Quantum Physical Symbol Systems

Kathryn Blackmond Laskey KLASKEY@GMU.EDU
Department of Systems Engineering and Operations Research
MS4A6
George Mason University
Fairfax, VA 22030, USA

Abstract

Today's theories of computing and machine learning developed within a nineteenth-century mechanistic mindset. Although digital computers would be impossible without quantum physics, their physical and logical architecture is based on the view of a computer as an automaton executing pre-programmed sequences of operations exactly as instructed. Recent innovations in representations and algorithms suggest that a shift in viewpoint may be occurring. In the newly emerging view, a computer program executes a stochastic process that transforms inputs and internal state into a sequence of trial solutions, as it evolves toward an improved world model and better task performance. A full realization of this vision requires a new logic for computing that incorporates learning from experience as an intrinsic part of the logic, and that permits full exploitation of the quantum nature of the physical world. Knowledge representation languages based on graphical probability and decision models have now attained sufficient expressive power to support general computing applications. At the same time, research is progressing rapidly on hardware and software architectures for quantum computing. It is hypothesized that a sufficiently expressive probabilistic logic executing on quantum hardware could perform Bayesian learning and decision-theoretic reasoning with efficiency far surpassing that of classical computers. Moreover, a computer architecture based on graphical models can support multi-level representations spanning the sub-symbolic through the cognitive levels, opening the possibility for human-computer interfaces unprecedented in their adaptability and flexibility.

Keywords: Bayesian networks, decision theory, graphical models, quantum computing, quantum measurement

1 Introduction

The challenge of Thagard (1988) on which this special issue is based was presaged by Newell and Simon's 1976 paper urging an empirically grounded approach to theories of computation, cognition and intelligence. Intelligence, Newell and Simon argue, is a property possessed by physical agents operating in a physical environment, whose cognition is performed by a physical brain and nervous system. Therefore, the means by which an agent identifies, evaluates, selects among, and implements policies to control its motor system is constrained by the laws of physics. A scientifically valid theory of intelligent behavior must be formulated in a manner that supports assessment of its degree of empirical confirmation. In particular, theories of intelligence must be tested by implementing them in physical devices that can carry out the required computations and actions rapidly and accurately enough to allow the agent to operate successfully in its environment. Thus, Newell and Simon regard intelligence as a property possessed to greater or lesser degree by a *physical symbol system*, which they define as:

> "...a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called an expression (or symbol structure).

> [Symbols in a structure] are related in some physical way… A physical symbol system … produces through time an evolving collection of symbol structures. Such a system exists in a world of objects wider than just these symbolic expressions themselves."

Two essential capabilities a physical symbol system must possess are *designation* and *interpretation*. Symbol structures can designate objects in the world external to the system, thus allowing the system to affect and/or be affected by the designated object. Symbol structures designating a sequence of actions can be interpreted, thus allowing the system either to act in the world or to issue instructions to control the actions of an external system. Given this definition, Newell and Simon articulated the *physical symbol system hypothesis*:

> "A physical symbol system has the necessary and sufficient means for intelligent action."

They offered the physical symbol system hypothesis as a rigorously formulated, empirically falsifiable scientific hypothesis about the nature of intelligence. They challenged the community to develop the hypothesis further and subject it to empirical test. The intervening years have seen a substantial literature on the physical symbol system hypothesis, the correctness of which is viewed by many as a fundamental tenet of artificial intelligence. Considerable progress has been achieved in both theory and applications of computational intelligence. Attempts to build devices that behave intelligently have resulted in partial success, in vastly improved understanding of the nature of the challenge, and in a better appreciation of what we do and do not understand about intelligence and its role in Nature. The resulting cross-fertilization between artificial intelligence, biology and cognitive psychology has led to advances in all these disciplines.

The search for empirically grounded computational theories of intelligence raises the possibility of a new and far-reaching integration of our theories of the physical and the informational aspects of Nature. Prior to the twentieth century, science focused on constructing and evaluating theories of a physical world viewed as external to the theorizing agent. The relationship of the domain under investigation to the mind that formulates and tests scientific theories was regarded as outside the province of science. This separation has been rendered untenable by the advent of quantum theory, the most stunningly successful scientific theory humanity has yet achieved. In a radical departure from classical physics, an observer in quantum theory enters into the dynamics of a physical system in a fundamental way. According to the orthodox interpretation of the theory, the observable behavior of a quantum system depends on whether an agent decides to observe it, and if so, on when it is observed and what features the agent chooses to observe. Quantum theory also departs from the strict determinism of classical physics. Although statistics is important in classical thermodynamics, it is regarded as an approximation due to imperfect knowledge of the details of systems whose underlying dynamics are deterministic. According to quantum theory, uncertainty is a fundamental: there are aspects of the physical world that are intrinsically unpredictable. Moreover, the very act of acquiring knowledge about a system alters the system under observation in ways that can be predicted only imperfectly. Despite considerable effort, no attempt to remove uncertainty and observer dependence from quantum theory has been fully successful. This state of affairs suggests strongly that a unified science of cognition and information processing must include theories not just of how an agent represents and reasons about the world it inhabits, but also of the physical processes by which the agent acquires information to create and revise its representations, and of the relationship between the physical and cognitive aspects of intelligent systems. Furthermore, it is not inconceivable that a better understanding of the processes by which cognitive agents acquire and process information will lead to a better understanding of some of the more puzzling aspects of quantum theory.

The ability to form, manipulate and evolve representations is vital to intelligence. An intelligent agent maps percepts from the environment into action policies that are likely to achieve its goals. The transformation from percepts to actions is mediated by an internal representation of task and environment that enables the agent to evaluate different candidate policies and select one for execution. Learning is a core capability that allows an agent to improve its performance over time, to modify its policies to apply to new kinds of situations, and to adapt to changing environments. Learning, representation and action are intimately connected, in that the purpose of learning is to produce a representation that is effective for the target class of tasks. The development of computationally feasible architectures for knowledge representation, learning and action selection depends in an essential way on the physical constraints governing the interface between the physical world and agents that acquire information about and act upon the world.

Recent years have seen great strides in theory and methods for representing knowledge in computational form, as well as for using observations to learn improved representations. The vast majority of research in knowledge representation and machine learning takes as given the view of computation as algorithms running on digital computers. There are several major problems with this assumption. First, a fundamental aspect of intelligence as we know it is the ability to learn about the world by deliberately manipulating the environment and observing the consequences, as well as by observing and communicating with other agents. All operations executed by an algorithm are programmed by its designer. There is no provision in classical computing theory for how a system can make *its own choices* that involve deliberate manipulation of its environment. Second, the correspondence of symbol structures in a classical computer to what the symbols represent in the world is external to the theory, added in by the designer of a computer program. Classical computing theory provides no account of how representational content, or "aboutness," arises *within the agent itself*. Third, classical computing theory is inherently deterministic, with no principled means for accruing evidence and drawing plausible conclusions in an uncertain world. Finally, classical computing theory is sequential and single-agent, while an intelligent agent must perceive and act in parallel with other agents inhabiting its world.

Recently, there has been a surge of interest in randomized algorithms. Randomized algorithms can find highly accurate solutions to some problems for which no efficient deterministic algorithm is known (e.g., Solovay and Strassen, 1976). Randomization steps can permit escape from local optima (e.g., Kirkpatrick, et al., 1983) and achieve good statistical coverage of large, high-dimensional search spaces (e.g., Gilks, et al., 1996). The right amount of randomness can ensure that a system escapes local optima in finite time, yet still permit a reasonable degree of predictability and control. A useful feature of many randomized algorithms is their characterization as stochastic processes evolving toward a target stationary distribution (e.g., Gilks, et al., 1996). One selects a target stationary distribution that places high probability on good solutions to the problem at hand, and attempts to identify an evolution dynamic that yields trajectories converging as rapidly as possible to the target. The theory of stochastic processes provides methods for analyzing convergence rates and bounds on performance metrics.

Many of the most successful randomized algorithms have been imported from or inspired by statistical physics. Their theoretical basis is the principle of action minimization in physical systems. All physical systems evolve in time according to a trajectory that minimizes a quantity called action, having the units of energy×time. To exploit the action minimization analogy, one formulates an inference or optimization problem as a fictitious physical system in which better solutions map to lower action in the physical system. Computation then simulates an action-minimizing trajectory perturbed by random fluctuations corresponding to thermal noise. The rapid proliferation of general-purpose inference and optimization methods in this class suggests the possibility of formulating a theory of intelligent information processing and decision-making in terms of stochastic processes modeled after physical systems. In this view, an information processing system is conceived not as an automaton deterministically executing steps someone has programmed into it, but as a stochastic process that transforms inputs and internal state into a sequence of trial solutions, and evolves over time to improved solutions. Software engineering in this view is the design of dynamic systems that evolve over time to solutions of better quality, where quality is inversely related to action in the corresponding physical system. A new theory of computing would apply the theory of stochastic processes and Bayesian statistical inference to analyze asymptotic solution quality, the rate of improvement in solution quality over time, and the process by which a system learns better representations and improved policies. Such a theory could be used to develop statistical tests for evaluating the cost of further computation against the benefit of an improved solution. A new axiomatic semantics of computational processes would be based on decision theory, dynamic programming, and partially observable Markov decision processes (e.g., Cassandra, et al, 1994).

Randomized algorithms provide a solution to the problem of determinism and escape from local optima. In itself, however, randomization provides no way to address the problem of "aboutness," nor does it provide a satisfactory account of how a physical symbol system can make choices in pursuit of goals. We hypothesize that the advent of quantum computing may open up new ways of addressing these issues.

Although it is common to equate computation with algorithms running on classical computers, interest in quantum computation (Aharanov, 1999; Deutsch, 1985; Nielson and Chuang, 2000) is increasing rapidly. As the relentless pursuit of miniaturization is pushing computer hardware into the realm in which quantum effects cannot be ignored, the classical computing model may soon encounter physical limitations on its applicability. At the same time, indications are appearing that quantum computers may be intrinsically more powerful than Turing machines with randomization operations. It has been demonstrated that quantum computers can solve some important classes of problems with far less computational resources than the best known classical algorithms (e.g., Grover, 1995; Shor, 1994). It has been hypothesized that the advantage of quantum computation lies in the ability to perform Fourier transforms with much greater efficiency than any known classical algorithms (Aharanov, 1999). Although research in quantum computing hardware is in its infancy, there has been rapid progress in the ability to construct and control coherent multi-qbit systems.[1] Progress in quantum computing hardware and software may also spark new insights regarding some of the thorny philosophical issues associated with quantum theory.

