

From Goals to Actions and Viceversa: An Analysis of the Ideomotor Principle and the TOTE

Giovanni Pezzulo¹, Gianluca Baldassarre¹, Martin Butz², Gianguglielmo Calvi¹, Cristiano Castelfranchi¹, and Joachim Hoffmann²

¹ ISTC-CNR, Via S. Martino della Battaglia, 44 - 00185 Roma, Italy

² University of Wurzburg, Rontgenring 11 - 97070 Wurzburg, Germany

Abstract. What does it mean for a system to be goal oriented? In this paper we investigate how goals are represented and how goals activate actions. We review the main philosophical and psychological assumptions about the ideomotor principle and we compare it with the TOTE model in cybernetics. We also present three computational architectures which implement in different way goal orientedness, discussing their main peculiarities and limitations.

Keywords: ideomotor, TOTE, goal, representation, anticipation

1 Introduction

Intelligence of more complex organisms such as humans and other apes resides in the capacity conceived as the capacity to solve a problem by working on internal representations of problems, i.e. by acting upon “images” or “mental models” with simulated actions (“reasoning”), before acting in the world. Recently, many converging evidences in psychology and neurobiology indicate a crucial role of anticipatory representations for many cognitive functionalities such as visual and motor control [7]. As suggested by the discovery of mirror neurons [22], representations are mainly action-oriented and deeply based on the motor apparatus [2, 6] try to provide unitary accounts of these phenomena and anticipatory functionalities now begin to be explored from a computational point of view.

We think that by *conceiving representations as mainly anticipatory* it is possible to reframe many of the central claims of AI. In fact, the ability that characterizes and defines a “true mind”, as opposed to a merely adaptive systems, is that of building representations of the non-existent, of what is not currently (yet) “true” or perceivable. A real mental activity begins when the organism is able to endogenously (i.e. not as the consequence of current perceptual stimuli) produce an internal perceptual representation of the world (“simulation” of perception) [3, 7]. For example, the organism can generate internal “images” for matching them against perceptual inputs while actively searching for a given object or stimulus while exploring the environment.

In this paper we discuss how endogenously generated, anticipatory representations can thus be used as *goals* driving the behavior in a teleonomic way: in

Sec. 2 we introduce the *ideomotor principle* and in Sec. 3 the *TOTE*. We discuss the main peculiarities of the two models and we compare them in Sec. 4. In Sec. 5 we introduce three computational architectures implementing in different ways the above discussed principles.

2 The Ideomotor Principle

According to the ideomotor principle [8, 9, 11], *action planning takes place in terms of anticipated features of the intended goal*. [1] underlines the role of anticipation in action selection: *Beginning the movement brings an anticipatory arousal of the [perceptual] trace, and the feedback from the ongoing movement is compared with it*. This is also the position of [5]: *a current response is selected on the basis of its own anticipated sensory feedback*. The *Theory of Event Coding* [10] proposes a common coding in perception and action, suggesting that the motor system plays an important role in perception, cognition and the representation of goals. It focuses on learning action-effect relations which are used to reverse the linear stage theory of human performance (from stimulus to response) afforded by the sensorimotor view. Neurobiological evidences for common mechanisms in perception and action are reported in [12, 23]. [4] suggests that *the goal is represented as a goal-state, namely, as a successfully terminated action pattern*.

In the ideomotor view, in a sense, causality is reversed, since a mental representation of the intended effect of an action is the cause of the action: it is not the action that produces the effect, but the effect that produces the action. [17, par. 21.5] describes an “automatic mechanism” mechanism realizing this principle (see Fig 1): when the features of, say, an apple are endogenously activated, an automatic mechanism is oriented toward (seeing or grasping) apples in a teleonomic way.

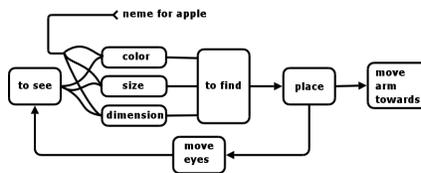


Fig. 1. The “automatic mechanism” in [17, par. 21.5].

[TO BE DONE: 2 paragraphs, Gianluca]: General description of the ideomotor principle, based on Figure 2].

