

# Knowledge Representation, Learning, and Problem Solving for General Intelligence

Seng-Beng Ho    Fiona Liausvia

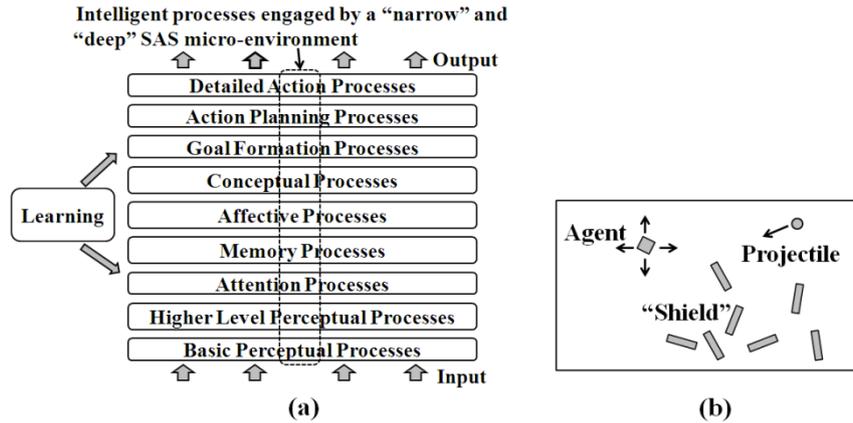
National University of Singapore, Singapore  
hosengbeng@nus.edu.sg    tsflf@nus.edu.sg

**Abstract.** For an intelligent agent to be fully autonomous and adaptive, all aspects of intelligent processing from perception to action must be engaged and integrated. To make the research tractable, a good approach is to address these issues in a simplified micro-environment that nevertheless engages all the issues from perception to action. We describe a domain independent and scalable representational scheme and a computational process encoded in a computer program called LEPS (Learning from Experience and Problem Solving) that addresses the entire process of learning from the visual world to the use of the learned knowledge for problem solving and action plan generation. The representational scheme is temporally explicit and is able to capture the causal processes in the visual world naturally and directly, providing a unified framework for unsupervised learning, rule encoding, problem solving, and action plan generation. This representational scheme allows concepts to be grounded in micro-activities (elemental changes in space and time of the features of objects and processes) and yet allow scalability to more complex activities like those encountered in the real world.

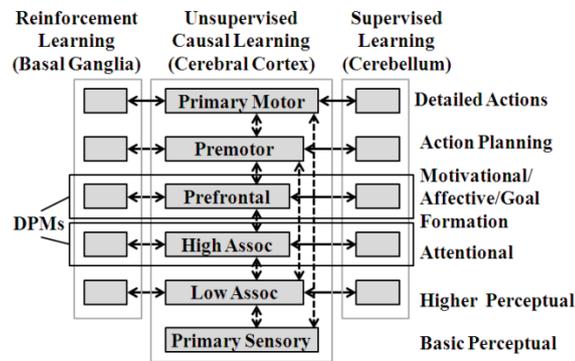
## 1 Introduction

In a previous paper [7] we laid out a general approach for creating fully autonomous and adaptive artificial general intelligence. The idea, illustrated in Fig. 1(a), is that firstly one must engage the various aspects of intelligent processes from that of the perceptual to that of the attentional, the memory, the affective, the conceptual, the planning, and the action. Secondly, to make the approach tractable, one can study a thin slice of the processes – i.e., study these processes in a simplified micro-environment - but one must engage the entire depth of processes from perception to action. In [7] a Shield-and-Shelter (SAS) micro-environment was proposed for this purpose - a “narrow” and “deep” engagement of the various processes from perception to action as shown in Fig. 1(a). Examples of a narrower width of processes at the perceptual level could be processes that handle just 2D space, instead of 3D space, and ignore certain visual properties such as texture. A narrower width of processes at the affective level could be processes that handle a handful of built-in internal signals such as pain and hunger but exclude signals like temperature, reproductive needs, etc. Fig. 1(b) shows the SAS micro-environment as proposed in [7]. It consists of an agent, a “projectile” that can hurt the agent, and some rectangular objects that can be

used as shields or reconfigured into a “shelter.” The previous paper [7] described the detailed specifications of the SAS micro-environment but suffices it here to say that it can engage all the 10 aspects of processes as depicted in Fig. 1(a) and also, for the affective domain, it can engage the basic value signal “pain” as well as higher level emotions such as “anxiousness,” “relief,” “stressed,” and “comfort.”



