Action Networks: A Framework for Reasoning about Actions and Change under Uncertainty

Moisés Goldszmidt and Adnan Darwiche
Rockwell International Science Center
444 High Street, Suite 400
Palo Alto, CA 94301, U.S.A.
{moises, darwiche}@rpal.rockwell.com

Abstract

This work proposes action networks as a semantically well founded framework for reasoning about actions under uncertainty. Action networks extend probabilistic causal networks to allow the representation of actions as directly controlling specific events in the domain (i.e., setting the value of nodes in the network) subject to preconditions. They also introduce models of time and persistence, and the capability of specifying uncertainty at different levels of abstraction. This paper describes both recent results and work in progress.

1 Introduction

The work reported in this paper is part of a project that proposes a decision support tool for plan simulation and analysis [6]. The objective is to assist a human/computer planner in analyzing plan trade-offs and in assessing properties such as reliability, robustness, and ramifications under uncertain conditions. The core of this tool is a framework for reasoning about actions under uncertainty called *action networks.* Action networks are extensions to probabilistic causal networks (Bayes networks [20]) that allow the modeling of actions, their preconditions, time, and persistence. They also allow the symbolic and numeric quantification of uncertainty in the cause-effect relations at different levels of abstraction.

The work on action networks was motivated by the need to extend the capabilities of current frameworks for reasoning about plans to include notions and models of uncertainty. Reasoning about plans in AI has been traditionally tackled by the symbolic logic community, which does not emphasize the inherent uncertainties in real-world domains. Only recently did this problem of uncertainty become a focus of planning research, but, unfortunately, most new formalisms introduce minimal abilities for handling uncertainty and seem to inherit the difficulties associated with their symbolic logic origins, such as the frame, ramification, and concurrency problems.

In the same way that logic-based approaches for planning have stopped short of addressing the uncertainties inherent in real-world applications, probability-based approaches have stopped short of providing adequate representations of time and action. The probabilistic literature, for example, does not seem to provide an agreed upon model of time and action. Moreover, most existing proposals for action do not account for action preconditions — which seems to be necessary in realistic AI applications — and require that actions be given a special status, as in influence diagrams.

The work on action networks is an attempt to combine the best of the two worlds. In particular, action networks aim at combining uncertainty techniques from probabilistic causal networks with action and time representations from logic-based approaches for planning. Action networks also insist on representing the causality of a domain using causal structures and uses such information in dealing with some of the key obstacles in reasoning about actions, including the frame, ramification, and concurrency problems [18, 7].

The development of action networks involves three key tasks, which correspond to the representation of uncertainty, action, and time. We will now briefly describe these tasks.

Traditionally, causal networks have been quantified probabilisticly by providing the probabilities of effects given their causes. But action networks as proposed here are based on causal networks that are not necessarily probabilistic. Instead, a causal network will consist of two parts: a directed graph representing a "blueprint" of the causal relationships in the domain and a quantification of these relationships. The quantification introduces a representation of the uncertainty in the domain because it specifies the degree to which causes will bring about their effects. Action networks will allow uncertainty to be specified at different levels of abstraction: point probabilities, which is the common practice in causal networks [20], order-of-magnitude probabilities, also known as ε-probabilities [13, 14], and symbolic arguments, which allow one to explicate logically the conditions under which causes would bring about their effects [3]. In this paper, we will concentrate on order-of-magnitude probabilities as proposed in [13]. Other quantifications are described elsewhere [6, 7].

Action networks add to causal networks a missing link that is necessary for reasoning about action, namely, the ability to reason about changes in beliefs due to actions as opposed to observations. In particular, action networks introduce the notion of a controllable variable and the associated notion of a precondition arc. Controllable variables are ones that can be influenced directly by an agent, and precondition arcs connect controllable variables to other variables that set preconditions for controllability.

When reasoning about plans, the representation of time becomes inevitable since the execution of plans takes place and causes changes across time. Therefore, in simulating the execution of plans, one must resolve a number of issues that are related to time [16, 9, 8]. At the heart of these issues is modeling persistence: how would variables in a given domain persist over time when they are not influenced by actions? For this purpose, action networks support a model of persistence that is based on viewing a non–deterministic causal network as a parsimonious encoding of a more elaborate deterministic one in which suppressors (exceptions) of causal influences are explicated. The basic intuition here is that suppressors persist over time, and that things tend to persist when causal influences on them are deactivated by these suppressors.

This paper is organized as follows. Section 2 describes action networks. It starts with a brief review of network based representations, and the κ-calculus of plain beliefs (Section 2.1). Section 2.2 presents a functional expansion of a causal network, and Section 2.3 introduces the suppressor model, which are the technical basis for the model of persistence explained in Section 2.4. The representation of actions can be found in Section 2.5. Finally, Section 3 summarizes the main results — including an implementation of action networks on top of CNETS [4] — and describes future work.

