A Preference-based Interpretation of Other Agents' Actions

Jérôme Lang
IRIT/UPS
F-31062 Toulouse Cedex, France
lang@irit.fr

Abstract

What can we infer about the beliefs, the goals and the intended next actions of an agent by observing him acting? We start from the same intuitive principle as in (Brafman & Tennenholtz 1997), namely, that an agent normally follows some optimal plan (with respect to some preference relation) and chooses actions that are part of this plan. This principle allows us to infer beliefs about the agent's plausible next actions, about his goals and about his initial and current beliefs. We give several possible definitions and study their respective meanings, especially with respect to the assumptions about the acting agent's rationality.

Introduction

Understanding the mental state of an agent (his goals, his beliefs) and being able to predict his future actions is a crucial issue in many situations, from cooperative problem solving to man-machine interaction. In this article we start from a similar perspective as in (Brafman & Tennenholtz 1997), who interpret an agent's actions by ascribing him a mental state (beliefs and goals), using the decision-theoretic assumption that agents follow optimal plans (in a given sense to be defined); this mental state can then be used to predict the actor's next actions. We then depart from (Brafman & Tennenholtz 1997), whose level of generality is very high, and study mental state ascription and action prediction in a much more specific case where agents are goal-seeking (i.e., they have binary preferences) and their decision strategy merely consists in choosing a nondominated plan in the set of solution plans according to a given preference relation over plans (this difference in the level of generality allows us to focus on different problems).

To state things more precisely, we consider two rational agents: the first one, called the *observer* (to make things simpler, the observer will often be referred to as *we*, that is, we put ourselves in the observer's place), observes the behaviour of the second one, called the *actor*, without interfering with him. We assume furthermore, so as to avoid considering strategic interactions, that the actor is indifferent to the observer (and *vice versa*), so his behaviour is not influenced by the fact that he is observed. We can assume, to make these questions irrelevant, that the actor *does not know* he is observed. Observing the actor performing a subplan allows us to infer some plausible beliefs about his goals, intended next actions, and beliefs (the latter in the case where he may have an incomplete knowledge of the environment).

This paper studies several different ways of inferring such beliefs, compares them and discusses their meaning. To make the presentation simpler, we first consider the simpler case where the actor has complete knowledge of the environment; in this case, inference bears on goals and intended actions (not on beliefs). Then we extend the framework to the incomplete knowledge case, and conclude by pointing to related work and further issues.

The complete knowledge case

Domains and contexts

In this part, the environment is completely known by the actor (and the observer is aware of that): he knows the state of the world at each step and actions are deterministic. To make things simple, the observer is also assumed to have a complete knowledge of the environment (relaxing this assumption would lead to interesting issues where the observer learns some properties of the environment by observing the agent, but this is left for further study). Therefore, two types of information may be ignored or ill-known by the observer:

- the actor's goals;
- his intended future actions, i.e., the continuation of his plan.

We start by fixing the domain and its dynamics are fixed (which are perfectly known by the actor and the observer).

Definition 1 (domains and plans)

- *a* domain D *is composed of*
  - *a finite set of states S;*
  - *a finite set of actions* $ACT = \{\alpha_1, \ldots, \alpha_n\}$*, containing the void action* λ_{ACT}*; their dynamics is described by means of a transition function* $next : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$*, where* $next(s, \alpha)$ *is the successor state of* s *when action* α *has been performed (for the sake of simplicity, we do not consider the case where actions can be inexecutable in some states). The dynamics of the void action* λ_{ACT} *is described by* $next(s, \lambda_{ACT}) = s$ *for all* s.

- *a plan π is a finite sequence (possibly empty) of actions, that is, a member of ACT^*:*

$$\pi = \langle \pi(t) \mid 0 \leq t \leq T \rangle$$

$\pi; \pi'$ denotes the concatenation of π and π'. λ is the empty plan. The function $next$ is extended in the usual way to plans by $next(s, (\alpha; \pi)) = next(next(s, \alpha), \pi)$.
- *A plan π_{done}* starts *a plan π if and only if there exists $\pi' \in ACT*$ such that $\pi = \pi_{done}; \pi'$.*
- *$Develop(\pi_{done}) = \{\pi \in ACT^* \mid \pi_{done}$ starts $\pi\}$ is the set of plans starting by π_{done} (called the* developments *of π_{done}).*

Definition 2 (contexts)

- *Given a fixed domain D, a* context *over D is a triple*

$$C = \langle s_0, g, \pi_{done} \rangle$$

where $s_0 \in \mathcal{S}$ is the initial state of the world, $g \subseteq \mathcal{S}$ is the actor's set of goal states, and $\pi_{done} \in ACT^*$ is the partial plan already performed by the actor.
- *$\pi \in ACT^*$ is a* solution plan *for C if and only if $next(s_0, \pi) \in g$. $SolPlans(C)$ denotes the set of solution plans for C.*

Saying that the goal of the actor is g means that he wants to reach whatever state of g (the goal is satisfied as soon as this is realized) – the goal *does not consist* in reaching *all* states of g in sequence.

For the sake of simplicity, we chose to define goals as a crisp (all-or-nothing) notion. A more general model would consist in considering an ordering relation or a utility function on states, rather than a simple partition between goal states and non-goal states. However, as developed in the following Section, noncrisp preference will bear on plans (rather on states).

Preference between plans

Decision theory views action selection as a maximization process: agents choose actions that are non-dominated with respect to some preference relation. Accordingly, we assume that solution plans are ordered by a *preference relation*.

Notice however that while the usual assumption in classical decision theory is that the preference relation underlying the choices of a rational agent are not only transitive but complete, we do not commit to the completeness assumption here and thus allow for incomparabilities between plans. There are two possible possible interpretations of incomparability: the first one is a *(partial) ignorance of the actor about his true preferences*, which may be explained by the fact that his limited resources do not let him enough deliberation time to decide which one of both plans he prefers; the second interpretation is a *(partial) ignorance of the observer about the actor's preferences*. Both interpretations will be discussed further in the Section "Relevance of an action in a context".

Definition 3 (preference relations) *A preference relation $\succeq$ is a reflexive and transitive relation (not necessarily complete). By convention, $\pi \succeq \pi'$ means that π is at least as good as π'. $\succ$ is the strict preference relation induced by $\succeq$: $\pi \succ \pi'$ if and only if $\pi \succeq \pi'$ and not $(\pi' \succ \pi')$. Given a set of plans P, $\pi \in P$ is* non-dominated *in P w.r.t. $\succeq$ if and only if there is no $\pi' \in P$ such that $\pi' \succ \pi$. $Max_{\succeq}(P)$ denotes the set of non-dominated plans in P w.r.t. $\succeq$.*

Given a goal g and a preference relation $\succeq$, we define a *preferred* plan as an non-dominated solution plan, that is, π is preferred if and only if π is non-dominated in $SolPlans(C)$.

Since, for the sake of simplicity, goals have been defined as an all-or-nothing notion, we assume that the actor prefers any solution plan to any non-solution plan (if the actor's preferences were described more generally by a utility function over S, as in (Brafman & Tennenholtz 1997), $\geq$ would have to be defined accordingly).

