Fridges, Elephants, and the Meaning of Autonomy and Intelligence

Ricardo Sanz, Fernando Matía and Santos Galán

Universidad Politécnica de Madrid
sanz@disam.upm.es

December 15, 1999

Abstract

Intelligence is the exploitation of information to perform better. From this perspective I propose a definition of *autonomy* as a foundation for an integrated interpretation of existing perspectives of intelligent control system. This definition will be used to gather to a common ground different approaches to autonomy trying to combine some of their research topics under a common point of view. Some approaches to intelligent systems construction are evaluated and the sterile discussion on representation is addressed. The paper concludes with some comments on the future of the discipline of artificial intelligence in of real settings.

Keywords: Intelligent control, robotics, autonomous systems, intelligence, GOFAI, physically grounded systems, constructiveness.

1 Introduction

Using the words of James Thurber, I can say that *"The Anatomy of Confusion is a large subject, and I have no intention of writing the standard treatise on it, but I offer to whoever does, the most singular of all my cases, the Case of the Cockeyed Spaniard"* [Thurber 1953].

This paper is about a confusion regarding the source and nature of intelligent behavior. This confusion has emerged mainly in the area of behavior-based robotics [Brooks 1990, Arkin 1998, Brooks et al. 1998], in relation with the unnecessary and even pernicious use of representation-based and cognition-oriented approaches in the field of artificial intelligence (AI). This movement opposing mental models tries to debunk the classical AI approach, searching for new foundations for intelligent behavior.

Classical approaches to intelligence were based on representational models where a generic processor exploited the knowledge stored in a symbolic representation. Paradigmatic examples of this approach were generalized problem solvers as GPS, STRIPS or Soar [Laird and Newell 1993]. These knowledge representations were usually built using logic approaches like predicate calculus.

Some more recent approaches tend to diverge from this representational paradigm, based on the belief that they are inadequate as a basis for real intelligent behavior. This misconception is grounded in the demonstrated failures of classical AI when dealing with real settings. Robot planners in toy or controlled worlds worked well enough, but when used in real settings they demonstrated an absolute incompetence.

I use the word *misconception* when due to this failure to accomplish real tasks, researchers conclude that representation and reasoning can not be the basis of intelligence. They ground this interpretation in biological analogies that are a resemblance of Skinnerian behaviourism. Making an analogy, we can conclude that cargo ships are not a good basis for transport of merchandise because they cannot deliver a pizza to my home. The conclusion here is that generalized reasoners based on deep representations are not good for carrying pizzas, not that they are not good for transport.

Put in other words, the types of problems that artificial systems should solve and the type of solutions suitable for them are quite heterogeneous; both in transport systems and in intelligent systems.

Representation systems and reasoning mechanisms based on them are demonstrably very useful to achieve intelligent behavior. Behavior-based approaches, on the other side, are also extremely useful in other situations (for example to meet real-time requirements). As we will see in this paper, they are not ontologically different; behavior-based systems are only a particular case of representation based systems.

2 The nature of intelligence

It is well known that *intelligence* is considered an elusive term [Bellman 1978]. This is a broadly accepted belief, even when most people do agree on what is a proper use of it; i.e. we can conclude that there exist a generalized consensus –at least unconsciously– about what is a shared meaning of the term.

As with many other things the problem of intelligence is a problem of fuzziness. Paul is *very intelligent*, his dog is *quite intelligent* and his armchair is *no intelligent at all* . These are fuzzy terms. This is fuzzy terminology because intelligence is a matter of grades. The black/white consideration of intelligence is pure manichaeism. We say that a system is intelligent if behaves properly when confronted with a task.

IQ tests measure intelligence because *intelligence* is information processing put on act. We say that a system is intelligent if it can exploit what it knows to achieve better levels of performance. AI can be considered part engineering (solving problems) and partly science (understanding how can it be done) [Winston 1992].

This view of intelligence matches that classical sentence that says that *intelligence is what you use when you do not know what to do*, i.e. what you do beyond pure reflexes. Intelligence is the basis of high performance behavior, as clearly both classic AI and the new AI approaches believe.

