

Symbols and embodiment: The perspective of a neural modeler

Andreas Knoblauch

2005

Preprint:

This is an accepted article published in Proceedings of the Garachico Workshop on Symbols, Embodiment and Meaning.. The final authenticated version is available online at: [https://doi.org/\[DOI not available\]](https://doi.org/[DOI not available])

Symbols and Embodiment: The perspective of a neural modeler *

Andreas Knoblauch

Honda Research Institute Europe GmbH
Carl-Legien-Str.30, D-63073 Offenbach/Main, Germany

Abstract. We try to contribute to the current debate about symbols and embodiment by pointing out the perspective of a neural modeler. We illustrate the default definitions of 'symbol', 'embodiment', and 'grounding' in the context of detailed neural network models, i.e., on a level more detailed than common connectionist approaches. Our arguments are based on cortical models of Hebbian cell assemblies. These models have been employed to implement a large-scale cortical architecture to enable a robot to perform simple tasks such as understanding and reacting to simple spoken commands. In the discussion we focus mainly on the first debate about the relation of Searle's Chinese room argument to embodiment.

1 Definitions

1.1 Symbols

For a neural network modeler, one simple possible way to discern symbols from non-symbols is to look at the inner structure of the representational units. Sub-symbols have an inner structure which can be used to define a similarity metric relevant for the represented entity. In contrast, symbols have no relevant inner structure (i.e., symbols are abstract and arbitrary). For example, in object recognition systems, a non-symbol or sub-symbol may be a vector of sensory features, while a symbol may correspond to a single node representing an object category. However, these definitions do not seem sufficient for the current debate which is about language and the representation of meaning.

According to the workshop's default definition a symbol is a "theoretical element that is arbitrary, abstract, and amodal" [1]. Before we proceed by discussing and adapting that definition, it may be useful to be aware of the different contexts in which we will use the word "symbol". The situation is illustrated in Fig. 1. We live in a physical world W where systems or subjects S are part of that world and interact with the world. Some of the systems (namely we, the subjects) somehow are able to generate a usually unique psychological or phenomenological space P , which we can employ, for example, to generate ideas or theories T about all kinds of issues on all levels W,S,P,T . In particular, we can make theories about S (predominantly done by biology, neuroscience, and AI), P (psychology), or T (metamathematics or logics). Ideas or theories

*early draft(!) for the Garachico Workshop on Symbols, Embodiment and Meaning; Tenerife, Spain, 2005

T consist essentially of a set of symbols (as defined above) and additional rules determining how the symbols can be “manipulated”.




Fig. 1: Different modeling levels: We live in a physical world W . Systems (or subjects) S are part of that world and interact with the world. The subjects somehow are able to generate a (usually unique) psychological or phenomnological space P , which we can use, for example, to generate theories T about all kind of issues on all levels W,S,P,T .

Since symbols are part of theories or ideas this implies two aspects of a symbol, on the one hand side in T , on the other hand the implementation of the symbol in S . As a neural modeler one is predominantly interested in the S -level implementation of P in W i.e., in reducing the observable psychological and behavioral phenomenons as good as possible to detailed neural and synaptic processes and finally to the physical laws. (The underlying preliminary naturalistic working hypothesis assumes that this goal is actually possible.) Since models about S are finally also T theories, the neural modeler (and any other kind of S -modeler such as many AI researcher) has obvioulsy to discern between the two kind of symbols *within* his theory: Symbols to model the T -symbols of S , and symbols to model the implementation of the T -symbols of S . Thus, we will refer to these two kind of symbols as T -symbols and S -symbols, respectively. For example, for a cognitive system S capable of understanding language, a T -symbol is the representation of a word, while S -symbols are finer-grained entities used for implementing the word representation. For example, an S -symbol could be a node in a connectionist network (or alternatively, a state or band variable in a Turing machine) while a T -symbol could be implemented by a set of distributed S -symbols (and possibly further dynamic processes).

The current debate is about the question whether cognitive systems (such as we) are or have to be either symbolic or “embodied”. At least the debate question 2.2 seems to imply this dichotomy where the use of “embodied” appears to be equivalent to “non-symbolic”. Of course, any kind of cognitive system must be symbolic in the trivial sense that we have symbolic language and theories T . Thus the critical question is not about the reality of T symbols, but are about the implementation of T -symbols in S -symbols. Note that, depending on how we define “symbol” and “embodiment”, it might be thinkable that T -symbols could be implemented by non-symbolic processes on the S level (although this seems counterintuitive since we actually aim at developing a symbolic theory T about S).

