Edinburgh Research Explorer

Local Perspectives on Actions

Citation for published version:
Fourman, M 2007, Local Perspectives on Actions. in *Workshop Logic, Rationality and Interaction* .

Link:
Link to publication record in Edinburgh Research Explorer

Document Version:
Peer reviewed version

Published In:
Workshop Logic, Rationality and Interaction

Local Perspectives on Actions

Michael P. Fourman

April 18, 2007

Abstract

Giving an account of agents acting in the world—sensing, planning, communicating, doing—requires a coordinated account of, at least, three different kinds of action: ontic, epistemic, and communicative, which focus, respectively, on fact, knowledge and communication.

In this note we are concerned primarily with ontic actions. The motivating example is the STRIPS approach to the frame problem, where actions are restricted to change only specified fluents.

We present an algebraic setting for ontic actions modeled as relations describing controlled state change. We start from a standard model that encodes STRIPS updates to address the frame problem in logical terms. This model is a *system*, in the sense of Resende and Baltag.

We describe a structure that introduces a notion of local perspective or experiment. This provides a novel treatment of causal relations (which are closely related to integrity constraints and domain axioms in the AI-planning literature).

We show how this local structure arises naturally from the semantic structure of the set of possible states, and suggest that it may also help in modeling agents with different perspectives.

1 Introduction

Algebras of actions can be traced back to Tarski's calculus of relations [26], and beyond (see [22]). This provided the foundation for fifty years of fertile study (see, for example, [21, 10, 22, 14]). More recently, there has been revived interest, placing these algebraic ideas in a more abstract setting, and extending them—for example, to encompass semantics of knowledge and give an account of epistemic actions [23, 2].

Meanwhile, the AI planning community has worked on syntax and semantics of languages for specifying planning domains and actions, and on algorithms for solving for generating plans—sequences of actions that lead from presumed initial (or pre-)conditions to desired outcomes, or post-conditions.

The *frame-problem* [24] has been a recurrent issue for practical planning. Early syntactic approaches are now superseded, or subsumed, by more robust semantic accounts. However, these assume that all fluents are independent.

Inspired by [9], [20], and [2], we introduce a local structure on the set of fluents that can be derived from dependencies inherent in the world model, or can be established beforehand to impose such dependencies.

2 Properties and Actions

We consider *systems* in the sense of Resende [23] (see also Baltag *et al.* [3, 1, 2, 4]). A complete $\bigvee$-lattice of *properties*, $\varphi, \psi, \ldots$, is operated on by a quantale of *actions*, $\alpha, \beta, \ldots$.

$$(\varphi \bullet \alpha) \bullet \beta = \varphi \bullet (\alpha \, ; \beta) \tag{1}$$

$$\left(\bigvee_i \varphi_i\right) \bullet \alpha = \bigvee_i (\varphi_i \bullet \alpha) \qquad \varphi \bullet \bigvee_i \alpha_i = \bigvee_i (\varphi \bullet \alpha_i) \tag{2}$$

3 The Standard Model

Let I be a set of *state-variables*, with each $x \in I$ ranging over a non-empty set $\mathcal{S}_x$ of possible values. States are *valuations*, **v**—functions that assign a value to each state-variable. So, valuations are elements of the product *state-space*, $\mathcal{S}$:

$$\mathbf{v} \in \mathcal{S} = \prod_{x \in I} \mathcal{S}_x \qquad \mathbf{v} = \langle \mathbf{v}_x \in \mathcal{S}_x \mid x \in I \rangle \tag{3}$$

We often consider only a subset $\mathcal{P} \subseteq \mathcal{S}$ of states. A *property*, $\varphi \subseteq \mathcal{S}$, is a subset of the state space. ($\mathcal{P}$ will be an invariant property.) The set of properties is a complete lattice—indeed, in this case, a complete boolean algebra.

