On abstraction: psychological, neural, and computational perspectives

Igor Farkaš
Centre for Cognitive Science
Faculty of Mathematics, Physics and Informatics, Comenius University
Mlynská dolina, 84248 Bratislava
farkas@fmph.uniba.sk

Abstract

Abstraction is a core concept in cognitive science, representing a challenge for all theories of cognition. Conceptualization of abstraction is also complicated by the fact that it is an entity with several potential meanings and involved mechanisms. Abstraction occupies the agenda of many disciplines, including psychology, linguistics, artificial intelligence, and more recently, neuroscience. In this paper, we attempt to shed light on this topic, by summarizing evidence accumulated in these disciplines.

1 Introduction

Abstraction is one of the core concepts in cognitive science. It is an important building mechanism in the cognition of intelligent beings, especially humans, helping them to efficiently cope with the complex and dynamic world. However, abstraction is not only interesting from the perspective of cognition; it appears in the agenda of many disciplines: psychology, computer science, artificial intelligence, mathematics, linguistics, music, philosophy, art, and more recently, also neuroscience. Despite evident progress in research, the concept of abstraction is not still completely understood in empirical sciences, nor sufficiently well grasped in computational models. One reason for this is that abstraction is not a single homogeneous entity, but rather offers several interpretations and underlying mechanisms.

Abstraction can be primarily seen as a conceptual process by which general rules and concepts are derived from the usage and classification of specific examples. Conceptual abstractions may be formed by filtering the information content of a concept or an observable phenomenon, selecting only those aspects which are relevant for a particular purpose. In this paper, we do not attempt to provide a complete survey of this complex topic. Instead, we try to summarize evidence (not exhaustively) coming from the disciplines crucial for cognitive science (psychology, linguistics, neuroscience and computer science).

Abstraction typically refers to representing abstract concepts that can be mostly expressed by words in natural language(s). It can also refer to symbols (in Peirce's semiotic view) that have meanings in certain worlds (e.g. mathematics) or cultures. Linguistic symbols (words) can refer not only to objects (expressed by nouns), but also actions (verbs), or properties (adjectives). Within cognitive science, abstraction is differently viewed in classical theories and in embodied theories of cognition.

1.1 Theories of cognition

According to the *classical view of cognition*, the mind is considered a symbol system and cognition relates to processes of symbol manipulation (Fodor, 1998). Cognition and action/perception are separate and independent systems that work according to different principles (Barsalou, 1999). Hence, conceptual representations are non perceptual and unrelated to the body. In this framework, concepts are generated by combining and manipulating abstract, arbitrary and amodal symbols for which their internal structures are unrelated to the perceptual states and actions that produced them (Fodor, 1998).

The problem of classical approach to assign meanings to symbols, articulated by the Chinese room argument (Searle, 1980), and subsequently by the symbol grounding problem (Harnad, 1990), triggered the development of alternative theories of *grounded cognition* (see the overview, e.g. in Barsalou 1999). In these views, conceptual knowledge is represented with (multi)modal symbols related to the perceptual states that produce them. The concept of grounding embraces the other two crucial concepts – embodiment and situatedness. *Embodiment* provides the pathways towards learning concepts via sensorimotor interaction with the environment, enabled by sensory-motor features (perceptual and proprioceptive) originating in agent's sensors and actuators. *Situatedness* provides rich repertoire of features from the environment about which particular knowledge is being acquired (situated learning).

Within the grounded cognition paradigm, cognitive linguistic theories deny the presence of a separate and autonomous language module in the brain responsible for language acquisition and refuse the separation of semantic representations from the rest of cognition (Lakoff and Johnson, 1980). They claim that even

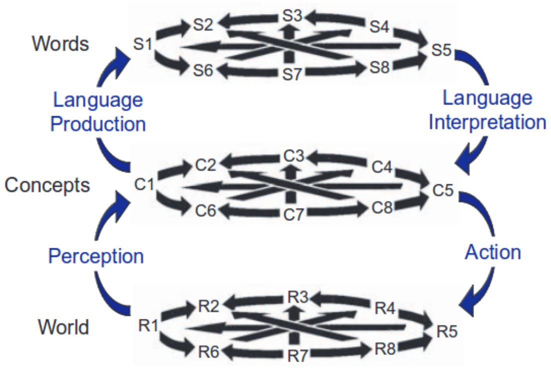

Fig. 1: General scheme of the relationship between the (multimodal) network of concepts and (symbolic) word networks, consistent with grounded cognition paradigm (from Roy 2008).

abstract concepts are grounded metaphorically in embodied and situated knowledge (for instance, love can be understood as eating, consuming the beloved person). Cognitive simulation theories focus on the role of modal simulation, situated action and bodily states in the grounding of cognitive processes (Barsalou, 1999). Social simulation theories propose that the understanding of mental states in other people requires simulations of our own mind (Goldman, 2006) and typically it requires the activation of the mirror neuron system (Rizzolatti and Craighero, 2004).

