


**IEEE COMPUTER SOCIETY
PRESS REPRINT**

**THE PERCEIVING ROBOT: WHAT DOES IT SEE?
WHAT DOES IT DO?**

**Oliver G. Selfridge
Judy A. Franklin**

Reprinted from PROCEEDINGS OF THE 5th IEEE INTERNATIONAL
SYMPOSIUM ON INTELLIGENT CONTROL 1990,
Philadelphia, Pennsylvania, September 5-7, 1990



IEEE Computer Society
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264


Washington, DC • Los Alamitos • Brussels • Tokyo



THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.



IEEE COMPUTER SOCIETY



**The Perceiving Robot: What Does It See?
What Does It Do?**

by

Oliver G. Selfridge and Judy A. Franklin



The Perceiving Robot: What Does It See? What Does It Do?

Oliver G. Selfridge & Judy A. Franklin

GTE Laboratories

We examine the nature of robots in the future, and propose that their role is fundamentally to be *responsible agents* for people, and not mere programmed artifacts. That means that besides extended powers of perception, they will need to deal with their own purposes—embedded in *purpose structures*—and the ways to modify and optimize their purposes in parallel.

The primary purpose of robotic perception is to see how well the robot is performing on a current task (or subtask).

Being responsible means being able to take care of more than one task at once: for example, a *cleaning robot* must be able not only to clean, but also to do it efficiently, and to conserve power and itself; and also to be able to modify what it tries to do, in order to take care of new requirements. That means that the robot must be able to modify its own understanding of what clean means in various circumstances.

This viewpoint of robotics extends the current approaches towards a new vision.

1 Introduction—The Perceiving Robot as a Responsible Agent

As the robot becomes more intelligent and more self-conscious, a throng of other questions emerges: how does its perception of itself change? How the robot perceives its interaction with and relationship to the environment will determine how it acts. The robot perceives in several modes and for several reasons. We hold that it should be expressing purposes and acting as a responsible agent—responsible to us. How can it handle responsibility? We need terminology in order to communicate these ideas and develop the perceptual robot.

What does a robot need to perceive and how can what it is trying to do be matched to its perception of both itself and its world? We consider this question with the idea that a robot is more than a mere mechanism; we want to enable it to be an assistant to us. We see the need to extend the underlying model of what a robot is and what it does by incorporating in it a *purpose structure*, which is a set of extendable and modifiable purposes.

The point is that for a robot to be a useful and powerful assistant it must also be a *responsible agent*. Here we mean the term *responsible* to imply that the robot must be in a position to make judgments, and more importantly, to have to exercise choices in what it does by, *inter alia*, adjusting itself to meet changing conditions and to respond to changing demands from us. That is to say, we want our robot to learn from experience. Otherwise it is a mere algorithm. This is not the place to explore the philosophical ramifications of that conclusion, and we merely accept it.

To be responsible, it seems to us, is to accept the concerns that are assigned, and to integrate them in such a way that they lead to behavior that balances those concerns in the right way. That is not merely circular argument, for it emphasizes that the concerns—the tasks, perhaps—are not spelled out in absolute detail.

To be specific about purposes and responsibility, let us consider Asimov's famous Three Laws of Robotics [Asimov-72]:

- I. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.
- II. A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.
- III. A robot must protect its own existence, as long as such protection does not conflict with the First or Second Law.

Those rules carry with them a burden of capabilities that are worth examining: How is a robot to perceive a human being? Does humanity have a concrete definition? Is a *human being* merely an object that claims that it is one? Nowadays, industrial robots have big cages around them. How will the robot know about injury? How can it learn about injury without injuring? Or do we believe that we can define *human being-ness*, *injury*, and *harm*?

We believe that the only way a robot can handle such generalities, and in real world contexts, is by learning from its own experience what they mean.

A responsible agent must be perceptive, have purpose, and be able to learn. These capabilities result in emergent properties that characterize intelligence. We discuss these capabilities, the idea of emergent properties, and present some visions of the possibilities.

