

Agents in the Real World

Computational Models in Artificial Intelligence and Cognitive Science

Matthew Stone

mdstone@cs.rutgers.edu

March 26, 2001

1 Motivation and Overview

The mid-twentieth century saw the introduction of a new general model of processes, COMPUTATION, with the work of scientists such as Turing, Chomsky, Newell and Simon.¹ This model so revolutionized the intellectual world that the dominant scientific programs of the day—spearheaded by such eminent scientists as Hilbert, Bloomfield and Skinner—are today remembered as much for the way computation exposed their stark limitations as for their positive contributions.² Ever since, the field of Artificial Intelligence (AI) has defined itself as the subfield of computer science dedicated to the understanding of intelligent entities as computational processes.

Now, drawing on fifty years of results of increasing breadth and applicability, we can also characterize AI research as a concrete practice: an ENGINEERING APPROACH to constructing COMPUTATIONAL ARTIFACTS that act in the REAL WORLD. Such artifacts are known as AGENTS, so we can also describe AI as the practice of agent design.³ Today, it is the strength and usefulness of this practice—in domains from human-computer interaction to robotics to electronic commerce and services—that motivates study in AI and informs the design of the AI curriculum. By contrast, the view of human intelligence as a computational process now guides empirical and theoretical research throughout academia, in philosophy, linguistics, psychology, and even neuroscience. As this volume attests, these disparate fields continue to achieve increasingly full accounts of the stunning complexity of human biology, of human environments, and of human abilities. Yet AI can still offer unique insights into the computational mechanisms and the real-world knowledge that must underlie our own complex, adaptive behavior; and many AI researchers are proud to offer what contributions they can to this overarching and imposing effort.

In this chapter, I will first set out, in broad strokes, the activities of modeling, implementation and evaluation that distinguish the practice of agent design. I continue by illustrating these activities in more depth as they play out in the most straightforward case of decision analysis for agent design. This decision analysis show-

cases both the basic concepts of modeling, implementation and evaluation in AI and the fundamental and inherent tradeoffs that must be reconciled in an effective agent. The remainder of the chapter explores an array of techniques for agent design that provide a standard repertoire of alternatives on these tradeoffs. I encourage you to view these techniques simultaneously as a toolkit for innovative practical system-building and as a inspiration to new research into the computational problems of acting in our world.

2 A General Description of Agents

2.1 Agents Take Real-world Actions Based on Real-world Data

In thinking of AI as an engineering approach to building computational artifacts that act in the real world, it is easy at first to overlook the dramatic constraints that the real world imposes on agents, and indirectly on engineers who build agents. But consider how much freedom in design we find in more typical computer programming—in structuring classes and methods in object-oriented programming, for instance.

On the one hand, we can often expect to stipulate what inputs the program will accept, in an unambiguous and restrictive specification. Thus, a compiler or interpreter works with texts described by a formal grammar; operating-systems modules work with the data structures of a particular protocol or programming interface; and even a user interface channels its communication with its human user through a small and predefined set of conventional widgets. In circumscribing input this way, we presuppose that other people will respect these artificial limits: that programmers must take responsibility to write syntactic code, to implement applications that conform to a systems protocol, to learn to use a graphical interface.

On the other hand, we can typically expect to frame the output of the program in terms of extremely flexible suites of low-level actions, carefully designed and robustly implemented to achieve useful effects reliably and easily. Look at the sequences of microprocessor instructions delivered by a compiler, the storage and communications primitives beneath an operating system, or the graphics commands of a windowing system for a user interface. Such resources are never perfect, but in practice, programmers can expect to overcome any unexpected results or missing functionality with simple workarounds, not global analysis.

These expectations make for an attractive and useful division of labor. The design of a computer program becomes an exercise in internal organization—a matter of carving up an overall task into subproblems in an elegant, comprehensible, exten-

sible way and of realizing lean, encapsulated modules to solve these subproblems. Such tasks are challenging but sharply delineated.

These idealizations collapse with a commitment to the real world, to an input environment and output effects that precede our implementation and exhibit no overarching design. In an agent, a designer can no longer stipulate any artificial invariants that inputs must respect; the inputs must be those that nature provides. In an agent, a designer can no longer presume a general repertoire of actions with robustly-designed links to desired effects; the consequences of actions unfold naturally, through the inherent dynamics of a broader world. As if this weren't enough, the complexity and unpredictability of the real world seems to complicate the internal organization of an agent, too. It is fantastically difficult to carve small and interesting tasks out of the world and solve them in a constrained way, the way you would carve up the problems induced in typical computer programs and write separate modules to handle them. Indeed, some AI researchers have suggested that real-world problems are so complex and interdependent that any sufficiently interesting problem will effectively involve a solution to all of them. (This leads to the gag concept of an AI-complete problem.)

The difficulties of agent design come as no surprise with the picturesque tasks that sometimes motivate AI research; after all these are tasks that are only now starting to afford suggestive implementations. One such benchmark is an agent design to control a physical robot running tours or deliveries in a populated space.⁴ Here, inputs might include the readings of real sensors from the environment: optical cameras, laser and sonar rangefinders, perhaps a global-positioning receiver. These readings are always more or less noisy, and sometimes fail altogether. Outputs, meanwhile, might take the form of voltages to motors (on wheels, for example); the success or failure of such an action will depend on such ephemera as the location of the robot itself, the location of people and other obstacles in the robot's surroundings, and even the condition of the floor. The gap between inputs and outputs calls for an implementation whose every action responds to an imposing range of considerations. The agent must draw on background information, such as maps of the environment; it must interpret its current sensed data as mirroring or diverging from the real state of the world; it must predict and weigh the range of effects a choice might have in its specific situation.

But don't be misled by ambitious tone of such benchmarks. The use of real-world actions and real-world data is now an increasingly common requirement of

cutting-edge computer systems. When internet software assesses the content of web pages written by people for people, it must deal with the human idiosyncracies of style, sloppiness and error. When ubiquitous computation is embedded not in a pre-defined computational infrastructure but an open-ended physical environment, inputs and outputs involve inevitable uncertainties and indirections. Any commitment to the real world, no matter how constrained, brings with it the ensemble of perspectives and challenges to which AI research responds.

Real-world tasks give AI a focus on MODELING, describing the real world by a mathematical approximation. A model, as an approximation of its object, acknowledges the open-ended texture of reality; there are always finer distinctions to be drawn, wider-ranging interactions to capture. An effective design allows an agent to respond robustly across a range of possible environments despite these approximations—a good model can capture just those features and regularities that an agent depends on in the environment for its success. But even a good model suggests possible refinements. Pursuing them not only guides the development of more successful and flexible agents, but also, in the best case, can challenge our understanding of the rich world in which we too must act.

Mathematics imbues a model with precision. With the methods of logic, probability and statistics, we can definitively articulate the information about the world our model can offer, and the assumptions about the world our model embodies. In so doing, we lay essential groundwork for the implementation and evaluation of an agent that exploits just this information, but depends on just these assumptions, to solve its real-world tasks.

2.2 Agents Are Computational Processes

We live in a world so pervaded by silicon, in all its ramifications, that it is easy to forget that there are alternatives to computational processes. But of course until recently, these alternatives were the mainstay of invention. At one extreme, we counted on assemblies of gears and escapements and pendulums to coordinate real-world data, the passage of time, with a real-world action, the motion of a dial. Now we are happy to conduct our schedules according to contraptions of piezo-electric crystal, LCDs, and silicon wafers instead.

Many fields closer to AI than this mechanical engineering example—notably control theory in electrical engineering—originate in the description of analog and continuous systems that generate sophisticated real-world behavior in response to

messy real-world data. These fields have developed fundamental mathematical tools which Moore's law increasingly brings within the sphere of a computational approach: a case in point is the discrete-time Kalman filter, a probabilistic framework for optimally reconciling a noisy prediction from the past and a noisy datum from the present.⁵ Such contrasts impel us to be precise about computation; it is, after all, perhaps the oldest and strongest constant of AI practice. I follow Allen Newell's answer.⁶

To say that an agent is a computational artifact means that it is a PHYSICAL SYMBOL SYSTEM. Most generally, we can count as a symbol any object that has an arbitrary conventional meaning; this concept originates with the nineteenth-century American philosopher Peirce.⁷ For AI in particular, then, the symbol consists of a discrete, atomic element which by stipulation is interpreted as standing for some generalized constituent of the real world: perhaps an individual object, or a property that an object may or may not enjoy, or a state of affairs that may or may not obtain. In a concrete implementation, symbols are realized as tokens of sequences of charged and neutral bits in memory cells; their correspondence with reality is mediated by a symbolic programming language and determined by a mathematical model of reality framed by a human designer.

In a physical symbol system, occurrences of symbols are assembled together into REPRESENTATIONS: complexes of symbols that correspond to statements about the real world. The correspondence endows the state of the agent—otherwise just an array of ones and zeros—with intelligible meaning. The agent's representations may spell out facts that the agent knows, goals that the agent has for its environment, or intentions to which the agent is committed. These representations give the agent meaningful reasons to act in its environment in a specified way—reasons to act that can themselves take the form of representations. Moreover, where this correspondence originates in a model of the world, our understanding of the model informs our understanding of the agent's behavior. An agent's representations, and its behavior, can fit the world only as much as the model that guides its design.

Here is the paradox and the opportunity of computational processes: though WE understand representations in terms of their MEANING, the process itself manipulates representations purely on the basis of FORM. An agent undergoes a deterministic sequence of discrete states; a finite specification ultimately spells out how each state is extended, revised, or streamlined in its successor, based just on the readings of the agent's sensors and the brute identity of the state as a binary array. The

paradox, that the meaning of representations is simultaneously fundamental and irrelevant, is palpable; the opportunity, founded in the basic results of computational logic and probability, is the fact that we can design computational processes in an impressively wide range of cases to serve up representations exactly when those representations follow from our mathematical models. Such design epitomizes the AI implementation.

2.3 Agents Must Be Engineered

AI is an engineering methodology, dedicated to understanding the empirical performance of agents as a guide to agent design. That means that after the modeling and implementation for an agent comes a thoroughgoing evaluation, spanning the concerns of AI research. For starters, we must determine how useful the agent proves for the overall task for which it was designed. In this assessment, nothing substitutes for observed runs of the agent, both in realistic field tests and across specifically controlled protocols of execution.

A successful design is not established merely with a satisfactory product; we must also evaluate our agent as a computational process. In particular, we can ask how well our algorithms perform not in idealized or worst-case conditions, but when running on specific hardware under a specific distribution of problem instances. Such questions become particularly important when algorithms only approximately realize the predictions of an underlying model, so that the very behavior of the agent depends on the empirical properties of its computations.

Most abstractly, but equally importantly, we must undertake an evaluation of our models themselves; the construction of our agent enables new investigation of genuinely scientific questions. How well do the models that we have constructed to describe the environment of our agent actually fit the situations our agent encounters? Are the approximations and idealizations about the world that we have adopted in our model appropriate for practical problem-solving? At long last, is the world as we thought it was?

