Capturing Context in Causal Propagation

Mikhail Prokopenko[1], Maurice Pagnucco[2], Pavlos Peppas[3], Abhaya Nayak[2]

[1] CSIRO Mathematical and Information Sciences, Australia
[2] Computational Reasoning Group, Department of Computing, Macquarie University, Australia
[3] Department of Business Administration, Patras University, Greece

Abstract

Conciseness and intuitive representation are two primary objectives motivating the work on semantics for various logics of action. In addition, a unifying semantic framework for different reasoning approaches provides an ideal tool to compare these competing alternatives. It has been shown recently that a pure preferential semantics alone is not capable of providing such a unifying framework for a class of action theories dealing with causality. On the other hand, variants of preferential semantics augmented by additional structures on the state space have been successfully used to characterise some influential approaches to causal reasoning about action. The primary aim of this paper is to provide a general unifying semantics for Sandewall's causal propagation semantics [Sandewall, 1996] and McCain-Turner causal fixed-points approach [McCain and Turner, 1995].

1 Introduction

One particular trait emerging in recent literature on Reasoning about Action attempts to explicitly embody a notion of causality in logic-based action theories. It is often argued that such an extension would not only enhance the expressive power of theories of action, but may also provide more concise representations. What seems to be lacking so far is a general semantic framework that covers this particular class of action theories. It appears that a pure preferential semantics alone is not capable of providing such a unifying framework.

For example, McCain and Turner's causal theory of action [McCain and Turner, 1995] was recently characterised by an *augmented preferential semantics*, using an appropriately constructed binary relation on states in addition to a preference relation [Peppas *et al.*, 1999]. This additional relation captured causal context of action systems by translating individual causal laws into (causality-driven) state transitions. On the other hand, another rather general semantical approach — the *causal propagation semantics* proposed by Sandewall [1996] — deals with causal ramifications without explicitly relying on the principle of minimal change.

Recently, another causal theory of action — that of Thielscher [1997] — has been characterised by a variant of the augmented preferential semantics [Prokopenko *et al.*, 1999]. This time the minimality component was complemented by a binary relation on states of higher dimension. The standard state-space of possible worlds was extended to a hyper-space, and action effects (including indirect ones) were traced in the hyper-space. Again, the purpose of these hyper-states was to supply extra context to the process of causal propagation. The hyper-space semantics [Prokopenko *et al.*, 1999] can be clearly seen to employ a component of minimal change coupled with causality. Moreover, another variant of the augmented preferential semantics capturing Thielscher's approach — the power-space semantics — has been unified with the Sandewall's causal propagation semantics under certain conditions [Prokopenko *et al.*, 2000].

This work introduces a preferential style semantics augmented with a causal transition relation on states, that is general enough to unify three mentioned frameworks to reasoning about action and causality — Sandewall's causal propagation semantics [Sandewall, 1996], Thielscher's causal relationships approach [Thielscher, 1997], and McCain-Turner causal fixed-points approach [McCain and Turner, 1995]. This is achieved by observing that the principle of minimal change is hidden behind action invocation and causal propagation in all proposals.

In particular, we focus on conditions and approximations required for the unification of the causal propagation semantics and the causal fixed-points approach. The variant of the augmented preferential semantics needed for capturing Thielscher's causal relationships approach is described elsewhere [Prokopenko *et al.*, 2000]. The main contribution of this paper is identification of a specific semantical component — a family of choice functions — used in scaling *context-sensitive causal propagation* employed by our motivating approaches.

In the following section we briefly sketch Sandewall's proposal and McCain-Turner's causal theory of action. In Section 3 we describe a general augmented preferential semantics. Then, in Section 4, we investigate conditions required to characterise the causal propagation semantics and causal fixed-points within given generalisation. In particular, sections 4.2 and 4.4 establish the representation theorems. Section 5 discusses the importance of these results.

2 Technical Background

2.1 Causal Propagation Semantics

The causal propagation semantics introduced by Sandewall [1996], uses the following basic concepts. The set of possible states of the world, formed as a Cartesian product of the finite sets of a finite number of state variables, is denoted as $\mathcal{R}$. E is the set of possible actions. The causal propagation semantics extends a basic state transition semantics with a *causal transition relation*. The causal transition relation C is a non-reflexive relation on states in $\mathcal{R}$. A state r is called *stable* if it does not have any successor s such that $C(r, s)$. Another component, $\mathcal{R}_c$, is a set of admitted states chosen as a subset the set of stable states.

Another important concept, introduced by Sandewall, is an *action invocation* relation $G(e, r, r')$, where $e \in E$ is an action, r is the state where the action e is invoked, and r' is "the new state where the instrumental part of the action has been executed" [Sandewall, 1996]. In other words, the state r' satisfies direct effects of the action e. It is required that every action is always invokable, that is, for every $e \in E$ and $r \in \mathcal{R}$ there must be at least one r' such that $G(e, r, r')$ holds.

A finite (the infinite case is omitted) transition chain for a state $w \in \mathcal{R}_c$ and an action $e \in E$ is a finite sequence of states $r_1, r_2, ..., (r_k)$, where $G(e, w, r_1), C(r_i, r_{i+1})$ for every $i, 1 \leq i < k$, and where r_k is a stable state The last element of a finite transition chain is called a result state of action e performed in state w.

These basic concepts define an *action system* as a tuple $\langle \mathcal{R}, E, C, \mathcal{R}_c, G \rangle$. The following definition strengthens action systems based on the causal propagation semantics.

