# DISTRIBUTIONAL VALUE GRADIENTS FOR STOCHASTIC ENVIRONMENTS

**Anonymous authors**Paper under double-blind review

000

001

003 004

010 011

012

013

014

016

017

018

019

021

025

026027028

029

031

032

033

034

035

037

038

040

041

042

043

044

045

046

047

048

051

052

#### **ABSTRACT**

Gradient-regularized value learning methods improve sample efficiency by leveraging learned models of transition dynamics and rewards to estimate return gradients. However, existing approaches, such as MAGE, struggle in stochastic or noisy environments, limiting their applicability. In this work, we address these limitations by extending distributional reinforcement learning on continuous stateaction spaces to model not only the distribution over scalar state-action value functions but also over their gradients. We refer to this approach as *Distributional* Sobolev Training. Inspired by Stochastic Value Gradients (SVG), our method utilizes a one-step world model of reward and transition distributions implemented via a conditional Variational Autoencoder (cVAE). The proposed framework is sample-based and employs Max-sliced Maximum Mean Discrepancy (MSMMD) to instantiate the distributional Bellman operator. We prove that the Sobolevaugmented Bellman operator is a contraction with a unique fixed point, and highlight a fundamental smoothness trade-off underlying contraction in gradient-aware RL. To validate our method, we first showcase its effectiveness on both a controlled supervised-learning task and a simple reinforcement-learning toy problem, then benchmark its performance on several MuJoCo environments.<sup>1</sup>

#### 1 Introduction

Reinforcement learning (RL) tackles sequential decision-making by training agents to maximize cumulative rewards. Off-policy actor-critic algorithms pair an actor, generating the control policy, with a critic, estimating expected returns (i.e., the Q-function). This mapping from state-action pairs to expected returns, known as credit assignment, is typically learned via temporal-difference (TD) methods (Sutton, 1988) and is critical for policy optimization in continuous-action settings. **This paper is motivated by two lines of work aimed at improving credit assignment:** 

- The use of action-gradients: In value-based continuous control, the critic's value is not used directly to select actions but to provide action-gradients for policy optimization (Lillicrap et al., 2016; Fujimoto et al., 2018; Haarnoja et al., 2018; D'Oro & Jaskowski, 2020). Conventional TD learning implicitly learns these gradients via value prediction, but this relies on smoothness assumptions of the true value function that can degrade performance. To address this limitation, D'Oro & Jaskowski (2020); Garibbo et al. (2024) incorporate gradient information (Czarnecki et al., 2017) into critic training by learning a transition–reward world model (i.e. a differentiable proxy for the environment) and backpropagating through it (Heess et al., 2015).
- **Distributional RL:** Many environments exhibit irreducible uncertainty in transitions and rewards. Distributional RL (Morimura et al., 2010; Bellemare et al., 2017; 2023) captures this by modeling the return distribution rather than just its expectation. Categorical (Barth-Maron et al., 2018) and quantile-based (Dabney et al., 2018b;a) approaches have provided rich and stable learning signals, yielding performance gains in a variety of tasks (Barth-Maron et al., 2018; Dabney et al., 2018a; Hessel et al., 2017).

We argue that randomness also affects action gradients of returns, which can have a detrimental effect, especially in high-dimensional action spaces. As our experiments show (Section 6), existing

<sup>&</sup>lt;sup>1</sup>We will make the code publicly available upon acceptance of the paper.

methods that use gradient information deterministically (Czarnecki et al., 2017; D'Oro & Jaskowski, 2020) may struggle once the gradient to model becomes noisy or stochastic, losing some of the sample-efficiency benefits of gradient modeling.

Paper contributions We extend distributional modeling to capture both returns and their gradients, coining the framework **Distributional Sobolev Reinforcement Learning**. At its core is a novel Sobolev Bellman operator that bootstraps both return and gradient distributions. By marrying gradient-based training with uncertainty modeling, we aim to boost policy and value learning. This necessitates a generative model that supports differentiation of outputs and their input gradients, so we introduce **Distributional Sobolev Training** and detail its implementation. Since most environments are non-differentiable, we employ a conditional VAE (cVAE (Sohn et al., 2015)) to model transitions and rewards. This enriches SVG (Heess et al., 2015) with a more expressive neural architecture. Finally, we extend previous works on value gradient and introduce the framework of Sobolev Temporal Difference. We provide the first contraction proofs in this scheme. In this context, we introduce the maximum-sliced MMD metric, as a practical metric that is both a contraction and tractable to approximate.

## 2 BACKGROUND

#### 2.1 NOTATION AND RL OBJECTIVE

We consider a Markov Decision Process (MDP) with continuous state and action spaces,  $\mathcal{S}$  and  $\mathcal{A}$ , transition kernel  $P \colon \mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathcal{S})$ , reward law  $R \colon \mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R})$ , and initial distribution  $\mu \in \mathcal{P}(\mathcal{S})$ . A deterministic policy  $\pi_{\theta} \colon \mathcal{S} \to \mathcal{A}$  induces the  $\gamma$ -discounted occupancy  $d_{\mu}^{\pi_{\theta}} = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^{t} \Pr(s' | \pi_{\theta}, \mu)$  (Silver et al., 2014; D'Oro & Jaskowski, 2020). The Q-function  $Q^{\pi}(s, a)$  is the expected future return starting from state s and action s, i.e.,  $Q^{\pi}(s, a) = \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t} r(s_{t}, a_{t}) \mid s_{0} = s, a_{0} = a\right]$ . It yields the objective

$$J(\theta) = \mathbb{E}_{s \sim \mu} \left[ Q^{\pi_{\theta}}(s, \pi_{\theta}(s)) \right]. \tag{1}$$

Under mild conditions (Silver et al., 2014), the Deterministic Policy Gradient theorem gives

$$\nabla_{\theta} J(\theta) = \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d_{\mu}^{\pi_{\theta}}} \left[ \nabla_{\theta} \pi_{\theta}(s) \nabla_{a} Q^{\pi_{\theta}}(s, a) \right]_{a = \pi_{\theta}(s)} \tag{2}$$

#### 2.2 Temporal-difference as an affine operator

In practice the true  $Q^{\pi}$  is unknown and is approximated by a parameterized critic  $Q_{\phi}$ . More generally, any value-like mapping  $V \colon \mathcal{S} \times \mathcal{A} \to \mathcal{Y}$  (where  $\mathcal{Y} = \mathbb{R}$  or a space of probability distributions Bellemare et al. (2017)) admits a temporal-difference update written as a single affine operator:

$$(\mathcal{T}_{\pi}V)(s,a) = b(s,a) + \mathcal{L}[V](s,a). \tag{3}$$

Here b(s, a) injects the immediate-reward term and  $\mathcal{L}$  linearly transforms the successor estimate. The following recovers the Bellman expectation operator where Q is the state-action value function

$$b(s,a) = \mathbb{E}[R(s,a)], \qquad \mathcal{L}^{\text{Exp}}[Q](s,a) = \gamma \,\mathbb{E}[Q(s',\pi(s')) \mid s,a] \tag{4}$$

In distributional RL Bellemare et al. (2017), let  $Z^{\pi}(s,a)$  be the random return with distribution  $\eta^{\pi}(s,a)$ . One gets

$$b(s,a) = \operatorname{Law}[R(s,a)], \qquad \mathcal{L}^{\operatorname{Dist}}[\eta](s,a) = \mathbb{E}_{s' \sim P(\cdot \mid s,a)}[(x \mapsto \gamma x)_{\#} \eta(s', \pi(s'))], \quad (5)$$

where  $(x \mapsto \gamma x)_{\#} \eta$  is the *law of*  $\gamma X$  when  $X \sim \eta$ , which yields

$$\left(\mathcal{T}_{\pi}^{\text{Dist}}\eta\right)(s,a) = \text{Law}\left[R(s,a) + \gamma Z(s',\pi(s'))\right], \quad \text{where} \quad s' \sim P(\cdot \mid s,a). \tag{6}$$

Using off-policy samples  $(s, a, r, s') \sim B$  from a replay buffer (Mnih et al., 2013) together with delayed target networks  $\theta'$ ,  $\phi'$  (Lillicrap et al., 2016; Fujimoto et al., 2018; Haarnoja et al., 2018), we define the one-step targets

$$\delta_{\text{tgt}}(s, a, s') = r + \gamma Q_{\phi'}(s', \pi_{\theta'}(s')), \qquad \eta_{\text{tgt}}(s, a) = \text{Law}[r + \gamma Z_{\phi'}(s', \pi_{\theta'}(s'))]. \tag{7}$$

The critic  $V_{\phi}$  (scalar  $Q_{\phi}$  or distribution  $Z_{\phi}$ ) is then trained by minimizing

$$\mathcal{L}(\phi) = \mathbb{E}_{(s,a,r,s')\sim B} \left[ d\left(V_{\phi}(s,a), T(s,a)\right) \right], \tag{8}$$

where  $T = \delta_{tgt}$  in the expected-value case or  $T = \eta_{tgt}$  in the distributional case, and d is either a regression loss (e.g. squared error) or a distributional metric such as the Wasserstein distance (Bellemare et al., 2017; Sun et al., 2024).