Evaluation stands on its own as a feature of AI practice, but it revolves around modeling every bit as much as does design. Whereas design involves models of the agent's environment, in evaluation, we model the ENSEMBLE of the agent acting in its environment. Our evaluation model distinguishes a particular hypothesis about the running agent from a family of plausible alternative descriptions of it; each description maps out the distribution of behavior we should expect from a suite of

events in which our agent runs.

Now, we are rarely fortunate enough in trials of our agent to observe exactly the behavior that our evaluation hypothesis predicts, and rarely unfortunate enough to observe results that completely diverge. The difference between what we expect in trials and what we see is a matter of degree. Modeling in evaluation serves as the background for the judgments we report in assessing this actual observed behavior. By modeling the agent in its environment, we normalize the difference between observation and prediction—we quantify how surprising such a discrepancy would be in view of our specific hypothesis about the running agent and alternative models that could better match the experimental observations.

In evaluation, we can easily be proved wrong, but we will never discover that our agent works, that our computations are accurate, that our model fits. The best we can hope for is a provisional go-ahead—an invitation to attempt a richer model, a more complex computation, even just a more strenuous evaluation. No research practice rewards those who seek definitive answers. The appeal of AI is the ability to ask with absolute freedom, “what if?”—what does it mean to view the world, and ourselves, against a specific set of assumptions—and the ability to flesh out the question into a model, to realize the model in a real artifact, and set that thing loose into a world where, in the end, we can all judge “what if” for ourselves.

3 Decision Models: A Concrete Illustration of AI Practice

For the rest of this chapter, we will explore this general view of agent design. We start by considering a simplistic but general way of describing how an agent’s observations and decisions can lead to successful or unsuccessful real-world outcomes. These descriptions are known as `DECISION MODELS`, and the process of constructing such models is known as `DECISION ANALYSIS`. Decision models do not originate within AI; decision models offer a perspicuous and principled tool for a wide range of fields, particularly business management, medical care and public policy, whose practice involves choice under uncertainty.⁸ But here we will emphasize how computer scientists in particular can use decision models, as a mathematical approximation to the `DESIGN SPACE` for their agent implementations.

3.1 Modeling

To apply decision analysis to agent design, a designer starts by finding a set of situations that the agent will enter in a distinctive computational state, and that the agent

could face repeatedly without any repetition substantially affecting another. I will call this the `CONTEXT` under investigation in the design. The problem for design is to determine what the agent will do in this context.

Observe that the context of investigation is characteristic of AI practice in setting up a link to the real world, and an approximation. When we link a distinctive state of the agent to a specific set of real-world situations, we help ensure that we have a meaningful subproblem for agent design, because we ensure that the computations we plan for the agent can and will determine the agent's responses to a clear range of circumstances. Indeed, it is common in robotics to call such a collection of coherent responses a `BEHAVIOR`, especially among those who emphasize the importance of equipping a successful agent with an appropriate repertoire of behaviors.⁹ At the same time, in an open-ended world, we must expect that the agent's computational states across this set of real-world situations will be only approximately equivalent; we must expect that such states will arise not under just these situations exactly, but only approximately so; and we must admit that some of these situations exert some influence on some others, even if that influence can be ignored for practical purposes. We must be mindful of these approximations in implementing and evaluating our design.

Let us consider the implementation of a robot-pet consumer toy an illustration. Its task, informally, will be to entertain its owners, by performing tricks such as fetching a ball thrown past it or dodging a ball rolled towards it. Our toy will have legs for locomotion, and a simple camera and touch receptors to sense its environment; we would doubtless want it to take the form suggestive of a familiar pet—a dog say.

A designer might devote special consideration to those situations where one of these tricks takes the dog from a brightly-illuminated area to a darker one. What if, for example, a slightly malicious child gets the dog to scoot through a doorway from a sunny room into a dark closet? The dog faces particular uncertainty and risks in such situations; the dog will have limited visual information until the aperture of its camera opens mechanically, and the camera can capture crisp images at the lower light level. Fortunately, since the dog can be provided with a trigger that identifies cases of insufficient illumination robustly (if never perfectly), the dog can be designed to handle these situations strategically.

Once we have a context of investigation, we can frame the design problem precisely. Our final agent will have some algorithm that it follows in this context; this

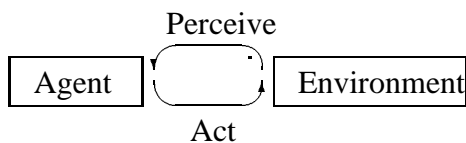


Figure 1: The agent’s cycle of perception and action in its environment.

algorithm continues until the agent’s task is complete or until other conditions prevail and other algorithms take over. We suppose that the execution of this algorithm involves a sequence of possible states or CYCLES; each cycle offers the agent an opportunity first to PERCEIVE its environment and then to ACT to change it. Figure 1 offers a mnemonic depiction of the origin of this cycle of execution in the agent’s interaction with its real-world environment.

In each cycle, the agent begins with expectations about its environment informed by its history, including the distinctive initial state of the algorithm, any sensor readings it has obtained during prior cycles of execution, and any past actions it has taken. The agent can then reconcile these expectations with the new readings about the environment available in this cycle. Finally, the agent selects an appropriate action, as dictated by the algorithm of our design.

A decision model offers perhaps the simplest possible description of the space of such algorithms. A decision model represents the stages of execution of possible algorithms in terms of a history of observations and choices that the agent has made from its initial state. It represents the agent’s expectations at each stage of execution statistically, in terms of the probability of different observations. At any stage where the agent will act, the decision model lays out each of the possible actions that we might design the agent to take. Finally, for stages where the algorithm terminates, the model provides an assessment of the outcome that the agent has achieved with the complete alternative history.

The formal structure of a decision model is a TREE—that is, a mathematical object containing a set of NODES and a set of EDGES that lead from one node to another, where there is exactly one path from any one node to any other node. (A path follows a succession of distinct edges forwards or backwards.) Each node in a decision model represents a possible stage of computation for our agent, which the agent would reach in a subset of possible situations in the context of investigation. Each edge in a decision model represents a possible way that the computation could

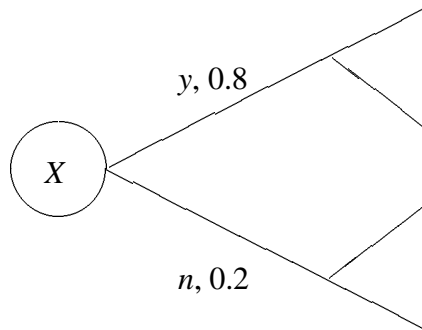


Figure 2: An observation node.

evolve in one step of perception or action.

Internal nodes in a decision model are of two kinds. The first kind, OBSERVATION NODES or CHANCE NODES, represent a stage of execution of the algorithm which obtains information from sensors. From the agent's point of view, these observations are unpredictable; the agent must be prepared for any possible result. At the same time, the agent may have some evidence about which results are more likely and which results are less likely.

Mathematically, we describe such an observation using a RANDOM VARIABLE associated with the observation node. This variable takes on one of a finite set of different VALUES, corresponding to the different possible results of the observation. Since the decision model must represent all possible alternative histories of computation, each possible value determines an edge to a new node, called a CHILD of the observation node; this child represents a further stage of execution in which the agent has observed this value. In addition, to formalize the evidence available to the agent about how its history meshes with the real-world situations of the context, each value gets a PROBABILITY between 0 and 1. Here, this probability is measured relative to the situations in this context which match the agent's history; the probability reports the fraction in which an agent would observe the specified value. Of course, the sum of the probabilities for alternative values of the variable must sum to 1.

Figure 2 illustrates the graphical convention that we will use to display observation nodes in decision models. The node is drawn with a circle with the random variable within; here it is X . (The use of a circular node will provide our standard indication for a random variable.) The edges to the children of the node are labeled

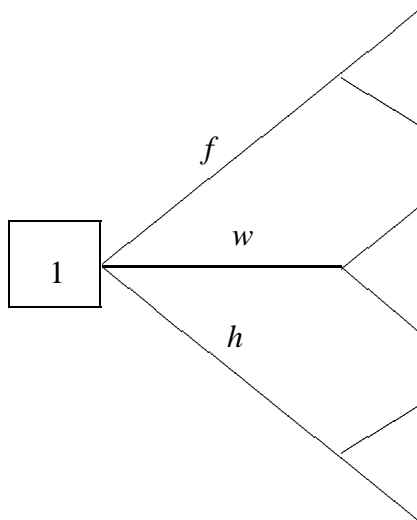


Figure 3: A decision node.

by value and probability; here we see $X = y$ with probability 0.8 and $X = n$ with probability 0.2. For our robot, we can understand X as a boolean observation, derived from automatic analysis of the latest image, which reports whether the agent's field of view contains a certain kind of blob (possibly indicative of a nearby toy ball and so possibly suitable for guiding a following behavior). Thus, the reading $X = y$ is evidence there is a ball to follow; the reading $X = n$ is evidence there is not. The submodels associated with these different outcomes are not explicitly represented; triangles mark where the subtrees have been suppressed for exposition.

The second kind of internal nodes, DECISION NODES or CHOICE NODES, represent a stage of execution of the algorithm at which an action can be performed. Our design space can include a range of different algorithms that involve alternative actions for each choice point. We assume that our design may freely select any of these actions at this stage of execution, whichever we think best. Accordingly, each decision node has a unique child for each alternative action that we are prepared to consider.

Figure 3 illustrates the graphical convention we will use for decision nodes. The node is drawn with a square; an index within the square, here 1, labels the choice for further discussion. (Again, a square will provide our standard notation for a choice.) The label of an outgoing edges indicates the action taken at the decision node for that child's history. Here we consider three possible actions that might be available for

our robotic dog: running a primitive following behavior, represented by f , to track a visual object; waiting for its cameras to accommodate to the ambient illumination, represented by w ; and returning to home base by some robust behavior, represented by h (we might assume this behavior uses a different sensory modality from vision, such as a wireless beacon).

Finally, a leaf node in a decision model represents a state in which our algorithm has terminated and the agent has realized some distribution of possible outcomes. We will represent the degree of success that an agent achieves with a course of action in terms of a real number between 0 and 1, called a UTILITY VALUE. A utility value of 0 represents total failure; a value of 1 represents complete success. Intermediate utility values indicate partial success, the higher the better.

Utility is a human abstraction; it records both the preferences by which we as designers rank alternative real-world results, and our willingness to let our agent take risks whose results are only possibly favorable. As such, utility is not an arbitrary score but rather a precise tool, whose use requires methodological and technical care.

