An Outline of ALX3, a Multi-Agent Action Logic

Zhisheng Huang and Michael Masuch
Center for Computer Science in Organization and Management (CCSOM)
University of Amsterdam
{huang,michael}@ccsom.uva.nl

Abstract

ALX3 is a multi-agent version of ALX with a first-order description language. ALX, a modal action logic, combines ideas from H.A. Simon's *bounded rationality*, S. Kripke's *possible world semantics*, G. H. von Wright's *preference logic*, Pratt's *dynamic logic*, and Stalnaker's *conditional logic*. ALX avoids important drawbacks of other action logics, especially the counterintuitive necessitation rule for goals (every theorem must be a goal) and the equally counterintuitive closure of goals under logical implication. ALX3 is sound and complete, and is already proving its practical use in the formal representation of modern organization theory.

1 Introduction

Action logics are usually developed for the (hypothetical) use by intelligent robots [CL90, GS87, RG91] or as a description language of program behavior [Har84]. Our effort is motivated by a different concern. We want to develop a formal language for social science theories, especially for theories of organizations. The difference in motivation leads to a new approach to action logic. It combines ideas from various strands of thought, notably H.A. Simon's notion of *bounded rationality*, Kripke's *possible world semantics*, V.R. Pratt's *dynamic logic,* Stalnaker's notion of *minimal change*, G. H. von Wright's approach to *preferences*, and J. Hintikka's approach to *knowledge* and *belief*. ALX3 has a first order description language with multiple agents, and four modal operator types.[1] We outline the language, and discuss some of its important properties. A fuller presentation is given in [HM94]. In a companion paper, ALX3's potential for knowledge representation is extensively demonstrated in the formalization of an important organization theory, J.D. Thompson's *Organizations in Action* [MH94].

2 ALX's Background

Most social science theories are expressed in natural language, but natural language does not provide a formal scaffold for checking a theory's logical properties. As a

[1] ALX stands for the *x*'s *A*ction *L*ogic. ALX1, the first version, had a propositional description language and a (backward-looking) update operator instead of the conditional; it was a single-agent language [HMP95]. ALX2, the intermediate version, is not multi-agent [HMP93].

consequence, the social sciences have acquired a reputation for "softness" — a soft way of saying that the logical properties of their theories are often dubious. Reformulating a social theory in a formal language with known logical properties would facilitate the tasks of consistency checking, disambiguation, or the examination of other important logical properties, such as contingency (whether or not the theory is falsifiable).

We focus on action logic as a formal language, because actions are key to the understanding of social phenomena. In fact, most social scientist agree that action theory provides the underlying framework for the social sciences in general [Blu69, Gid79, Par37]. Yet actions involve attitudes and engender change, and both phenomena are notoriously hard to grasp in the extensional context of first order languages [Gam90]. This explains our attempt to develop a new logic, rather than taking First Order Logic off the shelf.

Herbert A. Simon's conceptualization of *bounded rationality* [Sim55] serves as a point of departure. His approach is intuitively appealing, and had great impact on the postwar social sciences. Simon wanted to overcome the omniscience claims of the traditional conceptualizations of rational action. He assumed (1) an agent with (2) a set of behavior alternatives, (3) a set of future states of affairs (each such state being the outcome of a choice among the behavior alternatives), and (4) a preference order over future states of affairs. The omniscient agent, endowed with "perfect rationality", would know all behavior alternatives and the exact outcome of each alternative; the agent would also have a complete preference ordering for those outcomes. An agent with bounded rationality, in contrast, might not know all alternatives, nor need it know the exact outcome of each; also, the agent might lack a complete preference ordering for those outcomes.

Kripke's *possible world semantics* provides a natural setting for Simon's conceptualization. We assume a set of possible worlds with various relations defined over this set (we may also call those possible worlds *states*). One can see a behavior alternative as a mapping from states to states, so each behavior alternative constitutes an accessibility relation. An accessibility relation, in turn, can be interpreted as an opportunity for action, that is, as an opportunity for changing the world by moving from a given state to another state. Accessibility relations are expressed by indexed one-place modal operators, as in dynamic logic [Har84]. For example, the formula $\langle a_i \rangle \phi$ expresses the fact that the agent has an action a at its disposal such that effecting a in the present situation would result in the situation denoted by ϕ.

