

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 GRL-SNAM: GEOMETRIC REINFORCEMENT LEARNING WITH DIFFERENTIAL HAMILTONIANS FOR NAVIGATION AND MAPPING IN UNKNOWN ENVIRONMENTS

Anonymous authors

Paper under double-blind review

## ABSTRACT

We present GRL-SNAM, a geometric reinforcement learning framework for Simultaneous Navigation and Mapping in unknown environments. GRL-SNAM differs from traditional SLAM and other reinforcement learning methods by relying exclusively on local sensory observations without constructing a global map. Our approach formulates navigation and mapping as coupled dynamics on generalized Hamiltonian manifolds: sensory inputs are translated into local energy landscapes that encode reachability, obstacle barriers, and deformation constraints, while policies for sensing, planning, and reconfiguration evolve stagewise under Differential Policy Optimization (DPO). A reduced Hamiltonian serves as an adaptive score function, updating kinetic/potential terms, embedding barrier constraints, and continuously refining trajectories as new local information arrives. We evaluate GRL-SNAM on 2D deformable navigation tasks, where a hyperelastic robot learns to squeeze through narrow gaps, detour around obstacles, and generalize to unseen environments. We evaluate GRL-SNAM on procedurally generated 2D deformable-robot tasks (hyperelastic ring) with narrow gaps and clutter, comparing against *local reactive* baselines (PF, CBF, staged DWA) and *global* A\* references (rigid, clearance-aware) under identical stagewise sensing constraints. GRL-SNAM matches near-CBF path quality while using the minimal map coverage of PF, preserves clearance, generalizes to unseen layouts, and demonstrates that Hamiltonian-structured RL enables high-quality navigation through *minimal exploration* via local energy refinement rather than global mapping.

## 1 INTRODUCTION

Reinforcement learning has achieved remarkable successes in high-dimensional control, yet its application to real-world continuous navigation remains fundamentally limited. Long-horizon reasoning, multi-scale decision making, and online adaptation pose challenges that overwhelm existing methods. Model-free RL consumes millions of interactions, while hierarchical variants introduce brittle complexity. In simultaneous navigation and mapping (SNAM), where agents must traverse and construct evolving environmental representations, these limitations become prohibitive.

At its core, the difficulty arises because conventional RL policies are *structureless*. They treat navigation as black-box optimization, ignoring the geometric and physical principles that make locomotion stable, adaptive, and safe. Without inductive bias, policies overfit training environments, fail under distribution shift, and collapse during long rollouts.

### 1.1 BEYOND BELLMAN OPTIMIZATION: PURELY FEEDFORWARD CONTROL:

Our framework does not optimize a value function via the Bellman equation. Standard RL algorithms hinge on recursive bootstrapping for estimating returns, propagating value updates, and iteratively improving policies. This induces high sample complexity, instability, and delayed credit assignment, especially in navigation with long horizons.

In contrast, our approach is **purely feedforward**: policies emerge as direct gradient flows of Hamiltonian energies, without value iteration. Navigation decisions are computed in a single pass from

054 local sensory input and the reference Hamiltonian, bypassing dynamic programming altogether.  
 055 This eliminates the need for rollout-based value propagation, yielding stable training, low variance  
 056 adaptation, and sample-efficient online updates.  
 057

## 058 1.2 KEY INSIGHT: HAMILTONIAN STRUCTURE AS NAVIGATION INDUCTIVE BIAS 059

060 We propose addressing these limitations by grounding RL in **Hamiltonian mechanics**. Our central  
 061 insight is that navigation can be framed as learning energy functionals:  
 062

$$063 \quad \mathcal{H}(q, p) = K(p) + P(q) \quad (1)$$

064 where kinetic and potential energies encode control objectives, constraints, and adaptation strategies.  
 065 This formulation introduces three structural advantages:  
 066

067 (1) **Energy conservation** stabilizes long-horizon rollouts by preventing accumulation of numerical  
 068 errors. (2) **Symplectic geometry** naturally separates fast reactive dynamics from slow strategic  
 069 planning, addressing multi-scale temporal coordination. (3) **Barrier encoding** integrates safety and  
 070 collision avoidance directly into potential functions, eliminating fragile reward shaping.

071 **Hamiltonian structure transforms policy optimization into Differential Policy Optimization (dfPO)**  
 072 Nguyen & Bajaj (2025), where policies emerge as gradient flows of learned energies that respect  
 073 geometry, conserve invariants, and generalize across environments. We not only compare our al-  
 074 gorithm against standard deep reinforcement learning baselines but also focus our comparisons on  
 075 task-specific navigation and mapping baselines such as PF, CBF, and A\*.

## 077 1.3 OFFLINE-ONLINE HAMILTONIAN SYNERGY 078

079 We distinguish between complementary learning regimes that exploit this geometric structure:  
 080

081 **Offline learning** discovers reference Hamiltonians  $h^{\theta^*}$  trained on trajectory data, capturing funda-  
 082 mental multi-scale navigation dynamics in local frames. These provide stable geometric priors  
 083 encoding essential coupling between sensing, planning, and deformation.

084 **Online adaptation** fuses new environmental context into learned Hamiltonians through contextual  
 085 corrections:  $h^{\text{adapted}} = h^{\text{ref}} + \Delta h^{\text{context}}$ . This creates conservative adaptation: systems default to  
 086 learned physics-based behaviors while adding minimal corrections for environmental variations.

087 The synergy transforms every offline policy into a *reference Hamiltonian* and every online update  
 088 into a *geometric alignment step*. Navigation emerges from meta-policies that parse environments,  
 089 assemble energy landscapes, and integrate them through symplectic dynamics.

## 091 1.4 CONTRIBUTIONS: 092

093 This work establishes **GRL-SNAM** as a new approach beyond structureless policy learning. Our  
 094 contributions are:  
 095

- 096 1. **Hamiltonian RL framework:** Allows adaptable integration of classical mechanics into  
 097 RL for navigation, treating rewards as energies and policies as symplectic flows.
- 098 2. **Multi-scale geometric coordination:** Differential policies for sensing, planning, and adap-  
 099 tation unified through shared energy formulations, achieving temporal scale separation  
 100 without manual hierarchy design.
- 102 3. **Physics-grounded adaptation:** Principled offline-online decomposition where stable ref-  
 103 erence dynamics adapt through geometric alignment rather than catastrophic relearning.
- 104 4. **Theoretical guarantees:** Symplectic structure preservation ensures stability, while inde-  
 105 pendent policy learning achieves linear sample complexity scaling.
- 106 5. **Empirical validation:** Hyperelastic ring navigation demonstrates superior sample effi-  
 107 ciency and generalization compared to A\* and CBF baselines.

## 2 RELATED WORK

We focus on structure-preserving, deployable navigation with deformable bodies. Our work intersects advances in geometric learning, safety-critical control, and deformable robot navigation.

**Mathematical Foundations for RL Navigation.** Most navigation RL methods operate in Euclidean spaces using standard PPO Schulman et al. (2017) or TD3 Fujimoto et al. (2018) formulations without geometric constraints. Geometric approaches include SE(3) equivariant policies Hoang et al. (2025) for manipulation and Riemannian safe navigation Klein et al. (2023) using tangent space projections. Hamiltonian neural networks Desai et al. (2021) demonstrate superior learning dynamics through symplectic structure but remain limited to simple control tasks.

**Safety-Critical Navigation.** Control Barrier Function (CBF) integration with RL achieves formal safety guarantees Li et al. (2023), but treats constraint satisfaction as orthogonal to navigation optimality, often resulting in conservative behaviors. Our Hamiltonian formulation integrates safety constraints directly within the energy structure.

**Deformable Robot Navigation.** Recent work demonstrates ring-like navigation through pre-programmed strategies: aerial gap navigation via fixed Liquid Crystal Elastomer responses Qi et al. (2024) and HAVEN Mulvey & Nanayakkara (2024) using predetermined shape-changing sequences. These approaches rely on offline parameter optimization followed by deterministic execution—they cannot adapt deformation strategies online as environmental conditions change.

**Neural Scene Representations.** NeRF-based SLAM methods like NICE-SLAM Zhu et al. (2022) provide rich environmental representations that complement our energy-based navigation formulation by supplying obstacle and free-space information for barrier and goal potential computation.

**Simultaneous Navigation and Mapping:** Most SNAM approaches prioritize building detailed maps before navigation. SGoLAM Kim et al. (2021) couples goal localization with occupancy mapping, CMP Gupta et al. (2019) integrates a differentiable planner into learned mapping, and CL-SLAM Vödisch et al. (2023) maintains maps for long-term adaptability. In contrast, our GRL-SNAM framework aims to *reach goals via high-quality, well-weighted paths while mapping as little of the environment as possible*. To our knowledge, no prior work explicitly targets minimal exploration; our method introduces progressive path refinement, continually improving least-cost trajectories as new observations arrive.

**Positioning.** GRL-SNAM addresses key gaps by extending Hamiltonian mechanics from simple control to complex navigation requiring sensing, planning, and deformation. Unlike existing methods that require manual task decomposition or rely on pre-programmed strategies, our differential multi-policy architecture learns specialized policies naturally coupled through shared Hamiltonian energy formulations. The symplectic structure ensures stable coordination across temporal scales with formal convergence guarantees, bridging theoretically principled geometric methods with practical navigation frameworks.

### 3 METHODOLOGY

We present GRL-SNAM (Geometric Reinforcement Learning for Spatial Navigation and Manipulation): a Hamiltonian-structured navigator that unifies offline physics learning with online adaptive correction through black-box modular policies. The code for this paper is available at: [Code](#)

We formulate navigation in unknown environments as iterative energy minimization with active response of dynamical constraints. Consider a deformable robot with state  $q_t = (c_t, \theta_t, y_t, \psi_t) \in Q$  navigating from  $x_0$  to  $x_g$  through unknown obstacles characterized by binary occupancy  $I : \mathbb{R}^2 \rightarrow \{0, 1\}$ .  $c_t$  and  $\theta_t$  are robot frame coordinates and axis angle w.r.t. world coordinate, while  $y_t$  and  $\psi_t$  are robot sensory and object configuration state respectively. We denote each fixed local scenario as an environment configuration  $\mathcal{E}$ . The goal of GRL-SNAM is to i) identify response from three offline solvers ii) build a digital-twin style surrogate and give feedback of each black-box across stages (at the time scale when  $\mathcal{E}$  updates). To achieve both goals, we first address the Hamiltonian energy setup attached to each fixed scenario  $\mathcal{E}$ , then one explain what could be the connection between sensory, movement and reconfiguration solvers after integration of expected Hamiltonians. Lastly,

162 we explain what feedback one can provide via interaction with offline solvers that online correction  
 163 can inject partially or fully by surgery of energy functionals.  
 164

165 **3.1 ONLINE STAGEWISE ADAPTATION**  
 166

167 Before turning to the experimental evaluation, we briefly summarize how the Hamiltonian structure  
 168 (§3.2), the modular policy architecture (§3.3), and the meta-learning framework (§C.4) are combined  
 169 at test time into a unified navigation system.  
 170

171 **Overview.** At each navigation stage, the navigator sequentially queries the three offline-trained  
 172 policies  $\pi_y$  (sensor),  $\pi_f$  (frame/FPE), and  $\pi_o$  (object/reconfig) to obtain state-dependent control  
 173 proposals. The meta-policy  $g_\xi$  maps the current environment and policy responses to environment-  
 174 dependent energy weights and friction, from which we assemble a surrogate Hamiltonian and inte-  
 175 grate the corresponding port-Hamiltonian dynamics with dissipation and port correction. Observable  
 176 quantities such as clearance, goal progress, and speed are then used to perform a short Jacobian-  
 177 based update of the active energy weights and friction coefficients, yielding a stagewise adaptation  
 178 loop. The full pseudo-code, including initialization, query protocols, energy assembly, integration  
 179 details, and adaptation rules, is deferred to Algorithm 3 in Appendix E.  
 180

181 **3.2 NAVIGATION AS HAMILTONIAN OPTIMIZATION**  
 182

183 **From optimal control to Hamiltonian via Legendre–Fenchel conjugacy (fixed  $\mathcal{E}$ ).** Fix an  
 184 environment  $\mathcal{E}$  and consider the control-affine dynamics  $\dot{q} = f(q) + A(q)u$  with stage cost  
 185  $L(q, u; \mathcal{E}) = -\mathcal{R}(q; \mathcal{E}) + \varphi(u)$ , where  $\mathcal{R}$  encodes goal/deflection/barrier terms and  $\varphi$  penalizes  
 186 effort. Pontryagin’s principle introduces a costate  $p$  and the *control Hamiltonian*  $\mathcal{H}(q, p, u; \mathcal{E}) :=$   
 187  $p^\top(f(q) + A(q)u) - L(q, u; \mathcal{E})$ . Eliminating  $u$  amounts to taking the Legendre–Fenchel conjugate  
 188 of  $\varphi$ :

$$\begin{aligned} \mathcal{H}(q, p; \mathcal{E}) &= \sup_u \{p^\top A(q)u - \varphi(u)\} + p^\top f(q) + \mathcal{R}(q; \mathcal{E}) \\ &= \varphi^*(A(q)^\top p) + p^\top f(q) + \mathcal{R}(q; \mathcal{E}), \end{aligned} \quad (2)$$

191 provided  $\varphi$  is proper, closed, and strictly convex. The optimal feedback is  $u^*(q, p) =$   
 192  $\nabla \varphi^*(A(q)^\top p)$ . In the common quadratic case  $\varphi(u) = \frac{1}{2}u^\top \Phi u$  (with  $\Phi \succ 0$  as a kinetic “term”),  
 193 we have  $\varphi^*(u) = \frac{1}{2}u^\top \Phi^{-1}u$ , hence

$$\mathcal{H}(q, p; \mathcal{E}) = \frac{1}{2}p^\top(A(q)\Phi^{-1}A(q)^\top)p + p^\top f(q) + \mathcal{R}(q; \mathcal{E}). \quad (3)$$

194 Identifying the *inverse mass* as  $M(q)^{-1} := A(q)\Phi^{-1}A(q)^\top$  and (optionally) absorbing  $p^\top f(q)$  into  
 195 a gauge term (or set  $f(q) \equiv 0$ ) yields the mechanical form  $\mathcal{H}(q, p; \mathcal{E}) = \frac{1}{2}p^\top M(q)^{-1}p + \mathcal{R}(q; \mathcal{E})$ .  
 196 The canonical equations,  $\dot{q} = \nabla_p \mathcal{H}$  and  $\dot{p} = -\nabla_q \mathcal{H}$ , are therefore the *Hamiltonian outcome* of  
 197 the fixed- $\mathcal{E}$  optimal control problem. Soft constraints (barriers) simply contribute additively to  $\mathcal{R}$ ;  
 198 nonconservative effects can be modeled as port inputs without altering the conjugate construction  
 199 (esp. friction). Thus, for each scenario  $\mathcal{E}$ , the inner motion law is Hamiltonian with kinetic energy  
 200 induced by the control penalty via conjugacy and potential shaped by the environment. Note that  
 201 motion planning offline policy may or may not follow the surrogate Hamiltonian one wish to align  
 202 but we calibrate the surrogate by interaction with response (and in this paper we particularly focus  
 203 on discrete  $(q, p)$  dynamics via aggregations).  
 204

205 **Search space for the Hamiltonian.** The goal of navigator is to learn to search in the energy space  
 206 of  $\mathcal{H}$ . The Hamiltonian defined in 3 is a function on the cotangent bundle  $T^*Q$ :

$$H \in \mathcal{H} := \{H(q, p; \mathcal{E}) = \frac{1}{2}p^\top M(q)^{-1}p + \mathcal{R}(q; \mathcal{E}) \mid M : \mathcal{Q} \rightarrow \mathbb{S}_{++}^2, \mathcal{R} \in \mathcal{R}\}.$$

207 We regard  $\mathcal{R}$  as a Hilbert space of admissible potentials on  $\mathcal{Q}$  (and w.r.t. environmental configu-  
 208 ration) (e.g.  $L^2(\mathcal{Q} \times \mathcal{E})$ ). For planar navigation we *restrict* the search to the environment-indexed  
 209 linear cone generated by task energies, where each policy governs a distinct energy term and a joint  
 210 dynamical barrier term. We model the meta navigator by a parametrized map  $\mathcal{E} \rightarrow \eta_\xi(\mathcal{E})$  producing  
 211 nonnegative dual weights that shape the primal potential. In general, each energy term may itself be  
 212 parametrized:  
 213

$$H(q, p; \omega, \xi, \mathcal{E}) = \frac{1}{2}p^\top M(q; \omega_M)^{-1}p + \mathcal{R}(q; \omega, \eta_\xi(\mathcal{E})),$$

$$\begin{aligned}
216 \quad \mathcal{R}(q; \omega, \eta_\xi(\mathcal{E})) &= E_{\text{sensor}}(q; \mathcal{E}, \omega_y) + \beta(\mathcal{E}) E_{\text{goal}}(q; \mathcal{E}, \omega_g) + \lambda(\mathcal{E}) E_{\text{obj}}(q; \omega_d) \\
217 \quad &+ \sum_{i \in \mathcal{C}_t(\mathcal{E}, q)} \alpha_i(\mathcal{E}, t) b(d_i(q; \mathcal{E}); \omega_b),
\end{aligned} \tag{4}$$

221 with  $\eta_\xi(\mathcal{E}) = (\beta(\mathcal{E}), \lambda_{\text{obj}}(\mathcal{E}), \{\alpha_i(\mathcal{E}, t)\}) \in \mathbb{R}_+^{m(\mathcal{E}, t)}$ . Here  $\omega = (\omega_y, \omega_M, \omega_g, \omega_d, \omega_b)$  are *intra-term* parameters (e.g. metric, goal shape, deformation model, barrier template), while  $\eta_\xi$  learns the *inter-term* tradeoffs by mapping the environment  $\mathcal{E}$  to dual weights. The cardinality  $m(\mathcal{E}, t) = 2 + |\mathcal{C}_t(\mathcal{E}, q)|$  is environment/active-set dependent, so  $\eta_\xi$  is implemented with a permutation-invariant set encoder that outputs per-constraint scores  $\alpha_i(\mathcal{E}, t) \geq 0$ , together with scalars  $\beta(\mathcal{E}), \lambda(\mathcal{E}) \geq 0$ . The active set  $\mathcal{C}_t(\mathcal{E}, q) := \{i \mid d_i(q, \mathcal{E}) \leq \hat{d}\}$  is discovered online by sensing.

### 3.3 NAVIGATOR'S SUBMODULAR ARCHITECTURE

230 Rather than learning monolithic navigation policies, we decompose the problem into three independent 231 score functions, each dedicated to a specific navigation aspect (we refer readers Figure 1 for 232 details):

233 **Definition 3.1** (Independent Score Functions). *Let  $\mathcal{K} = \{y, f, o\}$  denote the set of policy indices 234 corresponding to sensor, frame, and object domains respectively. For each  $k \in \mathcal{K}$ , define:*

- 236 •  $z_k \in \mathcal{Z}_k$ : the phase space state for policy  $k$ , where  $\mathcal{Z}_k = \mathcal{Q}_k \times \mathcal{P}_k$  with configuration 237 space  $\mathcal{Q}_k$  and momentum space  $\mathcal{P}_k$
- 238 •  $\theta_k \in \Theta_k$ : the learnable parameters for policy  $k$ , where parameter sets satisfy disjointness: 239  $\Theta_i \cap \Theta_j = \emptyset$  for  $i \neq j$
- 240 •  $h_k^{\theta_k} : \mathcal{Z}_k \times \mathcal{E} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ : a learned energy functional parameterized by  $\theta_k$

243 Each policy  $\pi_k$  is defined as an independent score function:  $s_k^{\theta_k}(z_k, \mathcal{E}, t) = S_k^{\theta_k}(\nabla_{z_k} h_k^{\theta_k}(z_k, \mathcal{E}, t))$

244 The parameter disjointness ensures independence:  $\frac{\partial s_k^{\theta_k}}{\partial \theta_j} = 0$  for all  $j \neq k$

245 allowing parallel training while maintaining coordination through shared constraints  $\mathcal{C}_t$ .

248 **Policy Abstraction.** Each policy is treated as a black box that:

- 250 • **Sensor Policy** ( $\pi_y$ ): Adapts perception parameters  $\rightarrow$  energy gradients for information 251 gathering
- 252 • **Frame Policy** ( $\pi_f$ ): Plans collision-free paths  $\rightarrow$  energy gradients for goal attraction
- 253 • **Shape Policy** ( $\pi_o$ ): Controls robot deformation  $\rightarrow$  energy gradients for obstacle navigation

255 The key insight is that our Navigator is agnostic to policy implementation—our contribution is the 256 Hamiltonian structure binding them together through dynamic constraint sets  $\mathcal{C}_t$ .

258 Algorithm 3 details the online adaptation procedure, where the navigator issues sequential queries 259 to the sensor, frame, and reconfig policies, integrates their energy gradients into a Hamiltonian 260 update, and applies meta-corrections for contextual alignment to generate stable trajectories in novel 261 environments.

263 **Hamiltonian of modular sub-systems.** Let  $\mathcal{K} = \{y, f, o\}$  index three Hamiltonian submodules 264 with local states  $z_k = (q_k, p_k) \in T^*Q_k$  and local Hamiltonians

$$265 \quad H_k(q_k, p_k; \xi, \mathcal{E}) = \frac{1}{2} p_k^\top M_k(q_k)^{-1} p_k + \underbrace{\mathcal{R}_k(q_k; \mathcal{E})}_{\text{module potential}}.$$

268 The *navigator* represents the whole stack by a *surrogate* Hamiltonian on  $T^*Q$ , where  $\eta(\mathcal{E}) =$  269  $g_\xi(\mathcal{E}) \in \mathbb{R}_+^{m(\mathcal{E}, t)}$  is the meta policy(explorer) output for potential functional. For navigation, in

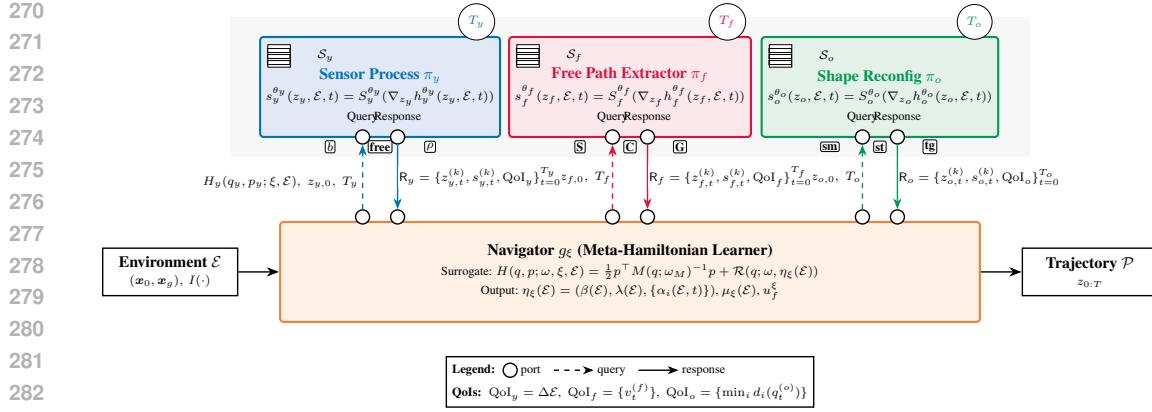


Figure 1: Independent score function architecture and query–response interface. The Navigator  $g_\xi$  issues queries containing local Hamiltonians  $H_k$ , initial states  $z_{k,0}$ , and time horizons  $T_k$  to each policy  $\pi_k$  ( $k \in \{y, f, o\}$ ). Each policy computes score functions  $s_k^{\theta_k}$  via energy gradients from learned Hamiltonians  $h_k^{\theta_k}$ , backed by spatial indices  $\mathcal{S}_k$  for efficient neighbor queries. Policies return standardized responses  $R_k$  containing state trajectories, score sequences, and QoIs. The Navigator aggregates these to update the surrogate Hamiltonian and generate meta-corrections  $\eta_\xi(\mathcal{E})$ ,  $\mu_\xi(\mathcal{E})$ , and  $u_f^\xi$  for adaptive navigation.

particular the reference potential for each submodule is:

$$\mathcal{R}_y(q_y; \xi, \mathcal{E}) = E_{\text{sensor}}(q_y; \mathcal{E}, \omega_y) + \sum_{i \in \mathcal{C}_t(\mathcal{E}, q)} \alpha_i(\mathcal{E}, t) b(d_i(q; \mathcal{E}); \omega_b) \quad (5)$$

$$\mathcal{R}_f(q_f; \xi, \mathcal{E}) = \beta(\mathcal{E}) E_{\text{goal}}(q_f; \mathcal{E}, \omega_g) + \sum_{i \in \mathcal{C}_t(\mathcal{E}, q)} \alpha_i(\mathcal{E}, t) b(d_i(q; \mathcal{E}); \omega_b) \quad (6)$$

$$\mathcal{R}_o(q_o; \xi, \mathcal{E}) = \lambda(\mathcal{E}) E_{\text{obj}}(q_o; \omega_d) + \sum_{i \in \mathcal{C}_t(\mathcal{E}, q)} \alpha_i(\mathcal{E}, t) b(d_i(q; \mathcal{E}); \omega_b) \quad (7)$$

**Remark.** A black-box setup assume one cannot observe the potential components  $\{\mathcal{R}_k\}$  yet can only observe kinetic terms by the integrated dynamics via policy  $\pi_k$ . A gray-box setup can allow navigator reshape potential (and possibly kinematics) so one do not need to calibrate surrogate energy equation 3 with local Hamiltonian aggregation  $\hat{H} := \sum_{k \in \mathcal{K}} H_k^{\theta_k}$ .

