2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

Sloppy motors, flaky sensors, and virtual dirt: Comparing imperfect ill-informed robots

Jason M. O'Kane and Steven M. LaValle

Abstract— **Robots must complete their tasks in spite of unreliable actuators and limited, noisy sensing. In this paper, we consider the *information requirements* of such tasks. What sensing and actuation abilities are needed to complete a given task? Are some robot systems provably "more powerful" than others? Can we find meaningful equivalence classes of robot systems? This line of research is inspired by the theory of computation, which has produced similar results for abstract computing machines. The basic idea is a dominance relation over robot systems that formalizes the idea that some robots are stronger than others. We show that this definition is directly related to the robots' ability to complete tasks. Our prior work in this area assumes perfect control and sensing, requires that the robot begin with a single fixed initial condition within a known environment, and models of time as a sequence of variable-length discrete stages, rather than as a continuum. In this paper, we substantially improve upon that earlier work by addressing these problems.**

I. INTRODUCTION

Suppose we want a robot to complete some task, such as navigating to a goal, manipulating an object, or localizing itself within its environment. Many different combinations of sensing and motion modalities can be (and have been) used to complete each of these tasks. Indeed, much of the robotics literature is concerned with finding *sufficient conditions* on the sensing and actuation capabilities needed to complete such tasks. In this paper we take a different approach. For a given task, we are interested in determining the *necessary conditions*: What sensors and actuators are needed? What are the *information requirements* of robotic tasks? The long term goal of this research is to develop a theory of robots and sensing that helps in answering such questions.

A. Robots, sensors, and the theory of computation

This work is inspired in part by the theory of computation, which begins with precisely defined models of abstract machines, such as finite automata, pushdown automata, Turing machines, and so on [8]. In this context, a *problem* is usually a language of strings; to solve the problem is to accept strings in this language and reject all others. The theory of computation gives answers several kinds of basic questions about these machines and problems.

Solvability: Can a given machine can solve a given problem?

This work is supported by ONR Grant N00014-02-1-0488 and by DARPA grant #HR0011-07-1-0002. J. M. O'Kane and S. M. LaValle are with the Department of Computer Science, University of Illinois at Urbana-Champaign, 201 North Goodwin Avenue, Urbana, IL 61801, USA. {jokane, lavalle}@cs.uiuc.edu

Complexity: If the machine can solve the problem, how efficiently (in terms of time or space, for example) can it do so?

Comparison: Are some machines strictly more powerful, in terms of the problems they can solve, than others? It is known, for example, that pushdown automata can accept a strictly larger set of languages than can finite automata. Likewise, Turing machines are more powerful than pushdown automata.

Equivalence: Are there apparently dissimilar machines that can solve the same set of problems? For example, it is a standard result that a Turing machine with multiple tapes is functionally equivalent to an ordinary single tape Turing machine.

These ideas are well understood. In the sense that they form the formal foundation of the discipline, they are part of the core of computer science. Current robotic science lacks a comparable foundation; the field needs a unified theory in which meaningful statements can be made about the complexity of robotic tasks and the robot systems we build to complete these tasks.

Can we adapt standard models of computation to the robotics context? Unfortunately, these models are fundamentally ill suited for studying robotics problems. They assume that all of the relevant information is supplied ahead of time on the machine's tape. Sensing and uncertainty are central defining issues in robotics. This structure is destroyed by an *a priori* encoding of the problem on a machine's tape.

The aim of this paper is to develop a "sensor-centered" theory for analyzing and comparing robot systems. Our contribution is to develop such a theory more completely than in prior work and to illustrate its usefulness with examples.

B. Organization

Section II is a brief survey of related work. In Section III we give a basic problem definition. Our definition of robot dominance and its properties are in Section IV. Section V relates the continuous-time model we introduce in this paper to our prior work that models time as a sequence of discrete stages. We make concluding remarks and discuss open problems in Section VI. We illustrate with examples throughout.

II. RELATED WORK

We partially address issues issues of robot comparison and dominance in prior work [12], in which we establish a *dominance* relation over robot systems. That work has several important shortcomings that limit its applicability.

1) *Perfect control* – In [12], we assumed that the robot can execute all of its actions with perfect precision

1-4244-0602-1/07/$20.00 ©2007 IEEE.

and complete reliability. The motions of real robots are imprecise and unpredictable.

2) *Perfect sensing* – Although [12] accounts for the importance of sensing by assuming that the robot is uncertain of its current state and must rely on sensing, it assumes that sensor readings are uncorrupted by noise. A more realistic sensor model would allow information from sensors to be subject to error.
3) *Modeling of time* – In [12], time is managed in discrete *stages*. The robot makes a single decision at each stage. Continuous-time models have a more direct correspondence with reality.
4) *Fixed, known environment* – In [12], we assumed (tacitly) that the robot operates in a fixed, known environment. This assumption is unsatisfactory in all but the most structured contexts.
5) *Identical state spaces* – The dominance relation in [12] is only able to compare robots that share the same state space. To compare robots that are truly dissimilar, the framework must allow each robot to have a distinct state space.