Preferences – not goals – provide the basic rationale for rational action in ALX3. Following von Wright [vWr63], a preference statement is understood as a statement about situations. For example, the statements that "I prefer oranges to apples" is interpreted as the fact that "I prefer the states in which I have an orange to the states in which I have an apple." Following von Wright again, we assume that an agent who says that she prefers oranges to apples should prefer a situation where she has an orange but *no* apple to a situation where she has an apple but *no* orange. We call this principle *conjunction expansion principle* and restrict attention to preference statements that obey it. Preferences are expressed via two-place modal operators; if the agent prefers the proposition ϕ to the proposition ψ, we write $\phi \mathbf{P}_i \psi$.

Normally, the meaning of a preference statement is context dependent, even if this is not made explicit. An agent may say to prefer an apple to an orange – and actually mean it – but she may prefer an orange to an apple later – perhaps because

Figure 1: Simon's Bounded Rationality

then she already had an apple. To capture this context dependency, we borrow the notion of minimal change from Stalnaker's approach to conditionals [Sta68]. The idea is to apply the conjunction expansion principle only to situations that are minimally different from the agent's present situation – just as different as they really need to be in order to make the propositions true about which preferences are expressed. We introduce a binary function, cw, to the semantics that determines the set of "closest" states relative to a given state, such that the new states fulfill some specified conditions, but resembles the old state as much as possible in all other respects.

The syntactic equivalent of the closest world function is the wiggled "causal arrow". It appears in expressions such as $\phi \rightsquigarrow \psi$ where it denotes: in all closest worlds where ϕ holds, ψ also holds. The causal arrow expresses the conditional notion of a causal relation between ϕ and ψ: if ϕ *were* the case, ψ *would* also be the case.

The last primitive operator of ALX3 is the indexed belief operator. In a world of bounded rationality, an agent's beliefs do not necessarily coincide with reality, and in order to make this distinction, we must be able to distinguish between belief and reality; $\mathbf{B}_i(\phi)$ will denote the fact that agent i believes ϕ. As the logical axioms characterizing the belief operator show, $\mathbf{B}$ represents a sense of "subjective knowledge", not metaphysical attachment, or epistemic uncertainty.

3 Syntax and Semantics

3.1 Formal Syntax

ALX3 has the following primitive symbols:

(1) For each natural number $n(\geq 1)$, a countable set of n-place predicate letters, $p_i, p_j, ...$

(2) A countable set of regular variables, written $x, x_1, y, z, ...$; a countable set of action variables, written $a, a_1, b, ...$; a countable set of agent variables, written $i, i_1, j, ...$

(3) A countable set of regular constants, written $c, c_1, c_2, ...$; a countable set of actions constants, written $ac, ac_1, ac_2, ...$; a countable set of agent constants, written $ag, ag_1, ag_2, ...$

(4) The symbols $\neg$(negation), $\wedge$(conjunction), $\boldsymbol{B}$(belief), $\exists$(existential quantifier), $\boldsymbol{P}$(preference), $\rightsquigarrow$(conditional), ;(sequence), $\cup$(choice), $\langle, \rangle$, (, and).

Furthermore, ALX3 has the following syntax rules:

$$
\begin{array}{ll}
\langle \text{Variable} \rangle & ::= \langle \text{Regular variable} \rangle | \langle \text{Action variable} \rangle | \langle \text{Agent variable} \rangle \\
\langle \text{Constant} \rangle & ::= \langle \text{Regular constant} \rangle | \langle \text{Action constant} \rangle | \langle \text{Agent constant} \rangle \\
\langle \text{Term} \rangle & ::= \langle \text{Variable} \rangle | \langle \text{Constant} \rangle \\
\langle \text{Action term} \rangle & ::= \langle \text{Action variable} \rangle | \langle \text{Action constant} \rangle \\
\langle \text{Agent term} \rangle & ::= \langle \text{Agent variable} \rangle | \langle \text{Agent constant} \rangle \\
\langle \text{Atom} \rangle & ::= \langle \text{Predicate} \rangle (\langle \text{Term} \rangle, \cdots, \langle \text{Term} \rangle) \\
\langle \text{Action} \rangle & ::= \langle \text{Action term} \rangle_{\langle \text{Agent term} \rangle} | \langle \text{Action} \rangle ; \langle \text{Action} \rangle | \langle \text{Action} \rangle \cup \langle \text{Action} \rangle \\
\langle \text{Formula} \rangle & ::= \langle \text{Atom} \rangle | \neg \langle \text{Formula} \rangle | \langle \text{Formula} \rangle \wedge \langle \text{Formula} \rangle | \exists \langle \text{Variable} \rangle \langle \text{Formula} \rangle | \\
 & \langle \langle \text{Action} \rangle \rangle \langle \text{Formula} \rangle | \langle \text{Formula} \rangle \rightsquigarrow \langle \text{Formula} \rangle | \\
 & \langle \text{Formula} \rangle \mathbf{P}_{\langle \text{Agent term} \rangle} \langle \text{Formula} \rangle | \mathbf{B}_{\langle \text{Agent term} \rangle} \langle \text{Formula} \rangle
\end{array}
$$