Definition 2.1 *If three states w, p, q are given, we say that the pair p, q respects w, denoted as $\lhd_w(p, q)$, if and only if $p(f) \neq q(f) \rightarrow p(f) = w(f)$ for every state variable f that is defined in $\mathcal{R}$, where $r(f)$ is a valuation of variable f in state r.*

An action system $\langle \mathcal{R}, E, C, \mathcal{R}_c, G \rangle$ is called respectful if and only if, for every $w \in \mathcal{R}_c$, every $e \in E$, w is respected by every pair r_i, r_{i+1} in every transition chain for the state w, and the last element of the chain is a member of $\mathcal{R}_c$.

According to Sandewall [1996], respectful action systems are intended to ensure that in each transition there cannot be changes in state variables which have changed previously upon invocation or in the causal propagation sequence. This requirement, of course, guarantees that a resultant state is always consistent with the direct effects of the action (which cannot be cancelled by indirect ones), and that there are no cycles in transition chains.

As with many other state transition action systems, the intention is to characterise a result state in terms of an initial state w and action e, without "referring explicitly to the details of the intermediate states" [Sandewall, 1996]. In other words, it is desirable to define a selection function $Res(w, a)$. For a respectful action system $\langle \mathcal{R}, E, C, \mathcal{R}_c, G \rangle$, the selection function can be given as

$$Res_{C\mathcal{R}_cG}(w, e) = \{r_k \in \mathcal{R}_c : \\ G(w, e, r_1),\ \ C(r_i, r_{i+1}),\ \ \lhd_w(r_i, r_{i+1}),\ \ 1 \leq i < k\}.$$

2.2 Causal Systems with Fixed-Points

In this section we sketch McCain and Turner's [1995] causal theory of actions. We shall be working with a fixed finitary propositional language $\mathcal{B}$ whose propositional letters we call *fluents*. A *literal* is a fluent or the negation of a fluent. The set of all fluents is denoted by $\mathcal{F}$, and the set of all literals by $L_{\mathcal{F}}$. A *state* is defined as a maximal consistent set of literals. By $[\phi]$ we denote all states consistent with the sentence $\phi \in \mathcal{B}$ (i.e., $[\phi] = \{w \in W : w \vdash \phi\}$). Domain constraints are sentences which have to be satisfied in all states.

McCain and Turner introduce a new connective $\Rightarrow$ to denote a causal relationship between sentences ϕ and ψ of the underlying language $\mathcal{B}$. This allows for expressions of the form $\phi \Rightarrow \psi$ (where $\phi,\ \psi \in \mathcal{B}$) which are termed *causal laws* (or *casual rules*). Nesting of $\Rightarrow$ is not permitted. A set of causal laws $\mathcal{Q}$ is referred to as a *causal system*. Given any set of sentences $\Lambda \subseteq \mathcal{B}$ and a causal system $\mathcal{Q}$, the (causal) *closure* of Λ in $\mathcal{Q}$ is denoted $C_{\mathcal{Q}}(\Lambda)$ and defined to be the smallest superset of Λ closed under classical logical consequence such that for any $\phi \Rightarrow \psi \in \mathcal{Q}$, if $\phi \in C_{\mathcal{Q}}(\Lambda)$, then $\psi \in C_{\mathcal{Q}}(\Lambda)$. We also say that Λ *causally implies* ϕ with respect to $\mathcal{Q}$ if and only if $\phi \in C_{\mathcal{Q}}(\Lambda)$ and denote this as $\Lambda \mid\!\sim_{\mathcal{Q}} \phi$.

Another important notion is that of a *legitimate state* with respect to a causal system $\mathcal{Q}$. Any state r is legitimate with respect to $\mathcal{Q}$ if and only if $r = C_{\mathcal{Q}}(r) \cap L_{\mathcal{F}}$. That is, a state is legitimate if and only if it does not contravene any causal laws of $\mathcal{Q}$.

McCain and Turner's aim is to determine the set of possible resultant (or successor) states $Res_{\mathcal{Q}}(w, e)$ given an initial state w and the direct effects (or post-conditions) of an action represented by the sentence e. We shall refer to actions only through their direct effects as they play no direct role in McCain and Turner's framework. Formally speaking, we have for any causal system $\mathcal{Q}$ a function $Res_{\mathcal{Q}}$ mapping a legitimate (initial) state w and sentence e (direct effects) to the set of states $Res_{\mathcal{Q}}(w, e)$ according to the definition [McCain and Turner, 1995]:

$$r \in Res_{\mathcal{Q}}(w, e) \quad \text{if and only if}$$

$$r = \{p \in L_{\mathcal{F}} :\ (w \cap r) \cup \{e\} \mid\!\sim_{\mathcal{Q}} p\}$$

We often refer to the elements of $Res_{\mathcal{Q}}(w, e)$ as *causal fixed-points*. Intuitively, a causal fixed-point is an outcome incorporating direct action effects, where all other changed fluents (properties) are causally justified (in a certain sense). In other words, every detail in the outcome must be "explained" either as persisting through the action, or as a direct effect, or as a causal ramification of *other properties contained in the outcome* — hence the fixed-point flavour. Obviously, a possible outcome that contains at least one detail without such justifications is rejected, even if it does not violate domain constraints.

Note that it follows from this definition that if $r \in Res_{\mathcal{Q}}(w, e)$, then $r \in [e]$ (i.e., r must satisfy the direct effects of the action). Intuitively speaking, the elements of $Res_{\mathcal{Q}}(w, e)$ are simply those e-states where all changes with respect to w can be justified by the underlying causal system.