**Hamiltonian Dynamics of modular sub-systems.** Each submodule  $k \in \mathcal{K}$  integrates its local dynamics for a short horizon and returns a standardized response  $R_k$ . Let the *effective module Hamiltonian* be  $h_k^{\theta_k}(q_k, p_k, t)$  with initial condition  $h_k^{\theta_k}(q_k, p_k, 0) = H_k(q_k, p_k; \mathcal{E})$ . The local (port-)Hamiltonian flow with dissipation and remaining nonconservative input is

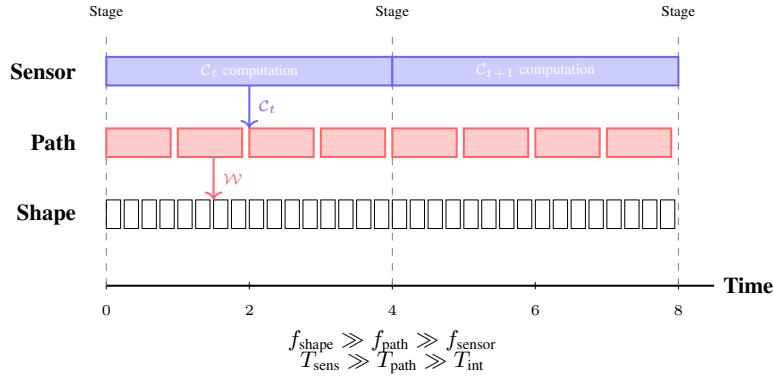
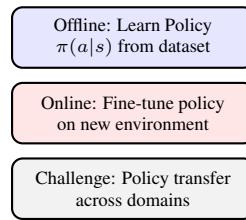
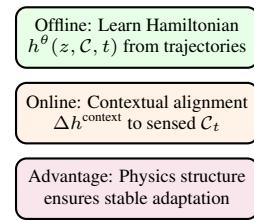
$$\dot{q}_k = \nabla_{p_k} h_k^{\theta_k}(q_k, p_k, t), \quad (8)$$

$$\dot{p}_k = -\nabla_{q_k} h_k^{\theta_k}(q_k, p_k, t) - \Gamma_k^\xi(q_k; \mathcal{E}) \nabla_{p_k} h_k^{\theta_k}(q_k, p_k, t) + G_k^\xi(q_k; \mathcal{E}) u_k^\xi(q_k, p_k, t, \mathcal{E}), \quad (9)$$

with  $\Gamma_k \succeq 0$  a Rayleigh/viscous damping and  $G_k^\xi u_k^\xi$  the *nonconservative* (non-potential) external input. We define the *score of dynamics* recorded by module  $k$  as the deterministic drift,

$$s_k(z_k, t) := \begin{bmatrix} \nabla_{p_k} h_k^{\theta_k} \\ -\nabla_{q_k} h_k^{\theta_k} \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ -\Gamma_k^\xi \nabla_{p_k} h_k^{\theta_k} + G_k^\xi u_k^\xi \end{bmatrix}}_{\text{non-Hamiltonian contributions (friction/ports)}}, \quad z_k = (q_k, p_k).$$

**Navigator meta-learning details.** For clarity of exposition, we defer the full formulation of the navigator as a meta-Hamiltonian learner including the construction of  $\mathcal{R}(q; \eta_\xi(\mathcal{E}))$ , the training objective in equation 48, and the QoI-based online adaptation scheme to Appendix C.4.

324 3.4 MULTI-SCALE TEMPORAL COORDINATION  
325326 The policies operate at natural temporal hierarchies, creating stable multi-scale coordination:  
327341 Figure 2: Temporal hierarchy. Sensor policy operates at low frequency (once per stage), estab-  
342 lishing environmental constraints  $\mathcal{C}_t$ . Path policy operates at medium frequency within each stage,  
343 computing waypoints  $\mathcal{W}$ . Shape policy operates at high frequency, continuously adapting at each in-  
344 tegration step. This creates a natural hierarchy where slow sensor updates provide stable constraints  
345 for faster path and shape adaptations.346  
347 This temporal separation enables a **nested quasi-static approximation**: the fastest dynamics (re-  
348 configuration) equilibrate within each frame update, and frame dynamics settle before the slower  
349 sensor policy evolves. This hierarchy prevents destabilizing interactions across timescales while  
350 preserving the necessary coupling for coherent, coordinated behavior.351 3.5 OFFLINE PHYSICS LEARNING VS ONLINE ADAPTIVE CORRECTION  
352353 Our approach resolves the fundamental tension between learning complex dynamics and real-time  
354 adaptation through principled decomposition:355 356 **Standard RL**357 **Our GRL-SNAM**366 Figure 3: Comparison between standard RL offline/online adaptation and our physics-grounded  
367 approach. Standard methods learn arbitrary policies and struggle with transfer, while our approach  
368 learns physically meaningful Hamiltonians that naturally adapt to environmental variations.371 3.6 THEORETICAL PROPERTIES  
372

373 Our framework provides three key theoretical guarantees:

374 **Theorem 3.2** (Multi-Policy Stability). *Under temporal scale separation  $T_{\text{sens}} \gg T_{\text{path}} \gg T_{\text{int}}$  and  
375 bounded parameter updates, the coupled system maintains stability with error bound  $\mathcal{E}_{\text{total}} \leq \epsilon$ .*  
376377 **Theorem 3.3** (Symplectic Preservation). *Each score function generates symplectic dynamics pre-  
378 serving the canonical structure  $\omega_k(z_{k,t+1}) = \omega_k(z_{k,t})$ .*

378 **Theorem 3.4** (Linear Sample Complexity). *Independent training achieves total sample complexity*  
 379  $N_{\text{total}} = \sum_{k \in \{y, f, o\}} O(\epsilon_k^{-(2d_k+4)})$ , *linear in the sum of policy dimensions rather than exponential*  
 380 *in joint dimensionality.*

382 We defer the proof of theorems in appendix. The system thinks in physics during offline training but  
 383 adapts through energy corrections during online execution, combining principled dynamics stability  
 384 with real-world deployment flexibility.

## 387 4 EXPERIMENTAL EVALUATION

389 We evaluate GRL-SNAM across multiple dimensions that highlight the unique capabilities of our  
 390 geometric approach compared to standard reinforcement learning and classical navigation methods.  
 391 Our evaluation encompasses task performance, safety guarantees, and learning efficiency under min-  
 392 imal sensing constraints. For more detailed results and analysis, refer to Appendix I

393 **Experimental Setup:** We evaluate GRL-SNAM in procedurally generated 2D deformable navi-  
 394 gation tasks, where a hyperelastic ring must traverse cluttered environments with narrow gaps and  
 395 varying obstacle densities. The robot perceives only a local window of size  $2\hat{d} \times 2\hat{d}$ , from which we  
 396 construct a Hamiltonian energy functional with goal-directed potential  $F_g$ , barrier potentials  $F_{bs}$ ,  
 397 and adaptive coefficients  $(\beta, \gamma, \alpha)$  modulated by context encoders.

399 **Baselines.** We compare against two classical categories under matched information constraints:  
 400 **Global planning:** rigid A\* (obstacle inflation) and deformable A\* (clearance-aware penalty), and  
 401 **Local reactive:** Potential Field (PF), Control Barrier Functions (CBF), and staged DWA using iden-  
 402 tical local windows and stage management as GRL-SNAM. In addition, we include **learning-based**  
 403 **deep RL baselines** (PPO, TRPO, SAC) trained on the same short-rollout dataset and observation  
 404 space as GRL-SNAM.

406 **Metrics.** Success Rate, Success-weighted Path Length (SPL), Detour Ratio, Minimum Clearance,  
 407 Path Smoothness, Collisions, and Mapping Ratio (fraction of environment observed).

### 409 4.1 MAIN RESULTS

411 **Q1. How efficiently does GRL-SNAM trade mapping for navigation quality?** Table 4 demon-  
 412 strates that GRL-SNAM achieves CBF-level navigation quality (SPL = 0.95, Detour = 1.09) while  
 413 using essentially the same minimal map coverage as PF (10.7% vs. CBF’s 11.2%). This validates  
 414 that our stagewise Hamiltonian refinement extracts maximum value per sensed unit of the environ-  
 415 ment.

416 For each deep RL algorithm (PPO, TRPO, SAC) we train three control parameterizations under  
 417 the *same* short-rollout distribution and local observation space as GRL-SNAM: (i) a kinematic controller  
 418 (policy outputs velocities), (ii) a dynamic controller (policy outputs forces integrated by a  
 419 damped point-mass model), and (iii) a coefficient controller (policy outputs the Hamiltonian force-  
 420 field coefficients  $(\alpha, \beta, \gamma)$  used by GRL-SNAM). The aggregated PPO/TRPO/SAC rows in Table 4  
 421 summarize the best-performing configuration for each family; even then, the best TRPO/SAC vari-  
 422 ants reach at most SPL = 0.57 with almost grazing clearances ( $\text{MinClear} \approx 0$ ) and require larger  
 423 mapping ratios ( $\approx 14\text{--}15\%$ ), while PPO collapses to SPL = 0.07 with negative effective clearance.  
 424 In contrast, GRL-SNAM attains high SPL and positive clearance under a strictly smaller sensing  
 425 budget.

427 **Q2. Does GRL-SNAM outperform classical, reactive, and RL planners in complex environ-  
 428 ments?** Yes. Figure 9 shows GRL-SNAM achieves near-perfect success rates ( $\approx 100\%$ ) across  
 429 both in-distribution and out-of-distribution test cases, while all baselines degrade significantly. GRL-  
 430 SNAM consistently maintains high SPL ( $\approx 1.0$ ) with low variance and produces the smoothest tra-  
 431 jectories with lowest turning angles. The Pareto frontier analysis confirms GRL-SNAM uniquely  
 432 dominates the safety-performance trade-off.

Table 1: Navigation quality comparison (success-only runs). GRL-SNAM achieves near-CBF efficiency with minimal mapping budget, while deep RL baselines trained on the same short-rollout distribution and local observations yield lower SPL, larger detours, and smaller clearances.

Method	SPL $\uparrow$	Detour $\downarrow$	Min. Clearance (m) $\uparrow$	Mapping Ratio (%) $\downarrow$
PF	0.77	1.42	0.18	10.3
CBF	0.96	1.04	<b>0.32</b>	11.2
GRL-SNAM	<b>0.95</b>	<b>1.09</b>	0.26	<b>10.7</b>
PPO	0.07	1.65	-0.09	14.7
TRPO	0.57	1.44	0.004	14.3
SAC	0.57	1.53	0.004	14.6

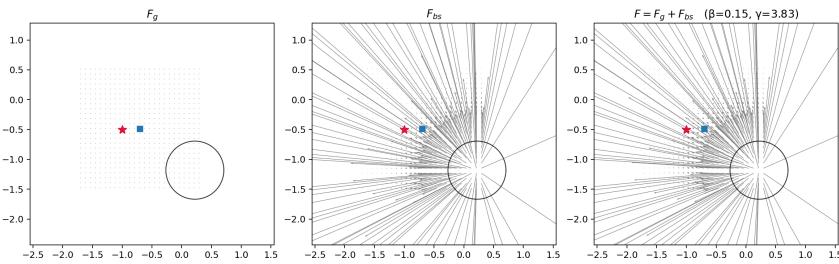


Figure 4: Hamiltonian force field composition. Left: goal force  $F_g$ ; Middle: barrier forces  $F_{bs}$ ; Right: adaptive combination yielding safe, goal-directed trajectories.

**Q3. How does the Hamiltonian formulation enable coherent navigation?** Figure 4 illustrates how GRL-SNAM unifies goal attraction  $F_g$  and barrier repulsion  $F_{bs}$  into a coherent navigation field through adaptive coefficients. Unlike reactive methods that treat forces independently, our differential composition  $F = \beta F_g + \gamma F_{bs}$  creates contextually balanced dynamics that simultaneously pursue goals and avoid obstacles, which aligns with the high SPL and positive clearances observed in Table 4.

**Q4. What distinguishes GRL-SNAM’s online adaptation from standard RL approaches?**  
 Unlike standard deep RL policies that learn a fixed mapping from observations to actions (or coefficients), GRL-SNAM modifies the *entire local energy landscape* as new obstacles are sensed. Figure 10 demonstrates that the coefficients  $(\beta, \gamma, \alpha)$  evolve dynamically to redefine the reduced Hamiltonian itself, ensuring energy-consistent posterior updates rather than heuristic reactive adjustments. This online reshaping of the Hamiltonian explains why GRL-SNAM maintains high success, SPL, and clearance under the same sensing budget where PPO/TRPO/SAC (in all three control variants) either collide, stall, or take inefficient detours.

**Deep RL baselines.** To contextualize the algorithmic contribution of GRL-SNAM, we additionally evaluate strong deep RL baselines (PPO, TRPO, and SAC) implemented under the *same* sensing pipeline, observation structure, action space, and Transformer encoder as our method. For fairness, all baselines operate on the identical short-rollout stagewise dataset derived from the dungeon environment of Liang et al. (2023), with horizon  $H \in [2, 6]$ , stage-exit goals, and locally reconstructed obstacles. Each agent outputs continuous 2D velocity actions and is trained with identical shaped rewards (goal progress, smoothness, and terminal imitation).

Table 2 summarizes performance under these matched conditions: despite millions of interaction steps, PPO/TRPO/SAC achieve only 18–26% success, whereas GRL-SNAM attains 87.5% with an order-of-magnitude fewer updates. This highlights that, even under identical data, sensing, and architectures, GRL-SNAM’s Hamiltonian force-learning offers a substantially more reliable and sample-efficient mechanism for local navigation.

486	Method	Success (%) $\uparrow$	Mean State Error (m) $\downarrow$	Mean Goal Dist. (m) $\downarrow$
487	PPO	26.1	1.8	1.2
488	TRPO	21.7	2.1	1.5
489	SAC	18.4	2.4	1.9
490	GRL-SNAM	<b>87.5</b>	<b>0.3</b>	<b>0.1</b>
491				

492 Table 2: Short-rollout navigation performance under identical sensing, rollouts, and architecture.  
493494  
495 4.2 ADDITIONAL EVALUATIONS  
496497 We conducted comprehensive ablation studies on loss components ( $\mathcal{L}_{\text{friction}}$ ,  $\mathcal{L}_{\text{multi}}$ ) confirming that  
498 friction matching is critical for stability while multi-start robustness prevents over-conservatism. Ro-  
499 bustness evaluations under sensor noise and dynamics perturbations show graceful degradation (87%  
500 success under severe noise vs 99% nominal) due to our adaptive Hamiltonian framework. Sam-  
501 ple efficiency analysis demonstrates faster convergence than RL baselines due to physics-informed  
502 structure.503  
504 4.3 KEY INSIGHTS  
505506 **Minimal mapping suffices:** GRL-SNAM achieves optimal navigation quality using  $\sim 10\%$  environ-  
507 ment coverage, validating the core SNAM principle that local geometric structure contains sufficient  
508 information for global navigation tasks.509 **Hamiltonian unification:** The differential geometric formulation naturally balances competing ob-  
510 jectives (goal-seeking, obstacle avoidance, smoothness) through principled energy minimization  
511 rather than heuristic weight tuning.512 **Principled online adaptation:** By modifying the energy landscape itself rather than just policy  
513 outputs, GRL-SNAM maintains physical consistency while adapting to new sensory information,  
514 enabling robust performance across diverse environments.515 **Superior performance:** GRL-SNAM consistently outperforms classical planning and reactive con-  
516 trol methods across all metrics (success, efficiency, safety, smoothness) while requiring minimal  
517 computational overhead and sensing budget.518 These results establish GRL-SNAM as the first method to successfully unify global navigation objec-  
519 tives with local safety constraints in hyperelastic navigation through principled geometric learning.  
520521  
522 5 CONCLUSION  
523524 We introduced GRL-SNAM, a reinforcement learning framework that leverages Hamiltonian struc-  
525 ture to couple sensing, planning, and deformation into a unified energy-based policy. Our formu-  
526 lation enables stable, feedforward navigation updates and achieves near-optimal path quality with  
527 minimal mapping effort in challenging deformable-robot tasks. The results highlight that incorpo-  
528 rating geometric priors into RL can yield both efficiency and robustness, even under noisy sensing  
529 and out-of-distribution layouts. Future work will extend the approach to richer sensing modalities  
530 and more complex environments, with the goal of validating its scalability to real robotic systems.531  
532 REFERENCES533 J. I. Alora, Moses C. Beard, Thomas Libby, Philipp Rothemund, et al. Discovering dominant dy-  
534 namics for nonlinear continuum robot control. *npj Robotics*, 3(1):5, 2025.  
535536 Karl Johan Åström and Björn Wittenmark. *Adaptive Control*. Courier Corporation, 2010.  
537538 Atılım Güneş Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind.  
539 Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*,  
18(153):1–43, 2018.

540 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn,  
 541 Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian  
 542 Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalash-  
 543 nikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deek-  
 544 sha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez,  
 545 Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi,  
 546 Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vin-  
 547 cent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu,  
 548 and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale, 2023. URL  
 549 <https://arxiv.org/abs/2212.06817>.

550 Brandon Caasenbrood, Alexander Pogromsky, and Henk Nijmeijer. Control-oriented models for  
 551 hyperelastic soft robots through differential geometry of curves. *Soft Robotics*, 9(2):346–361,  
 552 2022.

553 Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topo-  
 554 logical slam for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer*  
 555 *Vision and Pattern Recognition*, pp. 12875–12884, 2020.

556 Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential  
 557 equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

558 Mohammad Dehghani Tezerjani et al. A survey on reinforcement learning applications in slam.  
 559 *arXiv preprint arXiv:2408.14518*, 2024.

560 Shaan A Desai, Marios Mattheakis, David A Roberts, and Pavlos Protopapas. Port-hamiltonian  
 561 neural networks for learning explicit time-dependent dynamical systems. *Physical Review E*, 104  
 562 (3):034312, 2021.

563 Aditya S Ellendula and Chandrajit Bajaj. Self-balancing, memory efficient, dynamic metric space  
 564 data maintenance, for rapid multi-kernel estimation, 2025. URL <https://arxiv.org/abs/2504.18003>.

565 Zhan Feng et al. Safer gap: A gap-based local planner for safe navigation with nonholonomic mobile  
 566 robots. *arXiv preprint arXiv:2303.08243*, 2023.

567 Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-  
 568 critic methods. In *International conference on machine learning*, pp. 1587–1596, 2018.

569 Saurabh Gupta, Varun Tolani, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra  
 570 Malik. Cognitive mapping and planning for visual navigation, 2019. URL <https://arxiv.org/abs/1702.03920>.

571 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy  
 572 maximum entropy deep reinforcement learning with a stochastic actor. In *International confer-  
 573 ence on machine learning*, pp. 1861–1870, 2018.

574 Tai Hoang, Huy Le, Philipp Becker, Vien Anh Ngo, and Gerhard Neumann. Geometry-aware rl for  
 575 manipulation of varying shapes and deformable objects. *arXiv preprint arXiv:2502.07005*, 2025.

576 Junwoo Jang and Maani Ghaffari. Social zone as a barrier function for socially-compliant robot  
 577 navigation, 2024. URL <https://arxiv.org/abs/2405.15101>.

578 Eshagh Kargar and Ville Kyrki. Macrpo: Multi-agent cooperative recurrent policy optimization,  
 579 2021. URL <https://arxiv.org/abs/2109.00882>.

580 Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer,  
 581 et al. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *Proceedings of the*  
 582 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

583 Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Efficient and accurate gradients for  
 584 neural sdes. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

594 Junho Kim, Eun Sun Lee, Mingi Lee, Donsu Zhang, and Young Min Kim. Sgolam: Simultaneous  
 595 goal localization and mapping for multi-object goal navigation, 2021. URL <https://arxiv.org/abs/2110.07171>.

596

597 Donald E Kirk. *Optimal Control Theory: An Introduction*. Dover Publications, 2004.

598

599 Holger Klein, Noémie Jaquier, Andre Meixner, and Tamim Asfour. On the design of region-avoiding  
 600 metrics for collision-safe motion generation on riemannian manifolds, 2023. URL <https://arxiv.org/abs/2307.15440>.

601

602

603 Hengyuan Lai et al. Roboballet: Planning for multirobot reaching with graph neural networks and  
 604 reinforcement learning. *Science Robotics*, 2025.

605

606 Kywoon Lee, Seongun Kim, and Jaesik Choi. Adaptive and explainable deployment of navigation  
 607 skills via hierarchical deep reinforcement learning. In *International Conference on Robotics and*  
 608 *Automation*, 2023.

609

610 Chengshu Li et al. Hrl4in: Hierarchical reinforcement learning for interactive navigation with mo-  
 611 bile manipulators. In *Conference on Robot Learning*, 2020.

612

613 Junjie Li et al. Learn with imagination: Safe set guided state-wise constrained policy optimization.  
 arXiv preprint arXiv:2308.13140, 2023.

614

615 Jingsong Liang, Zhichen Wang, Yuhong Cao, Jimmy Chiun, Mengqi Zhang, and Guillaume Adrien  
 616 Sartoretti. Context-aware deep reinforcement learning for autonomous robotic navigation in un-  
 617 known area. In *Conference on Robot Learning*, pp. 1425–1436. PMLR, 2023.

618

619 Elisabetta Liu and Cosimo Della Santina. Physics-informed neural networks to model and control  
 620 robots: a theoretical and experimental investigation. *Advanced Intelligent Systems*, 2024.

621

622 David Martínez-Rubio and Sebastian Pokutta. Accelerated riemannian optimization: Handling con-  
 623 straints with a prox to bound geometric penalties. In *The Thirty Sixth Annual Conference on*  
 624 *Learning Theory*, pp. 359–393. PMLR, 2023.

625

626 Nicholas Mohammad and Nicola Bezzo. Soft actor-critic-based control barrier adaptation for robust  
 627 autonomous navigation in unknown environments, 2025. URL <https://arxiv.org/abs/2503.08479>.

628

629 Barry W Mulvey and Thrishantha Nanayakkara. Haven: haptic and visual environment navigation  
 630 by a shape-changing mobile robot with multimodal perception. *Scientific Reports*, 14(1):27018,  
 2024.

631

632 Minh Nguyen and Chandrajit Bajaj. A differential and pointwise control approach to reinforcement  
 633 learning, 2025. URL <https://arxiv.org/abs/2404.15617>.

634

635 Tariq Patanam, Eli Shayer, and Younes Bensouda Mourri. Deep dagger imitation learning for indoor  
 636 scene navigation.

637

638 L.S. Pontryagin, V.G. Boltyanskii, R.V. Gamkrelidze, and E.F. Mishchenko. *The Mathematical*  
*Theory of Optimal Processes*. Interscience, 1962.

639

640 Doina Precup. *Temporal abstraction in reinforcement learning*. University of Massachusetts  
 Amherst, 2000.

641

642 F. Qi, C. Zhou, H. Qing, H. Sun, and J. Yin. Aerial track-guided autonomous soft ring robot.  
 Advanced Science, 2024.

