

Introduction

Traditional path planning under uncertainty, such as the **Canadian Traveller Problem (CTP)** and **Informative Path Planning (IPP)**, often overlooks **heterogeneous visibility**.

In real-world environments, vantage points offer varying levels of information. We propose a framework that:

- **Exploits High-Visibility Nodes:** Leverages non-local sensing to observe distant edges from specific vantage points.
- **Balances Exploration/Exploitation:** Uses an *observation award* combined with traversal costs to balance both.

Our approach yields more robust paths in high-uncertainty environments by prioritizing informative vantage points.

Problem Setting

The environment is modeled as a graph $G = (\mathcal{V}, \mathcal{E})$ with source v_s and target v_t .

Key Components:


- **Traversal Cost:** Each edge e has a cost $c(e)$.
- **Visibility Mapping:** $\text{vis}(v) \subseteq \mathcal{E}$ defines edges observable from node v .
- **Edge Blockages:** A subset $\mathcal{E}_b \subset \mathcal{E}$ is statically blocked. No prior knowledge of probability distribution of blockages is assumed.

Objective: Navigate from v_s to v_t with minimum cumulative cost $\sum_{e \in P} c(e)$.

Since minimum path cost is a random variable, the policy should balance information gain with path cost minimization.

Short Diverse Path Sampling

We generate a **path tree** $\mathcal{P}$ of depth r by recursively branching n times, where each branch represents a shortest path computed after removing all edges within an m -hop radius of a node sampled from the parent path.



Path Generation

We define a path reward $R(P)$ that balances **cost-efficiency** with the **information gain** of visiting high-visibility nodes.

1. Edge Utility (U): An edge's utility is based on its frequency in the diverse path set $\mathcal{P}_1$:

$$U(e) = \sum_{P_{ij} \in \mathcal{P}: e \in P_{ij}} \frac{1}{n \cdot r} \quad (1)$$

2. Observation Reward (R_{obs}): The sum of utilities for all currently unobserved edges e visible from node v :


$$R_{obs}(v) = \sum_{e \in \text{vis}(v)} U(e) (1 - 1_{obs}(e)) \quad (2)$$

3. Overall Path Reward:

$$R(P) = \lambda \sum_{v \in P} R_{obs}(v) - \sum_{e \in P} c(e) \quad (3)$$

The optimal path P^* is selected from the path tree $\mathcal{P}_2$ such that $P^* = \arg \min_{P \in \mathcal{P}_2} R(P)$.

Experiment: Plateau



We test on a grid with *plateaus* (high visibility, high entry cost). Chokepoint edges have an independent blockage probability p (not used during planning). We compare our method against a Shortest Path (SP) baseline and RPP.

Method	Mean	Std Dev	Runtime(ms)
SP	43.75	16.05	2.5
RPP	34.00	0.00	11083
Ours ($\lambda = 3$)	36.00	6.00	94.2

Result on Plateau Environment ($p = 0.5$)


On a different plateau environment we evaluate the performance of our algorithm with different reward ratios λ .

	SP	$\lambda = 0$	$\lambda = 3$	$\lambda = 10$
$p = 0$	22.00	22.00	22.00	22.00
$p = 0.1$	30.10	30.10	29.98	30.51
$p = 0.2$	35.35	35.35	34.77	34.85
$p = 0.3$	39.20	39.20	37.62	37.63
$p = 0.4$	42.46	42.46	39.71	39.63
$p = 0.5$	45.11	45.11	40.98	40.93

Mean Path Cost (Plateau)

Experiment: Natural Terrain

We test on a natural terrain environment with mountains (high visibility, high traversal cost). Three obstacles are randomly placed with probability p .



The shortest path baseline (blue) and our proposed method (red) are compared; dashed lines show the initial intended path and the solid lines show the actual path used.

	SP	$\lambda = 1$	$\lambda = 3$
$p = 0$	147.62	150.49	150.80
$p = 0.2$	153.37	151.31	151.61
$p = 0.4$	157.41	152.07	152.40
$p = 0.6$	159.40	152.76	153.19
$p = 0.8$	159.01	153.38	153.97

Mean Path Cost (Natural Terrain)

Discussion

Baseline Equivalence: At $\lambda = 0$, the proposed agent's behavior is equivalent to that of the Shortest Path (SP) agent.

Observation Weighting: Increasing λ assigns higher priority to edge observations, yielding superior performance in environments with a high probability of edge blockage.