Methodologically, the appeal to utility presumes that we are explicitly committed in design to specific objectives for our agent. Many objectives are reasonable: to minimize the cost in dollars of the resources consumed by the agent; to maximize work achieved by the agent, as measured in the natural units of that work whether they be bytes transmitted in an embedded network or packages delivered in an augmented office; to maximize the satisfaction of the users or owners of the agents, as indicated by reported scores on a subjective rating scale; or several such objectives, in any weighted combination. Typically, our inspiration (or employment) in building an agent impels the choice. For our robotic dog, for example, we can expect performance to be measured by the reports that members of its target market give about how much they enjoy observing and interacting with the toy.

Technically, the appeal to utility demands that such fundamental measures be used indirectly, in order to account for our attitude towards risk. We stipulate that the utility of a risky situation is determined by a summation that counts the utilities of the different possible outcomes in proportion to their probabilities. For example, if we measure utility as dollar cost savings, we assess a risk that sometimes saves \$.75 but is equally likely to save \$.25 as worth \$.50; that is an improvement on a definite outcome that saves \$.49. If we measure utility directly on a subjective scale of user satisfaction, we assess a risk that sometimes pleases the user to degree 0.75 but is equally likely to please the user to degree 0.25 as worth 0.50; that is an improvement

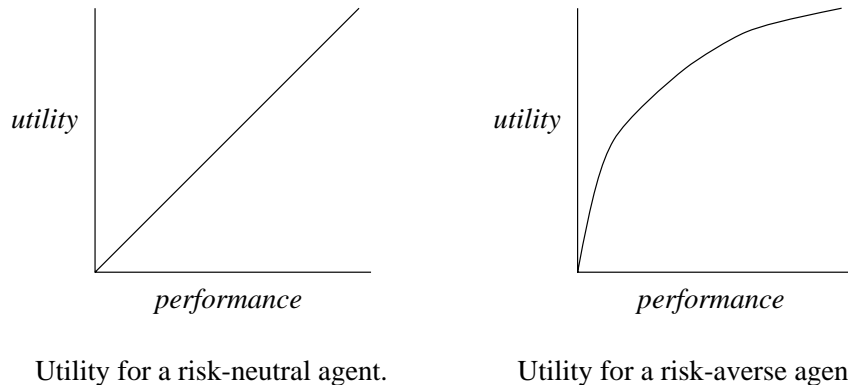


Figure 4: Performance, utility and risk.

on a definite outcome that pleases all users to degree 0.49. Equating measure and utility makes sense in some design problems, such as those where many agent runs contribute independently to the overall benefit of deploying the agent or agents.

However, some problems call for more conservative strategies, called RISK-AVERSE, which instead prioritize surer outcomes. Our robotic dog example could easily prove to be such a problem. For instance, we might want our robotic dog to develop a long, engaging, and consistently satisfying relationship with its principal user, and we might find that an episode in which the dog upsets this person jeopardizes this objective in a way that one equally happy outcome cannot offset. Accordingly, we would want the dog to prefer a strategy that leads to sure but mild satisfaction over one that shoots for thrills but risks disaster.

Utility can accommodate risk-averse strategies through transformations of performance measures that discount higher values relative to lower ones. Figure 4 contrasts the relationship between performance and utility for risk-neutral and risk-averse agents. Thus, to get our dog to prefer to satisfy users less thoroughly but more of the time, we might transform a user rating of 0.25 to a utility of 0.50; a rating of 0.49, to a utility of 0.70; and a rating of 0.75, to a utility of 0.87. (If p is the performance measure and u is the utility, our transformation is $u = \sqrt{p}$.) That original risk now stands at 0.68 and fails to surpass the definite outcome in utility.

Figure 5 illustrates the graphical convention we will use for leaf or REWARD nodes. The diamond is the standard notation for an assessment of outcome; the number within reports the utility achieved.

In our discussion thus far, we have already presented the key ingredients of a decision model for our robotic dog. In summary, we have identified a context for



Figure 5: A reward node.

agent design in which the robotic dog chases a ball into a darkly-lit region; the low illumination of the agent’s cameras in these situations sets up a computational context for specialized decision-making. In this context, the agent has three actions at its disposal: to continue the chase however it can (f), to wait for its cameras to accommodate (w), and to return home (h). It can schedule any of these actions at any stage of execution. However, we rule out two consecutive steps of waiting, because once the agent has waited for its cameras to accommodate, further waiting will not improve the images it obtains but will make the agent less likely to catch the ball or otherwise entertain its owner. Meanwhile, as events develop, the agent has access to a binary observation X . The reading $X = y$ provides excellent evidence that the agent sees the ball it was chasing; the reading $X = n$ provides excellent evidence that it does not.

These considerations determine a design space for the agent as illustrated in Figure 6. The agent begins at node 1 with an initial choice of action. If we design the agent to follow the ball or to return home, we immediately determine the agent’s performance in this situation. If, however, we design the agent to wait, we offer the agent a second perception-action cycle; the agent’s next action is conditional on its reading for X . At this stage, at decision points 2 and 3, we call for the agent to commit to follow or to return home. Consequently an outcome is assured.

Figure 6 describes our alternative designs, but does not describe the success that those designs will achieve in the actual circumstances in which the robot will act. For that, we must provide the probabilities and utilities that Figure 6 as yet omits.

In special cases, we will be able to supply these parameters as part of our specification of the agent. Typically these are cases where the design team can base any values they provide on substantial experience with the task context. Even with such expertise, specification is contraindicated when the exact parameter values greatly affect the design, either because small changes in those values can affect which actions the agent should take or because they can induce substantial differences in the utility the model predicts for those actions.

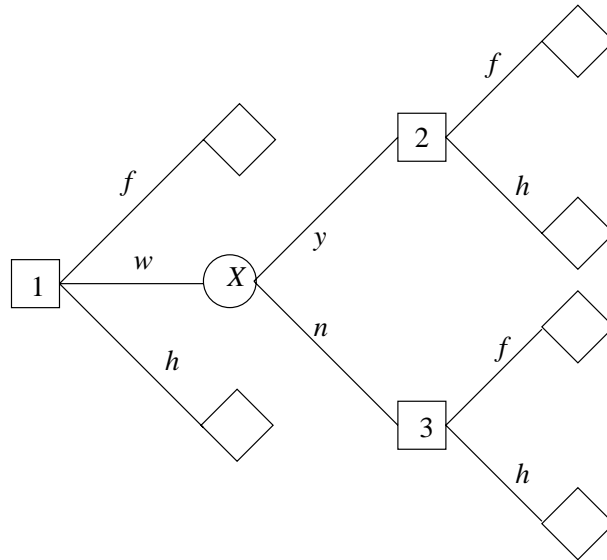


Figure 6: Design space for the robotic dog context.

The alternative is to determine parameter values by empirical investigation. For our robotic dog, that would mean building a prototype that could exhibit the full gamut of possible behavior for the context under investigation. We then estimate parameter values by collecting a range of possible histories with this prototype and assessing the outcome the prototype obtains in each. In other words, this is a problem of MACHINE LEARNING, a topic we return to in Section 4.1.

For the moment, though, we proceed by speculation. As Figure 7 shows, we speculate that users find a successful chase quite entertaining (0.7). On the other hand, for the robot to chase after a ball it can't see is hardly any fun (0.2), and it is all the more boring when the robot waits before doing so (0.1). In between are cases where the robot just returns home, either immediately (0.3) or after some delay (0.25); while these outcomes are not themselves fun, they at least afford another game. Finally, as Figure 7 again shows, we speculate that our robot is able to reacquire the ball with probability 0.8 after waiting to get a better image.

By this point, you may be quite keen to know what it is that our robot should actually do. The answer is the algorithm—or as we shall soon say, the POLICY—presented informally in (1).

- (1) Wait. Then, if $X = y$, follow; if $X = n$, return home.

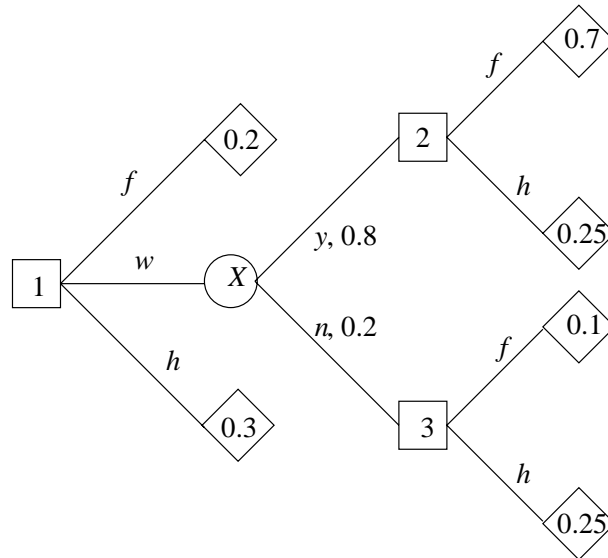


Figure 7: Full decision model for the robotic dog context.

You should be able to convince yourself that by implementing this policy, the robot can achieve a utility of 0.61.¹⁰ You can also check that no alternative choice for our agent at any stage of execution will yield higher utility than what (1) provides at that stage.¹¹ In fact, then, our design space does not offer any policies that improve on (1).

3.2 Implementation

In building our robotic dog, or any other AI agent, we have the opportunity not only to implement specific policies such as (1) but also to work with classes of models and classes of policies in a general way. It is these general computations that offer the most compelling demonstration of the flexibility and elegance of physical symbol systems. This section offers an illustration: here we consider a general implementation of representation and reasoning with decision models.

In presenting this implementation, I will adopt the practice, common in computer science, of abstracting from the details of any particular programming language. All programming languages provide the resources to describe a useful and general inventory of basic types of representation; all programming languages allow collections of these representations to be assembled and organized into regular DATA STRUCTURES. Similarly, all programming languages provide useful and gen-

eral ways to perform basic calculations and to sequence complex computations by examining assemblies of structured representations. In view of these commonalities of implementations across programming languages, presentations of representation and reasoning can proceed in an abstract PSEUDO-CODE that focuses on the essentials, and allows the reader to adapt the discussion easily to the programming language of their choice.¹²

3.2.1 Representation

We begin by describing data structures for decision models and for policies. Our pseudo-code assumes basic types for INTEGERS, as typically realized as binary numerals of fixed length; REAL NUMBERS, as typically realized in accord with standards for floating-point arithmetic operations; and a further basic type, for SYMBOLS, which is rather distinctive to AI. The symbol type provides a set of arbitrary elements for which a programmer maintains an arbitrary and open-ended correspondence with constituents of a real-world environment; programmers specify symbols as character strings and the computational realization of symbols is typically closely related to these strings, but may streamline the implementation to speed up computations.

Our pseudo-code also assumes three ways of packaging data elements together into larger structures. A LIST groups together any number of data elements of the same kind into a sequence whose elements can be accessed only by stepping forward one element at a time. To specify a list, I will give the elements in brackets, as $[1, 2, 3]$; when describing computations that apply uniformly across a list, I will also use ellipses, as $[a_1 \dots a_n]$, and indices, as a_i . A TUPLE groups together a specified number of data elements of specified kinds into a structure each of whose elements can be accessed directly. To specify a tuple, I will give the elements in parentheses, as $(1, 2, 3)$. Finally, a NAMED STRUCTURE provides a way of grouping together heterogeneous collections of data when a representation takes qualitatively different forms in different cases. Each instance of a named structure is written fT where f is a symbol that identifies which form the data takes and T is a tuple that specifies the specific data elements for this instance; the definition of the named structure describes which such cases are possible.

