lizzie and Rosie. Group Two received minimal assistance prior to the learning phase. They received a single pre-exposure to each twin by being given an unlabelled photograph that they could refer to throughout the learning phase. This allowed them the possibility of making a visual comparison in order to assist them in their visual differentiation of the two girls. This is analogous to the participants receiving a set of feature inputs for each of the girls. Group Three received maximal assistance prior to the learning phase. They received the same photographs of each girl, but each was labelled with the appropriate name. Thus, these participants had reference to a visual means of discriminating between the girls and the opportunity to learn which name went with which arrangement of features. This is analogous to being given a name and associated feature set for each individual.

Across all groups, the computer recorded the number of trials taken to reach the criterion level of learning, as well as the identification accuracy on each of these trials.

Results and Discussion

Examination of the data confirmed that all participants in all conditions were able to learn the identities of each twin. Consequently, even given a more complex and difficult learning task involving the acquisition of two highly similar faces and their associated names, participants were able to succeed within a relatively short period of training and feedback.

Formal examination of the task difficulty, as expressed by the number of trials taken to complete this learning, suggested that performance was substantially affected by the experimental condition. Table 2 summarises the number of trials required to reach the criterion level of learning, from which it is evident that participants given no help learned slowest and participants given face and name assistance learned quickest. Interestingly, participants given visual help only, in the form of photographs to permit visual discrimination of the two faces, did not show an intermediate rate of learning but rather, showed a level of performance markedly similar to the control group who had no help at all.

(Please insert Table 2 about here)

Analysis by means of a one way Analysis of Variance (ANOVA) across experimental groups, using the number of trials to criterion as the dependent variable, confirmed a main effect of experimental condition ($F(2, 42) = 3.004$, $p_{(1\text{-tailed})} < .05$). Further analysis of the a-priori pairwise contrasts confirmed that this was due to the groups given no help (group 1) and visual information (group 2) showing an equivalent rate of learning ($t(28) < 1$, $p > .05$), whilst being significantly slower than those given visual and naming information (group 3 vs no help: $t(28) = 2.54$, $p < .025$; group 3 vs visual information: $t(28) = 2.07$, $p < .05$).

Another way of expressing the facilitation effect that face and name information had on person acquisition was by examination of the averaged learning curves across each of the three conditions (see Figure 5). From this it was evident that the learning curve displayed by those participants given face and name information to assist them was substantially steeper, and reached its asymptote after fewer trials, than the learning curve associated with the other conditions. Consequently, for IACAPA to be able to simulate human learning of complex and realistic associations, learning must be achieved, and must be achieved in fewer network cycles when multiple inputs are provided compared to when unitary inputs are provided. Simulation 2 addresses this issue.

(Please insert Figure 5 about here)

Simulation 2: The influence of a name when learning new faces

In the above learning situation, participants were presented with a face and a name input belonging to the same person. The learning task therefore involved the development of a set of associations so that the face and name excited the same mental representations. In terms of the performance of a network, such a learning task requires that the novel pattern of features recruits a set of units at the FRU, PIN and NOU layers, and moreover, that the name excites the same units. In this sense, the architecture must learn that the face and the name belong to the same person.

To investigate this issue, an IACAPA architecture similar to that used in Simulation 1 was created. The only difference was that the present architecture could provide a name input, via word recognition units (WRUs) and then name recognition units (NRUs) (see Burton & Bruce, 1993), as well as a face input via the feature units (FTUs). As before, the network was initialised to have one known face and 50 free units available for learning at each layer of the network. Units representing the face and name of the known person were connected by links of maximal excitatory strength. Consequently, activation of the first name (i.e., Stan) or the family name (i.e., Laurel) of a familiar person would be propagated from the relevant WRUs down maximal excitatory links to activate a NRU associated with that name ('Stan Laurel'). Activation at this NRU level is thus associated with the <u>presentation</u> of a name rather than with its retrieval. Successful activation at the NRU level is then propagated forward to the PIN to allow a sense of familiarity with the person and subsequent retrieval of semantic information and name production.

As in Simulation 1, however, free units which had not previously been recruited to represent a known individual were connected to units in previous and subsequent network layers by links that carried small random weights. Consequently, application of a Hebbian Weight Update rule was required to simulate the learning of an association of a word with a name, and then of that name with the features and face of that new person. All parameters and initial network settings were identical to those used in Simulation 1, and the architecture is illustrated in Figure 6 below.

(Please insert Figure 6 about here)

<u>Procedure</u>

Simulation 2 consisted of the presentation of (i) a new face and (ii) a new face and new name combination. This was achieved by manual excitation of (i) the feature units, and (ii) the feature units and word recognition units associated with the new person. As in Simulation 1, alternate activation and weight update functions were applied such that weights within the network were updated only when a stable state of activation had been reached. Again, stability was defined as a constant level of activation across fifteen consecutive network cycles. This process of activation and weight update was repeated until a free FRU, (NRU), PIN and NOU had been recruited, at which time the network could be said to have created a representation for the face, person, and name of the newly acquired pattern. The number of network cycles required to reach this point, and the final activation levels of the winning units at each layer were recorded as indicators of the relative difficulty of the learning tasks.

Results of Simulation 2

Figure 7 (below) summarises the resultant activation at the FRU, PIN and NOU (and where relevant, NRU) levels when this extended architecture was presented with (i) a face only and (ii) a face and name combination. As can be seen, IACAPA has successfully recruited a free FRU, PIN, NOU (and NRU) to represent the new face, person, and name.

(Please insert Figure 7 about here)

Again, for each learning situation, this pattern of results was replicated across four further new individuals to demonstrate that the learning was generalisable and was not the product of the development of a certain set of connections to a certain pattern of excitation (see Table 3). The results confirmed that, in each case, a new FRU, PIN NOU (and NRU) was recruited and thus the complete cycle of learning was achieved.

(Please insert Table 3 about here)

Examination of these results also allows determination of the influence of presentation of the name alongside the new face compared to the situation where the face alone was presented. In this respect, it was clear that both the speed and final activation level of the recruited units were influenced by the concurrent presentation of the name. The network reached a stable state of activation of a new FRU, (NRU), PIN and NOU within fewer cycles, and the final activation levels of the PIN and NOU were higher, when face and name information was provided than when the face alone was provided. The activation of the newly recruited FRU remained unaffected and at ceiling level, although this was achieved after slightly fewer network cycles when both the face and name information were provided (281 cycles) than when the face alone was provided (298 cycles). In addition, it was notable that, as a consequence of the higher activation levels of the recruited units, links between units were acquired faster, and required fewer (larger) weight updates to reach a maximum strength when the face and name were presented compared to when the face was presented alone. This confirms the picture provided by the activation data.

In summary, the results of this simulation mirrored those from the human learning situation. In both Experiment 2 and Simulation 2, provision of the name as well as the face resulted in greater ease of acquisition of a new person as demonstrated by PIN and NOU level activation. Consequently, it was clear that the network was able to simulate the fact that a concurrent learning task in which there was potential for associations between related facts made the overall learning task easier rather than more difficult.

