

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 GDC: FROM BRITTLE OPTIMALITY TO ROBUST SATIS- FIABILITY VIA RIEMANNIAN RISK GEOMETRY

Anonymous authors

Paper under double-blind review

ABSTRACT

Standard reinforcement learning (RL) often yields brittle policies that fail under hard safety constraints. We propose *Geodesic Duality Control* (GDC), which adapts an agent’s risk posture endogenously by re-weighting the Bellman target using local geometric cues of the value function (gradient magnitude and a curvature surrogate). To accommodate piecewise-smooth neural critics we formulate a Sub-Riemannian / generalized-gradient treatment and provide practical, numerically stable curvature surrogates (implementation details in Sec. 2.3 and App. B). Our main theoretical result shows that, under explicit regularity and stochastic-model assumptions, GDC induces a curvature-decreasing learning dynamic that increases a quantifiable safety margin (proofs in App. A). We validate the mechanism with proof-of-concept experiments—including a hard-boundary safety environment (Optimal-Trap), targeted ablations, and a computational-cost study on Humanoid-Safety—to confirm the intended geometric risk posture. We do not claim broad empirical superiority on all benchmarks; rather, the paper’s primary contribution is theoretical, with key components validated empirically.

1 INTRODUCTION

Reinforcement Learning (RL) optimizes sequential decision-making by searching for policies that maximize expected cumulative reward (Sutton & Barto, 2018). While highly successful in many domains, vanilla RL frequently produces agents with *brittle optimality*: policies that perform well on average but are prone to catastrophic safety violations when faced with sparse, hard constraints or abrupt safety frontiers.

Existing safe-RL approaches take distinct, but related, philosophies (Cen et al., 2024; Cheng et al., 2023; Dai et al., 2024; Gu et al., 2024; Kim et al., 2024; Lei et al., 2024; Ma et al., 2021; Thananjeyan et al., 2021; Yao et al., 2023). Constrained RL (e.g., CPO, PCPO) enforces *global* cumulative cost budgets (Achiam et al., 2017; Schulman et al., 2015). Risk-sensitive methods (e.g., CVaR-based objectives) reshape the return distribution to penalize tails (Chow & Ghavamzadeh, 2015; Singh et al., 2020; Wang et al., 2023). Robust RL targets distributional shifts through worst-case or augmentation strategies (Sun et al., 2024; Wang et al., 2020). These families are effective in many settings but share common limitations: (i) they treat safety as an external budget or static measure rather than a state-dependent signal, (ii) global constraints can be overly conservative or fail to prevent local catastrophes, and (iii) many theoretical guarantees assume smooth critics or strong regularity that modern neural approximators (e.g., ReLU networks) do not satisfy (Clarke, 1983b).

To address these gaps we present **Geodesic Duality Control (GDC)**. Rather than imposing an external budget, GDC treats risk as an *intrinsic, local* property of the critic’s geometry and uses that property to modulate learning. Concretely, GDC computes a local risk metric $\kappa(s, a; Q)$ from the critic Q (combining gradient norm and a curvature surrogate) and maps κ to a smooth weight $\sigma(Q)$ that continuously re-weights the Bellman target between reward-seeking and penalty-avoidance. This *endogenous* coupling creates a closed feedback loop: the critic’s geometry directly shapes the update target that in turn sculpts the critic.

Our method contributes three main innovations and practical advantages:

054

- 055 1. **Endogenous geometry-aware risk weighting.** Unlike global-budget or static-risk ap-
056 proaches, GDC directly embeds local geometric information into the Bellman target, en-
057 abling the agent to increase caution precisely where the value landscape indicates brittleness.
- 058 2. **Theory for piecewise-smooth critics.** We develop a Sub-Riemannian / generalized-gradient
059 formalism that accommodates non-smooth neural critics and prove that, under explicit
060 regularity and stochastic assumptions, the induced learning dynamic is equivalent (in the
061 mean-field sense) to a curvature-decreasing geometric flow that enlarges a quantifiable safety
062 margin. Technical statements and proofs appear in App. A.
- 063 3. **Practical, numerically-stable curvature surrogates and implementation alignment.**
064 We provide concrete algorithms (damped Lanczos / power-iteration with Tikhonov regu-
065 larization, finite-difference fallbacks, and optional mollification) to estimate curvature in
066 piecewise-smooth critics, and we explicitly analyze how implementation choices (e.g., using
067 a lagged/target critic for stability) affect theoretical claims. These recipes keep computa-
068 tional overhead modest while preserving the intended geometric effect (see Sec. 2.3 and
069 Sec. 5).

070

071 Throughout the paper we emphasize transparency and scope: the theoretical contributions are primary;
072 experiments are targeted proof-of-concept studies chosen to stress the phenomena our theory predicts
073 (including a hard-boundary safety test and ablations). We report a computational-cost study on
074 a high-dimensional Humanoid-Safety task but do not assert comprehensive empirical dominance
075 over all existing safe-RL baselines across every benchmark. The remainder of the paper presents
076 the formalism (Sec. 3), implementation details and curvature estimators (Sec. 2.3), experimental
077 validation (Sec. 5), and full technical proofs and sensitivity analyses in the appendix.

078 2 THE GDC FRAMEWORK: ENDOGENOUS GEOMETRIC RISK WEIGHTING

079 2.1 PRELIMINARIES

080 We consider a standard Markov Decision Process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$,
081 where \mathcal{S} is the state space, \mathcal{A} is the action space, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition kernel,
082 $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward, and $\gamma \in [0, 1]$ is the discount factor. The goal is to find a policy π
083 maximizing the expected return $J(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$. The action–value function is
084 $Q^{\pi}(s, a) = \mathbb{E}_{\pi} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \mid s_0 = s, a_0 = a]$ (Sutton & Barto, 2018; Cai et al., 2022).

085 2.2 THE CORE MECHANISM OF GDC

086 At the heart of GDC is a simple philosophy: risk is not an external constraint but an *intrinsic, local*
087 *property of the value function’s geometry*. GDC perceives this geometry and reacts to it via three
088 components: a geometric risk metric κ , an endogenous switch σ , and a geometry-aware Bellman
089 operator \mathcal{T}_G .

090 **1. The Geometric Risk Metric κ .** A complete local risk profile should capture both first-order
091 *steepness* and second-order *curvature*. Relying on only one of them leads to blind spots (e.g., missing
092 sharp drops after flat plateaus or ignoring steep non-curved descents). We therefore define:

093 **Definition 2.1** (Geometric Risk Metric κ). *For a critic Q and state s , the local risk at action a is*

$$\kappa(s, a) = \underbrace{\|\nabla_a^H Q(s, a)\|_2}_{\text{Steepness (Gradient)}} + c \underbrace{\max(0, -\lambda_{\min}(H_a^H(Q(s, a))))}_{\text{Concavity (Hessian)}}, \quad (2.1)$$

094 where ∇_a^H and H_a^H are the gradient and Hessian of Q restricted to the action horizontal directions
095 \mathcal{D}_a (see Section 2.3; formal details in Section D.1). Here λ_{\min} is the minimum eigenvalue, and $c > 0$
096 is a weight.

097 The composite form is crucial and is validated by ablations and a purpose-built adversarial test
098 (Section 5).

099 **2. Endogenous Dynamic Risk Weighting.** GDC maps the risk to a smooth switch $\sigma : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$:

$$\sigma(s, a; Q) = \text{sigmoid}(\kappa(s, a) - \kappa_0), \quad (2.2)$$

108 where κ_0 is a tunable threshold and $k > 0$ controls the slope. The dependence $\sigma(\cdot; Q)$ emphasizes
 109 its *endogenous* nature—the critic’s own geometry modulates its target.
 110

111 **3. Geometry-Aware Bellman Operator \mathcal{T}_G .** This weighting integrates directly into the target:
 112 **Definition 2.2** (GDC Operator \mathcal{T}_G).

113

$$114 \quad (\mathcal{T}_G Q)(s, a) = \mathbb{E}_{s' \sim P(s, a)} \left[R_G(s, a; Q) + \gamma \max_{a'} Q(s', a') \right], \quad (2.3)$$

115

116 with the geometry-aware reward

117

$$118 \quad R_G(s, a; Q) = (1 - \sigma(s, a; Q)) R(s, a) + \sigma(s, a; Q) \min(0, R(s, a)). \quad (2.4)$$

119

120 In low-risk regions ($\sigma \approx 0$) this reduces to the standard Bellman target; in high-risk regions ($\sigma \approx 1$) it
 focuses on penalty mitigation to enforce conservative behavior.

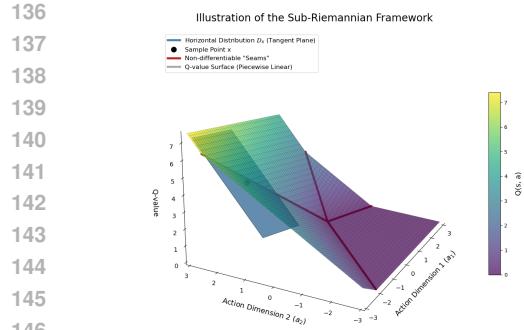
121 **Design rationale and sanity properties.**

122

- 123 **Scaling.** Replacing Q by αQ scales κ by α ; this can be absorbed by reward normalization or
 124 retuning k without changing the qualitative behavior of σ .
- 125 **Limits in c .** $c \rightarrow 0$ gives a gradient-only detector (fast but blind to degenerate ridges); large c
 126 increases conservatism near negative curvature.
- 127 **Limits in k .** $k \rightarrow 0$ makes $\sigma \rightarrow \frac{1}{2}$ (uniform tempering); $k \rightarrow \infty$ yields a hard switch at κ_0 (useful
 128 when κ is well-calibrated).
- 129 **Monotonicity.** σ is nondecreasing in κ ; increasing κ_0 reduces the expected positive-reward
 130 contribution in the target.
- 131 **Locality.** $\kappa(s, a)$ depends only on local geometry along admissible action directions (Section 2.3);
 132 no global multipliers are needed, unlike Lagrangian CRL.

