Dominance and equivalence for sensor-based agents[*]

Jason M. O'Kane and **Steven M. LaValle**
Department of Computer Science
University of Illinois
Urbana, IL 61821
{jokane,lavalle}@cs.uiuc.edu

Abstract

This paper describes recent results from the robotics community that develop a theory, similar in spirit to the theory of computation, for analyzing sensor-based agent systems. The central element to this work is a notion of dominance of one such system over another. This relation is formally based on the agents' progression through a derived information space, but may informally be understood as describing one agent's ability to "simulate" another. We present some basic properties of this dominance relation and demonstrate its usefulness by applying it to a basic problem in robotics. We argue that this work is of interest to a broad audience of artificial intelligence researchers for two main reasons. First, it calls attention to the possibility of studying belief spaces in way that generalizes both probabilistic and nondeterministic uncertainty models. Second, it provides a means for evaluating the information that an agent is able to acquire (via its sensors and via conformant actions), independent of any optimality criterion and of the task to be completed.

Introduction

Across the artificial intelligence literature, there is strong interest in understanding and reasoning about uncertainty. Such issues are especially important in robotics. Actuators are often subject to noise. Sensors are generally error-prone and provide only limited information. Useful robots must overcome these obstacles by making good use of the information available to them, and by making good decisions in spite of the incompleteness of this information.

Unfortunately, sensing and movement capabilities are often available only at significant cost. This reality motivates us determine what information is truly necessary to complete a particular task. We believe that understanding a task's information requirements is an important part of understanding the task itself. More practically, this insight would also allow robot designers to make more informed decisions about the tradeoffs associated with the inclusion or exclusion of each sensor.

[*]This work is supported by DARPA award #HR0011-07-1-0002.

To that end, we are working to develop a theory for analyzing robot systems, including assessing the solvability of various problems with various robot systems, measuring the complexity (execution time, computation time, or the use of some other resource) of those solutions, and finding meaningful notions and comparison and equivalence between apparently dissimilar robot systems. The primary contribution we present in this paper is a formulation of *dominance* of one robot system over another. Our inquiry is inspired in part by the theory of computation, which provides similar machinery for analyzing abstract computing machines. For example, the hierarchy of robot systems induced by our dominance relation is functionally similar to the hierarchy of degrees induced by Turing reductions.[1]

This work originally appeared in the robotics community (O'Kane & LaValle 2006a; 2007; 2006b). However, the ideas are equally applicable to any agent that must explicitly collect information to inform its decisions. In the remainder of this paper, we review our results and discuss their relevance to a general AI audience. For greater technical detail, we direct the interested reader to the original publications.

Summary of results

In this section, we briefly summarize the previously published results.

Basic definitions Place a robot in a *state space* X, which need not be finite or even countable. Time proceeds in discrete stages indexed with consecutive integers starting with 1, and the robot's state at stage k is denoted x_k. At each state, the robot executes an action from its *action space* U. Let u_k denote the robot's action at stage k. The robot's state changes according to a *state transition function* $f : X \times U \to X$, so that $x_{k+1} = f(x_k, u_k)$. The state is unknown to the robot, but at each stage k, the robot receives an observation y_k from its *observation space* Y.

[1]For a general introduction to the theory of computation, see (Sipser 1997) or (Hopcroft, Motwani, & Ullman 2007). Details on Turing degrees appear in (Soare 1987).

The observations are governed by an observation function $h : X \times U \to Y$, under which $y_k = h(x_k, u_k)$.

Note that this model includes several simplifying assumptions. The robot's transitions and its observations are deterministic. Time is modelled discretely, rather than as a continuum. The results we present remain valid but are more notationally cumbersome if these assumptions are removed.

Information spaces What information is available to the robot as it chooses actions to execute? Other than any *a priori* knowledge it may have, the robot must make its decisions based only on the history of actions it has executed and the corresponding history of observations it has received. This fact leads us to define the robot's *history information state*[2] (history I-state) η_k, formed from these two histories:

$$\eta_k = (u_1, y_1, \ldots, u_k, y_k).$$

The *history information space* (history I-space) $\mathcal{I}_{hist}$ is the set of all history information states.

Unfortunately, history I-states are unwieldy (since they grow with each passing stage) and generally uninformative (since they merely record the raw histories). This motivates us to choose a *derived information space* (derived I-space) $\mathcal{I}$ and an information mapping (I-map) $\kappa : \mathcal{I}_{hist} \to \mathcal{I}$. Informally, the derived I-space represents an "interpretation" or "compression" of the history I-state. In this context, a *task* for the robot is a goal region $\mathcal{I}_G \subset \mathcal{I}$ and a *plan* is a function from $\mathcal{I}$ to U that leads the robot into $\mathcal{I}_G$. Although, in principle, $\mathcal{I}$ may be an arbitrary set, we emphasize that selecting $\mathcal{I}$ is a crucial modelling choice, because it affects the tasks and plans that can be expressed, and strongly influences the dominance relation we state later.

