Running head: Decision-making in robotics and psychology

Decision-making in Robotics and Psychology: A Distributed Account

Benoit Hardy-Vallée

Department of philosophy, University of Toronto

Corresponding author:

Benoit Hardy-Vallée
Posdoctoral Fellow
Department of Philosophy
University of Toronto
Jackman Humanities Building
170 St. George St., 4th Floor,
Toronto, ON, M5R 2M8
Phone: (416) 946-5982
Fax: (416) 978-8703
Academic email: ben.hardy.vallee@utoronto.ca
Homepage: http://www.hardyvallee.net

Abstract

Abstract- Decision-making is usually a secondary topic in psychology, relegated to the last chapters of textbooks. The psychological study of decision-making assumes a certain conception of its nature and mechanisms that has been shown wrong by research in robotics. Robotics indicates that decision-making is not—or at least not only—an intellectual task, but also a process of dynamic behavioral control, mediated by embodied and situated sensorimotor interaction. The implications of this conception for psychology are discussed.

Keywords: decision-making, rationality, robotics, psychology, cognition

PsycINFO classification: 2630 Philosophy 4140 Robotics

6565 words (without references)

1. Introduction

Psychology can be roughly divided into eight branches: abnormal, behavioral, cognitive, comparative, developmental, neuro-, personality and social psychology (Colman, 1999). Psychology thus studies *brains*, *behaviors* (in humans and other animals) and *psychological processes* such as information processing, learning and development. Each branch, however, tends to neglect—in one way or another—an important subject: decision-making. Decision-making figures in psychology textbooks, but most of the time chapters devoted to this topic are limited to a presentation of the *Homo economicus* model, an acknowledgement of its failure and a list of heuristics and biases that explain human irrationality. In a recent article, H.A. Gintis documents this neglect of decision-making:

> *(...) a widely used text of graduate- level readings in cognitive psychology, (Sternberg & Wagner, 1999) devotes the* ninth *of eleven chapters to "Reasoning, Judgment, and Decision Making," offering two papers, the first of which shows that human subjects generally fail simple logical inference tasks, and the second shows that human subjects are irrationally swayed by the way a problem is verbally "framed" by the experimenter. A leading undergraduate cognitive psychology text (Goldstein, 2005) placed "Reasoning and Decision Making" the* last *of twelve chapters. This includes one paragraph describing the rational actor model, followed by many pages purporting to explain why it is wrong. (...) in a leading behavioral psychology text (Mazur, 2002), choice is covered in the* last *of fourteen chapters, and is limited to a review of the literature on choice between concurrent reinforcement schedules and the capacity to defer gratification (Gintis, 2007, pp. 1-2)*

The standard conception of decision-making in psychology, I suggest, consist in two claims, one methodological, one empirical. The methodological one holds that decision-making is a *separate* topic: it is *one* of the subjects that psychologists may study, together with categorization, inference, perception, emotion, personality, etc. As Gintis showed, decision-making has its own chapters (usually the last ones) in psychology textbooks. On the empirical side, the standard conception construes decision-making as an explicit deliberative process, such as reasoning: "our inner deliberations, argues Billig (1996, p. 35), are silent arguments conducted within a single self".

The standard account implicitly assumes that making choices is similar to playing chess: an activity we engage in occasionally and voluntarily. It therefore deserves a separate chapter since, like chess, it is not a fundamental psychological phenomenon. Moreover, decision-making is characterized as an intellectual faculty. For instance, in a special edition of *Cognition* on decision-making (volume 49, issues 1-2, Pages 1-187), one finds the following assertions:

> *Reasoning and decision making are high-level cognitive skills [...]*(Johnson-Laird & Shafir, 1993, p. 1)
>
> *Decisions . . . are often reached by focusing on reasons that justify the selection of one option over another* (Shafir *et al.*, 1993, p. 34)

By 'intellectual' or 'high-level', psychologists imply that making a decision depend upon cognitive mechanisms distinct from perception, action of emotion. For instance, in Eysenck et al.'s textbook (2005), the chapter *Judgement and Decision-making* figures in the "Thinking and Reasoning" category, along with working memory and language comprehension.

In this article, I first show how the standard account of decision-making in psychology is embedded is a complex framework of folk-psychological, philosophical, economic and psychological concepts. I then suggest that research in robotics provides significant evidence that some assumptions concerning the cognitive architectures of deciding agents—as the standard conceptual framework depict them—are flawed. The development of control architectures in robotics, I argue, recommends another conception of decision-making. Finally, I show how this conception is supported in other domains and discuss its theoretical and practical implications for psychological science.

2. Decision-Making: philosophy, economics and psychology

Rational agents display their rationality mainly in making decisions. Certain decisions are more basic (turn left or turn right), others are crucial issues (“to be or not to be”). In any case, being an agent entails making choices. Our folk-psychology inclines us to believe that making a decision implies a deliberation and that this process is inherently tied to human characteristics like language, free will and complex mental representations. When subjects are asked to explain or predict actions, their intuitions lead them to describe observed actions as the product beliefs, desires and intentions (Malle & Knobe, 1997, 2001). The folk-psychological model of rationality construes decision-making as the outcome of a practical reasoning by which an agent infers, from her beliefs and desires, the right action to do.

An important duty of philosophy if to make things explicit, that is, to show what justifies (or fail to justify) assertions we implicitly take for granted (Brandom, 1994, 2000). Conceptual analysis, as it is usually practiced, unravel the semantic norms that make certain inferences valid and revise them if necessary. Hence philosophy of mind classically analyzed concepts such as consciousness, beliefs, motivation, etc. The standard philosophical conception of decision-making equates *deciding* and *forming an intention before an action* (e.g., Davidson, 1980, 2004; Hall, 1978; Searle, 2001). According to different analyses, this intention can be equivalent to, inferred from or accompanied by, desires and beliefs. Beliefs and desires are constitutive of rationality because they *justify* decisions: there is a logical coherence between beliefs, desires and intentions. Actions are irrational when their causes do not justify them. Beliefs and desires are thus embedded in our interpretations of rational agents *as* rational agents: “[a]nyone who superimposes the longitudes of desire and the latitudes of belief is already attributing rationality" (Sorensen, 2004, p. 291). Philosophy of mind therefore made explicit the inferential structure of folk-psychology. Theoretical economics and rational-choice theory formalize the abstract structure of this conceptual scheme.