2 Perception

Such a robot will need a broad repertoire of sensory modalities and capabilities. In the long tale of life, sensory capabilities of great complexity arose long before complex actions were achievable. Only a few animals use tools, but even insects and lower animals have extremely complex sensory filters and recognizers. These processing capabilities were not produced by following specifications, but in response to the changing requirements of evolution—even though we do not know exactly what evolution is or how it works. Can we really specify now how to recognize a person?

Beyond the problems of recognition lie the morasses of prediction and self-awareness—"except where such order would conflict ... " and "must protect its own existence ... " Those we shall not consider here in detail, but it is important to note the complexity imposed by a set of three "simple" laws.

Most robots are conceived and constructed these days with very specific sensors that are supposed to perform very specific sensing; these specifications are considered part of the design process. Certainly it is important to talk about sensory experiences. It is not clear whether we should cast them in terms of the types of things they measure, such as change, time, properties of pieces of the environment, or in terms of why they are making the measurements. Or a (possibly hierarchical) combination of them—including types and functions—all as alluded to by Meystel [Meystel-88].

A big question indeed is how sensorial data (perceptual experiences) relate to its thought processes and goal analyses. What is the organization of robotic perception? Perception is related to capabilities and these are part of the robot state. The robot needs to know the state of the world and the state of itself, well enough to do its jobs. These states must be perceived with respect to the robot's purposes.

A more general design should be fundamentally concerned with being *improvable*, and we would expect the robot itself to do the improving, though perhaps under our guidance or training. In any event, the robot itself must continue to adapt and improve its perception of its sensory data while it is working for us. Consider a cleaning robot. Should it be trained to pick up and discard any object it visually experiences that it finds? Or perhaps it should be trained—or instructed—to *discard waste*, and to *keep the floor clean*, but to *sort metal waste* and to *save metal parts*; so that it then has to focus on what it perceives. It may not be easy or possible to prescribe what waste is, or what metal is, or how it should be detected; rather we propose that the robot uses its senses in a learning way to improve its detection.

A robot should be given the tools to organize its perception. How much should be hardwired? The trade-offs between learning vs. hardwiring will depend on what the robot is being used for, our own capabilities in designing it, and the risks associated with its learning.

Franklin [Franklin-88] proposes the idea of refining the controller of a single link robot arm, through learning. The idea is that the robot is given some hardwired capability that allows it to function and that it can then improve those functions to perform its tasks better. These ideas can be extended to more esoteric risks and imposed limitations. The robot should also be able to *extend and modify its functions so as to be able to perform modified tasks or new ones*. For example, giving a robot the tools to recognize the face of a person may take away its ability to modify its perception of that person. But giving it sufficient tools with the capacity to perceive and change may enable both recognition and the necessary modifications of perception.

A robot is often perceived as being designed and built for a given fixed task, but consider an example of a robot changing jobs. How many jobs are built in? How many can be learned? How much can or should a human be in the loop to change the purposes and the tools of the robot? Must the "actions" that it is to try be hardwired? Or can it learn the kinds of powers it has to try?

The robot perceives the results of its doings. In other words it has mechanisms for feedback at many levels of abstraction. Albus well understands the requirement that a complex system must be (at least) hierarchically controlled:

"At this point the problem of controlling a sensory-interactive robot becomes similar to that of controlling any complex system such as an army, a government, a business, or a biological organism. Introducing the type of hierarchical command and control structure that has historically proven successful in controlling such systems then becomes necessary ... [Albus-81, p. 261]

3 Purposes

Current notions of "purpose" in artificial intelligence treat a purpose as a goal state: which is a particular instance of the universe as described by certain specified variables. Implicit in this approach is the assumption that a robot is a finite state automaton, and that what it knows and perceives can be properly and conveniently represented by finite state variables. Some modern research is trying to avoid the limitations inherent in such attitudes:

"Autonomous mobile robots ... are experimental testbeds for research in intelligent machines. In accomplishing given objectives, they manage their resources and maintain their integrity ... " [Weisbin-89, p. 29]

Just as with the laws of robotics, the researchers suppose that the objectives can be precisely defined, as can also managing resources and integrity. They note that

“... precise positioning of the obstacles is unlikely because of the intrinsic uncertainties and resolution limits of the sensors ...” [ibid, p. 30]

And they turn, as we do, to machine learning, to aid in “World modeling and path planning.” But the emphasis is on a very limited set of choices, like a so-called Blocks World.