2.1 The main constituents of the ideomotor principle

The comparison of the presentations of the ideomotor principle by the various authors presented in the previous section allows identifying the three main con-

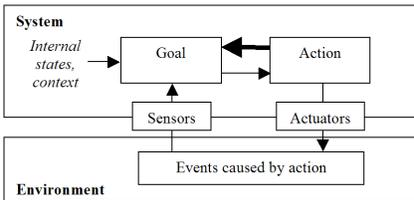


Fig. 2. A scheme that represents the main features of the ideomotor principle. Thin arrows represent information flows, whereas the bold arrow represents the direction of the internal association between goal and action following the causality. See text for further explanations.

stituents of the principle itself. These form the core of the principle and abstract over minor details and different emphases stressed by the various authors. The three constituents are now analyzed in detail.

- *Perceptual-like coding of goals.* An important characteristic of the ideomotor principle is that it has been developed within a vision of intelligence seen as closely related to the sensorimotor cycle (for an example from the psychology literature see [14], whereas for an example from embodied artificial intelligence see [6]). As a consequence, the authors proposing the ideomotor principle usually stress the fact that the system’s internal representation of the goal is similar, or the same, as the internal representations activated by perception. This feature of the principle has also an important “corollary”: the source of goals is usually assumed to be experience, that is goals tend to correspond to previously perceived states. This aspect of the principle is also at the origin of some of its limitations, as shown in the in Sec. 2.2.
- *Learning of action-expectation relations.* Another important constituent of the principle is that the authors tend to stress that experience allows the system to create associations between the execution of actions (e.g., due to exploration, motor babbling, etc.) and the perceived consequences. This requires a learning process that is based on the co-occurrence and contingency of actions and their effects observed in the environment [9]).
- *Goals are used to select actions.* Another core constituent of the principle is the fact that the system exploits the learned association between actions and the resulting perceived states of the world to select actions. As mentioned, the idea is that the activation of the representation of a previously experienced state allows the system to select the action that led to it. When this process occurs, the representation of the state assumes the function of goal because it has an anticipatory nature with respect to the states that the environment will assume in the future, and because it guides behavior so that those states become more likely.

It is important to note that the selection of actions with this process requires an “inversion” of the direction of the previously learned association, from

“actions \rightarrow resulting states” to “resulting states \rightarrow actions”. This inversion is particularly important because it implies that the system passes from the causal association that links the two elements, as resulting from experience, to the teleonomic association between them, as needed to guide behavior.

2.2 Limitations of the Ideomotor Principle

The ideomotor principle has important “limitations”, or, better, “underspecifications” with respect to some important elements that one has to tackle is one wants to use it to design a complete functioning intelligent system.

- *Absence of goal checking system.* The first limitation has to do with the absence of a goal-checking sub-function: what happens if the goal is already achieved? How does the system “know” that it has achieved the goal with the action selected through it? This information is important for the successive correct selection of actions depending of the fact that the pursued goal has been achieved or not. The ideomotor principle is mute on this.
- *Goal or perceived state?* The second limitation originates from the fact that the ideomotor principle postulates that the goal is represented in the same format as the percept generated from the sensation of the state corresponding to it. Often aspect of the principle is so important that the authors supporting it tend to claim the physical machinery use do represent the goals and that used at the higher levels of perceptual processing are the same (e.g. [22]). This raises a problem: how does the system distinguish between the activation of the representation corresponding to a pursued goal and the activation of it caused by the perception of the corresponding state? This information is clearly needed by the system to control action (e.g., see previous point).
- *No abstract goals?* The fact that the literature referring to the ideomotor principle tends to emphasize the role of experience in the formation of potential goals, and the fact that it assumes a common type of representation of goals and percepts, has a consequence on the abstraction potentialities of the system. With this respect the system cannot pursue abstract goals defined in terms of comparisons, for example required by goals of the type: “find the biggest object in the scene”, or “find the farthest object”. In fact in these cases the goal cannot be represented by representations such as templates, prototypes, etc. to be matched, but requires complex processes such as “find an object, store it in memory, find a second object, compare it with the previous one”.

3 The TOTE and cybernetic principles

The **TOTE** (Test, Operate, Test, Exit) was introduced by [16] as the basic unit of behavior, as opposed to the stimulus-response principle. In a TOTE unit, inspired by cybernetics [24], firstly a goal is tested to see if it has been

achieved; if not, the operation is executed until a test is successful and the goal achieved. One of the examples of a TOTE unit is a plan for hammering a nail; in this case, the test consists in verifying if the nail touches the surface and the operation consists in hitting the nail. In this case, the representation used for the test is in the sensory format, and the operation is always the same, even if the TOTE cycle can involve many steps. However, TOTE units can be composed and used hierarchically for achieving more complex goals, also including any kind of representation for the test and any kind of action. The TOTE also inspired many subsequent theories such as the general problem solver [18].