Economics, according to a standard definition by Lionel Robbins, is the "science which studies human behavior as a relationship between ends and scarce means which have alternative uses" (Robbins, 1932, p. 15). This definition shows the centrality of decision-making in economic science: since means are scarce, behavior should use them efficiently. The two branches of rational-choice theory, decision theory and game theory, specify the formal constraints on optimal decision-making in individual and interactive contexts. An individual agent facing a choice between two actions can make a rational decision if she takes into account two parameters: the probability and utility of the consequences of each action. By multiplying the subjective probability by the subjective utility of an action's outcomes, she can select the actions that have the higher subjective expected utility (see Baron, 2000, for an introduction). Game theory models agents making decisions in a strategic context, where the preferences of at least another agent must be taken into account. Decision-making is represented as the selection of a strategy in a *game*, that is, a set of rules that dictates the range of possible actions and the payoffs of any conjunct of actions. Thus, economic decision-making is mainly about computing probabilities and utilities (Weirich, 2004). The folk-psychological account and the philosopher's beliefs-desire model are hence reflected in the economist's probability-utility model: probabilities represent beliefs while utilities represent desires. Note that nothing in game or decision theory recommends this interpretation: it became standard in rational-choice theory to map propositional attitudes onto numeric values. This is coherent with the intellectual conception of decision-making.

Rational-choice theory can be construed as a normative theory (what agents *should* do) or as a descriptive one (what agents do). On its descriptive construal, rational-choice theory is a framework for building predictive models of choice behavior: which lottery an agent would select, whether an agent would cooperate or not in a prisoner's dilemma, etc. Experimental economics, behavioral economics, cognitive science and psychology (I will refer to these empirical approaches as 'psychology') use these models to study how subjects make decisions and which mechanisms they rely on for choosing. These patterns of inference and behavior can then be compared with rational-choice theory. In numerous studies, Amos Tversky and Daniel Kahneman showed that subjects deviate markedly from normative theories (Kahneman, 2003; Kahneman *et al.*, 1982; Tversky, 1975). They tend to make decisions according to their 'framing' of a situation (the way they represent the situation, e.g. as a gain or as a loss), and exhibit loss-, risk- and ambiguity-aversion (Camerer, 2000; Kahneman & Tversky, 1979, 1991, 2000; Thaler, 1980). In most of their experiments, Tversky and Kahneman asked subjects to choose among different options in fictive situations in order to assess the similarity between natural ways of thinking and normative decision theory. For instance, subjects were presented the following situation (Tversky & Kahneman, 1981):

Imagine that the United States is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

- If Program A is adopted, 200 people will be saved

- If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved.

Which of the two programs would you favor?

Most of the respondent opted for A, the risk-averse solution. When respondents were offered the following version:

- If Program A is adopted, 400 people will die

- If Program B is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die

Although Program A has exactly the same outcome in both versions (400 people die, 200 will be saved), in the second version Program B is the most popular. Thus subjects' risk-attitude depends on the framing of the situation: they are risk-prone in loss, and risk-averse in gain. The study of decision-making is thus the study of the heuristics and biases that impinge upon human judgment. The main explanatory target is the discrepancies between rational-choice theory and human psychology. Just like the psychology of perception explains visual illusions (e.g., the Muller-Lyer illusion), the psychology of decision explains cognitive illusions: why agents prefer systematically prospect X to Y when rational-choice theory recommends Y. Loss-aversion, for instance, can be explained by the shape of the value function: it is concave for gains and convex for losses. Thus losing $100 hurts more than winning $100 makes one happy.

Proponent of the ecological rationality approach suggested nonetheless that these heuristics and bias might be adaptive in certain contexts and that failures of human rationality can be lessen in proper ecological conditions (Gigerenzer, 1991; Gigerenzer *et al.*, 1999). For instance, when probabilities are presented as frequencies (6 out of 10) instead of subjective probabilities (60%), performance of subjects tend to be much better, partly because we encounter more sequences of events than degrees of beliefs. Heuristics might be 'fast and frugal' procedures tailored for certain tasks that lead to suboptimal outcomes in other contexts. Or they could be vestigial adaptations to ecological and social environments where our hunters-gatherers ancestors lived. Thus heuristics may not completely ineffective.

Whether heuristics are effective or not, the methodology for deciphering human decision-making capacities is the same in both case. The problems are presented in linguistic forms, and it is supposed that subjects would represent mentally the whole situation before taking action. From a list of possible actions coupled with the probability and utility of these actions' consequences, it is assumed that an optimal decision-maker would rank possible actions according to their subjective expected utility and choose the best one. Without much surprise, psychologists discovered that humans are not very good at following normative decision theory, and proposed alternative models (Tversky and Kahneman's Prospect Theory, for instance). Yet, the main assumptions of these experiments are rarely, if ever, discussed: why should decision-making be represented as an internal process of symbolic inference? Why questionnaires are considered as reliable tools for exploring decisions? Why linguistic inferences should be the medium of decision-making? All these assumptions support an *intellectual* conception of decision-making.

Since decision-making is a reasoning process, it is studied mostly by multiple-choice tests using the traditional paper and pen method. Psychological research assumes that the subjects' competence in probabilistic reasoning—as revealed by these tests—mirrors their decision-making capacities. The standard conception of decision-making is thus an empirical hypothesis about the nature of decision-making and a methodological norm for its study.

This conception also has implications outside psychology: if one wants to assess the empirical plausibility of this conception, the efforts of roboticists to build autonomous agents constitute an important source of evidence (or lack thereof), since robots *have* to make decision. Mataric defines robot control as

> *the process of taking information about the environment, through the robot's sensors, processing it as necessary in order to make decisions about how to act, and then executing those actions in the environment. (Mataric, 2002, p. 25)*

The implementation of a cognitive architecture in a robot is a partial test for the psychological theory behind it. Epistemologically, robots are *models*, that is, a simplified representation of a complex phenomenon generated from the assumptions of a theory. However, as I will argue in the next section, the design of control architectures in robotics suggests that this model is flawed. “Classical” robotics—roughly, from its inception to the 90’s; see Nilsson, 1984, Turk et al, 1988—endorsed a cogitative conception of decision-making that led to control architectures of a limited success: they were ineffective in many real-world settings and lack adaptability. “New” robotics—research and development that followed Rodney Brooks’ revolutionary suggestions such as perception-action coupling and behavior-based robotics (see Mataric,1997a; Michaud & Mataric,1998; Pfeifer & Scheier, 1999)—showed how decision emerges from the dynamic coordination of multiple sensorimotor modules. Instead of adding more knowledge “in the head”, New Robotics suggested that intelligence is not a matter of internal representations and complex deductions, but a matter of action and interaction.