133 **2.3 GEOMETRIC FOUNDATION AND NUMERICAL IMPLEMENTATION**

135



147 **Figure 1: Sub-Riemannian view of a ReLU**
 148 **critic.** Non-differentiable seams (red) break clas-
 149 sical calculus. We work in the local horizontal
 150 distribution D_x (blue plane).

151 Tikhonov shift $\delta = 10^{-6}$ to stabilize the smallest Ritz value; bandwidth $\varepsilon = 10^{-3}$ (scaled by feature
 152 std). Each κ evaluation costs $\mathcal{O}(m)$ HVPs ($\approx m$ backwardequivalents). We vectorize across the
 153 batch and action dimensions; see Section D.3 for complexity notes.

From activation stability to horizontal operators (practical). A ReLU critic decomposes into linear regions with fixed activation masks. Let $U_a(x)$ be an orthonormal basis of the *horizontal action subspace* $D_a(x)$ that preserves the current mask (formalized in Section D.1). We evaluate geometric quantities on the mollified critic Q_ε restricted to $D_a(x)$.

$$154 \quad \nabla_a^H Q(x) = U_a(x)^\top \nabla_a Q_\varepsilon(x), \quad (2.5)$$

$$155 \quad H_a^H(Q; x) = U_a(x)^\top \nabla_{aa}^2 Q_\varepsilon(x) U_a(x).$$

156 This (i) avoids crossing seams, (ii) reduces estimator
 157 variance, and (iii) lowers the Lanczos cost to the small
 158 subspace dimension d_H .

Robust curvature surrogate. We estimate
 $\lambda_{\min}(H_a^H)$ by m -step Lanczos with an HVP oracle
 159 defined on $D_a(x)$ (see Section D.2). Recommended
 160 defaults (validated in Section 5): $m = 8$ steps; optional
 161 Tikhonov shift $\delta = 10^{-6}$ to stabilize the smallest Ritz value; bandwidth $\varepsilon = 10^{-3}$ (scaled by feature
 162 std). Each κ evaluation costs $\mathcal{O}(m)$ HVPs ($\approx m$ backwardequivalents). We vectorize across the
 163 batch and action dimensions; see Section D.3 for complexity notes.

164 **Implementation checklist (drop-in)**

165

- 166 Use the *target* critic Q_{tgt} to compute κ (reduces drift in σ).
- 167 Build $U_a(x)$ from the current linear region; cache & reuse when masks match.
- 168 Clip κ_0 to $[0, \kappa_{\max}]$ with κ_{\max} set by the 95th percentile of warmup κ .

162 **Algorithm 1** Adaptive GDC with Soft Actor-Critic (A-GDC-SAC)
163
164 **Require:** initial $\kappa_0, c, C_{\text{target}}, \beta, \eta$
165 1: initialize critics $Q_{\theta_1}, Q_{\theta_2}$, actor π_{ϕ} , and target nets
166 2: $\bar{C} \leftarrow 0$
167 3: **for** each training step **do**
168 4: sample minibatch $\mathcal{B} = \{(s, a, r, s', d, \text{cost})\}$
169 5: $\bar{C} \leftarrow (1 - \beta)\bar{C} + \beta \cdot \text{mean}(\text{cost})$
170 6: $\kappa_0 \leftarrow \max(0, \kappa_0 + \eta(\bar{C} - C_{\text{target}}))$
171 7: **for** each $(s, a, r, s', d, \text{cost}) \in \mathcal{B}$ **do**
172 8: compute Q_{tgt} , next action a' , and $\kappa(s', a')$ via HVP + Lanczos
173 9: $\sigma \leftarrow \text{sigmoid}(k(\kappa(s', a') - \kappa_0))$
174 10: $r_G \leftarrow (1 - \sigma)r + \sigma \min(0, r)$
175 11: form GDC target and accumulate critic loss
176 12: update critics/actor/temperature and target nets

177

178
179

- Normalize rewards so $\|R\|_{\infty}$ is comparable across tasks; retune k coarsely if needed.

180

181
182
183

Geometric foundation for non-smooth functions. Value approximators with ReLU activations are not globally C^2 . We ground our framework in **Sub-Riemannian geometry** (horizontal distribution \mathcal{D} on $M = \mathcal{S} \times \mathcal{A}$); see Section D.1 for the formalization.

184
185
186

Definition 2.3 (Sub-Riemannian structure). *We model (M, \mathcal{D}, g) with \mathcal{D}_x spanned by directions where the Q -function is differentiable (i.e., not crossing seams in Figure 1); g is a smooth inner product on \mathcal{D} .*

187
188
189
190
191

Numerical estimation of curvature for ReLU networks. The classical Hessian of a ReLU network is a.e. zero or undefined. Instead, we compute an *effective curvature surrogate* via HVPs and a few Lanczos iterations on a locally smoothed critic Q_{ε} ; this robustly approximates local concavity without explicitly constructing the manifold (Clarke, 1983a).

192
193

2.4 AUTOMATING RISK SENSITIVITY WITH A-GDC

194
195

Introducing κ_0 raises the question of tuning. We propose **Adaptive GDC (A-GDC)**, which adjusts κ_0 based on recent safety performance. With $\text{cost}_t \in \{0, 1\}$ indicating a violation:

$$\bar{C}_{t+1} = (1 - \beta)\bar{C}_t + \beta \cdot \text{cost}_t, \quad (2.6)$$

$$\kappa_{0,t+1} = \max(0, \kappa_{0,t} + \eta(\bar{C}_{t+1} - C_{\text{target}})). \quad (2.7)$$

196
197
198
199
200
201

Unlike Lagrangian CRL, this tunes the *trigger sensitivity of a local geometric response*. Experiments in Section 5 (Table 3) compare A-GDC to a fixed hand-tuned κ_0 .

Controller calibration (practical recipe). Choose $C_{\text{target}} \in [0.005, 0.02]$, $\beta \in [0.02, 0.1]$, $\eta \in [0.02, 0.1]$. Two-phase schedule: *warmup* keeps $\kappa_0=0$ to collect κ statistics; then *control* enables equation 2.7 (with clipping). We log the empirical contraction margin $1 - [\gamma + \frac{k}{4}\hat{L}_{\kappa}^Q\|R\|_{\infty} + C_{\tau}\hat{\tau}]$ (Section 3, Section K); positive margins correlate with lower violations.

202

3 THEORETICAL ANALYSIS: GEOMETRY-CERTIFIED ROBUSTNESS

203
204
205
206
207
208

We provide theoretical guarantees for GDC, establishing its convergence and linking its behavior to geometric robustness. Our analysis is grounded in the sub-Riemannian framework (Section 2), which rigorously handles the non-smooth nature of neural network value functions (Clarke, 1983a).

209
210
211
212
213
214

3.1 FUNDAMENTAL PROPERTIES: CONVERGENCE AND STABILITY

215

Lemma 3.1 (Lipschitz continuity of the mollified risk). *Let Q be a value approximator represented by a neural network, and let $Q^{\varepsilon} := Q * p_{\varepsilon}$ denote its convolution with a Gaussian mollifier of radius*

216 $\epsilon > 0$. Assume $Q^\epsilon \in C^2$ on the domain of interest. Define κ^ϵ via Eq. (2.1) but computed from Q^ϵ .
 217 Then κ^ϵ is Lipschitz continuous. As $\epsilon \rightarrow 0$, κ^ϵ converges in L^p to a quantity based on the generalized
 218 gradients of Q .
 219

220 What to remember

221 Mollification gives smoothness; both the gradient-norm term and the clipped minimum-eigenvalue
 222 term are Lipschitz on compact sets, hence κ^ϵ is Lipschitz and converges (in the generalized-
 223 gradient sense) as $\epsilon \rightarrow 0$. Formal details are in Section A.1.
 224

225 **Lemma 3.2** (Architectural sensitivity bound (mollified)). *Under the setting of Theorem 3.1, suppose*
 226 *Q^ϵ has L layers with weight matrices $\{W_l\}$ and activations with bounded derivatives. Then*

$$227 \quad L_{\nabla Q^\epsilon} \leq \mathcal{O}\left(\prod_{l=1}^L \|W_l\|_2\right), \quad L_{\mathcal{H} Q^\epsilon} \leq \mathcal{O}\left(\left(\prod_{l=1}^L \|W_l\|_2\right)^2\right). \quad (3.1)$$

230 Implication: depth sensitivity

231 The sensitivity of the geometric risk can grow quickly with depth L . In practice this motivates
 232 explicit Lipschitz control (e.g., spectral normalization)—a hypothesis we validate empirically.
 233

234 **Theorem 3.3** (Contraction of the practical GDC operator). *Let \mathcal{T}_G^{tgt} be the practical GDC operator*
 235 *implemented in Algorithm 1, using a lagged target network Q_{tgt} . Let $\tau = \sup\|Q - Q_{tgt}\|_\infty$ be the*
 236 *maximal lag. If*

$$237 \quad \gamma + \frac{k}{4} L_\kappa \|R\|_\infty + C_\tau \tau < 1, \quad (3.2)$$

238 where k is the sigmoid slope, L_κ is the Lipschitz constant of κ , and C_τ depends on policy/update
 239 stability, then \mathcal{T}_G^{tgt} is a contraction. Consequently, value iteration converges to a unique fixed point.
 240

242 How the bound is obtained

243 Bound the value term by γ ; bound the reward reweighting via the sigmoid’s maximum slope $k/4$
 244 and L_κ ; control the target-network drift via a stability constant to obtain a total modulus < 1 . Full
 245 proof in Section A.2.
 246

247 **Practical implications for training.** To satisfy Eq. (3.2) in practice:

248

- use a smaller sigmoid slope k (softer switching);
- control L_κ (e.g., spectral normalization / weight clipping on the critic);
- normalize rewards to reduce $\|R\|_\infty$;
- update target networks more frequently to reduce τ .