2 Action Networks

The specification of an action network requires three components: (1) the causal relations in the domain with a quantification of the uncertainty of these relations, (2) the set of events that are "directly" controllable by actions with the events that constitute their respective preconditions, (3) events that persist over time.

Once the domain is modeled using this network-based representation (including uncertainty), the action network will expand this information into a functional representation of the causal relations, and will then expand over time those events marked to be persistent.

In this paper, propositions will be denoted by lowercase letters e, we will assume that propositions can take values from $\{false, true\}$,[1] which will be denoted as e^- and e^+ respectively. An instantiated proposition (or set of propositions) will be denoted by $\vec{e}$, and $\neg\vec{e}$ denotes the "negated" value of $\vec{e}$. .

2.1 Network Representations

As action networks are extensions of probabilistic causal networks, we briefly review some of the key concepts behind a network representation. A *causal network* consists of a directed–acyclic graph Γ and a quantification Q over Γ. Nodes in Γ correspond to domain variables, and the directed edges correspond to the causal relations among these variables. We denote the set of parents $\{c_1, \ldots, c_n\}$ of a node e in a belief network by $\pi(e)$. $\vec{\pi}(e)$ will denote a state of the conjunction of propositions that constitute the parent set of e. The set $\pi(e)$ represents the direct causes for e, and the set conformed by e and its causes is usually referred to as the causal family of e (see for example Figure 1 representing the causal family of *wet_grass*).

The quantification of Γ over the families in the network encodes the uncertainty in the causal influences between $\pi(e)$ and e. In Bayesian networks, this uncertainty is encoded using numerical probabilities [20]. There are, however, other ways to encode this uncertainty that do not require an exact and complete probability distribution. Two recent approaches are the κ-calculus where uncertainty is represented in terms of plain beliefs and degrees of surprise [14, 13], and argument calculus where uncertainty is represented using logical sentences as arguments [2, 3]. These approaches are regarded as abstractions of probability theory since they retain the main properties of probability including Bayesian conditioning [5, 14]. Although action networks allows any of these versions of uncertainty, for the purposes of this paper we focus on the κ-calculus. We provide next a brief summary of κ and how it can quantify over the causal relations in a network Γ.

Let $\mathcal{M}$ be a set of worlds, each world $m \in \mathcal{M}$ being a truth-value assignment to a finite set of atomic propositional variables $(e_1, e_2, \ldots, e_n)$. Thus any world m can be represented by the conjunction of $\vec{e_1} \wedge \ldots \wedge \vec{e_n}$. A belief *ranking function* $\kappa(m)$ is an assignment of non-negative integers to the elements of $\mathcal{M}$ such that $\kappa(m) = 0$ for at least one $m \in \mathcal{M}$. Intuitively, $\kappa(m)$ represents the degree of surprise associated with finding a world m realized, worlds assigned $\kappa = 0$ are considered serious possibilities, and worlds assigned $\kappa = \infty$ are considered absolute impossibilities. A proposition $\vec{e}$ is believed iff $\kappa(\neg\vec{e}) > 0$ (with degree k iff $\kappa(\neg\vec{e}) = k$). $\kappa(m)$ can be considered an order-of-magnitude approximation of a probability function $P(m)$ by writing $P(m)$ as a polynomial of some small quantity ε and taking the most significant term of that polynomial, i.e., $P(m) \cong C\varepsilon^{\kappa(m)}$. Treating ε as an infinitesimal quantity induces a conditional ranking function $\kappa(e|c)$ on propositions and wffs governed by properties derived from the ε-probabilistic interpretation [12].

A causal structure Γ can be quantified using a ranking belief function κ by specifying the conditional belief for each proposition $\vec{e}$ given every state of $\vec{\pi}(e)$, that is $\kappa(\vec{e}|\vec{\pi}(e))$. Thus, for example, Table 1 shows the κ-quantification of the *wet_grass* family (Figure 1).

Table 1 is called the "κ-matrix" for the *wet_grass* family. It represents the following pair of default causal rules: "If $rain^+$ or $sprinkler^+$ then wet_grass^+" and "If $rain^-$ and $sprinkler^-$ then wet_grass^-". [2]

[1] Propositions need not be binary; in fact suppressors (Sec. 2.3) will have more than 2 values.

[2] In fact, there is no need in general to specify the 2^n values in the matrix. There is a compilation process that would allow rules as input and will fill in the unspecified values

rain	*sprinkler*	*wet_grass* = true	*wet_grass* = false
true	true	0	2
true	false	0	1
false	true	0	1
false	false	3	0

Table 1: κ-Quantification for the *wet_grass* family.