Although it is not a complete refutation of the standard conception of decision-making, it supports the exploration of alternative models.

3. The Standard Conception and Its Flaws

Classical robotics adopted a cognitive model analogous to our intuitive representation of decision-making:

> *Perception is commonly cast as a process by which we receive information from the world. Cognition then comprises intelligent processes defined over some inner rendition of such information. Intentional action is glossed as the carrying out of commands that constitute the output of a cogitative, central system.* (Clark, 1997, p. 51)

This *sense-model-plan-act* (SMPA) account of cognition and decision-making is explicit in many control architectures. In these architectures, (e.g. Nilsson, 1984), decision-making is a logical problem. Sensors or cameras represent the whole perceptible environment while internal processors convert sensory inputs in first-order predicate calculus. From this explicit model of its environment, the robot's central planner transforms a symbolic description of the world into a sequence of actions (see Hu & Brady, 1996, for a survey). Decision-making is taken in charge by an expert system or a similar deductive engine. Hence intelligence lies in the planning step. Take for instance the autonomous land vehicle Alvin (Turk *et al.*, 1988). Alvin was designed for autonomous outdoor navigation, on and off-road. It comes equipped with five modules: sensors (laser sensor and a CCD camera), vision (description of the road), reasoning (building trajectory), knowledge base (a priori road map) and a pilot (motor control).

INSERT FIGURE 1 HERE

Making a decision such as turning left vs. turning right, in Alvin's case, means that it must solve the following problem: which sequence of operators enable me to reach some target state from the current one? Once the problem is solved, actions are merely motor implementations of the solution. Research and development in robotics, however, suggests that decision-making should not be construed as a deliberative process, for at least three reasons: 1) SMPA control architectures lack robustness 2) adaptive control is better achieved by distributed architectures 3) SMPA architectures fail to deal efficiently with uncertain and unreliable information.

3.1 Robustness and control

Mataric (2000) distinguishes four classes of robotic control architectures:

1. Reactive control ("don't think, react")
2. Deliberative control ("think, then act")
3. Hybrid control ("think and act independently in parallel")
4. Behavior-based control ("think the way you act").

Reactive architectures are constituted of collections of finite state machines that "fire" when a particular stimulus (their target, or *affordances*) is present. In a mobile robot, for instance, one module may specialize in avoiding obstacle while another achieves the goal of reaching some home destination. These architectures can be effective in simple tasks (e.g. garbage collecting), but perform poorly when learning and memory are required. Although Brook's initial proposal of "intelligence without representation" caused a dramatic paradigm shift in robotics, roboticists soon realized that minimal representations are mandatory for fluent interactions. Smart reflexes can be useful in stationary environment, but sometimes the robots need to model some aspect of its environment or its body (Mac Dorman, 1999, p. 21):

> *To deal effectively with new situations a robot needs to model its affordances so that it can test its actions against a model before testing them against the world. In this way, the robot doesn't have to jump off a cliff before discovering that this is dangerous; it can recognize the affordance and let its hypothesis about moving toward the cliff action die in its place (...).*

Too much representation, however, is not optimal. Deliberative control architectures (Mataric, 1997a), based on the SMPA model also failed to achieve robust and adaptive control (Brooks, 1999; Pfeifer & Scheier, 1999). They could be effective, but only in environment carefully designed for the robot. The colors, lightning and objects disposition were optimally configured for simplifying perception and movement:

> *The walls were of a uniform color and carefully lighted, with dark rubber baseboards, making clear boundaries with the lighter colored floor. (...) The blocks and wedges were painted different colors on different planar surfaces. (....) Blocks and wedges were*

> *relatively rare in the environment, eliminating problems due to partial obscurations*
> (Brooks, 1999, p. 62)

Thus, outside its engineered environment, the robot would be unable to cope with simple task such as finding a route or avoiding an obstacle. Just like industrial robots, these robots were competent only in pre-specified, limited environments. In uncertain, dangerous or non-stationary environments, they would not survive. A complete, explicit model of the environment, manipulated by a central planner, is not useful for robots. Re-*action* depends upon re-*modeling*: if anything unexpected happen, they must sense, model, and plan *before* acting.

Brooks (1999, pp. 107-109) gives a simple example that illustrates the failure of the central planner assumption: manipulator arms. The classical way to control robotic arms is known as *position control*. In these schemes, control is based on an explicit model of the environment and a representation of the arm's desired position. From the actual position and the goal position, the planner devises a series of motor commands that allow the arm to move from the former to the latter. Interacting with the world requires extremely precise models and sensors: since the only source of information about what to do is the model and the discrepancies between the model and the world, the representational and sensory devices should be both complete and reliable. Much time and energy is spent in updating the model and sensing. The reliance on the model, however, comes at another cost. When manipulator arms grasp their payloads, the grasping substantially changes the arm's dynamics, making the model inaccurate. In order to avoid these supplementary discrepancies, designers started to build more massive robots carrying smaller loads: large arms are 'insensible' to small loads and thus their dynamics is less affected by variations induced by grasping. The manipulator-to-payload ratio is usually between 100:1 and 1000:1, which contrasts neatly with any animal or human arm. Recent control architectures, however, implement force control or hybrid force/position control instead of pure position control. The idea is to control the balance of all the forces (gravity, friction, pressure, acceleration), not only the position. Instead of representing the object to be picked, the system is guided by feedback mechanisms that signal when the appropriate force is applied, without concern for the position. Consequently, robotic control is efficient and effective when the system is dynamically coupled with its environment through sensorimotor interaction and can readjust its behavior in real time.

Reactive and deliberative control architectures have been shown to be highly limited. Hybrid control architectures perform well, but they inherit the problems of reactive and deliberative (Mataric, 2000). The greatest advance in autonomous robotics is behavior-based control ("think the way you act").