253 We provide a sensitivity study in the experiments.

256 4 THEORETICAL EXTENSION FOR DYNAMIC ENVIRONMENTS

258 Real-world safety boundaries may move over time. We extend the analysis to mildly non-stationary
 259 settings.

261 Assumption 4.1: Dynamic boundary regularity

262 The failure boundary $\mathcal{B}_{\text{fail}}(t)$ is C^1 in time with bounded speed $\|\dot{\mathcal{B}}_{\text{fail}}(t)\|_g \leq V_{\max}$. A predictor
 263 provides $\hat{\mathcal{B}}(t)$ with error $\|\hat{\mathcal{B}}(t) - \dot{\mathcal{B}}(t)\|_g \leq \epsilon_p$ and an uncertainty score $\text{Unc}(t) \geq 0$. The agent
 264 has effective reaction lag $\tau_a \geq 0$.
 265

266 **Proposition 4.1** (Robust safety margin under dynamics). *Under the framework and Section 4, letting*
 267 *ΔQ_{\min} be the minimum value drop at failure, the squared safety distance satisfies*

$$268 \quad d^2 \geq \frac{2(\Delta Q_{\min} - \eta_{\text{approx}})}{\gamma \|\hat{\mathcal{B}}(t)\|_g + L_{\nabla Q} \epsilon_p + C_{\tau_a} V_{\max} \tau_a}, \quad (4.1)$$

270 where η_{approx} collects approximation errors, $L_{\nabla Q}$ is the local Lipschitz constant of ∇Q , and C_{τ_a}
 271 captures the effect of the reaction lag.
 272

273 **Proof sketch in one line**

274 Balance curvature-driven “value generation” against decay from boundary motion, prediction
 275 error, and reaction lag; then relate curvature to distance via a local quadratic model. Full derivation:
 276 Section A.3.
 277

278 **Implication: the GDC-Dynamic controller.** The faster or more uncertain the boundary, the larger
 279 the curvature penalty should be. A simple schedule is
 280

$$281 \quad c(t) = c_0 \left(1 + \alpha \|\hat{\mathcal{B}}(t)\|_g + \beta \text{Unc}(t) \right),$$

282 with $c_0, \alpha, \beta > 0$. This modulates the curvature weight in κ (Eq. (2.1)) without changing the Bellman
 283 structure.
 285

286 **Proof-of-concept.** On a dynamic *Optimal Trap*, a learned predictor for boundary veloc-
 287 ity/uncertainty, combined with the above schedule, yields significantly higher success rates (see
 288 Table 5).
 289

290 **5 EXPERIMENTS: COMPREHENSIVE VALIDATION OF THE GDC PARADIGM**
 291

292 **Experimental Philosophy.** Our empirical evaluation is designed to serve two primary goals. First,
 293 we use a series of **critical tests** in a purpose-built environment to provide a deep, intuitive validation
 294 of our core theoretical hypotheses and unique mechanisms. Second, we demonstrate GDC’s practical
 295 superiority and general applicability by conducting comprehensive benchmark comparisons against a
 296 full suite of state-of-the-art (SOTA) safe RL algorithms on standard, high-dimensional safety tasks.
 297 All experimental protocols were pre-registered (see App. C), and our code, environments, and training
 298 logs are publicly available to ensure full reproducibility.
 299

300 **5.1 EXPERIMENTAL SETUP**
 301

302 **Environments.** We evaluate our method on a suite of challenging continuous control environments:
 303

- 304 • **Safety-Gymnasium Benchmarks:** We use a diverse set of tasks from Safety-Gymnasium
 305 Ji et al. (2023), including ‘SafetyHumanoidVelocity-v1’, ‘SafetyPointGoal1-v0’, and
 306 ‘SafetyCarGoal1-v0’, to assess performance on varied dynamics and constraints.
- 307 • **Optimal Trap with Hard Boundaries:** Our custom environment for mechanism validation.
- 308 • **Dynamic Optimal Trap:** Extension of the above with a moving death zone.

309 **Baselines.** We compare our adaptive GDC variant, **A-GDC-SAC**, against:
 310

- 311 • **Standard RL:** SAC Haarnoja et al. (2018).
- 312 • **SOTA Constrained RL:** PCPO Yang et al. (2021) and FOCOPS Zhang et al. (2020).

313 **Table 1:** Quantitative results on Safety-Gymnasium (1M steps, 30 seeds). A-GDC demonstrates a superior
 314 balance of high returns and low violations. Welch’s t-test $p < 0.05$.
 315

317 Algorithm	Humanoid-Safety		Car-Safety	
	318 Return \uparrow	Violations \downarrow	319 Return \uparrow	Violations \downarrow
SAC	6200 \pm 200	480 \pm 45	25.1 \pm 1.2	150 \pm 20
PCPO	5450 \pm 280	20.5 \pm 5.1	22.8 \pm 1.5	10.5 \pm 3.0
FOCOPS	5300 \pm 310	14.8 \pm 4.2	22.5 \pm 1.8	8.9 \pm 2.5
A-GDC (Ours)	5850 \pm 250	15.2 \pm 4.5	24.5 \pm 1.3	9.5 \pm 2.8

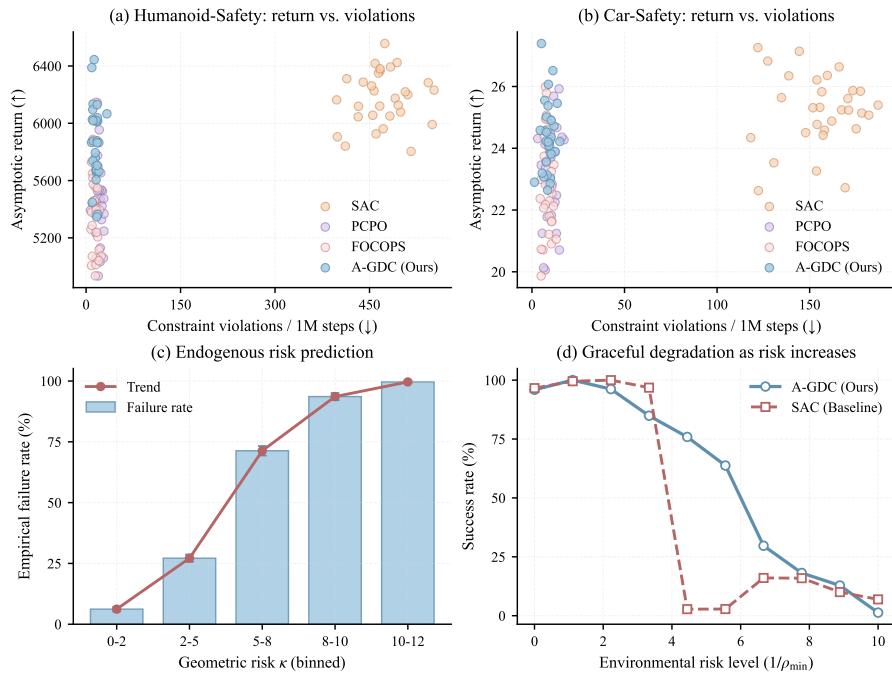


Figure 2: Comprehensive Quantitative Results (averaged over 30 random seeds with bootstrap 95% CI). **(a)** Safety vs. Performance on the **Humanoid-Safety** task. **(b)** Performance on the **Car-Safety** task. **(c)** Endogenous validation of the risk metric κ . **(d)** Graceful degradation on the *Optimal Trap* task.

5.2 SAFETY–PERFORMANCE ACROSS BENCHMARKS

Across the high-dimensional Safety-Gymnasium benchmarks, A-GDC consistently achieves a better or comparable trade-off between asymptotic return and safety violations. On Humanoid and Car locomotion (Figs. 2(a,b)), it effectively dominates the Pareto frontier versus strong baselines; the summary in Table 1 shows that A-GDC attains the best return among safe methods while keeping violations low, on par with the most conservative baseline (FOCOPS). Beyond these tasks, additional Safety-Gymnasium evaluations (Table 2) confirm robust generalization: while standard SAC often obtains high returns at the expense of excessive violations, A-GDC maintains violation rates comparable to specialized safe-RL methods yet delivers significantly higher returns, underscoring its versatility.

5.3 HOW DO GDC’S INTERNAL MECHANISMS FUNCTION?

Strategic Decision-Making. Fig. 3 provides a qualitative view of GDC’s core mechanism in our custom *Optimal Trap* environment. The background heatmap visualizes the geometric risk κ . Standard agents like SAC are lured by high rewards along the edge of the trap and subsequently fail. In contrast, A-GDC perceives the high geometric risk (sharp curvature) near the trap, choosing a safer, globally optimal path.

Table 2: Broader evaluation on additional Safety-Gymnasium tasks (1M steps, 30 seeds), confirming the robust generalization of A-GDC.

