Reasoning About Probabilistic Actions At Multiple Levels of Granularity

Lonnie Chrisman
School of Computer Science
Carnegie Mellon University
Pittsburgh, Pa 15213
chrisman@cs.cmu.edu

Abstract

Models of physical actions are abstractions of the physical world and can never represent reality precisely. There are two fundamental choices anyone modeling physical action must face: The choice of what granularity to represent actions at and the choice of how to represent the state description. These two choices are closely coupled and form a basic tradeoff. Finer grain actions require more detail in the state description in order to discern possible interactions between the actions. This extra detail can result in large and complex models that are hard to acquire and tedious to use. Coarser grain actions alleviate the need to discern interactions and thus enable a more abstract state description to be used just as effectively; however, by modeling only coarse grain actions, the models are less flexible and less general.

To reap the benefits of coarse-grain actions without sacrificing the flexibility of fine-grain actions, we advocate the use of models at multiple levels of granularity. Few (if any) existing frameworks support action models at multiple levels of granularity in a coherent fashion. We demonstrate this fact for the case of standard probabilistic models, and then introduce an approach that can handle multiple granularities and use it to demonstrate the fundamental tradeoff between abstraction and granularity.

1 Introduction

A planner that is to construct a course of action to be executed in the physical world must possess a set of action models describing the possible effects of its actions. In a decision-theoretic planner, these models may also specify explicit measures of outcome uncertainty. However, the physical world is very complex and doing a good job of modeling physical actions can be very difficult [Shafer, 1991]. Some have suggested that it is so difficult that it may be better to forget about modeling actions altogether (and therefore, forget about planning) and instead just tell the system how to respond to different situations [Brooks, 1991].

In most domains, we usually have considerable freedom in choosing what to encapsulate and model as an action. Consider, for example, a robot manipulator. At the lowest level (finest granularity), one could identify the action "Apply I amps of current to the motor for T milliseconds." At a slightly higher level, this action can be wrapped in a servo-loop that implements the action "Move hand to location (X, Y)." At a very high level (coarsest granularity), these actions might be used to implement a robust "Grasp Object" action. At the highest level, the action might implement the complete behavior of an agent. By encapsulating actions appropriately, it is often possible to create higher level actions with outcomes that are easier to describe (at least for certain variables of interest), and a greater level of fidelity may be achieved than that obtainable from modeling lower level actions alone. On the other hand, high-level action models lack the flexibility that models at finer granularity have. For example, a model of a high-level action is only applicable when the sequence of steps within that action are followed precisely. There may also be some variables that are easier to predict directly from lower level models. For this reason, it can be beneficial for a planning system to maintain action models at multiple levels of granularity, and to mix these as necessary during reasoning.

There are some examples of systems that maintain action models at multiple levels of granularity, but their motivations are different from those proposed here. Several Explanation-Based Learning systems, for example, derive models of action sequences by compiling the effects of their low-level actions ([Blythe and Mitchell, 1989], [Turney and Segre, 1989], [Minton, 1989]). In these systems, the higher level models serve to save computation the next time that particular action sequence is considered. However, these systems differ from the fundamental idea discussed here in that the compiled models in these systems are identical in terms of information content to the lower level models. In this paper, we are pointing out that higher level models are useful for encoding information specific to the action sequence that cannot be derived from the lower level component action models.

Standard probabilistic approaches to modeling action outcome uncertainty do not allow for the mixing of different models at multiple levels of granularity. By "stan-

dard probabilistic approaches", we are referring to the practice of representing an action model using the conditional probability: $Pr(outcome|action, prestate)$, the approach utilized by most existing planners that handle probabilistic information ([Chrisman and Simmons, 1991], [Dean *et al.*, 1993], [Drummond and Bresina, 1990], [Kushmerick *et al.*, 1993], etc.). The problem with mixing models at multiple levels of granularity arises because the information content at different levels of granularity can be different, making the different probabilistic models mutually incoherent. For example, if a "mega-action" A consists of executing A_1 followed by A_2, the net outcome can be predicted using the composition of the models for A_1 and A_2, or the outcome can be predicted using only the higher level model for A. In the first case, we might obtain $Pr(Outcome_1|A, state_0) = 0.2$, and in the second case we might obtain $Pr(Outcome_1|A, state_0) = 0.4$. The fact that these are different is a logical contradiction (probability literature uses the term incoherence [Lindley, 1987]), and is the reason that standard probabilistic approaches are not applicable. The standard Bayesian dogma is that an agent satisfied with incoherent models is irrational, but in this case there is nothing irrational about maintaining different models at different levels of granularity — it may well be the best way of modeling the actions' outcomes. The incoherence that arises in this example is not the result of irrationality, but an indication that the practice of using a single probability number to encode uncertainty is not powerful enough to properly capture the fact that every model is an abstraction of reality.

