Peter König, Kai-Uwe Kühnberger, and Tim C. Kietzmann

# A Unifying Approach to High- and Low-Level Cognition

## Introduction

### Cognitive science on low- and high-level – A divided land

From its early beginnings to today, the interdisciplinary endeavor of cognitive science has led to a fundamentally improved understanding of many aspects of cognition. Some of this is due to the multi-leveled approach, as researchers have adopted a wide variety of techniques to understand cognitive phenomena at various levels of description. One way of distinguishing these different levels is by separating them into high- and low-level cognitive processes. Whereas the former includes cognitive abilities like planning and reasoning, the latter is generally seen as including the various modalities of sensory processing.

There are many reasons for such a seemingly principled division. For instance, low-level cognition, such as sensory processing, exists in virtually all animal-species, whereas high-level cognition, as described in more detail below, is mostly ascribed to human cognitive processing. In terms of bandwidth, vision, a low-level cognitive function, is the most dominant sensory modality in humans. Vision can be found in most species, specifically in all chordates (Land and Fernald 1992). In many of the latter, e.g. birds of prey, the spatial acuity of the visual system even surpasses human performance by a factor of 2 and more (Reymond 1985). Similar statements can be made with respect to other sensory modalities, such as audition. Sophisticated auditory systems are found in all chordates (Alexander 1981) and many species outperform human capabilities with respect to frequency range or sensitivity. Importantly, similarities across species can also be found with respect to the structures supporting sensory processing. For instance, although many different forms of receptors for optic signals can be found, lens-bearing eyes, as present in vertebrates, have evolved several times (Land and Fernald 1992; Nilsson 1989). Moreover, more proximal structures that support sensory processing exhibit similar organizational structures (Kaas 1997). From this it can be concluded that high performance sensory processing is a general capability, performed by most living species and that it is mostly based on related principles.

With regard to high-level cognition, there is no general definition or classification available and the typical assignment is mostly based on intuitions. Yet, researchers agree that logical reasoning, planning and language belong to its core capabilities (Thagard 2008). On a broader scope, decision-making, memory skills, creativity, general intelligence and social interactions are also mentioned in this context. Contrary to low-level cognitive processes, these capabilities are mostly thought of as being uniquely human. As studies comparing human and animal performance are still scarce, reports of intelligent animal behavior are greeted with great attention (Watanabe ND Huber 2006; Blaisdell et al. 2006). Summing up, the present state of research ascribes high-level cognitive processes primarily to human cognition.

A second case for the view of a principled division can be made by investigating in how far low- and high-level cognitive processes are approachable with modern information processing techniques. Whereas visual processing in artificial systems, again classified as a low-level cognitive function, can greatly benefit from our increased understanding of the cortical visual system (Pinto et al. 2008; Kietzmann et al. 2009), high-level cognition poses more challenging problems. In the realm of sensory processing, ideas flow back and forth between the two disciplines and the performance of artificial systems can be quantified and compared to human performance. Importantly, common belief holds that there are no principled obstacles to achieving near-perfect performance with artificial systems. In contrast to this, state-of-the-art computer systems targeting high-level cognitive capabilities, as defined above, do not (yet) resemble neuronal structures in the least. Doctor Dostert predicted that "five, perhaps three years hence, interlingual meaning conversion by electronic process in important functional areas of several languages may well be an accomplished fact"[1] (IBM 1954). After more than 50 years, with machines that are 100.000.000 times faster[2] problems originating from the domain of high-level cognition are still considered as extremely difficult, even though many of them are in fact considered simple by human standards. Impressive advances have been made in the context of well-defined artificial settings, e.g. chess playing, but artificial systems still perform poorly in high-level cognitive tasks that require a combination or integration of many specific high-level abilities. A good example is the usage of natural language, an ability that requires the integration of background knowledge, linguistic knowledge, reasoning, pragmatic aspects, gestures etc. Hence, the widely disparate progress approaching low- and high-level cognitive tasks in artificial systems underlines the view of a principled division.

---

[1] http://www-03.ibm.com/ibm/history/exhibits/701/701_translator.html.

[2] http://www.tomshardware.com/reviews/core-i7-990x-extreme-edition-gulftown,2874-6.html, http://www.ibm.com/ibm100/us/en/icons/ibm700series/.

Finally, cognitive phenomena can be divided into high- and low-level on the basis of their relation to our body and actions in the real world. Sensory processing seems to be necessarily connected to actions performed by natural agents of all developmental levels in real environments (Varela et al. 1991). Moreover, it has been proposed that knowledge of changes of sensory signals contingent on performed actions is constitutive of perceptual consciousness (O'Regan and Noe 2001). In contrast to this, high-level cognition, at least in humans, allows for abstract reasoning processes that are completely decoupled from concurrent or subsequent actions. This line of reasoning does of course not imply that the development of abstract reasoning abilities can evolve without grounding in real world actions. Nevertheless, reasoning may be performed without direct reference to actions in the world, once they are developed. Rodin's "the thinker", being deep in thoughts while completely immobilized, can be seen as a symbolization of this decoupling of high-level cognitive processing and actions in the world (Figure 1a). Contrary to this, the Braitenberg vehicle (Braitenberg 1984, Bach et al. 2007) represents a direct interaction of sensory processing and corresponding actions (Figure 1b)[3] and is therefore an example of pure low-level cognition.

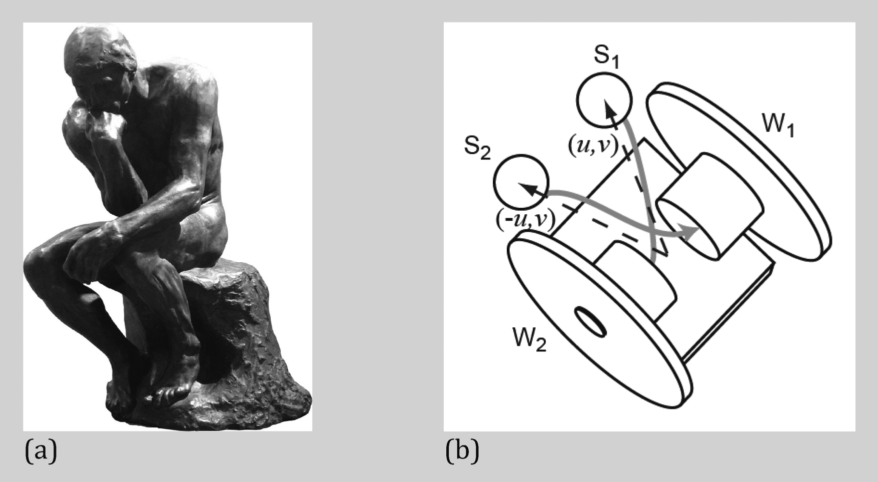

**Figure 1:** The difference between high- and low-level cognition can be symbolized by the contrast of "the thinker (Rodin)" and a Braitenberg vehicle.

