

Perception as Abduction

Turning Sensor Data into Meaningful Representation

Murray Shanahan

Department of Electrical and Electronic Engineering,
Imperial College,
Exhibition Road,
London SW7 2BT,
England.

m.shanahan@imperial.ac.uk

+44 (0)207 594 6307

Abstract

This paper presents a formal theory of robot perception as a form of abduction. The theory pins down the process whereby low-level sensor data is transformed into a symbolic representation of the external world, drawing together aspects such as incompleteness, top-down information flow, active perception, attention, and sensor fusion in a unifying framework. In addition, a number of themes are identified that are common to both the engineer concerned with developing a rigorous theory of perception like the one on offer here, and the philosopher of mind who is exercised by questions relating to mental representation and intentionality.

1. Introduction

Philosophers who subscribe to a computational theory of mind and AI researchers who work in the “classical” symbolic style are both vulnerable to awkward questions about how their symbols and representations acquire meaning. The issue dates back at least as far as Franz Brentano, whose concern was how mental states gain their *intentionality*, their “directedness” towards things (Brentano, 1874). Related philosophical spectres have surfaced more recently in the form of the Chinese Room argument (Searle, 1980) and the symbol grounding problem (Harnad, 1990).

A typical way to parry such arguments is by appeal to the notion of embodiment (Chrisley & Ziemke, 2003). A system that only interacts with

the outside world via a keyboard and screen is said to be *disembodied*. The champions of embodiment are willing to concede that the symbols deployed in such a system lack true intentionality, that they inherit whatever meaning they have from the human designer. By contrast, a robot that enjoys genuine sensory-motor interaction with the same physical world as ourselves is said to be *embodied*. Symbolic representations built up by a robot through sensory-motor interaction with the physical world are *grounded*, and can be properly said to have intentionality.¹

This is a promising line of argument, as far as it goes. But it opens a great many questions. In particular, it needs to be filled out with a theoretically respectable account of the process involved in constructing the symbolic representations in question (Barsalou, 1999).² The present paper attempts to plug this gap by offering a formal abductive account of the means by which low-level sensor data is transformed into meaningful representation. The account incorporates detailed treatments of a variety of issues, including incomplete and uncertain sensor data, active perception, and top-down information flow, and it briefly addresses a number of other topics including attention and sensor fusion.

The methodological context here is robotics, and one of the strengths of the theory on offer is that it has been implemented on a number of robot platforms. No explicit claim is made about the extent to which it applies in the human case. However, at the close of the paper various philosophical questions relating to meaning and intentionality are placed, for comparison, alongside answers an engineer might give if the same questions had been asked about a robot built according to the theory of perception being proposed.

While the paper draws heavily on work carried out in the area of logic-based artificial intelligence, its primary focus is not the challenge of building programs that use explicitly logical forms of representation and inference. Rather, it is to supply an account of robot perception that is

¹ It can also be argued that symbols can be grounded through interaction with a virtual environment, or through disembodied interaction with the real world (through the Internet, for example). The concern of the present paper, however, is grounding through embodiment.

² Barsalou's own way of tackling this problem (for humans) is very different from the classical AI solution for robots offered here. Both approaches can be valid for their respective domains of application.

couched in a precise theoretical language, and that is capable of engaging with a variety of contemporary issues. Naturally the theory *will* appeal to strictly logicist AI researchers, and the paper does present work in which logic is used as an explicit representational medium. However, the account also serves as a theoretical yardstick for roboticists working with symbolic AI but outside the logicist tradition.³ In addition, it can be viewed as supportive of philosophers of mind who favor a computationalist theory, and who are open to the charge that they have failed to supply such an account (Barsalou, 1999, p.580).⁴

The paper is structured into three main parts. In the first part, the abductive theory of perception is outlined, and some of its methodological underpinnings are discussed along with such topics as incompleteness, top-down information flow, active perception, and sensor fusion. In the second part, a number of experiments are sketched in which the abductive theory of perception is put to the test with several real robots ranging from simple mobile platforms to an upper-torso humanoid with stereo vision. These experiments supply the context for a series of worked examples that are presented in more technical detail. In the third part, the paper presents a few philosophical reflections on such subjects as the representational theory of mind and intentionality.

2. An Abductive Account of Perception

We begin with a very basic logical characterization of perception as an *abductive* task (Shanahan, 1996a). Given a stream of low-level sensor data, represented by the conjunction \square of a set of observation sentences, the task of perception is to find one or more explanations of \square in the form of a consistent logical description \square of the locations and shapes of hypothesized objects, such that,

$$\square \square \square \models \square$$

³ This distinction between logic as a representational medium and logic as a specification tool was first clearly articulated by Newell (1982).

⁴ For the record, this author's view is that the logicist approach is a viable way to engineer good robots. However, this does not entail endorsement of a computational theory of the *human* mind. From a philosophical point-of-view, the present paper shows that if the computational theory of mind

where \mathcal{T} is a background theory describing how the robot's interactions with the world impact on its sensors. The form of \mathcal{T} , that is to say the terms in which an explanation must be couched, is prescribed by the domain. Typically there are many \mathcal{T} s that explain a given \mathcal{D} , so a preference relation is introduced to order them. When deployed on a robot, a perceptual system conforming to this theoretical description is a perpetual process, embedded in a closed control loop, that outputs a constant stream of (sets of) \mathcal{D} s in response to a constant stream of incoming \mathcal{S} s from the robot's sensors, which could be anything from bump switches to a stereo camera.

Observation sentences in \mathcal{D} will not, in general, describe sensor data in its raw state. The sensor data in \mathcal{D} is only "low-level" in the following sense. Whatever pre-processing the raw data has undergone to produce it, \mathcal{D} is never taken as a description of the external world, but rather as a more refined description of the sensor data. For example, \mathcal{D} might be a description of edges extracted from an image using the Sobel operator, or it might be a sequence of sonar data points that has been through a Kalman filter. Even if the pre-processing applied to the raw sensor data generates information that more obviously purports to describe the world, such as the depth map produced by a stereo matching algorithm, the abductive theory of perception treats this information simply as a rich form of low-level data. It is strictly the responsibility of the abductive process to decide how to interpret each item of low-level data — whether it should be accepted or rejected, and what it says about the presence, identity, and configuration of objects in the external world.

The real utility of any implementation of perception as abduction will be dictated by the form and content of the background theory \mathcal{T} — the underlying formalism in which it is expressed, the ontology manifest in its choice of sorts and predicates, and the causal knowledge captured by its constituent formulae. Most of the experiments carried out to date have deployed an underlying formalism capable of representing action, change, and causal relations, and with a rudimentary capability for representing space and shape. As far as content goes, in general \mathcal{D} has two parts — a set

is flawed, this is not for want of an account of the transition from sensor data to meaningful

of formulae describing the effect of the robot's actions on the world, and a set of formulae describing the impact of changes in the world on the robot's sensors. In an ideal implementation of the theory, \square would incorporate a full axiomatization of naïve physics (Hayes, 1985), either hand-coded or acquired by a compatible learning method, such as inductive logic programming (Muggleton, 1991).

2.1 Incompleteness and Uncertainty

Shortly, we'll see how the basic abductive framework above can be extended to accommodate active perception, attention, and top-down expectation. But even this simple formal description is sufficient to demonstrate several methodological points. In particular, we can lay to rest the false notion that a representational approach to perception assumes the need to build and maintain "a complete and accurate model of the environment" (Brooks, 1991).

Vision, according to David Marr, is "the process of discovering from images what is present in the world and where it is" (Marr 1982, p.3). In the past, researchers have taken this to mean the construction of "a complete and accurate representation of the scene", and in order to live up to this definition have been forced to make gross idealizations about the environment. For example, an introductory chapter on constraint-based vision from a 1980s AI textbook invites the reader to consider a world of "crack-free polyhedra with lighting arranged to eliminate all shadows" (Winston, 1984, p. 47). This is patently an unsatisfactory starting point if we want to build robots that work in realistic environments. Moreover, it sidesteps two issues central to understanding intelligence, namely incompleteness and uncertainty, which are unavoidable features of what John McCarthy calls the "common sense informatic situation" (McCarthy, 1986).

Uncertainty arises because of noise in a robot's sensors and motors, and is inevitable because of the physical characteristics of any electrical device. Incompleteness, on the other hand, arises for two reasons, both of which run deeper than the physics. First, incompleteness results from having a limited

representation.

window on the world. No robot can ever sense more than a small portion of its environment because of the inherently limited reach of any sensor and because of its singular spatial location. Second, incompleteness is an inevitable (and welcome) consequence of the individuality of each robot’s goals and abilities. So there is no need for a robot’s internal representations to attempt to capture every facet of the external world.

In the same way, an animal’s perceptual system is adapted to survival within its particular ecological niche. The subjective visual world of a bee — its *Umwelt*, in the terminology of biologist Jakob von Uexküll — might be thought of as an otherwise meaningless blur in which pollen-bearing flowers stand out like beacons (von Uexküll, 1957). In a similar vein, the *Umwelt* of a robot designed to fetch coke-cans might be characterized as a meaningless chaos punctuated by can-shaped objects. So, although it may be a methodological error to presuppose a world of “crack-free polyhedra,” it still makes sense to build a robot whose perceptual system is tailored to pick out nearly polyhedral objects (such as books), so long as it can do so in a normal, messy working environment.

Mathematical logic is ideally suited to expressing each of the above forms of incompleteness and uncertainty, while at the same time possessing a precise semantics. The trick, in the context of the abductive treatment of perception, is to carefully design the permitted form of \square to take advantage of this. To illustrate, consider the following formula, which constitutes a partial description of an object in the robot’s workspace.

$$\square c,x,w,v [Boundary(c,x) \square FromTo(c,w,v) \square \\ Distance(w,[\text{−}5,5]) < 1 \square Distance(v,[\text{5},5]) < 1]$$

The formula $Boundary(c,x)$ represents that line c is part of the boundary of object x , the formula $FromTo(c,w,v)$ represents that line c stretches from point w to point v , while the term $Distance(v,w)$ denotes the distance between points v and w . Both incompleteness and uncertainty are expressed in this formula, thanks to the use of existential quantification. Incompleteness is expressed because the formula only talks about part of the boundary of a single object. It says nothing about what other objects might or might not exist. Uncertainty is expressed because the end-points of the

boundary are only specified within a certain range. Moreover, no commitment to an absolute co-ordinate system is inherent in such a formula, so the locations of these end-points could be taken as relative to the robot’s own viewpoint.

Of course, this formula is purely illustrative, and the design of a good ontology for representing and reasoning about shape and space — perhaps one that is more qualitative and less numerical — is an open research issue. For comparison, here is an example of the kind of formula generated by the implemented abductive system for visual perception which is described in Section 3.

$$\exists x,a,s,r [Solid(x,a) \wedge Shape(x,Cuboid) \wedge FaceOf(r,x) \wedge SideOf(E_0,r)]$$

The formula $Solid(x,a)$ represents that a solid object x is presenting aspect a to the viewer, the formula $Shape(x,s)$ represents that the solid object has shape s , the formula $FaceOf(r,x)$ represents that the two-dimensional surface r is one face of object x , and the formula $SideOf(e,r)$ represents that visible edge e is on the boundary of surface r . The term E_0 denotes a particular visible edge. This is a purely qualitative representation of the simple fact that a certain kind of object is visible and “out there” somewhere.