3.2 Distributed control

In classical architecture, a central planner infers the next move from a set of “if-then” propositions, a complex world model and sensory data. Perception and action are separated from this control faculty. To the contrary, as many roboticists found, the best way to achieve effective control is to build a modular architecture, where many processes interact together in order to produce behavior. Each module, also called “behavior”, is a *control law* (in the control-theoretic sense):

> *each can take inputs from the robot's sensors (for example, cameras and ultrasound, infra-red or tactile sensors) and/or from other behaviors, and send outputs to the robot's effectors (such as wheels, grippers, arms or speech) and/or to other behaviors in the system. Consequently, a behavior-based robot is controlled by a structured network of interacting behaviors*. (Mataric, 1998, p. 82)

The modules are not perceptive, cognitive and motor faculty, but complete input-output routes, endowed with learning and mnemonic capacities. Each module is simple, fast and relies on simple cues. The architecture thus manipulate multiple, partial models of the environment and control the robot's behavior without a central planner or an explicit model.

Control systems of autonomous robots are partitioned in various subsystems, each having a tacit knowledge of a target field (walk, vision, other agents, etc., see Dorigo & Colombetti, 1998; Mataric, 1997b). Modules are also hierarchically assembled in layers so that modules in higher-level layers can override modules in lower-level layers (Brooks's "subsumption architecture", see Bonasso *et al.*, 1997; Kortenkamp *et al.*, 1998). In order to design robots able to imitate people, for instance, roboticists build systems that control their behavior through multiple modules. Mataric's (2002) robots rely on the following modules:

1) a selective attentional mechanisms that extract salient visual information (other agent's face, for instance)

2) a sensorimotor mapping system that transforms visual input in motor program

3) a repertoire of motor primitives

4) a classification-based learning mechanism that learns from visuo-motor mapping

Robots need to learn mappings between their appearance and that of others, their behavior and that of others, their bodily behaviors and their internal states, and between the external states of others and their one own internal states (see also Barsalou et al, in press). There are neither complete representations of the external world nor a central cognitive module, but a complex coordination of sensorimotor modules. Hence there is no need to integrate models in a unified representations or a common code: distributed architectures, were many processes runs in parallel, achieve better results.

In distributed architectures, intelligence does not lie in the planning, but in the coordinating. Mataric's robots learn to forage efficiently their environment by adjusting different modules, such as avoidance, safe-wandering, following, aggregation, dispersion, and homing

(1995). Robots start by dispersing. Once they obtain food, they return to their initial location (homing). If they encounter another robot, they can be informed (through radio signals) whether the other agent also carries food. They could then follow it, and eventually foraging as a flock if enough robots found food in the same area. Thus individual and group foraging emerges out of the coordination between building blocks (avoiding, following, etc.). The series of choices these agents had to make were not the product of a detached faculty, but of the coordination of many sensorimotor faculties.

3.3 Coping with uncertainty

Finally, SMPA architectures suppose that information about possible courses of action is already available and reliable: but what if information is absent, incomplete, uncertain, or if the environment is rapidly changing? Classical Robotics was plagued with *encodingism*, the assumption that internal symbols stand for something else outside the system and that cognitive processes are equivalent to predefined symbols manipulation (Bickhard & Terveen, 1995). It neglected the acquisition , grounding and updating of representations. Classical robotics represented all decision-making situations as contexts of *first-order uncertainty*, that is, situations where the agents know the probabilities. In many situations however, uncertainty can be *radical*, or *second-order uncertainty*: even the probabilities are unknown. Agents in this setting must, at the same time, explore their environment in order to gather information about its payoff structure and exploit this information to obtain reward. They face an important problem—known as the exploration/exploitation trade-off—because they cannot do both at the same time: one cannot explore all the time, one cannot exploit all the time and exploration must be reduced but cannot be eliminated. A well-known example of this trade-off is *K-armed bandit problem*.

Suppose an agent has *n* coins to spend in a slot machine with K arms (here K=2 and we will suppose that one arm is high-paying and the other low-paying, although the agent does not know that). The only way the agent has access to the arms' rate of payment—and obtains reward—is by pulling them. Hence she must find an optimal trade-off when spending its coins: trying another arm just to see how it pays or staying with the one who already paid? The goal is not only to maximize reward, but also to maximize reward *while obtaining information about the arm's rate*. Two types of error may appear: false negative (a low-paying sequence of the high-paying arm) or false positive (a high-paying sequence of the low-paying paying arm).

To solve this problem, the optimal solution is to compute an index for every arm, updating this index according to the arm's payoff and choosing the arm that has the greater index (Gittins, 1989). In the long run, this strategy amounts to following decision theory after a learning phase. But as soon as switching from one arm to another has a cost, as Banks & Sundaram (1994) showed, the index strategies cannot converge towards an optimal solution. A huge literature in optimization theory, economics, management and machine learning addresses this problem (Kaelbling *et al.*, 1996; Sundaram, 2003; Tackseung, 2004). These researches look for the normative, optimal policy to adopt in K-armed bandit problems. Roboticists are interested not only in finding optimal policies, but finding also the optimal control architecture that will allow robot to behave efficiently in these situations. Again, deliberative architectures similar to the SMPA model are not optimal. Robots must take decisions every second, for everything (Kimura *et al.*, 2001; Urmson *et al.*, 2003; Volpe *et al.*, 2001). Looking left or right, moving, grasping, exploring, etc., are different actions among which robotic architectures must select. There is no logical solution to decision problems under certainty. In order to solve the exploration/exploitation trade-off, it is better to learn and guess, something distributed architectures excel at (Michaud & Mataric, 1998).

Thus the deliberative conception of decision-making, and its SMPA implementations, had to be abandoned. If it did not work for mobile robots, it is justified to argue that for cognitive agents in general the standard conception also has to be abandoned. In other words, If robots are models of a psychological theory of decision-making, they showed that the standard model is wrong and hence that the theory could also be faulty. In the following section I discuss how these findings could support another psychological approach of decision-making.

4. Psychology and the Distributed Conception of Decision-Making

> *conventional wisdom has long modeled our internal cognitive processes, quite wrongly, as just an inner version of the public arguments and justifications that we learn, as children, to construct and evaluate in the social space of the dinner table and the marketplace. Those social activities are of vital importance to our collective commerce, both social and intellectual, but they are an evolutionary novelty, unreflected in the brain's basic modes of decision-making* (Churchland, 2006, p. 31).