**Fig. 1.** (a) Various aspects of intelligent processes [3] and the width of engagement of the proposed Shield-and-Shelter (SAS) micro-environment. (b) The SAS micro-environment. From [7].



**Fig. 2.** A brain-inspired general cognitive architecture that incorporates, in a hierarchical structure, general learning modules (DPMs) that are applicable to a wide range of cognitive processes and that is hence scalable. The cerebral cortical areas in the Unsupervised Causal Learning section are typical cortical areas performing various functions from perception to action. From [7].

The previous paper [7] also reviewed neuroscience literature and presented the findings that general learning modules (called the Distributed Processing Modules – DPMs [10]) each of which consisting of a reinforcement learning sub-module (the basal ganglia), an unsupervised causal learning sub-module (the cerebral cortex), and

a supervised learning sub-module (the cerebellum) are present in the brains (of humans and animals) for the processing of the different domains of information from that of the motivational, to that of the executive, the motor and the visual. This idea is illustrated in Fig. 2 (to avoid clutter, only 2 DPMs are highlighted). I.e., no matter the nature of the information processing from the perceptual stage to the action stage, similar general modules – the DPMs - are used. This provides an inspiration for artificial general intelligence. Can we construct artificial general intelligence systems likewise with general modules of learning and processing that are applicable to the various aspects of processing depicted in Fig. 1(a)?

The contention of the previous paper [7] is that if a general solution can be discovered for the SAS micro-environment, that solution would be easily scalable “width-wise” to encompass a wider range of processes that allow an adaptive autonomous agent (AAIA) to behave intelligently in the real world.

[6] and [8] presented a novel, general, and unified representational scheme based on explicit spatiotemporal representations that can be used in the various levels of processing from that of the perceptual to that of the action. This allows the various learning mechanisms depicted in Fig. 2 to operate in a similar manner across the various aspects of intelligent processing. This representational scheme allows concepts to be grounded in micro-activities (elemental changes in space and time of the features of objects and processes) and yet allow scalability to more complex concepts and activities. Explicit spatiotemporal representations have also been used in cognitive linguistics to represent higher level concepts [12]. Our method thus converges with cognitive linguistics. The current paper attempts to flesh out part of the SAS micro-environment solution using this general and unified spatiotemporal representational scheme. It will use similar but altered parts of the SAS micro-environment shown in Fig. 1(b) to address the various issues.

An AAIA must have the ability to learn about how objects behave, including how its actions may influence these objects by observing the activities and interactions through the visual input, and then use the learned knowledge for problem solving and various cognitive processes. A lot of research has been carried out in the area of computer vision in which the information in the visual world is processed for the purposes of object recognition, scene description, and activity classification (e.g., [1], [16], [17], [20], [22]) but these efforts have not emphasized characterizing the conceptual and causal aspects of the visual world based on the activities and interactions observed. In qualitative physics or naïve physics ([9], [13], [21]), knowledge of the physical behavior of objects and causal agents in the physical world is encoded and represented for reasoning and problem solving purposes but this area of research does not emphasize the learning of the knowledge directly from the visual input. Traditional AI problem solving and action planning as reported in the early days [2] and recently [15] also do not emphasize the learning of the knowledge directly from the visual input, visual observation, and visual experience.