Even this qualitative representation makes some commitment to the nature of the object in view, through the $Shape$ predicate. Yet, as Pylyshyn argues, it is sometimes necessary to provide “a direct (preconceptual, unmediated) connection between elements of a visual representation and certain elements in the world [that] allows entities to be referred to without being categorized or conceptualized” (Pylyshyn, 2001). Once again, mathematical logic facilitates this, simply by allowing an existentially quantified variable to be substituted for the term $Cuboid$.

The open-ended versatility of mathematical logic also enables the abductive account of perception to accommodate the idea that to perceive the environment is to perceive what it can afford, “what it provides or furnishes, either for good or ill” (Gibson 1979, p.127).⁵ Indeed, the

⁵ Gibson himself would hardly approve of the present theory. But this doesn’t prevent us from lifting this particular insight from his work.

representations generated by an abductively characterized perceptual process are only there as intermediate structures to help determine the most appropriate action in any given situation. Their use is only justified to the extent that they enhance rather than hinder the robot's ability to respond rapidly and intelligently to ongoing events.

Respecting this point is, once again, down to the design of the right ontology for the background theory \square . On the one hand, the design of the background theory \square might reflect the assumptions of a classical approach to machine vision, and describe the relationship between completely described 3D objects and the edges and features they give rise to in the visual field. But on the other hand, it can be designed to reflect the concerns of a more biologically-inspired approach to perception, and capture the relationship between, say, graspable containers and a small number of easily detected visual cues for such objects. In itself, the theoretical framework is entirely neutral with respect to this sort of choice.

2.2 Top-Down Information Flow

Cognitive psychology has long recognized that the flow of information between perception and cognition is bidirectional in humans (Cavanagh, 1999; Peterson, 2003), and the presence of many re-entrant neural pathways in the visual cortex reinforces this view (Ullman 1996, ch.10; Grossberg, 1999). Similarly, in the field of machine vision, although great emphasis was placed on top-down processing in the 1970s (eg: Waltz, 1975) and on bottom-up techniques in the 1980s (eg: Marr, 1982), the need for a two-way flow of information has always been acknowledged. Gestalt effects illustrate the point dramatically, as when a subject is presented with a picture that looks like a collection of random spots and blotches until it is pointed out verbally that it depicts, say, a cow, whereupon the spots and blotches resolve themselves, as if by magic, into indisputably animal form (Ullman, 1996, p.237). In such cases, where linguistically acquired knowledge causes the subject to perceive objects and boundaries where none were seen before,

it is clear that high-level cognitive processes have influenced low-level perception.⁶

In the abductive theory of perception on offer here, top-down information flow is accommodated by making use of *expectation*. As with all the proposals advanced here, the aim is not necessarily to achieve psychological or biological plausibility, but to produce a well-engineered mechanism. But in this case, as so often, there is a strong engineering rationale for imitating Nature. Presented with a cluttered, poorly lit scene full of objects occluding each other and camouflaged by their surface patterns, the use of expectations derived from past experience can reduce the task of distinguishing the important objects from the rest to computationally manageable proportions. Indeed, it makes engineering sense to permit the influence of high-level cognition to percolate all the way down to low-level image processing routines, such as edge detection.

The way expectation is exploited is as follows. Suppose a set of ϕ s has been found that could explain a given ψ . Then, in addition to entailing ψ , each ϕ will, when conjoined with the background theory Γ , entail a number of other observation sentences that might not have been present in the original sensor data ψ . These are ϕ 's expectations. Having determined these, each ϕ can be confirmed or disconfirmed by checking whether or not its expectations are fulfilled. This might involve passively processing the same low-level sensor data with a different algorithm, or it might involve a more active test to gather new sensor data, perhaps in a different modality. In Section 3, we'll see how these ideas are applied to vision.

2.3 Active Perception

The use of expectation in the way just described can be thought of as a form of *active perception* (Aloimonos, *et al.*, 1987; Ballard, 1991). The concept of active perception can be widened from its original application in the work of such pioneers as Aloimonos and Ballard to include any form of action that leads to the acquisition of new information via a robot's sensors, ranging from short duration, low-level actions, such as adjusting the

⁶ The precise nature of the interplay between cognition and perception is controversial (Pylyshyn, 1999). That such an interplay exists is less so.

threshold of an edge detection routine or rotating a camera to give a different viewpoint, to longer duration, high-level actions, such as climbing to the top of a hill to see what's on the other side. More subtle “knowledge-producing” actions (Scherl & Levesque, 1993), such as looking up a phone number, or asking another agent a question, can be thought of as lying on the same spectrum.

In general, the varieties of active perception can be classified along two axes. Along one axis, we have the kind of informative action involved — low-level or high-level, short duration or long duration — as detailed above. Along the other axis, we have the mechanism that triggers the informative action in question. In each of the above examples, the informative action could have been triggered by a reactive condition-action rule whose side-effect is the acquisition of knowledge, it could have been performed as a deliberate, reasoned effort to acquire knowledge in the face of ignorance, or it could have been an automatic attempt to confirm an expectation about the current situation. In other words, the triggering mechanism could be habit, intention, or expectation.

As well as providing a theoretical framework for handling uncertainty, incompleteness, and expectation, the basic abductive characterization of sensor data assimilation can be extended to encompass all the above forms of active perception. Each is an example of an action or series of actions with the side-effect of generating new sensor data, that is to say new \square s. For each type of active perception, the new sensor data can be interpreted via abduction to yield useful knowledge, that is to say new \square s. The purpose of action here is to control the flow of \square s. However, the role of action in perception isn't only to *maximize useful* information, but also to *minimize useless* information. The basic abductive formula makes this clear. If \square comprises a large volume of data, the problem of finding \square s to explain it will be intractable.

Part of the responsibility for managing incoming sensor data so that \square is always small but pertinent belongs to the *attention* mechanism (Niebur & Koch, 1998). In most animals, turning the head towards a loud noise and saccading to motion are instinctive behaviors with obvious survival advantages. An animal needs to gather information from those portions of

the environment most likely to contain predators, prey, or potential mates. Analogous strategies have to be found for robots, regardless of their design aims.

In general, a full abductive account of perception has to incorporate three components — attention, explanation, and expectation. The flow of information between these components is cyclical. Thanks to the attention mechanism, the robot’s sensory apparatus is directed onto the most relevant aspects of the environment. The explanation mechanism turns the resulting raw sensor data into hypotheses about the world. And in response to expectation, ways of corroborating or refuting those hypotheses are found, thereby influencing the attention mechanism.

2.4 Sensor Fusion

The abductive account of perception can also be thought of as supplying a theoretical foundation for addressing the problem of *sensor fusion*, that is to say the problem of combining potentially conflicting data from multiple sources into a single model of some aspect of the world (Clark & Yuille, 1990). The low-level sensor data in \square can come from different sensors, or from the same sensor but pre-processed in different ways, or from the same sensor but extracted at different times. If the background theory \square successfully captures how a configuration of objects in the world gives rise to each kind of low-level sensor data, then the basic abductive characterization of perception also takes care of the data fusion issue. Any \square conforming to the definition of an explanation has implicitly integrated each item of data in \square , whatever its source.

However, \square must be carefully designed to ensure this works, especially in the presence of potentially conflicting data. This is facilitated by the use of *noise* and *abnormality* terms. In addition to describing the “normal” causes of incoming sensor data, namely the robot’s interaction with objects in the world, \square can include formulae that can “explain away” unexpected items of sensor data as mere noise, and formulae that can dismiss the absence of expected sensor data as abnormal. The causal laws expressed in \square will then include noise and/or abnormality conditions, and \square ’s must be permitted to include corresponding noise and abnormality terms. It then

remains for the preference relation over \square s to prioritize explanations with the fewest such terms, thus ensuring that the explanations that best fit the data rise to the top of the pile.

Looked at in this way, the sensor fusion problem is just an instance of the general problem of finding an explanation for a body of sensor data. Indeed, with a rich sensory modality like vision, the same issue arises even with a single snapshot of a scene pre-processed by a single algorithm, such as edge detection. Both phantom and occluded edges can occur, and the abductive account of perception needs to be able to deal with them. This topic is treated in more detail in Section 3.

3. Elaborating and Applying the Abductive Account

This section describes a number of experiments, carried out over a period of several years, that were designed to evaluate the abductive account of perception using real robots. In each experiment, the basic abductive characterization is elaborated in a different way. So a series of detailed worked examples is also presented to illustrate how the theoretical framework is applied in each experimental context.

3.1 The Rug Warrior Experiment

The earliest and simplest of these experiments (Shanahan, 1996a; Shanahan, 1996b) used a low-cost two-wheeled mobile robot (called a “Rug Warrior”), whose only sensors were bump switches that detected collisions with walls and obstacles (Jones & Flynn, 1993). The robot used a reactive exploration strategy, but concurrently ran a sensor data assimilation process which, as a side-effect of the robot’s exploratory movements, built a map of the environment.⁷ Although the algorithm for sensor data assimilation was implemented as a straightforward C program, without any form of explicit logical representation, the behavior of the algorithm provably conformed to a precise abductive specification expressed in terms of abduction and predicate logic.

⁷ In this experiment, the robot’s exploratory movements are not the result of planning, and are not influenced by what the robot already knows. In the next section, a map-building experiment is described in which the robot reasons about its own ignorance to decide where to move next.

Since the only way for a sensor event to occur with this robot was as a result of its own motion, the map-building process can be viewed as a form of active perception, in the broad sense defined above. To encompass this form of active perception, the basic abductive account needs to be augmented in such a way that the background theory Σ describes the *effects of actions* and must be conjoined to a description of the *narrative* of robot actions that gave rise to the sensor data Δ . In the work cited, action and change are described using the *event calculus*, a well-known formalism for reasoning about action that incorporates a robust solution to the frame problem (Shanahan, 1997a; 1999).

To get a fuller flavor of how this sort of active perception works, let's make the example of the simple Rug Warrior robot more precise. Suppose we let,

- Δ be a conjunction of observation sentences, each describing an event in which one of the robot's bump switches is tripped,
- Σ_N be a conjunction of sentences, each describing a robot action, such as "turn left", "move forward", and so on, and,
- Σ be a background theory describing both the effect of the robot's actions on the world (specifically the effect of movements on the location of the robot itself) and the impact of the world on the robot's sensors (specifically how the presence of an obstacle causes bump events).

Then, the task of sensor data assimilation is to find a consistent formula Σ_M describing an arrangement of walls and other obstacles (a map), such that,

$$\Sigma \wedge \Sigma_N \wedge \Sigma_M \models \Delta$$

Clearly there will, in general, be many Σ_M s that conform to this prescription, and a major challenge for any abductive account of perception is to manage multiple explanations in a principled way. This topic will be dealt with in more detail towards the end of this section, in the context of vision.