1.2 Embodiment

Also embodiment comes in several different flavors (cf. [2]). The strongest claim would be that embodiment could extend the qualitative or at least quantitative computational capabilities of a system S by exploiting the properties of W , e.g., described by the physical laws. If true, this would essentially negate the Church-Turing thesis that (symbolic) Turing machines can compute any “naturally” (or physically) computable function. (For example, “embodied” analog computers might be better in simulating physical systems, or in computing real numbers with infinite precision).

A less strong idea of embodiment is the dichotomy of embodied versus symbolic cognitive systems addressed by the current debate which attempts to classify cognitive systems according to the interface between the system S itself and the external world W . Obviously, any cognitive system S that deserves that name will have to interact with its environment (percept and act) and is therefore embodied in a trivial sense. Similarly, any cognitive system S must be symbolic in a trivial sense since it must explain our capabilities to use language and think in symbols (e.g., to develop theories T within our psychological space P). Thus, in this trivial sense any model of a cognitive system will be both embodied and symbolic.

Note that any cognitive system can be divided into sensors, actors, and internal machinery. Our default definition of an “embodied” system demands that the meaning of a symbol must depend on activity in systems also used for perception, action, and emotion and reasoning must require the use of those systems. This form of *weak embodiment* is stronger than the trivial version of embodiment, but obviously addresses merely the high-level structure of the internal machinery (e.g., in Marr’s terms, the algorithmic or computational levels, but not the implementation level). In particular, this idea of embodiment will probably be neutral to such questions as whether “ideas are the sole province of biological systems” (as discussed in [1], p.4). This is since on the implementation level an “embodied” system then can be translated into a “symbolic” system (e.g., a computer program or Turing machine) with the same sensor/actor interface and vice versa. And, of course, the property of embodiment will be a gradual property. Nevertheless, this idea of embodiment might still prove useful, e.g., in building more efficient artificial cognitive systems, or in guiding the analysis of the brain.

2 A neural modeler’s perspective

When words referring to actions or visual scenes are presented to humans, distributed networks including areas of the motor and visual systems of the cortex become active (e.g., [3, 4]). The brain correlates of words and their referent actions and objects appear to be strongly coupled neuron ensembles in defined cortical areas. The theory of *cell assemblies* [5, 6, 7, 8] provides one of the most promising frameworks for modeling and understanding the brain in terms of distributed neuronal activity. It is suggested that entities of the outside world

(and also internal states) are coded in groups of neurons rather than in single (“grandmother”) cells, and that a neuronal cell assembly is generated by Hebbian coincidence or correlation learning where the synaptic connections are strengthened between co-activated neurons. Models of neural (auto-) associative memory have been developed as abstract models for cell assemblies.

2.1 Local cell assemblies, associative memory, and “neural S-symbols”

The notion of cell assemblies as strongly coupled neurons leads to the concept of neural (*auto-*) *associative memory* [9, 10, 11]. A particular simple model of neural associative memory is the one proposed by Steinbuch and Willshaw [9, 12, 10, 13] consisting of simple McCulloch Pitts like threshold units and recurrent binary synapses (Fig. 2). Thus the activity pattern of the cell population can be described by a binary vector (i.e., we identify these activity patterns with the cell assemblies). After learning a number of cell assemblies the network can be described by a connection matrix A corresponding to a graph, where the nodes correspond to the neurons, and cell assemblies correspond to k -cliques of neurons (a k -clique is a subset of size k consisting of completely connected neurons). *Hetero-association* works similar to auto-association except that the “memory matrix” describes the synaptic connections between *two* different neuron populations. Hetero-associative connections can map assemblies of the first population (or sets or parts of them) to cell assemblies of the second population (or sets or parts of them).




Fig. 2: Left: Neural (auto-)associative memory where two cell assemblies of size $k = 4$ have been stored (corresponding to the activation patterns u^1 and u^2) in the “memory matrix” of synaptic connections. Right : A more realistic implementation of an associative memory [14] modeling a small patch of cortical tissue. The model comprises several spiking populations of excitatory and inhibitory neurons. The “memory matrices” are employed in several afferent and recurrent connections (c^H, c^A).