In this paper, we present substantial revisions and extensions to the framework of [12] to remedy these shortcomings. These extensions illuminate several issues and subtleties not evident in the former paper.

Our goals are similar to those of Donald [4]. The reductions in that work are similar to our dominance relation; Donald's notion of calibration is related to our idea of initial conditions. The most fundamental difference is that our analysis is rooted in the information space. We claim that for robotic problems for which sensing is a crucial issue, the information space is the space in which the problem can most naturally be posed.

A third line of related research is the work of Erdmann [6], which is itself grounded in the preimage planning ideas due Lozano-Perez, Mason, and Taylor [10]. In Erdmann's work, sensors are modeled by giving a partition of state space. The problem of sensor design is choose a partition so that from each region in the partition, the robot knows what action to select in order to make progress toward its goal.

Others in artificial intelligence [2] and control theory [5], [7] have addressed related issues.

III. BASIC DEFINITIONS

This section contains basic definitions for planning with uncertainty in the robot's current state.

A. State spaces and environment spaces

The robot moves in a state space X, which must be sufficiently expressive to encode all of the relevant information about the condition of the world. In a simple case, X might be defined as the configuration space [11] of the robot in a certain environment. Time proceeds continuously starting at $t = 0$ and continuing indefinitely. The robot's state at time t is denoted $x(t)$.

What happens when the robot begins with limited or no knowledge about its environment, in the sense that positions and geometry of obstacles, map topology, navigability of terrain, and so on are unknown? Imperfect knowledge about the environment is a more drastic instance of the general issue of state uncertainty. If the state is defined to include a description of the environment in addition to the robot's configuration, then uncertainty in the environment can be represented as an additional dimension of state uncertainty.

Concretely, choose an *environment space* $\mathcal{E}$ of which each element $E \in \mathcal{E}$ is a potential environment for the robot. Possibilities for $\mathcal{E}$ (with varying degrees of realism, interest, practicality, and amenability to analysis), include:

1) the set of simple polygons in the plane, and
2) the set of compact regions in $\mathbb{R}^2$ or $\mathbb{R}^3$ with connected interiors and piecewise analytic boundaries.
3) the set of terrain maps from $\mathbb{R}^2$ to $\mathbb{R}$, giving the elevation or navigability at each point in the plane.

The state space is formed by combining the robot's configuration space $\mathcal{C}$ with $\mathcal{E}$, so that $X = \mathcal{C} \times \mathcal{E}$. In our models, the true environment $E \in \mathcal{E}$ affects the robot by influencing the state transitions that the robot makes and the observations that the robot receives.

B. Actions and transitions

The robot influences its current state by choosing actions from some action space U. At each instant t, the robot chooses some $u(t) \in U$. Let $\widetilde{U}_t$ denote the space of all functions from $[0, t)$ into U, and let $\widetilde{U} = \bigcup_{t \in [0,\infty)} \widetilde{U}_t$. For simplicity of notation, adopt the convention that $[0, 0) = \emptyset$. Define $\widetilde{u} : [0, \infty) \to U$ as the robot's complete action history, and let $\widetilde{u}_t \in \widetilde{U}$ denote the robot's action history up to (but exclusive of) time t.

We include a special *termination action* $u_T \in U$. The robot selects u_T to indicate that it has finished its task and intends to terminate execution. We require that if $u(t) = u_T$, then $u(t') = u_T$ for all $t' > t$.

How do these actions influence the state? We model disturbances and unexpected events as interference from a fictitious external decision maker we call "nature". Choices made by both the robot and by nature affect changes in the state. Let Θ denote a *nature action space*. Let $\widetilde{\Theta}_t$ denote the space of all functions mapping $[0, t)$ into Θ, and let $\widetilde{\Theta} = \bigcup_{t \in [0,\infty)} \widetilde{\Theta}_t$. Let $\widetilde{\theta} : [0, \infty) \to \Theta$ denote the complete history of nature actions and $\widetilde{\theta}_t \in \widetilde{\Theta}_t$ the nature action history up to (and including) t.

We describe changes in the state with a *state transition function*, $\Phi : X \times \bigcup_{t \in [0,\infty)} (\widetilde{U}_t \times \widetilde{\Theta}_t) \to X$. The intuition is that, given a starting state $x(0)$, and action histories $\widetilde{u}_t$ and $\widetilde{\theta}_t$ of equal duration for the robot and nature respectively, the state transition function computes the resulting state

$$x(t) = \Phi(x(0), \widetilde{u}_t, \widetilde{\theta}_t). \tag{1}$$

This notation of a "black box" state transition function follows notation employed in control theory, for example by Chen [3].