A many-sorted first-order predicate calculus with a sort corresponding to each state variable provides a convenient language for defining properties. For any first-order formula θ whose free variables are contained in I, a standard definition of satisfaction defines an associated property, $[\![\theta]\!]$:[1]

$$[\![\theta]\!] = \{\mathbf{v} \mid \mathbf{v} \vDash \theta\} \tag{4}$$

A non-deterministic *action* is represented by a relation $\alpha \subseteq \mathcal{S} \times \mathcal{S}$ from states to states. The relationship, $x \; \alpha \; y$, has an operational reading as, *"doing α in state x may lead to state y"*. Actions, composed sequentially by relation composition, and ordered by inclusion, form a *quantale*.

An alternative, but isomorphic, representation of the quantale of actions, as sup-preserving functions $\wp(S) \to \wp(S)$, equipped with function composition and point-wise ordering, has a logical reading: the result, $\varphi \bullet \alpha$, is the strongest post-condition including all states that *may be reached by doing α in a state satisfying* φ.

---

[1]Context is usually sufficient to distinguish θ from $[\![\theta]\!]$, so we may abuse notation by omitting the *Scott brackets*, $[\![\ ]\!]$. We will also let a boolean variable stand for itself: $b \equiv b = \top$.

We call this the *standard model*, since it has emerged as a common foundation for much work on actions and planning (sometimes explicitly, often implicitly: see, e.g. [8, 19, 17, 18, 11, 16, 6, 15, 7, 13]). In this setting, we define some basic actions.[2]

For every property, φ, we have an action $?\varphi$, called *require φ*:

$$\psi ?\varphi = \psi \wedge \varphi \tag{5}$$

For every map, $\sigma : I \to I$, we have an action $[\sigma(x)/x]$ *(substitute $\sigma(x)$ for x):*

$$\varphi[\sigma(x)/x] = \{\mathbf{w} \mid \mathbf{w} \circ \sigma \in \varphi\} \tag{6}$$

For every set of state variables, $U \subseteq I$, we have an action, $\||U$:

$$\psi\big\|\big| U = \{\mathbf{w} \mid \mathbf{w} \restriction U \in \psi \restriction U\} \tag{7}$$

this action, *preserve U*, is the largest action that preserves the values of state variables in U. It allows other variables to change non-deterministically. We also write $\updownarrow V$, called *change V*, for $\||(I \setminus V)$. This action may be expressed using quantification over state variables: $\varphi \updownarrow V \equiv \exists V. \phi$. A more refined version, permitting change in one direction only, can be expressed using bounded quantification.

3.1 Local Change

The standard solution to the frame-problem is to confine attention to actions that leave most observations unchanged.

The STRIPS representation of actions was originally defined operationally, for a setting with only boolean state variables. A state is represented by the set of *true* variables. The STRIPS action α, given by a triple, (*pre*, *add*, *delete*), of sets of variables, can be applied to a set $U \subseteq I$ of state variables, iff *pre* $\subseteq U$; if applied, it produces the result $(U \setminus \textit{delete}) \cup \textit{add}$.

In the present setting, we can define the action corresponding to a STRIPS triple by

$$\text{STRIPS}(\textit{pre}, \textit{add}, \textit{delete}) =$$
$$(? \bigwedge \textit{pre}; \updownarrow (\textit{add} \cup \textit{delete}); ?(\bigwedge \textit{add} \wedge \neg \bigvee \textit{delete})) \tag{8}$$

It is straightforward to calculate that applying this action to the property corresponding to the single state in which U is the set of *true* variables, simulates the STRIPS procedure.

---

[2]We often omit the $\bullet$ and use simple postfix notation in place of $- \bullet \alpha$.

4 Causal Models

The treatment of domain axioms or integrity constraints remains an outstanding problem. The issue is, how to model a world where the values of state variables may be inter-related? When one variable is changed, others may require adjustment to accommodate some universal constraint. It has seemed hard to restrict the ability to change to the appropriate variables in a semantically principled way. Here, we suggest an approach to this problem that may initially appear *ad hoc*. In the following section we derive our treatment from purely semantic principles.

Rather than considering all state variables as independent atoms, we identify a structure that encodes functional dependencies. This structure places the atomic observations of fluents within a poset of perspectives, which is a complete $\bigvee$-lattice with distributive joins—a *frame*.

