Proceedings of the 1996 IEEE International
Symposium on Intelligent Control
Dearborn, MI • September 15-18, 1996

A Satisficing Approach to Intelligent Control of Nonlinear Systems

M. A. Goodrich, W. C. Stirling, and R. L. Frost

Electrical and Computer Engineering Department
Brigham Young University, Provo, UT 84604

Abstract

The need to control complex and uncertain systems has motivated the development of new ways to represent knowledge of the system, leading to the development of intelligent controllers. Since complexity and uncertainty may render the superlative paradigm (optimality seeking) inappropriate, alternative decision paradigms may be appropriate for such systems. Satisficing control is compatible with the limited rationality associated with such systems. Strongly satisficing controllers represent a rigorous, systematic synthesis procedure for the design of satisficing controls.

1. Introduction

There are two essential components of virtually any controller: a mechanism for knowledge representation and a mechanism for decision-making. With conventional optimal controllers, knowledge is represented with a mathematical model (such as a set of differential equations) and decision-making is based on the paradigm of self-interested rational choice, whereby the decision-maker seeks its own greatest good or greatest preference. We term this notion of rationality the *superlative* decision-making paradigm. This interpretation of rational choice is usually expressed, for conventional control problems, in terms of a utility function to be maximized.

As systems either become more complex or less precisely known, however, it becomes difficult to characterize the knowledge adequately with conventional means, and considerable attention has been focused on finding alternative ways to encode the available knowledge about the system. This search has resulted in the development of intelligent controllers based on analogies to biological systems. Fuzzy logic, knowledge bases, neural networks, and genes (for genetic algorithms) provide mechanisms for expressing knowledge via natural language, production rules, highly parallel adaptive structures, and measures of fitness for survival. These architectures for knowledge encoding have been demonstrated to possess significant virtues, especially for the control of nonlinear, uncertain, and adaptive systems.

The decision-making paradigm for these intelligent control mechanisms, however, is basically the same as for conventional control–that of seeking a superlative decision. Defuzzifiers are typically designed to obtain the best (according to a given criterion) crisp decision, production rules are at least implicitly designed to produce the best response, neural networks weights are typically computed to minimize a given performance index, and genetic algorithms operate according to the "survival of the fittest" rule.

When dealing with complex or uncertain systems, however, the issue of finding the best solution becomes clouded and should not be uncritically accepted. If our state of knowledge is incomplete or uncertain, it is not clear that a notion of "best" can even be defined. Even if it can be defined, it is not obvious that it can be found using the available information. Furthermore, even if it can be solved, it it not certain that available computational resources will permit the solution to be obtained.

Just as alternative means of knowledge representation have been motivated by the inadequacy of conventional knowledge representation techniques, it may be argued that alternative means of decision-making must be sought due to the inadequacy of the

0-7803-2978-3/96/$5.00 © 1996 IEEE

conventional notion of rationality–the superlative decision-making paradigm–in complex and uncertain environments. In this paper we present an alternative paradigm for decision-making that admits a cognitive analogy parallel to the biological analogies associated with "intelligent" knowledge representation procedures. Section 2 introduces the satisficing concept and reviews epistemic utility theory. Section 3 discusses the problem of reducing the set of seriously possible local control decisions identified by epistemic utility theory to a single control, and introduces the notion of strongly satisficing controls. Section 4 provides an illustrative example of satisficing control.

2. Theory of Satisficing Decisions

Simon has addressed the problem of making intelligent decisions when superlative paradigm is not appropriate: "broadly stated, the task is to replace the global rationality ... with a kind of behavior that is compatible with the access to information and the computational capabilities that are actually possessed ..." [3]. According to this perspective, solutions that meet a minimum standard, or aspiration level, possibly obtained under constraints of partial information or restricted computation, are termed *satisficing* solutions.

The notion of a minimum standard suggests a comparative, rather than a superlative, paradigm of decision-making. A epistemological basis for a comparative paradigm has been developed by Levi in the form of an *epistemic utility function*, which is composed of two simpler utility functions [2]. Levi's theory was originally developed as a model of cognitive decision-making. In that context, one utility represents the subjective probability of a proposition being true, and the other utility (also expressed as a normalized measure, or probability) represents the informational value that accrues to the decision-maker if the proposition is rejected. The decision-maker then rejects all propositions for which the utility of being true is dominated by the utility of being rejected. We have adapted this theory to the control context by reinterpreting truth-value as *accuracy*, meaning conformity to a given standard, and interpreting the informational value of rejection, or *rejectability*, as the cost of control. Each control may then be evaluated by comparing the accuracy of the control with its cost and rejecting all controls whose accuracy is dominated by its cost. In this way, only controls that achieve the goal at reasonable cost will be considered.

