TITLE: **MODELS, CONTROLS, AND LEVELS OF SEMIOTIC AUTONOMY**

RECEIVED
OCT 0 5 1998
OSTI

AUTHOR(S): Cliff Joslyn

Los Alamos

Los Alamos National Laboratory
Los Alamos New Mexico 87545

DISCLAIMER

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Models, Controls, and Levels of Semiotic Autonomy

Cliff Joslyn *

Abstract

In this paper we consider forms of autonomy, forms of semiotic systems, and any necessary relations among them. Levels of autonomy are identified as levels of system identity, from adiabatic closure to disintegration. Forms of autonomy or closure in systems are also recognized, including physical, dynamical, functional, and semiotic. Models and controls are canonical linear and circular (closed) semiotic relations respectively. We conclude that only at higher levels of autonomy do semiotic properties become necessary. In particular, all control systems display at least a minimal degree of semiotic autonomy; and all systems with sufficiently interesting functional autonomy are semiotically related to their environments.

1 Introduction

A crucial aspect of the modern Information Sciences, from robotics to self-organizing databases, is the manner and extent to which information systems can be said to be "autonomous". Of course, there are many different sense of autonomy, from mere physical separation to the establishment of complex functional relations sufficient to allow the identification of the systems as "independent actors". But in any event, some sense of *closure* of the system's activities, or identification of processes which are closed loops, is required. These closures can also take many forms, including physical boundaries, dynamic self-organization like autocatalytic chemical cycles, and informational "boundaries" as in elementary sensor-affector functioning loops.

For many classes of information systems, closures are of a **semiotic** nature, where there is now a closure of *meaning* within the system through interaction with the environment. Here issues arise concerning the use and interpretation of **symbols** ("symbol grounding"), **representations**, and/or **internal models** (whether explicit or implicit) by the system; and the syntactic, semantic, and pragmatic relations among the sign tokens, their interpretations, and their use or function for the systems in question. This is the case, for example, with biological systems, whose autonomy is enabled by their own genetic information as used in the context of their environments; and with control systems, where autonomy is enabled by the information in their engineered construction and use by human operators.

In this paper, we consider various classes of autonomous and semiotic systems, and consider the question whether there is any form of mutual necessity among these classes. In particular, we will consider the question as to whether autonomous systems of a certain class have any necessary semiotics properties, and *vice versa*.

2 Autonomy and Identity

Autonomy literally means "self-governing", and is derived from the Greek for "self-law". In linguistic history, autonomy typically refers to social governing bodies, for example autonomous regions within nations. Generally, it has implications of independence, separatedness, and apartness. In modern scientific usage, autonomy has come to be associated with the self-regulation of smaller systems, and in particular how on the one hand biologically evolved organisms, and on the other mechanically designed robotic systems, develop and maintain autonomous *functioning* in variable environments.

New classes of systems and environments, which we can broadly classify as Distributed Information Systems (DIS), have attracted a great deal of attention. Examples of these computer-aided virtual environments include Nets, Webs, and Multi-User Virtual Environments (MUVE's) [9, 15]. The systems which inhabit these spaces include humans and software agents. The nature of the possible autonomies in these classes of systems is therefore also of great interest.

---

*Computer Research and Applications Group (CIC-3), Mail Stop B265, Los Alamos National Laboratory, Los Alamos, NM 87545, USA, joslyn@lanl.gov, http://www.c3.lanl.gov/~joslyn, (505) 667-9096.

2.1 Autonomy as Closure and Identity

In systems science, we try to produce generalized models of system-environment couplings in a variety of modalities. In tracing the relations within and between systems, we can generally recognize two forms of relations: **input/output relations**, where linear influences flow *through* systems (or components of systems, or system-environment couplings); and **closures**, where circular influences flow *within* systems (or components, or couplings).

We therefore recognize autonomy as a form of closure. Many types of closures are also recognized in systems science, the most important being **self-reference** (referential closure), **self-orgnization** (functional closure), **autopoeisis** (organizationl closure), **control** (causal closure), **autocatalytic cycles** (dynamic closure), etc.

Each form of closure introduces a form of hierarchical scaling. In particular, **boundaries** mark those processes which are included in the closure (and are therefore inside the closure), from those which are excluded from the closure (and thus outside). There is a corresponding **stability**, in that processes involved in the closure can exist at a relatively larger spatial or temporal scaling than those outside. It is through this hierarchy that system **identity** itself is established in terms of those boundaries and stabilities.

