Towards a Computational Account of Context Mediated Affective Stimulus-Response Translation

Pascal Haazebroek (phaazebroek@fsw.leidenuniv.nl)
Saskia Van Dantzig (sdantzig@fsw.leidenuniv.nl)
Bernhard Hommel (hommel@fsw.leidenuniv.nl)
Cognitive Psychology Unit & Leiden Institute for Brain and Cognition, Wassenaarseweg 52
Leiden, 2333 AK The Netherlands

Abstract

Affective stimuli can elicit fast and automatic responses. Traditional accounts assume a dual route mechanism with one route underlying automatic processing and one route underlying controlled processing. Recent studies show that automatic processes may be affected by top-down influences. To account for this interaction between automatic and controlled processing, we present a computational model. In this model, the cognitive representation of the task influences how stimuli and responses are encoded. As a result, the model can perform fast, automatic (affective) stimulus-response translation. This automatic processing is mediated by cognitive control.

Keywords: affective; stimulus-response compatibility; Simon effect; task context; Ideomotor theory; action effect; Theory of Event Coding; computational modeling; PDP.

Introduction

Affectively connotated stimuli, such as the crying face of your baby child or the hungry eyes of a wild predator during a sudden encounter on a safari trip, have the tendency to elicit fast and strong responses: before you realize it, you are caressing your child or running away from that tiger.

In order to account for these fast and automatic responses to affective stimuli, LeDoux (1996) proposed a dual route system. Within this system, the ‘low road’, associated with the amygdala, automatically translates stimuli to responses. In parallel with this subcortical pathway there is a ‘high road’, associated with the cortical structures of the brain. This pathway analyzes the stimulus in a more fine-grained, but slower way. Together, these routes enable someone to respond quickly to affective stimuli *and* to process these stimuli in more detail in order to adjust behavior at a later point in time. The ability to respond quickly to affectively connotated stimuli clearly has advantages for survival.

Empirical findings suggest that affective stimuli can automatically activate action tendencies related to approach and avoidance (e.g., Chen and Bargh, 1999). Such automatic response tendencies are not restricted to affective processing, but have also been found in studies on non-affective stimulus-response translation. Stimuli may facilitate responses that share features (e.g., location) with the stimuli. Such non-affective compatibility effects have also been explained by dual route models. (e.g., Kornblum, Hasbroucq, & Osman, 1990). In these models, one route typically reflects the controlled processing of task relevant features (e.g., stimulus color) and the other route reflects the automatic translation of task irrelevant features (e.g., stimulus location). Presenting a stimulus on a certain location automatically primes - via the automatic route - the compatible response location, yielding a stimulus-response compatibility effect known as the Simon effect (Simon & Rudell, 1967).

However, both in affective processing and in non-affective stimulus-response translation, there is evidence that automatic action tendencies are not completely impermeable to top-down influences. Recent studies on affective processing have demonstrated that automatic action tendencies are receptive to influences on a cognitive level, such as task relevance (e.g., Lavender and Hommel, 2007; Rotteveel and Phaf, 2004) and spatial reference frame of responses (Markman & Brendl, 2005). In a similar vein, findings on non-affective stimulus-response translation show that automatic translation depends on the exact task rule specification (Valle-Inclán & Redondo, 1998), stimulus coding, and response coding (Hommel, 1993; for an overview see Hommel (2000a).

To account for various types of interaction between perception and action, including stimulus-response compatibility effects, Hommel, Müsseler, Aschersleben, and Prinz (2001) formulated the Theory of Event Coding (TEC). Most notably, they proposed a level of common representations, where stimulus features and action features are coded by means of the same representational structures: ‘feature codes’. Feature codes refer to distal features of objects and events in the environment, such as distance, size and location, but on a remote, descriptive level, as opposed to the proximal features that are registered by the senses. Second, at this common codes level, stimulus perception and action planning are considered to be similar processes; both involve activating and integrating feature codes into complex structures called ‘event files’. Third, action features refer to the perceptual consequences of a motor action; when an action is executed, its perceptual effects are integrated into an event file, an action concept. Following the Ideomotor theory (James, 1890), one can plan an action by anticipating the features belonging to this action concept. As a result, actions can be planned voluntarily by intending their perceptual effects. Finally, TEC stresses the role of task context in stimulus and response coding. In particular, feature codes are “intentionally weighted” according to the action goal at hand (Fagioli, Hommel, & Shubotz, 2007).