3.2 Reinventing Shakey

In the Rug Warrior experiment, the gap between specification and implementation was considerable. Although the map-building algorithm conformed to a specification expressed in terms of abduction, no explicit abductive reasoning was carried out by the robot, and no explicit logical representations were built. The aim of the next set of experiments, involving miniature Khepera robots (Mondada, *et al.*, 1993), was to demonstrate that a practical robot control system could be made in which the basic unit of computation is a proof step and the basic unit of representation is a sentence of logic (Shanahan, 2000; Shanahan & Witkowski, 2001). In such a system, computation is *transparent*, in the sense that intermediate computational states have declarative meaning — each being a logical sentence in an ongoing derivation — and each step of computation can be justified by the laws of logic.

The attraction of such a foundation from an engineering point of view is clear.⁸ As in so many branches of engineering, elegance in design goes hand-in-hand with an elegant theoretical basis, and leads to artefacts that are easy to understand, to modify, and to maintain. However, one of the finest early attempts to use logic and theorem proving in the way described, namely Green's resolution-based planner (Green, 1969), fell foul of computational complexity problems. This led the Shakey project to compromise the ideal of using logic and theorem proving to control a robot (Nilsson, 1984), and no research group took up the challenge again until the cognitive robotics revival of the mid 1990s (Lespérance, *et al.*, 1994). So the question at the top of the agenda was whether a robot control system based on logic is computationally viable.

The Khepera experiments demonstrated this through the use of a logic programming meta-interpreter with *anytime* properties (Zilberstein & Russell, 1993), in other words one capable of suspending computation at any time while still yielding useful partial results (Kowalski, 1995). This capacity facilitates a design that departs from the classical AI approach

⁸ The arguments for using logic as an explicit representational medium are well known (Hayes, 1977; Moore, 1982; Nilsson, 1991). They include its possession of a mathematically precise *semantics*, its

exemplified by Shakey in a very important way. In an architecture like Shakey's, logical inference is the driving force behind the selection of actions. By contrast, in the Khepera experiments, logical inference is inserted into a closed-loop control mechanism which only allows it to carry out a small, bounded amount of computation in each cycle before the robot's sensors are consulted again. So the robot can never get stuck in a thinking rut while the world changes around it.

The miniature Khepera mobile robots used in the experiments were each equipped with two drive-wheels and a suite of infra-red proximity sensors. The robots were placed in a miniature office-like environment comprising walls, rooms, and doorways. The only low-level behaviors the robots were programmed with were wall-following and corner-turning. The first experiment centred on a navigation task (Shanahan, 2000). Given a description of the layout of rooms and doorways, and knowing its initial location, the robot had to find its way to a different specified room. But as the robot proceeded, the experimenter could unexpectedly block any of the doorways to hinder its progress.

Once again, an abductive approach to perception was adopted, although this time true abductive inference was carried out in real-time, using the sort of anytime meta-interpreter described above. The job of the abductive sensor data assimilation process was to explain the ongoing stream of sensor events in terms of what might be going on in the robot's environment. When sensor events conformed to expectations, this was a trivial task. But sometimes anomalous sensor events could occur, such as the detection of an unexpected obstacle. If the only explanation for the anomaly was a blocked doorway, a process of replanning was triggered.

Using this technique, the robot was able to navigate successfully from room to room, without perceptible computational delays, in spite of unexpected and unfavorable changes to its environment. So the experiment showed that, in principle, a robot might be able to deal with a changing world using the technology of classical AI adapted for closed-loop control, albeit in a simple environment.

putative *universality* as a representational formalism, and the *reusability* of (declarative) logical sentences in multiple contexts.

The second experiment with Khepera robots was another map-building task (Shanahan & Witkowski, 2001). Once again, the crucial difference from the earlier Rug Warrior experiment was that actual abductive inference was carried out in real-time. The abductive characterization of the map-building task was similar to that for the Rug Warrior. But for the Khepera experiment, a more sophisticated map-building strategy was deployed, in which the robot reasoned about its own ignorance, and carried out a planned sequence of exploratory actions to fill in gaps in its knowledge. To have accurate knowledge of its own ignorance, the robot had to be able to recognize novel locations and to re-identify locations it had previously visited, which is also an important aspect of the “symbol anchoring problem” (Coradeschi & Saffiotti, 2000).

3.3 A Simple Worked Example of Active Perception as Abduction

Based on the navigation task in the Khepera experiments, this section presents a worked example of how a simple form of active perception can be formalized as an abductive process. We begin with a short introduction to the event calculus, the formalism for reasoning about action deployed in all the work reported in this paper. The version of the event calculus used here has been widely adopted, and has become absorbed into the Artificial Intelligence canon (Russell & Norvig, 2002, ch. 10). In addition to robotics, it has found application in such areas as natural language processing (Lévy & Quantz, 1998), intelligent agents (Jung, 1998), and common sense reasoning (Morgenstern, 2001; Shanahan, 2004; Mueller, 2004). A more detailed presentation can be found in (Shanahan, 1999).

The ontology of the event calculus includes events (or actions), time points, and fluents. A fluent can be anything whose value changes over time, such as the location of a robot. The basic language of the calculus is presented in Table 1.

Formula	Meaning
<i>Initiates</i> (e,f,t)	Fluent f starts to hold after action e occurs at time t
<i>Terminates</i> (e,f,t)	Fluent f ceases to hold after action e occurs at time t
<i>Initially</i> (f)	Fluent f holds from time 0
<i>Happens</i> (e,t)	Action or event e occurs at time t

$ HoldsAt(f,t) $	Fluent f holds at time t
$ Clipped(t_1,f,t_2) $	Fluent f is terminated between times t_1 and t_2

Table 1: The Language of the Event Calculus

The relationship between these predicates is governed by the following three axioms, whose conjunction will be denoted EC.

$$HoldsAt(f,t) \sqcap Initially(f) \sqcap \sqcap Clipped(0,f,t) \tag{EC1}$$

$$HoldsAt(f,t_2) \sqcap \tag{EC2}$$

$$Happens(e,t_1) \sqcap Initiates(e,f,t_1) \sqcap t_1 < t_2 \sqcap \sqcap Clipped(t_1,f,t_2)$$

$$Clipped(t_1,f,t_3) \sqcap \tag{EC3}$$

$$\sqcap e,t_2 [Happens(e,t_2) \sqcap t_1 < t_2 < t_3 \sqcap Terminates(e,f,t_2)]$$

The frame problem is overcome using circumscription (McCarthy, 1986; Lifschitz, 1994). The predicates *Happens*, *Initiates*, and *Terminates* are minimized, corresponding to the default assumption that the known event occurrences are the only event occurrences, and the known effects of actions are the only effects of actions. By circumscribing separate parts of a theory independently, the difficulties that beset early attempts to apply circumscription to the frame problem are avoided (Shanahan, 1997a; 2003). In what follows, given a formula \square and a set P of predicates, the circumscription of \square minimising P will be denoted $CIRC[\square ; P]$.

In practise, the deductive application of these axioms performs prediction, that is to say reasoning from causes to effects. Given a narrative of events and actions in the form of a set \square of *Initially* and *Happens* formulae, and a description of the effects of actions in the form of a set \square of *Initiates* and *Terminates* formulae, the consequences of,

$$CIRC[\square ; Initiates, Terminates] \sqcap CIRC[\square ; Happens] \sqcap EC$$

will include a set of *HoldsAt* formulae describing which fluents hold when, in other words capturing the effects of the events and actions in \square .

However, the concern of the present paper is not with prediction, but with explanation, that is to say reasoning from effects to causes, where the effects are the raw sensor data and the causes are features of the world external to the robot. In event calculus terms, given,

- a set Σ of *HoldsAt* formulae describing sensor data,
- a set Σ_N of *Initially* and *Happens* formulae describing the robot's initial state and subsequent actions, and
- a description of the effects of actions in the form of a set Σ of *Initiates* and *Terminates* formulae,

the basic abductive task is to find a set Σ_M of formulae describing features of the world such that,

$$\Sigma_M \sqcup \text{CIRC}[\Sigma; \textit{Initiates}, \textit{Terminates}] \sqcup \text{CIRC}[\Sigma_N; \textit{Happens}] \sqcup \text{EC} \models \Sigma.$$

Many variations on this basic formulation are possible. For instance, in the example to follow, Σ includes *Happens* formulae describing sensor events.

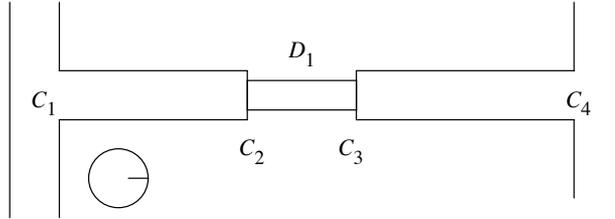


Figure 1: A Simple Mobile Robot Scenario

Now, consider Figure 1. Suppose the robot possesses a topological map and knows the layout of rooms and doorways in its world, but it has no way to keep track of whether doors are open or shut. Since it also has no knowledge of the distances between the various corners, as it proceeds along the wall, the robot cannot possibly know that D_1 is in fact shut until it encounters corner C_4 .

Let the robot's knowledge of the corners and doorways be captured by the following set of formulae, whose conjunction will be denoted Σ_{M1} . The predicate names should be self-explanatory.

$$\begin{array}{ll} \textit{NextCorner}(C_1, C_2) & \textit{NextCorner}(C_2, C_3) \\ \textit{NextCorner}(C_3, C_4) & \textit{Inner}(C_1) \\ \textit{Outer}(C_2) & \textit{Outer}(C_3) \\ \textit{Inner}(C_4) & \textit{Door}(D_1, C_2, C_3) \end{array}$$

Next, we need a set of formulae describing the effects of the robot's actions. The only action we will consider here is *FollowWall*, which takes the robot to the next detectable corner. In the present example, that corner will be C_2 or C_4 depending on whether door D_1 is open or shut.

$$\begin{aligned}
& \text{Initiates}(\text{FollowWall}, \text{AtCorner}(c_2), t) \sqcap \\
& \quad \text{HoldsAt}(\text{AtCorner}(c_1), t) \sqcap \text{NextVisibleCorner}(c_1, c_2) \\
& \text{NextVisibleCorner}(c_1, c_2) \sqcap \\
& \quad \text{NextCorner}(c_1, c_2) \sqcap \sqcap \text{InvisibleCorner}(c_2) \\
& \text{NextVisibleCorner}(c_1, c_3) \sqcap \\
& \quad \text{NextCorner}(c_1, c_2) \sqcap \text{InvisibleCorner}(c_2) \sqcap \text{NextVisibleCorner}(c_2, c_3) \\
& \text{InvisibleCorner}(c_1) \sqcap \sqcap d, c_2 [\text{Door}(d, c_1, c_2) \sqcap \text{DoorShut}(d)] \\
& \text{DoorOpen}(d) \sqcap \sqcap \text{DoorShut}(d)
\end{aligned}$$

Let \sqcap be the conjunction of the above formulae. Now we require a set of formulae that captures the causes of robot sensor events. In the present example, the only sensor events we'll consider are *GoesHigh(Front)*, which occurs when the robot's front infra-red proximity sensors detect an obstacle ahead, and *GoesLow(Left)*, which occurs when the robot's sensors encounter a gap to its left.

$$\begin{aligned}
& \text{Happens}(\text{GoesHigh}(\text{Front}), t) \sqcap \\
& \quad \text{Happens}(\text{FollowWall}, t) \sqcap \\
& \quad \text{Initiates}(\text{FollowWall}, \text{AtCorner}(c), t) \sqcap \text{Inner}(c) \\
& \text{Happens}(\text{GoesLow}(\text{Left}), t) \sqcap \\
& \quad \text{Happens}(\text{FollowWall}, t) \sqcap \\
& \quad \text{Initiates}(\text{FollowWall}, \text{AtCorner}(c), t) \sqcap \text{Outer}(c)
\end{aligned}$$

Let \sqcap_N be the conjunction of the above *Happens* formulae with the following narrative formulae, which describe the relevant episode in the robot's history.

$$\begin{aligned}
& \text{Initially}(\text{AtCorner}(C_1)) \\
& \text{Happens}(\text{FollowWall}, T_1)
\end{aligned}$$

The sensor data to be explained comprises a single event.