We envision a different tack. There are purposes behind the tasks that the robot takes on, and the robot must have a broad range of choices in which to experiment and act in order to achieve those purposes. How are purposes to be represented? This leads to questions about the relationship of higher level goals to lower level goals and parallels in perception [Albus-81, Fu-70, Meystel-88, and Saridis-80].

A high level purpose must be able to generate goals and subgoals, or have them generated for it. Those sub-purposes themselves will have to be optimized, to take care of the changing environment that the high level purpose is working in.

Consider a collie herding sheep. This intelligent animal acts as though it knows its purpose is to herd the sheep through the gate. Clearly it has been trained to achieve this goal. But it does seem unlikely to us that it has any kind of manipulable model of that purpose; that is to say, a robot does not necessarily have to have models of its purposes.

This brings us at once to the need for modification and learning of purposes. That is all because of the changing world and because of our ignorance in knowing how to express our desired purposes for the robot exactly, and how to instruct it precisely on how to carry them out. How does a robot learn purposes by itself?

Well, there are three related possibilities:

- A purpose can be supplied by, or imposed by somebody else, like a person. Example: *robot, when you pick up paper trash, save the green rectangular pieces, for they might be dollar bills. Robot, that's not a dollar bill, it's too small, but this is.*
- A purpose, or subpurpose, can be derived from new purposes. Example: if the robot must separate dollar bills, it must learn to perceive them.
- A purpose can be invented by the robot by modifying an old purpose. If old techniques are not working then modifications of old purposes should be tried to see if they can.

The perception of the results of a robot's actions must be *with respect to* its purposes. It must be able to perceive the *differences* that result from its trying. The robot will not be a Markov process. In other words, it must have a memory that will keep track of what has been tried and what happened as a result. It is important that the robot be able to try something again, especially in light of the fact that the world does change and that its actions may produce very different results at different times.

Current robots are limited in their perceptual capabilities. They do not perform learning. We would like the robot to notice what it does, how it does it, and to make associations with the “rightness” or “wrongness” of the results and its actions. At the lower levels these are associations of senses with respect to *tasks* being performed and how well the robot performed the tasks.

Franklin and Selfridge [Franklin-90] propose a direction of work in adaptive control in which the robot is given a far wider concept of the purpose of control, beyond mere stability; that concept includes multiple and changing goals. For example, a responsible robot does more than this. It “worries” about other things at the same time that it performs its primary tasks.

Consider a robot running (part of) a production line. It may be required sometimes to optimize output rate, and sometimes to minimize breakage, while maintaining some minimum rate of processing. Moreover, it must be able to undertake dynamic planning without knowing the control parameters or even their range ahead of time. It may learn appropriate planning behavior at the same time as it learns about these characteristics of the production line. It must perceive the effects of control actions on the state of the line with respect to its current goals. It must be able to change its subgoals efficiently and learn to plan its actions accordingly.

An important concept is that the robot learns what to attend to, i.e. about focus of attention and how to direct it. Herman *et al.* [Herman-90] have done some work on the *coordinating* of sensors, using a sensor to determine where to focus another sensor.

One practice that a robot must have is to try new things. That is in order that it may learn to optimize its components outside the previously prescribed limits. Optimizing at one level is learning to place a lower level at the appropriate state. Moreover, many of the dealing with purposes in the previous section required that the robot learn to form and improve new purposes. Learning is the expression of purpose.

