What Sort of Computation Mediates Best between Perception and Action?

Murray Shanahan

Department of Computer Science,
Queen Mary and Westfield College,
Mile End Road,
London E1 4NS,
England.
mps@dcs.qmw.ac.uk

Abstract

This paper addresses the question of what sort of computation should be used to mediate between perception and action. Drawing on recent work in the area of Cognitive Robotics, the paper argues for the viability of an answer based on a rigorous, logical account of interleaved perception, planning and action. A number of common criticisms of approaches in this style are reviewed.

Introduction

This paper addresses a fundamental methodological question in Artificial Intelligence, namely what sort of computation mediates best between perception and action in a mobile robot. The answer offered here has a traditional feel, and appeals to a conceptual framework which is inherited from the time of the field's inception in the Fifties. In particular, I support the view that robot design can be based on the manipulation of sentences of logic. One of the paper's main concerns is to defend this idea from criticisms which rest on a faulty view of logic and its limitations. Rather than relying solely on abstract argument to justify my views, the paper appeals to work that has been carried out with a real robot based on a rigorously logical formalism for representing and reasoning about action, continuous change, shape and space.

Before proceeding with the argument of the paper, the question under discussion requires clarification. Two points come to mind.

First, determining the sort of computation that mediates best between perception and action is, as far as this paper is concerned, an engineering matter, albeit one which is made more interesting by its deployment of terms which have currency in contemporary philosophy of mind. That is to say, while we are free to seek inspiration from biology, psychology or philosophy, the only criteria on which an answer to the question should be judged are engineering criteria, such as whether our choice of computational medium facilitates the construction of more capable, more robust, easier-to-modify, and easier-to-maintain machines.

Second, if we accept the Church-Turing thesis then, from a reductionist point of view, all computation is the same. But this observation is unhelpful in the present context. Our interest here is in the choice of the fundamental units of computation and (if relevant) of representation, and how these fit in to the architecture of an embodied agent situated in a world like our own.

We could opt for neural computation, in which the unit of computation is the artificial neuron, and representation corresponds to a pattern of activation over a network of such units. Or we could opt for the logical formula as the unit of representation, with rewrites of these formulae as the units of computation.

We could altogether reject the assumption that representation needs to play a part in what mediates between perception and action. Indeed, we might even reject the whole idea of computation. According to the argument for this position, the human brain (on which the argument assumes our research should be based) is a complex dynamical system, and it is misleading to suggest that it performs computation. But whatever choices we make, they will impinge on the engineering concerns voiced above.

The paper opens by rehearsing some of the arguments for the importance of embodiment in AI research, and then moves on to survey some of the advantages of a logic-based approach. This leads to the presentation of the paper's main argument, which is that a logical approach to embodied cognition is feasible. A rigorously logical account of perception, reason and action is presented which confronts the usual criticisms levelled at representational approaches to AI.

1 Embodiment

In the late Eighties and early Nineties, traditional approaches to Artificial Intelligence were frequently lambasted for working in disembodied, abstract domains.

> The only input to most AI programs is a restricted set of simple assertions deduced from the real data by humans. The problems of recognition, spatial understanding, dealing with sensor noise, partial models, etc. are all ignored.
>
> [Brooks, 1991a, page 143]

> Traditional Artificial Intelligence has adopted a style of research where the agents that are built to test theories of

> intelligence are essentially problem solvers that work in a symbolic abstract domain.
>
> [Brooks, 1991b, page 583]

Taking their cue from Rodney Brooks, many researchers placed a new emphasis on *autonomous* systems *situated* in dynamic environments.

> Rather than working on computer programs that appear to mimic some limited aspect of high-level human intelligence . . . the new approach concentrates instead on studying complete autonomous agents.
>
> [Cliff, 1994, page 800]

For many researchers, the issue of *embodiment* is also crucial because,

> . . . unless you saddle yourself with all the problems of making a concrete agent take care of itself in the real world, you will tend to overlook, underestimate or misconstrue the deepest problems of design.
>
> [Dennett, 1994, page 143]

The “deepest” problems of design often concern incompleteness and uncertainty.

> Sensors deliver very uncertain values even in a stable world. . . . The data delivered by sensors are not direct descriptions of the world as objects and their relationships. . . . Commands to actuators have very uncertain effects.
>
> [Brooks, 1991c, page 5]

Underlying the concerns of these researchers is the belief that,

- the isolated study of different aspects of intelligence leads to the development of incompatible sub-theories, and
- the temptation to idealise away the imperfection of a robot’s connection to the world leads to the development of theories which are useless in practice.