---

[3] Figure 1b adapted from Bach (2009) with author's permission.

These arguments match the common intuitions of a fundamental gap between the stream of sensory input and the conceptual or symbolic level of interpretations. In fact, it still seems to be rather magic how analog and distributed sensory input can be “lifted“ to a symbolic level on which many high-level cognitive processes operate and, back again, how reasoning processes can be propagated to the actuator level. Taken together, these properties argue for a principled division of labor between two different cognitive systems employing qualitatively different algorithms suited for their respective problem domain.

A question that arises naturally from this division of labor is whether the two systems, high- and low-level cognition, share common neural substrates or whether their function is in fact reflected in distinct cortical systems. Evidence for the latter is based on the long tradition of investigations of functional losses after localized cortical lesions. Here it was demonstrated that lesions to different parts of the cerebral cortex result in selective loss or changes in functionality of high- or low-level cognitive processing capabilities. For instance, damage to the cortical region in the occipital pole was found to lead to a loss of visual function and perceptual awareness. Additionally, to name just a few, blindsight results from a lesion of early visual areas (Weiskrantz 1968, Stoerig and Cowey 1997); visual agnosia was shown to be the result of lesions to the occipitotemporal cortex (Farah 2000); prosopagnosia can be elicited by lesions to the inferior occipital cortex or fusiform gyrus (Steeves et al. 2009); akinetopsia by lesioning paroetal cortex (Zihl et al. 1983); achromatopsia by lesions to the ventral occipital cortex (Zeki 1990); and personality and behavior were shown to be affected by lesions to the frontal lobe (Barker 1995). Today, this view is complemented by studies applying modern imaging techniques, such as fMRI, PET, EEG, and MEG that demonstrate a functional compartmentalization in far reaching areas of cognition. Indeed, a fair part of the research effort focuses on the localization of cognitive functions and the remaining white spots on the cortical map are shrinking quickly. Thus, in addition to above observations, also the rapidly growing number of experiments that identify various functional specializations of different cortical regions speak in favor of a cortical dissociation between high- and low-level cognition.

## Re-unification by a statistical approach and embodiment

From the apparent differences in high- and low-level cognitive processes it could be proposed that both need to also rest on qualitatively different cortical operations and that the respective functional modules therefore exhibit structural differences. Indeed, individual areas have been delineated based on cortical structure in the form of lamination (Brodman 1909). Complementing this structural approach, more recent imaging techniques have been used to provide *functional* definitions of cortical modules (Felleman and van Essen 1991; Hilgetag and Barbas 2009). By taking both types of information together, it is now possible to validate the claim of a direct structure-to-function mapping. As a result, mostly early sensory areas and the primary motor areas were shown to be dissociable based on functional as well as structural information. However, many regions that can be functionally separable cannot be distinguished based on their anatomical structure. Among others, this holds for the large variety of areas in the intraparietal lobulus (LIP, VIP, MIP, PRR, AIP), which were shown to be functionally distinguishable despite all being situated within Brodman area 39. In addition to this structural similarity across different functionally defined areas, larger scale structures exhibit functional restructuring. This highly impressive capability of the human cortex was demonstrated for auditory information that can successfully be rerouted to visual cortex (Sur et al. 1988). It is also evident in blind subjects for whom the visual cortex seems to fulfill detailed sensory information processing during Braille reading (Sadato et al. 1996; Hamilton and Pascual-Leone 1998; Merabet et al. 2009). From this we can conclude that cortical modules are not limited to their primary function but can adapt to a wide variety of tasks. Importantly, it can be hypothesized that the quantitative properties typically used for a structural separation are more related to a fine-tuning of function, but not to qualitative differences of operations in the respective areas. That is, different functional specializations do not necessarily match one-to-one on different structural specializations. This allows for the proposition that operations performed in different cortical modules, including both high- and low-level cognition, are also not as distinct as the supported functions might suggest, but that they are in fact rather comparable. This resonates well with the concept of a canonical microcircuit (Douglas and Martin 2004), which holds that the structure of neocortical circuits is general and that neuronal circuits in neocortex can therefore be considered canonical.

Given these observations, we hypothesize that a similar approach can be taken for the description of high- and low-level cognition: Although both are based on different networks of functionally defined cortical regions, both types of cognitive processing may in fact implement comparable operations. As a result, *functional differences arise solely from different statistical properties of afferent signals and different context of those networks and not from inherently different structural properties*. Put differently, we argue that low-level sensory processing has many more similarities with high-level cognitive reasoning than previously assumed. To illustrate this admittedly bold claim, we concentrate on two central examples of low- and high-level cognition in the remainder of this article: invariant object recognition and analogical reasoning.

## An example of a low-level cognitive process: object recognition

### The hard problems of object recognition

Cognitive tasks of diverse complexity rely on a successful and reliable recognition of objects. For instance, consider the recognition of your car in a parking lot. Without problems you can identify it immediately in a large array of similar objects, despite different light-conditions, occlusions, different viewpoints depending on the direction in which you approach it, different retinal sizes that arise from different distances, as well as largely diverse background colors and clutter. Importantly, object identification is fast. Thus, even if timing is more crucial and the environment is more dynamic, as in the case of being part of a soccer game, we are immediately able to recognize the ball, independently of its color or texture, together with the goal and other players albeit their dramatic changes in shape upon movement. Finally, consider the more general case of object classification, as for instance in the case of classifying an animal as a dog. Despite the large variety of sizes, colors and types of dogs, we are very well able to successfully complete this task.

As these examples illustrate, both types of object recognition (identification and classification) belong to the most essential capabilities of the human visual system and prepare the grounds for many higher-level cognitive processes. Although we perform this task constantly and seemingly without effort, it is an extremely difficult problem from a computational point of view, as exemplified in above examples.