This paper argues that a marriage of quantum theory with Bayesian decision theory provides a unifying account of the physical and informational aspects of physical symbol systems. Section 2 provides an introduction to computational Bayesian decision theory and graphical decision theoretic models. Graphical models have become increasingly popular as a language for expressing logically sound and tractable computational theories. Section 3 describes an interpretation for quantum theory that maps directly to Bayesian decision theory and presents a graphical model representation of quantum evolution and quantum measurement. With no changes to the mathematical structure of existing physics, the ontology proposed here connects

---

[1] In quantum mechanics and quantum computing, the term *coherent* refers to a system in which the elements have a definite phase relationship to each other. In decision theory, the term refers to an agent who follows the principle of maximizing subjective expected utility (or minimizing subjective expected loss). In this paper, the term is used in the former sense. Coherence in the quantum mechanical sense is essential to the ability of quantum computers to improve on the efficiency of classical computers.

physical reality in a plausible way to efficacious choices by agents. Section 4 concludes with a summary and discussion.

2 Graphical Probability and Decision Models

Intelligence requires the ability to reason and act in the presence of uncertainty. One of the most difficult challenges in artificial intelligence has been the development of tractable and logically sound methods for plausible reasoning and decision making in the presence of uncertainty and incomplete information. Although once controversial, Bayesian decision theory is now regarded as a foundational theory for computational inference and decision making under uncertainty, and a standard against which proposed alternative approaches should be compared (c.f., Russell and Norvig, 2002). Game theory, or multi-agent decision theory, is becoming standard as a foundation for systems of multiple interacting agents (e.g., Kearns, 2002).

2.1 Bayesian Decision Theory

Bayesian decision theory is a mathematical theory of rational decision making under uncertainty. The theory provides a sound way to combine beliefs with values to arrive at logically consistent, value-driven decisions. It applies to situations in which an agent must choose an action or series of actions from among a set of alternatives. The consequence of the choice depends on both the selected action(s) and the state of the world. The state of the world, which may be unknown to the agent at the time the choice is made, belongs to a set of mutually exclusive and collectively exhaustive possible states. The agent expresses uncertainty about the state of the world by attaching a probability to each possible state.[2] The agent expresses preferences among consequences by assigning a numerical utility to each consequence (alternatively, the agent may assign losses, which are negative utilities). Taken all together, the possible actions, possible states, possible consequences, the probability function and utility (or loss) function comprise a decision theoretic model for the agent's decision problem. According to the model, the agent's optimal choice is to select the action for which the mathematical expectation of the utility is maximized (or for which the loss is minimized). When the agent acquires information about the world, probability assignments are updated according to a mathematical formula called *Bayes rule*. The revised probabilities are used for subsequent predictions and decisions.

Many arguments have been put forward both for and against the principle of maximum expected utility as a model of rational decision making under uncertainty (e.g., Howsen and Urbach, 1993). A number of authors (e.g., Savage, 1972; Pratt, et al, 1965) have developed axiomatic systems that capture intuitive notions of rational behavior, and demonstrated that the axioms imply the optimality of subjective utility maximization. Such optimality proofs are invoked to argue that that rational decision making requires utility maximization. A common objection to axiomatic arguments is that computing optimal policies for realistically complex decision problems is intractable.[3] Great strides have been made in recent years (e.g., Pearl, 1988; Jensen, 2001; Neapolitan, 2003; Korb and Nicholson, 2003) in tractable exact or approximate algorithms for computational probabilistic inference and decision making. Decision theoretic methods have found their way into numerous successful applications (e.g., Levitt, et al., 1995; Parker and Miller, 1987). Among pragmatists, successful applications provide a much stronger argument in favor of Bayesian decision theory than axiomatic arguments. It has been argued that

[2] If the set possible states is uncountably infinite, then probabilities must be specified via a probability density function.

[3] Other counter-arguments question some of the standard axioms, such as those implying a simple ordering of all options. Relaxing the objectionable axioms has led to alternative decision theories, such as theories of interval probabilities and utilities. These theories are not discussed here, except to note that they typically pose even more challenging computational issues than standard Bayesian decision theory.

intelligence requires the functional equivalent of approximate Bayesian inference and decision theory (e.g., Lee and Mumford, 2003). Many heuristic methods proposed as alternatives to decision theory can be shown to result in approximate decision theoretically optimal behavior within their domain of applicability (e.g., Martignon and Laskey, 1999). When computational limits are taken into account, decision theory itself would recommend an approximately optimal heuristic strategy over an optimal solution that cannot be computed rapidly enough to be applied in a real situation (c.f., Gigerenzer, et al, 1999).

There has been a great deal of controversy over the interpretation of the probabilities that appear in a decision theoretic model (e.g., Howson and Urbach, 1993; Fine, 1973). The dominant view in artificial intelligence is subjectivism, in which probability is viewed as a measure of the degree of belief of a rational agent about uncertain hypotheses. Subjectivists assign probabilities to any hypotheses about which they are uncertain. To a subjectivist, reasonable individuals may assign different probabilities to the same outcome. The only requirements are that beliefs must conform to the mathematical properties of the probability calculus, and may not contradict evidence known to the agent. Although the subjectivist view is gaining favor in statistics, until recently the frequentist view has dominated. To a frequentist, probabilities are limiting frequencies in long sequences of outcomes generated by intrinsically stochastic systems. Unlike subjectivists, frequentists regard it as illegitimate to assign probabilities to anything except chance set-ups that give rise to sequences of random events. It is meaningless to assign probabilities to individual events in a sequence or to hypotheses with a definite but unknown truth-value (except trivially to assign probability one if an event happens and zero if it does not). Frequentists view probability as an objective property of a chance set-up. If two individuals assign different probability distributions to a chance set-up, at least one of them must be wrong. Although frequentists criticize subjectivists for lack of objectivity, the ability to represent and reason with subjective information is a necessary aspect of intelligent behavior. Furthermore, some of the most successful applications of probability theory (including thermodynamics and classical statistical mechanics) are to problems involving incomplete knowledge about a system whose underlying dynamic is assumed to be deterministic. Another strength of the subjectivist approach is its ability to handle small data sets and large numbers of parameters, a situation that occurs frequently in the types of problems encountered in artificial intelligence and machine learning. The theory of precise measurement (deGroot, 1970; von Winterfeldt and Edwards, 1987) identifies conditions under which subjectivist agents beginning with different prior probabilities will converge to nearly identical posterior probabilities. The conditions under which these results hold are also characteristic of situations for which objective propensities might be hypothesized. Thus, the frequentist and subjectivist views can be reconciled on problems to which both are willing to apply probability, but subjectivists extend the domain of applicability of probability far beyond what frequentists consider legitimate. Much of machine learning and data mining falls outside the zone of applicability of frequentist statistics, and applications of probability theory require a subjectivist interpretation.

Subjectivist Bayesian decision theory demands conformance to rationality principles that seem empirically questionable and may be computationally unachievable. Recently developed game-theoretic interpretations (Dawid and Vovk, 1999; Nau and McCardle, 1991; Shafer and Vovk, 2001) regard probability as arising out of the behavior of interacting agents who receive rewards for correctly forecasting events. Agents participate in an economic system in which they can announce forecasts, make bets, and/or buy and sell contingent options whose values depend on the outcomes of uncertain events. If the market is sufficiently liquid and the rules of interaction permit opportunities for arbitrage[4] to be exploited, then consistent probability

[4] Arbitrage means executing a sequence of trades leading to a riskless profit. Efficient markets evolve prices that eliminate opportunities for arbitrage.

forecasts can be expected to emerge from the prices at which contingent options are traded. De Finetti (1974) showed that any agent who violates the axioms of decision theory would agree to a sequence of transactions resulting in a sure loss. Agents violating the rationality axioms present arbitrage opportunities that can be exploited by other agents, and if not corrected, will lead to bankruptcy. The agents remaining in the market will tend to behave as utility maximizers. Prices for contingent options in such a market measure a market consensus probability for the contingencies on which the options depend. There is evidence that markets for contingent options can provide more accurate probability estimates than standard methods of eliciting probabilities from experts (Berg, et al., 2001). Unlike standard axiomatic decision theory, market-based evolutionary theories do not impose rationality axioms as constraints. Rather, selective pressure for rationality is only one of the "forces" operating on a bounded rational agent[5] engaging in trades. Thus, game-theoretic probability is in many ways a more satisfying foundation for probabilistic knowledge representation than either the subjectivist or frequentist interpretations.

2.2 Graphical Probability Models

Explicitly representing and reasoning with all possible exceptions and contingencies results in a combinatorial explosion of possibilities. Graphical probability models have become popular because they can tractably represent and perform inference on reasonably faithful models of the uncertainties involved in realistically complex tasks. In a graphical probability model, a directed or undirected graph is used to represent qualitative information about probabilistic dependencies. Nodes in the graph represent *random variables*, or sets of mutually exclusive and collectively exhaustive hypotheses. Edges in the graph represent direct dependencies of the probabilities of the possible values of a random variable on the value of its neighboring random variables. Quantitative information about the strength of dependency is represented by local probability distributions associated with the nodes in the graph. Whereas the resources required to store and/or compute with a general probability distribution are exponential in the number of random variables, knowledge representation in a graphical probability model with a bounded number of neighbors per node scales linearly in the number of random variables. When the graph is singly connected (i.e., there is only one path between any two random variables), inference also scales linearly with the number of random variables. Although there are special cases in which a singly connected graph is adequate, realistic tasks often require more complex connectivity. Exact inference algorithms have been developed for multiply connected graphs. Although their worst-case complexity is exponential, there are interesting classes of problems for which exact methods are tractable. In the general case, approximate inference is required. A number of general-purpose methods have been developed for approximating Bayesian inference in multiply connected dependency graphs.

