Available online at www.sciencedirect.com
SCIENCE DIRECT®
NEUROCOMPUTING
Neurocomputing 58–60 (2004) 599 – 605
www.elsevier.com/locate/neucom

Increasing number of objects impairs binding in visual working memory

Gwendid T. van der Voort van der Kleij[a,*], Marc de Kamps[b], Frank van der Velde[a]

[a] *Cognitive Psychology, Leiden University, Wassenaarsweg 52, Leiden 2333, AK, The Netherlands*
[b] *Department of Informatics, Robotics and Embedded Systems, TU München, Garching bei München, Germany*

Abstract

The number of objects that can be maintained in visual working memory without interference is limited. We present simulations of a model of visual working memory in ventral prefrontal cortex that has this constraint as well. One layer in ventral PFC represents all objects in memory. These representations are used to bind the features of the objects. If there are too many objects, their representations interfere and therefore the quality of the representations degrades. Consequently, it becomes harder to bind the shape to location for an object that is maintained.

Keywords: Binding; Visual working memory; Prefrontal cortex; Blackboard architecture; Attention

1. Introduction

Investigations [5] have shown that humans have the ability to maintain a number of visual objects in visual working memory. A remarkable characteristic of this finding is that the number of objects that can be maintained in working memory without interference (i.e., loss of information) is limited (to about four), but the number of object features (e.g., shape, color, location, motion, etc.) is unlimited for each of the objects. We presented a model of visual working memory in prefrontal cortex (PFC) that theoretically can explain this characteristic [1]. A basic characteristic of this model is a 'blackboard' that links different 'processors' to one another [2]. Objects in working memory are represented in the blackboard. One layer in ventral PFC functions as the

* Corresponding author. Tel./fax: +31-71-52-73-783.
E-mail address: gvdvoort@fsw.leidenuniv.nl (G.T. van der Voort van der Kleij).

doi:10.1016/j.neucom.2004.01.101

blackboard, containing representations that consist of conjunctions of (partial) 'identity' information and location information. This blackboard serves to bind the information processed in each of the specialized processors. The processors in this case are networks for feature identification. When too many objects are put in working memory, their representations in the blackboard interfere. Consequently, an object's representation in the blackboard muddles and the blackboard's performance to bind the features of an object degrades.

After getting deeper into this model of visual working memory, we present simulations in which information about the location of an object is used to bind its shape. In line with previous simulations, which explored the opposite binding route from shape to location [3], the results reflect our expectations that the model is limited in the number of visual objects that it can maintain without interference complicating correct binding.

2. Blackboard architecture of visual working memory in PFC

Our model of visual working memory in PFC is based on a neural blackboard architecture that is used in a simulation of object-based attention in the visual cortex [4]. We assume that the neural blackboard architecture is located in the ventral prefrontal cortex (V-PFC) [1]. This is in line with human neuroimaging studies and recent monkey studies [6]. Activation in V-PFC is sustained (reverberating) activation, characteristic of working memory activation in the cortex.

In the model (Fig. 1, left), the V-PFC has a layered structure with representations similar to the representations in the visual (temporal) cortex. First, the posterior infero-temporal cortex (PIT) connects to the blackboard. As in PIT itself, the representations in this layer of V-PFC consist of conjunctions of location and (partial) identity

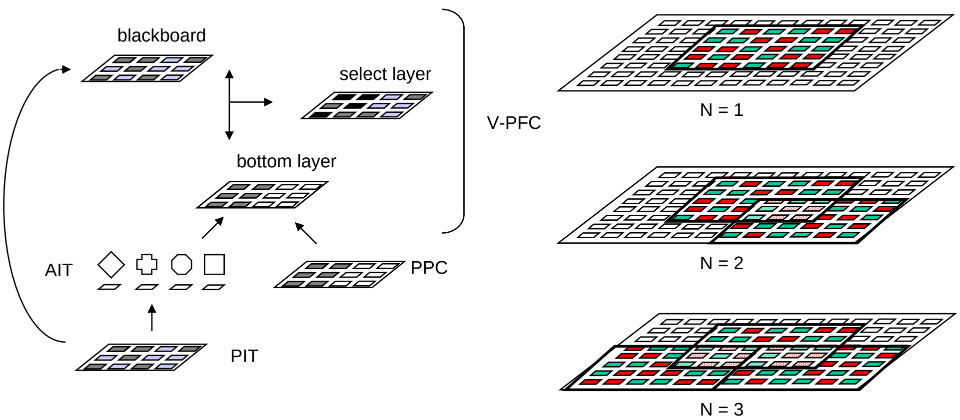

Fig. 1. (Left) A blackboard architecture in prefrontal cortex (PFC). PIT = posterior infero-temporal cortex; AIT = anterior infero-temporal cortex; V-PFC = ventral prefrontal cortex. (Right) Interference between object representations in the blackboard. For explanation see text.