According to TEC, actions are coded in terms of their effects, and compatibility effects may arise when there is congruence between a stimulus and an action effect. This was demonstrated by Beckers, De Houwer and Eelen (2002). During the first phase of their experiment, participants learned that one particular action consistently resulted in a negative effect (a mild electroshock) while another action had no such consequence. In the subsequent test phase participants categorized affective words according to their grammatical category (noun or verb). Participants responded by performing the same two actions as used in the first phase. Responses associated with a negative action effect were performed faster in response to negative words than to positive words, even though word valence was irrelevant for the task at hand. Based on this finding, the authors concluded that action concepts in TEC include affective features. In order to account for their results, they suggest that the affective connotations of both stimuli and action effects are automatically represented on a semantic level, where they interact and yield a compatibility effect.

In this paper we re-examine the findings of Beckers et al. and discuss how their results can be reconciled with a more task-oriented account (see also Hommel, 2000b). More specifically, we demonstrate how the cognitive representation of the task biases the cognitive system to encode stimuli and responses in terms of valence. As a result of this top-down influence, the system automatically processes the (task-irrelevant) affective features of stimuli and responses, which leads to an affective stimulus-response compatibility effect

In order to computationally specify the mechanisms proposed in TEC and to validate its principles and assumptions by means of simulations, we are developing the HiTEC architecture (Haazebroek & Hommel, submitted). HiTEC is a generic architecture that can be used to define more specific computational models of human perception and action control and that can serve as a starting point for a novel control architecture for cognitive robots in the PACO-PLUS project (www.paco-plus.org).

In the following, we will first describe the HiTEC architecture in terms of its structures and processes. Next, we show how a specific HiTEC model gives rise to the affective Simon effect as reported by Beckers et al. Finally, we discuss how this approach compares to traditional dual route accounts of affective stimulus-response translation.

HiTEC

The Theory of Event Coding provides a number of constraints on the structure and processes of the HiTEC architecture. First, we describe the general structure of HiTEC and its representations. Next, we elaborate on the processes operating on these representations, following the two-stage model for the acquisition of voluntary action control (Elsner and Hommel, 2001).

HiTEC's Structure and representations

HiTEC is architected as a connectionist network model that uses the basic building blocks of parallel distributed processing (PDP; Rumelhart, Hinton, & McClelland, 1986). In HiTEC, the elementary units are codes which can become associated. As illustrated in Figure 1, codes are organized into three main systems: the sensory system, the motor system and the common coding system. Each system will now be discussed in more detail.

Sensory System The primate brain encodes perceived objects in a distributed fashion: different features are processed and represented across different cortical maps (e.g., DeYoe & Van Essen, 1988). In HiTEC, different perceptual modalities (e.g., visual, auditory, tactile, proprioceptive) and different dimensions within each modality (e.g., visual color and shape, auditory location and pitch) are processed and represented in different sensory maps. Each sensory map is a module containing a number of sensory codes that are responsive to specific sensory features (e.g., a specific color or a specific pitch).

Note that Figure 1 shows only those sensory maps relevant for our current modeling purposes: (complex) visual shapes, tactile intensity and a proprioceptive direction map. However, other specific models based on the HiTEC architecture may include other sensory maps as well (e.g., auditory maps, visual color map, etc.).

Motor System The motor system contains motor codes, referring to proximal aspects of specific movements. Although motor codes could also be organized in multiple maps, in the present version of HiTEC we consider only one basic motor map with a set of motor codes.

Common Coding System According to TEC, both perceived events and action generated events are coded in one common representational domain (Hommel et al., 2001). In HiTEC, this is implemented in a common coding system that contains common feature codes. Feature codes refer to distal features of objects as opposed to the proximal features coded by the sensory codes and motor codes.

Feature codes may be associated to both sensory codes and motor codes and are therefore truly sensorimotor. They can combine information from different modalities and are in principle unlimited in number. Feature codes are not given but they evolve and change. In HiTEC simulations, however, we usually assume a set of feature codes to be present initially, to bootstrap the process of extracting sensorimotor regularities in interactions with the environment.

Feature codes are contained in feature dimensions. Feature dimensions may be enhanced as a whole. This makes each feature code within such a dimension more sensitive to stimulation originating from sensory or motor codes.