Intelligence is a measure of the capability of the control subsystem of any behaving entity. This perspective identifies *intelligence* as the proper study of –obviously– artificial intelligence, but also of *cybernetics* [Wiener 1961], *automatic control* [Kuo 1991] or *robotics* [Poole 1989]. All these disciplines have been focused in specific areas of research, but, in the end, all these activities are always confronting the same problems and orbiting around the same solutions. They can, in fact, be considered one unique discipline: they are the sciences of the artificial behavior.

3 A concept of Autonomy

I've had some discussions and heard expressions of dismay related with the possibility of reaching a consensus about the meaning of the adjective *"autonomous"*. This is, however, strongly used in our area of work, mainly due to the inherent objective of artificiality of making it independent of human resources (like human hands or brains). I strongly believe it is time to clarify the use of the adjective, and I think it if possible to reach a consensus.

I propose to employ an interpretation in terms of a three place predicate: `Autonomous (SYSTEM, TASK, CONTEXT)`.

The meaning of this predicate is that the system `SYSTEM` is autonomous if it can fulfill the task `TASK` in the context `CONTEXT`.

I'm pretty sure that this proposed definition of the meaning of *autonomy* will not be accepted by all researchers. The objections will be mostly based on the type of systems researchers want to get out of it. An example is a refrigerator.

From my point of view, a refrigerator (`SYSTEM` = *"refrigerator of Alex"*) is *autonomous* because it can fulfill its task (`TASK` = *"keep the interior temperature at* 5^oC*"*) in a specific context (`CONTEXT` = *"Interior of a house in Philadelphia"*).

For many people a refrigerator (or a toaster, or a lightbulb) is not the best example of an autonomous system. They like to find in an autonomous system at least the level of intelligence of a cockroach, and a fridge cannot wander through the kitchen and escape when children appear. But, this is another discussion that is focused on the intelligence of the systems, and not in its autonomy (even when they are strongly related).

Fridges aren't of much interest because they are not quite intelligent. Cockroaches are intelligent, at least to some level that we would like to reproduce.

This confusion appears because the relation between intelligence and autonomy is very strong. What we want form our artificial creations is *autonomy* (we do not want to keep one eye on the fridge) and to achieve this autonomy, what we need to put on them is intelligence. Intelligence is information processing on act. So, our machines can use it to be better performing beings.

The three topics mentioned in the definition (system, task and context) compose the universe of an artificial being. Artificial systems are built with a purpose [Simon 1981] that is expressed in the form of a task to do. Capability of task fulfillment –autonomy– is affected by all three factors, and what we try to build are systems that are robust to uncertainty in all these aspects. This is a more common ground for researchers: autonomous systems are systems that are intelligent enough to solve problems that appear when they are trying to achieve an objective.

The three factors (`SYSTEM`, `TASK` and `CONTEXT`) form a dependence network that we handle in our search for autonomous behavior.

If we take the refrigerator to another –more harsh– *context*, say a desert with 60 Celsius degrees or my house with two small children, it is not easy at all for the refrigerator to fulfill the task of keeping the interior temperature at the setpoint. In the last case I can help the fridge accomplish the task by continuous human supervision (like human operators in automated power plants). The refrigerator cannot be considered autonomous in this setting. To regain autonomy we can modify the `SYSTEM` (add more refrigerating power or a door lock) or modify the `TASK` (for example change to *store food*).

If we change the *task* from *"keep the interior temperature at 5^{o}C"* to the apparently similar *"keep the food inside in good condition"* the refrigerator cannot be considered autonomous any longer (think about what will happen to oranges in some weeks or fish in some days). To achieve autonomy for this task we change the `SYSTEM` from a refrigerator to a freezer or change the `CONTEXT` putting the fridge in an sterile room.

Reconfiguration is an example of a way to handle uncertainty related with the *system* itself. If we provide the refrigerator with more intelligence, it will be able to detect a failure in the compressor and lock the door to reduce heat incoming and keep inner temperature.

