CoGeL: Constraint-to-Geometry Learning from Demonstrations and Failures for Safe Long-Horizon Robot Navigation

Jumman Hossain[1], Nirmalya Roy[1]

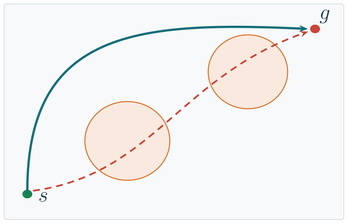
CoRL 2026

Learn risk from local failures; build safer routes over much longer horizons.

01 Motivation & Core Idea

CoGeL learns risky actions from demonstrations, failed attempts, and safety interventions, then plans safer long-range routes without a predefined hazard map.

1 Robot Evidence



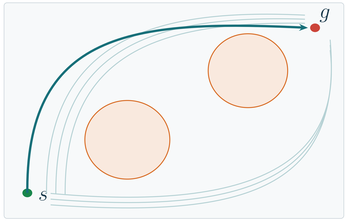
Safe demonstrations
Failures / interventions

2 Soft Safety Signal



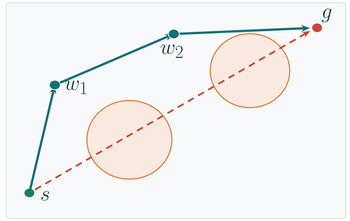
$\hat{\chi}_\psi(s,a) \in [0,1]$
Learned violation likelihood

3 Directed Geometry



$d_\theta(s,g) \neq d_\theta(g,s)$
in general

4 Waypoint Planning



Edge cost: $d_\theta(w_i, w_j)$
Local actor follows waypoints

A Practical Example

Situation: A robot encounters loose debris along a shorter route.

Experience: Previous failures suggest that crossing the debris is risky.

Decision: CoGeL favors a longer, safer detour to reach the goal.

02 Learning and Planning

Learn a local safety signal, embed it into a directed distance, then compose safe goal-reaching behaviors.

1 Infer risk from robot evidence

Fit a soft violation predictor from successful rollouts, failed attempts, and interventions.

$$\hat{\chi}_\psi(s,a) \;=\; P_\psi(y=1 \mid s,a)$$

2 Add risk to the travel cost

$$\tilde{c}_\psi(s,a,s') = c(s,a,s') + \lambda \hat{\chi}_\psi(s,a)$$

Unsafe actions become expensive *before* a longer route is planned.

3 Learn a compositional distance

$$y_{\mathrm{B}} = \tilde{c}_\psi(s,a,s') + \gamma d_{\bar{\theta}}(s',g)$$

Bellman regression uses the observed next state s' from replay. A route-consistency loss encourages valid waypoint composition:

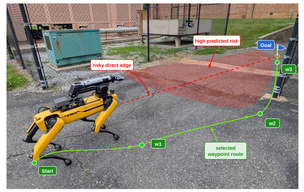
$$\mathcal{L}_\triangle = \mathbb{E}[[d_\theta(s,g) - d_\theta(s,w) - d_\theta(w,g) - m]_+]$$

The resulting $d_\theta(s,g)$ is directed: $d_\theta(s,g)$ and $d_\theta(g,s)$ may differ.

What is new?
Risk shapes the planning distance, rather than filtering routes afterward.

Why directed?
84% vs 54%
CoGeL vs symmetric-distance ablation.

At deployment: plan, then execute



1 Score directed edges
Edge cost: $d_\theta(w_i, w_j)$

2 Select waypoints
Graph search using $d_\theta(w_i, w_j)$

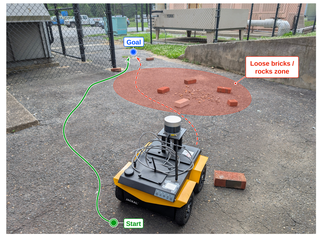
3 Execute segments
Local actor: $\pi_\phi(\cdot \mid s_t, w_{i+1})$

Directed route cost $C(W) = d_\theta(w_0,w_1) + \cdots + d_\theta(w_{N-1},w_N)$

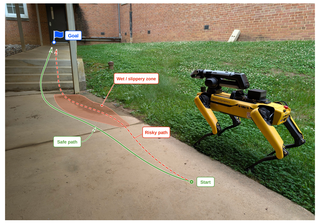
03 Results and Deployment

Learned costs favor detours around local hazards.

Jackal: loose debris



Spot: wet / slippery path



Green: CoGeL safe route. Red dashed: unsafe shortcut.

Physical trials: CoGeL vs Safety+Plan

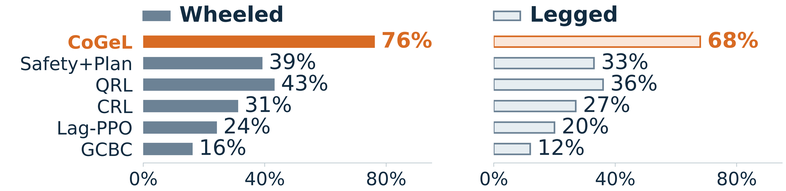
Jackal
90% vs 80%
Safety events/trial: **0.15** vs 0.40

Spot
85% vs 75%
Safety events/trial: **0.20** vs 0.45

Success above; 20 trials per method per robot, under matched conditions.

Longer missions: success at $8\times$ horizon ↑

Test horizon relative to training; simulation means over five seeds.



Fewer constraint violations in simulation

74% wheeled / 69% legged

Reduction per episode vs QRL, from the main simulation evaluation.

Takeaway & Limitation

CoGeL learns from local failures to plan safer routes over long horizons. However, its effectiveness depends on representative, well-calibrated safety evidence.

Acknowledgments NSF #2233879, #2509680, #1750936, #2502886; U.S. Army #W911NF2120076, #W911NF2410367; ONR #N00014-23-1-2119.

[1]Department of Information Systems, University of Maryland, Baltimore County (UMBC)
{jumman.hossain, nroy}@umbc.edu

Project Page
Scan for details and paper updates