Associations In HiTEC, codes can become associated, both for short term and for long term. Short term associations between feature codes reflect that these codes 'belong together in the current task or context' and that their binding is actively maintained in working memory. In Figure 1, these temporary bindings are depicted as dashed lines. Long term associations can be interpreted as learned connections reflecting prior experience. These associations are depicted as solid lines.

Event Files Another central concept in the theory of event coding is the event file (Hommel, 2004). In HiTEC, the event file is modeled as a structure that temporarily associates to feature codes that 'belong together in the current context' in working memory. An event file serves both the perception of a stimulus as well as the planning of an action. When multiple events are present in working memory, choosing between these events (e.g., deciding between different action alternatives) is reflected by competition between the associated event files. This competition is computationally modeled by means of inhibitory associations between event files, depicted as solid lines with filled disk ends in Figure 1.

HiTEC's processes

How do associations between codes emerge? What mechanisms result of their interactions? And how do these mechanisms give rise to stimulus-response compatibility effects? Elsner and Hommel (2001) proposed a two-stage model for the acquisition of voluntary action control. For both stages, we now describe how processes take place in the HiTEC architecture. Next, we discuss how HiTEC allows for task preparation and fast stimulus-response translation.

Stage 1: Acquiring Action – Effect Associations Feature codes are perceptually grounded representations as they are derived by abstracting regularities in activations of sensory codes. However, associations between feature codes and motor codes reflect acquired knowledge of action-effect contingencies: motor codes m_i are activated, either because of some already existing action-effect associations or simply because of network noise (cf. motor babbling behavior of newborns). This leads to a change in the environment (e.g., the left hand suddenly touches an object) which is registered by sensory codes s_i. Activation propagates from sensory codes towards feature codes f_i. Eventually, these feature codes are integrated into an event file e_i which acts as an action concept. Subsequently, the cognitive system learns associations between the feature codes f_i belonging to this action concept and the motor code m_i that just led to the executed motor action. The weights of these associations depend on activation of the motor code and the feature code. Crucially, this allows the task context to influence the learning of action effects, by moderating the activation of certain feature codes. Due to this top-down moderation, task-relevant features are weighted more strongly than task-irrelevant features. Nonetheless, this does not exclude task-irrelevant but very salient action effects to become involved in strong associations as well.

Stage 2: Using Action – Effect Associations Once associations between motor codes and feature codes exist, they can be used to select and plan voluntary actions. Thus, by anticipating desired action effects, feature codes become active. Now, by integrating the feature codes into an action concept, the system can treat the features as constituting a desired state and propagate their activation towards associated motor codes. Initially, multiple motor codes m_i may become active as they typically fan out associations to multiple feature codes f_i. However, some motor codes will have more associated features that are also part of the active action concept and some of the $m_i – f_i$ associations may be stronger than others. In time, the network will converge towards a state where only one code m_i is strongly activated, which will lead to the selection of that motor action.

In addition to the mere selection of a motor action, feature codes also form the actual action plan that specifies (in distal terms) how the action should be executed: namely, in such a way the intended action effect features are realized. This action plan is kept active in working memory, allowing the system to monitor, evaluate and adjust the actual motor action.

Task Preparation In reaction-time experiments, participants typically receive a verbal instruction of the task. In HiTEC, a verbal task instruction can directly activate the respective feature codes by means of verbal labels. The cognitive system integrates these feature codes into an event file that is actively maintained in working memory. When the model receives several instructions to respond differently to various stimuli, different event files e_i are created and maintained for the various options. Due to the mutual inhibitory links between these event files, they will compete with each other during the task.

Stimulus-Response Translation When a stimulus in an experimental trial is presented, its sensory features will activate a set of feature codes, allowing activation to propagate towards one or more event files, already associated during task preparation. Competition takes place between these event files. Subsequently, activation propagates from event files to action effect features and motor codes, resulting in the execution and control of motor action.

Note that task preparation already sensitizes feature codes both for the to-be-perceived stimuli and for the to-be-planned responses. Therefore, the cognitive system is biased in perceiving stimuli and anticipating responses in terms of these feature codes. When feature codes for expected stimuli and anticipated responses overlap, stimulus-response compatibility effects can arise: when a feature code activated by the stimulus is also part of the event file of the correct response, planning this response is facilitated, yielding faster reactions. If, on the other hand, the feature code activated by the stimulus is part of the incorrect response, this increases the competition between action events, resulting in slower reactions.