3 Augmented Preferential Semantics

3.1 Minimal Change vs Causal Change

The general augmented preferential semantics presented in [Prokopenko *et al.*, 2000] included a concept of the power-state space needed to faithfully trace ramifications produced by Thielscher's causal relationship approach. Components of this system can be viewed as a tuple

$$\mathcal{H} = \langle \mathcal{W},\ \mathcal{D},\ \mathcal{E}, \Gamma,\ \mathcal{O},\ \mathcal{M}, \mathcal{P}, Res \rangle$$

where $\mathcal{W}$ is a set of states; $\mathcal{D}$ a set of legitimate (admitted) states; $\mathcal{E}$ a set of actions; Γ a set of power-states; $\mathcal{O}$ a set of (non-strict) orderings $<_\gamma$, each with respect to some power-state $\gamma \in \Gamma$; $\mathcal{M}$ a binary relation on power-states; $\mathcal{P}$ a projection function from Γ to $\mathcal{W}$; and Res a selection function defined on $\mathcal{W} \times \mathcal{E}$ and responsible for producing successor states in $\mathcal{W}$.

A particularly simple variant of this semantics is obtained under the approximation $\mathcal{W} = \Gamma$ and the trivial projection function $\mathcal{P}(\iota) = \iota$. In this case, the power-states (playing the role of meta states) are not needed, and both the preferential structure $\mathcal{O}$ and the binary relation $\mathcal{M}$ are defined on normal states $\mathcal{W}$. This variant was used in unifying the proposed semantics with Sandewall's causal propagation semantics under certain conditions [Prokopenko *et al.*, 2000] .

In this paper we take this simplified variant of the general augmented preferential semantics as a "baseline", and enhance it further with a component responsible for scaling context-sensitive propagation observed in McCain and Turner theory of actions [Peppas *et al.*, 1999] (described in detail in section 4.3). Components of this system can be viewed as a tuple

$$\mathcal{H}_\Sigma = \langle \mathcal{W},\ \mathcal{D},\ \mathcal{E},\ \mathcal{O},\ \mathcal{M}, \Sigma, Res \rangle$$

where the new component Σ is a family of choice functions defined on $\mathcal{W} \times \mathcal{E} \times \mathcal{W}$.

The principle of minimal change can be identified with $\mathcal{O}$ and that of causality with $\mathcal{M}$. In other words, both principles are represented and play a clear and distinct role. They are not inter-reducible but go hand-in-hand to capture causal action theories. We maintain that both principles are required to solve the frame and ramification problems in a *concise* fashion.

Let us describe informally how a successor state can be obtained, given an initial state and an action, if we wish to stay within simple and clear state transition semantics. Intuitively, a successor state is an admitted state which is reachable (by means of some transition relation) from states nearest (to the initial one) among states satisfying post-conditions of the performed action. In other words, tracing all action effects (both direct and indirect) may involve two steps. First, we find states satisfying the action's post-conditions while staying as close as possible to the initial state according to some preference relation. Then, we propagate along the transition relation from all such minimal states all the way to some stable (and admitted) state — our result state.

However, sometimes there is a causal context present in the domain that may require additional checks and balances, potentially obscuring this simple and clear view.

A starting point of propagation, for example, can be associated with a group of states that capture a direct effect (an immediate causal context) of an action. This can be done using an ordering $<_w \in \mathcal{O}$ and a choice function $\sigma \in \Sigma$, providing necessary context to the process of causal propagation. An appropriately constructed transition relation $\mathcal{M}$ on states would then allow us to propagate in a very simple way, resulting in a clearly defined "final" set of successor states.

Let $[e]$ denote a set of states satisfying the post-conditions of an action e. We also define a set $min(<_w, [e])$ as a subset of $[e]$ containing states nearest to the state w in terms of the ordering $<_w$, in other words,

$$min(<_w, [e]) = \{\beta \in [e], \neg\exists\alpha \in [e], \alpha \neq \beta, \alpha <_w \beta\}.$$

Sometimes, we will refer to an element of $min(<_w, [e])$ as a state $<_w$-minimal in $[e]$.

Let us denote by $\mathcal{K}_\mathcal{M}$ the set of stable states

$$\{p \in \mathcal{W} : \neg\exists q \in \mathcal{W}, \mathcal{M}(p, q)\},$$

an obvious counterpart of the set of stable states in the causal propagation semantics. In spirit of the latter, we require that

$$\mathcal{D} \subseteq \mathcal{K}_\mathcal{M}.$$

In other words, some domain constraints may eliminate more illegitimate states than causal propagation alone. Arguably, any state which is not admitted, should be excluded by some causal laws—and the set $\mathcal{D}$ can be *made* equal to the set $\mathcal{K}_\mathcal{M}$ of stable states. However, the distinction between admitted and (causally-)stable states may be useful in providing some flexibility to domain constraints. For example, stable states may reflect hard constraints, while admitted states may correspond to soft constraints. Varying selection requirements on successor states (stable and admitted, or merely stable) would allow us to capture an additional degree of state eligibility if required.

3.2 Selection Function and Gradient Choice Functions

Before specifying a selection function Res based on our augmented preferential semantics, we need to define a few more constructs derived from elements of action system $\mathcal{H}_\Sigma$.

Firstly, we define *e-predecessors* of a given state — a set of states preceding the given state with respect to an ordering from $\mathcal{O}$. Formally:

Definition 3.1 *Given any two states γ, β and any action e, the set of e-predecessors of β with respect to γ is defined to be the set*

$$\langle\!\langle \beta, e \rangle\!\rangle_\gamma = \{\alpha :\ \alpha \in [e] \ \ \textit{and} \ \ \alpha <_\gamma \beta\}.$$

The e-predecessors of β with respect to γ are just the $[e]$ states which are closer to γ than β. If the ordering $<_\gamma$ is reflexive, then any $\beta \in [e]$ is an e-predecessor of itself with respect to γ (in other words, $\beta \in \langle\!\langle \beta, e \rangle\!\rangle_\gamma$).

