A Neural Network Model for Explaining the Asymmetries between Linguistic Production and Linguistic Comprehension

Marco Mirolli (marco.mirolli@istc.cnr.it)
Institute of Cognitive Sciences and Technologies, CNR,
44 Via S. Martino della Battaglia, 00185, Roma, Italy

Federico Cecconi (federico.cecconi@istc.cnr.it)
Institute of Cognitive Sciences and Technologies, CNR,
Via S. Martino della Battaglia 44, 00185, Roma, Italy

Domenico Parisi (domenico.parisi@istc.cnr.it)
Institute of Cognitive Sciences and Technologies, CNR,
Via S. Martino della Battaglia 44, 00185, Roma, Italy

Abstract

Several kinds of empirical evidence point to the existence of an asymmetry between linguistic production and linguistic comprehension: in general, understanding words seems to be easier than producing them. In this contribution we propose a neural model of the relationships between the semantic and the lexical systems. Our model explains the asymmetry between language comprehension and production as an effect of the difference between the dimensions of the brain areas which process semantic and lexical information. In fact, the model's performance in lexical recall is worse than the performance in semantic recall due to the fact that the semantic network is constituted of more computational units (neurons) then the lexical network.

Introduction

A considerable number of empirical evidences of various kind point to the presence of an asymmetry between linguistic production and linguistic comprehension. In general, comprehension seems to be easier than production.

This asymmetry takes several forms. First, the asymmetry is well documented in developmental linguistics: it is in fact a very well known fact that children learn to understand words far earlier than to produce them (see, for example, Bates, Thal, Finlay, & Clancy, 2002).

Another field in which the asymmetry is well documented is the psychology of aging. As people get older, in fact, it gets more and more difficult to retrieve known words (the so called 'tip of the tongue' or TOT phenomenon). On the other hand, there seems to be no decrease in the language understanding capacity whatsoever (see, for example, Burke & MacKay, 1997; Burke, MacKay, & James, 2000).

Furthermore, the fact that word production is more difficult than word comprehension seems to hold in the whole lifetime, and not only in elderliness. In fact, everyone experiences sometimes the tip of the tongue state, and empirical evidences confirm that such a phenomenon just increases with age, but is present also in all normal adults (see, for example, Brown, 1991; Heine, Ober, & Shenaut, 1999).

Finally, there seems to be also some spare evidence that the asymmetry between language production and comprehension exists even in neuropathologies. In this case, the asymmetry takes the form of a larger number of patients with linguistic production deficits than patients with linguistic comprehension ones. For example, it seems that *all* the various kinds of aphasias imply some deficit in word production (anomia), while a dysfunction in linguistic production is not necessarily correlated with comprehension problems (see Bates & Goodman, 1997; Dick, Bates, Wulfeck, Utman, & Dronkers, 2001).

Notwithstanding the fact that all the above mentioned empirical evidences regard very different phenomena, the phenomenon of an asymmetry between linguistic production and linguistic comprehension is common. This suggest that it could have a common cause. In this paper we model the relationships between semantics and the lexicon as the coupling (bi-directional connection) between two associative, Hopfield-like (Hopfield, 1982) networks. With our model we try to explain the asymmetry between production and comprehension of language – in particular, the fact that linguistic production is more difficult than linguistic comprehension – as an effect of the quantitative difference in size of the computational spaces which are devoted to the processing of semantics (word meanings) and the lexicon (word forms), respectively. In particular, our model assumes that the parts of the brain which are devoted to the processing of word forms are smaller, in terms of the number of recruited neurons, than the parts of the brain which are devoted to the processing of word meanings. Our hypothesis is that the production-comprehension asymmetry is determined by this difference in size between the semantical and the lexical areas.

The rest of the paper is structured as follows. In the next section we describe the model which we have used for modelling the semantic and the lexical systems, that is, Hopfield networks, and then we describe the details of our simulations. In the following section we describe the principal results of the simulations, and finally, in the final section, we discuss the strengths and weaknesses of the proposed model together with possible future work.

The Model

We model the linguistic system as the coupling (reciprocal connection) between two auto-associative neural networks, that is, two Hopfield networks (Hopfield, 1982; Rolls & Treves, 1998): the semantic networks and the lexical network.