These concerns can be addressed without appealing to embodiment, by conducting research using complete agents situated in some artificial environment, such as the Internet. But a further argument mitigates against this approach, insofar as we are interested in one day achieving human level intelligence in a machine as well as in designing useful products for today.

This is how the argument goes. The primary purpose of cognition is to enhance an agent’s ability to interact with a world of *spatio-temporally located objects* given only *incomplete* and *uncertain* information. The incompleteness arises because of an agent’s limited window on the world, and the uncertainty arises because of sensor and motor noise. In other words, the agent is confronted with what McCarthy calls the *common sense informatic situation* [McCarthy, 1989]. A capacity to deal with the common sense informatic situation is the substrate on which other cognitive skills rest. Only by building on such a substrate will we be able to duplicate human-level cognitive ability in a machine.[1]

This doesn’t set a limit, in principle, to the kind of research that can be carried out with agents situated in a simulated world. The argument is methodological rather than metaphysical. The suggestion is only that human-level cognitive skills are tied to the human epistemic predicament — the need to act in the presence of incomplete and uncertain information about a world of spatio-temporally located objects — and that this predicament is a consequence of our embodiment.

2 Logic

The logicist agenda in AI dates back to the Fifties [McCarthy, 1959]. According to the logical approach to AI, knowledge is represented by sentences in a formal language, and intelligent behaviour is mediated by the proof of theorems with those sentences. In the late Sixties, Green presented his classical account of planning [Green, 1969], and in the early Seventies the logical approach was applied to robotics in the form of the Shakey project [Nilsson, 1984]. Sadly, further progress was slow, and the popularity of logic in the robotics community subsequently declined. However, in spite of frequent premature obituaries, the logicist research programme has been vigorously pursued by a substantial minority of enthusiasts in the AI community ever since McCarthy’s 1959 proposal. Today, as we shall see, the logic-based approach to robotics is enjoying a renaissance.

It’s important to note that the logicist prescription does not demand a one-to-one correspondence between the data structures in the machine and the sentences of the chosen formal language. In other words, representations in the machine do not have to be explicitly stored sentences of logic. Similarly, the logicist prescription does not demand the use of algorithms whose state transitions correspond exactly to the steps of a proof. In other words, the machine does not have to implement a theorem prover directly. Between the abstract description of a logic-based AI program and the actual implementation can come many steps of transformation, compilation, and optimisation. The final product’s functional equivalence to the abstract specification counts more than anything else.[2]

The chief advantages of a logic-based approach to AI in general, and to robotics in particular, are threefold. First, if a robot’s design is logic-based, we can supply a rigorous, mathematical account of its success or otherwise at achieving its goals. Second, because its knowledge and goals are expressed in a universal declarative language, such a robot is easily modified and maintained. Third, it is relatively clear how to incorporate high-level cognitive skills in a logic-based robot, such as the ability to plan, to reason about other agents, or to reason about its own knowledge.

---

[1] This includes linguistic ability. Perhaps natural language will even turn out to be a relatively straightforward phenomenon once we understand how to cope with the common sense informatic situation. After all, linguistic skills are evolution’s most recent innovation, the final fold in the human cortex.

[2] However, it should be emphasised that, for many researchers, the explicit storage of sentences of logic and their manipulation by theorem proving techniques is an important ideal. The argument here often appeals to the idea that declaratively represented knowledge can be used in many different ways.

The remaining two sections of this paper report recent logic-based work on perception and action in robots. In doing so, they address some common criticisms levelled at traditional approaches to AI. The first of these criticisms concerns the supposed inability of the traditional symbolic approach to AI to deal with incomplete information, the hallmark of the common sense informatic situation. Here are two representative quotes.