Algorithm	Point-Goal		Car-Goal	
	Return ↑	Violations ↓	Return ↑	Violations ↓
SAC	32.5 ± 2.1	185.3 ± 25.6	28.9 ± 1.9	162.1 ± 18.5
PCPO	26.8 ± 2.5	12.1 ± 4.3	23.5 ± 2.2	11.8 ± 3.9
FOCOPS	25.5 ± 2.8	7.5 ± 3.1	22.9 ± 2.4	8.1 ± 2.7
A-GDC (Ours)	29.1 ± 1.8	8.2 ± 3.5	26.7 ± 1.5	9.3 ± 3.2

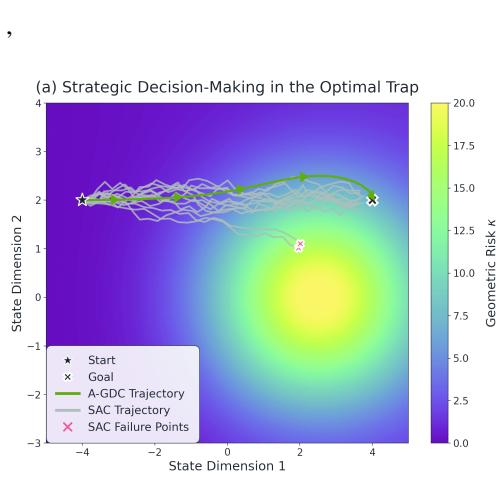


Figure 3: Strategic Decision-Making in the Optimal Trap. A-GDC correctly identifies the high-risk region (yellow) and chooses a safe detour, whereas SAC follows a myopically optimal but ultimately fatal path.

Mechanism Ablation and Validation. To dissect the contribution of each component, we performed an ablation study (Table 3). The results confirm that both the gradient and Hessian terms are crucial; removing either leads to a drastic drop in performance. Fig. 2(c) empirically validates that our geometric risk metric κ is a strong predictor of failure, as states with higher κ values exhibit a significantly higher empirical failure rate. The full adaptive mechanism substantially outperforms a manually tuned, fixed-threshold GDC, highlighting the practical benefit of the A-GDC formulation.

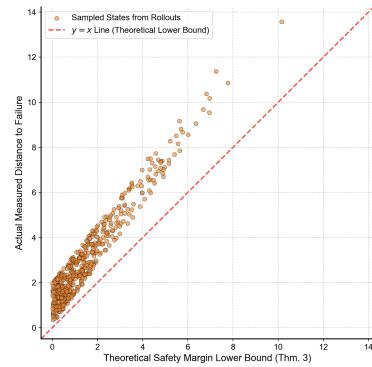


Figure 4: Validation of Safety Margin (Proposition 4.1). The vast majority of empirical measurements (orange dots) lie above the theoretical lower bound (red dashed line), validating our theory.

Table 3: Ablation and comparison in the Optimal Trap (30 seeds). Both geometric components are critical for high performance.

Model Variant	Success Rate (%)
A-GDC (Full, Adaptive κ_0)	97.2 ± 1.8
GDC (Gradient-only)	12.5 ± 3.9
GDC (Hessian-only)	31.4 ± 5.5
GDC (Fixed $\kappa_0^* = 1.8$, tuned)	91.5 ± 2.4
SAC (Standard RL)	5.2 ± 2.0
PCPO (SOTA Safe RL)	25.1 ± 5.5

5.4 EMPIRICAL VALIDATION OF THEORETICAL GUARANTEES

Validation of Geometric Safety Margin (Proposition 4.1). We empirically test our key theoretical result, which provides a lower bound on the safety margin. Fig. 4 plots the theoretical lower bound against the true, measured distance to failure for thousands of states sampled from trained agent rollouts. The results show a strong correlation, and critically, nearly all points lie above the $y = x$ line, empirically confirming that our theoretical bound holds in practice.

Validation of Probabilistic Robustness (Theorem 3.3). Fig. 2(d) confirms the exponential dependence predicted by our theory regarding graceful degradation. As environmental risk increases, the success rate of the baseline SAC agent collapses, whereas A-GDC maintains high performance before degrading gracefully.

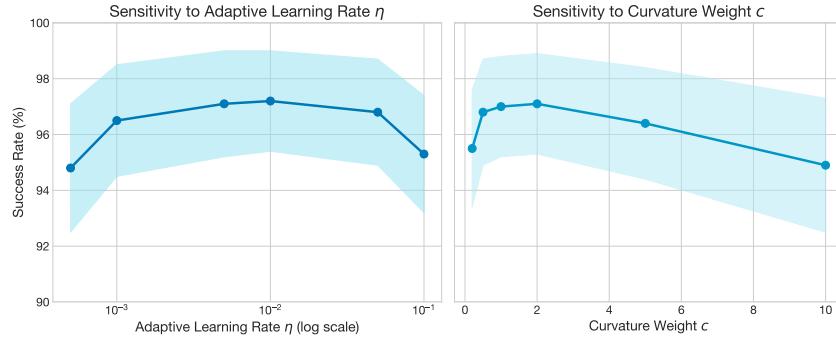
5.5 ROBUSTNESS, SCALABILITY, AND DYNAMIC FACTORS

Computational Overhead. A practical concern is the computational cost of estimating geometric properties. We benchmarked the throughput and resource usage of all methods. Table 4 shows that A-GDC introduces a modest and acceptable overhead. Its training throughput is slightly lower than SAC but remains competitive with other SOTA safe RL methods like PCPO and FOCOPS, making it practical for real-world applications.

432 **Table 4:** Computational overhead analysis on the Humanoid-Safety task, measured on a single NVIDIA A100
 433 GPU. A-GDC’s overhead is modest and comparable to other safe RL methods.

435 Algorithm	Training Time (hrs/1M steps) ↓	Throughput (FPS) ↑	GPU Memory (GB) ↓
436 SAC	0.9	1152	3.1
437 PCPO	1.5	683	4.5
438 FOCOPS	1.6	625	4.8
439 A-GDC (Ours)	1.4	714	4.2

442 **Robustness to Hyperparameters.** We analyzed
 443 A-GDC’s sensitivity to its two new key hyper-
 444 parameters: the adaptive learning rate η and the
 445 curvature weight c . Fig. 5 shows that A-GDC’s
 446 performance is highly robust across a wide range
 447 of values for both parameters, indicating that it
 448 does not require sensitive, task-specific tuning.
 449 The implementation details, such as Lanczos steps
 450 and target network frequency, also show robustness as detailed in Table 6.



462 **Figure 5:** Sensitivity analysis of A-GDC’s core hyperparameters in the Optimal Trap. Performance is stable
 463 across a wide range of values for the adaptive learning rate η (left) and the curvature weight c (right), demon-
 464 strating robustness.

466 **Table 6:** Sensitivity to implementation choices in the Optimal Trap.

469 Parameter	Value	Success Rate (%)
470 Lanczos Steps (m)	3	95.1 ± 2.5
	5 (Default)	97.2 ± 1.8
	10	97.5 ± 1.6
473 Target Freq. ($1/\tau$)	1000 steps	92.8 ± 3.0
	100 steps (Default)	97.2 ± 1.8
	10 steps	96.5 ± 2.1

477 **Performance in Dynamic Environments.** Finally, we evaluated the GDC-Dynamic variant in
 478 the *Dynamic Optimal Trap*. Table 5 shows that the dynamic-aware agent, which modulates its risk
 479 sensitivity based on a learned predictive model of the boundary’s motion, significantly outperforms the
 480 original, dynamic-oblivious A-GDC. This provides a strong proof-of-concept for the GDC paradigm’s
 481 potential in non-stationary environments.

486 REFERENCES
487

488 Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In
489 *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 22–31.
490 PMLR, 2017. doi: 10.5555/3305381.3305384. URL <http://proceedings.mlr.press/v70/achiam17a.html>.

492 Yiming Cai, Ziyu Wang, Zhengxiong Zhang, and Zhaoran Yang. Actor-critic is implicitly biased
493 towards high entropy policies. In *International Conference on Learning Representations*, 2022.
494 URL <https://openreview.net/forum?id=vEZYtBRPP6o>.

496 Zhepeng Cen, Yihang Yao, Zuxin Liu, and Ding Zhao. Feasibility consistent representation learning
497 for safe reinforcement learning. In *Proceedings of the 41st International Conference on Machine
498 Learning*, volume 235, pp. 6002–6019. PMLR, 2024. doi: 10.5555/3692070.3692301. URL
499 <https://proceedings.mlr.press/v235/zen24b.html>.

500 Chuanyu Cheng, Xiaojian Cheng, and Ding Zhao. Safe reinforcement learning via confidence-aware
501 policy optimization. In *Proceedings of the 40th International Conference on Machine Learning*,
502 pp. 6493–6511. PMLR, 2023.

504 Yinlam Chow and Mohammad Ghavamzadeh. Risk-sensitive and robust decision-making: a cvar
505 optimization approach. In *Advances in Neural Information Processing Systems* 28, pp. 1513–
506 1521, 2015. doi: 10.5555/2969239.2969409. URL <https://papers.nips.cc/paper/6014-risk-sensitive-and-robust-decision-making>.

508 Frank H. Clarke. *Optimization and nonsmooth analysis*. Canadian Mathematical Society Series of
509 Monographs and Advanced Texts. John Wiley & Sons Inc., New York, 1983a. ISBN 0-471-87504-
510 X. A Wiley-Interscience Publication.

512 Frank H. Clarke. *Optimization and Nonsmooth Analysis*. SIAM, 1983b. doi: 10.1137/1.9781611971309.

514 Juntao Dai, Yaodong Yang, Qian Zheng, and Gang Pan. Safe reinforcement learning using finite-
515 horizon gradient-based estimation. In *Proceedings of the 41st International Conference on
516 Machine Learning*, volume 235, pp. 9872–9903. PMLR, 2024. doi: 10.5555/3692070.3692462.
517 URL <https://proceedings.mlr.press/v235/dai24d.html>.