3.2 Semantics

Definition 1 (ALX3 Model)

$$Call \;\; M = \langle O, PA, AGENT, W, cw, \succ, \mathcal{R}, \mathcal{B}, I \rangle$$

an ALX3 model, if O is a set of objects, PA is a set of primitive actions, $AGENT$ is a set of agents, W is a set of possible worlds, $cw : W \times \mathcal{P}(W) \to \mathcal{P}(W)$ is a closest world function, $\succ$: $AGENT \to \mathcal{P}(\mathcal{P}(W) \times \mathcal{P}(W))$ is a function that assigns a comparison relation for preferences to each agent, $\mathcal{R} : AGENT \times PA \to \mathcal{P}(W \times W)$ is a function that assigns an accessibility relation to each agent and each primitive action, $\mathcal{B} : AGENT \to \mathcal{P}(W \times W)$ is a function that assigns a serial and transitive accessibility relation for the belief operation to each agent, I is a interpretation function, and if $cw, \succ$, satisfy the following conditions respectively:

$(\boldsymbol{CS1})$: $cw(w, X) \subseteq X$.
$(\boldsymbol{CS2})$: $w \in X \Rightarrow cw(w, X) = \{w\}$.
$(\boldsymbol{CS3})$: $cw(w, X) = \emptyset \Rightarrow cw(w, Y) \cap X = \emptyset$.
$(\boldsymbol{CS4})$: $cw(w, X) \subseteq Y$ and $cw(w, Y) \subseteq X \Rightarrow cw(w, X) = cw(w, Y)$.
$(\boldsymbol{CS5})$: $cw(w, X) \cap Y \neq \emptyset \Rightarrow cw(w, X \cap Y) \subseteq cw(w, X)$.
$(\boldsymbol{NORM})$: $(\emptyset \not\succ_i X), (X \not\succ_i \emptyset)$, where $\succ_i = \succ (i)$ for each agent $i \in AGENT$.
$(\boldsymbol{TRAN})$: $cw(w, X \cap \overline{Y}) \succ_i cw(w, Y \cap \overline{X})$ and $cw(w, y \cap \overline{Z}) \succ_i cw(w, Z \cap \overline{Y})$ $\Rightarrow cw(w, X \cap \overline{Z}) \succ_i cw(w, Z \cap \overline{X})$, where $\overline{X} = W - X$.

The (CS#) constrain the closest world function. They constitute the standard requirements for a closest world function as established by Stalnaker[Sta68]. (NORM) and (TRAN) constrain the semantic preference relation. (NORM) is required in support of the logical axiom (N) (normality), which, in turn, protects the preference logic against counterintuitive consequences. (TRAN) guarantees the soundness of the logical axiom (TR) which, in turn, assures transitivity for preferences. The seriality and transitivity of the relation $\mathcal{B}$ are standard requirements for the semantics of beliefs.

Definition 2 (Meaning function)

$[\![p(t_1, ..., t_n)]\!]^v_M = \{w : \langle v_I(t_1), ..., v_I(t_n) \rangle \in I(p, w)\}$
$[\![\neg\phi]\!]^v_M = W \setminus [\![\phi]\!]^v_M$
$[\![\phi \wedge \psi]\!]^v_M = [\![\phi]\!]^v_M \cap [\![\psi]\!]^v_M$
$[\![\exists x \phi]\!]^v_M = \{w : (\exists d \in D)(w \in [\![\phi]\!]^{v(d/x)}_M)\}$
$[\![\langle a \rangle \phi]\!]^v_M = \{w : (\exists w')(R^a w w' \text{ and } w' \in [\![\phi]\!]^v_M)\}$
$[\![\phi \rightsquigarrow \psi]\!]^v_M = \{w : cw(w, [\![\phi]\!]^v_M) \subseteq [\![\psi]\!]^v_M\}$
$[\![\phi \boldsymbol{P}_i \psi]\!]^v_M = \{w : cw(w, [\![\phi \wedge \neg\psi]\!]^v_M) \succ_{v_I(i)} cw(w, [\![\psi \wedge \neg\phi]\!]^v_M)\}$
$[\![\boldsymbol{B}_i \phi]\!]^v_M = \{w : (\forall w')(\langle w, w' \rangle \in \mathcal{B}_{v_I(i)} \Rightarrow w' \in [\![\phi]\!]^v_M)\}$