> The key problem I see with [all work in the style of Shakey] is that it relied on the assumption that a complete world model could be built internally and then manipulated.
>
> [Brooks, 1991b, page 577]

> [In traditional AI] the key issue on which emphasis is laid is a complete, correct internal model, a perfect copy of the world (with all its object and relationships) inside the system, which the system can rely on to predict how the problem can be solved.
>
> [Maes, 1993, page 4]

Unfortunately, these claims are based on an incorrect view of the nature of the representational approach to AI. More specifically, they betray a lack of understanding of the nature of predicate logic, which underpins the symbolic paradigm. At a foundational level, the problem of incomplete information was solved by Frege and Tarski, who gave us a good formal account of disjunction and existential quantification. Naturally, a good deal of work is still required to translate their mathematical insights into robotics practice. Section 3 reports an attempt to do this for robot perception.

The second criticism to be addressed relates to robot action. According to Brooks, because it worked in carefully engineered, static domain, the planner in Shakey,

> . . . could ignore the actual world, and operate in the model to produce a plan of action for the robot to achieve whatever goal it had been given.
>
> [Brooks, 1991b, page 570]

Maes makes a similar point.

> [In traditional AI] the central system evaluates the current situation (as represented in the internal model) and uses a search process to systematically explore the different ways in which this situation can be affected so as to achieve the desired goal.
>
> [Maes, 1993, page 4]

Accordingly, a robot like Shakey is slow and intolerant to changes in its environment. This view is now uncontroversial, and robot architectures which incorporate a degree of *reactivity* are the norm. However, although systems which combine reactive and deliberative elements have been around for some time [Georgeff & Lansky, 1987], [Mitchell, 1990], [Gat, 1992], a rigorous logical account of interleaved planning, sensing and acting has only recently been achieved. This account is outlined in Section 4.

3 A Logical Account of Robot Perception

This section provides a technical overview of the abductive account of robot perception presented in [Shanahan, 1996a] and [Shanahan, 1996b]. In this account, a mobile robot's sensor data is abductively explained by postulating the existence of objects with suitable locations and shapes. The account is the product of the following three steps.

1. Design a logic-based formalism for reasoning about action and space.
2. Using this formalism, construct a theory which captures the robot's relationship to the world.
3. Consider the task of sensor data assimilation as a form of abduction with respect to this theory.

The logic-based formalism for reasoning about action and change must be able to cope with a variety of phenomena. First, because the robot's motion is continuous, it must be able to represent *continuous change*. Second, since events in the world can occur at any time, it must be able to represent *concurrent events*. Third, in order to deal with noise it needs to be able to handle actions and events with *non-deterministic effects*.

The formalism employed in [Shanahan, 1996a] is adapted from [Shanahan, 1995b]. It is based on the *event calculus* of Kowalski and Sergot [1986], but is expressed in full predicate calculus augmented with *circumscription* [McCarthy, 1986]. Circumscription is deployed to overcome the frame problem, using a technique inspired by [Kartha & Lifschitz, 1995] and explored more fully in [Shanahan, 1997].

Space, in this formalism, is represented as the plane $\mathbb{R}^2$. Objects, including both the robot and the obstacles in its workspace, occupy open, path-connected subsets of $\mathbb{R}^2$. A fluent for spatial occupancy is employed, and for reasons set out in [Shanahan, 1995a], spatial occupancy is minimised using circumscription. The formalism includes a number of axioms about continuous change and spatial occupancy, which are gathered together in the theory Σ_B.

The theory Σ_E, which is expressed in the language of this formalism, describes the effects of the robot's motor activity on the world, and the consequent effect of the world on the robot's sensors. For example, consider a wheeled mobile robot equipped with bump sensors. Σ_E captures the fact that, if the robot executes a move forward command, its location in $\mathbb{R}^2$ will start to vary continuously. Σ_E also captures the fact that this continuous variation in location will cease if the robot collides with an obstacle, and that the robot's bump sensors will be tripped as a result of the collision.

Sensor data assimilation can now be considered as abduction. The deterministic, noise-free case is presented first. If a stream of sensor data is represented as the conjunction Ψ of a set of observation sentences, the task is to find an explanation of Ψ in the form of a logical description (a map) Δ_M of the initial locations and shapes of a number of objects, such that,

$$\Sigma_B \land \Sigma_E \land \Delta_N \land \Delta_M \vDash \Psi$$

where,

- Σ_B is a background theory, comprising axioms about change (including continuous change), action, space, and shape,

- Σ_E is a theory relating the shapes and movements of objects (including the robot itself) to the robot's sensor data, and
- Δ_N is a logical description of the movements of objects, including the robot itself.

The incompleteness of the robot's knowledge due to its limited window on the world is reflected in the fact that there are always many explanations Δ_M for any given collection Ψ of sensor data.