## 3 A NEW BELLMAN OPERATOR

#### 3.1 Learning a useful critic

Many value-based methods (Lillicrap et al., 2016; Fujimoto et al., 2018; Haarnoja et al., 2018) rely on a learned critic to provide the actor's training signal, implying that "an actor can only be as good as allowed by its critic" (D'Oro & Jaskowski, 2020). Unfortunately, typical critics — predicting only mean returns — cannot capture inherent return uncertainty. Distributional RL addresses this by modeling the return distribution. However, as noted in D'Oro & Jaskowski (2020), another fundamental issue is that minimizing TD-error does not guarantee the critic will be effective at steering policy optimization. We therefore propose a more principled approach to distributional temporal-difference learning, which explicitly incorporates the critic's action-gradient into its training objective. This aligns critic optimization with policy improvement rather than just fitting returns or their distribution.

**Proposition 1.** Let  $\pi$  be an  $L_{\pi}$ -Lipschitz continuous policy, and let  $G(s) = \operatorname{Law}[\nabla_a Z^{\pi}(s,a) \mid_{a=\pi(s)}]$  and  $\hat{G}(s) = \operatorname{Law}[\nabla_a \hat{Z}(s,a) \mid_{a=\pi(s)}]$  denote the true and estimated distributions of the action-gradients at  $a=\pi(s)$ , respectively. Define the p-Wasserstein distance between two probability measures  $\mu, \nu$  by

$$W_p(\mu,\nu) = \left(\inf_{\gamma \in \Pi(\mu,\nu)} \mathbb{E}_{(X,Y) \sim \gamma} \left[ \|X - Y\|^p \right] \right)^{1/p}.$$

Then, specializing to p=1, the error between the true policy gradient  $\nabla_{\theta}J(\theta)$  and its estimate  $\nabla_{\theta}\hat{J}(\theta)$  satisfies

$$\|\nabla_{\theta} J(\theta) - \nabla_{\theta} \hat{J}(\theta)\| \leq \frac{L_{\pi}}{1 - \gamma} \mathbb{E}_{s \sim d_{\pi}^{\mu}} \Big[ W_1 \big( G(s), \, \hat{G}(s) \big) \Big].$$

Proof is in Appendix D. This result generalizes Proposition 3.1 from D'Oro & Jaskowski (2020) to a distributional setting. The Lipschitz continuity of  $\pi$  typically holds when using neural-network function approximation.

Following D'Oro & Jaskowski (2020), we induce a critic optimization objective from Proposition 1, showing that we can approximate the true policy gradient by matching the action gradients in the *distributional sense*. Using bootstrapping to approximate the true distribution leads to the optimization problem

$$\widehat{Z} \in \underset{\widehat{Z} \in \mathcal{Z}}{\arg \min} \underset{\substack{s \sim d_{\mu}^{\pi}(s) \\ (s',r) \sim p(s',r|s,\pi_{\theta}(s))}}{\mathbb{E}} \left[ W_{1} \left( \nabla_{a} \widehat{Z}(s,a), \nabla_{a} r(s,a) + \gamma \nabla_{a} \widehat{Z}(s',\pi_{\theta}(s')) \right) \right]. \tag{9}$$

We note that Eq. 9 assumes a known and differentiable dynamics p. We maintain this assumption for the time being and will relax it in Section 5. Mirroring previous work (D'Oro & Jaskowski, 2020; Garibbo et al., 2024), we introduce this assumption upfront and then lift the constraint. In the next section we formalize the notions necessary to instantiate a working implementation of this optimization problem.

#### 3.2 DISTRIBUTIONAL SOBOLEV TRAINING

In this section, we introduce a novel Bellman operator for learning the joint distribution of the discounted cumulative reward and its action-gradient, and then express it in the affine-transform form presented earlier.

Random action Sobolev return We extend the random return Z(s,a) to a *joint random variable* that captures both the return and its action-gradient. Formally, the **random action Sobolev return** is

$$Z^{S_a}(s,a) = \left[\sum_{t=0}^{\infty} \gamma^t \, r(s_t, a_t); \, \nabla_a \sum_{t=0}^{\infty} \gamma^t \, r(s_t, a_t)\right], \quad s_0 = s, \, a_0 = a.$$
 (10)

**Sobolev distributional temporal difference** Next, we define the Sobolev distributional Bellman operator  $T_{\pi}^{S_a}$  over these  $(|\mathcal{A}|+1)$ -dimensional random variables. Let  $\eta^{S_a}(s,a) = \operatorname{Law}[Z^{S_a}(s,a)]$ . We borrow notation from Zhang et al. (2021); Rowland et al. (2019) and extend the classical distributional operator as follows. Under policy  $\pi$ , sample

$$s' \sim P(\cdot \mid s, a), \quad a' = \pi(s'), \quad r \sim R(\cdot \mid s, a), \quad X' \sim \eta^{S_a}(s', a').$$

Define the full affine transform

$$\mathbf{f}^{S_a}(x; r, s', \gamma) = \left[ f^{\text{return}}(x); f^{\text{action}}(x) \right]. \tag{11}$$

For readability we hereafter write  $\mathbf{f}^{S_a}(x)$ , implicitly carrying the dependence on  $(r,s',\gamma)$ . We define  $T_{\pi}^{S_a}$  as the operator that, at each (s,a), pushes the next-step law  $\eta^{S_a}(s',a')$  forward through this pointwise affine map:

$$\left(T_{\pi}^{S_a} \eta^{S_a}\right)(s, a) := \operatorname{Law}\left[\mathbf{f}^{S_a}(X')\right]. \tag{12}$$

Its components are

$$f^{\text{return}}(x) = r + \gamma x^{\text{return}},$$
 (13)

$$f^{\text{action}}(x) = \frac{\partial r}{\partial a}(s, a) + \gamma \left(\frac{\partial f}{\partial a}(s, a)\right)^T \left[\partial_s x^{\text{return}} + (\partial_s \pi(s'))^T x^{\text{action}}\right]. \tag{14}$$

This action-gradient component is novel: it arises by differentiating the Bellman target, capturing how the return's gradient transforms under  $P(s',r\,|\,s,a)$ . The derivation of Eqs. equation 13–equation 14 appears in Appendix C. Notably, since  $f,\,r,\,$  and  $\pi$  are differentiable, these updates are implemented automatically via backpropagation through the reparameterized simulator (Baydin et al., 2018; Paszke et al., 2019; Bradbury et al., 2018). It is worth noting that exactly the same chain-rule derivation extends to also bootstrap the state-gradient (not just the action-gradient), yielding the *complete* Sobolev Bellman operator that no longer requires the on-the-fly differentiation in Eq. 14. We discuss this operator in more details in Appendix C.

**Affine form of the Sobolev Bellman backup** As shown in Appendix C, the Sobolev Bellman backup can be written in a single affine-operator form:

$$Z^{S_a}(s,a) = b(s,a) + \mathcal{L}^{Sob}(s,a) [Z^{S_a}(s',a')],$$
(15)

where  $b(s,a)=(r(s,a),\partial_a r(s,a))\in\mathbb{R}^{1+|\mathcal{A}|}$  collects the immediate reward and its action-gradient, and  $\mathcal{L}^{\operatorname{Sob}}(s,a)$  is a state-action-dependent *linear operator* (not merely a matrix) encapsulating the Jacobian blocks of the transition f and policy  $\pi$ . Importantly, Equation 15 is a strong departure from recent works using critic gradient information (D'Oro & Jaskowski, 2020; Garibbo et al., 2024). While these works considered the target action-gradient for regularization, we consider it as a part of the bootstrapped quantity at equal level with the scalar value. This formalism allows to hold contraction statements as described in the next section.

#### 4 THEORETICAL RESULTS

The most natural distance for distributional RL is the Wasserstein metric, and we therefore begin by establishing contraction under this choice. See Appendix D for proof and details.

**Theorem 1** (Action-gradient Sobolev contraction). *Assume bounded Jacobians for the transition f and policy*  $\pi$ *, and a Lipschitz coupling relating state- to return-gradients. Then* 

$$\bar{W}_p(T_\pi^{S_a}\eta_1, T_\pi^{S_a}\eta_2) \leq \gamma \kappa \bar{W}_p(\eta_1, \eta_2),$$

where  $\bar{W}_p(\eta_1, \eta_2) = \sup_{(s,a)} W_p(\eta_1(s,a), \eta_2(s,a))$ . If  $\gamma \kappa < 1$ ,  $T_{\pi}^{S_a}$  is a strict contraction.

Towards a tractable metric Both the classical and *Sobolev* distributional Bellman operators (Eqs. 6, 12) involve intractable push-forward integrals. While the supremum–p–Wasserstein distance  $\bar{W}_p$  is a natural theoretical choice (Bellemare et al., 2017), exact multivariate optimal transport costs  $O(m^3 \log m)$  with m samples and, more broadly, Wasserstein distances are difficult to estimate and to use directly for training (Appendix E). This motivates a shift to metrics that remain faithful to distributional structure yet are easier to compute in practice. One such candidate is the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012), a kernel-based divergence that is tractable, sample-based, and already explored in distributional RL (Nguyen et al., 2020; Killingberg & Langseth, 2023; Wiltzer et al., 2024). For laws  $P, Q \subset \mathbb{R}^{d'}$ , the squared Maximum Mean Discrepancy between P and Q is

$$MMD^{2}(P,Q) = \mathbb{E}_{x,x'\sim P}[k(x,x')] + \mathbb{E}_{y,y'\sim Q}[k(y,y')] - 2\,\mathbb{E}_{x\sim P,\,y\sim Q}[k(x,y)], \quad (16)$$

where k denotes the kernel function. Further properties and empirical estimators are discussed in Appendix F.2.