When the task gets harder, the context uncertain or the body weak, the system can in many cases still fulfill the task if it is more intelligent. Autonomy is characterized by robust behavior under uncertainty and intelligence is a key component for it.

4 GOFAI and beyond

GOFAI (good old-fashioned artificial intelligence) was focused on hard `TASKS` because they carried the sense of the core meaning of intelligence. Focused on the –then considered– mechanisms of thought, they worked on abstract logical representations of entities, neglecting the importance of handling real `CONTEXTS`. They also neglected the importance of the embodiment of the intelligent `SYSTEM`, considered that abstract reasoning processes were independent of the underlying hardware (mechanical or biological).

They focused on abstract machines working on toy worlds and developed advanced mechanisms for task accomplishment in these worlds. All these techniques were based on abstract inference procedures that operated on representations of the world they were dealing with.

When people –like us, control engineers– tried to apply these *artificial intelligence* technologies in real settings, they mostly failed because the complexity of things in the real world is quite higher than in toy worlds.

They focused on real tasks, ignoring that the tasks should be done in real contexts by real bodies.

5 Brooksian robotics and more

The behavior-based robotics movement rejects the GOFAI approach arguing that it does not work in real settings. They are true. Behavior based robotics have demonstrated that there are alternative approaches to intelligent behavior that can achieve

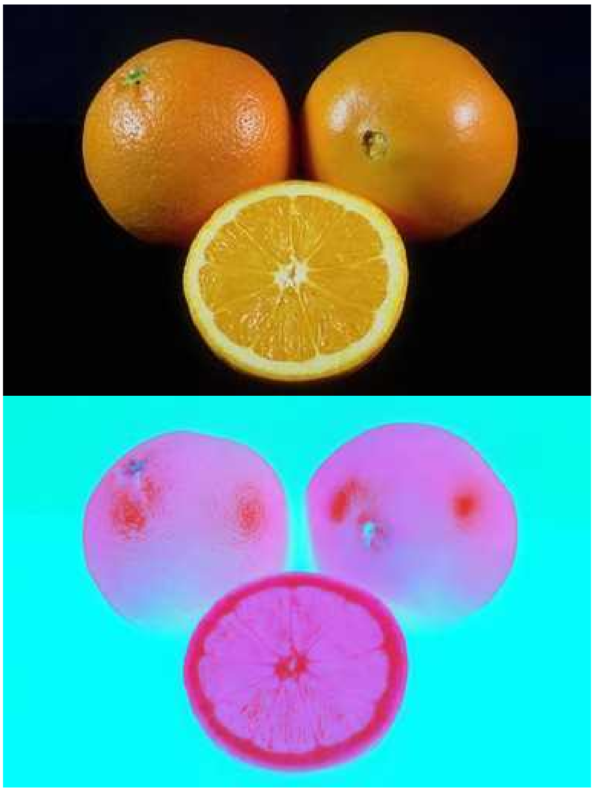

Figure 1: *Autonomy is a three legged concept: being, task and context. Research effort should be put into achieving unified results than can be exploited to achieve autonomy to do complex tasks in complex settings. Simple animals like insects can survive in varying environmental conditions because its genome provides them with evolution-time generated robustness to changing conditions. Faster adaptation to a changing environment can only be done using memes, and this means the exploitation of culturally transmitted representations. Behavior-based animals are condemned to extinction in a world of nuclear winter or genetically engineered new crops.*

results that are unthinkable using classical AI technology. Using their words, we can say that there are *alternative essences of the intelligence*. This praise of the Brooskian approach is not without counterpart. First of all it is not new. During decades, control engineering in real plant has been based on the use of finite state machines that control working equipment. In the area of mobile robots, in a classic work around 1950 -before the boom of AI-Walter developed a collection of light-seeking tortoises (using the technology of the fifties).