643

644 Mohammad Roshanfar, Javad Dargahi, and Amir Hooshiar. Hyperelastic modeling and validation  
 645 of hybrid-actuated soft robot with pressure-stiffening. *Micromachines*, 14(5):1001, 2023.

646

647 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy  
 648 optimization algorithms. 2017.

648 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 649 Poole. Score-based generative modeling through stochastic differential equations. In *International  
 650 Conference on Learning Representations (ICLR)*, 2021.

651

652 Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and po-  
 653 sitioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer  
 654 Vision*, pp. 6229–6238, 2021.

655 Richard S Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: Learning,  
 656 planning, and representing knowledge at multiple temporal scales. Technical report, Technical  
 657 Report 98-74, University of Massachusetts, Amherst, 1998.

658

659 Hamid Taheri et al. Deep reinforcement learning with enhanced ppo for safe mobile robot naviga-  
 660 tion. *arXiv preprint arXiv:2405.16266*, 2024.

661

662 Lei Tai, Jingwei Zhang, Ming Liu, and Wolfram Burgard. Socially compliant navigation through  
 663 raw depth inputs with generative adversarial imitation learning, 2018. URL <https://arxiv.org/abs/1710.02543>.

664

665 Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David  
 666 Silver, and Koray Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. In  
 667 *International conference on machine learning*, pp. 3540–3549. PMLR, 2017.

668

669 Niclas Vödisch, Daniele Cattaneo, Wolfram Burgard, and Abhinav Valada. *Continual SLAM: Be-  
 670 yond Lifelong Simultaneous Localization and Mapping Through Continual Learning*, pp. 19–35.  
 671 Springer Nature Switzerland, 2023. ISBN 9783031255557. doi: 10.1007/978-3-031-25555-7\_3.  
 672 URL [http://dx.doi.org/10.1007/978-3-031-25555-7\\_3](http://dx.doi.org/10.1007/978-3-031-25555-7_3).

673

674 Tianhao Wang et al. Pinn-ray: A physics-informed neural network to model soft robotic fin ray  
 675 fingers. *arXiv preprint arXiv:2407.08222*, 2024.

676

677 Weizheng Wang et al. Multi-agent llm actor-critic framework for social robot navigation. *arXiv  
 678 preprint arXiv:2503.09758*, 2025.

679

680 Chi Yan, Delin Qu, Dan Wang, Dan Xu, Zhigang Wang, et al. Gs-slam: Dense visual slam with 3d  
 681 gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern  
 682 Recognition*, 2024.

683

684 Songyuan Zhang, Zhangjie Cao, Dorsa Sadigh, and Yanan Sui. Confidence-aware imitation learn-  
 685 ing from demonstrations with varying optimality, 2022. URL <https://arxiv.org/abs/2110.14754>.

686

687 Chao Zheng et al. Semantic slam system for mobile robots based on large visual model in complex  
 688 environments. *Scientific Reports*, 15(1):1–15, 2025.

689

690 Jun Zhu, Zihao Du, Haotian Xu, Fengbo Lan, Zilong Zheng, Bo Ma, Shengjie Wang, and Tao  
 691 Zhang. Navi2gaze: Leveraging foundation models for navigation and target gazing, 2024. URL  
 692 <https://arxiv.org/abs/2407.09053>.

693

694 Zihan Zhu, Songyou Peng, Viktor Laehner, Weiyang Xu, Michael Niemeyer, et al. Nice-slam:  
 695 Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on  
 696 Computer Vision and Pattern Recognition*, pp. 12786–12796, 2022.

697

698

699

700

701

702 A EXTENDED INTRODUCTION AND MOTIVATION  
703704  
705 This section provides expanded context for the challenges addressed by GRL-SNAM and detailed  
706 justification for our geometric approach.  
707708 A.1 COMPREHENSIVE ANALYSIS OF RL LIMITATIONS IN NAVIGATION  
709710 Contemporary reinforcement learning methods face several critical limitations that become particu-  
711 larly pronounced in continuous navigation tasks:  
712713 **Sample Efficiency Bottlenecks.** Standard RL algorithms like SAC Haarnoja et al. (2018) and PPO  
714 Schulman et al. (2017) require millions of environment interactions to learn effective navigation  
715 policies. This inefficiency stems from the curse of dimensionality in continuous control settings  
716 where the action space is infinite-dimensional and policies must simultaneously master fine-grained  
717 motor control and high-level strategic reasoning. In real-world deployment scenarios where data  
718 collection is expensive and potentially dangerous, this sample complexity becomes prohibitive.  
719720 The problem is exacerbated by the need for exploration in high-dimensional spaces. Unlike discrete  
721 control problems where systematic exploration strategies like  $\epsilon$ -greedy or UCB can provide theo-  
722 retical guarantees, continuous control requires sophisticated exploration mechanisms that often rely  
723 on injected noise or entropy bonuses. These mechanisms frequently lead to unsafe or inefficient  
724 exploration behaviors that are unsuitable for real-world navigation tasks.  
725726 **Generalization Failures.** Policies trained in specific environments exhibit catastrophic performance  
727 degradation when deployed in novel settings, even when new environments share similar structure.  
728 This brittleness stems from the lack of inductive bias in standard neural network architectures. With-  
729 out explicit encoding of physical principles or geometric structure, learned policies tend to memorize  
730 environment-specific features rather than discovering generalizable navigation principles.  
731732 The generalization problem is particularly acute in navigation because environmental variations can  
733 affect multiple aspects of the task simultaneously: obstacle configurations change collision con-  
734 straints, surface properties affect dynamics, and lighting conditions influence perception. Standard  
735 RL approaches learn monolithic mappings that cannot decompose these variations into their con-  
736 stituent factors, leading to brittle behaviors that fail when any component deviates from training  
737 conditions.  
738739 **Temporal Decomposition Challenges.** Navigation inherently requires coordination across multiple  
740 timescales: immediate obstacle avoidance operates on millisecond timescales, local path planning  
741 unfolds over seconds, and strategic goal-directed behavior spans minutes or hours. Standard RL al-  
742 gorithms struggle to learn policies that reason effectively across these scales, often getting trapped in  
743 locally optimal behaviors that satisfy short-term objectives while failing to make long-term progress.  
744745 Existing approaches to multi-scale reasoning such as hierarchical RL Sutton et al. (1998), options  
746 frameworks Precup (2000), or feudal networks Vezhnevets et al. (2017), typically require manual  
747 decomposition of the task space and careful engineering of reward functions for different levels.  
748 These methods introduce additional complexity without fundamentally addressing the structural  
749 issues that make multi-scale learning difficult.  
750751 A.2 THE SNAM CHALLENGE: WHY STRUCTURE MATTERS  
752753 Simultaneous Navigation and Mapping (SNAM) represents a particularly challenging instance of the  
754 navigation problem where agents must build environmental representations online while traversing  
755 unknown spaces. This challenge amplifies the limitations of conventional RL approaches in several  
756 ways:  
757758 **Memory and Representation Learning.** SNAM requires policies to maintain and update spatial  
759 representations based on sensory observations. This places enormous demands on the policy’s mem-  
760 ory architecture, requiring it to simultaneously master memory management, spatial reasoning, and  
761 motor control. Standard recurrent architectures like LSTMs or GRUs struggle with this multifaceted  
762 learning problem, often failing to maintain coherent spatial representations over long episodes.  
763

756 **Exploration-Exploitation Tradeoffs.** In SNAM, exploration serves dual purposes: gathering information about the environment for mapping and discovering navigation strategies. This creates complex exploration-exploitation tradeoffs that standard RL exploration mechanisms cannot handle effectively. Random exploration may discover new regions but fails to systematically map environmental structure, while directed exploration based on current maps may miss critical environmental features.

762 **Dynamic Environmental Coupling.** Unlike traditional navigation where environments are static, SNAM requires reasoning about how the agent’s actions affect both its position and its knowledge of the environment. This creates a coupled learning problem where navigation decisions influence future mapping accuracy, and mapping quality affects navigation performance. Standard RL frameworks treat these as separate problems, missing the critical coupling that enables efficient SNAM.

767 Recent approaches in simultaneous navigation and mapping (SNAM) have coupled local mapping with policy learning to improve navigation performance. For example, SGOLAM Kim et al. (2021) interleaves goal localization with occupancy mapping to enable point-goal navigation, while Cognitive Mapping and Planning (CMP) Gupta et al. (2019) integrates a differentiable planner into a learned mapping framework. Continual SLAM (CL-SLAM) Vödisch et al. (2023) further emphasizes long-term adaptability by maintaining and updating maps during navigation. However, these methods rely on progressively constructing detailed maps of the environment before exploiting them for navigation. In contrast, our objective is to *reach the goal along high-quality, well-weighted paths while mapping as little of the unknown environment as possible*. To the best of our knowledge, no prior work explicitly formulates navigation with minimal exploration as the central goal. Our proposed GRL-SNAM framework achieves this by progressively refining paths: from observed environmental variations, the policy differentially learns to identify the least-cost trajectory, such that the path improves continuously as new local information is revealed.

### 780

### 781

### 782 A.3 GEOMETRIC STRUCTURE: THE INEVITABLE SOLUTION

### 783

784 The limitations outlined above are not merely implementation details but fundamental consequences of treating navigation as unstructured optimization. Several lines of evidence suggest that geometric structure is not just helpful but inevitable for solving complex navigation problems:

788 **Physical Realizability.** Real robotic systems operate under physical constraints imposed by conservation laws, kinematic limitations, and actuator dynamics. Policies that violate these constraints 789 cannot be implemented on physical systems, yet standard RL approaches have no mechanism to 790 enforce such constraints during learning. Geometric formulations naturally incorporate physical 791 constraints through the mathematical structure of the problem.

793 **Stability Requirements.** Long-horizon navigation requires numerical stability over extended rollouts. Standard neural network policies accumulate errors over time, leading to unstable behaviors 794 in long episodes. Hamiltonian formulations with symplectic structure preserve important invariants 795 (energy, momentum) that ensure stability over arbitrarily long rollouts.

797 **Compositionality Needs.** Complex navigation tasks require composing simpler behaviors: obstacle 798 avoidance, path following, goal seeking, and environmental adaptation. Standard RL approaches 799 learn monolithic policies that cannot decompose into interpretable components. Geometric 800 formulations enable natural decomposition through energy terms that can be composed, weighted, and 801 adapted independently.

### 802

### 803

### 804 A.4 DIFFERENTIAL POLICY OPTIMIZATION: BEYOND FIXED POLICIES

### 805

807 Traditional RL optimizes fixed policy parameters  $\theta$  to maximize expected returns over discrete 808 timesteps. Differential Policy Optimization Nguyen & Bajaj (2025) approach fundamentally reconceptualizes 809 this by learning dynamics operators through a continuous-time differential dual formulation.

**Mathematical Foundation.** Rather than directly learning policies, DPO reformulates RL through continuous-time optimal control. By approximating discrete reward sums with time integrals:

$$\max_{\pi} \mathbb{E} \left[ \sum_{k=0}^{H-1} r(s_k, a_k) \right] \approx \max_{\pi} \mathbb{E} \left[ \int_0^T r(s_t, a_t) dt \right] \quad (10)$$

Applying Pontryagin’s Maximum Principle introduces adjoint variables  $p$  and defines the Hamiltonian function:

$$HF(p, s, a) := p^T f(s, a) - r(s, a) \quad (11)$$

The key insight is that optimal actions can be implicitly represented through the stationarity condition  $\frac{\partial HF}{\partial a} = 0$ , yielding the reduced Hamiltonian:

$$hf(s, p) := HF(s, p, a^*(s, p)) \quad (12)$$

**Score Function Learning.** DPO learns a score function  $g(x) \approx hf(x)$  where  $x = (s, p)$  combines state and adjoint variables. The dynamics operator is constructed as:

$$G(x) = x + \Delta S \nabla g(x) \quad (13)$$

where  $S = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$  is the canonical symplectic matrix and  $\Delta$  is the discretization step.

**Stagewise Learning Advantages.** Unlike methods requiring backward-in-time adjoint calculations (as in Pontryagin’s Maximum Principle), DPO enables feedforward learning where each stage  $t$  defines a local Hamiltonian  $\mathcal{H}_t$  integrated forward in time:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{H}_t \quad (14)$$

This avoids the computational complexity and numerical instability of adjoint methods while maintaining theoretical guarantees through the geometric structure of the Hamiltonian formulation.

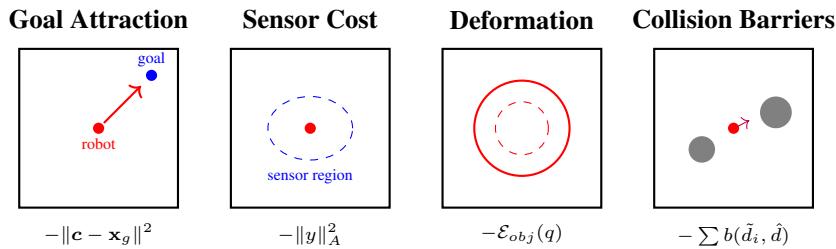


Figure 5: Policy-aligned energy decomposition. Each policy governs a distinct energy component: the Sensor Policy minimizes sensor cost, the FPE balances goal attraction and collision barriers, and the Reconfig Policy adapts size through deformation energy. Together these terms define the Hamiltonian reward  $\mathcal{R}$ .

### A.5 MULTI-POLICY ARCHITECTURE DETAILS

Our multi-policy decomposition addresses temporal scale separation through three specialized components operating at different timescales:

**Sensor Policy ( $\pi_y$ ):** Operates at slow timescales to adapt perception strategies based on stagewise environmental feedback. This policy learns to focus attention on relevant environmental features, adjust sensor parameters for optimal information gain, and filter sensory noise. The sensor policy outputs constraints  $\mathcal{C}_t$  that inform slower planning processes.

**Frame Policy ( $\pi_f$ ):** Operates at medium timescales to plan collision-free trajectories in local coordinate frames. This policy takes constraints from the sensor policy and generates waypoints  $\mathcal{W}_t$  for shape control. The frame policy handles local obstacle avoidance and path optimization within a limited spatial horizon.

864 **Shape Policy** ( $\pi_o$ ): Operates at fast timescales to control robot morphological adaptation. For  
 865 deformable robots, this includes shape changes, stiffness modulation, and configuration updates. For  
 866 conventional robots, this might include gait transitions, tool selection, or behavioral mode switches.  
 867

868 The key insight is that these policies are not manually designed hierarchies but emerge naturally  
 869 from the temporal structure of the Hamiltonian dynamics. Fast variables (sensor adaptation) reach  
 870 quasi-equilibrium before slower variables (shape changes) evolve significantly, creating natural scale  
 871 separation without manual decomposition.

872 This extended analysis demonstrates that geometric structure is not merely a useful inductive bias but  
 873 a necessary foundation for solving complex navigation problems that require multi-scale reasoning,  
 874 online adaptation, and long-horizon stability.

## 876 B EXTENDED RELATED WORK SURVEY

### 877 B.1 GEOMETRY AND MECHANICS PRIMER

878 Navigation learning methods can be categorized by their underlying mathematical spaces, with sig-  
 879 nificant implications for performance and theoretical guarantees:

880 **Euclidean Space Methods** ( $\mathbb{R}^n$ ): Standard RL treats navigation as optimization in flat spaces using  
 881 Euclidean distance metrics. Enhanced PPO Taheri et al. (2024) demonstrate improved collision  
 882 avoidance but ignore inherent geometric structure of robotic systems. Sample efficiency remains  
 883 poor, typically requiring millions of environment interactions Dehghani Tezerjani et al. (2024).

884 **Lie Group Methods:** Recognition of orientation constraints has led to SE(2) and SE(3) formula-  
 885 tions using equivariant neural architectures. These preserve rotational and translational symmetries  
 886 but remain primarily limited to manipulation rather than navigation tasks.

887 **Riemannian Manifold Approaches:** Advanced geometric formulations employ differential geo-  
 888 metry for constraint handling through tangent space projections. Martínez-Rubio & Pokutta  
 889 (2023) demonstrates constraint satisfaction through geometric structure rather than penalty meth-  
 890 ods, achieving superior theoretical properties but limited practical deployment.

891 **Hamiltonian and Symplectic Methods:** Port-Hamiltonian neural networks show significant per-  
 892 formance improvements through symplectic integrators, proving that respecting geometric structure  
 893 fundamentally improves learning dynamics. However, applications remain confined to simple con-  
 894 trol problems.

### 895 B.2 SAFETY-CRITICAL NAVIGATION TAXONOMY

896 **External Safety Projection:** Control Barrier Functions create safe action spaces through constraint  
 897 projection. Neural Network Zeroing Barrier Functions Feng et al. (2023) enable collision-free nav-  
 898 igation, while adaptive safety constraints Mohammad & Bezzo (2025) handle dynamic environ-  
 899 ments. Social navigation approaches Jang & Ghaffari (2024) extend CBFs to human-robot interac-  
 900 tion. These methods achieve formal safety guarantees but often exhibit conservative behaviors due  
 901 to the separation between safety and optimality.

902 **Energy-Integrated Safety:** Our approach incorporates safety directly within the Hamiltonian en-  
 903 ergy structure via barrier potentials. This enables aggressive navigation while maintaining formal  
 904 guarantees through symplectic structure preservation, avoiding the conservatism of external projec-  
 905 tion methods.

### 906 B.3 DEFORMABLE AND SOFT ROBOT NAVIGATION

907 **Hyperelastic Material Models:** Recent advances include pressure-stiffening control with 6.40%  
 908 maximum error validation Roshanfar et al. (2023) and passivity-based control using differential  
 909 geometry of curves Caasenbrood et al. (2022). Spectral Submanifold Reduction Alora et al. (2025)  
 910 achieves computational speedup for real-time hyperelastic control with stability guarantees.

911 **Ring and Circular Robots:** Liquid Crystal Elastomer responses enable aerial gap navigation Qi  
 912 et al. (2024) through predetermined actuation patterns. HAVEN Mulvey & Nanayakkara (2024)

navigates constrained spaces via fixed shape-changing sequences based on multimodal perception. These approaches use offline parameter optimization with deterministic execution, lacking online adaptation capabilities.

**Physics-Informed Learning:** PINN-Ray Wang et al. (2024) achieves state-of-the-art hyperelastic displacement prediction, while extensions to non-conservative effects Liu & Della Santina (2024) provide experimental validation. However, these remain primarily modeling tools rather than adaptive control frameworks.

#### B.4 NEURAL SCENE REPRESENTATIONS FOR NAVIGATION

**NeRF-Based SLAM:** Real-time dense reconstruction through NICE-SLAM Zhu et al. (2022) and keyframe-free tracking via iMAP Sucar et al. (2021) provide rich environmental representations. Neural Topological SLAM Chaplot et al. (2020) combines learning with classical planning, while semantic approaches Zheng et al. (2025) integrate large vision models.

**3D Gaussian Splatting:** GS-SLAM Yan et al. (2024) and SplatAM Keetha et al. (2024) demonstrate state-of-the-art reconstruction quality with real-time performance, offering dense 3D representations suitable for navigation applications.

**Integration with Energy Terms:** Scene representations feed our energy formulation through:

$$\text{Barrier Energy: } \mathcal{U}_{\text{barrier}} = \sum_{\text{obstacles}} b(\text{SDF}(\mathbf{x})) \quad (15)$$

$$\text{Free-Space Energy: } \mathcal{U}_{\text{free}} = - \sum_{\text{free regions}} w(\mathbf{x}) \quad (16)$$

$$\text{Goal Energy: } \mathcal{U}_{\text{goal}} = \|\mathbf{x} - \mathbf{x}_{\text{goal}}\|^2 \quad (17)$$

#### B.5 MULTI-SCALE AND HIERARCHICAL METHODS

**Hierarchical RL:** Task decomposition approaches like HRL4IN Li et al. (2020) handle heterogeneous navigation phases, while Lee et al. (2023) learns specialized policy families with high-level coordination. These require manual decomposition and struggle with principled coordination, often leading to ad-hoc design choices without theoretical guarantees.

**Multi-Agent Coordination:** RoboBallet Lai et al. (2025) achieves coordination for 8 robots across 40 tasks using graph neural networks. MACRPO Kargar & Kyrki (2021) enhances information sharing beyond parameter sharing. However, these approaches lack the geometric structure preservation critical for deformable robot coordination.

#### B.6 IMITATION LEARNING FOR NAVIGATION

**Behavioral Cloning:** RT-1 Brohan et al. (2023) demonstrates impressive generalization across 700+ tasks using 130k demonstration episodes with transformer architectures achieving significant zero-shot performance improvements.

**Inverse Reinforcement Learning:** GAIL for Safe Navigation Tai et al. (2018) combines generative adversarial imitation with safety constraints. DAgger for Continuous Navigation Patanam et al. iteratively improves policies through expert querying.

**Sub-Optimal Demonstrations:** Confident Imitation Learning Zhang et al. (2022) handles demonstration uncertainty through confidence-aware training, addressing distribution shift in novel environments.

These approaches excel with high-quality demonstrations but assume expert availability and struggle with the full behavioral range needed for adaptive deformation strategies.

#### B.7 FOUNDATION MODEL INTEGRATION

Large-scale models for navigation reasoning Zhu et al. (2024); Wang et al. (2025) focus on high-level semantic understanding and multi-agent coordination at the symbolic level. Foundation models

972 excel at reasoning and semantic understanding, while our GRL-SNAM provides principled low-level  
 973 geometric control.

974 **Integration Pathway:** Foundation models could generate high-level objectives encoded as potential  
 975 energy terms in our energy functional  $\mathcal{R}(q_t)$ . The geometric structure preservation ensures high-  
 976 level semantic goals translate into physically consistent behaviors, addressing the critical gap where  
 977 foundation model outputs often lack grounding in physical dynamics.

## 979 B.8 PARADIGM COMPARISON

981 Table 3: Extended paradigm-level comparison of learning frameworks. Scoring:  $\checkmark$  = comprehensive  
 982 support,  $\triangle$  = limited support,  $\times$  = not supported.

984 Capability	985 GRL-SNAM	986 Standard RL	987 Geometric RL	988 Imitation Learning	989 Semi/Unsupervised	990 CBF Methods	991 Hierarchical RL	992 Foundation Models
Energy Conservation	$\checkmark$	$\times$	$\triangle$	$\times$	$\times$	$\times$	$\times$	$\times$
Geometric Structure	$\checkmark$	$\times$	$\checkmark$	$\times$	$\triangle$	$\triangle$	$\times$	$\times$
Constraint Integration	$\checkmark$	$\triangle$	$\checkmark$	$\times$	$\times$	$\checkmark$	$\triangle$	$\triangle$
Online Adaptation	$\checkmark$	$\triangle$	$\triangle$	$\times$	$\checkmark$	$\triangle$	$\checkmark$	$\checkmark$
Multi-Scale Coordination	$\checkmark$	$\times$	$\times$	$\times$	$\times$	$\times$	$\checkmark$	$\triangle$
Sample Efficiency	$\checkmark$	$\times$	$\triangle$	$\checkmark$	$\triangle$	$\triangle$	$\triangle$	$\checkmark$
Zero-Shot Generalization	$\checkmark$	$\times$	$\triangle$	$\times$	$\checkmark$	$\times$	$\times$	$\checkmark$
Real-World Deployment	$\checkmark$	$\checkmark$	$\triangle$	$\checkmark$	$\triangle$	$\checkmark$	$\checkmark$	$\triangle$
Deformable Robot Support	$\checkmark$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$

### Scoring Criteria:

- **Energy Conservation:** Explicit conservation laws in dynamics
- **Geometric Structure:** Preservation of manifold properties
- **Constraint Integration:** Safety/task constraints within optimization
- **Online Adaptation:** Real-time policy modification during deployment
- **Multi-Scale Coordination:** Principled coordination across temporal scales
- **Sample Efficiency:** Learning with minimal environment interaction
- **Zero-Shot Generalization:** Performance in unseen environments
- **Real-World Deployment:** Practical implementation feasibility
- **Deformable Robot Support:** Explicit modeling of shape change

1003 This comprehensive survey positions GRL-SNAM as uniquely addressing the intersection of geo-  
 1004 metric structure preservation, multi-scale coordination, and deformable robot control—capabilities  
 1005 that existing approaches handle separately or incompletely.