Starting from any poset, the downsets constitute the frame of open sets of a point-set topology on the poset. The poset is embedded, in the frame, by taking each element, x to the down-set $\downarrow x$, composed of elements $y \leq x$. This embedding preserves existing meets ($\wedge$), but not joins ($\bigvee$).

The original poset has a natural map, ν, to any quotient of this frame (by $\wedge$, $\bigvee$-congruences). Indeed, a quotient may be specified by stipulating *covering conditions*, of the form $\nu(y) \leq \bigvee_i \nu(x_i)$ (algebraically $\nu(y) \wedge \bigvee_i \nu(x_i) = \bigvee_i \nu(x_i)$).

Consider an example. Pearl [20] models causal relationships as functional dependencies, and defines a *causal model* (*op cit.* p27.) to be

> a set of equations of the form
>
> $$x_i = f_i(pa_i, u_i), \quad i = 1, \ldots, n, \tag{9}$$
>
> where pa_i (connoting *parents*) stands for the set of variables judged to be immediate causes of X_i and where the U_i represent errors (or "disturbances") due to omitted factors.

Such a model is represented as a directed graph, with a node for each variable, and an arrow $x \rightarrow y$ if x is a parameter of an equation defining y. In this note, we confine our attention to causal models represented by a directed acyclic graph (DAG). We take this graph as a pre-order on variables, and take the reflexive, transitive closure of $x \rightarrow y$, to give a partial order.

Consider a specific example: suppose that we want to include in our model a particular dependency — that $y = f(x_1, \ldots, x_n)$. In general, if one of the x_i changes, then y may change. So, the set of fluents that may be changed by an action cannot be entirely arbitrary. Moreover, y cannot change unless at least one of the x_i does.

We define a (point-set) topology on the set of fluents such that any closed set of fluents a is a suitable candidate for changes.

First, we equip our poset with a topology in which down-sets are open. So any open set containing y contains all of the x_i. Second, we choose the topology a

coarser topology, with fewer open sets, generated by letting the family $\langle x_i \to y_i \mid i = 1, \ldots, n \rangle$ cover y. In other words, we enforce the condition that any open set containing all of the x_i contains y.

A closed set contains y iff it contains at least one of the x_i; We allow changes to a closed set—equivalently, we allow the action $|||U$ only where U is an open set.

4.1 Grothendieck Topology

Let P be a poset, and ν a $\wedge$-preserving map from P to a frame. For $y \in P$, we say a family $x_i \leq y$ ν-*covers* y iff $\nu(y) \leq \bigvee_i \nu(x_i)$. The collection of ν-covers is a *Grothendieck topology* on the poset. Every Grothendieck topology arises in this way—we can construct a suitable frame, and the map, ν, from the topology. An intersection of Grothendieck topologies is a Grothendieck topology—since a product of frames is a frame. The coarse topology, in which every family is covering, is a Grothendieck topology—since the singleton lattice is a frame. So any collection of covering families generates a Grothendieck topology.

For details of Grothendieck topologies on posets, see [25], and, for a more general account, [12].

5 Canonical Causal Dependencies

We revisit the example of the previous section. Instead of imposing a partial order and topology, we derive these from the semantic structure. To do this we introduce new structure, conservatively extending the standard model. In the next section we introduce an abstract presentation of this extension, which is in general substantive.

We start from the standard model, but consider only those valuations that satisfy the equations of the causal model: $\mathcal{P} = \{\mathbf{v} \mid \mathbf{v} \text{ satisfies the causal constraints.}\}$

We introduce a family $U \in \mathcal{O}$ of *perspectives*—these may be thought of as experiments. In the standard model, these are just sets of variables, $U \subseteq I$. The intuition is that these "perspectives" represent a local focus on some aspects of the world, an isolated environment within which we may experiment independently of outside influences. We write $\mathcal{P}(U)$ for the local states that may be hypothesized from this perspective U. In the standard model these are just the restrictions of global states to U.