To better understand their claims let's quote some marketing propaganda of the last book of the new AI business [Brooks 1999]:

> *Until the mid-1980s, AI researchers assumed that an intelligent system doing high-level reasoning was necessary for the coupling of perception and action. In this traditional model, cognition mediates between perception and plans of action. Realizing that this core AI, as it was known, was illusory, Rodney A. Brooks turned the field of AI on its head by introducing the behavior-based approach to robotics. The cornerstone of behavior-based robotics is the realization that the coupling of perception and action gives rise to all the power of intelligence and that cognition is only in the eye of an observer. Behavior-based robotics has been the basis of successful applications in entertainment, service industries, agriculture, mining, and the home. It has given rise to both autonomous mobile robots and more recent humanoid robots such as Brooks'* Cog.

Behavior-based robotics propose a new approach, based on what they call grounded, embodied systems that map directly perception to action without mediating representations and reasoning. They say that they do not use internal models because the world is the best model of itself. Perhaps they do not make reasoning in the traditional sense of rule based systems, but for sure they do representation because perception is, in essence, a representation process. They argue that they do not have centralized representations, using purely distributed architectures; but the description of what they do to deal with maps (a pure representation) is very embarrassing [Maes 1990].

They are successful in cockroaching and other real tasks, but at the end, their approach is so scaled-down that it is hard to find a real application of this technology (cleaning television screens is a proposal of Brooks [Brooks 1997]).

They concentrate on real contexts forgetting about real problems, and if they say that picking empty soda cans is a real task they are not very ambitious for this technology.

I can understand that some people need publish papers to make a living; and perhaps, the best way to do this is being polemic and getting funds by building big toys that can grasp the interest of the general public.

Cog is pretty cute and Kismet can perhaps be your best friend, but if you need a robotic fireman what you need is a fire extinguisher with the intelligence enough to recognize the type of fire and assess the biological impact of the smoke in persons. Determining the type of smoke can be considered just a matter of specialized sensing, but imagine the network of finite state machines that you need to connect to reach this complex behavior, not to speak of the difficulties of guaranteeing a correct behavior [Butler and Finelli 1993].

6 A representationism vindication

The movement opposing mental models tries to debunk the classical AI approach, searching for new foundations for intelligent behavior. Debunking approaches are only reasonable if you are fighting for their place or their money. This is, from my point of view, what behavior-based robotics has tried to do. Now that behavior-based robotics is dead as a global foundation for intelligence [Moravec 1999], it is time to re-vindicate the role of models in complex behavior generation, because the image of model-based systems is somewhat deteriorated (that's the reason I have used the term *representationism* –i.e. it is an *ism*– in the title of this section).

Behavior-based engineering has suffered more or less the history of behaviourism in psychology. Grounded on lack of quality of the results of classic *consciousness-based* approaches it tried to reject any mental model foundation and change to an input-output psychology. Now its importance has been reconsidered due to the advances in the objective study of mental content.

This confrontation is pretty stupid to me, because it is, in essence, a discussion between input-output representations and state-based representations. Any systems engineer knows that both approaches are not opposing things but two sides of the same coin. A good engineer uses what fits better.

If the failure of model-based approaches was not in the method itself, where it was? After years of developing large, heterogeneous intelligent control systems, we have reached the conclusion that model based approaches have failed due to complexity.

This may sound now as a truism, but it was not so clear some years before. In 1992 we started a project called HINT, with the objective of integrate heterogeneous AI technologies in real-plant, industrial application [Alarcón et al. 1994]. A demonstration system was built to provide human operator assistance in controlling a dewaxing unit in a refinery. The system was based on the use of expert systems, fuzzy control technology, neural networks and model-based diagnosis.

Model-based diagnosers had demonstrated good results in the diagnosis of electronic equipment and so we thought that we were able to scale them to chemical processes. The conclusion was that after a lot of effort, the model-based diagnoser consumed most of the computational resources giving almost no result. A fuzzy system doing a similar task was more that a thousandfold effective diagnoser (with the obvious differences between model-based and rule-based diagnosis).