The interpretation of the atomic formulas, the boolean connectives and the existential quantifier is straightforward. The interpretation of $\langle a \rangle \phi$ yields the set of worlds from where the agent can access at least one ϕ-world via action a. The interpretation of $\phi \rightsquigarrow \psi$ yields the set of worlds in which the closest ϕ-worlds are also ψ-worlds. So, $\phi \rightsquigarrow \psi$ is true at a world if ϕ would *ceteris paribus* entail ψ. This is the standard counter-factual conditional used to express a causal relation between ϕ and ψ. Note that our wiggled arrow does not require actual counter-factuality, so ϕ may be true

in the actual world. The interpretation of $\phi \boldsymbol{P}_i \psi$ assures the conjunction expansion principle.

Definition 3 (ALX3 inference system)

$(\boldsymbol{BA})$: *all tautologies of the first order logic*
$(\boldsymbol{A1})$: $\langle a \rangle \bot \leftrightarrow \bot$
$(\boldsymbol{A2})$: $\langle a \rangle (\phi \vee \psi) \leftrightarrow \langle a \rangle \phi \ \vee \langle a \rangle \psi$
$(\boldsymbol{A3})$: $\langle a ; b \rangle \phi \leftrightarrow \langle a \rangle \langle b \rangle \phi$
$(\boldsymbol{A4})$: $\langle a \cup b \rangle \phi \leftrightarrow \langle a \rangle \phi \vee \langle b \rangle \phi$
$(\boldsymbol{AU})$: $[a] \forall x \phi \leftrightarrow \forall x [a] \phi$
$(\boldsymbol{ID})$: $\psi \rightsquigarrow \psi$
$(\boldsymbol{MPC})$: $(\psi \rightsquigarrow \phi) \rightarrow (\psi \rightarrow \phi)$
$(\boldsymbol{CC})$: $(\psi \rightsquigarrow \phi) \wedge (\psi \rightsquigarrow \phi') \rightarrow (\psi \rightsquigarrow \phi \wedge \phi')$
$(\boldsymbol{MOD})$: $(\neg \psi \rightsquigarrow \psi) \rightarrow (\phi \rightsquigarrow \psi)$
$(\boldsymbol{CSO})$: $[(\psi \rightsquigarrow \phi) \wedge (\phi \rightsquigarrow \psi)] \rightarrow [(\psi \rightsquigarrow \chi) \leftrightarrow (\phi \rightsquigarrow \chi)]$
$(\boldsymbol{CV})$: $[(\psi \rightsquigarrow \phi) \wedge \neg(\psi \rightsquigarrow \neg \chi)] \rightarrow [(\psi \wedge \chi) \rightsquigarrow \phi]$
$(\boldsymbol{CS})$: $(\psi \wedge \phi) \rightarrow (\psi \rightsquigarrow \phi)$
$(\boldsymbol{CEP})$: $\phi \boldsymbol{P}_i \psi \leftrightarrow (\phi \wedge \neg \psi) \boldsymbol{P}_i (\neg \phi \wedge \psi)$
$(\boldsymbol{N})$: $\neg(\bot \boldsymbol{P}_i \phi), \neg(\phi \boldsymbol{P}_i \bot)$
$(\boldsymbol{TR})$: $(\phi \boldsymbol{P}_i \psi) \wedge (\psi \boldsymbol{P}_i \chi) \rightarrow (\phi \boldsymbol{P}_i \chi)$
$(\boldsymbol{PC})$: $(\phi \boldsymbol{P}_i \psi) \rightarrow \neg((\phi \wedge \neg\psi) \rightsquigarrow \neg(\phi \wedge \neg\psi)) \wedge \neg((\psi \wedge \neg\phi) \rightsquigarrow \neg(\psi \wedge \neg\phi))$
$(\boldsymbol{KB})$: $\boldsymbol{B}_i \phi \wedge \boldsymbol{B}_i (\phi \rightarrow \psi) \rightarrow \boldsymbol{B}_i \psi$
$(\boldsymbol{DB})$: $\neg \boldsymbol{B}_i \bot$
$(\boldsymbol{4B})$: $\boldsymbol{B}_i \phi \rightarrow \boldsymbol{B}_i \boldsymbol{B}_i \phi$
$(\boldsymbol{BFB})$: $\forall x \boldsymbol{B}_i \phi \leftrightarrow \boldsymbol{B}_i \forall x \phi$