## 1006 B.9 KEY INSIGHTS

1008 Our framework builds upon a set of Hamiltonian and reinforcement learning principles, unifying  
 1009 offline reference dynamics with online adaptive updates. Below, we summarize the six key insights  
 1010 that form the backbone of GRL-SNAM.

### 1012 1. Hamiltonian energy as task reward.

$$1013 \quad \mathcal{H}(q, p) = K(p) + P(q), \quad (18)$$

1015 with kinetic energy  $K$  and task-specific potential  $P$ . In our setup,  $P$  encodes navigation objectives  
 1016 (goal attraction, barrier avoidance, deformation penalties). Following Pontryagin et al. (1962); ?,  
 1017 the Hamiltonian coincides with the surrogate objective in policy gradient methods, i.e.

$$1018 \quad \nabla_{\theta} J(\pi_{\theta}) \approx \nabla_{\theta} \mathbb{E}_{\pi_{\theta}}[-\mathcal{H}(q, p)], \quad (19)$$

1019 linking task reward to the Hamiltonian gradient flow. This equivalence grounds the DPO surrogate  
 1020 in a physical structure.

1022 **2. Offline Hamiltonian vs. Online task reward.** In offline training, the agent minimizes trajec-  
 1023 tories under a fixed  $\mathcal{H}$  constructed from synthetic local patches. Online, the environment is sensed,  
 1024 and task rewards  $\mathcal{R}_{\text{env}}$  are parsed into Hamiltonian subtasks. By interpreting

$$1025 \quad \mathcal{H}_{\text{online}} = \mathcal{H}_{\text{offline}} + \Delta \mathcal{R}_{\text{env}}, \quad (20)$$

1026 we align local sensory updates with the reference offline Hamiltonian. This mirrors the adaptive  
 1027 control interpretation in Åström & Wittenmark (2010).  
 1028

1029 **3. Offline policy as reference Hamiltonian.** Every offline policy  $\pi_{\text{ref}}$  is equivalent to a reference  
 1030 Hamiltonian  $\mathcal{H}_{\text{ref}}$ , where the score function  $s^\theta = \nabla \mathcal{H}_{\text{ref}}$  defines canonical dynamics:  
 1031

$$1032 \quad \dot{q} = \frac{\partial \mathcal{H}_{\text{ref}}}{\partial p}, \quad \dot{p} = -\frac{\partial \mathcal{H}_{\text{ref}}}{\partial q}. \quad (21)$$

1034 Online adaptation then minimizes the divergence  
 1035

$$1036 \quad D(\pi_{\text{online}} \parallel \pi_{\text{ref}}) \propto \mathbb{E}[\|\nabla \mathcal{H}_{\text{online}} - \nabla \mathcal{H}_{\text{ref}}\|^2], \quad (22)$$

1037 a structure exploited in score-based models (Song et al., 2021).  
 1038

1039 **4. Advantages of stagewise updates.** Rather than solving adjoint equations as in Pontryagin’s  
 1040 Maximum Principle, we adopt a stagewise decomposition. Each stage defines a local  $\mathcal{H}_t$  and is  
 1041 integrated feedforward:  
 1042

$$\theta_{t+1} = \theta_t - \eta \nabla_\theta \mathcal{H}_t. \quad (23)$$

1043 This avoids backward-in-time adjoint calculations and recovers the efficiency noted in adjoint-free  
 1044 feedforward networks (Chen et al., 2018; Kidger et al., 2021).  
 1045

1046 **5. Universality of the pipeline.** Our pipeline  
 1047

$$1048 \quad \text{Environment} \xrightarrow{\text{Encoder}} \text{Context} \xrightarrow{\text{Setup}} \mathcal{H}_{\text{adapted}} \quad (24)$$

1050 is universal. As long as  $\mathcal{H}$  is differentiable, adaptation reduces to evaluating its gradients, regardless  
 1051 of whether the system is white-box (explicit potentials) or black-box (sensor-level inputs). This  
 1052 follows from the variational formulation of differentiable programming (Baydin et al., 2018).  
 1053

1054 **6. Navigator as meta-controller.** The navigator policy  $\pi_{\text{nav}}$  interacts with three black boxes: the  
 1055 offline Hamiltonian  $\mathcal{H}_{\text{ref}}$ , the online sensed reward  $\mathcal{R}_{\text{env}}$ , and the adaptive fusion  $\mathcal{H}_{\text{adapt}}$ . Its role is to  
 1056 formulate and solve  
 1057

$$\mathcal{H}_{\text{adapt}} = \alpha \mathcal{H}_{\text{ref}} + (1 - \alpha) \mathcal{R}_{\text{env}}, \quad (25)$$

1058 where  $\alpha$  is dynamically updated by the context encoder (e.g., LSTM). This positions the navigator  
 1059 as a meta-controller that continually reforms the Hamiltonian problem, a principle consistent with  
 1060 adaptive RL formulations in Kirk (2004).  
 1061

## 1062 B.10 SECANT GAUSS-NEWTON CONTROLLER WITH IMPLICIT OBSERVABLE TRANSITION

1063 In the online setting, policy parameters ( $\Theta = [\beta, \gamma, \{\alpha_i\}]$ ) do not act directly on target observables  
 1064  $y_{\text{tgt}}$ . Instead, there exists an implicit transition chain:  
 1065

$$1066 \quad \Theta \xrightarrow{\text{policy-induced energy}} \mathcal{E}_\Theta \xrightarrow{\text{dynamics rollout}} z \mapsto f(\Theta) \xrightarrow{\text{task map}} y_{\text{tgt}}, \quad (26)$$

1068 where  $f(\Theta)$  denotes observables (e.g., clearance, progress, admissible speed) produced after sym-  
 1069 plectic updates under the reshaped Hamiltonian  $\mathcal{E}_\Theta$ .  
 1070

1071 **Implicit Jacobian.** The true Jacobian of observables with respect to policy parameters is  
 1072

$$1073 \quad \frac{\partial f}{\partial \Theta} = \frac{\partial f}{\partial z} \cdot \frac{\partial z}{\partial \Theta}, \quad (27)$$

1075 which is expensive to evaluate through full rollouts. Instead, we maintain a rank-1 secant estimate  
 1076

$$1077 \quad \hat{J}_t \approx \frac{f(\Theta^t) - f(\Theta^{t-1})}{\Theta^t - \Theta^{t-1}}, \quad (28)$$

1078 smoothed via exponential moving average to reduce noise. This captures the \*implicit effect\* of  
 1079 parameter changes on observables.  
 1080

1080 **Energy Reshaping Step.** With the secant Jacobian, the Gauss–Newton update becomes  
 1081

$$1082 \Delta\Theta = -\widehat{G}^{-1}\widehat{J}^\top W(f(\Theta) - y_{\text{tgt}}), \quad \widehat{G} = \widehat{J}^\top W\widehat{J} + \varepsilon I, \quad (29)$$

1083 followed by projection and per-head learning rates:  
 1084

$$1085 \Theta^{t+1} = \text{Proj}_{\Theta \geq 0}(\Theta^t + \text{diag}(\eta_b, \eta_g, \eta_\alpha) \Delta\Theta). \quad (30)$$

1087 Here,  $W$  is a weighting matrix over observables, and  $\varepsilon I$  stabilizes inversion.  
 1088

1089 **Meta-Learning View.** The Navigator maintains a meta-objective  
 1090

$$1091 \mathcal{L}_{\text{nav}}(\Theta; \phi) = \ell(f(\Theta; z), y_{\text{tgt}}), \quad (31)$$

1092 with context-dependent parameters  $\phi$  (e.g., encoder weights for  $\alpha$ ). The implicit gradient is  
 1093

$$1094 \nabla_\Theta \mathcal{L}_{\text{nav}} = \widehat{J}^\top W(f(\Theta) - y_{\text{tgt}}), \quad (32)$$

1095 which coincides with the secant Gauss–Newton step above. Thus, meta-learning is implemented not  
 1096 by direct regression on  $\Theta$ , but by observable alignment through the implicit policy  $\rightarrow$  observables  
 1097  $\rightarrow$  targets chain.  
 1098

1099 **Interpretation.** This formulation clarifies that the controller does not operate in parameter space  
 1100 alone. Instead, it continuously reshapes Hamiltonian parameters so that the *induced observables*  
 1101 approach task targets, effectively coupling multiple energy components (safety, progress, speed)  
 1102 without additional rollouts.  
 1103

1104 **Sequential Query–Response (as ports).** At each  $t$  the Navigator issues queries  $\mathcal{Q}_k^t$  and receives  
 1105 responses  $\mathcal{R}_k^t$ , which determine  $y_{\text{tgt}}(\mathcal{R}_k^t)$  and any weights in  $\mathbf{W}$ ; the update equation ?? (with the  
 1106 secant  $\widehat{\mathbf{J}}$ ) is then applied to each block  $k \in \{y, f, o\}$ :

$$1107 \theta_k^{t+1} = \theta_k^t + h \Delta\theta_k^t, \quad \Delta\theta_k^t = -\widehat{\mathbf{G}}_k^{-1}\widehat{\mathbf{J}}_k^\top \mathbf{W}_k(f_k(\theta_k^t) - y_{\text{tgt},k}(\mathcal{R}_k^t)).$$

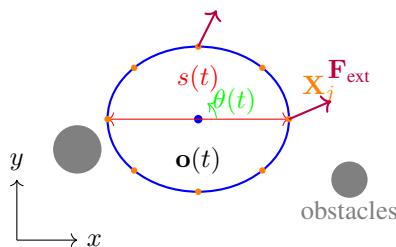
1109 **State Evolution.** With updated parameters, each policy advances its state as before,  
 1110

$$1111 z_k^{t+1} = z_k^t + \tau_k J_k s_k^{\theta_k^{t+1}}(z_k^t, \mathcal{C}_t^{\text{updated}}, t).$$

1112 *Remark.* If a strict port-Hamiltonian view is desired, the  $\mathbf{J}^\top \mathbf{W}(y_{\text{tgt}} - f(\Theta))$  term enters  $\dot{\Pi}$  as  
 1113 an external port (input) rather than being baked into the potential; the resulting discrete update is  
 1114 identical to the secant Gauss–Newton step above.  
 1115

## 1116 C HYPERELASTIC RING ROBOT MODEL

1117 We model the deformable robot as a closed hyperelastic ring with reduced-order dynamics to enable  
 1118 efficient navigation while capturing essential deformation behaviors.  
 1119



1120 Figure 6: Hyperelastic ring robot model showing generalized coordinates  $(s, o, \theta)$ , spline sample  
 1121 points  $\mathbf{X}_j$ , and external force fields.  
 1122

1134 C.1 GEOMETRIC REPRESENTATION  
11351136 The robot boundary is defined by a periodic cubic B-spline curve with  $n_{\text{ctrl}}$  control points:  
1137

1138 
$$\mathbf{S}(u) = \sum_{i=1}^{n_{\text{ctrl}}} N_{i,3}(u) \mathbf{P}_i, \quad u \in [0, 1] \quad (33)$$
  
1139

1140 where  $N_{i,3}(u)$  are degree-3 B-spline basis functions with  $C^2$  continuity. The base shape is a unit  
1141 circle:  
1142

1143 
$$\mathbf{P}_{0,i} = r_{\text{base}} \begin{bmatrix} \cos(2\pi i / n_{\text{ctrl}}) \\ \sin(2\pi i / n_{\text{ctrl}}) \end{bmatrix} \quad (34)$$
  
1144

1145 World coordinates are computed via similarity transformation:  
1146

1147 
$$\mathbf{P}_i(t) = \mathbf{o}(t) + s(t) \mathbf{R}(\theta(t)) \mathbf{P}_{0,i} \quad (35)$$
  
1148

1149 where  $s(t)$  is uniform scale,  $\mathbf{o}(t) \in \mathbb{R}^2$  is center position,  $\theta(t)$  is orientation.  
11501151 For physics computation, we sample  $K$  points on the curve using B-spline evaluation matrix  $B \in \mathbb{R}^{K \times n_{\text{ctrl}}}$ :  
1152

1153 
$$\mathbf{X}_j = \sum_{i=1}^{n_{\text{ctrl}}} B_{ji} \mathbf{P}_i, \quad j = 1, \dots, K \quad (36)$$
  
1154

1155 **Algorithm 1** Hyperelastic Ring Deformation Policy

---

1156 1: **Input:** State  $(s, \dot{s}, \mathbf{o}, \dot{\mathbf{o}}, \theta, \omega)$ , obstacles  $\{\mathbf{c}_k, r_k\}$ , target  $\mathbf{x}_{\text{target}}$   
 1157 2: **Output:** Updated state  $(s', \dot{s}', \mathbf{o}', \dot{\mathbf{o}}', \theta', \omega')$   
 1158 3: Update geometry:  $\mathbf{X}_j \leftarrow$  sample curve at current state  
 1159 4: Compute distances:  $d_{jk} \leftarrow \|\mathbf{X}_j - \mathbf{c}_k\| - r_k$ , clearance:  $d_{\min} \leftarrow \min_{j,k} d_{jk}$   
 1160 5: ▷ — Conservative Forces —  
 1161 6: IPC barriers:  $\mathbf{g}_j \leftarrow \sum_k \frac{\partial b_{\text{IPC}}(d_{jk})}{\partial \mathbf{X}_j}$   
 1162 7: Adaptive bulk:  $F_{s,\text{bulk}} \leftarrow -\frac{\partial \mathcal{U}_{\text{bulk}}}{\partial s}$  with  $A_{\text{target}}(d_{\min})$   
 1163 8: ▷ — Non-Conservative Forces —  
 1164 9: Stage forces:  $\mathbf{F}_{\text{stage},j} \leftarrow$  goal + radial + tangential components  
 1165 10: Friction:  $\mathbf{F}_{\text{friction},j} \leftarrow -\mu \text{contact\_pressure} \cdot \text{tangent\_velocity}$   
 1166 11: ▷ — Generalized Force Mapping —  
 1167 12: Map to coordinates:  $F_s, \mathbf{F}_o, \tau \leftarrow$  virtual work from  $\{\mathbf{g}_j + \mathbf{F}_{\text{stage},j} + \mathbf{F}_{\text{friction},j}\}$   
 1168 13: ▷ — Integration —  
 1169 14: Update velocities:  $\dot{s}' \leftarrow \dot{s} + \Delta t \cdot F_s / M_s$ , etc.  
 1170 15: Update positions:  $s' \leftarrow \text{clamp}(s + \Delta t \cdot \dot{s}'), \mathbf{o}' \leftarrow \mathbf{o} + \Delta t \cdot \dot{\mathbf{o}}'$ , etc.  
 1171 16: **return** updated state

---

1172 C.2 ENERGY FORMULATION  
11731174 The total Hamiltonian combines kinetic and potential components:  
1175

1176 
$$\mathcal{H} = \frac{1}{2} M_s \dot{s}^2 + \frac{1}{2} M_o \|\dot{\mathbf{o}}\|^2 + \frac{1}{2} I \omega^2 + \mathcal{U}_{\text{barrier}} + \mathcal{U}_{\text{bulk}} \quad (37)$$
  
1177

1178 **IPC Barrier Energy:** Collision avoidance using Incremental Potential Contact barriers:  
1179

1180 
$$\mathcal{U}_{\text{barrier}} = \sum_{j=1}^K w_j \ell_j \sum_{k=1}^{N_{\text{obs}}} b_{\text{IPC}}(d_{jk}) \quad (38)$$
  
1181

1182 where  $w_j = 1/K$ ,  $\ell_j = \|\mathbf{X}'_j\|$ ,  $d_{jk}$  is distance from sample  $j$  to obstacle  $k$ :  
1183

1184 
$$b_{\text{IPC}}(d) = \begin{cases} -(d - \hat{d})^2 (\log d - \log \hat{d}) & \text{if } 0 < d < \hat{d} \\ 0 & \text{if } d \geq \hat{d} \\ V_{\text{penalty}} & \text{if } d \leq 0 \end{cases} \quad (39)$$
  
1185  
1186  
1187

1188 **Adaptive Bulk Energy:** Area conservation with clearance-dependent target:  
 1189

$$1190 \quad \mathcal{U}_{\text{bulk}} = \frac{k_{\text{bulk}}}{2} (A(s) - A_{\text{target}})^2 \quad (40)$$

1192 where  $A(s) = s^2 A_{\text{ref}}$  and:  
 1193

$$1194 \quad A_{\text{target}} = [\alpha + (1 - \alpha) \tanh(\beta \cdot \max(d_{\text{min}}, 0))] A_{\text{ref}} \quad (41)$$

1195 with  $\alpha = 0.25$ ,  $\beta = 2.5$  encouraging compression in tight spaces.  
 1196

### 1197 C.3 GENERALIZED FORCE MAPPING

1199 Forces on spline samples map to generalized coordinates via virtual work:  
 1200

$$1201 \quad F_s = -\frac{\partial \mathcal{U}}{\partial s} + \sum_{j=1}^K \mathbf{F}_{\text{ext}}(\mathbf{X}_j) \cdot \frac{\partial \mathbf{X}_j}{\partial s} - \gamma_s \dot{s} \quad (42)$$

$$1204 \quad \mathbf{F}_o = -\frac{\partial \mathcal{U}}{\partial \mathbf{o}} + \sum_{j=1}^K w_j \mathbf{F}_{\text{ext}}(\mathbf{X}_j) - \gamma_o \dot{\mathbf{o}} \quad (43)$$

$$1207 \quad \tau = -\frac{\partial \mathcal{U}}{\partial \theta} + \sum_{j=1}^K w_j \mathbf{F}_{\text{ext}}(\mathbf{X}_j) \cdot (\mathbf{J}(\mathbf{X}_j - \mathbf{o})) - \gamma_{\theta} \omega \quad (44)$$

1210 where  $\mathbf{J} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$  generates rotation and  $\frac{\partial \mathbf{X}_j}{\partial s} = \mathbf{R}(\theta) \mathbf{P}_{0,j}$ .  
 1211

### 1213 C.4 NAVIGATOR AS META-HAMILTONIAN LEARNER

1215 The GRL-SNAM operates as a meta-learning system that coordinates multi-scale policies by learning  
 1216 how to update their Hamiltonian energy functions rather than directly manipulating phase space  
 1217 states.  
 1218

1219 **Navigation specialization.** Throughout the paper we further simplify our navigator’s Hamiltonian  
 1220 surrogate by assuming fixed  $\omega$  throughout all environment when we primarily focus on learning to  
 1221 optimize the environment-indexed linear cone generated by task energies. The potential energy is  
 1222 denoted as follows and is shown in Figure 5:  
 1223

$$1224 \quad \mathcal{R}(q; \eta(\mathcal{E})) = \underbrace{\|\mathbf{y}\|_S^2}_{\text{Sensor Cost } (E_{\text{sensor}})} + \underbrace{\beta \|\mathbf{c} - \mathbf{x}_g\|_2^2}_{\text{Goal Attraction } (E_{\text{goal}})} \\ 1225 \\ 1226 \quad + \underbrace{\lambda E_{\text{obj}}(q(t))}_{\text{Deformation Energy } (E_{\text{obj}})} + \underbrace{\sum_{i \in \mathcal{C}_t(\mathcal{E}, q)} \alpha_i b(d_i, \hat{d})}_{\text{Collision Barriers } (\text{FPE})}. \quad (45)$$

1230 The *Sensor Policy* contributes the sensor cost  $E_{\text{sensor}}$ , regularizing information acquisition. The  
 1231 *FPE (Free Path Extractor)* governs goal attraction and path-planning in current environment  $\mathcal{E}$  via  
 1232 goal attraction force  $E_{\text{goal}}$  and barrier potentials, balancing reachability with safety. The *Reconfig*  
 1233 *Policy* governs deformation, enabling radius modulation for narrow passages. This decomposition  
 1234 highlights that the total Hamiltonian is not a monolithic reward but a structured sum of physically  
 1235 interpretable energies, each attached to a specialized policy. In addition, we assume friction and  
 1236 additional forces (e.g. derived from safety constraints) are happened in FPE submodule only, and  
 1237 we pick a particular parametrization as:  
 1238

$$1239 \quad \Gamma_y^\xi \equiv 0, \quad G_y^\xi \equiv 0, \quad u_y^\xi \equiv 0, \\ 1240 \quad \Gamma_f^\xi = \mu_\xi(\mathcal{E}) \mathbf{I}, \quad G_f^\xi = \mathbf{I}, \quad u_f^\xi \neq 0, \quad (46) \\ 1241 \quad \Gamma_o^\xi \equiv 0, \quad G_o^\xi \equiv 0, \quad u_o^\xi \equiv 0,$$

1242 Thus the meta policy  $g_\xi$  *directly* produces the cone coordinates of potentials and non-conservative  
 1243 forces correction in FPE module as a friction and a port correction term:  
 1244

$$1245 \quad g_\xi(\mathcal{E}) = [\eta_\xi(\mathcal{E}), \mu_\xi(\mathcal{E}), u_f^\xi],$$

1246 which defines the potential—and therefore the generalized force—for the stagewise motion-  
 1247 planning Hamiltonian rollouts. Learning  $g_\xi$  gives a *meta-policy* that maps environments to energy  
 1248 weights, i.e. a stagewise bilevel scheme where the inner layer optimizes motion under  $H_k$  and the  
 1249 outer layer trains  $\xi$  so that scenario-level QoIs/constraints are satisfied across environments.  
 1250

1251 **Learning  $\eta_\xi(\mathcal{E}), \mu_\xi(\mathcal{E})$  throughout “Module→surrogate reduction”.** Each submodule exposes  
 1252 a response map that provides integrated dynamics rollout and additional quantity of interests (QoI)  
 1253 as the feedback:  
 1254

$$1255 \quad R_k : (H_k; \theta_k, \mathcal{E}, \xi) \mapsto \{z_t^{(k)}, s_t^{(k)}, \text{QoI}_k\}_{t=0}^{T_k}, \quad k \in \mathcal{K},$$

1256 where  $T_k$  refers time scale for each submodule. To be more specific:

$$1257 \quad \text{QoI}_y = \Delta\mathcal{E}, \quad \text{QoI}_f = \{v_t^{(f)} := M_f^{-1} p_t^{(f)}\}_{t=0}^{T_f}, \quad \text{QoI}_o = \{\min_{i \in \mathcal{C}_t(\mathcal{E}, q)} d_i(q_t^{(o)}; \mathcal{E})\}_{t=0}^{T_o}. \quad (47)$$

1260 Namely, feedback QoIs are: environmental update, velocity observation, and min distance clearance  
 1261 (and thus collision violation). Additional QoIs, that can be deduced from  $z_t^{(k)}, s_t^{(k)}$ , are not stated  
 1262 here explicitly. To train a policy that output  $\eta(\mathcal{E})$  and  $\mu_\xi(\mathcal{E})$  with different environments, we propose  
 1263 to minimize

$$1264 \quad \mathcal{L}(\eta_\xi, \mu_\xi) = \mathbb{E}_{\mathcal{E}} \left[ w_q \|\mathbf{q} - \mathbf{q}_{ref}\|_2^2 + w_v \|\mathbf{v} - \mathbf{v}_{ref}\|_2^2 + w_\mu \|\mu - \mu_{ref}\|_2^2 + w_d \mathcal{L}_{\text{multi}} \right]. \quad (48)$$

1266 where  $\mathcal{L}_{\text{multi}}$  is a short multi-start robustness penalty that re-rolls from perturbed  $q_t^{(k)}$  seeds near  
 1267 obstacles to discourage brittle  $\eta_\xi(\mathcal{E})$  and  $\mu_\xi(\mathcal{E})$  (details are addressed in Algorithm 2), and  $w_\bullet \geq 0$   
 1268 are user-input hyperparameters. The training via equation 48 can be conducted offline, component-  
 1269 wise, or even fine-tuned online, but we state that it is important to fully utilize the instantaneous  
 1270 response from a real navigation scenario which provides the scheme of per-scenario online cor-  
 1271 rection even when  $\eta_\xi(\mathcal{E})$  and  $\mu_\xi(\mathcal{E})$  are properly trained under large-scale simulated dataset with  
 1272 reference potentials.  
 1273