Perspectives are partially-ordered: $V \to U$ if V represents a narrower focus than U, and we require that this be evidenced by a *restriction map* $\restriction V : \mathcal{P}(U) \to \mathcal{P}(V)$. For the standard model this order is just set inclusion, and the restriction maps are immediate.

We say a family of hypotheses $a_i \in \mathcal{P}(U_i)$ is *coherent* iff, whenever there is a perspective W such that $U_i \leftarrow W \to U_j$ then $a_i \restriction W = a_j \restriction W$. A family $U_i \to W$ *P-covers* W iff for every coherent family $a_i \in \mathcal{P}(U_i)$, there is a unique hypothesis $\bar{a} \in \mathcal{P}(W)$ such that, for all i, $\bar{a} \restriction U_i = a_i$. We call $\bar{a}$ an extension of the a_i.

Now suppose we have a functional relationship $y = f(x_1, \ldots, x_n)$. We have a inclusions $\{x_i\} \to \{x_1, \ldots, x_n, y\}$. Any selection of values $a_i \in \mathcal{P}(\{x_i\})$ is a coherent family and it has a unique extension $\bar{a} = \langle a_1, \ldots, a_n, f(a1, \ldots, a_n) \rangle \in \mathcal{P}(x_1, \ldots, x_n, y)$.

The collection of all covering families generates a Grothendieck topology on the poset of perspectives. The sets $\mathcal{P}(U)$ together with restriction maps form a presheaf whose sheafification is our substitute for the presentation-derived construction of the previous section.

In the simple example with a single functional dependence, $\{x_1, \ldots, x_n, y\}$, in the sheafification, plays the rôle of y in our earlier construction; $\{y\}$ turns out to be covered by the empty family—and so behaves as a minimal element accommodating a trivial, singleton-set of possibilities in the sheafification.

6 Perspectives

Algebraically, perspectives form a partially ordered set, equipped with a Grothendieck topology; we take the state space $\mathcal{P}$ to be a sheaf over this site. Equivalently, let $\mathcal{X}$ be a locale, whose open sets we call *perspectives*, and $\mathcal{S}$ a sheaf on $\mathcal{X}$. We say that the space of perspectives is *irredundant* if, whenever $U_i \mathcal{P}$-covers W then $\bigvee_i U_i = W$.

Global states arise as global sections of $\mathcal{P}$—which correspond to consistent sets of hypotheses from each local perspective. Each perspective U induces an equivalence relation on global states: $\mathbf{v} \equiv \mathbf{w}$ iff $\mathbf{v} \upharpoonright U = \mathbf{w} \upharpoonright U$ these states appear equivalent from this perspective. There may also be states, elements of $\mathcal{P}(U)$, that appear possible from this perspective, but that do not extend to global states. It may be interesting to link the appearance maps of [1] to these structures.

7 Conclusions

This is a preliminary report on ongoing work.

We have shown how a causal model, in the sense of Pearl, can arise naturally from semantic analysis of a set of possible states. This leads to an account of local perspectives as the open sets of a space of observations.

The structures introduced to this end suggest the beginnings of an account of local perspective, hypothesis and experiment.

References

[1] A. Baltag. A logic of epistemic actions. In *Proceedings of the Workshop on 'Foundations and Applications of Collective Agent Based Systems*, 1999. 11th European Summer School on Logic, Language and Information, Utrecht University, Utrecht.

[2] Alexandru Baltag, Bob Coecke, and Mehrnoosh Sadrzadeh. Algebra and sequent calculus for epistemic actions. *Electronic Notes in Theoretical Computer Science*, 126:27–52, March 2005.

[3] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements, common knowledge, and private suspicions. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 43–56, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

[4] Alexandru Baltag and Mehrnoosh Sadrzadeh. The algebra of multi-agent dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 157(4):37–56, May 2006.

[5] Anthony G. Cohn, Lenhart Schubert, and Stuart C. Shapiro, editors. *KR'98: Principles of Knowledge Representation and Reasoning*, San Francisco, California, 1998. Morgan Kaufmann.