$(\boldsymbol{MP})$: $\vdash \phi \ \& \ \vdash \phi \rightarrow \psi \Rightarrow \vdash \psi$
$(\boldsymbol{G})$: $\vdash \phi \Rightarrow \vdash \forall x \phi$
$(\boldsymbol{NECA})$: $\vdash \phi \Rightarrow \vdash [a] \phi$
$(\boldsymbol{NECB})$: $\vdash \phi \Rightarrow \vdash \boldsymbol{B}_i \phi$
$(\boldsymbol{MONA})$: $\vdash \langle a \rangle \phi \ \& \ \vdash \phi \rightarrow \psi \Rightarrow \vdash \langle a \rangle \psi$
$(\boldsymbol{MONC})$: $\vdash \phi \rightsquigarrow \psi \ \& \ \vdash \psi \rightarrow \psi' \Rightarrow \vdash \phi \rightsquigarrow \psi'$
$(\boldsymbol{SUBA})$: $\vdash (\phi \leftrightarrow \phi') \Rightarrow \vdash (\langle a \rangle \phi) \leftrightarrow (\langle a \rangle \phi')$
$(\boldsymbol{SUBC})$: $\vdash (\phi \leftrightarrow \phi') \ \& \ \vdash (\psi \leftrightarrow \psi') \Rightarrow \vdash (\phi \rightsquigarrow \psi) \leftrightarrow (\phi' \rightsquigarrow \psi')$
$(\boldsymbol{SUBP})$: $\vdash (\phi \leftrightarrow \phi') \ \& \ \vdash (\psi \leftrightarrow \psi') \Rightarrow \vdash (\phi \boldsymbol{P}_i \psi) \leftrightarrow (\phi' \boldsymbol{P}_i \psi')$

Most axioms are straightforward. As usual, we have the tautologies (BA). Since ALX3 is a normal modal logic, the absurdum is not true anywhere, so it is not accessible (A1). The action modalities behave as usual, so they distribute over disjunction both ways (A2) (they also distribute over conjunction in one direction, but the corresponding axiom is redundant). (A3) characterizes the sequencing operator ';' and (A4) does the same for the indeterminate choice of actions. (AU) establishes the Barcan formula for universal action modalities. We have the Barcan formula because the underlying domain D is the same in all possible worlds.

The next seven axioms characterize the intensional conditional. Informally speaking, they syntactically specify the meaning of "ceteris paribus" in ALX3. They are

fairly standard, and, with the exception of (CC), they already provide a characterization of Lewis' system **VC**, which, in turn, is an adaptation of Stalnaker's conditional logic to a system for non-unique closest worlds. (ID) establishes the triviality that ψ is true in all closest ψ-worlds; (MPC) relates the intensional and the material conditional in the obvious way: so if ϕ would hold given ψ, then, if ψ actually does hold, ϕ must also hold. Conjunction distributes over the "wiggled arrow" in one way (CC). (MOD) rules out the eventuality of closest absurd worlds; (CSO) gives an identity condition for closest worlds, (CV) establishes a cautious monotony for the intensional conditional, and (CS) relates the conjunction to the intensional conditional. Replacing (CS) by $(\phi \rightsquigarrow \psi) \vee (\phi \rightsquigarrow \neg\psi)$ yields Stalnaker's original system, as the new axiom would require the uniqueness of the closest possible world).

The next four axioms characterize the preference relation. (CEP) states the conjunction expansion principle. (N) establishes "normality" and (TR) transitivity. As noted before, (TR) would go if its semantic equivalent, (TRAN), goes, so we could have non-transitive preferences. (CEP) and (N) together imply the irreflexivity of the **P** operator[Hua94]. The axiom (PC) says that if an agent i prefers ϕ to ψ, then both $\phi \wedge \neg\psi$ and $\psi \wedge \neg\phi$ are possible.