In this paper we motivate the need for multiple models at different levels of granularity and present an approach for abstract probabilistic modeling that supports probabilistic models at multiple levels of granularity. The framework is based on the laws of probability, but differs from the Bayesian-like approaches in that it drops the assumption that uncertainty must be encoded by a single probability number and replaces it with the assumption that every model is an abstraction of a more detailed reality. Many papers in the literature on Bayesian reasoning do acknowledge that models are abstractions of reality, and use this fact to motivate the use of explicit representations of uncertainty (e.g., [Horvitz *et al.*, 1988]), but this fact is not used to derive the representation of uncertainty, and these papers make a leap of faith that a single probability number can capture the essence of uncertainty that arises from abstraction. In the new framework, by incorporating more directly the fact that every models is an abstraction of reality, we find that there is a sense in which mutually coherent models at multiple levels of granularity can exist even when the different levels encode different information. The framework highlights a fundamental tradeoff when modeling physical systems: The tradeoff between the level of granularity and the level of detail.

2 Basic Modeling Choices

It is trivial to perfectly model the moves in the games of Chess. From a given board position and move specification, the model predicts the resulting board position. The set of actions (legal moves) are obvious, and the predictions are error-free. This situation is much different from what we face when modeling physical actions — i.e., actions that change the physical real world. It is seldom (if ever) possible to perfectly model the outcomes of physical actions, and it is rarely obvious what units of behavior we should be distinguishing as "actions". Examples of physical actions include those used to accomplish tasks such as tying a shoelace, throwing a ball, washing dishes, winding up a hose, alpine skiing, planning plays in a football game, picking up an object, opening a drawer, shoveling snow, sweeping dirt, making a bed, folding laundry, and so on.

Consider the issues faced if we are to model the actions involved in the task of folding a towel. There are two very fundamental choices that must be confronted: 1. What are the basic actions to be modeled? and 2. What parameters should be used to describe the world state?

An example of an action might be to apply an inward force to both the left hand and the right hand for 300 milliseconds (causing the hands and their contents to move closer together). Alternatively, we might decide to use slightly higher level actions, such as the action of bringing two corners of the towel together. It is even possible to identify and use still higher level actions, such as folding the towel in half, or even folding the towel completely (i.e., one action accomplishes the entire task). These actions are each at an increasingly coarse level of granularity. The finer grain actions might actually serve as subroutines in the implementation of the coarser grained actions. But there is no correct or obvious level of granularity here, any of the above levels are perfectly legitimate. This situation is indicative of all physical domains — the unit of behavior that we should encapsulate as an "action" is a fundamental part of the modeling problem.

There is also the problem of choosing what parameters to include in the state description. To describe the towel, we might impose a rectangular grid on the towel with points spaced one millimeter apart, and thus describe the towel's configuration by listing the x, y, z coordinates of each grid point in space. Alternatively, we might take the same approach using a lower resolution grid with points spaced 10 centimeters apart. We might take a very different approach and divide the towel up into anywhere from two to fifty triangles, or we might just provide qualitative descriptions of the major folds in the towel. Yet a different approach would be to classify the towel's configuration at any time into one of a small number (≈ 10) of categories, such as "wadded up", "spread out", "folded in half", "folded in quarters", etc. These potential state descriptions progress from highly detailed to more and more abstract. As was the case with the choice of granularity, there is no correct or obvious level of abstraction that we should use. Once again, this is indicative of all physical domains and is a fundamental part of the task of modeling actions.