Figure 1 shows a directed graphical model, called a *Bayesian network*, for a notional problem of reasoning under uncertainty. This model is highly oversimplified, but suffices to illustrate the essential features of graphical probability models. A model such as this might be used by an agent (the *reasoning agent*) to reason about the behavior of a second agent (the *behaving agent*). This model is a static snapshot. It will be extended later to a model that represents sequences of

[5] Fienberg and deGroot (1982) showed that proper scoring rules (rules that reward correct probability assessments) can be decomposed into components measuring *coherence* (conformance to the laws of probability), *calibration* (fit to empirical frequencies), and *refinement* (the ability to make fine distinctions). While it is true that an incoherent agent can always improve its score by finding and eliminating inconsistencies, it may be the case that the potential for improvement by improving calibration or refinement is far greater. When resource costs are taken into account and the goal is global task performance rather than accurate probability forecasts, it may be optimal to sacrifice strict coherence for better performance. See also the discussions of rationality in Russell and Norvig (2002).

actions carried out over time. According to the model, the probability that the behaving agent will engage in an activity depends on its goal, which in turn depends probabilistically on what type of agent it is. Different activities give rise (probabilistically) to characteristic spatial configurations of the physical elements of the behaving agent (e.g., an outstretched arm would provide evidence that the agent was picking up an object), and are likely to occur at characteristic locations (e.g., eating would be likely to occur in a kitchen or dining room; driving would tend to occur on a road). Evidence may be acquired about location, configuration, activity and type. The agent's goal is unobservable, but inferences can be drawn about the goal from evidence about variables related to the goal. Evidence about location and configuration are derived from sensors directly observing the behaving agent. Raw sensor inputs must be transformed into values for the evidence random variables (e.g., intensities of pixels on an image are transformed into a measure of how closely the pictured agent's arm resembles the typical configuration for an agent raising its arm). Although not pictured in the figure, the relationship between the evidence random variables and the raw pixel data is also probabilistic, and can itself be represented in the language of graphical models (e.g., Levitt and Binford, 2003; Grenander, 1993). Thus, graphical models provide a consistent, theoretically justified theory of knowledge representation and evidential reasoning that spans the subsymbolic through the cognitive levels.

The example of Figure 1 can be extended to a *decision graph* (also called an influence diagram) as shown in Figure 2. Two new types of node have been introduced: *decision* nodes, represented as rectangles, and *utility* nodes, represented as hexagons. Decision nodes represent choices open to the reasoning agent. The model of Figure 2 represents a decision of whether the reasoning agent should intervene to prevent an undesirable action by the behaving agent (e.g., in a building security application, the reasoning agent might intervene to prevent theft or terrorist acts). The utility nodes measure how well the agent's objectives are satisfied. In this example, the agent's utility function depends on the cost of intervening and the damage if the undesirable action succeeds. The arcs entering a decision node represent information available to the reasoning agent at the time the decision is made. In this example, the reasoning agent knows the reported location and configuration, but not the reported type or activity. The reasoning agent can choose as its policy any function of the information available at decision time. The optimal policy according to the model maximizes the sum of the mathematical expectations of the utility nodes (or minimizes the sum of expected losses) given the available information. In this example, a rational agent would choose a function of the observed location and observed configuration that minimizes the total expected costs from intervention and/or damage.

A natural question to ask about Figure 2 is why, if a decision graph represents a mathematical model for value-driven decision making, there is no arc from the utility nodes into the decision node. A naïve reading of the graph might give the impression that the agent's choice of whether to intervene is not affected by the intervention and damage costs. To understand why this is not the case, it is necessary to examine the semantics of the diagram more closely. The diagram and corresponding numerical information specify a mathematical model of task-relevant aspects of the world viewed from the reasoning agent's perspective. The arcs represent three kinds of influences. Arcs into world state nodes represent deterministic or stochastic relationships. For cause and effect relationships, the convention is to draw the arc from cause to effect; for correlations, the arc can go in either direction. Arcs from world state nodes into action nodes reflect information available to the agent at the time of choice. Arcs into utility nodes represent mappings from situations in the world to degrees of satisfaction the agent experiences on the corresponding dimensions of value. Arcs may not exit utility nodes. The agent's experience of satisfaction occurs *as a result* of the choice, and *at a later time* than the agent's action. Thus, the actual experienced satisfaction can neither cause the agent's action, nor be available as information on which the agent can base the decision.

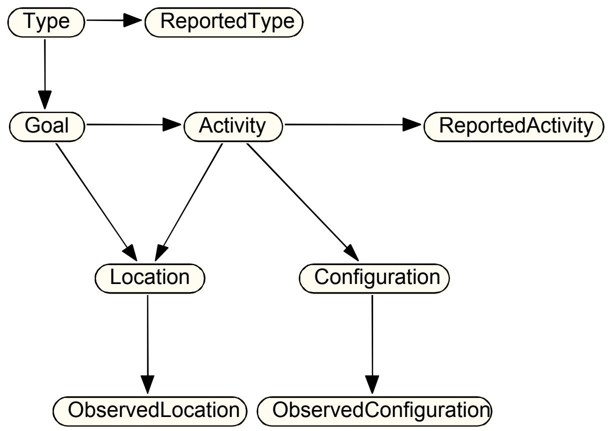

Figure 1: Bayesian Network

The effect of values on decisions is brought about by *solving* the decision graph. Conceptually, solving the diagram proceeds in two steps. The first step is to compute the mathematical expectation of the total cost (damage plus intervention) for each possible intervention action, given each of the possible information states of the reasoning agent at the moment of choice. The expectation is a function mapping values of *ObservedLocation*, *ObservedConfiguration*, and *Intervene?* to real numbers representing the total expected cost if the reasoning agent observes the given location and configuration and takes the indicated action. The second step in solving the graph is to select an optimal decision policy that maps information states to actions. The optimal policy maps an information state (value of *ObservedLocation* and *ObservedConfiguration* to the action for which the expected cost is lowest. After the diagram has been solved, the optimal policy is stored with the *Intervene?* node. The agent then receives observations on location and configuration, looks up the optimal action corresponding to these observations, and executes that action.

The graph of Figure 1 is simple and easily solved. Decision graphs for realistically complex problems can easily become intractable. There is an extensive and rapidly growing literature on exploiting independence relationships among random variables and decomposability of value functions to perform efficient computation of optimal or approximately optimal decision policies (e.g, Jensen, 2001; Neapolitan, 2003; Boutilier, et al., 1999).

The model of Figure 2 was constructed from the point of view of the reasoning agent. To the reasoning agent, the behaving agent's goal and activity are represented as uncertain world state

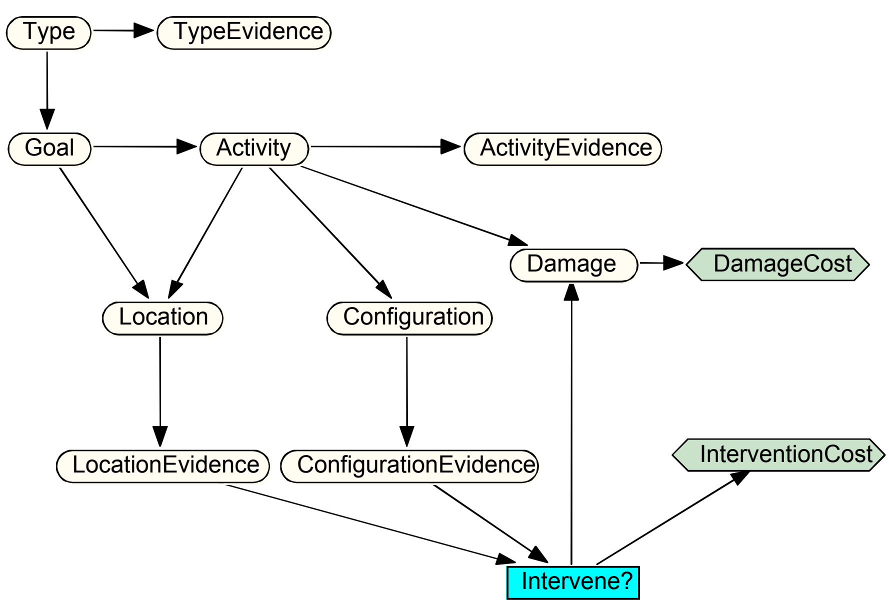

Figure 2: Decision Graph

random variables, whereas the reasoning agent's intervention decision and costs are represented as decision and utility nodes, respectively. If the model were being constructed from the point of view of the behaving agent, the behaving agent's goal would be represented as a utility node, the behaving agent's activity would be represented as a decision node, and the reasoning agent's intervention act (if any) and the damage and intervention costs would be represented as uncertain world state nodes. From the point of view of an external agent observing both the reasoning and behaving agents, all nodes in Figure 2 would be uncertain world state nodes. Formulating and solving a model of another agent's decision problem can be a useful aid to predicting the other agent's behavior, because the agent can be expected to select actions that serve his or her objectives given the information he or she has available, taking into account cognitive limitations.

Standard Bayesian networks and decision graphs are limited to problems with a fixed number of uncertain hypotheses in which all the relevant variables and relationships can be specified in advance of problem solving. This restriction is inadequate for complex real-world problems involving an unspecified number of objects of different types interacting in varied ways. A rapidly growing research area is the development of extensions to the language of graphical models that allow a graph to contain repeated sub-structures representing instances of classes of similar hypotheses. For example, all photographs of the same subject from the same angle under the same lighting conditions using the same type of camera produce the same probability distribution for image pixels. If a decision graph were being constructed for a problem involving more than one behaving agent, being observed under similar viewing conditions, it would include multiple copies of the Configuration → ObservedConfiguration portion of Figure 2, one for each observation event.

Figure 3 shows an extension of the decision model of Figure 2 to a *multi-entity decision graph* (Laskey, 2003; Laskey, et al., 2001). A multi-entity decision graph (MEDG, pronounced "medge") consists of a collection of partially specified graphical models, called MEDG fragments, that collectively specify a decision model involving a variable, possibly uncertain, and possibly unbounded number of interacting entities. The nodes in a multi-entity decision graph have arguments referring to entities, thus allowing multiple copies to refer to different entities. For example, different instances *Activity*(*G*1,*T*1) and *Activity*(*G*2,*T*2) of the random variable *Activity*(*g*,*t*) represent the activities of two different behaving agents *G*1 and *G*2 at two different times, *T*1 and *T*2. Figure 2 illustrates how multi-entity decision graphs can be used to represent composite systems evolving in time. The same basic structure of the example of Figure 1 is retained, but in this case the model is built to reason about groups of vehicles engaging in a common activity, such as driving down the road together or converging toward a rendezvous point. MEDG fragments may contain *resident* random variables, whose distributions are defined in the fragment, *input* and *context* random variables that condition the distributions of resident random variables, and *decision* and *utility* nodes representing choices open to a decision maker and the objectives used to guide the choice.