Interim Discussion

The findings of both Experiment 2 and Simulation 2 are somewhat counterintuitive and require explanation. Both sets of results seem to show that when required to learn two things rather than just one, the learning task is completed in less time. The simulation results provide some way of explaining these findings through appealing to the notion of back activation or bi-directionality within a highly interconnected system.

Essentially, in the situation where a face and name input are provided, the PIN receives activation from two sources – visual and linguistic. The net input to the system as a whole is thus increased compared to the situation where the face alone is presented. As a consequence, a newly recruited PIN has a greater direct input and can attain a higher level of activation (or can reach an arbitrary threshold in fewer network cycles) when face and name are presented than when the face alone is presented.

This facilitation at the PIN level is propagated forward to the NOU with the result that the NOU also attains a higher resultant level of activation (or reaches a threshold in fewer network cycles). This corresponds in human performance terms to a faster successful name generation. In addition to this forward propagation of activation, the facilitation at the PIN level is also propagated backwards, via the bi-directional links, to the FRU level. Activation is maintained at a high (or ceiling) level, but in addition, because the net activation at both ends of the PIN-FRU link is now greater, the link has the capacity to become stronger. The result is that not only is the level of activation facilitated throughout the network, but also the strength of the newly created links is boosted when face and name are presented compared to when the face alone is presented. Performance advantages result and are expressed through the fact that conscious attainment of a sense of familiarity or the recall of a name is achieved more quickly. What initially appeared as a counterintuitive result can therefore be explained with reference to the *association* of the to-be-learned information and the interconnectivity that results.

Conclusions

The research presented here has explored the pattern of human learning of the faces and names of new people, and has presented simulation data to model these learning patterns. Through development of an IAC system tailored for familiar face recognition, the current IACAPA network is capable of modelling both recognition of familiar faces (and names) and the learning of new faces (and names). This represents a novel contribution in itself. More than this, however, the development of the present neural network has provided insight into the process that humans may use to achieve complex learning tasks. In particular, the development of IACAPA has highlighted the importance of an economic, parsimonious and flexible learning system, and the use of a biologically plausible bi-directional network has allowed the positive effects of interconnectivity during learning to become apparent.

The strength of using IAC as a basis for the present development is that the resultant model retains the abilities and capabilities of IAC to simulate a range of face-related phenomena. In addition, the biological plausibility and cognitive economy of the IAC system has been preserved. IACAPA represents a system of excitation between, and inhibition within, pools of related units connected by bi-directional links and in this way it represents a truly interactive system that is sympathetic to the neural connectivity and topographical organisation of the human brain. Indeed, Medler (1998) has hailed the success of such a network because of its ability to display human-like qualities in its performance. It can, for example, retrieve information via several routes. It can indicate identity from a facial input, a semantic input (i.e., through the semantic category label 'teacher') and from a name input. In this sense, IAC and IACAPA, like humans, have what is termed a 'content-addressable memory'. The network also operates in a way in which the activation state of the system is readily interpretable at any stage. We can stop the flow of activation and examine the activation states of the units in the network to determine the system's best guess, just as with humans. Third, the network operates by using a process of iterative retrieval, which again is characteristic of human information processors. It retrieves information by trying each unit in a pool until the appropriate one is found (via its strong link). Finally, the network displays a gradual activation decline in which the activation of the units within the network gently decreases when the stimulus is removed. This, again, mirrors the performance of a human information processor.

Incorporation of Semantic Information

The present implementation of IACAPA is, however, limited to the case in which faces and names are represented in the absence of semantic information. Clearly, to be able to capture the fuller picture of person representation, IACAPA requires development to incorporate a semantic system in the form of semantic information units (SIUs). Brédart, Valentine, Calder and Gassi (1995) suggest that semantic information should be represented in the form of a series of SIU pools to describe facts such as occupation, hobbies, nationality, etc. Inhibition within each of these pools is appropriate to prevent the retrieval (or acquisition) of the wrong occupation or the wrong hobby, but inhibition is not required between the separate SIU pools. In this sense, retrieval (or acquisition) of the label 'teacher' does not interfere with the retrieval (or acquisition) of the hobby 'likes golf'. This modification to Burton et al.'s (1990) IAC architecture was driven by demonstration of the fact that retrieval of semantic information was assisted when many facts were known about a person as opposed to when few facts were known.

Incorporation of a system of SIUs would also allow IACAPA to address the robust evidence regarding the relative ease of retrieving (and learning) names and semantic information. A great deal of empirical work supports the intuition that names are more difficult to retrieve than information such as occupations (see Valentine , Brennen & Brédart, 1996, for a review), and several explanations have been put forward to account for this. The traditional IAC explanation rests on the assumption that full names (first name-surname combinations) are very unlikely to be shared across two individuals and so are connected to the PIN layer by one-to-one links. In contrast, occupation labels such as ‘teacher’ are likely to be associated with several people (we are likely to know several teachers) and are thus connected to the PIN layer by many-to-one links. This relative uniqueness of names compared to occupations was considered to render the name-face link more fragile. Damage to this link would lead to the irrevocable failure to retrieve the name through any other route with the consequence that naming failures were more likely than occupation retrieval failures. Evidence in support of this uniqueness hypothesis is, however, mixed (see Stanhope & Cohen, 1993). In particular, the fact that some anomic patients can retrieve unique semantic information about an individual whilst remaining unable to name them is problematic for the uniqueness explanation (see Semenza & Zettin, 1988; 1989; Hanley, 1995).

One[1] alternative explanation for the difficulty when retrieving (or learning) names is that names are meaningless (Cohen, 1990). Two consequences arise. First, no synonyms exist for the name. Instead, the name has to be recalled precisely for it to ‘work’ (Brédart, 1993; Cohen & Faulkner, 1986). Second, the proper noun representing a name is not embedded within a semantic context in the same way that a common noun representing an occupation is. Interestingly, designation of a word as a name seems to strip the word of its potential to be embedded into a semantic context with the result that it is more difficult to remember the name ‘Mr Baker’ than to remember the occupation ‘baker’ (McWeeny, Young, Hay & Ellis, 1987). This explanation of naming difficulties can be incorporated into a network account of person information by allowing SIUs associated with common nouns, such as ‘teacher’ to be associated with a host of propositions such as ‘works with children’, ‘works at a school’, ‘teaches mathematics’, etc., The activation level of the SIU representing the semantic concept ‘teacher’ would thus rise, through these associations, to a level above that of a name unit which lacks such associations. Development of the network in this way would allow the modelling of naming difficulties when both retrieving and learning names without appealing to the notion of information uniqueness.

Information Salience

One situation that would, at present, provide a problem for IACAPA involves the presentation of ambiguous inputs to the system. For example, consider the case of learning two very similar faces where the proportion of shared feature units is great. At present, once the first face has been learned, the successful acquisition of the second face depends critically on the selection and recruitment of a second FRU. If, however, the second face shared more than half its feature units with the first face, IACAPA would always respond by activating the FRU for that first face to a greater degree than any remaining free FRUs. As a consequence, the FRU for that first face would successfully inhibit the selection and recruitment of any new FRUs. The second face would, in network terms, be indistinguishable from the first face and the network would have failed to learn. While visual similarity effects of this kind are evident in the literature, and can be simulated using neural networks (see Valentine &

Ferrara, 1991; Humphreys, Lamote & Lloyd-Jones, 1995) it is notable here that humans are capable of learning even the faces of identical twins when given enough training (see also Stevenage, 1998). The introduction of the notion of <u>information salience</u> as a way of modelling top-down processing may help the network to perform in a more human fashion in such ambiguous situations.