[6] Patrick Doherty, Witold Łukasziewicz, and Ewa Madalińska-Bugaj. The PMA and relativizing change for action update. In Cohn et al. [5], pages 258–269.

[7] Enrico Giunchiglia, G. Neelakantan Kartha, and Vladimir Lifschitz. Representing action: Indeterminacy and ramifications. *Artificial Intelligence*, 95(2):409–438, 1997.

[8] Joakim Gustafsson and Patrick Doherty. Embracing occlusion in specifying the indirect effects of actions. In *Principles of Knowledge Representation and Reasoning*, pages 87–98, 1996.

[9] Joseph Y. Halpern and Moshe Y. Vardi. Model checking vs. theorem proving: A manifesto. In J. Allen, R. E. Fikes, and E. Sandewall, editors, *Proceedings 2nd Int. Conf. on Principles of Knowledge Representation and Reasoning, KR'91, ???*, pages 325–334. Morgan Kaufmann Publishers, San Mateo, CA, 1991.

[10] D. Harel. Dynamic logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic Volume II — Extensions of Classical Logic*, pages 497–604. D. Reidel Publishing Company: Dordrecht, The Netherlands, 1984.

[11] A. Herzig, J. Lang, P. Marquis, and Th. Polacsek. Updates, actions and planning. In Bernhard Nebel, editor, *IJCAI'01*, pages 119–124, Seattle, Washington, USA, 2001. Morgan Kaufmann.

[12] P. T. Johnstone. *Stone spaces*, volume 3 of *Cambridge studies in advanced mathematics*. Cambridge University Press, London, 1982.

[13] G. Neelakantan Kartha and Vladimir Lifschitz. Actions with indirect effects (preliminary report). In Jon Doyle, Erik Sandewall, and Pietro Torasso, editors, *KR'94: Principles of Knowledge Representation and Reasoning*, pages 341–350. Morgan Kaufmann, San Francisco, California, 1994. see also [7].

[14] Dexter Kozen. A representation theorem for models of *-free pdl. In *Proceedings of the 7th Colloquium on Automata, Languages and Programming*, pages 351–362, London, UK, 1980. Springer-Verlag.

[15] J. Lang and P. Marquis. Complexity results for independence and definability in propositional logic. In Cohn et al. [5], pages 356–367.

[16] Jérôme Lang, P. Liberatore, and P. Marquis. Propositional independence: Formula-variable independence and forgetting. *JAIR*, 18:391–443, 2003.

[17] Fangzhen Lin and Raymond Reiter. Forget it! In Russell Greiner and Devika Subramanian, editors, *Working Notes, AAAI Fall Symposium on Relevance*, pages 154–159, Menlo Park, California, 1994. American Association for Artificial Intelligence.

[18] John McCarthy. *Formalization of common sense, papers by John McCarthy edited by V. Lifschitz*. Ablex, 1990.

[19] Erik T. Mueller. Event calculus reasoning through satisfiability. *J Logic Computation*, 14(5):731–745, 2004.

[20] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, March 2000.

[21] Vaughan Pratt. Dynamic algebras: Examples, constructions, applications.

[22] Vaughan Pratt. Action logic and pure induction. In J. van Eijck, editor, *Proc. Logics in AI: European Workshop JELIA '90*, volume 478, pages 97–120. Springer-Verlag Lecture Notes in Computer Science, 1990.

[23] P. Resende. Quantales and observational semantics. In B. Coecke, D. Moore, and A. Wilce, editors, *Current Research in Operational Quantum Logic: Algebras, Categories and Languages*, volume 111, pages 263–288. Kluwer Academic Publishers, 2000.

[24] Murray Shanahan. *Solving the Frame Problem: a mathematical investigation of the common sense law of inertia*. Artificial Intelligence. The MIT Press, Cambridge, Mass, 1997.

[25] Isar Stubbe. The canonical topology on a meet-semilattice. *International Journal of Theoretical Physics*, 2005.

[26] A. Tarski. On the calculus of relations. *J. Symbolic Logic*, 6:73–89, 1941.