Example 1: A familiar special case of (1) occurs if $\widetilde{u}$ and $\widetilde{\theta}$ are smooth functions and there exists a function f such

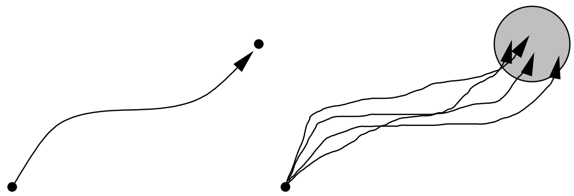

Fig. 1. [left] The robot in Example 2 gives velocity inputs that determine a nominal trajectory. [right] Nature interferes with this trajectory, but error bounds ensure that the final state is contained in a circle of radius $t\theta_{max}$.

that

$$\Phi(x(0), \widetilde{u}_t, \widetilde{\theta}_t) = x(0) + \int_0^t f(x(s), u(s), \theta(s)) \mathrm{d}s. \quad (2)$$

In this case, the system dynamics can be described by the differential equation $\dot{x} = f(x, u, \theta)$. □

Example 2: Consider a point in the plane with velocity input, for which the motion is subject to noise. Let u_{max} denote a bound on the magnitude of the commanded velocity, and let θ_{max} denote a bound on magnitude of the error in the velocity. Let $X = \mathbb{R}^2$, $U = \{u \in \mathbb{R}^2 \mid ||u|| \leq u_{max}\}$, $\Theta = \{\theta \in \mathbb{R}^2 \mid ||\theta|| \leq \theta_{max}\}$, and

$$\Phi(x(0), \widetilde{u}_t, \theta_t) = x(0) + \int_0^t (u(s) + \theta(s)) \mathrm{d}s. \quad (3)$$

At every time t, the robot can be certain that its state lies within a closed ball of radius $t\theta_{max}$, centered at the nominal (error free, i.e. $\widetilde{\theta} \equiv (0, 0)$) final point. See Figure 1.

□

C. Observations

As time passes, the robot's sensors provide feedback in the form of *observations* drawn from an observation space Y. Let $\widetilde{Y}_t$ denote the space of functions mapping $[0, t]$ into Y and let $\widetilde{Y} = \bigcup_{t \in [0,\infty)} \widetilde{Y}_t$. The robot's complete observation history is $\widetilde{y} : [0, \infty) \to Y$. The observation history up to t (inclusive) is $\widetilde{y}_t \in \widetilde{Y}_t$.

Nature interferes with the observations by choosing a *nature observation action* from a space Ψ. Let $\widetilde{\Psi}_t$ denote the space of functions mapping $[0, t)$ into Ψ and let $\widetilde{\Psi} = \bigcup_{t \in [0,\infty)} \widetilde{\Psi}_t$. The robot's complete nature observation action history is $\widetilde{\psi} : [0, \infty) \to \Psi$; the nature observation action history up to time (but not including) t is $\widetilde{\psi}_t \in \widetilde{\Psi}_t$. The observations received by the robot are governed by the *observation function* $h : X \times \Psi \to Y$.

Example 3: Suppose the mobile robot has a sensor that detects the distance to some landmark. Let $X = \mathbb{R}^2$ and $Y = \mathbb{R}$. Without loss of generality, position the landmark at the origin. Assume that the sensor has bounded additive disturbance, so that $\Psi = [-\psi_{max}, \psi_{max}]$ and $h(x, \psi) = ||x|| + \psi$. See Figure 2. At each instant, the robot knows with certainty that its state is within an annulus of width $2\psi_{max}$ centered at the origin. □

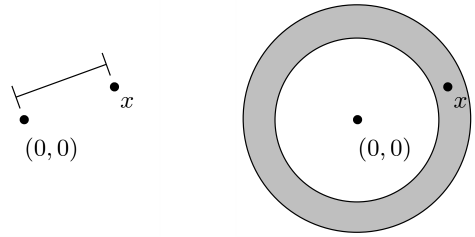

Fig. 2. [left] The robot in Example 3 has a sensor that reports a noisy estimate of the distance to the origin. [right] Accounting for noise bounded by ψ_{max}, the observation confines the robot's state to an annulus of width $2\psi_{max}$.

D. Information spaces and information mappings

To inform its decisions, the robot has access only to the histories of actions it has selected and observations it has received so far. That is, to select $u(t)$, the robot can use $\widetilde{u}_t$ and $\widetilde{y}_t$. This motivates our definition of the *history information space*, $\mathcal{I}_{hist} = \bigcup_{t \in [0,\infty)} \widetilde{U}_t \times \widetilde{Y}_t$. The tuple $\eta(t) = (\widetilde{u}_t, \widetilde{y}_t) \in \mathcal{I}_{hist}$ containing the robot's action and sensing histories is the robot's *history information state*.