3 Analyzing Sequences of Behavior

We use the term "action" to mean a unit of behavior that obtains the distinction of being called an action when we identify it as something that can be initiated by the agent, named, and modeled. An action model is a description of the effects that the unit of behavior has on the world state as a function of the state it is initiated from. The purpose of action models is to enable the analysis of behavioral sequences, such as predicting what the final outcome would be if the behavioral sequence were initiated. A planner performs such an analysis on a collection of candidate behavioral sequences in order to locate a course of action with favorable results.

Suppose an agent wishes to analyze a given behavioral sequence. If this sequence happens to correspond exactly to one of the modeled units of behavior, the corresponding action model can be recalled and used for the analysis. Typically, however, the agent will not possess an individual action model corresponding precisely to the sequence of interest. In this case, the behavioral sequence must be decomposed into units of behavior that are modeled, and the analysis from each of these models must be composed to analyze of the complete sequence. Not every arbitrary behavioral sequence can be analyzed by a given agent — those sequences that cannot be decomposed into smaller units that are modeled cannot be analyzed. This dimension of flexibility/generality is determined by the collection of action models available and is often a critical concern.

If a behavioral sequence can be decomposed into actions and the agent possesses models for each of these actions, there is still no inherent reason why we should be able to compose action models to obtain any meaningful information about the larger sequence. Here is the roll of the state description. The state description serves as a language for discerning the interactions among the component actions in the sequence[1], and it is this discernment of interactions that allows the results of smaller models to be composed. Let us return to the example of folding a towel. Suppose we are holding an unfolded towel by two adjacent corners, and we wish to predict where the midpoint of the top edge will be after first flipping the top edge outward and then immediately bringing the two corners together. The prediction will be accomplished using two models: One for the flipping action, and one for the action of bringing corners together. These two actions clearly interact — without the initial flip we usually end up with the towel bunched up between the two corners, but with the flip the towel is usually folded more or less in half. If both possible intermediate states are assigned equivalent representations by the state description, the interaction is not discerned, and we cannot meaningfully compose the models to analyze the situation.

A case of particular interest is when we decompose behavioral sequences into sequential units of behavior. In this case the state description discerns the interactions when it has the so called *Markov Property*.

[1]This is one of the two functions of the state description. Its other function is to supply us with a language for expressing the final outcomes we are interested in.

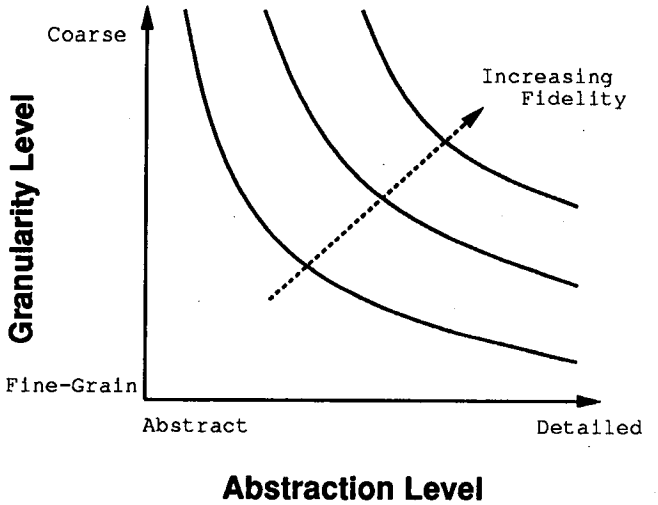

Figure 1: The Fundamental Tradeoff.

Although the state description should discern all potential interactions between the modeled units of behavior, in reality this is never possible. The infinite detail of the physical world means that there is no limit to the potential number of ways in which actions might interact. Thus, the discernment of interactions (including the Markov Property) is an unobtainable ideal. The ramifications of this on our ability to meaningfully compose action models is significant. Even if we have very good individual component models, their composition may be a poor model of the composed behavior if the state description does not discern the most relevant interactions. Because in a physical domain the state description can never realistically discern all interactions, there is always some accuracy lost whenever we analyze a behavioral sequence by decomposing it to finer grain models.