**Max–Sliced MMD.** To obtain *provable contraction* for our Sobolev operator, we lift MMD via the max–sliced divergence framework (Deshpande et al., 2019; Nadjahi et al., 2020). For  $\theta \in \mathbb{S}^{d'-1}$  and  $P_{\theta}(x) = \langle \theta, x \rangle$ , we define

$$\mathbf{MS}\mathrm{MMD}(\mu,\nu) = \sup_{\theta \in \mathbb{S}^{d'-1}} \mathrm{MMD}((P_{\theta})_{\#}\mu, (P_{\theta})_{\#}\nu).$$

Approximation. MSMMD can be approximated by gradient-based optimization of the direction  $\theta$  on the unit sphere; see Algorithm 1.

#### Contraction under Max-Sliced MMD.

**Theorem 2** (Action–gradient Sobolev contraction under MSMMD). Assume the conditions of Theorem 1 and the mild additions in Theorem 7. Then

$$\overline{\mathbf{MS}\mathrm{MMD}}\!\!\left(T_{\pi}^{S_a}\eta_1,\;T_{\pi}^{S_a}\eta_2\right)\;\leq\;\gamma\,\kappa\;\overline{\mathbf{MS}\mathrm{MMD}}\!\!\left(\eta_1,\eta_2\right),$$

where  $\overline{\mathbf{MSMMD}}(\eta_1, \eta_2) = \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbf{MSMMD}(\eta_1(s,a), \eta_2(s,a))$  and  $\kappa < \infty$  depends on the Jacobian bounds; see Appendix G. If  $\gamma \kappa < 1$ ,  $T_{\pi}^{S_a}$  is a strict contraction with a unique fixed point.

**Trade-off interpretation** The condition  $\gamma \kappa < 1$  makes the trade-off explicit: either enforce smoothness (reduce  $\kappa$  via bounded Jacobians and Lipschitz couplings) or shorten the effective horizon (reduce  $\gamma$ ). A similar quantity  $\kappa$  governs our contraction results for both Wasserstein (Theorem 1) and max–sliced MMD (Theorem 2), so the trade-off arises in each case. Importantly,  $\kappa$  is dictated by the environment's dynamics and policy sensitivities; when the underlying physics yields large gradients, this effect cannot be eliminated and the only remedy is to lower  $\gamma$ . **This observation is a core contribution of our work and is enabled by the Sobolev Temporal Difference framework we introduce.** 

#### 5 APPROACH

To turn our theoretical Sobolev-distributional Bellman operator into a practical algorithm, we need (i) a critic that produces joint return–gradient samples, (ii) Bellman backups without intractable integrals, (iii) a tractable metric for comparing predicted and target distributions, and (iv) to relax the assumption that our environment is differentiable. We address each of these below.

**Sobolev inductive bias.** We coin the term *Sobolev inductive bias* for the simple idea of modeling a gradient by a gradient: using the gradient of our approximator to stand in for the gradient of the true function. Concretely, let  $F: \mathbb{R}^a \to \mathbb{R}^b$  be a differentiable target and  $f_{\varphi}$  a neural network parameterized by  $\varphi$ . Sobolev training is one concrete instance of this principle (e.g. (Czarnecki et al., 2017; D'Oro & Jaskowski, 2020)), which implements  $\mathcal{L}^S(\varphi;x) = \|F(x) - f_{\varphi}(x)\|^2 + \lambda_S \|\nabla_x F(x) - \nabla_x f_{\varphi}(x)\|^2$ . Further mathematical motivation for this inductive bias can be found in Appendix B.1.

**Reparameterized Sobolev critic.** We model the joint return-and-gradient law  $\eta_{\pi}^{S_a}(s,a)$  by a generator

$$Z_{\phi}^{S_a}: (s, a, \xi) \longmapsto \Big(Z_{\phi}(s, a, \xi), \nabla_a Z_{\phi}(s, a, \xi)\Big), \qquad \xi \sim \mathcal{N}(0, I).$$

such that  ${\rm Law}\big[Z_\phi^{S_a}(s,a)\big]=\eta_\phi^{S_a}(s,a)\approx\eta_\pi^{S_a}(s,a)$ . This purely sample-based critic sidesteps intractable likelihoods and scales naturally to high-dimensional actions. It is similar in spirit to Singh et al. (2022); Freirich et al. (2019), and is structured as a generative model that deterministically maps noise to samples (Li et al., 2015; Goodfellow et al., 2014). We discuss further this choice and why most alternative parametrizations would not fit our scenario in Appendix B.2.

Overestimation bias. Value estimates routinely exhibit overestimation (Hasselt, 2010; van Hasselt et al., 2015), and even gradient-regularized critics inherit this issue (D'Oro & Jaskowski, 2020; Garibbo et al., 2024). It biases the policy toward overvalued actions, undermining its performance. TD3 addresses it by training two critics and setting the target in Equation. 8 to the minimum of the two (Fujimoto et al., 2018). In our sample-based distributional setting, we follow TQC (Kuznetsov et al., 2020): train two distributional critics, draw N samples from each, discard the top p% by magnitude, and concatenate the rest to form the target distribution.

One-step world model. Now we relax the assumption that the environment is differentiable. Unlike (D'Oro & Jaskowski, 2020; Garibbo et al., 2024), which fit only the conditional expectation  $(\hat{s}',\hat{r}) = \mathbb{E}[s',r \mid s,a]$ , we learn a stochastic, differentiable simulator g whose push-forward law approximates the true transition—reward distribution  $P(s',r \mid s,a)$ . Concretely, we posit

$$(\hat{s}', \hat{r}) = g(s, a, \varepsilon), \quad \varepsilon \sim \rho_w(\varepsilon), \quad \text{Law}[g(s, a)] \approx \text{Law}[s', r \mid s, a].$$

Using g in place of the true environment, our Sobolev Bellman update from Equation 11 becomes

$$Z^{S_a}(s,a) = \mathbf{f}^{S_a}(Z^{S_a}(\hat{s}',\hat{a}'); \hat{r}, s', \gamma), \quad \hat{a}' = \pi(\hat{s}').$$

We implement g via a conditional VAE (Sohn et al., 2015), which has proven effective for modeling dynamics in RL (Ha & Schmidhuber, 2018; Zhu et al., 2024), with learned prior  $p_v(\varepsilon \mid s, a)$ , encoder  $q_{\zeta}(\varepsilon \mid s, a, s', r)$ , and decoder  $p_{\psi}(s', r \mid s, a, \varepsilon)$ .

**Algorithm summary.** We replace the critic in standard off-policy actor–critic (Fujimoto et al., 2018; Barth-Maron et al., 2018) with our Sobolev distributional critic and train it and the cVAE world model  $g(s, a, \varepsilon)$  jointly; this yields the **Distributional Sobolev Deterministic Policy Gradient (DSDPG) algorithm**. At each step we sample Sobolev-returns and  $(\hat{s}', \hat{r})$ , form bootstrapped Sobolev-Bellman targets, minimize the MSMMD loss on value–gradient pairs, and update the actor by ascending the gradient of the critic's estimated expected return. *Even though plain MMD is not showed to be contractive, we also instantiate our framework using it thus proposing two metrics to train the critics*. Fig. 1 (left) illustrates the overall DSDPG workflow, while the right panel shows how Sobolev-return distributions are sampled and the MMD loss is estimated. The full pseudo-code can be found in Appendix J.

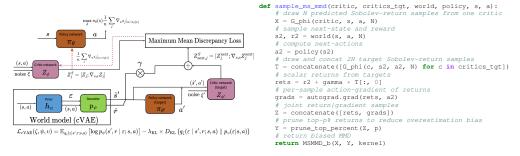


Figure 1: Left: block diagram of our DSDPG algorithm, where the critic  $Z_{\phi}$  maps noise  $\xi$  and (s,a) to Sobolev-return samples, a cVAE world model generates next-state–reward samples  $(\hat{s}',\hat{r})$ , MSMMD or MMD compares predicted and bootstrapped Sobolev-return distributions, and the policy is updated via the critic's mean (gradient flows shown as dashed arrows; inspired by Singh et al. (2022)). Right: pseudocode for estimating the biased MSMMD between predicted and bootstrapped Sobolev returns.

## 6 RESULTS

#### 6.1 Toy Reinforcement Learning

We introduce a simple continuous-state, continuous-action 2D point-mass task within a square bounding box (Figure 2a). At each time step, the agent controls the 2D acceleration of the mass. At the start of each episode, we randomly draw one of M possible bonus locations arranged uniformly around the center. Reaching this location ends the episode with a terminal reward. The agent does not know which location is correct. Instead, each time it visits a location, a binary memory flag is set. This partial observability makes the task sparse and forces exploration. By varying M, we tune the number of modes in the return distribution: small M yields near-deterministic returns, large M yields multimodal, high-variance returns. This toy task offers a clean, adjustable benchmark for gradient-aware distributional critics under controlled multimodal uncertainty.

As shown in Figure 2b, Distributional Sobolev (MSMMD Sobolev in pink and MMD Sobolev, orange) consistently outperforms all baselines as we increase the number of bonus locations. This indicates that our method effectively leverages gradient information and remains robust to the growing multimodality of the return distribution. By contrast, deterministic Sobolev (Huber Sobolev, red) offers no clear advantage over gradient-agnostic critics. Interestingly, the provably contractive MSMMD has a mild advantage over plain MMD. For brevity, we defer the full experimental details and baseline descriptions to the next section; exact empirical settings can be found in Appendix L.