One important example derived I-space is the nondeterministic information space $\mathcal{I}_{ndet} = \mathrm{pow}(X)$, in which the derived I-state $\kappa(\eta_k)$ is a minimal subset of X guaranteed to contain the robot's true state. This corresponds to worst case or set-based reasoning about uncertainty. Another reasonable choice is to let $\mathcal{I}$ be a set of probability distributions over X, and choose κ to compute the posterior distribution given η_k. A more complete discussion of information spaces and their use in robotics appears in Chapters 11 and 12 of (LaValle 2006).

Comparing robot systems The central contribution of this work is a definition of *dominance* of one robot system over another. The intuition of the definition is that one robot system (as defined above) dominates another if the former can "simulate" the latter.

To formalize this notion, equip $\mathcal{I}$ with a partial order *information preference relation* $\preceq$, under which we write $\kappa(\eta_1) \preceq \kappa(\eta_2)$ to indicate that $\kappa(\eta_2)$ is "more informed"

[2]Equivalent terms include *belief state*, *knowledge state*, and *hyperstate*.

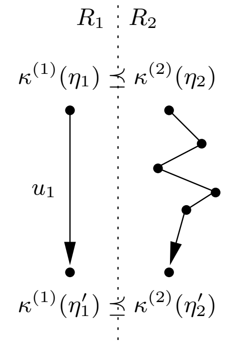

Figure 1: A illustration of Definition 1. Robot R_2 dominates robot R_1 if R_2 can always reach an information state at least as good as the one reached by R_1.

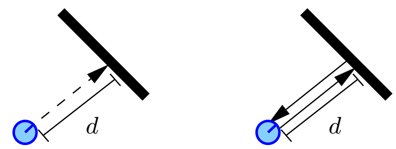

Figure 2: Comparing a robot, R_1, equipped with a unidirectional range sensor to another robot R_2, equipped with perfect odometry. R_2 dominates R_1 because the former can simulate the latter. [left] A distance measurement made directly by R_1. [right] Distance is measured indirectly by R_2 using its odometer.

than $\kappa(\eta_1)$. We require only that preference is preserved across transitions, so that for any $\eta_1, \eta_2 \in \mathcal{I}_{hist}$, $u \in U$ and $y \in Y$, we have

$$\kappa(\eta_1) \preceq \kappa(\eta_2) \implies \kappa(\eta_1, u, y) \preceq \kappa(\eta_2, u, y).$$

For example, in $\mathcal{I}_{ndet}$, one might define $\kappa(\eta_1) \preceq \kappa(\eta_2)$ if and only if $\kappa(\eta_2) \subseteq \kappa(\eta_1)$.

Now we can define the dominance relation.

Definition 1 (Robot dominance) *Consider two robots R_1 and R_2 along with a derived I-space $\mathcal{I}$ and two I-maps $\kappa^{(1)} : \mathcal{I}^{(1)}_{hist} \to \mathcal{I}$ and $\kappa^{(2)} : \mathcal{I}^{(2)}_{hist} \to \mathcal{I}$. If, for all $\eta_1 \in \mathcal{I}^{(1)}_{hist}$, $\eta_2 \in \mathcal{I}^{(2)}_{hist}$ for which $\kappa^{(1)}(\eta_1) \preceq \kappa^{(2)}(\eta_2)$, and all $u_1 \in U^{(1)}$, there exists a policy $\pi_2 : \mathcal{I}^{(2)}_{hist} \to U^{(2)}$ for R_2 that, in a finite number of stages, reaches an information state η_2', such that*

$$\kappa^{(1)}(\eta_1, u_1, h^{(1)}(x_1, u_1)) \preceq \kappa(\eta_2')$$

then R_2 dominates R_1, 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 $R_1 \equiv R_2$.

Figure 1 illustrates the definition. Figure 2 shows a simple example of dominance, in which a range sensor is simulated by odometry.

Note that the definition depends heavily on using the same derived I-space for both robots. As a result, the choices one makes for $\mathcal{I}$, $\kappa^{(1)}$, and $\kappa^{(2)}$ are extremely important, to the extent that different choices will generate different dominance hierarchies. Observe also that the definition is concerned only with which derived I-states can be reached, rather than with the cost or efficiency of reaching those I-states; it compares *feasibility* rather than *optimality*. Finally, note that this dominance relation is task independent; it makes no mention of a goal region, but instead considers reachability across all of $\mathcal{I}$.