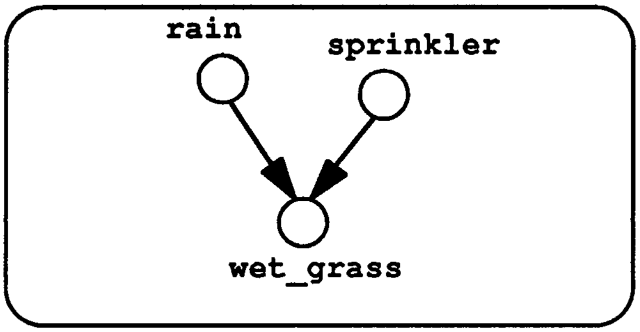

Figure 1: The *wet_grass* causal family.

The κ-calculus does not require commitment about the belief in e^+ or e^-. Thus, for example in the example in Figure 4 we may be ignorant about the status of *alive* given that the victim is shot with an unloaded gun. In such a case the user may specify that both $\kappa(alive^+|fired\_gun^+ \wedge loaded\_gun^-) = \kappa(alive^-|fired\_gun^+ \wedge loaded\_gun^-) = 0$ indicating that both alternatives are plausible.

Once similar matrices for each one of the families in a given network Γ are specified, the complete ranking function can be reconstructed by requiring that

$$\kappa(m) = \sum_i \kappa(\vec{e_i}|\vec{\pi}(e_i)), \qquad (1)$$

and queries about belief for any proposition and wff can be computed using any of the well-known distributed algorithms for belief update [20, 4]. The class of ranking functions that comply with the requirements of Eq. 1 for a given network Γ are called ***stratified rankings***. These rankings present analogous properties about conditional independence to those possessed by probability distributions quantifying a network with the same structure [13].

2.2 Functional Characterization of Causal Networks

The uncertainty recorded in the quantification of each family in a network Γ expresses the incompleteness of our knowledge in the causal relation between e and its set of direct causes $\pi(e)$. This incompleteness arises because e interacts with its environment in a complex manner, and this interaction usually involves factors which are exogenous to $\pi(e)$. Furthermore, these factors are usually unknown, un-observable or too many to enumerate. Thus, we can view a nondeterministic causal family as a parsimonious representation of a more elaborate, deterministic causal family, where the quantification summarizes the influence of other factors on e. In principle, we could expand the original network by adding another parent $\mathcal{U}(e)$ representing these exogenous influences on e. The minimum requirements for this new representation to qualify as an expansion of the original network are (let κ represent degree of belief in the original network, and κ' represent degree of belief in the expanded functional network): (1) For a fixed value of $\mathcal{U}(e)$ the conditional degree of belief $\kappa'(\vec{e}|\mathcal{U}(e) \wedge \vec{\pi}(e))$ should be functional[3] and, (2) When we marginalize over the states of $\mathcal{U}(e)$ the degree of belief in $\vec{e}$ given its causes $\vec{\pi}(e)$ should be equal to $\kappa(\vec{e}|\vec{\pi}(e))$, the degree of belief in the original network.

This model is intuitively appealing in that it encodes the causal relation of a family as a set of functions between the direct causes $\pi(e)$ and its effect e, where the state of the exogenous causes $\mathcal{U}(e)$ selects the "active" function that specifies the current causal relation. The likelihood that any of these functions is active depends entirely on the likelihood of the state of $\mathcal{U}(e)$. Notice that persistence in this model has a distinctive and precise flavor, since we can now model not only the tendency of states to remain unchanged but also the functional specification that determines their causal relation. A similar characterization was used by Pearl and Verma [22] for discovering causal relationships (and causal models) from observations, by Druzdzel and Simon [10] in their study about the representation of causality in Bayes networks, and by David Heckerman to represent persistence in a diagnosis/repair application involving a probabilistic quantification of uncertainty.[4]

The properties of such an expansion including the simulation of persistence and change are explored in depth by Dagum, Darwiche and Goldszmidt [1]. One important characteristic of this expansion is that it is not unique. That is, given a quantified causal network, there are many functional expansions that comply with conditions (1) and (2) above. An immediate consequence of this non–uniqueness is that in order to complete the specification of the functional expansion into a quantified network we must make some additional assumptions about the functional behavior of the families. The assumptions we propose are standard in knowledge representation in AI, and especially in default reasoning. It essentially assumes that the set $\mathcal{U}(e)$ is constituted by a set of exceptions that try to defeat the causal influence of $\pi(e)$ on e [11, 19]. This model is called the suppressor model and it is explained next.