It is interesting at this stage to compare the states $<_\gamma$-minimal in $[e]$, and the e-predecessors of some state β with respect to γ. Intuitively, the $<_\gamma$-minimal states in $[e]$ compose a boundary separating the set $[e]$ from the initial state γ. For example, as we shall show, any $<_\gamma$-minimal state may be used to form a starting point of propagation in Sandewall's approach. In other words, the boundary $min(<_\gamma, [e])$ is the earliest "zone", where we may start causal propagation. In some sense, the states on the boundary "support" the propagation, and can be thought of as "collaborators" in producing successor state(s). On the other hand, the e-predecessors in $\langle\!|\beta, e|\!\rangle_\gamma$ "compete" against each other because any of them could be eventually selected as a successor state. For example, in the McCain-Turner approach, we will request that all predecessors of a causal fixed-point r have to be eliminated during the state transition process leading to r — so, in a sense, the e-predecessors in $\langle\!|\beta, e|\!\rangle_\gamma$ frame a potentially final point (or a horizon) of propagation. Intuitively, the sets $min(<_\gamma, e)$ and $\langle\!|\beta, e|\!\rangle_\gamma$ represent the minimality-driven component of an action system, and indicate a possible "gradient" of causality-driven propagation. By definition, some e-predecessors of a state β with respect to γ, lie on the boundary $min(<_w, [e])$:

$$min(<_w, e) \cap \langle\!|\beta, e|\!\rangle_\gamma \neq \emptyset.$$

To accommodate seemingly different selection functions discussed in section 2, we will need to employ different choice functions from Σ aimed at identifying (narrowing down) the minimality-driven "gradient" of state transitions. A choice function is defined for a potential successor state $\beta \in \mathcal{W}$, an action $e \in \mathcal{E}$, and a state $\gamma \in \mathcal{W}$. We shall call states chosen by a choice function $\sigma(\beta, e, \gamma)$, the σ-chosen states in the (β, e, γ)-gradient.

Our first choice function $\sigma_F(\beta, e, \gamma)$, called a "full-meet gradient", merely returns all e-predecessors of some state β with respect to γ:

$$\sigma_F(\beta, e, \gamma) = \langle\!|\beta, e|\!\rangle_\gamma$$

Second choice function $\sigma_M(\beta, e, \gamma)$, called a "mini-choice gradient", chooses one $<_w$-minimal state in $[e]$, or in other words, an element of $min(<_w, [e])$:

$$\sigma_M(\beta, e, \gamma) = \{\alpha\}, \text{ where } \alpha \in min(<_\gamma, [e])$$

The "full-meet gradient" defines a set $\sigma_F(\beta, e, \gamma)$ that identifies potential challengers for a successor state place. Intuitively, the states in this set compete to become a successor state and the state β is being tested as a potential winner.

The "mini-choice gradient", on the contrary, does not test any potential winner or its challengers (in fact, it is independent of first argument). Instead, this choice function attempts to make subsequent causal propagation as much flexible as possible — by picking just one of the $<_\gamma$-minimal states in $[e]$. Intuitively, the "mini-choice gradient" supports all possibilities in terms of intermediate states consistent with direct action effects.

We introduce two more choice functions. One, $\sigma_P(\beta, e, \gamma)$, is called a "partial-choice gradient"[1], and chooses a $<_\gamma$-minimal state in $[e] \cap \mathcal{D}$ or in other words, a minimal element among legitimate states consistent with direct effects of the action e:

$$\sigma_P(\beta, e, \gamma) = \{\alpha\}, \text{ where } \alpha \in min(<_\gamma, [e] \cap \mathcal{D})$$

[1]It can be easily guessed that names of some choice functions that we use are loosely based on well-known *full-meet*, *partial meet* and *maxi-choice* belief revision functions [Gärdenfors, 1988].

It is important to point out that the "partial-choice gradient" $\sigma_P(\beta, e, \gamma)$ is weaker than a more demanding choice function

$$\sigma_L(\beta, e, \gamma) = \{\alpha\}, \text{ where } \alpha \in min(<_\gamma, [e]) \cap \mathcal{D}$$

which could be referred to as "legitimate-choice gradient".

Intuitively, the "legitimate-choice gradient" chooses such a $<_\gamma$-minimal state in $[e]$ that is, in addition, legitimate (admitted) with respect to $\mathcal{D}$. It may result, sometimes, in an empty set — when all states on the boundary $min(<_\gamma, [e])$ are not legitimate with respect to $\mathcal{D}$. On the contrary, the "partial-choice gradient" function focuses directly on states in $min(<_\gamma, [e] \cap \mathcal{D})$. Clearly,

$$min(<_\gamma, [e]) \cap \mathcal{D} \;\subseteq\; min(<_\gamma, [e] \cap \mathcal{D})$$

The set on the right-hand side of the containment can be thought of as another state-space boundary, "located" further from the initial point γ than the boundary $min(<_\gamma, [e])$, and the "partial-choice gradient" function targets these states.

It is worth pointing out again that all the gradient choice functions narrow down the search and selection of potential (context-sensitive) successor states in the state-space, based on the preference relation $<_\gamma$, and without the causality component $\mathcal{M}$ of an action system. In other words, the gradient choice functions scale potential causal propagation towards context-sensitive ramifications.

Now we are ready to define our selection function. Let $\mathcal{M}^*$ be a transitive closure of the relation $\mathcal{M}$. We shall say that a state β is $\mathcal{M}$-reachable from a state α, if $\mathcal{M}^*(\alpha, \beta)$.