Hopfield Networks

An Hopfield network is a neural network constituted by a single group of reciprocally connected processing units (neurons). Neuron's activation is bipolar, that is can be either –1 or +1. Every neuron is connected with all the others but not with itself, and connections are symmetrical: in other words, the weight of the connection which links neuron i to neuron j is equal to the weight which links neuron j to neuron i. This symmetry in connection weights is guaranteed by the learning rule which is used for training the network, which is a simple Hebbian rule (Hebb, 1949). Basically, the Hebb rule is such that connections which link neurons with correlated activation (that is, which are usually either both active or both inactive) increase their strength, while connections which link neurons whose activation is uncorrelated decrease their strength. Formally, the Hebb rule is expressed by the following formula:

$$\Delta w_{ij} = r\, a_i\, a_j$$

where Δw_{ij} is the change to the weights connecting unit i to unit j, a_x is the activation of unit x, and r is a constant, called the learning rate.

The network learns to memorize in its connection weights a certain number of activation patterns across its units, with a pattern consisting in a vector of N bipolar (–1 or +1) values, where N is the number of network's nodes. In our case, patterns of activation stand either for internal representations of words' meaningsor for internal representations of words' forms, in the cases of the semantic and lexical networks, respectively.

Weights learning happens in the following way. At the beginning, weights are all set to zero. Then, for each activation pattern, network's nodes are all activated to the bipolar values of that pattern, and weights are modified according to the Hebb rule.

After this kind of training, Hopfield networks have an interesting property. If we present any, even random, pattern as the network's input, and we let the network re-calculate the activation of each node as a consequence of the input coming from other nodes through the connection weights, after a certain number of cycles (usually less than 15), the network converges to a stable state, called an attractor, which doesn't change any more.

To understand Hopfield network's functioning, the activation state of its neurons can be conceived as a surface in a multi-dimensional space, with $N + 1$ dimensions: one dimension for each neuron, plus one dimension representing network's *energy* (Figure 1). Each point of the surface is a possible state of the network, and the local minima of energy (the valleys) are the network's attractors: given any given pattern the network will tend to follow the energy gradient until it reaches one of its valleys, where it will stop indefinitely unless external interference is applied.

The property of having attractors is guaranteed just by the network's rules of connectivity: that connection weights are symmetrical and that there are no self-connections. The role of the above mentioned Hebbian learning procedure is that of shaping the network's activation space so that the patterns on which the network is trained become its attractors. Hence, after successful training, if we present to the network one of the learned patterns, it will be indefinitely maintained. Furthermore, if we present a *partial* pattern, that is a learned pattern in which some percentage of the nodes are set to 0, in a few activation cycles the network is able to perfectly reconstruct the learned pattern, which will be maintained constant in the following cycles.

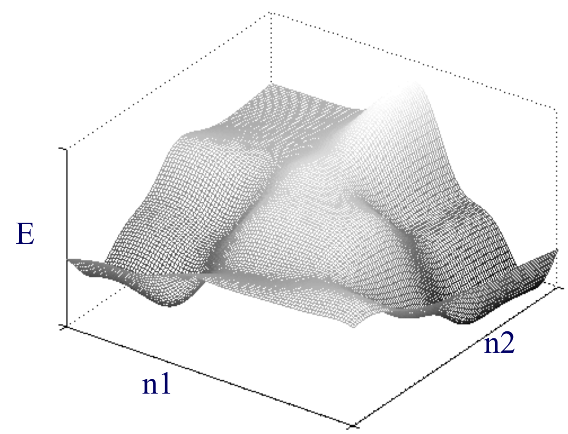

Figure 1: Hypothetical activation space of an Hopfield with just two neurons (n1 and n2). The third dimension (vertical axis) represents the energy (E) of the point which stands for the corresponding state of the network. Valleys represent network's stable states (attractors), towards which the activity of the network converges.