2.3 The Suppressor Model

The suppressor model, as a functional expansion of a nondeterministic family in a causal network, assumes

with zeros that works for a significant number of rule bases. This process is beyond the scope of this paper and will be reported elsewhere. To maintain the parallel with probabilistic networks we will continue to use matrices throughout the paper.

[3]In the κ-calculus, functional or deterministic means that either $\kappa'(\vec{e}|\mathcal{U}(e) \wedge \vec{\pi}(e)) = \infty$, or $\kappa'(\neg\vec{e}|\mathcal{U}(e) \wedge \vec{\pi}(e)) = \infty$, which is parallel to the probabilistic case of either $P(\vec{e}|\mathcal{U}(e) \wedge \vec{\pi}(e)) = 1$ or $P(\neg\vec{e}|\mathcal{U}(e) \wedge \vec{\pi}(e)) = 1$.

[4]Personal communication.

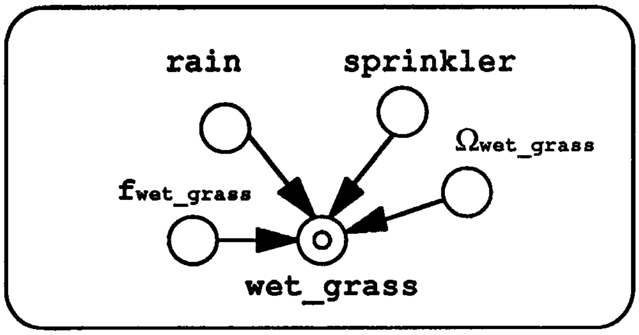

Figure 2: The functional expansion of the *wet_grass* family using the suppressor model.

that the uncertainty in the causal influences between $\pi(e)$ and e is a summary of the set of exceptions that attempt to defeat or suppress this relation. For example, "a banana in the tailpipe" is a possible suppressor of the causal influence of "turn key" on "car starts."

The expansion into the suppressor model makes the uncertain causal relation between $\pi(e)$ and e functional by adding a new parent to the family Ω_e that corresponds to the suppressors representing the possible exceptions to the causal relation. In addition to Ω_e, we add another parent f_e that will set the state of e in those cases in which the suppressors manage to defeat the causal influence of $\pi(e)$ on e. In these cases we say that the suppressors are "active".

When reasoning about change in a dynamic context, such as planning, f_e will usually represent the previous state of e. The intuition being that in those cases when the suppressors manage to prevent the natural causal influences on e, the state of e should simply persist and follow its previous state.

Suppressors will have different states corresponding to the specific causal influence they are trying to defeat. Moreover, the degree of beliefs in these states will correspond to the different degrees of belief in $\kappa(\vec{e}|\vec{\pi}(e))$. Figure 2 shows the expansion of the *wet_grass* family into a deterministic family under the suppressor model. The double circle for the node corresponding to *wet_grass* indicates that once the state of the parents, $\pi(wet\_grass)$, f_{wet_grass}, and Ω_{wet_grass} is known, the state of *wet_grass* is functionally determined.

Let Ω_e take values out of the set $\{\omega_e^0, \omega_e^1, \omega_e^2, \ldots\}$ where $\Omega_e = \omega_e^s$ stands for "a suppressor of strength s is active."[5] The function $\mathbf{F}$ relating e with its direct causes $\pi(e)$ and the expanded set including the suppressors Ω_e and f_e is given by

$$\mathbf{F}(\vec{\pi}(e), \omega_e^i, \vec{f}_e) = \begin{cases} e^+, & \text{if } \kappa(e^- \mid \vec{\pi}(e)) > i; \\ e^-, & \text{if } \kappa(e^+ \mid \vec{\pi}(e)) > i; \\ \vec{f}_e, & \text{otherwise.} \end{cases} \tag{2}$$

The translation of $\mathbf{F}$ into a κ matrix is given by Eq. 3 below:

$$\kappa'(\vec{e} \mid \vec{\pi}(e), \omega_e^i, \vec{f}_e) = \begin{cases} 0, & \text{if } \vec{e} = \mathbf{F}(\vec{\pi}(e), \omega_e^i, \vec{f}_e); \\ \infty, & \text{otherwise.} \end{cases} \tag{3}$$

Intuitively, the function $\mathbf{F}$ and its realization in the κ-calculus establishes that if the strength of the active suppressor ω_e^i is less than the causal influence of $\pi(e)$ on e, then the state of e is dictated by the causal influence. Otherwise, the suppressor is successful and the state of e is the same as the state of f_e, which in the case of our model of persistence corresponds to the previous state of e (see Section 2.4).[6]

The prior distribution of beliefs on Ω_e is given by

$$\kappa'(\omega_p^i) = i. \tag{4}$$

The functional specification in Eq. 3 and Theorem 1 below verify that the expansion of a causal family into the suppressor model complies with conditions (1) and (2) in Sec. 2.2. This guarantees soundness with respect to the original belief encoding $\kappa(\vec{e}|\vec{\pi}(e))$ of the causal family of e.