$Happens(GoesHigh(Front), T_2)$

$T_2 > T_1$

Now, if we let \square be the conjunction of the above two formulae, the abductive task is to find some \square_{M2} comprising only *DoorOpen* and *DoorShut* formulae, such that,

$\square_M \square CIRC[\square ; Initiates, Terminates] \square CIRC[\square_N ; Happens] \square EC \models \square$

where \square_M is $\square_{M1} \square \square_{M2}$. In this case, the obvious solution is simply to let \square_{M2} be the formula *DoorShut(d)*. The inclusion of this formula on the left-hand-side of the entailment ensures that the next detectable corner after C_1 is C_4 , which in turn explains the occurrence of a *GoesHigh(Front)* event after the robot executes its *FollowWall* action.

This example shows the robot actively probing its environment to discover more about it, and it demonstrates the need for inference to interpret the resulting inflow of sensor data. It's just about the simplest such example that can be presented in a few pages, but it illustrates a basic theoretical apparatus that can be used to couch much more complex forms of perception in abductive terms. For example, by making the whole of \square_M abducible, the general problem of map-building is reduced to a kind of active perception characterized in abductive terms. Indeed, this is precisely how map-building was treated in the Khepera experiments described above, and presented in detail in (Shanahan & Witkowski, 2001). A more substantial extension of the basic framework is required to deal with the multiple explanations problem, which will be addressed in the next worked example.

3.4 Visual Perception Through Abduction

Although the Rug Warrior and Khepera experiments served as a useful “proof of concept” for an updated attempt to achieve robot control through logic, they left open the question of whether the ideas under test would scale up to robots with richer sensors in more complex environments. To address the issue of scaling up, a much more ambitious project has been launched,

involving a bench-mounted, upper-torso humanoid robot with stereoscopic vision, known as Ludwig (Figure 2).

Ludwig has two arms, articulated at the shoulder and elbow, with three degrees-of-freedom each. The forearm is mechanically constrained to move in the same plane as the work-bench. The arms terminate with simple prods, rather than grippers. The stereoscopic camera is mounted on a pan-and-tilt head. Ludwig is mechanically very simple, having fewer degrees-of-freedom than other recently constructed humanoid robots, such as Cog at MIT (Brooks, *et al.*, 1999). This cuts down on low-level motor control problems, and allows the project to concentrate on the central issues of cognition and perception.

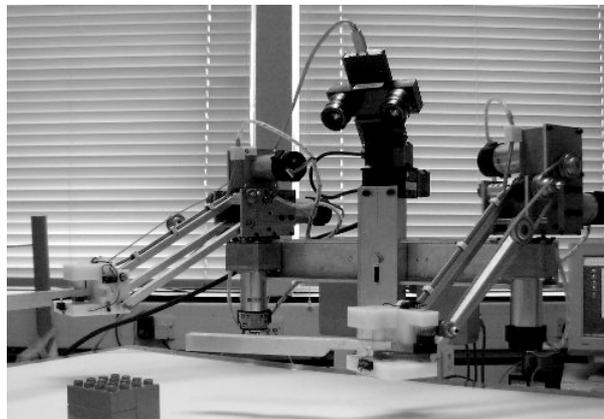


Figure 2: Ludwig, an Upper-Torso Humanoid

Like a human infant, Ludwig's main occupation is play. The specific challenge for Ludwig is to learn about the objects in its workspace through physical interaction with them. The primary function of the visual system is to pick out potential objects of interest — those that comprise Ludwig's Umwelt — and to direct the camera's gaze on them. Ludwig can then use visual servoing in an attempt to nudge them. The resulting changes in orientation and relative location of these and other objects on the work-bench lead to new and useful information about their shapes, sizes, and layout.

The Ludwig project embraces both classical AI technology and ideas from biologically-inspired robotics. On the one hand, the order in which research topics are tackled isn't top-down, as in classical AI, but bottom-up,

as in biologically-inspired robotics, mirroring ontogeny and phylogeny. Accordingly, rather than viewing human-level intelligence as a disembodied phenomenon and coming at it through language and problem solving, the project considers human-level intelligence to rest on bodily skills such as nudging and grasping, and to be grounded in (active) perception.

On the other hand, Ludwig's play has an end-product that will be familiar to anyone with a classical AI background, namely declarative representations with a denotational semantics — in particular, sentences of logic. (These should not, of course, be confused with sentences of natural language. It is not anticipated that Ludwig will be equipped with a capacity for natural language in the near future.) And as in the Rug Warrior and Khepera experiments, these representations are constructed by an abductive process, which is described below.

To meet the demands of vision — an altogether richer source of sensor data than bump switches or infra-red proximity sensors — a much more sophisticated abductive account of perception has been developed (Shanahan, 2002). At the core of this treatment of visual perception are two advances on the basic abductive prescription. First, a precise measure of the *explanatory value* of a hypothesis is defined, based on probability. Crudely put, the key property of this measure is that the more data items a hypothesis explains, the greater its explanatory value. By calculating their explanatory values, competing hypotheses can be compared and unpromising ones discarded.

The precise definition of explanatory value is as follows. Let $\square_1 \dots \square_n$ be the set of all hypotheses that explain the sensor data \square according to the usual abductive definition. Given that one and only one of these hypotheses can be the true explanation, and if none of the hypotheses is subsumed by any other, we can assume $\square_1 \oplus \dots \oplus \square_n$.⁹ Now, consider any hypothesis \square_k , and let R be,

$$\square_1 \oplus \dots \oplus \square_{k-1} \oplus \square_{k+1} \oplus \dots \oplus \square_n.$$

Then, from the laws of probability, we have,

⁹ In other words, one and only one hypothesis is true. Strictly speaking, this is an abuse of notation, since the \oplus (exclusive or) operator does not compose in the intended way.

$$P(\square_k | \square_k \oplus R) = \frac{P(\square_k)}{P(\square_k) + P(R)} \quad (1)$$

where $P(\square_k)$ is the prior probability of \square_k and $P(R)$ is the prior probability of R . From the form of R , we have,

$$P(R) = \prod_{i=1}^n P(\square_i) \prod_{k=1}^m P(\square_k). \quad (2)$$

For any hypothesis \square of the form $\square_1 \square \cdots \square_m$, we have,

$$P(\square) = \prod_{j=1}^m P(\square_j) \quad (3)$$

From equations (1) to (3), the posterior probability of any hypothesis can be calculated, given the set of all hypotheses, and this is taken to be its explanatory value. Now, recall that our measure of explanatory value is supposed to reward a hypothesis for explaining more data. Despite the fact that the abductive definition insists that \square is explained in its entirety, this criterion can be met by allowing a hypothesis to include *noise terms* (Poole, 1995; Shanahan, 1997b). A noise term “explains away” an item of sensor data as mere noise. If noise terms are assigned a low prior probability, equation (3) ensures that a hypothesis that properly explains an item of sensor data has a higher explanatory value than one that simply explains it away as noise.

The second advance on the basic abductive story allows for the performance of actions to confirm or disconfirm the *expectations* of competing hypotheses, and incorporates the results into an updated measure of explanatory value. A key property of the updated measure is that, other things being equal, the fewer unfulfilled expectations a hypothesis has, the greater its explanatory value.

Overall, the abductive processing of a set of low-level sensor data consists of the the following four steps.

1. Using abduction, establish a set of hypotheses that can explain the sensor data.

2. Select the best of these hypotheses according to their explanatory values.
3. Determine the expectations of each selected hypothesis, and re-consult the raw data to check each of these expectations.
4. Recalculate the explanatory values of the selected hypotheses taking into account the results of step 3, and re-order the hypotheses accordingly.

To see how these ideas work in the context of vision, consider Figure 3. The central image shows the results of applying a simple edge detection algorithm, based on the Sobel operator, to the left-hand image. The threshold of the edge detection algorithm is set high to reduce the volume of spurious data. Now consider the block circled in the upper image. Because of its dark half, only some of this block's edges have been picked out by the algorithm. This leads to two competing initial hypotheses to explain the edge data. According to Hypothesis A, which is incorrect, Ludwig is looking at a short block, the faint edges to the left being noise. According to Hypothesis B, which is correct, Ludwig is looking at a long block, but lots of its edges are missing.



Figure 3: An Ambiguous Block

These two hypotheses have different expectations. For example, a consequence of Hypothesis A would be the presence of a vertical edge from the top left-hand corner of the proposed block. A consequence of Hypothesis B is that the topmost edge of the block should extend further to the left than it apparently does. Using an edge checking algorithm that is much more sensitive than the algorithm originally applied to the whole scene, these expectations can be tested. In the present example, this leads to a reduction in the explanatory value of Hypothesis A, whose expectations

are not fulfilled, and an increase in the explanatory value of Hypothesis B, whose expectations are fulfilled.

3.5 A Worked Example of Top-Down Information Flow

More formally, for the purposes of abduction, low-level image data such as that in Figure 3 will be represented by a set of predicate calculus formulae. The low-level data in question could be the product of any off-the-shelf image processing technique, such as color-based segmentation, optical flow, stereo matching, or any combination of these. But for the worked example of the present section we'll stick to edge detection. The edge data in Figure 3 can be represented by a set of formulae that includes the following exemplars.

$Line(1,[238,157],[241,147])$

$Line(3,[249,157],[253,145])$

The formula $Line(w,p_1,p_2)$ represents that there is an edge in the image, labelled w , from point p_1 to point p_2 . (The labels are arbitrarily assigned by the low-level edge detection procedure.)

An abductive process can be structured into multiple layers (Josephson & Josephson, 1994; Shanahan, 2002). In effect, each explanation \square at the n^{th} layer becomes data \square to be explained at the $n+1^{\text{th}}$ layer. Accordingly, formulae that are abducible at layer n can become observable at layer $n+1$. In the present context, it's convenient to adopt a two-layer abductive process. The task of the first layer is to explain one-dimensional edges in terms of two-dimensional visible regions. The task of the second layer is to explain the visible regions in terms of spatially-located three-dimensional objects.