Among the many possible choice of preference relation, we give a few (that are representative of broader classes).

- for all π, π', $\pi \succeq^b \pi'$ is false (*basic* preference).
- for all π, π', $\pi \succeq^c \pi'$ if and only if $|\pi| \leq |\pi'|$;
- for all π, π', $\pi \succeq^o \pi'$ if and only if $\forall \alpha \in ACT$, $occ(\alpha, \pi) \leq occ(\alpha, \pi')$, where $occ(\alpha, \pi)$ is the number of occurrences of α in π ;
- for all π, π', $\pi \succeq^i \pi'$ if and only if there is a function $\sigma : \{1, \ldots, |\pi|\} \rightarrow \{1, \ldots, |\pi'|\}$, strictly increasing (for all i, $\sigma(i+1) > \sigma(i)$) such that $\pi(i) = \pi'(\sigma(i))$.

Theferore:

- preferred plans w.r.t. $\succeq^b$ are *all* solution plans; one interpretation of $\succeq^b$ is that we do not have any idea of the actor's preference between solution plans.
- preferred plans w.r.t. $\succeq^c$ are *shortest* solution plans (this preference relation can be generalized in a straightforward way to plans of minimal cost, where a positive cost is associated with each action);
- preferred plans w.r.t. $\succeq^i$ are those which do not contain useless or redundant actions: indeed, $\pi \succ^i \pi'$ iff π' contains all actions of π in the same order, plus at least another action. $\succeq^o$ is an intermediate between $\succeq^i$ and $\succeq^c$. For instance, $\pi = \langle a, b, a, c \rangle$ is preferred to $\pi' = \langle b, a, c, b, b, a, c \rangle$, while $\pi = \langle a, b, a, c \rangle$ and $\pi' = \langle a, a, b, b, c, c \rangle$ are incomparable.
- preferred plans w.r.t. $\succeq 0$ are those minimizing the number of action occurrences action by action.

Notice that $\succeq^i$ and $\succeq^o$ are generally not complete. One easily checks that the following implications hold:

$$\pi \succ^i \pi' \Rightarrow \pi \succ^o \pi' \Rightarrow \pi \succ^c \pi'$$
$$\pi \succeq^i \pi' \Rightarrow \pi \succeq^o \pi' \Rightarrow \pi \succeq^c \pi'$$

Other preference relations could be defined. In particular, we may know some partial preferences between actions (or action sequences) *ceteris paribus*: for instance, saying that *ceteris paribus*, the actor prefers action a to action b means that any two plans π and π' such that $\pi = \pi_1; a; \pi_2$ and $\pi' = \pi_1; b; \pi_2$ are such that $\pi_1 \succ \pi_2$. This resembles the *ceteris paribus* interpretation of goals as in (Doyle & Wellman 1991), with the noticeable difference that

Table 1: Synthesis of results for Example 1

situation	goal	s_0	π_{done}	R-relevance $\succeq^c$	R-relevance $\succeq^i$	U-relevance $\succeq^c$	U-relevance $\succeq^i$
1	G	A	λ	D	C, D, F	D	C, D, F
2	G	A	`go to` D	G	A, C, G	G	G
3	E	A	`go to` F	A	A, H	none	H
4	G	A	`go to` B	A	A	none	none
5	G	A	`go to` D; `go to` C	D, E	D, E	none	none

preference here bears on action sequences (in which the order counts) and not on static goals. Notice that the $\succeq^i$ preference relation is actually a particular case of *ceteris paribus* preference relation between plans induced by the constraints $\lambda_{ACT} \succeq a$ for each action $a \neq \lambda_{ACT}$.

We now define *two* notions of non-dominated plans given a context, which will be crucial for the rest of the article.

Definition 4 *Let C be a context over a domain D.*

relativized non-dominance

$$\begin{aligned} & Pref_R(C, \succeq) \\ = \; & Max_{\succeq}(SolPlans(C) \cap Develop(\pi_{done})) \end{aligned}$$

unrelativized non-dominance

$$\begin{aligned} & Pref_U(C, \succeq) \\ = \; & Max_{\succeq}(SolPlans(C)) \;\cap\; Develop(\pi_{done}) \end{aligned}$$

Therefore, relativized non-dominated plans for C (w.r.t. $\succeq$) are non-dominated solution plans *among the developments of* π_{done} while unrelativized non-dominated plans *are non-dominated solution plans that start by* π_{done}. Obviously, we have $Pref_U(C, \succeq) \subseteq Pref_R(C, \succeq)$, that is, the "unrelativized" notion filters out more plans than the "relativized" one.

Relevant actions in a context

Intuitively, an action is a plausible (or *relevant*) in a context C w.r.t. a preference relation $\succeq$, if there exists a non-dominated solution plan starting by π_{done} and whose action following π_{done} is α. When the actor's set of goal states is known, relevant actions can be identified with the actor's plausible next actions in the given context. Now, the two notions of non-dominance induce two notions of relevance.

Definition 5 (relevance)
An action α is R-relevant *(resp.* U-relevant*) in C w.r.t. $\succeq$ if and only if there exists a solution plan π in $Pref_R(C, \succeq)$ (resp. in $Pref_U(C, \succeq)$) such that $\pi_{done}; \alpha$ starts π .*
For any context C and any preference relation $\succeq$, let $Rel^{\succeq}_R(C)$ and $Rel^{\succeq}_U(C)$ be the set of R- (resp. U-) relevant actions in C for $\succeq$.

Since there are generally several (R- or U-) non-dominated plans for a context, there may be *several* (R- or U-) relevant actions for C w.r.t. $\succeq$. There may also be *no* (R- or U-) relevant action for C w.r.t. $\succeq$, since there may be no solution plan starting by π_{done} .

Before investigating the meaning of both definitions, let us see how they work on an example.

Example 1 *Consider the following non-oriented graph. We assume that the actor wants to reach a single node of the graph (the goal is therefore a singleton) and that he has complete and accurate knowledge of the graph and of his current position.*

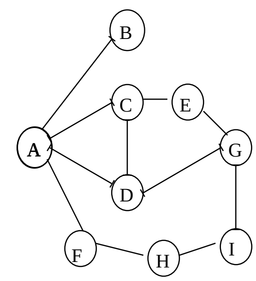

We consider several different situations. In all of them we let S= $\{A, \dots, I\}$ and A= \{ `go to` *A, ...,* `go to` *I\}*[1].

Situation 1 *$g = \{G\}$ and $\pi_{done} = \lambda$*
(the actor has not started to move from A).