The choice of modeling the semantical and the lexical system through Hopfield networks is due to the fact that this kind of networks represent (reasonably) good models of the functioning of single brain areas (Rolls & Treves, 1998). In fact, beyond possessing all those bio-mimetic properties which are shared by all kinds of parallel-distributed-processing neural networks (Plaut & Shallice, 1993; Rumelhart, McClelland & the PDP Research Group, 1986), such as robustness to noise, graceful degradation, and pattern completion, Hopfield networks have at least two other properties which make them biologically plausible. First, the Hebbian learning rule, of which the neural implementation has been found in the phenomena of synaptic long term potentiation and lond term depression (Kelso, Ganong & Brown, 1986; Stanton, & Sejnowski, 1989). Second, connections' recurrencies, that is the fact that neurons within the same group are reciprocally interconnected, represents an important anatomic characteristic of the brain. Furthermore, it is the presence of these recurrencies that makes it possible for the network to show temporal dynamics which permits to model important empirical phenomena like semantic priming or, more generally, differences in reaction times (see, for example, Masson, 1995; Sharkey & Sharkey, 1992).

The Semantic and the Lexical Networks

Our model consists in two Hopfield networks which represent, respectively, the semantic and the lexical systems. The two networks function in just the same way. The only difference between the two lies in their dimensionality. In fact, the semantic network is constituted by 2500 nodes, while the lexical network is constituted by 500 nodes. This difference correspond to the assumption that the parts of the brain which are devoted to the processing of word forms are much smaller than the parts of the brain which are devoted to the processing of word meanings. The two networks are coupled, that is mutually interconnected (Figure 2). But while nodes which belong to the same network are fully connected between each other, the probability p for two nodes belonging to different networks is very low (in between 0.02 and 0.001, see below).

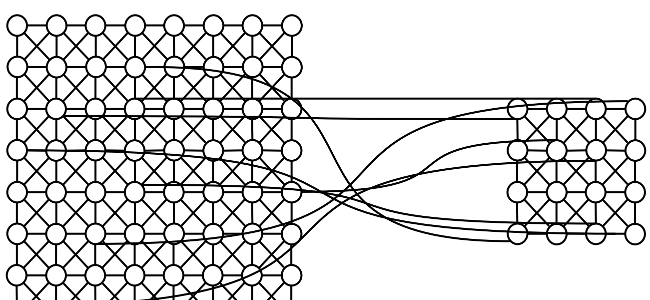

Figure 2: Schematic representation of the model. Circles represent nodes (neurons), while lines represent bidirectional connection between nodes. Within the same network neurons are fully interconnected, while the connectivity between the two networks is much lower. See text for details.

Each sub-network is trained with 50 different patterns: the semantic network learns to memorize 50 patterns of 2500 values, which each pattern representing a 'meaning'. The lexical network learns to memorize 50 patterns of 500 values, with each pattern representing a 'word' (form). All patterns are generated randomly, with a uniform probability, for each node, to be either in the active (1) or inactive (–1) state.

Each word form is associated with one word meaning, and the whole network, formed by the two semantical and lexical networks, learns such association in the following way. For each word form-word meaning pair, the corresponding word form pattern is presented to the lexical network, while, simultaneously, the corresponding word meaning pattern is presented to the semantic network. Connection weights which link the nodes of the two sub-networks are learned through the Hebb rule described above, in a way which is completely analogous to the learning of the intra-network connection weights.

Results

In our model (correct) linguistic production consists in activating the pattern in the lexical network (the word form) which corresponds to the pattern which is present in the semantic network (the word meaning). Viceversa, linguistic comprehension consists in activating, the semantic network, the meaning which corresponds to the word form which is present in the lexical network.

We have network's performance under two conditions: lexical recall and semantic recall. The lexical recall test, which is meant to test linguistic production, is done in the following way. For each pattern representing a meaning, that pattern is presented to the semantic network while the activations of all the nodes of the lexical network are set to zero. The whole network is then let relax, that is, nodes' activations are updated until the net has reached an attractor. At this point, we verify whether the pattern which is present in the lexical network is the correct one, that is, the pattern representing the word forms which corresponds to the given meaning. The semantic recall test, which is meant to test linguistic comprehension, is done in an analogous way: given a word form as the input to the lexical network, we check whether the recalled semantic pattern is the correct one.

In order to check for the presence of an asymmetry between linguistic production and linguistic comprehension in our model we have studied the behavior of the whole network with various levels of connectivity between the two sub-networks: in particular, with inter-networks connectivity in between 0.020 and 0.001.