Thanks to these assumptions, we can specify decision models using named structures as in (2).

- (2) A MODEL is a named structure consisting of one of the following:

$$\begin{aligned}
& dcsn(1, [(f, rwd(0.2)), \\
& \quad (w, obsv(X, [(y, 0.8, dcsn(2, [(f, rwd(0.7)), (h, rwd(0.25))])), \\
& \quad \quad (n, 0.2, dcsn(3, [(f, rwd(0.1)), (h, rwd(0.25))]))])), \\
& \quad (h, rwd(0.3))])
\end{aligned}$$

Figure 8: Data structure for the model of Figure 7.

- $obsv(X, os)$ where X is a symbol and os is a nonempty list of tuples (v_i, p_i, m_i) with v_i a symbol, p_i a real, and m_i a model data structure;
- $dcsn(n, as)$ where n is an integer and as is a nonempty list of tuples (a_i, m_i) with a_i a symbol and m_i a model data structure;
- $rwd(u)$ where u is a real.

The three cases correspond to the three kinds of nodes introduced earlier. The symbol $obsv$ keys an observation node; the parameter X represents a variable that the agent can observe, and the list os represents both the observations that the agent might make, and the consequences of those observations. The symbol $dcsn$ keys a decision node; n indexes a specific possible choice for the agent, and the list as determines both what the actions the agent can take and what consequences those actions have. The symbol rwd keys a reward node; the parameter u represents the utility of outcomes that the agent has achieved.

Figure 8 shows how we write the model of Figure 7 as a data structure according to the conventions culminating in (2). Figure 8 doubtless helps motivate the graphical format of Figure 7 as a means for people to communicate decision models! At the same time, however, it underscores that any depicted model also corresponds to a concrete data structure that can be recreated on a machine.

By the same token, the interpretation we sketched for graphical decision models in Section 3.1 now allows us to interpret any model data structure as embodying a precise claim about the world. We start from our design context. This context gives us a correspondence between the action symbols in a data structure and the real actions our agent can take; it gives us a correspondence between observation symbols and the real computations our agent could do to obtain a reading from a specific sensor; and it gives us a correspondence between value symbols and the real readings that our agent could derive from its sensors. Moreover, the context determines the distribution of real-world situations which a model structure as a whole must de-

scribe, for the model proceeds from a computational state triggered in just these situations and spells out the possible consequences of the agent’s subsequent actions.

In general, then, each substructure of the whole model describes the result of particular observations and actions in the context and so will make a CLAIM about a corresponding subset of these situations.

Drawing on this background, we can formalize such claims recursively.

- (3) a A structure $rwd(u)$ claims about situations S that these situations give the agent an expected utility u , and no opportunity for further action.
- b A structure $obsv(X, [(v_1, p_1, t_1) \dots (v_n, p_n, t_n)])$ claims about situations S that these situations all provide the agent a reading corresponding to observation symbol X ; that each result v_i occurs with probability p_i throughout S , determining corresponding situations S_i ; and that the claims of each substructure t_i hold about situations S_i .
- c A structure $dcsn(n, [(a_1, t_1) \dots (a_n, t_n)])$ claims about situations S that these situations all provide the agent an opportunity to act in which the agent can be designed to do any of the actions corresponding to $a_1 \dots a_n$, determining corresponding situations $S_1 \dots S_n$; and that the claims of each structure t_i hold about situations S_i .

A body of definitions such as (3) that relates an agent’s data structures recursively to conditions on its environment is called a COMPOSITIONAL SEMANTICS. A semantics allows us to interpret an agent’s data structures as representations of its environment. The payoff for this interpretation comes in the operations that we define on the agent’s representations. We will be able to show that these operations allow the agent to act as it should if its environment matches what it represents.

To proceed with this development, we also need a data structure that can represent an algorithm by which the agent acts in its environment. As a computer scientist, one learns to take this very deep step quite casually—we can design data structures that represent anything we want, so we can even design data structures that represent rules for carrying out computations! (4) defines a named structure for such an algorithm.

- (4) A POLICY is a named structure consisting of one of the following:
- $switch(X, ps)$ where X is a symbol and ps is a nonempty list of tuples (v_i, π_i) where v_i is a symbol and π_i is a policy data structure;

- $do(a, \pi)$ where a is a symbol and π is a policy data structure;
- a case *end* for which no further data is required.

Informally, $switch(X, [(v_1, \pi_1) \dots (v_n, \pi_n)])$ represents a conditional algorithm in which the agent first carries out the observation corresponding to X and then continues with the policy π_i for whichever reading v_i is obtained; $do(a, \pi)$ represents a sequential algorithm where the agent first performs the action corresponding to a and then continues with the policy π ; *end* represents the final stage of execution at which the policy ends. Thus the policy described in (1) can be implemented in the policy structure of (5).

$$(5) \quad do(w, switch(X, [(y, do(f, end)), (n, do(h, end))]))$$

As with Figure 8, this example serves as a reminder that we could at any point recast our human descriptions in an explicitly computational formalism (and as an illustration of why we may prefer not to).

A given model represents an environment in which only select policies are guaranteed to be executable; we will call such policies **FEASIBLE** in the model.

- (6) a In a model $rwd(u)$, *end* is the one and only policy that is feasible.
- b In a model $obsv(X, [(v_1, p_1, t_1) \dots (v_n, p_1, t_n)])$, a policy is feasible if and only if it takes the form $switch(X, [(v_1, \pi_1) \dots (v_n, \pi_n)])$, where each π_i is feasible in t_i .
- c In a model $dcsn(n, [(a_1, t_1) \dots (a_n, t_n)])$, a policy is feasible if and only if it takes the form $do(a_i, \pi)$, where π is feasible in t_i .

A model determines the expected utility of any policy that is feasible in it, by describing the probability of alternative outcomes under the policy and the utility associated with each outcome. The overall utility (as always) sums the utility of each outcome weighted by its probability. You should be able to construct a recursive definition by cases, analogous to that in (6), which calculates this utility.¹³

3.2.2 Reasoning

Now we turn to the main problem: given a decision model t , compute a feasible policy that has the highest expected utility in t . Such a policy is called an **OPTIMAL** policy for t . The significance of this problem derives from the compositional semantics we have provided for decision models. Suppose that t is an accurate model, so that the claim associated with t truly describes our agent's environment. Then

when the agent acts in this environment, it must carry out a policy that is feasible in t . Moreover, each feasible policy actually has the utility predicted for it by t . Thus, when t is an accurate model, an optimal policy obtains results for our agent that cannot be improved. This section describes a procedure `OPTIMUM` that solves this key problem, and justifies the details of this procedure in a mathematical way.

In order to streamline the mathematical development, we will avail ourselves of the certain constructs in our pseudo-code for procedures. First, we specify procedures by cases, depending on the form of representation of the available data. For example, for `OPTIMUM`, the data is the model structure t ; t may take the form of an observation, a decision, or a reward, and in each of these cases a different course of action is indicated. Be aware that we are abstracting away from the alternative mechanisms programming languages have to achieve such definitions: some do it directly, in procedure definitions; some do it indirectly, by packaging data together with special-case `METHODS` of operation into `OBJECTS`; others can only do it explicitly, with conditional statements.

Second, we assume that procedures can construct an arbitrary structure of data, and `RETURN` it as a result. For `OPTIMUM`, the result will be a tuple (π, u) reporting both the optimal policy π for t and the maximum utility u that the agent can achieve in t . (Evidently, u must be the utility of π in t .) Again, you would do well to remember that programming languages do not all make it equally convenient to deliver such complex results.

Third, we will use general assignment statements of the form in (7) to indicate that the values in the specified *structure* are determined through carrying out the specified *computation*.

$$(7) \quad \textit{structure} \leftarrow \textit{computation}$$

Our definitions will assign a value to any variable at most once. The convention of assignment may therefore be regarded as an instruction to reuse a computed value (like a kind of abbreviation); it does not require the use of program variables with mutable values, which can be changed repeatedly over the course of a computation.

Finally, we will use mathematical notation to specify such routine calculations as summation, finding the largest element in a list, applying a common operation to all the elements of a list, or constructing a new list in a uniform way.

Pause, then, to consider the procedure `OPTIMUM` as defined in the three cases of (8).

- (8) a $\text{OPTIMUM } rwd(u)$
 $\text{RETURN } (end, u)$
- b $\text{OPTIMUM } obsv(X, [(v_1, p_1, t_1) \dots (v_n, p_n, t_n)])$
 for each i : $(\pi_i, u_i) \leftarrow \text{OPTIMUM}(t_i)$
 $\text{RETURN } (switch(X, [(v_1, \pi_1) \dots (v_n, \pi_n)]), \sum_i p_i * u_i)$
- c $\text{OPTIMUM } dcsn(n, [(a_1, t_1) \dots (a_n, t_n)])$
 for each i : $(\pi_i, u_i) \leftarrow \text{OPTIMUM}(t_i)$
 $d \leftarrow$ any i that maximizes u_i
 $\text{RETURN } (do(a_d, \pi_d), u_d)$

We wish to show the CORRECTNESS of this definition; that is, we wish to show, for any tree t , that $\text{OPTIMUM } t$ returns a pair with the optimal policy for t and the maximum utility possible in t . In keeping with the form of Definition (8), which sets out the result of OPTIMUM on a larger model in terms the result of OPTIMUM on smaller model, we will argue for correctness by STRUCTURAL INDUCTION.

Each model structure has a maximum length in the number of steps of perception and action that it offers the agent; call this the HEIGHT of the model. We group together models by height, and consider what happens as height gradually increases.

The simplest models, with height 0, take the form $rwd(u)$. By (8a), here OPTIMUM returns (end, u) . This is the only feasible policy, so it must be optimal, and of course u is its utility.

We now hypothesize for the purposes of argument that OPTIMUM is correct for models of height h or less, and consider a model t of height $h + 1$, for which OPTIMUM returns (π, u) . Now u will be the utility of π , as you can easily verify,¹⁴ so it suffices to show that no feasible policy π' has higher utility than u in t . We prove this by contradiction, considering separately the two possible forms that t could take: $obsv(X, [(v_1, p_1, t_1) \dots (v_n, p_n, t_n)])$, where (8b) applies; and $dcsn(n, [(a_1, t_1) \dots (a_n, t_n)])$, where (8c) applies.

In the first case, π and π' are conditional policies; they spell out a course of action to take for each alternative reading for the observation. According to the clause (8b) which computes π , the overall model t includes a submodel t_i of how the world evolves with any particular reading v_i , and in each case π directs the agent to follow a policy π_i that is optimal for this submodel. Now if π' has higher utility than π , it must have higher utility than π for some possible reading v_i and submodel

t_i , where π says to take π_i for u_i but π' says to take π'_i for u'_i . But by hypothesis, because of the first step of (8b) where we compute OPTIMUM for a model of height h or less, π_i is optimal and such a π'_i is impossible.