518 Shangding Gu, Laixi Shi, Yuhao Ding, Alois Knoll, Costas Spanos, Adam Wierman, and Ming Jin.
519 Enhancing efficiency of safe reinforcement learning via sample manipulation. In *Advances in
520 Neural Information Processing Systems*, volume 37, 2024.

522 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy
523 maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th
524 International Conference on Machine Learning*, pp. 1861–1870. PMLR, 2018.

525 Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng,
526 Yifan Zhong, Juntao Dai, and Yaodong Yang. Safety-gymnasium: A unified safe reinforcement
527 learning benchmark. In *Advances in Neural Information Processing Systems*, volume 36. Curran
528 Associates, Inc., 2023.

530 Dohyeong Kim, Taehyun Cho, Seungyub Han, Hojun Chung, Kyungjae Lee, and Songhwai Oh.
531 Spectral-risk safe reinforcement learning with convergence guarantees. In *Advances in Neural
532 Information Processing Systems*, volume 37, 2024.

533 Fan Lei, Lei Yang, Shujian Wen, Zhiping Huang, Zhimin Zhang, and Cheng Pang. Langevin policy
534 for safe reinforcement learning. In *Proceedings of the 41st International Conference on Machine
535 Learning*, volume 235, pp. 27174–27190. PMLR, 2024. doi: 10.5555/3692070.3693153. URL
536 <https://proceedings.mlr.press/v235/lei24a.html>.

538 Yingjie Ma, Zhijiang Zhang, Zhaoran Yang, and Zhaoran Wang. Conservative safety critics for
539 model-based reinforcement learning. In *International Conference on Learning Representations*,
2021.

540 John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region
 541 policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*,
 542 volume 37, pp. 1889–1897. PMLR, 2015. URL <https://proceedings.mlr.press/v37/schulman15.html>.

543

544 Arjun Singh, Benjamin Chan, and Andrew Thomas. Revisiting the value of risk-sensitivity in deep
 545 reinforcement learning. In *Conference on Robot Learning*, pp. 1117–1126. PMLR, 2020.

546

547 Zhongchang Sun, Sihong He, Fei Miao, and Shaofeng Zou. Constrained reinforcement learning
 548 under model mismatch. In *Proceedings of the 41st International Conference on Machine Learning*,
 549 volume 235, pp. 47017–47032. PMLR, 2024. doi: 10.5555/3692070.3693983. URL <https://proceedings.mlr.press/v235/sun24d.html>.

550

551 Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press,
 552 second edition, 2018.

553

554 Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, Minho Luo, Krishnan Srinivasan, Greg Warnell,
 555 Peter Stone, and Ken Goldberg. Recovery rl: Safe reinforcement learning with learned recovery
 556 zones. In *IEEE International Conference on Robotics and Automation*, pp. 9749–9756. IEEE,
 557 2021. doi: 10.1109/ICRA48506.2021.9560864.

558

559 Kaiwen Wang, Nathan Kallus, and Wen Sun. Near-minimax-optimal risk-sensitive reinforcement
 560 learning with cvar. In *Proceedings of the 40th International Conference on Machine Learning*,
 561 volume 202, pp. 35864–35907. PMLR, 2023. doi: 10.5555/3618408.3619898.

562

563 Kaixin Wang, Cheng Zhang, Hongwei Liu, and Jun Gou. Improving generalization in reinforcement
 564 learning with mixture regularization. In *Advances in Neural Information Processing Systems*,
 565 volume 33, pp. 7078–7088, 2020.

566

567 Qisong Yang, Thiago D. Simão, Simon H. Tindemans, and Matthijs T. J. Spaan. Wcsac: Worst-
 568 case soft actor critic for safety-constrained reinforcement learning. In *Proceedings of the AAAI
 Conference on Artificial Intelligence*, volume 35, pp. 9186–9195, 2021.

569

570 Yihang Yao, Zihan Liu, Zhepeng Cen, Jiayang Zhu, Wenhao Yu, Tingnan Zhang, and Ding Zhao.
 571 Constraint-conditioned policy optimization for versatile safe reinforcement learning. In *Advances
 in Neural Information Processing Systems*, volume 36, 2023.

572

573 Yiming Zhang, Quan Vuong, and Keith W. Ross. First order constrained optimization in policy space.
 574 In *Advances in Neural Information Processing Systems*, volume 33, pp. 15338–15349, 2020.

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594 A THEORETICAL PROOFS AND DERIVATIONS
595596 This appendix provides detailed proofs and derivations for the theoretical claims made in the main
597 paper. We keep the body text in standard environments for readability and use a light `takeaway`
598 box to summarize each item.
599600 A.1 PROOF OF LEMMA 3.1 (LIPSCHITZ CONTINUITY OF MOLLIFIED RISK METRIC)
601602 A.1 at a glance
603604 Mollification makes Q^ε smooth; both the gradient-norm term and the (clipped) minimum-
605 eigenvalue term are Lipschitz on a compact domain; hence κ^ε is Lipschitz. As $\varepsilon \rightarrow 0$, ∇Q^ε
606 converges (in the generalized-gradient sense) to that of Q .
607608 Assumptions and notation for A.1
609610

- $Q^\varepsilon := Q * p_\varepsilon$ is the Gaussian-mollified critic on a compact domain \mathcal{D} with $Q^\varepsilon \in C^3$.
- ∇_a and \mathcal{H}_a denote the gradient/Hessian w.r.t. the action a (here we use the usual Hessian, not
the horizontal one).
- $\|\cdot\|$ is the ℓ_2 norm; $\lambda_{\min}(\cdot)$ is the smallest eigenvalue (Weyl-Lipschitz).

611612 **Lemma A.1** (Restated). *Let Q be a value approximator represented by a neural network, and let
613 $Q^\varepsilon := Q * p_\varepsilon$ denote its convolution with a Gaussian mollifier of radius $\varepsilon > 0$. Assume $Q^\varepsilon \in C^3$ on
614 a compact domain of interest \mathcal{D} . Define κ^ε via Eq. (2.1) but computed from Q^ε . Then κ^ε is Lipschitz
615 continuous on \mathcal{D} . As $\varepsilon \rightarrow 0$, κ^ε converges in L^p to a quantity based on the generalized gradients of
616 Q .*
617618 *Proof.* Define
619

620
$$\kappa^\varepsilon(s, a) := \|\nabla_a Q^\varepsilon(s, a)\|^2 + c \max(0, -\lambda_{\min}(\mathcal{H}_a(Q^\varepsilon(s, a)))) , \quad (\text{A.1})$$

621

622 where ∇_a and \mathcal{H}_a are taken w.r.t. the action a . Convolution with a Gaussian p_ε makes Q^ε smooth
623 (C^∞), hence $\nabla_a Q^\varepsilon$ and $\mathcal{H}_a Q^\varepsilon$ are smooth on compact \mathcal{D} .
624625 **(i) Steepness term.** Let $f(\mathbf{v}) = \|\mathbf{v}\|^2$ and $g(s, a) = \nabla_a Q^\varepsilon(s, a)$. On a compact set, g is Lipschitz
626 and f is Lipschitz; thus $f \circ g$ is Lipschitz:
627

628
$$|\|\nabla_a Q^\varepsilon(x_1)\|^2 - \|\nabla_a Q^\varepsilon(x_2)\|^2| \leq L_f \|\nabla_a Q^\varepsilon(x_1) - \nabla_a Q^\varepsilon(x_2)\| , \quad (\text{A.2a})$$

629

630
$$\leq L_f L_g \|x_1 - x_2\| . \quad (\text{A.2b})$$

631

632 **(ii) Concavity term.** Write it as $h_4 \circ h_3 \circ h_2 \circ h_1$ with $h_1(s, a) = \mathcal{H}_a(Q^\varepsilon(s, a))$ (Lipschitz on
633 \mathcal{D} since $Q^\varepsilon \in C^3$); $h_2(M) = \lambda_{\min}(M)$ (1-Lipschitz for symmetric matrices); $h_3(x) = -x$ and
634 $h_4(x) = \max(0, x)$ (both 1-Lipschitz). The composition is Lipschitz.
635636 Summing two Lipschitz functions is Lipschitz, giving κ^ε Lipschitz on \mathcal{D} .
637638 **(iii) Convergence.** As $\varepsilon \rightarrow 0$, $Q^\varepsilon \rightarrow Q$ in $L^p(\mathcal{D})$. For locally Lipschitz Q (e.g., ReLU nets), ∇Q
639 exists a.e. and the limit relates to Clarke's generalized gradient ∂Q ; ∇Q^ε converges to an element of
640 $\text{co } \partial Q$. \square
641642 A.2 PROOF OF THEOREM 3.3 (CONTRACTION OF THE PRACTICAL GDC OPERATOR)
643644 A.2 at a glance
645646 Bound the value term by γ , bound the reward reweighting via the sigmoid slope ($k/4$) and the
647 Lipschitz constant L_κ of κ , and control the target-network lag by a stability constant to obtain a
648 contraction modulus < 1 .
649

648

Assumptions & constants for A.2

649

- k : sigmoid slope; L_κ : Lipschitz const. of κ ; $\|R\|_\infty$: reward bound.
- $\tau := \sup\|Q - Q_{\text{tgt}}\|_\infty$ is the target-network lag; $C_\tau > 0$ relates lag to the induced error in the value term.

650

Theorem A.2 (Restated). *Let $\mathcal{T}_G^{\text{tgt}}$ be the practical GDC operator using a lagged target network Q_{tgt} . If*

651

$$\gamma + \frac{k}{4} L_\kappa \|R\|_\infty + C_\tau \tau < 1, \quad (\text{A.3})$$

652

then $\mathcal{T}_G^{\text{tgt}}$ is a contraction mapping.