2.2 Forms of Autonomy

Generally, we can recognize a continuum of classes of autonomous systems. One extreme is complete adiabatic isolation, and its corresponding complete boundary: no energy or information can flow across that barrier. No processes of organization or development can occur, but rather there is only an inexorable collapse to thermodynamic equilibrium.

The other extreme can never actually be reached. As a system becomes less closed from its environment, it becomes more involved in linear flows, and has weaker boundaries. In the limit, it actually loses all identity, and ceases to exist as a distinct system.

It is clear that real, complex systems (the systems which are of most interest to us, like organisms, complex mechanisms, and DIS), cannot exist at either extreme of autonomy. Rather, they are all both autonomous in certain modes with respect to their environments, and simultaneously involved, in other modes, in throughput relations with their environments.

For example, common objects have a form of **physical autonomy**, maintaining their identity by physical closure and separation, but interacting through energetic exchanges. Classical self-organizing systems have a form of **dynamic autonomy**, maintaining their identity by closures of flows and structures, but interacting through input/output flows of matter and energy. Organisms have a form of **functional autonomy**, where their identity is maintained by a constantly shifting organizational closure, but interacting through informational flows.

We introduce here the concept of **semiotic autonomy**. Semiotically autonomous systems (similar to what Pattee calls semantically or semiotically closed systems [11]) maintain cyclic relations of perception, interpretation, decision, and action with their environments. These are semiotic processes, involving the reference and interpretation of sign tokens maintained in coding relations with their interpretants.

Semiotic closures produce corresponding *intentional* closures and boundaries of *meaning*, allowing complex functional relations to exist within and between them and their environments. These are sufficient to establish forms of identity based on self-reference, allowing the identification of these systems as "independent actors". It has been hypothesized that these systems are equivalent to the class of organisms [7]. It is the nature of these semiotic relations, and therefore the nature of semiotic autonomy, which we wish to consider below.

3 Models and Controls as Semiotic Systems

In systems theory and cybernetics the modeling relation and the control relation serve as two fundamental and distinct classes of relations between a system and its environment, or "the world".

3.1 Simple Models and Controls

Consider first a classical control system. All knowledge of the environment by the system is mediated through the measurement (perception) process, which provides a (partial) representation of the environment to the system. Based on this, the system then chooses a particular action to take in the world, which has consequences for the change in state of the world and thereby states measured in the future.

To be in good control, the overall system must form a negative feedback loop, so that disturbances and other external forces from "reality" (noise or the actions of other external control systems) are counteracted by compensating actions so as to make the measured state (the representation) as close as possible to some desired state, or at least stable within

some region of its state space. If rather a positive feedback relation holds, then such fluctuations will be amplified, ultimately bringing some critical internal parameters beyond tolerable limits, or otherwise exhausting some critical system resource, and leading to the destruction of the system as a viable entity.

The left side of Fig. 1 is a functional view of a simple control system, representing the logical relations among certain components of the system and the world: the nodes are logical constructs and the arrows are labeled by the kind of relations which hold between them, or the nature of the constraint one places on the other.

Alternatively, a structural version of the same diagram is shown in the right side of Fig. 1, representing now the physical entities in the system and the world and how they are structurally related: the nodes are subsystems which perform certain physical processes, and the arrows are labeled by how they interact. Thus the physical sensors interact with the state of affairs in the world to produce a representation (token), which is passed to the agent, which executes a decision to choose a particular action taken in the world. Note how generally the functional and structural views are dual: nodes in one are generally arrows in the other, and vice versa.

Now consider the canonical modeling relation as shown in the left side of Fig. 2. As with the control relation, the processes of the world are still represented to the system only in virtue of measurement processes. But now the decision relation is replaced by a prediction relation, whose responsibility is to produce a new representation which is hypothesized to be equivalent (in some sense) to some future observed state of the world. To be a good model, the overall diagram must commute, so that this equivalence is maintained.

As with the control system, this is a functional representation, and a structural version is also shown on the right side of Fig. 2. Here as well the sensors enter into relations with states of affairs in the world and create representations, but these are now sent only to a comparator. There is no relation back from the system to the world.

Controls and models, as systems, have fundamentally different topological relations with their environments, and these map to the previous discussion of fundamental processes in systems science. Consider a system S in relation to its environment E. Naturally there are two sets of relations $g: E \mapsto S$ from the environment to the system and $f: S \mapsto E$ back from the system to the environment. In the modeling relation, only g is present as a measurement function, and thus the structure of a model is fundamentally

linear, from the world to the model. But in a control relation, g is present as measurement, but f is also present as the action relation from the system back to the world. Thus control is fundamentally circular, from the system to the world and back again.