- *the unique shortest plan from A to G being ADG, only the action* go to D *is (both U- and R-) relevant in this context w.r.t. $\succeq^c$;*
- *going to C, to D, or to F are (both U- and R-) relevant w.r.t. $\succeq^i$, because the solution plans $ACEG$, ADG and $AFHIG$ are all nondominated;*
- *any action is relevant w.r.t. $\succeq^b$, since after performing any action, it is still possible to reach G. Notice that if the graph were oriented and if there were an extra node $sink$ in the graph, accessible from A and such that no other node than $sink$ is accessible from $sink$, then going to $sink$ would not be relevant for $\succeq^b$.*

Situation 2 *$g = \{G\}$ and $\pi_{done} =$* `go to` *D*
In this context, `go to` *G is both the only R-relevant and the only U-relevant action w.r.t. $\succeq^c$. W.r.t. $\succeq^i$, going to G is still the only U-relevant action, since among the three $\succeq^i$-nondominated plans leading from A to G, only ADG starts by going to D, and the next action in this plan is going to G*[2]*. However, R-relevance gives three relevant actions: G, but also C and even going back to A, since*

[1] An alternative choice would consist in distinguishing two actions with the same destination but different origins. This would lead to different results when $\succeq^i$ or $\succeq^o$ is used.

[2] Notice that, if instead of considering that the set of actions is *goto* A, *goto* B, etc., we consider going to a same node from two different nodes as two separate actions – i.e., that there is one action for each *edge* of the graph (*go from A to B*, etc.), then there would be a fourth nondominated plan leading from A to G, namely

$ADAFHIG$ is undominated in the set of solution plans starting by AD.

Situation 3 $g = \{E\}$ *and* $\pi_{done} = \texttt{go to } F$.
W.r.t. $\succeq^i$, going to H is the only R-relevant action while going to back to A is U-relevant as well (for similar reasons as above). As for $\succeq^c$, the only U-relevant action is going back to A, since once the actor is in F, the shortest plan to reach E is $AFACE$; since this plan is dominated (by ACE) in the set of all solution plans, there is no R-relevant action *in this context.*

Situation 4 $g = \{G\}$ *and* $\pi_{done} = \texttt{go to B}$.
The only G-relevant action w.r.t. $\succeq^i$ or $\succeq^c$ is going back to A; there is no R-relevant action w.r.t. $\succeq^i$ nor $\succeq^c$.

The results are synthesized in Table 1.

We now give intuitive explanations about how relevance works; by default, the case chosen for illustration is situation 3, that is, $g = \{E\}$ and π_{done} = $\texttt{go to } F$.

Let us start by the difference between $\succeq^i$ and $\succeq^c$; even if other preference relations could have been chosen, comparing these two "typical" preference relations is interesting as such. Choosing $\succeq^c$ means that the actor prefers shorter plans *and that he is rational enough to find them*. There are at least two interpretations for $\succeq^i$: it may mean that the actor is perfectly rational but that he has no preference for shorter plans; or it may mean that although he prefers shorter plans, his limited rationality implies that he does not always find them.

The latter interpretation would be even better modelled with a sophisticated preference relation integrating plan length (or cost) and inclusion, thus intermediate between $\succeq^i$ and $\succeq^c$. Here is one, where $0 < \rho < 1$ is fixed.

$$\pi \succeq \pi' \text{ if and only if } \pi \succeq^i \pi' \text{ or } \tfrac{|\pi|}{|\pi'|} \leq \rho$$

Note that fixing $\rho = 0$ leads back to $\succeq^i$, whereas fixing $\rho = 1$ leads back to $\succeq^c$. We do not pursue this here in this paper.

There exists a third possible interpretation of $\succeq^i$: edges have costs, the actor prefers plans of minimal cost, but *we* do not know these costs – we only know that they are strictly positive. This interpretation explains why we may consider more plausible next actions with $\succeq^i$ than with $\succeq^c$: since we know less about the agent's state of mind, more possibilities must be considered as to his future behaviour. This ignorance becomes maximal with $\succeq^b$; accordingly, any action which does not make the goal unreachable is considered(both R- and U-) relevant with $\succeq^b$.

Let us now explain the difference between relativized and unrelativized relevance. The assumption underlying R-relevance is that the actor may deliberate during plan execution, which in some cases allows him to discover that his previously intended plan was suboptimal, as when he first goes to F: with $\succeq^c$, we believe that he will discover his mistake and then go back to A. On the contrary, U-relevance assumes that the deliberation phase and the acting phase are totally disjoint: once the actor has started to execute his plan, he sticks to it without questioning it. This difference appears clearly with $\succeq^i$: R-relevance allows for the actor to go back to A, which is not suboptimal *given that he has already moved to F*, while U-relevance excludes A: the actor initially chose $AFHIGE$ as his intended plan (her intended plan cannot be $AFACE$ nor $AFADGE$, which are dominated respectively by ACE and $ADGE$ *in the set of all solution plans*), and he will stick to it and go to H. Lastly, with $\succeq^c$, no action is U-relevant because (a) we assumed that the actor was rational enough to compute the best plan(s) and (b) we realize that he does not follow any of these. This is even clearer when the goal is G and $\pi_{done} = \texttt{go to } B$: w.r.t. both $\succeq^c$ and $\succeq^i$, we refuse to consider the subplan π_{done}, because it does not start any optimal plan (we therefore give up and conclude that the agent is not rational, and are not willing to predict anything).

Relativized relevance shows a nonmonotonic behaviour, in the sense that later actions may lead us to reconsider the agent's followed plan, as with $\succeq^c$ when the goal is H: before the actor starts to move ($\pi_{done} = \lambda$), we expect him to go to F because we believe his intended plan to be AFH; we revise this after observing him go to D and then to G, and we now expect him to go to next to I. On the contrary, unrelativized relevance is truly monotonic, because the set of plans we consider to be possible intended plans shrinks (until becoming eventually empty in some cases) as we observe more actions. Looking back to the example, it seems that U-relevance gives results that are more intuitive, but its drawback is that it may give empty sets of plausible actions while R-relevance does give some *intended* answers. We may thus think in using U-relevance except in this case, in which we use R-relevance:

$$Rel^{\succeq}_{UR}(C) = \begin{cases} Rel^{\succeq}_{U}(C) & \text{if } Rel^{\succeq}_{U}(C) \neq \emptyset; \\ Rel^{\succeq}_{R}(C) & \text{otherwise} \end{cases}$$

The following properties help to better understand how the notions work.

Proposition 1

1. *for any C and any $\succeq$, $Rel^{\succeq}_{U}(C) \subseteq Rel^{\succeq}_{R}(C)$;*
2. *for any C and any $\succeq$, $Rel^{\succeq}_{R}(C) = \emptyset$ if and only if there is no plan leading from $next(s_0, \pi_{done})$ to g (notice that this property does not hold for R-relevance);*
3. *if $\succeq_1$ and $\succeq_2$ are such that for any π, π', $\pi \succ_2 \pi'$ implies $\pi \succ_1 \pi'$, then for any C, $Rel^{\succeq_1}_{R}(C) \subseteq Rel^{\succeq_2}_{R}(C)$ and $Rel^{\succeq_1}_{U}(C) \subseteq Rel^{\succeq_2}_{U}(C)$.*

This may be summarized by the following inclusion graph, when edges represent inclusion. Note that generally, $Rel^{\succeq_1}_{R}(C)$ and $Rel^{\succeq_2}_{U}(C)$ are incomparable.

$$\begin{array}{ccc} Rel^{\succeq_1}_{R}(C) & \rightarrow & Rel^{\succeq_2}_{R}(C) \\ \uparrow & & \uparrow \\ Rel^{\succeq_1}_{U}(C) & \rightarrow & Rel^{\succeq_2}_{U}(C) \end{array}$$

Corollary 1 *For any $X \in \{G, R\}$ and any context C, we have*

$$Rel^{\succeq^c}_{X}(C) \subseteq Rel^{\succeq^i}_{X}(C) \subseteq Rel^{\succeq^o}_{X}(C) \subseteq Rel^{\succeq^b}_{X}(C)$$

$ADCEG$, and in this case going to C would be both R-relevant and U-relevant for $\succeq^i$.