1274 **Online Adaptation of  $g_\xi(\mathcal{E})$  via QoIs** We state how response map for each environment can yield  
 1275 a correction term under online navigation scenario. Given response  $R_k$  at time  $t$  we construct an  
 1276 *observable* measurement vector and its reference goal:

$$1277 \quad y_t = \begin{bmatrix} -\text{clr}_t \\ \text{dist}_t \\ -\text{speed}_t \end{bmatrix} \in \mathbb{R}^3, \quad y_t^* = \begin{bmatrix} -m_{\text{safe}} \\ \text{dist}_t - \varepsilon_{\text{prog}} \\ -\max(\text{speed}_t, \mathbf{1}_{\{\text{clr}_t \geq m_{\text{safe}}\}} v_{\text{min}}) \end{bmatrix},$$

1281 where  $\text{clr}$  is the minimum clearance to inflated obstacles,  $\text{dist}$  is goal distance (to the global goal  
 1282 point), and  $\text{speed} = \|\mathbf{v}\|$ . We update only the *active* barrier weights by selecting an index set  $\mathcal{I}_t$  that  
 1283 represents nearby obstacles. Define the parameter vector

$$1284 \quad \tilde{\eta}_t = \begin{bmatrix} \beta_t \\ \lambda_t \\ \alpha_{t, \mathcal{I}_t} \end{bmatrix} \in \mathbb{R}_+^{2+|\mathcal{I}_t|}, \quad \zeta_t = [\tilde{\eta}_t, \mu_t].$$

1288 We denote an estimator of Jacobian  $J_t = \frac{\partial y}{\partial \zeta_t}$  as

$$1290 \quad \tilde{J}_t = \frac{(y_t - y_{t-1})(\zeta_t - \zeta_{t-1})^\top}{\|\zeta_t - \zeta_{t-1}\|_2^2 + \varepsilon}, \quad J_t = \rho J_{t-1} + (1 - \rho) \tilde{J}_t,$$

1293 with smoothing  $\rho \in [0, 1)$  and  $\varepsilon > 0$ . Then, given the desired observable change  $\Delta y_t^{\text{des}} := y_t^* - y_t$ ,  
 1294 one can update meta-policy parameter via a Tiknhoov-regularized least square steps:

$$1295 \quad \Delta \zeta_t = \arg \min_{\Delta \zeta} \|J_t \Delta \zeta - \Delta y_t^{\text{des}}\|_2^2 + \lambda \|\Delta \zeta\|_2^2, \quad \Rightarrow \quad \Delta \zeta_t = (J_t^\top J_t + \lambda I)^{-1} J_t^\top \Delta y_t^{\text{des}}.$$

---

1296   **Algorithm 2** Meta-policy training of  $\eta(\mathcal{E})$  and  $\mu_\xi(\mathcal{E})$

1297   **Require:** dataset of different environments  $\{\mathcal{E}, \alpha_{ref}, \beta_{ref}, \lambda_{ref}, \mu_{ref}\} \in \mathcal{D}$ , step range  $T \in$   
1298     $\{2, 3, 4, 5, 6\}$ , short rollout trials  $M$ , time step  $\Delta t$ , weights  $w_\bullet$

1299

1300   1: Initialize  $\xi$

1301   2: **for** epoch = 1, 2, ... **do**

1302    3:   **for** batch  $\mathcal{B} \subset \mathcal{D}$  **do**

1303    4:     Sample  $(q_0, p_0)$  per scene  $\mathcal{E}$  in  $\mathcal{B}$

1304    5:     Build **tokens** from active constraints  $\mathcal{C}_t(\mathcal{E}, q_0)$ , goal  $x_g$ , and current  $(q_0, p_0)$

1305    6:      $(\{\alpha_j\}, \beta) \leftarrow \eta_\xi(\mathbf{tokens})$ ,  $\mu \leftarrow \mu_\xi(\mathbf{tokens})$

1306    7:     **for**  $m = 1, 2, \dots, M$  **do**

1307    8:       Sample a near obstacle point  $\tilde{q}_0$ , initialize  $\tilde{p}_0$  towards nearest obstacle at  $\tilde{q}_0$ .

1308    9:       Integrate equation 8 and equation 9 for  $T$  steps from  $(\tilde{q}_0, \tilde{p}_0)$ .  
1309    10:      ▷ e.g. Symplectic Euler

1310    11:      Compute  $clr \leftarrow \min_{i,t} d_i(\tilde{q}_t; \mathcal{E})$

1311    12:      Compute  $\mathcal{L}_{multi} \leftarrow \mathcal{L}_{multi} + \frac{1}{M} \text{softplus}\left(\frac{r_{min} - clr}{\hat{d}}\right)$ .  
1312    13:      ▷  $r_{min}$  refers minimal radius of robot object.

1313    14:      **end for**

1314    15:      Integrate equation 8 and equation 9 for  $T_k$  steps from  $(q_0, p_0)$

1315    16:      Evaluate loss  $\mathcal{L}(\eta_\xi, \mu_\xi)$  in equation 48

1316    17:      Update  $\eta_\xi, \mu_\xi$  using  $\nabla_{\eta_\xi} \mathcal{L}$  and  $\nabla_{\mu_\xi} \mathcal{L}$

1317    18:      E.g. Adam Optimizer

1318    19:      **end for**

20:   **end for**

This yields updated  $\beta_{t+1}$ ,  $\gamma_{t+1}$ , and  $\alpha_{t+1,i}$  for  $i \in \mathcal{I}_t$  (inactive weights keep their previous values).

$$\zeta_{t+1} = \Pi_{\mathbb{R}_+}((1-\kappa)\zeta_t + \kappa \Delta \zeta_t), \quad \kappa_i \in [0, 1).$$

Since the update is a result of least square step, there exists residual:

$$r_t = \Delta y_t^{\text{des}} - J_t \Delta \zeta_t.$$

The online port correction term can amend the energy change by solving another least square problem given port-observable sensitivity  $P_t \approx \partial y / \partial u_f$ :

$$u_{f,t}^\xi = \arg \min_{u \in \mathcal{U}} \|P_t u - r_t\|_2^2 + \lambda_u \|u\|_2^2 = (P_t^\top P_t + \lambda_u I)^{-1} P_t^\top r_t,$$

followed by componentwise clipping to a feasible box  $\mathcal{U}$ . A simple choice is to use  $P_t = \text{diag}(0, 0, \kappa_v)$  so that the port primarily regulates speed while the energy weights steer clearance and goal progress; richer  $P_t$  can be learned online by the same secant recipe as  $J_t$ .

## D DOMAIN-SPECIFIC POLICY IMPLEMENTATIONS

## D.1 SENSOR POLICY ( $\pi_y$ ) DETAILS

The sensor policy maintains spatial index  $\mathcal{T}_y$  of observations  $(\mathbf{x}_i, \text{type}_i, \text{attr}_i)$  and derives three energy components from single neighbor queries:

## Barrier Potential: Repulsion from obstacles

$$b_{\Sigma}(z_y, \mathcal{C}_t) = \sum_{i \in \mathcal{N}_{\text{obs}}} w_i \exp \left( -\frac{\|\mathbf{x}_i - \mathbf{c}_y\|^2}{2\sigma_b^2} \right) \quad (49)$$

### Free-Space Potential: Attraction to open regions

$$\mathcal{V}_{\text{free}}(z_y, \mathcal{C}_t) = - \sum_{j \in \mathcal{N}_{\text{enc}}} w_j \exp \left( - \frac{\|\mathbf{x}_j - \mathbf{c}_y\|^2}{2\sigma_f^2} \right) \quad (50)$$

1350 **Density Potential:** Information-theoretic density measure  
 1351

$$1352 \quad 1353 \quad 1354 \quad \rho(z_y, \mathcal{C}_t) = - \sum_{k \in \mathcal{N}_{\text{all}}} w_k \log \left( 1 + \frac{n_k}{|\mathcal{N}_{\text{all}}|} \right) \quad (51)$$

1355 The complete sensor score function is:  
 1356

$$1357 \quad 1358 \quad 1359 \quad s_y^{\theta_y}(z_y, \mathcal{C}_t, t) = \nabla_{z_y} \left[ \frac{1}{2} \|p_y\|_{M_y^{-1}}^2 + \alpha_b b_{\Sigma} + \alpha_f \mathcal{V}_{\text{free}} + \alpha_d \rho \right] \quad (52)$$

1360 **D.2 FRAME POLICY ( $\pi_f$ ) DETAILS**  
 1361

1362 The frame policy uses  $\mathcal{T}_f$  storing path samples with safety/contact distances and goal influence:  
 1363

1364 **Safety Field:** Distance-based safety measure  
 1365

$$1366 \quad 1367 \quad S(z_f, \mathcal{C}_t) = \sum_{i=1}^{N_s} w_i \max(0, d_{\text{safe}}^{\text{threshold}} - d_{\text{safe}}^i)^2 \quad (53)$$

1368 **Contact Field:** Proximity to obstacles  
 1369

$$1370 \quad 1371 \quad 1372 \quad C(z_f, \mathcal{C}_t) = \sum_{i=1}^{N_c} w_i \exp \left( -\frac{(d_{\text{contact}}^i)^2}{2\sigma_c^2} \right) \quad (54)$$

1373 **Goal Field:** Directional bias toward target  
 1374

$$1375 \quad 1376 \quad 1377 \quad G(z_f, \mathcal{C}_t) = -\|c_f - \mathbf{x}_g\|^2 + \sum_{i=1}^{N_g} w_i g_i \cos(\theta_i) \quad (55)$$

1378 The frame score function integrates these fields:  
 1379

$$1380 \quad 1381 \quad 1382 \quad s_f^{\theta_f}(z_f, \mathcal{C}_t, t) = \nabla_{z_f} \left[ \frac{1}{2} \|p_f\|_{M_f^{-1}}^2 + \alpha_s S + \alpha_c C + \alpha_g G \right] \quad (56)$$

1383 **D.3 SHAPE POLICY ( $\pi_o$ ) DETAILS**  
 1384

1385 The shape policy controls deformation through reduced coordinates  $z_o = (s, \dot{s}, \mathbf{o}, \dot{\mathbf{o}}, \theta, \omega)$ :  
 1386

1387 **Smoothness Energy:** Curvature regularization  
 1388

$$1389 \quad 1390 \quad 1391 \quad \mathcal{E}_{\text{smooth}} = \int_0^1 \|\kappa(u)\|^2 du \approx \sum_{j=1}^K w_j \|\kappa_j\|^2 \quad (57)$$

1392 **Stretching Energy:** Arc length preservation  
 1393

$$1394 \quad 1395 \quad 1396 \quad \mathcal{E}_{\text{stretch}} = \int_0^1 (\|\mathbf{S}'(u)\| - \ell_{\text{ref}})^2 du \approx \sum_{j=1}^K w_j (\ell_j - \ell_{\text{ref}})^2 \quad (58)$$

1397 **Target Energy:** Configuration constraints  
 1398

$$1399 \quad \mathcal{E}_{\text{target}} = \|\mathbf{P} - \mathbf{P}_{\text{target}}\|_F^2 + \|(s, \mathbf{o}, \theta) - (s_{\text{target}}, \mathbf{o}_{\text{target}}, \theta_{\text{target}})\|^2 \quad (59)$$

1400 The complete shape score function is:  
 1401

$$1402 \quad 1403 \quad s_o^{\theta_o}(z_o, \mathcal{C}_t, t) = \nabla_{z_o} \left[ \frac{1}{2} M_s \dot{s}^2 + \frac{1}{2} M_o \|\dot{\mathbf{o}}\|^2 + \frac{1}{2} I \omega^2 + \alpha_{sm} \mathcal{E}_{\text{smooth}} + \alpha_{st} \mathcal{E}_{\text{stretch}} + \alpha_{tg} \mathcal{E}_{\text{target}} \right] \quad (60)$$

---

1404 **E ALGORITHM IMPLEMENTATION DETAILS**
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414

This section provides comprehensive implementation details for Algorithm 3 presented in §3.1. We explain each algorithmic component, including parameter initialization, policy query protocols, Hamiltonian composition, integration schemes, observable extraction procedures, and online adaptation mechanisms. These details are essential for reproducible implementation but are deferred to the appendix to maintain focus on the core methodological contributions in the main text.

---

1415 **Algorithm 3** Online GRL-SNAM: Navigator-Driven Hamiltonian Composition
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

1: **Input:** Goal  $\mathbf{x}_g$ , initial  $z_0 = (q_0, 0)$ , step  $\tau$ , horizons  $(T_y, T_f, T_o)$   $\triangleright \S E.1$   
2: **Init:**  $t \leftarrow 0$ ,  $\mathcal{C}_0 \leftarrow \emptyset$ , meta-policy  $g_\xi$ , Jacobians  $J_0, P_0$   $\triangleright \S E.2$   
3: **while**  $\neg \text{REACHEDGOAL}(c_t, \mathbf{x}_g)$  **and**  $t < T_{\max}$  **do**  $\triangleright \S E.3$   
4:     **Query policies & collect responses:**  $\triangleright \S E.4$   
5:     **if**  $t \equiv 0 \pmod{T_y}$  **then**  
6:          $\mathcal{R}_y^t \leftarrow \pi_y(H_y; \theta_y, \mathcal{E}, \xi)$  returns  $\{z_{0:T_y}^{(y)}, s_{0:T_y}^{(y)}, \Delta\mathcal{E}\}$   $\triangleright \S E.5$   
7:          $\mathcal{E} \leftarrow \mathcal{E} \cup \Delta\mathcal{E}$ ;  $\mathcal{C}_t \leftarrow \{i \mid d_i(q_t; \mathcal{E}) \leq \hat{d}\}$   
8:     **end if**  
9:     **if**  $t \equiv 0 \pmod{T_f}$  **then**  
10:          $\mathcal{R}_f^t \leftarrow \pi_f(H_f; \theta_f, \mathcal{E}, \xi)$  returns  $\{z_{0:T_f}^{(f)}, s_{0:T_f}^{(f)}, \{v_\ell^{(f)}\}_{\ell=0}^{T_f}\}$   $\triangleright \S E.6$   
11:         **end if**  
12:          $\mathcal{R}_o^t \leftarrow \pi_o(H_o; \theta_o, \mathcal{E}, \xi)$  returns  $\{z_{0:T_o}^{(o)}, s_{0:T_o}^{(o)}, \{\min_i d_i(q_\ell^{(o)}; \mathcal{E})\}_{\ell=0}^{T_o}\}$   $\triangleright \S E.7$   
13:         **Meta-policy proposal:**  $\triangleright \S E.8$   
14:         Build tokens  $\mathcal{T}_t$  from  $(\mathcal{C}_t, q_t, \mathbf{x}_g)$   
15:          $[\eta_\xi^t, \mu_\xi^t, u_f^{\xi, t}] \leftarrow g_\xi(\mathcal{T}_t; \xi)$  with  $\eta_\xi^t = (\beta^t, \lambda^t, \{\alpha_i^t\}_{i \in \mathcal{C}_t})$   
16:         **Compose surrogate Hamiltonian:**  $\triangleright \S E.9$   
 $\mathcal{R}(q_t; \eta_\xi^t, \mathcal{E}) = E_{\text{sensor}} + \beta^t E_{\text{goal}} + \lambda^t E_{\text{obj}} + \sum_{i \in \mathcal{C}_t} \alpha_i^t b(d_i)$   
 $H(q_t, p_t; \omega, \xi, \mathcal{E}) = \frac{1}{2} p_t^\top M(q_t)^{-1} p_t + \mathcal{R}(q_t; \eta_\xi^t, \mathcal{E})$   
17:         **Integrate dynamics:**  $\triangleright \S E.10$   
18:          $\nabla_p H|_t = M(q_t)^{-1} p_t; \nabla_q H|_t = \nabla_q E_{\text{sensor}} + \beta^t \nabla_q E_{\text{goal}} + \lambda^t \nabla_q E_{\text{obj}} + \sum_{i \in \mathcal{C}_t} \alpha_i^t \nabla_q b(d_i)$   
19:          $p_{t+\tau} = p_t - \tau \nabla_q H|_t - \tau \mu_\xi^t \nabla_p H|_t + \tau u_f^{\xi, t}$   
20:          $q_{t+\tau} = q_t + \tau M(q_t)^{-1} p_{t+\tau}$   
21:         **Extract observables:**  $\triangleright \S E.11$   
22:          $\text{clr}_t = \min\{\min_i d_i(q_{t+\tau}; \mathcal{E}), \min_{\ell, i} d_i(q_\ell^{(o)}; \mathcal{E})\}; \text{dist}_t = \|c_{t+\tau} - \mathbf{x}_g\|; \text{speed}_t = \|M^{-1} p_{t+\tau}\|$   
23:          $y_t = [-\text{clr}_t, \text{dist}_t, -\text{speed}_t]^\top; y_t^* = [-m_{\text{safe}}, \text{dist}_t - \varepsilon_{\text{prog}}, -\max(\text{speed}_t, \mathbf{1}_{\{\text{clr}_t \geq m_{\text{safe}}\}} v_{\min})]^\top$   
24:         **Parameter adaptation:**  $\triangleright \S E.12$   
25:         Select  $\mathcal{I}_t \subset \mathcal{C}_t$ ;  $\zeta_t = [\beta^t, \lambda^t, \{\alpha_i^t\}_{i \in \mathcal{I}_t}, \mu_\xi^t]^\top$   
26:          $\tilde{J}_t = (y_t - y_{t-\tau})(\zeta_t - \zeta_{t-\tau})^\top / (\|\zeta_t - \zeta_{t-\tau}\|^2 + \varepsilon); J_t = \rho J_{t-\tau} + (1 - \rho) \tilde{J}_t$   
27:          $\Delta y_t^{\text{des}} = y_t^* - y_t; \Delta \zeta_t = (J_t^\top J_t + \lambda_\zeta I)^{-1} J_t^\top \Delta y_t^{\text{des}}$   
28:          $\zeta_{t+\tau} = \Pi_{\mathbb{R}_+}((1 - \boldsymbol{\kappa}) \odot \zeta_t + \boldsymbol{\kappa} \odot (\zeta_t + \Delta \zeta_t))$   
29:         Unpack to  $[\beta^{t+\tau}, \lambda^{t+\tau}, \{\alpha_i^{t+\tau}\}_{i \in \mathcal{I}_t}, \mu_\xi^{t+\tau}]$   
30:         **Port correction:**  $\triangleright \S E.13$   
31:          $r_t = \Delta y_t^{\text{des}} - J_t \Delta \zeta_t; u_f^{\xi, t+\tau} = \text{clip}((P_t^\top P_t + \lambda_u I)^{-1} P_t^\top r_t, \mathcal{U})$   
32:          $\tilde{P}_t = (y_t - y_{t-\tau})(u_f^{\xi, t} - u_f^{\xi, t-\tau})^\top / (\|u_f^{\xi, t} - u_f^{\xi, t-\tau}\|^2 + \varepsilon); P_t = \rho_P P_{t-\tau} + (1 - \rho_P) \tilde{P}_t$   
33:          $t \leftarrow t + \tau$   
34:     **end while**  
35:     **Return:** Trajectory  $\{z_\ell\}_{\ell=0}^T$ , parameters  $\{\eta_\xi^\ell, \mu_\xi^\ell, u_f^{\xi, \ell}\}_{\ell=0}^T$   $\triangleright \S E.14$

---

1458  
1459

## E.1 INITIALIZATION

1460  
1461  
1462

**Robot State.** Initialize phase-space coordinates  $z_0 = (q_0, p_0)$  where  $q_0 = (c_0, \theta_0, s_0, \psi_0)$  contains the starting position  $c_0$ , orientation  $\theta_0$ , scale  $s_0 = 1.0$ , and shape parameters  $\psi_0$  (e.g., B-spline control points initialized as a circle). The initial momentum is  $p_0 = 0$ .

1463  
1464  
1465  
1466  
1467

**Time Step and Horizons.** The integration time step  $\tau$  is typically set to 0.01–0.05 seconds. Query horizons  $(T_y, T_f, T_o)$  determine how frequently each policy is queried: sensor policy every  $T_y$  steps, frame policy every  $T_f$  steps, and shape policy every step (or every  $T_o$  steps). Typical values:  $T_y = 10, T_f = 5, T_o = 1$ .

1468  
1469  
1470

**Maximum Time Horizon.**  $T_{\max}$  is the maximum number of integration steps allowed before timeout, typically 5000–10000 steps corresponding to 50–500 seconds of simulated time.

1471  
1472

## E.2 PARAMETER STRUCTURE AND INITIALIZATION

1473  
1474

**Meta-Policy Network.** The meta-policy  $g_\xi : \mathcal{E} \rightarrow [\eta_\xi, \mu_\xi, u_f^\xi]$  can be either:

1475  
1476  
1477  
1478

- A learned neural network trained offline via Algorithm 2
- Fixed constant values:  $\beta^0 = 2.0, \lambda^0 = 1.0, \alpha_i^0 = 1.5, \mu^0 = 0.1$

1479  
1480

If using a learned network,  $\xi$  are the network parameters. If using fixed values,  $g_\xi$  simply returns constants.

1481  
1482  
1483  
1484

**Jacobian Initialization.** Initialize  $J_0 \in \mathbb{R}^{3 \times (2+|C_0|+1)}$  as a small random matrix or identity-scaled matrix. Initialize port Jacobian  $P_0 \in \mathbb{R}^{3 \times \dim(u_f)}$  similarly. These will be refined online via secant updates.

1485  
1486  
1487

**Environment and Context.** Initial environment  $\mathcal{E}_0$  contains known static obstacles. Active constraint set  $\mathcal{C}_0 = \emptyset$  starts empty and will be populated by the sensor policy’s first query.

1488  
1489

## E.3 TERMINATION CONDITIONS

1490  
1491

The algorithm terminates when any of the following conditions is met:

1492  
1493

**Success: Goal Reached.**  $\|c_t - \mathbf{x}_g\| < \epsilon_{\text{goal}}$  where  $\epsilon_{\text{goal}} = 0.05$  m (5 cm tolerance).

1494  
1495

**Failure: Timeout.**  $t \geq T_{\max}$  without reaching the goal.

1496  
1497

**Failure: Collision.**  $\min_i d_i(q_t; \mathcal{E}) < 0$ , indicating penetration into an obstacle.

1498  
1499  
1500

**Failure: Stuck.**  $\|c_t - c_{t-50\tau}\| < \epsilon_{\text{stuck}}$  where  $\epsilon_{\text{stuck}} = 0.01$  m, indicating the robot has not moved more than 1 cm in the last 50 steps (indicating entrapment in a local minimum).

1501  
1502  
1503

## E.4 HIERARCHICAL QUERY FLOW

1504  
1505

Policies are queried in a specific sequence with information flowing forward:

**Sensor Policy Query (Line 5).** Executed every  $T_y$  steps. Takes as input the previous context  $\mathcal{C}_{t-T_y}$ , current state  $z_y^t = (c_t, \dot{c}_t)$ , and goal  $\mathbf{x}_g$ . Performs ray-casting from the robot’s position to detect new obstacles within sensing range.

1509  
1510  
1511

**Frame Policy Query (Line 8).** Executed every  $T_f$  steps. Takes as input the updated context  $\mathcal{C}_t$  from the sensor policy, current full state  $z_f^t = (q_t, p_t)$ , and goal  $\mathbf{x}_g$ . Plans a short-horizon path considering known obstacles.