3.2 Complex Modeling and Control Relations

Of course, the relations described here are a great deal more complex in real control systems. In general, modeling and control are combined in real systems in a variety of complex ways [6].

Computation in Control: It is possible to augment the control system with a computation relation between the representation of the measured state and another which is then passed to the decider. This "computation" plays the *role* of cognition, information processing, or knowledge development. Typically, extra or external knowledge about the state of the world or the desired state of affairs is brought to bear, and provided to the agent in some processed form, for example as an error condition or distance from optimal state. In real systems computations take such forms as the more abstract or combined perceptions in neural organisms or the results of a real computation in machines.

Hierarchical Control: Classical linear control systems theory for hierarchical control emerges here Fig. 3 [13, 14].[1] The computer is viewed as a comparator between the measured state and a hypothetical set point or reference level (goal). This then sends the second representation of an error signal to the agent. Such systems are hierarchically scalable, in that representations can be combined to form higher level representations, and actions of one control system can be the determination of the set-point of another, thus allowing goals to decomposed as a hierarchy of sub-goals. While ultimately the lower level is responsible for taking action in the world, it is doing so under the control of the comparison of a high-level goals against a high-level representation.

3.3 Hybrid Modeling and Control

It is clear that control can be done without computation, modeling, or planning, based strictly on feedback. The difference is that in control the representation of what *is* is compared to what is *wanted*, while

---

[1]http://www.ed.uiuc.edu/csg

in modeling it is compared to what is *expected* (based on the model's predictions). Ashby called this form of prediction-less control "error control" [1, 3], which he distinguished from "cause control", which involved the prediction of future events to guide actions. We also know cause control as anticipatory [16] or feed-forward [8] control. Ashby actually favored cause control, since in principle it could be made perfect (with a perfect model of the world), while error control can only be improved in the limit at infinitesimal lag.

Various hybrid control architectures are possible.

Mixtures: The simplest is to modify computation to include a further measurement, used corroborate the results of the computational step. The world is now the source of both the initial sensory input and the corroboratory measurement, but these steps are separated in time. This allows the introduction of the relation between the agent and the representation, and it becomes apparent that functionally this relation actually is the control relation itself. Thus we notice, with Powers, that it is not, in fact, the state of the world which is being controlled, but rather the *perception* of the state of the world by the system. This fundamental result of constructivism falls out naturally from our descriptions here.

Anticipatory: Usually when we think about cause or anticipatory control there is an embedded model which is used to make a decision as to which action to take, acting in the role of the agent. Thus we replace the agent with an inner system which is *both* a model and a control system. This inner system is a control system in the sense that there are states of its "world", its "dynamics", and an "agent" making decisions. However, it is also a model in that the states of its "world" are in fact representations, and its "dynamics" is actually a prediction function. The inner system is totally contained within the outer system, and runs at a much faster time scale in a kind of modeling "imagination". The representation from the sensors is used to instantiate this model, which takes imaginary actions resulting in imaginary stability within the model. Once this stability is achieved, then that action is exported to the real world.

3.4 Semiotic Relations in Controls and Models

Models and control systems are frequently both cast in the semiotic context, thereby evoking the distinc-

tions among three distinct classes of semiotic concepts [2, 4]:

Syntactic: Concerning the formal properties of symbol tokens as used in symbol systems.

Semantic: Concerning the interpretation of tokens as their meanings.

Pragmatic: Concerning the use of symbol tokens and their meanings for the overall purposes or survivability of the system.

We can identify the forms of relations present in models and control, and distinguish them in virtue of the fundamental concepts of variety and constraint.

Computation: *The constraint placed on tokens by themselves.* Codings are an expression of syntax, and are usually deterministic string replacement. As Pattee has commented at length [12], coding substitutions are computational, memory-dependent, and rate-independent.

Measurement: *The constraint placed on tokens by the world.* In sharp contrast to coding or computation, measurement provides the "grounding" of the symbol tokens.

Decision: *The constraint placed on actions by tokens.* Given the presence of a certain representation, either as the result of measurement or of computation, a particular action results.

Dynamics: *The constraint placed on the world by itself.* The rate-dependent dynamical structure of the universe, deterministic at *some* level.

