

Qualitative and Probabilistic Uncertainty

joint work with: Luca Iocchi, Thomas Lukasiewicz, Riccardo Rosati

- **qualitative models for handling uncertainty**: e.g., situation calculus and recent formalisms inspired by the action language \mathcal{A} .
- **probabilistic models for handling uncertainty**: e.g., probabilistic extensions of the situation calculus and \mathcal{A} as well as Poole's ICL
- only little work on **orthogonal combinations of qualitative and probabilistic models**: nondeterminism and probabilistic uncertainty in a game-theoretic framework (Halpern and Tuttle, 1993), and probabilistic extension of $\mathcal{C}+$ (Eiter and Lukasiewicz, 2003)

Robotic Soccer Example

The action “align to ball” may succeed (resp., fail) with probability .7 (resp., .3), while the goalkeeper’s action “open legs” may save the goal or not.

In the latter case, the effect is nondeterministic rather than probabilistic, since it is not possible to assign probabilities to the possible effects, which in fact depend on external factors (e.g., speed and kind of kick performed by an opponent robot) and cannot be evaluated a priori.

Action Descriptions: Syntax

- **static** and **dynamic fluents F** and fluent formulas
- **effect** and **sensing actions A**
- **precondition axiom: executable α if ϕ** , where ϕ is a fluent formula, and α is an action. Informally, α is executable in every state that satisfies ϕ .
- **conditional effect axiom: caused ψ after α when ϕ** , where ϕ is a fluent formula, ψ is a fluent conjunction, and α is an action. Informally, if the current state satisfies ϕ , then executing α produces the effect ψ .

- **sensing effect axiom: caused to know ω or $\neg\omega$ after α** , where ω is a fluent conjunction, and α is a sensing action. Informally, after executing α , the agent knows that ω is either true or false.
- **default frame axiom: inertial ϕ after α** , where ϕ is a fluent conjunction, and α is an action. Informally, if ϕ holds in the current state, then ϕ holds also after the execution of α , if it is consistent with the effects of α .
- **domain constraint axiom: caused ψ if ϕ** , where ϕ and ψ are fluent formulas. It represents background knowledge invariant to actions.
- **action description KB** : finite set of precondition, conditional effect, sensing effect, default frame, and domain constraint axioms.

Extended Action Descriptions: Syntax

- We divide the set of effect actions into *deterministic*, *nondeterministic*, and *probabilistic actions*.
- A *nondeterministic dynamic context formula* has the form

$$\mathbf{caused} \psi_1, \dots, \psi_n \mathbf{after} \alpha \mathbf{when} \phi, \quad (1)$$

where ϕ is a fluent formula, ψ_1, \dots, ψ_n are fluent conjunctions, α is an action, and $n \geq 2$.

- A *probabilistic dynamic context formula* has the form

$$\mathbf{caused} \psi_1 : p_1, \dots, \psi_n : p_n \mathbf{after} \alpha \mathbf{when} \phi, \quad (2)$$

where additionally $p_1, \dots, p_n > 0$ and $p_1 + \dots + p_n = 1$.

- An *extended action description* $D = (KB, C)$ consists of an action description KB and a finite set C containing one *nondeterministic* (resp., *probabilistic*) dynamic context formula for each action in KB .

Extended Action Descriptions: Semantics

Extended action description (KB, C) encodes a system of deterministic, nondeterministic, and probabilistic transitions between e-states:

- *Each nondeterministic (resp., probabilistic) action α in KB , with (1) (resp., (2)) in C , is associated with a set of **contexts** $V_\alpha = \{v_1, \dots, v_n\}$, where each v_i has the probability $Pr_\alpha(v_i) = p_i$, if α is probabilistic.*
- *If α is executable in S , then the successor e-state of S after executing α in $v \in V_\alpha$, denoted $\Phi_v(S, \alpha)$, is the e-state $\Phi(S, \alpha)$ under $KB_\alpha(v)$.*
- *If α is nondeterministic, then the set of successor e-states of S under α is defined as $F_\alpha(S) = \{\Phi_v(S, \alpha) \mid v \in V_\alpha\}$.*
- *If α is probabilistic, then the **probability distribution on the successor e-states** of S under α , denoted $Pr_\alpha(\cdot | S)$, is defined by $Pr_\alpha(S' | S) = \sum_{v \in V_\alpha, S' = \Phi_v(S, \alpha)} Pr_\alpha(v)$.*

Extended Action Descriptions: Example

executable bodykick *if* ballclose

...

caused goalsaved *after* openlegs *when* alignedtoball

caused to know ballclose *or* \neg ballclose *after* senseballclose

...

inertial l *after* α (for every fluent literal l and action α)

caused ballinpenalty *if* ballclose

caused goalsaved, \neg goalsaved *after* openlegs

caused ballclose:0.8, \neg ballinpenalty:0.1, \neg ballclose:0.1 *after* gotoball

...

Conditional Plans

- Conditional planning problem: Given an extended action description, an initial state description ϕ_I , and a goal description ψ_G , compute the best conditional plan to achieve ψ_G from ϕ_I .
- A *conditional plan* Π is either (i) the empty conditional plan λ , or (ii) of form $\alpha; \Pi'$, or (iii) of form $\beta; \text{if } \omega \text{ then } \{\Pi_\omega\} \text{ else } \{\Pi_{\neg\omega}\}$, where α is an effect action, β is a sensing action of outcomes ω and $\neg\omega$, and Π' , Π_ω , and $\Pi_{\neg\omega}$ are conditional plans.






Probabilistic conditional plans

When there are probabilistic actions we need to characterize the **best** plan based on a notion of **goodness**.

The *goodness* of Π for achieving ψ under ϕ is defined through a *bottom-up propagation rule*, which:

- *starts from the leaves with 1 for success (0 for failure)*
- *takes the minimum in case of non probabilistic actions*
- *sums up the joint probabilities of the success nodes*

Example: Goodness of a Conditional Plan



Summary and Outlook

- *formalism for reasoning about actions with sensing under qualitative and probabilistic uncertainty*
- *formal semantics in terms of deterministic, nondeterministic, and probabilistic transitions between epistemic states*
- *notion of a conditional plan and its goodness under qualitative and probabilistic uncertainty*
- *sound and complete algorithm for conditional planning*

Topics for future research:

- *sensing actions with noisy outcomes; extension of conditional plans to cyclic execution structures*
- *exploit the notion of goodness to implement heuristics for efficient search in the space of epistemic states; detailed complexity analysis for conditional planning in our framework*