2 Psychology and linguistics

Within the relevant disciplines, perhaps most information and knowledge related to abstraction has been accumulated in cognitive psychology and psycholinguistics. The focus has been put on abstract words and conceptual knowledge regarding abstract concepts as categories. The relationship between concepts and words is sketched in Fig. 1, consistently with grounded cognition paradigm. Concepts are mental constructs created through sensory-motor interaction. On the other hand, words are surface forms (phonological or orthographic) embedded in a concrete langage, arbitrarily linked to their meanings.

Three major classical *theories of categorization* have been proposed: In *rule-based* theory, members of a category share common (perceptual) properties (e.g. colour, shape, etc.) and the membership for a category is based on satisfying established rules that permit to verify the common properties of the category. Hence, categories have strict boundaries which only works well for some categories, e.g. mathematical objects, not so much for natural categories (Bruner and Austin, 1986).

In *prototype-based* theory, categories are represented by "prototype" stimuli, which are used for judging the membership of other items. This approach assumes a more continuous way of categorization with less strict boundaries between categories which works well for natural categories (Rosch, 1973).

In *exemplar-based* theory, concepts are represented by the exemplars of the categories stored in the memory. A new item is classified as a member of the category if it is similar to one of the stored exemplars in that category (Nosofsky et al., 1992).

2.1 Meanings of abstraction in psychology

In cognitive psychology, Barsalou (2003) suggested that abstraction is not a unitary concept but has several meanings:

1. *Categorical knowledge*, which has been abstracted from experience (e.g. 'chairs'). Various accounts of knowledge are comfortable with this sense, including rule-based, prototype, exemplar, and connectionist theories.
2. *Behavioural ability to generalize* (without exceptions) across category members (e.g. using generic claims "cats have fur", or quantifications, such as "some mammals swim".
3. *Summary representations* of category members in long-term memory (in some theories). In this sense, it is not necessary to produce the behavioural abstractions. For example, exemplar models do not store summary representations and produce behavioural abstractions by scanning and summarizing exemplars online (Hintzman, 1986).
4. *Schematic representations* of categories in memory, where "schematic" refers to summary representations being sparser than exemplars (due to extracting the critical properties of a category's exemplars and discarding irrelevant properties).
5. *Flexible representations* as a result of flexible application of summary representations to many different tasks (e.g. categorization, inference, language comprehension, reasoning). From this perspective, increasing abstractness allows a representation to become increasingly flexible.
6. *Abstract concepts* that are typically detached from physical entities and more associated with mental events (Paivio, 1986; Barsalou, 1999; Wiemer-Hastings et al., 2001).

Barsalou (2003) claims that meanings 3, 4 and 5 are controversial, and focuses on one of them, summary representations, within his theory of perceptual symbol systems (Barsalou, 1999). In this paper, we focus on senses 1 and 6 that, despite not being controversial, still pose a challenge for (computational) cognitive science.

Categorical knowledge refers to concepts that can be organized *hierarchically*. This applies to concepts expressed by nouns but also by action words.

2.2 Abstraction in object words

Regarding nouns, three levels of categorization have been proposed: subordinate, basic and superordinate levels (e.g. 'a rocking chair'/'chair'/'furniture'). *Subordinate* level categories are characterized by a low degree of generality and by clearly identifiable, detailed and specific features. *Basic* level provides (ecologically) the most relevant conceptual information. *Superordinate* level typically implies a high degree of generality and allows to store general information. As a result, subordinate categories are more concrete than basic categories, which in turn are more concrete than superordinate categories. Interestingly, despite the order of increasing abstractness, the acquisition of knowledge in children, does not proceed bottom up. Instead, it starts from the basic level, moving to subordinate and superordinate levels of categorization in parallel Bloom (2000).