The virtue of the binary model is that it can be easily understood and analysed, but the main results apply also to more realistic gradual models. Neural associative memories have a couple of nice features. They achieve pattern completion, i.e., a cell assembly can be activated not only by the very same inputs that have been used for learning, but also by modified patterns that are “sufficiently” similar to the original address pattern. For example, assembly u2 in Fig. 2 will already be activated by addressing an arbitrary subset of size ≥ 3 (and setting the firing threshold equal to the number of active neurons). It can be shown that the number of storable patterns scales almost with the number of synapses if the patterns are sparse and have random character (i.e., a population of n neurons can store almost n^2 sparse cell assemblies with $k \ll n$). Access time is essentially independent of the number of stored patterns. The overlaps of different cell assemblies can be used to express the similarities of the represented entities. Cell assemblies thereby provide a very natural associative way of grounding new representations in the sensory inputs by means of bidirectional associative connections.

Associative memories have been used to model small volumes of cortical tissue (e.g., 1mm^3 , corresponding to a macrocolumn or the range where dense local recurrent connections between any cell pairs are possible). A step towards more biological realism is to replace the single McCulloch-Pitts population by more realistic spiking neurons. Indeed, one can show that a network consisting of several spiking neuron population (one excitatory, two inhibitory) can effectively implement the simple model. Moreover, this model will inherently make use of spike timings according to a latency code in that early spikes (relative to an external event or an underlying oscillation) are much more relevant than late spikes for activating an assembly [14, 13].

Local cell assemblies can be seen as elementary neural (S-) symbols which can be “allocated” or learned to represent the inputs for further processing in down-stream target populations. The symbolic character is most apparent in two special cases: (i) if the assembly size is $k = 1$, i.e., if the assembly consists only of a single cell then the resulting assembly networks are essentially equivalent to corresponding node-based common connectionist networks. Or (ii) if the neurons that constitute a cell assembly are chosen randomly, e.g., by noise. Then the correlations (or overlaps) between two cells are minimal which is required, for example, to store a maximal number of activity patterns. Due to their singular or random character the cell assemblies can be said to be abstract and arbitrary, whereas the property of amodality depends on the location of the neuron population. E.g., a local population receiving visual inputs will develop visual perceptual symbols (cf. [15]).

2.2 Global cell assemblies, language, and T-symbols

We have designed a large-scale model of many interconnected cortical areas employing spiking associative memories. The model was implemented and tested on a robotic platform enabling the robot to understand and react to simple commands (such as “Bot show plum!”) [16]. The language part of the model is

illustrated by Fig. 3, the action part by Fig. 4.

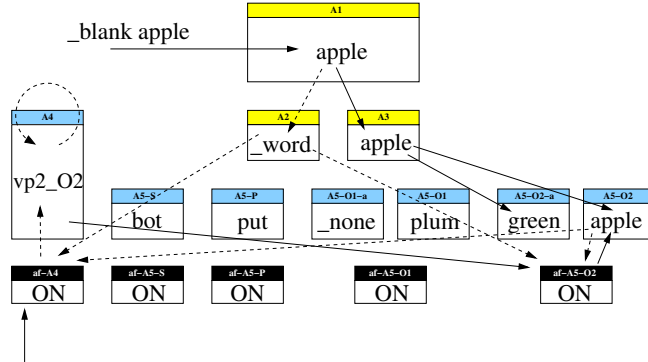


Fig. 3: The language part of an associative cortex model (see [16]) during processing the sentence “Bot put plum to green apple”. Each box corresponds to a cortical (or subcortical) area modelled as a neural associative memory. The meaning of the sentence is represented by distributed cell assemblies comprising “slot areas” for different grammatical roles to implement elementary productivity and systematicity. Auditory input enters via areas A1 and A3 the central areas and is distributed across the grammatical slots according to a logic controlled by a grammatical sequence memory (A4) (where basic sentence types are stored) and subcortical “activation fields” (small boxes).