Information salience means that some inputs carry greater importance than others. In the extreme example provided by twin faces one might consider that the two faces are described by a number of feature units - let us say 10 feature units - and that the twins, while being genetically identical, and physically very similar, are in fact distinguishable by a very small number of features. In network terms, the twins might share 9 out of the 10 feature units, but might differ with respect to one descriptor. One of the challenges when learning twin faces is to identify this descriptor - the thing that reliably tells one twin apart from the other - the salient information. This descriptor is unlikely to be one of the most common features, and the difference in this descriptor across the twins is likely to be small. Hence, twin learning takes time. In essence, the task of twin identification then becomes a task of salient descriptor identification. With feedback, the importance attached to the nine common features might reasonably be expected to reduce, while the importance of the one discriminating or salient feature might reasonably increase.

We can model information salience by means of the strength of the links connecting the input unit to the corresponding unit in the next layer (here, between FTU and FRU). It may well be that while the strength of all FTU-FRU weights might start off equivalent, hence the twin confusion, increasing exposure and feedback may mean that the strength of the weights between salient features and the FRUs for each twin increases to a greater degree while the strength of the weights between nondiscriminatory features and the FRUs for each twin increases to a lesser degree, or even gradually declines. This calls for the introduction of an ‘unlearning’ mechanism along the lines of a ‘use it or lose it’ rule (see Nakisa & Plunkett, 1998). This could be achieved by the introduction of a *weight decay* mechanism.

A weight decay mechanism would improve the capacity of IACAPA to account for several aspects of person perception such as the modelling of information salience and the forgetting of faces, whilst retaining a biologically plausible explanation of the well-known effects of recognition and priming. Future research is aimed at exploring this issue and thus improving the accountability of IACAPA even further.

References

Brédart, S. (1993). Retrieval failures in face naming. Memory, 1, 351-366.

Brédart, S., Valentine, T.R., Calder, A., & Gassi, L. (1995). An interactive activation model of face naming. Quarterly Journal of Experimental Psychology - A, 48(2), 466-486.

Brennen, T. (1993). The difficulty with recalling people's names The plausible phonology hypothesis. Memory, 1, 409-431.

Bruce, V., & Humphreys, G.W. (1994). Recognising objects and faces. Visual Cognition, 1(2/3),

Bruce, V., & Valentine, T.R. (1985). Identity priming in the recognition of familiar faces. British Journal of Psychology, 76, 373-383.

Burke, D.M., MacKay, D.G., Worthley, J.S., & Wade, E. (1991). On the top of the tongue: What causes word finding failures in young and older adults? Journal of Memory and Language, 30, 542-579.

Burton, A.M. (1994). Learning new faces in an interactive activation and competition model. Visual Cognition, 1(2/3), 313-348.

Burton, A.M., & Bruce, V. (1992). I recognise your face but I can't remember your name: A simple explanation? British Journal of Psychology, 83, 45-60.

Burton, A.M., & Bruce, V. (1993). Naming faces and naming names: Exploring an interactive activation model of person recognition. Memory, 1(4), 457-480.

Burton, A.M., Bruce, V., & Johnston, R.A. (1990). Understanding face recognition with an interactive activation model. British Journal of Psychology, 81, 361-380.

Burton, A.M., Young, A.W., Bruce, V., Johnston, R.A., & Ellis, A.W. (1991). Understanding covert recognition. Cognition, 39, 129-166.

Cohen, G. (1990). Why is it difficult to put names to faces? <u>British Journal of Psychology</u>, <u>81</u>, 287-298.

Cohen, G., & Faulkner, D. (1986). Memory for proper names: Age differences in retrieval. <u>British Journal of Developmental Psychology</u>, <u>4</u>, 187-197.

Ellis, A.W., Burton, A.M., Young, A.W., Flude, B.M. (1997). Repetition priming between parts and wholes: Tests of a computational model of familiar face processing. <u>British Journal of Psychology</u>, <u>88</u>, 579-608.

Goddard, N.H., Lynne, K.J., Mintz, T., & Bukys, L. (1989). Rochester connectionist simulator. Technical report 233 (revised). Department of Computer Sciences, University of Rochester, New York.

Hanley, J.R. (1995). Are names difficult to recall because they are unique? A case study of a patient with anomia. <u>Quarterly Journal of Experimental Psychology</u>, <u>48A</u>, 487-506.

Hebb, D.O. (1949). <u>The Organisation of Behaviour</u>. New York: John Wiley & Sons.

Humphreys, G.W., Lamote, C., & Lloyd-Jones, T.J. (1995). An interactive activation approach to object processing: Effects of structural similarity, name frequency and task in normality and pathology. <u>Memory</u>, <u>3-4</u>, 535-586.

McClelland, J.L., & Rumelhart, D.L. (Eds.), (1988). <u>Explorations in Parallel Distributed Processing</u>. Cambridge, MA: MIT Press.

McWeeny, K.H., Young, A.W., Hay, D.C., & Ellis, A. (1987). Putting names to faces. <u>British Journal of Psychology</u>, <u>78</u>, 143-149.

Medler, D.A. (1998). A brief history of connectionism. <u>Neural computing surveys</u>, <u>1,</u> 61-606.

Nakisa, R.C., & Plunkett, K. (1998). Evolution of a rapidly learned representation for speech. <u>Language and Cognitive Processes</u>, <u>13(2/3)</u>, 105-127.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, J.L. McClelland & the PDP Research Group (Eds.), Parallel Distributed Processing, Vol 1: Foundations, Cambridge, MA: MIT Press.

Semenza, C., & Zettin, M. (1988). Generating proper names: A case of selective inability. Cognitive Neuropsychology, 5, 711-721.

Semenza, C., & Zettin, M. (1989). Evidence from aphasia for the role of proper names as pure referring expressions. Nature, 342, 678-679.

Stanhope, N., & Cohen, G. (1993). Retrieval of proper names: Testing the models. British Journal of Psychology, 84, 51-65.

Stevenage, S.V. (1998). Which twin are you? A demonstration of induced categorical perception of identical twin faces. British Journal of Psychology, 89, 39-57.

Valentine, T., Brennen, T., & Bredart, S. (1996). The Cognitive Psychology of Perper Names: On the Importance of Being Ernest. Routledge: London, New York.

Valentine, T., & Ferrara, A. (1991). Typicality in categorization, recognition and identification: Evidence from face recognition. British Journal of Psychology, 82, 87-102.

Young, A.W., McWeeny, K.H., Ellis, A.W., & Hay, D.C. (1986). Naming and categorising faces and written names. Quarterly Journal of Experimental Psychology - A, 38, 297-318.