## Divide and conquer: processing in the visual hierarchy

How is the visual system set up in order to solve this complicated task? What cortical structures enable the system to fulfill the requirement of highly specific and at the same time robust classifications, i.e. to solve the specificity-vs-invariance problem? In the human brain, invariant object recognition is largely accomplished in the ventral visual stream (Haxby et al. 1991), which exhibits a hierarchical structure (Felleman and van Essen, 1991). Starting from retinal input, which passes through the Lateral Geniculate Nucleus (LGN), information enters striate cortex (V1), in which neurons are selective to bars of light and basic colors. Further downstream, information passes through areas in which neurons exhibit receptive fields of increasing complexity and size. These include the areas V2 and V4, in which color constancy is accomplished, and the lateral occipital complex, which is selective to spatially congruent, informative object parts (Lerner et al. 2008). Finally, information enters the cortical structures in the inferotemporal cortex. Here, cells exhibit selectivity for complex shapes including selective object views and view-invariant object representations (Tanaka 1996). Moreover, recent work demonstrates that neurons in the medial temporal lobe combine high selectivity for individuals with impressive invariance operations (Quiroga et al. 2005) and it is indeed possible to reliably classify and identify visual objects based on small populations of neurons in inferotemporal cortex (Hung et al. 2005). These results, and many more that cannot be covered in the scope of this article, paint the picture of a systematic division of labor within the ventral stream of the visual system. Following the hierarchy of visual areas, neurons exhibit more and more complex and at the same time increasingly robust response properties that lead to representations suitable for explicit object recognition.

Despite our increased understanding of the different neuronal selectivity across the ventral stream, however, it remains largely an unsolved question which principles underlie the development of these hierarchical representations and underlying cortical structures. Here, the class of normative approach is particularly promising, as detailed below.

## Optimality as a general statistical principle: from sparseness to stability

In recent years, an increasing number of studies explicitly addressed the variability of neuronal response properties by a normative approach. This notion dates back to Barlow's fundamental principle that neuronal representations should comply with the relevance for the animal, be suitable for decoding by downstream areas, and allow for efficient encoding by virtue of redundancy reduction (Barlow 1961). Specifically the latter endorses the approach that sensory processing should optimize mathematically defined criteria. These criteria are optimized for a given set of inputs, i.e. natural sensory stimuli. This rather different approach towards understanding sensory processing has mostly been studied in the visual domain, where computational models have successfully demonstrated the emergence of receptive field types exhibiting neuronal properties that are comparable to the ones found in the visual cortex. Notably, the normative approach is complementary to the experimental approaches: rather than measuring response properties of neurons in individual cortical areas, they are understood as the effect of unsupervised learning from natural input and its statistics.

The normative approach, which presupposes that neuronal representations optimize an objective function, requires a definition of the target properties. Following the requirement of efficient coding, optimality was first formalized on the basis of sparse representations. This implies that each representational unit specifically codes only for a small subset of the typical stimuli; i.e. neuronal receptive fields should be shaped in a way such that they lead to action potentials for only a small set of effective stimuli. Indeed, it has been found that the neuronal selectivity in the area V1 can be understood as adhering to a sparse code, given natural input (Olshausen and Field 1996). It was shown experimentally that natural stimuli evoke sparse activity patterns not only in V1, but also in higher visual cortices. Moreover, the application of a sparse coding scheme to intraareal interactions leads to functional coupling that is compatible with the lateral connectivity in V1 (Garrigues and Olshausen 2008). Finally, optimally sparse representations are closely related to independent component analysis, a statistical method suitable to infer the independent sources of a superposition of signals (Bell and Sejnowski 1997; Hyvärinen and Oja 2000).

Sparse coding leads by definition to high levels of specificity. This is due to the fact that sparseness enforces representations that react to only a small fraction of possible input. As seen in our earlier examples, however, specificity does not suffice for successful and robust object recognition, as invariance to sensory fluctuations and viewing conditions is equally important. This idea is picked up in another family of coding principles, which is based on the temporal continuity of natural stimuli. Despite changing implementations and names (stability, slowness, temporal coherence, etc.), the underlying assumption of these approaches is that relevant properties typically vary on a slower time-scale than irrelevant ones (Földiak 1991; Körding and König 2001; Wiskott and Sejnowski 2002). Thinking back to our first object recognition examples above (your car in a parking lot), what is common to all of the described complications is the fact that although sensory sampling differs largely from one situation to another, the identity of the object remains constant. Exactly this observation is capitalized upon with temporal coding schemes, which imply that the identity of an object changes on a slower timescale than the associated sensory information. Again, corresponding computational simulations targeting at striate cortex proved to be rather successful (Wiskott and Sejnowski 2002, Einhäuser et al. 2002; Körding et al. 2004, Berkes and Wiskott 2005). It was shown that stability does not only lead to simple-cell-like receptive field structures, but also that it can explain the phase invariance of complex cells. Although many questions are still open (Olshausen and Field 2005) the normative approach has led to significant progress in a principled understanding of primary visual cortex.

As an obvious next step, the normative approach was extended to higher visual cortices further down the ventral stream. For instance, it was shown that a stability-optimizing neural network increased the rotation invariance of the emerging representations, thereby enhancing recognition capabilities in a set of readout neurons (Einhäuser et al. 2005). Moreover, simulations of hierarchical networks based on the visual input of an artificial agent in a natural environment demonstrated the emergence of increasingly complex, yet stable visual representations. At the upmost hierarchical level, higher-level representations were shown to emerge that were selective to the position of the agent in space, but invariant with respect to its orientation (Wyss et al. 2006; Franzius et al. 2007). These matched properties of place cells as observed in the hippocampus (O'Keefe and Dostrovsky 1971). Finally, cells responsive for head-direction and spatial view-cells can be explained by the same set of principles (Franzius et al. 2007, Sprekeler and Wiskott 2011).

In addition to the computational work, important support for the stability approach was provided by electrophysiological experiments in which it was demonstrated that targeted changes of the temporal contiguity of objects lead to changes in response properties in inferotemporal neurons – a direct prediction of a neuronal coding scheme that is based on the stability principle (Li and DiCarlo 2010). These important results demonstrate that the normative approach does not only give a faithful description of neuronal response properties throughout the ventral stream, but that it also predicts the consequences of experimental manipulations. Moreover, the hierarchical application of the stability principle is a promising candidate in the attempt to close the gap to invariant object recognition. Thus, invariant object recognition and the development of neuronal response properties can be partly understood as a consequence of optimal sensory representations.

Besides to the well-studied visual domain, the normative approach has also been applied to other sensory modalities, such as auditory and somatosensory processing (Klein et al. 2003; Hipp et al. 2005; Duff et al. 2007). Taken together, these studies indicate that the wide variety of response properties on different levels of the visual system and of other modalities are fully compatible with a single set of principles governing sensory processing: Sparseness, slowness and decorrelation.