## **2 From Visual Learning to Action Plan Generation**

Our paper attempts to address the entire process of learning from the visual world to the use of the learned knowledge for problem solving and action plan generation. This

is divided into two parts. The first part (this section - Section 2) shows how this can be done for *physical* learning and problem solving – learning of the knowledge of physical behaviors of objects and the subsequent moving of physical objects into desired physical locations using the learned and encoded physical causal rules. (This addresses the topmost two levels and the bottommost two levels of the brain-inspired general cognitive architecture in Fig. 2). The second part (Section 3) shows how this can be done for the learning of *motivational* factors and motivational problem solving (which addresses one of the middle levels as well as the two bottommost levels in Fig. 2) using similar learning, representational, and computational devices as those in the first part.

## 2.1 Importance of Micro-Activities

We define micro-activities (elemental changes in space and time of the features of objects and processes) as the most elemental steps activities can be decomposed into in space and time. Given a certain spatial and temporal resolution of a visual capture device, the most elemental changes that can be characterized would be changes over a pixel of space and over the smallest time step that the visual capture device is able to capture and process the incoming dynamic information.

A macro-activity that involves a more complex object such as a walking person can be decomposed into myriads micro-activities at every part of the body, including the minute movements of the hair or finger tips. In the current paper, we deal with much simpler objects than human beings – simple shapes with no sub-parts that move independently. The basic method here can be extended to handle these more complex objects with complex subparts. However, it is important to capture and represent micro-activities as these are necessary to fully characterize the functioning and temporal characteristics of various objects, events, and processes – examples would be the flapping of cloth in the wind, the operations of minute mechanical parts in a watch, etc. Micro-activities are hence the foundation of all arbitrarily complex macro-activities. We have a prescription for how to scale up the micro-activity representations to handle much more complex activities that we intend to explore in a future paper. The scalability of the method can also be seen in cognitive linguistics [12].

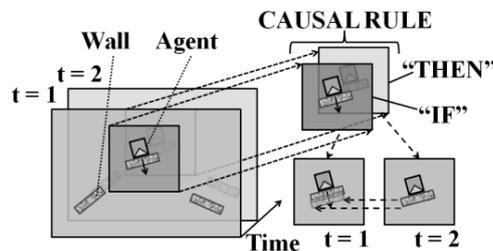
To facilitate the learning of micro-activities and the characterizations of macro-activities based on micro-activities, we propose the use of “Temporally Explicit Micro-Activity” (TEMA) representations ([6], [8]). This is described in the next section.

## 2.2 Unsupervised Causal Learning for Micro-Activities

Taking actions and causing things to change in the physical world is a complex process. Sometimes, the same action may not generate the same effect, and sometimes, more than one action is needed to generate a desired effect. To disentangle complex situations of causality requires a sophisticated causal analysis process ([4], [11], [14]) but once the causality can be established, it can be used beneficially for problem solving processes and to enhance the “survivability” of the AAIA. In our micro-environment, we have simplified the causal analysis process – we assume that a given

action always leads to a given effect and that the agent, being the generator of the action, assumes that its action is the cause of the effect.

Fig. 3 shows the use of a temporally explicit micro-activity (TEMA) representation to capture an elemental causal event in the micro-environment. The event involves the agent applying a forward (relative to its “face”) force and it and the object – a “wall” - it is in touch with in the forward direction then move by an elemental step as a consequence. (In order to simulate the physical effects in our micro-environment for the learning system/agent to learn about physical behavior, we have implemented a Physics Simulator that encodes various physical behaviors and updates the micro-environment accordingly.)



**Fig. 3.** The use of TEMA representations to capture an elemental causal event. The agent has a “face” represented by the little triangle within it and the rectangular objects can be thought of as “walls.”