We say that an admitted (and therefore, stable) state r is a successor state, $r \in Res(w, e)$, if and only if $r \in [e]$ and is $\mathcal{M}$-reachable from *all* σ-chosen states in the (r, e, w)-gradient. More precisely, a selection function of the action system $\mathcal{H}_\Sigma$ is given as

$$Res(w, e) = \{r \in \mathcal{D} \cap [e] : \forall\, \alpha \in \sigma(r, e, w),\;\; \mathcal{M}^*(\alpha, r)\}.$$

The notions of "stable" and "reachable" states are understood in terms of causal transition relation $\mathcal{M}$, and the gradient is given by one of the choice functions, using an ordering $<_w$ from $\mathcal{O}$.

There is an important sub-class of action systems, where the selection function can be modified; more precisely, strengthen as follows:

$$Res(w, e) = \{r \in \mathcal{D} \cap [e] : \forall\, \alpha \in \sigma(r, e, w),\;\; \mathcal{M}^*(\alpha, r), \text{ and there is a Hamiltonian path through states in } \sigma(r, e, w)\}.$$

In other words, not only a stable successor state should be $\mathcal{M}$-reachable from *all* states lying on the chosen gradient, but there must be a Hamiltonian path through these states. When a successor state is itself in the gradient (for example, "full-meet gradient"), the latter requirement (a Hamiltonian path) is clearly stronger than the former ($\mathcal{M}$-reachable state) — the condition that a successor state should be stable ensures that such a Hamiltonian path ends in this state which is, therefore, reachable from all gradient states. We shall refer to such action systems as *Hamiltonian action systems*. A Hamiltonian action system will be required to show selection-equivalence with McCain-Turner action systems.

Following [Peppas *et al.*, 1999], we shall say that an action system incorporating a function Res_1 is *selection-equivalent* to an action system with a function Res_2 if and only if

$$Res_1(w, e) = Res_2(w, e),$$

for every action e and state w. Now, our main goal becomes clear: we intend to identify under what conditions it is possible to achieve a selection-equivalence between a generalised action system and action systems based on our motivating frameworks. More precisely, we wish to identify conditions when

$$Res(w, e) = Res_{C\mathcal{R}_c G}(w, e),$$
$$Res(w, e) = Res_{\mathcal{Q}}(w, e).$$

Before we demonstrate how the desired selection-equivalence can be achieved for our motivating approaches, we consider a most obvious simplification of action systems, captured by the augmented preferential semantics $\mathcal{H}_\Sigma$. Setting $\mathcal{M} = \emptyset$ and taking the "partial-choice gradient" $\sigma_P(\beta, e, \gamma)$ as our choice function, we may produce a traditional preferential semantics with a variety of suitable preference relations $\mathcal{O}$. More precisely, the selection function is given then as

$$Res_P(w, e) = \{r \in min(<_w, [e] \cap \mathcal{D})\}.$$

Obviously, taking a more demanding "legitimate-choice gradient" $\sigma_L(\beta, e, \gamma)$ as our choice function, produces a preferential semantics tending to disqualify more successor states:

$$Res_L(w, e) = \{r \in \mathcal{D} \cap min(<_w, [e])\}.$$

4 Representation Results

4.1 Invoking Minimal Change

In this section, we intend to describe under what conditions it is possible to represent the Sandewall's causal propagation semantics as an instance of the augmented preferential semantics $\mathcal{H}_\Sigma$. Following [Prokopenko *et al.*, 2000] [2], this reduction will be carried out while staying within the same sets of states and actions ($\mathcal{W} = \mathcal{R}$, $\mathcal{D} = \mathcal{R}_c$, and $\mathcal{E} = E$) and transforming the causal transition relation C into the binary relation $\mathcal{M}$. In other words, our primary focus will be discovering and capturing the nature of minimality subsumed by the invocation relation G.

Motivated by a preferential-style semantics, one may be tempted to suggest an ordering on states such that the invocation relation G can be simply realised by selecting nearest states satisfying action post-condition. However, this does not appear to be possible without restricting the relation G.

Lemma 4.1 *There is no ordering $<_w$ such that for every action e and state r, $G(e, w, r)$ if and only if $r \in min(<_w, e)$.*

Consequently, our intention at this stage is to restrict the invocation relation G in such a way that, given an initial state and an action, the invoked states can be characterised precisely as states nearest to the initial one in terms of some appropriate minimality ordering. Before we identify required restrictions on the invocation relation G, we introduce some more abbreviations. If $V \subseteq [a]$, for a set of states V and an action $a \in E$, we call the action a a *V-covering* action. Furthermore, if there exists a V-covering action a such that $G(a, w, x)$ for states $w, x \in \mathcal{R}$, we say that the state x is *V-cover accessible* from state w. Also, we say that a state x is *not V-cover accessible* from state w, if there is no V-covering actions a such that $G(a, w, x)$.

Importantly, it follows that all states in a set V satisfy post-conditions of a V-covering action. It is worth pointing out that, given two states p and q satisfying post-conditions of some action a (that is, the action a is a $\{p, q\}$-covering action), the state p may be $\{p, q\}$-cover accessible from some state w, while state q is not $\{p, q\}$-cover accessible from w. This is the case when the following two conditions are satisfied:

$$\exists\, a \in E,\ p, q \in [a],\ G(a, w, p),$$
$$\forall\, e \in E,\ p, q \in [e],\ \neg G(e, w, q).$$

The first restriction on invocation relation is given as

(G_1)

if p is $\{p, q\}$-cover accessible from w but q is not, and
q is $\{q, x\}$-cover accessible from w but x is not
then p is $\{p, x\}$-cover accessible from w and x is not,
for arbitrary states w, p, q, x.

The premise of the implication is that, considering all actions whose post-conditions are satisfied by two states p and q, state p is chosen at least once by the invocation relation and state q is never chosen; and considering all actions whose post-conditions are satisfied by two states q and x, state q is chosen at least once, while state x is never chosen. This then necessitates that, considering all actions whose post-conditions are satisfied by states p and x, invocation of the state p must eventualise at least once, but state x cannot be invoked at all. The condition (G_1) is intended to capture transitivity of a corresponding preferential relation. Another condition is given as

(G_2) Given any two $\{p, q\}$-covering actions e' and e'', if $G(e', w, p)$ and $G(e'', w, q)$ then $G(e', w, q)$.

