Interaction between Task Oriented and Affective Information Processing in Cognitive Robotics

Pascal Haazebroek, Saskia van Dantzig, and Bernhard Hommel

Cognitive Psychology Unit, Leiden University
Wassenaarseweg 52. 2333 AK, Leiden, The Netherlands
{phaazebroek,sdantzig,hommel}@fsw.leidenuniv.nl

Abstract. There is an increasing interest in endowing robots with emotions. Robot control however is still often very task oriented. We present a cognitive architecture that allows the combination of and interaction between task representations and affective information processing. Our model is validated by comparing simulation results with empirical data from experimental psychology.

Keywords: Affective, Cognitive Architecture, Cognitive Robotics, Stimulus Response Compatibility, Psychology.

1 Introduction

An uplifting beep tone in moments of despair, a pair of artificial eyebrows showing an expression of genuine concern or a sudden decision to 'forget the rules' and 'save the girl' are common in Hollywood blockbuster movies that feature robots, but are currently not that realistic in everyday robot life. Typically, research in robot control focuses on the successful execution of tasks, such as grasping cups or playing the drums. The main goal of such research is to optimize task execution and to achieve reliable action control [1]. Increasingly, roboticists are also concerned with the social acceptance [2] of robots. A lot of effort is being put in the appearance of robots and their capability to display expressions that we may recognize as emotional. One may wonder, however, to what extent emotions (or affective information in general) may contribute to actual decision making [3].

In traditional machine learning approaches, such as reinforcement learning, affective information is usually treated as additional information that co-defines the desirability of a state (i.e., as a ‘reward’) or action alternative (i.e., as part of its ‘value’ or ‘utility’). By weighting action alternatives with this information, some can turn out to be more desirable than others, which can aid the process of decision making (e.g., [4]). In psychological literature, however, there is also evidence that affective information can influence how people respond to stimuli, by producing so-called compatibility effects. Empirical findings suggest, for example, that affective stimuli can automatically activate action tendencies related to approach and avoidance (e.g., Chen and Bargh [5]). The ability to respond quickly to affective stimuli clearly has advantages for survival, for humans and possibly for robots too.

M.H. Lamers and F.J. Verbeek (Eds.): HRPR 2010, LNICST 59, pp. 34–41, 2011.

In an empirical study by Beckers, De Houwer and Eelen [6], participants had to classify positive and negative words according to their grammatical category (noun or verb) by performing one of two actions (moving a response key up or down). Crucially, one of the responses systematically resulted in a mild but unpleasant electroshock. 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. This shows that actions are selected or executed more quickly when their effects are compatible with the affective valence of a stimulus than when they are incompatible.

In this paper we show how this experiment can be simulated in our computational HiTEC cognitive architecture [7] and thereby make it accessible for robot control. The general HiTEC architecture is described in section two. In section three we present the simulation results and finally, in section four, we discuss our findings and their implications for cognitive robotics.

2 HiTEC

2.1 Theory of Event Coding

The HiTEC cognitive architecture is based on the Theory of Event Coding (TEC), which was formulated by Hommel, Müsseler, Aschersleben and Prinz [8] to account for various types of interaction between perception and action, including stimulus-response compatibility effects. 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, stimulus perception and action planning are considered to be similar processes, as they both involve activating feature codes. Third, action features refer to the perceptual consequences of a motor action; when an action is executed, its perceptual effects are encoded by feature codes. Following the Ideomotor Theory of William James [9], actions can be planned voluntarily by intending their perceptual effects.

2.2 HiTEC’s Structure and Representations

HiTEC is implemented as a connectionist network model that uses the basic building blocks of parallel distributed processing (PDP) [10]. In HiTEC, the elementary units are codes that may be connected and are contained within maps. Codes within the same map compete for activation by means of lateral inhibitory connections. As illustrated in Figure 1, maps 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.