The uncertainty in the robot's knowledge that arises from the inevitable presence of sensor and motor noise can be considered as non-determinism [Shanahan, 1996b]. Here we'll consider only motor noise, but the technical issues are the same for sensor noise. Instead of including in Σ_E a formula describing an exact trajectory for the robot as it moves, we can include a formula which describes an ever-increasing circle of uncertainty within which the robot's location is known to fall.

This modification motivates the deployment of a new form of abduction. The non-monotonic nature of Δ_M, which uses circumscription to minimise spatial occupancy, entitles us to use a consistency-based form of abduction similar to that described by Reiter [1987]. Given a stream of sensor data Ψ, the task is now to find conjunctions Δ_M such that,

$$\Sigma_B \wedge \Sigma_E \wedge \Delta_N \wedge \Delta_M \not\models \neg \Psi.$$

Unlike the symbols used in disembodied systems, the symbols appearing in Δ_M are *grounded* in the robot's interaction with the world (see [Harnad, 1990]), as well as acquiring *meaning* via Tarski-style model theory. The dual notions of groundedness and meaning allow us to appeal to the *correctness* of the robot's representations and reasoning processes in explaining the success or failure of its behaviour.

For further details of the material presented in this section, the reader should consult [Shanahan, 1996a] and [Shanahan, 1996b].

4 A Logical Account of Robot Action

Work in the logicist tradition on planning and acting dates back to Green's seminal contribution in the late Sixties [Green, 1969]. However, it's only recently that the spirit of this work has been revived in the form of the Cognitive Robotics research programme [Lespérance, *et al.*, 1994]. This section reports recent efforts to supply a more up-to-date logical account of the interplay between planning, sensing and acting.

One approach to this issue is to preserve as much of classical Green-style planning as possible. This is the policy adopted by Levesque [1996]. Green's characterisation of the planning task, which is based on the situation calculus [McCarthy & Hayes, 1969] is as follows: given a description Σ of the effects of actions, an initial situation S0, and a goal Γ, find a sequence of actions σ, such that Σ logically entails that Γ is the case after the execution of σ in S0.

Green's account, in which planning and execution are sharply delineated, assumes that it's appropriate to attempt to plan an entire course of action given only knowledge of the initial situation (see the two quotes at the end of Section 2). Levesque's modified account goes beyond this by allowing plans which incorporate sensing actions, as well as familiar programming constructs such as repetition and conditional action.

A potential drawback of Levesque's work is that it still assumes a sharp division between planning and execution. A complete plan for achieving the goal has to be constructed before any action is executed. In this respect, the logic programming account offered by Kowalski [1995] departs more radically from classical planning. Kowalski's presentation, which is the basis of [Shanahan, 1996c], interleaves planning, sensing and execution. To achieve this interleaving, the planner has to be able to recommend an action to be executed using only a bounded amount of computation, because the planning process is subject to constant suspension while actions are performed and new sensor data is acquired.

In Kowalski's proposal, the need for such a planner motivates the introduction of a style of formula for describing the effects of actions which is incompatible with that used in Σ_E in [Shanahan, 1996a]. This style of formula also undermines one of the most pleasing properties of classical planning, namely the tight logical relationship between the goal and the plan. In [Shanahan, 1996c], Kowalski's account of interleaved planning, sensing and execution is adopted, but the event calculus style of effect formula used in Σ_E is restored, along with this property.

Following [Eshghi, 1988], planning in the event calculus can be considered as an abductive process in a way which resembles Section 3's treatment of sensor data assimilation. Given a goal Γ, the task is to find a sequence of robot actions Δ_N such that,

$$\Sigma_B \wedge \Sigma_E \wedge \Delta_N \wedge \Delta_M \models \Gamma$$

where Σ_B, Σ_E and Δ_M are defined as in Section 3.

Planning in the presence of noise can be given a treatment which is symmetrical to that supplied for perception in the presence of noise. Given a goal Γ, the task is to find a sequence of robot actions Δ_N such that,

$$\Sigma_B \wedge \Sigma_E \wedge \Delta_N \wedge \Delta_M \not\models \neg \Gamma.$$

For more details, the reader is referred to [Shanahan, 1996c], where this form of planning is married to Kowalski's account of interleaved planning, sensing and execution, and applied to robot path planning using a trapezoidal cell decomposition technique taken from [Latombe, 1991].