Suppose our cleaning robot has been instructed or trained to sweep dust and to pick up trash in the machine tool room, separating it all into metal and non-metal components. Now suppose that the room becomes infested with beetles, and that such a possibility had not been anticipated in the design or

implementation. Our robot should then first recognize that the previous techniques are not being especially effective; second (say) notice that attention needs to be paid to picking up the moving trash—which was previously not a problem—; and third, improve and optimize the new techniques needed for that. Notice that how to do that is not specified; it may have to try several techniques; such as rolling over (treading on) the beetles, or being provided with an insecticide.

5 The Concept of Emergent Properties

Anyone who has already experimented with primitive robotics has discovered unexpected behaviors because of combinations of unanticipated sensory readings and the robot program's reaction to them. It is not easy to see that the mere activity of a single feedback loop can be interpreted as active learning, but it is clear that deep hierarchies of them can be most easily understood by considering what they do as learning.

For an example, consider a cruise control system for an automobile: the basic system must have its gain set at an appropriate level, and the changes in that must be set as well; furthermore, the integration time of the velocity changes must be set appropriately—all these are part of the control. The easiest way to look at that is to consider the system as learning how to do cruise control appropriately. We regard that kind of learning as *emerging* from the hierarchy of the control loops.

Just as learning is an emergent property of feedback loops, so responsibility and self-consciousness are emergent properties of a complex perceptual robot. They emerge from the requirements of the robot interacting with other constraints and with human supervisors.

Must curiosity develop naturally (with respect to something)? How much curiosity is an infant born with and how abstract is it? In fact, the naive robot should be equipped with curiosity about itself and how it can move and act. At this level curiosity may be just the capability, referred to earlier, of the robot to try new things. This curiosity will evolve into abstract curiosity about its interaction with the environment. This is another *emergent* property.

In fact, we believe that curiosity, by now innate in human beings—and in many other animals as well—arises as a natural advantage in the kind of task domain that life is. This is not to say that it may not be profitable to find a way to program it in, so to speak. Note again that the particular actions undertaken by the robot need not be especially complex, but the sensory interpretations have to be rich and powerful enough to express and remember the appropriate discoveries about the environment.

Free will itself may be viewed as an eventually emergent property of the kinds of robotic systems that we are concerned with here, but we confess that we have very little notion of what it can mean beyond our own self-consciousness, or of how we can tell when they are exhibiting it. In fact, it may just be our interpretation of their complex behavior.

6 The Vision of Robotics

In the late seventies and early eighties, the field of robotics consisted mainly of studies of kinematics and the dynamics of a multi-link robot arm [Paul-81]. Robot control was in the form of position control, often using independent PID joint controllers. Then research began in force control and in dextrous robot hands [Salisbury-80, Salisbury-82, Venkataraman-88]. Computer vision and motion had been studied in parallel [Lawton-84]. More "exotic" robotic research began to appear such as mobile robots [Arkin-87] and sensor fusion [Herman-90]. Also in parallel is the work being done on teleoperation [NASA-89] that involves man-machine systems, telerobot architectures, navigation, planning, planetary rovers, multi-arm control and robotic vision and tactile sensation.

Learning has been introduced in at least two areas, that of learning control [Franklin-89] and that of learning simple navigation tasks for mobile robots [Mataric-90]. Ignoring the separate work in artificial intelligence and machine learning, the current state of robotics is probably reflected by the technical committees of the Robotics and Automation Society of the IEEE:

- 1 Computer-Aided Production Management
- 2 Future Directions
- 3 Intelligent Instrumentation & Measurement Systems
- 4 Robotics Motion Planning
- 5 Robotics Dynamics & Control
- 6 Computer Vision
- 7 Manufacturing Automation
- 8 Micro Robot
- 9 International

Having a vision of what a robot *might be* may not provide immediate suggestions about the next steps to take now. In artificial intelligence, the temptation is apparently to select advanced human tasks that use complex and intellectually stimulating thought processes. We believe that that is a trap. Although we too are discussing futuristic behavior, we believe that we should begin with simple tasks that reside in the real world. Just like Rossum's Universal Robots [Capek-23], robots should start with menial tasks and duties, even if they are self-aware.