The history information state, since it is composed of functions of time, is unwieldy in isolation. As a result, we select a *derived information space* $\mathcal{I}$ and an information mapping $\kappa : \mathcal{I}_{hist} \to \mathcal{I}$. Informally, a derived information space represents "compression" or "interpretation" of the histories.

We say that a state x is *consistent* with an information state $\eta(t) = (\widetilde{u}_t, \widetilde{y}_t)$ if and only if there exists some starting state $x(0)$ and nature histories $\widetilde{\theta}$ and $\widetilde{\psi}$ such that $\Phi(x(0), \widetilde{u}_t, \widetilde{\theta}_t) = x$ and $h(x(t'), \widetilde{\psi}(t')) = y(t')$ for $t' < t$. The next example is an information mapping that arises directly from the notion of consistent states.

Example 4 (Nondeterministic information mapping): Let $\mathcal{I}_{ndet} = \mathrm{pow}(X) - \emptyset$. The relevant information mapping is $\kappa_{ndet} : \mathcal{I}_{hist} \to \mathcal{I}_{ndet}$, under which each history information state maps to the minimal subset of X consistent with it. The intuition is that $\eta(t)$ gives a set of "possible states" for the robot at time t. □

E. Information feedback strategies

How does the robot decide which actions to select? We describe the robot's strategy as a feedback strategy $\pi : \mathcal{I}_{hist} \to U$ that specifies an action for history information state. As the robot executes π, the actions are given by $u(t) = \pi(\eta(t))$. We call π an *information feedback strategy*.

Even though we define π as a feedback strategy over the history information space, the next two examples illustrate that feedback over a derived information space can sometimes be a natural way to express familiar kinds of strategies.

Example 5 (Open loop strategy): Let $\mathcal{I}_{time} = [0, \infty)$ and consider the information map $\kappa_{time}(\eta(t)) = t$. In this case, the derived information state is simply the time elapsed. Then if the robot has an intended open loop action trajectory $\omega : [0, t_f) \to U$, a strategy to execute γ is $\pi(\eta(t)) = \omega(\kappa_{time}(\eta(t))$ if $t < t_f$ and $\pi(\eta(t)) = u_T$ otherwise. □

Example 6 (Memoryless strategy): Another possibility is that it is enough to know the "most recent" observation, so $\mathcal{I}_{obs} = Y$ and $\kappa_{obs}(\eta(t)) = y(t)$. Given a memoryless plan $\gamma : Y \to U$, the composed function $\kappa_{obs} \circ \gamma : \mathcal{I}_{hist} \to U$ is a memoryless information feedback strategy.[1] □

We assume that a given strategy is executed until it selects u_T. The time when this occurs, the resulting final state, and the observations received along the way are all affected by the strategy itself π, the starting state $x(0)$, and the actions of nature $\widetilde{\theta}$ and $\widetilde{\psi}$. Assuming that the robot executes π, the termination time is $T(\pi, x(0), \widetilde{\theta}, \widetilde{\psi}) = \inf\{t \in [0, \infty) \mid \pi(\eta(t)) = u_T\}$, and the final state is $F(\pi, x(0), \widetilde{\theta}, \widetilde{\psi}) = \Phi(x(0), \widetilde{u}_{t_f}, \widetilde{\theta}_{t_f})$, in which $t_f = T(\pi, x(0), \widetilde{\theta}, \widetilde{\psi})$.

Example 7 (Concatenating strategies): Given two strategies π_1 and π_2, a new strategy that concatenates them (that is, executes them in sequence) is expressed by $\pi(\eta(t)) = \pi_1(\eta(t))$ if $\pi_1(\eta(t)) \neq u_T$ and $\pi(\eta(t)) = \pi_2(\eta(t))$ otherwise. By nesting this construction, arbitrarily many strategies can be chained together. □

F. Tasks and solutions

A *task* (or problem) is defined by a goal region $\mathcal{I}_G \subset \mathcal{I}_{hist}$ in history information space. This notion is a generalization of the traditional idea of a goal state or goal region in state space. An information feedback plan is a *guaranteed solution* to a problem if there exists some time t_g such that, for any $\widetilde{\theta}_{t_g}$ and any $\widetilde{\psi}_{t_g}$, $\eta(t_g) \in \mathcal{I}_G$.

IV. COMPARING ROBOT SYSTEMS

In this section, we show that the basic results of [12] still hold in our generalized framework. We define a dominance relation between robot systems to formalize the informal idea that some robots are "more powerful" than others, in the sense of having richer sensing and motion abilities. This relation has direct implications on the ability of robot systems to complete tasks.