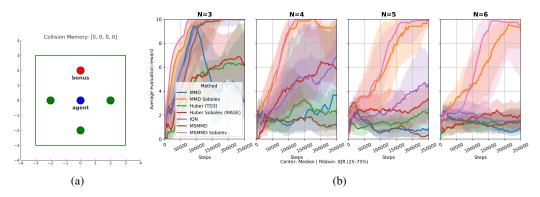


Figure 2: (a) Illustration of the 2D point-mass environment with N possible bonus locations. (b) Evaluation curves (10 means the agent reached the bonus) for our Distributional Sobolev (MMD Sobolev and MSMMD Sobolev, in orange and pink resp.), deterministic Sobolev (MAGE D'Oro & Jaskowski (2020), in red), and other baselines as N varies (median over 5 seeds and 25%-75% IQR).

#### 6.2 Reinforcement learning

In this section, we evaluate the complete solution, including the learned world model, on several standard Mujoco environments (Todorov et al., 2012) from the Gymnasium library Towers et al. (2024). Our deep learning framework is JAX (Bradbury et al., 2018). Similarly to MAGE (D'Oro & Jaskowski, 2020), for every exploration step 10 critic updates are taken (UTD ratio = 10). All experiments are ran in *Dyna* setting (Deisenroth & Rasmussen, 2011) where new samples are drawn from the world model for every critic update. Thus we evaluate the ability of the Distributional Sobolev training framework against deterministic Sobolev (Czarnecki et al., 2017; D'Oro & Jaskowski, 2020; Garibbo et al., 2024). Hence, the results should not be understood as looking for state-of-the-art performance in the given environments but rather a showcase that in some difficult settings, distributional methods can perform better.

We compare against four baselines: (i) a TD3 variant trained with a Huber loss (Fujimoto et al., 2018; D'Oro & Jaskowski, 2020); (ii) deterministic Sobolev (MAGE) trained on both return and action-gradient via Huber loss (D'Oro & Jaskowski, 2020); and (iii-iv) distributional baselines, IQN (Dabney et al., 2018a) and standard MMD training (Nguyen et al., 2020). All variants involving MMD use a multiquadric kernel with h=100. We use an identical base architecture for every method, with the distributional approaches sampling by reparameterization (Section 5). Each distri-

 butional variant generates N=10 samples per transition. Further architectural and hyperparameter details are in Appendix M.

To evaluate robustness under stochastic dynamics, we use two complementary perturbations in Mu-JoCo. (i) Multiplicative observation noise: at the start of each episode we sample  $n \sim \mathcal{U}[0.8, 1.2]$  and scale the observation  $s \mapsto n \, s$ , following Khraishi & Okhrati (2023); this induces partial observability and substantially increases task difficulty. (ii) Additive Gaussian dynamics noise: we inject zero-mean Gaussian noise directly into positions and velocities, adopting the standard deviations reported by Khraishi & Okhrati (2023). This second perturbation makes control harder while preserving the qualitative structure of the original tasks. Results are summarized in Figure 3, where the left panels report the final evaluation performance after 150,000 iterations (250,000 for Humanoid-v2), and the right panels show the normalized area under the evaluation curve (nAUC). All results are reported as medians with bootstrapped 95% confidence intervals. The full learning curves are provided in Figure 10.

In the noise-free setting, DSDPG (MSMMD Sobolev and MMD Sobolev) match the performance of all baselines across five MuJoCo tasks. Under multiplicative observation noise, it outperforms every competitor in three of the five environments, most notably Ant-v2 and Humanoid-v2, while deterministic Sobolev (MAGE) suffers severe drops on Walker2d-v2 and Humanoid-v2 and shows larger variance in the simple InvertedDoublePendulum-v2. Under Gaussian noise, DSDPG again outperforms the baselines in three of the five environments. Notably, on Ant-v2 the noise scale borrowed from Khraishi & Okhrati (2023) makes the task substantially more difficult than previously reported. Additional experiments (varying noise scale, kernel bandwidth, number of samples N, and world model size) are presented in Appendix N. We provide a wall-clock time comparison in Table 3.

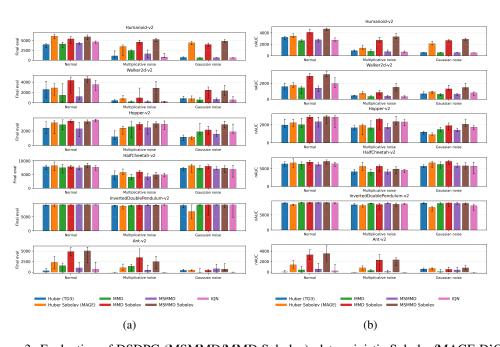


Figure 3: Evaluation of DSDPG (MSMMD/MMD Sobolev), deterministic Sobolev/MAGE D'Oro & Jaskowski (2020), TD3-Huber Fujimoto et al. (2018), IQN Dabney et al. (2018a), standard MMD Nguyen et al. (2020); Killingberg & Langseth (2023) and MSMMD on six MuJoCo tasks. Results are reported over 10 random seeds with median and 95% bootstrap confidence intervals. We compare three settings: normal environment, multiplicative observation noise, and Gaussian dynamics noise (Luo et al., 2021). (a) Final evaluation performance. (b) Normalized AUC over the entire training curve. Our DSDPG variants (MMD Sobolev and MDMMD Sobolev, red and brown) are on par or better than competing methods, shining especially on harder tasks and under noisy environments.

## 7 RELATED WORKS

Gradient informed training Our work explicitly models gradients in a stochastic manner using neural networks. As such, it can be seen as a distributional extension of Sobolev training Czarnecki et al. (2017) which was already adapted to reinforcement learning in D'Oro & Jaskowski (2020); Garibbo et al. (2024). Gradient modeling in value-based RL dates back to Fairbank (2008). Additionally, our approach shares connections with Physics Informed Neural Networks (PINNs) (Raissi et al., 2017), which approximate physical processes via differential constraints. In particular, uncertainty-aware PINNs (Yang et al., 2020; Zhong & Meidani, 2023) treat random processes and their derivatives as random variables, akin to our approach of leveraging Sobolev inductive bias and generative modeling (through reparameterization) while enforcing consistency via a tractable distributional distance.

**Distrbutional RL** We extend distributional RL (Bellemare et al., 2017; 2023) by modeling return gradients for deterministic policies in continuous action spaces, building on distributional DDPG (Barth-Maron et al., 2018; Lillicrap et al., 2016). Because these gradients are multi-dimensional, our method aligns with distributional multivariate returns (Zhang et al., 2021; Freirich et al., 2019; Wiltzer et al., 2024), highlighting the need for tractable discrepancy measures over multi-dimensional distributions. We treat our distributional critic as a generative model capable of producing actual samples of the underlying distribution (Freirich et al., 2019; Singh et al., 2022; Doan et al., 2018), in contrast to approaches that generate only pseudo-samples (Zhang et al., 2021; Nguyen et al., 2020) or rely on summary statistics (Bellemare et al., 2017; Barth-Maron et al., 2018; Dabney et al., 2018b;a). We measure distributional discrepancy via the Mean Maximum Discrepancy (MMD) (Gretton et al., 2012; Li et al., 2015; Bińkowski et al., 2021; Oskarsson, 2020), which has proven effective in distributional RL (Nguyen et al., 2020; Killingberg & Langseth, 2023; Zhang et al., 2021).

**Model-based RL** Finally, because environment dynamics and rewards are unknown and non-differentiable, we adopt a world model akin to SVG(1) (Heess et al., 2015), instantiated as a cVAE (Kingma, 2013; Sohn et al., 2015). Generating new data through this model places our approach within model-based RL (Chua et al., 2018; Feinberg et al., 2018), particularly the *Dyna* family (Sutton, 1991). More specifically, our method is related to approaches that rely on variational techniques (Ha & Schmidhuber, 2018; Hafner et al., 2020; Zhu et al., 2024) or backpropagate through world models (Hafner et al., 2019; Clavera et al., 2020; Amos et al., 2021; Henaff et al., 2017; 2019; Byravan et al., 2020).

#### 8 CONCLUSION, LIMITATIONS AND FUTURE WORK

In this work, we introduced Distributional Sobolev Deterministic Policy Gradient (DSDPG). Our main contributions involve modeling a distribution over both the output and the gradient of a random function, and deriving a tractable computational scheme to achieve this using Maximum Mean Discrepancy (MMD) and its max-sliced variant MSMMD. We demonstrated the effectiveness of the method in leveraging gradient information (Appendix K). We further extended the approach to reinforcement learning by leveraging a differentiable world model of the environment that infers gradients from observations. Building on this, we proposed the framework of Sobolev Temporal Difference, allowing us to provide the first contraction results in gradient-aware RL. We established contraction for both Wasserstein and the tractable MSMMD divergences, and highlighted smoothness assumptions required for contractive gradient-aware training. Finally, we showed that distributional gradient modeling improves stability in a controlled toy problem and enhances robustness to noise in MuJoCo environments, with stronger benefits in high-dimensional tasks.

While our approach shows promise, there remain challenges to address. A primary consideration is the high computational cost, as both policy evaluation and improvement require multiple samples from the distributional critic and their input gradients. Future work should explore more efficient inductive biases.

Finally, we believe that the ideas introduced in this work could benefit other fields where aleatoric uncertainty in gradient modeling is important, such as Physics-Informed Neural Networks (Yang & Perdikaris, 2019), Neural Volume Rendering (Lindell et al., 2021) and more generally applications of double backpropagation (Drucker & Le Cun, 1991).