(object-feature) representations (shape, color). In turn, the bottom layer of V-PFC is connected to higher-level areas in the visual cortex like the anterior infero-temporal cortex (AIT) and the posterior parietal cortex (PPC), which process, respectively, the shape and location information of an object.

The connections from these higher-level areas to the bottom layer of V-PFC are similar to the connections in the feedback network of the visual cortex [4]. They associate all possible representations that are selective for an activated feature (e.g., shape, location). For example, if one shape is selected in AIT, then all representations in the bottom layer of V-PFC that are consistent with that shape (on every possible position) are activated. Note that these connections have a fan-out structure. Likewise, an attended location in PPC activates all possible representations (e.g., for any shape) in the bottom layer of V-PFC on that location in (visual) space. The bottom layer of V-PFC thus represents the current focus of attention, whether this is based on location or (location-invariant) feature information. Consequently, interaction between the bottom layer of V-PFC and the blackboard can select the object representation that is consistent with the current attentional focus. The resulting activation in the select layer can be used to bind the features of this object [1].

The interaction between the bottom layer and the blackboard requires some discussion. As the nature of attentional modulation is being debated, the model does not include a clear perspective on this part. Instead, we have taken a more pragmatic stand to simulate, approximately, two competing hypotheses. Attention may either increase the sensitivity for attended features by providing an extra input to neurons representing those, or may boost the response strength for attended features without changing the sensitivity to them. We will refer to the former mechanism as additive and to the latter as multiplicative. Logically, though this is not simulated here, attention may involve a combination of both mechanisms as well.

3. Feature binding in working memory

The nature of the representations in V-PFC and the connections with the higher-level areas in the visual cortex produces the behavioral effects described before. The blackboard architecture of V-PFC results in a binding of the feature representations of the objects maintained in memory. Therefore, the features of an object can be retrieved (selected) in working memory as long as the representations of the objects stored in V-PFC do not interfere. However, when too many objects are present in a display, their representations in V-PFC will interfere, which results in loss of information (Fig. 1, right). As more objects are present in a display, the amount of interference increases, and it can be expected that the quality of the representation of an object in V-PFC becomes less. As a consequence, it becomes harder to correctly bind the feature representations of the objects that are maintained in memory. V-PFC might end up binding wrong feature representations for an object that is attended to. Following simulations tested whether our model of the visual working memory shows this behavior.

4. Simulations

For the simulations, we linked the V-PFC model with a (trained) neural network model of the ventral pathway in the visual cortex that is used in the simulation of object-based attention in the visual cortex [4]. This model consists of a feed-forward network that includes the areas V1, V2, V4, PIT and AIT, and a feed-back network that carries information about the identity of the objects to the lower areas in the visual cortex (V1-PIT). The model shares the basic architecture and characteristics (i.e., the nature of the representations) of the visual cortex. The feed-forward neural network was trained to identify 9 different objects on 9 possible positions (using backpropagation). This was done successfully five times, each time resulting in slightly different connection weights between the layers. The simulations explored the selection process in the V-PFC model that involves location information. We expected that information about the location of an object becomes less adequate to bind the object's shape as the number of objects stored in memory increases.

During simulations, displays consisting of N (different shaped) objects, with N ranging from 2 to 9, are presented to V1. The objects, presented in separate, non-overlapping, positions, are processed in the visual cortex, and their PIT representations also activate the representations in the blackboard in V-PFC. The location of one of the objects is selected (attended) in PPC (e.g., due to a competition between all object locations). The activation coding for this location in PPC activates its corresponding location in the bottom layer of V-PFC. As a result, the interaction between the bottom layer and the blackboard modulates the object representation in the select layer that is selective for the attended location. The activation in the select layer is processed further by AIT to identify the object's shape.