We now present a few basic results related to Definition 1, starting with the connection between dominance and solvability.

Lemma 2 *Suppose $R_1 \trianglelefteq R_2$. With certain technical conditions*[3]*, if R_1 can reach a given goal region $\mathcal{I}_G \subset \mathcal{I}$, then R_2 can also reach $\mathcal{I}_G$.*

Corollary 3 *If $R_1 \equiv R_2$, then with certain technical conditions, R_1 and R_2 can complete the precisely the same tasks.*

Definition 1 applies to arbitrary pairs of robot models. Alternatively, we can form robot models as collections of *robotic primitives*, each of which represents a self-contained "instruction set" for the robot that may involve sensing, motion, or both. In this context, a robot model is a set of primitives, so that ordinary set operations like union and intersection apply normally. The following results show how the dominance relation forms a hierarchy over the set of robot models formed in this way.

Lemma 4 *For any three robots R_1, R_2 and R_3 for which $R_1 \trianglelefteq R_2$:*

(a) $R_1 \trianglelefteq R_1 \cup R_3$ *(Adding primitives never hurts)*
(b) $R_2 \equiv R_2 \cup R_1$ *(Redundancy doesn't help)*
(c) $R_1 \cup R_3 \trianglelefteq R_2 \cup R_3$ *(No unexpected interactions)*

These results may be reminiscent of the axioms of rationality used in utility theory (Robert 2001; DeGroot 1970).

Active global localization To demonstrate the usefulness of our ideas, we applied them to the problem of active global localization, in which a robot with has an accurate map of its environment but no knowledge of its location within that environment. The task is to choose motions and sensing commands in order to eliminate this uncertainty. What sensors are required to solve this problem?

Using the concepts described above, we give a partial answer. Consider four very simple sensors: a compass (which reports the robot's orientation relative to a global reference direction), an angular odometer (which allows rotations relative to the robot's current orientation), a linear odometer (which measures the robot's translations), and a contact sensor (which reports when the robot comes in contact with the boundary). The 15 nonempty combinations of these sensors generate 15 distinct robot models. We grouped these models into 8 equivalence classes, found the dominance hierarchy, and determined the solvability of the localization problem for the robots in each equivalence class. See Figure 3.

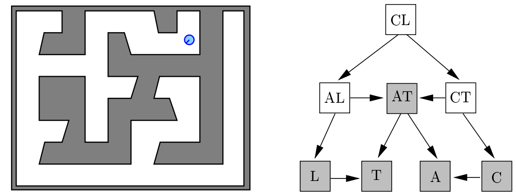

Figure 3: [left] A robot localizing itself within its environment. Localization is a fundamental problem for mobile robots, but relatively little had been known about what sensing this problem truly requires. [right] A hierarchy of simple robot models, variously equipped with compasses (C), angular odometers (A), linear odometers (L), and contact sensors (T). Arrows indicate dominances. Only the unshaded models can solve the active global localization problem.

Connections within AI

In this section, we discuss connections from our work to problems closer to the core of AI.

General information spaces In recent years, probabilistic methods for reasoning about uncertainty have gained significant momentum (Kaelbling, Littman, & Cassandra 1998; Zhang & Lin 1997). Such methods offer robustness, but also carry the heavy burden of modelling or learning the relevant probability models. In contrast, situations such as conformant planning (Cimatti, Roveri, & Bertoli 2004), adversarial games (Russell & Wolfe 2005), or nondeterministic domains (Amir & Russell 2003) are better suited to crisper set-based representations. Other lines of research (that focus more on reasoning than on acting) use nonmonotonic (McDermott & Doyle 1980) or modal logics (Blackburn, van Benthem, & Wolter 2006) to deal with incomplete or uncertain information.

It seems likely that general purpose autonomous agents will need to employ some combination of these techniques, depending on the situation. Is there any common ground between these methods that appear on the surface to be so dissimilar? We claim that the answer, at least on a fundamental level, is yes. The history I-state fully expresses the information available to any agent that interacts with the world by executing actions and receiving observations. The general approaches mentioned above differ only in how the history

[3] Precisely, we require that $\mathcal{I}_G$ be *preference closed*. A set $s \subseteq \mathcal{I}$ is preference closed if, for any $\eta_1 \in \mathcal{I}$ and $\eta_2 \in s$ with $\eta_1 \preceq \eta_2$, we also have $\eta_1 \in s$.