#### REFERENCES

- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Verlag, Berlin, Heidelberg, 3 edition, 2006. ISBN 978-3-540-29586-0. doi: 10.1007/3-540-29587-9. URL http://dx.doi.org/10.1007/3-540-29587-9.
- Brandon Amos, Samuel Stanton, Denis Yarats, and Andrew Gordon Wilson. On the model-based stochastic value gradient for continuous reinforcement learning. In *Learning for Dynamics and Control*, pp. 6–20. PMLR, 2021.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.
- Philip J Ball and Stephen J Roberts. Offcon 3: What is state of the art anyway? *arXiv preprint* arXiv:2101.11331, 2021.
- Gabriel Barth-Maron, Matthew W. Hoffman, David Budden, Will Dabney, Dan Horgan, Dhruva TB, Alistair Muldal, Nicolas Heess, and Timothy Lillicrap. Distributed distributional deterministic policy gradients. In *International Conference on Learning Representations (ICLR)*, Vancouver, Canada, 2018.
- Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18(153):1–43, 2018. URL http://jmlr.org/papers/v18/17-468.html.
- Marc G. Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 449–458. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/bellemare17a.html.
- Marc G Bellemare, Will Dabney, and Mark Rowland. *Distributional reinforcement learning*. MIT Press, 2023.
- Mikołaj Bińkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans, 2021.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1613–1622, Lille, France, 07–09 Jul 2015. PMLR. URL https://proceedings.mlr.press/v37/blundell15.html.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/jax-ml/jax.
- Arunkumar Byravan, Jost Tobias Springenberg, Abbas Abdolmaleki, Roland Hafner, Michael Neunert, Thomas Lampe, Noah Siegel, Nicolas Heess, and Martin Riedmiller. Imagined value gradients: Model-based policy optimization with tranferable latent dynamics models. In *Conference on Robot Learning*, pp. 566–589. PMLR, 2020.
- Shreyas Chaudhari, Srinivasa Pranav, and José MF Moura. Gradient networks. *arXiv preprint arXiv:2404.07361*, 2024.
- Xinyue Chen, Che Wang, Zijian Zhou, and Keith Ross. Randomized ensembled double q-learning: Learning fast without a model. *arXiv preprint arXiv:2101.05982*, 2021.
- Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper\_files/paper/2018/file/3de568f8597b94bda53149c7d7f5958c-Paper.pdf.

- Ignasi Clavera, Violet Fu, and Pieter Abbeel. Model-augmented actor-critic: Backpropagating through paths. *arXiv preprint arXiv:2005.08068*, 2020.
  - Wojciech M. Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Swirszcz, and Razvan Pascanu. Sobolev training for neural networks. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper\_files/paper/2017/file/758a06618c69880a6cee5314ee42d52f-Paper.pdf.
  - Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning*, pp. 1096–1105. PMLR, 2018a.
  - Will Dabney, Mark Rowland, Marc Bellemare, and Rémi Munos. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018b.
  - Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472, 2011.
  - Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, and Alexander G Schwing. Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10648–10656, 2019.
  - Thang Doan, Bogdan Mazoure, and Clare Lyle. Gan q-learning. *arXiv preprint arXiv:1805.04874*, 2018.
  - Pierluca D'Oro and Wojciech Jaskowski. How to learn a useful critic? model-based action-gradient-estimator policy optimization. In *34th Conference on Neural Information Processing Systems* (*NeurIPS 2020*), Vancouver, Canada, 2020.
  - Harris Drucker and Yann Le Cun. Double backpropagation increasing generalization performance. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*, volume 2, pp. 145–150. IEEE, 1991.
  - Michael Fairbank. Reinforcement learning by value gradients. arXiv preprint arXiv:0803.3539, 2008.
  - Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I Jordan, Joseph E Gonzalez, and Sergey Levine. Model-based value estimation for efficient model-free reinforcement learning. arXiv preprint arXiv:1803.00101, 2018.
  - Dror Freirich, Ron Meir, and Aviv Tamar. Distributional multivariate policy evaluation and exploration with the bellman gan. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 2361–2370, Long Beach, California, 2019. PMLR. \*Equal contribution.
  - Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.
  - Michele Garibbo, Maxime Robeyns, and Laurence Aitchison. Taylor td-learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Aude Genevay, Lénaic Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of sinkhorn divergences. In *The 22nd international conference on artificial intelligence and statistics*, pp. 1574–1583. PMLR, 2019.
  - Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.

- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. URL http://jmlr.org/papers/v13/gretton12a.html.
  - Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.
  - David Ha and Jürgen Schmidhuber. World models. arXiv preprint arXiv:1803.10122, 2018.
  - Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
  - Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
  - Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
  - Hado Hasselt. Double q-learning. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta (eds.), *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010. URL https://proceedings.neurips.cc/paper\_files/paper/2010/file/091d584fced301b442654dd8c23b3fc9-Paper.pdf.
  - Nicolas Heess, Gregory Wayne, David Silver, Timothy Lillicrap, Tom Erez, and Yuval Tassa. Learning continuous control policies by stochastic value gradients. *Advances in neural information processing systems*, 28, 2015.
  - Mikael Henaff, William F Whitney, and Yann LeCun. Model-based planning with discrete and continuous actions. *arXiv preprint arXiv:1705.07177*, 2017.
  - Mikael Henaff, Alfredo Canziani, and Yann LeCun. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. *arXiv preprint arXiv:1901.02705*, 2019.
  - Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning, 2017. URL https://arxiv.org/abs/1710.02298.
  - Max Jaderberg, Wojciech Marian Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, and Koray Kavukcuoglu. Decoupled neural interfaces using synthetic gradients. In *International conference on machine learning*, pp. 1627–1635. PMLR, 2017.
  - Raad Khraishi and Ramin Okhrati. Simple noisy environment augmentation for reinforcement learning, 2023. URL https://arxiv.org/abs/2305.02882.
  - Ludvig Killingberg and Helge Langseth. The multiquadric kernel for moment-matching distributional reinforcement learning. *Transactions on Machine Learning Research*, 2023.
  - Diederik P Kingma. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
  - Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
  - Anders Krogh and John Hertz. A simple weight decay can improve generalization. In J. Moody, S. Hanson, and R.P. Lippmann (eds.), *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann, 1991. URL https://proceedings.neurips.cc/paper\_files/paper/1991/file/8eefcfdf5990e441f0fb6f3fad709e21-Paper.pdf.
  - Arsenii Kuznetsov, Pavel Shvechikov, Alexander Grishin, and Dmitry Vetrov. Controlling overestimation bias with truncated mixture of continuous distributional quantile critics, 2020. URL https://arxiv.org/abs/2005.04269.
    - Yujia Li, Kevin Swersky, and Rich Zemel. Generative moment matching networks. In *International conference on machine learning*, pp. 1718–1727. PMLR, 2015.

- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa,
   David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, 2016.
  - David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14556–14565, 2021.
  - Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
  - Yudong Luo, Guiliang Liu, Haonan Duan, Oliver Schulte, and Pascal Poupart. Distributional reinforcement learning with monotonic splines. In *International Conference on Learning Representations*, 2021.
  - Anton Mallasto, Guido Montúfar, and Augusto Gerolin. How well do wgans estimate the wasserstein metric? *arXiv preprint arXiv:1910.03875*, 2019.
  - Takeru Miyato, Daisuke Okanohara, Shin-ichi Maeda, and Masanori Koyama. Synthetic gradient methods with virtual forward-backward networks. 2017.
  - Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
  - Tetsuro Morimura, Masashi Sugiyama, Hisashi Kashima, Hirotaka Hachiya, and Toshiyuki Tanaka. Nonparametric return distribution approximation for reinforcement learning. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 799–806, 2010.
  - Kimia Nadjahi, Alain Durmus, Lénaïc Chizat, Soheil Kolouri, Shahin Shahrampour, and Umut Simsekli. Statistical and topological properties of sliced probability divergences. *Advances in Neural Information Processing Systems*, 33:20802–20812, 2020.
  - Thanh Tang Nguyen, Sunil Gupta, and Svetha Venkatesh. Distributional reinforcement learning with maximum mean discrepancy. *Association for the Advancement of Artificial Intelligence (AAAI)*, 2020.
  - Joel Oskarsson. Probabilistic regression using conditional generative adversarial networks, 2020.
  - Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
  - Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017.
  - Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv* preprint arXiv:1710.05941, 2017.
  - Mark Rowland, Robert Dadashi, Saurabh Kumar, Rémi Munos, Marc G Bellemare, and Will Dabney. Statistics and samples in distributional reinforcement learning. In *International Conference on Machine Learning*, pp. 5528–5536. PMLR, 2019.
  - Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *The annals of statistics*, pp. 2263–2291, 2013.
  - David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pp. 387–395. Pmlr, 2014.