Unsurprisingly, more plausible inferences are allowed with $\succeq^c$ than with $\succeq^i$ and with $\succeq^i$ than with $\succeq^b$: indeed, the assumption that the actor always follows shortest plans amounts to assuming that he has a perfect rationality (because he never fails finding these shortest plans). Choosing one of the other preference relations amounts to assuming that the actor has a limited rationality (more and more limited as we tend towards $\succeq^b$). The extreme case would be that of an irrational agent for whom *all* actions would be considered as relevant, even those which prevent him for reaching his goal. $\succeq^i$ is particularly interesting: by eliminating plans containing obvious redundancies or obviously useless actions *but* not all suboptimal plans, we ascribe to the actor a *reasonably good, but limited rationality*: he may fail to find a shortest plan, but he is rational enough so as to avoid performing a "nonsensical" plan. Therefore, *choosing a preference relation and a notion or relevance (relativized or unrelativized) both involve implicit assumptions about the rationality of the actor*: perfect if we choose $\succeq^c$, more limited with $\succeq^o$ and with $\succeq^i$, and extremely limited with $\succeq^b$.

When we (the observer) know the goal of the actor, we can reasonably identify the plausible next actions and the relevant actions for C w.r.t. $\succeq$; thus, both notions of relevance allow for plausible inferences (more cautious with the unrelativized notion of relevance than with the relativized one) about the actor's intended actions. The next Section investigates the case where the actor's goal is only partially known by the observer.

Inferring plausible goals

We assume now that the actor's goal is ill-known (possibly totally unknown) by the observer. Let $\pi_{done} = \langle \pi_{done}(0), \pi_{done}(now-1)$, where $now = |\pi_{done}|$. The last action of π_{done}, denoted by $last(\pi_{done})$, is $\pi_{done}(now-1)$. For any integer $t \leq now$ we note $Restr(\pi_{done}, t)$ the plan consisting of the t first actions of π_{done}, that is,

$$Restr(\pi_{done}, t) = \langle \pi_{done}(0), \pi_{done}(t-1) \rangle$$

An *objective context OC* is a context in which, instead of the goal, we have a set of possible goals Γ (that is, a set of nonempty sets of states) expressing a partial knowledge of the actor's goal: $OC = \langle s_0, \Gamma, \pi_{done} \rangle$. Objective contexts differ from (intentional) contexts as defined in the previous Section; the terminology "objective" is based on the assumption the state and steps taken are visible to observers but the goal is not. In this Section, to avoid any ambiguity, we refer to contexts as defined in the previous Section, that is, triples $\langle s_0, g, \pi_{done} \rangle$ as *intentional contexts*.

In Example 1, if all the observer knows is that the actor wants to reach a single node, then Γ is the set of all singletons. In practice, Γ may be compactly described by a propositional formula.

If OC is an objective context and $g \in \Gamma$, then we define the intentional context $PC + g = \langle s_0, g, \pi_{done} \rangle$. The previous notions of relevance can be used to make plausible inferences about the actor's goal. A first definition consists in saying that a set of states g is a plausible goal if and only if any action performed until now is relevant for the intentional context corresponding to s_0, π_{done} and g.

Definition 6 *g is a R-plausible (respectively U-plausible) goal given an objective context $OC = \langle s_0, \Gamma, \pi_{done} \rangle$ (w.r.t. $\succeq$) if and only if the following two conditions are satisfied:*

1. *$g \in \Gamma$;*
2. *for any $t \in \{0, \ldots, now-1\}$, the action $\pi_{done}(t)$ performed by the actor at time t is R-relevant (respectively U-relevant) for the intentional context $\langle s_0, g, Restr(\pi_{done}, t) \rangle$ (w.r.t. $\succeq$).*

Let $Goals_R(OC, \succeq)$ (resp. $Goals_U(OC, \succeq)$) be the set of R-plausible (respectively U-plausible) goals for OC w.r.t. $\succeq$.

Since the two notions of relevance are different, we may expect that the definitions of R-plausible and U-plausible goals are radically different. Somewhat unexpectedly however, even if the two notions of relevance differ, R-plausible and U-plausible coincide when $\succeq$ is a *complete* preference relation:

Proposition 2 *If $\succeq$ is complete then $Goals_R(OC, \succeq) = Goals_U(OC, \succeq)$).*

In the general case, we only have

$$Goals_U(PC, \succeq) \subseteq Goals_R(PC, \succeq)$$

Moreover, if $\succ_2 \subseteq \succ_1$ then $Goals_U(OC, \succeq_1) = Goals_R(OC, \succeq_2)$.

The following result gives an intuitive equivalent formulation for U-plausible goals:

Proposition 3 *$g \in \Gamma$ is an U-plausible goal given OC (w.r.t. $\succeq$) if and only if there exists a plan $\pi \in SolPlans(OC + g) \cap Develop(\pi_{done})$ such that there is no plan $\pi' \in SolPlans(OC + g)$ verifying $\pi' > \pi$.*

Thus, g is U-plausible iff there is a non-dominated plan for g starting by π_{done}. It may be the case that there exists no R- or U-plausible goal for a given partial context (see for instance Example 1 when the plan already performed consists in going from A to D and then to C). This happens in particular when the actor has a suboptimal behaviour with respect to the chosen preference relation (*i.e.*, when his rationality is more limited then we expected).

These definitions are strong, since they do not consider "nearly" optimal plans for g; but on the other hand, this avoids introducing complications the technicities of which we would like to avoid in this paper (and moreover, the assumption that the rationality of the actor is limited is already expressed by the choice of the preference relation). We just sketch here the principle of the solution. Clearly, the notion of plausible goal has to become a *relative* notion, *i.e.*, a relation "is at least plausible as a goal as (given OC)" over $2^{\mathcal{S}} \setminus \emptyset$. *However,* so as to derive such a relation, a simple preference relation $\succeq$ on plans is not enough: one has to consider a preference relation on pairs (plan, goal) – (π, g) preferred to (π', g') meaning that π

Table 2: Synthesis of results for Example 2

	R/U-plausible π_{done} $\succeq^c$	R-plausible $\succeq^i$	U-plausible $\succeq^i$	weakly R-plausible $\succeq^c$	weakly R-plausible $\succeq^i$
$A \to D$	D, G, I	D, E, F, G, H, I	D, E, F, G, H, I	D, G, I	D, G, I
$A \to D \to C$	none	E	none	C, E	C, E
$A \to B \to A$	none	none	none	anything except B	anything except B

is more relevant with respect to reaching g than π' with respect to reaching g'. For instance, in the case where preference between plans is measurable, i.e., there is a cost function $K : ACT^* \to I\!R^+$ such that $\pi \succeq \pi'$ if and only if $K(\pi) \leq K(\pi')$ (which the case with $\succeq^c$ for $K(\pi) = |\pi|$), then we may adopt the following definition: let $\pi^*(g)$ a solution plan for g preferred w.r.t. $\succeq$, then we define $Q(\pi, g) = \frac{K(\pi)}{K(\pi^*(g))}$ (with the convention $\frac{0}{0} = 1$ and for all $x > 0$, $\frac{x}{0} = +\infty$); then, $R(g, C) = \min\{Q(\pi, g) \mid \pi \in SolPlans(C) \cup Develop(\pi_{done})\}$ and lastly: g is at least a plausible goal as g' (given C and w.r.t. $\succeq$) if and only if $R(g, C) \leq R(g', C)$. This direction will not be explored further in this paper.