Appendix

Unit Functions:

The net input to a unit i at time t is calculated as a sum of the weighted output of the units connected to unit i:

$$net_i = a \sum_j w_{ij} a_j(t), \quad \forall a_j(t) > 0$$

where a is a global strength parameter, w_{ij} is the weight on the connection between unit i and unit j, and $a_j(t)$ is the output activation of unit j at time t. Only units with a positive output $a_j(t)$ are entered into this sum.

The change in the activation of unit i in response to the input is given by the equation:

$$\Delta a_i(t) = \begin{cases} (max - a_i(t))net_i - decay(a_i(t) - rest), & net_i > 0 \\ (a_i(t) - min)net_i - decay(a_i(t) - rest), & net_i \leq 0 \end{cases}$$

where max and min are the maximum and minimum activation values for all units, $rest$ is the resting activation for all units, $decay$ is the global decay on all unit activations, and $a_i(t)$ is the current activation level of unit i. The new output activation of unit i, at time $t+1$, is then calculated as the sum of the current activation, $a_i(t)$, and the change in activation, $\Delta a_i(t)$:

$$a_i(t+1) = a_i(t) + \Delta a_i(t)$$

Link Functions:

The Hebb-like weight update function adopted by Burton (1994) was also used here:

$$\Delta w_{ij} = \begin{cases} l\, a_i a_j (1 - w_{ij}), & a_i a_j > 0 \\ l\, a_i a_j (1 + w_{ij}), & a_i a_j < 0 \end{cases}$$

where l is the global learning rate, but this update function was applied according to the following rules:

1. Stepped Update Rule:

Only weights on connections from a pool having a winning unit ($a_i(t) > 0$) and in all pools above this one are modified.

2. Exclusive Update Rule:

Only weights on connections to, or from, a winning unit are modified.

3. Flexible Update Rule:

Only weights on connections to a winning FRU from FTU pools that have an active unit are modified.

Global Parameters:

The global parameters were set as follows:

$max = 1.0$,

$min = -0.2$,

$rest = -0.1$,

$a = 0.2$,

$decay = 0.09$, and

$l = 0.5$.

In all simulations, between-pool excitatory connections for known people had weight 1.0, between-pool inhibitory connections for known people had weight –1.0, and all within-pool inhibitory connections had weight –1.0. Initial between-pool connections for free units had random weights in the range [-0.1, 0.1].

Table 1: Asymptotic activation levels achieved by winning FRU, PIN and NOU for each of five novel face presentations, together with the number of cycles taken to reach complete learning. (Mean refers to the mean level of activation across the limits between which the unit finally oscillated.)

	FRU Activation	PIN Activation	NOU Activation	
Activation	950.5	746	581	
	FRU Unit No.	PIN Unit No.	NOU Unit No.	No. cycles
Face 1	42	107	149	1207
Face 2	62	130	147	1329
Face 3	57	116	173	1199
Face 4	76	134	178	1229
Face 5	88	100	175	1253

Table 2: Mean number of trials taken to reach the criterion level of learning (and standard deviations) across different levels of input from Experiment Two.

	No Help	Face Input only	Face and Name Input
Mean no. trials to criterion	21.07	20.73	13.53
sd	9.45	11.79	6.57

Table 3:

	FRU Activation	PIN Activation	NOU Activation	NRU Activation	
Activation	950.5	747	582		
	FRU Unit No.	PIN Unit No.	NOU Unit No.	NRU Unit No.	No. cycles
Face 1	75	123	188		1218
Face 2	82	92	153		1234
Face 3	87	94	164		1208
Face 4	80	123	153		1241
Face 5	46	99	165		1231
	FRUActivation	PIN Activation	NOU Activation	NRU Activation	
Activation	950.5	820	606	780	
	FRU Unit No.	PIN Unit No.	NOU Unit No.	NRU Unit No.	No. cycles
Face/Name 1	68	124	166	195	1028
Face/Name 2	89	95	143	225	1019
Face/Name 3	60	106	163	200	1025
Face/Name 4	68	136	150	216	996
Face/Name 5	41	133	167	213	1035

Figure Legends

Figure 1: Interactive Activation and Competition model presented by Burton, Bruce and Johnston (1990). Reproduced with kind permission of the authors and the British Psychological Society.

Figure 2: Pattern of learning of the faces of identical twin girls by human observers.

Figure 3: Diagrammatic representation of the structure and connectivity within IACAPA. All links are bi-directional. For simplicity, only excitatory links to the known face, and only a sample of the random links between free units at each layer are shown.

Figure 4: Activation of the newly recruited FRU, PIN and NOU in response to the presentation of a novel face.

Figure 5: Averaged learning curves shown by human observers when learning the faces of identical twin girls under conditions of (i) no help, (ii) face help and (iii) face and name help.

Figure 6: Diagrammatic representation of the structure and connectivity within IACAPA when provided with a face and name input. All links are bi-directional. For simplicity, only excitatory links to the known face, and only a sample of the random links between free units at each layer are shown.

Figure 7: Activation of the newly recruited FRU, NRU, PIN and NOU in response to the presentation of (i) a novel face and (ii) a novel face and name.

Figure 1:

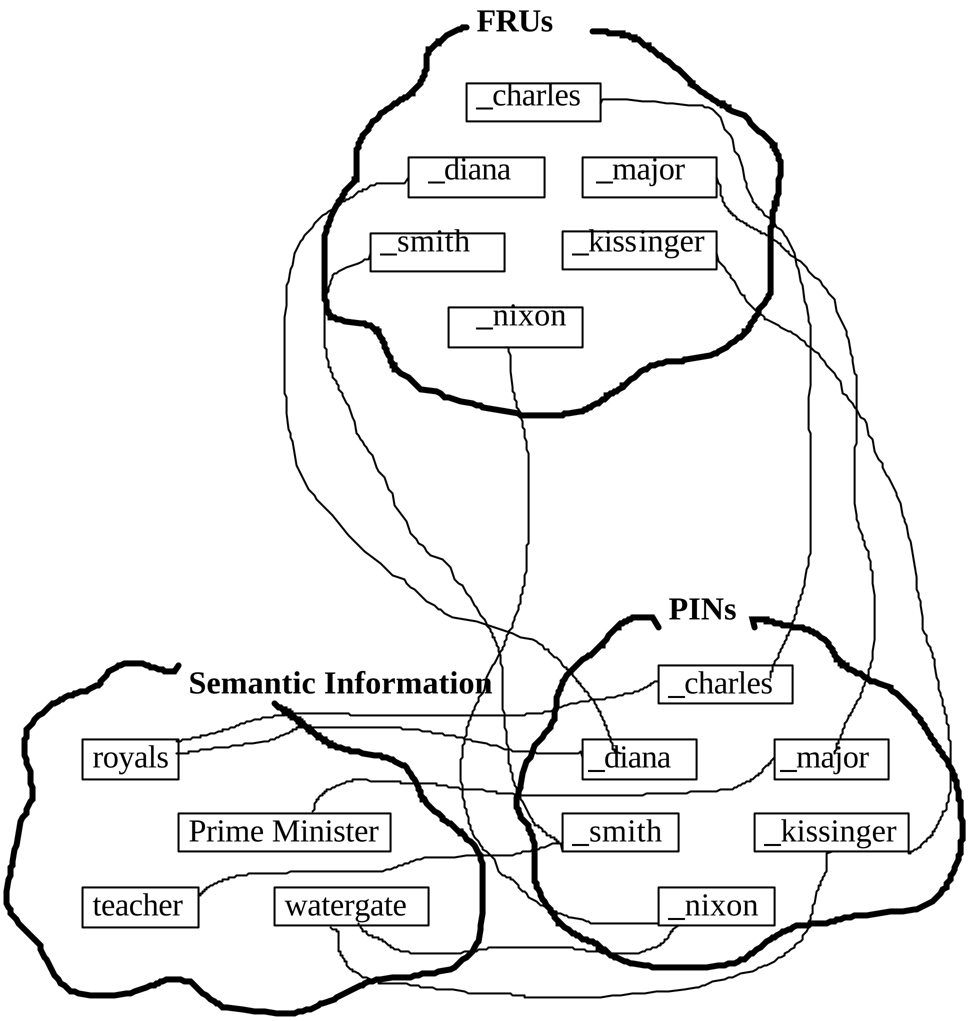
FRUs
_charles
_diana
_major
_smith
_kissinger
_nixon
PINs
Semantic Information
_charles
royals
_diana
_major
Prime Minister
_smith
_kissinger
teacher
watergate
_nixon