Theorem 1 $\kappa(\vec{e} \mid \vec{\pi}(e)) = \kappa'(\vec{e} \mid \vec{\pi}(e))$.

The proof of this theorem relies on marginalizing $\kappa'(\vec{e}|\vec{\pi}(e), \Omega_e, \vec{f}_e)$ over all the states of Ω_e and $\vec{f}_e$.

The combination of Eqs. 3 and 4 captures the intuitions behind the suppressor model. When the suppressors are inactive, ω_e^0, the functional specification that prevails is the one originally intended for the relation between e and its direct causes $\pi(e)$ with no exceptions. This is also the most plausible state, since according to Eq. 4, $\kappa'(\omega_e^0) = 0$. As $\Omega_e = \omega_e^i$, suppressors start disrupting causal influences whose strength of belief are less or equal to i. Yet, these suppressors become more and more unlikely since their degree of disbelief $\kappa'(\omega_e^i)$ increases proportionally. Table 2 shows the expansion of the causal relation between *wet_grass* and $\pi(wet\_grass)$ into the suppressor model.

Ω_{wet_grass}	Function
$\omega^0_{wet\_grass}$	*wet_grass* is true IFF either *rain* or *sprinkler* are true
$\omega^1_{wet\_grass}$	*wet_grass* is true IFF 1. both *rain* and *sprinkler* are true 2. either *rain* or *sprinkler* are true and f_{wet_grass} is true
$\omega^2_{wet\_grass}$	*wet_grass* is true IFF either *rain* or *sprinkler* are true and f_{wet_grass} is true
$\omega^3_{wet\_grass}$	*wet_grass*is true IFF f_{wet_grass} is true

Table 2: Suppressor model for the *wet_grass* family.

[5] As will be seen, the strengths on the suppressors are directly related to the belief strengths in the causal network.

[6] We remark that this model just corresponds to a standard default assumption behind the notion of persistence. The node f_e may take other values depending on the nature of e.

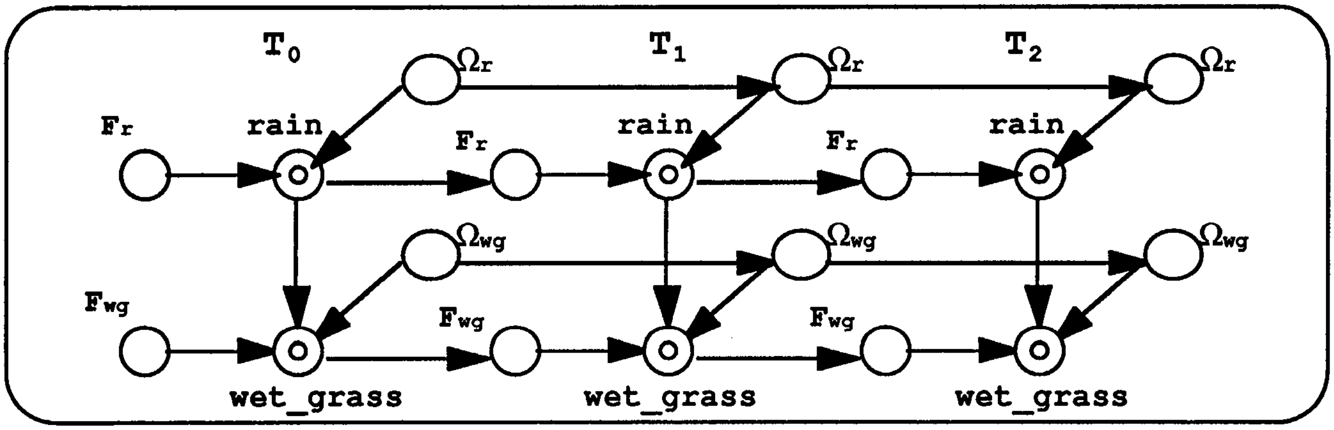

Figure 3: The temporal expansion of the simplified *wet_grass* network.

2.4 Time and Persistence

To represent the dynamics of beliefs and plausibility as we gather information about actions and observations, we simply use duplicates of the functionally expanded networks; each network referring to a specific time point.[7] This expansion of the network over time is depicted in Figure 3 where a simplified version of the *wet_grass* family (Figure 2) is expanded over three time points T_0—T_2.