Each box in the figures corresponds to a spiking neural associative memory storing local cell assemblies as described above. For illustration purposes, each area has been labelled according to the current activity pattern. (In general a superposition of several stored local assemblies can be activated, e.g. to represent uncertainty or to represent new entities to be learned; the labels correspond to the cell assembly most similar to the current activation pattern). The resulting *global assembly*, for example representing the T-symbol “plum”, stretches over many cortical areas (involving visual, auditory, action, and goal-related areas) and changes dynamically during the process of “understanding” and reacting to the command. Thus, the global cell assembly as a whole works as a sign in Peirce’s sense, i.e., as a mediator between the idea of a “plum” and the real plum in the external world. The global assembly consists of parts some of which can be attributed as “abstract” and “amodal”, e.g., the lexical representation of “plum” in A3. But these symbolic parts are naturally grounded in the synaptically connected perceptual and action-related parts of the global assembly.

2.3 Cortical macrocolumns, prediction, and embodiment

Cognitive processes must be able to distinguish between different representational modi. For example, representational states may refer to present, future (or prediction), reality, wish, signal (detailed and concrete), or symbol (abstract,




Fig. 4: Action part of an associative cortex model (see [16]) during performing the command “Bot put plum to green apple!”. The goal areas (G1-G4) received their inputs from the grammatical role areas (A5-S, A5-O1, etc.) and divide the goals into a sequence of subgoals (i.e., seek plum, pick plum, move to the apple, drop plum). (High-level) motor areas receive inputs from the goal areas in order to perform the current subgoal. The completion of subgoals and switching to the next subgoal is controlled by “evaluation fields” checking, for example, the consistency of visual perceptual activity patterns (e.g., in V4) with goal representations (e.g., in G3). At the shown system state the robot is about to finish the subgoal of seeking the plum.

amodal), perception or action. Many cognitive architecture take a modular approach where these different representational modi are segregated into different cognitive subsystems or modules. (In general, an architecture can be said to be modular if it can be divided into subsystems such that there is much more communication between processes inside a subsystem than between processes of different subsystems.) For example, we could segregate a cognitive system into different modules for perceptions, actions, goals, memory, rule-based prediction systems, etc.

In the last section we have indeed argued how global cell assemblies can implement and ground T-symbols (e.g., words of a language) by distributed activation stretching over many sensory, motoric, and associative cortical areas [3]. Thus, although the brain appears to have a modular character in that sense, hints to possible complementary strategies of grounding may be found when looking at the microstructure of a single cortical macrocolumn [17].

Although, it is well known for a long time that neocortical anatomy exhibits a 6-layered structure, modelers have often neglected this fact when modeling a cortical patch by a single “monolithic” neuron population (e.g., [7, 18, 14]).

This may be attributable to the wish to focus on a single layer or the lack of adequate computational resources to simulate more detailed models, but also to doubting or underestimating the functional significance of discrete within- or between-layer synaptic connections which appear to have a rather “fuzzy” character [19, 20].

In accordance with ideas developed earlier in [21] (cf., [22, 23, 24]), our working hypothesis assumes that the *basic function of a cortical column* is to adequately represent and predict (or generate) its sensory inputs. To achieve this in a self-organizing, autonomous way, it is necessary to have access to (at least some of) the different representational modi described above. We propose that different representational modi of the same entity are located in different layers within the same macrocolumn rather than monolithically in different columns or areas.

How can such a *generative model* look like? We can assume that the model represents external states that produce the observed sensory inputs. Thus at each time the model must represent a state v from the state space V (or more generally, a probability distribution on the state space describing in which state the columnar system “believes” to be in). Then the system should be able to use sensory input s to update the state v according to a function f ,

$$v(t + \Delta t) = f(v(t), s(t)) \quad (1)$$

It makes sense to divide the state variable $v = (w, a)$ into two rather independent entities, one variable w describing “external” entities from the outside world, and another variable a describing a local “internal actor”. In addition to updating a state, the system should also be able to predict a future state w' (without accessing sensory input) and sensory inputs s' ,

$$w(t + \Delta t) = f(w'(t), s(t)) \quad (2)$$

$$a(t) = f_a(w(t), \dots) \quad (3)$$

$$w'(t) = g(w(t), a(t)) \quad (4)$$

$$s'(t) = h(w'(t)) \quad (5)$$

Fig. 5 illustrates this functional model and our current implementation employing spiking associative networks similar as discussed in section 2.1. Note that the proposed circuitry provides the basic ingredients for simulating (or predicting over) larger time intervals.

By comparison with known anatomical facts we can match our functional model (Fig. 5) with the layered organization of neocortex ([21, 24, 17, 20]; see Fig. 5). We believe that the forward recognition function f is located in the middle and upper layers, while the remaining functionality, related to behavior and predictions, is located in the lower layers 5 and 6. Furthermore, we believe that the recognition system of the middle and upper layers is split up into two subsystems, one for fast bottom-up recognition (A system, layer 4 and upper 3) and another for refined recognition employing feedback (B system, layers 2 and 3).