At a higher level, for example, imagine a robot that could look at you with a "glint" in its "eye," point its finger and say "I'm as smart as you!" Well, a robot could be made to do that today, but it would neither understand what it was saying, nor have a reason for saying it. Part of an actual glint is the knowledge of the object it is glinting at and the feedback it receives in the reaction of the glint recipient. Also, why would a robot point its finger like that? As a robot is now envisioned, it would point with a physically-oriented task in mind, such as picking up one part with its finger to assemble with another part. A robot might be envisioned as a gesticulating being, but there is a long way to go before gesticulations are an implicit part of "robotic nature."

The cleaning robot was mentioned in section 2 will have to embody [*sic!*] some of these solutions to the complexity of the real world. Our cleaning robot is more blue-collar than managerial. Yet it is clearly not adequate merely to specify a number of "cleaning actions," and program the robot to perform them. Rather, the robot's goal must be set to have the floor clean, and the selection of the actions performed, and in what way, must be left up to the robot. For that purpose, the cleanliness of the floor cannot be the sole factor for judgment. For it must also be trying to expend as little energy as possible, to not get in people's way; and how clean is clean enough?—a clean room for building silicon chips has very different cleanliness needs from a machine shop. That is, the robot must not only perceive what is trash or dirt, it must also know what standards are going to be applied, and perceive how well it is meeting them. On top of that, the robot must know how to change those standards, or to have them changed for it, so that it can be subsequently tasked to take care of a clean room for chip manufacture.

For another more complex example, we may need a robot that assists elderly and handicapped people in airports. This free-roaming robot must perceive the need for its assistance and be able to interact with the person by asking questions and responding to their questions. It may need to use gestures that they understand and as aids in expressing itself. It may need to be able to move with the person and guide them through a twisting morass of travellers (without harming the travellers). It should carry their bags. It might need an expressive face and certainly a nonmonotonic voice. An important point is that it must be expressive enough for the human to know *its* desires and purposes.

But such interpretations today are necessarily anthropomorphic. We researchers ought to be paying attention to our own model of the robot as the technology grows towards learning and self-awareness. For the natural language and feelings of a person are not likely to be mirrored soon in robots.

What does this mean for perception? It implies a wide variety of ways in which sensor inputs are interpreted and it implies many types of "outputs" and actions. Importantly, it implies that the robot should be able to understand the desires and needs of the people it is assisting. It must be responsible in interacting with the human. That means that it must inherently have some degrees of learning from its experiences.

We ought not to believe that we can vault into a robotics that can at once provide versatility, multiple dynamic task satisfaction, self-awareness, learning, and free will. Rather we should focus on simple learning, useful mental skills, so that we ourselves can learn how to assign responsibility. Then we can begin the real education of the robot.

Bibliography

- [Albus-81]Albus, J.S., 1981, *Brains, Behavior, and Robotics*, BYTE Publications, Inc., Peterborough, New Hampshire.
- [Arbib-72]Arbib, M.A., 1972, *The Metaphorical Brain. An Introduction to Cybernetics as Artificial Intelligence*, Wiley-Interscience, New York.
- [Arkin-87]Arkin, R.C., 1987, "Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior," *Proceedings of the 1987 IEEE International Conference on Robotics and Automation*, pp. 264-271.
- [Asimov-82]Asimov, I., "Mirror Image," first appeared in *Analog: Science Fiction - Science Fact*, May 1972, as quoted in *The Complete Robot*, Nightfall, Inc. U.S.A.
- [Capek-23]Capek, K., *R.U.R. (Rossum's Universal Robots)*, drama published in Czech in 1921, first translated into English in 1923.
- [Franklin-88]Franklin, J.A., "Refinement of Robot Motor Skills Through Reinforcement Learning," in *Proceedings of the 27th IEEE Conference on Decision and Control*, Austin, Texas, December, 1988.
- [Franklin-89]Franklin, J.A., "Historical Perspective and State of the Art in Connectionist Learning Control," *Proceedings of the 28th IEEE Conference on Decision and Control*, December, 1989, Tampa, Florida.
- [Franklin-90]Franklin, J.A. and Selfridge, O.G., 1990, "Some New Directions for Adaptive Control Theory in Robotics," in *Neural Networks for Control*, Miller, W.T., III, Sutton, R.S., and Werbos, P., eds., MIT Press, in press.
- [Fu-70]Fu, K-S, 1970, "Learning Control Systems—Review and Outlook," *IEEE Trans. Auto. Control*, April 1970, AC-15, 2.
- [Herman-90]Herman, M., Albus, J.S, and Hong, T.-H., 1990, "Intelligent Control for Multiple Autonomous Undersea Vehicles," in *Neural Networks for Control*, Miller, W.T., III, Sutton, R.S., and Werbos, P., eds., MIT Press, in press.
- [Iberall-84]Iberall, Y. and Lyons, D., "Towards Perceptual Robotics," in *Proceedings of the 1984 IEEE International Conference on Systems, Man, and Cybernetics*, Halifax, Nova Scotia, Oct. 9-12, pp. 147-157.
- [Lawton-84]Lawton, D.T., 1984, "Processing Dynamic Image Sequences From a Moving Sensor," *COINS Technical Report 84-05*, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts.