Model-based systems suffer from complexity in several ways, mainly combinatorial explosion and the lack of good models of complex systems.

Non representational systems -if they do exist- are of two classes: extremely simple technology or castrated technology. Extremely simple technology is what you do in elevators or bread toasters. Castrated technology is what you get when you can do it better and you don't. Representation is of key importance to intelligent behavior like that of complex controllers [Samad 1998b]. In fact, if we think for a minute, we should conclude that non-representational systems do not exist. Let's quote the very beginning of the PhD thesis of Larsson

[Larsson 1992]:

> *Carelessly, we humans may think that we are in contact with the world itself, when we see things and people and feel the wind on our faces. And all the time it is a matter of models, generated to explain and foresee our own impressions and perceptions. What would you say, Prince Hamlet, about the very nature of existence? "Model, models, models".*

But, to what extent is it true that non-representation systems (simple or castrated) do not employ representations? Sensing and perception are always representation processes. Transducers receive this name because they translate –represent– one physical magnitude into a representation of it using a different magnitude.

Obviously an intelligent being does not need a global, common, unified representation of the world, the task to do and its own capabilities to perform most tasks. The main criticism should be directed towards totalitarian views of representation (theories of everything) instead to representation itself. These are the views that centered research in classical AI, with the final result of being unable to use them for some tasks.

From the field of control systems engineering we can discover that distributed and centralized representations can coexist. Each one is used for a specific purpose, achieving performance levels suitable for real time use and for optimal behavior. Take for example a typical, modern controller for a distillation column in a refinery. Several levels of control are piled up to build the global controller, and they are based on a hierarchy of representations. At the lower levels, reactive systems are employed mapping sensor inputs to actuator outputs in the same type of basic structure of behavior-based systems. These reactive behaviors, however are codified forms of representations of the plant under control. Representations simple enough to be mapped to a controller that can be used in real time. These representations (models of the plant) can be obtained by engineers based on physical theories or can be obtained on-line by the control system itself. This process of perception and world modeling is termed system identification and is a clear example of increased intelligence levels in control systems.

Distributed representations are in use in any complex control system. In fact in industrial control we employ the term *distributed control system* (DCS) to refer to them. But this distribution do not preclude the existence of unified representations that let the intelligent control system perform activities that are not easy to do with distributed representations. Top level layers of the refinery distillation column control are based on unified, global representations of the column that are used to perform global optimization of the column. These representations are based on differential algebraic equations constituting what is the most advanced cognitive model of a real system.

Model-based control fail because the models used are not good. The main reason for this is that we are still lacking tools for model building. As Åström commented in relation with the origin of his work on system identification [Wittenmark and Rantzer 1999], models build based on physical theories sometimes fail and the effective approach is completely different: build a model based on experiencing the real system behavior. This is what is done in systems identification or in neural-network modeling.

For behavior-based systems thinking however, uncertainty is not the real problem, but the key problem is in trying to build a model of the world using uncertain information. Using his words [Brooks 1991]:

> *If there are no models built, the problem of uncertainty is inherently reduced. This alternative is to operate in a tight coupling with the world through a sensing-acting feedback loop.*

But after that he says:

> *Instead of relying on inaccurate values returned by noisy sensors, we can rely on the time averaged derivative of these signals as the creature actively changes its state within the world in a way which forces larger changes in the sensor readings than those contributed by noise.*

At the end result that even behavior-based systems are based on the same type of models used to control systems based on systems identification. The difference here is the lack of a systematic, sound method to obtain these models. They generate ad-hoc solutions expecting they will provide good results. It is widely accepted that one of the best ways of estimating the behavior of a system with uncertainty is to use Kalman filtering.

It is obvious that all this discussion on cognition-based or behavior-based is mostly a matter of taste, of personal preferences of the system builder. In the same sense that you can build systems based on hardware or software, there exist dualities in any aspect of system design. Classical tradeoffs are memory/speed, the mentioned hardware/software, data/procedure, function-based/behavior-based or distributed/centralized. For a comment on the relation between functions and intelligence see Section **??**.