We may conjecture that psychology, like classical robotics, espoused a folk-psychological conception of the mind and uses psychological intuitions to build models of decision-making. This naïve theory, like many naïve theories, may be the source of numerous misunderstandings (Churchland & Churchland, 1998). For instance, if folkbiology would be the sole foundation of biology, whales would still be categorized as fish. The nature of the biological world is not explained by our (faulty and biased) folkbiology (Medin & Atran, 1999), but by making explicit the mechanism of natural selection, reproduction, cellullar growth, etc. There is no reason to believe that our folk-psychology is a better description of mental mechanisms.

Concepts borrowed from New Robotics should pave the way for a paradigm shift in the psychological study and theorizing of decision-making. Research in developmental (Terry & Nigel, 2003; Weng *et al.*, 2001), evolutionary (Nolfi, 1997) and multi-agents (Mataric, 2001) robotics all suggest that behavioral control, action selection and decision-making emerge from the tight coupling of a distributed control architecture with its environment. There is no decision-maker inside the robots: the situated robot as a whole *is* the decision-maker.

This suggests a “Distributed Conception of Decision-making” (DCDM), according to which decision-making is:

> *Embodied*: the mechanisms for decision-making are not only and not necessarily intellectual, high-level and explicit. Decision-making is the whole organism’s sensorimotor control.
>
> *Situated*: a decision is not a step-by-step internal computation, but also a continuous and dynamic adjustment between the agent and its environment that develop in the whole lifespan. Decision-making is always physically and (most of the time) socially situated: ecological situatedness is both a constraint on, and a set of informational resources that help agents to cope with, decision-making.

Psychology should do more than documenting our (in)ability to follow Bayesian reasoning in paper-and-pen experiment, but study our embodied and situated control capacities. In making a decision, an agent allocates sensorimotor resources to the evaluation, anticipation and selection of possible courses of action and their consequences. The ideal explanation of subjects’s performance should therefore be a flow-chart type modules and their interactions.

The DCDM is also supported by research in other domains. Neuroscience, and especially neuroeconomics—the merging of the neuroscientific study of decision-making with psychology and experimental economics—supports the *embodied* aspect of the DCDM (Glimcher, 2003; McCabe, 2005; Zak, 2004). There is no brain area, circuit or mechanisms specialized in decision-making, but rather a collection of neural modules. While certain area specializes in visual-saccadic decision-making (Platt & Glimcher, 1999), imaging studies revealed that individual and social decisions depend upon many different valuations mechanisms. Certain mechanisms compute the cost/benefit of outcomes, other encode the expected hedonic value, and other codes for preferences over actions (Rushworth et al, 2007). Social neuroeconomics indicates that decision in experimental games are mainly affective computations: choice behavior in these games is reliably correlated to neural activations of social emotions such as the 'warm glow' of cooperation (Rilling *et al.*, 2002), the 'sweet taste' of revenge (de Quervain *et al.*, 2004) or the 'moral disgust' of unfairness (Sanfey et al., 2003). Subjects without affective experiences or affective anticipations are unable to make rational decisions, as Damasio and his colleagues discovered. Damasio found that subjects with lesions in the ventromedial prefrontal cortex had problems in coping with everyday tasks (Damasio, 1994). They were unable to plan meetings; they lose their money, family or social status. They were, however, completely functional in reasoning or problem-solving task. Moreover, Damasio and its collaborators found that these subjects had lower affective reactions. They did not feel sad for their situation, even if they perfectly understood what "sad" means, and seemed unable to learn from bad experiences. The researchers concluded that these subjects were unable to use emotions to aid in decision-making, a hypothesis that also implies that in normal subjects, emotions *do* aid in decision-making.

Concerning the *situated* dimensions of decision-making, recent experiments requiring subjects to control dynamical systems via human—computer interfaces (e.g. Chhabra & Jacobs, 2006) indicate that when subjects do not have to make their cognitive process explicit or to reason about them, they are able to follow an optimal control policy (as computed by dynamical programming). Subjects embedded in a sensorimotor dynamics were not asked to make explicit their policy, only to *learn* one. Behavioral ecology (Krebs & Davies, 1997; Krebs & Kacelnik, 1991) also suggests that models of decision-making should incorporate *ecological* parameters. Optimal foraging theory, for instance, (Giraldeau & Caraco, 2000; Stephens & Krebs, 1986) models agents, their choices, the currency to be maximized (most of the time a caloric gain) and a set of constraints (e.g. time, energy, etc.) Most researches study where to forage (patch choice), what to forage (prey choice) and for how long (optimal time allocation). It is supposed that the individual animal makes a series of decisions in order to solve a problem of sequential optimization. An animal looking for nutrients must maximize its caloric intake while taking into account those spent in seeking and capturing its prey; to this problem one must also add, among others, the frequency of prey encounter, the time devoted to research and the calories each prey type afford. All these parameters can be represented by a set of equations from which numerical methods such as dynamic programming allow biologists to derive algorithms that an optimal forager would implement in order to optimize the objective function. These algorithms are used afterward for the prediction of the behavior, and meta-analysis showed their predictive efficacy (Sih & Christensen, 2001). Similar models are also predictive of human behavior. Human behavioral ecology (Smith, 1991; Winterhalder & Smith, 2000) applies the same bio-economic logic with the same success to humans. Agents are modeled as optimal forager subjects to a multitude of constraints. Given available resources in the environment of a community, one can generates a predictive model of optimal resources allocation. These models are of course more complex than animal models since they integrate social parameters like local habits, technology or economic structures, but human behavioral ecology is more predictive than traditional microeconomics. Models of human foraging were able for instance to explain differences in foraging style between tribes in the Amazonia, given the distance to be traversed and the technology used (Hames & Vickers, 1982). Food sharing, labor division between men and women, agricultural cultures and even Internet browsing (where the commodity is information) can be modeled and aptly described by human behavioral ecology (Jochim, 1988; Kaplan *et al.*, 1984; Pirolli & Card, 1999). Hence, formally speaking, optimal foraging theory is the translation of decision theory axioms into tractable calories-maximization algorithms that incorporate situated and embodied constraints.