For both conditions of semantic and lexical recall, and for each connectivity degree, we take two measures: (a) the percentage of correctly recalled patterns, and (b) the average recall time of the correctly recalled patterns (this is measured as the number of cycles which are necessary for the network to reach a stable point).

Figure 3 shows the percentage of correctly recalled patterns in the two experimental conditions as a function of the inter-networks connectivity, while Figure 4 shows average recall time. Reported data of both figures represent average results of 10 replications of each test with different random values.

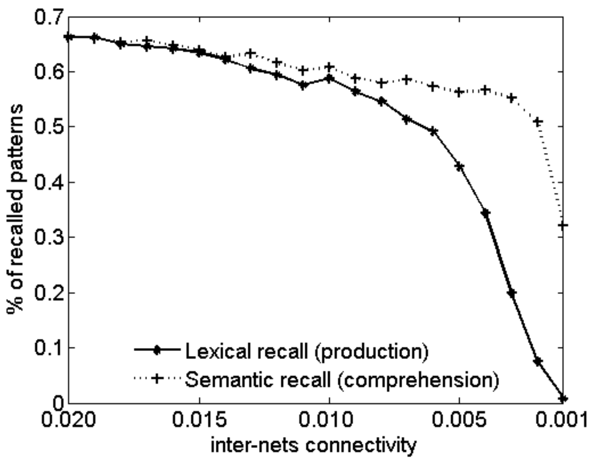

Figure 3: Percentage of correctly recalled patterns in both the lexical and the semantic recall tests as a function of inter-networks connectivity.

Generally speaking, the results show a clear asymmetry between production and comprehension, with a pattern simular to that which is found in reality: in fact, word production results to be more difficult than word comprehension, both in terms of number of correctly recalled patterns (Figure 3) and in terms of recall times (Figure 4).

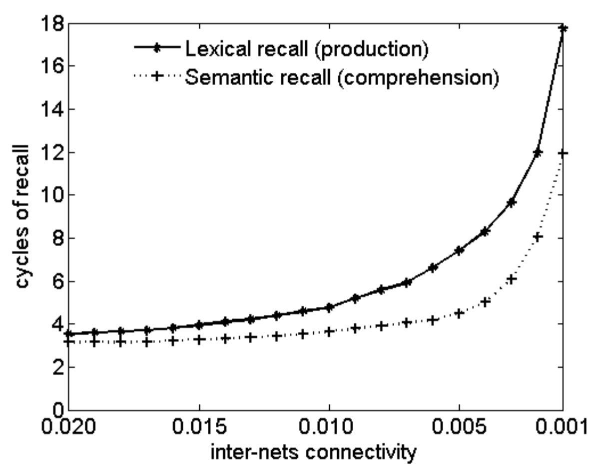

Figure 4: Average recall times of correctly recalled patterns in both the lexical and the semantic recall tests as a function of inter-networks connectivity.

Let's compare, for example, the results of the two conditions with a inter-nets connectivity degree of 0.002. In this case, while given a word form the correct meaning is recalled in about 50% of cases, given a meaning the probability of recalling the corresponding word form is just about 8%. Furthermore, with the same connectivity degree, while the average recall time in the semantic recall test is about 8 activation cycles, the average recall time in the lexical recall test is about 12 cycles.

Another interesting result of our simulations regards the different pattern which is shown by the production and comprehension curves as inter-nets connectivity decreases. In fact, the differences between lexical and semantic recall seem to increase as the connectivity degree decreases. This is in line with the fact, mentioned in the introduction, that while the asymmetry between linguistic production and linguistic comprehension is present during the whole lifetime, the asymmetry sharpens with aging, with a progressive impairment of linguistic production capabilities paralleled by an apparent preservation of linguistic comprehension capabilities. Our model suggests that this sharpening of the production-comprehension asymmetry could be due to the progressive decrease in the connectivity between the semantic and the lexical systems due to degenerative mechanisms related to aging.