A. Information preference relation

The first ingredient we need is some notion of when one derived information state is "better than" another. Fix a derived information state $\mathcal{I}$ and an information mapping $\kappa : \mathcal{I}_{hist} \to \mathcal{I}$. Equip $\mathcal{I}$ with a partial order *information preference relation* $\preceq$, under which $\eta_1 \preceq \eta_2$ means that η_2 is "more informed than" η_1. The only constraint on $\preceq$ is that is must be a partial order satisfying the following consistency property for any $t \in [0, \infty)$, $\widetilde{u}_t \in \widetilde{U}_t$ and $\widetilde{y}_t \in \widetilde{Y}_t$:

$$\kappa(\eta_1) \preceq \kappa(\eta_2) \implies \kappa(\eta_1, \widetilde{u}_t, \widetilde{y}_t) \preceq \kappa(\eta_2, \widetilde{u}_t, \widetilde{y}_t), \quad (4)$$

in which the concatenation on the right side indicates that the additional history information from $\widetilde{u}_t$ and $\widetilde{y}_t$ is appended to η_1 and η_2. The intuition is that information preference must be preserved if the same actions are selected and the same observations received from both η_1 and η_2.

Example 8: Recall κ_{ndet} from Example 4. Define $\preceq$ so that $\eta_1 \preceq \eta_2$ if and only if $\kappa_{ndet}(\eta_2) \subseteq \kappa_{ndet}(\eta_1)$. It is easy to verify that the consistency property holds. □

B. Definition of dominance

Our goal is a formal way to compare the power of robot systems. Consider two robot systems R_1 and R_2 defined as in Section III:

$$R_1 = (X^{(1)}, U^{(1)}, Y^{(1)}, \Theta^{(1)}, \Psi^{(1)}, \Phi^{(1)}, h^{(1)}) \quad (5)$$
$$R_2 = (X^{(2)}, U^{(2)}, Y^{(2)}, \Theta^{(2)}, \Psi^{(2)}, \Phi^{(2)}, h^{(2)}) \quad (6)$$

Because $U^{(1)}$ need not have any special relationship to $U^{(2)}$, and likewise $Y^{(1)}$ need not be related to $Y^{(2)}$, the comparison cannot be made directly in the history information space, which simply records actions and observations. Instead, map the two history information spaces to the same derived information space. The corresponding information mappings are $\kappa^{(1)} : \mathcal{I}^{(1)}_{hist} \to \mathcal{I}$ and $\kappa^{(2)} : \mathcal{I}^{(2)}_{hist} \to \mathcal{I}$.

To compare distinct robot systems (perhaps with distinct state spaces) operating in the same family of environments, use the environment space construction described in Section III-A with R_1 and R_2 in the same environment space, so that $X^{(1)} = \mathcal{C}^{(1)} \times \mathcal{E}$ and $X^{(2)} = \mathcal{C}^{(2)} \times \mathcal{E}$.

Now we can state the dominance relation between robot systems.

Definition 1 (Robot dominance): Consider two robots R_1 and R_2. If, for all $\eta^{(1)}(t_1) \in \mathcal{I}^{(1)}_{hist}$, $\eta^{(2)}(t_2) \in \mathcal{I}^{(2)}_{hist}$ with $\kappa^{(1)}(\eta^{(1)}(t_1)) \preceq \kappa^{(2)}(\eta^{(2)}(t_2))$, $t'_1 \in [0, \infty)$, and $\widetilde{u}^{(1)}_{t'_1} \in \widetilde{U}^{(1)}_{t'_1}$, there exists an information feedback strategy $\pi_2 : \mathcal{I}^{(2)}_{hist} \to U^{(2)}$, such that for all $x^{(1)} \in X^{(1)}$ consistent with $\eta^{(1)}(t_1)$ and $x^{(2)} \in X^{(2)}$ consistent with $\eta^{(2)}(t_2)$, there exists $t'_2 \in [0, \infty)$ such that for all $\widetilde{\theta}^{(1)}_{t'_1} \in \Theta^{(1)}_{t'_1}$, $\widetilde{\theta}^{(2)}_{t'_2} \in \Theta^{(2)}_{t'_2}$, $\widetilde{\psi}^{(1)}_{t'_1} \in \widetilde{\Psi}^{(1)}_{t'_1}$, and $\widetilde{\psi}^{(2)}_{t'_2} \in \widetilde{\Psi}^{(2)}_{t'_2}$, if R_1 executes $\widetilde{u}^{(1)}_{t'_1}$ from time t_1 to t'_1 and R_2 executes $\pi^{(2)}$ from time t_2 to t'_2, we have

$$\kappa(\eta^{(1)}(t'_1)) \preceq \kappa(\eta^{(2)}(t'_2)) \quad (7)$$

then R_2 *dominates* R_1 under $\mathcal{I}$, κ, and $\preceq$, denoted $R_1 \trianglelefteq R_2$. If $R_1 \trianglelefteq R_2$ and $R_2 \trianglelefteq R_1$, then R_1 and R_2 are *equivalent*, denoted $R_1 \equiv R_2$. If $R_1 \ntrianglelefteq R_2$ and $R_2 \ntrianglelefteq R_1$ then R_1 and R_2 are *incomparable*, denoted $R_1 \boxtimes R_2$. ∘

Informally, Definition 1 means that, regardless of the transitions made by R_1 (and regardless of the interference from nature R_1 receives), there exists some strategy for R_2 to reach an information state at least as good, in the sense of information preference, as that reached by R_1. Figure 3 illustrates this intuition.