- Rahul Singh, Keuntaek Lee, and Yongxin Chen. Sample-based distributional policy gradient. In
   Learning for Dynamics and Control Conference, pp. 676–688. PMLR, 2022.
  - Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015.
    - Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. URL http://jmlr.org/papers/v15/srivastava14a.html.
    - Jan Stanczuk, Christian Etmann, Lisa Maria Kreusser, and Carola-Bibiane Schönlieb. Wasserstein gans work because they fail (to approximate the wasserstein distance). *arXiv* preprint *arXiv*:2103.01678, 2021.
    - Ke Sun, Yingnan Zhao, Wulong Liu, Bei Jiang, and Linglong Kong. Distributional reinforcement learning with regularized wasserstein loss. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
    - Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3:9–44, 1988.
    - Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.
    - Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
    - Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, et al. Gymnasium: A standard interface for reinforcement learning environments. *arXiv preprint arXiv:2407.17032*, 2024.
    - Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. *CoRR*, abs/1509.06461, 2015. URL http://arxiv.org/abs/1509.06461.
    - Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
    - Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
    - Harley Wiltzer, Jesse Farebrother, Arthur Gretton, and Mark Rowland. Foundations of multivariate distributional reinforcement learning, 2024. URL https://arxiv.org/abs/2409.00328.
    - Liu Yang, Dongkun Zhang, and George Em Karniadakis. Physics-informed generative adversarial networks for stochastic differential equations. *SIAM Journal on Scientific Computing*, 42(1): A292–A317, 2020.
    - Yibo Yang and Paris Perdikaris. Adversarial uncertainty quantification in physics-informed neural networks. *Journal of Computational Physics*, 394:136–152, 2019.
    - Pushi Zhang, Xiaoyu Chen, Li Zhao, Wei Xiong, Tao Qin, and Tie-Yan Liu. Distributional reinforcement learning for multi-dimensional reward functions. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, 2021. \*Equal contribution, †Corresponding author.
  - Ruijie Zheng, Xiyao Wang, Huazhe Xu, and Furong Huang. Is model ensemble necessary? model-based rl via a single model with lipschitz regularized value function. *arXiv preprint arXiv:2302.01244*, 2023.
- Weiheng Zhong and Hadi Meidani. Pi-vae: Physics-informed variational auto-encoder for stochastic differential equations. *Computer Methods in Applied Mechanics and Engineering*, 403:115664, 2023.
  - Ting Zhu, Ruibin Ren, Yukai Li, and Wenbin Liu. A model-based reinforcement learning method with conditional variational auto-encoder. *Journal of Data Science and Intelligent Systems*, 2024.

#### APPENDIX CONTENTS A Derivatives of reparameterizable random variables **B** Sobolev Inductive Bias and Distribution Parametrization C Sobolev Bellman operator D Results with Wasserstein Practical difficulties with Wasserstein for training F Background on MMD G Results with Max Sliced MMD H Results with Max-Sliced MMD (MSMMD) **Background on the world-model** Algorithm design, pseudo-code and baseline J.1 J.2 J.3 K Toy supervised learning **Toy Reinforcement Learning**

M I	Rein	forcement learning experiments		
N	M.1 Full curves and wall-clock time			
N A	Additional experiments - RL			
N	N.1	Multiquadric kernel bandwidth		
N	N.2	Noise scale		
N	N.3	Number of samples		
N	<b>V</b> .4	World model capacity		

#### A DERIVATIVES OF REPARAMETERIZABLE RANDOM VARIABLES

In several places (e.g. Eq. equation 9, Section 3.2) we need to talk about "the gradient of a random variable with respect to its conditioning variable." To make this precise and intuitive, we rely on the common *reparameterization trick*. We explain that defining a random variable this way leads to a well-defined random variable. We explain that, under appropriate continuity and measurability assumptions, this construction yields a well-defined random variable.

**Reparameterization view.** Suppose  $Y \mid x$  is a conditional random variable, where  $x \in \mathcal{X} \subset \mathbb{R}^n$ . We assume there exists a measurable "noise" variable  $\Xi \sim p(\xi)$  on a space  $\mathcal{Z}$  and a deterministic map

$$q: \mathcal{X} \times \mathcal{Z} \longrightarrow \mathcal{Y}$$

that is  $C^1$  in its first argument, such that

$$Y = g(x, \Xi)$$
 has the same law as  $Y \mid x$ .

In other words, for each sample  $\xi$ , the realization

$$y(x;\xi) = g(x,\xi)$$

is a deterministic function of x.

**Pathwise derivative.** Because  $g(\cdot, \xi)$  is differentiable in x, we define the *pathwise* or *reparameterization* derivative by

$$\frac{\partial}{\partial x} Y \stackrel{\text{def}}{=} \frac{\partial}{\partial x} g(x, \Xi),$$

i.e. for each fixed  $\xi$ ,  $\partial_x y(x;\xi) = \partial_x g(x,\xi)$ . Since  $\Xi$  is random,  $\partial_x g(x,\Xi)$  is itself a random variable.

**Measurability.** Because  $g \in C^1$  in its first argument, for each fixed  $\xi$  the map

$$x \longmapsto \partial_x g(x,\xi)$$

is continuous in x. Moreover, by assumption  $\xi \mapsto g(x,\xi)$  is measurable for each x. It follows from Lemma 4.51 from Aliprantis & Border (2006) ("Carathéodory functions are jointly measurable") that

$$(x,\xi) \longmapsto \partial_x g(x,\xi)$$

is jointly Borel-measurable on  $\mathcal{X} \times \mathcal{Z}$ . Consequently, when  $\Xi \sim p(\xi)$  the composition  $\partial_x g(x,\Xi)$  is a well-defined  $\mathbb{R}^n$ -valued random variable.

**Practical computation.** In modern autodiff frameworks, once you write your generative model as  $Y=g(x,\xi)$ , calling  $\operatorname{grad}(Y,x)$  automatically returns the sample-wise Jacobian  $\partial_x g(x,\xi)$ . Hence all "random variable gradients" in our Sobolev Bellman operator are tractable to compute using standard reparameterization and automatic differentiation techniques.

## B SOBOLEV INDUCTIVE BIAS AND DISTRIBUTION PARAMETRIZATION

#### B.1 SOBOLEV INDUCTIVE BIAS

We discussed the Sobolev inductive bias in Section 5. Recall that given a differentiable target  $F \colon \mathbb{R}^a \to \mathbb{R}^b$ , Sobolev training (Czarnecki et al., 2017) (originally introduced as "double back-propagation" by Drucker & Le Cun (1991)) minimizes both value- and gradient-mismatch:

$$\mathcal{L}^{S}(\varphi; x) = \|F(x) - f_{\varphi}(x)\|^{2} + \lambda_{S} \|\nabla_{x} F(x) - \nabla_{x} f_{\varphi}(x)\|^{2}.$$

The first term aligns network outputs to F(x), while the second enforces that the input-derivatives of  $f_{\varphi}$  match those of F. Below, we build on this bias to derive conditions on valid gradient fields and their relation to Jacobian symmetry.

Using the gradient of a neural network to predict gradients is not the only way to do so. Indeed, Jaderberg et al. (2017) proposed to generate *synthetic gradients* directly via the outputs of a neural network or a single linear layer. However, Czarnecki et al. (2017); Miyato et al. (2017) observed that predicting synthetic gradients as gradients of a neural network directly improved performance on decoupled trainings experiments from Jaderberg et al. (2017). A first hint for a justification of this observation was provided in Proposition 4 from Czarnecki et al. (2017) that gives a necessary condition for a function  $f_{\varphi}(x): \mathbb{R}^n \to \mathbb{R}^n$  to produce a valid vector gradient field. It states that if  $f_{\varphi}(x)$  produces a valid gradient vector field of some potential function  $\phi(x)$  then the Jacobian matrix  $\frac{\partial f_{\varphi}}{\partial x}(x): \mathbb{R}^n \to \mathbb{R}^n$  must be symmetric.

This can be readily seen from the observation that if  $\frac{\partial \phi}{\partial x}$  is the gradient vector field to be approximated by  $f_{\varphi}(x)$  then  $\frac{\partial^2 \phi}{\partial \phi^2}(x)$  is the Hessian matrix that is known to be symmetric.

A more in depth discussion is provided in Chaudhari et al. (2024) where a **necessary and sufficient** condition is given in Lemma 1 that states that a differentiable function  $f: \mathbb{R}^n \to \mathbb{R}^n$  has a scalar-valued anti-derivative if and only if its Jacobian is symmetric everywhere, i.e.,  $\forall x \in \mathbb{R}^n \frac{\partial f}{\partial x}(x) = \left(\frac{\partial f}{\partial x}(x)\right)^T$ . **This suggests a good inductive bias should be to restrict**  $f_{\varphi}(x)$  **to the class of function with symmetric Jacobian.** This condition is directly satisfied via the Sobolev inductive bias as well as gradient networks proposed in Chaudhari et al. (2024). As pointed out in Czarnecki et al. (2017), this is a very unlikely condition to be met by a regular feedforward neural network that would predict the gradient directly as its output.

#### B.2 DISTRIBUTION PARAMETRIZATION

The Sobolev inductive bias introduced in Section 5 and B.1 requires a generative model where both the output and its input-gradient are treated as random variables. While the reparameterization trick (Kingma, 2013) could be used with a conditional Gaussian distribution, determining how to distribute the gradient of the samples with respect to the conditioners is less straightforward. To address this, we employ a method that relies solely on sampled data and does not assume tractable density estimation of Sobolev random variables.

Moreover, we found that the most common distribution parametrizations, namely discrete categorical (Bellemare et al., 2017; Barth-Maron et al., 2018) and quantile-based (Dabney et al., 2018b;a) do not scale tractably to higher dimensions, specifically in the way they instantiate Eq.8. These considerations further motivated us to adopt a sample-based approach for the distributional Sobolev critic, similar in spirit to Singh et al. (2022); Freirich et al. (2019). As a result, our distributional critic is structured as a generative model that deterministically maps noise to samples (Li et al., 2015; Goodfellow et al., 2014).

