



Universiteit
Leiden
The Netherlands

Commonalities of feature integration processing in and across perception and action planning

Colzato, L.S.

Citation

Colzato, L. S. (2005, June 14). *Commonalities of feature integration processing in and across perception and action planning*. Retrieved from <https://hdl.handle.net/1887/2697>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/2697>

Note: To cite this publication please use the final published version (if applicable).

Commonalities of feature integration processing in and across perception and action planning

Lorenza S. Colzato

Cover design by Aernout Tas & Lorenza S. Colzato
Photography “binding” by Lorenza S. Colzato

Commonalities of feature integration processing in and across perception and action planning

Proefschrift

ter verkrijging van

de graad van Doctor aan de Universiteit Leiden,
op gezag van de Rector Magnificus Dr.D.D. Breimer,
hoogleraar in de faculteit der Wiskunde en
Natuurwetenschappen en die der Geneeskunde,
volgens besluit van het College voor Promoties
te verdedigen op dinsdag 14 juni 2005
klokke 15.15 uur

door

Lorenza Serena Colzato

geboren te Bolzano (Italië)
in 1974

Promotiecommissie

Promotor: Prof. Dr. A. H. C. van der Heijden

Co-promotor: Dr. G. Wolters

Referent: Prof. Dr. K. R. Ridderinkhof

Overige leden: Prof. Dr. G.A.M. Kempen
Dr. F. v.d. Velde
Dr. S. Nieuwenhuis

Een mens kan nooit weten wat hij wil,
omdat hij maar een leven heeft dat hij niet aan zijn
voorgaande levens kan toetsen,
noch in zijn volgende levens kan herstellen.
Er bestaat geen mogelijkheid
om na te gaan welke beslissing beter is,
want er is geen vergelijking.
Wij maken alles zomaar voor het eerst en onvoorbereid mee,
net als een acteur die voor de vuist een stuk speelt.

Non si può mai sapere che cosa si deve volere
perchè si vive una vita soltanto e non si può
né confrontarla con le proprie vite precedenti,
né correggerla nelle vite future.
Non esiste alcun modo di stabilire quale decisione sia la migliore,
perchè non esiste alcun termine di paragone.
L'uomo vive ogni cosa subito per la prima volta,
senza preparazione, come un attore che entra
in scena senza aver mai provato.

De ondraaglijke lichtheid van het bestaan,
L'insostenibile leggerezza dell'essere, Milan Kundera

In loving memory of my grandmother Adriana,
for,
with,
because of

Contents

Introduction	11
Chapter 1	
Visual attention and the temporal dynamics of feature integration	17
Chapter 2	
Moderate alcohol consumption impairs feature binding in visual perception but not across perception and action	51
Chapter 3	
Caffeine, but not nicotine enhances visual feature binding	57
Chapter 4	
Priming and binding in and across perception and action: A correlational analysis of the internal structure of event files	65
Chapter 5	83
What do we learn from binding features? Evidence for multilevel feature integration	
Chapter 6	107
Conclusions	
References	111
Summary in Dutch (Samenvatting)	121
Acknowledgments	125
Curriculum Vitae	127

Introduction

Feature binding

One of the basic characteristics of the primate cortex is that representations of the external world are distributed. For example, a visually perceived external object like a red ball will not be represented by a single code, but by a multitude of feature-related codes in different representational maps, such as a color code in a color map, a shape code in a shape map, a location code in a location map (or even many location maps, each representing a different reference frame) and so forth (for overviews, see Cowey, 1985; DeYoe & Van Essen, 1988). If people would be confronted with only one object at any given moment, this would not lead to any problem—the object features only need to activate their corresponding codes and the activated ensemble would then correctly represent the feature conjunction that characterizes the object. In everyday life, however, our visual environment is relatively complex and we often see, and seem to be able to perceive, more than one object at a time. This introduces the so-called binding problem, the question of how our brain is able to integrate correctly the feature codes that belong to the same object.




Figure 1. Feature integration (see text for further explanation).

A possible solution to the binding problem requires the distinction between two representational modes: the activation of feature codes and their integration. Let us assume two external objects are perceived, one being coded by the features F1 and F2 (e.g., ORANGE and ROUND) and the other by features F3 and F4 (e.g., GREEN and RECTANGULAR). As soon as the features activate their corresponding codes, the cognitive system “knows” that it is confronted with something red, something round, something green, and something rectangular, but it has no information about which color belongs to which shape. It would be unable to determine whether one of the external objects is, say, an orange (implying the combination of ORANGE + ROUND) or an apple (implying GREEN + ROUND). If the system had means to bind together and integrate the features belonging to the same objects, however, no confusion can arise and the objects would be easily identified (see Figure 1).

The currently most plausible candidate for such a binding mechanism is the temporal synchronization of all those cell populations that represent the

different features of a given object (Abeles, 1991; Singer, 1994; Von der Malsburg, 1981). In fact, there is ample evidence from single-cell studies on cats and monkeys (for overviews, see Engel, Roelfsema, Fries, Brecht, & Singer, 1997; MacKay, 1997) and EEG and MEG studies on humans (for an overview, see Tallon-Baudry & Bertrand, 1999) supporting the idea that coherence between different parts of a cognitive representation is achieved by (or at least associated with) synchronizing the firing rates of the underlying neuronal populations.

From a psychological point of view, the question is whether these hypothetical neuronal underpinnings have behavioral implications and, if so, whether and how they can be demonstrated and investigated. In the domain of visual perception, Kahneman, Treisman and Gibbs (1992) demonstrated that task-irrelevant stimuli of a complex prime display were particularly effective if they matched an upcoming target stimulus with respect to both identity and location, hence there was a specific benefit for feature conjunctions. These authors proposed that the codes of the features belonging to the same object are integrated into, what they call, an object file—a temporary cognitive structure containing all the perceptual information about a given object and perhaps even more (e.g., semantic information).

Recent studies by Hommel (2004) have extended these findings, showing that the spontaneous binding of the visual features can even be demonstrated in very simple tasks, where the target stimulus (the probe) requiring a binary decision is preceded by a single, irrelevant prime. Moreover, it turned out that performance with complete repetition of a feature conjunction (Prime A → Probe, see Figure 2) was about as good as performance with a complete alternation of features (Prime B → Probe), as compared to partial repetitions (Prime C → Probe). This suggests that repetition of a feature of the prime in the probe leads to the automatic activation of the “fellow feature” of the prime.




Figure 2. Hommel (2004) integration study (see text for further explanation).

Feature binding in and across perception and action

Further investigations have shown that binding is not restricted to visual features, and even not to perceptual tasks. Hommel (1998) obtained evidence that features of accidentally paired stimuli and responses are spontaneously bound. In each trial he precued the first of two responses (R1), so that this response could be already prepared. Then subjects waited for a trigger

stimulus (S1), which could be red or green, an X or O, and appear at a top or bottom location, and then carried out R1. As R1 was already known, color, shape, and location of S1 was completely irrelevant. A second later, a binary-choice task followed, which required a speeded response R2 to the shape, say, of a second stimulus S2. Choice performance was good if the relationship between the features of S2 and R2 either completely matched that between S1 and R1 (e.g., RED-LEFT and RED-LEFT) or completely mismatched (e.g., RED-LEFT and GREEN-RIGHT), as compared to conditions with partial matches (e.g., RED-LEFT and RED-RIGHT)—again a binding cost, but this time between stimulus and response.

Another demonstration of interactions between stimulus and response domains stems from Stoet and Hommel (2002). They presented subjects with an object characterized by a particular combination of shape, color, and location, and asked them to remember the object for later recall. In the retention interval they presented a speeded left-right keypressing task. Responses were slowed if they shared the location feature with the memorized object (e.g., left-side object → left-hand keypress). This suggests that memorizing the object led to a binding of the location feature with the other features of the object, so that this feature was not easily available for planning a spatially corresponding action. In fact, this logic seems to work both ways. Muesseler and Hommel (1997) could show that planning a spatially defined action and holding this plan in working memory impairs the perception and even the detection of a feature-overlapping visual event, such as a masked left- or right-pointing arrow. Stoet and Hommel (1999) further demonstrated that planning a left-right action specifically impairs the concurrent planning of a spatially corresponding (i.e., feature-overlapping) action, suggesting that binding a spatial feature to one plan makes it less available for the construction of another plan involving the same feature.

To summarize, there is ample experimental evidence for specific, predictable effects of feature binding. This evidence is not restricted to the domain of visual perception, for which the binding problem has been first formulated (see Treisman, 1996 for an overview), but spans perception, action planning, and stimulus-response relationships.

Is feature binding necessary?

The discussion of possible roles about the binding processes and the neural codes that might mediate them is still going on. Some authors have emphasized the need for integration processes in distributed representational systems like the human brain and argued that the synchronization of cell populations might play a major role in binding features belonging to the same object (e.g., Singer, 1994; Treisman, 1996), while others have questioned the very necessity of feature binding (e.g., Cisek & Turgeon, 1999; van der Heijden, 1995) and/or the involvement of neural synchronization in it (e.g., Jellema & Perrett, 2002; van der Velde & de Kamps, 2002). In this thesis a pragmatic, empirical stance is taken. As indicated above, there is ample suggestive evidence for feature binding. Consequently in this thesis it is investigated when, under which circumstances, evidence for binding can be found.

Thesis question

The available evidence suggests that binding phenomena can be demonstrated in perception, in action planning, and across perception and action. The questions addressed in this thesis are when, under which circumstances, can these phenomena be observed.

With regard to the latter question, there are two viable possibilities: First, these binding phenomena may all follow the same processing logic but, nevertheless, represent distinct phenomena produced by behaviorally distinguishable and neuroanatomically separable mechanisms. Second, it may be that all feature binding phenomena are realized through the same control mechanism—whether the features are perceptual or related to an action plan (i.e., are coded in visual areas or in the premotor cortex).

The evidence from behavioral experiments (Hommel, 1998) and physiological studies (e.g., Roelfsema, Engel, Koenig & Singer, 1997), that binding seems to help coordinating cognitive representations across domains as different (and cortically distant) as vision and manual action, suggests that the binding mechanism itself is not domain specific. That is, there may be one single system controlling or mediating all kinds of feature binding in the cognitive system. If so, and this is the guiding idea for this thesis, the different phenomena indicative of feature binding should show common characteristics.

Outline of thesis

This thesis contains five chapters reporting empirical work on feature integration.

Chapter 1 investigates the temporal dynamics of feature integration. In this chapter two experiments study the emergence of bindings between stimulus features (object files) and between stimulus and response features (event files) over time. The results indicate that bindings emerge quickly and remain intact for at least four seconds and that integration reflects the current attentional set, that is, which features are considered depends on their task-relevance. Features are not integrated into a single, global superstructure, but enter independent local bindings presumably subserving different functions.

Chapter 2 reports the effect of alcohol on feature integration. In an experiment it is investigated whether suppressing cholinergic activity through moderate alcohol consumption in healthy humans affects behavioral measures of feature binding in visual perception and across perception and action. The experiment reveals a dissociation between local feature binding in visual perception and cross-domain binding between visual features and manual responses: Intake of alcohol impairs only binding of visual features bindings and not across perception and action.

Chapter 3 presents the effect of caffeine and nicotine on feature integration. In this study the experiment reported shows a specific link between the visual system and the muscarinic cholinergic system. It appears that the binding of shape and color, and of shape and location of visual objects in healthy humans is increased by stimulating the muscarinic

cholinergic system (caffeine consumption) but not by stimulating the nicotinic cholinergic system (nicotine consumption). Feature binding across perception and action is unaffected by either manipulation, suggesting again a dissociation between purely visual and visuomotor integration.

Chapter 4 explores the commonalities between binding effects across different domains. Individual performance was compared across three different tasks that tap into the binding of stimulus features in perception, the binding of action features in action planning, and the emergence of stimulus-response bindings (“event files”). Correlations between the size of binding effects were found within visual perception (e.g., the strength of shape-location binding correlated positively with the strength of shape-color binding) but not between perception and action planning, suggesting different, domain-specific binding mechanisms. To some degree, binding strength was predicted by priming effects of the respective features, especially if these features varied on a dimension that matched the current attentional set.

Chapter 5 investigates the relationship between the binding of visual features (as measured by their after-effects on subsequent binding) and the learning of feature-conjunction probabilities. Both binding and learning effects were obtained but they did not interact. Our findings suggest that the creation of a neurocognitive representation of feature conjunctions is a multi-component process involving several time scales and levels of integration. We propose that the interaction between top-down attentional processes and automatic binding processes is dynamic and adaptive to task constraints.

The five empirical chapters have either been published, are under revision or are submitted in international psychological journals. They have been inserted in this thesis in their original, submitted or published form. To acknowledge the important contributions of several co-authors to each of these articles, a list of references is here presented.

Chapter 1: Hommel, B., & Colzato, L. S. (2004). Visual attention and the temporal dynamics of feature integration. *Visual Cognition*, 11, 483-521.

Chapter 2: Colzato, L. S., Erasmus, V., & Hommel, B. (2004). Moderate alcohol consumption impairs feature binding in visual perception but not across perception and action. *Neuroscience Letters*, 360, 103-105.

Chapter 3: Colzato, L. S., Fagioli, S., Erasmus, V., & Hommel B. (2005). Caffeine, but not nicotine enhances visual feature binding. *European Journal of Neuroscience*, 21, 591-595.

Chapter 4: Colzato, L. S., Warrens, M. J., & Hommel B. (2004). Priming and binding in and across perception and action: A correlational analysis of the internal structure of event files. Submitted to *Quarterly Journal of Experimental Psychology. Part A*.

Chapter 5: Colzato, L. S., Raffone, A., & Hommel B. (2004). What do we learn from binding features? Evidence for multilevel feature integration. Submitted to *Journal of Experimental Psychology: Human Perception and Performance*.

Chapter 1

Visual Attention and the Temporal Dynamics of Feature Integration

Abstract

Two experiments studied the emergence of bindings between stimulus features (object files) and between stimulus and response features (event files) over time. Choice responses (R2) were signalled by the shape of a stimulus (S2) that followed another stimulus (S1) of the same or different shape, location, and colour. S1 did not require a response (Experiment 1) or trigger a precued simple response (R1) that was or was not repeated by R2 (Experiment 2). Results demonstrate that the mere co-occurrence of stimulus features, and of stimuli and responses, is sufficient to bind their codes. Bindings emerge quickly and remain intact for at least four seconds. Which features are considered depends on their task-relevance; hence, integration reflects the current attentional set. There was no consistent trend towards higher order interactions as a function of time or of the amount of attention devoted to S1, suggesting that features are not integrated into a single, global superstructure, but enter independent local bindings presumably subserving different functions.

Introduction

When an object appears before our eyes, its perceivable features are registered and coded in various areas in our brain—and yet, what we commonly perceive is not a mosaic bundle of attributes but a single, homogeneous object. This suggests the existence of some kind of feature-binding mechanism that keeps track of which feature goes with which, in such a way that features belonging to the same object can be integrated and cross-referenced in the process of internally reconstructing an observed external object (e.g., Allport, Tipper, & Chmiel, 1985; Singer, 1994; Treisman, 1996). In the visual domain, there is converging evidence for spontaneous feature integration from several lines of research.

First, Allport et al. (1985) had participants work through lists or sequences of superimposed pairs of letters or pictures. One member of a pair was the to-be-named target printed in a particular target colour (e.g., red), and the other member was a distractor printed in another colour (e.g., green). It turned out that reaction time (RT) substantially increased when, in a trial, the current target matched the preceding distractor as compared to trials with no match. One interpretation of this so-called negative priming effect (for overviews see Fox, 1995; May, Kane, & Hasher, 1995; Tipper, 2001) that was offered by Allport et al. is in terms of feature integration (for a revival of this theme see Neill, 1997; Park & Kanwisher, 1994). Assume that, in a given trial, both target and distractor features are integrated (i.e., linked across feature domains) separately, thus forming two different object representations. If then in the following trial the features appear in different combinations (e.g., when a

formerly green letter now appears in red), integration is more difficult than for exact repetitions of feature combinations because it requires additional time to undo the already formed, and now misleading, cross-domain links. Although it seems clear by now that negative priming also involves processes unrelated to feature integration (such as inhibition of S-R links: Houghton & Tipper, 1994), there are various demonstrations of the unwanted retrieval of spontaneously integrated stimulus episodes (Kane, May, Hasher, Rahhal, & Stoltzfus, 1997; Lowe, 1985; Neill, 1997; Waszak, Hommel, & Allport, 2003).

Second, Kahneman, Treisman, and Gibbs (1992) presented participants with two displays in a sequence, a brief multiletter preview or prime display requiring no response (S1) and a single-letter probe display requiring verbal identification (S2). If the probe letter had already been presented somewhere in the preview display, probe identification was facilitated (a repetition benefit), but only slightly so and not in each experiment. However, if the previewed letter matched the probe both in identity and (absolute or relative) location, pronounced and stable identification benefits were observed. According to Kahneman et al., attending to a visual object establishes what they call an “object file”, an integrated episodic trace containing information about the relationship between object features and their location, possibly enriched by object-related knowledge from long-term memory. If an object file is constructed for a previewed object, and if this object re-appears at the same location, object perception does not require constructing a new file, but an update of the old one will do. That is, performance should not so much depend on the repetition of one or more stimulus features per se, but rather on whether the particular feature conjunction (e.g., of shape and location) is repeated or not. Only if the same conjunction reappears, the old object representation is used another time, thus speeding up the identification process. If, however, feature repetition is only partial or absent altogether, a new representation needs to be constructed, just as without a preview.

Third, Hommel (1998) had participants perform a binary-choice task in response to the shape of a stimulus (S2) that was preceded by another stimulus (S1) the features of which were irrelevant. In contrast to previous studies, the experimental design allowed for an independent manipulation of the shape, location, and colour of S1 and S2, so that performance could be compared across several degrees and combinations of feature repetitions, ranging from a complete match of S1 and S2 (i.e., repetition of all possible features and feature conjunctions) to mismatch (i.e., no repetition of any feature or feature conjunction). Interestingly, the effects of feature repetitions were not independent of each other: Repeating shape produced better performance than alternation if colour was also repeated, but worse performance than alternation if colour alternated; and the same relationship was observed between shape and location. Given that complete matches yielded about the same performance as mismatches, these results do not so much point to a benefit of repeating a particular feature conjunction—as one would expect from Kahneman et al.'s (1992) approach—but rather suggest a cost of partial repetitions—as implied by Allport et al.'s (1985) account. Assume that being confronted with S1, say, a horizontal line at the bottom of a display, results in a spontaneous binding of the codes representing its shape and location, as depicted in Figure 1 (panel A). If S1 and S2 share either both

shape and location (complete match) or none (mismatch), integrating S2 features should not represent any particular problem. However, if only one but not the other feature overlaps (partial match), reactivating the code of the matching feature may spread activation to the code it has just been integrated with, thus impairing its integration with the actual feature¹.




Figure 1. A binding account of partial-match costs. (A) The basic assumption is that feature codes representing Stimulus 1 (S1) are temporarily integrated into a coherent event representation. If relations between features of Stimulus 2 (S2) completely match or mismatch with these bindings, performance is unaffected. Partial matches, however, prime incorrect and misleading feature codes, which induces conflict and, thus, slows down performance. (B) The same logic applies to stimulus and response features.

Altogether, the available evidence strongly suggests that seeing an object results in the more or less spontaneous integration or binding of its features. Once bound together, these features (or their codes) apparently can no longer be separately addressed, so that perceiving a new combination of the same features requires another time-consuming rebinding process and/or the resolution of the conflict induced by the previous binding. Interestingly, these kinds of binding effects are not restricted to stimulus features.

¹ As pointed out by an anonymous reviewer, the logic underlying this account bears an interesting similarity to Kingstone's (1992) crosstalk interpretation of the combined effects of multiple cues on stimulus processing. Kingstone cued his subjects with regard to two features of an upcoming stimulus, such as spatial location and shape, or shape and colour. Unsurprisingly, valid cues sped up responses considerably but the cuing effects were not independent. In particular, performance was impaired if the stimulus matched one expectation but not the other, such as when an unexpected target form appeared in an expected colour or an expected form appeared in an unexpected colour. Kingstone suggests that people had created a "combined expectancy" that, if one part of the expectation is matched by the upcoming stimulus, primes the other, related part—which again facilitates processing stimuli that fully match the expectations but hamper the processing of partial matches. One may speculate that the cognitive structure people create when building a "combined expectancy" is the same as the "object file" that is left by integrating the features of a stimulus. In other words, anticipating an event may have the same effect as just having seen it before.

In the Hommel (1998) study, participants were precued, in each trial, whether the first response (R1) should be a left-hand or a right-hand key press. R1 was then triggered by the next upcoming stimulus (S1) without depending on any particular feature of it. One second after S1, S2 would appear, and participants were instructed to respond to its shape (or, in another experiment, to its colour) by pressing the left or right key (R2). Hence, participants performed sequences of a simple RT task followed by a binary-choice RT task, and what varied was the identity of R1 and R2 and the shape, colour, and location of S1 and S2. The results showed that the repetition or alternation of stimulus features did not only interact with other stimulus-feature effects, they also interacted with response repetition. For example, response repetitions were faster and more accurate if stimulus shape was also repeated than if shape alternated, whereas response alternations were faster and more accurate if shape alternated than if shape was repeated.

These findings imply that the binding logic introduced above also applies to combinations of stimulus and response features, along the lines sketched in Panel B of Figure 1: The mere co-occurrence of a stimulus feature and a response (feature) may lead to the creation of a binding between their codes, so that reactivating one will tend to prime the other. Indeed, there is converging evidence in support of this idea. For instance, Hommel (2003) found that, in a free-choice task, repeating the shape, colour, or location of the stimulus increases the likelihood that subjects repeat the previous response. Likewise, Dutzi and Hommel (2003) observed that producing a particular stimulus by pressing a particular key increases the likelihood that this key is pressed again if the same stimulus appears during the next trial. These findings suggest that feature integration may not be restricted to object perception but cross borders between perception and action to create what Hommel (1998) called “event files”.

Purpose of the study

The available evidence points to the existence of object or event files, but the mechanisms underlying their creation, maintenance, and possible decay remain to be explored. The present study was motivated by three open questions that all in one or the other way refer to the temporal dynamics and the attentional preconditions of feature integration.

How complete is feature integration?

According to the original concept of an object file (Kahneman et al., 1992) one would assume that in the case of processing an object integration is well captured by the following scenario: Spatial attention is directed or attracted to a particular location in space and/or to an object occupying this location, and then all features of that object are integrated into a coherent object representation (e.g., Luck & Beach, 1998; Treisman, 1988, 1998). But are they? There is evidence that repeating conjunctions of letters and locations affect behaviour even if the letter case changes (Gordon & Irwin, 1996; Henderson, 1994). This suggests that feature integration may not be complete (e.g., may not necessarily include shape or shape details) but may be

mediated by the task context. Likewise, Hommel (1998) found interactions between shape and location repetition only if shape was task relevant (by virtue of signalling R2) but not if colour was task relevant; and the opposite tendency was observed for colour-location interactions. Effects of task relevance were also obtained by Hommel (2003), who found evidence for location response bindings if the responses were defined in terms of location (left vs. right key) but not if they were defined in terms of number (single vs. double press).

To account for the impact of task relevance and context one may assume that feature codes enter more enduring representations only if, or to the degree that, they pass a kind of relevance or pertinence filter (e.g., Bundesen, 1990; Norman, 1968). That is, spatial attention may (or may not) preselect the features of an attended location or object, these features may then be weighted according to their relevance to the task at hand (in addition to possible bottom-up saliency factors), and the feature codes surviving these procedures will enter an object file. However, even this scenario does not appear to fully account for the available findings. For instance, the Hommel (1998) study revealed several indications of bindings between shape and location and between shape and colour, while colour and location were independent. Or, with respect to the integration of stimulus and response features, colour was integrated with the response only if colour but not if shape was task relevant, whereas the signs of shape-colour integration were independent of whether shape or colour was relevant. Thus, not all features that have an effect (suggesting that they passed whatever filter had been applied) interact with each other, at least not in the form of a higher order interaction that would point to a comprehensive object or event file.

However, the reported studies used a very limited range of temporal intervals (or stimulus-onset asynchronies; SOAs) between the first, inducing display (S1) and the second, probe display (S2); e.g., all SOAs in the studies of Hommel (1998, 2003) employed SOAs of 1 s. Yet, the integration processes that presumably underlie the observed interactions between repetition effects might be rather time consuming, which implies that the construction of object or event representations is a temporally extended operation. If so, the findings reported so far may be just static snapshots of a dynamic binding process and, thus, represent arbitrarily chosen sessions of this process only. To get a better idea of the temporal characteristics and possible limitations of feature integration we therefore manipulated SOA across a wide range of 200-4100 ms. One possibility would be that features are rapidly integrated into rather short-lived, transient bindings, so that signs for complete integration may be found with short, but not with long SOAs. Alternatively, integration may take time, which would imply that complete integration is found with long, but not with short SOAs.

Another variable that we thought might impact the completeness or depth of integration is the amount of endogenous attention allocated to S1 processing. In Hommel's (1998) original design S1 is entirely irrelevant, except that its presence needs to be noticed to trigger R1. On the one hand, this means that S1 cannot go unnoticed and, thus, must attract some degree of attention—the more so as it is the only event that occurs within a relatively long interval of 2.5 s. On the other hand, however, there is no need to fully process or even integrate the features of that stimulus, so that one might

expect that processing is rather superficial. That feature repetition effects, and interactions between them, were nevertheless obtained indicates that the underlying binding processes do not strongly depend on the need or intention to integrate the particular features (although spatial attention may well be necessary in any case). However, integration may be deeper and more complete if it is really needed. Hence, it may well be that the lack of complete integration is merely a result of not requiring subjects to endogenously attend to S1 and perform operations that require the integration of its features. We tested this hypothesis by comparing an endogenously “unattended” condition designed after Hommel (1998) with an “attended” condition, where we required subjects to report S1 at leisure after R2 was completed. Apart from drawing more attention to S1, this manipulation is likely to require the consolidation of S1 features in short-term memory (Jolicœur, Tombu, Oriet & Stevanovski, 2002), which has been claimed to be associated with feature integration (Luck & Vogel, 1997; Raffone & Wolters, 2001).

To summarize, we were interested to see whether higher order interactions of feature repetition effects (i.e., effects involving more than two features and/or the response) could be obtained by allowing more time for integration to proceed (i.e., at longer SOAs) and/or by increasing the attentional resources devoted to processing S1 (i.e., in the “attended” condition).

Are feature bindings addressed by location?

A second question that motivated our study concerns the way object or event files are addressed. According to the original suggestion of Kahneman et al. (1992), object files are addressed by location. That is, encountering an object leads to the retrieval of that object file that includes spatial codes that match the location of the present object to at least some degree. However, developmental research provides evidence that infants and children often use (changes in) nonspatial features to individuate objects and spatiotemporally extended events, suggesting that object representations can be addressed in ways that are not mediated by location codes (e.g., Leslie & Kaldy, 2001; Leslie, Xu, Tremoulet, & Scholl, 1998). Moreover, the addressing-by-location assumption implies that information about object location must be a basic ingredient of object files, which does not fit with Hommel's (1998, 2003) observations of feature interactions not involving location repetition.

One possible conclusion from these findings is that the outcome of binding processes is not so much a single representational entity collecting all information about a particular event, as the object-file metaphor would suggest, but rather a distributed network of local bindings presumably subserving diverse functions (Hommel, 1998, 2003). Alternatively, to take up the argument developed above, the available findings may represent only a snapshot of a more dynamic integration process. For instance, integration may begin locally, connecting pairs of feature codes, and only gradually build up a global event representation. If so, we might find changes in the order of interactions obtained across SOA, lower order interactions dominating at short SOAs and increasingly higher order interactions (including stimulus location) at longer SOAs. And, again, the amount of attentional resources devoted to processing the integrated event (i.e., S1) might play a mediating role.

How are feature priming and feature integration related?

A third question underlying our study has to do with the relationship between feature priming and feature integration. Apart from evidence of integration Kahneman et al. (1992) were also interested in what they called nonspecific effects, that is, effects due to the repetition or alternation of a single stimulus feature, independent of any interaction with another feature. Little evidence for such effects was found by Kahneman et al. or Hommel (1998). However, substantial priming effects were obtained in the studies of Gordon and Irwin (1996), Henderson (1994), and Henderson and Anes (1994), where repeating nonspatial stimulus features significantly improved performance even if the stimulus changed location in between two appearances. Gordon and Irwin, for instance, had subjects make word-nonword judgements to target stimuli that randomly appeared in one of two vertically arranged boxes. Each stimulus was preceded by two prime words, and in some cases one of these primes matched the target stimulus (e.g., “doctor” + “bread” → “doctor”). Matching primes sped up reaction times substantially, in particular if prime and target appeared in the same box (i.e., shared location). This supports the assumption that processing the prime was accompanied by some sort of integration of its identity and its location, and that the product of this integration was maintained at least until target presentation. However, priming effects were smaller but still reliable even if the matching prime had appeared in the box opposite to the target, suggesting that retrieving prime information did not require the repetition of location. Hence, nonspecific priming does exist, at least under some circumstances. Kahneman et al. attributed the absence of nonspecific effects in their study to the small number of stimulus alternatives they had used: The same items were presented over and over again, so that their codes may have been primed to ceiling. However, given that Henderson and colleagues obtained nonspecific priming with even smaller stimulus sets, this is a rather unlikely explanation.