3 Comments on the debate topics

In the following I would like to focus on the first debate about “Searle’s Chinese room versus embodiment”, where most of my arguments also apply to the second debate about “Embodiment and the brain”. In general I am inclined to accept the major claims of the embodiment proponents, for example that understanding the meaning of a sentence requires the ability to simulate or predict the situation described by the sentence (but I see no reason why a symbolic system such as a Turing machine should not do that job) Also I believe that many so-called symbolists would finally also agree with these positions, and that probably much of the remaining disagreement results rather from imprecise definitions about what exactly is a symbol and what is embodiment.

In my opinion the whole embodiment debate is neutral to Searle’s arguments. This is because Searle’s argument is finally about Turing machines, i.e., the implementation level, while embodiment vs. symbol systems as intended to be discussed in the upcoming debate appears to be about the higher-level structure of cognitive systems. Of course, any kind of higher-level structures (e.g., the particular way T-symbols are finally implemented in S, see Fig. 1) can be implemented on a Turing machine. In particular, all ingredients required by the embodiment proponents can, of course, also be incorporated in the translation book used by the Englishman in the Chinese room. For example, there is no reason to doubt that the implementation of index-like symbols and affordances (as proposed to explain the on-the-fly understanding of sentences such as “The woman crutched the goalie the ball.” containing innovative denominal verbs; see [27]) could be implemented on a Turing machine, and the canny translation book is essentially equivalent to a Turing machine. Nevertheless, it may well be that the embodiment approach turns out to be the most efficient way (in terms of memory, speed, and learning on the implementation level) to design cognitive systems which can be said to represent meaning and understanding.

The debate aims at elucidating to what degree particular kinds of (T-) symbols should be embodied, e.g., by referring to action and simulation structures. As to my knowledge the current experiments rather demonstrate that understanding the meaning of a sentence interferes to some degree with the generation of actions (e.g. as in the ACH experiments of [27]). But it appears unclear whether this is merely a side effect of associative connections. (For a cell assembly modeler it seems quite natural that understanding a sentence will activate many related associated structures such as action representations, and therefore cause interferences). Thus the critical questions is whether these effects are really functional relevant, and if so, in what respect. I.e., would we still understand the meaning of a sentence if all action related representations could somehow be switched off. Even we could do such experiments (and had a precise definition what means “action related”) the question would remain whether action representations are a necessary part for understanding meaning in any cognitive system or just a side effect of the fact that some brain systems (such as the mirror system [28] in monkeys which has been brought in close relationship to Broca’s

area in humans) are used both for perception and action. E.g., it might then be possible to construct a cognitive system where the multiply used modules are duplicated such that one twin is used just for perception and understanding and the other one just for action (see Fig. 6). This new cognitive system might be behaviorally equivalent to the old one, but we would not detect any meaning/action interferences as described by [27], and the processes in the new twin are purely (S-)symbolic. Thus, would you call such a system still embodied?




Fig. 6: Left: a presumed embodied system: There is a meaning (“mirror”) system that tightly connects the understanding of language (or also of a visual scene) with corresponding actions. The concurrent use of the M-system by perceptual, understanding-related, and action-related processes explains experimentally observable effects such as the ACE. Right: By duplicating a sufficient portion of this system, we get a behaviorally equivalent cognitive system that do not show the interference effects such as ACE. Is this new system still to be called embodied?

If no, then, under the hypothesis that meaning requires embodiment, it is not clear why we should attribute meaning to the first system, but not to the second. If yes, then we have to modify the definition of embodiment from “... meaning [...] depends on activity in systems also used for perception, action, and emotion...”, to something like “... meaning [...] depends on activity in systems which are structurally similar to systems used for action, and emotion...”. This would considerably weaken the definition of “embodiment”, in particular, if the “structurally similar” processes could also be implemented in a symbolic way.

Another flaw to the current

Furthermore, the experimental findings of interference in systems apparently used both for understanding (or perception) and action have apparently another flaw. For example, the interference effects of understanding and action [27] can have two different interpretations with respect to the embodiment hypothesis: (a) understanding hinders action; (b) action hinders understanding. Note that only (b) would support the embodiment hypothesis. Thus, further experiments should contribute to clarify this question.