Discussions about what is better are usually fruitless because it is pretty easy to find enough arguments to defend any position, and people tend to defend their positions due to our biological background. The proper approach is gathering what each alternative has to offer.

7 Complexity Limits

In a recent symposium, Doyle presented a paper [Doyle 1999] analyzing fundamental tradeoffs in robustness in complex systems. His analysis implies that, perhaps, there are fundamental laws that limit our capability to build arbitrary complexity systems with arbitrary robustness. It is just not a matter of putting more effort to guarantee an utterly robust behavior; this double objective simply cannot be achieved.

If we look at the many efforts of using intelligent control to get safer systems (i.e. systems that respond better to some undesired events), we can discover that they can achieve their objective with the cost of increasing the complexity of the system and sacrificing the robustness against other types of events.

This if the fundamental tradeoff Doyle refers to, but it is usually a non explicit decision what leads to sacrifice some for of robustness in favor of other. Robustness issues should be make explicit during the construction of the system, and the way to do this is to handle inherent uncertainties that appear in the triad system-task-context. The only effective way to be able to perform this analysis is a model-

Figure 2: *Systems intelligence can be increased by the use of advanced information processing, but these leads to increased complexity and, as Doyle points out, perhaps there are fundamental laws that limit the robustness of our complex designs. Formula 1 authorities have banned most active control technology in what used to be the most advanced road vehicles. Now they are second to commercial vehicles in control technology.*

based unifying approach to uncertainity and related risk management in complex controller engineering.

Quoting [Doyle 1999, p. 261] we can say that:

> *While it is quite natural to distinguish between parametric uncertainty, noise and unmodeled dynamics, it is also important to treat them in an unified way.*

8 Conclusions: The future of intelligent control

Shakey is sometimes considered the paradigm of the failure of GOFAI and the success of behavior based-robotics. The accepted explanation is because the construction of a model of its world and the reasoning about it took too much time for the task the robot was supposed to do.

This failure, however is only in the eye of the failure-interested beholder. What Shakey demonstrated was not only that it was able to navigate -perhaps worse than a behavior based robot- but that it was possible to extract information from the world that could serve for any other purpose, *with so few amount of computing power*. This information-for-all idea is perhaps what Shakey designers had in mind (even if they didn't verbalize this). Models of the world can be used not only to navigate avoiding collisions, but to calculate the number of objects in the room, and the length of a corridor (for example to calculate the amount of paint needed to paint it, or cable to wire it). Think about extracting this information from a behavior-based robot.

Systems intelligence levels are raised by the proper integration of subsystems. While decentralization and downsizing is a way to real-time behavior, more complex intelligent behavior can be achieved by the proper integration of subsystems. Proper integration means not only action or control integration but knowledge sharing and interchange between subsystems. Without this knolwedge interchange, social behavior (the global behavior of the collection of subsystems) can only present limited structural forms.

If something is the hallmark of intelligent behavior is the capability of exploit information. This information should be exploited to enhance the autonomy of the systems we build handling uncertainty in the three aspects: *system*, *task* and *context*. If we analyze the present status of the technology this is what is already being done: intelligent control focus is handling of system uncertainty (for example in fault-tolerant control [Izadi-Zamanabadi 1999]), of environment uncertainty (for example in rough terrain navigation[**?**]) and of task uncertainty (for example handling natural language orders [**?**]).

Some effort should be put however, in creating a common viewpoint to this uncertainty management. The construction of robust autonomous systems will lead –inevitably– to complex organizations with emerging behaviors and we will have the need of incorporating into our discipline all these emerging technologies of complexity.

Generalized complexity will be the main engineering problem for future intelligent controllers [Samad 1998a], and systems to tackle generalized complexity will have the need of integrating heterogeneous control mechanisms. For example [Ali and Goel 1996] shows an example of an integrated approach, with a multistrategy qualitative navigational planner and a reactive-control mechanism.

The final conclusion is that we should start a search for a unifying theory of artificial behavior, somewhat in the line Wiener proposed in his *Cybernetics* [Wiener 1961].