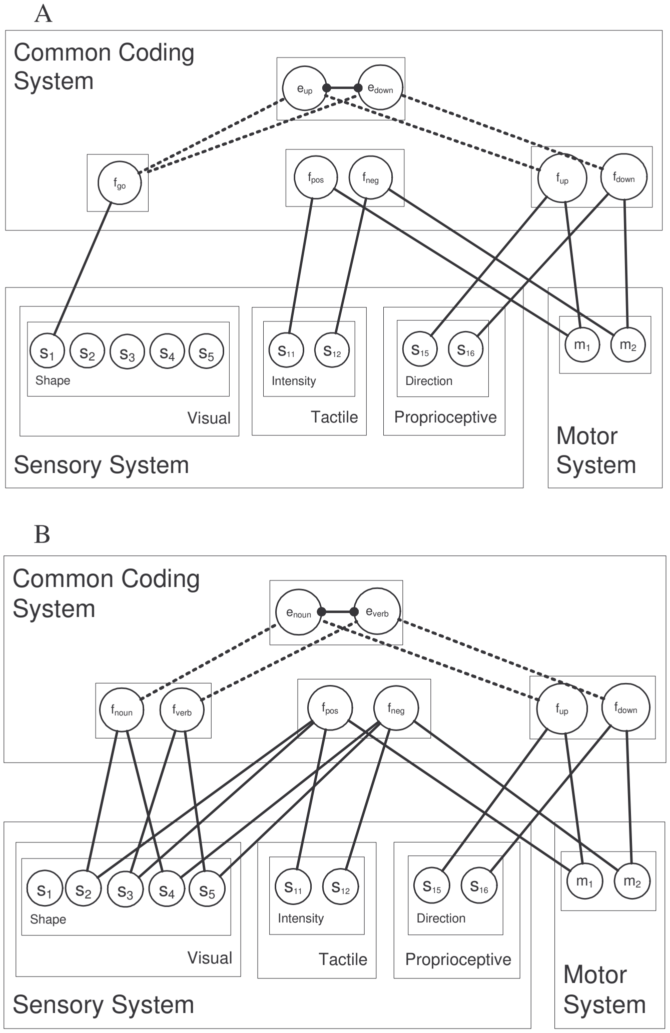

Figure 1: Experiment 1 modelled in the HiTEC architecture. (A) Training phase. (B) Test phase.

Affective Stimulus-Response Compatibility

In this section, we discuss how the results of Beckers et al (2002) could be replicated in a HiTEC model. In particular, we show how the task preparation configures the cognitive system to naturally yield the reported effects.

Experiment 1

The first experiment in Beckers et al. (2002) started with a training phase, during which participants responded to a ‘go’ signal by performing one of two actions (moving a response key up or down). One of the actions was consistently followed by a mild electroshock, whereas the other action was never followed by an electroshock. By repeating the actions many times, participants learned to associate them with their consequences. The response that was associated with the electroshock hence received a negative connotation, while the response associated with the absence of an electroshock received a positive connotation. In the subsequent test phase, participants had to classify positive and negative words according to their grammatical category (noun or verb). They responded by performing the same actions as in the training phase. Word valence, even though irrelevant for the grammatical judgment task, influenced response times. The ‘negative’ response (resulting in an electroshock) was performed faster in response to negative words than to positive words. In contrast, the ‘positive’ response (associated with the absence of a shock) was performed faster in response to positive words than to negative words.

In HiTEC, the verbal instruction that one of the responses yields an unpleasant electroshock will activate the valence feature dimension. In addition, the system is primed to expect a visual ‘go’ signal. In response to this ‘go’ signal, two actions can be performed. Therefore, two event files e_{up} and e_{down} are created, associated with the f_{up} and f_{down} feature codes, respectively. Note that we assume that these feature codes already exist, because the system has prior experience with upward and downward motion. Together, these codes and bindings form the task set for the training phase, as depicted in Figure 1.

During a trial of the training phase, a visual stimulus is presented as the ‘go’ signal, which is registered by sensory code s_1. Activation is propagated towards f_{go} and to both event files. The event files will start to compete. In the experiment by Beckers et al. participants were instructed to randomly choose one of the responses. This is modeled by adding random noise at the event file level, resulting in one strongly activated event file. Subsequently, activation propagates towards the f_{up} and f_{down} feature codes and to m_1 and m_2, resulting in the selection of one the motor actions.