The basic method employed in Fig. 3 is a “cookie-cutter” mechanism akin to that used in pattern recognition [19] in which parts of a pattern are “cut-out” to form potentially useful features to characterize the pattern for future recognition purposes, except that instead of a “spatial slice” that is being cut-out in the pattern recognition case, here a “spatiotemporal slice” is “cut-out,” along with the action and activities that occurs in the spatiotemporal slice. Once an elemental spatiotemporal slice is cut-out as shown in Fig. 3, it basically becomes a causal rule. The “action” (generated by the agent) in the slice is being labeled as the “cause,” and this rule is encoded in the form of the spatiotemporal pattern as depicted in Fig. 3 (and not transformed into any propositional or verbal form). The rule has a “IF” part (if this pattern appears in the world at  $t=1$  – in this case, a “forward” force generated by the agent and the agent is touching the “wall”) and a “THEN” part (then this pattern will ensue at  $t=2$  – in this case, an elemental displacement of both the agent and the “wall”) and for subsequent reasoning and problem solving tasks, the “IF” and “THEN” parts are simply pattern-matched to the situations and events perceived in the physical world or to the “goal” pattern in the agent’s internal memory for the purpose of forward or backward reasoning. The rule can be encoded directly as a bitmap pattern or in a vector form. For the efficiency in storage, we choose the vector form encoding for the causal rule.

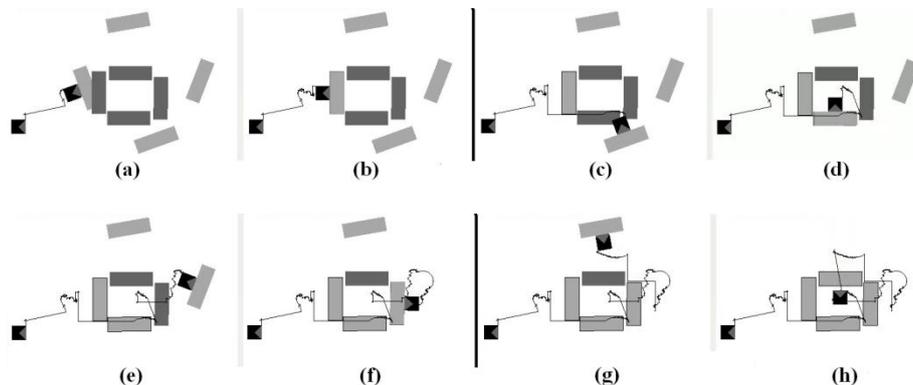
The agent, having interacted with the environment and extracting a variety of elemental causal rules (including twisting itself and causing the “wall” to turn in unison, attaching itself to the “wall” and pulling it, etc.), would store these causal rules in a Causal Rule Base. (In our current implementation, elemental rules capture elemental movements - for translational movements they are 1 pixel movements and for rotational movements, they are 1 degree movements.) Another kind of causal rules

are Chunked Causal Rules that derive from a series of elemental actions/movements discovered in some problem solving/search processes. These are also stored in the Causal Rule Base along with the elemental causal rules. The chunked sequence of micro-activities can be thought of as a “learned macro-activity” that may be of use in future problem solving situations.

### 2.3 Simulations

We have coded the unsupervised learning process described above in a program called LEPS (Learning from Experience and Problem Solving). For problem solving, we use an A\* search process [5]. The “problem” is presented as a goal consisting of various objects at some desired positions and orientations. The cost of each state reached is computed based on the sum of the number of elemental steps needed to reach that state and a closeness measure of that state to the goal state.

After a learning/developmental phase in which the agent learned about elemental actions as well as acquired some chunked rules given some relatively simple problems – in which the goals were not too far away - we posed a harder problem to the agent as shown in Fig. 4 – to push a number of walls from some initial positions (shown in lighter gray) into the final positions (shown in darker gray) and construct a “shelter.” Figs. 4(a) – (h) show the sequence of actions for constructing the “shelter.” A video of the entire process can be seen at <http://www.youtube.com/watch?v=W0YVSOu1xbo>.

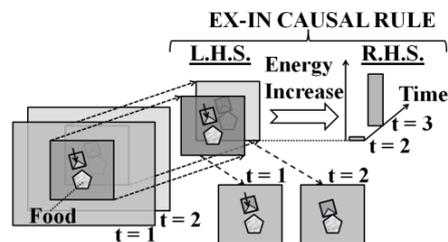


**Fig. 4.** A “shelter building” problem in which the agent is to move the walls in the initial positions (shown in lighter gray) into the final positions (shown in darker gray). A video of the process can be seen at <http://www.youtube.com/watch?v=W0YVSOu1xbo>.