[TO BE DONE (2 paragraphs, Gianluca): General description of TOTE, based on Fig. 3.]

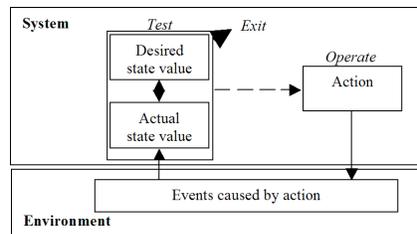


Fig. 3. A scheme that represents the main features of the TOTE. Words in *Italics* represent the main processes composing the principle. Thin arrows represent information flows. The double-headed arrow represents a process of comparison between the desired and the actual state value. The dashed arrow represents the fact that an action is selected and executed in the case the *Test* fails, but not how it is selected. The bold arrow represents a switch in the sequence of processes implemented by the system. See text for further explanations.

3.1 The main constituents of the TOTE

The three main constituents of the TOTE are now analyzed in detail on the basis of the comparison of the positions of the authors presented in the previous section.

- *Test*. A first fundamental constituent of the principle is the internal representation of the desired value(s) of the state of the environment. The representation of this value is a key element of the *Test* sub-process composing the principle. This is the sub-process through which the system repeatedly checks if the current state of the environment matches the goal.
- *Abstract goal*. The desired state value of the system, that is the goal, can be abstract. Indeed, the TOTE is underspecified with this respect, and the literature working with it has used several different types of encodings of goals, from perceptive-like encodings to more abstract, symbolic ones. The principle can manage this type of goals as the *Test* sub-process can be as complex

as needed, from a simple matching to pattern to a more sophisticated process of logical comparison of several features. This (possibly) abstract nature of the definition of goals has also an important implication on the origin of the goal itself, which can derive from previous experience but also from other sources such as other systems (communication) and “imagination” processes.

- *Multiple steps.* An important aspect of the TOTE is the fact that it is naturally suited to implement a course of action formed by multiple steps, as suggested by the repetition of the Test sub-process in its acronym.

3.2 Limitations of the TOTE

As the ideomotor principle, also the TOTE has some limitations and underspecification.

- *Origin of goals.* The first limitation is that the principle does not indicate how the goal representations are generated in the system. Indeed, the literature on TOTE tends to generically assume that goals derive from experience or from outside the system (e.g., from other intelligent systems, conspecific, etc.), but does not give specific indications on this.
- *How are actions selected?* The link between goals and actions executed in the *Operate* sub-process is not specified by the principle. For example, the origin of the knowledge needed to select the suitable actions in correspondence to goals is not specified. This is in line with the fact that the literature on TOTE tends to overlook the role that learning and experience might have in goal directed behaviour. Given this underspecification, the models working on the basis of the principle have adopted various solutions. The most common solution assumes that the controlled state is quantitative, and uses a mechanism that selects and executes actions so as to diminish the difference between the current and the desired values of the state itself.

4 Comparison

[TO BE DONE: Comparison between the ideomotor principle and the TOTE.]

5 Implementation of the Ideomotor Principle in Artificial Systems

After reviewing the literature about the ideomotor principle and the TOTE, in this section we present and discuss some implementations of goal orientedness.

5.1 Case Study I: An Architecture for Visual Search

A hierarchical architecture [20] inspired by the ideomotor principle and by the “automatic mechanism” in [17] is tested in a *Visual Search* task [27]. The goal is to find the red T in a picture containing also many distractors, i.e. green Ts and red Ls. The system can not see all the picture at once, but has a movable spotlight, consisting in three concentric spaces having good, mild and bad resolution (a simplified model of the human fovea).

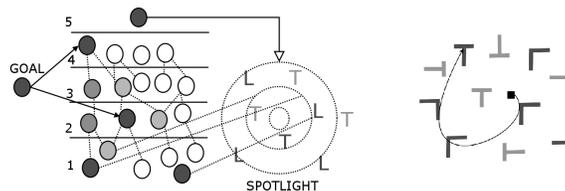


Fig. 4. *Left:* the components of the simulation: the goal, the spotlight and the modules, whose layers are numbered. Light and dark nodes represent more or less active modules. Modules learn to predict the activity level of some modules in the lower layer, which they receive in input (dotted lines). *Right:* a sample trajectory in the visual field, starting from the center (red letters are dark grey, green letters are light grey).