2.3 Abstraction in action words

The organization of action words is also hierarchical. The most concrete action words may correspond to motor primitives, such as 'push', 'pull', 'grasp', 'release', etc. More abstract action words, related to physical actions, require the concatenation of motor primitives; e.g. we could consider 'keep'='grasp'+'stop', or 'give'='carry'+'release' (Stramandinoli et al., 2011). Action verbs can also be differentiated in their level of concreteness and motor modality; e.g. 'push' is uniquely linked with the action of pushing by using the hand, while 'give' implies multiple motor instances of the process of passing an object by using one hand, two hands, mouth, etc. (Cangelosi and Schlesinger, 2014).

The above examples also apply to abstract actions. For example, we can grasp or keep an idea, wrap up a meeing, etc. (as explained by cognitive lingustics). Many action verbs are inherently abstract and refer to mental operations. For instance, in the educational setting, the well-known revised Bloom's taxonomy categorizes action verbs in cognitive, affective and psychomotor domains (Anderson and Krathwohl, 2001). For instance, in the cognitive domain, the verbs are categorized hierarchically (in ascending order) according to the abilities they involve: remembering, understanding, applying, analyzing, evaluating and creating. In general, verbs are clearly the most complex word category with verb semantics posing a challenge for cognitive science(Levin, 1993).

2.4 Differences between concrete and abstract words

Concrete and abstract words can be differentiated according to the following factors (Kousta et al., 2011):[1]

- *Perceivability* As opposed to concrete, abstract words do not have a physical referent and hence refer to entities that are distant from immediate perception, and typically represent entities that are not spatially constrained (e.g. 'truth', 'happiness').
- *Imageability and context availability* According to the dual-coding theory (Paivio, 1986), concrete concepts, activating both verbal and non-verbal systems, require a major involvement of memory. On the contrary, abstract concepts are represented in the verbal system only, with inferior involvement of memory (Barsalou, 2008). In addition, they evoke less imagery (Wiemer-Hastings and Xu, 2005), and activate a narrower contextual verbal support than concrete concepts, as suggested in the context-availability theory (Schwanenflugel, 1991). These can be reasons why acquisition of abstract concepts is more complex than that of concrete concepts.
- *Hierarchical categorization* Conceptual knowledge can be organized in categories hierarchically, as already mentioned in Section 2.2. This works well for concrete objects, but even there the concepts at the superordinate level can become very abstract, unifying subcategories that may not share any perceptual features (e.g. 'game').
- *Mode of Acquisition* In has been shown in experiments with elementary school children that MOA ratings change with the school age, shifting gradually from mainly perceptual to mainly linguistic MOA (Wauters et al., 2003).

Abstract concepts pose a bigger challenge for embodied (grounded) theories of cognition, because the grounding of abstract concepts, deprived of their physical referents, seems less straightforward than that of concrete concepts. Symbolic cognition treats all concepts as symbols, so the burden is shifted to accounting how all concepts acquire their meanings (Harnad, 1990).

Abstract concepts are also hiearchically organized (e.g. 'love' is less abstract than 'democracy'), depending on how well they can be grounded. Even abstract concepts appear to depend heavily on situations and situated action (Schwanenflugel, 1991). The processing of abstract concepts is facilitated when a background situation contextualizes it (Barsalou and Wiemer-Hastings, 2005). We can also add that for concrete objects, one usually abstracts over individuals, whereas for abstract objects one abstracts over situations or events.

Categorizing entities is a useful approach in science, but consistently with cognitive theories, it is more accurate to think of the distinction between concrete and abstract words as that of a *concreteness–abstractness continuum*, along which all entities can be varied. It

[1] Altarriba et al. (1999) proposed that words refering to emotions should be categorized as distinct from concrete and abstract words, because these three categories received different ratings in term of concreteness, imageability and context availability.

also applies to cognitive theories that range from purely symbolic (disembodied) accounts to purely embodied ones (Wilson, 2002). In addition, some have argued that both disemboded and embodied approaches are needed for complete account how humans acquire knowledge (Louwerse, 2010; Dove, 2011).