For simplicity, the activity in PPC that represents a certain location after competition between all object locations, its one-to-one connections to the bottom layer of V-PFC, and the interaction between the blackboard and the bottom layer are simulated altogether in one step by modulating the object representation in the blackboard at the attended location. To implement the last step regarding the binding of the object's shape, the blackboard layer served as input to area AIT, which is trained to identify shape information. A winner-takes-all mechanism in AIT selects the identified shape.

Location information modulated the representation in the blackboard in two qualitatively different ways during separate runs. In multiplicative runs, the activity of neurons representing the attended location in the blackboard was multiplied by a certain factor. Alternatively, in additive runs, these neurons were given extra input, and new activation values were accordingly computed. To ensure results that are sufficiently robust, multiplicative and additive runs were done with varying modulation strength from, respectively, 1 to 2 and 0 to 0.5, with a similar step size of 0.05. In additive runs, the range of extra input was chosen to balance apparent levels of sensory input. For each N, 90 random displays are presented to each instance of the model.

5. Results

Fig. 2 shows the probability of successful binding over the number of objects in visual working memory and modulation strength, for both additive and multiplicative runs. For each number of objects in working memory, data of all 5 instances of the neural network model are averaged over all relevant trials. Note that a modulation strength of 0 in the additive runs and of 1 in the multiplicative runs actually means that there is no selection by location information at all. Hence, the proportion of correct binding for each N should equal chance level. The figure indeed reflects this fact. Interestingly, we see that a slight increase in modulation strength immediately improves binding. Nevertheless, there appears to be a limit in the benefit of increasing the modulation strength. This makes sense as modulated neurons reach their maximum firing rate at some point. Moreover, modulation strength also affects unattended, overlapping object representations. Both for additive and multiplicative runs, binding is better when the number of objects held in working memory is low, even for quite high values of modulation strength. In other words, as the number of objects increases, it becomes less reliable to select an object's shape based on its location information. Hence, the binding process starts breaking down. Comparing the additive and multiplicative runs, we see that the latter shows slightly better binding (i.e., boosting the output of neurons enables better binding than increasing the input). This makes sense as multiplication amplifies the representation in the blackboard without affecting its structure, while adding does modify the structure of the representation to some extent.

So far we have assumed that the representation in the blackboard is identical to the one in PIT. However, this is not likely to be true. It is possible that the representation in the blackboard is reduced compared to PIT. New simulations explored the binding power of the model given a sparse and reduced representation in the blackboard.

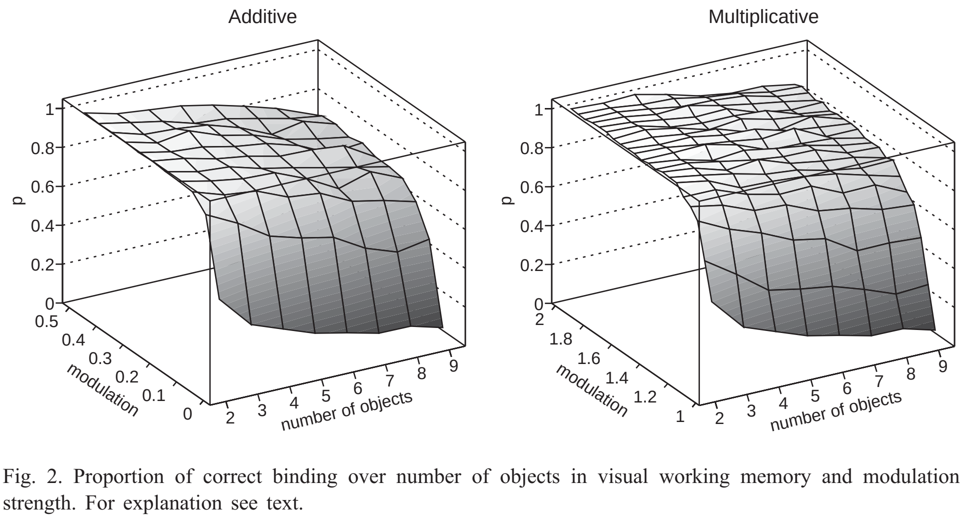

Fig. 2. Proportion of correct binding over number of objects in visual working memory and modulation strength. For explanation see text.

