

Bridging Adaptivity and Safety: Learning Agile Collision-Free Locomotion Across Varied Physics

Yichao Zhong

Carnegie Mellon University

YICHAOZ@ANDREW.CMU.EDU

Chong Zhang

ETH Zurich

CHOZHANG@ETHZ.CH

Tairan He

Carnegie Mellon University

TAIRANH@ANDREW.CMU.EDU

Guanya Shi

Carnegie Mellon University

GUANYAS@ANDREW.CMU.EDU

Abstract

Real-world legged locomotion systems often need to reconcile agility and safety for different scenarios. Moreover, the underlying dynamics are often unknown and time-variant (e.g., payload, friction). In this paper, we introduce BAS (Bridging Adaptivity and Safety), which builds upon the pipeline of prior work Agile But Safe (ABS) (He et al., 2024b) and is designed to provide adaptive safety even in dynamic environments with uncertainties. BAS involves an agile policy to avoid obstacles rapidly and a recovery policy to prevent collisions, a physical parameter estimator that is concurrently trained with agile policy, and a learned control-theoretic RA (reach-avoid) value network that governs the policy switch. Also, the agile policy and RA network are both conditioned on physical parameters to make them adaptive. To mitigate the distribution shift issue, we further introduce an on-policy fine-tuning phase for the estimator to enhance its robustness and accuracy. The simulation results show that BAS achieves 50% better safety than baselines in dynamic environments while maintaining a higher speed on average. In real-world experiments, BAS shows its capability in complex environments with unknown physics (e.g., slippery floors with unknown frictions, unknown payloads up to 8kg), while baselines lack adaptivity, leading to collisions or degraded agility. As a result, BAS achieves a 19.8% increase in speed and gets a 2.36 times lower collision rate than ABS in the real world. Videos: <https://adaptive-safe-locomotion.github.io>

Keywords: Reinforcement Learning, Adaptive Safe Control, Legged Locomotion

1. INTRODUCTION

Legged robot locomotion in cluttered and dynamic environments requires adaptivity to varying physics and environmental changes while simultaneously ensuring agility for efficient navigation and safety for reliable deployment. And such adaptivity to varied environments is claimed to be crucial for real-world tasks such as disaster response in forests (Sun et al., 2020), evacuation in fire-prone areas (Panahi et al., 2023), and rescue operations (Arabboev et al., 2021). Despite recent progress in legged locomotion (Brunke et al., 2022; Hwangbo et al., 2019; Kumar et al., 2021; Lee et al., 2020; Li et al., 2024b; Xue et al., 2024; He et al., 2024a; Zhang et al., 2024a), there remains a significant gap in methodologies that effectively integrate adaptivity, safety, and agility. In this work, we enable the robot to jointly achieve agility and safety with adaptivity, maintaining strong performance in challenging environments.

Striking a good balance of adaptivity, safety, and agility in legged locomotion remains a significant challenge, as focusing on one aspect often comes at the expense of the others. Recent pioneer legged / wheeled locomotion works use reinforcement learning (RL) (Levine et al., 2020; Silver

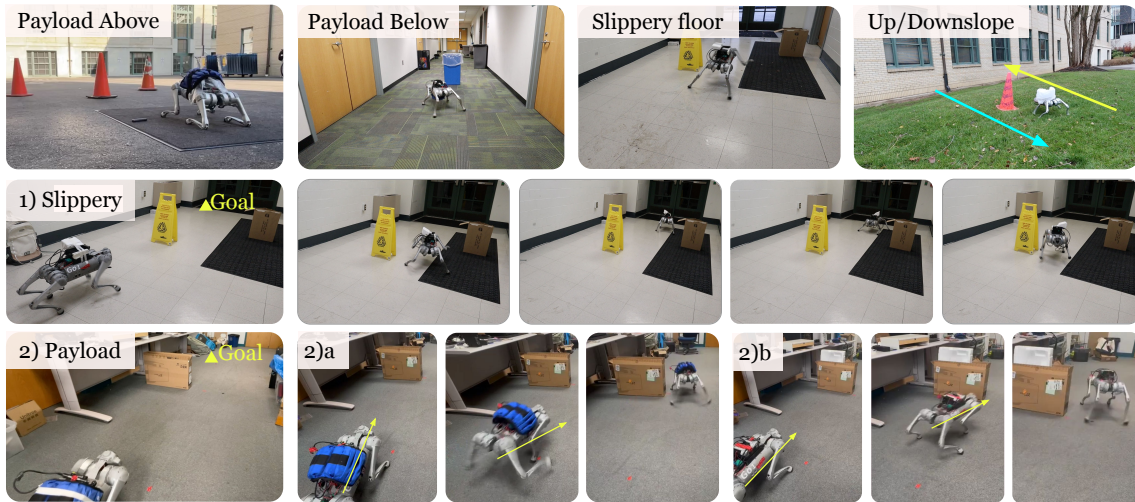


Figure 1: 1) The robot can handle collision-free locomotion in even super slippery terrain condition (soap water on both floor and robot feet), and also can adapt to rough terrain (dry carpet) suddenly. 2) Adaptive recovery triggering of the robot in different circumstances, such as a) early recovery with 8kg payload and b) late recovery with no payload.

et al., 2017) to prioritize adaptability in agility to handle environmental changes (Kumar et al., 2021; Lee et al., 2020; Wang et al., 2024; Long et al., 2024; Zhang et al., 2024b; Yang et al., 2023; Luo et al., 2024; Xiao et al., 2024b). However, these approaches somehow neglect safety considerations. On the other hand, the prior work ABS (He et al., 2024b) jointly pushes the safety and agility limits in nominal environments but is not adaptive to varying physics, and agility and safety performance can drop severely in challenging environments, as shown in our results in Figure 6.

Other studies (Xiao et al., 2024a; Chiu et al., 2022; Yun et al., 2024; Borquez et al., 2023; Gao et al., 2024) focus on adaptivity and safety, but sacrifice agility or need precise dynamics in the real world. For example, Yun et al. (2024) solves safe-legged locomotion using reduced-order dynamics models, so it sacrifices much on speed. Moreover, on the theoretical side, Borquez et al. (2023) proves that it is possible to achieve adaptive safety by parameter-conditioned reachability analysis, but the ground truth physical parameters are not accessible in the real world.