On the sentence level, comprehension and production processes focus, within the grounded cognition paradigm, mainly on the roles of *mental simulation* (e.g. Barsalou 1999; Decety and Grèzes 2006; Goldman 2006). Simulation is the reenactment of perceptual, motor, and introspective states acquired during experience with the world, body, and mind. The question is how even very abstract sentences can be simulated. Zwaan (2015) argues that sensorimotor and symbolic representations mutually constrain each other in discourse comprehension, and proposes that abstract concepts serve two roles in this process. They serve as pointers in memory, used (1) *cataphorically* to integrate upcoming information into a sensorimotor simulation, or (2) *anaphorically* to integrate previously presented information into a sensorimotor simulation. In either case, he concludes, the sensorimotor representation is a specific instantiation of the abstract concept.

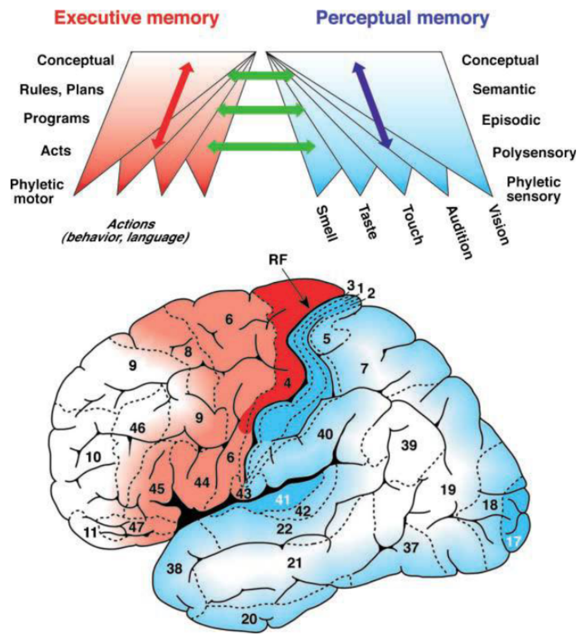

Fig. 2: General scheme of the hierarchical organization of memory in the lateral cerebral cortex of the left hemisphere (Fuster, 2009). Lighter shade of colors indicates the direction in the hierarchy towards more abstract memory and knowledge.

3 Neural perspective

Neurophysiological and neuropsychological evidence suggests hierarchical organization of knowledge (long-term memory) in the brain that can be observed mainly over the cortex (Fuster, 2009). Fuster's theory of memory organization is based on the existence of bidirectionally connected executive and perceptual memories, interacting at multiple levels, both forming the so-called *cognits*, formed by distributed neural networks. Cognits are items of knowledge, hierarchically organized in terms of semantic abstraction and complexity. As shown in Figure 2, frontal and posterior cortices provide different spatial layout of organization. Whereas in the frontal lobe, the apex of the hiearchy rests at the very anterior part of dorsolateral prefrontal cortex, the top of hierarchy in the posterior cortex rests "inside" the lobes. Fuster's theory departs from the idea of a memory as a passive storage, and instead sees the memory as operating in perception–action cycles (Fuster, 2004), which is a core concept in biology (Uexküll, 1926).[2]

3.1 Posterior cortex

In the temporal cortex, the increasing abstraction can be related to the growing invariance of the neurons in response to various perceptual stimuli. The visual system of mammals is organized hierarchically in terms of feature extraction (from the simplest to the more complex features), which is associated with increased radius of the receptive field of neurons (via afferent connections) towards the higher (more abstract) layers. On top of the pyramid stands the area IT (inferior temporal), responsible for invariant object recognition with respect to position, rotation, and scale, and in the case of biological objects also deformation (Jellema and Perrett, 2006). Likewise, area STS (Superior Temporal Sulcus) has its division in the context of the recognition of biological motion. An interesting property of STS is that it contains many neurons that are sensitive to viewpoint from which the object is observed (e.g. front view, side view, etc.) in its posterior part (STSp), but also neurons that are invariant to it, in anterior part (STSa).

Parietal cortex also demonstrates abstraction properties. For instance, intraperietal sulcus has been argued to be the seat of (abstract) numerical cognition (Dehaene et al., 1998) (but see the challenging view in Kadosh and Walsh 2009). Among other things, posterior parietal cortex is known to be involved in spatial cognition (the superior parietal lobule) and in non-spatial cognition (the inferior parietal lobule). The latter is typically related to abstract information processing. For instance, Yamazaki et al. (2009) proposed several mechanisms on non-spatial representations processed by the posterior parietal cortex, ranging from percpetial space (low abstraction), via temporal and social spaces to conceptual space (high abstraction), with corresponding brain areas involved (collected from previous papers).

[2] This is consistent with a psychological model of memory which is postulated to subserve intelligent action Glenberg (1997).