C. Dominance and solvability

Now we can establish the relationship between dominance and solvability. Fist, we define a class of "well-formed" tasks based on the information preference relation.

[1] In [12], we used a slightly different observation model, in which $h : X \times U \to Y$. In this context, the time period over which observations are available is the half-open interval $[0, t)$; $\widetilde{y}_t$ is undefined at t itself. As a result, the closest we can come to a memoryless strategy is to use the left-hand limit of $\widetilde{y}_t$ at t, $\kappa_{obs}(\eta(t)) = \lim_{t' \to t^-} y(t')$, provided the limit exists. This technicality is part of the motivation for preventing y from depending directly on u, as we have done in this paper.

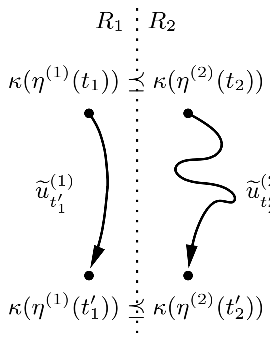

Fig. 3. An illustration of Definition 1.

Definition 2: Consider a set $I \subset \mathcal{I}$ of derived information states. If, for any $\eta_1 \in I$ and $\eta_2 \in \mathcal{I}$ with $\eta_1 \preceq \eta_2$, we have $\eta_2 \in I$, then I is *preference closed.* ○

For any preference closed goal region, we have the following result. A similar, but weaker (because of the limitations in robot models) result appeared in [12].

Lemma 1 (Solution by imitation): Consider two robot systems R_1 and R_2 with $R_1 \trianglelefteq R_2$ and a preference-closed goal region $\mathcal{I}_G$. If there exists a guaranteed solution for R_1 to reach $\mathcal{I}_G$, then also there exists a guaranteed solution for R_2 to reach $\mathcal{I}_G$.

D. Dominance examples

This section presents a few examples to illustrate the implications of Definition 1.

Example 9 (Omniscient sensing and perfect control): Consider a degenerate case with $Y = X$, and $h(x, \psi) = x$. Let $\Theta = \Psi = \{0\}$ be dummy singleton sets with no effect on state transitions or observations. This situation gives the robot perfect control and complete information about its state. Choose $\kappa(\eta(t)) = y(t) = x(t)$. Let $\eta_1 \preceq \eta_2$ if and only if $\eta_1 = \eta_2$. In this context, Definition 1 becomes a statement about the regions of state space reachable by different control systems.

Suppose three such systems R_1, R_2, and R_3 differ only in their action spaces $U^{(1)}$, $U^{(2)}$, and $U^{(3)}$. Let $Z(A)$ denote the subset of state space reachable by a robot with action space A. Suppose $R_1 \trianglelefteq R_2$. R_3 need not be comparable to either R_1 or R_2. Note that additional robot models can be constructed from unions of $U^{(1)}$, $U^{(2)}$ and $U^{(3)}$. We have the following results:

$$Z(U^{(1)}) \subseteq Z(U^{(2)} \cup U^{(3)}) \tag{8}$$

$$Z(U^{(1)}) = Z(U^{(1)} \cup U^{(2)}) \tag{9}$$

$$Z(U^{(1)} \cup U^{(3)}) \subseteq Z(U^{(2)} \cup U^{(3)}) \tag{10}$$

These results are somewhat analogous to Lemmas 2-4 in [12]. Note that in combining action spaces in this way, we allow the robot to choose *sequentially* the action set from which to choose its action. The results fail if the robot is somehow allowed to choose actions from each constituent set in parallel. □

Example 10 (Varying error bounds): Recall the incompletely specified models in Examples 2 and 3. Consider two robot systems R_1 and R_2 with state transitions as in Example 2 and observations as in Example 3; R_1 and R_2 differ only in the error bounds $\theta^{(1)}_{max}$, $\psi^{(1)}_{max}$, $\theta^{(2)}_{max}$, and $\psi^{(2)}_{max}$. We will compare these robots under κ_{ndet}.