The last four axioms give a characterization of the belief operator. As pointed out above, our belief operator is designed to represent subjective knowledge. (KB) is standard in epistemic logic, but it is often criticized, since it requires logical omniscience with respect to the material conditional. On the other hand, one would expect to draw correct logical inferences when necessary, so not having (KB) may be worse. (DB) rules out the belief in absurdities, (4B) establishes positive self-introspection for beliefs, and (BFB) is the Barcan formula for beliefs. These four axioms give a standard characterization of subjective knowledge. Together with the inference rules (MP) and (NECB), they turn the belief operation into a weak S4 system. As shown in [Hua94], we could weaken the belief operator considerably, but these weaker alternatives have their own problems that would overload ALX. (One radical alternative would be an empty belief operator.)

The remaining expressions characterize ALX3's inference rules. We have the modus ponens and generalization for obvious reasons. By the same token, we have the necessitation rule for the universal action modality: if indeed, ϕ is true in all worlds, then all activities will lead to ϕ-worlds; by the same token, we have the necessitation rule for beliefs. (MONA) connects the meaning of the action modality with the meaning of the material conditional. We have right monotonicity for the intensional conditional but *not* left monotonicity. Furthermore, logically equivalent propositions are substitutable in action-, conditional-, and preference formulae (SUBA), (SUBC), (SUBP). Note that we do *not* have monotonicity for preferences. Because of this, we are able to avoid the counterintuitive deductive closure of goals that mars other action logics. In [HM94], we prove that ALX3 is sound and complete.

4 Defined Operators

ALX provides considerable flexibility via the definition of additional modal operators. The standard alethic operators "Necessity" and "Possibility" are (inter)definable via

the causal conditional [Sta68]:

$$\Box\phi \stackrel{\text{def}}{\Longleftrightarrow} \neg\phi \rightsquigarrow \phi \qquad \Diamond\phi \stackrel{\text{def}}{\Longleftrightarrow} \neg(\Box\neg\phi)$$

These alethic operators have the conventional properties; in particular, they support the traditional (K), (T), and (D) axioms. We define the knowledge operator along classical lines as true belief:

$$\boldsymbol{K}_i\phi \stackrel{\text{def}}{\Longleftrightarrow} \phi \wedge \boldsymbol{B}_i\phi$$

The knowledge-operator has the familiar properties. See [Hua94, HM94]. It it sometimes relevant whether agent i can directly access a particular state via an action. Define direct accessibility as follows:

$$\mathbf{DA}_i(\phi) \stackrel{\text{def}}{\Longleftrightarrow} \exists a\langle a_i\rangle\phi$$

so a state ϕ is directly accessible if the agent has an action that can bring about ϕ. A state may not be directly accessible, even though it may be accessible via another, directly accessible state. Define accessibility:

$$\mathbf{A}_i(\phi) \stackrel{\text{def}}{\Longleftrightarrow} \mathbf{DA}_i(\phi) \vee (\mathbf{DA}_i(\psi) \wedge (\psi \rightsquigarrow \phi))$$

so a state ϕ is accessible if it is either directly accessible or if another state is directly accessible that leads to ϕ. Following von Wright, we define a "good" state ϕ as a state that agent i prefers to its negation, and conversely for a bad state:

$$\mathbf{GO}_i(\phi) \stackrel{\text{def}}{\Longleftrightarrow} (\phi\mathbf{P}_i\neg\phi) \qquad \mathbf{BA}_i(\phi) \stackrel{\text{def}}{\Longleftrightarrow} (\neg\phi\mathbf{P}_i\phi)$$

5 Goals

ALX allows for the definition of various goal-operators. Based on the notion of preferences, our operators eschews the counter-intuitive behavior that primitive goal-operators tend to exhibit (e.g., necessitation and closure under implication). We present four definitions in this section.

Agents might opt for a state simply because it is better than its negation, particularly if only few alternatives are considered. A "good" goal can be defined by using the "good" operator $\mathbf{GO}$. Let $\mathbf{G}_i^g\phi$ denote the fact that ϕ is a good goal of i. Define: $\mathbf{G}_i^g\phi \stackrel{\text{def}}{\Longleftrightarrow} \mathbf{GO}_i\phi$. Thus a good goal is a situation that is preferred to its negation.