Discussion and Conclusion

The model we have proposed in this paper represents an attempt to explain the asymmetry between production and comprehension of language which seems to take various forms: several different lines of evidence point to the fact that producing words is more difficult than understanding them, either for children, for normal adults, for elders, and, probably, also for patients with cerebral lesions. All these lines of evidence regard quite different phenomena and it is very likely that the various forms of production-comprehension asymmetries depend, in part, on different factors. Just to make an example, some degree of asymmetry between production and comprehension in child linguistic development is certainly due to the fact that words are socially learned. In fact, the capacity of comprehending a word, that is the ability to understand in which context that word is correctly applied, is of course a pre-requisite for the ability to produce it correctly, that is in the appropriate context. In other words, you cannot appropriately produce words which you do not understand. Hence, it seems inevitable that the number of understood words represent the superior limit with respect to the number of produced ones, and that linguistic production follow chronologically linguistic comprehension.

Notwithstanding the different factors which may underly the various forms which the production-comprehension asymmetry can take, it is also possible that these various forms have also a common cause.

In this contribution we proposed a neural model of the relationships between the semantic and the lexical systems which explains the greater difficulty in producing words than in comprehending them as a an effect of the size difference between the computational spaces which are devoted to the processing of word forms and word meanings, respectively. In fact, by just assuming that the brain regions devoted to semantic processing are considerably bigger (in terms of number of neurons involved) than those devoted to lexical processing, we have shown that lexical recall is more difficult than semantic recall, and that a decrease in the connections between the two areas leads to a decrease in lexical recall, while semantic recall remains more stable.

Obviously, our model suffer several limitations. A first limitation lies in the fact that in our model the asymmetry can be observed only when the network is able to recall, in both directions, only about 70% of the word forms or the word meanings of its repertoire. In reality, the asymmetry can be observed earlier: for example, elderly people begin to suffer word-finding problems when they are still able to comprehend all the words that they know. We are currently experimenting various ways in which this limitation might be overcome: for example, one possibility is to increase the size difference between the semantic and the lexical networks; another one is to decrease intra-networks connectivity (in our model, as in standard Hopfield networks, there is a full intra-net connectivity, but in the real brain the connectivity is certainly lower); still another possibility consists in abandoning the requirement of symmetrical connections which holds for standard Hopfield networks (in other words, a possibility is to make it possible for a neuron i to be connected with a neuron j without requiring that neuron j be connected to i: this is another change which would render the model more biologically realistic).

There are also more general limitations of our model. For example, in the current model the semantic system has no internal structure whatsoever. Meanings are represented, as it is common in neural network research, as patterns of activation distributed on a single group of neurons, which stands for the 'semantic system'. We already know that this is a strong oversimplification. For example, we know that different kinds of words, like names and verbs, are represented, in the real brain, in distinct groups of neurons (see, for example, Caramazza & Hillis, 1991). And even words of the same kind seem to be processed, at least partially, in different parts of the brain depending on their meaning (Martin, Wiggs, Ungerleider & Haxby, 1996; Plaut, 2002; Pulvermüller, 1999; Tettamanti, Buccino, Saccuman, Gallese, Danna, Scifo, Fazio, Rizzolatti & Cappa, 2005). Furthermore, the connectivity between different sub-groups of neurons devoted to processing different kinds of words it is certainly not uniform: not every group is connected to every other, or, at least, the connectivity degree between different pairs of groups is certainly different. As our model have already shown that the connectivity between groups influences inter-networks dynamics, an important way to improve the model is to consider this kind of empirical constrains. This possible line of future research might also allow us to explain other important phenomena, like the various forms of double-dissociations which are found language neuropsychology (see, for example, Caramazza & Mahon, 2005).

Another, even more general, limitation of our model consists in the fact that semantic representations (and also linguistic ones) are provided by us, the researchers, instead of being learnt (developed) by the network itself, as happens both in reality and in many other connectionist simulations. In order to solve this kind of limitation the model needs to be modified considerably, in that Hopfield networks, being constituted of a single group of neurons, cannot develop internal representations by themselves which would mediate between network's input and output, as happens in layered networks. Consequently, Hopfield networks require that memorized patterns be generated and provided directly by the researcher.