Again, the time interval between the first and the second presentation of the stimuli may be an important factor. Indeed, the studies where priming effects were weak or absent all used rather long SOAs (Hommel, 1998: 1000 ms; Kahneman et al., 1992, Exps. 1 and 2: 400-950 ms), whereas studies where reliable effects were observed employed short SOAs (Gordon & Irwin, 1996: SOAs of 1500 ms but interstimulus intervals of only 250 ms; Henderson, 1994: the latency of a saccade). It is therefore possible that the priming of codes of individual features is a rather short-lived phenomenon that is observable with very brief SOAs only (cf. Hommel, 1994). If so, we would expect priming effects with short, but not with longer, SOAs.

Experiments 1 and 2

We conducted two very similar experiments to address our three guiding questions. Experiment 2 was closely modelled after Hommel (1998): Subjects were cued to prepare a left- or right-hand key press (R1), which they carried out as soon as S1—the priming stimulus—was presented (see Figure 2). Even though it was only the presence but not the identity of S1 that counted for performing R1, we varied its shape, colour, and (vertical) location. After a

variable SOA (1100, 2100, or 4100 ms) S2 appeared to signal R2. The two alternative shapes of S2 were mapped onto the two R2 alternatives, while colour and location of S2 were entirely irrelevant to the task—which was pointed out to the subjects. In one half of the sessions (the attended sessions), subjects were also to report one randomly chosen (i.e., unpredictable) feature of S1 after R2 was completed, a manipulation that we considered to draw (more) attention to S1 and to motivate if not require the integration of its features.




Figure 2. Overview of the displays and the timing of events in Experiments 1 and 2.

We were particularly interested in three types of effects and their dependencies on our manipulations of attention and SOA. First, we wanted to see whether priming effects, i.e., effects of the repetition of an individual feature, would occur and, if so, whether they might be more pronounced at short than at long SOAs. Note, however, that even the shortest SOA of Experiment 2 was longer than our above considerations suggest is optimal for finding priming effects, which was the main reason for us to conduct Experiment 1 (see below). Second, we were interested to see whether the interactions between effects of stimulus-feature repetitions (e.g., Shape x Colour) obtained by Hommel (1998) can be replicated and, even more important, whether they would be affected by the amount of attention devoted to S1 and change across SOA. Of particular theoretical relevance were interactions between more than two stimulus features (which would point to complete integration) and/or of interactions not involving stimulus location (which would speak to the addressing-by-location issue), and possible changes of these interactions as a function of attention (which might create more complete bindings) and SOA (which might allow for the creation of increasingly global bindings). Third, we sought to replicate the interactions between stimulus features and response obtained by Hommel (1998). And, again, we were interested in whether these interactions remain stable across attentional manipulations and SOA or, rather, whether they would enter higher order interactions as attentional investment and SOA increases.

As pointed out, Experiment 2 with its long SOAs was unlikely to provide an optimal platform for priming effects, which can be expected to occur in the range of 0–500 ms. However, using that short SOAs would create a dual-task

situation in which the S2-R2 component of the task would temporally overlap with the S1-R1 component. This would be likely to create unpredictable and complicating side effects, such as dual-task costs or S1-R2 and S2-R1 integration (cf. Dutzi & Hommel, 2003), which we wanted to avoid. To do so we restricted the whole first part of each trial to the presentation of S1 (see Figure 2), which now, at least in unattended sessions, had no function at all. That is, people were presented with two stimuli in a row, separated by a variable SOA (200–4100 ms), and responded to the second stimulus (S2) by pressing a left or right key (R2—which in the absence of R1 was the only response!). As this modification eliminated R1, Experiment 1 did not speak to the integration of stimulus and response features. However, including a short SOA increased our chances to detect short-lived phenomena in the priming and integration of stimulus features.

Method

Participants

Seventeen students of the Leiden University took part for pay in Experiment 1 and 16 participated in Experiment 2. All reported having normal or corrected-to-normal vision, and were not familiar with the purpose of the experiment.

Apparatus and stimuli

The experiments were controlled by a Targa Pentium III computer, attached to a Targa TM 1769-A 17-inch monitor. Participants faced three grey square outlines, vertically arranged, as illustrated in Figure 2. From viewing distance of about 60 cm, each of these frames measured $2.6^\circ \times 3.1^\circ$. A thin vertical line ($0.1^\circ \times 0.6^\circ$) and a some what thicker horizontal line ($0.3^\circ \times 0.1^\circ$) served as S1 and S2 alternatives, which were presented in red or green in the top or bottom frame. Response cues (in Experiment 2 only) were presented in the middle frame (see Figure 2), with rows of three left-or right-pointing arrows indicating a left and right key press, respectively. Responses to S1 (in Experiment 2 only) and to S2 were made by pressing the left or right shift key of the computer keyboard with the corresponding index finger.

Procedure and design

Experiment 1. This consisted of six 1 hour sessions: Three S1 unattended and three S1 attended. In (endogeneously) unattended sessions participants made a single response on each trial, a binary-choice reaction to the second of two successive stimuli. Half of the participants responded to the vertical and the horizontal line by pressing the left and right key, respectively, while the other half received the opposite mapping. In attended sessions participants were, after each binary-choice reaction to a S2, probed for their memory of a feature of S1. They were presented with one of three questions, asking for the shape, colour, or location of S1 (e.g., “What was the colour of Stimulus 1?” in Dutch). Two words indicating the two response alternatives (“horizontal-vertical”, “red-green”, or “top-bottom”, in Dutch) were presented below the question, with their relative position indicating the mapping of alternatives onto

the left and right shift key. The six combinations of the three stimulus dimensions and two alternative key mappings were presented in pseudorandom sequence but equally often within one session. Half of the participants began with the unattended sessions; the other half began with the attended sessions.

The sequence of events is shown in the upper row of Figure 2. In unattended sessions, the intertrial interval of 2000 ms was followed by a 100 ms appearance of S1. The duration of the next, blank interval depended on the SOA condition: 100, 1000, 2000, or 4000 ms. Then S2 appeared and stayed until the response was given or 2000 ms had passed. If the response was incorrect auditory feedback was presented. In attended sessions, this sequence of events was followed by the memory probe question, which stayed until the response was given or 4000 ms had passed.

Each session comprised 256 trials, composed by a factorial combination of the two shapes (vertical vs. horizontal line), colours (red vs. green), and locations (top vs. bottom) of S2, the repetition vs. alternation of shape, colour, and location, and the four SOAs ($2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 4 = 256$). Thus, taken together, the three attended and three unattended sessions of Experiment 1 amounted to 1536 trials. Participants were allowed to take a short break during each session.

Experiment 2. This consisted of six 90 minute sessions: Three unattended and three attended sessions. The procedure was as in Experiment 1, with the following exceptions. In unattended sessions participants carried out two responses per trial. R1 was a simple reaction with the left or right key, as indicated by the response cue. It had to be carried out as soon as S1 appeared, independent of its shape, colour, or location. Participants were informed that there would be no systematic relationship between S1 and R1, or between S1 and S2, and they were encouraged to respond to the onset of S1 only, disregarding the stimulus' attributes. As in Experiment 1, R2 was a binary-choice reaction to the shape of S2 and attended sessions required the identification of a randomly selected feature of S1.

The sequence of events in each trial is shown in the lower row of Figure 2. Next to the intertrial of 2000 ms a response cue signalled R1 for 1500 ms, followed by a blank interval of 1000 ms. Then S1 appeared for 100 ms, followed by a further blank interval the duration of which depended on the SOA condition: 1000, 2000, or 4000 ms. If R1 was incorrect or not given within 600 ms the trial started again. After the respective SOA, S2 appeared and stayed until R2 was given or 2000 ms had passed.

Each session comprised 384 trials, composed by a factorial combination of the two shapes (vertical vs. horizontal line), colours (red vs. green), and locations (top vs. bottom) of S2, the repetition vs. alternation of shape, colour, location, and the response, and the three SOAs ($2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 3 = 384$). Thus, taken together, the six sessions of Experiment 2 amounted to 2304 trials.

Results and discussion

Analytical procedures. Trials with missing or anticipatory responses (1.4% in Experiment 1 and 1.8% in Experiment 2) were excluded from the analysis. We also excluded trials in which the memory probe response was incorrect. From the remaining data, mean RTs and proportions of errors (PEs) for R2 (i.e., the response to S2) were further analysed, as well as PEs for responses in the memory probe task (available from attended sessions only).

In Experiment 1, means were computed as a function of Attention (S1 unattended vs. attended), the four SOAs, and the three possible relationships between the two stimuli in each trial, that is, repetition vs. alternation of stimulus shape, colour, or location (see Table 1 for means). ANOVAs were performed by using a four-way design (in case of the memory data) and a five-way design for repeated measures. The significance criterion for all analyses was set to $p < .05$.

In Experiment 2, means were computed as a function of Attention (S1 unattended vs. attended), the three SOAs, and the four possible relationships between the two responses (R1 and R2) and the two stimuli in each trial, that is, repetition vs. alternation of response, stimulus shape, colour, or location (see Table 2 for means). ANOVAs were performed by using a five-way design (in case of the memory data) and a six-way design for repeated measures.

We first analysed the memory-probe data. Experiment 1 revealed only two reliable effects, one indicating that errors increased with SOA (3.7%, 4.3%, 4.8%, 5.4%), $F(3, 48) = 3.18$, $p < .05$, and the other that fewer errors were made with repetitions of stimulus location than alternations (4.1 % vs. 5.0%), $F(1, 16) = 8.15$, $p < .05$. These effects were not replicated in Experiment 2 ($F_s < 1.1$) where, however, Shape was involved in two effects: In a main effect, $F(1, 15) = 12.73$, $p < .005$, and an interaction with SOA, $F(2, 30) = 10.74$, $p < .001$. When shape was repeated, the errors followed the same pattern as in Experiment 1, i.e., they increased with SOA (3.4%, 3.9%, and 5.1%). When shape alternated, however, this pattern was not observed (5.9%, 5.2%, and 5.2%). Taken together, the only replicable outcome seems to be the increase in errors as SOA increases. As in our design SOA is confounded with the effective retention interval of the memory task, this is an unsurprising observation. More important, however, the overall memory performance was very good, demonstrating that our attention manipulation worked.

TABLE 2

Experiment 2: Means of mean reaction times for responses to Stimulus 2 (RT_{R2} ; in Ms) and percentages of errors on $R2$ (PE_{R2}) and on the report of S1 in attended sessions (PE_M), as a function of attention to Stimulus 1, stimulus-onset asynchrony (in Ms), the match between Response 1 and Response 2, and the feature match between Stimulus 1 and Stimulus 2

1100				2100				4100							
Response															
Repeated				Alternated				Repeated				Alternated			
RT _{R2}	PE _{R2}	PE _M		RT _{R2}	PE _{R2}	PE _M		RT _{R2}	PE _{R2}	PE _M		RT _{R2}	PE _{R2}	PE _M	
Stimulus I unattended															
Neither	577	6.72		539	7.23			552	4.99			532	6.31		
L(location)	551	5.56		578	6.22			536	5.14			526	6.29		
S(shape)	543	2.91		587	3.03			529	4.03			528	6.71		
C(colour)	557	4.93		574	4.26			537	4.66			526	3.95		
SL	531	3.49		596	4.04			519	3.09			520	4.08		
SC	542	1.98		590	3.70			524	6.37			521	5.48		
LC	551	5.68		585	6.55			527	4.49			523	5.09		
SLC	531	3.31		605	4.58			515	1.83			516	5.41		
Stimulus I attended															
Neither	687	14.58	6.72	627	5.46	7.23		641	13.28	4.99		596	11.72	6.31	
L	647	10.93	5.56	672	3.90	6.22		620	10.15	5.14		583	10.15	6.29	
S	643	8.59	2.91	706	10.41	3.03		615	8.86	4.03		589	9.11	6.71	
C	661	14.06	4.93	672	5.46	4.26		623	11.98	4.66		590	9.63	3.95	
SL	622	3.12	3.49	712	19.01	4.04		603	9.11	3.09		578	7.55	4.08	
SC	642	10.93	1.98	714	9.89	3.70		605	8.85	6.37		683	8.85	5.48	
LC	643	8.33	5.68	676	7.03	6.55		616	8.59	4.49		578	10.15	5.09	
SLC	620	6.77	3.31	722	22.65	4.58		593	4.42	1.83		567	3.64	5.41	

Tables 3 and 4 provide an overview of the ANOVA outcomes for RTs and PEs obtained for R2 in Experiments 1 and 2, respectively. To facilitate access to the relatively complex data pattern we sort, present, and discuss the outcomes according to their theoretical implications, attempting to integrate the findings from Experiments 1 and 2 as far as possible.

TABLE 3
Results of analysis of variance on mean reaction time of correct responses (RT) and percentage of errors (PE) for Experiment 1

<i>Effect</i>	<i>df</i>	<i>RT_{R2}</i>		<i>PE_{R2}</i>	
		<i>MSE</i>	<i>F</i>	<i>MSE</i>	<i>F</i>
Attention (Att)	1,16	172,690.22	18.95**	137.21	1.88
Soa	3,48	9,125.63	19.38**	36.77	8.17**
Shape (Shp)	1,16	4,674.04	0.49	67.29	1.20
Colour (Col)	1,16	633.68	2.15	14.67	2.83
Location (Loc)	1,16	3,271.31	6.60*	36.24	2.35
Att × Soa	3,48	8,477.65	12.75**	21.72	6.74**
Att × Shp	1,16	2,439.05	1.77	43.96	1.10
Att × Col	1,16	992.95	1.69	9.28	0.21
Att × Loc	1,16	1,499.61	7.94*	13.32	0.03
Soa × Shp	3,48	1,589.84	9.38**	33.95	3.00*
Soa × Col	3,48	956.16	0.49	12.56	0.81
Soa × Loc	3,48	1,263.26	1.19	19.73	0.67
Shp × Col	1,16	740.39	3.99	31.73	5.97*
Shp × Loc	1,16	1,553.19	5.32*	29.28	0.03
Col × Loc	1,16	894.73	1.97	9.38	9.06**
Shp × Col × Loc	1,16	924.86	0.04	32.15	0.02
Att × Soa × Shp	3,48	1,077.24	4.85**	17.36	1.51
Att × Soa × Col	3,48	751.51	1.83	20.60	0.64
Att × Soa × Loc	3,48	551.69	2.52	17.94	0.25
Att × Shp × Col	1,16	920.68	0.24	16.76	0.34
Att × Shp × Loc	1,16	940.20	2.99	21.98	0.32
Att × Col × Loc	1,16	299.16	10.05**	27.26	2.18
Att × Shp × Col × Loc	1,16	180.18	0.96	15.20	0.88
Soa × Shp × Col	3,48	897.30	0.54	19.28	1.78
Soa × Shp × Loc	3,48	1,088.54	2.52	17.74	2.30
Soa × Col × Loc	3,48	536.03	1.66	21.02	1.03
Att × Soa × Shp × Col	3,48	982.75	0.25	18.53	0.74
Att × Soa × Shp × Loc	3,48	916.00	1.78	26.72	0.35
Att × Soa × Col × Loc	3,48	740.09	0.59	11.86	0.52
Soa × Shp × Col × Loc	3,48	778.49	0.93	23.75	1.85
Att × Soa × Shp × Col × Loc	3,48	1,066.25	0.21	18.72	0.57

Soa = stimulus-onset asynchrony. * $p < .05$, ** $p < .01$.

TABLE 4
Results of analysis of variance on mean reaction time of correct responses (RT) and percentage of errors (PE) for Experiment 2

<i>Effect</i>	<i>df</i>	<i>RT_{R2}</i>		<i>PE_{R2}</i>	
		<i>MSE</i>	<i>F</i>	<i>MSE</i>	<i>F</i>
Attention (Att)	1,15	87,301.63	28.65**	484.85	3.08
Soa	2,30	8,029.66	50.22**	28.34	10.51**
Shape (Shp)	1,15	2,870.28	7.02*	70.50	2.58
Colour (Col)	1,15	970.91	0.60	56.07	0.69
Location (Loc)	1,15	1,272.91	2.45	41.17	0.01
Att × Soa	2,30	3,889.28	12.39**	36.97	0.26
Att × Shp	1,15	2,528.90	9.62**	21.75	7.08*
Att × Col	1,15	1,295.73	0.42	24.66	0.57
Att × Loc	1,15	751.230	0.12	14.30	0.95
Soa × Shp	2,30	1,172.57	1.83	40.10	3.99*
Soa × Col	2,30	984.09	0.39	26.31	1.38
Soa × Loc	2,30	650.91	5.07*	38.14	1.20
Shp × Col	1,15	1,482.39	0.16	35.51	0.03
Shp × Loc	1,15	1,101.04	25.49**	45.20	8.40*
Col × Loc	1,15	873.06	0.02	34.93	1.46
Shp × Col × Loc	1,15	1,147.65	8.38*	18.35	0.01
Att × Soa × Shp	2,30	905.82	0.06	21.61	0.88
Att × Soa × Col	2,30	1,110.85	0.24	19.53	2.06
Att × Soa × Loc	2,30	690.31	4.72*	25.35	1.06
Att × Shp × Col	1,15	827.05	0.36	20.47	0.68
Att × Shp × Loc	1,15	956.80	0.75	46.14	0.07
Att × Col × Loc	1,15	527.00	0.08	14.04	0.07
Att × Shp × Col × Loc	1,15	1,444.94	0.49	16.25	0.97
Soa × Shp × Col	2,30	662.73	2.10	32.10	2.09
Soa × Shp × Loc	2,30	975.07	2.87	24.37	9.46**
Soa × Col × Loc	2,30	699.83	1.39	29.97	0.98
Att × Soa × Shp × Col	2,30	851.50	0.43	20.07	1.68
Att × Soa × Shp × Loc	2,30	1,024.38	0.11	32.49	0.38
Att × Soa × Col × Loc	2,30	694.14	1.81	19.71	0.12
Soa × Shp × Col × Loc	2,30	811.98	2.99	27.27	1.76
Att × Soa × Shp × Col × Loc	2,30	773.77	0.50	18.75	5.02*
Response (Rsp)	1,15	9,560.97	0.02	157.79	0.01
Att × Rsp	1,15	3,802.59	1.10	24.10	1.97
Soa × Rsp	2,30	1,167.12	5.56**	36.13	7.38**
Shp × Rsp	1,15	3,945.80	76.38**	90.20	84.80**
Col × Rsp	1,15	517.22	36.56**	16.50	13.20**
Loc × Rsp	1,15	1,334.64	50.99**	97.38	33.44**
Shp × Col × Rsp	1,15	740.11	7.15*	15.34	0.26
Shp × Loc × Rsp	1,15	1,047.71	0.11	22.22	11.03**
Col × Loc × Rsp	1,15	779.90	2.65	23.90	0.73
Shp × Col × Loc × Rsp	1,15	1,164.77	2.54	22.41	4.34

TABLE 4
(Continued)

Effect	df	RT_{R2}		PE_{R2}	
		MSE	F	MSE	F
Att × Shp × Rsp	1,15	1,675.06	20.23**	62.03	1.57
Att × Col × Rsp	1,15	1,183.53	0.04	51.04	0.62
Att × Loc × Rsp	1,15	1,128.56	6.59*	40.21	1.25
Att × Shp × Col × Rsp	1,15	584.54	8.88**	19.62	0.16
Att × Shp × Loc × Rsp	1,15	459.34	1.38	19.55	6.89*
Att × Col × Loc × Rsp	1,15	762.97	0.21	19.05	13.20**
Att × Shp × Col × Loc × Rsp	1,15	964.95	0.13	41.15	0.24
Att × Soa × Rsp	2,30	978.28	2.06	31.92	0.83
Soa × Shp × Rsp	2,30	1,347.20	24.61**	37.41	29.81**
Soa × Col × Rsp	2,30	756.06	1.77	20.07	1.77
Soa × Loc × Rsp	2,30	1,750.25	2.49	34.29	7.34**
Soa × Shp × Col × Rsp	2,30	715.33	1.77	21.71	1.91
Soa × Shp × Loc × Rsp	2,30	525.65	7.81**	26.66	1.71
Soa × Col × Loc × Rsp	2,30	637.76	0.91	16.30	1.16
Soa × Shp × Col × Loc × Rsp	2,30	1,008.73	1.51	29.15	1.77
Att × Soa × Shp × Rsp	2,30	739.84	1.30	20.75	3.84*
Att × Soa × Col × Rsp	2,30	970.00	0.12	19.10	0.64
Att × Soa × Loc × Rsp	2,30	1,021.53	1.93	32.86	0.15
Att × Soa × Shp × Col × Rsp	2,30	598.73	0.12	32.73	1.71
Att × Soa × Shp × Loc × Rsp	2,30	774.01	0.75	24.44	1.82
Att × Soa × Col × Loc × Rsp	2,30	482.65	0.79	24.79	0.03
Att × Soa × Shp × Col × Loc × Rsp	2,30	572.13	0.45	12.66	0.45

Soa = stimulus-onset asynchrony. * $p < .05$, ** $p < .01$.

First, we address effects that are not specific to the repetition or alternation of particular stimulus or response features, that is, main effects of, and interactions between the attention factor and SOA. As these effects reflect the impact of task overlap, we call them *multiple-task effects*.

Second, we address effects that are restricted to the repetition or alternation of a single stimulus or response feature, either in form of a main effect or in interaction with Attention or SOA. These effects are likely to reflect some kind of priming, i.e., leftover activation of a feature code, or some action triggered by that (e.g., inhibition of return with location repetitions). We thus call them *priming effects*.

Third, we consider interactions between effects of stimulus-feature repetitions or alternations. Such effects show that the impact of repeating a particular feature depends on the repetition or alternation of another feature, which implies that the corresponding feature codes act as a unit. As we take this to reflect the integration of feature codes we call those effects *stimulus-integration effects*.

Finally, we discuss interactions between the effects of repeating or alternating one or more particular stimulus feature(s) on the one hand and the effect of repeating or alternating the response. To the degree that such effects

can be observed (which is only possible in Experiment 2) they can be taken to imply the integration of features across perception and action, which is why we call them *stimulus—response-integration effects*.

Multiple-task effects. Figure 3 gives an overview of the impact of our attentional manipulation (i.e., the memory probe task) and of SOA on RTs and PEs in Experiments 1 and 2. Introducing the memory task produced pronounced RT costs without increasing the error rates reliably—even though a numerical trend is obvious in the errors from Experiment 2. SOA had a strong impact as well by increasing both RTs and errors at shorter SOAs. This impact was modified by attention-SOA interactions, which affected both measures from Experiment 1 and RTs from Experiment 2. As Figure 3 shows, the interference from the memory task is particularly strong at the shortest SOA of Experiment 1.




Figure 3. Reaction times and percentage of errors in Experiments 1 (straight lines) and 2 (broken lines), as a function of stimulus-onset asynchrony (SOA) and attention (S1 unattended vs. attended).

Similar effects have been observed in a couple of recent studies by Jolicoeur and colleagues, summarized in Jolicoeur, Dell'Acqua, and Crebolder (2000). For instance, Jolicoeur and Dell'Acqua (1998) found that having subjects encode between one and three masked letters for later report delays

a binary-choice response to a tone the more letters are encoded and the shorter the SOA between letter and tone is. They attribute this effect to the need to consolidate stimulus information into some short-term store before a concurrent task can be taken on or pursued. Even though our stimuli were not masked it is reasonable to assume that S1 was also consolidated for the later memory probe, which delayed responding to S2 in attended conditions if SOA was short.

However, consolidation is unlikely to account for all aspects of our findings. In particular, RTs from both experiments and the errors in Experiment 2 provide evidence of performance costs in the attended condition that do not disappear at longer SOAs, that is, performance in this condition reaches its asymptote at a level that is considerably lower (or higher, in terms of RT and PE) as that reached in unattended conditions. Hints towards similar differences in asymptote were also obtained in the Jolicoeur and Dell'Acqua (1998) study, but only with memory loads of more than one item. One explanation for this difference might be that Jolicoeur and Dell'Acqua's task required the report of only one feature per item (e.g., the letter name) whereas we required subjects to maintain three features. If so, we would need to compare our findings with Jolicoeur and Dell'Acqua's three-item condition, and here even these authors found differences in asymptote. The only problem with this interpretation is that findings by Luck, Vogel, and colleagues (Luck & Vogel, 1997; Vogel, Woodman, & Luck, 2001) suggest that what matters for memory performance is the number of items but not the number of their features. However, the main focus of these authors was on memory limitations rather than on the impact of memory processes on performance in concurrent tasks, and their results do not rule out that this impact increases as a function of features. Also, they took care to prevent subjects from verbally encoding the items, whereas verbal encoding was certainly an option in our experiments. If our subjects had used this option, maintaining three features would in fact have implied the storage of three different items, which again would fit with the observation that our findings compare well with those of Jolicoeur and Dell'Acqua's three-item condition.

In summary, our findings reflect two types of intertask interference. One is restricted to short SOAs, where the memory task creates particularly visible performance deficits in the RT task, presumably due to the consolidation of S1-related codes. The other type of interference is also induced by the memory task but affects performance across the whole SOA range tested. These dual-task costs are likely to stem from processes responsible for the maintenance of feature-related information. Most important for our present purposes, the memory probe task produced considerable effects, which suggests it was successful in inducing increased attention to S1.

Priming effects. Figure 4 gives an overview of the effects of repeating versus alternating single stimulus features, as a function of attention to S1 and SOA. Let us first focus on *shape*, the only stimulus feature that was nominally relevant for the S2-R2 task. As expected, repeating shape produced benefits at the shortest SOAs, and this benefit was further boosted by increasing attention to S1. This pattern is reflected in the reliable interactions of shape repetition with attention and SOA obtained in Experiment 1 and in the two-way interactions of shape with attention and with SOA in Experiment 2. But there is

a second, negative effect that is confined to the attention condition and the longer SOAs (and, with regard to errors, to Experiment 2). Such reversals from positive to negative repetition effects are a common observation (e.g., Kirby, 1980; Kornblum, 1973). The received view is that positive and negative effects are due to different processes: While the former reflect automatic priming from leftover activation of the codes of the preceding stimulus or response, the latter represents a more strategic expectation bias towards stimulus (or response) alternation (e.g., Soetens, Boer, & Hueting, 1985). If so, one would indeed predict that such “later”, negative repetition effects would be restricted to conditions where the event the alternation bias is based on was attended.

For *colour*, no reliable main effect or interaction involving attention or SOA was obtained, even though Figure 4 hints to a possible priming effect at the shortest SOA. As the following sections will provide evidence that S1 colour was processed, we attribute the absence of colour-related priming effects to the fact that colour was not task relevant (cf. Hommel, 1998), neither directly nor, as we will explain below, indirectly.

The *location* stimulus was involved in several RT effects. In Experiment 1, there was an overall cost of location repetitions that was more pronounced in the unattended condition. This pattern likely reflects inhibition of return (IoR), the widespread observation that attending to an actually irrelevant stimulus impairs later responses to relevant stimuli appearing in the same location (e.g., Maylor, 1985; Posner & Cohen, 1984). Experiment 2 shows a different pattern resulting in an interaction of location and SOA, modified by a three-way interaction with attention. The former reflects the transition of a positive into a null or even negative effect as SOA increases, while the latter indicates that this tendency was restricted to the attended condition. In the absence of further evidence we hesitate to interpret these numerically very small effects. However, it is interesting to note that both attention conditions of Experiment 2 yielded results that are similar to those from the attended condition of Experiment 1. This might indicate that having people to respond to a stimulus releases it from producing IoR even though neither the identity of the stimulus nor its location matters for the task at hand. Another interesting observation is that location repetition effects affected RTs only but not error rates. Such a finding is consistent with claims that IoR does not impair the processing of the stimuli that appear at a previously cued location but only slows down responding to them (Fuentes, Vivas & Humphreys, 1999; Taylor & Klein, 2000).

Finally, response repetition (which was involved in Experiment 2 only) did not yield a main effect or an interaction involving attention, but it did interact with SOA. Whereas the 1100 ms SOA produced benefits of response repetitions in RTs (610 and 619 ms, for repetition and alternation, respectively) and PEs (8.2% and 9.8%), the 2100 ms SOA yielded no difference (584 vs. 583 ms and 8.2% vs. 7.9%, respectively), and the longest SOA produced a disadvantage of repetitions (561 vs. 556 ms and 8.1% vs. 6.9%, respectively). As in the case of shape repetitions, this pattern is consistent with the assumption of an automatic priming component of repetition effects, predominant at short SOAs, and a strategic expectation bias that emerges at longer SOAs (Soetens et al., 1985).




Figure 4. Repetition benefits ($RT_{\text{alternation}} - RT_{\text{repetition}}$, $PE_{\text{alternation}} - PE_{\text{repetition}}$) in Experiments 1 (straight lines) and 2 (broken lines), as a function of stimulus-onset asynchrony (SOA) and attention (S1 unattended vs. attended).