Above considerations are mostly based on the case of object recognition in which variations in sensory sampling originate from one object. However, it can also be argued that the resulting networks exhibit a most crucial new property: they can generalize from invariant object identification to the case in which different objects are associated with one label (object classification, our third example above). If object classification is understood as requiring invariance over object identities (as opposed to sensory variation), then the task could in principle be accomplished by the same normative approach as that lead to an increasing invariance over sensory sampling only. If this is indeed possible, then this implies that different objects of the same category share similar aspects of cortical representation. It has to be noted, however, that despite our ability to classify objects, we are still well able to recognize individual object instances under a great variety of conditions and viewpoints.

## Combining supervised and unsupervised learning schemes for successful object recognition

The renaissance of neural networks in the ’80s of the 20th century is tightly linked with the discovery and re-discovery of training methods for hierarchical neural networks (Werbos 1974; LeCun 1986; Rumelhart et al. 1986). How does an unsupervised training scheme, such as the normative approach described above, match with the typically utilized supervised algorithms of artificial neural networks? For the latter, the parameters and connection weights are iteratively tuned to match the output to the desired result. By now this work has expanded to a huge field and excellent reviews and textbooks are available (Bishop 2006). For the present purpose, however, we want to highlight a single specific problem only. These supervised learning procedures require labeled data, which are scarce and expensive in real life and thus might hinder proper convergence and generalization of the network structures. Hence, it is attractive to combine these methods with unsupervised learning, i.e. a normative approach described above. Indeed, applying unsupervised training to all layers of a hierarchical network but the last, and complementing this approach with supervised training of the output layer significantly reduces complexity of learning at a small price in performance only (Einhäuser et al. 2005, Franzius et al. 2008). Hence, the normative approach is fully compatible and fosters object recognition in hierarchical networks.

## From optimal sparseness, and optimal slowness to optimal predictability

The family of temporal coherence/slowness/stability approaches has been shown to explain many aspects of receptive field properties found in the visual hierarchy and thereby provides a principled approach for understanding invariant object recognition. However, if stability based on the statistics of natural input was the only objective function that is optimized in the mammalian cortex then the question arises why different species exhibit radically different sensory representations. Previously, we have put forward the hypothesis that sensory systems optimize the capabilities to predict and support sensory consequences of actions (König and Krüger 2006). Moreover, sensory selectivity should be shaped in a way such that they optimally support the potential actions of the agent. Because of this, neuronal representations should also be tuned to address those features that are optimally predictable with respect to the agent's actions. This entails the crucial step that the sensory predictability is integrated into the previously defined objective function (Weiler et al. 2010).

With the reference to different actions, the principle of predictability refers implicitly to the behavioral repertoire of the agent. Compared to the previously mentioned principles, this is a decisive step. Given that the visual systems of humans, non-human primates and carnivores differ in profound ways, relating visual processing to the behavioral repertoire opens a new avenue to understanding these differences.

## Towards optimal high-level processes: the example of analogies

How can we bring together the normative approach, which has been successfully applied in the domain of visual processing, a low-level cognitive function, with a high-level cognitive process such as the formation of analogies? In this section we will present our central claim that the principle of optimal action predictability and invariant actions supplies a unified framework of low-level and high-level cognitive functions. With this concept, we move from investigating purely sensory features to active representations that jointly address sensory information and the agent's action repertoire.

As an illustrative example, consider the case of the soccer ball from above. Kicking the ball requires the player to first recognize the individual ball, an item which was trained earlier, before any aiming or kicking can be accomplished (object recognition). However, if we were to swap the ball with a different one, it would nevertheless be possible for the player to immediately recognize the item to be kicked and to perform the appropriate action. Although this example might seem trivial at first, the performed computation is more complicated. This is because the player did not only generalize from one ball to a different exemplar (classification based on afforded actions), but also performed a generalized prediction of the consequences of the action. It is therefore an example of invariant actions. Importantly, despite the simplicity of the example, what has happened through the described generalization in sensory-motor space can in fact be seen as the drawing of an analogy. In the following sections, we will first describe the general research on (predictive) analogies before we describe the details on how both research areas can be understood on the basis of a unifying approach.

## (Predictive) analogies

Analogies are in the intense focus of research addressing high-level cognitive processes. Although an important topic in many disciplines for a long time, their scientific study in the context of cognitive science started with the seminal paper (Gentner 1983) introducing the *Structure-mapping theory*. This theory is based on the idea that establishing an analogical relation is a structural comparison of two domains, such that an "interesting" substructure in the source domain is aligned to an "interesting" substructure of the target domain. In other words, the forming of analogies relies on identifying commonalities of the two substructures. The structure-mapping theory has been proven its remarkable potential and is a de facto standard in cognitive models of analogy-making. Furthermore, the technical realization of the structure-mapping engine provides a standard computational model (Falkenhainer et al. 1998).

There are at least three classical domains from which typical examples of high-level cognition involving analogical reasoning are drawn: geometry (Evans 1968), naïve physics (Falkenhainer et al. 1998), and formal languages (Hofstadter et al. 1995). Besides these classical domains, however, many other domains have been discussed in the literature (intelligence tests, metaphorical expressions of natural language, problem solving, didactics of mathematics, sketch recognition etc.). Accompanying the variety of domains, researchers proposed a variety of different frameworks to account for the observed phenomena, ranging from symbolic models, like *Structure-mapping theory* or *Heuristic-driven-theory-projection* (Schwering et al. 2009), to neurally inspired frameworks, like *Learning and* *inference with schemas and analogies* (Hummel and Holyoak 1997) and hybrid approaches, like *Associative memory-based reasoning* (Kokinov and Petrov 2001). Despite its symbolic nature, *Heuristic-driven-theory-projection* and *Structure mapping theory* explicitly distinguish between low-level and high-level processes. *Associative memory-based reasoning* models all cognitive levels, but explicitly distinguishes between symbolic (reasoning-related) representations and neurally inspired activation spreading for attention and priming mechanisms. A similar separation holds in our opinion for the concept of *Learning and inference with schemas* and analogies. Hence, all of these frameworks accept the principled division of low-level and high-level cognitive processes.