Notwithstanding the limitations we have just discussed, we claim that our model represents an important first attempt to find an explanation to the very well documented asymmetry between linguistic production and linguistic comprehension. In this respect, still another line of possible future research consists in trying to give a more detailed model of the various forms which the asymmetry can take. For example, we could differentiate between two possible different causes (beyond the size difference between the semantic and the lexical systems) which might underly the Tip-Of-the-Tongue phenomenon in young and elder adults. The first cause, common to both young and elder people, would be the presence of noise in neural transmission, which we know to be largely present in real brains and which could be easily simulated in various ways, for example by adding some noise in the information transfer between simulated neurons and/or in neurons' activation function. The second cause, which might underly the progressive deterioration of linguistic production in elderly people would lie, as already shown in our current model, in the decrease of connectivity both within and between the semantic and the lexical networks due to neuro-degenerative processes associated with aging.

Acknowledgments

The research presented in this paper has been supported by the ECAGENTS project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract IST-2003-1940. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that may be made of data appearing in this publication.

References

Bates, E., & Goodman, J.C. (1997). On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia, and real-time processing. *Language and Cognitive Processes*, 12(5-6), 507-584.

Bates, E., Thal, D., Finlay, B.L., & Clancy, B. (2002). Early language development and its neural correlates. In I. Rapin, & S. Segalowitz (eds.), *Handbook of neuropsychology*, Vol. 7: Child neurology. Amsterdam: Elsevier.

Brown, A.S. (1991). A review of the tip-of-the-tongue experience. *Psychological Bulletin*, 109(2), 204-223.
Burke, D.M., & MacKay, D.G. (1997). Memory, language and ageing. *Philosophical Transactions of the Royal Society: Biological Sciences*, 352, 1845-1856.
Burke, D.M., MacKay, D.G., & James, L.E. (2000). Theoretical approaches to language and aging. In T. Perfect, & E. Maylor (Eds.), *Models of cognitive aging*. Oxford: Oxford University Press.
Caramazza, A.H., & Hillis, A.E. (1991). Lexical organization of nouns and verbs in the brain. *Nature*, 349, 788-790.
Caramazza, A., & Mahon, B. (2005). The organisation of conceptual knowledge in the brain: the future's past and some future directions. *Cognitive Neuropsychology*, 22, 1-25.
Dick, F., Bates, E., Wulfeck, B., Utman, J., & Dronkers, N. (2001). Language deficits, localization, and grammar: Evidence for a distributive model of language breakdown in aphasics and normals. *Psychological Review*, 108 (4), 759-788.
Hebb, D. O. (1949). *The Organization of Behavior*. New York: John Wiley.
Heine, M., Ober, B., & Shenaut, G. (1999). Naturally occurring and experimentally induced tip-of-the-tongue experiences in three adult age groups. *Psychology and Aging*, 14(3), 445-457.
Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 9, 2554–2558.
Kelso, S., Ganong, A., & Brown, T. (1986). Hebbian synapses in hippocampus. *Proceedings of the National Academy of sciences of the USA*, 83, 5326-5330.
Plaut, D.C. (2002). Graded modality-specific specialization in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*, 19(7), 603-639.
Plaut, D.C., & Shallice, T. (1993) Deep dyslexia: A case study of connectionist neuropsychology, *Cognitive Neuropsychology,* 10(5), 377-500.
Rolls, E. T., & Treves, A. (1998). *Neural Networks and Brain Function*. Oxford: Oxford University Press.
Rumelhart, D., McClelland, J. & the PDP Research Group (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1 & 2, Cambridge, MA: MIT Press.
Sharkey, A.J.C., & Sharkey, N.E. (1992) Weak contextual constraints in text and word priming. *Journal of Memory and Language*, 31, 543-572.
Stanton, P., & Sejnowski, T. (1989). Associative long-term depression in the hippocampus induced by hebbian covariance, *Nature*, 339(6221), 215-218.
Martin, A., Wiggs, C., Ungerleider, L., & Haxby, J. (1996). Neural correlates of category-specific knowledge. *Nature*, 379(6566), 649-52.
Pulvermüller, F. (1999). Words in the Brain's Language. *Behavioral and Brain Sciences*, 22, 253-279.
Tettamanti, M., Buccino, G., Saccuman, M., Gallese, V., Danna, M., Scifo, P., Fazio, F., Rizzolatti, G., & Cappa, S.P. (2005). Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience*, 17, 273-281.