The visual search task is performed by many feature-specific modules, such as color-detectors and line-detectors, organized hierarchically, as illustrated in the left part of Fig 4. Modules have a variable level of activation; more active modules can act more often and, as we will see, influence more the overall computation. Modules in layers 1 and 2 obtain an input from a simulated fovea; the other ones have no access to the fovea, but use as input the activation level of some modules in the immediately lower layer (dotted lines in Fig 4). The architecture has five layers:

- **Full Points Detectors** receive input from portions of the spotlight, e.g. the left corner, and match full or empty points. Modules are more numerous in the inner spotlight than in the central and outer spotlight.
- **Color Detectors** monitor the activity of Full Points Detectors and recognize if full points have the color they are specialized to find (red or green).
- **Line Detectors** categorize sequences of points having the same color as lines; they do not store positions and can only find sequences on-the-fly.
- **Letter Detectors** categorize patterns of lines as Ls or Ts; they are specialized for letters having different orientations.
- **The Spotlight Mover** is a single module; as explained later, it receives asynchronous motor commands from all the other ones (e.g. go to the left) and consequently moves the center of the spotlight.

In the **learning phase**, by interacting with a simulated environment, each module learns *action-expectation pairs*. Modules learn the relations between their

actions and their successive perceptions (the activation level of some modules in the lower layers), as in *predictive coding* [21]. In this way they also learn which actions produce successful matching; for example, a line-detector learns that by moving left, right, up or down the fovea its successive pattern matching operation will be successful (i.e. it will find colored points, at least for some steps), while by moving in diagonal its matching will fail; in this way the line-detector implicitly learns the form of a line by learning how to “navigate” images of lines. In a similar way, a T-detector learns how to find Ts by exploiting the information provided by line-detectors. There is also a second kind of learning: modules *evolve links* toward the modules in the lower layer which they use as input and can successfully predict; for example, T-detectors will link some line-detectors³. These top-down, *generative* links are used for spreading activation across the layers.

The **simulation phase** starts by setting a Goal module (*find the red T*) that spreads activation to the red-detector(s) and the T-detector(s). This introduces a strong goal directed pressure; at the beginning of the task some modules are more active than others and, thanks to the top-down links, activation propagates across the layers. During the search, each module in the layers 2, 3 and 4 tries to move the spotlight where it anticipates there is something relevant for its (successive) matching operation, by exploiting their learned action-perception associations. For example, if a red-detector anticipates something red on the left, it tries to move there the spotlight; a green-detector does the opposite (but with much less energy, since it does not receive any activation from the Goal module). Line- and letter- detectors try to move the spotlight for completing their “navigation patterns”. Modules which successfully match their expectations (1) gain activation, and thus the possibility to act more often and to spread more energy; and (2) send commands to the Spotlight Mover (such as *move left*); the controller dynamically blends them (with some inertia) and the spotlight moves, as illustrated in the right part of Fig 4. In this way the fovea movements are sensitive to both the goal pressures and the more contextually relevant modules, i.e. those producing good expectations, reflecting attunement to actual inputs. The simulation ends when the Goal module receives simultaneous success information by the two modules it controls; this means that the Goal module has only two functions: (1) to start the process by activating the features corresponding to the goal state and (2) to stop the process when the goal is achieved. As reported in [20], this model accounts for many evidences in the Visual Search literature, such as sensitivity to the number of distractors and “pop-out” effects [27].

The ideomotor principle and TOTE in play. According to the ideomotor principle, activity is preceded and driven by an endogenous activation of the anticipated (and desired) goal state. In this case, the goal “find the red T” can be reformulated as “center the fovea in a position in which there is a red T”; and the process starts by pre-activating the features of the desired state, i.e.

³ By learning different sets of action-prediction rules, modules can also specialize: for example, there can be vertical lines recognizers and horizontal lines recognizers

the modules for searching the color red (red-recognizer) and the letter T (t-recognizer); the “finding machine”, once activated, can only search for an object having these features. The key element of the model is the fact that modules embed action-effect rules and are self-fulfilling; when a module is endogenously activated, its effect becomes the *goal* of the system. It is worth noting that this system does not use any map of the environment, but only sensorimotor contingencies [19] and a close coupling between perception and action.