Comparing $\theta^{(1)}_{max}$ to $\theta^{(2)}_{max}$, and $\psi^{(1)}_{max}$ to $\psi^{(2)}_{max}$, there are three cases:

1) If $\theta^{(1)}_{max} \le \theta^{(2)}_{max}$ and $\psi^{(1)}_{max} \le \psi^{(2)}_{max}$, then $R_2 \trianglelefteq R_1$.
2) If $\theta^{(2)}_{max} \le \theta^{(1)}_{max}$ and $\psi^{(2)}_{max} \le \psi^{(1)}_{max}$, then $R_1 \trianglelefteq R_2$.
3) If $\theta^{(1)}_{max} \le \theta^{(2)}_{max}$ and $\psi^{(2)}_{max} \le \psi^{(1)}_{max}$ or $\theta^{(2)}_{max} \le \theta^{(1)}_{max}$ and $\psi^{(1)}_{max} \le \psi^{(2)}_{max}$, then $R_2 \boxtimes R_1$.

This implies that $\theta^{(1)}_{max} = \theta^{(2)}_{max}$ and $\psi^{(1)}_{max} = \psi^{(2)}_{max}$ if and only if $R_1 \equiv R_2$. These results follow in a straightforward manner from Definition 1. The intuition of this (perhaps unsurprising) example is that one robot system dominates the other if its error bounds are smaller. □

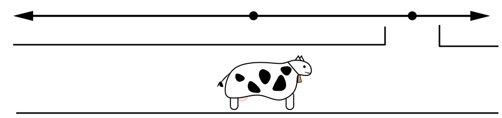

Fig. 4. The lost cow of Example 11 searching for a gate.

Example 11 (A Lost Cow): A well-known problem in online algorithms is the *lost cow problem* [1], [9] in which a near-sighted cow moves along a fence searching for a gate, as illustrated in Figure 4. The difficulty under the standard sensing model is that the cow must systematically search in both directions from its initial position without any information about the distance or direction to the gate. The interest in this problem derives from potential applications in (or at least the potential for better understanding of) exploration in unbounded environments.

We formulate the lost cow problem and consider how the sensing model affects the cow's searching ability. Let $X = \mathbb{R}$, in which $x(t)$ is the position of the gate relative to the cow at time t. For simplicity, assume perfect control and perfect sensing by setting $\Theta = \Psi = \{0\}$. The action space is $U = [-1, 1]$, with $\Theta = \{0\}$ and $\Phi(x(0), \widetilde{u}_t, \theta_t) = x(0) + \int_0^t u(s)\mathrm{d}s$. We compare three distinct models C_1, C_2, and C_3 under κ_{ndet}.

1) C_1: Let $Y^{(1)} = \mathbb{R}$ and $h^{(1)}(x, \psi) = x$. Here the cow can determine both the direction and distance to the gate.
2) C_2: Let $Y^{(2)} = \{-1, 0, 1\}$ and $h(x, \psi) = \mathrm{sign}(x)$. This allows the cow to determine the direction it must move to reach the gate, but not the distance.
3) C_3: Let $Y^{(3)} = \{0, 1\}$ and $h^{(2)}(x, \psi) = 1$ if $x = 0$ and $h^{(2)}(x, \psi) = 1$ otherwise. This is the standard lost cow sensing model, in which the cow cannot see the gate from a distance, but can detect the gate when it arrives.

Perhaps surprisingly, these three models are equivalent in the sense of Definition 1. This comes about as a result of the fact that each can eventually determine its state (by finding the gate) and after the state is known, the state uncertainty cannot recur. To simulate C_1 with C_3, first execute the algorithm of [1], then move to the state occupied by C_1. □

V. A DISCRETE-STAGE MODEL

This section describes how the continuous-time model given in Section III is related to the discrete-stage formulation of [12].

A. Transforming from continuous time to discrete stages

Consider a division of time into variable length stages, in which, in each stage, the robot executes a single information feedback strategy to completion. We require of each of these strategies the following special property:

Definition 3 (History invariance): If, for all $\eta(t) \in \mathcal{I}_{hist}$, all $x \in X$ consistent with $\eta(t)$, all $\widetilde{\theta} \in \widetilde{\Theta}$, all $\widetilde{\psi} \in \widetilde{\Psi}$, and all $y(0) \in Y$, we have

$$F(\pi, x, \eta(t), \widetilde{\theta}, \widetilde{\psi}) = F(\pi, x, \eta(0), \widetilde{\theta}, \widetilde{\psi}), \qquad (11)$$

then π is a *history-invariant* strategy. ◦

The intuition of the definition is that the robot executing π is free to use the observation and action history generated during its own execution, but it cannot peer into the past before its execution began in order to make decisions.