When m_2 is executed, an electroshock is applied, which is registered by the S_{12} tactile sensory code. As this was expected on the feature level (due to the task set based on the verbal instruction), the shock is encoded as a strong activation of the f_{neg} feature code in the valence feature dimension. Now, action-effect learning takes place resulting in additional strengthening of $m_1 - f_{up}$ and $m_2 - f_{down}$ associations and the creation (and subsequent strengthening during subsequent trials) of $m_1 - f_{pos}$ and $m_2 - f_{neg}$ associations. It is assumed that the absence of an electroshock will be coded as f_{pos}, the opposite of f_{neg}.

During the test phase of Experiment 1, words are presented as stimuli. Clearly, there exist more than four words, but in this task all words are either noun or verb and either positively or negatively valenced. Thus, for modeling purposes, it suffices to work with four word shapes, as depicted in Figure 2.

When a word shape is presented, activation propagates towards the feature codes f_{noun} and f_{verb} depending on the grammatical category of the word. Simultaneously, activation propagates towards the valence feature codes f_{pos} and f_{neg}. Activation propagates from the grammatical category feature codes towards the event files e_{noun} and e_{verb}. This results in their mutual competition and subsequent propagation of activation towards the f_{up} and f_{down} and m_1 and m_2 codes. Because m_1 and m_2 are also associated with f_{pos} and f_{neg}, the action-effect associations acquired in the train phase, their activation is also impacted by activation propagated through the valence feature codes.

When a positive noun word is presented, activation will subsequently propagate from s_2 to f_{noun} to e_{noun} to f_{up} to m_1 and from s_2 to f_{pos} to m_1. As both pathways activate m_1, this results in faster action selection. When a negative noun word is presented, activation will similarly propagate from s_4 through feature codes and event files to m_1, but the 'valence route' will propagate activation through f_{neg} to m_2. This hampers selection of the correct m_1 motor action. For negative verbs, the implemented task set results in facilitation of the selection of m_2, and for positive verbs, selection of m_2 is interfered.

The overall result, thus, would resemble the findings of first experiment of Beckers et al.: correspondence between the irrelevant affective connotation of a word stimulus and the affective valence of the action effect produced by the required response result in faster performance than non-correspondence.

Experiment 2

In a second experiment, Beckers et al. (2002) attempted to increase the size of the affective congruency effect by making word valence more task-relevant. On the critical trials, participants still responded to the grammatical category of the words. The critical trials were interspersed with a valence judgment task, in which participants responded to the valence of the words by saying 'POSITIVE' or 'NEGATIVE'. Due to this manipulation, word valence was made more relevant, not only in the intervening valence judgment trials, but also in the critical trials. As a result, the affective compatibility effect found in Experiment 2 was larger than the effect found in Experiment 1, both numerically and in terms of effect size.

In HiTEC, during the critical, grammatical category judgment trials in Experiment 2, the activation propagates as in Experiment 1. During the valence judgment trials, affective word connotation is processed attentively resulting in strong sensitization of valence features f_{pos} and f_{neg}.

As valence judgment trials and grammatical category judgment trials are more or less alternated, this sensitization will carry-over to the critical trials, resulting in a stronger affective stimulus-compatibility effect, as reported by Beckers et al.

Discussion

We have introduced HiTEC's three main interacting modules: the sensory system, the motor system, and the emergent common coding system. Crucially, because the model uses common codes for stimulus perception and response planning, stimulus-response compatibility effects follow naturally. Subsequently, we argued that action effects may include affective features, as has been suggested by Beckers et al. (2002) and Lavender and Hommel (2007).

In our efforts to account for the results of Beckers et al., we stressed the role of task implementation for both stimulus and response coding. We showed how these codings affect translation of both task-relevant and task-irrelevant features. According to our account, affective information is not extracted completely automatically, but is mediated by task context. For example, in Experiment 1, although word valence is considered task irrelevant, the valence dimension is primed by the verbal 'unpleasant' task instruction accompanied by the actual action effect, the electroshock. As a result, word processing is biased toward affective coding.

Under normal circumstances, enhancing feature dimensions for stimulus perception in order to translate quickly to action features is a very efficient mechanism to control our actions: the location of a cup is very relevant for the action plan to grasp it. In this experimental setup, however, this mechanism results in a stimulus-response compatibility effect.

In contrast to our model, traditional dual route models stress the distinction between controlled and automatic processing, thereby ignoring task context for the automatic translation. We have shown that by assigning a task-set preparatory role to cognitive processes, the automatic route can be cognitively controlled and fast at the same time.