References

- [1] A. Glenberg, M. De Vega, and A. Graesser. Framing the debate. *The Garachico Workshop on Symbols, Embodiment, and Meaning*, 2005.
- [2] M Wilson. Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4):625–636, 2002.
- [3] F. Pulvermüller. Words in the brain’s language. *Behavioral and Brain Sciences*, 22:253–336, 1999.
- [4] F. Pulvermüller. *The neuroscience of language: on brain circuits of words and serial order*. Cambridge University Press, Cambridge, UK, 2003.
- [5] D.O. Hebb. *The organization of behavior. A neuropsychological theory*. Wiley, New York, 1949.
- [6] V. Braitenberg. Cell assemblies in the cerebral cortex. In R. Heim and G. Palm, editors, *Lecture notes in biomathematics (21). Theoretical approaches to complex systems.*, pages 171–188. Springer-Verlag, Berlin Heidelberg New York, 1978.
- [7] G. Palm. *Neural Assemblies. An Alternative Approach to Artificial Intelligence*. Springer, Berlin, 1982.
- [8] G. Palm. Cell assemblies as a guideline for brain research. *Concepts in Neuroscience*, 1:133–148, 1990.
- [9] D.J. Willshaw, O.P. Buneman, and H.C. Longuet-Higgins. Non-holographic associative memory. *Nature*, 222:960–962, 1969.
- [10] G. Palm. On associative memories. *Biological Cybernetics*, 36:19–31, 1980.
- [11] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science, USA*, 79:2554–2558, 1982.
- [12] K. Steinbuch. Die Lernmatrix. *Kybernetik*, 1:36–45, 1961.
- [13] A. Knoblauch. Neural associative memory for brain modeling and information retrieval. *Information Processing Letters*, 95:537–544, 2005.
- [14] A. Knoblauch and G. Palm. Pattern separation and synchronization in spiking associative memories and visual areas. *Neural Networks*, 14:763–780, 2001.
- [15] L. Barsalou. Perceptual symbol systems. *Behavioral and Brain Sciences*, 22:577–609, 1999.

- [16] A. Knoblauch, R. Fay, U. Kaufmann, H. Markert, and G. Palm. Associating words to visually recognized objects. In S. Coradeschi and A. Saffiotti, editors, *Anchoring symbols to sensor data. Papers from the AAAI Workshop. Technical Report WS-04-03*, pages 10–16. AAAI Press, Menlo Park, California, 2004.
- [17] R.J. Douglas and K.A.C. Martin. Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.*, 27:419–451, 2004.
- [18] R. Ritz, W. Gerstner, U. Fuentes, and J.L. van Hemmen. A biologically motivated and analytically soluble model of collective oscillations in the cortex. II. Applications to binding and pattern segmentation. *Biol. Cybern.*, 71:349–358, 1994.
- [19] M. Abeles. *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge UK, 1991.
- [20] V. Braitenberg and A. Schüz. *Anatomy of the cortex. Statistics and geometry*. Springer-Verlag, Berlin, 1991.
- [21] E. Körner, M.-O. Gewaltig, U. Körner, A. Richter, and T. Rodemann. A model of computation in neocortical architecture. *Neural Networks*, 12:989–1005, 1999.
- [22] J. Hawkins. *On Intelligence*. Times Books Henry Holt, New York, 2004.
- [23] R.P. Rao and D.H. Ballard. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, 1999.
- [24] R.W. Guillery. Branching thalamic afferents link action and perception. *Journal of Neurophysiology*, 90:539–548, 2003.
- [25] A. Knoblauch, R. Kupper, M.-O. Gewaltig, U. Körner, and E. Körner. Design and simulation of a cortical control architecture for object recognition and representational learning. In *Proceedings of the 3rd HRI Global Workshop*. Honda Research Institute, Ito, Japan, 2005.
- [26] M.P. Young. The organization of neural systems in the primate cerebral cortex. *Proceedings of the Royal Society: Biological Sciences*, 252:13–18, 1993.
- [27] A. Glenberg and M. Kaschak. Grounding language in action. *Psychonomic Bulletin & Review*, 9(3):558–565, 2002.
- [28] G. Rizzolatti, L. Fadiga, L. Fogassi, and V. Gallese. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2:661–670, 2001.