Figure 2:

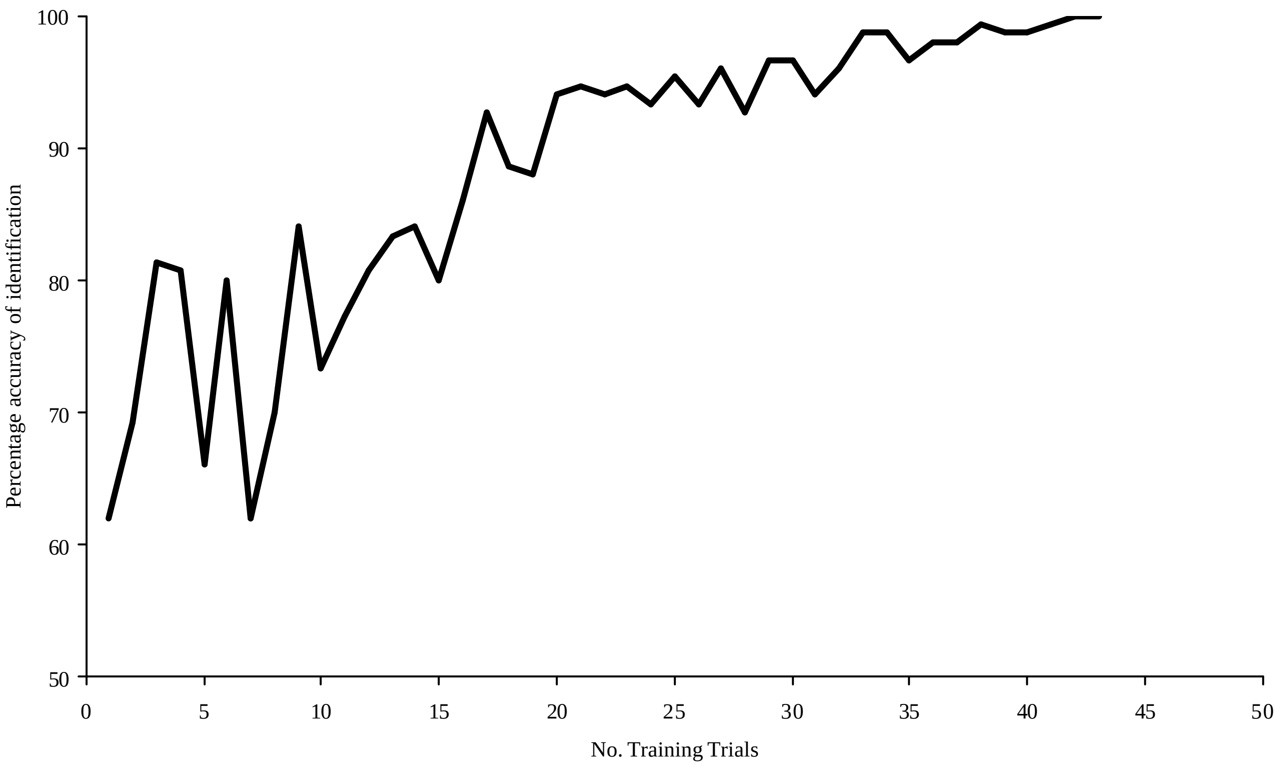
100
90
80
70
60
50
0
5
10
15
20
25
30
35
40
45
50
Percentage accuracy of identification
No. Training Trials

Figure 3:

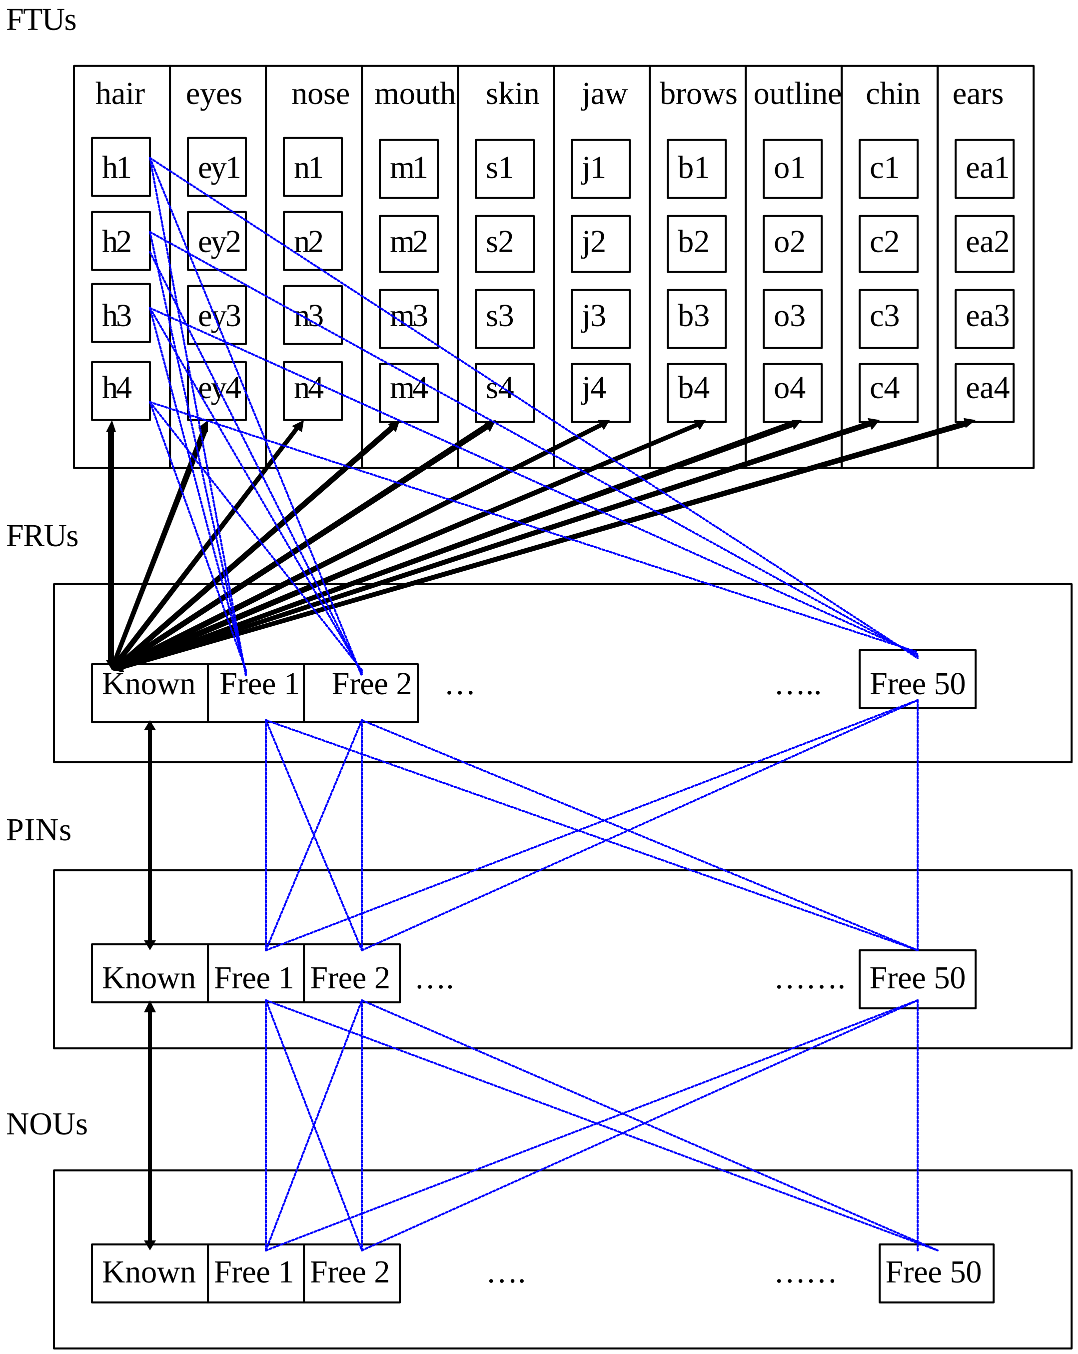

Key: = strong excitatory link, = small random link

Figure 4:

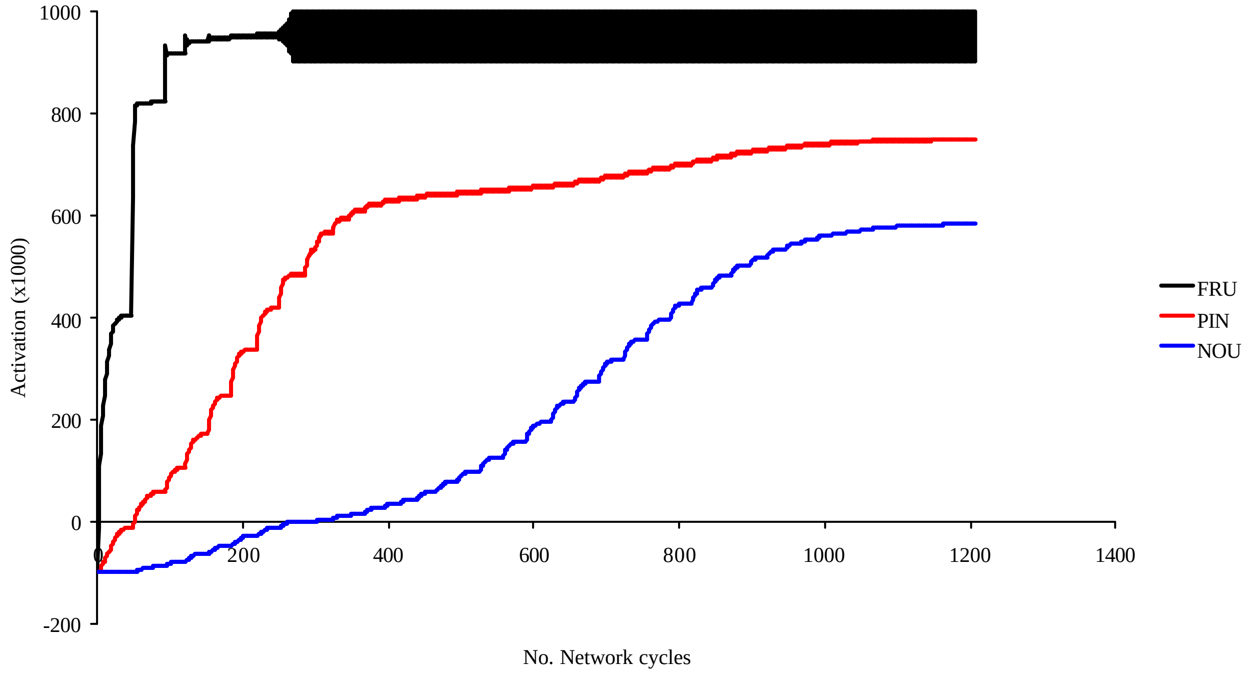
1000
800
600
400
200
0
-200
Activation (x1000)
0
200
400
600
800
1000
1200
1400
No. Network cycles
FRU
PIN
NOU

Figure 5:

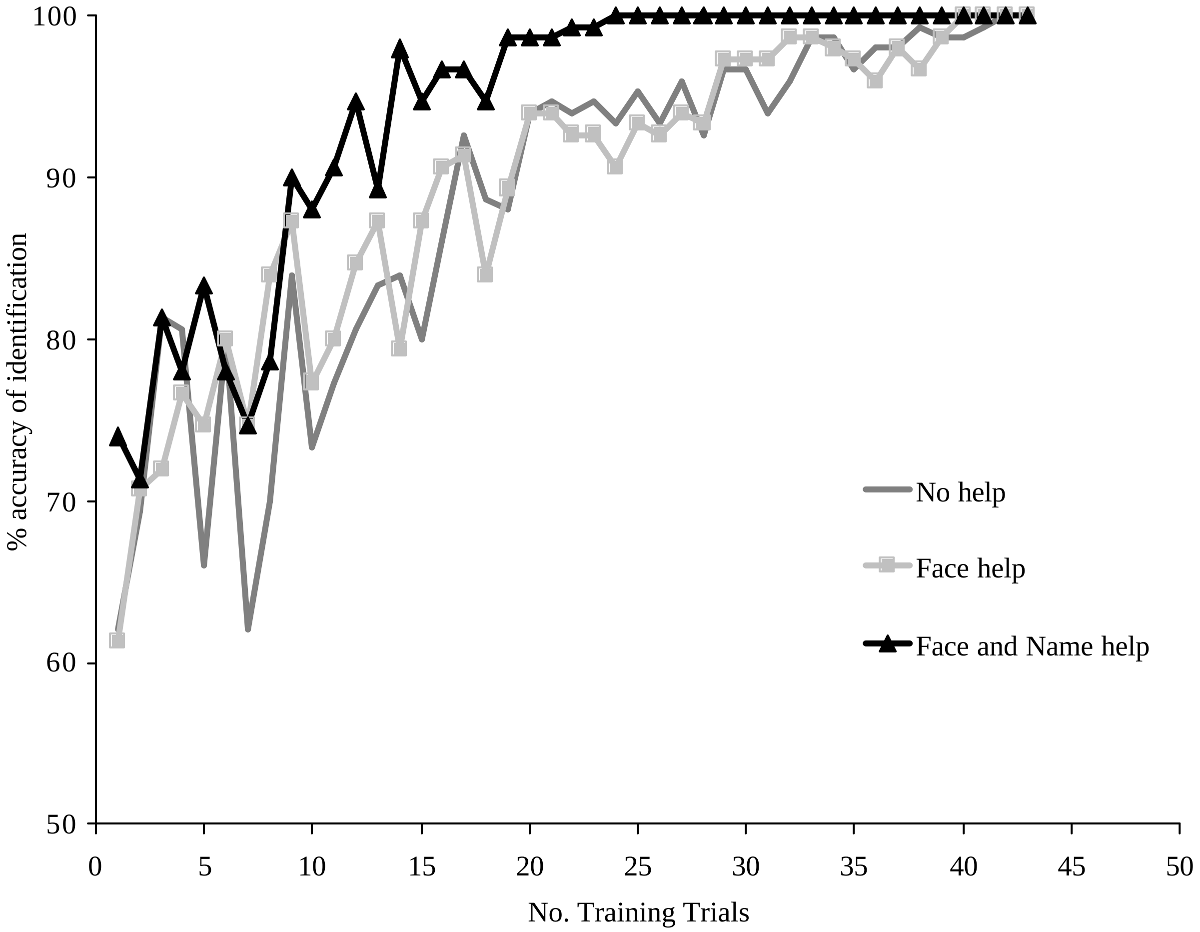
100
90
80
70
60
50
% accuracy of identification
0
5
10
15
20
25
30
35
40
45
50
No. Training Trials
No help
Face help
Face and Name help

Figure 6:

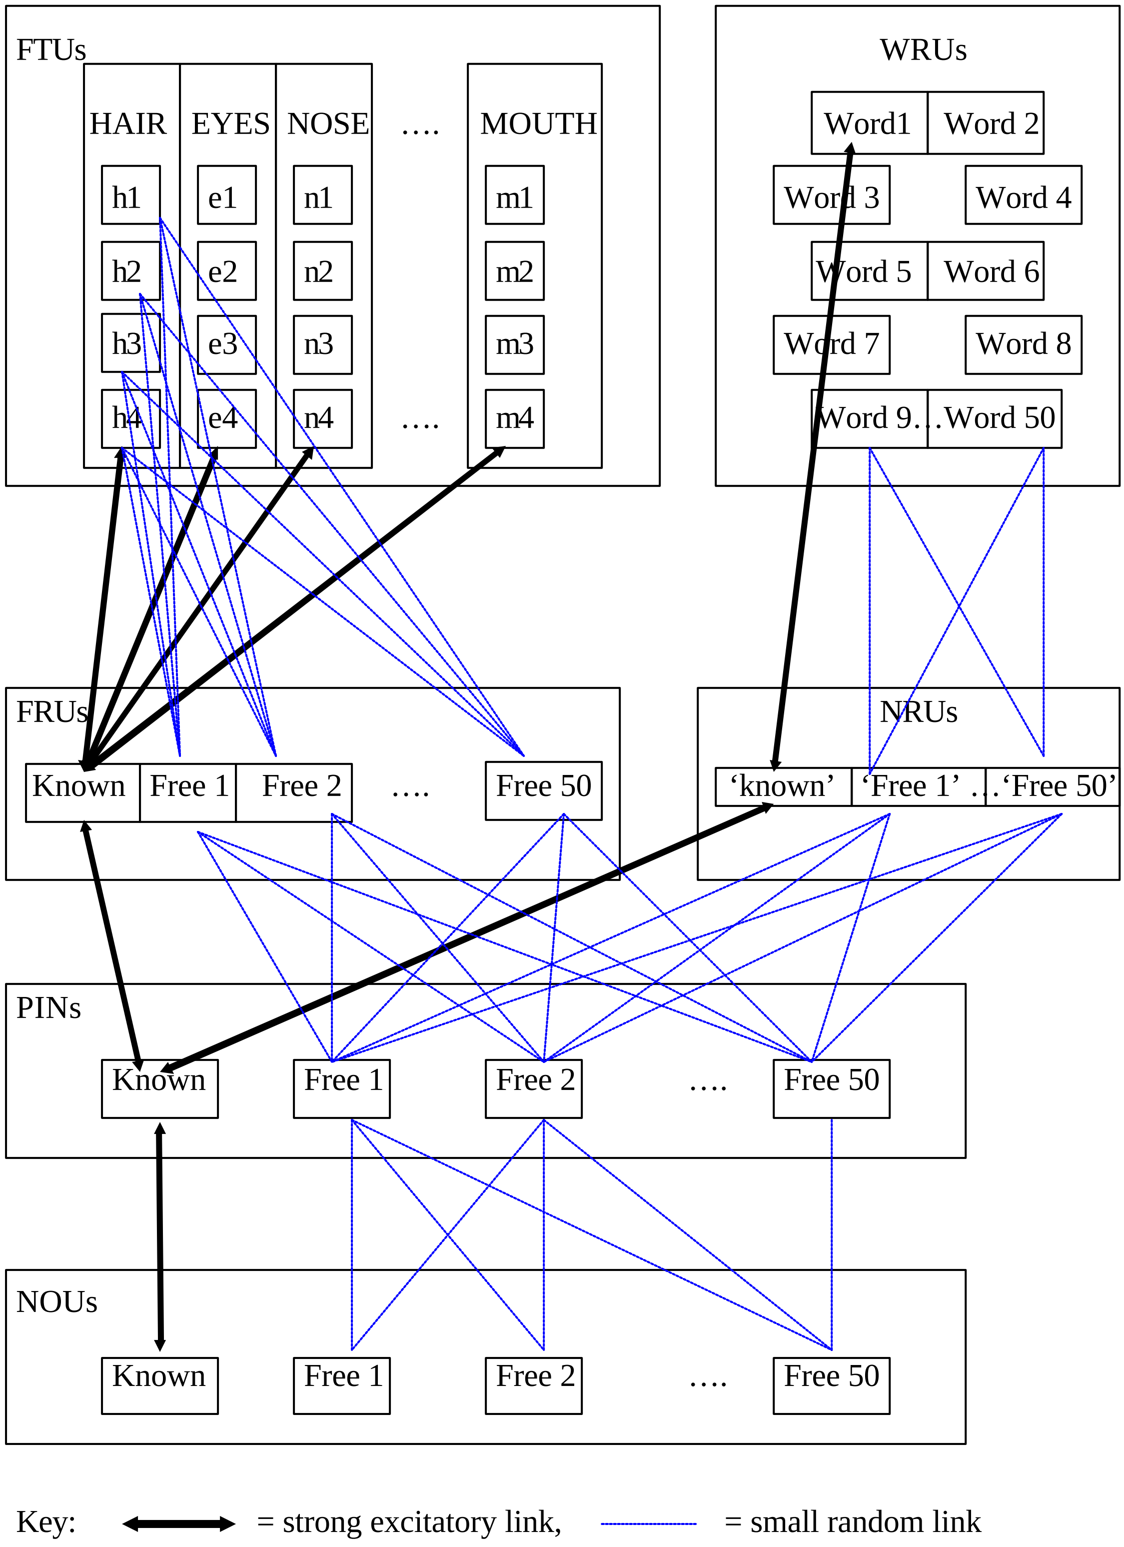
FTUs
HAIR
EYES
NOSE
….
MOUTH
h1
h2
h3
h4
e1
e2
e3
e4
n1
n2
n3
n4
m1
m2
m3
m4
….
WRUs
Word1
Word 2
Word 3
Word 4
Word 5
Word 6
Word 7
Word 8
Word 9…Word 50
FRUs
Known
Free 1
Free 2
….
Free 50
NRUs
'known'
'Free 1' …'Free 50'
PINs
Known
Free 1
Free 2
….
Free 50
NOUs
Known
Free 1
Free 2
….
Free 50
Key:
= strong excitatory link,
= small random link