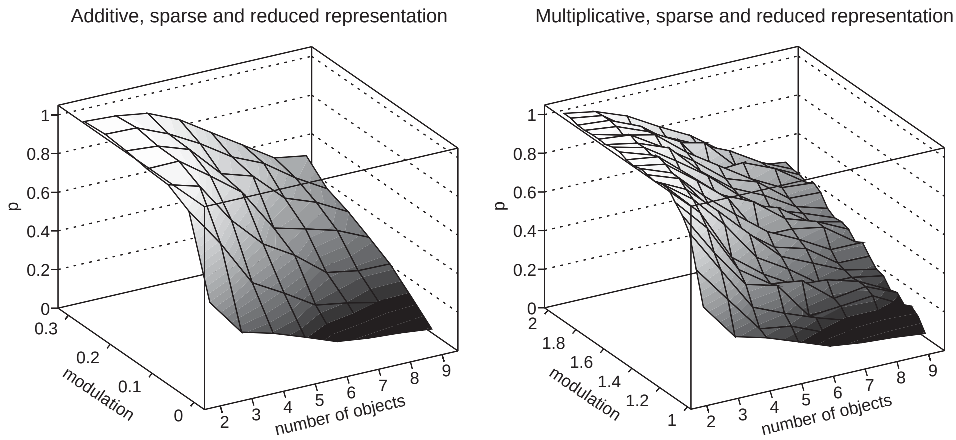

Fig. 3. Proportion of correct binding over number of objects in visual working memory and modulation strength, given a sparse and reduced representation in the blackboard. For explanation see text.

Before the location information of one object modulated the activity in the blackboard, a competition mechanism in the blackboard reduced its representation and made it sparse. Subtracting an inhibitory input from each neuron's input, which allows 30% of the neurons to be active, and computing new activation values, implemented this. In additive runs, the modulation strength now ranged from 0 to 0.3 to balance lower sensory input.

Fig. 3 shows the probability of successful binding over the number of objects in visual working memory and modulation strength, for these runs. We see that even when the representation in the blackboard is sparse and reduced compared to the one in PIT, it can still bind the shape to the location of an object considerably when the number of objects in memory is low. As expected, for higher number of objects the binding impairment already seen in former runs is amplified, as a higher number of objects leads to more competition and thus to a more reduced and sparse representation in the blackboard.

6. Discussion

The simulations point out that the model of visual working memory that we presented is limited in the number of objects that it can maintain in memory without interference (i.e., loss of information). Our model cannot successfully bind the feature(s) of the attended object anymore as it gets loaded with more objects. This is in accordance with findings about visual working memory and previous simulations. Naturally our simulations are of a qualitative nature. The fact that there is a limited number of visual objects that people can maintain in visual working memory is (probably) inherent to its architecture. The model that we presented shares this characteristic.

Our model predicts that this limit is also partly dependent on the distance between objects in a display [3]. Another prediction from our model is that the resolution of spatial attention is comparably limited in other tasks than visual working memory. Selection by location information is dependent on the amount of interference between object representations in the ventral pathway of the visual cortex. Note that it does not matter whether spatial attention (also) acts upon areas with a higher spatial resolution (e.g., V1 or V2), when areas like V4 and PIT, due to their conjunction representations, are still used to bind object's features. Selecting an object by a more centered focus (e.g., a Gaussian) of its location may overcome some interference between object representations. However, it also risks ignoring important information.

References

[1] F. van der Velde, M. de Kamps, A model of visual working memory in PFC, Neurocomputing, 52–54 (2003) 419–424.

[2] F. van der Velde, On the use of computation in modelling behavior, Network: Comput. Neural Systems 8 (1997) 1–32.

[3] G.T. van der Voort van der Kleij, M. de Kamps, F. van der Velde, A neural model of binding and capacity in visual working memory, in: Proceedings of the ICANN/ICONIP, Lecture Notes in Computer Science, Vol. 2714, Springer, Berlin, 2003, pp. 771–778.

[4] F. van der Velde, M. de Kamps, From knowing what to knowing where: modeling object-based attention with feedback disinhibition of activation, J. Cognitive Neurosci. 13 (4) (2001) 479–491.

[5] E.K. Vogel, G.F. Woodman, S.J. Luck, Storage of features, conjunctions, and objects in visual working memory, J. Exp. Psychol.: HPP 27 (2001) 92–114.

[6] F.A.W. Wilson, S.P.O. Scalaidhe, P.S. Goldman Rakic, Dissociation of object and spatial processing domains in primate prefrontal cortex, Science 260 (1993) 1955–1957.