In summary, standard priming effects with repetition benefits at short and alternation benefits (i.e., repetition disadvantages) at longer SOAs were observed for stimulus-shape and response repetition. Stimulus location merely showed evidence of an IoR-type pattern if S1 was not attended or relevant in any way, and stimulus colour showed no reliable effect at all.

Stimulus-integration effects. Across the two experiments, we obtained four clusters of results that involved interactions between stimulus-feature repetition effects. The first is actually a single finding from Experiment 1, showing that *shape* and *colour* had an interactive effect on PEs. This effect exhibited the typical crossover pattern with better performance for colour repetitions if shape was also repeated than if it was alternated (4.7% vs. 6.1%) but worse performance for colour alternations if shape was repeated than if it was alternated (5.9% vs. 5.6%). The corresponding RT effect followed a similar pattern but did not reach significance. It may be interesting to note that we have often observed this effect in both published (Hommel, 1998) and unpublished studies, and it often turns out to either just pass or just not pass the significance criterion. A possible explanation of this notorious unreliability may be that people integrate the irrelevant colour of a stimulus with its relevant shape to the degree that the colour is sufficiently salient—assuming that what counts as sufficient varies from subject to subject. This would suggest that which features are integrated depends on both top-down factors with a preference for task relevant information and bottom-up factors that attract attention in an automatic fashion (Dutzi & Hommel, 2003). We will get back to this issue below.

The second cluster involves interactions between *shape* and *location*. These factors produced two-way interactions in both experiments. The pattern was as expected—better performance for shape and location repeated, and shape and location alternated, than for partial matches; hence, combinations of one being repeated and the other alternated (see Figure 5). However, in case of the PEs in Experiment 2, the interaction was further modified by SOA. Separate ANOVAs confirmed that this was due to the two-way interaction being significant for the shortest SOA only. Moreover, the pattern of the interaction at the 1100 ms SOA was more or less opposite of that obtained for RTs, showing worse performance for shape and location repeated (11.5%) than for both alternated (8.9%) and shape-only (8.4%) or location-only repetitions (8.9%). Although this might indicate a speed-accuracy trade-off, we note that such an inversion is not only uncommon in experiments of this sort (e.g., Dutzi & Hommel, 2003; Hommel, 1998, 2003) but also runs counter the findings from the other SOAs and from Experiment 1—where RT crossover patterns were accompanied by either null effects or comparable effects in error rates.




Figure 5. Reaction times in Experiments 1 (straight lines) and 2 (broken lines), as a function of the repetition vs. alternation of stimulus location, stimulus shape, and stimulus colour.

The third cluster involves interactions between colour and location. Evidence of such interactions was only obtained in Experiment 1, where errors produced a two-way interaction and RTs a three-way interaction including attention. As shown in Figure 6, the patterns underlying these two effects are very similar: Colour repetitions had no impact if S1 was unattended, while attending it produced a crossover interaction of colour and location. Interestingly, this interaction does not show the “integration signature” of worse performance with partial matches but, on the contrary, better performance for colour-only or location-only repetitions than for the both-repeated or both-alternated conditions.

The fourth cluster involves interactions between shape, colour, and location—all three stimulus features. Such interactions occurred only in Experiment 2, where we obtained a three-way interaction in RTs and a five-way interaction involving attention and SOA in PEs. As Figure 5 indicates, the three-way interaction was due to a decrease of the shape-by-location interaction effect if colour was repeated. To figure out the effect underlying the five-way interaction we ran separate ANOVAs for all combinations of Attention and SOA on the error data from Experiment 2. The outcomes indicated that the three stimulus features interacted only in the 2100 ms SOA cell of the unattended condition. That interaction corresponds to what we see in RTs: Fewer signs of a disadvantage for shape-only or location-only repetitions over both-repeated and both-alternated if colour is repeated (6.8%, 6.2% vs. 8.3%, and 9.0%, respectively) than if colour is alternated (6.8%, 7.3% vs. 5.2%, and 6.6%, respectively).




Figure 6. Reaction times and percentage of errors in Experiment 1, as a function of the repetition vs. alternation of stimulus location and stimulus colour, and attention (S1 unattended vs. attended, white and black symbols, respectively).

To summarize, we find evidence of several interactions between stimulus features. From a theoretical point of view, a number of aspects of these findings are of relevance. First, most interactions are bilateral, hence, involve only two of the three manipulated stimulus features. Second, even the few hints towards an interaction of all three features do not suggest that complete integration took place. If it would have, repeating one more feature should have increased the impact of the other features; yet, the interaction of shape and location decreased if colour was repeated (see Figure 5). Third, there was no support of the idea that integrated feature compounds are addressed by location. If they were, the impact of feature repetitions and their interactions should have increased if, or even be restricted to situations where stimulus location is repeated; yet, a look at Figure 5 confirms that location repetitions did in no way boost the interactions between colour and shape repetition. Fourth, some two-way interactions between features seem to be more reliable and replicable than others. In particular, interactions between shape and location seem to belong to the more reliable effects while the two interactions involving colour seem to be less reliable. Interestingly, colour effects tended to come and go together, hence, all occurred in one but not the

other experiment. Finally, there was no evidence of any strong impact of attention or SOA on the interactions involving shape, the nominally task-relevant stimulus feature, and even the remaining interactions did not suggest any strong dependency on SOA.

Stimulus-response-integration effects. The effects falling into our last category all involve response repetition and, therefore, all come from Experiment 2. Let us first turn to interactions involving repetitions of the response and one stimulus feature. Figure 7 provides an overview of the two-way interactions in RTs as a function of attention. It is obvious that all three stimulus features interact with the response, and that they do so as expected: Repeating a response produces better performance than alternation, but only if the respective stimulus feature (shape, colour, or location) is also repeated. If it is not, the repetition effect turns into an alternation benefit. Some of these interactions were modified by attention and SOA. As evident from Figure 7, the interactions between *shape* and *response* and between *location* and *response* are substantial (and reliable) under both attention conditions but somewhat more pronounced if S1 is attended. SOA also matters, which can be seen in Figure 8. Both the interactions between shape and response and between location and response are most pronounced at the shortest SOA and then decrease as SOA increases. However, even at the longest SOA they are still highly reliable. The shape-by-response interaction in PEs is further modified by a four-way interaction involving attention and SOA, indicating that the decrease of the shape-by-response interaction across SOAs is more pronounced in the S1unattended than in the attended condition.

Let us now turn to interactions involving the response and two stimulus features. There were three clusters of interactions of that sort. First, *shape*, *colour*, and *response* produced a three-way interaction in RTs, which was modified by a four-way interaction with attention. As shown in Figure 9, the interaction between shape and response was slightly bigger if colour was also repeated (compare straight vs. broken lines), and this increase was more pronounced if S1 was attended (i.e., in the top part of the figure). Importantly, however, the shape-by-response interaction was reliable for all four combinations of colour repetition and attention.

Second, there was a three-way interaction of *shape*, *location*, and *response* in error rates, which was modified by attention and accompanied by a four-way interaction of shape, location, response, and SOA in RTs. The RT effect is shown in Figure 8. As confirmed by separate ANOVAs, shape, location, and response interact at the shortest SOA only, where the shape-by-response interaction is increased if location is repeated. The error-related effects are presented in Figure 10. They mirror the impact of colour on the shape-by-response in showing that repeating location increases the interaction between shape and response (compare straight vs. broken lines), and that it does more so if S1 is attended.

Third, we obtained a four-way interaction of *colour*, *location*, and *response* with attention in error rates. Figure 11 shows that part of this effect is due to that colour repetitions increase the location-by-response interaction if S1 is attended (see right panel). In contrast, in the unattended condition the location-by-response interaction is less pronounced for colour repetitions than alternations.




Figure 7. Reaction times and percentage of errors in Experiment 2 for the repetition vs. alternation of stimulus shape (top panel), stimulus location (middle panel), and stimulus colour (bottom panel), as a function of response repetition and attention (S1 unattended vs. attended, white and black symbols, respectively).




Figure 8. Reaction times in Experiment 2 for the repetition vs. alternation of stimulus shape (black vs. white symbols) and stimulus location (straight vs. broken lines), as a function of response repetition and stimulus-onset asynchrony (SOA, in ms).




Figure 9. Reaction times in Experiment 2 for the repetition vs. alternation of stimulus shape (black vs. white symbols) and stimulus colour (straight vs. broken lines), as a function of response repetition and attention.

To summarize, we find evidence that all three stimulus features were able to modify the effect of response repetitions or, depending on how one looks at it, that response repetitions modified the impact of repetitions of stimulus shape, colour, and location. The task-irrelevant colour dimension seemed to play a minor, more modifying role: The interaction between colour and response repetition was the by far least pronounced but a repeating colour in several cases increased the interactions of other stimulus features with the response. Shape and location repetitions interacted more strongly with the response, and these interactions were further boosted by attending to S1. Increased attention to S1 increased a number of interactions but there was no evidence that endogeneous attention was necessary for an interaction to occur. Also, SOA had no dramatic effects but its impact was more obvious than in the interactions between stimulus features.




Figure 10. Percentage of errors in Experiment 2 for the repetition vs. alternation of stimulus shape (black vs. white symbols) and stimulus location (straight vs. broken lines), as a function of response repetition and attention.




Figure 11. Percentage of errors in Experiment 2 for the repetition vs. alternation of stimulus location (black vs. white symbols) and stimulus colour (straight vs. broken lines), as a function of response repetition and attention.

Conclusions

The two experiments of this study aimed at addressing three questions regarding the integration of stimulus and response features: How complete is feature integration? Are feature bindings addressed by location? and How are feature priming and feature integration related? In particular, we were interested to see whether the completeness of integration, the integration of location codes, and the role of priming and integration would change over time—i.e., SOA—and depend on the amount of attention spend on the to-be-integrated stimulus. All in all, it is fair to say that the impact of attention and time was rather limited. But let us discuss the three guiding questions in turn.

How complete is feature integration? Importantly, we were able to replicate the main finding of Hommel (1998), namely, that the impact of repeating a stimulus feature depends on whether or not other stimulus features and/or the response are repeated as well. Of the binary interactions we obtained, those involving stimulus shape, stimulus location, and response were particularly pronounced and reliable. On the basis of our present data, we are unable to exclude that this reflects characteristics of these particular stimulus and response dimensions or modalities. However, there are two observations that speak against such an interpretation. One is that Hommel (1998) found the predominance of shape-response interactions to turn into a predominance of colour-response interactions when colour was made the relevant dimension for the S2-R2 task. Another is that Hommel (2003) was able to eliminate interactions involving stimulus location by using nonspatial responses (single vs. double key presses). In view of these findings we interpret the present preponderance of stimulus shape, stimulus location, and response as reflecting the impact of (RT-)task relevance. Indeed, shape was relevant for the RT task by virtue of signalling R2, responses were relevant by definition of the task, and location was—more indirectly—made relevant by defining the responses in terms of spatial location. From this perspective, the likelihood for a stimulus or response feature to enter binary interactions was determined by task relevance of the dimension on which the feature is defined (Hommel, 2003).

We speculated that integration may begin with creating binary bindings (that dominated in the studies of Hommel, 1998, 2003), which then over time enter a more comprehensive object or event file, and that this process may be boosted either in terms of time or outcome by attending S1. If so, we would have expected interactions among stimulus and/or response features to increase in order as SOA increases, especially in the attended condition. However, Experiment 1 did not produce any evidence of a more than two-way interaction between stimulus features in RTs or errors, be it in the form of a three-way interaction or a higher order interaction involving attention or SOA. Experiment 2 yielded some more evidence of this sort.

First, the three stimulus features were involved in a three-way interaction and in a five-way interaction with attention and SOA. If one considers the increasing completeness of feature integration to make the final product more specific (by virtue of describing the filed object or event more comprehensively), its impact on behaviour should become increasingly selective. Ideally, and ignoring any possible main effect of feature repetition,

one would thus expect that complete integration makes complete repetition special in producing considerably better performance than any other combination of repetitions and alternations. However, a look at Figure 5 shows that this is not what happens, which we think speaks against an interpretation in terms of complete integration. Moreover, such an interpretation would be difficult to bring in line with the observation that the five-way interaction locates the main action at the middle SOA of the unattended condition. A possible way to reconcile the idea of increasing integration with the three-way pattern in Figure 5 (though not as smoothly with the five-way interaction) would be to think of it as showing that the more stimulus features are repeated the less is the impact of each individual feature. That is, repetitions of feature conjunctions may be able to outweigh the impact of partial mismatches of other conjunctions to some extent, which does suggest some sort of higher order integration. But even then it would not be obvious why integration should have been less pronounced in Experiment 1, where we could not find any sign of a higher order interaction. Whichever interpretation one prefers it seems clear, however, that our findings do not suggest that integration comprises a transition from local, binary bindings to one global file where all information converges. Thus, object files seem to consist of a loosely connected, distributed network of bindings rather than one single superstructure (Hommel, 1998, 2003).

Second, there were a number of higher order interactions involving one stimulus feature and the response and two stimulus features and the response. In fact, all three binary combinations of stimulus features interacted with response, and all these interactions were modified by attention to S1. The general pattern of the resulting four-way interactions was rather consistent: Repeating one more stimulus feature increases the interaction between another stimulus feature and the response, but only or mainly in the S1-attended condition. A look at Figures 8, 10, and 11 reveals a hint to the possible mechanism underlying this pattern. In all cases the major contribution to the interaction comes from the condition where S2 shares most features with S1 (e.g., shape and location repeated in Figure 10) but a response alternation is required. It is as if a stimulus that is very similar to the previous one induces a tendency to repeat the response, which in case of response alternation needs to be inhibited. Such a tendency has already been considered to account for response repetition benefits in binary-choice tasks. For instance, Bertelson (1963) claimed that people check new stimuli against representations of previous ones and immediately proceed to repeating the last response if the two stimuli match. It is reasonable to assume that this matching strategy is more likely to be applied if the compared stimuli are both attended, which indeed is the standard condition in studies on response repetition benefits. If so, our finding of attentional modulations of higher level stimulus-response interactions may not have much to do with feature integration processes but, rather, reflect the fact that Bertelson's response selection strategy was restricted to attentional conditions. Such an interpretation would also fit with the four-way interaction depicted in Figure 8: While the shortest SOA shows a pattern indicative of the Bertelson strategy (i.e., slowest performance if a response alternation was required in the face of shape and location repetition) the longer SOAs do not. Indeed, applying the matching strategy presupposes an active memory trace of the previous

response, which according to the interaction of SOA and response repetition was effective at the shortest SOA only.

Thus, taken altogether, we find no strong evidence that having more time available and/or investing more attentional resources to process an event creates a single cognitive structure where information about all features of the event converges. Evidence is also sparse with respect to the less ambitious version of this question whether attention and/or time increase integration, that is, whether the resulting structure becomes more complex. There are some hints to higher order interactions among stimulus-related effects and to higher order interactions between multiple stimulus effects and response repetition, but the patterns of these interactions do not seem to fit the idea of (more) complete feature integration. In particular, the resulting representational structures do not get more specific or selective as a function of attention or time. That is, not all features of a given perception-action event are integrated with each other. What gets integrated seems to be determined by task relevance or, more precisely, by whether the given feature varies on a dimension that in the present task is explicitly or implicitly defined as relevant. In the present RT part of the task, this applied to shape, which was relevant for R2, and to location, which was relevant for the responses. However, it is likely that task relevance is only one factor that affects integration. Stimulus features that are sufficiently salient, such as tones, may enter integration processes even if they are not relevant at all (Dutzi & Hommel, 2003).

Are feature bindings addressed by location? According to Kahneman et al. (1992), object files can only be accessed via spatial information, so that information about the relative or absolute location of its object is an essential ingredient of every object file. If so, feature-binding effects could only be obtained if stimulus location is repeated, which implies that interactions between feature-related repetition effects should always be modified by a higher order interaction with stimulus location. Our results not only replicate previous demonstrations that this prediction is incorrect (Hommel, 1998, 2003; see also Gordon & Irwin, 1996; Henderson, 1994; Henderson & Anes, 1994), they also show that the picture these demonstrations suggest does not change much if attention and time come into play. In particular, a whole number of reliable interactions between effects of stimulus features and between effects of stimulus and response features were obtained in the absence of stimulus location repetition, and even though attending S1 increased some of these effects their existence did not depend on attention or time.

Thus, on the one hand, our findings do not support Kahneman et al.'s (1992) claim that object files are exclusively addressed by location. Rather, it seems that any match counts, which implies that a given stimulus or action event activates or primes all feature compounds that include at least one feature code of a matching value. This view fits nicely with the developmental evidence mentioned above (Leslie & Kaldy, 2001; Leslie et al., 1998) and with the assumption that episodic event representations can be addressed by a match with any feature code they include (Hommel, 1998; Hommel, Müsseler, Aschersleben & Prinz, 2001b; Leslie et al., 1998). On the other hand, however, our findings should not be taken to rule out all possible roles of stimulus location for feature integration. It may well be that integration is under

spatial control, as several authors have claimed (e.g., Treisman, 1988; van der Heijden, 1992; Wolfe, 1994). That is, the criterion for sampling information into the same event representation may well be defined in terms of the location the information is coming from, in addition to possible temporal constraints (Hommel & Akyürek, in press). And yet, this need not necessarily imply that location is coded in the emerging representation.

How are feature priming and feature integration related? With regard to the priming of single features, previous studies yielded a rather inconsistent picture: Some did find reliable effects (Gordon & Irwin, 1996; Henderson, 1994), while others did not (Hommel, 1998; Kahneman et al., 1992). We hypothesized that this apparent inconsistency might be due to the different SOA ranges used in these studies and thought that priming may show up at very short SOAs only. Indeed, Figure 4 and the corresponding analyses clearly indicate that most priming is restricted to the shortest SOA used here, i.e., 200 ms. If we assume that the amount of priming reflects the degree of activation of the respective feature code, this observation suggests that activation and integration do not necessarily go together. Thus, on the one hand, it is likely that what gets integrated is what is currently activated, which implies that the activation of a feature code precedes, and may even be the criterion for its integration (Hommel et al., 2001b; Hommel, Müsseler, Aschersleben & Prinz, 2001a). Once integration has taken place, however, activation is no longer necessary to impact processing (Hommel, 2002). For a concrete example, the temporal overlap of activation in the codes <vertical> and <bottom> creates a temporary link between them, as shown in Figure 1. Without activating these codes the link would not have been created, so that activation necessarily precedes integration. Once the link is established, however, activation is no longer needed: When <vertical> is activated again it will spread activation to <bottom>, and vice versa.

A rather surprising outcome of our study is the impact of attention—or the lack of it. On the one hand, explicitly attending to S1 and storing its features had a pronounced effect on performance, as revealed by a look at Figure 3. This assures us that the attentional part of our manipulation worked. On the other hand, however, the performance deficits produced by this manipulation are likely to reflect the consolidation of object information into working memory (Jolicoeur & Dell'Acqua, 1998; Jolicoeur et al., 2002), a process that is assumed to share capacity with and therefore delay response selection (of R2 in our case). Consolidation into working memory, though, has been likened to feature integration (Luck & Vogel, 1997; Vogel et al., 2001), which one would assume to boost interactions between feature-related effects. Accordingly, one might have expected a stronger impact of our memory probe task on the interactions between feature repetition effects. A tenable account for this might be to assume that our subjects used verbal coding strategies to retain the features of S1—an account that is also suggested by the high asymptote in the attended condition (see above). If so, it was verbal, not visual feature codes that were consolidated, which again is unlikely to provide any particular support for the type of feature integration that we believe to underlie the observed interactions between feature repetition effects. It may therefore well be that taking measures to exclude verbal strategies would allow for a greater impact of the probe task on feature inter-

actions. At any rate, however, it seems clear that these interactions do not require active consolidation to occur.

Further theoretical implications

Taken altogether, the present study confirms that the mere co-occurrence of stimulus features, and of stimuli and responses, is sufficient to induce bindings between their codes. These bindings emerge rather quickly—i.e., within the first hundreds of milliseconds—and remain intact for at least 4 seconds. Which features are bound depends to a large degree on the direct or indirect task relevance of their dimension. That is, feature integration within perception and across perception and action occurs spontaneously, but its outcome is codetermined by the current action goal and the attentional set implemented to achieve it. Let us conclude by pointing out some more general implications our findings, especially with regard to the processing and representation of events in and across perception and action.

It is fair to say that the discussion of possible roles of binding processes and the neural codes is still going on: Some authors have emphasized the need of integration processes in distributed representational systems like the human brain and argued that the synchronization of cell populations might play a major role in binding features belonging to the same event (e.g., Singer, 1994; Treisman, 1996), while others have questioned the very necessity of feature binding (e.g., Cisek & Turgeon, 1999; van der Heijden, 1995) and/or the involvement of neural synchronization in it (e.g., Jellema & Perrett, 2002; van der Velde & de Kamps, 2002). In this and previous studies we have taken a pragmatic, empirical stance and looked whether or under which circumstances evidence for binding can be found at all. As discussed, there is increasingly strong and converging evidence that binding does take place and that it does affect performance, and the present study extends this evidence by showing that it does so at least for a couple of seconds and even if the bound features need not be processed very deeply. Given the ongoing controversy it seems particularly important, however, to point out what our findings do not show.

First, there is no evidence that visual features can affect behaviour only if, or only after they have been integrated—as one would have expected on the basis of Kahneman et al. (1992; cf. Phaf, van der Heijden, & Hudson, 1990). If anything, our observations suggest the opposite: Main effects of feature repetitions—which represent the impact of a feature code independent from, or on top of any binding it may be involved in—were restricted to very short SOAs, whereas the interactions that we attribute to feature integration were found across the whole SOA range. As pointed out above, this suggests the existence of at least two functionally separable representations of an event: A map or system in which the features of an event are coded by (rather briefly) activating their respective feature codes and a network of links connecting the codes that are coactivated within a particular integration window (Hommel et al., 2001b; cf. Styles & Allport, 1986). It is tempting to relate the former to the increase of firing rates of cell populations within feature-specific representational maps and the latter to the synchronization of the firing patterns of cell populations across representational maps (as suggested by Singer, 1994; Treisman, 1996; and others), but our data do not

directly speak to the issue of how activation and binding of feature codes is neurally implemented.

Second, and relatedly, demonstrating aftereffects of feature integration does not necessarily rule out the possibility that the integrated features were encoded and processed independently, as predicted by parallel models of feature processing (e.g., Bundesen, 1990) and observed at least under some conditions (Bundesen, Kyllingsbæk & Larsen, 2003). Obviously, a feature code can only be integrated and retrieved if the respective feature has been successfully encoded, and there is no reason why the likelihood of encoding one member of an integrated event representation should depend on the likelihood of encoding another member of the same representation. However, once two or more feature codes are encoded and bound, they tend to prime each other and, thus, act as a functional unit—an event file. This observation does not challenge the basic assumptions most attentional models rely on, but it makes these models incomplete with respect to the outcome of feature integration processes (see Logan, 2002, for an attempt to overcome this shortcoming).

Third, our failure to find an effect of our attentional manipulation shows that binding does not require an explicit intention to integrate information about an event, nor does it presuppose that its outcome is of any use for the task at hand. Note that this does not demonstrate that integration is independent from the attentional set or the availability of attentional resources. To the contrary, the important role of set is obvious from the observation that features are integrated only if they are directly or indirectly related to the task (Hommel, 1998, 2003). That is, feature integration seems to be controlled by the task goal, which may be maintained in working memory and provide top-down support for information related to task-relevant feature dimensions (Desimone & Duncan, 1995; Pratt & Hommel, 2003). Once this support is provided integration seems to proceed automatically, however. How automatically it proceeds may well depend on the concrete visual situation, such as the presence and number of alternative targets and distractors. Whereas no explicit intention to integrate was necessary under the Spartan conditions in our experimental set-up—with only one, salient stimulus appearing in a rather wide time window—it seems likely that more complex visual situations, such as the multielement displays used by Kahneman et al. (1992), make selection and integration goals more relevant.

To conclude, irrespective of the eventual outcome of the “binding-problem” debate, the present study provides converging evidence that feature binding both within visual perception and across perception and action is a real empirical phenomenon that calls for theoretical consideration and further investigation.

Chapter 2

Moderate alcohol consumption impairs feature binding in visual perception but not across perception and action

Abstract

Animal studies suggest a relationship between activation of the cholinergic system and neural synchronization, which again has been suggested to mediate feature binding. We investigated whether suppressing cholinergic activity through moderate alcohol consumption in healthy humans affects behavioral measures of feature binding in visual perception and across perception and action. Indeed, evidence of the binding of shape and color, and of shape and location, of visual objects disappeared after alcohol consumption, whereas bindings between object features and the manual response were unaffected.

Introduction

Imagine that you are watching a red cat and a black dog. Given that colors, shapes, locations, and semantic features are processed in different cortical areas, how does our brain correctly integrate the features belonging to the same event but does not, say, make you perceive a red dog and a black cat? To solve this feature-binding problem the brain needs to employ some mechanism that interlinks and integrates the neural patterns coding the features of a given perceptual event (Treisman, 1996). One such mechanism might be the neural synchronization of cell populations (Abeles, 1991; Engel & Singer, 2001). That is, the firing rates of cells coding features of the same perceptual event may synchronize, which would provide a neural marker of "eventhood" and support the individual codes in their competition with other codes in their respective feature domains. Indeed, transient increases of synchronization in the gamma frequency range have been observed in perceptual tasks like figure-ground distinctions and feature binding (Engel & Singer, 2001), switching between bistable visual figures (Keil, Muller, Ray, Gruber & Elbert, 1999), or the retention of visual patterns in short term memory (Tallon-Baudry, Kreiter & Bertrand, 1999). Even motor tasks have revealed reliable links between neural synchronization and integrative cognitive processes such as the planning of multi-featured actions (Hari & Salenius, 1999).

Purpose of the study

If the idea that neural synchronization at least mediates the binding of perceptual and, perhaps, action features is correct, one would expect that factors that are known to impact synchronization in a particular fashion affect

behavioral measures of feature integration in the same way. One such candidate factor is alcohol, which is suspected to cause a hypoactivity of the cholinergic system. Apart from chronic alcohol consumption (Little, 1999), acute ethanol intake has been found to inhibit muscarinic receptors of the cholinergic system (Minami, Vanderah, Minami & Harris, 1997; Sanna, Dildy-Mayfield & Harris, 1994), which again is involved in driving at least visually induced synchronization (Rodriguez, Kallenbach, Singer & Munk, 2000; Rodriguez, Kallenbach, Singer & Munk, 2001). The aim of the present study was thus to test whether the intake of alcohol hampers the binding of features and whether this impact is specific to visual features.




Figure 1. Sequence of events in the present experiment [cf., (Hommel, 1998)]. A response cue signaled a left or right key press (R1) that was to be delayed until presentation of S1, a red or green, vertical or horizontal line in a top or bottom box. S2 appeared 1 s later - another red or green, vertical or horizontal line in the top or bottom box. S2 shape signaled R2, also a quick left or right key press. R2 speed and accuracy were analyzed as function of the repetition vs. alternation of stimulus shape, color, and location, and of the response.


We adopted the task from Hommel (1998), which involves the repetition of task-related and unrelated visual features and of the response (see Figure 1). The standard findings are interactions between (a) the task-related stimulus features (e.g., shape, if and only if it signals the response [Hommel, 1998], and location, if and only if the responses are spatially defined (Hommel, 2003); (b) the non-spatial stimulus features (e.g., shape and color); and (c) the relevant stimulus features and the response (for an overview, see Hommel et al., 2001). The patterns of these interaction all look alike: Performance is impaired in partial-repetition trials, that is, if one stimulus feature (or the response) is repeated while the other is not. This suggests that the mere co-occurrence of a feature-feature or feature-response conjunction is sufficient to create a temporary binding of the respective feature codes—an "event file" (Hommel, 1998; Hommel & Colzato, 2004; Hommel, Müsseler, Aschersleben, & Prinz, 2001a). Reactivating one member of this binding (as in the case of feature repetitions) reactivates other member(s) as well, which leads to confusion and requires a time-consuming re-binding process in partial-repetition trials. Importantly for present purposes, this partial-repetition cost can be taken to indicate feature-feature and feature-response binding, which

is why we chose it as behavioral marker. According to the reasoning outlined above we thus expected moderate alcohol consumption to decrease partial-repetition costs.

Method

Seventeen right-handed volunteers served in two experimental sessions. Informed consent was obtained from all participants after the nature and possible consequences of the study were explained to them, and the protocol was approved by the local ethics committee. Subjects were social drinkers (2-3 units per day on average) in the age range of 20-30 years, they were healthy non-smokers, not on medication or drugs, and without neurological or psychiatric history according to self-report. To minimize circadian-cycle influences, experimental sessions always started at 15:00 h, after subjects had had their regular lunch, but abstained from eating and from drinking caffeine-containing liquids for 2.5 h and abstained from alcohol consumption for 24 h. A double-blind, placebo-controlled, randomized cross-over design with counterbalancing of the order of conditions was used to avoid alcohol-expectancy effects. Sessions were separated by 3-7 days. Placebo and dose alcohol quantities corresponded to 0.00 and 0.45 g/kg, respectively. Body-weight dependent measures of vodka (containing 37.5% ethanol) were dissolved in orange juice such that total liquid volume amounted to 500cc. Adding peppermint oil (Ridderinkhof, de Vlugt, Bramlage, Spaan, Elton, Snel & Band, 2002) and serving the beverage in a sealed milkshake beaker effectively prevented that subjects tasted or smelled the presence of alcohol. Blood-alcohol concentration, recorded before the beginning (after 30 min of drinking) and after the end of the experiment, averaged 0.34‰ (S.D. 0.04) in dose sessions and 0.0‰ (S.D. 0.0) in placebo sessions.