We now introduce a third, weaker notion of plausible goals, based only on the *last* action performed by the actor.

Definition 7 *g is a* weakly *R-plausible goal given $OC = \langle s_0, \Gamma, \pi_{done}\rangle$ (w.r.t. $\succeq$) if and only if the following two conditions are satisfied:*

1. *$g \in \Gamma$;*
2. *the last action $last(\pi_{done}) = \pi_{done}(now - 1)$ performed by the actor is R-relevant for $\langle s_0, g, Restr(\pi_{done}, now - 1)\rangle$ (w.r.t. $\succeq$).*

Let $Goals_{WG}(OC, \succeq)$ be the set of wealy R-plausible goals for PC w.r.t. $\succeq$.

One might wonder why we did not introduce weak U-plausible goals. The reason is simple: applying the latter definition using U-relevance instead of R-relevance would lead to a notion of weak U-plausible goals equivalent to the (strong) notion of U-plausible goals (this is easy to check).

Example 2 *Let us consider again the graph of Example 1. Let $\Gamma = \{\ \{A\}, \{B\}, \ldots, \{I\}\ \}$ (all we know is that the actor wants to reach a single node).*

Situation 5 *$s_0 = A$ and $\pi_{done} =$ `go to` D.*
For $\succeq^c$, the R-plausible (and U-plausible since $\succeq^c$ is complete) goals are $\{D\}$, $\{G\}$ and $\{I\}$. These are also the only weakly R-plausible goals. We also infer plausible next actions conditioned by the goal, using relevance again: for instance, if the goal is $\{D\}$ (resp. $\{I\}$) then we expect the actor to stop (resp. to go to G). For $\succeq^i$, three other goals are R- and U-plausible, namely $\{E\}$, $\{F\}$ and $\{H\}$.

Situation 6 *$s_0 = A$ and $\pi_{done} =$ `go to` D; `go to` C.*
There is no U-plausible goal for neither for $\succeq^i$ nor for $\succeq^c$ (therefore, no R-plausible goal for $\succeq^c$ either). $\{E\}$ is R-plausible for $\succeq^i$. $\{C\}$ and $\{E\}$ are weakly plausible for $\succeq^c$ and $\succeq^i$.

Situation 7 *$s_0 = A$ and $\pi_{done} =$ `go to` B; `go to` A.*
There is no U-plausible nor any R-plausible goal for $\succeq^c$ and $\succeq^i$, while any goal except $\{B\}$ is weakly plausible for $\succeq^c$ and $\succeq^i$. If we choose $\succeq^c$ then three subsets of $\mathcal{S}$ are (R/U)-plausible goals: $\{E\}, \{F\}, \{E, F\}$. We also derive that if the goal is $\{E\}$ or $\{E, F\}$ then the actor intends to stop, and that is the goal is $\{F\}$ then he intends to go to F. Now, if we choose $\succeq^i$ then any subset of $\{C, D, E, F, G\}$ is a R-plausible goal.

The results are synthesized in Table 2.

Inference of (strong) R- and U-plausible goals is monotonic: with each new action performed by the actor, the set of his plausible goals gets no larger (until possibly becoming empty). Thus, with $\succeq^c$: any goal is initially plausible, after observing the actor going to D, the plausible goals are $\{D\}, \{G\}, \{I\}$, and the set of plausible goals eventually becomes empty after further observing the actor go to D. On the contrary, inference of *weak* R-plausible goals is nonmomotonic: when, still with $\succeq^c$, $\{D\}$, $\{G\}$ and $\{I\}$ (but not $\{C\}$ nor $\{E\}$) are weakly R-plausible after $\pi_{done} =$ `go to` D; when the actor is then seen to go to C, then $\{C\}$ and $\{E\}$ become weakly R-plausible.

If $\succeq$ is complete, a property of (R/U)-plausible goal inference is that the set of plausible goals for a given context is closed for set union:

Proposition 4 *Let $\succeq$ be a* complete *preference relation (i.e., for any π, π', either $\pi \succeq \pi'$ or $\pi' \succeq \pi$) and $g_1 \subseteq \mathcal{S}$, $g_2 \subseteq S$ two (R/U)-plausible goals for a given objective context OC w.r.t. $\succeq$. Then $g_1 \cup g_2$ is a (R/U)- plausible goal for OC w.r.t. $\succeq$.*

Therefore, if $\succeq$ is complete, the set of plausible goals for a given context can be expressed as the set of all unions of a set of elementary goals. This property does not hold if $\succeq$ is not complete. Here is a counterexample: $\mathcal{S} = \{a, b, c, d\}$, $\mathcal{A} = \{\alpha, \beta, \gamma\}$, and $next$ is described on Table .

	α	β	γ
a	b	d	c
b	b	c	d
c	c	c	c
d	d	d	d

Let $\succeq = \succeq^i$, $s_0 = a$, $\pi_{done} = \alpha$. $\{c\}$ is U-plausible, since $\alpha; \beta \in Pref(\succeq^i, SolPlans(\{c\}))$. $\{d\}$ is U-plausible, since $\alpha; \gamma \in Pref(\succeq^i, SolPlans(\{d\}))$. But $\{c, d\}$ is not U-plausible, because $Pref(\succeq^i, SolPlans(\{c, d\})) = \{\beta, \gamma\}$ and therefore contains no plan starting by α.

Another property that does not hold, even if $\succeq$ is complete, is

$$g_1 \text{ plausible and } g_2 \supseteq g_1 \Rightarrow g_2 \text{ plausible}$$

A counterexample is given using $\mathcal{S}$, $\mathcal{A}$ and $next$ as just above, $\succeq = \succeq^c$, $\pi_{done} = \alpha$: $\{b\}$ is a R/U-plausible but not $\{a, b\}$. Obviously, the property does not hold either if we replace $g_2 \supseteq g_1$ by $g_2 \subseteq g_1$ in the latter expression.

The incomplete knowledge case

We no longer assume that the actor has complete knowledge of the current state of the world; however, for the sake of simplicity and without loss of generality (see (Brafman & Tennenholtz 1997)), we still assume that actions are still deterministic. A context is now a triple $C = \langle b_0, g, \pi_{done} \rangle$ where b_0 is the initial *belief state* of the actor about the state of the world. We define a belief state as a nonempty subset of states: the set of belief states is $\mathcal{B} = 2^{\mathcal{S}} \setminus \{\emptyset\}$. (Notice that the rest of this Section could be rewritten so as to fit a quantitative view of incomplete knowledge, consisting typically in defining belief states as probability distributions. We omit it for keeping the exposition simple.)