In addition to being adaptive, another way to handle changing environments is to improve robustness by making the system more conservative (Buchanan et al., 2021; Kim et al., 2020). However, being conservative can be insufficient in certain scenarios (e.g., search and rescue tasks). Moreover, although the safety-related literature is rich (Achiam et al., 2017; Bansal et al., 2017; Liu et al., 2022; Xu et al., 2021; Margellos and Lygeros, 2011; Hsu et al., 2023; Liu et al., 2020), most of them are not tested in the real world. In summary, there is a missing space for adaptive, safe, and agile locomotion for the needs of real-world applications.

To address this, we propose BAS that builds on ABS (Agile But Safe) (He et al., 2024b) and manages to enhance the adaptivity to strike a balance. Previously, ABS involves an agile policy to avoid obstacles rapidly and a recovery policy to prevent failures, and a learned control-theoretic reach-avoid value network, which governs the policy switch, guides the recovery policy as an objective function and safeguards the robot in a closed loop. Yet, unlike ABS, BAS employs an explicit

physics-parameter estimator learned from proprioceptive history during policy training as an adaptation module and feeds forward the estimated parameters to the controller and the RA (Reach-Avoid) network to enhance the adaptivity. To mitigate the distribution shift caused by switching between agile and recovery policies, we further introduce an end-to-end on-policy fine-tuning strategy, improving the accuracy of the estimator during inference. Extensive evaluations demonstrate that BAS significantly outperforms ABS and other adaptive-and-safe baselines in both safety and agility metrics. In real-world experiments, BAS achieves a 19.8% advantage in speed and is 2.36× lower in collision rate than ABS in diverse and challenging environments.

Briefly, we identify our contributions as follows:

1. We propose an adaptive safety framework, BAS (Bridging Adaptivity and Safety), for legged locomotion.
2. We propose an on-policy fine-tuning method to enhance the robustness of the parameter estimator in dynamic environments.
3. We validate the adaptivity, safety, and agility of BAS through extensive evaluations in both simulation and real-world scenarios.
4. We provide theoretical insights for parameter-conditioned reach-avoid value functions, which support the practical algorithms.

2. Preliminaries and Problem Formulation

Dynamics The dynamics is defined by state $s \in \mathcal{S} \subset \mathbb{R}^{|s|}$ and action $a \in \mathcal{A} \subset \mathbb{R}^{|a|}$ and environmental physical parameters as $e \in \mathcal{E} \subset \mathbb{R}^{|e|}$: $s_{t+1} = s_t + f(s_t, a_t, e)$. For simplicity¹, in this paper, we denote e as the physical parameters, i.e., the combination of the mass of payload, the friction coefficient, the CoM shift, etc., which is assumed static within a trajectory in training sessions. The observations are from proprioceptive and exteroceptive sensors, denoted as $o = h(s)$ where h acts as the sensor mapping.

Goal Settings Given local position and goals $\mathcal{T} \in \Gamma$, we learn a goal-conditioned reaching policy $\pi : \mathcal{O} \times \Gamma \rightarrow \mathcal{A}$ to maximize the expected return: $J(\pi) = \mathbb{E}_{\pi, \mathcal{T}} [\sum_{t=0}^{\infty} \gamma_{RL}^t r(s_t, a_t, \mathcal{T})]$, where $r(\cdot)$ is the reward at time t and γ_{RL} is the discount factor.

Safety Settings First, we denote the system trajectory starting from state s while using control inputs from the policy π under environmental parameter e as $\xi_s^{\pi, e}(\cdot) : \mathbb{R} \rightarrow \mathcal{S}$. As in [Bansal et al. \(2017\)](#), we define several basic sets: The target set $\mathcal{T} \in \mathcal{S}$ which represents the area of the goal, the constraint set $\mathcal{K} \in \mathcal{S}$ which refers to the traversable areas for robots. And the failure set $\mathcal{F} = \mathcal{K}^C$, which is the complement of the constraint set and represents hazardous areas like obstacles.

Based on those basic sets above, we can define the following sets in the context of the reachability theory. The safe set is defined as the set of states from which the robot can start and has a positive probability of rolling out a trajectory without failure, expressed as: $\omega^{\pi, e}(\mathcal{F}) := \{s \in \mathcal{S} \mid \forall \tau \geq 0, \xi_s^{\pi, e}(\tau) \notin \mathcal{F}\}$. The backward reachable set is the collection of states from which the robot has a positive probability of reaching the target: $\mathcal{R}^{\pi, e}(\mathcal{T}) := \{s \in \mathcal{S} \mid \exists \tau \geq 0, \xi_s^{\pi, e}(\tau) \in \mathcal{T}\}$. And the reach-avoid set combines the safe set and the backward reachable set: $\mathcal{RA}^{\pi, e}(\mathcal{T}, \mathcal{F}) := \{s \in \mathcal{S} \mid \exists \tau \geq 0, \xi_s^{\pi, e}(\tau) \in \mathcal{T} \wedge \forall \tau \geq 0, \xi_s^{\pi, e}(\tau) \notin \mathcal{F}\}$, which represents states from which the robot can reach the target while avoiding failure.

1. In analysis, we assume the environment is stationary, and e is unknown but static. In experiments, e can be time-variant.

Reach-Avoid Value and Time-Discounted Reach-Avoid Bellman Equation (DRABE) Identical to the vanilla reach-avoid analysis (Bansal et al., 2017), we define two Lipschitz-continuous functions $l(\cdot), \zeta(\cdot) : \mathcal{S} \rightarrow \mathcal{R}$ which satisfy $\begin{cases} l(s) \leq 0 \iff s \in \mathcal{T} \\ \zeta(s) > 0 \iff s \in \mathcal{F} \end{cases}$ to illustrate if the robot has reached the target or collides with obstacles. Note that this function is only dependent on the state s and is environment-agnostic. Then we define the reach-avoid value function V_{RA}^π which satisfies $V_{RA}^\pi(s, e) \leq 0 \iff s \in \mathcal{RA}^{\pi, e}(\mathcal{T}, \mathcal{F})$:

$$V_{RA}^\pi(s, e) = \min_{\tau \in \{0, 1, \dots\}} \max \{l(\xi_s^{\pi, e}(\tau)), \max_{\kappa \in \{0, 1, \dots, \tau\}} \zeta(\xi_s^{\pi, e}(\kappa))\}. \quad (1)$$

Note that a negative reach-avoid value guarantees a successful trajectory without collision till now, and a positive possibility to reach the target in the future. However, the function above is not learnable because it's not a contraction mapping, which fails to guarantee convergence in value iteration. To learn this function, as introduced in Hsu* et al. (2021), we use Discounted Reach-Avoid Bellman Equation (DRABE) to make the value iteration a contraction mapping:

$$B_\gamma[V_{RA_\gamma}^\pi](s_t, e) = (1 - \gamma) \max \{l(s_t), \zeta(s_t)\} + \gamma \max \{\min \{V_{RA_\gamma}^\pi(s_{t+1}, e), l(s_t)\}, \zeta(s_t)\} \quad (2)$$

Hsu* et al. (2021) also gives a mathematical proof of the DRABE operator $B_\gamma[\cdot]$ is a contraction mapping. And trivially, having V_γ^π conditioned on a static physical parameter e does not alter the proof, maintaining the guarantee of convergence.