In practical terms, it means that studies based on paper-and-pen questionnaires should not be construed as informative about our decision-making *mechanisms*. It is illuminating, of course, about our *performances* in contexts where we explicitly manipulate information in probabilistic format, as in management or policy-making, for instance. But these intellectual faculties are the higher-level layers of our control architecture, not the whole of it.

I presented earlier two assumptions held by proponents of the standard conception of decision-making: that decision-making is a *separate* topic (methodological claim) and that decision-making is an explicit *deliberative* process (empirical claim). I take research in robotics to support the idea that the empirical claim ran into many problems that invalidate it. This undermining of the standard conception has empirical implications (e.g. its rejection and the adoption of the distributed account) but also methodological ones. The most important, I think, is that decision-making should not be a secondary topic for psychology but, following Gintis "the central organizing principle of psychology" (Gintis, 2007, p. 1). Since robotic decision-making is realized by the coordination of many modules, there is no deciding module. Decision-making is distributed in different modules and cannot be studied in isolation from other competence. Given the success of distributed architecture it is justified to revise the standard conception of decision-making and to study human decision-making as a concert of modules. In doing so, decision-making becomes the central topic of psychology, something that all chapters of psychology textbooks refer to. There might be a chapter on Bayesian reasoning because we sometimes engage in probabilistic reasoning, but the equivalence between that and decision-making should be abandoned. Abstract reasoning is *one* of the cognitive tools we have for making decision. Hence deciding should not be studied like a separate topic (e.g. perception), an occasional activity (e.g. chess-playing) or a high-level competence (e.g. logical inference), but like robotic control. Many process, mechanisms, modules and layers interact so that agents behave adaptively, prefer certain actions or outcomes and discard others.

Making decisions, and especially efficient ones, is a matter of survival. Even a bacteria like *E. Coli* is a decision-maker. With its rudimentary sensors and limited memory, the bacterium is able to avoid noxious substances, seek bioenergetic resources and even cooperate with other bacteria (Koshland, 1977). Although it does not entertain beliefs or desires, its behavior cannot be reduced to a simple action-reaction scheme. Research in nanoengineering showed that *E. Coli*'s behavor can be described as "real-time nonlinear optimization, robust adaptation and control" (Lyshevski, 2003, p. 690). Thus some internal computational mechanisms make the individual bacterium choosing to go up or down, left or right, based on its informational states.

Therefore, psychology should be the science of the mechanisms (normal and abnormal), development, individual and cultural variations, and neural implementation of decision-making in humans and animals. Perception, emotions, memory categorization, reasoning and consciousness are competence that aid and optimize decision-making. If decision-making is the central organizing principle of psychology, all the branches of psychology could be understood as research field that investigates different aspects of decision-making. Abnormal psychology explains how deficient mechanisms impair decision-making. Behavioral psychology focuses on choice behavior and behavioral regularities. Cognitive psychology describes the mechanisms and modules involves in valuation, goal representation, preferences and how they contribute to decision-making. Comparative psychology analyzes the variations in neural, behavioral and cognitive processes among different clades. Developmental psychology establishes the evolution of decision-making mechanisms in the lifespan. Neuropsychology identifies the neural substrates of these mechanisms. Personality psychology explains interindividual variations in decision-making, our various decision-making “profiles”. Social psychology can shed light on social decision-making, that is, either collective decision-making (when groups or institutions make decisions) or individual decision-making in context. Finally, we could also add environmental psychology (how agents use their environment to simplify their decisions) and evolutionary psychology (how decision-making mechanisms are—or are not—adaptations). Since the human psyche expresses itself through our decisions, a psychological science should be principally concerned with decisions.

5. Conclusion

Human life is one long decision tree. (Sterelny, 2006)

Decision-making is an important concept both for the study of natural agents (psychology) and the design of artificial agents (robotics). Agents have to choose among course of actions, outcomes, situations, object, etc. Intuitively, we take decision-making to be an intellectual process, a variety of theoretical reasoning with practical content. This picture, however, is of no help for building autonomous robots. Sense-model-plan-act architectures are effective only in a designed environment. Autonomous robotics shows that the key to autonomy, adaptivity and intelligence is not central planning, but the coordination of sensorimotor modules. As I suggested here, there is two lessons for psychology, one empirical, one methodological. Empirically, the standard conception of decision-making should be regarded as dubious: decision-making is a distributed, emergent competence, not the activity of “decision module”. Methodologically, since decision-making is not an isolated faculty, it should not be studied in isolation and relegated to the last chapters of psychology textbooks. Too often, the psychological study of probabilistic reasoning is confused with the study of decision-making: the former is an abstract, conscious and explicit variety of the latter.

Moreover, decision-making, as Gintis argued, should be the central organizing principle of psychology. All life—robotic or natural—is about decision; all cognitive faculties and mechanisms are aids for decision-making.

References

Banks, J. S., & Sundaram, R. K. (1994). Switching costs and the gittins index. *Econometrica: Journal of the Econometric Society, 62*(3), 687-694.

Baron, J. (2000). *Thinking and deciding* (3rd ed.). Cambridge, UK ; New York: Cambridge University Press.

Barsalou, L., Breazeal, C., & Smith, L. (*in press*). Cognition as coordinated non-cognition. *Cognitive Processing*.

Bickhard, M. H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science : Impasse and solution*. Amsterdam ; New York: Elsevier.

Billig, M. (1996). *Arguing and thinking : A rhetorical approach to social psychology* (New ed.). Cambridge ; New York, NY, USA: Cambridge University Press.

Bonasso, R. P., Firby, R. J., Gat, E., Kortenkamp, D., Miller, D. P., & Slack, M. G. (1997). Experiences with an architecture for intelligent, reactive agents. *Journal of Experimental & Theoretical Artificial Intelligence, 9*(2-3), 237-256.

Brandom, R. (1994). *Making it explicit : Reasoning, representing, and discursive commitment*. Cambridge, Mass.: Harvard University Press.

Brandom, R. (2000). *Articulating reasons : An introduction to inferentialism*. Cambridge, Mass. ; London: Harvard University Press.

Brooks, R. A. (1999). *Cambrian intelligence : The early history of the new AI*. Cambridge, Mass.: MIT Press.

Camerer, C. (2000). Prospect theory in the wild. In D. Kahneman & A. Tversky (Eds.), *Choice, values, and frames* (pp. 288-300). New York: Cambridge University Press.