I-state is interpreted. The I-map κ represents precisely this interpretation: It maps from the history of actions and observations into another space (the derived I-space) in which the implications of these histories is more clear.

In this regard, the novelty of our work is that it does not assume any particular I-map or derived I-space. Instead, we take an axiomatic approach, stating conditions on $\mathcal{I}$, κ, and $\preceq$ under which certain results hold. The advantage of this kind of I-space centered approach is that (ideally) one can seek results that are independent of the particular way that uncertainty is modelled. Failing that, our work suggests at least to state precisely the range of uncertainty models to which a given result applies.

Partial orders and sensor selection Much of the literature on sensor-based agents assumes that the agents have fixed action and sensing capabilities, instead focusing on the (also important) problems of learning those capabilities and developing plans to achieve goals. The research presented in this paper makes a case for thinking carefully about sensor selection. Since sensing and information gathering are costly, agent designers are wise to consider carefully which sensors to include.

A few researchers have considered sensor selection problems. One approach is to attempt to maximize the numerical utility of the information collected (Bian, Kempe, & Govindan 2006). Others use information-theoretic techniques (Zhang & Ji 2005) to assess the value of particular sensor combinations. This approach is unsatisfying because just a few bits of the right information are often more valuable for certain tasks than large quantities of less relevant information. Our work has similar goals, but omits the scalar optimality criterion. This leads directly to the partial order structure of dominances. Although the partial order admits the possibility that no meaningful comparisons can be made, we find this desirable: different sensing-action systems exhibit complex relationships and tradeoffs that can potentially defy meaningful linear ordering. This willingness to admit a partial order also allows us to make comparisons that are task independent. The partial order remains meaningful because Lemma 2 guarantees that the dominance hierarchy respects solvability.

Conclusion

This paper described recent work that seeks to construct an analog to the theory of computation that is sufficiently expressive to deal meaningfully with the sensing and uncertainty issues that are so central to robotics. Although the current work is only preliminary, we believe that the general framework we have presented – that of comparing sensor-based agents based on their progress through a derived information space – offers a unique perspective to several areas within AI.

References

Amir, E., and Russell, S. 2003. Logical filtering. In *Proc. International Joint Conference on Artificial Intelligence*.

Bian, F.; Kempe, D.; and Govindan, R. 2006. Utility based sensor selection. In *Proc. International Conference on Information Processing in Sensor Networks*, 11–18. New York, NY, USA: ACM Press.

Blackburn, P.; van Benthem, J.; and Wolter, F., eds. 2006. *Handbook of Modal Logic*. Amsterdam: Elsevier.

Cimatti, A.; Roveri, M.; and Bertoli, P. 2004. Conformant planning via symbolic model checking and heuristic search. *Artificial Intelligence* 159(1-2):127–206.

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.

Hopcroft, J. E.; Motwani, R.; and Ullman, J. D. 2007. *Introduction to Automata Theory, Languages, and Computation*. Reading, MA: Addison-Wesley, third edition.

Kaelbling, L.; Littman, M.; and Cassandra, A. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101:99–134.

LaValle, S. M. 2006. *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press. Available at http://planning.cs.uiuc.edu/.

McDermott, D., and Doyle, J. 1980. Nonmonotonic logic I. *Artificial Intelligence* 13:41–72.

O'Kane, J. M., and LaValle, S. M. 2006a. On comparing the power of mobile robots. In *Proc. Robotics: Science and Systems*.

O'Kane, J. M., and LaValle, S. M. 2006b. On comparing the power of robots. Submitted to *International Journal of Robotics Research*. Under review.

O'Kane, J. M., and LaValle, S. M. 2007. Sloppy motors, flaky sensors, and virtual dirt: Comparing imperfect ill-informed robots. In *Proc. IEEE International Conference on Robotics and Automation*. To appear.

Robert, C. P. 2001. *The Bayesian Choice, 2nd. Ed.* Berlin: Springer-Verlag.

Russell, S., and Wolfe, J. 2005. Efficient belief-state AND-OR search, with application to kriegspiel. In *Proc. International Joint Conference on Artificial Intelligence*.

Sipser, M. 1997. *Introduction to the Theory of Computation*. Boston, MA: PWS.

Soare, R. I. 1987. *Recursively enumerable sets and degrees*. Berlin: Springer-Verlag.

Zhang, Y., and Ji, Q. 2005. Sensor selection for active information fusion.

Zhang, N., and Lin, W. 1997. A model approximation scheme for planning in partially observable stochastic domains. *Journal of Artificial Intelligence Research* 7:199–230.