In the second case, π and π' are sequential policies that spell out initial actions a_i and $a_{i'}$ and policies π_i and $\pi'_{i'}$ to follow afterwards. According to the clause (8c) which computes π , the overall model t includes a submodel t_i of how the world evolves with any particular action a_i , and π is constructed first to perform that action a_d for which the optimal policy π_d for submodel t_d leads to the best results. (Of course, π continues with π_d .) Now if π' has higher utility than π , then $\pi'_{i'}$ has higher utility than π_d . In (8c), we chose π_d over a policy $\pi_{i'}$ we calculated for i' , based on its utility; that means that $\pi'_{i'}$ must have higher utility than $\pi_{i'}$. But, again, by hypothesis, $\pi_{i'}$ is in fact optimal, so such a $\pi'_{i'}$ is impossible.

We have thus established that OPTIMUM starts out treating models correctly, and that as models increase in height arbitrarily, OPTIMUM continues to treat them correctly. Thus, we have established that OPTIMUM is correct for all models.

I hope that the rigor of this discussion, which epitomizes the precision possible in studying computational operations on representations, has not obscured the central insight which lies behind the procedure of (8) (and its correctness proof, too). To decide how to act, an agent looks forward into the future as far as it can. In the case of decision models, that limit is set by reward nodes, which offer the agent no further actions. From there, the agent works backwards towards the present; at each step it gradually works out the best policy to guide its action in whatever decisions await, but at the same time eliminates from consideration a vast array of policies that will not fare as well. In this way, as the agent goes from step to step, it can focus its deliberation on just its good, live options.

The OPTIMUM procedure is just one of many useful algorithms for working with decision models and policies. An obvious further example is the procedure to EXECUTE a specified policy. Such a procedure must presume a primitive GET that obtains the reading of a specified variable, and a primitive RUN that carries out a specified action. With these resources, the procedure can be specified recursively by case analysis on the structure of policies, along similar lines to OPTIMUM—as you are invited to verify yourself.¹⁵

3.3 *Evaluation*

Let's survey the steps outlined thus far. We have framed our problem, by looking to the world for a coherent task and considering diverse approaches to solve it. We have refined this design space into a model of the agent's environment and task, by assessing the performance of an initial prototype. And we have drawn on general procedures to create a decision-making strategy for our agent that best fits this model of the world. We have, in short, done our very best to make sure that our agent is going to work. And why should we not have? Careful development may be painful, but real-world failure is worse.

Despite these efforts, however, success is never assured. Each step has its dangers. Fundamental approximations and basic design assumptions may not suit the open-ended environment in which the agent must finally act. Training data, too, may misrepresent that final environment, whether through changing circumstances or just bad luck. And a real implementation is not an idealized mathematical construction but a human and imperfect creation, in which frank errors cannot be ruled out.

It is fair to ask, then, does the system work? Does it in fact work for reasons we understand? How might such agents be improved in the future? The premise of evaluation is that any answers to such questions are ultimately subjective judgments, but that they are judgments for which general guidelines are nevertheless available.

The judgment of whether we understand the actual performance of a final agent (the second of our three questions) is often the clearest. We can base this judgment directly on a model of our agent's PERFORMANCE in its environment. A performance model describes the full distribution of outcomes that the agent will achieve by following its policy. As such, it may naturally extend the basic model of the environment that we use to construct the agent.

Take decision models. A decision model is an incomplete description of the distribution of an agent's outcomes in its environment, because each leaf in the model summarizes a whole distribution of outcomes by a single utility value. The model of Figure 7, for example, just says that on any trial our agent has a 0.8 probability of achieving outcomes that have utility 0.7 ON AVERAGE and a 0.2 probability of achieving outcomes that have utility 0.25 ON AVERAGE. This is useful for the agent: the agent really is indifferent between a certain outcome with utility 0.7 (say), and two equally likely outcomes, one with utility 1.0 and one with utility 0.4. But the difference matters for us for evaluation: in the first case, a trial where the agent sees

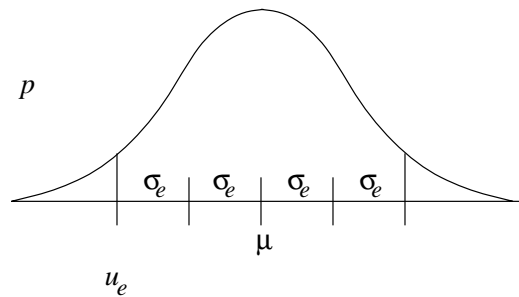


Figure 9: Idealized histogram showing relative probability of different experimental results u_e for average utility across n trials.

an outcome of 0.4 may be quite unexpected; in the second case, it is routine.

In the case of decision models, then, the performance model augments the world model by information about the distribution of possible utilities at each leaf. Naturally, this distribution can be estimated from the outcomes in training, much as the expected utility itself can be. (But you must plan for the evaluation!) In turn, the separate distributions determine the overall distribution of outcomes for a particular policy.

In practice, we just need to know how much variability the outcomes under a policy show for the observed utility u . We can measure this by a statistic called the VARIANCE of the distribution of u . If the mean value for u is μ_u (here the expected utility), the variance σ_u^2 is the value we expect for $(u - \mu_u)^2$. For the model of Figure 7, we will speculate that observed utilities for different outcomes cluster closely around their average values; this hypothesis would determine a value of σ_u^2 in this case of perhaps 0.036.

In order to use this model to assess the performance of the agent, we perform a substantial number n of representative runs with our agent and tabulate the results. Now, individual trials may have unpredictable results, but, if our model is good, our agent's average performance u_e across all n trials of an experiment is very likely to come close to the model average μ_u .

Figure 9 depicts the situation mathematically. It envisages an idealized histogram of the results that our model predicts we would obtain, across many experiments with n trials. The peak in the histogram occurs when the experiment sees outcomes whose average utility exactly matches the model mean μ_u . As the average observed utility diverges from the model mean μ_u , results become increasingly

rare, although any result remains possible in principle. The narrowness of the peak reflects the variance σ_e^2 of results across multiple potential experiments. It can be shown that σ_e^2 is approximately equal to σ_u^2/n , and that with probability approximately 0.95 the result of an experiment u_e will lie no further than twice σ_e from μ_u .

We can now regard our experimental trials as a test of the HYPOTHESIS that the performance of our agent in the real world matches our model of its performance. Our experimental trials have an average utility which diverges by some amount d from μ_u . A divergence as large as d or larger could always be due to chance. Figure 9 illustrates how our performance model quantifies the probability of such chance divergences. In particular, if $d > 2\sigma_e$, the odds of seeing such a divergence by accident are worse than one in twenty. Such unlikely results should raise our suspicions about the model.

In forming a judgment about our hypothesis, we must consider not only the fit between the hypothesis and the data but also the alternative hypotheses are available. Suppose that our results are suspicious. If no available hypothesis explains our results better, our results do not speak against our hypothesis; we simply know that our results truly are unlucky. On the other hand, when competing hypotheses about our agent's performance would fit our observed results better, the poor fit of our hypothesis to the data does provide evidence against our hypothesis.

A standard measure that reconciles these two considerations is the risk of being wrong if our performance model is correct, but we judge our experimental results to be in conflict with it. We calculate this by considering the probability of obtaining results that diverge from our prediction as much as those we have obtained and that some alternative model explains better. For most scientific purposes, we are willing to accept a risk of error of one in twenty. In most agent performance evaluation, we anticipate that our agent could achieve any utility as an alternative; the interactions we do not model could help or hurt the agent. Accordingly, we count an experiment as calling a performance model into question when the observed divergence d is greater than roughly twice the divergence σ_e expected in the experiment.

As a concrete illustration, suppose we run an experiment in which our robotic dog performs the accommodation-and-chasing procedure developed in Sections 3.1 and 3.2 40 times. In these trials, it achieves an average utility of 0.65. By eye it looks as though our agent is doing very well—almost too well. Is it? With this experiment, $\sigma_e^2 = \sigma_u^2/n = 0.036/40 = 0.0009$; so $\sigma_e = 0.03$. Thus, we reject the model only if the observed divergence exceeds 0.06. Here the divergence is 0.04. It appears that

our experimental result of 0.65 is not so unexpected.

We can evaluate other claims about our agent's performance along similar lines. For example, we often want to report that our agent is a success. Typically, we mean by this that our agent yields an improvement over what was previously possible, and we have to test the alternative that our agent performs the same as others, or worse.

Suppose that, in addition to experimental results from our own agent, we have experimental results from a benchmark agent, solving real-world tasks under identical conditions. For new problems, the comparison should be with a simple but reasonable agent; thus our robot dog suggests a benchmark design which always just returns home in cases of difficulty. For well-studied problems, of course, other researchers' work supplies the benchmark agents, and it may not even be necessary to obtain one's own experimental results for them. Agents' performance is often published for standard sets of training and test data made available by the research community.

We will have observed some improvement over the benchmark in our experiments with our agent; otherwise the experiments certainly provide no evidence that our agent is an improvement! As part of our judgment of evaluation, we have to pronounce our opinion as to whether the observed difference reflects a genuine fact about the agents' real-world performance, or whether the difference might have originated by chance. Here the key step is to devise a mathematical argument about the probability of obtaining an IMPROVEMENT as large as we did by chance, if our agent and the benchmark agent really have the same performance. It obviously bolsters our case if this probability is very small.

Where evaluation against a basic benchmark provides a check for the progress we have made with an agent, evaluation against a suitable upper bound gives insight into how much room our design leaves for improvement. Such bounds may be obtained experimentally by investigating how well expert humans are able to perform in the agent's task, or by investigating alternative agent designs whose training and test data sets include important task information that would not normally be available to an agent in practice. More generally, with suitable experimentation and comparison (always backed up by statistical argument), it becomes possible to defend precise claims about the situations, the representations, or the decisions that underlie an agent's strengths or contribute to its weaknesses. Such judgments can be an invaluable guide to further research.

4 Decision Models and Paradigmatic AI Problems

Decision analysis not only epitomizes the practice of agent design, but also manifests the enduring challenges of AI research. In this section, I use decision models as a starting point to delineate these directions for investigation. On the one hand, decision models illustrate all the general difficulties inherent in connecting representations with a complex and uncertain world. Chief among these is the problem of MACHINE LEARNING, exploiting an agent's experience in its environment to ground its models and improve its performance.

On the other hand, decision models turn out to be badly unsuited for most real-world tasks. This leads to a search for new models. The search responds not only to the diverse applications of AI—areas such as ROBOTICS and COMPUTER VISION; HUMAN-COMPUTER INTERACTION and NATURAL LANGUAGE DIALOGUE; MEDICAL and SCIENTIFIC INFORMATICS—but more generally to the diverse cognitive abilities which people manifest across different domains. Better models capture insights that we discover about the real world, in representations that are more general, more robust, or more concise than decision models. But new models also bring new kinds of problems; a case in point is PATTERN RECOGNITION, the computational investigation of discrete responses to a continuous environment.

