An Approach to Quantify Plans Robustness in Real-world Applications

The 34th International Joint Conference on Artificial Intelligence, IJCAI 2025

Francesco Percassi[1], **Sandra Castellanos-Paez**[1],
Romain Rombourg[2], Mauro Vallati[1]

[1]School of Computing and Engineering, University of Huddersfield, United Kingdom
[2]G2ELab, Grenoble INP, CNRS, Université Grenoble Alpes, Grenoble, France

University of Huddersfield

Grenoble INP UGA

Motivation and Context

Plans often fail due to uncertainty → costly replanning

Automated planning is used in real-world domains like traffic control, robotics, UAV navigation, etc.

- **Challenge**: Uncertainty and noise (sensor errors, actuator failures, environmental unpredictability) can make plans ineffective.
- **Costly fallback**: Replanning and plan repair are computationally expensive.

Key shift

In many real scenarios, *exact goal achievement* is unnecessary — reaching an **acceptable outcome region** is sufficient.

Core concepts

Executable vs. Valid

- **Executable vs. Valid Plans**:
  - **Executable plan**: Can be successfully executed from the initial state, satisfying all the necessary conditions at each step.
  - **Valid plan**: Executable **and** reaches a state satisfying the goal.

- **Execution-Invariant Plans and Tasks**:
  - **Execution-invariant plan**: Always executable even when certain initial numeric variables vary.
  - **Execution-invariant task**: All executable plans are invariant to some subset of numeric variables.

- Probability that a given plan π achieves the goal G under the distribution over the possible initial states.
- Measured via Bayesian estimation (Beta distribution confidence intervals).

Plan robustness $R_{\mathcal{I}}(\pi)$

Let Π be a planning task, let $\mathcal{I}$ be a random variable representing the possible initial states and $f_{\mathcal{I}}$ its distribution. The robustness of a plan π for Π with respect to $\mathcal{I}$ is defined as:

$$R_{\mathcal{I}}(\pi) = \mathbb{E}_{\mathcal{I} \sim f_{\mathcal{I}}} \left[[\![\pi \in \textsc{Plans}(\Pi[\mathcal{I}])]\!] \right]$$

where $[\![P]\!]$ is the Iverson bracket which returns 1 if proposition P is true and 0 otherwise.

A statistical approach for quantifying plan robustness

B robustness and B_{min} – Acceptable region

- Extends robustness to allow tolerance B around the goal.
- B_{min}: Minimum tolerance needed to reach a target robustness level $R^\star$.

B-Robustness

Given a tolerance factor B, the B-robustness of a plan π is defined as follows:

$$R_{\mathcal{I}}(\pi, B) = \mathbb{E}_{\mathcal{I}\sim f_{\mathcal{I}}}[[\![d_G(\gamma(\mathcal{I},\pi)) \leq B]\!]]$$

B_{min}

Let $R^\star \in [0,1]$ denote a desired robustness level. $B_{\min}$ is the minimum tolerance required for the plan to succeed with probability at least $R^\star$:

$$B_{\min}(\pi, R^\star) = \inf_{\mathbb{R}_+}\{B \mid R_{\mathcal{I}}(\pi, B) \geq R^\star\}$$

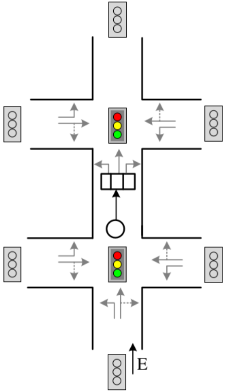

- Network as a directed graph:
  - Nodes = junctions
  - Edges = road links
- Each junction operates through predefined traffic signal configurations, which regulate vehicle flows between incoming and outgoing links.
- Traffic flows: continuous processes.
- Signal transitions: discrete events.