The MEDG represented in Figure 3 contains a random variable referring to the composition of a group, as well as random variables that relate the type and location of a vehicle to the type and location of the group of which it is a member. These random variables can be instantiated for any number of vehicles and groups. The model also specifies probabilistic rules by which activities and locations evolve in time, including how these random variables depend on random variables at previous times. Note that the parents of the node *Outcome*(*t*) refer to an argument *g* that is not an argument to *Outcome*(*t*). The parents of *Outcome*(*t*) consist of all copies of *Intervene*(*g*,*t*) and *Activity*(*g*,*t*) that match on the argument *t*. The node *Outcome*(*t*) thus has a variable and possibly uncertain number of parents. An *influence combination* rule (Laskey and Mahoney, 1997) specifies a function mapping the set of parents of *Outcome*(*t*) to a probability distribution. For example, in a counterterrorist application, the activities might include preparation and detonation of a bomb, and the outcome might indicate whether a bomb detonated by any of the groups under observation explodes at a given time. The influence combination function would map the activities of all groups under consideration and all intervention acts undertaken to a probability distribution on whether a bomb detonates.

Multi-entity decision graphs represent one of several emerging formal frameworks for extending graphical models to permit expression of repeated sub-structures. Other languages include pattern theory (Grenander, 1995), hidden Markov models (Elliott, et al., 1995), the plates language implemented in BUGS (Gilks, et al., 1994; Buntine, 1994; Spiegelhalter, et all, 1996), object-oriented Bayesian networks (Koller and Pfeffer, 1997; Bangsø and Wuillemin, 2001; Langseth and Nielsen, 2003), and probabilistic relational models (Getoor, et al., 2000, 2001; Pfeffer, 2001). Decision graphs can also be extended to multi-agent problems, in which each agent has its own *utility* and decision nodes, and each agent's optimal policy is to maximize the expectation of its *utility* nodes conditional on its available information (e.g., Kearns, 2002). Attractive features of graphical models as a language for representing knowledge are their principled treatment of uncertainty, their provision for specifying knowledge as modular components with well-defined interfaces, and the existence of general-purpose exact and approximate inference and learning algorithms.

3 Quantum Theory and Bayesian Decision Theory

Classical mechanics is a dynamically complete theory with no role for conscious thought and efficacious deliberate action. Once initial conditions are specified, a classical physical system follows a definite trajectory that, at least in principle, can be predicted with absolute precision indefinitely into the future. Of course, in practice this predictability is limited by approximation and measurement error in the specification of both the initial conditions and the parameters of the dynamical equations. Nevertheless, in principle, the evolution of a classical system is perfectly determined by initial conditions.

Early in the 20th century it was discovered that the classical picture of a world of perfectly deterministic physical systems evolving according to purely local influences was incorrect. The classical picture was replaced by the explicitly probabilistic quantum theory. The degree of accord between the theoretical predictions of quantum theory and empirical measurements performed on quantum systems is stunning. Nevertheless, many physicists remain uncomfortable with the theory. There are three major reasons for this discomfort. First, quantum theory makes only probabilistic predictions about the trajectory of a system. Many scientists are uncomfortable with a picture of Nature that has an intrinsically stochastic component. Second, the theory is nonlocal. That is, there are correlations between spacelike separated events that cannot be explained by a hidden variable theory with strictly local influences. Third, the theory contains a major explanatory gap known as the "measurement problem," in which deterministic evolution of the wave function is interrupted by "reduction events" for which current physics has no theory.

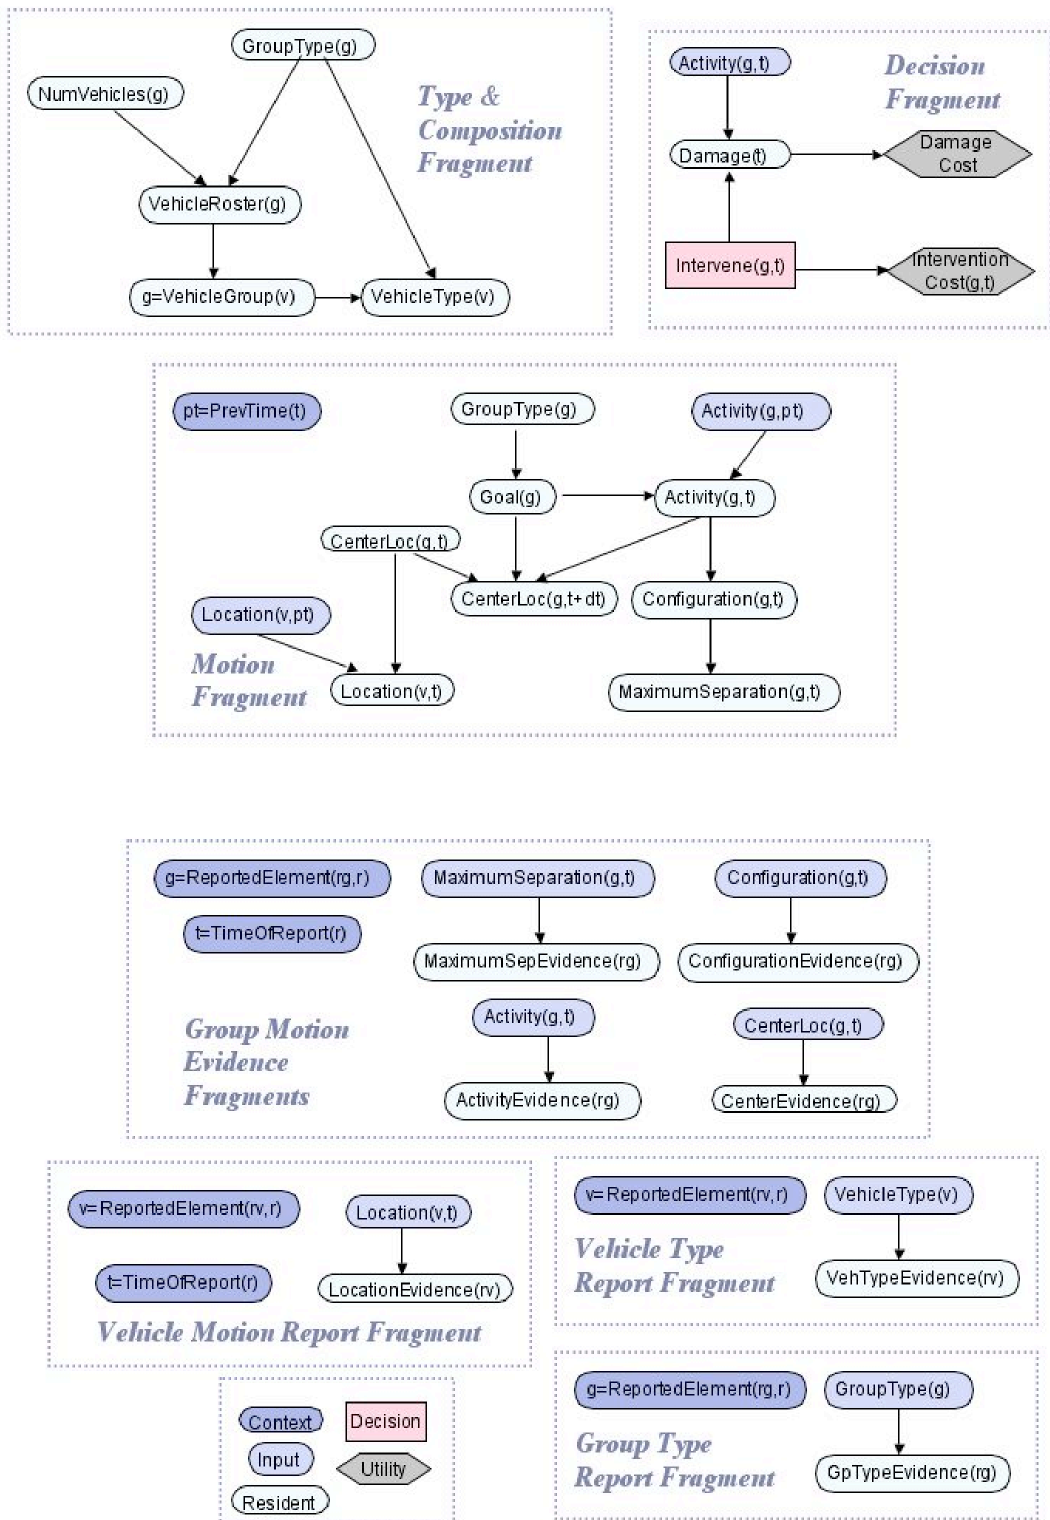

Figure 3: Multi-Entity Decision Graph

3.1 The Measurement Problem

The state of a quantum system at any time is described by a mathematical structure called the *quantum state*. According to quantum theory, the quantum state evolves in time according to three distinct processes (Stapp, 1999; Penrose, 1994; Shankar, 1994). The first process, Schrödinger evolution, concerns how the state evolves in the absence of interactions with its environment. The second and third processes concern an abrupt change called *state reduction*. The second process specifies a time at which reduction occurs and the set of possibilities for the quantum state after reduction. The third process selects one specific possibility to be actualized. Quantum theory describes the first and third processes to a high degree of accuracy. Between reduction events, Schrödinger evolution specifies that the system state evolves deterministically according to a differential equation called the Schrödinger equation. Although the state after reduction cannot be predicted with certainty, once the time of reduction and the set of possibilities have been specified, quantum theory provides a precise rule for calculating the probabilities of the different possibilities. As currently formulated, however, quantum theory has nothing at all to say about the manner in which the time of the next reduction and the set of possible outcomes is specified. Bohm (1951) states that although the quantum state has been called a "wave of probability," it is more accurately described as a "wave from which many related probabilities can be calculated." Different reduction operators applied at different times give rise to different probability distributions for outcomes, but they all can be calculated from the evolving quantum state.