4 Abstraction-Granularity Tradeoff

The granularity of action models and the level of abstraction we choose for the state description are intimately related and form what is perhaps the most fundamental tradeoff faced when modeling physical action. Although models of physical action are never exact, it is possible to obtain models with improved fidelity either by increasing the level of detail in the state description, or alternatively by choosing coarser units of behavior for the actions being modeled. In the first case, the increased detail enables the discernment of a greater number of interactions, while in the second case the coarser models alleviate the need to discern the interactions in the first place.

The tradeoff between abstraction and granularity is summarized by Figure 1. The curves trace out constant fidelity levels. Consider once again the towel folding example. At one extreme, suppose we model the very coarse action of folding the towel completely. To obtain an acceptable level of fidelity, it suffices for the state description to discern only two configurations: "Towel-

Folded" and "Towel-not-Folded." Thus, as the figure shows, it is possible to reach a desired fidelity by using very coarse actions with an extremely abstract state description. Alternatively, we can model actions at the slightly finer grain of "fold the towel in half." The previous state description is now grossly inadequate, but the same fidelity can be reached using the more detailed state description that discerns the towel configurations "Spread Out", "Folded in Half", "Folded in Quarters", etc. If we went to extremely fine grain motor control actions, such as "apply an inward force for 300 milliseconds," achieving the same fidelity would probably necessitate an excessively large state description capturing highly detailed information about geometric configuration, forces, and dynamics.

Although increased detail and coarser granularity are both advantageous in terms of information content and fidelity, each comes with corresponding disadvantages. Coarser models are less flexible and less general. For example, the "fold towel completely" action is of no help if we encounter the new task of hanging the towel on the clothesline, while a finer grain action like "fold towel in half" could be re-utilized. Likewise, increased detail leaves us with larger and more complex models, thus increasing the difficulty of model acquisition and computational complexity of using the models.

5 Motivations for Multiple Granularities

We would like to endow our agents with the flexibility that comes with using fine-grained models, but at the same time we desire the increased fidelity without excessive state detail that we get with coarse-grained actions. To accomplish this, we advocate the use of action models at multiple different granularities. We provide the agent with a general collection of fine-grained action models enabling it to analyze a wide range of behaviors, while at the same time providing coarse models of action sequences that are routinely useful.

The use of multiple granularities introduces redundant ways to model/decompose a given behavioral sequence. Although it may be decomposed in multiple ways, the results of each analysis are typically not identical. The fine-grain analyses accumulate greater error as a result of undiscerned interactions, but they may contribute some information that is not available from the coarser analyses (if there is epistemological ignorance in the coarse models or if the detailed models utilize additional state parameters). It is therefore highly advantageous to be able to combine and freely mix models at multiple levels of granularity. In fact, there are many advantages: combining multiple models may yield a more informed analysis than using either model individually, the decomposition of behavioral sequences does have to be into pieces at the same granularity, and no requirement or complications arise from having to distinguish some unique correct level of granularity. Obtaining these advantages requires the property termed *coherence*, which simply specifies that any analyses of the same behavior using different models must not be mutually contradictory. Amazingly, few (if any) existing frameworks can support models at multiple granularities in a coherent fashion. The general problem is simply that an agent's models comprise its knowledge and therefore the agent believes both answers to the same problem when answers result from analyses based on two different models. When two non-identical answers are obtained for the same quantity, this constitutes a contradiction in most frameworks.

6 Standard Probability Models

Probabilistic action models are attractive for dealing with the fact that models of physical actions are imperfect. However, the standard approaches for utilizing these models cannot coherently support multiple granularities. We are referring here to action models that specify a probability distribution over possible outcome states conditioned on the state before the action is executed. These action models also assume that the Markov Property holds (i.e., that the state description summarizes all relevant information), and are instances of Markov Decision Processes ([Howard, 1960]) or even more generally of Partially Observable MDPs ([Lovejoy, 1991], [Chrisman, 1992b]). They are sometimes also called Bayesian Action models due to their utilization of a point outcome probability distribution.