Each of these relations is a form of constraint. Given a particular model or control system, there is then very little freedom: a given state of affairs in the world will result in specific representations, predictions, decisions, and actions.

But that is not to say that all these relations are the same *kind* of constraint. In particular, we can distinguish *laws* from *rules* [11, 15, 7]. Laws are wholly (ontologically) necessary at all levels of analysis, but rules are necessary at one level, and contingent at another: once a particular rule or coding (set of interpretations) is established, then it must be followed, but in general from a perspective outside the system many such interpretations are possible.

This property of rules is the hallmark of semiotic systems: that the coding of their symbol tokens act as *contingent functional entailments*, and are thus dually contingent and necessary at complementary levels of analysis. From *within* the symbol system,

the token must necessarily be interpreted according to the code, but from *without* we are (or "evolution is") free to choose any coding we please. They are conventional, constructed and interpretable by a certain closed "linguistic community" [10]. The school of biosemiotics [5] is dedicated, in some sense, to the proposition that the classes of semiotic systems and living systems are equivalent, or at least coextensive [5, 7].

Thus the presence of rules (contingent functional entailments) in a "good" system, whether an "accurate" model or a "good" control system, implies a level of meta-constraint in addition to those identified above, namely the constraint on which rules themselves are viable. The making of *appropriate* choices is exactly the semantic function in a semiotic system. It is on this required "appropriateness" of the choice of the agent that the "intelligence" of the semiotic system rests: a certain action is "correct" in a given context, while another is not. It is only on this basis that meaning or semantics can be said to be present in a control system or a model.

This additional level of constraint is what Pattee calls *selection* [12].

Selection: *The constraint on measurement, computation, and decision by the world.* This new level of constraint is the constraint within the space of all possible rules, in particular of all possible measurements, all possible computations, and all possible actions.

Selection is an example of the pragmatic aspect of semiotic systems, and must be provided by a force acting outside of the system (control system or model) itself. The typical agents of this selection are either natural selection or the decisions provided by the designer. Thus in a system which has contingent entailments (rules) the pragmatics of the selection of those rules invokes semantic relations of meaning among the components.

4   Levels of Semiotic Autonomy

Finally we are able to consider the questions which are actually motivating us here: to what extent do semiotic relations imply autonomy, and *vice versa?*

First, it is clear that control systems are forms of closure, and thus imply boundaries and identity in virtue of that closure. Furthermore, that closure rests on exactly the pragmatic selection (whether by natural selection or by an engineer) and semantic decision-making which makes them semiotic systems. Therefore it is safe to assert that all viable control systems

of the form shown in Fig. 1 have a form of semiotic autonomy.

However, pure models of the form shown in Fig. 2 do not necessary show any form of closure. Of course, "naked" models never exist in isolation. Rather, they are built by humans or have evolved in organisms in order to serve the purposes of some embedding control system, as in Fig. 4.

This is actually a very important point: just by building models or representations of some part of the world, I do not become semiotically autonomous with respect to that environment. Rather, that model must be *used*, those symbols must be *interpreted*, in a cyclic process of interaction with that environment, in order to achieve some pragmatic purpose. Some would say that that system

The converse question then becomes, given a system which is autonomous in some aspects, need it be semiotic? The obvious general answer is no: certainly mere physical autonomy, and even the kinds of dynamical autonomous shown by complex self-organizing physical system, do not imply any form of rule-following semiotic phenomena.

However, consider that I observe a system which has autonomy of *action*. That is, it is able to produce behavior which is in some sense "contrary" to my expectations of a physically law-following system. For example, a ball does not roll off a hill, but tries to climb up it. This is evidence that rule-following, and thus semiotic, relations are present in that system.

References

[1] Ashby, Ross: (1956) *Introduction to Cybernetics*, Methuen, London

[2] Cariani, Peter A: (1989) *On the Design of Devices with Emergent Semantic Functions*, SUNY-Binghamton, Binghamton NY, PhD Dissertation

[3] Conant, Roger C and Ashby, Ross: (1970) "Every Good Regulator of a System Must Be a Model of that System", *Int. J. Systems Science*, v. 1:2, pp. 89-97

[4] Deely, John: (1990) *Basics of Semiotics*, Indiana UP, Bloomington IN

[5] Deely, John: (1992) "Semiotics and Biosemiotics: Are Sign-Science and Life-Science Coextensive?", in: *Biosemiotics: The Semiotic Web 1991*, ed. TA Sebeok, J Umiker-Sebeok, pp. 46-75, Mouton de Gruyter, Berlin/NY