This condition simply requires that if neither of two states p and q is chosen over the other in terms of invocation in one instance, then selection of either of them necessitates selection of the other in another instance as well. Finally, we reinforce the requirement that any action is invokable in principle.

(G_3) $\forall\, e \in E,\ w \in \mathcal{R},\ \exists\, p \in [e],\ G(e, w, p)$

As noted above, this condition does not guarantee that the invoked action will succeed—it may possibly be qualified by causal propagation ending in a non-admitted state. Now, we are ready to describe a set of orderings $\mathcal{O}$ corresponding to the invocation relation. Ideally, any ordering $<_w$ should satisfy only the transitivity property:

(M_1) if $p <_w q$ and $q <_w x$ then $p <_w x$.

However, it turns out that, given an action system, the related ordering has to satisfy, in addition, two other properties:

[2] For convenience, we partially reproduce here results reported in [Prokopenko *et al.*, 2000].

(M_2)

if p is $<_w$-minimal in $[a]$ for some $\{p,q\}$-covering action a

and

q is not $<_w$-minimal in $[e]$ for any $\{p,q\}$-covering action e

then $p <_w q$.

(M_3) if $p <_w q$ then

p is $<_w$-minimal in $[e]$ for some $\{p,q\}$-covering action e.

Basically, the second property (M_2) requires that any state p which is $<_w$-minimal in some set $[a]$ is preferred to any state q, where q belongs to the set $[a]$ as well, and which is not $<_w$-minimal in any set $[e]$ of post-conditions satisfied by both p and q.

The third property (M_3) posits that if a state p is preferred to a state q by a preference relation $<_w$, then there must exist an action e, whose post-conditions are satisfied by these two states, such that state p is $<_w$-minimal in $[e]$.

We intend to show at this stage that there is a way to define the invocation relation in terms of a preference relation and vice versa, while preserving respective selections of states, satisfying direct action effects. The following two definitions can be shown to ensure such an equivalence.

Definition 4.2 *A new invocation relation $G_<$ is defined as follows: $G_<(e,w,r)$ if and only if r is $<_w$-minimal in $[e]$, where $w,r \in \mathcal{R}, e \in E$.*

Put simply, the new relation $G_<(e,w,r)$ specifies states r that are nearest among all states in $[e]$ to the initial state w, where the action e was invoked.

Definition 4.3 *Given an invocation relation G, for each $w \in \mathcal{R}$ we define an ordering $<_{w,G}$ on states in $\mathcal{R}$ as follows: $p <_{w,G} q$ if and only if state p is $\{p,q\}$-cover accessible from w and state q is not $\{p,q\}$-cover accessible from w.*

This definition specifies a preference relation on states driven by a given invocation relation—state p is nearer to an initial state w than state q if and only if for all actions whose direct effects are satisfied by both states p and q, the state q is never selected by the invocation relation G, while state p is selected at least once.

4.2 Representation Results for Causal Propagation Semantics

The following lemma establishes the sought-after equivalence between invocation and preference relations.

Lemma 4.4 *If the relation G satisfies the conditions $(G_1) - (G_3)$, then for each $w \in \mathcal{R}$, the ordering $<_{w,G}$ satisfies conditions $(M_1) - (M_3)$.*
If each ordering $<_w$ for $w \in \mathcal{R}$ satisfies conditions $(M_1) - (M_3)$, then the relation $G_<$ satisfies the conditions $(G_1) - (G_3)$.

It is interesting to observe, at this stage, that the respectfulness requirement in terms of states is related to the notion of minimality as well. Let us recall that a state x is preferred to a state y in terms of the PMA ordering [Winslett, 1988], denoted $x \prec_w y$, if and only if $Diff(x,w) \subseteq Diff(y,w)$, where $Diff(p,q)$ represents the symmetric difference of p and q, i.e., $(p \setminus q) \cup (q \setminus p)$. Formally, the following observation holds.

Lemma 4.5 *The pair p,q respects w, $\lhd_w(p,q)$, if and only if $p \prec_w q$ in the PMA ordering $\prec_w$ associated with w.*

Lemma 4.5 indicates that in a respectful action system causality propagates from closer states to states which are more distant from the initial one, in terms of the PMA ordering. In order to capture respectful action systems, we would need the following two conditions.

(M_4) if $p \prec_w q$ then $p <_w q$.

(M_5) For every action e and state w, the set $min(\prec_w, [e])$ is a singleton.

The condition (M_4) ensures that an ordering $<_w$ incorporates the PMA ordering, or, in other words, includes all pairs p,q such that $p \prec_w q$. The unique minimality assumption (M_5) relates to connectivity of the set $[e]$ in terms of the ordering $\prec_w$. In other words, it ensures that the minimal element of the set $[e]$ is preferred to any other element of $[e]$.

The conditions (M_4) and (M_5) are sufficient to ensure that propagation along a transition chain starts from an element of $min(\prec_w, [e])$, and links states that pair-wise respect the initial state.

Having established these necessary connections, it is easy to observe that action systems based on the causal propagation semantics and the augmented preferential semantics $\mathcal{H}_\Sigma$ are selection-equivalent under certain conditions. The selection-equivalence can be achieved without extending the state-space and action domain. Formally, $\mathcal{R} = \mathcal{W}$, $\mathcal{R}_c = \mathcal{D}$ and $E = \mathcal{E}$. This indicates that minimal change is present in the causal propagation semantics subsumed by the invocation relation.