Control problems are often specified in terms of (a) the ultimate goal of the controller (for example, to drive the state to a fixed set-point), and (b) the design principles used to generate a specific control policy (for example, a performance index to be minimized). Let the set U denote the control space $(u_{\min}, u_{\max})$ where we have assumed that the accuracy and rejectability utilities assign all of their mass to this interval. Let $\mathcal{B}(U)$ be a σ-field in U. We define the accuracy of a set $G \in \mathcal{B}(U)$ as the probability measure $P_A : \mathcal{B}(U) \mapsto [0,1]$. Rejectability is a measure of how well a control decision obeys the design principles, independently of its effectiveness. The rejectability utility may be expressed through another probability measure, $P_R : \mathcal{B}(U) \mapsto [0,1]$. High rejectability means high cost of control, *independently* of accuracy considerations.

We restrict attention to measures for which density functions exist and denote these densities $f_A(u)$ and $f_R(u)$. The satisficing set is defined as

$$S = \{u \in U : f_A(u) \geq b f_R(u)\}, \qquad (1)$$

where b, the index of boldness, establishes the satisficing threshold. For $b \leq 1$, S will always be nonempty. We restrict attention to boldness values small enough to guarantee that S defined in (1) is nonempty. In contrast to many decision-making procedures, this approach relaxes the requirement for a unique best decision, and admits as satisficing all decisions that are good enough to qualify according to (1). Obviously, only one control can actually be implemented, but, from a strictly satisficing point of view, one may choose any of the unrejected control decisions with some confidence that the action will yield good, if not optimal, performance. Thus the designer has considerable latitude in the ultimate choice of the control to be implemented.

3. Strongly Satisficing Control

Satisficing, as we have defined it, is a weak notion of performance: broadly speaking, a proposition is satisficing if the good (characterized by accuracy) outweighs the bad (characterized by rejectability). Furthermore, the satisficing set, S, will generally not be a singleton set, and there may be many satisficing control possibilities. One way to select a single effective control from S is to choose, for a given level of rejectability, the proposition with maximum accuracy[1]. We will term such a proposition *strongly satisficing*. In the interest of clarity, we will deal with scalar valued controls.

Three strongly satisficing controls are immediately obvious. A *most accurate* satisficing control, $u_A = \arg\max_{z \in S}\{f_A(z)\}$ would be appropriate for cases with large variations in f_A relative to small variations in f_R. Such a decision represents a very aggressive stance to achieve the goal at the risk of excessive cost. A *least rejectable* satisficing control, $u_R = \arg\min_{z \in S}\{f_R(z)\}$ would be appropriate when there are large variations in f_R relative to changes in f_A. This procedure is very conservative, and reflects a willingness to compromise the goal in the interest of reducing cost. A *most discriminating* satisficing control, $u_D = \arg\max_{z \in S}\{f_A(z) - bf_R(z)\}$ reflects a desire to compromise between cost and achievement in a way that maximizes the difference between the two.

It is desirable to identify the set, S_s, of all strongly satisficing solutions. Clearly, $\{u_A, u_R, u_D\} \subset S_s$. Let $r \in [0,1]$ be a given rejectability level, and let ρ_r be the inverse image set of r under f_R, that is, $\rho_r = f_R^{-1}(\{r\})$. If both f_A and f_R are differentiable with derivatives f'_A and f'_R, respectively, the strongly satisficing set becomes

$$S_s = \{u \in U : \exists r \text{ s. t. } u = \arg\max_{z \in \rho_r}\{f_A(z) : f'_A(z) f'_R(z) \geq 0\}. \quad (2)$$

Strongly satisficing solutions enjoy an equilibrium property: the accuracy cannot be increased without also increasing the rejectability, and the rejectability cannot be decreased without also decreasing the accuracy.

4. Control of the Inverted Pendulum

We apply a satisficing receding horizon control to a problem that has proven to be surprisingly difficult: the control of an inverted pendulum in a vertical plane with full circular freedom by applying a lateral force to the cart to which the pendulum is attached, while simultaneously regulating the position of the cart.

Consider the apparatus illustrated in Figure 1. The problem is to bring the pendulum from a vertically downward to a vertically upward position by applying a force to the cart while keeping the cart regulated to the origin. This problem is prototypical of many nonlinear control problems, and thus makes a good test case for a satisficing control theory. Traditional controllers must linearize the dynamics model of the pendulum in a small region within, say 10 deg of the vertical. More recently, a fuzzy controller trained by a genetic algorithm has been shown to balance the pendulum 90 % of the time if the pendulum is given a random initial position within 80 deg of the vertical and a random initial velocity less that 80 deg/sec [1]. We design an satisficing controller to control the pendulum given any set of initial conditions while simultaneously regulating the cart at a desired point.