Subjects completed a version of the task adapted from Hommel (1998), (see Figure 1). They faced three gray, vertically arranged boxes in the middle of a monitor and carried out two responses per trial. R1 was a delayed simple reaction with the left or right key, as indicated by a 100% valid response cue (three left- or right-pointing arrows in the middle box) that preceded the trigger stimulus S1 by 3000 ms. S1 varied randomly in shape (a thin vertical or horizontal line), color (red or green), and location (top or bottom box). R1 was to be carried out as soon as S1 appeared, independent of its shape, color, or location; i.e., subjects were encouraged to respond to the mere *onset* of S1. R2 was a binary-choice reaction to the shape of S2 (vertical or horizontal orientation), which also appeared in red or green, and in the top or bottom box, 1000 ms after S1 onset. Responses to S1 and to S2 were made by pressing the left or right shift-key of the computer keyboard with the corresponding index finger. Each session was composed of a factorial combination of the two possible shapes, colors, and locations of S2, the repetition vs. alternation of shape, color, location, and the response, and three replications per condition.

**Fig. 2.**

Mean reaction times (symbols) and their standard errors (numbers) as a function of treatment (placebo vs. alcohol) and the repetition vs. alternation of stimulus shape and stimulus color (top left panel), of stimulus shape and stimulus location (top right panel), of response and stimulus shape (bottom left panel), and of response and stimulus location (bottom right panel). Note that apart from an interaction with shape, treatment affected only the interactions of shape and color and of shape and location.

Results

After excluding trials with missing or anticipatory responses (1.3%) mean reaction times (RTs) and proportions of errors for R2 (i.e., the response to S2) were analyzed as a function of treatment (placebo vs. alcohol) and the repetition vs. alternation of response (R1→R2), stimulus shape, color, and location (S1→S2). Replicating earlier findings (Hommel, 1998; Hommel & Colzato, 2004), RTs revealed significant interactions between shape and color, $F(1,16) = 6.01$, $P < 0.05$, and shape and location, $F(1,16) = 8.52$, $P < 0.01$ —repeating one but not the other feature slows down responding. However, apart from an interaction with shape, $F(1,16) = 5.69$, $P < 0.05$, treatment was involved in a three-way interaction with shape and location, $F(1,16) = 7.47$, $P < 0.05$, and in a four-way interaction that also comprised color, $F(1,16) = 4.93$, $P < 0.05$. To disentangle these effects, separate analyses of variance were run on the data from placebo and alcohol conditions. As obvious from Fig. 2 (top panels), the shape-by-color and shape-by-location interactions were reliable in the placebo condition, $F(1,16) = 6.53$, $P < 0.05$, and $F(1,16) = 21.83$, $P < 0.001$, but not in the alcohol condition, $P_s > 0.25$. This can be taken to indicate that the intake of alcohol prevents the binding of visual features to a degree that overt performance is no longer affected. Interestingly, alcohol did not impact effects reflecting stimulus-response bindings: Although we replicated the standard interactions between response and shape, $F(1,16) = 11.16$, $P < 0.005$, and response and location, $F(1,16) = 16.12$, $P < 0.001$, these effects were not involved in any interaction with treatment (see Fig. 2, bottom panels). The errors followed the same pattern: Significant interactions were obtained between shape and response, $F(1,16) = 11.47$, $P < 0.005$, and location and response, $F(1,16) = 4.65$, $P < 0.05$, due to fewer errors in conditions where the stimulus feature and the response were both repeated or both alternated, as compared to conditions where the stimulus feature but not the response was repeated, or vice versa. Again, these effects were not modified by treatment.

Conclusions

Our results suggest two conclusions: Moderate alcohol consumption impairs feature binding and it selectively affects the binding of visual features while sparing cross-domain bindings between visual features and manual responses. The finding that alcohol affects feature binding at all supports the hypothesized links between the cholinergic system and neural synchronization (Rodriguez et al., 2000; Rodriguez et al., 2001) on the one hand and between synchronization and feature integration (Abeles, 1991; Engel & Singer, 2001; Treisman, 1996) on the other. We can imagine at least two reasons for why the impact of alcohol might be restricted to local, in our case visual-visual binding. First, the cholinergic system may selectively drive neural synchronization in visual or, more generally, in perception-related areas of the cortex but not in areas involving motoric activity. Although we know of no studies that would definitely rule out this possibility, the central role of acetylcholine in voluntary motor control and its rather direct impact on neural activity in the motor cortex (Matsumara, Sawaguchi & Kubota 1990) renders

this possibility somewhat unlikely. A second reason considers the anatomical distance of the to-be-synchronized neural networks. Local integration processes in both visual and motor cortex are commonly associated with synchronization frequencies in the gamma band (Engel & Singer, 2001). In contrast, synchronization between more distant networks, such as in visuomotor integration, has been found to use the lower, beta frequency band (Roelfsema, Engel, Koenig & Singer, 1997). If we assume that alcohol intake impairs neural synchrony by increasing the variability of firing rates, and if we consider that this should affect higher frequencies more than lower frequencies (Kopell, Ermentrout, Whittington & Traub, 2000), local, short-range bindings should indeed be more vulnerable to alcohol-induced effects than long-range bindings.

Chapter 3

Caffeine, but not nicotine enhances visual feature binding

Abstract

The distributed organization of the human visual cortex calls for a mechanism that integrates and binds the features of a perceived event, and neural synchronization is a prime candidate to serve that purpose. Animal studies suggest that synchronization in the visual cortex is enhanced by the muscarinic-cholinergic system. Here we show that in healthy humans the binding of shape and color, and of shape and location of visual objects is increased by stimulating the muscarinic-cholinergic system (caffeine consumption) but not by stimulating the nicotinic cholinergic system (nicotine consumption). Binding across perception and action is unaffected by either manipulation, suggesting a specific link between the visual system and the muscarinic-cholinergic system.

Introduction

When we are facing multiple objects, a red apple and a yellow banana, say, the features of these objects are coded and processed in different cortical areas, and yet what we perceive are well integrated objects and not mere bundles of attributes. The human brain thus seems to employ some kind of feature-binding mechanism that integrates the neural patterns coding the features belonging to a given event and makes us validly perceive that bananas are yellow and not red (Treisman, 1996). One candidate for such a mechanism are conjunction detectors, that is, neural units that are selective for the presence of particular feature combinations (Mozer, 1991; Riesenhuber & Poggio, 1999). Coding by means of conjunction detectors makes sense for processing highly probable, evolutionary important feature conjunctions. However, excessive numbers of conjunction detectors would be necessary to code any arbitrary feature combination, suggesting that frequently changing feature relations are processed in a different way (Colzato et al., 2004; Hommel, 2004; Singer, 1994). Better suited for this case seems to be the neural coupling of cell populations (Abeles, 1991; Engel & Singer, 2001), which is assumed to be achieved by coordinating and synchronizing the firing rates of cells referring to the same event—i.e., feature conjunctions may be coded spontaneously through the temporal coherence of their neural codes.

The possible role of neural synchronization in binding features within the visual and the motor system has been implicated in many studies. Apart from numerous animal studies (Abeles, 1991; Engel & Singer, 2001; Roelfsema, Engel, Koenig & Singer, 1997), transient increases in synchronization in the gamma frequency range have been observed in

healthy humans in visual tasks like figure-ground discrimination (Engel & Singer, 2001), switching between bistable visual figures (Keil, Muller, Ray, Gruber & Elbert, 1999), and the retention of visual patterns in short-term memory (Tallon-Baudry, Kreiter & Bertrand, 1999). Motor tasks have revealed similar relationships between synchronization and integrative cognitive processes such as the planning of multi-featured actions (Hari & Salenius, 1999).

Although the available evidence points to a link between neural synchronization and feature binding, it remains unclear how and under which circumstances synchronization emerges, which makes it difficult to appreciate its true functional role. A potentially important observation in this context is the finding that, in the cat, neocortical synchrony in the gamma band (~30-70 Hz) is enhanced by muscarinic-cholinergic agonists and disrupted by antagonists (Rodriguez, Kallenbach, Singer & Munk, 2000, 2001; Rodriguez-Bermudez, Kallenbach, Singer & Munk, 2004). Empirical evidence for this link is consistent with Colzato, Erasmus & Hommel (2004), who observed in humans that alcohol—which is known to cause hypoactivity of the cholinergic system (Minami, Vanderah, Minami & Harris, 1997; Sanna, Dildy-Mayfield & Harris, 1994)—impairs feature binding in visual perception but not binding across perception and action. This fits with the hypothesis that the binding of visual features is driven by the muscarinic-cholinergic system (Metherate, Cox & Ashe, 1992; Rodriguez et al., 2000, 2001; Rodriguez-Bermudez et al., 2004).

Purpose of the study

The aim of the present study was to set up a specific test of this hypothesis in healthy human subjects: Can it be demonstrated that, first, cholinergic agonists increase the binding of visual features but not of other features and that, second, this increase is specifically driven by agonists of the muscarinic but not the nicotinic cholinergic system? The two agonists we compared were caffeine and nicotine. The behavioural effects of caffeine have been attributed to several neuromodulatory systems, including dopaminergic, GABA, serotonergic, cholinergic, and noradrenergic pathways (Nehlig, Daval & Debry, 1992). However, there is evidence for a direct muscarinic-cholinergic link between caffeine and the processing and short-term memory of visual information.

First, caffeine impacts the muscarinic-cholinergic but not the nicotinic-cholinergic system (Sorimachi, Yamagami & Nishimura, 1992). Cholinergic systems are under tonic inhibitory control by endogenous adenosine, as suggested by observations of an increase of adenosine extracellular concentrations in the basal forebrain cholinergic areas of the behaving cat during prolonged wakefulness (Porkka-Heiskanen, Strecker, Thakkar, Bjorkum, Greene & McCarley, 1997) and of dramatic decreases in waking induced by the perfusion of adenosine into the same regions (Portas, Thakkar, Rainnie, Greene & McCarley, 1997). Also, the state of prolonged sustained wakefulness can be mimicked by increasing adenosine levels in basal forebrain cholinergic regions but not by increasing those levels in noncholinergic areas (Porkka-Heiskanen et al., 1997). Given that caffeine (and other xanthines) are adenosine antagonists, it thus makes sense to assume that they unfold their arousing impact by reducing the amount of

inhibition adenosine exerts on cholinergic pathways. Moreover, evidence that the cholinergic impact on cortical states is mainly muscarinergic (Lamour, Dutar & Jobert, 1982) suggests a central role of muscarinic-cholinergic pathways in linking caffeine to information processing.

Second, a number of studies suggest that the impact caffeine exerts on the muscarinic-cholinergic system eventually targets visual processes. Monkey studies have shown that systematic injections of scopolamine, a muscarinic-cholinergic antagonist, impair the encoding of new visual objects but have little effect during recognition (Aigner & Mishkin, 1986; Aigner, Walker & Mishkin, 1991). In humans, the intake of caffeine attenuates the scopolamine-induced impairment of, among other things, perceptual sensitivity in visual search, visual short-term memory, and reading (Riedel, Hogervorst, Leboux, Verhey, van Praag & Jolles, 1995). At the same time, caffeine does not modulate the impact of scopolamine on simple and choice reaction time, suggesting that caffeine has a specific effect on visual processing via muscarinic-cholinergic pathways (cf., Smith, Brice, Nash, Rich & Nutt, 2003). Consistent with this interpretation, a recent PET study in humans using a visual task provides evidence that muscarinic-cholinergic effects modulate visual attribute processing (Mentis, Sunderland, Lai, Connolly, Krasuki, Levine, Friz, Sobti, Schapiro & Rapoport, 2001). In particular, muscarinic action was found to predominate in striate cortex (Brodmann Area 17) and lateral visual association areas (18 and 19), whereas nicotinic action predominated in the thalamus and inferior parietal regions (areas 39 and 40). In contrast to caffeine, nicotine is well known to induce cholinergic facilitation via nicotinic but not muscarinic receptors.

In view of this evidence, we assumed that, first, caffeine would act via muscarinic-cholinergic pathways and nicotine via nicotinic-cholinergic pathways and that, second, enhancing muscarinic but not nicotinic action would affect visual feature binding. Accordingly, we expected caffeine but not nicotine intake to affect the behavioral measure of feature integration we used, and that this effect is specific to visual binding.

We adopted the task from Hommel (1998), which involves the repetition of task-related and unrelated visual features and of the response (see Figure 1). The standard findings are interactions between (a) the task-related stimulus features (e.g., shape, if and only if it signals the response, and location, if and only if the responses are spatially defined); (b) the non-spatial stimulus features (e.g., shape and color); and (c) the relevant stimulus features and the response (for an overview, see Hommel, Müsseler, Aschersleben & Prinz, 2001a). These interactions all follow the same pattern: Performance is impaired in partial-repetition trials, that is, if one stimulus feature (or the response) is repeated while the other is not. This demonstrates that the mere co-occurrence of a feature-feature or feature-response conjunction is sufficient to create a temporary binding of the respective feature codes, a kind of "event file" (Hommel, 1998; Hommel & Colzato, 2004).




Figure 1. Sequence of events in the present experiments (cf., [Hommel, 1998]). A response cue signaled a left or right key press (R1) that was to be delayed until presentation of S1, a red or green, vertical or horizontal line in a top or bottom box. S2 appeared 1 sec later—another red or green, vertical or horizontal line in the top or bottom box. S2 shape signaled R2, also a speeded left or right key press. R2 speed and accuracy were analyzed as function of the repetition vs. alternation of stimulus shape, color, and location, and of the response.

Reactivating one member of this binding (as with feature repetition) spreads activation to other members, which calls for a time-consuming re-binding process in partial repetition trials. Most important for our purposes, these partial-repetition costs can be taken to indicate feature-feature and feature-response binding, which is why we chose them as behavioral markers. Along the lines described above we thus expected caffeine and nicotine consumption to decrease partial-repetition costs for shape-color conjunctions and for shape-location conjunctions (i.e., visual-visual bindings), but not for shape-response or location-response conjunctions.

Method

Eighteen volunteers took part in each study and served in two experimental sessions separated by 3-7 days. Informed consent was obtained from all participants after the nature and possible consequences of the study were explained to them; the protocol was approved by the local ethics committee (Leiden University, Faculty of Social and Behavioral Sciences). Subjects in the caffeine group (habitual coffee consumers: 3-5 units per day on average) and in the nicotine group were in the age range of 20-30, healthy non-smokers, not on medication or drugs, and without neurological or psychiatric history according to self-report. All experimental sessions were held in the morning to prevent time-of-day effects and the possible occurrence of withdrawal symptoms. Participants were asked to refrain from all caffeine containing foods and beverages for 12 hours prior to the experimental sessions, not to consume alcohol on the night before the experimental session and have a normal night rest. A double-blind, placebo-controlled, randomized cross-over design with counterbalancing of the order of conditions was used to avoid

expectancy effects. Treatments were deceptive: subjects were led to believe that they were drinking regular coffee or keeping a nicotine-patch each experimental session. Subjects' compliance was encouraged by taking a saliva sample (not further analyzed) at the beginning of each experimental session. Placebo and dose caffeine/nicotine quantities corresponded to 250mg lactose and 250mg caffeine (~3 cups of coffee) and 0mg and 7mg nicotine (~1 cigarette), respectively.


Subjects completed a version of the task adapted from Hommel (1998), (see Figure 1). They faced three gray, vertically arranged boxes in the middle of a monitor and carried out two responses per trial. R1 was a delayed simple reaction with the left or right key, as indicated by a 100%-valid response cue (three left- or right-pointing arrows in the middle box) that preceded the trigger stimulus S1 by 3,000 ms. S1 varied randomly in shape (a thin vertical or horizontal line), color (red or green), and location (top or bottom box). R1 was to be carried out as soon as S1 appeared, independent of its shape, color, or location; i.e., subjects were encouraged to respond to the mere onset of S1. R2 was a binary-choice reaction to the shape of S2 (vertical or horizontal orientation), which also appeared in red or green, and in the top or bottom box, 1,000 ms after S1 onset. Responses to S1 and to S2 were made by pressing the left or right shift-key of the computer keyboard with the corresponding index finger. Each session was composed of a factorial combination of the two possible shapes, colors, and locations of S2, the repetition vs. alternation of shape, color, location, and the response, and three replications per condition.

Results

After excluding trials with missing (< 1500 ms) or anticipatory responses (< 200 ms) mean reaction times (RTs) and proportions of errors for R2 (i.e., the response to S2) were analyzed as a function of treatment (placebo vs. treatment) and the repetition vs. alternation of response (R1→R2), stimulus shape, color, and location (S1→S2). Analyses of variance were performed by using a five-way design for repeated measures. Figure 2 shows the results from Caffeine and Nicotine groups comparing treatment with placebo conditions. We replicated the earlier finding (Hommel, 1998; Hommel & Colzato, 2004) of a shape-color and a shape-location interaction in both groups: repeating the shape of a stimulus but not its color or its location incurs a partial-repetition cost (see panels A and B). However, these two interactions were only affected by caffeine intake but not by nicotine. We also replicated the common interactions between response and shape, and between response and location in both groups (see C and D). Importantly, none of these effects was modified by treatment.

Conclusions

Our findings show that intake of caffeine, but not of nicotine, increases the binding of visual features, while both drugs spare cross-domain bindings between visual features and manual responses. This observation lends support to the hypothesis that feature binding in visual cortex is associated with neural synchronization enhanced by the muscarinic-cholinergic system (Rodriguez et al., 2000, 2001; Rodriguez-Bermudez et al., 2004). Whereas the apparently selective link between synchronization and muscarinic receptors is consistent with animal studies on both neocortical and hippocampal synchronization (e.g., Fellous & Sejnowski, 2000; Rodriguez-Bermudez et al., 2004), it remains to be determined why this link is selective for local visual feature binding. Interestingly, local integration processes within visual and motor cortex are commonly associated with synchronization frequencies in the gamma band (Engel & Singer, 2001), while interarea






Figure 2. Mean reaction time as a function of treatment and the repetition vs. alternation of stimulus shape and stimulus color (panels A and C, left), of stimulus shape and stimulus location (A and C, right), of response and stimulus shape (B and D, left) and of response and stimulus location (B and D, right). In each graph, white circles on the left and black circles on the right side represent partial-repetition conditions (one stimulus feature is repeated while the other feature, or the response, is not—or vice versa), which are expected to produce worse performance, i.e., partial-repetition costs. *Caffeine study* (A-B): Feature-feature bindings are indicated by interactions between shape and location and between shape and color, $ps < .001$, feature-response bindings by interactions between shape and response, location and response, and color and response, $ps < .001$. Importantly, caffeine intake increased shape-location and shape-color interactions, $ps < .05$, but not response-related interactions, $ps > .19$. Main effects were obtained for location, $p < .01$, and color, $p < .05$. Error rates (not shown) followed the same pattern: Reliable interactions were obtained for shape and location, $p < .05$, shape and response, $p < .005$, and location and response, $p < .003$, all showing the same partial-repetition-cost profile as reaction times. Apart from a location-repetition effect, $p < .05$, treatment produced a main effect, $p < .05$, but was not involved in any interaction. *Nicotine study* (C-D): Feature-feature bindings are indicated by interactions between shape and location, $p < .001$, and between shape and color, $p < .05$, feature-response bindings by interactions between shape and response, and location and response, $ps < .001$. Importantly, nicotine did not modify any of these interactions, $ps > .25$. Main effects were obtained for

location, $p < .001$, and shape, $p < .05$. Error rates followed the same pattern: Apart from a location-repetition effect, $p < .01$, reliable interactions were obtained for color and location, $p < .005$, shape and response and location and response, $ps < .001$, all showing the same partial-repetition-cost profile as reaction times.

synchronization, such as in visuomotor integration (Roelfsema et al., 1997) or visual target selection (Gross, Schmitz, Schnitzler, Kessler, Shapiro, Hommel, & Schnitzler, 2004), relies on lower frequencies in the beta band (~13-30 Hz). It thus seems possible that drugs like alcohol and caffeine exert (at least part of) their muscarinic-cholinergic influence on visual processing by increasing and decreasing the variability of firing rates, respectively (cf., Borchard, Singer & Munk, 2000); this again might affect higher synchronization frequencies more than lower frequencies (Kopell, Ermentrout, Whittington & Traub, 2000) and, hence, local, intra-area bindings more than long-range inter-area bindings.

Chapter 4

Priming and binding in and across perception and action: A correlational analysis of the internal structure of event files

Abstract

Individual performance was compared across three different tasks that tap into the binding of stimulus features in perception, the binding of action features in action planning, and the emergence of stimulus-response bindings (“event files”). Within a task correlations between the size of binding effects were found within visual perception (e.g., the strength of shape-location binding correlated positively with the strength of shape-color binding) but not between perception and action planning, suggesting different, domain-specific binding mechanisms. To some degree, binding strength was predicted by priming effects of the respective features, especially if these features varied on a dimension that matched the current attentional set.

Introduction

Perceiving a visual (but not only a visual) object involves registering, coding and processing its sensory features in numerous cortical areas (e.g., Cowey, 1985) and yet, what we perceive is a single, coherent event and not a bundle of attributes. This has been taken to suggest that there must be some kind of feature-binding mechanism that allows our brain to integrate the features belonging to the same event (see Treisman, 1996, for an overview). According to Kahneman, Treisman, and Gibbs (1992), attending to a visual object establishes what they call an “object file”, an integrated episodic trace containing information about the relationships between object features and their locations, possibly enriched by object-related knowledge from long-term memory. Indeed, a number of studies have provided evidence for the claim that perceiving a visual object involves creating an episodic assembly of object-related features codes (Gordon & Irwin, 1996; Henderson, 1994; Henderson & Anes, 1994; Hommel, 1998), even though the original approach seems to have over-estimated the importance of location for constructing and retrieving such object files (Hommel, 2002; Hommel & Colzato, 2004; Leslie & Kaldy, 2001; Leslie, Xu, Tremoulet, & Scholl, 1998; Pratt & Hommel, 2003).

Although the problem of integration in distributed representational systems has been discussed almost exclusively for the case of visual perception, the fact that the human cortex seems to represent almost all information in a distributed fashion suggests that “binding problems” exist and are apparently resolved in many representational and processing domains (Singer, 1994; Stoet & Hommel, 1999; Treisman, 1996; Wickens, Hyland, & Anson, 1994). In fact, a number of recent studies provide evidence of feature binding in action planning (Muesseler & Hommel, 1997; Stoet & Hommel,

1999, 2002) and across perception and action (Hommel, 1998; Hommel & Colzato, 2004). For instance, Stoet and Hommel (1999) showed that planning a speeded left-right keypressing action is delayed if it shares location-related features with (i.e., is carried out with an effector on the same side of the body as) another, already planned action held in memory. That is, planning an action might involve the integration or binding of those action features that specify the intended action, so that the integrated feature codes are temporarily not (or not that easily) available for the planning and control of other actions (Hommel, Muesseler, Aschersleben & Prinz, 2001a, 2001b).

Hommel and colleagues (Hommel, 1998; Hommel & Colzato, 2004; Hommel, Proctor & Vu, 2004) analyzed the combined effects of repeating versus alternating stimulus features and the response. The results showed that the repetition of stimulus features did not only interact with other stimulus-feature effects but also interacted with response repetition. With respect to stimulus-stimulus binding, repeating the shape of the stimulus produced better performance than alternation if the color or the location of the stimulus was also repeated, but worse performance than alternation if color or location alternated. This suggests that seeing an object results in the binding of its features; once bound together, these features can no longer be separately addressed, so that perceiving a new combination of the same features requires a time-consuming rebinding process and/or the resolution of the conflict induced by the previous binding. With respect to stimulus-response binding, repeating a stimulus feature (shape or position) facilitated performance only if the response was also repeated, otherwise stimulus repetition produced a cost. This implies that stimulus features are getting bound to the response they accompany, so that presenting the same stimulus feature again reactivates the associated response—thus creating a conflict in case of a response alternation.

We can conclude that feature-binding processes are not restricted to object perception, but cross borders between stimulus- and response-feature domains. This implies that the object-file concept introduced by Kahneman et al. (1992) is more general than anticipated, which led Hommel (1998) to suggest replacing it by the more universal concept of an “event file”. According to this idea, all the features belonging to an event, whether perceived, produced, or internally generated, might be integrated into episodic memory traces (cf., Logan, 1988). Question is, how do such traces look like? The perhaps most obvious possibility is that all the information a given event provides is lumped into one single master file, which would facilitate information exchange within a file. And yet, there is evidence that the structure of event files is more complex, more like a loosely connected network.

First, the studies looking into sequential effects of stimulus-response feature conjunctions provide ample evidence of binary interactions (e.g., shape X location, shape X response, location X response), which imply binary feature bindings, but hardly any higher-order interaction involving three stimulus and/or response features or more (Hommel, 1998; Hommel & Colzato, 2004; Hommel, Alonso & Fuentes, 2003). However, such higher-order interactions would be expected if all feature codes were integrated into the same file.

Second, interactions between stimulus features (e.g., shape and location) are affected by drugs that modulate the muscarinic-cholinergic

system (e.g., caffeine, an agonist, and alcohol, an antagonist), whereas stimulus-response bindings are unaffected by such drugs (Colzato, Erasmus & Hommel, 2004; Colzato, Fagioli, Erasmus & Hommel, 2005). The observed link between visual binding and cholinergic activity is consistent with the assumption that the integration of visual features is related to neural synchronization in the gamma frequency band (Engel & Singer, 2001; Keil, Muller, Ray, Gruber & Elbert, 1999), which is driven by the muscarinic-cholinergic system (Rodriguez, Kallenbach, Singer & Munk, 2000, 2001; Rodriguez-Bermudez, Kallenbach, Singer & Munk, 2004). If so, however, the observation that stimulus-response binding is not equally affected implies that these bindings are created by another neural mechanism, which presumably operates in the beta band (Kopell, Ermentrout, Whittington & Traub, 2000; see Roelfsema, Engel, Koenig & Singer, 1997).

And, third, the binding of perceptual features seems to be more automatic and to produce more stable traces than the binding of action features. For instance, the response-related costs of feature overlap between action plans that Stoet and Hommel (1999) observed were tightly linked to the planning process and disappeared a few hundred milliseconds after the planned action was carried out (Experiment 3). In contrast, stimulus-related overlap costs, as well as stimulus-response bindings, are largely unaffected by attentional manipulations and they last at least several seconds (Hommel & Colzato, 2004) if not minutes (Waszak, Hommel & Allport, 2003).

To summarize, feature integration seems to take place in and across perception and action, creating multi-modal episodic links between the codes that represent the features of the perceptual event and/or action plan at hand. However, a number of preliminary observations suggest that these integration processes do not create one single master file but, rather, a loosely connected associative network.

Purpose of the study

The present study was motivated by two questions, one concerned with the way event files are generated and the other with the internal structure of event files and the way different sub-components of an event file might be interrelated. To understand how we tackled the latter question, assume that we had evidence for the existence of a master event file, that is, let us assume that the codes of all available features of a given event were integrated into one coherent short-term memory structure. Let us further assume that people differ in the strength of feature integration, which suggests that there is a continuum ranging from “strong binders”, who create very strong temporary associations between the features they integrate, to “weak binders”, who create only weak associations. If so, we would expect interindividual variability in the sizes of binding effects, that is, of effects that are likely to reflect feature binding in and across perception and action. If all bindings would be created by the same binding mechanism, strong binders should show large binding effects whatever features are to be integrated, while weak binders should consistently show small effects. That is, the sizes of binding effects from different integration-requiring tasks should correlate positively: small (or large) binding effects in one task should go with small (or large) binding effects in the other. However, this should only be observed if all features are integrated by

the same mechanism—which in view of the available evidence is unlikely. Therefore, correlations between binding effects should occur only between those effects that were produced by the same mechanism. In other words, positive correlations between binding effects point to a common integration mechanism while the absence of correlations suggests different mechanisms.

Following this reasoning we had subjects carry out a number of tasks that all produce effects that can be assumed to reflect feature-integration processes. In particular, we used three tasks. The first was the “object-file task” (or “S-S task”, as we will call it here) modeled after Hommel and Colzato (2004), which is sensitive to sequential effects of conjunctions of stimulus features. This task measures after-effects of binding different features of the visual stimulus (here: shape, location, and color). It involves a prime stimulus (S1) followed by another stimulus (S2) that signals a speeded binary-choice response (R2). The standard outcome are cross-over interactions with repetition/alternation of different stimulus features (Hommel, 1998; Hommel & Colzato, 2004).

The second was the “event-file task” (or “S-R task”) introduced by Hommel and Colzato (2004). In addition to tapping into stimulus-stimulus integration (a partial replication of the S-S task), this task presumably taps into after-effects of binding stimulus features with the response. It involves a prime stimulus (S1) that triggers a precued prime response (R1), followed by another stimulus S2 that signals a speeded binary-choice response (R2). Apart from S-S interactions as described for the S-S task, the standard outcome in the S-R task are cross-over interactions between response repetition and the repetition of task-relevant stimulus features (Hommel, 1998; Hommel & Colzato, 2004).