Causal families are connected, across time points, through the f_e node and the suppressor Ω_e node. The conditional beliefs $k(\Omega_{e_{t+1}}|\Omega_{e_t})$ and $\kappa(\vec{f}_{e_{t+1}}|\vec{e_t})$ will formally determine the strength of persistence across time. Both conditional beliefs will encode a bias against a change of state, which captures the intuition that any change of state must be causally motivated. For the case of the suppressors, the strength of this assumption is proportional to the strength of the change in the state of the suppressor. The more unlikely this state, the more unlikely the change:

$$\kappa(\omega^i_{e_{t+1}}|\omega^j_{e_t}) = |j - i|. \quad (5)$$

This model of persistence is not only intuitive, but guarantees that suppressors maintain their prior degree of disbelief. The assumption of persistence for the f_e node responds to the following equation:

$$\kappa(\vec{f}_{e_{t+1}}|\vec{e_t}) = \begin{cases} p, & \text{if } \vec{f}_{e_{t+1}} \neq \vec{e_t}; \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Since $\mathbf{F}(e_t)$ determines the state of e_t when suppressors are active, the number p can be interpreted as the degree of surprise in a non-causal change of state of the proposition e_t.[8] We remark that not all the nodes in the network must be assumed to persist. Nodes can be divided (by the user) into *events*, and *facts*. Facts will be assumed to persist over time, and will therefore be temporally expanded using the persistence model above. Events on the other hand, will be assumed not to persist over time. Thus, for example, *fired_gun* in Figure 4 will be assumed to be an event while *alive* will be assumed to be a fact.

Example: Belief persistence and belief change. Consider the network in Figure 3, which encodes the belief that, by default, the grass is wet if and only if it is raining. Assume that initially both $rain^+$ and $rain^-$ are plausible. Observing at t_0 that the grass is not wet will yield the belief that it is not raining at t_0. Furthermore, these beliefs will persist over t_1 and t_2. If it is observed at t_1 that $rain^+$, then the belief in *wet_grass* will change to reflect the change in the state of *rain*. The change on *wet_grass*, however, will depend on the tension between the persistence of the state of *wet_grass* at t_0, and the causal influence of *rain* at t_1. If the belief on the causal relation is stronger than that of persistence, belief in *wet_grass* will flip from wet_grass^- to wet_grass^+ at t_1. If on the other hand, the belief on the causal relation is weaker than that of persistence, the belief on *wet_grass* will change to uncommitted (both wet_grass^+ and wet_grass^- are plausible).

Example: Persistence of the suppressor. Assume that we had a prior belief that $rain^+$. Then, the observation at t_0 that the grass is not wet will force the conclusion that the suppressor is active at t_0 (it is raining yet the grass is not wet). Since suppressors persist (there is no reason why an exceptional condition such as *grass_covered* shouldn't) the observation $rain^+$ at t_1 will not change the belief in *wet_grass*.

The suppressor model, coupled with the temporal expansion provides a powerful and rich formalism for reasoning about time under uncertainty. An in depth analysis is beyond the scope of this paper and will be provided elsewhere [1]. It is instructive to compare its behavior with the behavior of a more simplistic model of persistence proposed by Goldszmidt and Pearl in [13], and by Pearl [21]. This model of persistence will perform the temporal expansion directly on the original network,

[7] Time is assumed to be discrete, and furthermore, the intervals between one time point and the next are assumed to be big enough to accommodate the causal interactions in each network.

[8] This value does not need to be constant, although it will assumed to be so in the remainder of the paper.

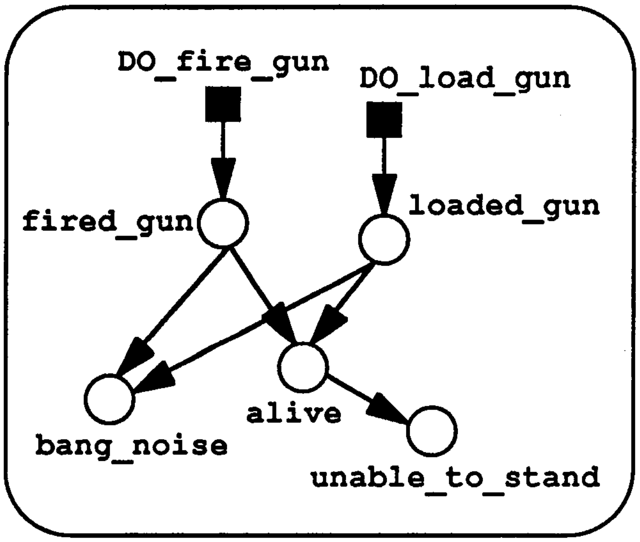

Figure 4: Causal family containing actions and ramifications (YSP).

without the use of the suppressor model or any other prior functional expansion. In this model, which we call "simple persistence" given the observation that at t_0 both $rain^+$ and wet_grass^+, the strength of belief in *rain* and *wet_grass* at T_1 will be determined simply by their causal relation. In contrast, in the suppressor model, the belief at T_1 in both propositions keeping their state will be determined be the strength of the persistence.