Lipschitz-continuity of V_γ^π We deduct comprehensive analysis on V_γ^π 's convergence and Lipschitz-continuity as presented in Appendix A, from which we imply that to guarantee Lipschitz-continuity of V_γ^π , π should be not too sensitive to s and e . To this end, we employ an L2-regularization and weight clipping on π to lower its sensitivity to s and e . Moreover, as Liu et al. (2021) notes, regularization also matters in policy optimization in the context of RL because it provides better sampling complexity and return distribution.

3. METHODOLOGIES

In this section, we present our proposed framework as shown in Figure 2, which has four training phases:(Section 3.1) training parameter estimator for adaptation; (Section 3.2) training RA network; (Section 3.3) on-policy fine-tuning estimator to address the history distribution shift, and real-world deployment. Here we denote ground-truth physical parameters as e_t , the estimated ones as \hat{e}_t , and with fusion interpolation, we feed the policy with $e_t' = \alpha e_t + (1 - \alpha)\hat{e}_t$ in training, where $\alpha = \min(2 * \text{training rate}, 1)$ is the clipped training rate. For simplicity, in the following explanations, agile and recovery policies are accordingly denoted as π_{agile} and $\pi_{recovery}$.

3.1. Phase 1: Joint-Train Agile Policy and Physical Parameter Estimator

Policy-Conditioned Physical Parameter Estimator Since the environmental factors are often inaccessible in the real world, we tackle this challenge by learning a concurrent estimator $\phi^{\pi_{agile}}(O_{t:t-49})$ conditioned on π_{agile} from robot proprioception history. which explicitly estimates the mass of the payload m , the position shift of CoM $\Delta x_c, \Delta y_c, \Delta z_c$, and the friction coefficient μ , which are critical for the daily use of autonomous robots. However, training a general state estimator with high accuracy is super challenging, so we first opt to train a policy-conditioned estimator to lower the

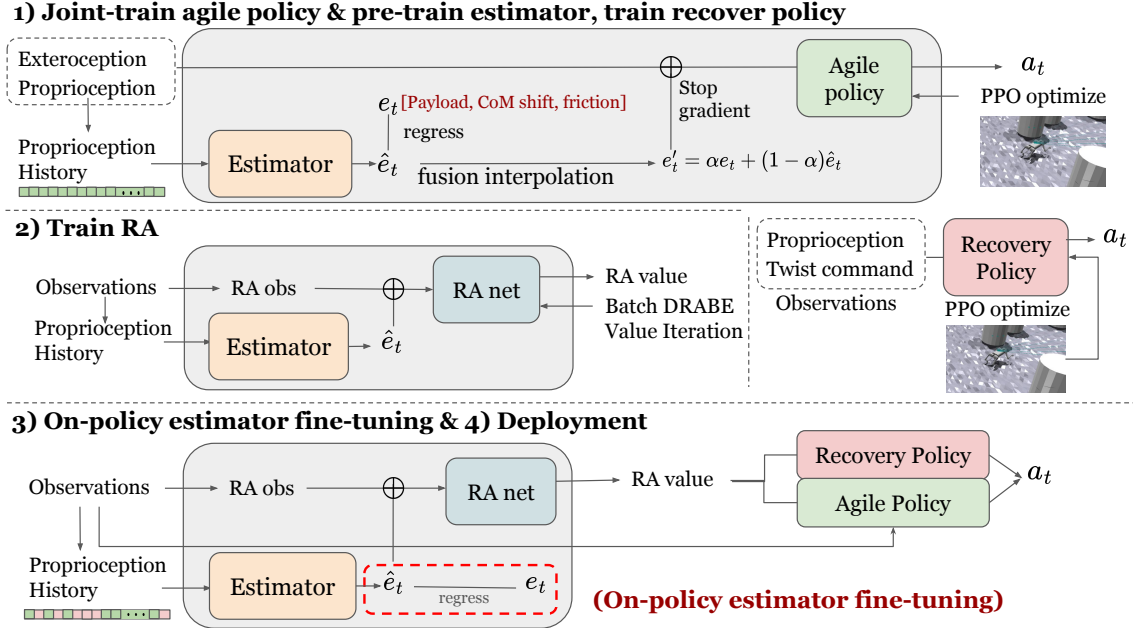


Figure 2: BAS Pipeline Overview.

challenges and then propose fusion interpolation in the joint-train pipeline to boost the accuracy further.

Additionally, note that physical parameters are policy-invariant variables. So, compared to predicting dynamics that tangle with policies, predicting physical parameters is more suitable for cases where multiple policies are used together like He et al. (2024b); Hoeller et al. (2023). What is more, estimating policy-invariant physical parameters partly lessens the potential issue of the history pattern misalignment when changing policies. To further reduce this effect, we also perform policy fine-tuning as described in Section 3.3 to maintain high accuracy at inference.

Policy Training Following ABS, we maintain the two policy switching structures: π_{agile} and $\pi_{recovery}$. π_{agile} is a goal-reaching policy and takes control most of the time, and we also retain $\pi_{recovery}$ designs from ABS, which tracks a given twist command. To make π_{agile} adaptive and aware of physics, we add the estimated physical parameters to its observation spaces as $a_t = \pi_{agile}(o_t, \hat{e}_t)$, and as for $\pi_{recovery}$, we train a robust tracking policy that works in all the environments with strong domain randomizations Table 3.