State reduction is thought to be associated with the interaction of a quantum system with its environment. The orthodox interpretation of quantum theory associates reductions with measurements performed on a system by scientists. Measurement involves an interaction of a measuring apparatus with the system, in which some feature of the microscopic quantum system is "amplified" to produce a macroscopically detectable change in the measurement device. The behavior of a system after measurement depends on the measurement outcome, and is different from what its behavior would have been if no measurement had occurred.

Although there exist experimental procedures, described in classical language, for effecting measurements on various types of physical systems, there are no fundamental physical laws governing how the agent, considered as a physical system, makes the choice of which measurement, if any, to perform, and the time at which this action occurs. Although there have been attempts to explain measurement as nothing but Schrödinger evolution on a larger system comprising both the observing agent and the observed system, none has been fully successful. Additional rules are required to reproduce the quantitative predictions of quantum theory. In the standard interpretation, these rules correspond to initiation of reduction events by agents and selection of outcomes by Nature. Our capacity to perform measurements is taken as a given empirical and phenomenological fact. The lack of fundamental physical theory for the timing and possible outcomes of state reduction is called the "measurement problem." The very means by which we are able to learn about the behavior of quantum systems, and thus construct a theory to predict a system's behavior between measurements and the probabilities of different measurement outcomes, is itself unexplained by quantum theory.

Thus, the dynamic behavior of a quantum system depends in macroscopically observable ways on a process for which physics has no theory. In particular, an observer can choose which of several distinct macroscopic effects to actualize by choosing which aspects of the system to observe. It is important to note that although quantum theory includes statistical laws governing Nature's choice of outcome when a measurement operator is applied, the known physical laws do not fix, *even statistically*, which measurement operators are applied under what conditions. Thus, the theory contains a contingent element. It specifies behavior of the system *given* the actions an agent external to the theory takes to observe the system. This dependence of the predictions of the theory on an aspect of reality for which there is no theory bothers some physicists. However, agents with brains and bodies built from the material particles studied by atomic physics have a demonstrated capacity to perform plausible reasoning. Moreover, there are aspects of the physical architecture of the brain that make it likely that quantum mechanical effects are important in its dynamical behavior (Schwartz, et al, 2004). Thus, a cognitive agents should be modeled not as a classical computer, but as a quantum Bayesian decision maker interacting with an environment that needs to be represented, reasoned about, and acted upon.

3.2 Ontologies for Quantum Theory

The orthodox interpretation for quantum theory is associated with Bohr (1934) and is called the Copenhagen interpretation. According to the Copenhagen interpretation, quantum theory replaces a classical theory that refers to an external material universe with a new theory that refers only to the experience of observers and not to the external universe itself. Conditional on a choice of experimental set-up that defines the macroscopically detectable possibilities available to the system, quantum theory predicts the probability that each of these classically describable possibilities will occur. Proponents of the Copenhagen interpretation make no ontological commitments regarding the entities that give rise to the experienced sequence of observations. It is sometimes asserted that it is meaningless to speak of the "actual state" of a quantum system. The quantum state is asserted to be nothing but a mathematical construct for organizing the experiences of observers and enabling the computation of accurate predictions of the outcomes of experiments. Thus, according to the orthodox interpretation, quantum theory represents a set of computational rules by which scientists can make predictions about which classically describable outcomes will occur as a result of the classically describable experiments they conduct. The quantum state is a mathematical construct used to make predictions about observables, but being inaccessible to direct observation, is not to be regarded as corresponding to any definite phenomenon in Nature.

Although the Copenhagen interpretation is the standard view, most physicists prefer, at least informally, to operate with an ontology that connects the terms in the theory to a physical reality that is reflected in the experience of observers. One popular ontology, known as *many worlds*, asserts that the system actually realizes *all* possibilities open to it, but each occurs in a separate reality inaccessible to the other realities. There is a copy of each of us in all the different realities, but only the copy in this particular reality has the experiences associated with this reality. The many worlds interpretation is common in the field of quantum computing. Another interpretation is the pilot wave ontology (Bohm and Hiley, 1993), a nonlocal deterministic theory that includes both classical-like particles and a wave function that guides their evolution. Yet another interpretation is due to Penrose (e.g., 1994), who hypothesizes that wave function reduction represents the singling out of an actual event to occur by a mechanism hypothesized to be related both to consciousness and gravitation.

Stapp (1999) and McFadden (2000) suggest an ontology closely related to the measurement theory first proposed by von Neumann (1932) and further elucidated by Wigner (1967). Their interpretation is similar to Penrose's in that it takes a realist view of state vector reduction, but does not specifically implicate gravity as a causal factor in reductions. The universe is hypothesized to include systems that can cause state reductions. These reducing agents can choose, within as yet to be determined physical limits, when the next reduction will occur and what are the possible outcomes of the reduction. A reducing agent's choice may depend, in an as yet to be determined way, on the quantum state of the reducing agent. Both Stapp and McFadden hypothesize that human beings are one type of reducing agent. Thus, complementary explanatory gaps in physics and psychology are filled by postulating an interaction between the informational structure represented by the quantum state and the informational structure of conscious experience. Stapp (1999) argues that such an interaction allows consciousness to become efficacious without disturbing any of the precepts or rules of quantum theory. McFadden suggests the induced electromagnetic field of the brain as the physical substrate corresponding to conscious experience. He notes the tendency of physicists to use terms connoting volition in describing wave function reduction and quantum measurement, and points out that textbooks on quantum theory assign to human agents the choice of what experiment to perform on a quantum system. Thus, the reducing agent ontology is consistent with the informal language used in most quantum theory texts and by most practicing physicists.

With no change to the mathematical machinery of quantum theory, the reducing agent ontology connects physical reality in a plausible way to conscious experience and deliberate choice. Because the evolution of a quantum system depends on the choice and timing of state reductions, ascribing reductions to the free choices of conscious players provides a means consistent with known physics for conscious agents to affect outcomes in the world. Moreover, deliberate conscious choice is inserted into physics at exactly the place where current physics lacks a theory. The commonly held view that quantum theory is nothing but a recipe for calculating probabilities and says nothing directly about physical reality is replaced by a realist theory with a role for deliberate choice. This ontology thus provides science with a theory of both the physical and informational aspects of nature, and describes how deliberate choices of conscious agents affect both the agents themselves and the world they inhabit. Moreover, the physical constraints necessary for deliberate choice to be operative in this way appear to be satisfied by the conditions occurring in live animal brains (Stapp, 1999).

Stapp (1998) and McFadden (2000) provide empirically verified examples of macroscopically detectable differences in behavior resulting from different policies for effecting state reductions in quantum systems. The quantum Zeno effect (Itano, et al., 1990; Gribbin, 1996) predicts that observations taken sufficiently rapidly can keep a quantum system within a constrained region of phase space. The inverse quantum Zeno effect (McFadden, 2000) induces a quantum system, via a sequence of rapidly repeated measurements, to follow a particular path in phase space. Stapp (1998) argues that an organism might use the quantum Zeno effect to keep its brain state within a given basin of attraction sufficiently long to trigger behaviors the organism desires to bring about. The quantum Zeno effect has been confirmed experimentally (Itano, et al., 1990) and is thought to occur at time and frequency scales consistent with patterns of electrochemical activity occurring in brains.

Stapp suggests that the quantum states for different outcomes of a reduction event caused by a player are associated with different psychological states, or *qualia* of experience. For human observers, qualia are multifaceted, highly complex gestalts that defy simple description. Nevertheless, Stapp argues, human observers may be able to select reduction policies that bring about brain states associated with qualia they prefer. Both Stapp and McFadden suggest that the choice of reduction policy occurs via a physical mechanism corresponding to what psychologists call *will* or *attention* (Anderson, 1999; James, 1890). That is, an organism complex enough to be labeled conscious can anticipate with some degree of accuracy the qualia associated with the different available reduction-causing policies, identify those that are most desirable, and focus attention on bringing about one of the most preferred policies. To do this, Stapp hypothesizes that the brain encodes a "body-world schema" that represents the body, the environment, and the predicted effects of alternative reduction-causing policies. The organism uses its body-world schema to direct its focus of attention to bring about desired qualia. McFadden (2000) cites Libet's theory of a conscious mental field that generating a unified subjective experience (Libet, et al., 1979). Invoking Occam's razor, he suggests that Libet's conscious mental field be identified with the induced electromagnetic field of the brain, and notes that electromagnetic fields are commonly used in computing and communication devices to encode and transmit information.

Although both Stapp and McFadden hypothesize that humans are reducing agents, there is no implication that humans are the only reducing agents or that reducing agents must be conscious. Reductions that occurred prior to the evolution of conscious organisms would have been caused by unconscious or proto-conscious reducing agents. Although it is conceivable that some form of the property we call consciousness at the human level exists throughout the natural world, the reducing agent ontology does not require it. McFadden raises the question of why, if consciousness is equated with electromagnetic fields, the electromagnetic fields generated by electronic devices are not also conscious. “The somewhat surprising answer,” he says, “is that we have no way of knowing whether or not any of these fields are indeed conscious. The only conscious minds that can state that they are conscious are those that can communicate their consciousness.” Although it would be unreasonable to suppose that computing devices constructed with current technology would possess a sophisticated form of consciousness, McFadden offers the opinion that a computer constructed with an appropriately sophisticated electromagnetic feedback loop would possess a primitive form of consciousness. Such a device would also be capable of making choices that were not programmed by its designer.

Whatever one's stance on the level of complexity at which consciousness can be said to exist, the question arises of why and how evolution would select for an increasingly sophisticated ability to direct actions via conscious intent. Stapp argues that because choice of reduction policy affects the evolution of a quantum system, the ability to predict outcomes of reduction events and use attention to effect those predicted to bring about desired outcomes would be expected to have survival value. In addition, agents that can communicate with each other can develop coordinated reduction policies that enhance collective survival. That is, it seems reasonable that evolution would tend to select for organisms that could form accurate representations of the choices available to them and select those options likely to lead to survival.