## C SOBOLEV BELLMAN OPERATOR

In this section we first specify our smoothness and boundedness assumptions on the return, transition, reward, and policy functions in the *Preamble and assumptions* (Section C.1). Then derive how differentiating the distributional Bellman equation yields the action-gradient update in the *Derivation of the action-gradient term* (Section C.1). Next, we show how to bundle the return and its

gradient into a single affine update rule in the *Distributional Sobolev Bellman update* (Eq. 17). We follow and give both the pathwise and integral forms of the full operator in the *Distributional Sobolev Bellman operator* (Eqs. 19 and 20). Finally, we explain how to include state-gradients for a fully Sobolev-compliant backup in the *Complete Sobolev Bellman operator* (Section C.5).

#### C.1 PREAMBLE AND ASSUMPTIONS

Unlike traditional Bellman operators acting on scalar- or fixed-dimensional vector-valued functions, our operator acts on the space of continuously differentiable, bounded functions with bounded first derivatives over the compact domain  $\mathcal{S} \times \mathcal{A}$ . Concretely, we assume the return-distribution admits a reparameterization

$$Z(s, a; \varepsilon_Z) \in C_b^1(\mathcal{S} \times \mathcal{A}),$$

where  $\varepsilon_Z \sim p(\varepsilon_Z)$  is exogenous noise, so that each sample  $Z(s,a;\varepsilon_Z)$  and its gradients  $\nabla_{s,a}Z(s,a;\varepsilon_Z)$  are bounded and continuous on  $\mathcal{S}\times\mathcal{A}$ .

Let  $\varepsilon_f \sim p(\varepsilon_f)$ ,  $\varepsilon_r \sim p(\varepsilon_r)$ , and  $\varepsilon_\pi \sim p(\varepsilon_\pi)$  be independent noise variables on  $\mathcal{Z}$ , and let

$$f: \mathcal{S} \times \mathcal{A} \times \mathcal{Z} \to \mathcal{S}$$
, (reparameterized transition)  $r: \mathcal{S} \times \mathcal{A} \times \mathcal{Z} \to \mathbb{R}$ , (reparameterized reward)

$$\pi: \mathcal{S} \times \mathcal{Z} \to \mathcal{A},$$
 (reparameterized policy)

be  $C^1$  maps such that, for each draw  $\varepsilon_f, \varepsilon_r, \varepsilon_\pi, \varepsilon_Z$ ,

$$s' = f(s, a; \varepsilon_f), \quad r = r(s, a; \varepsilon_r), \quad a' = \pi(s'; \varepsilon_\pi), \quad Z = Z(s, a; \varepsilon_Z).$$

Importantly, we assume that the gradients of all these  $C^1$  mappings  $(Z, f, r, and \pi)$  are tractable to compute in practice (e.g., via automatic differentiation).

We further assume these maps have uniformly bounded Jacobians: there exist constants  $L_{f,s}, L_{f,a}, L_r, L_\pi, L_Z < \infty$  such that

$$\sup_{(s,a,\varepsilon_f)} \left\| \frac{\partial f}{\partial s}(s,a;\varepsilon_f) \right\| \le L_{f,s}, \qquad \sup_{(s,a,\varepsilon_f)} \left\| \frac{\partial f}{\partial a}(s,a;\varepsilon_f) \right\| \le L_{f,a},$$

$$\sup_{(s,a,\varepsilon_r)} \left\| \frac{\partial r}{\partial a}(s,a;\varepsilon_r) \right\| \le L_{r}, \qquad \sup_{(s,\varepsilon_\pi)} \left\| \frac{\partial \pi}{\partial s}(s;\varepsilon_\pi) \right\| \le L_{\pi},$$

## $\sup_{(s,a,\varepsilon_Z)} \left\| \nabla_{s,a} Z(s,a;\varepsilon_Z) \right\| \le L_Z.$

#### C.2 DERIVATION OF THE ACTION-GRADIENT TERM

Starting from the distributional Bellman equation for each noise draw  $(\varepsilon_f, \varepsilon_r, \varepsilon_Z, \varepsilon_\pi)$ :

$$Z(s, a; \varepsilon_Z) = r(s, a; \varepsilon_r) + \gamma Z(s', a'; \varepsilon_Z'),$$

where  $s' = f(s, a; \varepsilon_f)$ ,  $a' = \pi(s'; \varepsilon_\pi)$ , and  $\varepsilon'_Z \sim p(\varepsilon_Z)$  independently. Differentiate w.r.t. a:

$$\frac{\partial}{\partial a}Z(s,a;\varepsilon_Z) = \frac{\partial}{\partial a}r(s,a;\varepsilon_r) + \gamma \frac{d}{da} \big[Z\big(s',a';\varepsilon_Z'\big)\big].$$

By the chain-rule,

$$\frac{d}{da}Z(s',a';\varepsilon_Z') = \underbrace{\left(f_a(s,a;\varepsilon_f)\right)^T}_{\partial s'/\partial a} \frac{\partial}{\partial s}Z(s',a';\varepsilon_Z') + \underbrace{\left(f_a(s,a;\varepsilon_f)\right)^T \left(\pi_s(s';\varepsilon_\pi)\right)^T}_{\partial a'} \frac{\partial}{\partial a}Z(s',a';\varepsilon_Z'),$$

where we have introduced the shorthand

$$f_a(s, a; \varepsilon_f) = \frac{\partial f}{\partial a}(s, a; \varepsilon_f), \qquad \pi_s(s'; \varepsilon_\pi) = \frac{\partial \pi}{\partial s}(s'; \varepsilon_\pi).$$

Plugging back (and evaluating at  $s' = f(s, a; \varepsilon_f), a' = \pi(s'; \varepsilon_\pi)$ ) yields

$$\begin{split} \frac{\partial}{\partial a} Z(s, a; \varepsilon_Z) &= \frac{\partial}{\partial a} r(s, a; \varepsilon_r) + \\ & \gamma \left( f_a(s, a; \varepsilon_f) \right)^T \left[ \frac{\partial}{\partial s} Z(s', a'; \varepsilon_Z') + \left( \pi_s(s'; \varepsilon_\pi) \right)^T \frac{\partial}{\partial a} Z(s', a'; \varepsilon_Z') \right], \end{split}$$

which becomes the action-gradient component in our subsequent block-affine formulation.

#### C.3 DISTRIBUTIONAL SOBOLEV BELLMAN UPDATE

In Section 3.2 we introduced a Bellman-style update that propagates both the random return and its action-gradient. We now rephrase it in our block-affine form.

Define the stacked random-vector, which by definition coincides with the joint Sobolev return introduced in the main text:

$$Z^{S_a}(s, a; \varepsilon_Z) := \begin{pmatrix} Z(s, a; \varepsilon_Z) \\ \frac{\partial}{\partial a} Z(s, a; \varepsilon_Z) \end{pmatrix} = H(s, a; \varepsilon_Z) \in \mathbb{R}^{1+m}.$$

where  $\varepsilon_Z \sim p(\varepsilon_Z)$ , and let  $\varepsilon_f, \varepsilon_r, \varepsilon_\pi$  be independent noise variables for transition, reward, and policy. For each draw  $(\varepsilon_f, \varepsilon_r, \varepsilon_\pi, \varepsilon_Z)$  define

$$s' = f(s, a; \varepsilon_f), \quad r = r(s, a; \varepsilon_r), \quad a' = \pi(s'; \varepsilon_\pi).$$

Then the sample-wise update is

$$H(s, a; \varepsilon_{Z}) = \underbrace{\begin{pmatrix} r(s, a; \varepsilon_{r}) \\ \frac{\partial}{\partial a} r(s, a; \varepsilon_{r}) \end{pmatrix}}_{b(s, a; \varepsilon_{r})} + \underbrace{\begin{pmatrix} \gamma & 0_{1 \times m} \\ 0_{m \times 1} & \gamma \left(\frac{\partial f}{\partial a}(s, a; \varepsilon_{f})\right)^{T} \left(\frac{\partial \pi}{\partial s}(s'; \varepsilon_{\pi})\right)^{T} \end{pmatrix}}_{A(s, a; \varepsilon_{f}, \varepsilon_{\pi})} H(s', a'; \varepsilon'_{Z})$$

$$+ \underbrace{\begin{pmatrix} 0_{1 \times n} \\ \gamma \left(\frac{\partial f}{\partial a}(s, a; \varepsilon_{f})\right)^{T} \end{pmatrix}}_{N(s, a; \varepsilon_{f})} \frac{\partial}{\partial s} Z(s', a'; \varepsilon'_{Z}), \tag{17}$$

where for the next-step noise we write  $(\varepsilon_f, \varepsilon_\pi, \varepsilon_Z') \sim p(\varepsilon_f) p(\varepsilon_\pi) p(\varepsilon_Z)$  independently. We abbreviate

$$b(s, a; \varepsilon_r) = \begin{pmatrix} r(s, a; \varepsilon_r) \\ \partial_a r(s, a; \varepsilon_r) \end{pmatrix}$$

as above.

Next, introduce the *state-gradient operator* 

$$\mathcal{D}_s: H \mapsto \frac{\partial}{\partial s} [e_1^T H] = \frac{\partial Z}{\partial s}.$$

where  $e_1^T = (1, 0, \dots, 0)$  selects the return component. For each  $(s, a; \varepsilon_f)$ , define the combined linear map

$$\mathcal{L}(s, a; \varepsilon_f, \varepsilon_\pi) = A(s, a; \varepsilon_f, \varepsilon_\pi) + N(s, a; \varepsilon_f) \mathcal{D}_s|_{(s', a'; \varepsilon_Z')}$$

Then the distributional Sobolev Bellman update can be written as the single (pseudo) affine transform

$$H(s, a; \varepsilon_Z) = b(s, a; \varepsilon_r) + \mathcal{L}(s, a; \varepsilon_f, \varepsilon_\pi) [H(s', a'; \varepsilon_Z')].$$
(18)

**Remark (i).** In this form, the computation of the next-state gradient  $\frac{\partial}{\partial s} Z(s', a'; \varepsilon'_Z)$  is not an external "on-the-fly" step but is captured by the linear operator  $\mathcal{L}(s, a; \varepsilon_f, \varepsilon_\pi)$  via  $\mathcal{D}_s$ . Thus computing the state-gradient of the distributional return is an integral part of the Bellman operator itself.