For the sake of simplicity, we also assume that the actor's beliefs are accurate (that is, they can be incomplete but never wrong.) In addition to physical actions meant to have effects on the state of the world, the actor may also perform sensing actions so as to gather more knowledge on the state of the world, and he incorporates observations in his belief state by belief expansion. More generally, if the actor's beliefs were not accurate, he would have to perform a *belief revision* step; see (Brafman & Tennenholtz 1997) for a discussion.

To make the exposition simpler, we assume that (a) sensing actions are "pure", that is, they do not change the state of the world, and (b) the only sensing actions considered are *binary tests*: $test(\varphi)$ returns the truth value of the propositional formula φ describing some property of the current state of the world. Because of the feedback given by sensing actions, plans are now conditional.

In the following, we let $ACT = ACT_P \cup ACT_S$, where ACT_P (resp. ACT_S) is the set of physical (resp. sensing) actions available to the actor.

The set $CondPlans(ACT)$ of conditional plans for a set of actions ACT is defined recursively by:

- λ is in $CondPlans(ACT)$;
- for any action $\alpha \in ACT_P$, α is in $CondPlans(ACT)$;
- for any π , π' in $CondPlans(ACT)$, $\pi;\pi'$ is in $CondPlans(ACT)$;
- for any sensing action $\alpha = test(\varphi)$ and for all π , π' in $CondPlans(ACT)$, the plan α ; $\texttt{if}\ \varphi\ \texttt{then}\ \pi\ \texttt{else}\ \pi'$] is in $CondPlans(ACT)$.

The plan π_{done} already performed by the actor is merely the sequence of actions performed so far, plus the observation collected after each sensing action; we do not have to consider the other branches leaving from nodes corresponding to sensing actions, since the plan *has already been performed*: π_{done} is therefore unconditional. $\pi \in CondPlans(ACT)$ is a solution plan (in the strong sense) for C if and only if *all* possible executions of π lead to a state of g. We now have to extend the notion of a plan starting another plan. Let π_{done} be an unconditional plan and π a conditional plan. We say that π_{done} starts π if and only if at least one if these conditions is satisfied:

1. $\pi_{done} = \lambda$;
2. $\pi_{done} = (\alpha; \pi')$, $\pi = (\alpha; \pi'')$, and π' starts π'';
3. $\pi_{done} = (test(\varphi); \pi')$,
   $\pi = (test(\varphi);\ \texttt{if}\ \varphi\ \texttt{then}\ \pi_1\ \texttt{else}\ \pi_2)$,
   and π' starts either π_1 or π_2.

Preference relations between plans are generalizations of the preference relations given in the complete knowledge case. $\succeq^b$ does not have to be reformulated. $\succeq^c$ can be generalized in (at least) two ways: by comparing the lengths of their respective longest branches of the plans, or by comparing their average length (assuming equiprobability of initial states):

- $\pi \succeq^{mcb} \pi'$ if and only if the longest branch of π is no longer than the longest branch of π' .
- $\pi \succeq^{al} \pi'$ if and only if $\overline{L}(\pi) \leq \overline{L}(\pi')$, where the average length $\overline{L}(\pi)$ of a plan π is defined as follows[3]:
  - $\overline{L}(\lambda) = 0$;
  - if $\alpha \in ACT_P$: $\overline{L}(\alpha; \pi') = 1 + l(\pi')$;
  - $\overline{L}(test(\varphi);\ \texttt{if}\ \varphi\ \texttt{then}\ \pi'\ \texttt{else}\ \pi'') = 1 + \frac{1}{2}\overline{L}(\pi') + \frac{1}{2}.\overline{L}(\pi'')$.

As to $\succeq^i$ it can be generalized into $\succeq^{ig}$, where $\pi \succeq^{ig} \pi'$ if and only if one of the following five conditions is satisfied:

1. $\pi = \lambda$;
2. $\pi = \alpha; \pi_1$, $\pi' = \alpha; \pi_2$, and $\pi_1 \succeq^{ig} \pi_2$;
3. $\pi' = \alpha; \pi''$ and $\pi \succeq^{ig} \pi''$;
4. $\pi = (test(\varphi);\ \texttt{if}\ \varphi\ \texttt{then}\ \pi_1\ \texttt{else}\ \pi_2)$
   and $\pi' \in \{\pi_1, \pi_2\}$;
5. $\pi = (test(\varphi);\ \texttt{if}\ \varphi\ \texttt{then}\ \pi_1\ \texttt{else}\ \pi_2)$,
   $\pi' = (test(\varphi);\ \texttt{if}\ \varphi\ \texttt{then}\ \pi'_1\ \texttt{else}\ \pi'_2)$,
   where $\pi_1 \succeq^{ig} \pi'_1$ and $\pi_2 \succeq^{ig} \pi'_2$.

Notice that $\succeq^{lm}$ and $\succeq^{lb}$ are complete, but not $\succeq^{ig}$. TWe have the following implications

$$\pi \succ^{ig} \pi' \Rightarrow \pi \succ^{al} \pi'$$
$$\pi \succ^{ig} \pi' \Rightarrow \pi \succ^{lb} \pi'$$

(but $\succeq^{al}$ and $\succeq^{lb}$ are generally incomparable). Furthermore, the restriction of $\succeq^{al}$ and $\succeq^{lb}$ to unconditional plans is equal to $\succeq^c$, while the restriction of $\succeq^{ig}$ is $\succeq^i$.

We now have to show how we draw plausible inferences on the actor's initial beliefs, goals and/or next actions. The principle is the same as in the complete knowledge case. We give here only the definition based on U-relevance (expressed in a form similar as in Proposition 3): given π_{done}, b_0 a belief state and $g \in \Gamma$, (b_0, g) is U-plausible for π_{done} w.r.t. $\succeq$ if and only if there exists a $\succeq$-undominated solution plan for $C = \langle b_0, g, \pi_{done} \rangle$ starting by π_{done}.

[3] This definition implicitly assumes that the test outcomes are equiprobable, which may look somewhat arbitrary.