3.2 Frontal cortex

Parietal cortex has bidirectional connections to the frontal cortex, which is crucial for sensory-motor integration. Various levels of abstraction are subserved by corresponding mirror neurons opeating simultaneously at various levels of a perceptual hierarchy (e.g. perspective-dependent neurons as opposed to goal-coding mirror neurons, that are insensitive to motor trajectories). Mirror neuron system has been labeled as the action-observation–action-execution matching circuit, supporting the direct matching hypothesis (Rizzolatti and Craighero, 2004). This circuit is erriched by its connections (from the parietal cortex) to the temporal cortex. Namely, a part of the focus area F5 (F5c) is connected with STSp through PFG (parietal frontal gyrus) forming a perspective variant path. Another part of F5 (F5a) is also connected with STSa through AIP, forming an invariant path emphasizing the actor and the object acted upon, rather than the viewpoint from which it is observed (Nelissen et al., 2011).

The frontal lobes are most expanded in humans, being hierarchically organized, all the way from the primary motor cortex towards the prefrontal cortex (PFC). Highest in the pyramid stands the apex of the dorsolateral PFC. It is noteworthy that different parts of the frontal cortex have a corresponding structure (with connections to the thalamus and some of the nuclei in the basal ganglia, which are involved in action selection (O'Reilly et al., 2012). It thus appears that similar (or identical) neural mechanisms operate at different levels of the hierarchy (abstraction).

Abstraction in PFC seems to comply with several different types:

- *Temporal abstraction* implies integration of behavior (Koechlin and Hyafil, 2007; O'Reilly, 2010), maintaining the organization of (sub)actions with increasing temporal spans. For example, the (more abstract) goal to make coffee can be sequentially decomposed to several subgoals (get a cup, pour water in a machine, etc.), which in turn can eventually be decomposed to a set of primitive actions.

- *Policy abstraction* (Badre and D'Esposito, 2007) implies the representation of a task, at any given moment of performance, as a set of choices over lower-level tasks. For instance, a certain action can be executed in various (motor) ways, achieving the same goal.

- *State abstraction* Christoff and Keramatian (2007) occurs when the subject treats non-identical stimuli or situations as equivalent. For instance, we ignore the font type during reading.

In sum, the brain seems to have two major abstracting pathways with revealed neural organization: one located posterior for (physical) object recognition, and the other one in frontal lobes, corresponding to astracting in execution and planning of actions. Probably the third pathway, corresponding to hierarchical organization of abstract objects (words) could also be searched for in the brain. The review and meta analysis of the fMRI literature provides a frame for this search (Binder et al., 2009). It was concluded that the semantic memory (covering the processing of both concrete and abstract words) is distributed in a left-lateralized network consisting of 7 distinct regions.[3]

Similar meta analysis of fMRI and PET studies (Wang et al., 2010) indicated consistent and meaningful differences in neural representations. Roughly speaking, abstract concepts elicit greater activity in the frontal areas, while concrete concepts elicit greater activity in the posterior brain areas. These results suggest greater engagement of the verbal system for processing of abstract concepts and greater engagement of the perceptual system for processing of concrete concepts.

4 Computational aspects

The role of computational modeling is increasingly growing[4], promising a huge potential for advancing knowledge in the interdisciplinary cognitive science (McClelland, 2009). In addition, the importance of computational modeling can be considered an indispensable ingredient allowing the mechanistic, neurally constrained understanding of the mind, extending the horizons of cognitive science (Farkaš, 2012).

4.1 Towards deciphering the brain code

The brain representations in IT area (mentioned in Section 3.1) seem quite reliable despite ever-present noise because the technology already allows their reading from the brain (Hung et al., 2005), using single-cell recordings in monkeys. However, even the non-invasive imaging methods (e.g. fMRI) provide means for deciphering the brain.

Mitchell et al. (2008) reported about a computational model that could not only be used for interpreting various spatial patterns of neural activation corresponding to different semantic categories (expressed by concrete nouns), but also for predicting (previously unrecorded) fMRI patterns for novel concrete nouns (via generalization). This was possible due to linking the neural data to the statistics of very large text corpora.

A nice global picture of the growing abstraction in the brain has been recently proposed as a result of a large-scale analysis of fMRI data (∼17,000 experiments and ∼1/4 of fMRI literature), i.e. the technology, that

[3] These are posterior inferior parietal lobe, middle temporal gyrus, fusiform and parahippocampal gyri, dorsomedial PFC, inferior frontal gyrus, ventromedial PFC, and posterior cingulate gyrus.