Each layer has a separate background theory \square , which includes three groups of logical formulae. Although every formula enjoys the same status from a declarative point-of-view, each of group plays a different role in the abductive process, and this will be reflected in any implementation. First, we have a set of formulae that can be used to *generate* explanatory hypotheses. Second, we have a set of formulae describing *constraints* on the

permissible hypotheses. Third, we have a set of formulae describing the *expectations* of each hypothesis.

Here are two sample first-layer formula belonging to the first of these three groups. Their style is typical. Each has the form $\square \square \square$, where \square is a description of sensor data and \square is a description of conditions under which that sensor data arises.

$$\square p_1, p_2 [Line(w, p_1, p_2)] \square \square r [Region(r) \square SideOf(w, r)]$$

$$\square p_1, p_2 [Line(w, p_1, p_2)] \square Noise(w)$$

In essence, these two formulae state that a visible edge might be the side of a region or it might simply be noise. The three predicates *Region*, *SideOf*, and *Noise* (whose meanings should be obvious) are declared to be abducible, so hypotheses can be expressed in their terms. Next, we have a sample first-layer formula from the second group.

$$[SideOf(w_1, r) \square SideOf(w_2, r) \square w_1 \neq w_2] \square \\ [Parallel(w_1, w_2) \quad Joins(w_1, w_2)]$$

Again, the style of the formula is typical. It has the form $\square \square \square$, where \square describes conditions that must be met by any hypothesis conforming to \square . In this case, the formula presents a constraint that must be met by a hypothesis describing trapezoidal regions, on the (artificial) assumption that the robot's *Umwelt* comprises cuboids whose visual aspect consists of such regions.

We'll postpone consideration of the third type of formula – those describing the expectations of hypotheses – until the second layer of the abductive process. In the mean time, it can be seen that the first-layer explanations of the low-level image data, according to the formulae above, include both the hypothesis that there is a region whose sides include edges 1 and 3 (where edge 7 is noise), and the competing hypothesis that there is a region whose sides include edges 7 and 3 (where edge 1 is noise). Let's call these, respectively, Hypothesis A* and Hypothesis B*. Hypothesis A* includes the following formulae.

$$Region(R_0) \qquad \qquad \qquad SideOf(1, R_0)$$

$SideOf(3,R_0)$

$Noise(7)$

Hypothesis B* includes the following formulae.

$Region(R_0)$

$SideOf(7,R_0)$

$SideOf(3,R_0)$

$Noise(1)$

According to Equations (1) to (3), these two hypotheses have equal explanatory value, since both postulate a single new region, and both contain the same number of noise terms. Note, however, that hypotheses containing more noise terms will have lower explanatory value. In other words, the best hypotheses are those that postulate only a few regions, each of which explains the presence of many edges.

Now let's consider the second layer of abduction, which moves from visible regions to solid objects. The second-layer formulae for generating hypotheses include the following.

$Region(r) \sqcap \sqcap x,a [Solid(x,a) \sqcap IsAspect(a) \sqcap FaceOf(r,x)]$

$IsAspect(3Face)$

The predicate $Solid(x,a)$ represents that a solid object x is in the visual field, presenting aspect a . The predicate $FaceOf(r,x)$ represents that the visible region r is a facet of object x that is turned towards the viewer. An *aspect* of an object is an equivalence class of possible views of that object that are invariant with respect to certain topological properties, such as the number of its corners or faces (Figure 4). A cube, for example, has three aspects. The one depicted in Figure 4 is denoted $3Face$.

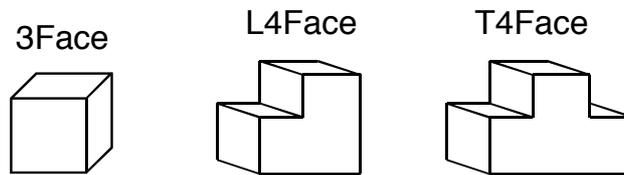


Figure 4: Aspects of a Cube, an L-Block, and a T-Block

The following formula is an example of a second-layer constraint, which states that every pair of faces in the aspect $3Face$ is abutting, and rules out any hypothesis for which this is not the case.

$[Solid(x,3Face) \sqcap FaceOf(r_1,x) \sqcap FaceOf(r_2,x) \sqcap r_1 \neq r_2] \sqcap Abuts(r_1,r_2)$

Now let's consider two competing second-layer hypotheses. Suppose Hypothesis A contains all of Hypothesis A*, and also includes the following formulae.

$$\begin{array}{ll}
 \textit{Region}(R_1) & \textit{SideOf}(4,R_1) \\
 \textit{SideOf}(6,R_1) & \textit{Solid}(X_0,3\textit{Face}) \\
 \textit{FaceOf}(R_0,X_0) & \textit{FaceOf}(R_1,X_0)
 \end{array}$$

Similarly, suppose Hypothesis B contains all of Hypothesis B*, and also includes the following formulae.

$$\begin{array}{ll}
 \textit{Region}(R_1) & \textit{SideOf}(4,R_1) \\
 \textit{Noise}(6) & \textit{Solid}(X_0,3\textit{Face}) \\
 \textit{FaceOf}(R_0,X_0) & \textit{FaceOf}(R_1,X_0)
 \end{array}$$

Hypothesis B is, of course, the correct one, as it treats edge 6 as noise rather than as the boundary of a solid object. However, Hypothesis A has a higher explanatory value according to Equations (1) to (3) because it contains more noise terms.

This brings us, finally, to the topic of expectation and top-down information flow. Hypothesis B is under-performing because of the poor quality of the original low-level data. In particular, the faintly visible face on the left of the block has been ignored by the edge detection algorithm. But Hypothesis B can be promoted to its proper place by actively seeking to confirm those of its expectations that are not already fulfilled by the original data. Here is an example of a second-layer "expectation" formulae.

$$\begin{aligned}
 & [\textit{Solid}(x,3\textit{Face}) \supset \textit{FaceOf}(r_1,x)] \supset \\
 & \supset r_2, r_3 [\textit{FaceOf}(r_2,x) \supset \textit{Abuts}(r_2,r_1)] \supset \\
 & \textit{FaceOf}(r_3,x) \supset r_2 \neq r_3 \supset \textit{Abuts}(r_3,r_1)
 \end{aligned}$$

This type of formula is superficially similar to a constraint. The difference is that the right-hand-side of the implication describes conditions which, if not already met, can be tested by referring back to the raw image data and re-processing it in a specialized way. In this case, if the posited cuboid is missing any of its expected faces, this formula will trigger a search

for a region having the right geometrical properties, namely that it abuts the cuboid's known faces. Because these properties are highly specific, the search can be very sensitive. In the example of Figure 3, this will enable the detection of the previously missed facet of the block.

The final step of the reasoning process is to incorporate the results of checking the expectations of each hypothesis into a newly calculated measure of explanatory value. Again, Equations (1) to (3) are used. The updated measure is obtained simply by widening the set of low-level sensor data subject to explanation to encompass both the presence *and* absence of expectations. To allow for the explanation of the absence of an expected item of low-level data, *abnormality terms* are introduced. These are analogous to noise terms – the more of them a hypothesis is forced to contain, the less its explanatory value. This results in an increase in the explanatory value of a hypothesis whose expectations are fulfilled, and a corresponding decrease in the explanatory value of a hypothesis with unfulfilled expectations.

In the present example, each hypothesis entails the expectation of a missing third region that abuts both R_0 and R_1 . This is a complex expectation involving several edges, which are actively sought out using a sensitive spatially guided edge-checking algorithm. For Hypothesis A this expectation is not met, while for Hypothesis B it is. Now both the presence of new edges on the left of the image and the absence of looked-for edges in the centre become additions to the low-level sensor data, and this results in Hypothesis B overtaking Hypothesis A in explanatory value.

A prototype abductive inference mechanism has been constructed that incorporates these ideas of explanatory value and top-down expectation. As with the Khepera experiments, the implementation is in the form of a logic programming meta-interpreter. The system has been applied to visual data, and has been shown to work for examples similar to that in Figure 3. The next step is to interface the system directly to Ludwig's camera, and to let the best hypotheses it generates lead the arm to nudge interesting objects on the work-bench. A suitable attention mechanism is also required to guide the camera and to select portions of the visual field for processing. The present solution is for the system to seek out motion. The experimenter is

able to draw attention to interesting objects by pushing them or waving them about, much as a parent does with a young infant.

3.6 Active Visual Perception

The preceding examples have shown how the abductive treatment can be given to a simple sensory modality in a dynamic setting (Section 3.3) and to a richer sensory modality, namely vision, in a static setting (Section 3.5). To complete the picture, this section gives an abductive account of vision in a dynamic setting (Shanahan & Randell, 2004). This is largely a matter of bolting the required logical apparatus for handling actions and change (from Section 3.3) to the abductive machinery for dealing with top-down information flow (from Section 3.5).

Consider Figure 5, which shows a two images from Ludwig’s camera alongside the results of applying edge detection. Between the two frames, the block has been nudged slightly. In the first frame, the bright highlight on the right of the block has resulted in the loss of detectable edges. From the edge information in this frame alone, there is no way to determine whether the object is an L-block or a T-block (Figure 4). However, the ambiguity is resolvable with the aid of the extra information obtained by nudging the block and thereby gaining a view of it from a different angle. The challenge now is to formalize this process of knowledge increase in abductive terms.

To meet this challenge, a third layer of abductive interpretation will be added in addition of the two layers described in the previous section. This layer has the responsibility of explaining sequences of momentary snapshots in terms of persistent solid objects. The first step that needs to be taken is to turn the predicates previously used to describe static images into fluents. This is achieved through “reification”, which amounts to wrapping them in the *HoldsAt* predicate. So instead of writing *Solid(x,a)* we now write *HoldsAt(Solid(x,a),t)* where *t* is a time point.

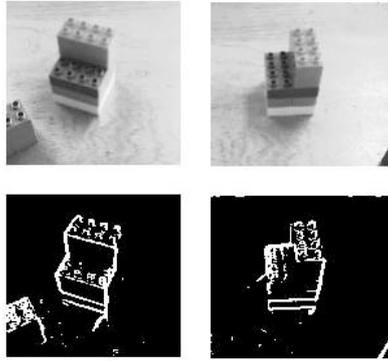


Figure 5: Nudging a Block

Now suppose that two competing second-layer hypotheses have been found to explain the edge data in the first frame of Figure 5, taken at time T_1 . According to Hypothesis A, the object is presenting an aspect of an L-block.

$$\text{HoldsAt}(\text{Solid}(X_0, L4\text{Face}), T_1)$$

According to Hypothesis B, the object is presenting an aspect of a T-block.

$$\text{HoldsAt}(\text{Solid}(X_0, T4\text{Face}), T_1)$$

Top-down information has helped to narrow down the candidate explanations to these two. But this is the best it can manage, and because the edges on the right of the block are completely obliterated, both hypotheses are accorded equal explanatory value.