Example 3 *Let x and y be two propositional variables, S the set of interpretations on $\{x, y\}$, $ACT = \{test(y), test(x \wedge y), switch(x), switch(y)\}$ where $switch(x)$ (resp. $switch(x)$) is an action switching the truth value of x (resp. y). Goals are expressed by logical formulas: saying that the goal is φ means that the goals states are the interpretations satisfying φ (for instance, we write $g = x$ instead of $g = Mod(x) = \{(x, y), (x, \neg y)\}$). We let $\succeq = \succeq^{ig}$.*

1. *$\pi_{done} = switch(y)$, $g = x$, actor's beliefs unknown.*
   The plausible initial belief states are $b_0 = \neg y$ and $b_0 = (x \leftrightarrow \neg y)$. In both cases, the R- and U-plausible next action is $test(x \wedge y)$.
2. *$\pi_{done} = test(y); [obs(y)]; switch(x)$; goal and initial belief state unknown.*
   There are five plausible initial belief states, each of which associated wit a set of plausible goals as seen on the following table. Plausible next actions do not figure on the table.

b_0	possible values of g
x	$\neg x \wedge y$
$\neg x$	$x \wedge y$
$x \vee \neg y$	$\neg x; \neg x \wedge \neg y; \neg x \wedge y$
$\neg x \vee \neg y$	$x; x \wedge \neg y; x \wedge y; x \leftrightarrow y$
$x \leftrightarrow y$	$\neg x; \neg x \wedge y; \neg x \wedge \neg y$
$x \leftrightarrow \neg y$	$x; x \wedge y; x \wedge \neg y$

3. *$\pi_{done} = test(b)$, $g = a$, initial belief unknown. The plausible initial belief states are $a \vee b$, $\neg a \vee b$, $a \vee \neg b$, $\neg a \vee \neg b$, $a \leftrightarrow b$, $a \leftrightarrow \neg b$ and $\top$ (tautology).*
4. *$\pi_{done} = test(a \wedge b)$, $g = a$, initial belief unknown. The plausible initial belief states are b, $a \vee b$, $\neg a \vee b$, $a \vee \neg b$, $a \leftrightarrow b$ and $\top$ (tautology).*

Related work

Action prediction

The most related approach is (Brafman & Tennenholtz 1997), from which we borrowed the intuitive principle for drawing plausible inference. Our article could be considered as the study of a specific case of (Brafman & Tennenholtz 1997); being more specific allows us for discussing details that were irrelevant to (Brafman & Tennenholtz 1997) because of it high level of generality. Note that, although (Brafman & Tennenholtz 1997) does not explicitly discuss the ascription of goals, they can do it easily in their framework – see Section 2.3 of (Brafman & Tennenholtz 1997). Thus, in our simpler framework in which we do not consider utility functions over states, nor consider decisions strategies from qualitative decision theory, we were able to focus on the precise way of defining plausible goals, belief and next actions, and study their meaning especially with respect to our assumptions on the actor's level of rationality.

Other approaches to action prediction can be found, such as (Davison & Hirsh 1998; Gorniak & Poole 2000), but unlike (Brafman & Tennenholtz 1997) and our approach, they do not consist in ascribing a mental state to the agent. In (Isozaki & Katsuno 1996), other agents' beliefs are updated when sequences of events are observed by the modelling agent, given some initial assumptions about the ability of others to observe events: agents are passive, just observe events and do not perform actions so as to reach a goal; one does not reason about the agent's goals or intended actions.

Plan and goal recognition

Another related line of work is *plan recognition* (e.g. (Kautz & Allen 1986) and many subsequent works). However, a plan recognition problem contains an history of plans already evoked by the agent in various circumstances, or a precomputed plan library; the problem then consists in determining the plan followed by the agent from an initial action sequence. Moreover, plan recognition rarely allows for reasoning about the mental state of the agent (except (Pollack 1990)). See (Brafman & Tennenholtz 1997) for a deep discussion on this. *Goal recognition* (Hong 2001), unlike plan recognition, does not require the data of a history or a library of plans but works with an implicit description of the available plans; but it still does not make use of mental models.

Other works

Our preference relations (especially $\succeq^i$) are related to evaluation criteria for plan quality in (Lin 1998); but the similarity between both works ends here.

A long-term further research direction would consist in relaxing the assumption of non-interaction between the agents and extend the present framework to signalling and cooperation, as suggested in (Levesque 2001).

References

Brafman, R., and Tennenholtz, M. 1997. Modeling agents as qualitative decision makers. *Artificial Intelligence* 94.

Davison, B. D., and Hirsh, H. 1998. Predicting sequences of user actions. In *Proceedings of the 1998 AAAI/ICML Workshop "Predicting the Future: AI Approaches to Time Series Analysis"*.

Doyle, J., and Wellman, M. P. 1991. Preferential semantics for goals. In *Proceedings of AAAI-91*, 698–703.

Gorniak, P., and Poole, D. 2000. Predicting future user actions by observing unmodified assumptions. In *Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*.

Hong, J. 2001. Goal recognition through goal graph analysis. *Journal of Artificial Intelligence Research* 15:1–30.

Isozaki, H., and Katsuno, H. 1996. A semantic characterization of an algorithm for estimating others' beliefs from observation. In *Proceedings of AAAI-96*, 543–549.

Kautz, H., and Allen, J. 1986. Generalized plan recognition. In *Proceedings of AAAI-86*, 32–37.

Levesque, H. 2001. Talk given at MFI-01 (Toulouse, May 2001).

Lin, F. 1998. On measuring plan quality (a preliminary report). In *Proceedings of KR'98*, 224–232.

Pollack, M. E. 1990. Plans as complex mental attitudes. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication*. MIT Press.

Acknowledgements

I wish to thank the participants of the French project "La pertinence dans le choix et l'interprétation d'actions" within the CNRS Programme "Cognitique". As well, thanks to the KR anonymous reviewers for many helpful suggestions, and to the participants of the NRAC-03 workshop. This paper has been partly written while I was visiting the University of Potsdam.

Appendix: proofs

Proof of Proposition 1:

1. let $\alpha \in Rel_U^{\succeq}(C)$. By definition, this means that there is a $\pi \in Max_{\succeq}(SolPlans(C)) \cap Develop(\pi_{done})$ such that $\pi_{done}; \alpha$ starts π. We have $\pi \in SolPlans(C) \cap Develop(\pi_{done})$. If π were dominated in $SolPlans(C) \cap Develop(\pi_{done})$, there would exist a π' such that $\pi' \succ \pi$ and $\pi' \in SolPlans(C) \cap Develop(\pi_{done}) \subseteq SolPlans(C)$, which would contradict the fact that π is undominated in $SolPlans(C)$. Therefore, we have $\pi \in Max_{\succeq}(SolPlans(C) \cap Develop(\pi_{done}))$, therefore $\alpha \in Rel_R^{\succeq}(C)$.