Figure 7:

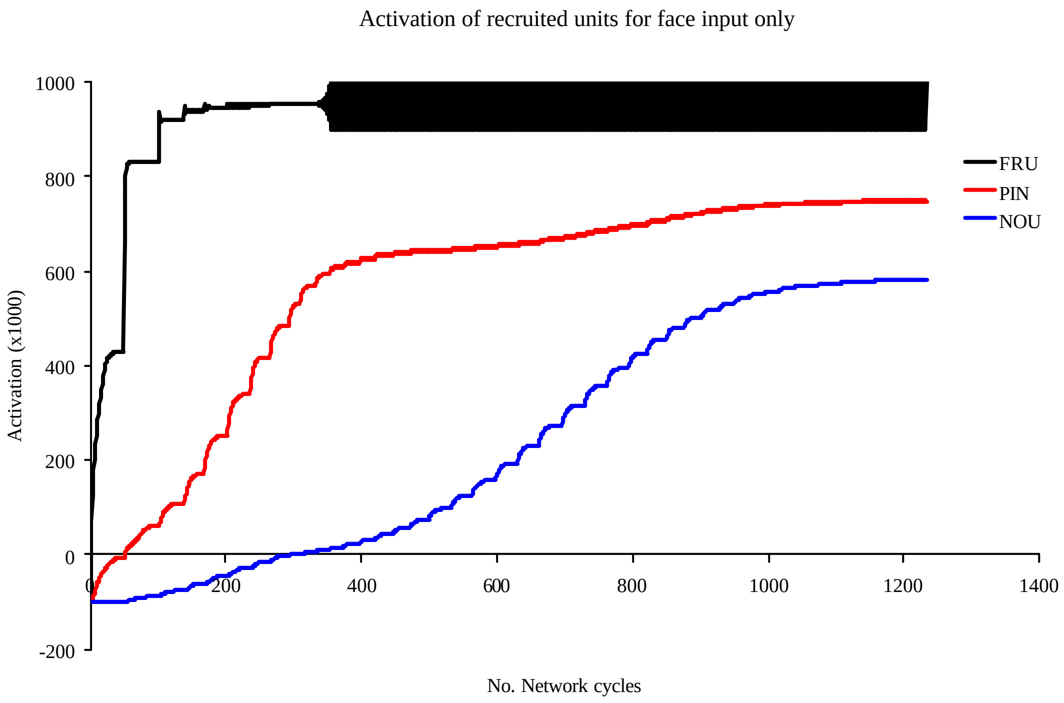
Activation of recruited units for face input only
1000
800
600
400
200
0
-200
Activation (x1000)
0
200
400
600
800
1000
1200
1400
No. Network cycles
FRU
PIN
NOU

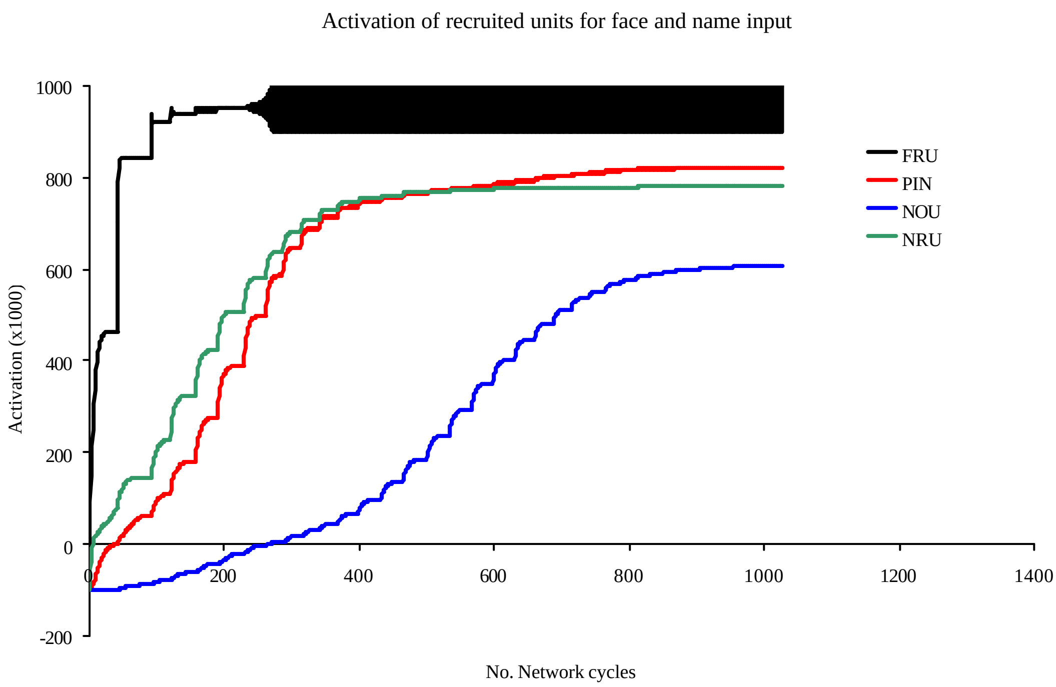
Activation of recruited units for face and name input
1000
800
600
400
200
0
-200
Activation (x1000)
0
200
400
600
800
1000
1200
1400
No. Network cycles
FRU
PIN
NOU
NRU

Footnotes

Footnote 1

Difficulty with naming may also be explained through appealing to either (a) the number of network links involved (assuming an extra node in the representation of proper over common nouns - Burke, MacKay, Worthley & Wade, 1991) or (b) the set size of plausible phonologies (assuming the set size to be greater and thus less easy to learn or to guess for proper names - Brennen, 1993). Both explanations make a useful contribution to the issue of naming difficulties and are reviewed in detail in Valentine, Brennen and Brédart (1996).