2.5 Action and Preconditions

For the representation of actions, we essentially follow the proposal in [13], which treats actions as external direct interventions that deterministically set the value of a given proposition in the domain. The semantic realization of this proposal cannot be the same as Bayesian conditioning. Intuitively, there is a difference between actions and observation. Consider the causal domain in the network of Figure 4. This network encodes an extension of the well known Yale Shooting Problem (YSP) [15]. It establishes that the consequences of shooting with a loaded gun are to produce a *bang_noise* and it also causes a victim to die. The state of the victim, in turn, will have a direct causal influence on the victim's ability to stand up. There is certainly a difference between hearing a loud noise (observation) and producing a loud noise (action). The first case constitutes evidence for a loaded gun being fired (and for the state of the unfortunate victim), while the second should not affect our beliefs about either of these propositions.[9]

The semantics of direct intervention for actions simulates the effect of directly setting the value of a node e. Actions are specified by indicating which nodes in the causal network are controllable and under what preconditions. Syntactically, we introduce a new arc $\blacksquare\!\!\longrightarrow$ for the representation of actions in the network. In Figure 4 for example, both *fired_gun* and *loaded_gun* are controllable propositions.[10] The advantage of using this proposal as opposed to others, such as STRIPS, is that the proposal of direct intervention takes advantage of the network representation capabilities for dealing with the indirect ramifications of actions and the related frame problem. In specifying the action shooting, for example, the user need not worry about how this action will affect the state of other related consequences such as *bang_noise* or *alive*.

Let the proposition do_e take values in $\vec{do_e} = \{\vec{e},\ idle\}$.[11] The new parent set of e after e is declared as controllable will be $\pi'(e) = \pi(e) \cup \{do_e\}$. The new ranking $\kappa'_{\vec{do_e}}(\vec{e}|\vec{\pi}(e))$ after e is set to $\vec{e}$ by the action is

$$\kappa'(\vec{e}|\vec{\pi}(e)) = \begin{cases} \kappa(\vec{e}|\vec{\pi}(e)) & \text{if } \vec{do_e} = idle \\ \infty & \text{if } \vec{do_e} \neq \vec{e} \\ 0 & \text{if } \vec{do_e} = \vec{e} \end{cases} \quad (7)$$

The effect of performing action do_e is to transform κ into a new belief ranking, governed by the equation above, guaranteeing that none of the direct influences $\pi(e)$ of e are active.

Example. Consider the example in Figure 4.[12] The relevant piece of causal knowledge available is that if a victim is shot with a loaded gun, then she/he will die. There are two possible actions, shooting and loading/unloading the gun. It is also assumed that both loaded and alive persist over time. Given this information the implementation of action networks will expand the network in Figure 4 both functionally and temporally.

In the first scenario, we observe at t_0 that the individual is alive and that the gun is loaded, and that there is an action at t_2 of shooting. The model will yield that alive at t_2 will be false. The reasons for this conclusion are due to the conditional independences assumed in the causal network representation.[13] In the second scenario, it is observed that at t_2 alive is true (the victim actually survived the shooting). The model will then conclude that: first, the gun must have been unloaded prior to the shooting (although it is not capable of asserting when),

[9]This example is essentially isomorphic to the one in the philosophic literature involving a barometer. Observing the state of the barometer change provides evidence about the weather; yet setting the barometer to a certain reading should provide no information about the weather.

[10]For simplicity we omitted preconditions for these actions. A suitable precondition for both nodes can be *holding_gun*, which can be represented as just another direct cause of these nodes. The corresponding matrices will then be constructed to reflect the intuition that the action do_e will be effective only if the precondition is true; otherwise, the state of a node e is decided completely by the state of its natural causes (that is, excluding do_e and the preconditions of do_e).

[11]Thus if e is binary $\vec{do_e} = \{e^+, e^-,\ idle\}$.

[12]There are already many of solutions to the YSP. The purpose of this example is not to add another solution to this collection, but to illustrate the properties of the action model, and the behavior of exceptions and persistence in the model proposed.

[13]They are explained in depth in [20][Chapter 10] and [13].

and furthermore, the belief of an action leading to unloading the gun increases (proportional to the degree of persistence in loading).