Finally, our third task (“R-R task”) was modeled after McDevitt and Fournier’s (2001) adaptation of Stoet and Hommel’s (1999) “action-file” paradigm, a task that arguably measures side- and after-effects of binding action-related features. It involves preparing a cued response (RA), making a speeded response (RB) to a following stimulus (SB), and carrying out the prepared prime response (RA). Standard outcomes are slower RTs on RB if it feature-overlaps with RA (presumably indicating feature integration) and faster RTs on RA if it feature-overlaps with RB (presumably indicating feature priming; Stoet & Hommel, 1999).

Every subject ran through all three tasks, so that we were able to calculate individual estimates for all task-specific binding effects. On the one extreme all these measures might correlate, suggesting one single master event file or, on the other extreme, no two measures might correlate, which would point to numerous different integration mechanisms. Our expectations lay in between: some measures are likely to correlate, such as those related to different features of the same stimulus, while those related to different domains were more likely to be uncorrelated (Colzato et al., 2004).

Our other research question refers to the process of event-file construction. Even though event files are apparently created automatically, that is, irrespective of whether they are useful or necessary or not (Hommel & Colzato, 2004), their structure is often affected by the task goal. In particular, stimulus features that vary on a task-relevant dimension are more likely to be integrated (i.e., produce stronger and more reliable interaction effects) than stimulus features varying on an irrelevant dimension (Hommel, 1998).

Interestingly, there is some evidence that feature priming effects—i.e., main effects of repeating versus alternating a stimulus feature—follow the same pattern in being stronger for task-relevant than irrelevant feature dimensions (Hommel, 1998; Hommel & Colzato, 2004).

This commonality might indicate the criterion underlying feature integration: Codes of features defined on dimensions that are primed by the task context might receive a stronger activation than codes of features defined on unprimed dimensions (Hommel et al.'s, 2001a, 2001b, intentional weighting principle). If feature codes are integrated only if they pass a particular activation threshold (Hommel, 2003), this would mean that codes related to context-primed dimensions are more likely to be integrated than other codes—just as our findings suggest. Simple feature-repetition or priming effects may thus reflect the degree of context-induced dimensional priming (i.e., more strongly primed codes leave more stable or more slowly decaying traces). If so, and if the degree of context-induced priming varies between subjects, priming effects and integration effects may correlate in such a way, that pronounced priming effects of two given features are associated with a pronounced integration effect, that is, with a strong interaction involving these two features.

To summarize, we were interested in two types of correlational patterns: correlations between the sizes of different binding effects, which we take to point at a common integration mechanism, and correlations between priming effects (main effects of feature repetition) and binding effects involving the respective feature(s), which we take to point to the way codes are integrated.

Method

Participants

Forty-four students of the Leiden University took part for pay in three sessions. All reported having normal or corrected-to-normal vision. They were not familiar with the purpose of the experiment.

Apparatus and Stimuli

The experiment was controlled by a Targa Pentium III computer, attached to a Targa TM 1769-A 17" monitor.

In S-S and S-R sessions, which were modeled after Hommel and Colzato (2004), participants faced three grey square outlines, vertically arranged, as illustrated in Figure 1. From viewing distance of about 60 cm, each of these frames measured $2.6^\circ \times 3.1^\circ$. A thin vertical line ($0.1^\circ \times 0.6^\circ$) and a some what thicker horizontal line ($0.3^\circ \times 0.1^\circ$) served as S1 and S2 alternatives, which were presented in red or green in the top or bottom frame. Response cues (in the S-R session only) were presented in the middle frame (see Figure 1), with a left- or right-pointing arrow indicating a left and right keypress, respectively. Responses to S1 (in the S-R session only) and to S2 were made by pressing the left or right shift-key of the computer-keyboard with the corresponding index finger.

In the R-R session modeled after McDevitt and Fournier (2001) and Stoet and Hommel (1999), illustrated in Figure 2, a white cross on black background that appeared at the center of the monitor served as the first

fixation mark. SA consisted of a white arrowhead appearing above the cross and an asterisk appearing above or below the arrowhead. From a viewing distance of about 60 cm., each character constituting SA (i.e., each asterisk and the arrowhead), measured about 0.3° in width and 0.4° in height. A white cross was used again as second fixation mark, which also appeared at screen center. Below the cross appeared, then, the symbol # or & to signal SB.




Figure1. Overview of the display and the timing of events in the S-S and the S-R task.




Figure 2. Overview of the display and the timing of events in the R-R task and display of the numbers of the computer keys used as response key.

Procedure and Design

S-S Task. In the S-S session participants carried out two responses per trial. First, they made a binary-choice reaction to the second of two successive stimuli. Half of the participants responded to the vertical and the horizontal line

by pressing the left and right key, respectively, while the other half received the opposite mapping. After each binary-choice reaction to S2, participants were probed for their memory of a feature of S1 (=R3)—a manipulation that encourages the processing of all features of S1 and that produces pronounced and stable effects (Hommel & Colzato, 2004). Participants were presented with one of three questions, asking for the shape, color, or location of S1 (e.g., "What was the color of Stimulus 1?", in Dutch). Two words indicating the two response alternatives ("horizontal-vertical", "red-green", or "top-bottom", in Dutch) were presented below the question, with their relative position indicating the mapping of alternatives onto the left and right shift key. The six combinations of the three stimulus dimensions and two alternative-key mappings were presented in pseudo-random sequence but equally often within one session. The sequence of events is shown in the upper row of Figure 1. The intertrial interval of 2000 ms was followed by a 500-ms appearance of S1. The duration of the next, blank interval was 2000 ms. Then S2 appeared and stayed until the response was given or 2000 ms had passed. If the response was incorrect auditory feedback was presented. This sequence of events was, then, followed by the memory-probe question, which stayed until the response was given or 4000 ms had passed. The S-S session comprised 256 trials, composed by a factorial combination of the two shapes (vertical vs. horizontal line), colors (red vs. green), and locations (top vs. bottom) of S2, and the repetition vs. alternation of shape, color, and location ($2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 4 = 256$). Participants were allowed to take a short break during each session.

S-R Task. The procedure in the S-R session was as in the S-S session, with the following exceptions. Participants carried out three responses per trial. R1 was a simple reaction with the left or right key, as indicated by the response cue. It had to be carried out as soon as S1 appeared, independent of its shape, color, or location. Participants were informed that there would be no systematic relationship between S1 and R1, or between S1 and S2, and they were encouraged to respond to the onset of S1. As in the S-S session, R2 was a binary-choice reaction to the shape of S2 and R3 required the identification of a randomly selected feature of S1. The sequence of events in each trial is shown in the lower row of Figure 1. Next to the intertrial of 2000 ms a response cue signaled R1 for 1500 ms, followed by a blank interval of 1000 ms. Then S1 appeared for 500 ms, followed by a further blank interval of 2000 ms. If R1 was incorrect or not given within 500 ms the trial started again. After the SOA, S2 appeared and stayed until R2 was given or 2000 ms had passed. The session comprised 384 trials, composed by a factorial combination of the two shapes (vertical vs. horizontal line), colors (red vs. green), and locations (top vs. bottom) of S2, the repetition vs. alternation of shape, color, location, and of the response ($2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 3 = 384$).

R-R Task. In the R-R session, the intertrial interval of 1500 ms was followed by the sequence of events shown in Figure 2. Subjects had to perform two tasks in each trial, Task A and Task B, and the second task was embedded in the first task. Stimulus A always appeared before Stimulus B, but the corresponding response had to be executed only after the response to Stimulus B was performed (ABBA design). This implied that the participants were forced to memorize Response A while Task B was performed. As shown

in Figure 2, numbers 1, 4 (home key) and 7 of the computer keyboard served as response keys for the left hand while numbers 3, 6 (home key), and 9 served as response keys for the right hand.

Task A was signaled by a left- or right-pointing arrowhead, accompanied by the symbol “*” above or below it. Arrowhead direction indicated whether Response A was to be performed with the index finger of the left or right hand. The asterisk indicated the direction of the response. When the asterisk appeared above, participants pressed the numeric key above the home key and then pressed again the home key. If the asterisk appeared below, participants pressed the numeric key below the home key and then pressed again the home key. The movement always started with pressing the home key with the hand specified by the arrowhead. For example, if the arrowhead pointed to the right and the asterisk was above, subjects had to use the index finger of the right hand by pressing 6 (home key), then 9 (numeric key above the home key) and then again 6 (back to the home key). Response B was signaled by presentation of the symbol “#” or “&”, which required a binary choice reaction to the shape of the symbol by pressing one of the home keys.

Figure 2 shows that after a 50 ms fixation cross, Stimulus A was presented with the fixation mark for 2 seconds. Following a 50 ms fixation cross, Stimulus B appeared with it for 200 ms. Stimulus B was to be responded to immediately by performing Response B, followed by the already planned Response A.

There were eight conditions resulting from the orthogonal variation of three within-participant factors with two levels each: side of RA (left or right), side of RB (left or right), and direction of first movement of RA (back or forth). Participants worked through a practice block of 40 trials (8 conditions x 5 replications) and an experimental block of 256 trials (8 conditions x 32 replications). The possible mappings of RB (binary choice reaction to the shape of the symbol # or &) were counterbalanced across participants.

We used a randomized cross-over design with counterbalancing of the order of sessions; that is, one third of the subjects began with the S-S session, one third with the S-R session, and one third with the R-R session.

Results and Discussion

Analytical procedures

To facilitate access to the rather complex data pattern we sort, present, and discuss the outcomes according to their theoretical implications. First, we present the data separately for each task, with particular emphasis on priming effects (main effects of the repetition or alternation of a single stimulus feature or the response) and binding effects (interactions between effects of repeating or alternating one stimulus-feature and the effect of repeating or alternating another stimulus feature or the response). Second, we present the correlations between the individually computed sizes of priming and integration RT effects from all three tasks. Table 6 summarizes most of the relevant findings: priming and binary binding RT effects (significant effects underlined) and their correlations (significant effects marked by asterisks).

Of the data from S-S and S-R sessions trials with missing or anticipatory responses (1.2% and 1.6%, respectively) were excluded from the

analysis. We also excluded trials in which the memory-probe response was incorrect. From the remaining data, mean RTs and proportions of errors (PEs) for R2 (i.e., the response to S2) were further analyzed. From S-S data, means and error rates were computed as a function of the three possible relationships between the two stimuli in each trial, that is, repetition vs. alternation of stimulus shape, color, or location (see Table 1 for means). ANOVAs were performed by using a three-way design for repeated measures. From S-R data, means and error rates were computed as a function of the four possible relationships between the two responses (R1 and R2) and the two stimuli in each trial, that is, repetition vs. alternation of response, stimulus shape, color, or location (see Table 3 for means). ANOVAs were performed by using a four-way design for repeated measures.

In the data from the R-R session, RTs refer to error-free trials only. For Response B, the first to-be-emitted reaction, RT was measured from the onset of Stimulus B to the pressing of the home key. For Response A, the second reaction, interresponse times (IRT) were measured from the release of the key for Response B, hence the first release of the home key, to the second release of the home key. Error rates for Response A refer to all incorrect responses following a correct Response B. All measures were analyzed with ANOVAs for repeated measures and the significance criterion was set to $p < .05$.

Results

S-S Task

Tables 1 and 2 provide an overview of the means and ANOVA outcomes for RTs and PEs obtained for R2. In the RTs analysis of R2 we found two main effects: one involving shape and the other location. For both stimulus features, repetition produces a cost. In the case of shape, this may be due to a strategic expectation bias towards stimulus (or response) alternation (Kornblum, 1973; Soetens, Boer, & Hueting, 1985). The benefit of location alternation reflects Inhibition of Return (IOR), the common observation that attending to an irrelevant stimulus impairs later responses to relevant stimuli appearing in the same location (Posner & Cohen, 1984).

Repeated	RT _{R2}	PE _{R2}
Neither	560	3.00
S(hape)	610	3.58
L(ocation)	598	3.19
C(olor)	572	3.10
SL	616	3.97
SC	597	4.10
LC	605	3.77
SLC	606	3.49

Table 1

S-S Task: Means of Mean Reaction Times for Responses to Stimulus 2 (RT_{R2}; in Ms) and Percentages of Errors on R2 (PE_{R2}), as a Function of the Feature Match Between Stimulus 1 and Stimulus 2.

Effect	df	RT _{R2}		PE _{R2}	
		MSE	F	MSE	F
Shape (S)	1,43	1,994.33	24.2**	14.40	1.54
Color (C)	1,43	1,025.75	0.08	10.94	0.23
Location (L)	1,43	1,582.66	25.6**	25.82	0.10
S x C	1,43	1,436.69	6.82*	12.18	0.20
S x L	1,43	1,056.05	17.3**	11.17	0.41
C x L	1,43	1,070.86	0.03	14.98	0.11
S x C x L	1,43	888.20	0.39	11.35	0.68

Note. * $p < .05$ ** $p < .01$

Table 2

Results of Analysis of Variance on Mean Reaction Time of Correct Responses (RT) and Percentage of Errors (PE) for S-S Task.




Figure 3

Reaction times in the S-S task, as a function of the repetition vs. alternation of stimulus shape and stimulus color (left panel) and of stimulus shape and stimulus location (right panel).

We obtained two interactions between stimulus-feature repetition effects: Shape, the most relevant stimulus feature, interacted with location and with color. The pattern of these interactions is shown in Figure 3: Taking into account the location main effect, we can see that performance was better if both of the respective features repeated or if they both alternated than if one repeated while the other did not.

The PE analyses did not yield significant effects.

Discussion. We were able to replicate the main findings of Hommel (1998) and Hommel and Colzato (2004): The impact of repeating a stimulus feature depended on whether or not other stimulus features repeated as well. That is, if only one but not the other feature overlaps (partial match), re-activating the code of the matching feature may spread activation to the code it has just been integrated with, thus impairing its integration with the actual

feature. We take these results to indicate the integration of feature codes in visual perception. Note that these integration effects were accompanied by significant priming (i.e., main) effects for shape and location, which were both negative.

S-R Task

Tables 3 and 4 provide an overview of the means and ANOVA outcomes for RTs and PEs obtained for R2. As in the S-S task we found a main (priming) effect for shape and location. There was also a main effect of the response. All three effects were due to better performance if the respective feature alternated than if it was repeated.

Response	Repeated		Alternated	
	RT _{R2}	PE _{R2}	RT _{R2}	PE _{R2}
Neither	591	9.75	519	1.42
L(ocation)	590	6.62	557	2.65
S(hape)	578	6.63	578	3.69
C(olor)	582	7.95	520	2.46
SL	560	2.74	599	8.61
SC	558	3.41	575	4.83
LC	585	8.61	566	5.21
SLC	543	1.13	600	13.8

Table 3

S-R Task: Means of Mean Reaction Times for Responses to Stimulus 2 (RT; in Ms) and Percentages of Errors on R2 (PE), as a Function of the Match Between Response 1 and Response 2, and the Feature Match Between Stimulus 1 and Stimulus 2.

Shape and location produced a two-way interaction of the same sort as in the S-S task, whereas the corresponding interaction of shape and color missed the significance criterion, $p < .08$. In addition, response repetition interacted with the repetition of all three stimulus features. Taking into account the significant negative priming effects, we can see in Figure 4 that performance was better if both response and respective stimulus feature repeated or alternated than if one repeated while the other did not.

The errors followed a similar pattern: Apart from main effects of location and color, significant interactions were obtained between location and color, and response repetition interacted with each of the three stimulus features. In addition, we found response being involved in two higher-order interactions with shape and color and with shape and location. Separate ANOVAs showed that color and shape interacted significantly if the response repeated, $p < .018$, but not if it alternated, $p < .272$, whereas shape and location interacted if the response alternated, $p < .001$, but not if it was repeated, $p < .07$.

Discussion. Again, we replicated the main findings of Hommel (1998) and of Hommel and Colzato (2004): The impact of repeating a stimulus feature depended on whether or not other stimulus features or the response repeated as well, suggesting that stimulus and response features were integrated.

Interestingly, most integration-related effects were restricted to binary interactions, a common observation in integration studies. Again, the integration effects were accompanied by significant priming (i.e., main) effects

Effect	df	RT _{R2}		PE _{R2}	
		MSE	F	MSE	F
Shape (S)	1,43	2,086.57	8.39**	28.71	0.04
Color (C)	1,43	1,330.08	3.69	14.10	5.49*
Location (L)	1,43	1,950.56	14.5**	33.60	7.05*
Response (R)	1,43	1,703.54	8.52**	44.86	1.06
S x C	1,43	987.78	3.20	16.07	0.88
S x L	1,43	1,079.10	13.3**	27.22	3.95
C x L	1,43	828.99	1.00	13.02	25.5**
S x L x C	1,43	1,049.06	0.13	21.84	0.02
S x R	1,43	2,284.95	106**	93.17	43.2**
L x R	1,43	1,017.26	70.0**	49.05	39.4**
C x R	1,43	1,067.59	8.63**	33.78	17.3**
S x L x R	1,43	748.34	0.01	25.30	20.2**
C x L x R	1,43	1,451.77	0.04	20.74	0.01
S x C x R	1,43	804.11	0.54	40.19	4.13*
S x L x C x R	1,43	1,203.28	0.05	18.59	3.32

Note. * $p < .05$ ** $p < .01$

Table 4

Results of Analysis of Variance on Mean Reaction Time of Correct Responses (RT) and Percentage of Errors (PE) for S-R Task.

for shape and location, as well as for the response, and again all priming effects were negative.

R-R Task

Table 5 provides an overview of the RTs, IRTs and PEs. The RT analysis for RB yielded a significant effect of overlap, $F(1,43) = 16.50$, $p < .001$, indicating that the latency of RB was longer with RB-RA feature overlap than with no overlap (612 vs. 587 ms). Likewise, overlap yielded more errors than non-overlap (4.9% vs. 3.1%), $F(1,43) = 15.45$, $p < .001$. The IRTs for RA were numerically faster for overlap than non-overlap (277 vs. 282 ms) but this effect did not reach significance.

Discussion. According to Stoet and Hommel (1999, 2002), integrated action-feature codes should be less available for other planning activities as long as the current plan is not executed or abandoned. The results indeed confirm that feature overlap between a planned and a to-be-performed action negatively affects the latter, whereas a previous, already executed action plan primes a following, feature-overlapping action (although this latter effect was unreliable here). We take these results to indicate the integration of feature codes in action planning.




Figure 4 Reaction times in the S-R task, as a function of the repetition vs. alternation of response and stimulus color (left panel), of response and stimulus shape (middle panel), and of response and stimulus location (right panel).

	Response B	Response A	
	RT	IRT	PE
Overlap	612	277	4.88
No Overlap	587	282	3.09

Table 5

R-R Task: Means of Mean Reaction Times (RTs), Interresponse Times (IRTs), and Proportion of Errors (PEs) as a function of Feature Overlap Between Response A and B.

Correlations

Having replicated all theoretically relevant effects in the three tasks, we ran an analysis using Pearson correlation coefficients on the individually calculated priming-effect and binding-effect sizes. The results are shown in Table 6.

Priming X Priming Effects. These correlations follow a consistent pattern that shows three characteristics. First, color and response repetition effects do not correlate with any other repetition effect, not even with corresponding effects in other tasks. Second, shape and location repetition effects correlate across experiments with their equivalents, that is, the relative sizes of shape-repetition and location-repetition effects in the S-S task correspond to those in the S-R task. Third, shape repetition correlates with location repetition, but only in the S-S task.

Binding X Binding Effects. There were only three reliable correlations, and all three relate within-task bindings to each other: In the S-S task, shape-location binding correlates with shape-color binding—hence, the individual sizes of the only two reliable binding effects covary. In the S-R task, shape-location binding correlates with color-location binding, which again correlates with shape-color binding. The relation between these correlations is, however, not transitive and the individual sizes of the only reliable binding effect covaries with only one (color-location) of the two (here both unreliable) within-object bindings. A clear-cut outcome is that binding effects in the S-S task do not correlate with binding effects in the S-R task.

Priming X Binding Effects. These effects fall into five clusters: First, the shape-repetition effect was correlated with both bindings involving shape in the S-S task but not in the S-R task. Second, the location-repetition effect was correlated with both location-shape and location-color bindings in the S-R task but not in the S-S task. Third, the response-repetition effect correlated with the shape-location binding in the S-R task. Fourth, the shape-location binding in the S-S task correlated with the shape-repetition effect in the S-R task. Finally, the color-repetition effect in the S-S task correlated with the color-response binding in the S-R task—note that the correlation between the latter and the color-repetition effect in the S-R task was also relatively high (.23).

Task	Effect category	S-S Task					S-R Task					R-R Task				
		Priming					Priming					Integration				
		C	L	S	CxL	SxL	C	L	S	CxL	SxL	R	S	CxR	LxR	SxR
S-S	Priming	C(olor)														
		L(ocation)														
		S(hape)														
	Integration	CxL	.23	.28	.20											
		SxL	.05	.42**	-.02											
S-R	Priming	SxC	.26	.44**	-.04	.32*										
		C	.15	-.10	-.08	.14	.18									
		L	.20	.42**	-.01	-.23	-.05	-.08								
		S	.22	.20	.45**	.08	.44**	.20	.25	.17						
		R(response)	-.08	.23	.04	.07	-.12	-.01	-.09	.23	.16					
	Integration	CxL	.11	.13	-.02	-.01	-.12	.08	.30*	.20	.12					
		SxL	-.02	.22	.16	-.13	.08	.16	-.12	.53**	.13	.31*	.42**			
		SxC	-.14	-.19	-.13	-.07	-.17	-.02	.19	-.20	-.03	.23	.31*	.09		
		CxR	.40**	-.23	-.17	-.08	.02	.05	.23	.06	.09	-.01	-.05	.02		
		LxR	.06	-.13	.06	.09	.13	.27	.16	.04	.06	-.15	.06	-.24	-.10	.09
R-R	Priming	SxR	.04	-.02	.21	.08	.03	-.13	.04	-.25	.13	-.16	.10	-.19	-.20	.15
	Integration	R _A	.04	-.14	.09	.01	.27	-.06	.06	.11	-.04	.10	-.12	.04	-.02	.06
		R _B	-.13	.03	.15	-.03	-.03	-.01	-.30	.03	.17	.10	-.05	.10	-.12	-.10

Note. * $p < .05$ ** $p < .01$

Table 6
Correlations among the priming and integration effects for the three tasks. Reliable main effects and two-way interactions are underlined.

Discussion

Our study sought for correlations between measures of feature-priming and feature-binding effects within and across three different but related tasks. In particular, we looked for two types of correlations: those between the sizes of different measures of bindings—which might point to a common integration mechanism—and those between priming effects on the one hand and binding effects on the other—which might reflect that the degree to which a feature that is contextually primed predicts the likelihood that it is integrated.

With regard to correlations between different binding measures the outcome is clear-cut: On the one hand, there is considerable evidence that the processes responsible for binding different features of a given stimulus within a task are related. Within the S-S and within the S-R task shape-location bindings were the most reliable, and they interacted with at least one other binding effect. Binding effects of the S-S task do, however, not correlate with binding effects of the S-R task. It therefore seems that, to some degree, feature binding is task specific. The only difference between the S-S and S-R task is the simple reaction (R1) that has to be carried out in the latter task. Why and how feature binding is dependent on task and task instructions is an open issue. We suppose that what differs in the two tasks is the internal representation of the context. It cannot be excluded that the information held in short-term memory for mediating the appropriate behavioral response of R1 in the S-R task determines the lack of correlations between the binding effects in the two tasks. This idea fits well with the suppositions of Cohen and Servan-Schreiber (1992) about the close relationship between selective attention and the internal representation of context and about the influence of context on the selection of the appropriate response.

The shape-color interaction was reliable in the S-S task but not in the S-R task, and it correlated with shape-location binding in the S-S task but with color-location binding in the S-R task. In contrast to these within-object correlations there was no evidence of any correlation crossing borders between perception and action. Given the equivocal status of null effects this lack of evidence should not be taken as a strong proof that such cross-border relations do not exist. And yet, the absence of such effects in the presence of strong within-object correlations is consistent with the hypothesis that within-object bindings are created by a different mechanism than bindings between stimulus and response features.

Converging support for this hypothesis comes from the already mentioned drug studies of Colzato and colleagues (2004, 2005), which employed a version of the present S-R task. Starting from the assumption that feature binding is mediated by neural synchronization, which in the visual cortex seems to be driven by the muscarinic-cholinergic system (Munk, 2003; Rodriguez et al., 2000, 2001, Rodriguez-Bermudez et al., 2004), Colzato et al. found that within-object bindings were significantly decreased by alcohol, a muscarinic-cholinergic antagonist (Colzato et al., 2004), significantly increased by caffeine, a muscarinic-cholinergic agonist, and unaffected by nicotine, which only affects the nicotinic-cholinergic system (Colzato et al., 2005). Most interestingly, none of these manipulations had any impact on bindings of stimulus and response features. Thus, if we assume that visual-

visual binding is mediated by cholinergically-driven neural synchronization in the gamma frequency band (Engel & Singer, 2001; Keil et al., 1999), stimulus-response bindings seem to be created in a different fashion—either in terms of the driving system or in terms of the synchronization frequency used. Indeed, there are a number of indications that longer-range bindings are associated more with the beta frequency band (Gross, Schmitz, Schnitzler, Kessler, Shapiro, Hommel, & Schnitzler, 2004; Kopell et al., 2000; Roelfsema et al., 1997).

Our second question was whether the degree to which a feature is contextually primed predicts the likelihood that it is integrated. If so, we would expect reliable correlations between priming effects and integration effects. The first thing to note is that all the reliable priming effects we obtained were negative (alternation was faster than repetition), whereas all reliable correlations between priming and integration effects were positive. The observation of negative priming effects as such is not uncommon, especially with interresponse times as long as the two seconds used in the S-S and S-R tasks (Kornblum, 1973; Soetens et al., 1985). They are likely to reflect the gambler's fallacy—the expectation that events are more likely to alternate than to repeat. However, the finding of a positive correlation with integration effects (e.g., the shape-location and shape-color interactions in the S-S task), means that integration was more pronounced the smaller the bias towards alternation. In other words, preparing for an expected alternation of stimulus or response works against integrating the expected stimulus or response feature into a more complex event representation. Indeed, positive repetition effects—which reflect the speedup of processing by left-over traces that we assume to support integration—and negative repetition effects have been argued to indicate antagonistic processes. It is generally supposed that the former are more automatic and the latter more strategic (Soetens et al., 1985). Along these lines we can conclude that individual variability in the more automatic component of the repetition effect is the more reliable predictor of integration. This outcome fits nicely with our hypothesis that the degree to which a feature is contextually primed predicts the likelihood that it is integrated.

An interesting observation is the double-dissociation of the impact of shape and location priming on integration: the shape-repetition effect predicted (statistically) shape-related bindings in the S-S task only, whereas the location-repetition effect predicted location-related bindings in the S-R task only. Given that the only difference between the S-S and the S-R tasks was the presence or absence of R1, the response to the prime, we consider this dissociation to reflect attentional task requirements. In the S-S task, subjects do not have anything to do before the appearance of S2, which suggests that the attentional set is optimized to process the relevant feature of S2—shape. To the degree that this set is optimized and maintained—a factor with respect to which people are known to differ (Duncan, Emslie, Williams, Johnson & Freer, 1996; Miyake, Friedman, Emerson, Witzki, Howerter & Wager, 2000)—processing stimulus shape will be facilitated and, thus, produce a stronger (positive) priming effect. The stronger this effect the more activated the respective shape code must be, which increases the likelihood to become integrated with other sufficiently activated feature codes. However, in the S-R task people are not able to prepare for S2 processing before having detected

S1 and having carried out the prepared R1. Accordingly, their attentional set should be related to the dimension that is relevant for this part of the task, which given the use of spatial responses must be location (see Hommel, 1996, for evidence that prepared responses are under spatial control until executed). If so, it will be the processing of location but not of shape of S1 that is facilitated, and the integration of location- but not shape-related bindings that is supported. This accounts for the correlation between location priming and the shape-location and color-location bindings and, given the spatial nature of the response, for the correlation between response repetition and the shape-location binding as well.

The only remaining observation relates to the correlation between the color-repetition effect in the S-S task and the color-response binding in the S-R task. Even though this particular cross-task effect was not expected, a closer look provides some suggestions as to why it may have occurred in this pattern. Comparing the effect sizes for color priming in the S-S and S-R task reveals a substantial degree of interindividual variability in both tasks but more in the former (range from -60 ms to +44 ms; mean = 1 ms) than in the latter (range from -46 ms to +41 ms; mean = 5 ms). This suggests that the S-S task might have been more sensitive in picking up interindividual differences in priming effects, so that the color-priming effects in the S-S task were more likely to produce a high correlation. The large ranges further show that mean priming effects are not particularly useful indicators of the degree to which stimulus color is processed—as far as priming effects do represent this degree. In any event, if the measure used is sufficiently sensitive, it can predict the strength of color-response bindings.