Given a continuous-time robot system $R = (X, U, Y, \Theta, \Psi, \Phi, h)$ as in Section III and a set Π of history-invariant information feedback strategies, construct a discrete-stage system $\overline{R} = (X, \overline{U}, \overline{Y}, \overline{\Theta}, \overline{\Psi}, \overline{f}, \overline{h})$ as follows:

1) The state space X is unchanged.
2) The action space is $\overline{U} = \Pi$.
3) The observation space is $\overline{Y} = \widetilde{Y}$.
4) The nature action space is $\overline{\Theta} = \widetilde{\Theta}$.
5) The nature observation action space is $\overline{\Psi} = \widetilde{\Psi}$.
6) The state transition function is $f : X \times \overline{U} \to X$, with $f(x, \pi) = F(\pi, x, \widetilde{\theta}, \eta(0))$.
7) The observation function is $h : X \times \overline{U} \times \overline{\Psi} \to \overline{Y}$.

The system starts at some (unknown) initial state $x_1 \in X$. Let $x_k \in X$, $u_k \in \overline{U}$, $y_k \in \overline{Y}$, $\theta \in \overline{\Theta}$, and $\psi_k \in \overline{\Psi}$ denote the appropriate values at stage k. These sequences are related to each other by $x_{k+1} = f(x_k, u_k, \theta_k)$ and $y_k = h(x_k, u_k, \psi_k)$. The history information state consists of the action and observation histories: $\eta_k = (u_1, y_1, \ldots, u_{K-1}, y_{K-1})$. We now argue that this discretized system faithfully represents the underlying continuous-time system.

Lemma 2: Any action sequence $u_1, \ldots, u_K$ executed by $\overline{R}$ reaches the same final state x and the analogous final history information state as does R.

Note, however, that in making this transformation, we may restrict the space of strategies that the robot can employ. If $\overline{U}$ does not contain a sufficiently rich selection of information feedback strategies, there may be regions of information space that are no longer reachable under the discretized model. It remains an open problem to find small (or at least succinctly described) sets of strategies that are complete or nearly complete in the sense of not eliminating any reachable regions in information space.

VI. CONCLUSION

Although the results we present here are a substantial improvement over those of [12], there are still important pieces missing.

A. Computational issues

We have focused mostly on the sensing and motion requirements of tasks. An important related question is to determine the kinds of computation power these tasks require. What are the tradeoffs between computation time, memory usage, sensing requirements and solution quality? Is there a satisfactory way to scalarize these competing objectives into a single-valued objective function, or should we expect a single problem will lead to many different Pareto optimal solutions?

B. Reductions and decision problems

One of the most powerful ideas in the theory of computation that we have not explored here is the idea of *reductions*, which hold promise for comparing robotic problems themselves. The resulting statements would have the form "Problem A is at least as hard as Problem B." To make things more concrete, we might consider *decision problems*, in which the robot must determine if its environment $E \in \mathcal{E}$ has a certain property. Such problems fit naturally as planning problems in information space.

REFERENCES

[1] R. A. Baeza-Yates, J. C. Culberson, and G. J. E. Rawlins, "Searching in the plane," *Information and Computation*, vol. 106, pp. 234–252, 1993.

[2] R. I. Brafman, J. Y. Halpern, and Y. Shoham, "On the knowledge requirements of tasks," *Artificial Intelligence*, vol. 98, no. 1-2, pp. 317–349, 1998.

[3] C.-T. Chen, *Linear System Theory and Design*. New York: Holt, Rinehart, and Winston, 1984.

[4] B. R. Donald, "On information invariants in robotics," *Artificial Intelligence*, vol. 72, no. 1-2, pp. 217–304, 1995.

[5] M. Egerstedt, "Motion description languages for multi-modal control in robotics," in *Control Problems in Robotics*, ser. Springer Tracts in Advanced Robotics, A. Bicchi, H. Cristensen, and D. Prattichizzo, Eds., 2002, pp. 75–90.

[6] M. A. Erdmann, "Understanding action and sensing by designing action-based sensors," *International Journal of Robotics Research*, vol. 14, no. 5, pp. 483–509, 1995.

[7] A. Girard and G. J. Pappas, "Approximation metrics for discrete and continuous systems," *IEEE Transactions on Automatic Control*, Mar. 2005, to appear.

[8] J. E. Hopcroft, J. D. Ullman, and R. Motwani, *Introduction to Automata Theory, Languages, and Computation*, 2nd ed. Reading, MA: Addison-Wesley, 2000.

[9] M.-Y. Kao, J. H. Reif, and S. R. Tate, "Searching in an unknown environment: An optimal randomized algorithm for the cow-path problem," in *SODA: ACM-SIAM Symposium on Discrete Algorithms*, 1993, pp. 441–447.

[10] T. Lozano-Pérez, M. T. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.

[11] T. Lozano-Pérez and M. A. Wesley, "An algorithm for planning collision-free paths among polyhedral obstacles," *Communications of the ACM*, vol. 22, no. 10, pp. 560–570, 1979.

[12] J. M. O'Kane and S. M. LaValle, "On comparing the power of mobile robots," in *Proc. Robotics: Science and Systems*, 2006.