An important class of analogies is given by so-called predictive analogies (Indurkhya 1992). Predictive analogies explain a new domain (target) by transferring information (knowledge) from the source to the target, such that non-trivial new conclusions can be drawn in the target domain. Because of this productive aspect, (predictive) analogies are often considered as a source of creativity and a mechanism for analogical, i.e. concept-guided, learning (Friedmann et al. 2009; Gust and Kühnberger 2006). For example, in the naïve physics domain, predictive analogies relate physical domains that are hardly accessible by our direct experience to domains that have perceivable properties. Due to the rich explanatory power supported by the source domain it is therefore possible to draw predictions in the target domain, which can then be experimentally evaluated. For example, let's consider an analogy between a water pipe system and an electric circuit. In the water pipe system, it can be observed that a "current" is triggered by "pressure" and that the system is necessarily closed. Another observation would be that "narrowing wires" influence the ongoing current. If we now apply these observations to previously learned concepts from the domain of electricity (the flow of electrons is triggered by a voltage difference, and a resistor influences the flow of electrons), the analogy is striking. Importantly, the analogy allows for the possibility of drawing new inferences. An example of such a prediction would be that adding a further resistance into the circuit should again reduce the flow of current. Notice that although these two exemplary domains were both chosen from the field of naïve physics, they do not show strong similarities, but are quite different from each other concerning the observable properties. Yet, the formation of analogies allowed for predictions from one domain to another.

## Object classification as predictive analogy

Now let us consider predictive analogies in the context of object recognition. First, it should be noted that analogies are already used in current visual sketch recognition systems and in systems designed for the recognition of geometric regularities in intelligence tests (Lovett et al. 2009; McLure 2010). Although such applications of analogy-making systems in the field of object recognition are rather new, some promising results for standardized tests in the geometry domain have already been achieved (e.g. Raven's progressive matrices; Lovett et al. 2010). Sketch-recognition systems for analogy-making typically work solely on the perceptual level, i.e. they identify important regions and features of such regions for an analogical comparison. In contrast to this, further aspects of cognition, like possible actions or functional properties of objects that are represented, do not play a role.

As an additional step, let us consider a case in which possible actions or action-outcomes can be included. Let us again consider the case in which we are presented with an object (e.g. a soccer ball) and need to classify it. Starting with the visual features of the object, it is possible to deduce properties that are relevant for an interaction with the object. We call this a 'predictive property' of an object. An example of such a property is "when kicked, it will roll". Furthermore, we might reasonably expect "when rolling on flat smooth ground, it will continue to roll for some time". However, this sequence of processing steps (from visual to predictive properties) is not necessarily required. Instead of starting with a visual analysis leading up to the recognition of an object and onwards to potential actions, we can also twist the argument and assign the primary importance to potential interactions with the object and thereby base object recognition on this set of afforded actions. Now, an object that satisfies the properties "when kicked, it will roll" and "when rolling on flat smooth ground, it will continue to roll for some time" is by (functional) definition a ball. With this, we have moved from a purely visual to a functional definition. Nevertheless, the provided functional definition of a ball can still be fulfilled via purely visual properties.

With the above case of a soccer ball, we have intentionally chosen an introductory example that is highly suitable for the classical approach that starts with a visual analysis that leads up to object recognition and only from there to functional predictions and we presented the functional definition as an alternative view. A visual definition can be seen in the well-known tradition in linguistics and logic of defining a concept by its intention, i.e. by the properties and attributes of the concept (Frege 1960). In the visual domain, such properties and attributes must be perceivable and as discriminative as possible. Important differences between this tradition and our proposal to define a ball functionally are the explicit emphasis of action-centered and manipulatory properties and the predictive character of such properties in the functional definition. Yet, already a slightly more complex example uncovers numerous problems with this purely visual approach. What is a chair? A quick look into Wikipedia gives a reasonable description: “a chair is a stable, raised surface used to sit on, commonly for use by one person”[4]. In normal circumstances, a raised surface can be defined based on visual features. A chair put upside down, however, has no raised surface anymore and thereby violates the visual definition – yet it is still a chair. On the other hand, a cube does have a raised surface, but it is usually not considered to be a chair. These problems leave us with the remainder of the definition: “to sit on”. This part puts the focus on the use of the object and is in essence a predictive analogy of the form “when you put your weight on it, you will not fall down”. In line with this, the Oxford Dictionary directly concentrates on the function of a chair, which is “a separate seat for one person, …”[5]. Notice that such problems occur necessarily with every intentional definition, because it is not possible to give a sufficient and necessary set of (visual) properties and attributes for classifying every potential instance correctly. Thus, as an alternative to a purely visual definition, we follow the approach that predictive analogies map functional connections and thereby are a vital part of the object definition.

While it has to be admitted that the original definition of *predictive* analogies does not perfectly fit into the domain of object recognition, successful performance in the visual domain nevertheless requires a transfer of (functional) knowledge from known examples to unseen ones in order to make the right classification and to select appropriate actions. Hence, although the original context of predictive analogies is in fact a different one, it seems unproblematic to call such analogies in the domain of object recognition to be predictive.

## How can we understand the emergence of analogies? – A unifying approach

Again, we start with a simple example: the case of driving your car to work in the morning and back again in the evening. Of course, while approaching your car in the morning, you recognize it albeit visually very dissimilar conditions. Moreover, (higher) cortical areas implement afforded actions (cyan). Once in the car, a specific action representation is activated (blue region) when you press the

---

[4] http://en.wikipedia.org/wiki/Chair.

[5] http://www.oxforddictionaries.com/definition/chair?view=uk.

right lever (gas pedal) with your foot and the sensory consequences of this action are predicted. This situation is visualized in Figure 2 in which part of the sensory representation (green area), as well as the action state (blue area, e.g. extension of the right leg and foot represented by changes in nodes 10 and 14) are altered by the active afforded actions (node 12 in cyan area). Because the now altered state, the afforded actions have changed as well (cyan area, pushing the break would now have an effect and opening the door is prohibited as represented by changes in activity of nodes 9 and 13). After work on the way home, it can be assumed that large parts, although not all, of active visual representations are identical to the ones activated in the morning[6], while some aspects might differ (e.g. no coffee in the cup holder (green node 7)). This has consequences on the sensory, afforded actions, and motor level. Yet, pressing the right pedal yields as expected the same effect and the sensory representation is transformed. In this case, the analogy is supported by largely overlapping sensory and action representations, which is in turn due to the performed invariant object recognition.