Training Pipeline To ensure that policy and estimator work well together, we propose a fusion interpolation on parameters such as $e'_t = \alpha e_t + (1 - \alpha)\hat{e}_t$ in Figure 2 in the joint-train pipeline inspired by Ji et al. (2022), where α is the training rate. Within this fusion mechanism, the agile policy receives $e'_t = \alpha e_t + (1 - \alpha)\hat{e}_t$ as input rather than e_t nor \hat{e}_t , because we expect the agile policy to converge fast with the aid of ground truth privileged observations for the beginning steps, and we want the agile policy and estimator to co-adapt to each other’s distribution when the policy converges. Moreover, such fusion mechanism reduces noise in e_t introduced by imperfect estimation in training time, helping to train a more stable π_{agile} . Furthermore, to lessen overfitting, we employ an MSE loss with L2 regularization on the estimator as well. We also perform an ablation analysis in Section 4.2 to validate this joint-training pipeline and fusion interpolation.

3.2. Phase 2: Learning Adaptive Reach-Avoid Network

As in ABS, the RA network learns a reach-avoid value function as described in Equation (2). To boost the RA network’s adaptivity, we also extend the RA observation space with the estimated physical parameters to make it aware of physics. To simplify the training, we opt to learn a policy-conditioned, normalized, and adaptive RA value function $V_\gamma^\pi(s, \hat{e})$ as a safety guard. The guarding is triggered when $V_\gamma^\pi(s, \hat{e}) > 0$, and then the system calls $\pi_{recovery}$ to take control. Typically, the RA network is learned through the MSE loss to the target value from DRABE as $L = \frac{1}{T} \|V_\gamma^\pi(s, \hat{e}) - B_\gamma[V_\gamma^\pi(s, \hat{e})]\|^2$ using $B_\gamma[\cdot]$ from Equation (2).

3.3. Phase 3: On-Policy Estimator Fine-Tuning

However, as our experiment shows in 4.2, the estimation isn’t accurate enough (e.g., $> 0.5\text{kg}$ error in mass). The cause probably lies in our structure which has two policies contribute to the same history buffer, potentially leading to a distribution shift on history, thus degrading the accuracy of the estimation. To cope with the distribution shift to ensure that the estimator performs well during policy switching, we fine-tune it end-to-end with supervision in a deployment where π_{agile} , $\pi_{recovery}$ and RA network work together. Note that this is different from the training session in Section 3.1 in that we generate the history rollouts solely with π_{agile} in phase 1, but with both policies taking effect in turn at this phase.

4. EXPERIMENTS

In this section, we present a series of simulation experiments in IsaacGym (Rudin et al., 2022; Makoviychuk et al., 2021) following the simulation setup and reward settings in ABS (He et al., 2024b) with domain randomization settings in Table 3. and real-world experiments on Unitree Go1 with onboard computations to investigate the following questions.

Q1: What are the most effective methodologies for achieving a balance between adaptive safety and agility in robotic systems?

Q2: What is the recipe for training the best estimation module in BAS?

Q3: How can we quantitatively assess the adaptivity and robustness of the BAS framework through in-depth analytical and experimental evaluations?

Q4: How well does BAS perform in real-world unseen scenarios, and how accurate can BAS’s parameter estimator be in the real world?

4.1. Safety and Agility Performance Analysis

To answer Q1 (*What are the most effective methodologies to achieve a balance between adaptive safety and agility?*), we compare the non-collision rates and average top speed within a trajectory in the simulation between BAS and other adaptive and/or safe locomotion baselines. To show BAS’s adaptivity, we introduce the following baselines: **1) ABS**, which has non-adaptive π_{agile} and RA network; **2) BAS w/o explicit estimator**, which adopts long-short term history structures and learns an encoder that maps history to latent space with end-to-end RL training (Li et al., 2024b); **3) RMA-RA**, which incorporates RMA (Kumar et al., 2021) and RA network with the latent environmental representation z_t as the additional inputs. **4) Action-Distillation**, which is similar to RMA and is inspired from Lee et al. (2020), where a student policy is distilled from an adaptive teacher policy by minimizing the difference between their actions. **5) BAS- π_{agile}** , which only uses π_{agile} ; **6) BAS-Lagrangian**, which learns π_{agile} with PPO-Lagrangian (Ray et al., 2019) with explicit estimation

Policy	Collision Rate(%) ↓	Reach Rate(%) ↑	Timeout Rate(%)	\bar{v}_{peak} of success (m/s) ↑
a) Adaptivity-wise				
BAS	1.11	93.84	5.06	2.70
BAS w/o explicit estimator	5.64	90.50	3.86	2.65
ABS	14.84	63.83	21.33	2.65
RMA-RA	12.51	80.12	7.37	2.70
Action-Distillation	15.72	68.99	15.29	2.63
b) Safety-wise				
BAS	1.11	93.84	5.06	2.70
BAS-Lagrangian	3.20	90.40	6.40	2.51
RMA-Lagrangian	13.69	76.33	9.98	2.48
BAS- π_{agile}	10.35	89.00	0.65	2.75
c) For adaptivity-robustness analysis				
BAS	1.11	93.84	5.06	2.70
BAS-random	100.00	0.00	0.00	/
RMA-RA	12.51	80.12	7.37	2.70
RMA-RA-random	19.37	74.59	6.04	2.49

Table 1: Simulation experimental results. Collision refers to trajectories with collision, Reach stands for reaching the target, and Timeout stands for being safe all over the trajectory without reaching the target. \bar{v}_{peak} of success refers to the peak speed in a trajectory on average of all the successful ones. Note that ABS values may differ from He et al. (2024b) because these experiments are done under larger domain randomizations, as shown in Table 3.

without RA network; **7) RMA-Lagrangian**, which learns a teacher PPO-Lagrangian policy with RMA and then distills it into a student policy.

As shown in Table 1 (a), BAS outperforms original ABS by 50% in reach rates in varied physics and distinctively stands out with the lowest collision rate and the highest reach rate throughout all of the adaptive methods. And in Table 1 (b), the RA safeguard structure also outperforms policies trained by PPO-Lagrangian, especially in agility. Moreover, validation of the effect of safeguarding on the entire system is observed through the comparison between BAS and BAS- π_{agile} , which indicates that adopting RA guard would transfer most of the failure cases to success cases or safe cases.