4.1 *Problems of Learning and the Complexity of Models*

In accounting for the limits of decision analysis—or any very flexible class of real-world models—the central concept is the COMPLEXITY of the model as a description of an agent's environment. A model is complex to the extent that it is free to represent arbitrary relationships that might hold in the environment. By this definition, we may argue that decision models are in fact as complex any model could be.

To see why, suppose we have a decision model t and a policy π that is feasible in t . For expository purposes, we will assume that all executions of policy π terminate after exactly $n + 1$ steps of action, and we will assume that each cycle of perception feeds the agent the value of a binary observation variable. We label these variables X_i with i ranging from 1 through n .

Familiar arguments show that the model t assigns 2^n possible outcomes to policy π .¹⁶ This structure is EXPLOSIVE in size, in the sense that small increases in the number of steps required to complete the agent's task will lead to enormous increases in the computational resources required to represent and reason about the agent's progress. For real-world action, however, it is important to emphasize another side

to this explosion. Small increases in the number of steps required to complete the agent's task will lead to enormous increases in the amount of REAL-WORLD EXPERIENCE required to construct a full and meaningful model.

The explosion originates in the fact that the model t and policy π are capable of determining an ARBITRARY joint distribution on the values of variables X_1 through X_n . To make this precise, let $P(X_1 = x_1 \dots X_n = x_n)$ denote the probability that the observations X_1 through X_n take on a specified combination of readings x_1 through x_n during a history with the agent. Then a joint distribution is any assignment of probability values for $P(X_1 = x_1 \dots X_n = x_n)$ across possible observations, and whatever this assignment may be, we can construct a model to realize it.¹⁷

Now, in Section 3.1, we saw that, except in those few cases where designers bring very precise and accurate background knowledge to the construction of a model, the free parameters of a model must be estimated from a set of training data. The field of MACHINE LEARNING¹⁸ offers diverse frameworks for this estimation problem. A simple statistical approach is to look for a model that maximizes the overall likelihood of the observed data; a more sophisticated one is to interpret the data as giving indirect evidence about a restricted set of likely parameter values that we supply in advance. Still other approaches recover parameters from data by more general optimization techniques that have proven effective throughout a range of practical distributions of problem instances; popular instances of this kind of strategy include SUPPORT VECTOR MACHINES¹⁹ and (despite their name) NEURAL NETWORKS.²⁰

Machine learning also affords numerous perspectives on how to process the evidence derived from data. The data may come in a single batch; in this case, learning algorithms can use any data at any stage of computation and can report just the final set of parameter values. By contrast, in INCREMENTAL LEARNING, the learning algorithm iteratively reconciles a provisional model with a new observation to derive a better model. Moreover, when action is involved, we may emphasize how the agent's policy changes with new observations (and refinements to parameter values); this is the perspective of REINFORCEMENT LEARNING.²¹

For all this diversity, no machine learning method can overcome the fact that training data is noisy. In general, empirical observations of the real world reflect underlying regularities of the world only partially. Our introduction to evaluation illustrated this point already in detail. We saw that repeated observations are necessary in order to derive an estimate for a parameter value that will be close to the true

value with high probability.

Thus, to build a decision model from data requires repeated observations of each of the 2^n possible outcomes. With a training set of practical size (and don't forget that the limit is often the real-world effort—even the real-world cash—required to obtain the training data) and with a decision model of substantial complexity, it is inevitable that many of these outcomes will not occur, or will occur so infrequently that any derived parameter estimates are quite likely to be inaccurate. Under such circumstances, decision models are unlikely to be useful. Readers familiar with linguistics or psychology will recognize this as a version of the Chomsky's POVERTY-OF-THE-STIMULUS ARGUMENT: an intelligent agent's open-ended behaviors must be informed not only by limited experience in the world but also by innately constrained models of the world.²² In this connection, it is worth reinforcing that the inference to restricted models of the world applies regardless of how learning is understood or how learning proceeds.

Sensor readings with a large space of possible values provide a complementary illustration of the complexity of decision analysis. Consider an observation D that approximates the distance to some obstacle. In a digital computer, the values for D will be discrete; the number of different possible values will be finite. In principle, a decision model can describe this, by inducing a distinct submodel for each alternative reading. When the number of alternative values for D reaches 256, 65536, or more, the amount of training data required also skyrockets; the model again becomes hopeless.

4.2 Problems of Modeling and the Complexity of Environments

We see how severe a drawback this complexity of decision models can be, once we distinguish between the complexity of an agent's models and the complexity inherent in the agent's environment. An ENVIRONMENT is COMPLEX to the extent that its outcomes inherently reflect intricate interrelationships among many causal processes. A complex environment offers little basis to generalize from one circumstance to another; in a complex environment, designers may have no alternative but to equip their agents with a model that is accurate and that therefore is complex, too. In a complex environment, designers must come to terms with the concomitant difficulties in computation and training.

However, a wide range of important environments are not complex in this sense. Indeed, a computational perspective on human intelligence suggests that none of

the behaviors that people naturally develop appeal to descriptively complex models of a causally complex environment. Instead, what frequently creates difficulties for agents is that the causal processes at play in these environments, though simple, cannot be observed directly. We often find, for example, that the state of the environment determines a straightforward distribution on the outcomes of actions, but that the observations of the agent can provide only fragmentary evidence about what this true state of the environment actually is.

For example, consider designing a behavior where our agent classifies and responds to a new object in its environment. Think of our robot pet reacting to an approaching object that is brought to its attention by the looming in its visual field: Is this a toy, which the agent might pick up or fetch? Is this (part of) a person, who the agent might attempt to interact with? Or is this something dangerous, which the agent must avoid? In choosing its response to the object, the agent's performance depends on whether it deals with the object as befits its category.

Unfortunately, the agent cannot directly observe this. The agent has only its observations—some perceptual features of the new object. Our robot pet, for example, might be able to estimate an approaching object's speed, size, or shape. A decision model could only attempt to connect these observations directly to the choice of action the agent must make.

However, a model that incorporates an explicit representation of the true but hidden state of the world may be simpler, if it concisely describes the underlying generalizations in the environment. In our example, such a model could spell out how the agent's observations of an object correlate with the object's category. Toys might tend to be small, and move at a medium speed; people might prove slower, larger and characteristically blobby; dangerous objects might appear fast, or large, or even just pointy. By describing the baseline frequency of these different cases—say toys and people are relatively common while danger is thankfully rare—the model allows the agent to interpret its observations as giving evidence about what category that the object belongs to, and to react appropriately. In effect, then, this model links an agent's behavior to its CATEGORIZATION of its environment. Of course, in designing agents, we need not always think of objects or categorization as concretely as this; we can apply these models, and our intuitions about categorization, in open-ended ways.

To bring home the generality of this idea, let us consider a class of models that represents the hidden state that the agent must infer and respond to as a random vari-

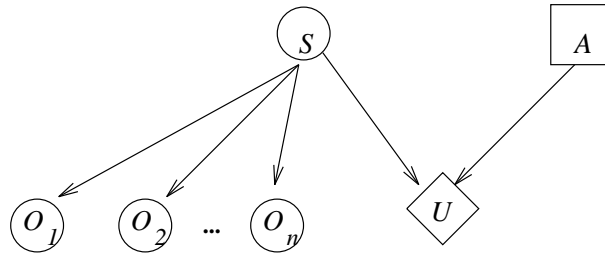


Figure 10: Naive Bayes Model

able S . The model specifies a constrained range of possible values for S , each of which corresponds to a distinct set of real-world circumstances that arise with a specified frequency in the agent’s experience. Mathematically, the model represents this by a PRIOR distribution $P_m(S = s)$ on the alternative values for S .

The model also describes these circumstances (approximately), by associating each value of S with a characteristic distribution of observations and outcomes for the agent. The model represents the agent’s observations as noisy, incomplete but independent clues to this state. Mathematically, for each observation variable O_i , the model specifies the LIKELIHOOD of seeing a particular reading o assuming the state has value s , or $P_m(O_i = o|S = s)$. Meanwhile, the model links the utility of each action just to the underlying state of the world. Mathematically, the model describes these outcomes with a table $U(s, a)$ of the utility expected when the state of the environment is s and the agent’s choice of action is a .

Diagrams such as that in Figure 10 use the standard convention of a directed arrow between nodes to depict these probabilistic dependencies. In Figure 10, the state is represented by the random variable S . The agent’s observations, represented by random variables O_1 through O_n , appear as nodes linked only from S . The agent’s utility, represented by the reward node U , is linked both from S and from the choice A that the agent makes. These representations amount to assumptions about the environment that we make in using such models to design an agent. When these assumptions fit, they allow an agent to reason successfully, to learn efficiently from limited training data, and to generalize reliably from it.

The agent’s behavior involves making a series of observations and then responding with an appropriate action; the model provides the agent with grounds for this choice as follows. The agent first estimates the probabilities of alternative states of the environment, by matching its observations against the model. The agent then

determines the most promising action, by calculating utilities in the model based on this inferred distribution. If the model describes the environment accurately, this policy allows the agent to act as successfully as possible.

More precisely, suppose that the agent has observed values o_1 through o_n for the variables O_1 through O_n . The model determines a probability distribution $P_m(S = s|O_1 = o_1 \dots O_n = o_n)$ describing how probable any state of the world is given the observations the agent has made. The probability $P_m(S = s|O_1 = o_1 \dots O_n = o_n)$, called the POSTERIOR, represents the agent's BELIEF that the environment is in state s ; thus the posterior distribution summarizes the information that the agent has about its environment. Using the posterior, the agent can determine the expected utility of any action a given its observations, which we write $U_m(a|O_1 = o_1 \dots O_n = o_n)$, by weighting the outcome for a in alternative possible states s by the agent's current belief in s . (9) formalizes the relationship.

$$(9) \quad U_m(a|O_1 = o_1 \dots O_n = o_n) = \sum_j U_m(s_j, a) P_m(S = s_j|O_1 = o_1 \dots O_n = o_n)$$

Of course, the agent's best strategy is to take the action with the highest expected utility.

The models depicted in Figure 10 are known as naive Bayes models. The Bayes part is because we use Bayes's Theorem²³ to calculate (9) from the parameters of the model. (10) shows the instance of Bayes's Theorem we need.

$$(10) \quad P_m(S = s|O_1 = o_1 \dots O_n = o_n) = \frac{P_m(O_1 = o_1 \dots O_n = o_n|S = s)P_m(S = s)}{P_m(O_1 = o_1 \dots O_n = o_n)}$$

In words, Bayes's Theorem states that our BELIEF, that the environment is in state s given our observations, trades off the LIKELIHOOD of obtaining our observations if the environment was in state s against the PRIOR probability that the environment is in state s . We obtain a probability by normalizing for the overall probability of making our observations—the normalization factor can be calculated by summing the probability of being in a specified state and making our observations, over all possible states.