We also tested the ability of LEPS to learn new physical rules. After the above experiments, we added a “large forward translational force” to the micro-environment. Whereas before, the translational force would move the agent and wall by one pixel in one elemental time frame, this large force would move them by 3 pixels in one elemental time frame. The corresponding causal rule is learned automatically through the unsupervised causal learning process (Fig. 3). The solution obtained for the shelter building were similar to those before (Fig. 4) except that now the large force is incorporated in many places among the solution steps and the agent is able to move itself and the wall to the goal positions faster.

### 3 Generalizations of the Representational Scheme to Motivational Problem Solving

In this section, we describe how the same TEMA representational scheme can be used in learning and problem solving in the *motivational* domain – e.g., to satisfy a certain energy need, the agent learns that certain items (food) in the real world can increase its energy level, and it encodes these experiences in spatiotemporal causal rules so that later these can be used in a similar problem solving process as that described in Section 2 to satisfy its energy need.



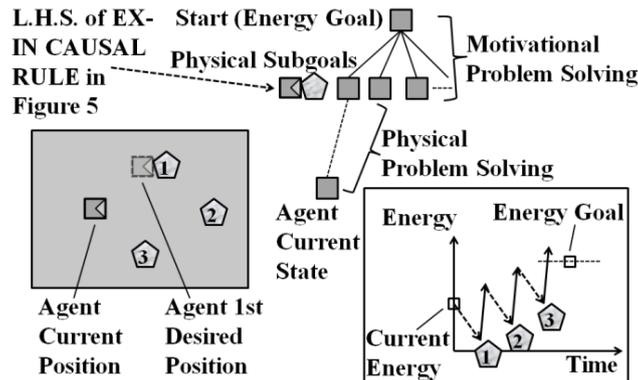
**Fig. 5.** At  $t=1$ , agent is one elemental distance away from food (pentagonal shape). At  $t=2$ , agent touches/eats food. At  $t=3$ , internal energy of the agent increases by a certain amount.

#### 3.1 Unsupervised Causal Learning for Energy-Enhancing Causes

Unlike in the case of the physical processes as depicted in Figs. 3 and 4, in which the effects of the causes take place in the outside world (the “physical” world), in the motivational domain, the effects (such as “pain,” “energy gain,” “depression,” “anxiousness,” etc.) of various causes (such as “pain-causing objects,” “food,” “impending doom,” etc.) are changes to some internal states of an agent. Fig. 5 shows what happens when an agent encounters a piece of food (the pentagonal shape). It moves one elemental step toward the food (from  $t=1$  to  $t=2$ ) and upon touching/eating the food (at  $t=2$ ), its internal energy goes up by a certain amount (at  $t=3$ ). (The energy change is represented using TEMA like that for the *physical* micro-activities above.) The top right corner of Fig. 5 shows how an “EX-IN CAUSAL RULE” (meaning a rule encoding “some changes in external states causing changes in internal states”) is then extracted from this visual/physical experience of the agent. The left-hand side (L.H.S.) of the rule shows the touching/eating event and the right-hand side (R.H.S.) shows the energy being increased. We have built an External-Internal Causal Effects Simulator, much in the spirit of the Physics Simulator, to simulate the effects of various kinds of food on the internal energy states of the agent. Some of this food will increase the energy of the agent by a different amount than that effected by the pentagonally shaped food in Fig. 5 and some will take more than one time frame to increase the energy. TEMA can represent any kind of energy change profile or pattern. An unsupervised causal learning “cut-out” mechanism much like what was described for the physical case in Fig. 3 encode the entire causal event in the EX-IN CAUSAL RULE form as shown in Fig. 5.