Conclusion

In response to McCarthy's appeal to Dreyfus to supply a well defined problem which he believes the logical approach to Artificial Intelligence will have difficulty solving [McCarthy, 1996], Dreyfus submits the following challenge to "McCarthy and his followers",

> How does one spell out propositionally our know-how concerning the body and its motor and perceptual systems?
>
> [Dreyfus, 1996]

The thrust of the work reported here is not *our* know-how, nor *human* motor and perceptual systems, but rather those of a mobile robot. However, I believe that Dreyfus is right to place emphasis on the issues of embodiment, perception and action. The present paper and the developments it reports offer a reply to Dreyfus's challenge.

The paper's main purpose, though, is to answer the question posed in the title. What sort of computation mediates best between perception and action? The foregoing abductive accounts of perception and planning licence a range of methodological options. At one extreme, the fundamental units of representation are taken to be sentences of formal logic, and the fundamental unit of computation is the proof step. Ideally, the path from specification to implementation is then a very short one, involving simply the application of a general purpose theorem prover to the theories Σ_B and Σ_E.

This approach preserves all the advantages of declarative representation. The same sentences of logic and the same theorem prover can perform both abductive sensor data assimilation and planning, as well as other reasoning tasks involving Σ_B and Σ_E. Unfortunately, no general purpose theorem prover exists which is up to the job, and it doesn't seem likely, at the time of writing, that future research will come up with one.

At the other end of the methodological spectrum, algorithms for interleaved planning, perception and action can be hand-designed, and proved correct with respect to formal specifications derived from the logical accounts supplied above. This approach has several drawbacks. First, there is no systematic process by which the implementation is derived from the specification. Second, the opaqueness of the implementation with respect to the specification renders it difficult to maintain and modify.

The most attractive option, in my opinion, is a logic programming approach, which lies in the middle of the methodological spectrum. Following Kowalski's slogan, "Algorithm = Logic + Control", the idea here is to preserve as much as possible of the logic of the specification, while rendering it into a computationally feasible form [Kowalski, 1979]. This is achieved by isolating a clausal fragment, subjecting the result to transformation employing the many tricks of the logic programmer's trade, and submitting the final product to an SLDNF theorem prover.

As well as maintaining a close relationship between specification and implementation, and retaining many of the advantages of declarative representation, the logic programming option has the attraction that it renders computation *transparent*. Intermediate computational states are representationally meaningful, since they correspond to sentences of logic expressed in the same language as Σ_B and Σ_E. One consequence of this is that the computational process can be interrupted at any time and still produce useful results, a feature taken advantage of in the approach to interleaved planning, sensing and acting presented in Section 4.

At the time of writing, the ideas in this paper are at various stages of implementation on actual robots. The abductive treatment of perception in Section 3 has been realised on a small two-wheeled robot with bump sensors. Section 4's account of robot action has been tested in simulation, but awaits implementation on a real robot.

References

[Brooks, 1991a] R.A.Brooks, Intelligence Without Representation, *Artificial Intelligence*, vol 47 (1991), pages 139-159.

[Brooks, 1991b] R.A.Brooks, Intelligence Without Reason, *Proceedings IJCAI 91*, pages 569-595.

[Brooks, 1991c] R.A.Brooks, Artificial Life and Real Robots, *Proceedings of the First European Conference on Artificial Life (1991)*, pages 3-10.

[Cliff, 1994] D.Cliff, AI and A-Life: Never Mind the Blocksworld, *Proceedings ECAI 94*, pages 799-804.

[Dennett, 1994] D.Dennett, The Practical Requirements for Making a Conscious Robot, *Philosophical Transactions of the Royal Society of London, Series A*, vol 349 (1994), pages 133-146.

[Dreyfus, 1996] H.L.Dreyfus, Response to My Critics, *Artificial Intelligence*, vol 80 (1996), pages 171-191.

[Eshghi, 1988] K.Eshghi, Abductive Planning with Event Calculus, *Proceedings of the Fifth International Conference on Logic Programming* (1988), pages 562-579.

[Gat, 1992] E.Gat, Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots, *Proceedings AAAI 92*, pages 809-815.

[Georgeff & Lansky, 1987] M.Georgeff and A.Lansky, Reactive Reasoning and Planning, *Proceedings AAAI 87*, pages 677-682.