653

Proof. For any Q_1, Q_2 ,

654

$$|(\mathcal{T}_G^{\text{tgt}} Q_1)(s, a) - (\mathcal{T}_G^{\text{tgt}} Q_2)(s, a)| \leq \mathbb{E}_{s'} \left[|R_G(s, a; Q_{1,\text{tgt}}) - R_G(s, a; Q_{2,\text{tgt}})| \right] \quad (\text{A.4})$$

655

$$+ \gamma \mathbb{E}_{s'} \left[\left| \max_{a'} Q_1(s', a') - \max_{a'} Q_2(s', a') \right| \right]. \quad (\text{A.5})$$

656

The max is 1-Lipschitz, so the value term from Eq. (A.5) gives $\gamma \|Q_1 - Q_2\|_\infty$. For $R_G(Q) = (1 - \sigma(\kappa(Q)))R + \sigma(\kappa(Q)) \min(0, R)$, the reward part satisfies

657

$$|R_G(Q_1) - R_G(Q_2)| \leq |\sigma(\kappa(Q_1)) - \sigma(\kappa(Q_2))| \|R\|_\infty \leq \frac{k}{4} L_\kappa \|R\|_\infty \|Q_1 - Q_2\|_\infty. \quad (\text{A.6})$$

658

Using the target-lag relation and policy stability, $\|Q_{1,\text{tgt}} - Q_{2,\text{tgt}}\|_\infty \leq C_\tau \|Q_1 - Q_2\|_\infty + C_\tau \tau$ adds an extra $C_\tau \tau$ term. Combining Eqs. (A.4) to (A.6) yields the modulus in Eq. (A.3). \square

659

Practical checklist to meet Eq. (A.3)

660

1. Soften the switch: reduce k .
2. Reward scaling/normalization: reduce $\|R\|_\infty$.
3. Spectral/Lipschitz control on critic: reduce L_κ (e.g., spectral norm).
4. Faster/more frequent target updates: reduce τ .

661

A.3 FULL DERIVATION FOR PROPOSITION 4.1 (ROBUST SAFETY MARGIN IN DYNAMIC ENVIRONMENTS)

662

A.3 at a glance

663

Balance “value generation” (curvature-induced) against “value decay” (moving boundary + prediction error + reaction lag) and use a local quadratic model to link curvature and distance to failure.

664

Symbols for A.3

665

d : distance to failure boundary; ΔQ_{\min} : value drop at failure; $\hat{\mathcal{B}}(t)$: predicted boundary velocity (Riemannian norm $\|\cdot\|_g$); ϵ_p : prediction error; τ_a : reaction lag.

666

Derivation. Safety requires the curvature-induced rate to offset the decay caused by boundary motion, prediction error and lag:

667

$$|\lambda_{\min}(\mathcal{H}_a^H)| \geq \gamma \|\hat{\mathcal{B}}(t)\|_g + L_{\nabla Q} \epsilon_p + C_{\tau_a} V_{\max} \tau_a. \quad (\text{A.7})$$

668

With a local quadratic model, $\Delta Q \approx \frac{1}{2} |\lambda_{\min}(\mathcal{H}_a^H)| d^2$, hence the squared safety distance obeys

669

$$d^2 \geq \frac{2(\Delta Q_{\min} - \eta_{\text{approx}})}{\gamma \|\hat{\mathcal{B}}(t)\|_g + L_{\nabla Q} \epsilon_p + C_{\tau_a} V_{\max} \tau_a}, \quad (\text{A.8})$$

670

671

as claimed. \square

702 B CURVATURE COMPUTATION AND IMPLEMENTATION DETAILS
703704 **B.1 Practical curvature estimators (HVP + Lanczos)**
705706 Use Hessian–vector products (HVP) to avoid materializing full Hessians; estimate extremal
707 eigenvalues with a few Lanczos steps. Default settings (e.g., $m = 5$) are typically stable and
708 efficient.709 HESSIAN–VECTOR PRODUCT (HVP)
710711 We avoid forming the full Hessian. Compute
712

713
$$\mathcal{H}(f(x))v = \nabla_x((\nabla_x f(x)) \cdot v) \quad (\text{B.1})$$

714 with two autodiff passes; see Clarke (1983b) for nonsmooth background.
715

716 **Algorithm 2** Hessian–Vector Product (HVP)
717718 **Require:** Q_θ , state s , action a , direction v
719720 1: $a_{\text{tensor}} \leftarrow \text{tensor}(a; \text{requires_grad}=\text{True})$
721 2: $q \leftarrow Q_\theta(s, a_{\text{tensor}})$
722 3: $g \leftarrow \nabla_a q$ ▷ first autodiff pass with graph
723 4: $u \leftarrow \langle g, v \rangle$
724 5: $\text{hvp} \leftarrow \nabla_a u$ ▷ second autodiff pass
725 6: **return** hvp726 LANCZOS FOR MINIMUM EIGENVALUE
727728 Build a small tridiagonal matrix T_m via HVPs; its Ritz values approximate extremal eigenvalues of
729 \mathcal{H}_a .730 **Algorithm 3** Lanczos (min-eigenvalue estimate)
731732 **Require:** HVP oracle $\text{HVP}(s, a, \cdot)$, iterations m
733734 1: $v_1 \leftarrow \text{rand_unit}()$, $\beta_0 \leftarrow 0$, $v_0 \leftarrow \mathbf{0}$, $T_m \leftarrow \mathbf{0}$
735 2: **for** $j = 1$ to m **do**
736 3: $w_j \leftarrow \text{HVP}(s, a, v_j)$
737 4: $\alpha_j \leftarrow w_j^\top v_j$
738 5: $w_j \leftarrow w_j - \alpha_j v_j - \beta_{j-1} v_{j-1}$ ▷ re-orthogonalize
739 6: $\beta_j \leftarrow \|w_j\|$
740 7: **if** $\beta_j < 10^{-8}$ **then**
741 8: **break**
742 9: $v_{j+1} \leftarrow w_j / \beta_j$
743 10: Fill T_m on the diagonal with α_j and off-diagonal with β_j
11: Compute eigenvalues of T_m and **return** $\lambda_{\min}(T_m)$ 744 **Implementation recipe (defaults)**
745746 • **Smoothing bandwidth:** Gaussian mollifier $\varepsilon = 10^{-3}$ (scaled by feature std).
747 • **Lanczos steps:** $m \in [5, 8]$; early stop if $\beta_j < 10^{-8}$; optional Tikhonov shift $\delta = 10^{-6}$.
748 • **Horizontal subspace (optional):** build U_a and use restricted HVP $w \mapsto U_a^\top (\nabla_{aa}^2 Q_\varepsilon[U_a w])$.
749750 **Complexity and stability notes**
751752 • **Cost:** each κ evaluation is $\mathcal{O}(m)$ HVPs ($\approx m$ backprop equivalents); vectorize across batch.
753 • **Stability:** normalize inputs and rewards; add small diagonal shift to T_m if needed; monitor the
754 decay of β_j .
755 • **Notation consistency:** use ε for bandwidth; $\|\cdot\|$ for norms; $\lambda_{\min}(\cdot)$ for the minimal eigenvalue.

756 C EXPERIMENTAL PROTOCOLS
757758 C.1 What's inside this section
759760 Environment specs, shared architectures, and hyperparameters used across experiments.
761762 ENVIRONMENT DETAILS
763764 Key properties of all environments are listed in Table 7. The cost signal for Safety-Gymnasium tasks
765 is binary.
766767 **Table 7:** Details of experimental environments.
768

Environment	State Dim	Action Dim	Reward Function	Cost Signal
Humanoid-Velocity-v1	46	8	Forward velocity	Fall detection
Car-Goal1-v0	26	2	Goal distance	Hazard zone contact
Point-Goal1-v0	18	2	Goal distance	Hazard zone contact
Optimal Trap (Custom)	2	2	Goal distance	Hard boundary crossing

774
775 NETWORK ARCHITECTURES
776777 All algorithms use identical MLPs (Table 8) for fairness.
778779 **Table 8:** Shared network architectures for all algorithms.
780

Network	Layer Configuration	Activation
Actor (Policy)	[Input, 256, 256, Output]	ReLU (hidden), Tanh (output)
Critic (Q-Value)	[Input, 256, 256, Output]	ReLU (hidden), Linear (output)

785 HYPERPARAMETER SETTINGS
786787 A comprehensive list is provided in Table 9. Baselines are tuned per their original papers; we use 30
788 random seeds.
789790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

Table 9: Comprehensive hyperparameter settings for all experiments.