The arrival of the next frame, taken at time T_2 , is an opportunity to resolve the ambiguity. However, thanks to the block's surface features, edge detection cannot distinguish a clear boundary at the top of the block. So the application of static abductive interpretation to this frame would again reveal two equally good competing hypotheses. According to one hypothesis, the block is presenting an aspect of a cuboid.

$$\text{HoldsAt}(\text{Solid}(X_0, 3\text{Face}), T_2)$$

According to the other hypothesis, the object is presenting a different aspect of an L-block.

$$\text{HoldsAt}(\text{Solid}(X_0, L2\text{Face}), T_2)$$

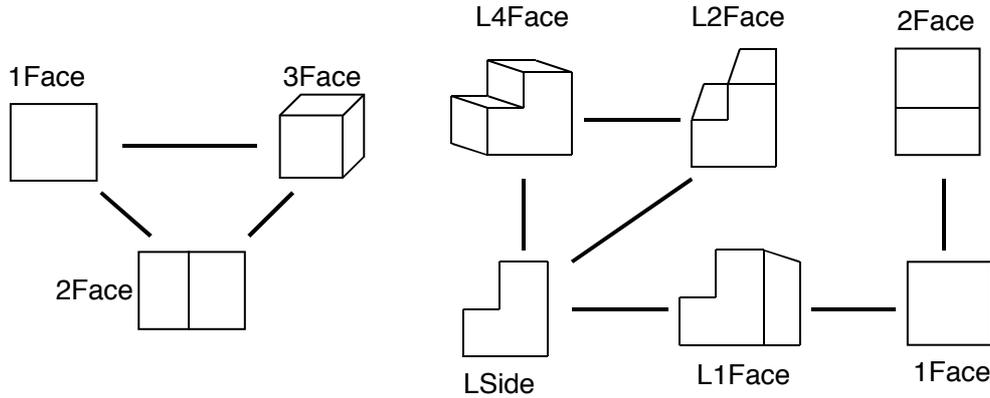


Figure 6: Partial Aspect Graphs of a Cube and an L-Block

But because the two frames are linked, a static analysis of the second one would be inappropriate. Instead, the abductive process must attempt to extend each of the ongoing hypotheses (in this case A and B) to accommodate the data in the new frame. To realise this, we can import a useful and well-established concept from computer vision, namely that of an *aspect graph* (Koenderink & van Doorn, 1979). The aspect graph of a three-dimensional shape is the set of all possible aspects it can present to the viewer linked by arcs to indicate legitimate transitions from one aspect directly to another (Figure 6).

In what follows, the formula $Shape(x,s)$ represents that object x has shape s , where a shape is defined by its aspect graph. In the third layer of abduction, $Shape$ formulae are made abducible. The formula $Arc(s,a_1,a_2)$ represents that there is an arc from aspect a_1 to aspect a_2 in the aspect graph of shape s . So the set of formulae required to represent the shapes illustrated in Figure 6 include the following.

$$Arc(LBlock,L4Face,L2Face)$$

$$Arc(Cuboid,3Face,2Face)$$

Now we can use the event calculus to describe the possible aspect transitions that can result from nudging an object. The results will be incorporated in the third-layer background theory \square . Let the term $Nudge(x)$ denote an act of nudging object x , and assume that no single nudge is dramatic enough to cause an object to pass through more than one aspect transition. We need to formalize the fact that a *Nudge* action might result in

no aspect transition, but it might result in the transition from the current aspect a_1 to any new aspect a_2 such that an arc exists from a_1 to a_2 in the relevant aspect graph.

$$\begin{aligned} & \text{Initiates}(\text{Nudge}(x), \text{Solid}(x, a_2), t) \sqcap \\ & \text{HoldsAt}(\text{Solid}(x, a_1), t) \sqcap \text{Shape}(x, s) \sqcap \\ & \text{Arc}(s, a_1, a_2) \sqcap \text{ByChance}(\text{TransTo}(x, a_2), t) \end{aligned}$$

The *Nudge* action is an example of an action with a non-deterministic effect, and this formula illustrates one way the event calculus can be used to represent such actions (Shanahan, 1997a). The formula *ByChance*(f, t) represents that fluent f holds at time t . Unlike *HoldsAt*, there are no axioms to constrain the *ByChance* predicate, so the fluent f , which is known as a *determining fluent*, can either hold or not hold at any time. In the context of abduction, *ByChance* formulae are abducible, but make no contribution to the calculation of explanatory value. For more details, see (Shanahan, 1997b) or (Shanahan & Randell, 2004).

The *Nudge* action that was carried out between the two frames is represented by a *Happens* formula.

$$\begin{aligned} & \text{Happens}(\text{Nudge}(X_0), T) \\ & T_1 < T < T_2 \end{aligned}$$

Because we have *Arc*(*LBlock*, *L4Face*, *L2Face*), Hypothesis A can now be extended to form a third-layer hypothesis that explains the two successive frames in terms of a persistent solid object. The extended hypothesis will include the following formula.

$$\text{Shape}(X_0, \text{LBlock})$$

By contrast, Hypothesis B cannot be extended to accommodate the data in the new frame. This is because there is no second-layer hypothesis that includes the formula *HoldsAt*(*Solid*(X_0, a), T_2) such that *Arc*($s, T4Face, a$), for any shape s . This is, of course, the desired outcome.

4. Related Work

The idea that perception can be viewed as a form of abduction can be traced back to the philosophical writings of C. S. Peirce (1903), who carried out an extensive study of abductive reasoning. (Indeed, the term “abduction” was coined by Peirce.) The essential insight is clearly present in (Peirce 1903), although his discussion is somewhat informal. Within psychology, Rock's theory of *indirect perception*, which in turn derives from the ideas of the nineteenth-century scientist Hermann von Helmholtz, is the most closely related to the abductive approach of this paper (Rock, 1983). Although Rock doesn't use the term “abduction”, that is precisely what he describes, as his theory of perception involves the generation of hypotheses about the external world to explain visual data and their testing through the acquisition of further data. Rock's theory is informally presented and unimplemented, but has the virtue of empirical support.

Returning to AI and computer science, the treatment of visual perception offered here has some points of contact with work on model-based vision dating back to the early 1980s (Brooks, 1981), and aspects of the abductive processing of image data are analogous to the matching operation at the heart of contemporary model-based object recognition systems (Jain, *et al.*, 1995, ch. 15). But, in general, these systems neither use logic as a representational formalism, nor carry out any form of logical inference. Moreover, the role of high-level knowledge in model-based vision is usually confined to the representation of the structure of the objects to be recognized.

The combination of abductive inference with the confirmation of expectations, here applied to visual perception, echoes the hypothetico-deductive model of scientific reasoning proposed by philosophers of science. In (Kakas & Evans 1992) we find a formal description of hypothetico-deductive reasoning in a logic programming context, although the authors don't apply their ideas to perception. The hypothetico-deductive model is applied to perception in (Josephson & Josephson 1994, ch.10), although the treatment there is informal, and the authors offer few hints about how their ideas might be realized on a computer.

A rare example of an attempt to apply logic directly to image interpretation can be found in (Reiter & Mackworth, 1989), where the authors provide a formal definition of the relationship between an image and the scene it depicts. Their definition, like its abductive counterpart in the present paper, relies on an axiomatic description of the way different objects in the environment give rise to raw image data. Reiter and Mackworth's treatment was purely theoretical, and their ideas were never implemented in a real vision system. Moreover, they largely neglected the issues of noise and incompleteness, which are a central concern here. However, their work is a clear precursor to the framework of this paper. To this author's knowledge, the only other contemporary attempt to bring logic to bear on the problem of perception is the ongoing work of Pirri and her colleagues (Pirri & Finzi, 1999; Pirri & Romano, 2002).¹⁰

Another example of a methodologically related approach to perception is the recent work of Chella, Frixione and Gaglio (Chella, *et al.*, 1997; Chella, *et al.*, 2000), who have developed an architecture for vision with many of the same features as the approach to perception advocated here. Their architecture employs off-the-shelf techniques for low-level vision, but uses declaratively represented knowledge and high-level reasoning to suggest hypotheses that might explain the low-level data, as well as to form the expectations of those hypotheses. These expectations influence an attention mechanism that can seek out components of a specified shape in designated parts of an image.

Although the efforts of Chella, Frixione and Gaglio perhaps bear the closest resemblance in the literature to the present work as far as perception is concerned, there are still many differences. In particular, logic and logical inference play a much smaller part in their approach than in the framework of this paper, and instead they make liberal use of neural networks. Their work lacks a theoretical counterpart to the abductive treatment of perception that forms the centrepiece of the research reported here. On the other hand,

¹⁰ Van der Does and Van Lambelgen (2000) present a "logic of vision" which draws on Marr's ideas. Although their main concern is to supply a semantics of perceptual reports rather than to formalize the perceptual process itself, it would be interesting to compare the relevant aspects of their logic to the present approach.

they use a wider range of techniques, from both AI and machine vision, and have a more sophisticated approach to shape.

5. Philosophical Afterthoughts

The motivation for developing the theory of robot perception presented in this paper is the belief that perception, action, and cognition form an inseparable trinity, and that success in AI will not result from studying any one of those topics in isolation. More specifically, the aim is to supply a solid theoretical foundation for the design of robot perceptual systems that incorporate action, in the form of active perception, and cognition, in the form of top-down information flow. As such, the context is engineering, not philosophy.

Yet it is engineering with a philosophical edge. At every turn, concepts are deployed that derive from centuries of thought in the philosophy of mind. Time and again, aspects of the abductive theory of perception echo well-known philosophical debates and conundrums. So it seems appropriate to make some philosophical remarks. But rather than confronting the philosophy head-on, what follows are brief caricatures of a number of issues set alongside the reflections of an imaginary engineer. To bring the issues into proper focus, it's necessary to envisage a more advanced humanoid robot than we have at present, albeit built using the principles set out in this paper, so the ensuing discussion has a science fiction flavor.

5.1 Representational Theories of Mind

The engineer in question, being a keen amateur reader of philosophy journals, has for years been interested in the debate surrounding “intentionality” (Searle, 1983; Crane, 2003). The philosophers, she has learned, are very struck by the fact that mental states are *directed* towards things — beliefs are always *about* something, as are other so-called *propositional attitudes* such as desires, intentions, and so on. In the 19th Century, Brentano’s challenge to the materialists was to find a theory of mind that could account for this property of intentionality. Not only would it have to account for the fact that propositional attitudes are directed at real, physical objects, it would also have to accommodate their capacity for

directedness at *non-existent* objects, such as unicorns. Moreover, beliefs can be false. So the same account must allow for *error*.

Many of the philosophers trying to meet Brentano's challenge adhere to some form of *representational theory of mind*, in which mental states are characterized by appeal to sentence-like representations (Fodor, 1975). Much is made of the *functional role* of these representations in the overall organization of the mind. In some variants of the theory, mental processes are identified with kinds of *computation* over the sentence-like representations.

The engineer also knows a fair amount of neuroscience. So she has noted the considerable gap between the philosophers' notion of representation and the messy grey stuff that is actually found inside people's heads. She is also aware that certain brands of philosopher have themselves remarked on this sort of discrepancy and have expressed their scepticism about the whole idea of a representational theory of mind (Clark, 1998). Some such philosophers claim that folk-psychological concepts like the propositional attitudes will, in due course, be *eliminated* altogether from our scientific understanding of the mind (Churchland, 1979).