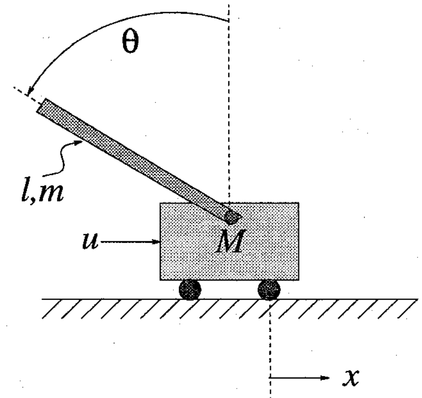

Figure 1: Inverted pendulum on a cart.

The discrete-time dynamical equation for this problem is

$$\mathbf{x}(t+1) = \mathbf{x}(t) + T\{\mathbf{f}[\mathbf{x}(t) + \mathbf{g}[\mathbf{x}(t)]u(t)\}, \quad (3)$$

for $t = 0, 1, \ldots$, where $\mathbf{x} = [\theta, x, \dot{\theta}, \dot{x}]^T$, T is the sampling interval, and $\mathbf{f} = [f_1, f_2, f_3, f_4)]^T$ and $\mathbf{g} = [g_1, g_2, g_3, g_4)]^T$ are given by

$$\begin{aligned} f_1(\mathbf{x}) &= x_3 \\ f_2(\mathbf{x}) &= x_4 \\ f_3(\mathbf{x}) &= \frac{(M+m)g \sin x_1 - ml \cos x_1 \sin x_1 x_3^2}{l(M + m \sin^2 x_1)} \\ f_4(\mathbf{x}) &= \frac{ml \sin x_1 x_3^2 - mg \cos x_1 \sin x_1}{M + m \sin^2 x_1}. \end{aligned}$$

[1]Such a decision is still a comparative decision; it is, in addition, however, a best comparative decision subject to a rejectability constraint. It may also be a superlative decision, but that status is not guaranteed.

$$
\begin{aligned}
g_1(\mathbf{x}) &= 0 \\
g_2(\mathbf{x}) &= 0 \\
g_3(\mathbf{x}) &= \frac{-\cos x_1}{l(M + m\sin^2 x_1)} \\
g_4(\mathbf{x}) &= \frac{1}{M + m\sin^2 x_1}.
\end{aligned}
$$

M is the mass of the cart, l is the length of the pendulum, m is the mass of the pendulum, θ is the angle from vertical (measured counterclockwise), x is the horizontal position of the cart, and u, the control input, is a lateral force applied to the cart.

Accuracy, for this problem, is a measure of how close a given control drives the pendulum and the cart to their respective origins. This measure may be specified, for example, via natural language expressions or by a mathematical expression. For this development, we adopt a mathematical expression, and define accuracy in terms of squared error for a one-step receding horizon controller:

$$\Phi(u) = x_1^2(k+2) + x_2^2(k+2) + x_3^2(k+1) + x_4^2(k+1). \tag{4}$$

We adopt, as the measure of rejectability, a positive-definite quadratic function of the form

$$
\begin{aligned}
\Lambda(u) = {} & q_{11}x_1^2(k+2) + q_{22}x_2^2(k+2) + q_{33}x_3^2(k+1) \\
& + q_{44}x_4^2(k+1) + q_{12}x_1(k+2)x_2(k+2) \\
& + q_{34}x_3(k+1)x_4(k+1) + ru^2,
\end{aligned} \tag{5}
$$

where q_{ij} and r are positive-definite weighting coefficients. From these equations, f_A and f_R may be determined as

$$
\begin{aligned}
f_A(u) &= \kappa_A\Big[\max_{z\in U}\{\Phi(z)\} - \Phi(u)\Big], \\
f_R(u) &= \kappa_R\Big[\Lambda(u) - \min_{z\in U}\{\Lambda(z)\}\Big],
\end{aligned}
$$

where κ_A and κ_R are the normalizing constants required to create probability densities. The resulting most discriminating controller is, after some calculations,

$$
\begin{aligned}
u_D = {} & -\left[G^T[\mathbf{x}(k)](I + b'Q)G[\mathbf{x}(k)] + b'r\right]^{-1} G^T \\
& \times[\mathbf{x}(k)](I + b'Q)F[\mathbf{x}(k)],
\end{aligned} \tag{6}
$$

where $Q = \{q_{ij}\}$, $b' = \frac{\kappa_A}{\kappa_R}b$, and

$$G[\mathbf{x}(k)] = \left[T^2 g_1[\mathbf{x}(k)] \;\; T^2 g_2[\mathbf{x}(k)] \;\; Tg_1[\mathbf{x}(k)] \;\; Tg_2[\mathbf{x}(k)]\right]^T$$

$$F[\mathbf{x}(k)] = \begin{bmatrix} T^2 f_1[\mathbf{x}(k)] + 2Tx_3(k) + x_1(k) \\ T^2 f_2[\mathbf{x};(k)] + 2Tx_4(k) + x_2(k) \\ Tf_1[\mathbf{x}(k)] + x_3(k) \\ Tf_2[\mathbf{x}(k)] + x_4(k) \end{bmatrix}.$$