The viewpoint is gaining favor (e.g., Pearl, 1988; Russell and Norvig, 2002; Lee and Mumford, 2003; Levitt and Binford, 2003) that intelligence requires the ability to perform the functional equivalent of approximate Bayesian reasoning. Graphical models are attractive as a logically consistent language for formulating theories of computational intelligence and developing computer implementations capable of approximately optimal inference and decision making. Improved understanding of the relationship between the sensory and cognitive levels of description has fostered advances in computer vision, robotics, and multi-source fusion. It seems reasonable to suppose that improving our understanding of the interface between the quantum and the classical realms will generate additional advances in the field of intelligent systems. The next subsection argues that the language of graphical models can be extended into the quantum realm. Bayesian decision theory and graphical models provide a unified theoretical foundation, spanning the quantum through the cognitive level, for the design of hardware and software architectures for intelligent agents.

3.3 A MEDG Representation of Quantum Evolution

Figure 4 shows a representation of the reducing agent ontology as a multi-entity decision graph. The model shown here is consistent with the mathematics of quantum measurement as typically presented in texts on quantum theory (e.g., Shankhar, 1994), and with the measurement ontology proposed by Stapp (1999).

The graph of Figure 4 represents an interaction between two systems, the reducing agent *RA* and the observed system *OS*. In quantum theory texts, the observed system is described in the language of quantum theory and the reducing agent is described in classical language. Textbook measurement situations assume that *OS* has been prepared in a manner that shields it from all environmental influences except those carefully controlled by the experimenter. An experiment begins with *OS* in a definite known quantum state. *OS* is then subjected to known forces that guide its evolution for a specified length of time, after which a measurement event is initiated by the experimenter. In Figure 4, *T*1 denotes the time at which the initial state is prepared, *QuantumState*(*OS*,*T*1+) denotes the initial state just after preparation, *H*(*OS*) denotes the "Hamiltonian" operator that summarizes the effects of the forces acting on *OS* between measurements, *T*2 denotes the time at which the measurement occurs, and *QuantumState*(*OS*,*T*2-) denotes the state of the system just prior to the measurement interaction. In introductory treatments, the Hamiltonian is usually assumed constant during the time between measurement events, although this assumption is not necessary for the validity of the theory. The pre-measurement state *QuantumState*(*OS*,*T*2-) is a deterministic function of the initial state *QuantumState*(*OS*,*T*1+), the Hamiltonian *H*(*OS*), and the time interval *T*2-*T*1. Its value can be calculated from these quantities using a differential equation called the Schrödinger equation. At time *T*2, a measurement operator, denoted in Figure 4 by *Operator*(*T*2), is applied to the system, and an outcome occurs. In *non-degenerate* measurement situations, this outcome, denoted in Figure 4 by *Outcome*(*T*2), completely determines the state *QuantumState*(*OS*,*T*2+) of *OS* just after the measurement event. In a more general situation, the quantum state of *OS* after the measurement event is a deterministic function of the outcome of the measurement and the pre-measurement quantum state of *OS*. Although the outcome of a quantum measurement typically cannot be predicted with certainty, quantum theory provides a rule for calculating the

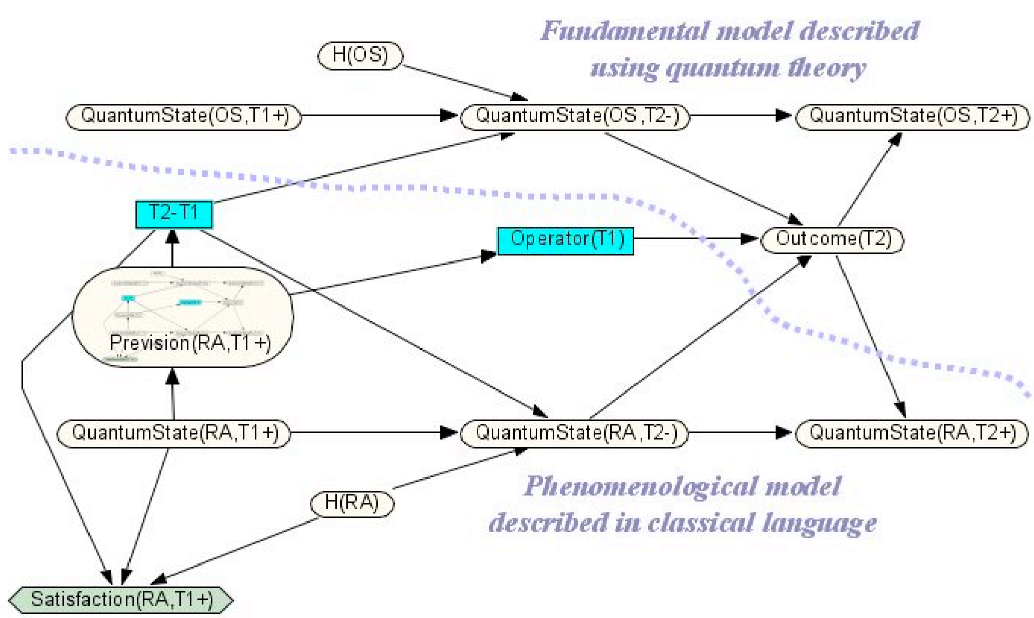

Figure 4: Decision Graph for Quantum Evolution

probabilities of the possible measurement outcomes given an experimental set-up and the state of *OS* just prior to the interaction. The outcome *Outcome*(*T*2) is one of the *eigenvalues* of the measurement operator *Operator*(*T*2). The post-measurement state *QuantumState*(*OS*,*T*2+) is an *eigenstate* corresponding to the eigenvalue that occurs. In non-technical language, an eigenstate of an operator is a state that is left unchanged by application of the operator, and an eigenvalue is the measurement outcome corresponding to a given eigenstate. Thus, if the system is measured again immediately after a previous measurement using the same operator, the same measurement outcome will occur and the post-measurement state will be the same as the state after the previous measurement.[6] In a non-degenerate measurement, the eigenvalues and eigenvectors are in one-to-one correspondence; in a degenerate measurement, there may be many eigenvectors for a given eigenvalue. The pre-measurement state *QuantumState*(*OS*,*T2-*) is represented as a sum in which each summand is a complex-valued weight (called an amplitude) multiplied by an eigenvector of *Operator*(*T*2). The amplitudes are normalized so that the sum of their squared magnitudes is 1. The probability that *Outcome*(*O*2) will equal a given eigenvalue is equal to the sum of the squared magnitudes of the complex amplitudes for its corresponding eigenvectors. The weighted sum of eigenstates of *Operator*(*T*2) is called a *superposition* of the corresponding eigenstates. After measurement, the amplitudes for eigenstates corresponding to outcomes that do not happen become zero, and the amplitudes are re-normalized so that the squared amplitudes for eigenstates corresponding to outcomes that happen sum to one. This re-normalization step is mathematically equivalent to applying Bayes rule to Bohm's "wave from which many related probabilities can be calculated." That is, after a measurement has occurred, the "wave of related probabilities" is conditioned on the measurement operator that was applied, the time of application, and the actual outcome that occurred. Predictions about the future behavior of the state are revised accordingly. In experimental tests, the probability forecasts produced by this recipe have proven to be stunningly accurate.

The model presented thus far simply restates in the language of decision graphs the standard method of predicting the evolution of quantum systems and the outcomes of measurement events.[7] The remaining parts of Figure 4 are discussed only informally in most quantum theory texts. The model presented here is consistent with standard informal accounts and with the mathematics of quantum theory. The parts of the model involving the thoughts and decisions of the reducing agent *RA* are quantum systems whose defining properties are expressed in the language of classical physics and psychology rather than in the language of quantum theory. The model is consistent with a physical realization of *RA* as a quantum system, but empirical predictions for this part of the model cannot be obtained to anything approaching the level of precision of the probabilistic predictions for *Outcome*(*T*2).

The language used in quantum theory texts to describe the measurement process clearly implies that the time of measurement and the operator applied are free choices of the experimenter. Many authors have speculated that because the bodies and brains of experimenters are presumed to obey the laws of physics, it might some day be possible to develop a fundamental theory of measurement in which the experimenter, the experimental apparatus, and the system undergoing measurement are all described in the language of quantum theory. Although there exist experimental protocols, described in classical language, that specify the processes a physicist carries out to apply a given operator at a given time, present-day physics has no

---

[6] This property of eigenstates, together with the fact that probabilities, unlike probability amplitudes, evolve non-linearly in the Schrödinger equation, provide the mathematical basis for the quantum Zeno effect that has been hypothesized as a mechanism by which conscious affects observable outcomes.

[7] The description is qualitative because it is not our purpose is to provide detailed calculations of quantum probabilities. Mathematical formulas for standard problems can be found in any quantum theory text.

fundamental theory governing the generation of the choices made by the agents. Our current knowledge of the allowable operators and the allowable temporal intervals between application of operators is entirely empirical.

The model of Figure 4 formalizes the assignment of *T*2-*T*1 and *Operator*(*T*2) as free choices of the experimenter by representing these variables as decision nodes in the graph. The possible values of these decision nodes are specified by means of existing phenomenological theories that depend on the particular type of system being analyzed and the physical parameters of the experimental apparatus.[8] The information available to inform *RA*'s choice is labeled *Prevision*(*RA*), borrowing a term used by de Finetti (1934) to refer to agents' subjective predictions of the outcomes of uncertain events. The prevision depends on *RA*'s state. As noted above, the defining properties of this state are expressed in classical language, but its physical realization is as a quantum system. *RA*'s prevision is described in psychological terms. Thus, the prevision mediates between the physical description of *RA*'s state and the psychological variables in a cognitive-level model of *RA*'s prediction and decision-making process. We can think of *QuantumState*(*RA*,*T*1+) as describing the hardware of *RA*'s sensory and computing apparatus, and *Prevision*(*RA*,*T*1) as describing *RA*'s mental state in a higher-level language. Indeed, we may think of *RA*'s prevision as a more or less faithful rendition of the decision graph of Figure 4. The job of the unconscious processes of *RA*'s brain is to solve the influence diagram: identify a set of action policies, compute their expected utilities given the currently available evidence, discard all but the few with the highest expected utility, and present the results to consciousness for selection of a policy to be implemented. The job of *RA*'s conscious attention-directing mental process is to select and implement a policy from among those presented by the unconscious processors.

