RISK-AWARE AGENT COMPOSITIONS

Guruprerana Shabadi, Rajeev Alur

University of Pennsylvania, US

AGENT GRAPHS

Agentic systems decompose complex objectives into a sequence of **subtasks** and choose a set of specialized **AI agents** to complete them.

Task	Agents π_i
Drone navigation	RL control policies
Information Retrieval	LLMs for retrieving and parsing data

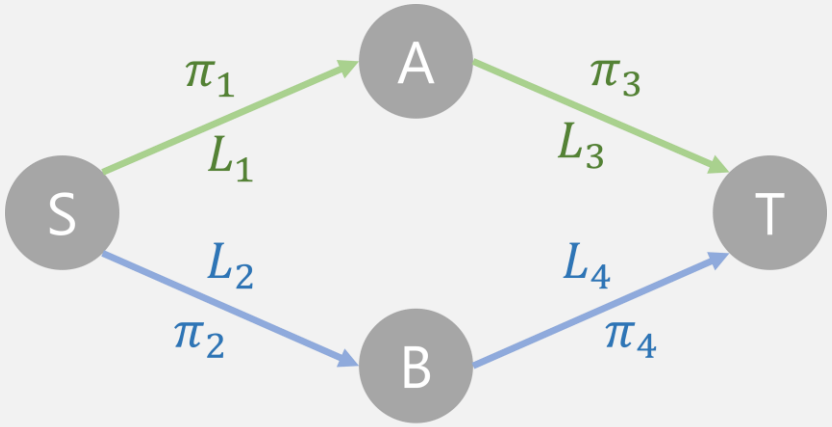

RISK MINIMIZATION

Risk is quantified by real-valued **loss functions** that measure **safety**, **fairness**, and **privacy** violations.

Agents π_i	**Losses L_i**
RL policies	Negative minimum distance to obstacles
LLM agents	Degree of hallucination

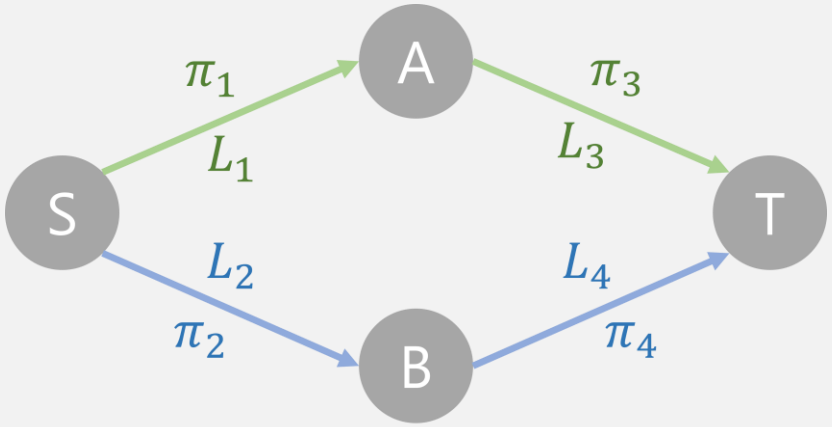

RISK MINIMIZATION

Objective: Determine composition of agents, i.e., a path in the graph, that minimizes **value-at-risk**. Let P be the set of paths.

$$\ell^* = \min_{p \in P} \mathrm{VaR}_\alpha \left[\max_{i \in p} L_i \right]$$

where VaR_α is the $(1-\alpha)$-quantile function.

In essence, with probability at least $1-\alpha$, each agent in the optimal composition achieves at most ℓ^* loss.

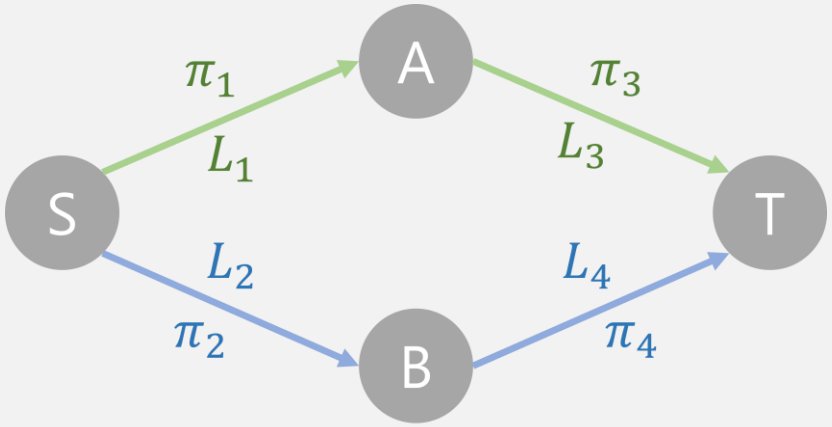

BUCKETED-VAR ALGORITHM

Polynomial time algorithm through **union bound!**

Let $\sum \alpha_i = \alpha$. If

$$\mathbf{Pr}[L_i < V_i] \geq 1 - \alpha_i$$

then,

$$\mathbf{Pr}[\max L_i < \max V_i] \geq 1 - \alpha$$

BucketedVaR works by:

- Sampling agent rollouts for VaR estimation
- Dynamic programming
- Optimizing allocation of α_i to each agent over a discrete set of buckets $\{0, \frac{\alpha}{d}, \frac{2\alpha}{d}, \dots, \alpha\}$

BUCKETED-VAR ALGORITHM

Polynomial time algorithm through **union bound!**

Let $\sum \alpha_i = \alpha$. If

$$\mathbf{Pr}[L_i < V_i] \geq 1 - \alpha_i$$

then,

$$\mathbf{Pr}[\max L_i < \max V_i] \geq 1 - \alpha$$

If the distributions of losses of the agents are independent, then

Theorem: As the sample size and discretization factor d grows towards ∞, the estimated quantile is at most the $(1 - \alpha + {}^{\alpha^2}/_{2})$-quantile of the optimal composition.

EVALS

Evaluate performance on a set of compositional RL environments where losses represent safety and resource consumption.

- Produces tight estimates of the value-at-risk
- Approximation improves with growing sample size and discretization
- Scales efficiently to big graphs