Following this consideration, and in analogy to empirical evidence in non-affective stimulus-response translation (Hommel, 1993), it would be interesting to conduct an alternative version of Experiment 1 of Beckers et al. In this alternative experiment, instead of an electroshock, a more ambiguous action effect could be presented, such as a moderately uncomfortable auditory tone. Half of the participants could be instructed that this tone is an 'unpleasant tone'. This instruction should yield the same results as the current experiment. The other half of the participants could be told that this is merely a 'high tone'. In this case, we expect that the affective stimulus-response compatibility effect will be reduced, if not eliminated, because the task is not cognitively represented in terms of valence. As a result, word processing is no longer biased toward valence coding.

Of course, in laboratory studies, task context is quite specific, which allows us to formalize it procedurally. In reality, there is typically not one specific task to solve. However, the environment may prime certain feature codes, based on our earlier experiences. This prepares us for particular stimuli we might perceive and actions we might have to perform. For example, walking through a dark, narrow alley primes our fearful anticipations of a sudden attack, making us jump at the slightest sound. The moment we enter a specific circumstance, sets of associated codes are primed. When we then actually perceive stimuli that resonate with these codes, the stimuli become cognitively represented in terms of these codes and are quickly translated into responses as if they are triggered completely automatically. While the ability to automatically translate stimuli into response tendencies is very useful in some contexts, it may lead to slightly less useful side-effects when applied in other domains: the automatic tendency to run away from a tiger facilitates our survival in the bushes, but lets us make a fool out of ourselves when applied in the movie theater.

Acknowledgments

Support for this research by the European Commission (PACO-PLUS, IST-FP6-IP-027657) is gratefully acknowledged.

References

Beckers, T., De Houwer, J., & Eelen, P. (2002). Automatic integration of non-perceptual action effect features: The case of the associative affective Simon effect. *Psychological Research*, *66*, 166–173.

Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin*, *25*, 215–224.

DeYoe, E. A., & Van Essen, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neuroscience, 11,* 219–226.

Elsner, B., & Hommel, B. (2001). Effect anticipation and action control. *Journal of Experimental Psychology: Human Perception and Performance*, *27*, 229–240.

Fagioli, S., Hommel, B., & Schubotz, R. I. (2007). Intentional control of attention: Action planning primes action-related stimulus dimensions. *Psychological Research*, *71*, 22–29.

Haazebroek, P., & Hommel, B. (submitted). HiTEC: A computational model of the interaction between perception and action.

Hommel, B. (1993). Inverting the simon effect by intention. *Psychological Research*, *55*, 270–279.

Hommel, B. (2000a). Intentional control of automatic stimulus-response translation. In Y. Rossetti & A. Revonsuo (eds.), *Interaction between dissociable conscious and nonconscious processes*. Amsterdam: John Benjamins Publishing Company.

Hommel, B. (2000b). The prepared reflex: Automaticity and control in stimulus-response translation. In S. Monsell & J. Driver (eds.), *Control of cognitive processes: Attention and performance XVIII*. Cambridge, MA: MIT Press.

Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences, 8*, 494-500.

Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, *24*, 849–937.

James, W. (1890). *The principles of psychology.* New York: Dover Publications.

Kornblum, S., Hasbroucq, T., & Osman, A. (1990). Dimensional overlap: Cognitive basis of stimulus-response compatibility—A model and taxonomy. *Psychological Review*, *97*, 253–170.

Lavender, T., & Hommel, B. (2007). Affect and action: Towards an event-coding account. *Cognition and Emotion*, *21*, 1270–1296.

LeDoux, J. (1996). The emotional brain: The mysterious underpinnings of emotional life. New York: Touchstone.

Markman, A. B., & Brendl, C. M. (2005). Constraining theories of embodied cognition. *Psychological Science*, *16*, 6–10.

Rotteveel, M., & Phaf, R. H. (2004). Automatic affective evaluation does not automatically predispose for arm flexion and extension. *Emotion*, *4*, 156–172.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, and the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Simon, J., & Rudell, A. (1967). Auditory s-r compatibility: The effect of an irrelevant cue on information processing. *Journal of Applied Psychology*, *51*, 300-304.

Valle-Inclán, F., & Redondo, M. (1998). On the automaticity of ipsilateral response activation in the Simon effect. *Psychophysiology*, *35*, 366–371.