Compared to your own car, driving a different car to work introduces some more changes. For instance, the color and the geometry of the seat might be different. Yet, pressing the right pedal does lead to an acceleration and the predictions of sensory changes based on the experience on the former car are correct. This again can be considered as an analogy, which is supported by invariant object classification (different individual, same class) and its associated predictive properties. Indeed, we argue that despite an overwhelming amount of variance of sensory signals, the basic functionality is identical in both cases. Only because of this can the two objects be considered to belong to the same class. In this case the problem of invariant action representation has been transferred to invariant object classification. Now again consider driving a pellet jack. This might yield unexpected results although it does have a steering wheel, foot pedal and is part of the general category car. Yet, some come with a left foot accelerator pedal and using it in the usual way leads to a mismatch of predictions of the sensory consequences of actions and reality. Although the object is obviously a vehicle allowing a partly overlapping set of afforded actions (including pressing the right pedal), the result is not the same and the analogy breaks down. This demonstrates that the predictive framework is in fact working on probabilities and thus does not always allow for literal logical inferences.

Although our examples might lead to this view, it is in general not possible to neatly divide the different representational areas into sensory representa-

---

[6] For the sake of the argument and visualization we assume a highly sparse representation of sensory signals (green). Please note, however, that a coarse population code does not change the principle of the argument.

tions, afforded actions, and motor representations (figure 3D). To the contrary, a gradual transformation of sensory representations to motor actions leads to a setup in which neurons at each level act as “sensory” representations feeding bottom-up input to higher levels, modulate other potential affordances by tangential interactions at the same level and predict “action” induced changes of sensory representations at the lower level. The label of sensory/affordance/motor representations is therefore relative to the viewpoint. Still, even with this view the approach is fully compatible with a hierarchical network composed of general processing units according to a small set of optimization principles.

In our car example, the common set of afforded actions directly mirrors the invariant processing in the bottom-up pathway and sensory representations can be assumed to largely overlap. Hence, the prediction of sensory changes induced by the afforded action applies to the whole set of similar sensory representations. This is, however, not a necessary precondition. In general the afforded action is dependent only on a small part of the sensory representation and invariant with respect to other parts. This property defines it as an invariant action, which is at the core of making a predictive analogy.

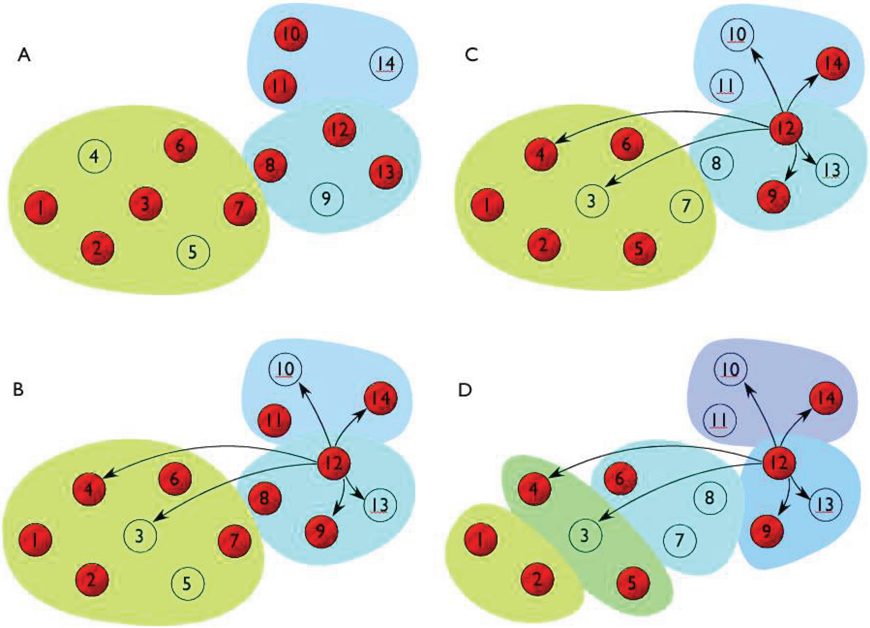

**Figure 2:** Schema of gradual transformation of sensory representation via affordances to motor actions. For detailed description see text.

## Easy and difficult analogies

In many everyday situations (and in the example above), the analogy comes quite natural such that, although being one, it is often not considered to be an analogy at all. In this section, we now consider more complex situations in which the overlap of sensory representations is not that large and the concept of invariant actions is more explicit. In classic examples, the basic constituents differ from each other in fundamental ways: To see this, consider, for example, the famous Rutherford analogy between the solar system, i.e. a system of planets revolving around a sun, and an analogous atom model, in which electrons are no longer homogeneously distributed as in the historically prior “plum pudding” model, but are revolving around a nucleus. In such examples, the overlap of sensory representations is minimal or even not existing and structural commonalities on a higher, i.e. abstract, level seem to be important. This brings us back to the origins of the scientific study of analogies in cognitive science in which rather abstract domains were considered. How are analogies emerging under such conditions that seem to be completely decoupled from any sensory representations? An explanation can be given by considering a situation in which the solar system – atom model analogy is visualized in form of diagrammatic representations (as in a scenario of a teaching situation in high school). In this form of representation, the constituents are in fact very similar to each other. There is a center, revolving objects represented as little circles and there are attracting and retracting forces etc., in short, the analogy is striking. It is rather uncontroversial that the emergence of the abstract conceptualization of a revolution movement is without any doubts grounded on a simpler, more concrete level and learned by using simpler, more concrete examples. Sensory representation, among other aspects as, *is* such a concrete level establishing a solid foundation of such generalizations.

## Object recognition, context, and actions

It is well-known that object recognition performances in psychological experiments change significantly, if the object in question is put into varied more or less prototypical contexts (flying eagle vs. sitting eagle, Zwaan et al. 2002). If object recognition has anything to do with establishing analogical relations, then context effects need to be considered also for analogies. How can context effects be transferred to the domain of analogy making? We suggest that contextual effects in the visual domain are quite often reducible to even more fundamental afforded actions and their representations. This is mainly due to the fact that perception and recognition tasks never occur without a dynamic environment and an active agent. For natural scenes, the coupling of the recognition of an object and action-related aspects is natural (eagles fly, planes fly as well, therefore they need wings etc.).