3 Conclusions and Future Work

Action networks can be viewed as an extension to probabilistic causal networks across the dimensions of uncertainty, action, and time. In particular, action networks allow uncertainty to be specified at different levels of abstraction, permit the representation of actions with preconditions, and provide models of persistence that are necessary for reasoning about plans.

This paper has focused on the κ–calculus instantiation of action networks. Future work includes allowing other quantifications of uncertainty, such as probabilities, and arguments — the first steps in this direction are reported in [7]. We also plan to add notions of utility and preferences on outcomes and explore other functional expansions of causal networks that go beyond the suppressor model. Finally, we plan to explore the possibility of interfacing action networks to probabilistic planners such as Buridan [17].

All the features of action networks described in this paper, including the suppressor model expansion, the temporal expansion, and the specification of actions, are fully implemented on top of CNETS [4].[14] All the examples described in this paper were tested using this implementation.

Acknowledgments

We wish to thank P. Dagum for discussions on the nature and properties of the functional transformation of a causal network, and to C. Boutilier and J. Pearl for discussions on the representation of actions. D. Draper and D. Etherington provided useful comments.

This work was partially supported by ARPA contract F30602-91-C-0031, and by IR&D funds from Rockwell Science Center.

References

[1] P. Dagum, A. Darwiche, and M. Goldszmidt. Functional specification of Bayes networks. Technical memorandum, Rockwell International Science Center, Palo Alto, CA, 1994. Submitted.

[2] A. Darwiche. A symbolic generalization of probability theory. Ph.D. dissertation, Computer Science Dept., Stanford University, Palo Alto, CA, 1992.

[3] A. Darwiche. Argument calculus and networks. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, pages 420–421, Washington DC, 1993.

[4] A. Darwiche. CNETS: A computational environment for causal networks. Technical memorandum, Rockwell International, Palo Alto Labs., 1994.

[5] A. Darwiche and M. Ginsberg. A symbolic generalization of probability theory. In *Proceedings of the American Association for Artificial Intelligence Conference*, pages 622–627, San Jose, CA, 1992.

[14] CNETS is an experimental environment for representing and reasoning with generalized causal networks, which allow the quantification of uncertainty using probabilities, κ degrees of belief, and logical arguments [2].

[6] A. Darwiche and M. Goldszmidt. A framework for plan simulation and analysis. Technical memorandum, Rockwell International, Palo Alto Labs., 1993.

[7] A. Darwiche and J. Pearl. Symbolic causal networks for reasoning about actions and plans. In this volume.

[8] T. Dean, L. P. Kaebling, J. Kirman, and A. Nicholson. Planning with deadlines in stochastic domains. In *Proceedings of the American Association for Artificial Intelligence Conference*, pages 574–579, Washington D.C., 1993.

[9] T. Dean and M. Wellman. *Planning and Control.* Morgan–Kaufmann, San Mateo, CA, 1991.

[10] M. Druzdzel and H. Simon. Causality in Bayesian belief networks. In *Proceedings of the 9^{th} Conference in Uncertainty in Artificial Intelligence Conference*, pages 3–11, Washington D.C., 1993.

[11] M. Ginsberg. *Readings in Nonmonotonic Reasoning.* Morgan Kaufmann, San Mateo, CA, 1987.

[12] M. Goldszmidt. Qualitative probabilities: A normative framework for commonsense reasoning. Ph.D. dissertation, University of California Los Angeles, Cognitive Systems Lab. TR-190, Los Angeles, CA, 1992.

[13] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*, pages 661–672, Boston, 1992.

[14] M. Goldszmidt and J. Pearl. Reasoning with qualitative probabilities can be tractable. In *Proceedings of the 8^{th} Conference on Uncertainty in AI*, pages 112–120, Stanford, 1992.

[15] S. Hanks and D. McDermott. Non-monotonic logics and temporal projection. *Artificial Intelligence*, 33:379–412, 1987.

[16] K. Kanazawa and T. Dean. A model for projection and action. In *Proceedings of the International Joint Conference on AI*, pages 985–990, 1989.

[17] N. Kushmerick, S. Hanks, and D. Weld. An algorithm for probabilistic planning. TR-93-06-03, Department of Computer Science, University of Washington, 1993.

[18] J. Pearl. From imaging and stochastic control to a calculus of action. In this volume.

[19] J. Pearl. Embracing causality in formal reasoning. *Artificial Intelligence*, 35:259–271, 1988.

[20] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[21] J. Pearl. From conditional oughts to qualitative decision theory. In *Proceedings of the 9^{th} Conference on Uncertainty in AI*, pages 12–20, Washington D.C., 1993.

[22] J. Pearl and T. Verma. A theory fo inferred causation. In *Proceedings of Principles of Knowledge Representation and Reasoning*, pages 441–452, Cambridge, MA, 1991.