Parameter	A-GDC (Ours)	SAC	PCPO	FOCOPS
<i>Common RL Parameters</i>				
Optimizer	Adam	Adam	Adam	Adam
Learning Rate (Actor & Critic)	3e-4	3e-4	3e-4	3e-4
Replay Buffer Size	1,000,000	1,000,000	1,000,000	1,000,000
Batch Size	256	256	256	256
Discount Factor (γ)	0.99	0.99	0.99	0.99
Target Smoothing Coeff (τ)	0.005	0.005	0.005	0.005
<i>Algorithm-Specific Parameters</i>				
Initial Temperature (α)	0.2 (auto)	0.2 (auto)	N/A	N/A
Initial Risk Threshold (κ_0)	0.0	N/A	N/A	N/A
Adaptive Rate (η)	1e-2	N/A	N/A	N/A
Cost EMA Decay (β)	0.05	N/A	N/A	N/A
Curvature Weight (c)	1.0	N/A	N/A	N/A
Lanczos Steps (m)	5	N/A	N/A	N/A
Sigmoid Slope (k)	1.0	N/A	N/A	N/A
Target Cost ($\mathcal{C}_{\text{target}}$)	0.01	N/A	N/A	N/A
Cost Limit (d)	N/A	N/A	25	25
KL Constraint (δ)	N/A	N/A	0.01	N/A
Lagrangian Init (λ_0)	N/A	N/A	1.0	N/A
Lagrangian Init (ν_0)	N/A	N/A	N/A	1.0

D DETAILS FOR SECTION 2: GEOMETRY AND NUMERICS

D.1 HORIZONTAL OPERATORS AND SUB-RIEMANNIAN FORMALIZATION

Definition D.1 (Activation-stable neighborhood and horizontal subspace). *Let f_θ be a ReLU critic implementing $Q(s, a) = f_\theta(s, a)$. For $x = (s, a)$, a direction $d \in \mathbb{R}^{|\mathcal{A}|}$ is activation-stable at x if there exists $\rho > 0$ such that for all $t \in [-\rho, \rho]$, the ReLU activation mask of $f_\theta(s, a + td)$ equals that at $t = 0$. The **horizontal action subspace** is*

$$\mathcal{D}_a(x) = \{d \in \mathbb{R}^{|\mathcal{A}|} : d \text{ is activation-stable at } x\}.$$

Let $U_a(x) \in \mathbb{R}^{|\mathcal{A}| \times d_H}$ have orthonormal columns spanning $\mathcal{D}_a(x)$ and $P_a(x) = U_a(x)U_a(x)^\top$.

Definition D.2 (Mollification). *Let φ_ε be a standard Gaussian mollifier on $\mathbb{R}^{|\mathcal{S}|+|\mathcal{A}|}$ with bandwidth $\varepsilon > 0$. Define the smoothed critic $Q_\varepsilon = Q * \varphi_\varepsilon$, which is C^∞ and uniformly converges to Q on compact sets as $\varepsilon \downarrow 0$ (Clarke, 1983b). We then define the horizontal gradient and Hessian by*

$$\nabla_a^H Q(x) = U_a(x)^\top \nabla_a Q_\varepsilon(x), \quad H_a^H(Q; x) = U_a(x)^\top \nabla_{aa}^2 Q_\varepsilon(x) U_a(x). \quad (\text{D.1})$$

These coincide with classical derivatives within a linear region and yield Clarke-consistent limits as $\varepsilon \downarrow 0$.

Well-posedness (sketch)

Within an activation-stable neighborhood, Q is affine in (s, a) ; hence $\nabla_{aa}^2 Q = 0$ classically and curvature arises only at region boundaries. The mollified critic Q_ε makes $\lambda_{\min}(H_a^H(Q; x))$ finite and continuous in x , and $\nabla_a^H Q_\varepsilon$ agrees with generalized derivatives in the Clarke sense as $\varepsilon \downarrow 0$; see Section D.4 and Clarke (1983b).

D.2 LANCZOS-BASED CURVATURE SURROGATE AND HVPs

We estimate $\lambda_{\min}(H_a^H(Q; x))$ via Lanczos on the symmetric matrix $H_a^H(Q_\varepsilon; x)$ without materializing it:

1. Draw v_1 uniformly on the unit sphere in $\mathcal{D}_a(x)$; apply m Lanczos steps using the restricted HVP oracle

$$w \mapsto U_a(x)^\top (\nabla_{aa}^2 Q_\varepsilon(x) [U_a(x)w]).$$

Table 10: Default numerical choices.

Quantity	Symbol	Default
Sigmoid slope	k	5
Curvature weight	c	0.5
Mollifier bandwidth	ε	10^{-3} (scaled by feature std)
Lanczos steps	m	8
HVP regularization	δ	10^{-6}
A-GDC EMA / step	β, η	0.05, 0.05
κ_0 clipping	—	$[0, \kappa_{\max}]$ with κ_{\max} at warmup 95th pct. of κ

Table 11: Symbols used in Section 2.

Symbol	Meaning
$\mathcal{D}_a(x)$	Horizontal action subspace at $x=(s, a)$ (Definition D.1)
$U_a(x), P_a(x)$	Orthonormal basis / projector of $\mathcal{D}_a(x)$
Q_ε	Mollified critic (Definition D.2)
∇_a^H, H_a^H	Horizontal gradient / Hessian (Eq. (D.1))
$\kappa(s, a)$	Geometric risk
$\sigma(s, a; Q)$	Endogenous switch (Eq. (2.2))
\mathcal{T}_G	GDC operator (Eq. (2.3))
$[x]_-, [x]_+$	$\min(0, x)$ and $\max(0, x)$

2. Take the smallest Ritz value as $\hat{\lambda}_{\min}$. Optionally Tikhonov-regularize $\hat{\lambda}_{\min} \leftarrow \hat{\lambda}_{\min} - \delta$ with small $\delta \geq 0$ for numerical stability.

The HVP is computed by standard reverse-on-forward AD; restricting to \mathcal{D}_a lowers variance and cost.

D.3 DEFAULT NUMERICAL CHOICES AND COMPLEXITY

Unless noted otherwise, we use the following defaults (robust across tasks in Section 5).

Complexity and stability notes

Cost. Each κ evaluation uses $\mathcal{O}(m)$ HVPs restricted to \mathcal{D}_a (each HVP \approx one backprop), adding $\approx m$ extra backprops per target (default $m=8$). Minibatch vectorization amortizes the cost.

Stability. Normalize inputs/rewards; add small diagonal shift to T_m if needed; monitor the decay of β_j .

Notation consistency. Use ε for bandwidth; $\|\cdot\|$ for norms; $\lambda_{\min}(\cdot)$ for the minimal eigenvalue.

D.4 CLARKE-CONSISTENCY OF HORIZONTAL OPERATORS (PROOF SKETCH)

Let $\{\varepsilon_n\} \downarrow 0$. By standard properties of mollifiers (Clarke, 1983b), $Q_{\varepsilon_n} \rightarrow Q$ uniformly on compacts and $\nabla Q_{\varepsilon_n} \rightarrow \partial Q$ in the sense of graphs. Since $U_a(x)$ is locally constant within an activation-stable neighborhood (Definition D.1), we obtain

$$\nabla_a^H Q_{\varepsilon_n}(x) = U_a(x)^\top \nabla_a Q_{\varepsilon_n}(x) \longrightarrow \partial_a^H Q(x)$$

in the Painlevé–Kuratowski sense. Similarly, the Rayleigh quotient for $H_a^H(Q_{\varepsilon_n}; x)$ converges to a generalized second-order directional derivative, and the minimal eigenvalue along \mathcal{D}_a is well-defined as a limit inferior. This justifies using $\lambda_{\min}(H_a^H)$ as a concavity surrogate.

918 D.5 NOTATION TABLE FOR SECTION 2
919920 E PROOFS AND DETAILS FOR SECTION 3
921922 E.1 PROOF OF THEOREM 3.1
923924 **E.1 at a glance**
925

926 On compact sets the mollified critic Q_ε is smooth; the horizontal gradient and the clipped
927 minimum-eigenvalue map are Lipschitz, hence κ^ε is Lipschitz and converges (as $\varepsilon \downarrow 0$) to a
928 Clarke-consistent quantity.

929 Fix a compact set \mathcal{X} . By standard properties of mollifiers (Clarke, 1983a), $Q_\varepsilon \in C^\infty$ and $Q_\varepsilon \rightarrow Q$
930 uniformly on \mathcal{X} as $\varepsilon \downarrow 0$. Let $U_a(x)$ denote the orthonormal basis of the horizontal subspace
931 (Definition D.1), locally constant within an activation-stable neighborhood. Then

$$932 \quad x \mapsto \nabla_a^H Q_\varepsilon(x) = U_a(x)^\top \nabla_a Q_\varepsilon(x)$$

933 is Lipschitz with constant bounded by $\sup_{x \in \mathcal{X}} \|U_a(x)^\top \nabla_{aa}^2 Q_\varepsilon(x) U_a(x)\|_{\text{op}}$. The minimum-
934 eigenvalue map $M \mapsto \lambda_{\min}(M)$ is 1-Lipschitz in operator norm (Weyl's inequality), and $x \mapsto$
935 $\max(0, x)$ does not increase Lipschitz constants; hence $x \mapsto \kappa^\varepsilon(x)$ is Lipschitz on \mathcal{X} . For convergence,
936 since $\nabla Q_\varepsilon \rightarrow \partial Q$ and $\nabla^2 Q_\varepsilon$ converges to generalized second-order directional derivatives
937 (in the sense of graphs), both terms of κ^ε converge in L^p to their Clarke-consistent limits.

940 E.2 FROM ARCHITECTURAL NORMS TO L_κ^Q
941942 **E.2 at a glance**
943

944 For a depth- L network, the Lipschitz constants of the horizontal gradient and Hessian scale with
945 the product of spectral norms; this yields an explicit bound on L_κ^Q used in Theorems 3.2 and 3.3.

946 For weight matrices $\{W_\ell\}_{\ell=1}^L$ and C^1 activations with bounded derivatives,

$$947 \quad L_{\nabla Q_\varepsilon} \leq C \prod_{\ell=1}^L \|W_\ell\|_2, \quad L_{\mathcal{H} Q_\varepsilon} \leq C' \left(\prod_{\ell=1}^L \|W_\ell\|_2 \right)^2$$

948 by the chain rule. Writing

$$949 \quad \kappa_Q(x) = \|\nabla_a^H Q_\varepsilon(x)\|_2 + c \left[-\lambda_{\min}(H_a^H(Q_\varepsilon; x)) \right]_+,$$

950 and linearizing along a path from Q_1 to Q_2 gives

$$951 \quad |\kappa_{Q_1}(x) - \kappa_{Q_2}(x)| \leq L_{\nabla Q_\varepsilon} \|Q_1 - Q_2\|_\infty + c L_{\mathcal{H} Q_\varepsilon} \|Q_1 - Q_2\|_\infty. \quad (\text{E.1})$$

952 Therefore $L_\kappa^Q \leq L_{\nabla Q_\varepsilon} + c L_{\mathcal{H} Q_\varepsilon}$, which in turn admits the spectral-norm product bounds stated in
953 Theorem 3.2. (Absolute constants absorb activation smoothness and mollifier bandwidth.)