[Mataric-90]Mataric, M.J., 1990, *A Model for Distributed Mobile Robot Environment Learning and Navigation*, M.S. Dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts.

[Meystel-88]Meystel, A., 1988, "Intelligent Control in Robotics," *Journal of Robotic Systems*, 5(4), pp. 269-308.

[NASA-89]*Proceedings of the NASA Conference on Space Telerobotics*, JPL Publication 89-7, Rodriguez, G. & Seraji, H., eds., January 31, 1989, JPL, California Institute of Technology, Pasadena, California.

[Paul-81]Paul, R.P., 1981, *Robot Manipulators: Mathematics, Programming, and Control*, MIT Press, Cambridge, Massachusetts.

[Salisbury-80]Salisbury, J.K., 1980, "Active Stiffness Control of a Manipulator in Cartesian Coordinates," *19th IEEE Conference on Decision and Control*, December 1980, Albuquerque, New Mexico.

[Salisbury-82]Salisbury, J.K., *Kinematic and Force Analysis of Articulated Hands*, Ph. D. Dissertation, July 1982, Department of Computer Science, Stanford University, Stanford, California.

[Saridis-89]Saridis, G.N., 1989, "On the Revised Theory of Intelligent Machines," Dept. of Computer and Systems Engineering, CIRSS Doc 23, Rensselaer Poly. Inst., Troy New York.

[Venkataraman-88]Venkataraman, S.T., *Task Dependent Dextrous Hand Control*, Ph. D. Dissertation, May 1988, Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, Massachusetts.

[Weisbin-89]Weisbin, C.R., de Saussure, G., Einstein, J.R., & Pin, F.G., "Autonomous Mobile Robot Navigation and Learning," *Computer*, IEEE Computer Society, June 1989.

The first part of the document discusses the importance of maintaining accurate records of all transactions. It emphasizes that every entry should be supported by a valid receipt or invoice. This not only helps in tracking expenses but also ensures compliance with tax regulations. The document further outlines the steps for recording these transactions, from identifying the nature of the expense to the final posting to the ledger.

Next, the document addresses the process of reconciling bank statements. It explains how to compare the bank's records with the company's books to identify any discrepancies. Common reasons for these differences include timing differences, such as deposits in transit or outstanding checks. The document provides a detailed guide on how to investigate these differences and adjust the books accordingly.

The third section focuses on the preparation of financial statements. It details the steps involved in calculating net income, preparing the balance sheet, and the income statement. The document stresses the importance of accuracy in these statements as they provide a clear picture of the company's financial health. It also discusses the role of these statements in decision-making and reporting to stakeholders.

Finally, the document concludes with a summary of the key points discussed. It reiterates the importance of diligent record-keeping and regular reconciliation to ensure the integrity of the financial data. The document serves as a comprehensive guide for anyone responsible for the financial management of a business.