The second definition involves a satisficing goal. "Satisficing" has been an important procedural procedural addendum to the conceptualization of bounded rationality. Whereas the declarative part of bounded rationality concerns incomplete knowledge, the procedural part concerns the question what to do when the knowledge is not complete enough. So, the definition should involve the action of satisficing. Let $\mathbf{S}_i\phi$ stand for an arbitrary situation ϕ satisficing agent i and relax the definition of action terms by allowing for mnemonic expressions: $\mathbf{S}_i\phi \Leftrightarrow \langle \textit{satisficing-search}\rangle\phi$.

Note that this definition does not exclude the possibility that the search action is void in cases that the satisficing solution is already at hand. Define a satisficing goal in terms of a satisficing state. Let $\mathbf{G}_i^s\phi$ denote the fact that ϕ is a satisficing goal for i. Define: $\mathbf{G}_i^s\phi \stackrel{\text{def}}{\Longleftrightarrow} \mathbf{S}_i\phi$.

Agents may try to maximize, or even optimize, if the context supports the search for extremal values. Both attitudes involve some total view of available alternative and thus require quantification over situations. In the general case, this quantification would push us beyond the syntactic boundaries of ALX3, and undermine ALX's completeness, since the set of situations need not be countable. However, we can avoid this complication via some technical means, at least as far as preference formulas are involved (See [Hua94]). Define an element of i's preference order as follows:

$$\mathbf{PO}_i\phi \Leftrightarrow \phi \boldsymbol{P}_i \psi \vee \psi \boldsymbol{P}_i \phi$$

Whether a solution is maximal or optimal depends, of course, on the structure of the preference order of an agent. If it is partial, but not total, the order may contain several maximal, incomparable elements. If, furthermore, more than one maximal element is accessible, then an optimal goal in the sense of a best overall solution cannot be defined. Conversely, if the order is total, then an optimal goal can be identified (under the assumption that at least one situation occurring in the order is accessible). Alternatively, a partial order may give rise to an optimal goal if only one maximal situation is accessible. In the light of this reasoning, we define a "best choice" as a maximal goal and specify the conditions under which such a best choice may, in fact, be optimal. Let $\mathbf{G}_i^{ma}\phi$ denote the fact that ϕ is a maximal goal. Define: $\mathbf{G}_i^{ma}\phi \stackrel{\text{def.}}{\Longleftrightarrow} \mathbf{PO}_i\phi \wedge \forall\chi(\chi\mathbf{P}_i\phi \rightarrow \neg\mathbf{A}_i\chi)$.

A best choice is an accessible situation to which no other accessible situation is preferred. A best choice ϕ is optimal, if ϕ is unique. Let $\mathbf{G}_i^{op}\phi$ denote the fact that ϕ is optimal. Define: $\mathbf{G}_i^{o}\phi \stackrel{\text{def.}}{\Longleftrightarrow} \mathbf{G}_i^{bc} \wedge \forall\psi((\mathbf{G}_i^{bc}\psi) \rightarrow (\phi \leftrightarrow \psi))$

Ironically, best choices need not be good nor satisficing. In a tight spot, an agent's best alternative might simply be the best among dubious alternatives.

6 Conclusions and Summary

ALX is the first action logic modeled on the decision cycle of potentially rational agents. Perhaps its most important feature is its preference operator [HMP92]. The preference operator has a closest-world semantics in combination with the conjunction expansion principle, so agents prefer ϕ to ψ if they prefer the closest ϕ-and-not-ψ worlds to the closest ψ- and-not-ϕ worlds. This facilitates the representation of situation-dependent preferences, but does not exclude the representation of stable preferences. The preference operator is "normal" in the sense that agents cannot have preferences with respect to absurd worlds. This normality protects the conjunction expansion principle against counterintuitive utilizations [CS66]. On the other hand, the preference operator is not "tarskian", i.e., it does not distribute over disjunction [Mar95] and this protects all intensional operators built on the preference operator against the necessitation rule and against undesired closure properties. For example, goal operators can be defined as preferred states subject to additional qualifications (e.g., the best accessible state, the best state not believed to be inaccessible). The preference operator is transitive, but a non-transitive version is easily generated by removing the constraint (TRAN) on the semantic preference relation. The conditional operator is adapted from Stalnaker's system but allows for non-unique closest worlds. It allows for an easy representation of the notion of "ceteris paribus", and hence for

the standard notion of causality. This, in turn, greatly facilitates the representation of causal effects, side effects, and similar non-monotonic relations that would have otherwise to be represented by the (monotonic) material conditional.