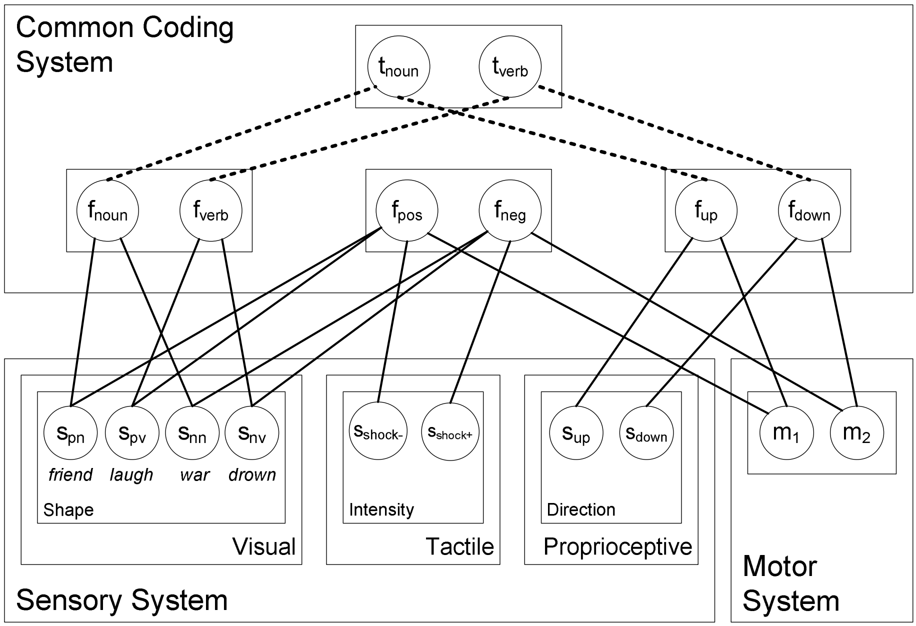

Fig. 1. HiTEC Architecture with Experiment 1 of Beckers et al. implemented. The smallest enclosing rectangles are maps (sensory, motor, feature and task maps) containing codes and lateral inhibitory connections (omitted from the figure for clarity).

Sensory System. The primate brain encodes perceived objects in a distributed fashion: different features are processed and represented across different cortical maps [11]. 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 format. In HiTEC, this is implemented in a common coding system that contains feature codes. Feature codes are perceptually grounded representations as they are derived by abstracting regularities in activations of sensory codes.

Task Codes. A task code is a structure at the common coding level that temporarily associates feature codes that 'belong together in the current context' in working memory. A task code serves both the perception of a stimulus as well as the planning of an action. When multiple task options are available, choosing between these options (e.g., deciding between different action alternatives) is reflected by competition between the task codes.

Associations. In HiTEC, codes can become associated, both for short term and for long term. In Figure 1, short-term task-related 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 in Figure 1.

2.3 HiTEC's Processes

Following the Ideomotor Theory [10], Elsner and Hommel [12] 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.

Stage 1: Acquiring Action – Effect Associations. In this stage, associations between feature codes and motor codes are explicitly learned. A random motor code is activated (comparable to the spontaneous '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. Activation propagates from sensory codes towards feature codes. Subsequently, the system forms associations between the active feature codes and the active motor code. The strength of these associations depends on the level of activation of both the motor code and the feature codes.

Stage 2: Using Action – Effect Associations. Once associations between motor codes and feature codes exist, they can be used to select and plan actions. Thus, by anticipating desired action effects, feature codes become active and propagate their activation towards associated motor codes. Initially, multiple motor codes may become active as they typically fan out associations to multiple feature codes. However, some motor codes will have more associated features and some of the associations between motor codes and feature codes may be stronger than others. In time, the network converges towards a state where only one motor code is strongly activated, which leads to the selection of that motor action.

Task Preparation. In reaction-time experiments, participants typically receive a verbal instruction of the task. In HiTEC, a verbal task instruction is assumed to directly activate the respective feature codes. The cognitive system connects these feature codes to task codes. When the model receives several instructions to respond differently to various stimuli, different task codes are recruited and maintained for the various options. Due to the mutual inhibitory links between these task codes, 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 task codes, already associated during task preparation.

Competition takes place between these task codes. Subsequently, activation propagates from task codes to action effect features and motor codes, resulting in the execution and control of motor action.

3 Affective Stimulus-Response Compatibility

In this section, we discuss how the results of Beckers et al. [6] can be replicated in a HiTEC model.