4.2. Ablation Studies on Estimator

To answer Q2 (*How to train the best estimation module in BAS?*), we investigate our two proposed methodologies to train the estimator: joint-train pipeline with fusion interpolation and on-policy fine-tuning.

Estimator Training Pipeline For ablation purposes, we test **BAS w/o fusion** (arbitrarily setting α in Figure 2 to a constant 0 or 1) and **BAS w/o joint-train** (first learn a privileged policy then learn estimation from rollout data and use estimation as privileged observation at inference) as in Table 2(a), which shows that the fusion interpolation offers a better accuracy on estimation and better overall agility-safety performance. Also, we demonstrated the tracking of mass of payload as in Figure 3, where the joint-trained estimator is much more accurate.

On-Policy Fine-Tuning Note that the estimator is only trained with the rollout of π_{agile} , which may not have seen the trajectories contributed by the agile and recovery policies. To this end, we implemented on-policy post-finetuning on the estimator to diminish this distribution shift in an end-to-end scheme. As can be seen in Table 2, BAS outperforms BAS w/o finetuning in both estimation accuracy and safety performance.

4.3. Adaptivity-Robustness Analysis




Figure 3: Mass estimation tracking of BAS, BAS w/o fusion and BAS w/o joint-train pipeline. Environment and the history buffer resets per 8s.

Entry	estimation loss	Collision Rate(%) ↓	Reach Rate(%) ↑	Timeout Rate(%)	\bar{v}_{peak} of success (m/s) ↑
a) Ablation: on training pipelines (before finetuning)					
BAS	0.570	3.10	92.48	4.42	2.69
BAS w/o fusion($\alpha \equiv 1$)	1.955	3.71	91.10	5.19	2.66
BAS w/o fusion($\alpha \equiv 0$)	5.008	16.31	52.30	31.39	2.63
BAS w/o joint-train	1.511	6.21	88.20	1.89	2.69
b) Ablation: on-policy finetuning					
BAS w/o finetuning	0.570	3.10	92.48	4.42	2.69
BAS	0.323	1.11	93.84	5.06	2.68

Table 2: Comparisons on estimators w/ and w/o fusion or joint-train and on-policy finetuning.

To Answer Q3 (*Can we identify adaptivity of BAS with deeper analysis?*), we visualize the heatmap of RA values under different physical conditions (see Figure 4). The trend in RA values aligns with common sense: heavier payloads correlate with greater danger.

Moreover, for further adaptivity analysis, we try to compare BAS to the classic adaptive baseline,

Hyperparameter Name	Value
Mass of Payload range(kg)	-2.0,12.0
Friction range	0.25,1.5
CoM shift-x(m) range	-0.05,0.05
CoM shift-y(m) range	-0.05,0.05
CoM shift-z(m) range	-0.05,0.15
External Force-x range	-15N,15N
External Force-y range	-15N,15N

Table 3: Domain Randomization Setting




Figure 4: Heatmaps of RA values under the different mass of payloads at the state of 3.0m/s base linear velocity right forward. The more reddish, the higher the RA values, the more dangerous; the more bluish, the lower the RA values, the safer.




Figure 5: Real Experiment for Adaptive Safety test Settings, where **yellow triangle** notes the starting point and **red triangle** notes the goal. Once the robot reaches the goal, we switch the goal and starting point. A trajectory from the start to the goal and then getting back to the start without collision is counted as success. **0) Vanilla test:** same as mass test settings, but without payloads. **1) Mass test:** carry a 5kg payload in a corridor and avoid boxes. **2) Friction test:** avoid box and a slip sign on very slippery floor and a dry carpet. **3) Slope test:** avoid a cone on a grass slope after rain, which is also very slippery.

RMA. As conservativeness can be identified by robustness to noise in adaptation modules, we test **BAS-random** and **RMA-RA-random**, where the output of the adaptation module (explicit estimation e_t in BAS, latent z_t in RMA) is replaced with random numbers in the same distributions. As Table 1 (c) shows, BAS deviates when the predicted mass is masked with random numbers, while masking RMA’s latent vector has a minor loss in performance, which means that BAS is less conservative and more adaptive than RMA. Explicit estimation also enhances the interpretability of the system by providing a clear understanding of the underlying physical significance of the estimations. Conversely, if environments are encoded to latent space, their meaning may remain obscure.

In summary, BAS outperforms other baselines in agility and safety metrics across our testing environments, demonstrating that its adaptivity, safety, and agility are all linked together.

4.4. Real-World Experiments

Experiment Setup To answer Q4 (*How well does BAS perform in unseen scenarios in the real world, and how accurate is the parameter estimator of BAS in the real world?*), we deploy our modules on a Unitree Go1 with onboard computations on NVIDIA Orin NX. We test three entries here: BAS, ABS and RMA+Lagrangian, as described in Section 4.1, among which ABS is our prior work and RMA+Lagrangian is also an adaptive-and-safe baseline which is worth comparing. In our experiment, the agility tests measure the agility of a policy under conditions as in Figure 5 but without obstacles, and the safety tests quantify the safety by statistics on success rates in different environments. As shown in Figure 5, we have different environment settings tailored for each physical factor that should be adapted, together with the vanilla test. Note that CoM shift is very hard to identify in real world, so as an alternative, we build an overall test which is the slope test on grass to cover it.

Real World Safety-Agility Performances As shown in Table 4, BAS outperforms ABS and RMA+lagrangian in both safety and agility across different physics and settings. We also find that ABS fails to turn 90 degrees with a 5kg payload and struggles to maintain safety on slippery terrains during experiments. During experiments, we also encountered some failures with BAS. And we analyzed some causes of failure: 1) Restricted by limited visual angle. 2) The robot only

encounters with mild collision with obstacles. 3) Indoor environments are obstacle-dense, and the ray-prediction network deviates due to out-of-distribution rays.

Policy	Adaptive Agility test(s)↓					Adaptive Safety test↑				
	Vanilla	Mass	Slope	Friction	Avg.	Vanilla	Mass	Slope	Friction	Avg.
BAS	1.39	1.67	1.50	1.09	1.41	8/8	7/8	5/8	6/8	81.25%
ABS	1.52	2.37	1.67	1.40	1.74	7/8	1/8	3/8	0/8	34.38%
RMA-Lag	1.76	1.92	1.85	2.02	1.89	6/8	5/8	0/8	2/8	40.63%

Table 4: Test results in real world. For pure agility tests, we compare the average time consumed to run 2.4m from stance in 3 trials. For safety-related tests, we compare the average success rate of 8 trials.