Chase, V. M., Hertwig, R., & Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Science, 2*(6), 206-214.

Chhabra, M., & Jacobs, R. A. (2006). Near-optimal human adaptive control across different noise environments. *J. Neurosci., 26*(42), 10883-10887.

Churchland, P. M., & Churchland, P. S. (1998). *On the contrary : Critical essays, 1987-1997*. Cambridge, Mass.: MIT Press.

Churchland, P. M. (2006). Into the brain: Where philosophy should go from here. *Topoi, 25*(1), 29-32.

Clark, A. (1997). *Being there : Putting brain, body, and world together again*. Cambridge, Mass.: MIT Press.

Colman, A. M. (1999). *What is psychology?* (3rd ed.). London ; New York: Routledge.

Damasio, A. R. (1994). *Descartes' error : Emotion, reason, and the human brain*. New York: Putnam.

Davidson, D. (1980). *Essays on actions and events*. Oxford: Oxford University Press.

Davidson, D. (2004). *Problems of rationality*. Oxford: Oxford University Press.

de Quervain, D. J., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., et al. (2004). The neural basis of altruistic punishment. *Science, 305*(5688), 1254-1258.

Dennett, D. C. (1987). *The intentional stance*. Cambridge, Mass.: MIT Press.

Dorigo, M., & Colombetti, M. (1998). *Robot shaping : An experiment in behavior engineering*. Cambridge, Mass.: MIT Press.

Eysenck, M. W., & Keane, M. T. (2005). *Cognitive psychology : A student's handbook* (5th ed.). New York: Psychology Press.

Gigerenzer, G. (1991). How to make cognitive illusions disappear: Beyond heuristics and biases. *European Review of Social Psychology, 2*(S 83), 115.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 62–88). Oxford: Blackwell.

Gigerenzer, G., Todd, P. M., & ABC Research Group. (1999). *Simple heuristics that make us smart*. New

York: Oxford University Press.

Gintis, H. (2007). A framework for the integration of the behavioral sciences (with open commentaries and author's response). *Behavioral and Brain Sciences, 30*, 1-61.

Giraldeau, L.-A., & Caraco, T. (2000). *Social foraging theory*. Princeton, N.J.: Princeton University Press.

Gittins, J. C. (1989). *Multi-armed bandit allocation indices*. New York: Wiley.

Glimcher, P. W. (2003). *Decisions, uncertainty, and the brain : The science of neuroeconomics*. Cambridge, Mass. ; London: MIT Press.

Goldstein, E. B. (2005). *Cognitive psychology : Connecting mind, research, and everyday experience*. Australia Belmont, CA: Thomson/Wadsworth.

Hall, J. W. (1978). Deciding as a way of intending. *The Journal of Philosophy, 75*(10), 553-564.

Hames, R. B., & Vickers, W. T. (1982). Optimal diet breadth theory as a model to explain variability in amazonian hunting. *American Ethnologist, 9*(2, Economic and Ecological Processes in Society and Culture), 358-378.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *J Exp Anal Behav., 4*(4), 267–272.

Hu, H., & Brady, M. (1996). A parallel processing architecture for sensor-based control of intelligent mobile robots. *Robotics and Autonomous Systems, 17*(4), 235-257.

Jochim, M. A. (1988). Optimal foraging and the division of labor. *American Anthropologist, 90*(1), 130-136.

Johnson-Laird, P. N., & Shafir, E. (1993). The interaction between reasoning and decision making: An introduction. *Cognition, 49*(1-2), 1-9.

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research, 4*, 237-285.

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *Am Psychol, 58*(9), 697-720.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty : Heuristics and biases*. Cambridge ; New York: Cambridge University Press.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263-291.

Kahneman, D., & Tversky, A. (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quartely Journal of Economics, 106*(4), 1039-1061.

Kahneman, D., & Tversky, A. (2000). *Choices, values, and frames*. Cambridge, UK: Cambridge University Press.

Kaplan, H., Hill, K., Hawkes, K., & Hurtado, A. (1984). Food sharing among ache hunter-gatherers of eastern paraguay. *Current Anthropology, 25*(1), 113-115.

Kimura, H., Yamashita, T., & Kobayashi, S. (2001). Reinforcement learning of walking behavior for a four-legged robot. *Decision and Control, 2001. Proceedings of the 40th IEEE Conference on, 1*.

Kortenkamp, D., Bonasso, R. P., & Murphy, R. (1998). *Artificial intelligence and mobile robots: Case studies of successful robot systems*: MIT Press Cambridge, MA, USA.

Koshland, D. E., Jr. (1977). A response regulator model in a simple sensory system. *Science, 196*(4294), 1055-1063.

Krebs, J. R., & Davies, N. B. (1997). *Behavioural ecology : An evolutionary approach* (4th ed.). Oxford, England ; Malden, MA: Blackwell Science.

Krebs, J. R., & Kacelnik, A. (1991). Decision making. In J. R. Krebs & N. B. Davies (Eds.), *Behavioural ecology: An evolutionary approach* (3rd ed., pp. 105-137). Oxford: Blackwell Scientific Publishers.

Lyshevski, S. E. (2003). Modeling, Simulation, Control and Optimization Paradigms for E.Coli Bacteria.

Proceedings of the Third IEEE Conference on Nanotechnology, 2, 690-693.

MacDorman, K. F. (1999). Grounding symbols through sensorimotor integration. *Journal of the Robotics Society of Japan, 17*(1), 20-24.

Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology, 33*, 101-112.

Malle, B. F., Moses, L. J., & Baldwin, D. A. (2001). *Intentions and intentionality : Foundations of social cognition*. Cambridge, Mass.: MIT Press.

Mataric, M. J. (1995). Designing and understanding adaptive group behavior. *Adaptive Behavior, 4*(1), 51.

Mataric, M. J. (1997a). Behaviour-based control: Examples from navigation, learning, and group behaviour. *Journal of Experimental & Theoretical Artificial Intelligence, 9*(2 - 3), 323-336.

Mataric, M. J. (1997b). Reinforcement learning in the multi-robot domain. *Autonomous Robots, 4*(1), 73.