Suppose we flip a towel outward and let go, and we wish to know the probability that it lands in a spread-out configuration. The flip action is likely to cause the towel to spread out from our grip in roughly a horizontal plane. Releasing it from a horizontal spread out configuration is likely to cause the towel to land on the floor spread out. Our fine grain models reflect these outcomes. But when the two actions occur in direct succession, a rather different phenomenon occurs: The far edge of the towel sometimes experiences a backlash at the final moments of the flip action causing the towel to come together in a wad before landing on the floor. Denote the flip action model by $P_{flip}(s_1|s_0)$ and the release action model by $P_{letgo}(s_2|s_1)$, but suppose that the intermediate state description s_1 does not discern the backlash interaction. We analyze the outcome of the total episode by multiplying P_{flip} and P_{letgo} and summing over all discerned intermediate states s_1. Using an actual MDP action model that fit the above description we obtained $P_{flipoletgo}(SpreadOut|s_0) = 0.8$.

Next, suppose the same agent possesses a model for the coarser action $A \equiv flipoletgo$. When the two actions are studied together, we find that the towel lands spread out only about half the time, so $P_A(SpreadOut|s_0) = 0.5$ (which is less than 0.8 since it compiles in the backlash effect). But P_A and $P_{flipoletgo}$ are the identical quantity, so by believing its models at both granularities, the agent encounters an incoherence. Because of this phenomenon, the standard probabilistic approaches cannot handle multiple granularities.

There is also a more general and useful way to understand the source of these problems that is summed up with the following equation:

$$abstr(D_{A_1} \circ D_{A_2}) \neq abstr(D_{A_1}) \circ abstr(D_{A_2}) \qquad (1)$$

One can conceptualize D_{A_1} and D_{A_2} as being perfect (possibly infinitely long) descriptions of the outcomes of actions A_1 and A_2 and $abstr(\cdot)$ as a function that returns an abstraction of the supplied model. The equation states that the abstraction of the composition of two actions is not the same as the composition of the abstraction of each action. The reason they are not equal is because the right hand side cannot account for interactions that are not discerned by the language of the abstraction, while these are compiled into the abstraction on the left hand side. Since every model is an abstraction of reality, we can let $M_{A_1} = abstr(D_{A_1})$, etc., and rewrite (1) as

$$M_{A_1 \circ A_2} \neq M_{A_1} \circ M_{A_2}$$

. By definition each model is providing information about reality, but the inequality demonstrates that distinct information is obtained from models at different levels of granularity.

The standard probabilistic framework (like many other frameworks) does not provide a realistic account of abstraction that accounts for (1), and as a result, it cannot coherently handle multiple granularities.

7 Approach

We have developed an approach for modeling physical action that supports models at multiple granularities. The approach is derived from a combination of the laws of probability and an explicit account of abstraction. The approach differs from standard probabilistic approaches by using lower and upper probability intervals over outcomes rather than point outcome distributions, where the intervals encode the only precision warranted by the level of abstraction of the state description.

We assume the existence of a detailed standard probabilistic model, such that the state description in the detailed model discerns all relevant interactions. We visualize this model as being infinitely large. Normally, this model is not explicitly represented in the computer — its existence simply provides the means for giving the intervals a precise semantics in terms of abstraction.

Consider a simple example where there are two (detailed) states ω_1 and ω_2 such that when A is executed from either state, two outcomes o_1 and o_2 are possible, with $P(o_1|\omega_1) = .3$, $P(o_2|\omega_1) = .7$, $P(o_1|\omega_2) = .2$ and $P(o_2|\omega_2) = .8$. The model of A is an abstraction of this more detailed description and does not discern ω_1 from ω_2. It utilizes a state $\theta = \{\omega_1, \omega_2\}$ and the following interval probabilities: $.2 \leq P(o_1|\theta) \leq .3$, $.7 \leq P(o_2|\theta) \leq .8$. These intervals represent the most one can possible infer about the outcome estimate of the detailed model by using only information contained within the abstract state description. Anything more specific would assume information that is not available. We say a detailed model is consistent with an abstract interval valued model if all possible probabilities derived from the detailed model are contained within the intervals derived from the abstract model [Chrisman, 1992a].

Interval probabilities allow us to utilize non-identical models of the same behavioral sequence. Two different models of the same behavior are coherent as long as both are an abstraction of some more detailed model that is consistent with both. In terms of probability intervals, this simply means that the intervals must intersect.