Graphical probability models provide a convenient language both for representing *RA*'s prevision and for describing its relation to *RA*'s physical state. A prevision model would include cognitive-level random variables such as those discussed in Section 2 above. For example, *RA* might reason about the likelihood of different macroscopically described consequences such as whether a dial on a measuring instrument will point in the up or down direction, or whether the ink cartridge needs to be replaced in a device that records measurement results. Also described in psychological terms is the degree of satisfaction *RA* experiences with the outcome of the experiment. In the model of Figure 4, the *utility* node *Satisfaction*(*RA*,*T*1) may depend on the entire trajectory of *RA*'s state until the next interaction event. Thus, *RA*'s satisfaction depends on *RA*'s state just after the last interaction with *OS*, the forces to which *RA* is subject, and the time until the next interaction -- these being the variables that affect the evolution of *RA*'s state between measurement events. Unlike the carefully controlled conditions under which the state of *OS* evolves, it may not be reasonable to suppose that *RA*'s state is completely shielded from all external influences during the time between measurements. Fortunately, no such assumption is required for the graph of Figure 4 to apply to *RA*'s evolution. We can model *RA* as being in a *mixed* rather than a *pure* state and include random fluctuations in the equation governing the evolution from *QuantumState*(*RA*,*T*1+) to *QuantumState*(*RA*,*T*2-). It also may be desirable to assume a time-varying Hamiltonian.

The utility node in Figure 4 represents *RA*'s actual experienced satisfaction with the outcome of the decision. As noted above, *RA*'s satisfaction is experienced after the choice has been made, and therefore cannot influence the choice *RA* makes. However, *RA*'s *anticipated* satisfaction is represented in the prevision. This is depicted in Figure 4 by a miniature copy of the decision graph inside the *Prevision*(*RA*,*T1+*) node representing *RA*'s model of the decision situation. A rational reducing agent makes a choice among the available policies by comparing the anticipated

[8] A more complete representation would include these parameters as context random variables affecting the possible values and probability distributions of the random variables in the classically described part of Figure 4.

satisfaction for the different policies and selecting the one with the highest expected utility. An agent whose decision graph is an accurate model of the decision situation and who chooses optimally according to the decision graph will, on average, achieve a higher level of satisfaction than agents with the same evidence and utilities who either have inaccurate models or choose non-optimally.

Schwartz, et al. (2004) present a model of the relative roles of brain and mind that is consistent with Figure 4:

> "[Schrödinger evolution] can do most of the necessary work of the brain. It can do the job of creating, on the basis of its interpretation of the clues provided by the senses, a suitable response, which will be controlled by a certain pattern of neural or brain activity that acts as a *template for action*. But, due to its quantum nature, the brain necessarily generates an amorphous mass of overlapping and conflicting templates for action. [Application of a reduction operator] acts to extract from this jumbled mass of possibilities some particular template for action."

A great deal is known about the physical mechanisms by which the brain operates. Schwartz, et al. (2004) state that nerve firings are triggered by the flow of calcium ions through ion channels that are less than a nanometer in diameter at their narrowest point. At these dimensions, they state, quantum effects cannot be ignored. The spatial constriction implies, by the Heisenberg uncertainty principle, a corresponding spread in the lateral velocity. They hypothesize that multiple patterns of nerve firings representing different macroscopic possibilities for the agent's plan of action evolve in superposition until the agent chooses a reduction operator that defines the possibilities (perhaps just two) between which Nature then chooses in accordance with the statistical rules. Nature's choice actualizes one of the possibilities: an eigenvalue and its set of eigenvectors. Schwartz, et al. suggest the quantum Zeno effect as a mechanism by which the agent's choice merely of the timings of its own actions, can hold a long sequence of Nature's choices in place and thereby strongly influence what eventually is actualized. The different possibilities, they hypothesize, would most likely be

> "…constructed out of oscillating states of macroscopic subsystems of the brain, rather than out of the states of individual particles. The states associated with [measurement outcomes would involve] large assemblies of particles of the brain… moving in a coordinated way that will lead on, via the mechanical laws, to further coordinated activities."

These oscillating states of macroscopic subsystems of the brain may be performing the functional equivalent of a quantum computer constructing a decision graph for the problem faced by the reducing agent and solving for an optimal policy. One important function of conscious control might be to mediate value conflicts, such as the conflict between the craving for a cigarette and the desire for a long and healthy life. Deliberate focusing of attention directs the system toward more nearly optimal choices and away from the ones that would be favored if purely physical automatic processors were left unmonitored.

The above description summarizes the decision graph representation of the process of evolution of quantum systems. The upper part of Figure 4 restates the standard account of Schrödinger evolution and the measurement process as described in texts on quantum theory. The lower part of the graph corresponds to standard informal accounts of the quantum measurement process. The time of measurement and the operator to be applied are represented as free choices of the reducing agent. This choice depends on the agent's quantum state through the reducing agent's prevision, which can be thought of as a rendering in psychological terms of those aspects of the quantum state that are relevant to the agent's anticipated satisfaction with the outcome.

An intelligent reducing agent's brain is organized to simulate task-relevant aspects of its environment and to improve the fidelity of its simulations in a manner that approximates Bayesian learning from observations. It is hypothesized that evolution selects for the ability to make accurate forecasts of outcomes, the ability to apply attention to effect reduction events that bring about desired outcomes, and the desire for outcomes conducive to survival. It is hypothesized that evolution also selects for systems of agents that co-evolve toward evolutionarily stable patterns of interaction.

4 Discussion: Quantum Computing Agents

The above sections provided a sketch of the basic elements of a graphical model based software architecture for physical symbol systems implemented in quantum hardware. To behave intelligently, an implementation of the reducing agent architecture described in Figure 4 would need to possess the basic attributes the artificial intelligence and machine learning communities have learned are essential to intelligent performance. Among these capabilities are the following:

- *RA*'s state would need to include a memory subsystem in which past events are recorded, and from which information about past events can be recalled.
- The physical architecture of *RA*'s memory should be such that recorded memory traces remain stable over time.
- *RA* would require a physical means for recording the outcomes of measurement events in its memory and retrieving outcomes of remembered events.
- *RA*'s memory would require a large storage capacity and its contents should be indexed for efficient retrieval. In particular, its contents should be organized according to high-level conceptual categories. Data organized purely as records of raw sensor inputs, no matter voluminous, may be worse than useless from the perspective of supporting task performance.
- *RA* would need to have a method for forecasting the outcomes of its actions, that includes flexible methods for reasoning under uncertainty and adapting to unanticipated situations. These methods should be functionally equivalent to approximate Bayesian inference.
- *RA* would require a method for identifying and selecting actions or policies predicted to lead to desired outcomes. These methods should be functionally equivalent to approximate maximization of subjective expected utility.
- The objective function built into *RA*'s hardware (what might be termed its instinctual desires) should be conducive to long-term survival of *RA* and others of its kind, should minimize negative impact on its natural environment, and should be supportive of human society.

The above characteristics have been treated extensively in the literature on intelligent systems, without any reference to quantum computing. Current classical architectures implement many of these features, and their power and capabilities are advancing rapidly. So why should researchers in machine learning and artificial intelligence pay attention to quantum computing?

First, the argument that current technology is advancing rapidly and there is no apparent need for a new, unproven, and as yet poorly developed technology was used against the development of decision theoretic technology prior to the 1980's. Yet in a few short decades, decision theory has moved from the fringes into the mainstream of machine learning and artificial intelligence. Arguments that real intelligent agents do not optimize and that the intrinsic impracticality of decision theoretic optimality necessitates abandoning it in favor of ad hoc heuristics have given way to acceptance of approximate optimality as a fundamental characteristic of intelligence and a view of decision theoretic optimality as a standard of comparison and foundation for the design of good heuristics. In a similar vein, if we are going to demand that our theories of intelligent agents be validated in physically realizable implementations, then our foundational theory for the construction of the physical hardware for our implementations should be based on accurate physical theory, and our foundational theory for software architectures should fully exploit the capabilities of the hardware. Most scientists, even those who believe quantum theory is fundamental to an understanding of intelligence and cognition, believe that classical physics will be adequate for many aspects of the architecture of intelligent agents. But there is a vocal and growing community arguing that quantum effects are likely to be fundamental for some aspects of intelligence and learning. A theory of computing built on a quantum theoretic foundation can address questions classical computing theory cannot, while reducing to classical theory in areas to which classical theory applies.

Second, quantum algorithms appear to be intrinsically more powerful for some types of problem than classical algorithms. This power appears to be related to the ability to perform Fourier analysis with great efficiency. Efficient spatio-temporal reasoning is among the most important and as yet intractable challenges in machine learning and artificial intelligence. Because Fourier analysis is an essential element of standard methods for spatial and temporal reasoning, it is reasonable to suggest that quantum computation might provide new and more powerful methods for reasoning and learning in problems involving space and time.

Finally, quantum theory with the reducing agent ontology provides a plausible physical mechanism for efficacious free choice. Standard theories of computing, cognition and learning are limited to two categories of dynamic behavior -- deterministic evolution and randomization. The reducing agent ontology treats genuine free choice as an intrinsically different category of behavior from either determinism or randomness. If indeed genuine free choice exists in the natural world, is fundamentally different from both deterministic causation and random effects, and produces observable effects in the world, then a science that excludes the possibility of genuine free choice is incomplete and fundamentally flawed as a foundation for theories of cognition and learning.

Acknowledgements

The ideas presented in this paper had their genesis in research conducted under a career development fellowship awarded to the author in 1994 by the Krasnow Institute at George Mason University. Many people have contributed to the evolution of the concepts presented here. Special appreciation is due to Henry Stapp for many long discussions and patient explanations, and for helpful suggestions on an earlier draft of this paper. Thanks are also due to Menas Kafatos, Paul Lehner, Tod Levitt, Laura Martignon, Harold Morowitz, Karl Pribram, and Jack Sarfatti for encouragement and helpful discussions. Last but not least, this paper is dedicated to the memory of Danny Pearl and to the potential of his father Judea Pearl's seminal research, if properly applied, to build a world in which people of all faiths and ethnic backgrounds live together in harmony. May those who till the earth Judea plowed keep Danny's memory alive and bring his dream to fruition.