[4] despite the dominant role of psychology (Gentner, 2010)

has increasingly become an important source of information over the last decades (Taylor et al., 2015). In the paper, the hypothesis was tested, using formal methods based on a new cortical graph metrics (network depth), that regions deeper in the brain (i.e. more remote from the sensory inputs) represent more abstract functions. Data-driven analyses defined a hierarchically ordered connectome, revealing a related continuum of cognitive function. The authors concluded that progressive functional abstraction over network depth may be a fundamental feature of brains.

It can be expected that the number of papers of this sort will grow, revealing the patterns represented by the brain. Most probably, desciphering abstract concepts from the brain will remain a challenge.

4.2 Neural network models

So far, four major frameworks have been proposed for representing conceptual knowledge (symbolic, connectionist, dynamic and probabilistic) (McClelland, 2009). Each of the frameworks has to be able to incorporate abstraction, in order to be able to account for the spectrum of cognitive processes, ranging from sensory-motor behaviors to abstract thoughts. Believing that connectionism has the best potential in this regard (Farkaš, 2011, 2013), we will henceforth focus on computational aspects of abstraction from the connectionist perspective.

Neural networks perspective involves the idea of brain operation as performing a cascade of nonlinear transformations of patterns across layers. These transformations go in both directions (i.e. also backwards), reflecting the fact that most of the connections between brain areas are bidirectional. Since their breakthrough a decade ago (Hinton et al., 2006), *deep neural networks* (DNNs) have become the influential approach in connectionist modeling, fueled by very successful applications in various tasks, such as object recognition/classification, speech recognition and natural language processing (Schmidhuber, 2015). DNNs lend themselves nicely to hierarchical modeling (see Fig. 3), since at each layer the features with increasing abstraction are formed during training.[5] Visual system of mammals is also known to have deep hierarchical structure (Felleman and Essen, 1991), starting at the retina, through the cascade of feature detecting layers with an increasing complexity, all the way up to IT cortex where (physical) object recognition takes place.

It is important to realize that this pathway only applies to concrete (imageable) objects which can be still be considered a rather low level of abstraction, compared to abstract objects (entities). In case of abstract objects, probably a different neural pathway has to be sought, because these do not have direct reference in

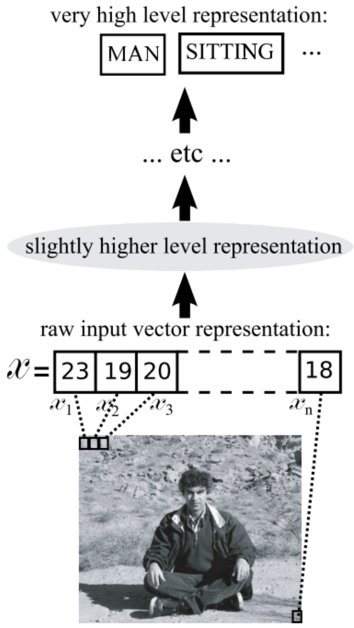

Fig. 3: Picture illustrating the "distance" between the raw data and high-level description of the image (Bengio, 2009). The established mapping will involve a number of intermediate levels of representation with an increasing abstraction.

the environment. Even though abstract concept representations can be triggered by visual inputs (visual context), they will be more linked to language. Linguistic context can hence as well serve as the cue. Representations of abstract objects are usually modeled not as deep networks but as widely distributed associative networks spread over distant cortical (and subcortical areas) (Pulvermüller, 1999). Shallice and Cooper (2013) suggest that networks for abstract words differ from those for concrete words, and depend more on frontal regions (left PFC or left inferior frontal gyrus) that are known to be involved in syntactic processing (and hypothesized as crucial for understanding abstract words).

5 Conclusion

Our understanding of abstraction is far from complete, regardless of the perspective taken. In cognitive psychology and linguistics, good evidence has been accumulated about how concrete and abstracts objects depend on the environment and how they differ regarding their properties and the role of language in their acquisition and use (more important for abstract concepts). Better understanding of semantic organization of content words (referring to concepts) in the brain could provide important material to be integrated in cognitively and biologically inspired modeling of this phenomenon. Organization of semantic knowledge on the sentence level and in discourse is yet more elusive.