[6] Joslyn, Cliff: (1997) "Semiotic Aspects of Control and Modeling Relations in Complex Systems", in: *Control Mechanisms for Complex Systems*, ed. Michael Coombs, Addison-Wesley, Redwood City CA, SFI Stud.in Sci. of Comp.toapp

[7] Joslyn, Cliff: (1998) "Are Meaning and Life Coextensive?", in: *Evolutionary Systems*, ed. G. van de Vijvier, Kluwer, in press

[8] Klir, George: (1991) *Facets of Systems Science*, Plenum, New York

[9] Landauer, Chris and Bellman, Kirstie L: (1998) "Integration and Modeling in MUVEs", in: *Proc. Virtual Worlds and Simulation Conf.*, ed. C. Landauer + K. Bellman, pp. 187-194, Soc. forComputer Sim, San Diego

[10] Lewis, David: (1969) *Convention*, Harvard U Press, Cambridge

[11] Pattee, Howard: (1982) "Cell Psychology: An Evolutionary View of the Symbol-Matter Problem", *Cognition and Brain Theory*, v. 5, pp. 325-341

[12] Pattee, Howard: (1997) "The Physics of Symbols and the Evolution of Semiotic Controls", in: *Proc. Workshop on Control Mechanisms for Complex Systems*, Addison-Wesley

[13] Powers, WT: (1973) *Behavior, the Control of Perception*, Aldine, Chicago

[14] Powers, WT, ed.: (1989) *Living Control Systems*, CSG Press

[15] Rocha, Luis M and Joslyn, Cliff: (1998) "Simulations of Evolving Embodied Semiosis: Emergent Semantics in Artificial Environments", in: *Proc. 1998 Conf. on Virtual Worlds in Simulation*, pp. 233-238, Society Comp. Sim., San Diego

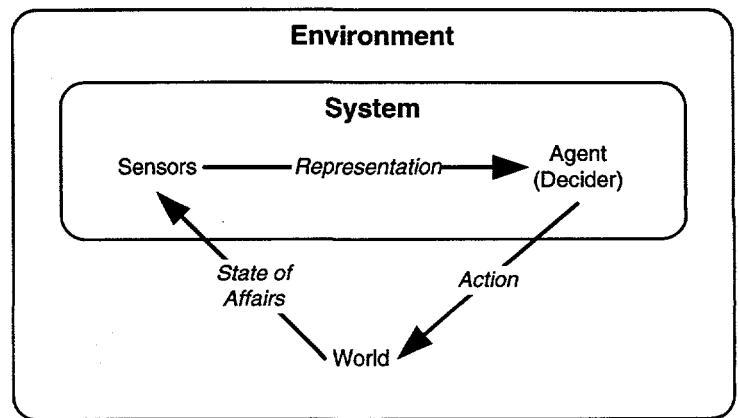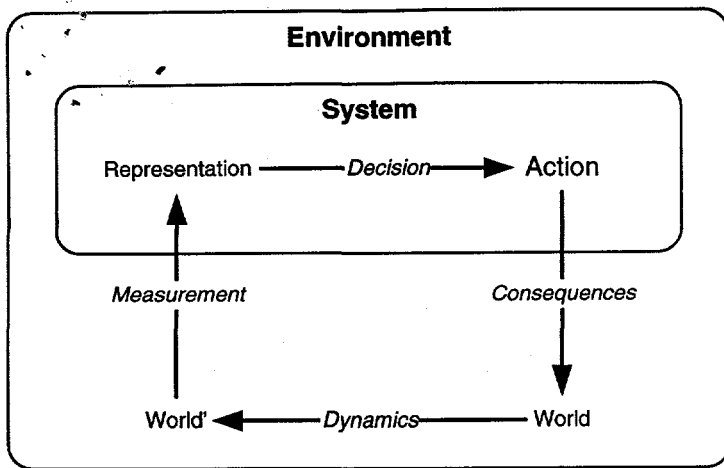[16] Rosen, Robert: (1985) *Anticipatory Systems*, Pergamon, Oxford