1512    **Shape Policy Query (Line 10).** Executed every step. Takes as input the context  $\mathcal{C}_t$  from sensor  
 1513    and the current shape state  $z_o^t = (s_t, \psi_t, \dot{s}_t, \dot{\psi}_t)$ . Computes deformation forces to navigate through  
 1514    constrained spaces.  
 1515

1516    **Information Flow.** 

1518    **E.5 SENSOR RESPONSE**

1520    **Response Structure.**  $\mathcal{R}_y^t = \{z_{0:T_y}^{(y)}, s_{0:T_y}^{(y)}, \Delta\mathcal{E}\}$  where:

1522    •  $z_{0:T_y}^{(y)} = \{(q_\ell^{(y)}, p_\ell^{(y)})\}_{\ell=0}^{T_y}$  is the integrated trajectory of the sensor policy's local Hamiltonian  $H_y$   
 1523    •  $s_{0:T_y}^{(y)} = \{s_\ell^{(y)}\}_{\ell=0}^{T_y}$  is the score function (dynamics drift) at each step  
 1524    •  $\Delta\mathcal{E}$  contains newly detected obstacles:  $\Delta\mathcal{E} = \{(\mathbf{x}_j, r_j, \text{type}_j)\}$  with positions, radii, and  
 1525    types  
 1526    •  $\Delta\mathcal{E}$  contains newly detected obstacles:  $\Delta\mathcal{E} = \{(\mathbf{x}_j, r_j, \text{type}_j)\}$  with positions, radii, and  
 1527    types  
 1528    •  $\Delta\mathcal{E}$  contains newly detected obstacles:  $\Delta\mathcal{E} = \{(\mathbf{x}_j, r_j, \text{type}_j)\}$  with positions, radii, and  
 1529    types

1530    **Environment Update (Line 6).** Merge new detections:  $\mathcal{E} \leftarrow \mathcal{E} \cup \Delta\mathcal{E}$ . Update active constraint  
 1531    set:  $\mathcal{C}_t \leftarrow \{i \mid d_i(q_t; \mathcal{E}) \leq \hat{d}\}$  where  $\hat{d}$  is the sensing/activation radius (typically 0.5–1.0 m).  
 1532

1533    **Policy Hamiltonian (not directly observed).** The sensor policy operates on  $H_y(q_y, p_y; \xi, \mathcal{E}) =$   
 1534     $\frac{1}{2}p_y^\top M_y(q_y)^{-1}p_y + \mathcal{R}_y(q_y; \xi, \mathcal{E})$  where  $\mathcal{R}_y = E_{\text{sensor}}(q; \mathcal{E}, \omega_y) + \sum_{i \in \mathcal{C}_t} \alpha_i b(d_i)$  as defined in the  
 1535    methodology.

1536    **E.6 FRAME RESPONSE**

1539    **Response Structure.**  $\mathcal{R}_f^t = \{z_{0:T_f}^{(f)}, s_{0:T_f}^{(f)}, \{v_\ell^{(f)}\}_{\ell=0}^{T_f}\}$  where:

1541    •  $z_{0:T_f}^{(f)} = \{(q_\ell^{(f)}, p_\ell^{(f)})\}_{\ell=0}^{T_f}$  is the short-horizon rollout under  $H_f$   
 1542    •  $s_{0:T_f}^{(f)}$  are the score functions  
 1543    •  $\{v_\ell^{(f)} = M_f^{-1}p_\ell^{(f)}\}_{\ell=0}^{T_f}$  are velocity observations used for extracting speed in observable  
 1544    computation  
 1545    •  $\{v_\ell^{(f)} = M_f^{-1}p_\ell^{(f)}\}_{\ell=0}^{T_f}$  are velocity observations used for extracting speed in observable  
 1546    computation

1547    **Policy Hamiltonian.**  $H_f(q_f, p_f; \xi, \mathcal{E}) = \frac{1}{2}p_f^\top M_f(q_f)^{-1}p_f + \mathcal{R}_f(q_f; \xi, \mathcal{E})$  where  $\mathcal{R}_f =$   
 1548     $\beta(\mathcal{E})E_{\text{goal}} + \sum_{i \in \mathcal{C}_t} \alpha_i b(d_i)$ .  
 1549

1551    **Dissipation and Port in Frame Policy.** The frame policy integrates with dissipation  $\Gamma_f^\xi = \mu_\xi(\mathcal{E})\mathbf{I}$   
 1552    and port input  $G_f^\xi u_f^\xi$  as per the port-Hamiltonian equations (5-6) in the methodology. This is re-  
 1553    flected in the Navigator's integration step (Line 16).  
 1554

1555    **E.7 SHAPE RESPONSE**

1557    **Response Structure.**  $\mathcal{R}_o^t = \{z_{0:T_o}^{(o)}, s_{0:T_o}^{(o)}, \{\min_i d_i(q_\ell^{(o)}; \mathcal{E})\}_{\ell=0}^{T_o}\}$  where:

1559    •  $z_{0:T_o}^{(o)} = \{(q_\ell^{(o)}, p_\ell^{(o)})\}_{\ell=0}^{T_o}$  is the deformation trajectory  
 1560    •  $s_{0:T_o}^{(o)}$  are score functions  
 1561    •  $\{\min_i d_i(q_\ell^{(o)}; \mathcal{E})\}_{\ell=0}^{T_o}$  tracks minimum clearance at each step of the shape rollout

1564    **Policy Hamiltonian.**  $H_o(q_o, p_o; \xi, \mathcal{E}) = \frac{1}{2}p_o^\top M_o(q_o)^{-1}p_o + \mathcal{R}_o(q_o; \xi, \mathcal{E})$  where  $\mathcal{R}_o =$   
 1565     $\lambda(\mathcal{E})E_{\text{obj}} + \sum_{i \in \mathcal{C}_t} \alpha_i b(d_i)$ .

1566     **Usage in Clearance.** The clearance values from the shape rollout are used in Line 18 to compute  
 1567      $\text{clr}_t$ , accounting for the robot's actual deformed geometry when assessing collision risk.  
 1568

1569     **E.8 META-POLICY PROPOSAL**

1570     **Token Construction (Line 12).** Build input tokens  $\mathcal{T}_t$  from current environment state:

- Obstacle features: positions, radii, types for each  $i \in \mathcal{C}_t$
- Robot state: position  $\mathbf{c}_t$ , velocity  $M^{-1}p_t$ , current scale  $s_t$
- Goal information: relative position  $\mathbf{x}_g - \mathbf{c}_t$ , distance  $\|\mathbf{x}_g - \mathbf{c}_t\|$

1571     These are typically encoded as a permutation-invariant set representation (e.g., using attention mech-  
 1572     anisms over obstacle features).

1573     **Meta-Policy Output (Line 13).**  $g_\xi(\mathcal{T}_t; \xi) = [\eta_\xi^t, \mu_\xi^t, u_f^{\xi, t}]$  produces:

- Energy weights:  $\eta_\xi^t = (\beta^t, \lambda^t, \{\alpha_i^t\}_{i \in \mathcal{C}_t}) \in \mathbb{R}_+^{2+|\mathcal{C}_t|}$
- Friction coefficient:  $\mu_\xi^t \in \mathbb{R}_+$
- Initial port suggestion:  $u_f^{\xi, t} \in \mathbb{R}^{\text{dim}(u_f)}$  (will be refined online)

1579     **Fixed vs Learned.** If  $g_\xi$  is not trained, simply return fixed values:  $\beta^t = 2.0, \lambda^t = 1.0, \alpha_i^t =$   
 1580      $1.5 \forall i, \mu^t = 0.1, u_f^{\xi, t} = 0$ . The online adaptation (Lines 22-28) will refine these values regardless  
 1581     of initialization.

1582     **E.9 HAMILTONIAN ASSEMBLY**

1583     **Energy Component Definitions.** The potential energy is decomposed as:

$$E_{\text{sensor}}(q; \mathcal{E}, \omega_y) = \|\mathbf{y}_t\|_S^2 \quad (\text{sensor configuration cost}) \quad (61)$$

$$E_{\text{goal}}(q; \mathcal{E}, \omega_g) = \|\mathbf{c}_t - \mathbf{x}_g\|_2^2 \quad (\text{goal attraction}) \quad (62)$$

$$E_{\text{obj}}(q; \omega_d) = \text{deformation energy of shape } \psi_t \quad (63)$$

$$b(d_i(q; \mathcal{E}); \omega_b) = \text{barrier function, e.g., } \frac{\hat{d}^2}{d_i^2} \text{ or IPC log-barrier} \quad (64)$$

1607     **Total Potential (Line 15).**  $\mathcal{R}(q_t; \eta_\xi^t, \mathcal{E}) = E_{\text{sensor}} + \beta^t E_{\text{goal}} + \lambda^t E_{\text{obj}} + \sum_{i \in \mathcal{C}_t} \alpha_i^t b(d_i)$

1608     This is the **surrogate potential** that aggregates contributions from all three policy domains,  
 1609     weighted by the meta-policy outputs  $\eta_\xi^t = (\beta^t, \lambda^t, \{\alpha_i^t\})$ .

1612     **Surrogate Hamiltonian (Line 15).**  $H(q_t, p_t; \omega, \xi, \mathcal{E}) = \frac{1}{2} p_t^\top M(q_t; \omega_M)^{-1} p_t + \mathcal{R}(q_t; \eta_\xi^t, \mathcal{E})$

1613     The kinetic energy uses a (possibly state-dependent) mass matrix  $M(q_t)$ . For planar navigation, this  
 1614     is often constant:  $M = \text{diag}(m, m, I_{zz}, m_s, \dots)$  for translation, rotation, scale, and shape degrees  
 1615     of freedom.

1618     **Fixed Parameters  $\omega$ .** As stated in the methodology section, the intra-term parameters  $\omega =$   
 1619      $(\omega_y, \omega_M, \omega_g, \omega_d, \omega_b)$  are assumed **fixed** throughout all environments. Only the dual weights  $\eta_\xi$   
 1620     are adapted online.

1620 **E.10 SYMPLECTIC INTEGRATION WITH PORT-HAMILTONIAN DYNAMICS**  
 1621

1622 **Gradient Computation (Line 16).** Compute the Hamiltonian gradients:  
 1623 
$$\nabla_p H|_t = M(q_t)^{-1} p_t \quad (65)$$
  
 1624 
$$\nabla_q H|_t = \nabla_q \mathcal{R}(q_t; \eta_\xi^t, \mathcal{E}) \quad (66)$$
  
 1625 
$$= \nabla_q E_{\text{sensor}} + \beta^t \nabla_q E_{\text{goal}} + \lambda^t \nabla_q E_{\text{obj}} + \sum_{i \in \mathcal{C}_t} \alpha_i^t \nabla_q b(d_i) \quad (67)$$
  
 1626

1627 Each gradient  $\nabla_q b(d_i)$  is computed using the chain rule:  $\nabla_q b(d_i) = \frac{\partial b}{\partial d_i} \cdot \nabla_q d_i(q; \mathcal{E})$   
 1628

1629 **Symplectic Euler with Dissipation and Port (Lines 17-18).** The integration follows the port-  
 1630 Hamiltonian structure from methodology equations (5-6):  
 1631

1632 
$$p_{t+\tau} = p_t - \tau \nabla_q H|_t - \tau \underbrace{\mu_\xi^t \mathbf{I}}_{\Gamma_f^\xi} \cdot \nabla_p H|_t + \tau \underbrace{\mathbf{I}}_{\mathbf{G}_f^\xi} \cdot u_f^{\xi, t} \quad (68)$$
  
 1633

1634 
$$q_{t+\tau} = q_t + \tau M(q_t)^{-1} p_{t+\tau} \quad (69)$$
  
 1635

1636 The dissipation term  $-\mu_\xi^t \mathbf{I} \cdot \nabla_p H|_t = -\mu_\xi^t M(q_t)^{-1} p_t$  acts as velocity-proportional damping. The  
 1637 port input  $u_f^{\xi, t}$  provides non-conservative forcing to correct for residual errors.  
 1638

1639 **Frame Policy Only.** As clarified, the dissipation  $\mu_\xi^t$  and port  $u_f^{\xi, t}$  apply **only** to the frame pol-  
 1640 icy dynamics. The sensor and shape policies have  $\Gamma_y^\xi \equiv 0, \Gamma_o^\xi \equiv 0$ , and no port inputs in their  
 1641 local dynamics. However, since the Navigator integrates the **full robot state** using the composed  
 1642 Hamiltonian, the dissipation and port appear in the full-state integration.  
 1643

1644 **E.11 OBSERVABLE EXTRACTION**  
 1645

1646 **Frame-Based but Multi-Module (Line 18-19).** The observable vector  $y_t$  is constructed from  
 1647 quantities that are **primarily** derived from the frame trajectory\*\* but require information from all  
 1648 three modules:  
 1649

1650 **Clearance** (uses sensor + shape):  
 1651

1652 
$$\text{clr}_t = \min \left\{ \min_i d_i(q_{t+\tau}; \mathcal{E}), \min_{\ell=0, \dots, T_o} \min_i d_i(q_\ell^{(o)}; \mathcal{E}) \right\} \quad (70)$$
  
 1653

1654 This takes the minimum over both the integrated Navigator state  $q_{t+\tau}$  and all states from the shape  
 1655 policy rollout  $\{q_\ell^{(o)}\}$ . Computing  $d_i(q; \mathcal{E})$  requires:  
 1656

1657 

- Environment  $\mathcal{E}$  from sensor policy
- Robot geometry (radius, shape) from shape policy

  
 1658

1659 **Distance** (frame position):  
 1660

1661 
$$\text{dist}_t = \|\mathbf{c}_{t+\tau} - \mathbf{x}_g\| \quad (71)$$
  
 1662 Directly from the frame component of  $q_{t+\tau}$ .  
 1663

1664 **Speed** (frame velocity):  
 1665

1666 
$$\text{speed}_t = \|M(q_{t+\tau})^{-1} p_{t+\tau}\| \quad (72)$$
  
 1667 Computed from the updated momentum  $p_{t+\tau}$ .  
 1668

1669 **Observable Vector and Target (Line 19).**  
 1670

1671 
$$y_t = \begin{bmatrix} -\text{clr}_t \\ \text{dist}_t \\ -\text{speed}_t \end{bmatrix}, \quad y_t^* = \begin{bmatrix} -m_{\text{safe}} \\ \text{dist}_t - \varepsilon_{\text{prog}} \\ -\max(\text{speed}_t, \mathbf{1}_{\{\text{clr}_t \geq m_{\text{safe}}\}} v_{\min}) \end{bmatrix} \quad (73)$$
  
 1672

1673 Target parameters:  $m_{\text{safe}} = 0.15$  m (desired minimum clearance),  $\varepsilon_{\text{prog}} > 0$  (desired progress per  
 1674 step),  $v_{\min} = 0.3$  m/s (minimum speed when safe).

1674  
1675**Physical Interpretation.**1676  
1677  
1678  
1679  
1680  
1681  
1682  
1683

- $y_t[0] = -\text{clr}_t$ : negative clearance (larger value = closer to obstacles)
- Target  $y_t^*[0] = -0.15$  means we want to maintain at least 15 cm clearance
- $y_t[1] = \text{dist}_t$ : current distance to goal
- Target  $y_t^*[1] = \text{dist}_t - \varepsilon_{\text{prog}}$  means we want to make progress (reduce distance)
- $y_t[2] = -\text{speed}_t$ : negative speed
- Target encourages minimum speed  $v_{\min}$  when clearance is safe

1684  
1685**E.12 PARAMETER ADAPTATION VIA OBSERVABLE FEEDBACK**1686  
1687  
1688

**Active Parameter Selection (Line 21).** Select a subset  $\mathcal{I}_t \subset \mathcal{C}_t$  of nearby obstacles (e.g., within 0.5 m) to limit the parameter vector size. Assemble:

1689  
1690

$$\zeta_t = [\beta^t, \lambda^t, \{\alpha_i^t\}_{i \in \mathcal{I}_t}, \mu_\xi^t]^\top \in \mathbb{R}_+^{2+|\mathcal{I}_t|+1} \quad (74)$$

1691  
1692

For obstacles not in  $\mathcal{I}_t$ , their  $\alpha_i$  values remain unchanged from the previous step.

1693  
1694

**Secant Jacobian Update (Line 22).** Approximate the Jacobian  $J_t \approx \partial y / \partial \zeta$  using a rank-1 secant update:

1695  
1696  
1697

$$\tilde{J}_t = \frac{(y_t - y_{t-\tau})(\zeta_t - \zeta_{t-\tau})^\top}{\|\zeta_t - \zeta_{t-\tau}\|^2 + \varepsilon} \in \mathbb{R}^{3 \times (2+|\mathcal{I}_t|+1)} \quad (75)$$

1698

Apply exponential moving average for smoothing:

1699  
1700

$$J_t = \rho J_{t-\tau} + (1 - \rho) \tilde{J}_t, \quad \rho \in [0, 1] \quad (76)$$

1701  
1702

Typical value:  $\rho = 0.9$ .

1703  
1704

**Tikhonov-Regularized Least Squares (Line 23).** Compute the desired observable change:

1705  
1706

$$\Delta y_t^{\text{des}} = y_t^* - y_t \quad (77)$$

1707  
1708

Solve the regularized least-squares problem:

1709  
1710

$$\Delta \zeta_t = \arg \min_{\Delta \zeta} \|J_t \Delta \zeta - \Delta y_t^{\text{des}}\|_2^2 + \lambda_\zeta \|\Delta \zeta\|_2^2 \quad (78)$$

1711

Closed-form solution:

1712  
1713

$$\Delta \zeta_t = (J_t^\top J_t + \lambda_\zeta I)^{-1} J_t^\top \Delta y_t^{\text{des}} \quad (79)$$

1714  
1715

Regularization parameter:  $\lambda_\zeta = 10^{-4}$  ensures numerical stability.

1716  
1717

**Parameter Update with Step Sizes (Line 24).** Apply the update with componentwise step sizes  $\kappa = [\kappa_\beta, \kappa_\lambda, \{\kappa_{\alpha_i}\}, \kappa_\mu]^\top$  where each  $\kappa_j \in [0, 1]$ :

1718  
1719

$$\zeta_{t+\tau} = \Pi_{\mathbb{R}_+}((1 - \kappa) \odot \zeta_t + \kappa \odot (\zeta_t + \Delta \zeta_t)) \quad (80)$$

1720  
1721

The projection  $\Pi_{\mathbb{R}_+}$  enforces non-negativity:  $[\zeta_{t+\tau}]_j = \max(0, [\zeta_{t+\tau}]_j)$ .

1722  
1723

Typical step sizes:  $\kappa_\beta = 0.05, \kappa_\lambda = 0.05, \kappa_{\alpha_i} = 0.1, \kappa_\mu = 0.02$ .

1724  
1725

**Unpacking (Line 25).** Extract the updated parameters:

1726  
1727

$$[\beta^{t+\tau}, \lambda^{t+\tau}, \{\alpha_i^{t+\tau}\}_{i \in \mathcal{I}_t}, \mu_\xi^{t+\tau}] \leftarrow \zeta_{t+\tau} \quad (81)$$

1728

For obstacles  $j \notin \mathcal{I}_t$ , retain previous values:  $\alpha_j^{t+\tau} = \alpha_j^t$ .

1728 E.13 PORT CORRECTION FROM RESIDUAL  
17291730 **Residual Computation (Line 27).** After the Jacobian-based parameter update, there may remain  
1731 a residual error:

1732 
$$r_t = \Delta y_t^{\text{des}} - J_t \Delta \zeta_t \in \mathbb{R}^3 \quad (82)$$

1733 This residual captures observable errors that cannot be corrected by reshaping the energy weights  
1734 alone (e.g., due to model mismatch or unmodeled dynamics).1735  
1736 **Port Correction via Least Squares (Line 27).** Solve a second least-squares problem to find a port  
1737 input that addresses the residual:

1738  
1739 
$$u_f^{\xi, t+\tau} = \arg \min_{u \in \mathbb{R}^{\text{dim}(u_f)}} \|P_t u - r_t\|_2^2 + \lambda_u \|u\|_2^2 \quad (83)$$

1740  
1741 Closed-form solution:

1742  
1743 
$$u_f^{\xi, t+\tau} = (P_t^\top P_t + \lambda_u I)^{-1} P_t^\top r_t \quad (84)$$

1744 Clip to feasible box:  $u_f^{\xi, t+\tau} \leftarrow \text{clip}(u_f^{\xi, t+\tau}, \mathcal{U})$  where  $\mathcal{U} = [-u_{\text{max}}, u_{\text{max}}]^{\text{dim}(u_f)}$ .1745  
1746 **Port-Observable Jacobian (Line 28).** Estimate  $P_t \approx \partial y / \partial u_f$  using secant method:

1747  
1748 
$$\tilde{P}_t = \frac{(y_t - y_{t-\tau})(u_f^{\xi, t} - u_f^{\xi, t-\tau})^\top}{\|u_f^{\xi, t} - u_f^{\xi, t-\tau}\|^2 + \varepsilon} \quad (85)$$

1749  
1750 Apply smoothing:

1751  
1752 
$$P_t = \rho_P P_{t-\tau} + (1 - \rho_P) \tilde{P}_t, \quad \rho_P \in [0, 1] \quad (86)$$

1753  
1754 **Simplified Port Jacobian.** If full  $P_t$  estimation is noisy, use a simple diagonal form:

1755  
1756 
$$P_t = \text{diag}(0, 0, \kappa_v) \quad (87)$$

1757 where  $\kappa_v > 0$  means the port primarily affects the speed observable (third component of  $y_t$ ), while  
1758 clearance and distance are controlled via energy weights.1759  
1760 **Role of Port Correction.** The port correction  $u_f^{\xi, t+\tau}$  provides **non-conservative forcing** that can-  
1761 not be represented by a potential function. It acts as a "corrective impulse" applied at each step to  
1762 handle:1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781  
1782  
1783  
1784  
1785  
1786  
1787  
1788  
1789  
1790  
1791  
1792  
1793  
1794  
1795  
1796  
1797  
1798  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1830  
1831  
1832  
1833  
1834  
1835  
1836  
1837  
1838  
1839  
1840  
1841  
1842  
1843  
1844  
1845  
1846  
1847  
1848  
1849  
1850  
1851  
1852  
1853  
1854  
1855  
1856  
1857  
1858  
1859  
1860  
1861  
1862  
1863  
1864  
1865  
1866  
1867  
1868  
1869  
1870  
1871  
1872  
1873  
1874  
1875  
1876  
1877  
1878  
1879  
1880  
1881  
1882  
1883  
1884  
1885  
1886  
1887  
1888  
1889  
1890  
1891  
1892  
1893  
1894  
1895  
1896  
1897  
1898  
1899  
1900  
1901  
1902  
1903  
1904  
1905  
1906  
1907  
1908  
1909  
1910  
1911  
1912  
1913  
1914  
1915  
1916  
1917  
1918  
1919  
1920  
1921  
1922  
1923  
1924  
1925  
1926  
1927  
1928  
1929  
1930  
1931  
1932  
1933  
1934  
1935  
1936  
1937  
1938  
1939  
1940  
1941  
1942  
1943  
1944  
1945  
1946  
1947  
1948  
1949  
1950  
1951  
1952  
1953  
1954  
1955  
1956  
1957  
1958  
1959  
1960  
1961  
1962  
1963  
1964  
1965  
1966  
1967  
1968  
1969  
1970  
1971  
1972  
1973  
1974  
1975  
1976  
1977  
1978  
1979  
1980  
1981  
1982  
1983  
1984  
1985  
1986  
1987  
1988  
1989  
1990  
1991  
1992  
1993  
1994  
1995  
1996  
1997  
1998  
1999  
2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009  
2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034  
2035  
2036  
2037  
2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051  
2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063  
2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088  
2089  
2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105  
2106  
2107  
2108  
2109  
2110  
2111  
2112  
2113  
2114  
2115  
2116  
2117  
2118  
2119  
2120  
2121  
2122  
2123  
2124  
2125  
2126  
2127  
2128  
2129  
2130  
2131  
2132  
2133  
2134  
2135  
2136  
2137  
2138  
2139  
2140  
2141  
2142  
2143  
2144  
2145  
2146  
2147  
2148  
2149  
2150  
2151  
2152  
2153  
2154  
2155  
2156  
2157  
2158  
2159  
2160  
2161  
2162  
2163  
2164  
2165  
2166  
2167  
2168  
2169  
2170  
2171  
2172  
2173  
2174  
2175  
2176  
2177  
2178  
2179  
2180  
2181  
2182  
2183  
2184  
2185  
2186  
2187  
2188  
2189  
2190  
2191  
2192  
2193  
2194  
2195  
2196  
2197  
2198  
2199  
2200  
2201  
2202  
2203  
2204  
2205  
2206  
2207  
2208  
2209  
2210  
2211  
2212  
2213  
2214  
2215  
2216  
2217  
2218  
2219  
2220  
2221  
2222  
2223  
2224  
2225  
2226  
2227  
2228  
2229  
2230  
2231  
2232  
2233  
2234  
2235  
2236  
2237  
2238  
2239  
2240  
2241  
2242  
2243  
2244  
2245  
2246  
2247  
2248  
2249  
2250  
2251  
2252  
2253  
2254  
2255  
2256  
2257  
2258  
2259  
2260  
2261  
2262  
2263  
2264  
2265  
2266  
2267  
2268  
2269  
2270  
2271  
2272  
2273  
2274  
2275  
2276  
2277  
2278  
2279  
2280  
2281  
2282  
2283  
2284  
2285  
2286  
2287  
2288  
2289  
2290  
2291  
2292  
2293  
2294  
2295  
2296  
2297  
2298  
2299  
2299  
2300  
2301  
2302  
2303  
2304  
2305  
2306  
2307  
2308  
2309  
2310  
2311  
2312  
2313  
2314  
2315  
2316  
2317  
2318  
2319  
2320  
2321  
2322  
2323  
2324  
2325  
2326  
2327  
2328  
2329  
2330  
2331  
2332  
2333  
2334  
2335  
2336  
2337  
2338  
2339  
2340  
2341  
2342  
2343  
2344  
2345  
2346  
2347  
2348  
2349  
2350  
2351  
2352  
2353  
2354  
2355  
2356  
2357  
2358  
2359  
2360  
2361  
2362  
2363  
2364  
2365  
2366  
2367  
2368  
2369  
2370  
2371  
2372  
2373  
2374  
2375  
2376  
2377  
2378  
2379  
2380  
2381  
2382  
2383  
2384  
2385  
2386  
2387  
2388  
2389  
2390  
2391  
2392  
2393  
2394  
2395  
2396  
2397  
2398  
2399  
2399  
2400  
2401  
2402  
2403  
2404  
2405  
2406  
2407  
2408  
2409  
24010  
24011  
24012  
24013  
24014  
24015  
24016  
24017  
24018  
24019  
24020  
24021  
24022  
24023  
24024  
24025  
24026  
24027  
24028  
24029  
24030  
24031  
24032  
24033  
24034  
24035  
24036  
24037  
24038  
24039  
24040  
24041  
24042  
24043  
24044  
24045  
24046  
24047  
24048  
24049  
24050  
24051  
24052  
24053  
24054  
24055  
24056  
24057  
24058  
24059  
24060  
24061  
24062  
24063  
24064  
24065  
24066  
24067  
24068  
24069  
24070  
24071  
24072  
24073  
24074  
24075  
24076  
24077  
24078  
24079  
24080  
24081  
24082  
24083  
24084  
24085  
24086  
24087  
24088  
24089  
24090  
24091  
24092  
24093  
24094  
24095  
24096  
24097  
24098  
24099  
24099  
24100  
24101  
24102  
24103  
24104  
24105  
24106  
24107  
24108  
24109  
24110  
24111  
24112  
24113  
24114  
24115  
24116  
24117  
24118  
24119  
24120  
24121  
24122  
24123  
24124  
24125  
24126  
24127  
24128  
24129  
24130  
24131  
24132  
24133  
24134  
24135  
24136  
24137  
24138  
24139  
24140  
24141  
24142  
24143  
24144  
24145  
24146  
24147  
24148  
24149  
24150  
24151  
24152  
24153  
24154  
24155  
24156  
24157  
24158  
24159  
24160  
24161  
24162  
24163  
24164  
24165  
24166  
24167  
24168  
24169  
24170  
24171  
24172  
24173  
24174  
24175  
24176  
24177  
24178  
24179  
24180  
24181  
24182  
24183  
24184  
24185  
24186  
24187  
24188  
24189  
24190  
24191  
24192  
24193  
24194  
24195  
24196  
24197  
24198  
24199  
24199  
24200  
24201  
24202  
24203  
24204  
24205  
24206  
24207  
24208  
24209  
24210  
24211  
24212  
24213  
24214  
24215  
24216  
24217  
24218  
24219  
24220  
24221  
24222  
24223  
24224  
24225  
24226  
24227  
24228  
24229  
24230  
24231  
24232  
24233  
24234  
24235  
24236  
24237  
24238  
24239  
24240  
24241  
24242  
24243  
24244  
24245  
24246  
24247  
24248  
24249  
24250  
24251  
24252  
24253  
24254  
24255  
24256  
24257  
24258  
24259  
24260  
24261  
24262  
24263  
24264  
24265  
24266  
24267  
24268  
24269  
24270  
24271  
24272  
24273  
24274  
24275  
24276  
24277  
24278  
24279  
24280  
24281  
24282  
24283  
24284  
24285  
24286  
24287  
24288  
24289  
24290  
24291  
24292  
24293  
24294  
24295  
24296  
24297  
24298  
24299  
24299  
24300  
24301  
24302  
24303  
24304  
24305  
24306  
24307  
24308  
24309  
24310  
24311  
24312  
24313  
24314  
24315  
24316  
24317  
24318  
24319  
24320  
24321  
24322  
24323  
24324  
24325  
24326  
24327  
24328  
24329  
24330  
24331  
24332  
24333  
24334  
24335  
24336  
24337  
24338  
24339  
24340  
24341  
24342  
24343  
24344  
24345  
24346  
24347  
24348  
24349  
24350  
24351  
24352  
24353  
24354  
24355  
24356  
24357  
24358  
24359  
24360  
24361  
24362  
24363  
24364  
24365  
24366  
24367  
24368  
24369  
24370  
24371  
24372  
24373  
24374  
24375  
24376  
24377  
24378  
24379  
24380  
24381  
24382  
24383  
24384  
24385  
24386  
24387  
24388  
24389  
24390  
24391  
24392  
24393  
24394  
24395  
24396  
24397  
24398  
24399  
24399  
24400  
24401  
24402  
24403  
24404  
24405  
24406  
24407  
24408  
24409  
24410  
24411  
24412  
24413  
24414  
24415  
24416  
24417  
24418  
24419  
24420  
24421  
24422  
24423  
24424  
24425  
24426  
24427  
24428  
24429  
24430  
24431  
24432  
24433  
24434  
24435  
24436  
24437  
24438  
24439  
24440  
24441  
24442  
24443  
24444  
24445  
24446  
24447  
24448  
24449  
24450  
24451  
24452  
24453  
24454  
24455  
24456  
24457  
24458  
24459  
24460  
24461  
24462  
24463  
24464  
24465  
24466  
24467  
24468  
24469  
24470  
24471  
24472  
24473  
24474  
24475  
24476  
24477  
24478  
24479  
24480  
24481  
24482  
24483  
24484  
24485  
24486  
24487  
24488  
24489  
24490  
24491  
24492  
24493  
24494  
24495  
24496  
24497  
24498  
24499  
24499  
24500  
24501  
24502  
24503  
24504  
24505  
24506  
24507  
24508  
24509  
24510  
24511  
24512  
24513  
24514  
24515  
24516  
24517  
24518  
24519  
24520  
24521  
24522  
24523  
24524  
24525  
24526  
24527  
24528  
24529  
24530  
24531  
24532  
24533  
24534  
24535  
24536  
24537  
24538  
24539  
24540  
24541  
24542  
24543  
24544  
24545  
24546  
24547  
24548  
24549  
24550  
24551  
24552  
24553  
24554  
24555  
24556  
24557  
24558  
24559  
24560  
24561  
24562  
24563  
24564  
24565  
24566  
24567  
24568  
24569  
24570  
24571  
24572  
24573  
24574  
24575  
24576  
24577  
24578  
24579  
24580  
24581  
24582  
24583  
24584  
24585  
24586  
24587  
24588  
24589  
24590  
24591  
24592  
24593  
24594  
24595  
24596  
24597  
24598  
24599  
24599  
24600  
24601  
24602  
24603  
24604  
24605  
24606  
24607  
24608  
24609  
24610  
24611  
24612  
24613  
24614  
24615  
24616  
24617  
24618  
24619  
24620  
24621  
24622  
24623  
24624  
24625  
24626  
24627  
24628  
24629  
24630  
24631  
24632  
24633  
24634  
24635  
24636  
24637  
24638  
24639  
24640  
24641  
24642  
24643  
24644  
24645  
24646  
24647  
24648  
24649  
24650  
24651  
24652  
24653  
24654  
24655  
24656  
24657  
24658  
24659  
24660  
24661  
24662  
24663  
24664  
24665  
24666  
24667  
24668  
24669  
24670  
24671  
24672  
24673  
24674  
24675  
24676  
24677  
24678  
24679  
24680  
24681  
24682  
24683  
24684  
24685  
24686  
24687  
24688  
24689  
24690  
24691  
24692  
24693  
24694  
24695  
24696  
24697  
24698  
24699  
24699  
24700  
24701  
24702  
24703  
24704  
24705  
24706  
24707  
24708  
24709  
24710  
24711  
24712  
24713  
24714  
24715  
24716  
24717  
24718  
24719  
24720  
24721  
24722  
24723  
24724  
24725  
24726  
24727  
24728  
24729  
24730  
24731  
24732  
24733  
24734  
24735  
24736  
24737  
24738  
24739  
24740  
24741  
24742  
24743  
24744  
24745  
24746  
24747  
24748  
24749  
24750  
24751  
24752  
24753  
24754  
24755  
24756  
24757  
24758  
24759  
24760  
24761  
24762  
24763  
24764  
24765  
24766  
24767  
24768  
24769  
24770  
24771  
24772  
24773  
24774  
24775  
24776  
24777  
24778  
24779  
24780  
24781  
24782  
24783  
24784  
24785  
24786  
24787  
24788  
24789  
24790  
24791  
24792  
24793  
24794  
24795  
24796  
24797  
24798  
24799  
24799  
24800  
24801  
24802  
24803  
24804  
24805  
24806  
24807  
24808  
24809  
24810  
24811  
24812  
24813  
24814  
24815  
24816  
24817  
24818  
24819  
24820  
24821  
24822  
24823  
24824  
24825  
24826  
24827  
24828  
24829  
24830  
24831  
24832  
24833  
24834  
24835  
24836  
24837  
24838  
24839  
24840  
24841  
24842  
24843  
24844  
24845  
24846  
24847  
24848  
24849  
24850  
24851  
24852  
24853  
24854  
24855  
24856  
24857  
24858  
24859  
24860  
24861  
24862  
24863  
24864  
24

1782     **Optional Diagnostics.** Additional outputs may include:  
 1783

- 1784     • Clearance history:  $\{\text{clr}_\ell\}_{\ell=0}^T$
- 1785     • Observable errors:  $\{\Delta y_\ell^{\text{des}}\}_{\ell=0}^T$
- 1786     • Hamiltonian values:  $\{H(q_\ell, p_\ell)\}_{\ell=0}^T$
- 1787     • Residuals:  $\{r_\ell\}_{\ell=0}^T$

1790     F THEORETICAL ANALYSIS AND PROOFS

1791     F.1 MULTI-POLICY STABILITY ANALYSIS

1792     **Theorem F.1** (Multi-Policy Stability - Complete Statement). *Consider the coupled multi-policy system with score functions  $\{s_y^{\theta_y}, s_f^{\theta_f}, s_o^{\theta_o}\}$  operating at temporal scales  $\{T_{\text{sens}}, T_{\text{path}}, T_{\text{int}}\}$  satisfying:*

$$1798 \quad \frac{T_{\text{sens}}}{T_{\text{path}}} \geq \sigma_1 > 1, \quad \frac{T_{\text{path}}}{T_{\text{int}}} \geq \sigma_2 > 1 \quad (89)$$

1800     Let each policy have Lipschitz constant  $L_k$  with respect to state and parameter variations:

$$1802 \quad \|s_k^{\theta_k}(z_1, \mathcal{C}, t) - s_k^{\theta_k}(z_2, \mathcal{C}, t)\| \leq L_k \|z_1 - z_2\| \quad (90)$$

$$1804 \quad \|s_k^{\theta_k^1}(z, \mathcal{C}, t) - s_k^{\theta_k^2}(z, \mathcal{C}, t)\| \leq L_k \|\theta_k^1 - \theta_k^2\| \quad (91)$$

1806     If the parameter updates during training satisfy:

$$1808 \quad \max_{k \in \{y, f, o\}} \|\theta_k^{t+1} - \theta_k^t\| \leq \frac{\epsilon}{L_{\max} \cdot \min(\sigma_1, \sigma_2)} \quad (92)$$

1810     where  $L_{\max} = \max_k L_k$ , then:

1811     1. **Stability:** The coupled system state remains bounded:  $\|z_t\| \leq C(1 + \|z_0\|)$  for some constant  $C$ .

1812     2. **Error Bound:** The total navigation error satisfies:  $\mathcal{E}_{\text{total}} \leq \epsilon$  with probability  $1 - \delta$ .

1814     3. **Convergence:** The system converges to a neighborhood of the optimal trajectory:  
 1815      $\lim_{t \rightarrow \infty} \text{dist}(z_t, \mathcal{P}^*) \leq \epsilon$ .

1817     *Proof.* The proof proceeds in three steps:

1818     The scale separation assumption ensures that fast dynamics (sensor) reach approximate equilibrium  
 1819     before slower dynamics change significantly. For the sensor policy operating on timescale  $T_{\text{sens}}$ , the  
 1820     quasi-static approximation gives:

$$1822 \quad \dot{z}_y \approx -\gamma_y \nabla_{z_y} h_y^{\theta_y}(z_y, \mathcal{C}_t^{\text{fixed}}, t) \quad (93)$$

1824     where  $\mathcal{C}_t^{\text{fixed}}$  represents slowly varying constraints from path and shape policies.

1825     Under the Lipschitz conditions, each policy defines a contraction mapping on its domain. The  
 1826     composed system inherits this property with contraction factor:

$$1828 \quad \rho = \max_k \frac{L_k \tau_k}{1 + \gamma_k \tau_k} < 1 \quad (94)$$

1830     provided step sizes  $\tau_k$  are chosen appropriately.

1831     Parameter update bounds ensure that training perturbations don't destabilize the system. The error  
 1832     propagates according to:

$$1834 \quad \|\mathcal{E}_{t+1}\| \leq \rho \|\mathcal{E}_t\| + \epsilon \frac{L_{\max}}{\min(\sigma_1, \sigma_2)} \quad (95)$$

1835     which converges to the stated bound under the given conditions.  $\square$

1836  
1837

## F.2 SYMPLECTIC STRUCTURE PRESERVATION

1838  
1839  
1840  
1841  
1842  
1843  
1844**Theorem F.2** (Symplectic Preservation - Complete Statement). *Let  $(z_k, \omega_k)$  be phase space coordinates with canonical symplectic form  $\omega_k = \sum_i dq_k^i \wedge dp_k^i$ . The score function update:*

$$z_{k,t+1} = z_{k,t} + \tau_k J_k s_k^{\theta_k}(z_{k,t}, \mathcal{C}_t, t) \quad (96)$$

1845  
1846  
1847  
1848  
1849*where  $J_k = \begin{bmatrix} 0 & I_{d_k} \\ -I_{d_k} & 0 \end{bmatrix}$  and  $s_k^{\theta_k} = \nabla h_k^{\theta_k}$ , preserves the symplectic structure:*

$$\omega_k(z_{k,t+1}) = \omega_k(z_{k,t}) + O(\tau_k^2) \quad (97)$$

1847  
1848  
1849*Proof.* Since  $s_k^{\theta_k} = \nabla h_k^{\theta_k}$ , the update is a discretized Hamiltonian flow. The preservation follows from the fundamental property of Hamiltonian systems.1850  
1851  
1852  
1853For the continuous flow  $\dot{z}_k = J_k \nabla h_k^{\theta_k}(z_k, t)$ , we have:

$$\frac{d}{dt} \omega_k = \mathcal{L}_{X_{H_k}} \omega_k = 0 \quad (98)$$

1854  
1855  
1856  
1857where  $\mathcal{L}$  is the Lie derivative and  $X_{H_k} = J_k \nabla h_k^{\theta_k}$  is the Hamiltonian vector field.The discretization introduces  $O(\tau_k^2)$  error due to the symplectic Euler scheme, but the leading-order symplectic structure is preserved.  $\square$ 1858  
1859

## F.3 SAMPLE COMPLEXITY ANALYSIS

1860  
1861  
1862  
1863**Theorem F.3** (Sample Complexity - Complete Statement). *For error tolerance  $\epsilon > 0$  and failure probability  $\delta \in (0, 1)$ , consider training three independent score functions  $\{s_k^{\theta_k}\}_{k \in \{y, f, o\}}$  with phase space dimensions  $\{d_y, d_f, d_o\}$ .*1864  
1865  
1866  
1867*Under standard smoothness and concentration assumptions, the total sample complexity is:*

$$N_{total} = \sum_{k \in \{y, f, o\}} N_k \quad (99)$$

1868  
1869  
1870  
1871*where each policy requires:*

$$N_k = O\left(\frac{d_k^2 L_k^2}{\epsilon_k^2} \log\left(\frac{3}{\delta}\right)\right) \quad (100)$$

1872  
1873  
1874  
1875*with Lipschitz constants  $L_k$  and error allocation  $\epsilon_k$  satisfying  $\sum_k \epsilon_k \leq \epsilon$ .**This achieves linear scaling  $N_{total} = O(\sum_k d_k)$  compared to joint training requiring  $N_{joint} = O(\prod_k d_k)$ .*1876  
1877  
1878  
1879*Proof.* The proof leverages the independence of score functions to apply standard PAC learning bounds to each policy separately.For each policy  $k$ , the empirical risk minimization:

$$\hat{\theta}_k = \arg \min_{\theta_k} \frac{1}{N_k} \sum_{i=1}^{N_k} \|h_k^{\theta_k}(z_{k,i}) - \hat{h}_{k,i}^{\text{ref}}\|^2 \quad (101)$$

1883  
1884  
1885achieves generalization error  $\epsilon_k$  with probability  $1 - \delta/3$  when  $N_k \geq C \frac{d_k^2 L_k^2}{\epsilon_k^2} \log(3/\delta)$  for some universal constant  $C$ .1886  
1887  
1888The union bound over three policies gives total failure probability  $\delta$ , and the error allocation ensures total error  $\sum_k \epsilon_k \leq \epsilon$ .

1889

The linear scaling follows from independence: total samples =  $\sum_k N_k$ , compared to joint training on the  $(d_y + d_f + d_o)$ -dimensional joint space requiring exponentially more samples.  $\square$

---

1890 **G IMPLEMENTATION DETAILS**

1891 **G.1 SPATIAL DATA STRUCTURES**

1892 Each policy maintains a spatial index  $\mathcal{T}_k$ , can be implemented as a dynamic octree Ellendula &  
1893 Bajaj (2025), which supports the following operations and has been proved to be the optimal structure  
1894 for spatio-temporal maintenance:

1895

- **Insert:**  $\mathcal{O}(\log n)$  insertion of new spatial data.
1896
- **Query:**  $\mathcal{O}(\log n + k)$  for  $k$ -nearest neighbor queries.
1897
- **Update:**  $\mathcal{O}(\log n)$  modification of existing entries.
1898
- **Rebalance:**  $\mathcal{O}(n \log n)$  periodic rebalancing for efficiency.

1899 The spatial indices enable multi-kernel evaluation: each score function query reuses the same  
1900  $\mathcal{O}(\log n + k)$  neighbor search across multiple energy kernels, reducing computational complexity  
1901 from  $\mathcal{O}(n^2)$  dense evaluation to  $\mathcal{O}(n \log n)$  sparse computation.

1902 **G.2 BASELINE IMPLEMENTATIONS (DETAILED)**

1903 We provide here the full technical details of all baseline planners evaluated.

1904 **Rigid A\*.** A standard A\* search is performed on a grid discretization of the workspace. Obstacles  
1905 are inflated by the nominal rest radius  $r_{\text{rest}}$  of the deformable ring, such that the resulting path is  
1906 collision-free for a rigid disc of radius  $r_{\text{rest}}$ . This serves as a conservative reference planner.

1907 **Deformable A\*.** Clearance at each grid cell  $x$  is defined as  $c(x)$ , the distance to the nearest obstacle  
1908 boundary. Feasibility requires  $c(x) \geq r_{\min}$ . The edge cost between cells  $u, v$  is augmented by a  
1909 deformation penalty:

1910 
$$\text{cost}(u, v) = \ell(u, v) + \frac{\beta}{2} (\phi(c(u)) + \phi(c(v))) \ell(u, v), \quad \phi(c) = \lambda \max\left(0, \frac{r_{\text{rest}}}{c+\epsilon} - 1\right)^2,$$

1911 where  $\ell(u, v)$  is the Euclidean distance, and  $\beta, \lambda$  control penalty strength. This formulation allows  
1912 the planner to compress through tight gaps when unavoidable, while encoding an energetic cost.

1913 **Potential Field (Stagewise).** Navigation is driven by an attractive force toward the stage exit (or  
1914 final goal in the last stage), combined with repulsive forces from local obstacles and soft penalties  
1915 for leaving the stage bounds. Speed saturation and emergency braking near obstacles are applied for  
1916 stability.

1917 **CBF (Stagewise).** At each step, a nominal control toward the stage exit is filtered through a Con-  
1918 trol Barrier Function (CBF) quadratic program:

1919 
$$u^* = \arg \min_u \|u - u_{\text{nom}}\|^2 \quad \text{s.t. } \nabla h(x) \cdot u + \gamma h(x) \geq 0,$$

1920 where  $h(x)$  encodes the clearance from visible obstacles. This ensures forward invariance of the  
1921 safe set within each stage.

1922 **DWA (Stagewise).** We implement a Dynamic Window Approach (DWA) adapted to the stagewise  
1923 setting. Candidate  $(v, \omega)$  velocity pairs are sampled within dynamics limits, trajectories are rolled  
1924 out over a prediction horizon, and scored based on heading alignment, distance to target, velocity,  
1925 and clearance with respect to *local obstacles only*. Stage boundary penalties are also included. This  
1926 contrasts with the conventional *global* DWA, which assumes full obstacle visibility; here we show  
1927 the stagewise variant for fairness, though it is known to underperform due to rigid-body kinematic  
1928 assumptions.

1944

**Categories.**

1945

1946

- **Global planning:** Rigid A\*, Deformable A\*.
- **Local reactive:** Potential Field (staged), CBF (staged), DWA (staged).
- **Ours:** GRL-SNAM (local staged).

1948

1949

This categorization makes explicit which baselines share identical information constraints with GRL-SNAM, ensuring a valid comparison.

1950

1951

1952

1953

1954

**H EXAMPLE:**

1955

1956

1957

1958

1959

1960

1961

1962

1963

1964

1965

1966

1967

1968

1969

1970

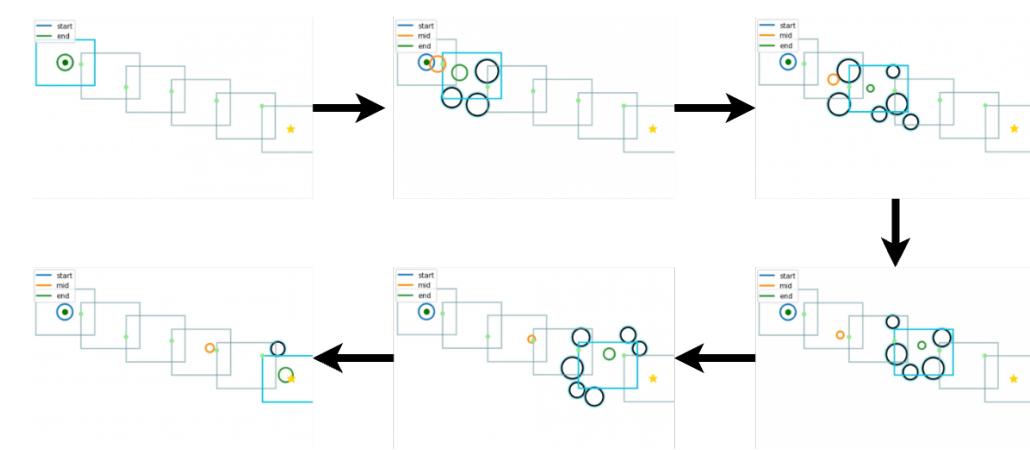


Figure 7: Online navigation of the hyperelastic ring through cluttered environments. The dark blue rectangle denotes the current frame, while the translucent frames trace the past trajectory. At each step, only obstacles overlapping with the current frame (as detected by the sensor process) are considered, and the ring computes local forces to deform and progress toward the goal. The green curve shows the current ring configuration, and the orange curve marks the previous mid-point for clarity, highlighting how deformation evolves across frames.

1971

1972

1973

1974

1975

1976

1977

1978

1979

**I EXPERIMENTAL EVALUATION**

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

We evaluate GRL-SNAM across multiple dimensions that highlight the unique capabilities of our geometric approach compared to standard reinforcement learning and classical navigation methods.