- **Goal**: Optimize traffic signals to maximise throughput and minimise congestion.
- **Variables**: Initial traffic occupancy of each link (Link occupancies), Average flow of vehicles in junctions between ingoing and outgoing links, during given green times (turn rates).
- **Data**: 90 real instances from Yorkshire corridor.
- **Findings**:
  - Robustness varies between days.
  - Conservative robustness requires ∼20% tolerance; most probable case ∼5%.

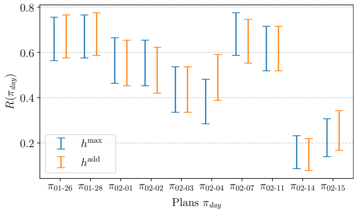

Figure: $[\underline{R}, \overline{R}]$: Robustness CI per day plan.

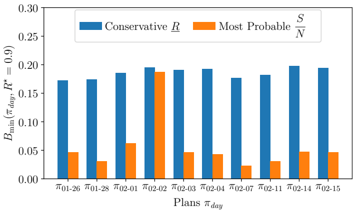

Figure: $B_{\min}$ per day plan.

A Baxter robot provided with two arms, tasked to manipulate an object into a desired final configuration. The manipulation involves a sequence of actions that allows the Baxter to grasp two links and modify the angle of the joint connecting these links together.

- **Initial state**: The initial object pose characterised by the orientation of each link l, which is described by two angles: θ_l^{xy} for the horizontal plane and θ_l^z for the z-axis.
- **Goal**: Numeric conditions imposed on the orientation variables of some links.
- **Variables**: Initial object pose angles.
- **Results**: Plans are generally not robust; tolerance levels vary less than in UTC.

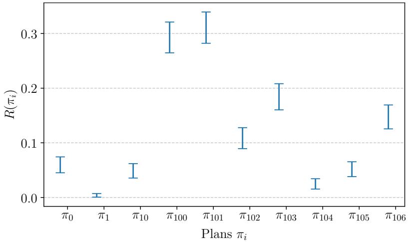

Figure: $[\underline{R}, \overline{R}]$: Robustness CI per plan.

Results and Insights

Key Findings (execution-invariance, tolerance–robustness trade-offs)

- Robustness varies significantly between plans, even in same domain.
- Execution-invariance naturally occurs in some domains (e.g., UTC, Baxter).
- Tolerance–robustness trade-offs are **domain-specific**.
- Statistical framework works with both historical and synthetic data.

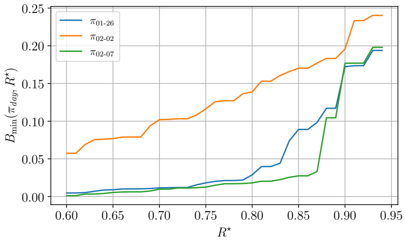

Figure: Tolerance vs desired robustness trade-off for UTC.

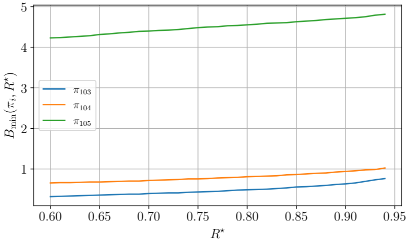

Figure: Tolerance vs desired robustness trade-off for Baxter.

Looking Forward

Future Work (bounds, repairing plans, guiding plan generation)

- **Proposed**:
  - Clear definition of execution-invariant problems.
  - Statistical framework to quantify plan robustness.
- Applicable to PDDL+ and other numeric planning formalisms.
- **Future directions**:
  - Automatic derivation of numeric variable bounds.
  - Plan repair for insufficient robustness.
  - Using robustness to guide plan generation.

Looking Forward

Final Takeaways (robustness > perfection, practical usability)

- In noisy environments:
  - Focus on executability + acceptable outcomes.
  - Quantify robustness to plan with confidence.
  - Framework supports informed plan selection without over-reliance on replanning.

Thanks for your attention! Questions?

Scan to get in touch :)