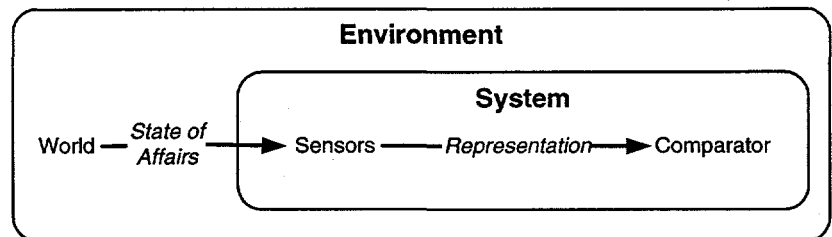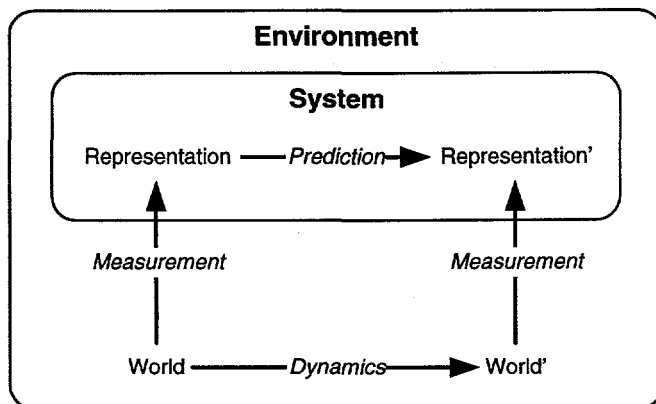Firgure 1: Simple control systems. Left: Functional; Right: Structural

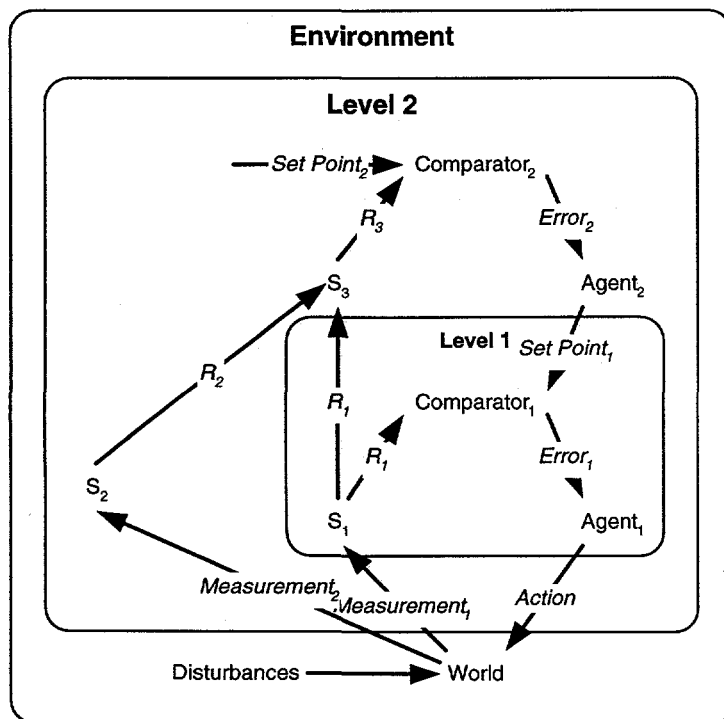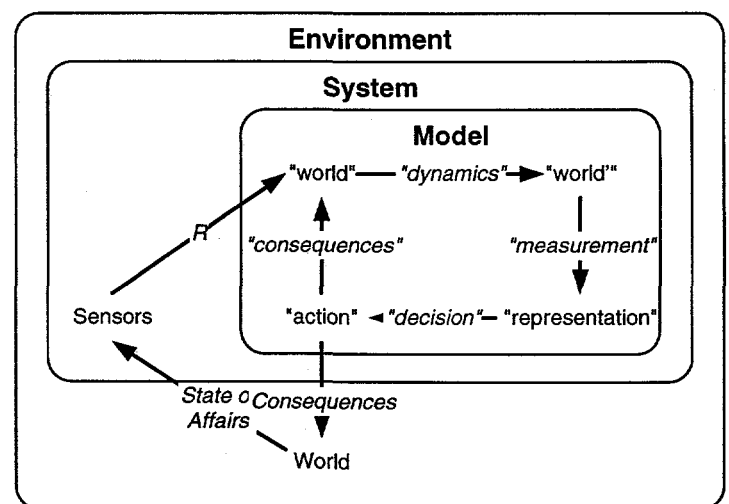Firgure 2: Simple models. Left: Functional; Right: Structural

Figure 3: Hierarchical Control

Figure 4: Anticipatory Control