Our approach supports models at multiple levels of granularity in the following fashion. Suppose $abstr(\cdot)$ is a function that abstracts a detailed probabilistic model to some abstract state description of interest, producing the appropriate probability intervals. If A_1, A_2, A_3 are fine-grain actions, (conceptually) described in the detailed model by D_{A_1}, D_{A_2}, D_{A_3}, then the agent models them as $M_{A_1} = abstr(D_{A_1})$, etc. Suppose $A \equiv A_1 \circ A_2 \circ A_3$ is the higher level action formed by composing the smaller units of behavior. Then $D_A = D_{A_1} \circ D_{A_2} \circ D_{A_3}$ (because the detailed model discerns all interactions), and the agent's model of A is $M_A = abstr(D_A) = abstr(D_{A_1} \circ D_{A_2} \circ D_{A_3})$. Even though the model M_A and the composition $M_{A_1} \circ M_{A_2} \circ M_{A_3}$ produce non-identical intervals, the models are coherent as long as there exists some more detailed model that discerns all interactions and gives rise to M_A, M_{A_1}, M_{A_2}, and M_{A_3} when abstracted. Information from M_A can be combined with $M_{A_1} \circ M_{A_2} \circ M_{A_3}$ simply by intersecting the resulting intervals.

Lower and upper probability intervals have been studied by many. The most relevant for our work include [Walley, 1991], [Fagin and Halpern, 1989], [Walley and Fine, 1982], [Chateauneuf and Jaffray, 1989], and [Shafer, 1981]; however, the semantics of our intervals as arising from abstraction is new.

8 A Demonstration

The framework from the previous section can be utilized to explicitly demonstrate the abstraction-granularity tradeoff. A number of advanced inference mechanisms for analyzing behavioral sequences were developed and implemented based on the general theory of lower probability. These enabled the agent to derive interval probability predictions for queries about the outcome of a behavior using a collection of action models at any level of abstraction and granularity. We also implemented routines to automatically abstract a detailed model to specified levels of granularity and abstraction. The details of these methods are beyond the scope of this paper, but were used to run the following experiment.

An eight state, two action Markov Decision Process was used as a detailed model. By conglomerating two states into one, a more abstract state description was obtained. This was repeated five more times to generate a totally ordered list of state abstractions, the most abstract discerning only two states. We decided to analyze the behavioral sequence: $A_1 \circ A_1 \circ A_2 \circ A_1 \circ A_2 \circ A_2$, and by composing two adjacent actions in this sequence, we obtained a coarser-grained action that was applicable to this problem. By repeating four more times, we obtained coarse actions A_{11}, A_{21}, A_{22}, A_{1121} and A_{112122}, which define a list of decompositions for the sequence that is totally ordered with respect to the granularity of the analysis. With seven abstraction levels and six levels of decomposition granularity, we could analyze the sequence in 42 different ways.

For each of the 42 combinations of abstraction and granularity, we wished to measure the fidelity achieved

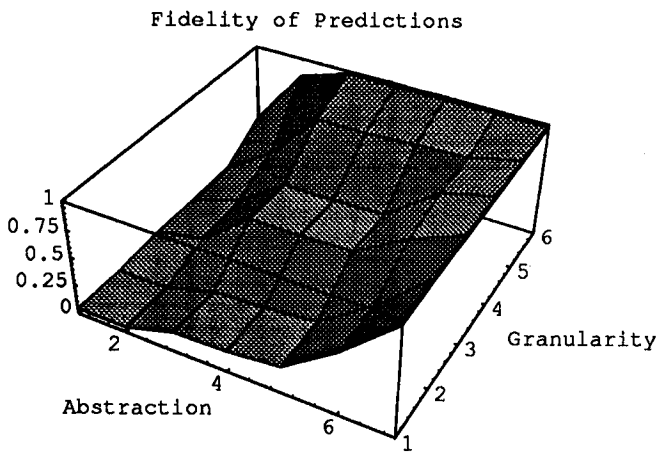

Figure 2: Empirical demonstration of the tradeoff between abstraction and granularity.