The model's "naive" independence assumptions entail (11).

$$(11) \quad P_m(O_1 = o_1 \dots O_n = o_n|S = s) = \prod_{i=1}^n P_m(O_i = o_i|S = s)$$

We can draw on the results so far to justify an algorithm for computing the OPTIMUM action after observations o_1 through o_n according to a naive Bayes model m .

(12) a OPTIMUM $m, o_1 \dots o_n$

for each j : $p'_j \leftarrow P_m(S = s_j) * \prod_i P_m(O_i = o_i | S = s_j)$
 $p_o \leftarrow \sum_j p'_j$
 for each k : $u_k \leftarrow \sum_j (U_m(a_k, s_j) * p'_j / p_o)$
 $d \leftarrow$ any k that maximizes u_k
 RETURN (a_d, u_d)

By (10) and (11), the algorithm ensures that the value of p'_j / p_o is $P_m(S = s_j | O_1 = o_1 \dots O_n = o_n)$. Thus by (9), the algorithm stores in u_k the expected utility in m of action a_k .

We can also show how we can eliminate much of the complexity of a decision model in a naive Bayes model. To illustrate, we will use $|S|$ to represent the number of states in the model and $|A|$ to represent the number of actions available to the agent; we will again assume that the observations are binary.

The naive Bayes model determines a utility for each combination of observations and actions. This is $|A|2^n$ values in this case. In a decision model, this would be prohibitive; each value would have to be represented, and learned, separately. But in the naive Bayes model, these $|A|2^n$ values are derived indirectly, using a number of parameters that is potentially much smaller. The naive Bayes model has a table of utilities with $|A||S|$ entries; it has a table of priors with $|S| - 1$ entries; and it has n tabled likelihood functions, each with $|S|$ entries. That is a total of $(|A| + n + 1)|S| - 1$ parameters.

We may be able to identify a reasonable number of meaningful states in the world as part of designing our agent, so that $|S|$ is very small by comparison to 2^n . Then the naive Bayes model offers a vast improvement in size; there is a corresponding improvement in the amount of memory required to represent the model and the amount of training data required to build it.

Of course, not all environments can be described in terms of a single variable that takes on a small number of values. Richer environments demand richer models which offer an agent ways to reason about the state of the world more systematically. Indeed, effective modeling may mean understanding at a high level how and why an environment works, and realizing that explanation in computational terms; an agent may need this high-level explanation to select the actions that it should take based on the observations, and more generally on the experience, available to it.

Fortunately, this modeling need not start from scratch. Designers have ready sources of causal hypotheses, from their scientific understanding of the real world,

from evaluations of prior agents, and even from intuitions about our own action in the world. At the same time, AI research offers an inventory of broadly useful models that implement common forms of explanation. I close with a brief description of two further classes of models, to suggest how current AI practice makes the design of computational artifacts that act in the world increasingly routine.

In many domains, an agent needs to recognize a pattern of change over time from a sequence of observations. The agent may be tracking a moving object through its visual field; at each step, the agent's current image of the object offers a noisy and incomplete evidence about the object's current location. The agent may be listening to a speech signal; at each step, the agent's acoustic observations gives evidence about the phones that the speaker is currently articulating. More generally, the agent may need to act to influence a dynamic environment; its observations give it noisy feedback not just about the current state of the environment but more generally about the actions that the agent should take in the future.

The naive Bayes model requires us to model the likelihood of any observation $P_m(O_i = o | S = s)$ directly in terms of the overall state $S = s$. In this case, since the environment is dynamic, the value s must determine the specific conditions in the environment at the step i when the observation is taken. Thus, to model all observations in combination, S must describe complete histories in the environment. Histories proliferate too explosively for this to be practical.

Models for dynamic domains must therefore describe the succession of states that the environment evolves through over time. The state in step i is modeled as a random variable S_i . The observation O_i at that step depends only on the current state, so the state is hidden. The state S_{i+1} in step $i + 1$ depends only on the previous state S_i and the previous action taken A_i ; the relationship does not depend on the value of i . This MARKOV ASSUMPTION ensures that the model need not grow in complexity in order to describe longer and longer runs of observation and action.

Figure 11 depicts one possible realization of this pattern of probabilistic dependencies. In the behavior modeled in Figure 11, the agent carries out three cycles of observation and action before obtaining a reward that depends just on the final action and the final state.

The Markov assumption allows us to collapse together complete alternative histories for the purposes of training; if we LABEL the states in the training data, we can fit a model simply based on the observations and transitions that we observe in each state. The Markov assumption also makes it efficient for an agent to calculate its

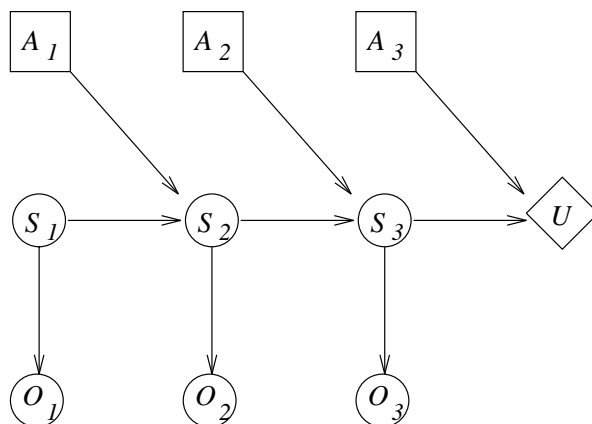


Figure 11: A model for three execution cycles, with a Markov assumption.

belief about the current state of the environment given the observations and actions that it has seen so far; this algorithm helps explain the ubiquity of Markov models in AI applications such as speech recognition, language processing and visual tracking. Unfortunately, the Markov assumption does not allow an agent to determine an optimal policy exactly without reasoning about complete alternative histories. Current research focuses instead on developing useful approximations that find good policies in practice.

More complex environments may require an agent to infer hierarchical relationships in the world to explain a collection of observations. Language use is one case—the meaning of a sequence of words depends on the syntactic relationships that group words hierarchically into constituents—but not the only one. In visual perception, for instance, inferring the real-world objects and relationships behind an image involves grouping image features together hierarchically into units that correspond to parts of objects, wholes of objects, and complexes of objects.

As you might expect, it can prove explosive to recognize hierarchical relationships without hierarchical models. We can account for the structure of a sentence by specifying the hierarchy of constituents to which each word contributes; we thereby associate the sequence of words we observe with a sequence of hidden states. Why not model this sequence of states atomically, using just a Markov assumption? In general, the number of possible states will increase as we consider progressively longer sentences; the more complex structure of longer sentences makes for states with a greater number of open constituents. The effect is familiar: training

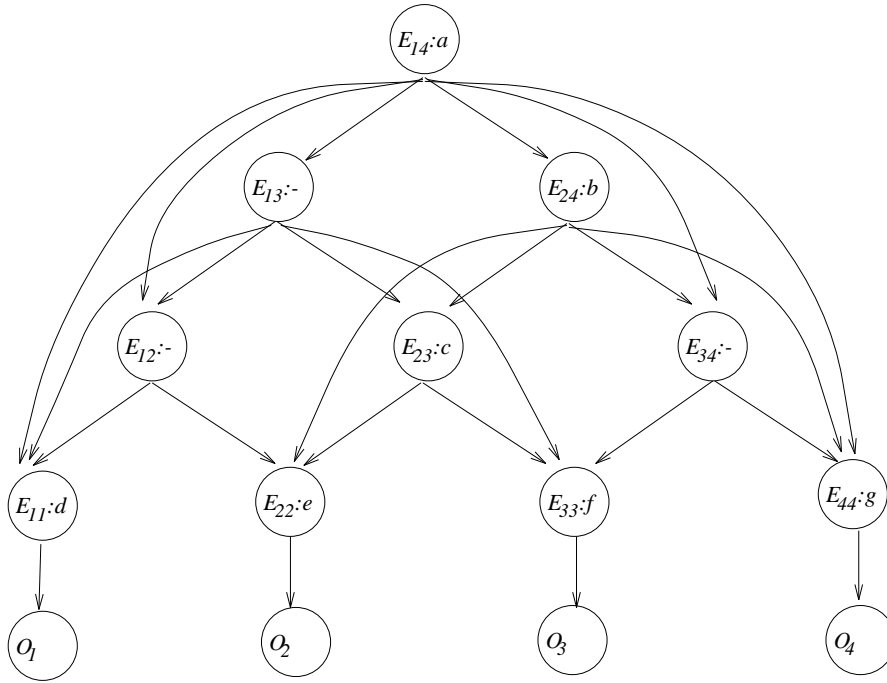
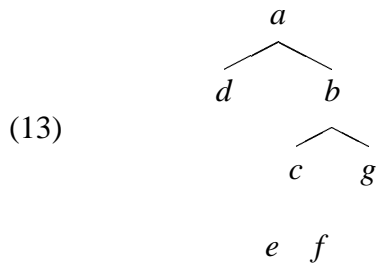


Figure 12: A PCFG model for four observations.

on shorter sentences may not generalize to longer ones.

The probabilistic context-free grammar (PCFG) is one class of model that does describe hierarchically-related processes that account for sequences of observations from the environment. Let us assume that the PCFG is BINARY BRANCHING, in that the lifetime of any extended process is split between a first child process active initially and a second child process active finally. This PCFG model sets up variables E_{ij} indicating what kind of process, if any, begins with observation i and ends with observation j . The variable E_{ij} may depend probabilistically on any variable E_{ik} with $k > j$, because a process active from i to k may have a process from i to j as its first child; E_{ij} may also depend on any variable E_{kj} with $k < i$, because a process active from k to j may have a process from i to j as its second child.

Figure 12 illustrates the situation. The four observations induce the six variables E_{12} , E_{13} , E_{14} , E_{23} , E_{24} and E_{34} for possible processes spanning multiple observations. The figure depicts a possible analysis of the observations, in which processes in E_{14} , E_{23} and E_{24} contribute to the derivation but no processes are active for E_{12} , E_{13} and E_{34} . The analysis corresponds to the tree in (13).



Perhaps surprisingly, even with such sophisticated analysis and such intricate dependencies, it is still efficient for an agent to compute its belief about the current state of the environment given its observations.

5 Directions, Connections, Distinctions

Better than any reasoned argument, these diverse examples attest to the open-ended enterprise of modeling the world, and the open-ended enterprise of AI itself. In a chapter such as this, I cannot hope to catalogue all possible problems, all possible models, even all possible approaches. Rather my aims have been to introduce the goals and perspectives of AI practice, in Section 2; to acquaint you in detail with a concrete case of AI practice, through Section 3; to hint at the problems and possibilities to which past AI research has led, in Section 4; and, perhaps with all this, to whet your curiosity for the additional results—indeed the future results—which AI research still offers.