At this point, the engineer can't help remarking to herself that her robots, if only they were a little more advanced, would make excellent candidates for a representational theory of mind — far more clear candidates, in fact, than human beings. Unlike the philosophers, she has little difficulty pinning down exactly what she means by an internal representation. These are sentences of first-order predicate calculus. Pressed for further details, she can easily explain how these are encoded in standard data structures familiar to any computer scientist, such as arrays, linked lists, and so on, which in turn are stored in the computer's memory in binary form. Among the various computational processes carried out by the computer are inferences from one such representation to another, and the algorithms executed by these processes provably respect the laws of logic.

5.2 Intentionality and Symbol Grounding

As she persists in her reading of the philosophical literature, the engineer soon discerns a number of overarching themes in the objections raised to

representational theories of the human mind. One of the overarching themes goes to the heart of the whole intentionality issue. How is it that the mind's supposed representations, being comprised of nothing but symbols, can actually be said to *represent* anything? The more attractive replies to this question all seem to appeal, in some way or other, to the mind's "connectedness" to the world. The symbols that go to make up a representation acquire their intentionality because they are *grounded* in the perception of objects in the physical world (Harnad, 1990; Barsalou, 1999).

Setting aside her wider concerns about representational theories of the *human* mind, the engineer is struck afresh by the comparative ease with which she can answer this sort of question for her robots, since they have all been built according to the abductive theory of perception. Take the case of the upper-torso humanoid robot recently constructed (on a pitifully small budget) in her lab. Where do its internal representations come from? Well, taking visual perception as exemplary, a very clean causal story can be recounted that begins with light entering the lenses of the robot's cameras and ends with the storage in the computer's memory of structures standing for sentences of predicate calculus.

This causal story can be told at the purely physical level. But thanks to the abductive theory of perception to which the humanoid's design conforms, a more valuable, parallel narrative can be told at the level of symbols, representations, and logic. Suppose a student places a white cup on the robot's workbench, and the robot's attention mechanism, drawn by the motion, has directed its gaze on this cup. The perpetual stream of pixel arrays delivered by the robot's camera is subjected to various forms pre-processing using off-the-shelf vision algorithms. These off-the-shelf techniques supply a symbolic description □ of many aspects of the image — edges, corners, patches of uniform color, texture, or depth, motion cues, tracked features, and so on.

What happens next is the most important part. This symbolic description □ is submitted to a computational process that carries out abduction. This process produces a prioritized list of explanations □ that conform to the abductive prescription given in the present paper. These are lodged in the computer's memory and their disjunction contributes to the robot's internal

representation of its workbench environment. At the top of the list of predicate calculus sentences is one describing the presence on the workbench of a white vessel with a handle.

$$\exists x [On(x, Workbench) \wedge Is(x, Cup)]$$

This sentence plays a vital role in the selection of the robot's next action. Working forwards from its updated model of the present situation, the robot's planner soon identifies an executable action that will bring it closer to one of its chief goals, namely the production of tea. It carefully positions the tea-pot it is holding over the cup, using visual servoing, and proceeds to pour. The onlookers applaud.

The engineer is then able to explain to the audience how the robot was able to build up a "correct" representation of the scene. The robot "knows there is a cup on the table" because its system of perception conforms to an abductive theory of perception. This theory pins down the process by which the base metal of raw sensor data is transmuted into the gold of meaningful representation, and it does so in such a way that the correctness (or otherwise) of that process can be logically verified. At the same time, it is clear how the process reduces to a causal story at the level of computer science, electronics, or physics, as preferred.

So when pressed to justify the ascription of meaning to the symbols in a predicate calculus sentence like the one above, the engineer can appeal to more than their suggestive English names. Ideally, she would be able to claim that the above sentence is inserted into the robot's model of the current situation if and only if there is indeed a cup on the table. In practise, of course, the robot is sometimes fooled into taking non-cups for cups and cups for non-cups, thanks to poor lighting, unusual surface patterns, and so on. But it generally recovers the correct \exists as soon as it tries to grasp the item in question.

So the engineer's description of the robot's representations makes it sound remarkably as if they have intentionality. The engineer can explain their apparent directedness towards things in the world, under usual operating conditions. She can point to pathological cases in which they are not, in this way, directed at the things they should be, and can legitimize the

sense in which these cases are indeed pathological and can thus be said to be examples of “error”. Finally, she can point to the hypothetical situations that are represented, using novel combinations the same predicates and terms, in the course of the robot’s planning processes, as examples of directness towards something non-existent.

5.3 Empiricism and Idealism

Unfortunately, during one such demonstration there is a philosopher in the audience, who asks the engineer how the terms and predicates in the background theory \square acquire their meaning. Surely any intentionality that can be ascribed to the sentences in \square piggybacks on the intentionality of the sentences in \square , and their intentionality in turn is parasitic on the intentionality of the designer’s mental states.

The engineer — after pointing out that, as a humble engineer, she isn’t really interested in intentionality as such, just in designing and building robots according to sound engineering principles — then reiterates the point that a predicate symbol like *Cup* appears in the robot’s model of the world as a causal consequence of there being a cup on the table, and this symbol is subsequently causally implicated in the robot’s treating the cup as a cup. It would fulfill the same role whether it had been put there by a human designer, or acquired by a learning algorithm, or had appeared through an evolutionary process.

But the philosopher isn’t entirely satisfied by this reply. He would love to ascribe intentionality to the robot’s representations, he says, but he is still bothered by the origin of \square . The engineer patiently explains that, with the current generation of humanoids, it is indeed the case that the background theory \square has been largely hand-coded by legions of hard-working Ph.D students. But work in, for example, inductive logic programming, has shown that similar theories can be learned, right down to the invention of new predicate and function symbols. Future generations of humanoid are likely to acquire more and more of their background theories through such learning techniques.

That’s all very well, says the philosopher, but surely there will be always be something in the original background theory, something to get the whole

process off the ground. The engineer cheerfully concedes that this is true. The background theory will include at least a basic set of sorts for time, objects, and spatial location, and basic predicates for representing the occurrence of events, cause and effect, spatial relations, and so on. Without this foundation at least, the learning algorithms will never work. She pulls a dog-eared commentary on Kant from the shelf, and concludes by quoting “the natural world as we know it ... is thoroughly conditioned by [certain] features: our experience is essentially experience of a spatio-temporal world of law-governed objects conceived of as distinct from our temporally successive experiences of them” (Strawson, 1966, p. 21).

6. Concluding Remarks

One way to think of the theory of perception on offer here is to contrast it with a strictly classical approach to the problem of sensing, in which some physical quantity needs to be measured using an imperfect instrument. From this standpoint, machine vision is analogous to doing graphics backwards. The initial challenge is the “graphics problem” — to find a mathematical model that takes a description of the solid objects in a scene, the lighting conditions and the position of the viewer, and renders the scene for the viewer in question. The vision problem is then just a matter of inverting the equations so the model works backwards, yielding a precise description of the solid objects in the scene given what the viewer sees. Under this approach, uncertainty is treated in the same way that any engineer trying to measure a physical quantity using a sensor would treat it, namely by factoring noise into the mathematical model.

The theoretical paradigm for robot perception supplied by the present paper is very different from this classical approach, but remains entirely compatible with it. Rather than assuming that the ideal final product of the perceptual process is a quantitative, numerical description of the scene, the abductive framework assumes the final product is a qualitative, symbolic description of the scene. The former type of description is now relegated to the intermediate stage between the raw sensor data and the final product. The bridge between the quantitative and the qualitative, between the numerical and the symbolic, is made by abduction, which, drawing on high-

level knowledge, is the final arbiter in the fixation of belief for the perceiving agent.

The result is an account of perception that should be attractive to engineers who want to build sophisticated robots working in realistic environments.¹¹ The main achievement is the provision of a precise theoretical framework in which the relationship between world and acquired representation is clear. Moreover, this theoretical framework can account for several aspects of robot perception in a uniform way. These include the processing of incomplete and uncertain data, the use of top-down information flow, sensor fusion, and the interplay of perception and action. With the aid of a number of logic-based knowledge representation techniques, each of these has been assimilated into the notion that the transformation of raw sensor data into meaningful representation can be cast as a form of logical abduction.

Whether or not the abductive theory of perception has genuine philosophical implications is hard to say. But it is striking that the theory echoes so many themes in the philosophical debate on representation and intentionality, and it may help to expose philosophers engaged in that debate to unfamiliar ideas, such as top-down information flow and active perception, that are crucial to filling out any theory of perception that purports to deliver sentence-like representations.

Acknowledgments

The work described in this paper was supported by EPSRC grants GR/L20023 “Cognitive Robotics”, GR/N13104 “Cognitive Robotics II”, and GR/S02891 “Spatial Reasoning and Perception in a Humanoid Robot”. Thanks to Mark Witkowski, Dave Randell, and Tim Guhl. Thanks also to the reviewers for many helpful suggestions.

¹¹ In addition to robotics, the theoretical framework of this paper has been influential in cognitive vision (Cohn, *et. al.*, 2003) and spatial reasoning (Hazarika & Cohn, 2002; Remolina & Kuipers, 2004).

References

- Aloimonos, J., Weiss, I. & Bandyopadhyay, A. (1987). Active Vision. In *Proceedings 1st International Conference on Computer Vision* (pp. 35–54).
- Ballard, D.H. (1991). Animate Vision. *Artificial Intelligence*, 48, 57–86.
- Barsalou, L.W. (1999). Perceptual Symbol Systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Brentano, F. (1874). *Psychology from an Empirical Standpoint*. Trans. A.C.Rancurello, D.B.Terrell and L.L.McAlister. London: Routledge (1995).
- Brooks, R.A. (1981). Symbolic Reasoning Among 3-D Models and 2-D Images. *Artificial Intelligence*, 17, 285–348.
- Brooks, R.A. (1991). Intelligence Without Reason. *Proceedings 1991 International Joint Conference on Artificial Intelligence (IJCAI 91)* (pp. 569–595).
- Brooks, R.A., Breazeal, C., Marjanovic, M., Scassellati, B. & Williamson, M. (1999). The Cog Project: Building a Humanoid Robot. In C.Nehaniv (Ed.), *Computation for Metaphors, Analogy, and Agents*. Lecture Notes in Artificial Intelligence 1562. Springer. pp. 52–87.
- Cavanagh, P. (1999). Top-Down Processing in Vision. In *The MIT Press Encyclopedia of the Cognitive Sciences*. pp. 844–845.
- Chella, A., Frixione, M. & Gaglio, S. (1997). A Cognitive Architecture for Artificial Vision. *Artificial Intelligence*, 89(1–2), 73–111.
- Chella, A., Frixione, M. & Gaglio, S. (2000). Understanding Dynamic Scenes. *Artificial Intelligence*, 123(1–2), 89–132.
- Chrisley, R. & Ziemke, T. (2003). Embodiment. In L.Nadel (ed.) *The Macmillan Encyclopedia of Cognitive Science*. Macmillan. Vol. 1. pp. 1102–1108.
- Churchland, P.M. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge University Press.
- Clark, A. (1998). *Being There: Putting Brain, Body, and World Back Together*. MIT Press.
- Clark, J.J. & Yuille, A.L. (1990). *Data Fusion for Sensory Information Processing Systems*. Kluwer.