This system can achieve only two kinds of goals: (1) *goal states which were experimented during learning*, such as “find the red T”; and (2) *goal states which are a combination of features*; for example, by combining a green-detector and an L-detector, the system can *find a green L* even if it has never experimented green Ls during learning, but only green Ts and red Ls. On the contrary, this system can not achieve other kinds of goals such as: (1) *The red T on the left*, since locations are not encoded; (2) *The biggest red T*, since there is no memory of past searches and different Ts can not be compared; (3) *The farthest red T*, since temporal features are not encoded. These goals require a more sophisticated procedure for testing and a more abstract encoding: two of the features of the TOTE. The system uses a feature of TOTE: a stopping condition, consisting in a matching between the goal and the activation level of the corresponding features.

6 Case Study II: An Architecture for Reaching

The second system used to illustrate the ideomotor principle and the TOTE in play has been used to control a simple 2D two-segment arm involved in solving sequence-reaching tasks by reinforcement learning. Here we present only the features of the system useful for the purposes of the paper and refer the reader to [12] for the details.

The system is mainly formed by two components, *postural controller* and a *reinforcement-learning component* (“RL component” for short). In a first learning phase, the postural controller learns how to execute sensorimotor primitives that lead the arm to assume a certain posture in space. In order to do so, while the system performs random actions (similarly to “motor babbling” in infants, see [15]), the postural controller learns to categorize the perceived arm’s angles in a 2D self-organizing map [13]. At the same time a 2 layer network (a “perceptron”) is trained, by a supervised learning algorithm [26], to associate the arm’s angles (desired output pattern) with the map’s representation of them (input pattern). This process allows the system: a) to develop semi-local representations of sensorimotor primitives within the self-organizing map, encoded in terms of the corresponding “goals” (i.e. postures, suitably encoded in the map); b) to develop weights between the map and the desired arm’s angles that allow selecting sensorimotor primitives by suitably activating the corresponding goals within the map.

In a second learning phase, the RL component learns to select primitives to accomplish reward-based reaching-sequence tasks, for example in order to reach

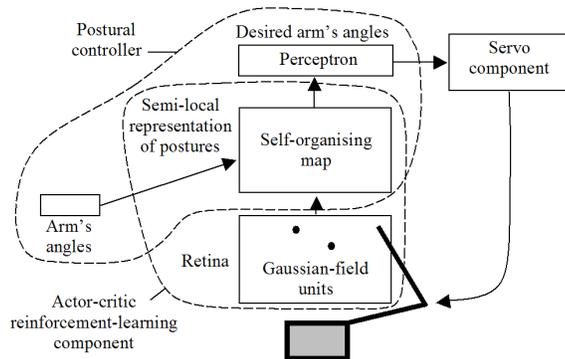


Fig. 5. The architecture of the model of reaching. Rectangular boxes indicate neural layers. Text in boxes indicates the type of neural-network model used. Text near boxes indicates the type of information encoded in the layers. Callouts indicate the two major components of the system. The graph also shows the controlled arm and two targets activating the retina (black dots). See text for further explanations.

two visible dot targets in a precise order (see Fig for an example; the RL component is an actor-critic model [25]). Each time the RL component selects an action, the desired arm's angles produced by it are used to perform detailed movements (variations of the arm's angles) through a hardwired servo-component so that the arm's angles progressively approach the desired ones. When this happens control is again passed to the RL component.

The ideomotor principle and TOTE in play. The reaching system has strong relations with both the ideomotor principle and the TOTE, and in so doing it emphasizes their complementarities. In line with the ideomotor principle, in the first phase of learning (motor babbling) the system performs (random) actions, in this case sensorimotor primitives, and learns to associate the resulting consequences, perceived via the proprioception of the arm's angles, to them. In the second phase of learning, the system uses the expected consequences of the primitives (expected in terms of proprioception) to select the same sensorimotor primitives so as to pursue rewarding states. However, a first important departure from the ideomotor principle is that the “goals” of the primitives (i.e., previously perceived postures), through which the system selects the primitives themselves, are not encoded in a “pure” perceptual-like format, but in terms of different, more abstract representations due to the categorization of the self-organizing map. This might represent a first step towards more abstract representations of goals in the spirit of TOTE.