References

Dorit Aharanov, Quantum Computation, In D. Stauffer (ed.), *Annual Reviews of Computational Physics VI*. Singapore, World Scientific, 1999.

Anderson, J.R. *Cognitive Psychology and Its Implications*, Worth Publishing, 1999.

Olav Bangsø and Pierre-Henri Wuillemin. Object Oriented Bayesian Networks: A Framework for Topdown Specification of Large Bayesian Networks and Repetitive Structures. Technical Report CIT-87.2-00-obphw1, Department of Computer Science, Aalborg University, 2000.

Berg, J., Nelson, F., & Rietz, T. (2001). Accuracy and Forecast Standard Error of Prediction Markets. Tech. rep., University of Iowa, College of Business Administration.

Thomas Binford and Tod S. Levitt. Evidential reasoning for object recognition, *IEEE Transactions Pattern Analysis and Machine Intelligence*, 25(7): 837-851.

David Bohm, *Quantum Theory*, Prentice-Hall, 1951.

Bohm, David and Hiley, Basil. *The Undivided Universe*. Rutledge, London, England, 1993.

Niels Bohr. , *Atomic Theory and the Description of Nature*. Cambridge: Cambridge University Press, 1934.

Craig Boutilier, Thomas Dean and Steve Hanks. Decision-Theoretic Planning: Structural Assumptions and Computational Leverage. *Journal of Artificial Intelligence Research* (*JAIR*) 11: 1-94, 1999.

Wray Buntine. Operations for Learning with Graphical Models. *Journal of Artificial Intelligence Research,* 2: 159-225, 1994.

A.R. Cassandra, L.P. Kaelbling, L.P., and M.L. Littman. Acting optimally in partially observable stochastic domains. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, Seattle, WA, 1994.

Dawid, A.P. and Vovk, V.G. (1999), Prequential Probability: Principles and Properties, *Bernoulli, 5: 125-162.*

David Deutsch, Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer, *Proceedings of the Royal Society of London A*, 400:97, 1985.

Morris deGroot. *Optimal Statistical Decisions*. McGraw-Hill, 1970.

Bruno de Finetti. *Theory of Probability: A Critical Introductory Treatment* (2 vols.), New York: Wiley, 1974-75.

Robert J. Elliott, Lakhdar Aggoun and John B. Moore. *Hidden Markov Models: Estimation and Control*. Springer-Verlag, New York, New York, 1995.

Fienberg, S.E. and DeGroot, M.H. (1982) Assessing probability assessors: Calibration and refinement, *Statistical Decision Theory and Related TopicsIII*, Vol. (S.S. Gupta and J.O. Berger, eds.), Academic Press, 291--314.

Fine, Terrence L. *Theories of Probability*. New York: Academic Press, 1973.

Gershenfeld, N. and Chuang, I. Quantum Computing with Molecules, *Scientific American*, June 1998.

Lise Getoor, Nir Friedman, Daphne Koller, and Avi Pfeffer. Learning Probabilistic Relational Models. In Saso Dzeroski and Nada Lavrac, editors. *Relational Data Mining*, Springer-Verlag, New York, New York, 2001.

Lise Getoor, Daphne Koller, Benjamin Taskar, and Nir Friedman. Learning Probabilistic Relational Models with Structural Uncertainty. In *Proceedings of the ICML-2000 Workshop on Attribute-Value and Relational Learning:Crossing the Boundaries*, Stanford, California, 2000.

Gigerenzer, G., Todd, P., and the ABC Group. *Simple Heuristics that Make us Smart.* Oxford University Press, 1999.

W. Gilks, S. Richardson, and D. Spiegelhalter (eds) (1996). *Markov Chain Monte Carlo in Practice*. London: Chapman & Hall.

W. Gilks, A. Thomas, and David J. Spiegelhalter. A language and program for complex Bayesian modeling. *The Statistician,* 43: 169-178, 1994.

Ulf Grenander. *Elements of Pattern Theory*. Johns Hopkins University Press, Baltimore, MD, 1995.

John Gribbin. *Schrödinger's Kittens and the Search for Reality*, London: Phoenix, 1996.

L. Grover. In Proceedings of the 28th Annual ACM Symposium on the Theory of Computation, pp. 212-219, ACM Press, New York, 1996.

Colin Howson and Peter Urbach. *Scientific Reasoning: The Bayesian Approach*. Open Court, 1993.

W. Itano, D. Heinzen, J. Bollinger, D. Wineland, Phys. Rev. 41A, 2295-2300, 1990.

William James. *Principles of Psychology*, Dover, 1890.

Finn V. Jensen. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, New York, 2001.

Michael Kearns and Y. Mansour. Efficient Nash Computation in Large Population Games with Bounded Influence. In *Uncertainty in Artificial Intelligence: Proceedings of the Eighteenth Conference*, Morgan Kaufmann Publishers, San Mateo, California, 2002.

S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. *Optimization by simulated annealing*. Science, 220:671--680, 1983.

Daphne Koller and Avi Pfeffer. Object-Oriented Bayesian Networks. In *Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference*, pages 302-313, Morgan Kaufmann Publishers, San Francisco, California, 1997.

Kevin B. Korb and Ann E. Nicholson, *Bayesian Artificial Intelligence*, Chapman and Hall, 2003.

Helge Langseth and Thomas Nielsen. Fusion of Domain Knowledge with Data for Structured Learning in Object-Oriented Domains. *Journal of Machine Learning Research*, 2003

Kathryn Laskey, *MEBN: A Logic for Open-World Probabilistic Reasoning*. George Mason University, 2003.

Kathryn B. Laskey and Suzanne M. Mahoney. Network Fragments: Representing Knowledge for Constructing Probabilistic Models. In *Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference*, Morgan Kaufmann Publishers, San Mateo, California, 1997.

Kathryn B. Laskey, Suzanne M. Mahoney, and Edward Wright. Hypothesis Management in Situation-Specific Network Construction. In *Uncertainty in Artificial Intelligence: Proceedings of the Seventeenth Conference*, Morgan Kaufmann Publishers, San Mateo, California, 2001.

T.S. Lee, and David Mumford, Hierarchical Bayesian inference in the visual cortex., *J.Opt Soc Am A Opt Image Sci Vis*. 20(7):1434-48, 2003.

Tod S. Levitt, C. Larrabee Winter, Charles J. Turner, Richard A. Chestek, Gil J. Ettinger, and Steve M. Sayre. Bayesian Inference-Based Fusion of Radar Imagery, Military Forces and Tactical Terrain Models in the Image Exploitation System/Balanced Technology Initiative. *International Journal of Human-Computer Studies,* 42, 1995.

JohnJoe McFadden *Quantum Evolution*, Harper-Collins, 2000.

Nau, R.F. and K.F. McCardle Arbitrage, Rationality, and Equilibrium. *Theory and Decision* **31** 199-240, 1991.

Benjamin Libet, E.W. Wright, Jr., B. Feinstein, and D. K. Pearl, Subjective referral of the timing for a conscious sensory experience. *Brain* 102, 193-224, 1979.

Laura Martignon and Kathryn B. Laskey, Taming Wilder Demons: Bayesian Benchmarks for Fast and Frugal Heuristics, in Gigerenzer, G., Todd, P., and the ABC Group. *Simple Heuristics that Make us Smart* Oxford University Press, 1999.

Richard Neapolitan. *Learning Bayesian Networks*, Prentice-Hall, 2003.

Alan Newell and Herbert Simon. Computer Science as Empirical Inquiry: Symbols and Search, *Communications of the ACM*, 19(3), 1976.

Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information*, Cambridge, UK, Cambridge University Press, 2000.

R. C. Parker and R. A. Miller. Using causal knowledge to create simulated patient cases: the CPCS project as an extension of Internist-1. *Proceedings of the Eleventh Annual Symposium on Computer Applications in Medica Care*. IEEE Comp Soc Press, 473-480, 1987.

Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, San Francisco, California, 1988.

Roger Penrose. *Shadows of the Mind*, Oxford University Press, 1994.

Avi Pfeffer. IBAL: A Probabilistic Rational Programming Language *International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.

John W. Pratt, Howard Raiffa, and Ronald Schlaifer. *The Foundations of Decision Under Uncertainty: An Elementary Exposition*. New York, NY: McGraw Hill, 1965.

Stuart Russell and Peter Norvig. *Artificial Intelligence*: *A Modern Approach*. (2nd edition) Prentice-Hall, New York, New York, 2002.

Jeffrey M. Schwartz, Henry P. Stapp, and Mario Beauregard. Quantum Physics in Neuroscience and Psychology: A New Model with Respect to Mind/Brain Interaction. To appear in *Behavior and Brain Science*, 2004.

Leonard J. Savage. *The Foundations of Statistics*. New York: Dover Publications, Inc., 1972.

Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game*: New York: Wiley, 2001.

Peter W. Shor, Algorithms for Quantum Computation: Discrete Logarithms and Factoring, In *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, IEEE Press, Los Alamitos, CA, 1994.

R. Solovay and V. Strassen, A Fast Monte-Carlo Test for Primality, *SIAM Journal of Computing,* 6:84-85, 1976.

David J. Spiegelhalter, Andrew Thomas, and Nicky Best. Computation on Graphical Models. *Bayesian Statistics,* 5: 407-425,1996.

Henry P. Stapp Science of Consciousness and the Hard Problem. *The Journal of Mind and Behavior* 18, 171-193, 1999.

Henry P. Stapp. Whiteheadian Process and Quantum Theory, conference on Whitehead, 1998.

Paul Thagard. *Computational Philosophy of Science*. Cambridge, MA: MIT Press, 1988.

John von Neumann. *Mathematical Foundations of Quantum Mechanics* (Ch. VI), Princeton, NJ., Princeton University Press, 1932

Detlof von Winterfeldt and Ward Edwards. *Decision Analysis and Behavioral Research*. Cambridge, UK: Cambridge University Press, 1986.

Wigner , E., Remarks on the Mind-Body Problem. in *Symmetries and Reflections*, Indiana University Press, Bloomington, 171-184, 1967.

R. Shankar, *Principles of Quantum Mechanics*, Plenum, 1994.