It is an interesting question why other stimulus-response bindings were not as well predicted by the respective stimulus-priming effects, i.e., shape-response binding by shape repetition or location-response binding by location repetition. Lack of variability cannot account for this difference, as repetition-effect sizes varied considerably for both stimulus shape (ranges from -91 to +102 ms and from -74 to +40 in the S-S and the S-R task, respectively) and location (from -82 to +68 ms and from -81 to +37 ms). However, it might be that the task requirements put so much more emphasis on shape (relevant S2 dimension) and location (relevant response dimension) that these dimensions were primed to a degree that individual variability no longer mattered. That is, shape and location codes might have been sufficiently activated for integration even in “weak primers”. Clearly, this issue needs to be investigated more closely.

In sum, then, the present study provides converging evidence that feature integration is a general phenomenon that however is accomplished by domain-specific mechanisms. In particular, binding effects within the same domain—visual feature integration in our case—tend to correlate with each other, suggesting that people differ in the degree to which they bind visual features and/or with respect to the “stickiness” these bindings exhibit over time. To some degree, the strength of a binding can be predicted based on the degree to which the respective features are activated—as indicated by priming effects. The most reliable predictors in this respect seem to be features varying on a dimension that matches the currently implemented attentional set.

Chapter 5

What do we learn from binding features? Evidence for multilevel feature integration

Abstract

Four experiments were conducted to investigate the relationship between the binding of visual features (as measured by their after-effects on subsequent binding) and the learning of feature-conjunction probabilities. Both binding and learning effects were obtained but they did not interact. Interestingly, (shape-color) binding effects disappeared with increasing practice, presumably due to the fact that only one of the features involved was task-relevant. However, this instability was only observed for arbitrary combinations of simple geometric features but not for real objects (colored pictures of a banana and strawberry), where binding effects were strong and practice-resistant. Findings are interpreted in a neurocognitive framework that makes a distinction between integration at low-level feature maps, short-term acquisition of frequency-based expectations, and Hebbian learning of object representations.

Introduction

Considerable evidence suggests that cortical networks encode the external environment in a distributed fashion. A striking example of spatially-distributed coding in cortical information processing is given by the primate visual cortex, where processing visual event features occurs in parallel in numerous cortical maps (Cowey, 1985; Felleman & van Essen, 1991). This coding scheme also applies to events in the auditory and other sensory modalities, and to multimodal event processing. Distributed coding creates numerous so-called “binding problems”, that is, difficulties in relating the codes of a given entity or processing unit (e.g., visual object) to each other. To solve these problems, the brain needs some sort of integration mechanism that binds together the distributed codes belonging to the same event, while keeping these codes separated from codes for other events (Treisman, 1996).

Mechanisms of Feature Integration

At a neural level, a possible solution to the binding problem may be given by high-order cardinal cells (Barlow, 1972), onto which signals from neurons coding for the to-be-bound features converge. However, given the high variability of objects belonging to a given category in terms of their instances and retinal projections, as well as the numerous ways in which discrete features can be potentially combined, the exclusive reliance on such a convergent mechanism would ultimately lead to a combinatorial explosion and is therefore not plausible.

Another potential solution to the binding problem is given by cell (neural) assemblies or sets of tightly-connected neurons. The identity of an assembly is defined in terms of higher firing rates or coactivation of the participating neurons (Amit, 1995; Braitenberg & Schuz, 1991; Hebb, 1949). In a representational scheme based on cell assemblies, individual neurons encode for simple features, and the associative connections between these neurons enable pattern encoding and completion within the assembly. This solution avoids the combinatorial explosion problem implied by cardinal cells. However, it does not enable discriminative responses in the presence of multiple objects within a given visual scene, if neurons coding for the features of these objects are in the same area or project onto neurons in the same area with large receptive fields (like the inferotemporal cortex). Given that neurons in the same area coding for different objects would all be characterised by an elevation of firing rate, the segregation between neural representations of objects would become impossible due to a ‘superimposition catastrophe’ (e.g., Engel, König, Kreiter, Schillen & Singer 1992; von der Malsburg, 1981, 1999).

This problem of ambiguous allocation of the responses of neural assemblies can be theoretically solved by a spatial spotlight or attentional window (Treisman, 1988; Treisman & Gelade, 1980), restricting processing to one object location and therefore one neural assembly at the time. However, there is evidence that attentional selection often takes place in terms of objects rather than locations (Duncan, 1984). Moreover, it is currently unclear how the attentional spotlight would be implemented in terms of neural mechanisms, and whether an attention-based binding of features into separate objects would be possible or guided in the absence of a prior definition of the objects in terms of neural responses.

Alternatively, it has been suggested that the perceptual binding problem may be solved at a neural level in terms of temporal coding, based on the selective synchronisation of time-resolved neuronal responses (Eckhorn, Bauer, Jordan, Brosch, Kruse, Munk & Reitboeck, 1988; Engel et al., 1992; Gray, König, Engel & Singer, 1989; von der Malsburg, 1981, 1999). According to this view, the action potentials of neurons coding the features of the same object are synchronised, while being uncorrelated to the responses of neurons coding for the features of other objects. This view is supported by the evidence that neurons act as coincidence detectors, since synchronous synaptic inputs are more effective than asynchronous ones in eliciting spikes of the neurons on which they converge (Abeles, 1991). Temporal neural codes may be readout in terms of these coincidence detection properties. Many recording studies from the visual cortex of cats and monkeys have shown that the selective synchronisation of oscillatory neuronal discharges may plausibly be involved in visual grouping and segregation (e.g., Castelo-Branco, Goebel, Neuenschwander & Singer, 2000; Eckhorn et al., 1988; Gray et al., 1989), as well as in binding sensory and motor responses (Murthy & Fetz, 1992; Roelfsema, Engel, König & Singer, 1997).

At a behavioral level, one way to study feature binding mechanisms is to put processing systems under conditions that render proper integration difficult or impossible, and then to look for the creation of incorrect bindings or “illusory conjunctions” (Treisman & Schmidt, 1982). Another way is to search for after-effects of feature integration, that is, for effects of previously created

feature bindings on performance in a subsequent task that requires a different binding. In a seminal study along these lines, Kahneman, Treisman, and Gibbs (1992) presented participants with two displays in a sequence, a brief multi-letter prime display (S1) followed by a single-letter probe display (S2) requiring verbal identification. Having just seen the probe letter somewhere in the prime display tended to facilitate probe identification. However, more reliable than this nonspecific repetition effect was the benefit of repeating the particular combination of letter and location, a finding that since then has been replicated many times in both the visual (e.g., Gordon & Irwin, 1996; Henderson, 1994; Henderson & Anes, 1994; Hommel, 1998; Park & Kanwisher, 1994) and the auditory modality (Mondor, Hurlburt & Thorne, 2003). Interestingly, binding after-effects are not restricted to stimulus shape and location but can also be found for other feature combinations, such as shape and color (Hommel, 1998; Hommel & Colzato, 2004)—the features the present study will focus on.

Apparently, perceiving an event automatically creates a kind of "object file" (Kahneman et al., 1992) or "event file" (Hommel, 1998), an integrated episodic trace containing information about the various features and bindings of that event. Re-viewing parts or aspects of that event automatically retrieves the file, which produces a benefit if previous and present event perfectly match (Kahneman et al., 1992) or confusion if they mismatch (Hommel, 1998, 2004). Indeed, incomplete repetitions (e.g., shape match combined with color mismatch) commonly produce worse performance than conditions in which no stimulus feature is repeated, whereas the latter yield performance comparable to complete repetitions (Hommel, 1998; Hommel & Colzato, 2004). That is, re-using an already created object file seems to be of little help but retrieving an old file that also includes mismatching codes apparently causes conflict (Hommel, 2004).

Taken altogether, the available evidence strongly suggests that perceiving an event results in the integration of its features, that is, in the binding of the individual codes representing them. Once bound, the feature codes can no longer be selectively addressed, so that perceiving some combination of the same features retrieves the whole file, a kind of pattern completion process. Feature binding is supposed to be a fast-acting process (simple bindings emerge after 300 ms or less: Hommel & Colzato, 2004) that creates transient representational structures. In the present study we asked how this process might be related to the learning of feature combinations—that is, the creation of relatively permanent memory changes.

Binding and Learning in Cortical Networks

On the one hand, one may consider conjunction learning being a direct consequence of binding—we call this the strong-dependence hypothesis. As suggested by Fell, Fernandez, Klaver, Elger and Fries (2003), synchronized neural activity may cause Hebbian learning (neurons that fire together, wire together), that is, learning through the long-term modification of synaptic efficacy induced by reverberation of neural activity in cerebral circuits (Hebb, 1949). Indeed, Miltner, Braun, Arnold, Witte and Taub (1999) demonstrated that associative learning in humans is accompanied by neural synchronization

between the brain areas representing the to-be-associated stimuli. Along these lines, one would expect that binding particular features is a first, preliminary step towards creating a more durable memory trace, suggesting that relevance and impact of binding decrease over time, to the degree that feature conjunctions approach their asymptotic association values. If binding and conjunction learning go hand in hand, learning a given feature conjunction should affect the way in which these features are bound, in terms of strength and speed of binding. In the neural domain, the strength of binding may be defined in terms of firing rates and synchronization of the neurons coding for the bound features. The speed of binding would correspond to the time to achieve a fixed proportion of the maximal firing rates or the synchronization state. The strength of learning of a given feature conjunction should be a function of corresponding binding occurrences. At a neural implementation level, the strength of conjunction learning would correspond to the synaptic strength between the neurons coding for the individual features. On a strong dependence hypothesis, the learning rate from transient binding to long-term associations, as well as the reciprocal top-down effect of existing associative conjunctions on a novel feature binding task, are supposed to be high.

On the other hand, however, binding and learning are expressed over different time-scales and mediate different kinds of neural representations, i.e., perceptual and active working memory representations versus latent long-term memory representations. Binding processes are thought to solve problems in perception resulting from distributed processing, whereas learning processes concern the long-term storage of information that is to be used on a later occasion. Although some combinations of features are more likely than others, many feature conjunctions are so arbitrary—just think of the color of a shirt or the font of a letter—that it makes little sense to store them permanently by creating a long-lasting memory. That is, not much of what binding processes integrate is worthwhile to maintain for much longer than the event in question is perceptually available, which leaves the possibility that binding and learning are less tightly connected than Hebb's account might be taken to suggest. If so, one would not expect a significant impact of long-term learning on the effects of short-term binding with arbitrary feature combinations. According to this weak-dependence hypothesis one would expect that the learning rate from short-term binding to long-term associations is relatively low, thus enabling to filter-out conjunction occurrences with a low behavioral salience.

A strong-dependence hypothesis would suggest that long-term learning factors such as object familiarity, repetition (frequency) of stimuli (i.e., feature-conjunctions) and the frequency of association between stimuli and responses in a given task setting, modulate short-term binding effects in terms of response times or response accuracy. By contrast, a weak-dependence hypothesis would predict a negligible modulation of transient binding effects by familiarity and other variables influencing conjunction learning. To test these alternative hypotheses, we conducted four experiments in which participants were presented with two stimuli in succession, S1 and S2. These stimuli varied on two dimensions, shape and color, thus creating a set of four possible feature combinations. To avoid confounding stimulus repetition effects with response repetition effects we used Hommel's (1998) experimental design, which comprises a pre-cued left or right response (R1)

to the mere onset of S1 (ensuring that no S1 feature was correlated with a particular response and S1 repetitions and R1 repetitions were independent) and a left-right response (R2) cued by the shape of S2. The general idea was to make two of the four feature combinations more likely than the other two, so to induce strong associations between the underlying codes. If binding effects strongly depend upon prior learning, a higher number of presentations of a particular feature conjunction (conjunction learning strength) or enhanced familiarity of a given conjunction (object) should affect the way the respective features are bound, by either reducing or boosting the impact of this binding on different aspects of performance.

Experiment 1

Experiment 1 was modeled after Hommel (1998): Participants were cued to prepare a left- or right-hand keypress (R1), which they carried out as soon as S1—the prime stimulus—was presented (see Figure 1). Even though the identity of S1 did not matter for the response, it varied in shape or orientation (horizontal vs. vertical line) and color (red vs. green). One second later S2 appeared to determine R2. The two alternative shapes of S2 were mapped onto the two R2 alternatives, while the color of S2 was entirely irrelevant to the task. Our focus was on interactions between shape (orientation) and color-repetition effects in S1 and S2. Based on earlier findings (Hommel, 1998; Hommel & Colzato, 2004), we expected that shape repetitions produce better performance on S2 than shape alternations if color is also repeated, but worse performance if color alternates. In other words, performance on S2 should be better with a complete S1-S2 match or mismatch than for partial matches—a pattern that we will call partial-repetition cost. We hypothesize that shape and color features of S1 are still bound when processing S2, so that repeating one feature of S1 would also re-activate the other one, causing an increased interference at S2-dependent response selection.

The crucial question was whether this interaction would vary as a function of conjunction learning or, more precisely, as a function of the relative frequency (i.e., probability) of a given feature conjunction. In Experiment 1 we manipulated the conjunction frequency by presenting two shape-color combinations of S1 (e.g., green-vertical and red-horizontal) four times as often as the other two (red-vertical and green-horizontal). This manipulation was assumed to induce stronger associations between the codes of the more frequent pairs of features, which should yield a main effect of frequency on partial repetition costs.

This frequency manipulation may also affect other than the targeted aspects of performance. In particular, introducing unbalanced frequencies will raise particular expectations leading to a higher degree of preparedness or bias of the cognitive system towards the more probable stimuli. To separate these context-bound short-term effects from the impact of a proper long-term learning, we ran two blocks: an acquisition block in which the conjunction-frequency manipulation was administered and a test block in which all feature combinations were equally probable. The critical test was whether running through the acquisition block would affect performance in the test block. However, working through an extended block of trials may also have effects that are unrelated to the frequency manipulation proper. To control for such

nonspecific effects we compared performance in the experimental group (where frequencies were unbalanced in the acquisition block as described) with that of a control group (where frequencies were balanced).

Method

Participants

Twenty-four students of the Leiden University took part for pay in Experiment 1. All reported having normal or corrected-to-normal vision, and were not familiar with the purpose of the experiment.

Apparatus and Stimuli

The experiments were controlled by a Targa Pentium III computer, attached to a Targa TM 1769-A 17" monitor. Participants faced three grey square outlines, vertically arranged, as illustrated in Figure 1. From a viewing distance of about 60 cm, each of these frames measured $2.6^\circ \times 3.1^\circ$. A thin vertical line ($0.1^\circ \times 0.6^\circ$) and a somewhat thicker horizontal line ($0.3^\circ \times 0.1^\circ$) served as S1 and S2 alternatives, which were presented in red or green in the middle frame. Response cues were presented in the middle frame (see Figure 1), with rows of three left- or right-pointing arrows indicating a left and right keypress, respectively. Responses to S1 and to S2 were made by pressing the left or right shift-key of the computer-keyboard with the corresponding index finger.




Figure 1. Overview of the displays and the timing of events in Experiments 1-2.

Procedure and Design

The experiment consisted of a one-hour acquisition session and a 30 min test session. In both sessions participants carried out two responses per trial. R1 was a simple reaction with the left or right key, as indicated by the response cue. It had to be carried out as soon as S1 appeared, independent of its shape or color. Participants were informed that there would be no systematic relationship between S1 and R1, or between S1 and S2, and they were encouraged to respond to the onset of S1 only, disregarding the stimulus' attributes. R2 was a binary-choice reaction to S2. Half of the participants responded to the vertical and the horizontal line by pressing the left and right key, respectively, while the other half received the opposite mapping. All the participants began with the acquisition session and after 5 minutes break continued with the test session. Half of the participants (control group) received a balanced acquisition session (in which every feature combination of S1 had the same probability to occur) while the other half (experimental group) received an unbalanced acquisition session (in which we manipulated the frequency of feature conjunctions in S1: half of these participants received as S1 the horizontal and red line and the vertical and green line each for 40% of the time and the horizontal and green line and the vertical and red line each for 10% of the time, while the other half received the opposite mapping). The test session was the same for both groups: every feature combination of S1 had the same probability to occur.

The sequence of events in each trial is shown in Figure 1. A response cue signaled a left or right key press (R1) that was to be delayed until presentation of S1, a red or green, vertical or horizontal line in the middle box. S2 appeared 1 sec later—another red or green, vertical or horizontal line in the middle box. The shape of S2 signaled R2, also a speeded left or right key press. R2 speed and accuracy were analyzed as function of the repetition vs. alternation of stimulus shape and color. If the response was incorrect auditory feedback was presented.

The acquisition session comprised 320 trials composed by a factorial combination of the two shapes (vertical vs. horizontal line), colors (red vs. green) of S2, the repetition vs. alternation of shape and color, and, only for the unbalanced session, the frequency (high = 80% vs. low = 20%) of S1. In the balanced acquisition session every feature combination was repeated 20 times while in the unbalanced acquisition session the high-frequency trials were repeated 32 times and the low-frequency trials only 8 times. The test session comprised 224 trials composed by the same factorial combination as in the acquisition session except for the frequency manipulation. In this session, every feature combination had the same probability to appear and was repeated 14 times. Thus, taken together, the two sessions amounted to 544 trials.

Results

S1-R1

The significance criterion for all analyses was set to $p < .05$. We first analyzed the R1 data, the prepared response to S1. In case of errors or anticipatory responses (RTs < 500 ms) subjects had to repeat R1 immediately. Mean correct RTs were analyzed as a function of conjunction frequency (high vs.

low, dummy-coded for the control group) and session (test vs. acquisition), the two within-participant factors, and group (control vs. experimental) as between-participant factor.

There were two significant effects, a main effect of session, $F(1,24) = 26.868$, $p < .001$, that was further modified by group and frequencies, $F(1,24) = 4.376$, $p < .05$. In the acquisition session, the experimental group exhibited faster responses for high- than for low-frequency combinations (399 vs. 411 ms), but frequency effects were absent in the test session of this group (337 vs. 334 ms). As could be expected, in the control group the dummy coded frequency variable showed no effect in the acquisition session (315 vs. 313 ms) nor in the test session (262 vs. 264 ms). Note that in the experimental group the frequency effect was restricted to the acquisition session and did not transfer to the test session.

S2-R2

After excluding trials with missing or anticipatory responses (1.4%), mean RTs and proportions of errors (PEs) for R2 (i.e., the response to S2) were analyzed as a function of group, frequency, session, and the possible relationships between S1 and S2, that is, repetition vs. alternation of stimulus shape or color (see Table 1 for means).

The RTs produced only two reliable effects: a main effect of session, $F(1,22) = 15.58$, $p < .001$, that interacted with color repetition, $F(1,22) = 4.49$, $p < .05$. Whereas the former indicated that responses became faster with practice, the latter was due to the fact that the negligible color-repetition effect in the acquisition session (-3 ms) turned into a small positive effect in the test session (9 ms, i.e., faster RTs with color repetition). However, as this latter effect is accompanied by a (unreliable) negative main effect of color in the error rates (-1.0%), it may reflect a speed-accuracy trade-off and thus should be treated with caution.

More interesting is the three-way interaction of shape, color, and session, which missed the significance criterion in RTs, $p < .1$, but passed it in the PEs, $F(1,22) = 12.11$, $p < .01$. Separate ANOVAs for the acquisition and the test session revealed that, as shown in Figure 2, the acquisition session exhibited typical partial-repetition costs: better performance for shape repetition than alternation if color is repeated but worse performance for shape repetition than alternation if color alternates. In the test session, however, this pattern was not observed. That is, inasmuch as shape-color interactions are produced by feature binding, these features are less likely to get bound as practice increases because the partial repetition costs seem to disappear with more practice. Importantly, this effect does not seem to be due or related to the frequency manipulation, as indicated by the absence of higher-order interactions involving group, $p < .36$ (shape x color x group), $p < .48$ (shape x color x frequency x group), and $p < .81$ (shape x color x frequency x group x session).

The only remaining effect was a somewhat opaque interaction of group, session, and shape repetition in the error rates, $F(1,22) = 9.21$, $p < .01$; it was due to the fact that shape repetition effects were negative (i.e., more errors with repetitions than alternations) in both sessions of the experimental group (-.9% and -.9%) and in the acquisition session of the control group (-1.2%) but not in the test session of the control group (.2%).

Experiment 1														Experiment 2			
Control Group								Experimental Group									
Frequent				Infrequent				Frequent				Infrequent					
Match	RT	PE	RT	PE	RT	PE	RT	PE	RT	PE	RT	PE	RT	PE			
Acquisition Session																	
Neither	512	5.8	520	5.6	558	3.9	567	1.6	442	4.7	468	6.2	456	3.2			
S(hape)	541	8.7	522	7.3	559	4.5	557	1.8	449	5.9	492	8.8	461	6.2			
C(olor)	535	8.1	524	7.1	570	5.3	575	3.6	450	5.7	476	8.5	458	5.8			
SC	522	9.4	520	10.4	556	5.3	560	2.1	435	5.5	458	8.5	436	4.2			
Test session																	
Neither	511	8.9	508	6.5	528	1.8	522	1.8	456	3.2	451	4.9	436	4.2			
S(hape)	505	4.5	502	4.8	519	1.5	530	2.7	461	6.2	444	3.9	458	5.8			
C(olor)	494	5.9	496	4.8	530	3.0	508	2.4	458	5.8	452	4.2	436	4.2			
SC	502	8.6	512	7.1	512	4.1	499	4.5	436	4.2	452	4.5	436	4.2			

Table 1. Acquisition and Test Sessions in Experiments 1 and 2: Means of Mean Reaction Times for Responses to Stimulus 2 (RT; in Ms) and Percentages of Errors on R2 (PE), as a Function of Group, Frequency of Feature Conjunctions, and the Feature Match Between Stimulus 1 and Stimulus 2.




Figure 2. Reaction times and percentage of errors in Experiment 1, as a function of the repetition vs. alternation of stimulus shape and stimulus color, frequency of feature conjunction, and session (acquisition vs. test).

Discussion

Experiment 1 produced six noteworthy results. First, the effect of the conjunction frequency manipulation on R1 shows that indeed a higher frequency of feature conjunctions affects the speed of feature binding, either in terms of a faster binding of perceptual features at stimulus detection, or in terms of a faster activation spreading from perceptual conjunctive codes to the relevant response code.

Second, the effect of frequency manipulation on R1 does not transfer to the test session, thus suggesting either a saturation of feature conjunction learning or the decay of its outcome at some point before the beginning of the test session.

Third, as expected on the basis of earlier observations (Hommel, 1998; Hommel & Colzato, 2004), there was evidence that shape and color were integrated more or less automatically, although this evidence was stronger in the error rates than in RTs. Thus, even if color was a response-irrelevant dimension in the task, color codes engage in shape-color binding, which showed up in partial repetition costs.

Fourth, this shape-color binding effect tended to disappear as practice increased, suggesting that it is unstable for some reason. One possibility is that practice is accompanied by a fine-tuning of selective attention to stimulus dimensions. As color is not relevant to the task, stimulus color may attract some attention in the beginning of the task but loose impact over time. Indeed, manipulations of task relevance have provided evidence that feature integration is modulated by attention to feature dimensions (Hommel, 1998), suggesting that the attentional set determines what gets integrated (Hommel, 2004).

Fifth, there was no evidence that feature combinations that are more frequent in S1-R1 increase or decrease (i.e., modulate) the after-effects of color-shape integration in S2-R2. This may be taken to provide evidence that conjunction learning and short-term feature binding are mediated by different mechanisms.

Finally, it should be noted that we observed better performance for frequent combinations on R1 but not on R2. We had expected that the frequency manipulation in terms of S1 would transfer to R2, given that S1 and S2 were made up of the same features, but this transfer did not take place.

This lack of transfer may suggest that conjunction learning is very context-sensitive, e.g., by integrating the conjunction with codes representing the particular task. If so, the absence of an interaction between binding and frequency-induced learning may be less informative than we had hoped. To test this possibility, in Experiment 2 we re-ran the experimental group but this time manipulated the frequency of feature conjunctions in S2.

Experiment 2

As pointed out, one reason for why we failed to find an interaction between learning and after-effects of binding in Experiment 1 might have to do with the lack of transfer of the frequency effect from S1-R1 to S2-R2. Accordingly, Experiment 2 was conducted, in which the frequency manipulation was applied to S2 instead of S1. This should produce a main effect of frequency on R2 and, if the lack of transfer was indeed the critical factor, it also should

produce the sought-for interaction of binding after-effects and learning in terms of biased conjunction frequencies.

Method

Twenty-four students participated. They all fulfilled the same criteria as in Experiment 1. The method was exactly the same as in the experimental group of Experiment 1, except that the frequency manipulation referred to S2 instead of S1.

Results

The data were analyzed as in Experiment 1. The data from the control group from Experiment 1 was included in the analysis to create the Group factor.

S1-R1

Mean correct RTs were analyzed as a function of session (test vs. acquisition) and group (control [from Experiment 1] vs. experimental). The only significant effect of session, $F(1,21) = 32.30$, $p < .001$, indicated that responses were faster in the test session (243 ms) than in the acquisition session (288 ms).

S2-R2

Trials with missing or anticipatory responses (1.6%) were excluded from the analysis. Table 1 provides an overview of the means for RTs and PEs obtained for R2. The errors yielded only a single reliable effect, an interaction between shape repetition and session, $F(1,21) = 12.18$, $p < .05$; it was due to a slight disadvantage for repeated as compared to alternated shapes in the acquisition session (7.1% vs. 6.0%) that disappeared in the test session (4.7% vs. 4.7%).

The RT analysis produced two clusters of effects. One cluster involved a main effect of session, $F(1,21) = 5.20$, $p < .05$, an interaction of group and frequency, $F(1,21) = 6.53$, $p < .05$, and a three-way interaction including all three factors, $F(1,21) = 8.97$, $p < .01$. Apart from a practice effect, the underlying pattern revealed that the frequency effect was restricted to the acquisition session of the experimental group (444 vs. 474 ms, i.e., a 30-ms benefit for frequent feature combinations) but absent in the test session (453 vs. 450 ms) and in both sessions of the control group (528 vs. 522 ms and 503 vs. 505 ms). That is, the frequency manipulation was successful in affecting R2 performance, even though this effect again failed to transfer to the test session.

The other cluster comprised the expected interaction of shape and color, $F(1,21) = 6.87$, $p < .05$, which was further modified by session, $F(1,21) = 10.01$, $p < .05$. As shown in Figure 3, the acquisition session produced the common partial-repetition-cost pattern whereas the test session did not—an observation that was confirmed by separate ANOVAs, $p < .001$ and $p = .535$, respectively. Thus, as in Experiment 1, the shape-color interaction and the feature binding process it indicates seem to disappear with practice. Importantly, frequency was again not involved in any higher-order interaction with shape and color, $p < .35$ (shape x color x frequency x session), $p < .61$ (shape x color x frequency x group), and $p < .12$ (shape x color x frequency x group x session).




Figure 3. Reaction times and percentage of errors in Experiment 2, as a function of the repetition vs. alternation of stimulus shape and stimulus color, frequency of feature conjunction, and session (acquisition vs. test).

Discussion

Experiment 2 yields evidence that the unbalanced frequency of feature combinations leads to an immediate adjustment that benefits performance. However, this benefit does not transfer to later performance on the same task (the test session in this experiment) and it does not seem to affect feature integration after-effects in terms of partial repetition costs. Accordingly, it seems unlikely that the lack of interaction between frequency manipulation and binding after-effects in Experiment 1 was merely due to a lack of transfer from S1-R1 to S2-R2. Rather, it seems that frequency-dependent conjunction learning does not bias short-term feature binding, which points to an independence between these processes. The fact that the conjunction-frequency manipulation of S2 affects R2 but not R1 supports our suspicion that conjunction learning is highly context-specific and, thus, is more complex than a simple Hebbian association of neural units coding for stimulus features. Finally, the fact that binding after-effects tend to vanish with practice from the acquisition to the test session, just as in Experiment 1, reinforces the idea of a progressive attentional tuning to the relevant stimulus dimensions.

Experiment 3

Based on Experiments 1 and 2, the lack of interaction between conjunction learning and binding after-effects may be due to two reasons. On the one hand, these two processes may work independently of each other, so that any further attempt to find the sought-for interactions would be doomed to fail. On the other hand, however, our manipulation of conjunction frequency may not have led to learning of a proper integrated representation at high levels of object representation (in inferotemporal cortex) but, rather, to a merely transient general bias towards the more likely conjunctions—by a kind of continuously updated situational model held in working memory (involving prefrontal cortical systems, see Duncan, 2001). To address this possibility, we designed Experiment 3 to study the effect of stably-learned feature combinations in object representations on transient binding. Instead of the arbitrary feature conjunctions used in Experiments 1 and 2, Experiment 3 employed images of real-life stimuli, a banana and a strawberry, which could appear in either their “natural” colors—i.e., yellow and red—or in the opposite colors—red and yellow—which participants were unlikely to have experienced frequently in combination with these objects. That is, we used well-known stimuli and presented them in either their standard color or in a color that is unlikely to be associated with them. We did not include an acquisition session but had participants to work through a test session only, where every feature combination of S1 and of S2 was equally probable. We hypothesized that familiar, i.e., stably-learned feature combinations, such as the yellow banana and the red strawberry, might affect binding differently than less familiar combinations, such as a red banana and a yellow strawberry.