[5] Typically in classification tasks, most layers are trained in unsupervised way, and are combined with supervised training at the top layer (Hinton et al., 2006).

Acknowledgment

This work was supported by VEGA project 1/0898/14 and KEGA project 017UK-4/2016.

References

Altarriba, J., Bauer, L. a Benvenuto, C. (1999). Concreteness, context availability, and imageability ratings and word associations for abstract, concrete, and emotion words. *Behavior Research Methods*, 31(4):578–602.

Anderson, L.W., a Krathwohl (Eds.). (2001). *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. New York: Longman.

Badre, D. a D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *Journal of Cognitive Neuroscience*, 19:2082–2099.

Barsalou, L. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22(04):577–660.

Barsalou, L. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 358(1435):1177–87.

Barsalou, L. (2008). Grounded cognition. *Annual Reviews of Psychology*, 59:617–645.

Barsalou, L. a Wiemer-Hastings, K. (2005). *Situating abstract concepts*, pp. 129–163. Cambridge University Press.

Bengio, Y. (2009). Learning deep architectures for ai. *Foundations and Trends in Machine Learning*, 2(1):1–127.

Binder, J., Desai, R., Graves, W. a Conant, L. (2009). Where is the semantic system? a critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, 19:2767–2796.

Bloom, P. (2000). *How Children Learn the Meanings of Words*. MIT Press, Cambridge, MA.

Borghi, A., Flumini, A., Cimatti, F., Marocco, D. a Scorolli, C. (2011). Manipulating objects and telling words: a study on concrete and abstract words acquisition. *Frontiers in Psychology*.

Bruner, J. a Austin, G. (1986). *A Study of Thinking*. Transaction Books.

Cangelosi, A. a Schlesinger, M. (2014). *Developmental Robotics: From Babies to Robots*. MIT Press, Cambridge MA.

Christoff, K. a Keramatian, K. (2007). Abstraction of mental representations: theoretical considerations and neuroscientific evidence. In *The Neuroscience of Rule-Guided Behavior*, pp. 107–127. Oxford University Press.

Decety, J. a Grèzes, J. (2006). The power of simulation: imagining one's own and other's behavior. *Brain Research*, 1079:4–14.

Dehaene, S., Dehaene-Lambertz, G. a Cohen, L. (1998). Abstract representations of numbers in the animal and human brain. *Trends in Neurosciences*, 21:355–361.

Dove, G. (2011). On the need for embodied and disembodied cognition. *Frontiers in Psychology*, 1(242).

Farkaš, I. (2012). Indispensability of computational modeling in cognitive science. *Journal of Cognitive Science*, 13(12):401–435.

Farkaš, I. (2011). Connectionism embraced by computational cognitive science (in Slovak). In *Umelá inteligencia a kognitv́na veda III*, pp. 19–62. Vydavateľstvo STU, Bratislava.

Farkaš, I. (2013). Is representational pluralism in cognitive science inevitable? (in Slovak). In *Kognitivní věda a umělý život*, pp. 107–114. Slezská univerzita v Opavě.

Felleman, D. a Essen, D. V. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1:1–47.

Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press.

Fuster, J. (2004). Upper processing stages of the perception–action cycle. *Trends in Cognitive Science*, 8:143–145.

Fuster, J. (2009). Cortex and memory: Emergence of a new paradigm. *Journal of Cognitive Neuroscience*, 21(11):2047–2072.

Gentner, D. (2010). Psychology in cognitive science: 1978-2038. *Topics in Cognitive Science*, 2:328–344.

Glenberg, A. (1997). What memory is for. *Behavioral and Brain Sciences*, 20:1–55.

Goldman, A. (2006). *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford University Press.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, pp. 335–346.

Hinton, G., Osindero, S. a Teh, Y. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554.

Hintzman, D. (1986). 'Schema abstraction' in a multiple-trace memory model. *Psychological Review*, 93:411–428.

Hung, C., Kreiman, G., Poggio, T. a DiCarlo, J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science*, 310:863–866.

Jellema, T. a Perrett, D. (2006). Neural representations of perceived bodily actions using a categorical frame of reference. *Neuropsychologia*, 44:1535–1546.

Kadosh, R. a Walsh, V. (2009). Numerical representation in the parietal lobes: Abstract or not abstract? *Behavioral and Brain Sciences*, 32:313–373.

Koechlin, E. a Hyafil, A. (2007). Anterior prefrontal function and the limits of human decision-making. *Science*, 318:594–598.