3.1 Model Configuration

The model, displayed in Figure 1, has three types of sensory codes; visual codes for detecting the words as (complex) visual shapes, tactile intensity sensory codes for registering the electroshock or its absence, and priopriоceptive sensory codes for detecting movement direction (action effect) of the response. Furthermore, it has three types of feature codes, representing grammatical category (noun or verb), valence (positive or negative) and direction (up or down). The task (respond to the grammatical category of a word by moving the key up or down) is internalized by creating a connectivity arc from the grammatical category feature codes through the task codes toward the direction feature codes (the dotted lines in Figure 1). The associations between valence feature codes and tactile codes are assumed, reflecting that the model already 'knows' that an electroshock is inherently experienced as unpleasant. The associations between word shapes and valence codes are also assumed, reflecting prior knowledge of the valence of certain words. In contrast to these fixed associations, the model has to learn the associations between valence codes and motor codes during the training phase. In other words, it has to learn which effects (shock or no shock) result from the different motor actions (moving the key up or down).

3.2 Training Phase

During a trial of the training phase, a motor action is randomly chosen and executed, which may result in a particular effect. For example, if m_2 is executed, an electroshock is applied, which is registered by the s_+ tactile sensory code. 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 strengthening of $m_1 - f_{up}$ and $m_2 - f_{down}$ associations and the creation (and subsequent strengthening during later trials) of $m_1 - f_{pos}$ and $m_2 - f_{neg}$ associations. It is assumed that the absence of an electroshock can indeed be coded as f_{pos}, the opposite of f_{neg}. In this way, over the course of 20 repetitions, the model learns the ideomotor assocations between motor codes and the activated feature codes.

3.3 Test Phase

The test phase consists of 40 experimental trials. In these trials, the model is presented a stimulus word (randomly one of the four possibilities) and has to give a motor response. Word 1 (e.g., "friend") and Word 2 ("laugh") are positive words, whereas

Word 3 (e.g., "war") and Word 4 (e.g., "drown") are negative. Word 1 and Word 3 were nouns and Word 2 and Word 4 were verbs.

During the test phase, 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 1.

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 task codes t_{noun} and t_{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}, through the action-effect associations acquired in the training phase, their activation is also influenced by activation propagated through the valence feature codes.

3.4 Simulation Results

When a positive noun (e.g., "friend") is presented, activation propagates from s_{pn} to f_{noun} to t_{noun} to f_{up} to m_1 , but also more directly from s_{pn} to f_{pos} to m_1. Because both the task-based pathway and the valence-based pathway activate m_1, this results in fast action selection. In contrast, when a negative noun (e.g., "war") is presented, activation propagates from s_{nn} through feature codes and task codes to m_1, while the valence-based pathway propagates activation through f_{neg} to m_2. Because both motor codes are now activated, competition arises, which hampers selection of the correct motor action. As a result, the model responds faster to positive nouns than to negative nouns. The reverse effect occurs for verbs. The selection of the correct motor code m_2 is facilitated by negative verbs (e.g., "drown"), and hampered by positive verbs (e.g., "laugh").

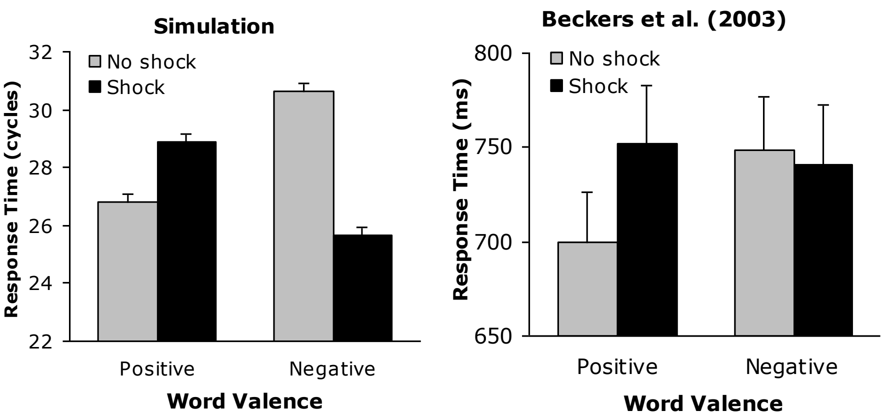

Fig. 2. Results of the HiTEC simulation (left) and the original results of Beckers et al. [6] (right)

The overall result, as can be seen in Figure 2, resembles the findings of the experiment of Beckers et al.: if the (task-irrelevant) affective valence of a word is compatible with the valence of the action-effect produced by the required response, performance is faster than if the word's valence is incompatible with the valence of the action-effect.