by using models at that level. For each combination, the computer automatically derived the best possible action models (i.e., tightest justified intervals) for that level of abstraction and granularity. We specified four queries pertaining to the final outcome, and used each combination to compute intervals for all four queries. For each combination, a measure of fidelity was obtained by averaging the widths of the interval for each of the four queries and subtracting this from one. A fidelity of 1.0 is the best possible and indicates that no information is lost as a result undiscerned interactions at that level of abstraction and granularity. A fidelity of zero indicates that the results are entirely vacuous, such that at that level of abstraction we are completely unjustified in concluding anything about the complete behavioral sequence from a composition of the component models. The results are graphed in Figure 2. In the graph, abstraction level one is the most abstract where only two states are discerned and abstraction level seven is the most detailed where every state in the detailed model is discerned. The graph demonstrates that for any level of granularity, an increase in detail improves the fidelity of the model. However, the improvement between abstraction levels 3, 4 and 5 is rather minor indicating that the interactions discerned by the extra detail in level five over level three were not significant for the four outcomes we analyzed. One should note that they could very well be important for other tasks. In granularity level one, we used the basic fine-grain actions from the MDP. Granularity six corresponds to coarsest action model, namely that in which the entire behavioral sequence of interest was modeled as a single action. It is worth comparing the tradeoff as measured in the experiment (Figure 2) to Figure 1.

9 Conclusion

The models of physical actions such as those used by a planner must necessarily be abstractions of reality. The real world is too complex and too intricate to model perfectly in complete detail. When we model physical actions, we face two very fundamental choices: The choice of what units of behavior to encapsulate as "actions" and the choice of how to represent the world state. These two choices are mutually dependent and form what is probably the most fundamental tradeoff encountered when modeling actions: The tradeoff between abstraction and granularity.

A state description should ideally discern all potential interactions between the actions that compose a larger behavior. Since the world presents us with unlimited detail, it is never possible in practice to completely discern all interactions. In the case of sequential action execution, we can never completely obtain the so called Markov Property. This limits our ability to meaningfully compose action models in order to analyze an action sequence — even if the individual action models themselves are in good agreement with reality, their composition may not be. It all depends on how well the relevant interactions in the sequence are discerned by the state description.

The burden on the state description to discern all potential interactions can be alleviated by modeling actions at a coarser granularity. The coarser action eliminates the need for the state description to discern the "intra-action" interactions, thus making it possible to obtain models with high fidelity using more abstract (and thus simpler) state descriptions. Despite this advantage, the use of coarser action models comes with a disadvantage as well: the models are less flexible. Finer grain models are applicable to the analysis of a greater number of possible behavioral sequences.

In order to obtain the benefits of improved fidelity and simpler models that we get with coarse grained while maintaining the flexibility that accompanies the use of fine-grain models, we advocate the use of action models at multiple levels of granularity. We provide the agent with models of both fine-grain and coarse-grain actions, giving the agent multiple ways of decomposing and analyzing any given behavioral sequence. Different levels of granularity provide an agent with some redundant information but also with some distinctly different information about the real world. In order to utilize all information obtained from models at multiple levels of fidelity, we desire models at different levels of fidelity to be mutually coherent. Unfortunately, very few existing frameworks provide a coherent means for utilizing models at different granularities.

We presented an approach based on lower and upper probability intervals that does allow action models at multiple levels of granularity to coexist in a coherent fashion. The framework was derived from the laws of probability augmented by an explicit account of abstraction. The width of probability intervals expresses the amount of precision that is justified by a model's level of abstraction and granularity. The degree to which the state description discerns the interactions relevant

to a given problem determines the robustness of the final answer that is obtained by using models based on that state description. The probability intervals provide an explicit measure of this robustness. We used the framework to empirically demonstrate the abstraction-granularity tradeoff.

Because our approach is based on the laws of probability and utilizes a type of probabilistic action model, we expect a planner that uses the framework to be of a decision-theoretic variety. Unlike pure decision-theoretic planning, however, there is an extra consideration: An agent may have to choose between (1) a coarse of action with an apparently high expected utility, but where the analysis of the expected utility was not robust due to unmodeled interactions, and (2) a coarse of action with a medium expected utility, but where this estimate is quite robust and can be trusted. It may be more rational to choose the latter.