Method

Twenty-four students participated. They all fulfilled the same criteria as in Experiment 1. The experiment consisted of one half-hour session. The procedure and the sequence of events were as in the test session of Experiments 1 and 2, with the following exceptions: Instead of a vertical and a

horizontal line we presented figures of a banana ($0.3^\circ \times 0.6^\circ$) and a strawberry ($0.5^\circ \times 0.6^\circ$, see Figure 4), appearing in red or yellow inside the middle frame. Half of the participants responded to the shape of the banana and the strawberry by pressing the left and right key, respectively, while the other half received the opposite mapping. As in the previous experiments, color varied orthogonally to the shape and was completely irrelevant to the task.



Figure 4. Bitmaps of the stimuli used in Experiments 3 and 4.

Results

S1-R1

Correct RTs were analyzed as a function of the familiarity of the feature combinations. Familiarity indeed affected performance by producing faster responses with familiar than unfamiliar combinations (304 vs. 313 ms), $F(1,22) = 4.46$, $p < .05$.

S2-R2

Data were analyzed as a function of familiarity, shape repetition, and color repetition. Trials with missing or anticipatory responses (1.3%) were excluded from the analysis. Table 2 provides an overview of the means for RTs and PEs obtained for R2. No reliable effect was obtained for PEs. In RTs, familiarity produced a main effect, $F(1,22) = 7.31$, $p < .05$, indicating that responses were faster to familiar than to unfamiliar combinations (498 vs. 511 ms). The only other significant effect was the interaction between shape and color repetition, $F(1,22) = 9.10$, $p < .01$, following the standard pattern indicative of partial-repetition cost (see Figure 5). Importantly, however, there was no sign of a three-way interaction involving familiarity, $F < 1$.

	Experiment 3				Experiment 4	
Familiarity	high		low		low	
	RT	PE	RT	PE	RT	PE
Match						
Neither	495	5.4	508	4.9	462	5.0
Shape	508	4.0	511	4.8	469	7.8
Color	507	4.8	520	5.5	468	5.8
Both	483	5.6	504	4.8	449	7.3

Table 2. Experiment 3 and 4: Mean Reaction Times for Responses to Stimulus 2 (RT-R2; in Ms) and Percentages of Errors on R2 (PE-R2), as a Function of Experiment, Familiarity, and the Feature Match Between Stimulus 1 and Stimulus 2.




Figure 5. Reaction times and percentage of errors in Experiment 3, as a function of the repetition vs. alternation of stimulus shape and stimulus color and familiarity (familiar vs. unfamiliar) of object-color conjunction.

Discussion

Experiment 3 was successful in demonstrating a frequency effect induced by life-long learning that meets two important criteria: it affected performance even under conditions where all feature combinations were equally probable and it affected both S1-R1 and S2-R2. This suggests that in this experiment a higher level of object representation and perceptual learning was involved. As another important result, we replicated the partial-repetition-cost pattern with real-life, object-like stimuli, which demonstrates that the previous observations with simple geometric elements (e.g., Hommel, 1998) are generalizable. As we will see further on, the shape-color interaction was even more pronounced than observed with simpler stimuli, an issue we will address in Experiment 4. And, yet, there is still no evidence of any interdependency of familiarity and binding after-effects. Even though our considerations are still based on a null effect, which necessarily renders them preliminary, we thus tend to conclude that short-term binding after-effects, like the partial repetition cost, and long-term learning, are independent processes.

Experiment 4

Introducing familiar objects in Experiment 3 produced both frequency effects and the common shape-color interaction without demonstrating the sought-for interaction between these variables. In this respect, using familiar objects did not change the outcome obtained in Experiments 1 and 2 or the conclusion they suggested. However, we have mentioned that something did change in Experiment 3: the shape-color interaction was more pronounced (see further on). One possible (though theoretically less interesting) reason for this observation may relate to the amount of practice. As briefly considered in Experiment 1, practice on the task may allow fine-tuning input selection processes and thus increasingly prevent the irrelevant color information from being processed. In Experiments 1 and 2, fine-tuning could begin in the acquisition session already, so that color would no longer be processed in the test session. Given that there was no acquisition session in Experiment 3 and the fact that the test block was somewhat shorter than the acquisition block, the stronger effect in Experiment 3 may thus simply reflect the fact that participants in Experiments 1 and 2 were tested after more extended practice.

However, the outcome of some post-hoc analyses render this account implausible. First, we divided the acquisition trials in Experiments 1 and 2 (i.e., the two experimental groups and the control group) in four equal miniblocks and re-ran ANOVAs on these data with miniblock as additional factor. It turned out that the shape-color interaction interacted with miniblock, $F(1,30) = 8.96$, $p < .005$. Separate analyses showed that shape and color interacted in the first miniblock only, $p < .001$, but not in the three other miniblocks, $p > .59$, $p > .16$, and $p > .74$. The same analysis of the test trials in Experiment 3 (block length equated) did not reveal any modification of the shape-color interaction by miniblock, $F < 1$. That is, in Experiments 1 and 2, 80 trials of practice were sufficient to eliminate the shape-color interaction, which in Experiment 3 survived 224 trials without any drop in size. This conclusion was further confirmed by a direct comparison of performance in the control group of Experiment 1 and Experiment 3, in which we included the first two acquisition miniblocks from the former and the first two test miniblocks from the latter.

The shape-color interaction was modified by miniblock and experiment, $F(1,33) = 4.67$. Separate analyses showed that experiment modified the shape-color interaction in the second, $p < .05$, but not the first miniblock, $p > .3$. That is, practice seems to eliminate shape-color integration—presumably by gating out color information—but only if the stimulus material consists of arbitrary geometric symbols.

Experiment 4 was designed to disentangle two possible interpretations of this outcome. Clearly, the stimulus material differed between Experiments 1 and 2 on the one hand and Experiment 3 on the other: the stimuli were more simple and arbitrary, and less “biological” in the former than the latter. This implies that the stimuli used in Experiment 3 may be cognitively and neurally represented in a different way than the lines used in the previous experiments. They are objects and are likely to be perceived and categorized as such, which among other things will involve the activation of conceptual traces in long-term memory—which, after all, was the reason to employ them. This fact may change the way these stimuli were processed. However, not only were the stimuli in Experiment 3 more object-like in terms of their more complex shapes and meanings, they also appeared in their standard colors—at least in 50% of the trials. To rule out that this color appearance was the responsible factor, we replicated Experiment 3 but replaced the two “biologically plausible” colors by colors that were unlikely to be closely associated with one or the other object.

Method

Twenty-four students participated. They all fulfilled the same criteria as in Experiment 1. The procedure and the sequence of events were as in Experiment 3, except that the colors were pink and blue instead of yellow and red.

Results

S1-R1

Valid responses were carried out in 280 ms on average.

S2-R2

Table 2 provides an overview of the means of RTs and PEs obtained for R2. The data were analyzed as in Experiment 3, except that the familiarity factor no longer applied. The errors yielded a main effect of shape repetition, $F(1,22) = 10.92$, $p < .01$, due to more errors being made with shape repetitions than alternations (7.6% vs. 5.4%). However, as the shape effect had the opposite sign in RTs (459 vs. 465 ms, an unreliable difference) this might reflect a mere speed-accuracy trade-off. The RTs produced a main effect of color repetition, $F(1,22) = 4.96$, $p < .05$, and an interaction of shape and color repetition, $F(1,22) = 13.66$, $p < .01$. Whereas the former indicated a 7-ms benefit for color repetitions, the latter followed the expected cross-over cost pattern shown in Figure 6. An additional analysis with miniblock as a factor did not provide any evidence that the shape-color interaction decreases with practice, $F < 1$.




Figure 6. Reaction times and percentage of errors in Experiment 4, as a function of the repetition vs. alternation of stimulus shape and stimulus color.

Discussion

In sum, we find the same outcome as in Experiment 3: a pronounced, stable interaction indicative of shape-color binding. As color was again an irrelevant dimension and as the colors were not pre-experimentally related to the shapes of the stimuli, this finding supports the idea that the objecthood of stimuli changes the way they are processed and the way their features are integrated. Once again, this suggests that higher-level neural units are involved in integration of real-world shapes with color features. Indeed, it is plausible to assume that long-term object representations are less context-specific than representations of conjunction probabilities of arbitrary features and, hence, less selective with respect to the feature dimensions related to the object at hand. In other words, even though one may learn to neglect the color of arbitrary shape-color conjunctions if the task at hand only requires attending to shape, it makes little sense to drop and neglect color information as a constituent of the long-term representation of an apple only because in one given situation the color of a particular apple played no role. We will

elaborate on the architectural implications of this consideration in the General Discussion.

GENERAL DISCUSSION

The four experiments of this study aimed at addressing the relation between short-term binding effects and long-term learning: to what extent are binding and learning processes independent, to what extent do they interact? We considered short-term binding (after-)effects and learning in terms of both experimentally induced feature-conjunction biases and stably-learned natural feature conjunctions. Although it is clear that more research on this issue is necessary, we take our findings to point to an independence of binding and learning, at least with respect to direct interactions. That is, we think it is justified to reject the strong-dependence hypothesis outlined in the introduction. The different result patterns we obtained in Experiments 1 and 2 on the one hand and in Experiments 3 and 4 on the other further suggest that it matters whether binding is restricted to visual features or whether overlearned objects are involved.

In our view, explaining the outcome of this study requires the consideration of three different systems or representational levels: a low-level representation of features in feature maps; a higher-level, long-term representation of objects; and a working-memory system in which situational contingencies are temporarily stored. As we will argue below, it makes sense to assume the location of these three systems or levels in visual cortex, inferotemporal cortex, and in prefrontal cortex, respectively—even though our considerations do not hinge on these particular localizations. On the one hand, we assume that the binding process proper is automatic and not directly impacted by higher-level signals from either inferotemporal or prefrontal cortex. On the other hand, however, which features are considered for binding does depend on attentional settings (Hommel, 2004), which in turn are affected by situational models in prefrontal cortex and, if available, object representations in inferotemporal cortex. In the following, we will elaborate this theoretical framework and consider its neural plausibility.

Higher-order object representation and low-level binding

Experiments 1 and 2 provide evidence that processing the shape (or orientation) and color of simple geometric stimuli leads to a binding of the neural codes representing the features. However, in contrast to the real objects used in Experiments 3 and 4, the geometric stimuli produced binding over a couple of trials only. This means that the presence or absence of long-term representations of the particular stimuli makes a difference for binding. However, we found no evidence for any direct interaction between short-term or long-term object representations and binding. How can that be explained? We think the key to understanding this rather complex relationship between learning and binding requires the distinction between at least two levels of representation and integration: a lower representational level at which features are temporarily linked and a higher representational level at which integrated feature assemblies are stored. Apparently, linking features at the proposed lower level does not directly translate into having stored feature

links at the higher level. And reversely, having a stored feature link at the higher level does not directly impact creating a link at the lower level.

Feature-binding problems are often discussed with respect to the visual system, where the strong evidence for a whole multitude of feature maps (Felleman & van Essen, 1991; Lamme & Roelfsema, 2000) makes the need for integration processes particularly obvious. As discussed in the introduction, a promising candidate mechanism for fast binding of low-level visual features like orientation or color, is the synchronization of neural responses (Gray et al., 1989; von der Malsburg, 1981, 1999). As compared to other mechanisms, synchronization would not only be a fast and flexible mechanism but would also enable the representation of any novel and arbitrary feature combination. Indeed, there is significant evidence in support of the idea that synchronization in the gamma frequency band (high frequency EEG activity above 30 Hz) plays a role in visual feature binding (Engel & Singer, 2001), visual working memory (Luck & Vogel, 1997; Raffone & Wolters, 2001), and consciousness (Engel & Singer, 2001).

However, the visual system does not only comprise of low-level feature maps with a high spatial resolution such as in V1 and V2/V3. Higher level neurons coding for more complex shapes and multi-feature objects can be found in the occipito-temporal (ventral) stream, like in V4 and IT. Converging feed-forward connections are likely to enable increased response selectivity and transmission of signals for fast bottom-up processing, and feedback diverging synapses to mediate attentional and learning-based modulation of neural responses. Numerous studies provide evidence that convergence plays a more important role at this higher representational level. For instance, some cells have been shown to be selective to stimuli as complex as faces (Young & Yamane, 1992) or, in the posterior inferior temporal cortex, that are selective for conjunctions of a striped patch and flanking black spots (Tanaka, Saito, Fukada & Moriya, 1991). Hence, even if we exclude convergence as the only integration mechanism, there are good reasons to believe that at least some feature conjunctions are encoded by assemblies of a limited number of selectively-tuned neurons in inferotemporal cortex, which are adapted and shaped by Hebbian learning (e.g., Amit, 1995). These neural assemblies would also be involved in visual short-term memory, in cooperation with prefrontal cortex.

Thus, we suggest that integration by synchronization and by convergence are not as exclusive as previously held (e.g., Jellema & Perrett, 2002; Singer, 1994) but, rather, may coexist to solve binding problems at different levels (Hommel, 2004; Singer, 1999). In particular, the recognition of familiar objects is achieved by assemblies of highly selective conjunction detectors that only emerge for behaviorally relevant, frequently occurring events and that change only slowly through Hebbian learning. In contrast, frequently changing or novel combinations of arbitrary visual features are coded by synchronizing relatively raw feature codes represented at feature maps.

Objects versus feature conjunctions: The role of frequency learning and attentional weighting

Let us now consider how low-level feature binding may be affected by the objecthood of the stimuli involved and the probability of particular feature

combinations. As we have pointed out, we found no evidence for any direct impact of frequency or familiarity manipulations on partial-overlap costs, our measure of feature binding. And yet, more probable combinations were processed faster than less probable combinations. This means that the system was biased towards more likely feature conjunctions but that this bias affected processing only after features were bound or, in the case of partial overlap, re-bound. One possibility is that the shape-color associations underlying this bias represent a first step in the emergence of a new object representation. However, this idea does not seem to fit with the lack of transfer of the frequency bias to the test session and with the different result patterns in Experiments 1 and 2 on the one hand and Experiments 3 and 4 on the other. Alternatively, frequency-based expectations may be incorporated into situational models held in working memory. For instance, low-level integration may run autonomously but its outcome may be registered and processed more quickly if it fits situation-specific expectations. In any case, however, it is important to note that the speed of what we attribute to low-level feature integration is unaffected by top-down expectations.

And yet, there was evidence for an indirect modulation of low-level integration by top-down processes. When we used stimuli made up of arbitrary feature conjunctions (combinations of shape and color) we found our measure of feature binding to be rather instable and it even disappeared over time, whereas life-long practiced stimuli yielded robust and stable effects. We attribute the first finding to adaptive feature weighting (Hommel, Muesseler, Aschersleben & Prinz, 2001a), that is, to the dynamic weighting of feature dimensions according to their contribution to task performance. Feature weighting is an attentional process that selectively prepares the cognitive system for the differential processing of relevant (i.e., to-be-attended) and irrelevant (i.e., to-be-ignored) features of anticipated perceptual events (cf., Bundesen, 1990). As the color of stimuli is rather salient and likely to be helpful in discriminating targets from nontarget stimuli, such as the fixation point or cues, the weight of the color dimension is unlikely to be zero, at least at the beginning of an experiment. With increasing practice, however, people will fine-tune the weights of the perceptual dimensions to better reflect their use for current performance. As color was irrelevant to the task, this is likely to have led to a continuous decrease of the color-dimension weight. (Note that the weak contribution from color is not due to a particularity of this dimension, as the contribution from shape is as weak if S2 is not defined by shape: Hommel, 1998.) If we further assume that the weight of a perceptual dimension determines the probability that the corresponding features are considered in perceptual binding (Hommel, 2004), it is easy to see why the decrease in color weighting eliminated partial-overlap costs: the color feature was activated to a degree that was insufficient for binding, so that its code was not involved anymore.

Using real-life objects brought long-term object representations (presumably located in inferotemporal areas) into play. These representations must have emerged from numerous encounters with the represented objects, which implies that they do not include situational particularities, such as the task-specific value of one or the other feature dimension. Accordingly, it makes sense to assume that the involvement of long-term object representations top-down primed all of the feature dimensions defining the

object to an equal degree. Indeed, there is considerable evidence that processing one feature of an object automatically opens the attentional gate to other features of this object (Baylis & Driver, 1992; Duncan, 1984; Kahneman & Henik, 1981). This top-down priming effect may prevent or overrule the practice-induced diminished weighting of nominally irrelevant feature dimensions, and thereby keep the contribution of features defined on such a dimension sufficiently strong to stay involved in binding.

Levels of binding and learning

Taken altogether, our experiments suggest the existence of different levels of integration and conjunction learning in the human brain. A first level flexibly combines entries in low-level feature maps (such as in V1, V2) probably mediated by neural synchronization. Binding at this level is highly context-sensitive, suggesting that features are linked to task or context information, perhaps via the occipito-parietal processing stream and multimodal neurons in parietal posterior cortex.

A second type of short-term learning expresses itself in rather local contingencies, such as the probability of particular feature conjunctions. This type of learning is also transient and task specific, which may point to the involvement of situational models maintained in prefrontal cortex. It may lead to a faster readout of expected feature conjunctions and/or the lowering of thresholds in responding to expected conjunctions. In any case, however, frequency-based expectations do not seem to speed up (re-)binding or access to object representations but seem to act on subsequent processing steps.

Third, known, overlearned objects are represented at a more integrated level, presumably by means of conjunctive neurons and reverberating neural assemblies in inferotemporal cortex. At this level, familiar objects are long-term encoded, and activating their codes provides top-down priming of object-related feature dimensions. In particular, coding familiar objects in terms of Hebbian cell assemblies with reverberatory (attractor) dynamics has the consequence that activating one part of an object representation spreads to the whole integrated assembly, in a kind of pattern completion process (Hommel, 2004). This makes it difficult to isolate the contributions of individual components (e.g., feature codes) in the assembly, which among other things overrules possibly differential attentional weights for object features provided by feedback signals from prefrontal cortex. These inferotemporal assemblies are likely to be only part of a larger inter-area neural representation of events in the cerebral cortex, i.e., of long-term integrated event files (Hommel, 2004). Experiment 4 suggests that these neural assembly representations need not be restricted to overlearned feature values but may also be updated by and “capture” context-specific features, like colors. We hypothesize that this capture process recruits neurons at different levels in the visual cortical hierarchy, like in areas IT (complex real-world object shape) and V4 (color).

The proposed neurocognitive architecture leads to interesting empirical predictions with respect to a whole number of tasks and phenomena. For instance, consider the recently developed change-detection paradigm for visual working memory (Luck & Vogel, 1997; Vogel, Woodman & Luck, 2001). Our architecture would predict that conjunctions between shapes of real-world objects (like bananas or strawberries) and arbitrary colors would be more stably retained in visual working memory than meaningless conjunctions of

features, due to the higher involvement of reverberatory activities in inferotemporal cortex. The stability of maintenance of such neural binding codes may affect response accuracy, response time, as well as robustness against interference and decay. Moreover, our model would suggest that in the presence of task-irrelevant shape changes, effects of changing the color may occur with shapes of real-world objects, but to a lower degree than with meaningless shapes (e.g., orientations).

To conclude, our findings suggest that the neurocognitive representation of feature conjunctions is a multi-component process involving several time scales and levels of integration. They also suggest that the interaction between top-down attentional processes and automatic binding processes is dynamic and adaptive to task constraints. It remains to be seen how integrated structures or event files of different nature behave over short and long time scales. For instance, intentional maintenance of binding codes in working memory probably depends upon the interactions between prefrontal cortex, posterior cortical cortices and premotor cortex, and perhaps involves long-range neural synchronization (Gross, Schmitz, Schnitzler, Kessler, Shapiro, Hommel, & Schnitzler, 2004; Tononi, Sporns & Edelman, 1992). Related neuroimaging and neurophysiological investigations, as well as large-scale neurocomputational modeling, will play a crucial role in answering these core neurocognitive questions.

Chapter 6

Conclusions

The so-called binding problem is the problem of how our brain is able to integrate correctly the feature codes that belong to one and the same event. For visual perception Kahneman, Treisman and Gibbs (1992) assumed that the codes of features belonging to the same object are integrated into what they call an object file, a temporary cognitive structure addressed by location containing all the perceptual information about the object and perhaps even more (e.g., semantic information). Further investigations have shown that feature integration is not limited to the domain of visual perception, for which the binding problem has been first formulated (see Treisman, 1996, for an overview), but spans perception, action planning, and stimulus-response relationships (Hommel, 2004). Hommel (1998) obtained evidence that features of accidentally paired stimuli and responses are spontaneously bound. Stoet and Hommel (1999) demonstrated that binding a spatial feature to one plan makes it less available for the construction of another plan involving the same feature.

A current discussion concerns the need of feature binding. Some authors have emphasized the need of integration processes in distributed representational systems like the human brain (e.g., Cowey, 1995). Others, however, have questioned the very necessity of feature binding. Phaf, van der Heijden, and Hudson's (1995) SLAM model, for instance, is a hierarchical, but parallel model of selective attention. In SLAM no binding problem exists because the model is based on conjunction detector modules, that code every possible combination of features. Jellema, Maassen and Perrett (2004) indeed found in the superior temporal cortex of the macaque monkey single cells that integrate information about the form, motion and location of animate objects.

There is also a discussion about the mechanisms underlying feature binding. Some investigators argue that the synchronization of activation of cell populations might play a major role in binding features belonging to the same event (Singer, 1994; Treisman, 1996). Other authors have questioned the involvement of neural synchronization in feature binding. Van der Velde and de Kamps (2002), for instance, suggested that models that rely on synchrony of activation lack the systematicity and productivity of true compositional systems.

In the present study a pragmatic stance with regard to these discussions is taken. Because in the literature a diversity of binding effects is reported, it is simply assumed that binding is a real phenomenon and the characteristics of this phenomenon are investigated and put to a test. The principal outcomes of this investigation are the following.

First, the data reported in chapter one and four show that *feature integration is not global* as assumed by Kahneman et al. (1992), but that *feature binding is dependent on task and task instruction*. Not all features of a given perception-action event are integrated with each other: what gets integrated seems to be determined by task relevance or, more precisely, by

whether the given feature varies on a dimension that is explicitly or implicitly defined as relevant in the task. How this is possible is an open issue. We speculate that a higher order integrated representation of the context plays for selecting the appropriate response an important role.

Second, in chapter one and four, we found that, to some degree, *the strength of a binding can be predicted from the degree to which the respective features are activated*. The most reliable predictor in this respect seems to be the relevance of the features for matching the currently implemented attentional set. That is, to the degree that this set is optimized and maintained, processing the task relevant stimulus feature is facilitated. The more a relevant feature code is activated, the more likely it is to become integrated with other sufficiently activated feature codes.

Third, in chapter two and three, we presented some indirect evidence that supports the idea that *feature binding in the visual cortex is associated with neural synchronization driven by the muscarinic cholinergic system*. Synchronization of cell populations driven by the muscarinic-cholinergic system might play a major role in binding features in the visual cortex (Munk, 2003; Rodriguez, Kallenbach, Singer & Munk, 2000, 2001; Rodriguez-Bermudez, Kallenbach, Singer & Munk, 2004). We found that within-object bindings were significantly decreased by alcohol and increased by caffeine (a muscarinic-cholinergic antagonist and agonist, respectively), and that they were unaffected by nicotine, which only affects the nicotinic-cholinergic system.

Fourth, in chapter two, three and four, it appeared that *feature binding is accomplished by domain-specific mechanisms* for perception and action. None of the drug manipulations had any impact on bindings of stimulus and response features. This suggests that bindings in visual perception are mediated by cholinergically-driven neural synchronization in the gamma frequency band (see above; Engel & Singer, 2001; Keil, Muller, Ray, Gruber & Elbert, 1999), whereas bindings across perception and action seem to be created in a different fashion, either in terms of the driving system or in terms of the synchronization frequency used. There are, indeed, a number of indications that longer-distance bindings are associated more with synchronization in the beta frequency band (Gross, Schmitz, Schnitzler, Kessler, Shapiro, Hommel, & Schnitzler, 2004; Kopell, Ermentrout, Whittington & Traub, 2000; Roelfsema, Engel, Koenig & Singer, 1997). In chapter four, also evidence against a possible common integration mechanism was obtained. We failed to find any correlations between feature-binding effects across tasks that measured binding in perception, in action planning and across perception and action.

Last, in chapter five, the *objecthood* (feature conjunction in real-world objects) *of stimuli changes the way they are processed and the way their features are integrated*. Processing the shape and color of simple geometric stimuli leads to a binding of the neural codes representing the features. However, in contrast to real world objects, the geometric stimuli produced binding over a couple of trials only. To account for these results we argued that neural synchronization and conjunction detectors are not two mutually exclusive mechanisms but solve the binding problem at different levels. We speculated that conjunction detectors are compelling candidates for coding behaviorally-relevant or familiar perceptual events, that is, for transforming the

repeated experience of the same feature combinations into more durable representations (long-term memory). By contrast, synchronization is primarily input-based, and is a prime candidate for flexible coding of arbitrary feature combinations (serving current perceptual processing and short-term memory).

In sum, the present thesis provides converging evidence that feature binding is a general phenomenon that takes place within and across domains. It is accomplished by domain-specific mechanisms. For instance, feature binding in the visual cortex is associated with neural synchronization driven by the muscarinic cholinergic system, but stability of binding over time seems to depend on the presence or absence of long-term representations. Importantly, binding is also dependent on task and on task instruction. Feature binding is not an automatic global mechanism, but it follows a task-relevance criterion: not all features of a given event are integrated with each other and what gets integrated depends on what is task relevant.

