Queen's University Belfast

PACES

Risk-aware planning for BDI systems

Lookahead planning for BDI agents

Belief-desire-intention (BDI) model for **reactive** and **scalable** multi-agent systems using predefined plan library

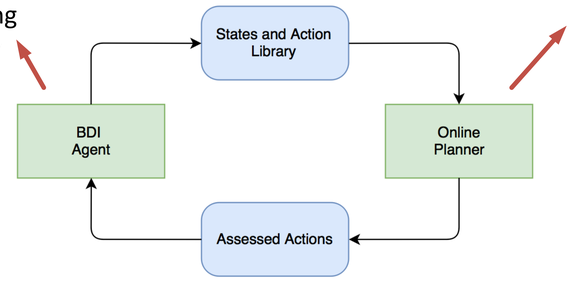

Handle **unforeseen events** and **stochastic effects** of actions with on-demand planning

Handle **information fusion, uncertain beliefs**, etc

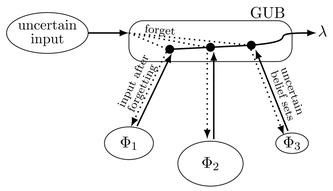

Retain reactive nature of BDI by finding "good enough" actions using **approximate solutions**

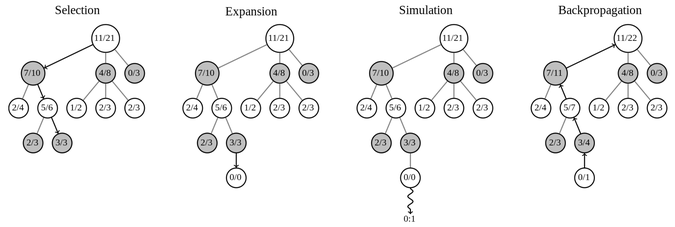

Risk-aware planning

Risk is the possibility of obtaining a utility/reward from an action which is lower than the expected utility for this action, due to the potential for undesirable outcomes

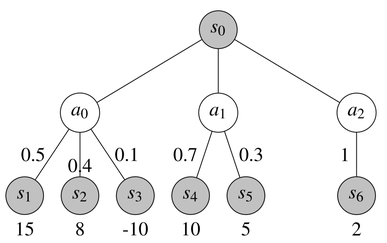

Variance can be used as an indication of risk

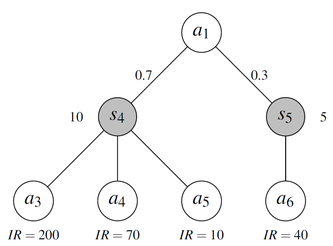

Must also consider potential of **future risk**

$$CMR(a_t) = \sum_{t=1}^{h-1} \left(\gamma^{t-1} \cdot \sum_{s' \in O} T(s, a_t, s') \cdot CRE(s', C_{t+1}) \right)$$

$$CRE(s, C) = \min_{a \in C} (CMR(a))$$

If the variance of a1 is 5.3, then the **cumlative minimum risk** (CMR) of a1 is 5.3 + (0.7 x 10) + (0.3 x 40) = 24.3

Anytime estimate of variance is needed to provide BDI agent with on-demand planning

$$\mu_i = \mu_{i-1} + \frac{x_i - \mu_{i-1}}{i}$$ ← Current mean

$$M_i = M_{i-1} + (x_i - \mu_{i-1})(x_i - \mu_i)$$ ← Sum of squares of differences from current mean

$$r = \frac{M_i}{i-1}$$ ← Variance estimate

Decision making

Balance **reward-seeking** and **risk-averse** behaviour

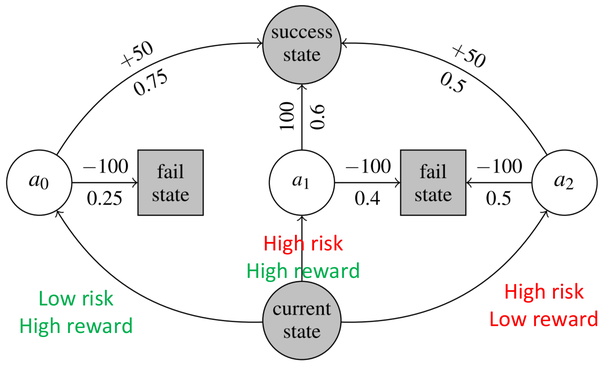

Given a **risk-aversion degree R**, we can define the **optimal R risk-averse actions**

$$\alpha^* = \arg\max_{\alpha \in \mathcal{A}} \left(u(\alpha) - R\left(\sqrt{r(\alpha)}\right) \right)$$

Principle 1: a rational agent will only consider an action with a lower utility when this also involves a lower risk

Principle 2: an action with a lower utility will only be adopted when it reduces the risk sufficiently according to the agent's level of risk aversion

Principle 3: the level of risk aversion increases as the number of resources decreases

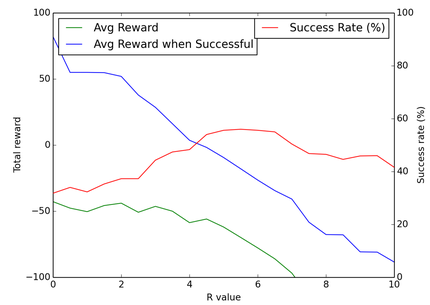

Illustrative scenarios

Navigation

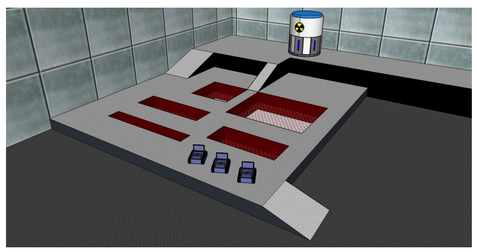

Goal: To reach the target within a time limit by navigating a sequence of bridges, where quicker routes increase the probability of failing to reach the target

Risk:
- Not reaching target on time
- Falling from bridge

Path clearance

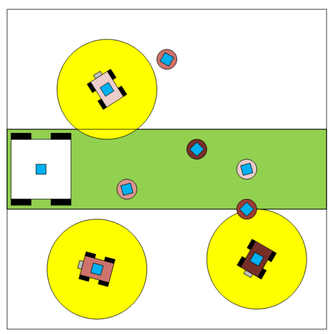

Goal: To clear a path through obstacles as part of a group of robots, with uncertainty over location of obstacles and other robots

Risk:
- Collisions

Applications: dealing with hazardous material