- Cohn, A.G., Magee, D., Galata, A., Hogg, D. & Hazarika, S. (2003). Towards an Architecture for Cognitive Vision Using Qualitative Spatio-Temporal Representations and Abduction. In C.Freska, C.Habel & K.F.Wender (eds.), *Spatial Cognition III*. Springer. pp. 232–248.
- Coradeschi, S. & Saffiotti, A. (2000). Anchoring Symbols to Sensor Data; Preliminary Report. *Proceedings 2000 American Association for Artificial Intelligence Conference (AAAI 2000)* (pp. 129–135).
- Crane, T. (2003). *The Mechanical Mind: A Philosophical Introduction to Minds, Machines, and Mental Representation* (2nd edition). Routledge.
- Fodor, J. (1975). *The Language of Thought*. Harvester.
- Gibson, J.J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.
- Green, C. (1969). Applications of Theorem Proving to Problem Solving. *Proceedings 1969 International Joint Conference on Artificial Intelligence (IJCAI 69)* (pp. 219–240).
- Grossberg, S. (1999). How Does the Cerebral Cortex Work? Learning, Attention, and Grouping by the Laminar Circuits of Visual Cortex. *Spatial Vision*, 12, 163–186.
- Harnad, S. (1990). The Symbol Grounding Problem. *Physica D*, 42, 335–346.
- Hayes, P.J. (1977). In Defence of Logic. In *Proceedings 1977 International Joint Conference on Artificial Intelligence (IJCAI 77)* (pp. 559–565).
- Hayes, P.J. (1985). The Second Naive Physics Manifesto. In J.R.Hobbs & R.C.Moore (eds.), *Formal Theories of the Commonsense World*. Ablex. pp. 1–36.
- Hazarika, S. & Cohn, A.G. (2002). Abducing Qualitative Spatio-Temporal Histories from Partial Observations. In *Proceedings 2002 Knowledge Representation Conference (KR 2002)* (pp. 14–25).
- Jain, R., Kasturi, R. & Schunk, B.G. (1995). *Machine Vision*. McGraw-Hill.
- Jones, J.L & Flynn, A.M. (1993). *Mobile Robots: Inspiration to Implementation*. A.K.Peters.
- Josephson, J.R. & Josephson, S.G. (1994). *Abductive Inference: Computation, Philosophy, Technology*. Cambridge University Press.

- Jung, C.G. (1998). Situated Abstraction Planning by Abductive Temporal Reasoning. In *Proceedings 1998 European Conference on Artificial Intelligence (ECAI 98)* (pp. 383–387).
- Kakas, A.C. & Evans, C. (1992). Hypothetico-Deductive Reasoning. In *Proceedings 1992 International Conference on Fifth Generation Computer Systems (FGCS 92)* (pp. 546–554).
- Koenderink, J.J. & van Doorn, A.J. (1979). The Internal Representation of Solid Shape with Respect to Vision. *Biological Cybernetics*, 32, 211–216.
- Kowalski, R.A. (1995). Using Meta-Logic to Reconcile Reactive with Rational Agents. In K.R.Apt & F.Turini (eds.), *Meta-Logics and Logic Programming*. MIT Press. pp. 227–242.
- Lespérance, Y., Levesque, H.J., Lin, F., Marcu, D., Reiter, R. & Scherl, R.B. (1994). A Logical Approach to High-Level Robot Programming: A Progress Report. In B.Kuipers (ed.), *Control of the Physical World by Intelligent Systems: Papers from the 1994 AAI Fall Symposium* (pp. 79–85).
- Lévy, F. & Quantz, J.J. (1998). Representing Beliefs in a Situated Event Calculus. In *Proceedings 1998 European Conference on Artificial Intelligence (ECAI 98)* (pp. 547–551).
- Lifschitz, V. (1994). Circumscription. In D.M.Gabbay, C.J.Hogger & J.A.Robinson (eds.), *The Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press. pp. 297–352.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman & Company.
- McCarthy, J. (1986). Applications of Circumscription to Formalizing Common Sense Knowledge. *Artificial Intelligence*, 26, 89–116.
- McCarthy, J. (1989). Artificial Intelligence, Logic and Formalizing Common Sense. In R.Thomason (ed.), *Philosophical Logic and Artificial Intelligence*. Kluwer Academic. pp. 161–190.
- Mondada, F., Franzi, E. & Ienne, P. (1993). Mobile Robot Miniaturisation: A Tool For Investigation In Control Algorithms. In *Proceedings of the Third International Symposium on Experimental Robotics* (pp. 501–513).

- Moore, R.C. (1982). The Role of Logic in Knowledge Representation and Commonsense Reasoning. In *Proceedings 1982 American Association for Artificial Intelligence Conference (AAAI 82)* (pp. 428–433).
- Morgenstern, L. (2001). Mid-Sized Axiomatizations of Commonsense Problems: A Case Study in Egg Cracking. *Studia Logica*, 67, 333–384.
- Mueller, E.T. (2004). Event Calculus Reasoning Through Satisfiability. *The Journal of Logic and Computation*, to appear.
- Muggleton, S. (1991). Inductive Logic Programming. *New Generation Computing*, 8(4), 295–318.
- Newell, A. (1982). The Knowledge Level. *Artificial Intelligence*, 18(1), 87–127.
- Niebur, E. & Koch, C. (1998). Computational Architectures for Attention. In R.Parasuraman (ed.) *The Attentive Brain*, MIT Press, pp. 163–186.
- Nilsson, N.J. (ed.). (1984). *Shakey the Robot*. SRI Technical Note no. 323. SRI, Menlo Park, California.
- Nilsson, N.J. (1991). Logic and Artificial Intelligence. *Artificial Intelligence*, 47, 31–56.
- Peirce, C.S. (1903). Pragmatism as the Logic of Abduction. In *The Essential Peirce: Selected Philosophical Writings*, vol. 2 (ed. The Peirce Edition Project), Indiana University Press (1998), pp. 226–242.
- Peterson, M.A. (2003). Vision: Top-Down Effects. In L.Nadel (ed.) *The Macmillan Encyclopedia of Cognitive Science*. Macmillan. Vol. 4. pp. 500–504.
- Pirri, F. & Finzi, A. (1999). An Approach to Perception in Theory of Actions: Part I. *Electronic Transactions on Artificial Intelligence*, 3, 19–61.
- Pirri, F. & Romano, M. (2002). A Situation-Bayes View of Object Recognition Based on SymGeons. In *Proceedings Third International Cognitive Robotics Workshop*.
- Poole, D. (1995). Logic Programming for Robot Control. In *Proceedings 1995 International Joint Conference on Artificial Intelligence (IJCAI95)* (pp. 150–157).

- Pylyshyn, Z.W. (1999). Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception. *Behavioral and Brain Sciences*, 22, 341–364.
- Pylyshyn, Z.W. (2001). Visual Indexes, Preconceptual Objects, and Situated Vision. *Cognition*, 80, 127–158.
- Remolina, E. & Kuipers, B. (2004). Towards a General Theory of Topological Maps. *Artificial Intelligence*, 152, 47–104.
- Reiter, R. & Mackworth, A.K. (1989). A Logical Framework for Depiction and Image Interpretation. *Artificial Intelligence*, 41(2), 125–155.
- Rock, I. (1983). *The Logic of Perception*. MIT Press.
- Russell, S. & Norvig, P. (2002). *Artificial Intelligence: A Modern Approach*. 2nd edition. Prentice Hall.
- Scherl, R. & Levesque, H. (1993). The Frame Problem and Knowledge Producing Actions. In *Proceedings 1993 American Association for Artificial Intelligence Conference (AAAI 93)* (pp. 689–695).
- Searle, J. (1980). Minds, Brains, and Programs. *Behavioral and Brain Sciences*, 3, 417-424.
- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Shanahan, M.P. (1996a). Robotics and the Common Sense Informatic Situation. In *Proceedings 1996 European Conference on Artificial Intelligence (ECAI 96)* (pp. 684–688).
- Shanahan, M.P. (1996b). Noise and the Common Sense Informatic Situation for a Mobile Robot. In *Proceedings 1996 American Association for Artificial Intelligence Conference (AAAI 96)* (pp. 1098-1103).
- Shanahan, M.P. (1997a). *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. MIT Press.
- Shanahan, M.P. (1997b). Noise, Non-Determinism and Spatial Uncertainty. In *Proceedings 1997 American Association for Artificial Intelligence Conference (AAAI 97)* (pp. 153–158).
- Shanahan, M.P. (1999). The Event Calculus Explained. In Wooldridge, M.J. & Veloso, M. (eds.) *Artificial Intelligence Today*. Springer-Verlag Lecture Notes in Artificial Intelligence no. 1600. Springer-Verlag. pp. 409-430.

- Shanahan, M.P. (2000). Reinventing Shakey. In J.Minker (ed.) *Logic-Based Artificial Intelligence*. Academic Press. pp. 233–253.
- Shanahan, M.P. (2002). A Logical Account of Perception Incorporating Feedback and Expectation. In *Proceedings 2002 Knowledge Representation Conference (KR 2002)* (pp. 3–13).
- Shanahan, M.P. (2003). The Frame Problem. In L.Nadel (ed.) *The Macmillan Encyclopedia of Cognitive Science*. Macmillan. Vol. 2. pp. 144–150.
- Shanahan, M.P. (2004). An Attempt to Formalise a Non-Trivial Benchmark Problem in Common Sense Reasoning. *Artificial Intelligence, 153*, 141–165.
- Shanahan, M.P. & Randell, D. (2004). A Logic-Based Formulation of Active Visual Perception. In *Proceedings 2004 Knowledge Representation Conference (KR 2004)* (pp. 64–72).
- Shanahan, M.P. & Witkowski, M. (2001). High-Level Robot Control Through Logic, In *Intelligent Agents VII*. Springer. pp. 104–121.
- Strawson, P. F. (1966). *The Bounds of Sense*. Methuen.
- Ullman, S. (1996). *High-Level Vision: Object Recognition and Visual Cognition*. MIT Press.
- Von Uexküll, J. (1957). A Stroll Through the Worlds of Animals and Men. In C.H.Schiller (ed.) *Instinctive Behaviour: The Development of a Modern Concept*. International Universities Press. pp. 5–80.
- Van der Does, J. & Van Lambalgen, M. (2000). A Logic of Vision. *Linguistics and Philosophy, 23*, 1–92.
- Waltz, D. (1975). Understanding Line Drawings of Scenes with Shadows. In P.H.Winston (ed.) *The Psychology of Computer Vision*. McGraw-Hill. pp. 19–91.
- Winston, P.H. (1984). *Artificial Intelligence* (2nd edition). Addison Wesley.
- Zilberstein, S. & Russell, S.J. (1993). Anytime Sensing, Planning and Action: A Practical Model for Robot Control. In *Proceedings 1993 International Joint Conference on Artificial Intelligence (IJCAI 93)* (pp. 1402–1407).