To achieve the desired reduction, we take the "mini-choice gradient" $\sigma_M(r,e,w)$ as our choice function — the one that chooses a $<_w$-minimal state in $[e]$:

$$Res_M(w,e) = \{r \in \mathcal{D} \cap [e] : \alpha \in min(<_w, [e]), \ \mathcal{M}^*(\alpha, r)\}.$$

This allows us to unify the causal propagation semantics and our augmented preferential semantics.

Theorem 4.6 *For every respectful action system $\langle \mathcal{R}, E, C, \mathcal{R}_c, G \rangle$ there exists a selection-equivalent action system $\langle \mathcal{W}, \mathcal{D}, \mathcal{E}, \mathcal{O}, \mathcal{M}, \Sigma, Res \rangle$, if the relation G satisfies conditions $(G_1) - (G_3)$.*

Conversely, for every action system $\langle \mathcal{W}, \mathcal{D}, \mathcal{E}, \mathcal{O}, \mathcal{M}, \Sigma, Res \rangle$ there exists a selection-equivalent respectful action system $\langle \mathcal{R}, E, C, \mathcal{R}_c, G \rangle$, if the orderings in $\mathcal{O}$ satisfy conditions $(M_1) - (M_5)$.

Basically, this observation establishes that, under considered conditions, we obtain

$$Res_{C\mathcal{R}_cG}(w,e) = Res_M(w,e).$$

for any action e and state w.

4.3 Propagating Minimal Change towards a Fixed-Point

One particular state-selection mechanism based on state elimination that captures causal fixed-points was described in [Peppas *et al.*, 1999]. A state rejected by a state elimination rule is one which contravenes a causal relationship deemed to hold in the resultant state (in fact, in the causal system as a whole).

Definition 4.7 *(State elimination rule)*
A state elimination rule *(or simply,* elimination rule*) is an expression of the form*

$$\{r_1,\, r_2,\, \cdots,\, r_k,\, r_{k+1},\, \cdots,\, r_n\} \rhd \{r_1,\, r_2,\, \cdots,\, r_k\},$$

where each r_i *is a state.*

An elimination rule functions by rejecting certain states from among those currently considered possible. Suppose that according to an agent's current beliefs it considers the states that are possible to be among $\{r_1, \cdots, r_n\}$. An elimination rule like that in Definition 4.7 allows the agent to reject states $r_{k+1}, \cdots, r_n$.

The underlying idea is to use *state elimination rules* to discard some states in $[e]$ from further consideration, given that in McCain and Turner's [1995] causal theory $Res_{\mathcal{Q}}(e, w) \subseteq [e]$. At any point we are working with the set of states currently being entertained (a subset of $[e]$). We repeatedly apply elimination rules to this set of states to reject the illegitimate ones (those not possible) focusing on the possible resultant states. All elimination rules need to be applied until no further states can be rejected to ensure that all illegitimate states have been purged and only definite possibilities remain. To put it another way, a state elimination system acts as a *filtering* mechanism; illegitimate states are successively filtered out through use of elimination rules. The states that cannot be eliminated by any combination of such rules are called *compact*.

The reduction described in [Peppas *et al.*, 1999] produces an elimination rule $[\phi]\rhd[\phi \wedge \psi]$ for every causal law $\phi \Rightarrow \psi$ in $\mathcal{Q}$. It has also been shown in [Peppas *et al.*, 1999] that all PMA-ordered e-predecessors of a causal fixed-point are eliminated by a repeated application of elimination rules.

4.4 Representation Results for Causal Fixed-Points

The elimination rules can be re-arranged, without loss of generality, in such a way that a single rule eliminates only one state. Subsequently, it is possible to construct a binary relation on states $\mathcal{M}(r, p)$ which traces the process of states elimination.

Consider an arbitrary set of states V with cardinality $n + 1$. The string of elimination rules $\sigma_1; \sigma_2; \cdots; \sigma_n$ *dissolves* V if and only if after applying these rules successively (in the order given), all but one of the states of V are eliminated and, furthermore, the one remaining state is compact [Peppas *et al.*, 1999]. Assume that $\sigma_1; \sigma_2; \cdots; \sigma_n$ dissolves V and for all $1 \leq i \leq n$, let r_i be the state of V that is eliminated by the rule σ_i; let us also call v the one state of V that is not eliminated. The sequence of states $r_1; r_2; \cdots; r_n; v$ is called a *trace* for V.

For any two states r and p, $\mathcal{M}(r, p)$ if and only if there is a dissolvable set of states V containing r and p, such that state p appears immediately after r in some trace of V. This binary relation $\mathcal{M}$ is exactly what is needed to fully characterise causal fixed-points $Res_{\mathcal{Q}}(w, e)$ according to the augmented preferential semantics.

As mentioned above, we focus on a Hamiltonian action system. In this case, causal context will be represented entirely by a Hamiltonian path through e-predecessors of a successor state. We shall set all orderings in $\mathcal{O}$ to be PMA orderings.

Importantly, we take the "full-meet gradient" $\sigma_F(\beta, e, \gamma)$ as our choice function. This results in the following reduction:

$$Res_F(w, e) = \{r \in \mathcal{D} \cap [e] : \text{there is a Hamiltonian path through states in } \langle\!|r, e|\!\rangle_w\}.$$

It is easy to observe now that action systems describing fixed-points and the augmented preferential semantics $\mathcal{H}_\Sigma$ are selection-equivalent. More precisely, the following result can be obtained.