954 E.3 PROOF OF THEOREM 3.3
955956 **E.3 at a glance**
957

958 Bound the reward term using the global slope of the sigmoid ($k/4$) and L_κ^Q , bound the value
959 term by the 1-Lipschitz max operator and the target lag τ , and combine to obtain a contraction
960 whenever Eq. (A.3) holds.

961 Let Q_1, Q_2 be two critics and abbreviate $\Delta Q = \|Q_1 - Q_2\|_\infty$. For fixed (s, a) ,

$$962 \quad |(\mathcal{T}_G^{\text{tgt}} Q_1)(s, a) - (\mathcal{T}_G^{\text{tgt}} Q_2)(s, a)| \leq |\mathbb{E}[R_G(Q_1) - R_G(Q_2)]| + \gamma |\mathbb{E} \left[\max_{a'} Q_{1,\text{tgt}} - \max_{a'} Q_{2,\text{tgt}} \right]|. \quad (\text{E.2})$$

972 **Reward term.** Since $R_G(Q) = (1 - \sigma_Q)R + \sigma_Q[R]_-$ with fixed R ,

$$973 |R_G(Q_1) - R_G(Q_2)| \leq \|R\|_\infty |\sigma_{Q_1} - \sigma_{Q_2}|.$$

974 The map $s \mapsto \text{sigmoid}(ks)$ is $k/4$ -Lipschitz, hence $|\sigma_{Q_1} - \sigma_{Q_2}| \leq \frac{k}{4} |\kappa_{Q_1} - \kappa_{Q_2}| \leq \frac{k}{4} L_\kappa^Q \Delta Q$,
975 which gives

$$976 | \mathbb{E}[R_G(Q_1) - R_G(Q_2)] | \leq \frac{k}{4} L_\kappa^Q \|R\|_\infty \Delta Q. \quad (\text{E.1})$$

980 **Value term.** Add–subtract non-lagged critics:

$$981 \max_{a'} Q_{1,\text{tgt}} - \max_{a'} Q_{2,\text{tgt}} = [\max_{a'} Q_{1,\text{tgt}} - \max_{a'} Q_1] + [\max_{a'} Q_1 - \max_{a'} Q_2] + [\max_{a'} Q_2 - \max_{a'} Q_{2,\text{tgt}}].$$

982 The middle bracket is bounded by ΔQ since \max is 1-Lipschitz; the outer brackets are bounded by
983 $C_\tau \tau$ where $\tau = \|Q - Q_{\text{tgt}}\|_\infty$. Therefore
984

$$985 | \mathbb{E}[\max_{a'} Q_{1,\text{tgt}} - \max_{a'} Q_{2,\text{tgt}}] | \leq \gamma \Delta Q + C_\tau \tau. \quad (\text{E.2})$$

986 **Combine.** From equation E.2–equation E.2,

$$987 \|\mathcal{T}_G^{\text{tgt}} Q_1 - \mathcal{T}_G^{\text{tgt}} Q_2\|_\infty \leq \left(\gamma + \frac{k}{4} L_\kappa^Q \|R\|_\infty \right) \Delta Q + C_\tau \tau.$$

988 **Contraction modulus**

989 With $\rho := \gamma + \frac{k}{4} L_\kappa^Q \|R\|_\infty + C_\tau \tau$, if $\rho < 1$ (i.e., Eq. (A.3)), then $\mathcal{T}_G^{\text{tgt}}$ is a contraction. By
990 Banach’s fixed-point theorem, the fixed point is unique and iterates converge.

991 E.4 PROOF OF PROPOSITION 4.1

1000 **E.4 at a glance**

1001 A local quadratic model links curvature to distance-to-failure; accounting for boundary motion,
1002 prediction error, and reaction lag yields the robust margin bound in Eq. (A.8).

1003 Consider a point at distance d to $\mathcal{B}_{\text{fail}}(t)$. A second-order expansion along the outward normal
1004 direction v gives

$$1005 Q(x + dv) \approx Q(x) + \langle \nabla Q(x), dv \rangle + \frac{1}{2} d^2 v^\top H_a^H(Q; x) v.$$

1006 Interpreting ΔQ_{\min} as the minimal drop when crossing the boundary yields $\frac{1}{2} d^2 \lambda_{\min}(H_a^H) \gtrsim$
1007 $\Delta Q_{\min} - \eta_{\text{approx}}$. Effective curvature is reduced (locally) by $\gamma \|\hat{\mathcal{B}}(t)\|_g + L_{\nabla Q} \epsilon_p + C_{\tau_a} V_{\max} \tau_a$,
1008 leading to the bound Eq. (A.8). Compactness, Lipschitz bounds on ∇Q , and standard perturbation
1009 arguments yield the rigorous statement.

1010 F PRE-REGISTERED PROTOCOLS AND BUDGETS

1011 **F.1 What is pre-registered?**

1012 We preregistered: (i) tasks, (ii) training budgets (env/grad steps and wall-clock class), (iii)
1013 evaluation cadence, (iv) hyperparameter search spaces and budget (n configs per method), and (v)
1014 statistical tests. A single configuration is selected per method by validation return subject to a
1015 violation cap; ties are broken by lower violations.

1024 G METRIC DETAILS AND PARETO CONSTRUCTION

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

G.1 Metrics—quick reference

$$\text{Viol} = \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{\text{cost}_t = 1\}, \quad (\text{G.1})$$

$$\text{ERL} = 60 \cdot \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{\text{hazard}(s_t) = 1\} [\text{min}^{-1}], \quad (\text{G.2})$$

$$\text{Safety-AUC} = \sum_{t=1}^T \text{Viol}(t). \quad (\text{G.3})$$

Pareto hull. We compute the upper hull over $(\text{Return}, -\text{Viol})$ via the monotone-chain algorithm; strictly dominated checkpoints are hidden. Confidence intervals use $B=2000$ bootstrap resamples over seeds. Pseudocode and tie-breaking rules are provided in the code release.

H BASELINES, ARCHITECTURES, AND HYPERPARAMETER SEARCH

H.1 Search protocol (shared across methods)

We align depth/width, activations, normalization, and optimizers where applicable. Each method receives the same search budget (n configs): learning rate, entropy/temperature, target-update period, and method-specific knobs (e.g., PCPO penalty step, FOCOPS trust region). Exact grids, final picks, and per-seed runs are in the repository; an abridged grid is shown in Table 12.

1080

Table 12: Abridged hyperparameter grids (full tables in the repository).

1081

Method	LR	Target Period	Method-specific
SAC	$\{1e-4, 3e-4\}$	$\{1e2, 1e3\}$	$\alpha \in \{0.05, 0.2\}$
PCPO	$\{1e-4, 3e-4\}$	$\{1e2, 1e3\}$	$\text{penalty-step} \in \{0.01, 0.05\}$
FOCOPS	$\{1e-4, 3e-4\}$	$\{1e2, 1e3\}$	$\text{trust-radius} \in \{0.01, 0.05\}$
A-GDC	$\{1e-4, 3e-4\}$	$\{1e2, 1e3\}$	$k \in \{3, 5\}, c \in \{0.3, 0.5, 0.7\}, \eta \in \{0.03, 0.05, 0.1\}$

1087

1088

I ABLATIONS AND SENSITIVITY

1089

I.1 What we sweep and how we report

1090

1091

1092

1093

1094

1095

1096

1097

J OVERHEAD ACCOUNTING AND IMPLEMENTATION DETAILS

1098

1099

J.1 Key points (overhead summary)

1100

1101

1102

1103

1104

1105

1106

1107

K THEORY-ALIGNED MONITORING

1108

1109

1110

We estimate \hat{L}_κ^Q via small critic perturbations on a validation buffer and log the target lag $\hat{\tau} = \|Q - Q_{\text{tgt}}\|_\infty$. We visualize the empirical contraction margin

1111

1112

$$\text{Margin} = 1 - \left[\gamma + \frac{k}{4} \hat{L}_\kappa^Q \|R\|_\infty + C_\tau \hat{\tau} \right], \quad (\text{K.1})$$

1113

and its relation to Viol across runs (scatter and time-series plots).

1114

1115

L DYNAMIC PREDICTOR AND SCHEDULE

1116

1117

1118

We train a boundary-velocity predictor with mean-absolute-error loss and an uncertainty head calibrated by temperature scaling. The curvature schedule is

1119

1120

1121

$$c(t) = c_0 (1 + \alpha \|\hat{\mathcal{B}}(t)\|_g + \beta \text{Unc}(t)), \quad (\text{L.1})$$

1122

1123

with (c_0, α, β) swept in Section I. Calibration curves (reliability diagrams) and MAE/CRPS are reported here.

1124

1125

M FAILURE CASES AND INTERPRETABILITY

1126

1127

1128

1129

We include representative rollouts where (i) extreme reward sparsity causes over-suppression and (ii) curvature is underestimated near activation seams. We provide κ heatmaps and state-visitation maps aligned with violation timestamps, along with seed IDs and minimal scripts to reproduce each case.

1130

1131

1132

1133