Mataric, M. J. (1998). Behavior-based robotics as a tool for synthesis of artificial behavior and analysis of natural behavior. *Trends in Cognitive Sciences, 2*(3), 82-86.

Mataric, M. J. (2001). Learning in behavior-based multi-robot systems: Policies, models, and other agents. *Cognitive Systems Research, 2*(1), 81-93.

Mataric, M. J. (2002). Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics. *Imitation in Animals and Artifacts*, 391–422.

Mataric, M. J. (2002). Situated robotics. *Encyclopedia of cognitive science, 4*, 25-30. McMillan, UK. s

Mazur, J. E. (2002). *Learning and behavior* (5th ed.). Upper Saddle River, N.J.: Prentice Hall.

McCabe, K. (2005). Neuroeconomics. In L. Nadel (Ed.), *Encyclopedia of cognitive science* (pp. 294-298). New York, N.Y.: Wiley InterScience.

Medin, D. L., & Atran, S. (1999). *Folkbiology*. Cambridge, Mass.: MIT Press.

Michaud, F., & Mataric, M. J. (1998). Learning from history for behavior-based mobile robots in non-stationary conditions. *Autonomous Robots, 5*(3), 335-354.

Nilsson, N. J. (1984). *Shakey the robot*: SRI International.

Nolfi, S. (1997). Using emergent modularity to develop control system for mobile robots. *Adaptive Behavior, 5*, 343-364.

Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, Mass.: MIT Press.

Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review, 106*(4), 643.

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature, 400*(6741), 238.

Popper, K. R. (1994). Models, instruments, and truth: The status of the rationality principle in the social sciences. In *The myth of the framework. In defence of science and rationality* (pp. 154-184). London: Routledge.

Rilling, J. K., Gutman, D. A., Zeh, T. R., Pagnoni, G., Berns, G. S., & Kilts, C. D. (2002). A neural basis for social cooperation. *Neuron, 35*(2), 395-405.

Robbins, L. (1932). *An essay on the nature and signifiance of economic science*. London Macmillan.

Rushworth, M. F. S., Behrens, T. E. J., Rudebeck, P. H., & Walton, M. E. (2007). Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. *Trends in Cognitive Sciences, 11*(4), 168-176.

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science, 300*(5626), 1755-1758.

Searle, J. (2001). *Rationality in action*. Cambridge, Mass.: MIT Press.

Selten, R. (2001). What is bounded rationality ? . In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 13-36). MIT Press: Cambridge, MA.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition, 49*(1-2), 11-36.

Sih, A., & Christensen, B. (2001). Optimal diet theory: When does it work, and when and why does it fail? *Animal Behaviour, 61*(2), 379.

Smith, E. A. (1991). *Inujjuamiut foraging strategies : Evolutionary ecology of an arctic hunting economy*. New York: A. de Gruyter.

Sorensen, R. (2004). Charity implies meta-charity. *Philosophy and Phenomenological Research, 26*, 290-315.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, N.J.: Princeton University Press.

Sternberg, R. J., & Wagner, R. K. (1999). *Readings in cognitive psychology*. Fort Worth, TX: Harcourt Brace College Publishers.

Sundaram, R. K. (2003). *Generalized bandit problems*: Working Paper, Stern School of Business.

Tackseung, J. (2004). A survey on the bandit problem with switching costs. *De Economist, V152*(4), 513-541.

Terry, E., & Nigel, R. S. (2003). Developmental robotics: Manifesto and application. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 361*(1811), 2187-2206.

Thaler, R. H. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization, 1*(1), 39-60.

Thrun, S. (2002). *Robotic mapping: A survey*: School of Computer Science, Carnegie Mellon University.

Turk, M. A., Morgenthaler, D. G., Gremban, K. D., & Marra, M. (1988). Vits-a vision system for autonomous land vehicle navigation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on, 10*(3), 342-361.

Tversky, A. (1975). A critique of expected utility theory: Descriptive and normative considerations. *Erkenntnis, V9*(2), 163-173.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and psychology of choice. *Science, 211*, 453-458.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *The Journal of Business, 59*(4), S251-S278.

Urmson, C., Simmons, R., & Nesnas, I. (2003). A generic framework for robotic navigation. *Aerospace Conference, 2003. Proceedings. 2003 IEEE, 5.*

Volpe, R., Nesnas, I., Estlin, T., Mutz, D., Petras, R., & Das, H. (2001). The CLARAty architecture for robotic autonomy. *Aerospace Conference, 2001, IEEE Proceedings., 1.*

Weirich, P. (2004). Economic rationality. In A. in Mele, & Rawlings, P. (Ed.), *Oxford handbook of rationality* (pp. 380–398). Oxford: Oxford University Press.

Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., et al. (2001). Artificial intelligence: Autonomous mental development by robots and animals. *Science, 291*(5504), 599-600.

Winterhalder, B., & Smith, E. A. (2000). Analyzing adaptive strategies: Human behavioral ecology at twenty-five. *Evolutionary Anthropology: Issues, News, and Reviews, 9*(2), 51-72.

Zak, P. J. (2004). Neuroeconomics. *Philos Trans R Soc Lond B Biol Sci, 359*(1451), 1737-1748.

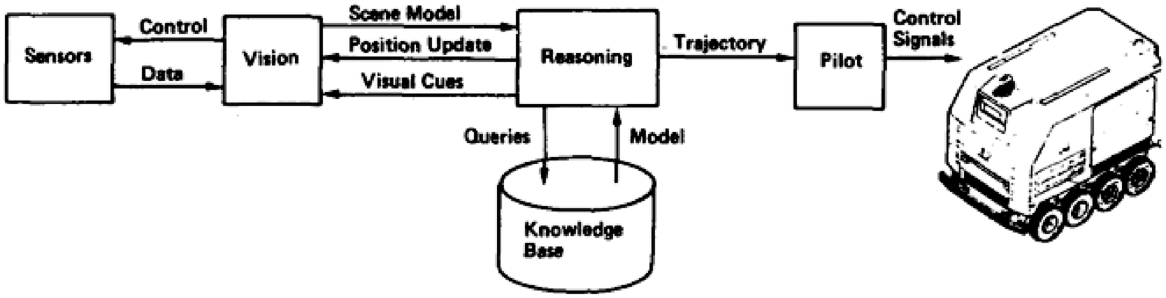

Fig. 1. The ALV system configuration.