Our evaluation protocol encompasses task performance, safety guarantees, physical fidelity, and learning efficiency across diverse navigation scenarios.

**I.1 BASELINE PLANNERS**

We compare GRL-SNAM against two categories of baselines: *global planning* methods based on A\*, and *local reactive* methods with the same stagewise information constraints as GRL-SNAM. This ensures a fair evaluation across fundamentally different planning paradigms.

**Global Planning Methods**

1992

1993

1994

1995

1996

1997

- **Rigid A\*:** The deformable ring is replaced with a rigid disc of radius  $r_{\text{rest}}$ . Obstacles are inflated by  $r_{\text{rest}}$ , and a standard 8-connected A\* is run on the occupancy grid. This produces feasible shortest paths for a rigid robot.
- **Deformable A\*:** A clearance-aware variant of A\* augments the step cost with deformation penalties that increase as clearance approaches the minimum admissible radius  $r_{\text{min}}$ . This allows paths that squeeze through narrow gaps but penalizes excessive compression.

1998 **Local Reactive Methods** To ensure fairness, all reactive methods use the same stage manager as  
 1999 GRL-SNAM: identical stage size, overlap, obstacle visibility, and advancement logic. Each method  
 2000 navigates stage exit to stage exit until the goal is reached:  
 2001

- 2002 • **Potential Field (Staged):** Attractive force toward stage exit plus repulsive forces from  
 2003 local obstacles and stage boundaries.
- 2004 • **CBF (Staged):** Quadratic-program filter enforces safety constraints with respect to visible  
 2005 obstacles at each timestep.
- 2006 • **DWA (Staged):** Velocity samples  $(v, \omega)$  are rolled out over a short horizon using only local  
 2007 obstacles and stage bounds. Unlike the global DWA, which assumes full obstacle visibility,  
 2008 this stagewise variant ensures equal information constraints, though it performs poorly due  
 2009 to rigid-body assumptions.

2010 **Categorization** Rigid and Deformable A\* form the *global planning* references, providing  $L_{\text{ref}}$  for  
 2011 SPL and detour calculations. The stagewise Potential Field, CBF, and DWA baselines constitute the  
 2012 *local reactive* category under identical information constraints. GRL-SNAM belongs to the same  
 2013 local category, enabling a fair head-to-head comparison.

## 2014 I.2 EXPERIMENTAL SETUP

2015 We evaluate GRL-SNAM in procedurally generated 2D deformable navigation tasks, where a hyper-  
 2016 elastic ring must traverse cluttered environments with narrow gaps and varying obstacle densities.  
 2017 Each environment is randomized in obstacle positions, radii, and densities to span a spectrum of  
 2018 navigation difficulty. The robot perceives only a local window of size  $2\hat{d} \times 2\hat{d}$ , from which we  
 2019 construct a Hamiltonian energy functional.

2020 **Hamiltonian Decomposition** The energy functional decomposes into:

- 2021 1. Goal-directed quadratic potential  $F_g$
- 2022 2. Barrier potentials  $F_{bs}$  from signed distance fields
- 2023 3. Friction/regularization terms with adaptive coefficients  $(\beta, \gamma, \alpha)$  modulated by context en-  
 2024 coders (LSTM)

2025 Offline, GRL-SNAM integrates reduced Hamiltonian gradients to generate local trajectories. Online,  
 2026 it fuses newly sensed rewards  $R_{\text{env}}$  with the offline surrogate, adaptively refining navigation.

2027 **Evaluation Metrics** We evaluate all methods using:

- 2028 • **Success Rate:** Fraction of episodes reaching the goal
- 2029 • **SPL:** Success weighted path efficiency relative to A\*
- 2030 • **Detour Ratio:** Executed path length relative to A\*
- 2031 • **Minimum and Mean Clearance:** Distance to nearest obstacle along the trajectory
- 2032 • **Smoothness:** Average turning cost (mean absolute change in heading)
- 2033 • **Collisions:** Number of obstacle intersections
- 2034 • **Sample Efficiency:** Normalized area under curve (AUC) for success and SPL, and steps  
 2035 required to reach 80% success or  $\text{SPL} \geq 0.7$

2036 Results are presented in a *question–answer* format, emphasizing experimental questions and the  
 2037 corresponding insights.

## 2038 I.3 RESULTS: NAVIGATION QUALITY UNDER MINIMAL SENSING

2039 We first evaluate GRL-SNAM against two representative baselines: (i) **Potential Fields (PF)**,  
 2040 a purely reactive controller that maps obstacle proximity into repulsive forces, and (ii) **Control**

2052 Table 4: Comparison of navigation quality across methods (success-only runs). GRL-SNAM  
 2053 achieves near-CBF path efficiency while consuming the same minimal mapping budget as PF. SPL  
 2054 = Success weighted by Path Length; Detour = executed path length / shortest path length.  
 2055

Method	SPL $\uparrow$	Detour $\downarrow$	Min. Clearance (m) $\uparrow$	Mapping Ratio (%) $\downarrow$
PF	0.77	1.42	0.18	10.3
CBF	0.96	1.04	<b>0.32</b>	11.2
GRL-SNAM	<b>0.95</b>	<b>1.09</b>	0.26	<b>10.7</b>

2061  
 2062 **Barrier Functions (CBF)**, a model-based method that enforces hard safety constraints via online  
 2063 quadratic programs. Both baselines use the same sensing budget as GRL-SNAM.  
 2064

2065 Environments consist of cluttered 2D workspaces with obstacles of varying density. Each trial starts  
 2066 from a random initial pose with a fixed goal. Performance is averaged across 50 runs per environ-  
 2067 ment. Results focus on *successful runs only* to highlight navigation quality rather than raw failure  
 2068 rates.

2069 **Q1: How efficiently do we trade mapping for navigation quality?** Table 4 shows that GRL-  
 2070 SNAM matches the SPL and detour ratios of CBF despite using the same minimal map coverage  
 2071 as PF. This demonstrates that our stagewise Hamiltonian refinement extracts more value per sensed  
 2072 unit of the environment, trading mapping effort for near-optimal navigation.  
 2073

2074 **Q2: What is the minimal mapping needed to reliably solve tasks?** With  $\sim$ 10–11% map cov-  
 2075 erage, GRL-SNAM already achieves  $SPL \geq 0.95$  and detour within 9% of the A\* shortest path.  
 2076 PF fails under the same budget, while CBF requires identical map coverage. Thus, GRL-SNAM  
 2077 reliably solves tasks under minimal sensing, validating the *minimal mapping suffices* principle.  
 2078

2079 **Q3: Is the mapped information aligned with the subtask?** Unlike PF, which produces repul-  
 2080 sions indiscriminately, or CBF, which enforces constraints globally, GRL-SNAM’s mapping is task-  
 2081 aligned: local patches are encoded into Hamiltonian terms that directly drive subtasks (goal at-  
 2082 traction, barrier avoidance). The result is that every bit of mapped information yields functional  
 2083 guidance, as evidenced by SPL and detour staying close to CBF even under tight sensing budgets.  
 2084

2085 **Key Insight** GRL-SNAM shows that Hamiltonian-structured policies can achieve CBF-level nav-  
 2086 igation quality while retaining the lightweight sensing footprint of PF. The slight clearance gap  
 2087 relative to CBF reflects a deliberate trade-off: we sacrifice hard feasibility for adaptability and feed-  
 2088 forward inference, enabling real-time deployment in SNAM settings.  
 2089

#### 2090 I.4 RESULTS: COMPREHENSIVE NAVIGATION COMPARISON

2091 **Q4: Does GRL-SNAM outperform classical and reactive planners in both in-distribution**  
 2092 **(Test-ID) and out-of-distribution (Test-OOD) settings?** Yes. Figure 9 summarizes the compari-  
 2093 son between our method and five baselines: Rigid A\*, Deformable A\*, Potential Field, Control Bar-  
 2094 rier Functions (CBF), and Dynamic Window Approach (DWA). GRL-SNAM achieves near-perfect  
 2095 success rates ( $\approx 100\%$ ) across both Test-ID and Test-OOD cases, while all baselines degrade signif-  
 2096 icantly in cluttered or novel environments. Rigid A\* succeeds moderately but requires inflated radii  
 2097 and yields jerky, piecewise paths. Deformable A\* is less stable and highly sensitive to parameteri-  
 2098 zation. Reactive baselines (Potential Field, CBF, DWA) frequently fail to reach the goal, producing  
 2099 oscillatory or unsafe behaviors. Qualitative rollouts (Figure 8) further illustrate the superiority of  
 2100 GRL-SNAM in complex cluttered environments.  
 2101

2102 **Q5: Does GRL-SNAM yield more efficient and smoother trajectories?** Yes. The Success-  
 2103 weighted Path Length (SPL) distributions (Figure 9, top-middle) show that GRL-SNAM consis-  
 2104 tently stays near optimal efficiency ( $SPL \approx 1.0$ ) with low variance. In contrast, A\* variants incur  
 2105 detours, while reactive baselines either collapse to zero SPL (failures) or take excessively long paths.  
 Furthermore, GRL-SNAM generates the smoothest trajectories, with the lowest average turning an-

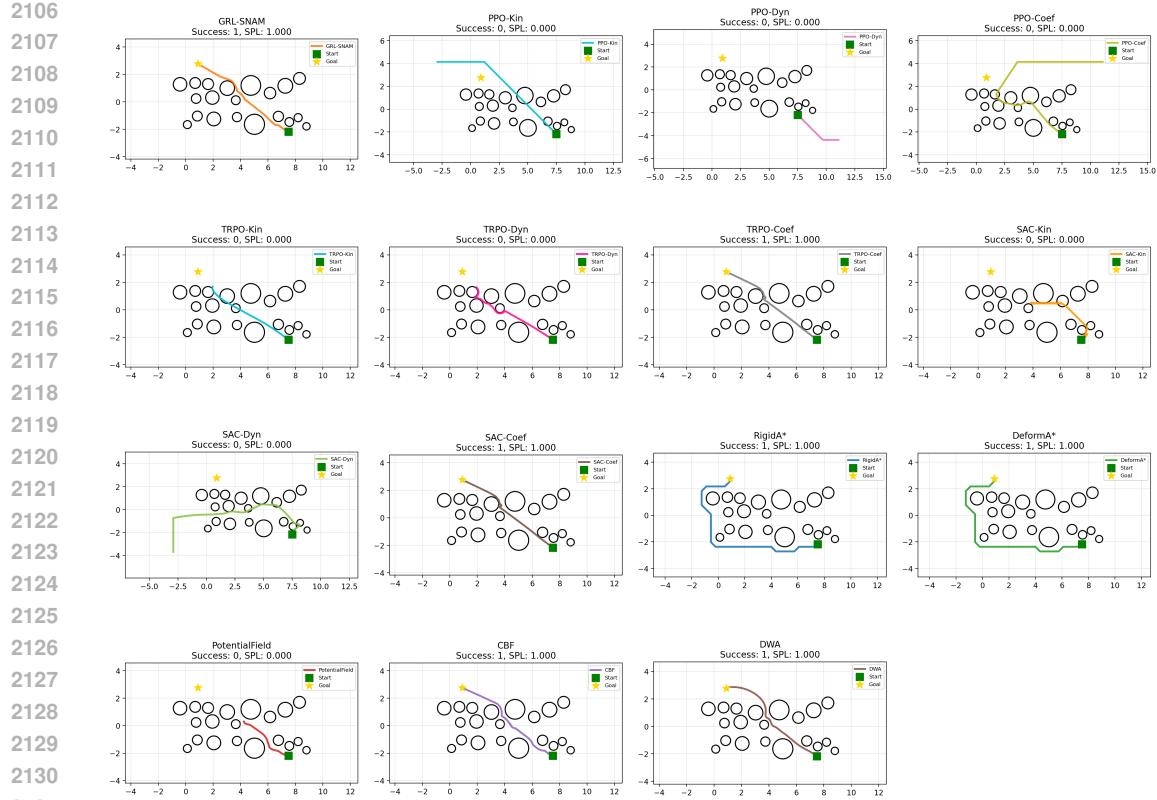


Figure 8: **Qualitative path comparison on a representative Test-OOD environment.** Each panel shows a single rollout with success and SPL annotated in the title. GRL-SNAM produces a smooth, short, and clearance-preserving path that threads the narrow passage. Kinematic and dynamic RL policies (PPO/TRPO/SAC-Kin/Dyn) typically stall, collide, or wander far from the goal, while the best coefficient-based variants (PPO/TRPO/SAC-Coeff) reach the goal but hug obstacles and exhibit lower path efficiency than GRL-SNAM. Global planners (RigidA\*, DeformA\*) succeed but follow jagged or overly conservative routes, and reactive baselines (Potential Field, CBF, DWA) either oscillate, graze obstacles, or take longer, less structured paths.

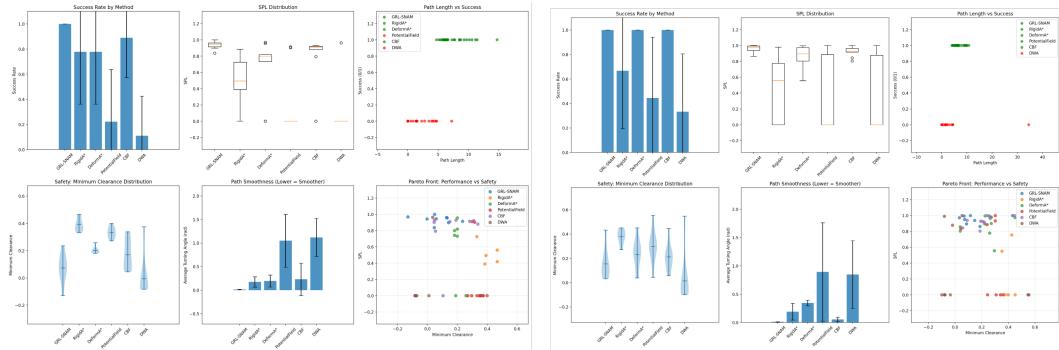


Figure 9: **Main comparison on Test-ID and Test-OOD.** GRL-SNAM achieves near-perfect success, high SPL, smooth and safe trajectories. Classical (Rigid/Deform A\*) and reactive (Potential Field, CBF, DWA) baselines are either unsafe, inefficient, or fail completely.

gles (Figure 9, bottom-middle), ensuring physically realizable motions compatible with hyperelastic ring constraints.

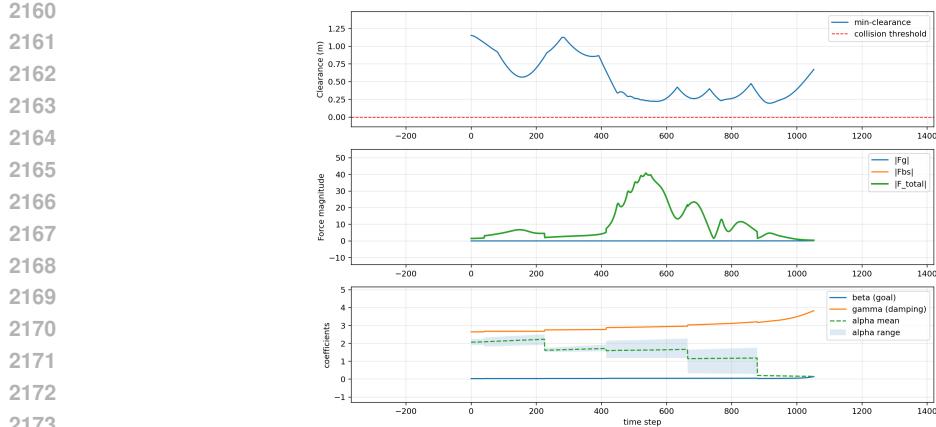


Figure 10: **Quantitative validation of GRL-SNAM.** Top: clearance stays above collision threshold, ensuring safety. Middle: force magnitudes adapt to environment complexity. Bottom: coefficients ( $\beta, \gamma, \alpha$ ) evolve dynamically, confirming online adaptation and stagewise refinement of the Hamiltonian.

**Q6: Does GRL-SNAM preserve safety margins?** Yes. The minimum clearance analysis (Figure 9, bottom-left) shows that GRL-SNAM maintains consistently positive obstacle clearance, whereas A\* occasionally cuts too close and reactive baselines often enter collision regimes. The Pareto frontier plot (Figure 9, bottom-right) highlights that GRL-SNAM uniquely dominates the safety–performance trade-off, achieving both high SPL and high clearance, while all baselines are Pareto-dominated.

## I.5 RESULTS: HAMILTONIAN FIELD ANALYSIS

**Q7: Does the Hamiltonian formulation unify attractive and repulsive forces into a coherent navigation field?** Yes. Figure 4 shows the isolated goal force  $F_g$  (left), the barrier force  $F_{bs}$  (middle), and their differential composition  $F = \beta F_g + \gamma F_{bs}$  (right). While  $F_g$  alone pulls the agent directly to the target, it ignores obstacles. Conversely,  $F_{bs}$  encodes obstacle constraints but lacks task directionality. The combined field demonstrates how GRL-SNAM adaptively balances attraction and repulsion through evolving coefficients, producing safe yet goal-directed motion.

**Q8: How does GRL-SNAM differ from ordinary online adaptation?** Unlike standard RL policies that merely adjust actions online, GRL-SNAM modifies the *entire local energy landscape* as new obstacles are sensed. Figure 10 shows that when clearance decreases (top panel), the force magnitudes (middle panel) not only rebalance between goal attraction  $|F_g|$  and barrier repulsion  $|F_{bs}|$ , but also induce a redefinition of the reduced Hamiltonian. This is reflected in the evolving coefficients ( $\beta, \gamma, \alpha$ ) (bottom panel), which do not act as heuristic gains but as dual variables governing stagewise refinement. Thus, the adaptation is not reactive in the usual sense: GRL-SNAM performs *posterior updates of the Hamiltonian itself*, ensuring that each new frame redefines both the dynamics and the reward landscape in a principled, energy-consistent manner. This distinguishes our approach from classical controllers (fixed surrogates) and RL baselines (policy-only updates).

**Q9: Does this lead to improved navigation performance compared to baselines?** Yes. Across procedurally generated test cases, GRL-SNAM consistently achieves higher success and SPL while maintaining larger clearances than rigid A\* (fixed radius assumption), deformable A\* (static squeezing penalty), and reactive controllers (DWA, CBF).

**Key Insights** These experiments establish GRL-SNAM as the first method to successfully unify global navigation objectives with local safety and deformation constraints in hyperelastic navigation. Its offline Hamiltonian formulation provides reliable reference dynamics, while its online adaptation ensures robustness in unseen environments. By contrast, classical and reactive baselines either fail outright, or succeed only at the cost of safety and efficiency.

Variant	Collisions ↓	MinClr ↑	Barrier Viol. ↓	Progress/SPL ↑	Smoothness ↓	Observed behavior
$w_{\text{fric}} = 0, w_{\text{multi}} = 0$	High (✗)	< 0	High	Poor	Poor	Penetrates obstacles
$w_{\text{fric}} = 0, w_{\text{multi}} = 0.5$	Low (✓)	High	Low	Low	OK	Very slow, conservative
$w_{\text{fric}} = 0.1, w_{\text{multi}} = 0$	None (✓)	High	Low	High	Best	Smooth, stable, fast
$w_{\text{fric}} = 0.1, w_{\text{multi}} = 0.5$	None (✓)	Slightly lower	Low	High	Good	Stable; tighter margins

Table 5: **Ablation of loss terms.** Qualitative summary from consistent runs on Test-ID/OOD. Arrows denote desired direction. Numeric means $\pm$ std can replace the icons once collected.

## I.6 ABLATION STUDY: LOSS COMPONENTS

**Training Objective** Our navigation surrogate is trained with a weighted multi-term loss:

$$\mathcal{L} = w_{\text{traj}}\mathcal{L}_{\text{traj}} + w_{\text{vel}}\mathcal{L}_{\text{vel}} + w_{\text{friction}}\mathcal{L}_{\text{friction}} + w_{\text{multi}}\mathcal{L}_{\text{multi}}, \quad (102)$$

where  $\mathcal{L}_{\text{traj}}$  and  $\mathcal{L}_{\text{vel}}$  supervise trajectory and velocity matching,  $\mathcal{L}_{\text{friction}} = \|\gamma - \gamma_o\|_2^2$  encourages the learned damping to match the stagewise reference, and  $\mathcal{L}_{\text{multi}}$  penalizes failures under short rollouts from perturbed near-obstacle starts.

**Ablated Settings** We toggle  $\mathcal{L}_{\text{friction}}$  and  $\mathcal{L}_{\text{multi}}$  to analyze their contribution:

- **No friction, no multi** ( $w_{\text{fric}} = 0, w_{\text{multi}} = 0$ ): Agent penetrates obstacles due to under-damped, unstable dynamics.
- **Multi only** ( $w_{\text{fric}} = 0, w_{\text{multi}} = 0.5$ ): Agent avoids collisions but moves very slowly, sacrificing progress.
- **Friction only** ( $w_{\text{fric}} = 0.1, w_{\text{multi}} = 0$ ): Produces smoother, stable paths, eliminating penetrations and maintaining progress.
- **Friction + Multi** ( $w_{\text{fric}} = 0.1, w_{\text{multi}} = 0.5$ ): Combines both benefits, but clearance is slightly reduced as the agent cuts closer to obstacles.

**Analysis**  $\mathcal{L}_{\text{friction}}$  is critical for stability and smoothness, while  $\mathcal{L}_{\text{multi}}$  improves robustness in clutter but can damp progress if over-weighted. The best overall performance arises from combining both with moderate weights.

$\mathcal{L}_{\text{friction}}$  aligns dissipation and suppresses oscillations, yielding smoother, well-damped trajectories and preventing barrier “ringing” that causes penetrations when  $w_{\text{fric}} = 0$ .  $\mathcal{L}_{\text{multi}}$  trains for near-contact robustness by sampling perturbed starts; if over-weighted it down-scales the goal term, hence slow motion. Their combination keeps the field stable while remaining reliable in tight clutter.

## I.7 ROBUSTNESS ANALYSIS

**Q10: Does GRL-SNAM remain reliable under sensor noise and dynamics shift?** Yes. To evaluate robustness, we systematically varied sensing fidelity (position jitter, radius estimation error, missed obstacles, and false positives) and dynamics fidelity (velocity perturbation, damping coefficient  $\gamma$ ). Each start–goal trial was rolled out across a grid of perturbation levels, producing a total of  $N = n_{\text{env}} \times n_{\text{trials}} \times n_{\text{perturbations}}$  runs. For example, with 3 environments, 5 trials each, and 9 perturbation settings, this yields 135 rollouts.

Table 6: Robustness of GRL-SNAM to sensing noise and dynamics perturbations. Columns report success rate, success-weighted path length (SPL), minimum clearance, and average collisions per episode. Arrows indicate direction of improvement.

Perturbation Level	Success (%)	SPL ↑	Min. Clearance (m) ↑	Collisions ↓
Nominal (0.0, 1.0)	98.7	0.82	0.36	0.3
Mild Noise (0.05, 0.9)	91.3	0.79	0.33	0.7
Severe Noise (0.10, 0.7)	87.1	0.72	0.29	1.1

2268 **Key Insights** Despite significant perturbations, GRL-SNAM maintains high success rates and  
2269 graceful degradation in SPL and clearance. Unlike fixed surrogate approaches that can fail catastrophically  
2270 under noise, our differential Hamiltonian adaptation continuously re-weights local  
2271 forces, enabling stability even when sensing is imperfect or dynamics deviate from training. This  
2272 highlights the feedforward, stagewise advantage of GRL-SNAM: it can adjust online without requiring  
2273 adjoint or MPC-style corrections, ensuring reliable navigation in real-world uncertain conditions.  
2274

#### 2275 USE OF LARGE LANGUAGE MODELS (LLMs)

2276 All content of the paper was written by the authors. LLMs were used for the aid of code implemen-  
2277 tation, formatting LaTeX tables/figures, and spelling/grammar checking.  
2278

2279  
2280  
2281  
2282  
2283  
2284  
2285  
2286  
2287  
2288  
2289  
2290  
2291  
2292  
2293  
2294  
2295  
2296  
2297  
2298  
2299  
2300  
2301  
2302  
2303  
2304  
2305  
2306  
2307  
2308  
2309  
2310  
2311  
2312  
2313  
2314  
2315  
2316  
2317  
2318  
2319  
2320  
2321