A second important departure from the ideomotor principle is that the system incorporates a “stop” mechanism for which, when it achieves the goal for which it selected the corresponding sensorimotor primitive, control passes again to the RL component. As we have seen, this is a typical feature of TOTE. Note

how we had to introduce this “stopping” condition to allow the system to accomplish a task requiring the execution of more than one “action” (in this case two sensorimotor primitives), a second typical feature of TOTE. With an opposed perspective, it is interesting to notice how by using some of the core ideas behind the ideomotor principle, the system overcomes some limitations of TOTE. First, it uses experience to encode goals and to associate them to sensorimotor primitives, an core idea of the ideomotor principle. Second, it uses motor babbling to create an association between goals and actions used to achieve them.

7 Case Study II: Anticipatory Classifier Systems

[TO BE DONE]

8 Conclusions

[TO BE DONE]

Acknowledgements This work is supported by the EU project **MindRACES**, FP6-511931.

References

1. J. A. Adams. A closed-loop theory of motor learning. *Journal of Motor Behavior*, 3:111–149, 1971.
2. L. W. Barsalou. Perceptual symbol systems. *Behavioral and Brain Sciences.*, 22:577–600, 1999.
3. Cristiano Castelfranchi. Mind as an anticipatory device: For a theory of expectations. In *BVAI 2005*, pages 258–276, 2005.
4. V. Gallese and T. Metzinger. Motor ontology: The representational reality of goals, actions, and selves. *Philosophical Psychology*, 13(3):365–388, 2003.
5. A. G. Greenwald. Sensory feedback mechanisms in performance control: With special reference to the ideomotor mechanism. *Psychological Review*, 77:73–99, 1970.
6. Rick Grush. The emulation theory of representation: motor control, imagery, and perception. *Behav Brain Sci*, 27(3):377–96, Jun 2004.
7. G. Hesslow. Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6:242–247, 2002.
8. J. Hoffmann. *Vorhersage und Erkenntnis: Die Funktion von Antizipationen in der menschlichen Verhaltenssteuerung und Wahrnehmung [Anticipation and cognition: The function of anticipations in human behavioral control and perception]*. Goettingen, Germany, 1993.
9. B. Hommel. Planning and representing intentional action. *TheScientificWorld JOURNAL*, 3:593–608, 2003.
10. B. Hommel, J. Musseler, G. Aschersleben, and W. Prinz. The theory of event coding (tec): a framework for perception and action planning. *Behavioral and Brain Science*, 24(5):849–78, 2001.

11. W. James. *The Principles of Psychology*. Dover Publications, New York, 1890.
12. G. Knoblich and W. Prinz. *Higher-order motor disorders*, chapter Linking perception and action: An ideomotor approach., pages 79–104. Oxford University Press., Oxford, UK, 2005.
13. T. Kohonen. *Self-Organizing Maps*. Springer-Verlag, Berlin Heidelberg, New York, 3rd edition, 2001.
14. S.M. Kosslyn. *Image and Brain: The Resolution of the Imagery Debate*. MIT Press, Cambridge, 1994.
15. A. N. Meltzoff and M. K. Moore. Explaining facial imitation: A theoretical model. *Early Development and Parenting*, 6:179–192, 1997.
16. G. A. Miller, E. Galanter, and K. H. Pribram. *Plans and the Structure of Behavior*. New York, 1960.
17. Marvin Minsky. *The Society of Mind*. Simon & Schuster, 1988.
18. A. Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA, 1990.
19. J.K. O'Regan and A. Noe. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5):883–917, 2001.
20. Giovanni Pezzulo, Gianguglielmo Calvi, Daniela Lalia, and Dimitri Ognibene. Fuzzy-based schema mechanisms in akira. In M. Mohammadian, editor, *Proceedings of CIMCA'2005*, Vienna, 2005.
21. Rajesh P. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci*, 2(1):79–87, January 1999.
22. G. Rizzolatti, L. Fadiga, V. Gallese, and L. Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3, 1996.
23. Giacomo Rizzolatti and Michael A. Arbib. Language within our grasp. *Trends in Neurosciences*, 21(5):188–194, 1998.
24. A. Rosenblueth, N. Wiener, and J. Bigelow. Behavior, purpose and teleology. *Philosophy of Science*, 10(1):18–24, 1943.
25. R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge MA, 1998.
26. B. Widrow and M.E. Hoff. Adaptive switching circuits. In *IRE WESCON Convention Record*, page 96104, 1960. Part 4.
27. J. M. Wolfe. *Attention*, chapter Visual search. London, UK: University College London Press, 1996.