## Summary and conclusions

We propose that a key to overcome the artificial separation of low-level and high-level cognition is the concept of invariant actions, which optimally predicts action-induced changes of sensory signals. This concept is rooted in the ideas of Gibson (1977), yet makes crucial extensions. (1) To a first order of approximation, cortical processing is based on cortical modules of homogeneous structure. The function of these modules is to transmit optimally predictable parts of afferent signals to higher levels and to make predictions of changes of lower-level representations. Hence, functional differences originate mostly in the differences in input/output connectivity. (2) The optimization process leads to the emergence of invariant representations. Afforded actions emerge gradually in a hierarchical processing scheme obliviating a strict separation in sensory and motor representations. Focusing on the bottom-up direction, these might be viewed as invariant object representations, focusing on the top-down direction, these are invariant action representations. (3) Invariant actions are the core of predictive analogies. In most situations, the invariance is so natural that we emphasize invariant object recognition and do not realize that the implied actions are based on predictive analogies. The more arcane invariant actions, the classical examples of predictive analogies, are at the heart of higher cognitive functions. Together, these three steps establish “optimally predictive active representations” as a unified description and postulate a uniform cortical substrate and functional mechanisms for low-level and high-level cognitive processes.

## References

Bach, J., Bauer, C., & Vuine, R. (2007). MicroPsi: Contributions to a Broad Architecture of Cognition. *KI 2006: Advances in Artificial Intelligence. Lecture Notes in Computer Science* 4314. 7–18.

Barker, F. G. (1995). Phineas among the phrenologists: the American crowbar case and nineteenth-century theories of cerebral localization. *Journal of Neurosurgery* 82. 672–682.

Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In: Rosenblith, W. A. (ed.). *Sensory Communication*. Cambridge: MIT Press. 217–234.

Bell, A. J. & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Research* 37. 3327–3338.

Berkes, P. & Wiskott, L. (2005). Slow feature analysis yields a rich repertoire of complex cell properties. *Journal of Vision* 5(6). 579–602.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.

Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science* 311(5763). 1020–1022.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Leipzig: Johann Ambrosius Barth Verlag.

Douglas, R. J. & Martin, K. (2004). Neuronal circuits of the neocortex. *Annual Review of Neuroscience* 27. 419–51.

Einhäuser, W., Hipp, J., Eggert, J., Körner, E., & König, P. (2005). Learning viewpoint invariant object representations using a temporal coherence principle. *Biological Cybernetics* 93(1). 79–90.

Einhäuser, W., Kayser, C., König, P., & Körding, K. P. (2002). Learning the invariance properties of complex cells from their responses to natural stimuli. *European Journal of Neuroscience* 15(3). 475–86.

Evans, T. (1968). A program for the solution of a class of geometric-analogy intelligence-questions. In: Minsky, M. (ed.). *Semantic Information Processing*. Cambridge, MA: MIT press. 271–353.

Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 20(41). 1–63.

Farah, M. J. (2000). *The Cognitive Neuroscience of Vision. Fundamentals of Cognitive Neuroscience*. Malden, MA: Blackwell Publishers.

Felleman, D. J. & Essen, D. C. van (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* 1(1). 1–47.

Földiak, P. (1991). Learning invariance from transformation sequences. *Neural Computing* 3. 194–200.

Franzius, M., Sprekeler, H., & Wiskott, L. (2007). Slowness and sparseness lead to place, head-direction, and spatial-view cells. *PLOS Computational Biology* 3(8). e166.

Franzius, M., Wilbert, N., & Wiskott, L. (2008). Invariant Object Recognition with Slow Feature Analysis. In: *Artificial Neural Networks – ICANN 2008. Lecture Notes in Computer Science* 5163. 961–970.

Frege, G. (1960). On Sense and Reference. In: Geach, P. & Black, M. (eds.). *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Basil Blackwell.

Friedman, S., Taylor, J., & Forbus, K. (2009). Learning Naïve Physics Models by Analogical Generalization. In: Kokinov, B., Holyoak, K., & Gentner, D. (eds.), *New Frontiers in Analogy Research. Proceedings of the Second International Conference on Analogy*. Sofia: New Bulgarian University Press. 145–154.

Freiwald, W. A. & Tsao, D. Y. (2009). Cingulate cortex: diverging data from humans and monkeys. *Trends in Neurosciences* 32(11). 566–574.

Garrigues, P. & Olshausen, B. A. (2008). Learning Horizontal Connections in a Sparse Coding Model of Natural Images. In: Platt, J. C., Koller, D., Singer, Y., Roweis, S. (eds.), *Advances in Neural Information Processing Systems* 20. Cambridge, MA: MIT Press. 505–512.

Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science* 7(2), 155–170.

Gibson, J. J. (1977). The Theory of Affordances. In: Shaw, R. & Bransford, J. (eds.), *Perceiving, Acting, and Knowing*. New York: Wiley & Sons.

Gust, H. & Kühnberger, K.-U. (2006). Explaining Effective Learning by Analogical Reasoning. In: Sun, R. & Miyake, N. (eds.). *28th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. 1417–1422.

Hamilton, R. H. & Pascual-Leone, A. (1998). Cortical plasticity associated with Braille learning. *Trends in Cognitive Sciences* 2(5). 168–174.

Haxby, J. V., Grady, C. L., Horwitz, B., Ungerleider, L. G., Mishkin, M., Carson, R. E., Herscovitch, P., Schapiro, M. B., & Rapoport, S. I. (1991). Dissociation of object and spatial visual processing pathways in human extrastriate cortex. *Proceedings of the National Academy of Sciences USA* 88(5). 1621–1625.

Hilgetag, C. C. & Barbas, H. (2009). Sculpting the Brain. *Scientific American* 300. 66–71.

Hipp, J., Einhäuser, W., Conradt, J., & König, P. (2005). Learning of somatosensory representations for texture discrimination using a temporal coherence principle. *Network: Computation in Neural Systems* 16. 223–238.

Hofstadter, D. & the fluid analogies research group (1995). *Fluid Concepts and Creative Analogies*. New York: Basic Books.

Hummel, J. & Holyoak, K. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review* 104. 427–466.

Hung, C. P., Kreiman, G., Poggio, T., & DiCarlo J. J. (2005). Fast Readout of Object Identity from Macaque Inferior Temporal Cortex. *Science* 310(5749). 863–866.

Hyvärinen, A. & Oja, E. (2000). Independent Component Analysis: Algorithms and Application. *Neural Networks* 13(4–5). 411–430.

Indurkhya, B. (1992). *Metaphor and cognition*. Dodrecht: Kluver.

Kaas, J.H. (1997). Topographic maps are fundamental to sensory processing. *Brain Research Bulletin* 44(2), 107–112.

Kietzmann, T. C., Lange, S., & Riedmiller, M. (2009). Computational Object Recognition: A Biologically Motivated Approach. *Biological Cybernetics* 100. 59–79.

Klein, D. J., König, P., & Körding, K. P. (2003). Sparse spectrotemporal coding of sounds. *EURASIP Journal on Applied Signal Processing* 3. 659–667.