An important property of ALX is the virtual absence of interaction between modal operators. This property may raise eyebrows in philosophical circles, but we designed ALX as a flexible knowledge-representation tool, and such a tool should, in our view, not preempt the structure of domains to be represented. As a flexible tool, ALX3 allows for the definition of various additional intensional operators, such as the alethic modalities, goal operators, intention-operators, and for the characterization of other action-related notions, such as ability, effect, side effect, intended effect, etc [HM94].

As demonstrated in a companion paper [MH94], ALX3 serves already as a versatile tool of knowledge representation. However, there are some desiderata left to be satisfied. ALX3 has no explicit notion of time in its semantics, and this complicates a direct representation of events. For example, we cannot define a "do"-operator, and hence no actions that are not deliberate. Second, we think that the notion of closest worlds needs closer inspection. The constraints on the closest world function are relatively weak; stronger constraints may be required, in particular if one wants to combine the notion of action with the notion of the closest world. Third, and related, one may want to strengthen ALX so that it allows for a *calculation* of causal outcomes. Such a calculation would be an answer to the frame problem, but it requires a more specific notion of possible worlds.

References

[Blu69] Blumer, H., *Symbolic Interactionism: Perspective and Methods*, (Englewood Cliffs, NJ, Prentice-Hall, 1969).

[CS66] Chisholm, R., and Sosa, E., Intrinsic preferability and the problem of supererogation, *Synthese* **16** (1966), 321-331.

[CL90] Cohen, P. R. and Levesque, H. J., Intention is choice with commitment. *Artificial Intelligence* **42** (1990) 213-261.

[Gam90] Gamut, L.T.F., *Logic, Language, and Meaning*, (The University of Chicago Press, 1991).

[Gid79] Giddens, A., *Central Problems in Social Theory: Action, Structures, and Contradiction in Social Analysis*, (Berkeley, CA, University of California Press, 1979).

[GS87] Ginsberg, M., and Smith, D., Reasoning about action I: a possible worlds approach, in: *Readings in Non-monotonic Reasoning*, (Morgan Kaufman, 1987).

[Har84] Harel, D., Dynamic logic, in: *Handbook of Philosophical Logic*, Vol.II, (D. Reidel Publishing Company, 1984) 497-604.

[HMP92] Huang, Z., Masuch, M., and Pólos, L., Een preference-logica voor rationele handelingen, in: H. de Swaan Arons, et al. (eds.), *NAIC'92 Conferentie Proceedings*, 17-28.

[HMP95] Huang, Z., Masuch, M., and Pólos, L., ALX, an action logic for agents with bounded rationality, *Artificial Intelligence* (forthcoming).

[HMP93] Huang, Z., Masuch, M., and Pólos, L., ALX2, a quantifier ALX logic, CCSOM Working Paper 93-99.

[Hua94] Huang, Z., *Logics for Agents with Bounded Rationality*, ILLC Dissertation series 1994-10, University of Amsterdam, (1994).

[HM94] Huang, Z., and Masuch, M., ALX3, a Multi-agent Action Logic, CCSOM Working Paper 94-102.

[MH94] Masuch, M., and Huang, Z., A Logical Deconstruction: Formalizing J.D. Thompson's *Organizations in Action* in a Multi-agent Action Logic, CCSOM Working Paper 94-120.

[Mar95] Marx, M., *Algebraic Relativization and Arrow Logic*, ILLC Dissertation series 1995-3, University of Amsterdam, 1995.

[Par37] Parsons, T., *The Structure of Social Action*, (Glencoe, IL, Free Press, 1937).

[RG91] Rao, A. and Georgeff, M., Modeling rational agents within a BDI- architecture, in: *Proceedings of the 1991 KR Conference*, (Morgan Kaufmann Publishers, 1991) 473-484.

[Sta68] Stalnaker, R., A theory of conditionals, in: *Studies in Logical Theory, American Philosophical Quarterly* **2** (1968) 98-122

[Sim55] Simon, H. A., A behavioral model of rational choice, *Quarterly Journal of Economics* **69** (1955) 99-118.

[vWr63] von Wright, G. H., *The Logic of Preference*, (Edinburgh, 1963).