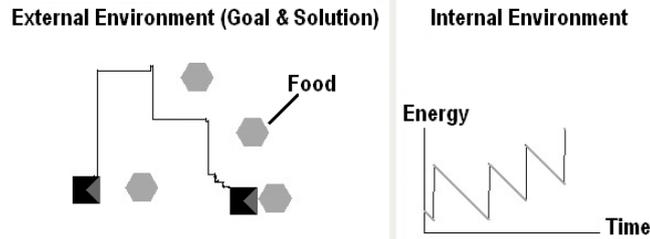
### 3.2 Motivational Problem Solving with Ex-In Causal Rules

A motivational goal can be specified for the agent to reach in a motivational problem solving process. An example is shown in Fig. 6. The micro-environment on the left side of the figure has 3 pieces of food labeled 1, 2, and 3 and the agent is currently not near anyone of them. The right bottom corner of the figure shows a specification of an energy goal that is higher than the current energy of the agent. In the top right corner of the figure it is shown that a Motivational Problem Solving process is begun by a backward search process looking for rules in the Causal Rule Base that will increase the energy. Looking for a rule that will increase the energy to reach the specified energy level is like applying the physical heuristic in Section 2 to look for rules that will bring the agent closer to the goal. The R.H.S. of the food consumption EX-IN rule as depicted in Fig. 5 matches the energy increase specification and the L.H.S. of the rule establishes a physical subgoal, which is the presence of the agent next to the food (and touching it before food consumption can take place). This subgoal will trigger a Physical Problem Solving process such as that described in Sections 2 to bring the agent from the current position to the desired position for food consumption.



**Fig. 6.** Energy problem solving. There are 3 pieces of food labeled 1, 2 and 3 and the energy level picture is on the right. See text for explanation.

Now, in an earlier learning process, the agent learned that by applying a force to move itself it would consume energy and the energy will decrease (in Sections 2 we did not consider this energy consumption factor). The learned causal rule for this was also stored in the Causal Rule Base. The system will forward simulate what happens to the energy level if the agent first moves to food number 1. There will be a decrease in energy level and then there will be an increase in energy level when food number 1 is consumed. If the decrease in the energy causes the agent to be energy exhausted (i.e., the energy falls below 0), then this is considered a search deadend and another piece of food will be considered next. In our example in Fig. 6 when all 3 food pieces are consumed, the agent is able to raise its energy level above the energy goal. LEPS' output of the energy problem solving process is shown in Fig. 7 and also a video of the process can be seen at <http://www.youtube.com/watch?v=exoVU0dX4RQ>.



**Fig. 7.** LEPS' output of the energy problem solving process described in Section 3. A video of the process can be seen at <http://www.youtube.com/watch?v=exoVU0dX4RQ>.

## 4 Discussion and Conclusions

In this paper we described a unified framework of intelligent processes, employing the same temporally explicit micro-activities (TEMA) representations, from unsupervised causal learning and encoding of causal rules of activities and interactions in the physical world to using the rules for problem solving/action plan generation, for both the physical level as well as the motivational level. The system is adaptive – there can be new physical behavior or new kinds of food made available, and it will learn about them and use them for problem solving accordingly.

To map the various components in our current system to the general biologically-inspired brain architecture consisting of the three modules of unsupervised causal learning, supervised learning, and reinforcement learning depicted in Fig. 3, firstly the unsupervised causal learning would correspond to that we have described both for the physical as well as for the motivational problem solving in Figs. 3 and 5 respectively. As mentioned in Section 2, the physical problem solving processes we described in Sections 2 would cover the topmost two and bottommost two levels of Fig. 2. The motivational problem solving described in Section 3 would cover the bottommost two levels as well as one of the two middle levels – the motivational processes – of Fig. 2. Secondly, the Causal Rule Base would be identified with the cerebellum as the cerebellum has been identified to be involved in encoding mental models of the world [18] which is essentially the function of the causal rules as we have described in this paper. As for reinforcement learning, the problem solving processes for both the physical as well as the motivational levels employ an intermediate reward signal – causal rules that result in a state closer to the goal are selected for consideration first.