Figure 6: Adaptation analysis in real with online changes in the environment. 1) BAS Accomplishing collision avoidance while carrying an 8kg payload at first and then no payload in one trajectory. 2) BAS accomplishes collision avoidance in terrains with different frictions(liquid soap and water on the floor and dry mattress), while 3) ABS fails due to the lack of adaptivity. BAS estimator functions well in both cases with a correct trend.

Real World Adaptation Analysis and Run-time Estimation Figure 6 shows that BAS maintains adaptive safety even under sudden environmental changes online, such as extracting the 8kg payload or sudden changes in terrain properties such as friction, while ABS fails with insufficient adaptivity to maintain safety in this case. Moreover, as shown in the estimation plots in Figure 6, the estimation remains accurate after the changes, and BAS accomplishes avoiding obstacles under all environmental conditions, confirming its adaptability. Note that the estimated values may differ from real-world ground truth because of imperfect simulation, especially in friction, so we mainly focus on the relative values and the trends here.

5. CONCLUSIONS, LIMITATIONS, AND FUTURE PROSPECTS

In this paper, we propose BAS, which achieves collision-free locomotion in real-world dynamic environments and strikes a balance between adaptivity, agility, and safety by learning a nominal physical parameter estimator. For future works, we have several interesting research topics based on

BAS: 1) BAS currently uses the 2D ray distances to the obstacles as the exteroceptive observation, and we may try to tackle 3D scenarios with VLAs (Zhen et al., 2024; Kim et al., 2024) in the future; 2) We only trained and tested with static obstacles in this work, and will try to avoid highly dynamic obstacles in the future; 3) The current framework is focused on local obstacle avoidance, and we will try to combine high-level planning with low-level safety adaptation to accomplish more complex navigation problems in dynamic and challenging scenarios.

Acknowledgments

We gratefully acknowledge the dedication and contribution of Guanqi He, who helped us repair the hardware and gave bunches of advice on hardware designs. We appreciate Haotian Lin, Wenli Xiao, and Jiawei Gao for their assistance in real-world experiments. Special thanks to Andrea Bajcsy and Toru Lin for their ideas in graphics design.

References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization, 2017. URL <https://arxiv.org/abs/1705.10528>.
- Mukhriddin Arabboev, Shohruh Begmatov, Khabibullo Nosirov, Alisher Shakhobiddinov, Jean Chamberlain Chedjou, and Kyandoghene Kyamakya. Development of a prototype of a search and rescue robot equipped with multiple cameras. In *2021 International Conference on Information Science and Communications Technologies (ICISCT)*, pages 1–5, 2021. doi: 10.1109/ICISCT52966.2021.9670087.
- Somil Bansal, Mo Chen, Sylvia Herbert, and Claire J. Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances, 2017. URL <https://arxiv.org/abs/1709.07523>.
- Javier Borquez, Kensuke Nakamura, and Somil Bansal. Parameter-conditioned reachable sets for updating safety assurances online. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, page 10553–10559. IEEE, May 2023. doi: 10.1109/icra48891.2023.10160554. URL <http://dx.doi.org/10.1109/ICRA48891.2023.10160554>.
- Lukas Brunke, Melissa Greeff, Adam W. Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P. Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5(Volume 5, 2022):411–444, 2022. ISSN 2573-5144. doi: <https://doi.org/10.1146/annurev-control-042920-020211>. URL <https://www.annualreviews.org/content/journals/10.1146/annurev-control-042920-020211>.
- Russell Buchanan, Lorenz Wellhausen, Marko Bjelonic, Tirthankar Bandyopadhyay, Navinda Kottege, and Marco Hutter. Perceptive whole-body planning for multilegged robots in confined spaces. *Journal of Field Robotics*, 38(1):68–84, 2021. doi: <https://doi.org/10.1002/rob.21974>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21974>.
- Jia-Ruei Chiu, Jean-Pierre Sleiman, Mayank Mittal, Farbod Farshidian, and Marco Hutter. A collision-free mpc for whole-body dynamic locomotion and manipulation, 2022. URL <https://arxiv.org/abs/2202.12385>.

- Jiawei Gao, Ziqin Wang, Zeqi Xiao, Jingbo Wang, Tai Wang, Jinkun Cao, Xiaolin Hu, Si Liu, Jifeng Dai, and Jiangmiao Pang. Coohei: Learning cooperative human-object interaction with manipulated object dynamics. *arXiv preprint arXiv:2406.14558*, 2024.
- Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J. Cree. Regularisation of neural networks by enforcing lipschitz continuity, 2020. URL <https://arxiv.org/abs/1804.04368>.
- Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. *arXiv preprint arXiv:2406.08858*, 2024a.
- Tairan He, Chong Zhang, Wenli Xiao, Guanqi He, Changliu Liu, and Guanya Shi. Agile but safe: Learning collision-free high-speed legged locomotion, 2024b. URL <https://arxiv.org/abs/2401.17583>.
- David Hoeller, Nikita Rudin, Dhionis Sako, and Marco Hutter. Anymal parkour: Learning agile navigation for quadrupedal robots, 2023. URL <https://arxiv.org/abs/2306.14874>.
- Kai-Chieh Hsu*, Vicenç Rubies-Royo*, Claire Tomlin, and Jaime Fisac. Safety and liveness guarantees through reach-avoid reinforcement learning. In *Robotics: Science and Systems XVII*, RSS2021. Robotics: Science and Systems Foundation, July 2021. doi: 10.15607/rss.2021.xvii.077. URL <http://dx.doi.org/10.15607/RSS.2021.XVII.077>.
- Kai-Chieh Hsu, Allen Z. Ren, Duy P. Nguyen, Anirudha Majumdar, and Jaime F. Fisac. Sim-to-lab-to-real: Safe reinforcement learning with shielding and generalization guarantees. *Artificial Intelligence*, 314:103811, January 2023. ISSN 0004-3702. doi: 10.1016/j.artint.2022.103811. URL <http://dx.doi.org/10.1016/j.artint.2022.103811>.
- Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), January 2019. ISSN 2470-9476. doi: 10.1126/scirobotics.aau5872. URL <http://dx.doi.org/10.1126/scirobotics.aau5872>.
- Gwanghyeon Ji, Juhyeok Mun, Hyeongjun Kim, and Jemin Hwangbo. Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion. *IEEE Robotics and Automation Letters*, 7(2):4630–4637, April 2022. ISSN 2377-3774. doi: 10.1109/lra.2022.3151396. URL <http://dx.doi.org/10.1109/LRA.2022.3151396>.
- D. Kim, D. Carballo, J. Di Carlo, B. Katz, G. Bledt, B. Lim, and S. Kim. Vision aided dynamic exploration of unstructured terrain with a small-scale quadruped robot. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2464–2470, 2020. doi: 10.1109/ICRA40945.2020.9196777.
- Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model, 2024. URL <https://arxiv.org/abs/2406.09246>.

- Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots, 2021. URL <https://arxiv.org/abs/2107.04034>.
- Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science Robotics*, 5(47), October 2020. ISSN 2470-9476. doi: 10.1126/scirobotics.abc5986. URL <http://dx.doi.org/10.1126/scirobotics.abc5986>.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems, 2020. URL <https://arxiv.org/abs/2005.01643>.
- Jingqi Li, Donggun Lee, Somayeh Sojoudi, and Claire J. Tomlin. Infinite-horizon reach-avoid zero-sum games via deep reinforcement learning, 2024a. URL <https://arxiv.org/abs/2203.10142>.
- Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control, 2024b. URL <https://arxiv.org/abs/2401.16889>.
- Anqi Liu, Guanya Shi, Soon-Jo Chung, Anima Anandkumar, and Yisong Yue. Robust regression for safe exploration in control. In *Learning for Dynamics and Control*, pages 608–619. PMLR, 2020.
- Zhuang Liu, Xuanlin Li, Bingyi Kang, and Trevor Darrell. Regularization matters in policy optimization, 2021. URL <https://arxiv.org/abs/1910.09191>.
- Zuxin Liu, Zhepeng Cen, Vladislav Isenbaev, Wei Liu, Zhiwei Steven Wu, Bo Li, and Ding Zhao. Constrained variational policy optimization for safe reinforcement learning, 2022. URL <https://arxiv.org/abs/2201.11927>.
- Junfeng Long, ZiRui Wang, Quanyi Li, Liu Cao, Jiawei Gao, and Jiangmiao Pang. Hybrid internal model: Learning agile legged locomotion with simulated robot response. In *The Twelfth International Conference on Learning Representations*, 2024.
- Shixin Luo, Songbo Li, Ruiqi Yu, Zhicheng Wang, Jun Wu, and Qiuguo Zhu. Pie: Parkour with implicit-explicit learning framework for legged robots, 2024. URL <https://arxiv.org/abs/2408.13740>.
- Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning, 2021. URL <https://arxiv.org/abs/2108.10470>.
- Kostas Margellos and John Lygeros. Hamilton–jacobi formulation for reach–avoid differential games. *IEEE Transactions on Automatic Control*, 56(8):1849–1861, 2011. doi: 10.1109/TAC.2011.2105730.

- Farzad H. Panahi, Fereidoun H. Panahi, and Tomoaki Ohtsuki. An intelligent path planning mechanism for firefighting in wireless sensor and actor networks. *IEEE Internet of Things Journal*, 10(11):9646–9661, 2023. doi: 10.1109/JIOT.2023.3235998.
- Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. Technical report, OpenAI, 2019. URL <https://openai.com/research/benchmarking-safe-exploration-in-deep-reinforcement-learning>.
- Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning, 2022. URL <https://arxiv.org/abs/2109.11978>.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, October 2017. ISSN 1476-4687. doi: 10.1038/nature24270. URL <https://doi.org/10.1038/nature24270>.
- Shang-jie Sun, Shu-hai Jiang, Song-he Cui, Yue Kang, and Yu-tang Chen. Path planning of forest fire-fighting robots based on deep learning. 36:51–57, 2020. ISSN 1006-8023.
- Zhicheng Wang, Wandi Wei, Ruiqi Yu, Jun Wu, and Qiuguo Zhu. Toward understanding key estimation in learning robust humanoid locomotion, 2024. URL <https://arxiv.org/abs/2403.05868>.
- Wenli Xiao, Tairan He, John Dolan, and Guanya Shi. Safe deep policy adaptation, 2024a. URL <https://arxiv.org/abs/2310.08602>.
- Wenli Xiao, Haoru Xue, Tony Tao, Dvij Kalaria, John M Dolan, and Guanya Shi. Anycar to anywhere: Learning universal dynamics model for agile and adaptive mobility. *arXiv preprint arXiv:2409.15783*, 2024b.
- Tengyu Xu, Yingbin Liang, and Guanghui Lan. Crpo: A new approach for safe reinforcement learning with convergence guarantee, 2021. URL <https://arxiv.org/abs/2011.05869>.
- Haoru Xue, Chaoyi Pan, Zeji Yi, Guannan Qu, and Guanya Shi. Full-order sampling-based mpc for torque-level locomotion control via diffusion-style annealing. *arXiv preprint arXiv:2409.15610*, 2024.
- Yuxiang Yang, Guanya Shi, Xiangyun Meng, Wenhao Yu, Tingnan Zhang, Jie Tan, and Byron Boots. Cajun: Continuous adaptive jumping using a learned centroidal controller, 2023. URL <https://arxiv.org/abs/2306.09557>.
- Kai S. Yun, Rui Chen, Chase Dunaway, John M. Dolan, and Changliu Liu. Safe control of quadruped in varying dynamics via safety index adaptation, 2024. URL <https://arxiv.org/abs/2409.09882>.
- Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts. *arXiv preprint arXiv:2406.06005*, 2024a.

Yuanhang Zhang, Tianhai Liang, Zhenyang Chen, Yanjie Ze, and Huazhe Xu. Catch it! learning to catch in flight with mobile dexterous hands. *arXiv preprint arXiv:2409.10319*, 2024b.

Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and Chuang Gan. 3d-vla: A 3d vision-language-action generative world model, 2024. URL <https://arxiv.org/abs/2403.09631>.

Appendix A. Detailed Proof

Since Li et al. (2024a) proves that $V_\gamma(s)$ is Lipschitz-continuous to s , we extend the proof to prove that $V_\gamma^\pi(s, e)$ is Lipschitz-continuous to both s and e . For s , introducing a static e trivially doesn't alter the Lipschitz continuity of V_γ^π to s . So in this section we try to prove the Lipschitz continuity of the value function $V_\gamma^\pi(s, e)$ to environment factor e in Theorem 1.