[Green, 1969] C.Green, Applications of Theorem Proving to Problem Solving, *Proceedings IJCAI 69*, pages 219-240.

[Harnad, 1990] S.Harnad, The Symbol Grounding Problem, *Physica D*, vol 42 (1990), pages 335-346.

[Kartha & Lifschitz, 1995] G.N.Kartha and V.Lifschitz, A Simple Formalization of Actions Using Circumscription, *Proceedings IJCAI 95*, pages 1970-1975.

[Kowalski, 1979] R.A.Kowalski, Algorithm = Logic + Control, *Communications of the ACM*, vol 22 (1979), pages 424-436.

[Kowalski, 1995] R.A.Kowalski, Using Meta-Logic to Reconcile Reactive with Rational Agents, in *Meta-Logics and Logic Programming*, ed. K.R.Apt and F.Turini, MIT Press (1995), pages 227-242.

[Kowalski & Sergot, 1986] R.A.Kowalski and M.J.Sergot, A Logic-Based Calculus of Events, *New Generation Computing*, vol 4 (1986), pages 67-95.

[Latombe, 1991] J.-C.Latombe, *Robot Motion Planning*, Kluwer Academic Press (1991).

[Lespérance, *et al.*, 1994] Y.Lespérance, H.J.Levesque, F.Lin, D.Marcu, R.Reiter, and R.B.Scherl, A Logical Approach to High-Level Robot Programming: A Progress Report, in *Control of the Physical World by Intelligent Systems: Papers from the 1994 AAAI Fall Symposium*, ed. B.Kuipers, New Orleans (1994), pages 79-85.

[Levesque, 1996] H.Levesque, What Is Planning in the Presence of Sensing? *Proceedings AAAI 96*, pages 1139-1146.

[Maes, 1993] P.Maes, Behavior-Based Artificial Intelligence, *Proceedings of the Second International Conference on the Simulation of Adaptive Behavior (SAB 93)*, pages 2-10.

[McCarthy, 1959] J.McCarthy, Programs with Common Sense, *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, Her Majesty's Stationery Office, London (1959), pages 75-91.

[McCarthy, 1986] J.McCarthy, Applications of Circumscription to Formalizing Common Sense Knowledge, *Artificial Intelligence*, vol 26 (1986), pages 89-116.

[McCarthy, 1989] J.McCarthy, Artificial Intelligence, Logic and Formalizing Common Sense, in *Philosophical Logic and Artificial Intelligence*, ed. R.Thomason, Kluwer Academic (1989), pages 161-190.

[McCarthy, 1996] J.McCarthy, Book Review of "What Computers Still Can't Do" by Hubert Dreyfus, *Artificial Intelligence*, vol 80 (1996), pages 143-150.

[McCarthy & Hayes, 1969] J.McCarthy and P.J.Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in *Machine Intelligence 4*, ed. D.Michie and B.Meltzer, Edinburgh University Press (1969), pages 463-502.

[Mitchell, 1990] T.Mitchell, Becoming Increasingly Reactive, *Proceedings AAAI 90*, pages 1051-1058.

[Nilsson, 1984] N.J.Nilsson, ed., *Shakey the Robot*, SRI Technical Note no. 323 (1984), SRI, Menlo Park, California.

[Reiter, 1987] R.Reiter, A Theory of Diagnosis from First Principles, *Artificial Intelligence*, vol 32 (1987), pages 57-95.

[Shanahan, 1995a] M.P.Shanahan, Default Reasoning about Spatial Occupancy, *Artificial Intelligence*, vol 74 (1995), pages 147-163.

[Shanahan, 1995b] M.P.Shanahan, A Circumscriptive Calculus of Events, *Artificial Intelligence*, vol 77 (1995), pages 249-284.

[Shanahan, 1996a] M.P.Shanahan, Robotics and the Common Sense Informatic Situation, *Proceedings ECAI 96*, pages 684-688.

[Shanahan, 1996b] M.P.Shanahan, Noise and the Common Sense Informatic Situation for a Mobile Robot, *Proceedings AAAI 96*, pages 1098-1103.

[Shanahan, 1996c] M.P.Shanahan, Planning, Sensing and Acting in the Common Sense Informatic Situation, in preparation.

[Shanahan, 1997] M.P.Shanahan, *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*, MIT Press (1997), to appear.