This general TEMA representational scheme ([6], [8]) and the associated learning and problem solving architecture that can be used across a number of domains of processing, as has been illustrated in this paper, can form the core and basis for further extension and scaling up to handle more complex processes often encountered in the real world. We will not belabor this except to refer to [6] and [8] for more detail on the scalability of the representational scheme. As our explicit spatiotemporal representations converge with that of cognitive linguistics, the scalability of the representational scheme can be seen in the higher level concepts represented in cognitive linguistics [12].

## References

1. Ali, S., Shah, M.: Human action recognition in videos using kinematic features and multiple instance learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(2), 288-303 (2010)
2. Fikes, R., Nilsson, N.: STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2, 189-208 (1971)
3. Gazzaniga, M.S., Ivry, R.B., Mangun, G.R.: *Cognitive Neuroscience*, 2<sup>nd</sup> ed. W.W. Norton, New York (2002)
4. George, D., Hawkins, J.: Towards a mathematical theory of cortical micro-circuits. *PLoS Computational Biology*, 5(10), e100532 (2009)
5. Hart, P.E., Nilsson, N.J., Raphael, B.: A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics SSC4*, 4(2), 100-107 (1968)
6. Ho, S.-B.: The Atoms of Cognition: A Theory of Ground Epistemics. In: *Proceedings of the 34<sup>th</sup> Annual Meeting of the Cognitive Science Society*, pp. 1685-1690, Cognitive Science Society, Austin, Texas (2012)
7. Ho, S.-B.: A Grand Challenge for Computational Intelligence. In: *Proceedings of the IEEE Symposium Series for Computational Intelligence*, pp. 44-53, Singapore (2013)
8. Ho, S.-B.: The Atoms of Cognition: Actions and Problem Solving. To appear as member abstract in: *Proceedings of the 35<sup>th</sup> Annual Meeting of the Cognitive Science Society*, Cognitive Science Society: Austin, Texas (2013)
9. Hobbs, J.R., Moore, R.C. (eds.): *Formal Theories of the Commonsense World*. Alex Publishing, Norwood, New Jersey (1985)
10. Houk, J.C.: Agents of the mind. *Biological Cybernetics*, 92, 427-437 (2005)
11. Kersten, D., Mamassian, P., Yuille, A.: Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271-304 (2004)
12. Langacker, R.W.: *Cognitive Grammar: A Basic Introduction*. Oxford University Press, Oxford (2008)
13. Liu, J., Daneshmend, L.K.: *Spatial Reasoning and Planning: Geometry, Mechanism, and Motion*. Springer-Verlag, Berlin (2004)
14. Pearl, J.: *Causality*. Cambridge University Press, Cambridge, UK (2009)
15. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Pearson, Boston (2010)
16. Shapiro, L.G., Stockman, G.C.: *Computer Vision*. Prentice Hall (2001)
17. Szeliski, R.: *Computer Vision: Algorithms and Applications*. Springer, Berlin (2010)
18. Ito, M.: Control of mental activities by internal models in the cerebellum. *Nature Reviews Neuroscience*, 9, 304-313 (2008)
19. Uhr, L., Vossler, C.: A pattern-recognition program that generates, evaluates, and adjusts its own operators. In: Feigenbaum, E.A., Feldman, J. (eds.): *Computers and Thought*. Robert E. Krieger Publishing Company, Inc., Malabar, Florida (1981)
20. Wang, Y., Mori, G.: Human action recognition by semi-latent topic models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(10), 1762-1774 (2009)
21. Weld, D.S., de Kleer, J. (eds.): *Readings in Qualitative Reasoning About Physical Systems*. Morgan Kaufmann Publishers, Inc., San Mateo, California (1990)
22. Yuan, J., Liu, Z., Wu, Y.: Discriminative video pattern search for efficient action detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(9), 1728-1743 (2011)