Theorem 4.8 *For every causal action system* $\mathcal{Q}$ *there exists a selection-equivalent Hamiltonian action system* $\langle \mathcal{W},\, \mathcal{D},\, \mathcal{E},\, \mathcal{O},\, \mathcal{M},\, \Sigma,\, Res\rangle$*.*

Conversely, for every Hamiltonian action system $\langle \mathcal{W},\, \mathcal{D},\, \mathcal{E},\, \mathcal{O},\, \mathcal{M},\, \Sigma,\, Res\rangle$ *there exists a selection-equivalent causal action system* $\mathcal{Q}$*.*

In short,

$$Res_{\mathcal{Q}}(w, e) = Res_F(w, e).$$

for any action e and state w.

5 Discussion and Conclusions

In this paper we considered an augmented preferential semantics $\mathcal{H}_\Sigma$ for reasoning about action and causality. As mentioned above, the principle of minimal change can be identified with the set of orderings, while causality is represented by the additional binary relation on states. Importantly, a new component — a family of choice functions — was introduced to capture context-sensitive nature of causal propagation. Varying all these components allows us to specify different instances of the framework. For example, a pure preferential semantics can be obtained by requesting that the causal relation is an empty set.

Under certain conditions, the causal propagation semantics, proposed by Sandewall [1996], is unified with the augmented preferential semantics $\mathcal{H}_\Sigma$, and can be found very similar to McCain-Turner's theory of action based on causal fixed-points [McCain and Turner, 1995]. The latter approach was characterised by the augmented preferential semantics relying on the PMA ordering and an appropriately constructed binary relation operating on standard state-space. The main difference appears to be in the additional condition requiring that there exist a *Hamiltonian* path through certain states in a state transition system leading to a McCain and Turner causal fixed-point. Essentially, such a Hamiltonian path serves as a contextual mechanism: the effects of causality are allowed to contribute in certain situations and not in others. Similarly, the additional conditions (M_4) and (M_5) capture the context present in respectful action systems, where causal propagation is not allowed to change "direction".

The well-known approach of Thielscher [1997], captured by the augmented preferential semantics [Prokopenko *et al.*, 2000], can be characterised within the general augmented preferential semantics by using an extended state-space and embedding causal laws into the causal relation. The distinction between Sandewall's and Thielscher's approaches to propagation-oriented ramification may be explained by the variance in the transition space dimensions and employment of different preference metrics in identifying states nearest to the initial one.

Sandewall's causal propagation semantics allows the propagation to start at any intermediate state — accounting, in particular, for disjunctive direct effects. Thielscher's approach also suggests to start state transitions at an intermediate state that is as close as possible to the initial state. Since the latter approach handles actions with conjunctive effects, there is only one such intermediate state. However, the approach reported by Thielscher [1998] extends the original one [Thielscher, 1997] towards *alternative* effect propositions, where disjunction of effects $e_1 \mid e_2$ is interpreted as exclusive, and inclusive disjunction is simply modelled as $e_1 \mid e_2 \mid e_1 \wedge e_2$. In this extended case, alternative effects lead to alternative intermediate (preliminary) states, and the original causal relationship approach is then applied to each of these preliminary states. In other words, in order to account for indirect effects, preliminary states are "taken as starting points for the successive applications of causal relationships" until overall satisfactory successor states are obtained [Thielscher, 1998]. This extension is easily handled by the "mini-choice gradient" function. In general, this function should be considered if an action system includes non-deterministic actions, or more precisely, actions with disjunctive direct effects.

We should point out, however, that restrictions imposed on the invocation relation G in causal propagation semantics may seem to limit actions preconditions. Arguably, an action direct effects should not be contingent on state of invocation. However, in some domains such specificity may be required. In these cases, instead of restricting the invocation relation G and orderings in $\mathcal{O}$, we may choose to extend our state-space by encoding relevant preconditions, achieving the desired selection-equivalence. Precise nature of a trade-off between restrictions on the preference relation and dimensionality of the propagation space remains a subject for future research.

References

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux. Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Massachusetts, 1988.

[McCain and Turner, 1995] Norman McCain and Hudson Turner. A causal theory of ramifications and qualifications. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1978–1984. Montreal, 1995.

[Peppas *et al.*, 1999] Pavlos Peppas, Maurice Pagnucco, Mikhail Prokopenko, Abhaya Nayak and Norman Foo. Preferential Semantics for Causal Systems. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 118–123. Stockholm, 1999.

[Prokopenko *et al.*, 1999] Mikhail Prokopenko, Maurice Pagnucco, Pavlos Peppas, Abhaya Nayak. Causal Propagation Semantics — A Study. In *Proceedings of the 12th Australian Joint Conference on Artificial Intelligence*, pages 378–392. Sydney, 1999.

[Prokopenko *et al.*, 2000] Mikhail Prokopenko, Maurice Pagnucco, Pavlos Peppas, Abhaya Nayak. A Unifying Semantics for Causal Ramifications. In R. Mizoguchi and J. Slaney (editors), *Proceedings of the Sixth Pacific Rim International Conference on Artificial Intelligence*, Springer Verlag Lecture Notes in AI, Volume 1886, 38 – 49, 2000.

[Sandewall, 1996] Erik Sandewall. Assessments of ramification methods that use static domain constraints. In L. Aiello, J. Doyle, and S. Shapiro, editors, *Proceedings of the Fifth International Conference on Knowledge Representation and Reasoning*, Morgan-Kaufmann, 1996.

[Thielscher, 1997] Michael Thielscher. Ramification and Causality. Artificial Intelligence 89: 317–364, 1997.

[Thielscher, 1998] Michael Thielscher. How (Not) To Minimize Events. In *Proceedings of the Sixth International Conference on Knowledge Representation and Reasoning*, Morgan-Kaufmann, 1998.

[Winslett, 1988] Marianne Winslett. Reasoning about actions using a possible models approach. In *Proceedings of the Seventh National Artificial Intelligence Conference*, San Mateo, CA., Morgan Kaufmann Publishers, 1988.