References

- Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.
- Aigner, T. G., & Mishkin, M. (1986). The effects of physostigmine and scopolamine on recognition memory in monkeys. *Behavioral Neuroscience*, 45, 81-87.
- Aigner, T. G., Walker, D. L., & Mishkin, M. (1991). Comparison of the effects of scopolamine administered before and after acquisition in a test of visual recognition memory in monkeys. *Behavioral Neural Biology*, 55, 61-67.
- Allport, D. A., Tipper, S. P., & Chmiel, N. R. J. (1985). Perceptual integration and postcategorical filtering. In M. I. Posner, & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 107-132). Hillsdale, NJ: Erlbaum.
- Amit, D. J. (1995). The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*, 18, 617-626.
- Barlow, H. B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology. *Perception*, 1, 371-394.
- Baylis, G. C., & Driver, J. (1992). Visual parsing and response competition: The effect of grouping factors. *Perception & Psychophysics*, 51, 145-162.
- Bertelson, P. (1963). S-R relationships and reaction times to new versus repeated signals in a serial task. *Journal of Experimental Psychology*, 65, 478-484.
- Borchard, F., Singer, W., & Munk, M. (2000). Electrical stimulation of the substantia innominata leads to an increase in fast oscillatory multiunit activity in cat visual cortex. *Abstract Society of Neuroscience*, 30, 736.20.
- Braitenberg, V., & Schuz, A. (1991). *Anatomy of the Cortex: Statistics and Geometry*. Berlin: Springer-Verlag.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, 97, 523-547.
- Bundesen, C., Kyllingsbæk, S., & Larsen, A. (2003). Independent encoding of colors and shapes from two stimuli. *Psychonomic Bulletin & Review*, 10, 474-479.
- Carter, A. J., O'Conner, W. T., Carter, M. J., & Ungerstedt, U. (1995). Caffeine enhances acetylcholine release in the hippocampus in vivo by a selective interaction with adenosine A1 receptors. *Journal of Pharmacology and Experimental Therapeutics*, 273, 637-462.
- Castelo-Branco, M., Goebel, R., Neuenschwander, S., & Singer, W. (2000). Neural synchrony correlates with surface segregation. *Nature*, 405, 685-689.
- Cisek, P., & Turgeon, M. (1999). Binding through the fovea: A tale of perception in the service of action. *Psyche*, 5, <http://psyche.cs.monash.edu.au/v5/psyche-5-34-cisek.html>.
- Cohen, J. D. & Servan-Schreiber, D. (1992). Context, cortex and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, 99, 45-77.
- Colzato, L. S., Erasmus, V., & Hommel, B. (2004). Moderate alcohol

- consumption impairs feature binding in visual perception but not across perception and action. *Neuroscience Letters*, 360, 103-105.
- Colzato, L. S., Fagioli, S., Erasmus, V., & Hommel B. (2005). Caffeine, but not nicotine enhances visual feature binding. *European Journal of Neuroscience*, 21, 591-595.
- Colzato, L. S., Raffone, A., & Hommel B. (2004). What do we learn from binding features? Evidence for multilevel feature integration. Manuscript under revision for publication.
- Colzato, L. S., Warrens, M. J., & Hommel B. (2004). Priming and binding in and across perception and action: A correlational analysis of the internal structure of event files. Manuscript submitted for publication.
- Cowey, A. (1985). Aspects of cortical organization related to selective attention and selective impairments of visual perception: A tutorial review. In M.I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 41-62). Hillsdale, NJ: Erlbaum.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193-222.
- DeYoe, E. A., & Van Essen, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neuroscience*, 11, 219-226.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113, 501-517.
- Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nature Reviews Neuroscience*, 2, 820-829.
- Duncan, J., Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: The organization of goal-directed behavior. *Cognitive Psychology*, 30, 257-303.
- Dutzi, I. B., & Hommel, B. (2003). Spontaneous but goal-dependent binding of actions and their effects. Manuscript submitted for publication.
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H.J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biological Cybernetics*, 60, 121-130.
- Engel, A. K., König, P., Kreiter, A. K., Schillen, T. B., & Singer, W. (1992). Temporal coding in the visual cortex: New vistas on integration in the nervous system. *Trends in Neuroscience*, 15, 218-226.
- Engel, A. K., Roelfsema, P. R., Fries, P., Brecht, M., & Singer, W. (1997). Binding and response selection in the temporal domain - a new paradigm for neurobiological research. *Theory in Biosciences*, 116, 196-221.
- Engel, A. K., & Singer W. (2001). Temporal binding and the neural correlates of sensory awareness. *Trends in Cognitive Science*, 5, 16-25.
- Fagioli, S., Paolieri, D., & Colzato, L. S. (2004). *Regressione acuta a stadio preadolescenziale*. Intern report, University of Leiden.
- Fell, J., Fernandez, G., Klaver, P., Elger, C. E., & Fries, P. (2003). Is synchronized neuronal gamma activity relevant for selective attention? *Brain Research Reviews*, 42, 265-272.
- Felleman, D. J., & van Essen, D. C. V. (1991). Distributed hierarchical processing in the primate visual cortex. *Cerebral Cortex*, 1, 1-47.
- Fellous, J.-M., & Sejnowski, T. J. (2000). Cholinergic induction of oscillations

- in the hippocampal slice in the slow (0.5-2 Hz), theta (5-12 Hz), and gamma (35-70 Hz) bands. *Hippocampus*, 10, 187-197.
- Fox, E. (1995). Negative priming from ignored distractors in visual selection: A review. *Psychonomic Bulletin & Review*, 2, 145-173.
- Fuentes, L. J., Vivas A. B., & Humphreys, G. W. (1999). Inhibitory tagging of stimulus properties in inhibition of return: Effects on semantic priming and flanker interference. *Quarterly Journal of Experimental Psychology*, 52A, 149-164.
- Gordon, R. D., & Irwin, D. E. (1996). What's in an object file? Evidence from priming studies. *Perception & Psychophysics*, 58, 1260-1277.
- Gray, C. M., König, P., Engel, A.K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334-337.
- Gross, J., Schmitz, F., Schnitzler, I., Kessler, K., Shapiro, K., Hommel, B., & Schnitzler, A. (2004). Long-range neural synchrony predicts temporal limitations of visual attention in humans. *Proceedings of the National Academy of Sciences USA*, 101, 13050-13055.
- Hari, R., & Salenius, S. (1999). Rhythmical corticomotor communication. *Neuroreport*, 10, R1-10.
- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
- Henderson, J. M. (1994). Two representational systems in dynamic visual identification. *Journal of Experimental Psychology: General*, 123, 410-426.
- Henderson, J. M., & Anes, M. D. (1994). Roles of object-file review and type priming in visual identification within and across eye-fixations. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 826-839.
- Hommel, B. (1994). Spontaneous decay of response code activation. *Psychological Research*, 56, 261-268.
- Hommel, B. (1996). S-R compatibility effects without response uncertainty. *Quarterly Journal of Experimental Psychology*, 49A, 546-571.
- Hommel, B. (1998). Event files: Evidence for automatic integration of stimulus response episodes. *Visual Cognition*, 5, 183-216.
- Hommel, B. (2002). Responding to object files: Automatic integration of spatial information revealed by stimulus-response compatibility effects. *Quarterly Journal of Experimental Psychology*, 55A, 567-580.
- Hommel, B. (2003). Feature integration across perception and action: Event files affect response choice. Manuscript under revision for publication.
- Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences*, 8, 494-500.
- Hommel, B., & Akyürek, E. (in press). Lag-1 Sparing in the Attentional Blink: Benefits and costs of integrating two events into a single episode. *Quarterly Journal of Experimental Psychology (A)*.
- Hommel, B., Alonso, D., & Fuentes, L. J. (2003). Acquisition and generalization of action effects. *Visual Cognition*, 10, 965-986.
- Hommel, B., & Colzato, L. S. (2004). Visual attention and the temporal dynamics of feature integration. *Visual Cognition*, 11, 483-521.
- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001a). The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24, 849-878.

- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001b). Codes and their vicissitudes. *Behavioral and Brain Sciences*, 24, 910-937.
- Hommel, B., Proctor, R. W., & Vu, K. P. L. (2004). A feature-integration account of sequential effects in the Simon task. *Psychological Research*, 68, 1-17.
- Houghton, G., & Tipper, S. P. (1994). A model of inhibitory mechanisms in selective attention. In D. Dagenbach & T. Carr (Eds.), *Inhibitory mechanisms in attention memory and language*. San Diego: Academic Press.
- Jellema, T., Maassen, G., & Perrett, D. I. (2004). Single cell integration of animate form, motion, and location in the superior temporal sulcus of the macaque monkey. *Cerebral Cortex*, 14, 781-790.
- Jellema, T., & Perrett, D. I. (2002). Coding of visible and hidden objects. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action: Attention and Performance XIX* (pp. 356-380). Oxford: Oxford University Press.
- Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology*, 36, 138-202.
- Jolicoeur, P., Dell'Acqua, R., & Crebolder, J. (2000). Multitasking performance deficits: Forging some links between the attentional blink and the psychological refractory period. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 309-330). Cambridge, MA: MIT Press.
- Jolicoeur, P., Tombu, M., Oriet, C., & Stevanovski, B. (2002). From perception to action: Making the connection. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action: Attention & performance XIX* (pp. 558-586). Oxford: Oxford University Press.
- Kahneman, D., & Henik, A. (1981). Perceptual organisation and Attention. In: M. Kubovy & J.R. Pomerantz (Eds.), *Perceptual organization*. Hillsdale, N.J.: Erlbaum, 181-211.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24, 175-219.
- Kane, M. J., May, C. P., Hasher, L., Rahhal, T., & Stoltzfus, E. R. (1997). Dual mechanisms of negative priming. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 632-650.
- Keil, A., Muller, M. M., Ray, W. J., Gruber, T., & Elbert, T. (1999). Human gamma band activity and perception of a Gestalt. *Journal of Neuroscience*, 19, 7152-7161.
- Kingstone, A. (1992). Combining expectancies. *Quarterly Journal of Experimental Psychology*, 44A, 69-104.
- Kirby, N. H. (1980). Sequential effects in choice reaction time. In A.T. Welford (Ed.), *Reaction Times* (pp. 129-172). London: Academic Press.
- Kopell, N., Ermentrout, G. B., Whittington, M. A., & Traub, R. D. (2000). Gamma rhythms and beta rhythms have different synchronization properties. *Proceedings of the National Academy of Science USA*, 97, 1867-1872.
- Kornblum, S. (1973). Sequential effects in choice RT: A tutorial review. In S. Kornblum (Ed.), *Attention & performance IV* (pp. 259-288). Hillsdale, NJ: Erlbaum.

- Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neuroscience*, 23, 571-579.
- Lamour, Y., Dutar, P., & Jobert, A. (1982). Excitatory effect of acetylcholine on different types of neurons in the first somatosensory neocortex of the rat: Laminar distribution and pharmacological characteristics. *Neuroscience*, 7, 1483-1494.
- Leslie, A. M., & Kaldy, Z. (2001). Indexing objects in infant working memory. *Journal of Experimental Child Psychology*, 78, 61-74.
- Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: Developing 'what' and 'where' systems. *Trends in Cognitive Sciences*, 2, 10-18.
- Levy R., Ashby P., Hutchison W.D., Lang A.E., Lozano A.M., & Dostrovsky J.O. (2002). Dependence of subthalamic nucleus oscillations on movement and dopamine in Parkinson's disease. *Brain*, 125, 1196-1209.
- Lisman, J. E., & Idiart, M. A. P. (1995) Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science*, 267, 1512-1515.
- Little, H. J. (1999). The contribution of electrophysiology to knowledge of the acute and chronic effects of ethanol. *Journal of Pharmacology and Experimental Therapeutics*, 84, 333-353.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.
- Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review*, 109, 376-400.
- Lowe, D. G. (1985). Further investigation of inhibitory mechanisms in attention. *Memory & Cognition*, 13, 74-80.
- Luck, S. J., & Beach, N. J. (1998). Visual attention and the binding problem: A neurophysiological perspective. In R. D. Wright (Ed.), *Visual Attention* (pp. 455-478). New York: Oxford University Press.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279-281.
- MacKay, X. (1997). Synchronized neuronal oscillations and their role in motor processes. *Trends in Cognitive Sciences*, 1, 176-183.
- Matsumura, M., Sawaguchi, T., & Kubota, K. (1990). Modulation of neuronal activities by iontophoretically applied catecholamines and acetylcholine in the primate motor cortex during a visual reaction-time task. *Journal of Neuroscience Research*, 8, 138-145.
- May, C. P., Kane, M. J., & Hasher, L. (1995). Determinants of negative priming. *Psychological Bulletin*, 118, 35-54.
- Maylor, E. A. (1985). Facilitatory and inhibitory components of orienting in visual space. In M.I. Posner, & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 189-204). Hillsdale, NJ: Erlbaum.
- McDevitt, J., & Fournier, L., (2001). Withholding a planned action delays execution of partially compatible actions. *Abstracts of the Psychonomic Society*, 6, 16.
- Mentis, M. J., Sunderland, T., Lai J., Connolly, C., Krasuki, J., Levine, B., Friz, J., Sobti, S., Schapiro, M., & Rapoport, M. D. (2001). Muscarinic versus nicotinic modulation of a visual task: a pet study using drug probes. *Neuropsychopharmacology*, 25, 555-564.

- Metherate, R., Cox, C. L., & Ashe, J. H. (1992). Cellular bases of neocortical activation: modulation of neural oscillations by the nucleus basalis and endogenous acetylcholine. *Journal of Neuroscience*, 12, 4701-4711.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, 16, 5154-5167.
- Miltner, W. H. R., Braun, C., Arnold, M., Witte, H., & Taub, E. (1999). Coherence of gamma-band EEG activity as a basis for associative learning. *Nature*, 397, 434-436.
- Minami, K., Vanderah, T. W., Minami, M., & Harris, R. A. (1997). Inhibitory effects of anesthetics and ethanol on muscarinic receptors expressed in *Xenopus* oocytes. *European Journal of Pharmacology*, 339, 237-244.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, 41, 49-100.
- Mondor, T. A., Hurlburt, J., & Thorne, L. (2003). Categorizing sounds by pitch: Effects of stimulus similarity and response repetition. *Perception & Psychophysics*, 65, 107-114.
- Morton, R. A., & Davies, C. H. (1997). Regulation of muscarinic acetylcholine receptor-mediated synaptic responses by adenosine in the rat hippocampus. *Journal of Physiology*, 502, 75-90.
- Mozer, M. C. (1991) *The perception of multiple objects: a connectionist approach*. Cambridge, MA: MIT Press.
- Munk M. H. J. (2003). The principle of controlling neuronal dynamics in neocortex: Rapid reorganization and consolidation of neuronal assemblies. In: R. H. Kluwe, G. Lüer, F. Rösler (Eds.), *Principles of learning and memory* (pp. 187-205). Basel: Birkhäuser Verlag.
- Murthy, V. N., & Fetz, E. E. (1992). Coherent 25 to 35 Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proceedings of the National Academy of Science USA*, 89, 5670-5674.
- Müsseler, J., & Hommel, B. (1997). Blindness to response-compatible stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 861-872.
- Nehlig, A., Daval, J., & Debry, G. (1992). Caffeine and the central nervous system: Mechanisms of action, biochemical, metabolic and psychostimulant effects. *Brain Research Reviews*, 17, 139-170.
- Neill, W. T. (1997). Episodic retrieval in negative priming and repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 1291-1305.
- Norman, D. A. (1968). Towards a theory of memory and attention. *Psychological Review*, 75, 522- 536.
- Park, J., & Kanwisher, N. (1994). Negative priming for spatial locations: Identity mismatching, not distractor inhibition. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 613-623.
- Phaf, R. H., van der Heijden, A. H. C., & Hudson, P. T. W. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology*, 22, 273-341.
- Porkka-Heiskanen, T., Strecker, R. E., Thakkar, M., Bjorkum, A. A., Greene, R. W., & McCarley, R.W. (1997): Adenosine: A mediator of the sleep-

- inducing effects of prolonged wakefulness. *Science*, 276, 1265-1268.
- Portas, C. M., Thakkar, M., Rainnie, D. G., Greene, R. W., & McCarley, R. W. (1997). Role of adenosine in behavioral state modulation: A microdialysis study in the freely moving cat. *Neuroscience*, 79, 225-235.
- Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. Bowhuis (Eds.), *Attention and Performance X* (pp. 531-556). Hillsdale, NJ: Erlbaum.
- Pratt, J., & Hommel, B. (2003). Symbolic control of visual attention: The role of working memory and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 835-845.
- Raffone, A., & Wolters, G. (2001). A cortical mechanism for binding in visual working memory. *Journal of Cognitive Neuroscience*, 13, 766-785.
- Ridderinkhof, K. R., de Vlugt, Y., Bramlage, A., Spaan, M., Elton, M., Snel, J., & Band, G. P. H. (2002). Alcohol consumption impairs detection of performance errors in mediofrontal cortex. *Science*, 298, 2209-2211.
- Riedel, W., Hogervorst, E., Lebourg, R., Verhey, F., van Praag, H., & Jolles, J. (1995). Caffeine attenuates scopolamine-induced memory impairment in humans. *Psychopharmacology*, 122, 158-168.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2, 1019-1025.
- Rodriguez, R., Kallenbach, U., Singer, W., & Munk, M. H. (2000). Effects of cholinergic stimulation on gamma oscillations during visual responses in cat visual cortex. *European Journal of Neuroscience Supplement*, 12, 221.15.
- Rodriguez, R., Kallenbach, U., Singer, W., & Munk, M. H. (2001). Reduction of gamma-oscillatory responses after scopolamine application in cat visual cortex. *Abstract Society of Neuroscience*, 27, 123.3.
- Rodriguez, R., Kallenbach, U., Singer, W., & Munk, M. H. (2004). Short- and long-term effects of cholinergic modulation on gamma oscillations and response synchronization in the visual cortex. *Journal of Neuroscience*, 24 (46), 10369-10379.
- Roelfsema, P. R., Engel, A. K., Koenig, P., & Singer, W. (1997). Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature*, 385, 157-161.
- Rumelhart, D. E., & Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science*, 9, 75-112.
- Sanna, E., Dildy-Mayfield, J. E., & Harris, R. A. (1994). Ethanol inhibits the function of 5-hydroxytryptamine type 1c and muscarinic-M1 G protein-linked receptors in *Xenopus* oocytes expressing brain messenger mRNA: role of protein kinase C. *Molecular Pharmacology*, 45, 1004-1012.
- Singer, W. (1994). The organization of sensory motor representations in the Neocortex: A hypothesis based on temporal coding. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 77-107). Cambridge, MA: MIT Press.
- Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24, 49-65.
- Smith, A., Brice, C., Nash, J., Rich, N., & Nutt, D. J. (2003) Caffeine and

- central noradrenaline: Effects on mood, cognitive performance, eye movements and cardiovascular function. *Journal of Psychopharmacology*, 17, 283-292.
- Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception & Performance*, 11, 598-616.
- Sorimachi, M., Yamagami K., & Nishimura S. (1992). A muscarinic receptor agonist mobilizes Ca^{2+} from caffeine and inositol-1, 4, 5-trisphosphate-sensitive Ca^{2+} stores in cat adrenal chromaffin cells. *Brain Research*, 571, 154-158.
- Stoet, G., & Hommel, B. (1999). Action planning and the temporal binding of response codes. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 1625-1640.
- Stoet, G., & Hommel, B. (2002). Interaction between feature binding in perception and action. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action: Attention & Performance XIX* (pp. 538-552). Oxford: Oxford University Press.
- Styles, E. A., & Allport, D. A. (1986). Perceptual integration of identity, location and colour. *Psychological Research*, 48, 189-200.
- Tallon-Baudry, C. & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*, 3, 151-162.
- Tallon-Baudry, C., Kreiter, A., & Bertrand, O. (1999). Sustained and transient oscillatory responses in the gamma and beta bands in a visual short-term memory task in humans. *Visual Neuroscience*, 16, 449-459.
- Tanaka, K., Saito, H., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology*, 66, 170-189.
- Taylor, T. L., & Klein, R. (2000). Visual and motor effects in inhibition of return. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 1639-1656.
- Tipper, S. P. (2001). Does negative priming reflect inhibitory mechanisms? A review and integration of conflicting views. *Quarterly Journal of Experimental Psychology*, 54A, 321-343.
- Tononi, G., Sporns, O., & Edelman, G. M. (1992). Reentry and the problem of integration of multiple cortical areas: Simulation of dynamic integration in the visual system. *Cerebral Cortex*, 2, 310-335.
- Treisman, A. (1988). Features and objects: The fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, 40A, 201-237.
- Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6, 171-178.
- Treisman, A. (1998). The perception of features and objects. In R. D. Wright (Ed.), *Visual Attention* (pp. 26-54). New York: Oxford University Press.
- Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14, 107-141.
- Van der Heijden, A. H. C. (1992). *Selective attention in vision*. London: Routledge.

- Van der Heijden, A. H. C. (1995). Modularity and attention. *Visual Cognition*, 2, 269-302.
- Van der Heijden, A. H. C. (2003). *Attention in vision: Perception, communication and action*. London: Routledge and Psychology Press.
- Van der Velde, F., & de Kamps, M. (2002). Synchrony in the eye of the beholder: An analysis of the role of neural synchronization in cognitive processes. *Brain and Mind*, 3, 291-312.
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 92-114.
- Von der Malsburg, C. (1981). The correlation theory of brain function. Internal report 81-2. Göttingen: Max Planck Institute for Biophysical Chemistry.
- Von der Malsburg C. (1999). The what and why of binding: The modeler's perspective. *Neuron*, 24, 95-104.
- Waszak, F., Hommel, B., & Allport, A. (2003). Task-switching and long-term priming: Role of episodic stimulus-task bindings in task-shift costs. *Cognitive Psychology*, 46, 361-413.
- Wickens, J., Hyland, B., Anson, G. (1994). Cortical Cell Assemblies: A possible mechanism for motor programs. *Journal of Motor Behavior*, 26, 66-82.
- Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, 1, 202-238.
- Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science*, 256, 1327-1331.

Samenvatting

De vraag hoe ons brein in staat is om de feature codes van één gebeurtenis correct te integreren staat bekend als het 'binding probleem'. Kahneman, Treisman en Gibbs (1992) veronderstelden dat bij visuele perceptie de feature codes van een object geïntegreerd worden in een 'object file', een tijdelijke cognitieve structuur, geadresseerd naar locatie, die alle perceptuele informatie over een object bevat en wellicht ook andere informatie (bijvoorbeeld semantische). Nader onderzoek heeft uitgewezen dat feature-integratie niet alleen optreedt bij visuele perceptie, het domein waarvoor het binding probleem oorspronkelijk geformuleerd werd (zie Treisman, 1996, voor een overzicht), maar ook een rol speelt bij andere vormen van perceptie, bij het plannen van acties en bij stimulus-respons associaties (Hommel, 2004). Hommel (1998) liet zien dat features van toevallig gepaarde stimuli en responses spontaan aan elkaar gekoppeld worden. Stoet en Hommel (1999) toonden aan dat een ruimtelijke feature, wanneer deze eenmaal gekoppeld is aan een bepaald plan, minder beschikbaar is om aan een ander plan gekoppeld te worden.

Een vraag die de gemoederen op dit moment bezighoudt is wat de noodzaak is van feature-binding. Sommige auteurs benadrukken dat integratieprocessen onontbeerlijk zijn in gedistribueerde representatieve systemen zoals het menselijke brein (bijv. Cowey, 1995). Anderen trekken de noodzaak van feature-binding in twijfel. Het SLAM-model van Phaf, van der Heijden en Hudson (1995) is bijvoorbeeld een hiërarchisch maar parallel model van selectieve aandacht. In het SLAM-model is geen sprake van een binding probleem omdat het model gebaseerd is op conjunctie detector modulen, modulen die elke mogelijke combinatie van features coderen. Jellema, Maassen en Perrett (2004) troffen in de superieur temporele cortex van de makaak-aap daadwerkelijk cellen aan die informatie over de vorm, beweging en locatie van levende objecten integreren.

Er is tevens een discussie gaande over de mechanismen die ten grondslag liggen aan feature binding. Sommige onderzoekers denken dat de synchronisatie van activiteit van celpopulaties mogelijk een rol speelt bij het koppelen van features die bij dezelfde gebeurtenis horen (Singer, 1994; Treisman, 1996). Andere onderzoekers plaatsen vraagtekens bij deze opvatting. Van der Velde en de Kamps (2002) stellen bijvoorbeeld dat modellen die gebaseerd zijn op synchronisatie van activiteit niet over de systematiciteit en productiviteit van echte compositionele systemen beschikken.

In het voorliggende onderzoek is een pragmatisch standpunt ingenomen ten aanzien van de bovengenoemde discussies. Omdat er in de literatuur een diversiteit aan binding-effecten is gerapporteerd, wordt er simpelweg van uit gegaan dat binding een bestaand fenomeen is. De kenmerken van dit fenomeen worden op empirische wijze onderzocht. Wat nu volgt zijn de belangrijkste bevindingen van dit onderzoek.

De data die in hoofdstuk één en vier besproken worden laten zien dat feature integratie niet globaal is, zoals Kahneman et al. (1992) veronderstelden, maar dat feature binding afhankelijk is van taak en taakinstructie. Niet alle features van een gegeven perceptie-actie gebeurtenis

worden met elkaar geïntegreerd. Wat geïntegreerd wordt lijkt bepaald te worden door taakrelevantie: varieert de gegeven feature op een dimensie die expliciet of impliciet gedefinieerd is als relevant voor de taak? Hoe dit mogelijk is, is een open vraag. Wij vermoeden dat een hogere-orde geïntegreerde representatie van de context belangrijk is voor het selecteren van de juiste respons.

Uit hoofdstuk één en vier blijkt verder dat de sterkte van een koppeling (binding) voorspeld kan worden uit de mate waarin de betrokken features geactiveerd zijn. De relevantie van de features voor het matchen van de geïmplementeerde aandachtsset lijkt hier de meest betrouwbare predictor: als de aandachtsset geoptimaliseerd wordt en gehandhaafd blijft, treedt er facilitatie op in de verwerking van de taakrelevante stimulus feature. Hoe sterker een relevante feature code geactiveerd wordt, des te groter is de kans dat deze geïntegreerd wordt met andere feature codes die voldoende geactiveerd zijn.

In hoofdstuk twee en drie wordt indirect bewijs geleverd voor het idee dat er een relatie bestaat tussen feature binding in de visuele cortex en neurale synchronisatie door het muscarinisch-cholinergische systeem. Synchronisatie van celpopulaties door het muscarinisch-cholinergische systeem speelt mogelijk een belangrijke rol bij het koppelen van features in de visuele cortex (Munk, 2003; Rodriguez, Kallenbach, Singer & Munk, 2000, 2001; Rodriguez-Bermudez, Kallenbach, Singer & Munk, 2004). Ons onderzoek wees uit dat koppelingen binnen objecten significant verzwakt worden door alcohol en versterkt worden door cafeïne (respectievelijk een muscarinisch-cholinergische antagonist en agonist). De koppelingen bleken niet beïnvloed te worden door nicotine, een stof die alleen effect heeft op het nicotinisch-cholinergische systeem.

Uit hoofdstuk twee, drie en vier blijkt dat feature binding berust op domein-specifieke mechanismen voor perceptie en actie. Geen van de drug-manipulaties was van invloed op de koppelingen tussen stimulus en respons features. Dit wijst erop dat koppelingen in visuele perceptie gevormd worden door cholinergisch-gedreven neurale synchronisatie in de gamma-frequentieband (zie hierboven; Engel & Singer, 2001; Keil, Muller, Ray, Gruber & Elbert, 1999), terwijl perceptie-actie koppelingen op andere wijze tot stand komen (d.w.z. via een ander aandrijvingssysteem of via een andere synchronisatie-frequentie). Er zijn inderdaad aanwijzingen dat lange-afstand koppelingen meer gepaard gaan met synchronisatie in de beta-frequentieband (Gross, Schmitz, Schnitzler, Kessler, Shapiro, Hommel & Schnitzler, 2004; Kopell, Ermentrout, Whittington & Traub, 2000; Roelfsema, Engel, Koenig & Singer, 1997). In hoofdstuk vier wordt tevens bewijs gepresenteerd tegen een mogelijk gemeenschappelijk integratiemechanisme. Er werden geen correlaties gevonden tussen feature-binding effecten in taken die binding maten in perceptie, in het plannen van acties en tussen perceptie en actie.

Tot slot wordt in hoofdstuk vijf aangetoond dat de objectheid (feature-conjunctie in bestaande objecten) van stimuli van invloed is op de manier waarop zij verwerkt worden en de manier waarop hun features geïntegreerd worden. De verwerking van vorm en kleur van simpele geometrische stimuli leidt tot een koppeling tussen de neurale codes die de features representeren. In tegenstelling tot bestaande objecten produceren de geometrische stimuli

slechts een koppeling gedurende enkele trials. Om deze resultaten te verklaren wordt betoogd dat neurale synchronisatie en conjunctie detectors geen wederzijds uitsluitende mechanismen zijn, maar mechanismen die het binding probleem op verschillende niveaus oplossen. Conjunctie detectors zijn in onze ogen serieuze kandidaten voor het coderen van gedragsrelevante of bekende perceptuele gebeurtenissen, ofwel voor het vastleggen van herhaaldelijk ervaren feature-combinaties in duurzame representaties (langetermijngeheugen). Synchronisatie is daarentegen vooral input-gebaseerd, en is daarmee een voorname kandidaat voor het flexibel coderen van willekeurige feature combinaties (ten behoeve van perceptuele processen en het kortetermijngeheugen).

Alles overziend levert deze dissertatie aanvullend bewijs voor de stelling dat feature binding een algemeen verschijnsel is dat zowel binnen als tussen domeinen optreedt. Feature binding wordt mogelijk gemaakt door domein-specifieke mechanismen. Feature binding in de visuele cortex gaat bijvoorbeeld gepaard met neurale synchronisatie door het muscarinisch-cholinergische systeem, maar stabiliteit van de koppeling over langere tijd lijkt vooral afhankelijk te zijn van de aan- of afwezigheid van langetermijn-representaties. Een belangrijke observatie is dat binding ook afhankelijk is van taak en taakinstructie. Feature binding is niet een automatisch, globaal mechanisme, maar hangt af van een taakrelevant criterium: niet alle features van een gegeven gebeurtenis worden met elkaar geïntegreerd en wat geïntegreerd wordt is afhankelijk van wat relevant is voor de taak.

Acknowledgments

Danke Baby per tutto il tuo amore, sei la luce della mia vita!
 Grazie Nanni, spero che dall'alto del cielo siate orgogliosi di me
 Grazie Ciccine (Dani e Sabri) per le serate al WU, per il film,
 il calendario, le foto, per avermi protetto dal Corvo, per avermi bene...
 Dank je Marina, jij bent zoals een moeder voor mij.
 Dank je Theo, jij bent me geleerd dat ik niet liefsle mag overkomen.
 Dank je Lisette, "it has done me good because of the color of the
 wheat field" ...
 Grazie Bob, Auto, Ale per essere le persone costanti della mia vita.
 Grazie Mara, Stefi, Tomy, Bay per saperne prendere per quella che sono...
 Dank je, mijn lieve en schatje mon Heertjes (Herd = Maxi / Helleke = Mich)
 Dank je Chris: je bent een lieverd!
 Thank you Teresa for teaching me that in life I have to take risks...
 Grazie Antonio per quel sogno, per il principio dell'impermanenza.
 Dank je Jous (Zik-zusje), Albertien, Wido, Sacha voor de steun.
 Danke Steffen (Ciccio) und Rico, ihr seid zwei Schätze!
 Grazie Laura, Graziano, Rosaria e Gioia per l'affetto... e i film...
 Dank je, labassistenten (vooral Rik), de pp en studenten (TJ & Kn).
 Grazie papà, mamma, Alex, fam. Corrado Bona e Zambotto a PD.

Grazie a tutti voi per aver creduto in me



Curriculum Vitae

Lorenza Serena Colzato was born on September 20, 1974, in Bolzano, Italy. In 1994, she took up her studies in psychology at the University of Padova (Italy), where she graduated in February 1999 in Experimental and Clinical Psychology (advisor: Niels Birbaumer). From March 1999 until September 2000, she worked as guest researcher at the Max Planck Institute for Psychological Research in Munich. From 1999 to 2003 she participated in a Master program in Humanistic Psychology in Rome (advisor: Alberto Zucconi), which she completed as recognized Client-Centered Psychotherapist. Since 2000 she works as a PhD student at the University of Leiden (promotor: Lex van der Heijden), where she also gives courses on the programming of on-line experiments.