4 Discussion

We were able to replicate the affective compatibility effect reported by Beckers et al. [6] in HiTEC. A crucial aspect of this architecture is the common-coding principle: feature codes that are used to cognitively represent stimulus features (e.g., grammatical category, valence) are also used to represent action features (e.g., direction, valence). As a result, stimulus-response compatibility effects can arise; when a feature code activated by the stimulus is also part of the effect features belonging to 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 motor actions, resulting in slower reactions.

In addition, the task preparation influences the learning of action effects, by moderating the activation of certain feature codes through the binding between task codes and 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. In these simulations, this is clearly the case for valence features representing/resulting from the electroshock. As affective connotations often carry important information relevant for survival it can be assumed that other existing mechanisms moderate the sensitivity of affect related features. The mechanisms discussed in this paper account for how this influence may be applied in actual information processing.

In conclusion, response selection in HiTEC is not only based on 'rational' task-specific rules, but also on 'emotional' overlap between stimuli and responses. A robot endowed with such architecture may -on some day- actually forget the rules and save the girl.

Acknowledgments. Our thanks go to Antonino Raffone who was instrumental in the implementation of HiTEC and in conducting our simulations. Support for this research by the European Commission (PACO-PLUS, IST-FP6-IP-027657) is gratefully acknowledged.

References

1. Kraft, D., Baseski, E., Popovic, M., Batog, A.M., Kjær-Nielsen, A., Krüger, N., et al.: Exploration and Planning in a Three-Level Cognitive Architecture. In: International Conference on Cognitive Systems, CogSys (2008)
2. Breazeal, C.: Affective Interaction between Humans and Robots. In: Kelemen, J., Sosík, P. (eds.) ECAL 2001. LNCS (LNAI), vol. 2159, pp. 582–591. Springer, Heidelberg (2001)

3. Picard, R.W.: Affective Computing. MIT Press, Cambridge (1997)
4. Broekens, J., Haazebroek, P.: Emotion & Reinforcement: Affective Facial Expressions Facilitate Robot Learning. Adaptive Behavior (2003)
5. Chen, M., Bargh, J.A.: Consequences of Automatic Evaluation: Immediate Behavioral Predispositions to Approach or Avoid the Stimulus. Pers. Soc. Psychol. B. 25, 215–224 (1999)
6. Beckers, T., De Houwer, J., Eelen, P.: Automatic Integration of Non-Perceptual Action Effect Features: the Case of the Associative Affective Simon Effect. Psychol. Res. 66(3), 166–173 (2002)
7. Haazebroek, P., Raffone, A., Hommel, B.: HiTEC: A computational model of the interaction between perception and action (manuscript submitted for publication)
8. Hommel, B., Müsseler, J., Aschersleben, G., Prinz, W.: The Theory of Event Coding (TEC): A Framework for Perception and Action Planning. Behav. Brain Sci. 24, 849–937 (2001)
9. James, W.: The Principles of Psychology. Dover Publications, New York (1890)
10. Rumelhart, D.E., Hinton, G.E., McClelland, J.L.: A General Framework for Parallel Distributed Processing. In: Rumelhart, D.E., McClelland, J.L., the PDP (eds.) Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1, pp. 45–76. MIT Press, Cambridge (1986)
11. DeYoe, E.A., Van Essen, D.C.: Concurrent Processing Streams in Monkey Visual Cortex. Trends Neurosci. 11, 219–226 (1988)
12. Elsner, B., Hommel, B.: Effect Anticipation and Action Control. J. Exp. Psychol. Human 27, 229–240 (2001)