Kousta, S., Vigliocco, G., Vinson, D., Andrews, M. a Campo, E. D. (2011). The representation of abstract words: why emotion matters. *Journal of Experimental Psychology: General*, 140(1):14–34.

Lakoff, G. a Johnson, M. (1980). *Metaphors We Live By*. University Chicago Press, Chicago.

Levin, B. (1993). *English Verb Classes and Alternations*. Chicago University Press, Chicago.

Louwerse, M. (2010). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, pp. 1–30.

McClelland, J. (2009). The place of modeling in cognitive science. *Topics in Cognitive Science*, 1(1):11–38.

Mitchell, T. et al. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, 320(5880):1191–1195.

Mussa-Ivaldi, F. a Bizzi, E. (2000). Motor learning through the combination of primitives. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 355(1404):1755–1769.

Nelissen, K. et al. (2011). Action observation circuits in the macaque monkey cortex. *Journal of Neuroscience*, 31(10):3743–3756.

Nosofsky, R., Kruschke, J. a McKinley, S. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(2):211–233.

O'Reilly, R. (2010). The what and how of prefrontal cortical organization. *Trends in Neurosciences*, 33(4):355–361.

O'Reilly, R., Munakata, Y., Frank, M., Hazy, T. a contributors (2012). *Computational Cognitive Neuroscience*. Wiki book, 2. vyd.

Paivio, A. (1986). *Mental Representation: A Dual Coding Approach*. Oxford University Press, Oxford.

Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22:253–336.

Rizzolatti, G. and Craighero, L. (2004). The mirror neuron system. *Annual Reviews of Neuroscience*, 27:169–192.

Rosch, E. (1973). Natural categories. *Cognitive Psychology*, 4(3):328–350.

Roy, D. (2008). A mechanistic model of three facets of meaning. In *Symbols and Embodiment: Debates on Meaning and Cognition*, pp. 195–222. Oxford University Press.

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117.

Schwanenflugel, P. (1991). Why are abstract concepts hard to understand? In *The Psychology of Word Meanings*, pp. 223–50. Lawrence Erlbaum.

Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3):417–424.

Shallice, T. a Cooper, R. (2013). Is there a semantic system for abstract words? *Frontiers in Human Neuroscience*, 7:1–10.

Stramandinoli, F., Cangelosi, A. a Marocco, D. (2011). Towards the grounding of abstract words: A neural network model for cognitive robots. In *Proc. of the Int. Joint Conf. on Neural Networks*, pp. 467–474.

Taylor, P., Hobbs, J., Burroni, J. a Siegelmann, H. (2015). The global landscape of cognition: hierarchical aggregation as an organizational principle of human cortical networks and functions. *Scientific Reports*, 5, doi:10.1038/srep18112

Thagard, P. (2005). *Mind: An Introduction to Cognitive Science*. MIT Press, second. vyd.

Uexküll, J. (1926). *Theoretical biology*. Harcourt, Brace and Co., New York.

Wang, J., Conder, J., Blitzer, D. a Shinkareva, S. (2010). Neural representation of abstract and concrete concepts: A meta-analysis of neuroimaging studies. *Human Brain Mapping*, 31(10):1459–1468.

Wauters, L., Tellings, A., van Bon, W. a van Haaften, A. (2003). Mode of acquisition of word meanings: The viability of a theoretical construct. *Applied Psycholinguistics*, 24(3):385—406.

Wiemer-Hastings, K., Krug, J. a Xu, X. (2001). Imagery, context availability, contextual constraint, and abstractness. In *Proc. of the 23rd Annual Conf. of the Cognitive Science Society*, pp. 1134–1139.

Wiemer-Hastings, K. a Xu, X. (2005). Content differences for abstract and concrete concepts. *Cognitive Science*, 29(5):719–736.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4):625/636.

Wintermute, S. (2012). Imagery in cognitive architecture: Representation and control at multiple levels of abstraction. *Cognitive Systems Research*, 19–20:1–29.

Yamazaki, Y., Hashimoto, T. a Iriki, A. (2009). The posterior parietal cortex and non-spatial cognition. *F1000 Biology Reports*, 1:1–6.

Zwaan, R. (2015). Situation models, mental simulations, and abstract concepts in discourse comprehension. *Psychonomic Bulletin and Review,* doi:10.3758/s13423-015-0864-x