Acknowledgements

The development of the work reported in this article has benefited considerably from discussions with my advisors Reid Simmons and Tom Mitchell. This research was sponsored by the NASA-Jet Propulsion Lab under Grant No. NGT-51039 and by the NASA Grant No. NAGW-1175. The views and conclusions contained in this article are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of NASA or the US government.

References

[Blythe and Mitchell, 1989] Jim Blythe and Tom Mitchell. On becoming reactive. In *Sixth International Workshop on Machine Learning*, Cornell Univ., June 1989.

[Brooks, 1991] Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[Chateauneuf and Jaffray, 1989] Alain Chateauneuf and Jean-Yves Jaffray. Some characterizations of lower probabilities and other monotone capacities through the use of Möbius inversion. *Mathematical Social Sciences*, 17:263–283, 1989.

[Chrisman and Simmons, 1991] Lonnie Chrisman and Reid Simmons. Sensible planning: Focusing perceptual attention. In *Ninth National Conference on Artificial Intelligence*, July 1991.

[Chrisman, 1992a] Lonnie Chrisman. Abstract probabilistic modeling of action. In J. Hendler, editor, *First International Conference on AI Planning Systems*, June 1992.

[Chrisman, 1992b] Lonnie Chrisman. Planning for closed-loop execution using partially observable markovian decision processes. In *AAAI Spring Symposium Series: Control of Selective Perception*, Stanford University, March 1992.

[Dean *et al.*, 1993] Thomas Dean, Leslie Pack Kaelbling, Jak Kirman, and Ann Nicholson. Planning with deadlines in stochastic domains. In *Eleventh National Conference on AI (AAAI)*, pages 574–579, Wash. D.C., 1993.

[Drummond and Bresina, 1990] Mark Drummond and John Bresina. Anytime synthetic projection: Maximizing the probability of goal satisfaction. In *Eight National Conference on Artificial Intelligence*, July 1990.

[Fagin and Halpern, 1989] Ronald Fagin and Joseph Y. Halpern. A new approach to updating beliefs. Technical Report RJ 7222, IBM, May 1989.

[Horvitz *et al.*, 1988] Eric J. Horvitz, John S. Breese, , and Max Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2:247–302, 1988.

[Howard, 1960] Ronald A. Howard. *Dynamic Programming and Markov Processes.* The M.I.T. Press, 1960.

[Kushmerick *et al.*, 1993] Nicholas Kushmerick, Steve Hanks, and Daniel Weld. An algorithm for probabilistic planning. Technical Report TR-93-06-03, Dept. of Comp. Sci. and Eng., University of Washington, Seattle, June 1993.

[Lindley, 1987] Dennis V. Lindley. The probability approach to the treatment of uncertainty in artificial intelligence and expert systems. *Statistical Science*, 2(1):17–24, 1987.

[Lovejoy, 1991] William S. Lovejoy. A survey of algorithmic methods for partially observable markov decision processes. *Annals of Operations Research*, 28:47–66, 1991.

[Minton, 1989] Steven Minton. Qualitative results concerning the utility of explanation-based learning. In *Seventh National Conference on AI (AAAI)*, pages 564–569, St. Paul, 1989.

[Shafer, 1981] Glenn Shafer. Constructive probability. *Synthese*, 48:1–60, 1981.

[Shafer, 1991] Steve Shafer. Why can't we model the physical world? *CMU Computer Science: A 25th Anniversary Commemorative*, pages 295–309, 1991.

[Turney and Segre, 1989] Jennifer Turney and Alberto Segre. A framework for learning in planning domains with uncertainty. Technical Report TR-89-1009, Dept. of Comp. Sci., Cornell Univ., May 1989.

[Walley and Fine, 1982] Peter Walley and Terrence L. Fine. Towards a frequentist theory of upper and lower probability. *The Annals of Statistics*, 10(3):741–761, 1982.

[Walley, 1991] Peter Walley. *Statistical Reasoning with Imprecise Probabilities.* Chapman and Hall, New York, 1991.