These additional results include the numerous approaches to intelligent behavior based purely on optimization to which I have given short shrift. These engineering approaches are effective and useful, and historically have been just as inspired by biological computation—and just as inspiring in the effort to understand biological computation—as the more severe model-based techniques I have emphasized.

I have also passed over a rich tradition of purely logical models of the world. The motivations for such models have always been diverse (and contentious); for me, such models contribute a worst-case analysis of intelligent behavior. Logic demands that we avoid risk altogether by formulating models and policies that must succeed, within limits, no matter how the world might turn out to be. This worst-case design provides the only principled way to back off from dependence on distributions derived from empirical data to strategies that apply in genuinely novel situations—that is, even in the most remote possibility.

Even within the probabilistic models that this chapter has surveyed—beginning with decision models, and continuing with naive Bayes models, Markov models

and PCFGs—I have had to omit a flourishing range of models for perception and action. Graphical representations of probabilistic dependencies, as illustrated in Figure 10, Figure 11 or Figure 12, can describe arbitrary causal interactions in an environment.²⁴ However, as models become more sophisticated, it becomes increasingly difficult to tie them to the training available in concrete domains or to algorithms that enable tractable decision-making. Research proceeds.

The diverse investigations of AI aim for agents that are sensitive to empirical regularities in their surroundings and that exploit these regularities efficiently to thrive in their environment. These are strengths that biological intelligence also exhibits. But AI research can base its strategies and computations on designs that depart radically from biology; as it does so, it increasingly delivers intelligent real-world behavior in ways that diverge from biological intelligence. For example, in its chess matchup with Kasparov,²⁵ IBM's Deep Blue relied on hardware and software that was capable of performing billions of operations per second and that in fact pursued extravagant search forward move-by-move through possible developments of the game; such a strategy seems impossible for a brain that makes up in parallelism what it lacks in speed. Mismatches proliferate among the many important problems which people master only with difficulty: medical diagnosis, investment, scientific data analysis, engineering design.

Such examples underscore the uniqueness of AI as an engineering approach to building computational artifacts that act in the real world. Practitioners of AI must take this engineering perspective to heart to have an impact on the way people design and build useful agents. Cognitive scientists must also come to grips with it if they are to exploit new ideas from AI research which promise to inform our scientific study of ourselves.

But all should welcome the rigor that the view inspires. In engineering, an *ad hoc* design—however ingenious—is easily discarded as problems and technology change; a principled design—a model of the world and the computational results to put the model to use—remains a lasting contribution. In cognitive science, theories of the functional value of human cognitive abilities will remain speculation until connected with an engineering model of the problems people face and the performance that people achieve in everyday circumstances. Inspired by such inherent connections, I for one expect the long and fruitful interchange between AI and cognitive science to continue.

Notes

¹ With his paper *On computable numbers, with an application to the Entscheidungsproblem* (Proceedings of the London Mathematical Society, 1937), Alan Turing offered his celebrated characterization of all mechanical procedures which can be executed according to a finite set of precise instructions.

Noam Chomsky's *Syntactic Structures* (published 1957) and the more incendiary *A review of B. F. Skinner's Verbal Behavior* (Language, 1959) decisively introduced linguistics, philosophy and psychology to a computational interpretation of the rules of linguistic structure.

Alan Newell and Herbert Simon's *general problem solver* (implemented 1957 with programmer J. C. Shaw, and described in a contribution to Feigenbaum and Feldbaum's *Computers and Thought*, 1963) invoked an explicitly symbolic conception of computation as a general way to derive goal-directed real-world behavior by algorithmic processes.

² David Hilbert stormed through mathematics at the end of the nineteenth century and beginning of the twentieth; he published groundbreaking work in algebra, number theory, geometry, and the foundations of physics and mathematics. But he concluded his career with investigations into proof theory whose goal was the demonstration of the consistency and completeness of mathematical method; the theory of computation, as developed by Kurt Gödel, Alan Turing and others, dissolved this goal into impossibility.

Leonard Bloomfield's text *Language* from 1933 set the agenda for American linguistics for the next twenty-five years. The book is remarkable both for its meticulous methodological advice, which helped a generation of fieldworkers document the words and sounds of the indigenous languages of North America, and for its frank admission that, to study such matters as linguistic meaning, linguistics awaited some new scientific tools (in retrospect, that tool was computation).

B. F. Skinner's *The Behavior of Organisms* (published 1938) advanced a particular kind of associative learning, the pairing of stimulus and response through reinforcement, as a uniform explanation for the results of experiments documenting the ability of animals to act successfully in problematic environments. Skinner's behaviorist approach left no place for psychological representation, or for computation, in the explanation—an omission that ultimately proved the theory's undoing.

³ One sometimes feels that the term AGENT has a more insidious working definition: any computer program that happens to be implemented by a researcher with AI training. I abhor this usage—and its implication that AI is nothing more than a community of seat-of-the-pants hackers with sloppy heuristics and cliquish attitude.

⁴ The landmark implementation is described in W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, “The Interactive Museum Tour-Guide Robot,” Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98), 1998.

⁵ Rudolf Kalman introduced this discrete-time filter in his paper “A new approach to linear filtering and prediction problems”, in *Journal of Basic Engineering—Transactions of the ASME, Series D*, 83(1), pages 35–45, 1960. It was used right away in aerospace, for example for navigation in the Apollo space program. But its heavy burden of computational operations and numerical precision meant that the filtering process typically had to be approximated and simplified to run on available hardware; see for example Peter Maybeck’s text *Stochastic Models, Estimation, and Control*, Academic Press, 1978.

⁶ Newell’s perspective, articulated for example in *Unified Theories of Cognition* (Harvard University Press, 1980), suggests an intimate relationship between AI and computational perspectives on human cognition.

⁷ With his theory of signs, Charles Sanders Peirce (1839–1914) offered an account of the parallels and diverges between linguistic meaning and natural meaning. Linguistic meaning is symbolic; the word *fire* for example is connected with fires in an arbitrary way. Smoke also means fire, but smoke is what Peirce called an INDEX: such signs are linked to what they mean by correlations in the natural world. See *The Essential Peirce, Volumes 1 and 2*, Indiana University Press, 1992 and 1998.

⁸ *Making Hard Decisions* (Robert Clemen, Duxbury Press, 1995) is a representative text on decision analysis—for the MBA curriculum.

⁹ See for example Ronald Arkin’s *Behavior-Based Robotics*, MIT Press, 1998.

¹⁰ A fraction of 0.8 of histories with this policy lead to an observation of $X = y$, to following, and then to a utility of 0.7; a fraction of 0.2 lead to an observation of $X = n$, to returning home, and then to a utility of 0.25. Weighting these outcomes by their probability gives $0.8 \times 0.7 + 0.2 \times 0.25 = 0.61$.

¹¹ At choice 1, the other options’ utilities of 0.2 and 0.3 are less than the 0.61 we get by waiting and continuing with (1). At choice 2, our policy of following for 0.7 beats the alternative of returning for 0.25. At choice 3, our policy of returning for 0.25 beats the alternative of following for 0.1.

¹² Indeed, after studying this section, you might profit from constructing your own definition of its two key data structures, the MODEL and the POLICY, and from accomplishing your own implementation of its two key algorithms, OPTIMUM and EXECUTE—either in a familiar language or even in one of the specialized languages

of AI, such as Prolog, Scheme or SML.

¹³ Here is one such definition:

- (14) a In a model $rwd(u)$, policy end has utility u .
 b In a model $obsv(X, [(v_1, p_1, t_1) \dots (v_n, p_1, t_n)])$, policy $switch(X, [(v_1, \pi_1) \dots (v_n, \pi_n)])$ has utility

$$\sum_i p_i * u_i$$

where u_i is the utility in model t_i of policy π_i .

- c In a model $dcsn(n, [(a_1, t_1) \dots (a_n, t_n)])$, policy $do(a_i, \pi)$ has utility u_i where u_i is the utility in model t_i of policy π .

¹⁴ Just refer to (14).

¹⁵ Here's one specification of EXECUTE.

- (15) a To EXECUTE end , nothing is required.
 b To EXECUTE $switch(X, [(v_1, \pi_1) \dots (v_n, \pi_n)])$

$o \leftarrow \text{GET } X$
 $j \leftarrow \text{the } i \text{ for which } o = v_i$
 EXECUTE π_j

- c To EXECUTE $do(a, \pi)$

RUN a
 EXECUTE π

¹⁶ At the first step of action, there is one outcome; after each step of observation, the number of outcomes doubles; by the time the policy terminates, the leaves number $2 * \dots * 2$ (n 2s); thus, 2^n .

¹⁷ This fact admits two kinds of explanations. Conceptually, we can see each layer i of observation in the model as specifying a distribution of variable X_i that is conditional on specific values of any earlier observations X_1 through X_{i-1} ; this distribution is written $P(X_i = x_i | X_1 = x_1 \dots X_{i-1} = x_{i-1})$. In general, any joint distribution on multiple variables (perhaps conditioned on some information C) can be factored as the distribution of the first (given C) times the distribution of the others given the first (and C), as in (16):

$$(16) \quad P(X_1 = x_1 \dots X_i = x_i | C) = P(X_1 = x_1 | C) P(X_2 = x_2 \dots X_i = x_i | X_1 = x_1, C)$$

Unfolding (16) $n - 1$ times, then, we see that the distribution provided at each layer is exactly what is needed. More algebraically, and more significantly for present purposes, the joint distribution is specified by $2^n - 1$ parameters and the model t also has exactly $2^n - 1$ free parameters for probabilities (the counts reflect the fact that probabilities always sum to 1). Matching these parameters is thus a matter of solving appropriate equations.

¹⁸ The standard introduction to machine learning is Tom Mitchell's *Machine Learning*, McGraw-Hill, 1997.

¹⁹ See Nello Cristianini and John Shawe-Taylor, *An Introduction to Support Vector Machines*, Cambridge University Press, 2000.

²⁰ See Christopher Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.

²¹ See Richard Sutton and Andrew Barto, *Reinforcement Learning: An Introduction*, MIT Press, 1998.

²² The famous formulation—in which Chomsky argues that people's productive use of language must be informed not only by experience with language but also by models of language embodying substantive universals—is in the collection *Language and Learning: The Debate Between Piaget and Chomsky* edited by Massimo Piatelli-Palmerini and published by Harvard University Press, 1980.

²³ The result, first explored by minister and sometime mathematician Thomas Bayes (1702–1761), is easy to rederive. Just start from the two ways to express the joint probability of two events, a possible cause c and its possible effect e :

$$(17) \quad P(c, e) = P(c|e)P(e) = P(e|c)P(c)$$

Divide by $P(e)$, to see how to infer the cause c given a model of its effects:

$$(18) \quad P(c|e) = \frac{P(e|c)P(c)}{P(e)}$$

²⁴ The seminal monograph is Judea Pearl's *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann, 1988.

²⁵ Documented on the internet at <http://www.research.ibm.com/deepblue/home/html/b.html>.