And what about the elephants? Paraphrasing the title of the, now classical, paper of Brooks, we can say that not only *elephants don't play chess*, they *cannot play chess*. We need other types of intelligence for our complex controllers than elephant intelligence.

References

Alarcón, M., Rodrí guez, P., Almeida, L., Sanz, R., Fontaine, L., Gómez, P., Alamán, X., Nordin, P., Bejder, H., and de Pablo, E. (1994). Heterogeneous integration architecture for intelligent control. *Intelligent Systems Engineering*.

Ali, K. S. and Goel, A. K. (1996). Combining navigational planning and reactive control. In *Proceedings AAAI 96*.

Arkin, R. C. (1998). *Behavior-Based Robotics*. Intelligent Robots and Autonomous Agents. MIT Press, Cambridge, MA.

Bellman, R. (1978). *Artificial intelligence: Can Computers Think?* Boyd and Frasier Publishing Co., Boston, MA.

Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6:3–15.

Brooks, R. A. (1991). How to build complete creatures rather than isolated cognitive simulators. In VanLehn, K., editor, *Architectures for Intelligence*, pages 225–239. Lawrence Erlbaum Associates, Hillsdale, NJ.

Brooks, R. A. (1997). From earwigs to humans. *Robotics And Autonomous Systems*, 20(2-4):291–304.

Brooks, R. A. (1999). *Cambrian Intelligence*. MIT Press, Cambridge, MA.

Brooks, R. A., Breazeal, C., Irie, R., Kkemp, C. C., Marjanovic, M., Scasellati, B., and Williamson, M. M. (1998). Alternative essences of intelligence. In *Proceedings of AAAI 98*.

Butler, R. W. and Finelli, G. B. (1993). The infeasibility of quantifying the reliability of life-critical real-time software. *IEEE Transactions on Software Engineering*, 19(1):3–12.

Doyle, J. (1999). Fundamental tradeoffs in robustness of complex systems. In Wittenmark, B. and Rantzer, A., editors, *The Åström Symposium on Control*, pages 241–261. Studentlitteratur, Lund, Sweden.

Izadi-Zamanabadi, R. (1999). *Fault-tolerant Supervisory Control-System Analysis and Logic Design*. PhD thesis, Department of Control Engineering, Aalborg University.

Kuo, B. (1991). *Automatic Control Systems*. Prentice-Hall, Upper Saddle River, NJ.

Laird, P. R. J. and Newell, A., editors (1993). *The Soar Papers: Research on Integrated Intelligence*. MIT Press, Cambridge, MA.

Larsson, J. E. (1992). *Knowledge-Based Methods for Control Systems*. PhD thesis, Lund Institute of Technology, Lund, Sweden.

Maes, P., editor (1990). *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. MIT Press, Cambridge, MA.

Moravec, H. P. (1999). *Robot: Mere Machine to Trascendent Mind*. Oxford University Press, Oxford-New York.

Poole, H. (1989). *Fundamentals of Robotics*. Van Nostrand-Reinhold, New York.

Samad, T. (1998a). Complexity management: Multidisciplinary perspectives on automation and control. Technical Report CON-R98-001, Honeywell Technology Center, Minneapolis, MI.

Samad, T. (1998b). Multi-modelling: Knowledge integration in complex systems. In Samad, T., editor, *Complexity Management: Multidisciplinary Perspectives on Automation and Control*, pages 55–64. Honeywell Technology Center.

Simon, H. A. (1981). *The Sciences of the Artificial*. MIT Press, Cambridge, MA, second edition.

Thurber, J. (1953). *Thurber Country*. Hamish Hamilton, London.

Wiener, N. E. (1961). *Cybernetics, or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA, second edition.

Winston, P. H. (1992). *Artificial Intelligence*. Addison-Wesley, Reading, MA.

Wittenmark, B. and Rantzer, A., editors (1999). *The Åström Symposium on Control*. Studentlitteratur, Lund, Sweden.