Kokinov, B. & Petrov, A. (2001). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In: Gentner, D., Holyoak, K., & Kokinov, B. (eds.). *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.

König, P. & Krüger, N. (2006). Symbols as self-emergent entities in an optimization process of feature extraction and predictions. *Biological Cybernetics* 94(4). 325–334.

Körding, K. P., Kayser, C., Einhäuser, W., & König, P. (2004). How are complex cell properties adapted to the statistics of natural stimuli. *Journal of Neurophysiology* 91(1). 206–212.

Körding, K. P. & König, P. (2001). Neurons with two sites of synaptic integration learn invariant representations. *Neural Computing* 13(12). 2823–2849.

Land, M. F. & Fernald, R. D. (1992). The Evolution of Eyes. *Annual Review of Neuroscience* 15. 1–29.

LeCun, Y. (1986). Learning Processes in an Asymmetric Threshold Network. In: Bienenstock, E., Fogelman-Soulié, F., & Weisbuch, G. (eds.). *Disordered systems and biological organization*. Les Houches, France: Springer. 233–240.

Lerner, Y., Epshtein, B., Ullman, S., & Malach, R. (2008). Class information predicts activation by object fragments in human object areas. *Journal of Cognitive Neuroscience* 20(7). 1189–1206.

Li, N. & DiCarlo, J. J. (2010). Unsupervised natural visual experience rapidly reshapes size-invariant object representation in inferior temporal cortex. *Neuron* 67(6). 1062–1075.

Lovett, A., Forbus, K., & Usher, J. (2010). A structure-mapping model of Raven's Progressive Matrices. *Proceedings of Cognitive Science* 10.

Lovett, A., Tomai, E., Forbus, K., & Usher, J. (2009). Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science* 33(7). 1192–1231.

McLure, M., Friedman, S., & Forbus, K. (2010). Learning concepts from sketches via analogical generalization and near-misses. In: Ohlsson, S. (ed.). *Proceedings of the 32nd Annual Conference of the Cognitive Science Society (CogSci)*. Portland, OR: Curran Associates, Inc.

McNeill Alexander, R. (1981). *The chordates*. 2nd edition. Cambridge: Cambridge University Press.

Merabet, L., Battelli, L., Obretenova, S., Maguire, S., Meijer, P., & Pascual-Leone, A. (2009). Functional Recruitment of Visual Cortex for sound encoded object identification in the Blind: A TMS Case Study. *NeuroReport* 20(2). 132–138.

Nilsson, D.-E. (1989). Vision Optics and Evolution. *BioScience* 39. 298–307.

Olshausen, B. A. & Field D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381. 607–609.

Olshausen, B. A. & Field, D. J. (2005) How close are we to understanding v1. *Neural computation* 17(8). 1665–1699.

O'Keefe, J. & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* 34(1). 171–175.

O'Regan, J. K. & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(5). 939–973.

Pinto, N., Cox, D. D., & DiCarlo, J. J. (2008). Why is real-world visual object recognition hard? *PLOS Computational Biology* 4(1). e27.

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature* 435(7045). 1102–1107.

Reymond, L. (1985). Spatial visual acuity of the eagle Aquila audax: a behavioural, optical and anatomical investigation. *Vision Research* 25. 1477–1491.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature* 323. 533–536.

Sadato, N., Pascual-Leone, A., Grafman, J., Ibañez, V., Deiber, M. P., Dold, G., & Hallett, M. (1996). Activation of the primary visual cortex by Braille reading in blind subjects. *Nature* 380(6574). 526–528.

Sprekeler, H. & Wiskott, L. (2011). A theory of slow feature analysis for transformation-based input signals with an application to complex cells. *Neural Computation* 23(2). 303–335.

Steeves, J., Dricot, L., Goltz, H., Sorger, B., Peters, J., Milner, D., Goodale, M.-A., Goebel, R., & Rossion, B. (2009). Abnormal face identity coding in the middle fusiform gyrus of two brain-damaged prosopagnosic patients. *Neuropsychologia* 47. 2584–2592.

Stoerig, P. & Cowey, A. (1997). Blindsight in man and monkey. *Brain* 120. 535–559.

Sur, M., Garraghty, P. E., & Roe, A. W. (1988). Experimentally induced visual projections into auditory thalamus and cortex. *Science* 242(4884). 1437–1441.

Schwering, A., Krumnack, U., Kühnberger, K.-U., & Gust, H. (2009). Syntactic Principles of Heuristic-Driven Theory Projection. *Cognitive Systems Research* 10(3). 251–269.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience* 19. 109–139.

Thagard, P. (2008). Cognitive Science. In: Edward N. Zalta (ed.). *The Stanford Encyclopedia of Philosophy*.

Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.

Watanabe, S. & Huber, L. (2006). Animal logics: decisions in the absence of human language. *Animal Cognition* 9(4). 235–245.

Weiller, D., Läer, L., Engel, A. K., & König, P. (2010). Unsupervised learning of reflexive and action-based affordances to model adaptive navigational behavior. *Front Neurorobotics* 4. 2.

Weiskrantz, L. (1986). *A Case Study and Implications*. Oxford: Oxford University Press.

Werbos, P. J. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Harvard University.

Wiskott, L. & Sejnowski, T. J. (2002). Slow feature analysis: unsupervised learning of invariances. *Neural Computation* 14(4). 715–770.

Wyss, R., König, P., & Verschure, P. F. M. J. (2006). A model of the ventral visual system based on temporal stability and local memory. *PLOS Biology 4*. e120.

Zeki, S. (1990). A century of cerebral achromatopsia. *Brain* 113(6). 1721–1777.

Zihl, J., Cramon, D. von, & Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain* 106. 313–340.

Zwaan, R., Stanfield, R., & Yaxley, R. (2002). Language comprehenders mentally represent the shapes of objects. *Psychological Science* 13(2).

**Prof. Dr. Peter König**
University of Osnabrück
Institute of Cognitive Science
Albrechtstraße 28
49069 Osnabrück
Germany
pkoenig@uos.de

**Prof. Dr. Kai-Uwe Kühnberger**
University of Osnabrück
Institute of Cognitive Science
Albrechtstraße 28
49069 Osnabrück
Germany
kkuehnbe@uos.de

**Tim C. Kietzmann, M. Sc.**
University of Osnabrück
Institute of Cognitive Science
Albrechtstraße 28
49069 Osnabrück
Germany
tkietzma@uos.de