2. We have $Rel_R^{\succeq}(C) \Leftrightarrow Max_{\succeq}(SolPlans(C) \cap Develop(\pi_{done})) = \emptyset$. (Notice that in the case there is a $\pi \in Max_{\succeq}(SolPlans(C) \cap Develop(\pi_{done}))$ such that $\pi = \pi_{done}$, then since $\lambda_{ACT} \in ACT$ then $\pi; \lambda_{ACT}$, which leads to the same final state as π, belongs to $Max_{\succeq}(SolPlans(C) \cap Develop(\pi_{done}))$ so that we have $\lambda \in Rel_R^{\succeq}(C)$, and the latter is not empty.) Now,

$$\begin{array}{ll} & Rel_R^{\succeq}(C) = \emptyset \\ \Leftrightarrow & Max_{\succeq}(SolPlans(C) \cap Develop(\pi_{done})) = \emptyset \\ \Leftrightarrow & SolPlans(C) \cap Develop(\pi_{done}) = \emptyset \\ \Leftrightarrow & \text{no solution plan starts by } \pi_{done} \\ \Leftrightarrow & \text{there is no plan leading from } next(s_0, \pi_{done}) \text{ to } g. \end{array}$$

3. let $\succeq_1$ and $\succeq_2$ such that for any π, π', $\pi \succ_2 \pi'$ implies $\pi \succ_1 \pi'$. Then for any plan π and any set X of plans, we have

$$\begin{array}{ll} & \pi \text{ undominated for } \succeq_1 \text{ in } X \\ \Rightarrow & \text{there is no } \pi' \in X \text{ such that } \pi' \succ_1 \pi \\ \Rightarrow & \text{there is no } \pi' \in X \text{ such that } \pi' \succ_2 \pi \\ & \pi \text{ undominated for } \succeq_2 \text{ in } X. \end{array}$$

This entails $Rel_R^{\succeq_1}(C) \subseteq Rel_R^{\succeq_2}(C)$ and $Rel_U^{\succeq_1}(C) \subseteq Rel_U^{\succeq_2}(C)$.

■

Proof of Proposition 2: Assume that $\succeq$ is complete. The inclusion $Goals_U(OC, \succeq) \subseteq Goals_R(OC, \succeq)$ is a straightforward consequence of point 1 of Proposition1, so we have to prove the inclusion $Goals_R(OC, \succeq) \subseteq Goals_U(OC, \succeq)$, that is, any R-plausible goal is U-plausible for OC. Assume that g is *not* U-plausible. Let us distinguish two cases:

Case 1 there is no solution plan at all starting by π_{done}, that is, $SolPlans(OC + g) \cap Develop(\pi_{done}) = \emptyset$. In that case, $last(\pi_{done})$ cannot be R-relevant for $Restr(\pi_{done}, now - 1)$, therefore g is not R-plausible.

Case 2 there is a solution plan starting by π_{done}. Since g is not U-plausible for OC, there is an integer $t < now$ such that $\pi_{done}(t)$ is not U-relevant for $Restr(\pi_{done}, t)$. Let t^* the smallest integer satisfying the latter condition, that is, $t^* = \min\{t \mid \pi_{done}(t) \text{not U-relevant for} Restr(\pi_{done}, t)\}$. (We have $0 \leq t^* \leq now - 1$). Saying that $\pi_{done}(t^*)$ is not U-relevant for $Restr(\pi_{done}, t^*)$ is equivalent to
$SolPlans(\langle s_0, g, Restr(\pi_{done}, t^*)\rangle)$
$\cap Develop(Restr(\pi_{done}, t^*)) = \emptyset$
which means that

(1) any solution plan starting by $Restr(\pi_{done}, t^*); \pi_{done}(t^*)$ is dominated in $SolPlans(\langle s_0, g, Restr(\pi_{done}, t^*)\rangle)$.

Notice that $Restr(\pi_{done}, t^*); \pi_{done}(t^*) = Restr(\pi_{done}, t^* + 1)$. Now, let $\pi^* \in Max(\succeq, SolPlans(\langle s_0, g, Restr(\pi_{done}, t^*)\rangle))$. The existence of such a π^* is guaranteed by the assumption that there is a solution plan starting by π_{done}, therefore *a fortiori* starting by $Restr(\pi_{done}, t^*)$. Now, (1) together with the fact that $\geq$ is complete imply

(2) for each solution plan π starting by $Restr(\pi_{done}, t^*); \pi_{done}(t^*)$ we have $\pi^* \succ \pi$.

Therefore, $\pi_{done}(t^* + 1)$ is not R-relevant for $\langle s_0, g, Restr(\pi_{done}, t^*)\rangle$. This implies that g is not R-plausible for OC.

■

Proof of Proposition 3:

$\Rightarrow$ Assume that $g \in \Gamma$ is U-plausible given $OC = \langle s_0, \Gamma, \pi_{done}\rangle$; then the last action $last(\pi_{done}) = \pi_{done}(now - 1)$ is U-relevant for $\langle s_0, g, Restr(\pi_{done}, now - 1)\rangle$, which means that there exists a plan $\pi \in Max(\succeq, SolPlans(\langle s_0, g, Restr(\pi_{done}, now - 1)\rangle)$ starting by $Restr(\pi_{done}, now - 1); \pi_{done}(now - 1)$. Now, noticing that $Restr(\pi_{done}, now - 1); \pi_{done}(now - 1) = \pi_{done}$, we have found a plan π such that $\pi \in Max(\succeq, SolPlans(\langle s_0, g, Restr(\pi_{done}, now - 1)\rangle)$, that is, $\pi \in Max(\succeq, SolPlans(OC + g) \cap Develop(\pi_{done})$; therefore there is no $\pi' \in SolPlans(OC + g)$ such that $\pi' \succ \pi$.

$\Leftarrow$ Let $g \in \Gamma$ and assume there is a plan $\pi \in SolPlans(OC + g) \cap Develop(\pi_{done})$ such that there is no plan $\pi' \in SolPlans(OC + g)$ verifying $\pi' \succ \pi$. Then for each $t < now$, π is an undominated solution plan starting by $Restr(\pi_{done}, t)$, hence, $\pi_{done}(t)$ is U-relevant in $\langle s_0, g, Restr(\pi_{done}, t)$. Therefore, since moreover $g \in \Gamma$, we have that g is U-plausible given OC.

■

Proof of Proposition 4: Assume $\succeq$ is complete (which entails that R-plausible and U-plausible goals coincide, due to Proposition 2). Let g_1 and g_2 be two (R/U)-plausible goals. By definition there exist a plan π'_1 such that $\pi_1 = \pi_{done};\pi'_1$ is nondominated w.r.t. $\succeq$ in $SolPlans(g_1)$ and a plan π'_2 such that $\pi_2 = \pi_{done};\pi'_2$ is nondominated w.r.t. $\succeq$ in $SolPlans(g_2)$. Assume that $g_1 \cup g_2$ is not a plausible goal (H): then there exists no plan $\pi_{done};\pi'$ nondominated w.r.t. $\succeq$ in $SolPlans(g_1 \cup g_2)$. Let $\pi_3 \in Pref(\succeq, SolPlans(g_1 \cup g_2))$; since $\pi_1 \in SolPlans(g_1) \subseteq SolPlans(g_1 \cup g_2)$, (H) together with the fact that $\succeq$ is complete imply that $\pi_3 \succ \pi_1$, and similarly $\pi_3 \succ \pi_2$. Now, since the environment is completely known, $SolPlans(g_1 \cup g_2) = SolPlans(g_1) \cup SolPlans(g_2)$. Therefore we have $\pi_3 \in SolPlans(g_1)$ or $\pi_3 \in SolPlans(g_2)$. If $\pi_3 \in SolPlans(g_1)$ then $\pi_3 \succ \pi_1$ contradicts the hypothesis that π_1 is nondominated in $SolPlans(g_1)$ and similarly, if $\pi_3 \in SolPlans(g_2)$ then $\pi_3 \succ \pi_2$ contradicts the hypothesis that π_2 is nondominated in $SolPlans(g_2)$. Therefore we have a contradiction. ■