Figures 2 and 3 illustrate the rotational (pendulum) and translational (cart) phase planes. The ∘ symbol represents the initial conditions (the cart at the origin with the pendulum in the vertical down position) and the × symbol represents the terminal conditions (the cart at the origin with the pendulum balanced in the vertical up position). The system achieves its desired objective of balancing the pendulum at the origin by swinging the pendulum back and forth while the cart oscillates around the origin. As the cart oscillates, the pendulum gathers momentum. In the translational and rotational phase planes, this motion is manifest as growing spirals. When the amplitude increases sufficiently, the oscillation ceases and the pendulum then converges to the vertical upright position. Finally, the cart returns slowly to the origin. Figure 4 illustrates the control history.

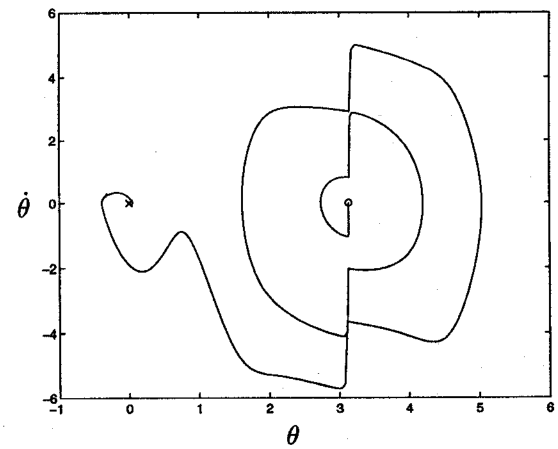

Figure 2: Rotational phase plane for the inverted pendulum on a cart (in radians and radians per second).

5. Discussion

Intelligent controllers are usually characterized by the way they represent knowledge. Equally important to the design of a controller is the way decisions are made. Both conventional control and intelligent control typically use the same superlative decision-making paradigm. Satisficing control, however, represents a *comparative* decision-making paradigm, and is therefore an alternative to the superlative approach. Satisficing decision-making may be applied to either conventional or intelligent control designs.

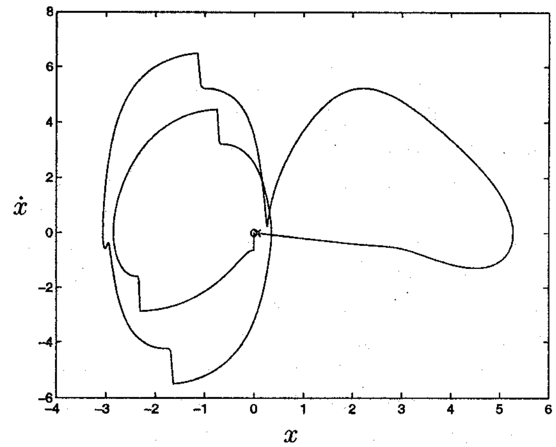

Figure 3: Translational phase planes for the inverted pendulum on a cart (in meters and meters per second).

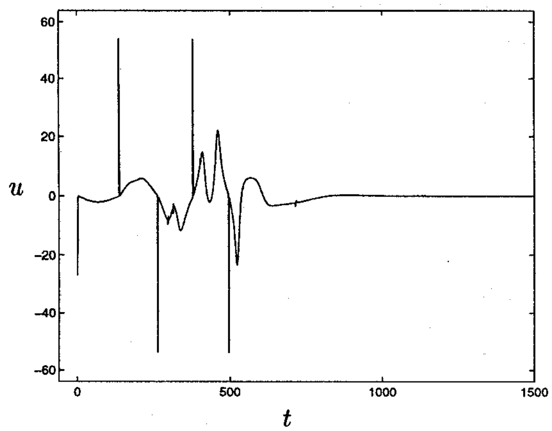

Figure 4: Control time-history for the inverted pendulum on a cart (newtons).

Since it is based on epistemic utility theory, satisficing control enjoys an intuitively appealing relationship to human cognition: accept as admissible all decisions for which the accuracy of the decision is as least as great as is the cost benefit that accrues if the decision were rejected. Strongly satisficing represents a systematic tie-breaking mechanism among the satisficing solutions.

References

[1] J. Kim, Y. Moon, and B. P. Zeigler. Designing Fuzzy Net Controllers Using Genetic Algorithms. *IEEE Control Systems Magazine*, 15(3):66–72, 1995.

[2] I. Levi. *The Enterprise of Knowledge.* MIT Press, Cambridge, Massachusetts, 1980.

[3] H. A. Simon. A Behavioral Model of Rational Choice. *Quart. J. Economics*, 59:99–118, 1955.