Theorem 1 (Lipschitz Continuity of V_γ^π to e) *The Learned Value Function $V_\gamma^\pi(s, e)$ Possesses Lipschitz Continuity w.r.t. Environmental Factors e under the following conditions:*

- The functions $l(s)$ and $\zeta(s)$ are defined as L_l - and L_ζ -Lipschitz continuous functions of the state s .
- $\gamma(1 + L_{f_\pi}) < 1$, which will be naturally introduced in the proof.
- Given a specific policy π , the transition dynamics defined as $f_\pi(s, e) := f(s, \pi(s, e), e)$ are L_{f_π} -Lipschitz continuous w.r.t. the tuple (s, e) .
 L_{f_π} is defined as, for any states $s_1, s_2 \in \mathcal{S}$ and environmental factors $e_1, e_2 \in \mathcal{E}$

$$\|f(s_1, \pi(s_1, e_1), e_1) - f(s_2, \pi(s_2, e_2), e_2)\| \leq L_{f_\pi}(\|e_1 - e_2\| + \|s_1 - s_2\|),$$

where L_{f_π} is conditioned upon the policy π .

Remark. As noted by Gouk et al. (2020), the sample complexity of neural network approximation can be enhanced if the function being approximated is continuous. Consequently, the Lipschitz continuity of the value function eq. (2) is a valuable property that leads to reliable empirical performance when using neural network approximations for Reach-Avoid values. which can be found in Appendix A.

Proof Here we note V as for V_γ^π because there's only one value function in this section. By definition, we got

$$V(s, e) := \min_{\tau \in \{0, 1, \dots\}} \max \left\{ \gamma^\tau l(\xi_s^{\pi, e}(\tau)), \max_{\kappa \in \{0, 1, \dots, \tau\}} \gamma^\kappa \zeta(\xi_s^{\pi, e}(\kappa)) \right\}$$

And define $P(s, e, t)$ as payoff at timestep t :

$$P(s, e, t) := \max \left\{ \gamma^t l(\xi_s^{\pi, e}(t)), \max_{\kappa \in \{0, 1, \dots, t\}} \gamma^\kappa \zeta(\xi_s^{\pi, e}(\kappa)) \right\}$$

For all $e_1, e_2 \in \mathcal{E}$ and $s \in \mathcal{S}$, and $\theta > 0$ we have:

$$\begin{cases} \forall t \in \mathcal{R}, P(s, e_1, t) > V(s, e_1) - \theta \\ \exists \bar{t} \in \mathcal{R}, P(s, e_2, \bar{t}) < V(s, e_2) + \theta \end{cases} \quad (3)$$

Combining the two inequalities:

$$\begin{aligned} V(s, e_1) - V(s, e_2) - 2\theta &< P(s, e_1, \bar{t}) - P(s, e_2, \bar{t}) \\ &\leq \max\{\gamma^{\bar{t}} L_l \|\xi_s^{\pi, e_1}(\bar{t}) - \xi_s^{\pi, e_2}(\bar{t})\|, \max_{\kappa \in \{0, 1, \dots, \bar{t}\}} \gamma^\kappa L_\zeta \|\xi_s^{\pi, e_2}(\kappa) - \xi_s^{\pi, e_1}(\kappa)\|\} \end{aligned}$$

We use $\Delta\xi(t) := \xi_s^{\pi, e_1}(t) - \xi_s^{\pi, e_2}(t)$, and by definition of f 's Lipschitz continuity, we have

$$\begin{aligned} \Delta\xi(\bar{t}) &\leq (1 + L_f) \|\Delta\xi(\bar{t} - 1)\| + L_f (\|p_1 - p_2\|) \\ &\leq \dots = ((1 + L_{f_\pi})^{\bar{t}} - 1) \|e_1 - e_2\| \end{aligned}$$

Because this holds for some \bar{t} , so it must be less than the maximum for all \bar{t} . Thus we got

$$V(s, e_1) - V(s, e_2) \leq 2\theta + \max\{L_l, L_\zeta\} \max_{\bar{t}} \left\{ \max_{t \in \{0, 1, \dots, \bar{t}\}} \gamma^t ((1 + L_{f_\pi})^t - 1) \right\} \|e_1 - e_2\| \quad (4)$$

As θ is an arbitrary variable, we can set it to infinitesimal. To guarantee a finite bound for Lipschitz continuity of V to e , it should be assured that $\gamma(1 + L_{f_\pi}) \leq 1$ which necessarily holds that the Lipschitz constant is finite. Then we got the Lipschitz constant for the Value Function V to environmental factor e :

$$\forall s \in \mathcal{S}, V(s, e_1) - V(s, e_2) \leq L_V \|e_1 - e_2\| ,$$

where

$$L_V = \max\{L_l, L_\zeta\} \max_{t=0, 1, \dots, T} \gamma^t ((1 + L_{f_\pi})^t - 1) ,$$

where T denotes the maximum time steps for a system trajectory. Assuming $T \rightarrow \infty$ for infinite-horizon cases, by calculating the maximum point of t in the right part, we got the upper bound of L_V :

$$UB(L_V) = \max\{L_\zeta, L_l\} \cdot L_{f_\pi} \gamma^{t^*} \frac{\log(1 + L_{f_\pi})}{-\log(\gamma(1 + L_{f_\pi}))} , \quad (5)$$

where

$$t^* := \frac{\log\left(\frac{\log(\gamma)}{\log(\gamma(1 + L_{f_\pi}))}\right)}{\log(1 + L_{f_\pi})} .$$

Similar to Equation (4), we can show that $V(s, e_2) - V(s, e_1) \leq 2\theta + L_V \|e_1 - e_2\|$. Combining these two inequalities together, it can be implied that $V(s, e)$ is L_V -Lipschitz-continuous to e . ■

Following the proof, we can observe that L_V is also bounded by $\max\{L_l, L_\zeta\}$ because the exponential term $\gamma^t((1 + L_{f_\pi})^t - 1)$ should be less than 1.

The assumption $\gamma(1 + L_{f_\pi}) < 1$ also gives a constraint that the dynamics with respects to the policy π shouldn't be too sensitive to environment factor e , i.e. π should be a robust policy to e .