**Remark** (ii). Although we write the update in a single "block-affine" line, strictly speaking it does not close on the 2-vector  $(Z, \partial_a Z)$  alone so we refer to it as **pseudo-affine**. Indeed, there is still an  $N \partial_s Z$  term lurking outside of H. We therefore call it affine only in a loose, not literal, sense; the genuine finite-dimensional affine form appears once we lift to the full Sobolev vector  $(Z, \partial_a Z, \partial_s Z)$  in Section C.5.

**Measurability**. By the same argument used in Appendix A, the map  $(s, a, \varepsilon) \mapsto \partial_s Z(s, a; \varepsilon)$  is jointly Borel-measurable.

#### C.4 DISTRIBUTIONAL SOBOLEV BELLMAN OPERATOR

Let  $\eta: \mathcal{S} \times \mathcal{A} \longrightarrow \mathcal{P}(\mathbb{R}^{1+m})$  be the law of  $H(s, a; \varepsilon_Z) = [Z(s, a; \varepsilon_Z), \partial_a Z(s, a; \varepsilon_Z)]$ , and assume three independent noise variables  $\varepsilon_f \sim p_f, \varepsilon_r \sim p_r, \varepsilon_\pi \sim p_\pi$  generate

$$s' = f(s, a; \varepsilon_f), \quad r = r(s, a; \varepsilon_r), \quad a' = \pi(s'; \varepsilon_\pi).$$

We define the Sobolev distributional Bellman operator  $T_{\pi}^{S_a}$  in two equivalent ways:

#### Pathwise (law-equality) form:

$$(T_{\pi}^{S_a} \eta)(s, a) = \text{Law} \Big[ b(s, a; \varepsilon_r, \varepsilon_{\pi}) + \mathcal{L}(s, a; \varepsilon_f) [X'] \Big],$$
 (19)

where  $X' \sim \eta(s', a')$  after sampling  $\varepsilon_f, \varepsilon_r, \varepsilon_\pi$  and setting  $s' = f(s, a; \varepsilon_f), a' = \pi(s'; \varepsilon_\pi)$ .

## **Explicit-noise integral (pushforward) form:**

$$\left(T_{\pi}^{S_a} \eta\right)(s, a) = \int_{\varepsilon_f} \int_{\varepsilon_{\pi}} \int_{\varepsilon_r} \left(b(s, a; \varepsilon_r) + \mathcal{L}(s, a; \varepsilon_f)\right)_{\#} \eta \left(f(s, a; \varepsilon_f), \pi(f(s, a; \varepsilon_f); \varepsilon_{\pi})\right) \times p_r(d\varepsilon_r) \times p_{\pi}(d\varepsilon_{\pi}) \times p_f(d\varepsilon_f).$$
(20)

## C.5 COMPLETE SOBOLEV BELLMAN OPERATOR

In Section 3.2 and above we developed the action-gradient Sobolev Bellman backup, in which our bootstrapped object is the pair

$$Z^{S_a}(s,a) = (Z(s,a), \partial_a Z(s,a)) \in \mathbb{R}^{1+m}.$$

Because the Bellman operator needs the next-step state-gradient  $\partial_s Z(s',a')$ , which is not included in  $Z^{S_a}$ , we must recover it "on-the-fly" by differentiating the map Z. Equivalently, we have to own the differentiable function Z itself at implementation time and thus rely on a Sobolev inductive bias as introduced in Appendix B.1.

**Lifting to a full Sobolev return.** Nothing prevents us from treating the state-gradient as a *third* component of our bootstrapped vector. Define

$$Z^{S_{s,a}}(s,a) = \begin{pmatrix} Z(s,a) \\ \partial_a Z(s,a) \\ \partial_s Z(s,a) \end{pmatrix} \in \mathbb{R}^{1+m+n}.$$

Then, just as before, we obtain a single affine update

$$Z^{S_{s,a}}(s,a) = b(s,a) + \mathcal{L}_{\text{full}}(s,a) \left[ Z^{S_{s,a}}(s',a') \right], \tag{21}$$

where  $b(s,a) \in \mathbb{R}^{1+m+n}$  collects  $(r, \partial_a r, \partial_s r)$  and  $\mathcal{L}_{\text{full}}(s,a)$  is a bounded linear operator on  $\mathbb{R}^{1+m+n}$ .

Using similar derivations as for the action-gradient version we got the vector b and linear operator  $L_{\rm full}$  (which is just a matrix in this case) such that

$$b(s, a; \varepsilon_r) = \begin{pmatrix} r(s, a; \varepsilon_r) \\ \partial_a r(s, a; \varepsilon_r) \\ \partial_s r(s, a; \varepsilon_r) \end{pmatrix}$$

and

$$\mathcal{L}_{\text{full}}(s, a; \varepsilon_f, \varepsilon_\pi) = \gamma \begin{pmatrix} 1 & 0 & 0 \\ 0 & f_a^\top(s, a; \varepsilon_f) \, \pi_s^\top(s'; \varepsilon_\pi) & 0 \\ 0 & f_s^\top(s, a; \varepsilon_f) \, \pi_s^\top(s'; \varepsilon_\pi) & f_s^\top(s, a; \varepsilon_f) \end{pmatrix}.$$

With  $s' = f(s, a; \varepsilon_f)$  and  $a' = \pi(s'; \varepsilon_\pi)$  we denote the complete Sobolev Bellman operator by

$$(T_{\pi}^{S_{s,a}}\eta)(s,a) = \operatorname{Law} \left[ b_{\text{full}}(s,a;\varepsilon_r,\varepsilon_{\pi}) + \mathcal{L}_{\text{full}}(s,a;\varepsilon_f) \left[ X' \right] \right], \qquad X' \sim \eta(s',a').$$
 (22)

Consequences for implementation. In this complete form, the second and third blocks of  $Z^{S_{s,a}}$  need not be tied to the derivatives of the scalar head. One may simply design a network with three output heads

$$(s,a) \mapsto (Z_{\theta}(s,a), G_{\theta}^{a}(s,a), G_{\theta}^{s}(s,a)),$$

and train it against the Sobolev-affine targets in Eq. equation 11. Of course, in practice on can still exploit the Sobolev inductive bias and thus ensuring  $G^a = \partial_a Z$  and  $G^s = \partial_s Z$ .

#### D RESULTS WITH WASSERSTEIN

**Proof of Proposition 3.1:** We derive the upper bound on the policy gradient error in Section D.2.

Theorem 3 shows that the *complete* Sobolev Bellman operator  $T_{\pi}^{S_a}$  contracts by the uniform operator norm  $\|\mathcal{L}\|_d$  in the supremum-p-Wasserstein metric.

Theorem 4 shows a similar result to the *incomplete* operator that backs up only  $(Z, \partial_a Z)$ , yielding the contraction coefficient  $\gamma \kappa_{\text{eff}}$ .

Corollary 5.1 then applies Banach's fixed-point theorem to conclude that the Sobolev iterates  $\eta_{n+1} = T_\pi^S \eta_n$  converge geometrically to the unique fixed point  $\eta^\pi$  identified in Lemma 8.

#### D.1 PROOFS OF CONTRACTION AND FIXED POINT

**Section overview.** In this section we prove that the Sobolev Bellman operator admits a unique fixed point and that iterating the operator converges to that solution. We begin by recalling the definitions of the *p*–Wasserstein distance (Def. 1) and its supremum–*p* variant (Def. 2), and then show in Lemma 2 that the distributional Sobolev backup can be written as an affine push-forward. Two key transport estimates follow: the affine push-forward contraction in any normed space (Lemma 3) and the mixture non-expansion property under conditioning (Lemma 4).

**Definition 1** (p-Wasserstein distance Villani et al. (2009)). Let (X,d) be a metric space and let  $\mathcal{P}_p(X)$  be the set of all probability measures on X with finite pth moment. For any  $\alpha, \beta \in \mathcal{P}_p(X)$ , the p-Wasserstein distance between  $\alpha$  and  $\beta$  is

$$W_p(\alpha, \beta) = \left(\inf_{\pi \in \Pi(\alpha, \beta)} \int_{X \times X} d(x, y)^p d\pi(x, y)\right)^{1/p},$$

where  $\Pi(\alpha, \beta)$  is the set of all couplings of  $\alpha$  and  $\beta$ , i.e. Borel probability measures on  $X \times X$  whose marginals are  $\alpha$  and  $\beta$ , respectively.

**Definition 2** (Supremum-p-Wasserstein distance Zhang et al. (2021), eq.(12)). Let (X,d) be a metric space, let S be the set of states and A the set of actions, and let  $\mu_1, \mu_2 \colon S \times A \to \mathcal{P}_p(X)$  assign to each state-action pair (s,a) a probability measure on X with finite pth moment. The supremum-p-Wasserstein distance between  $\mu_1$  and  $\mu_2$  is

$$\bar{W}_p(\mu_1, \mu_2) = \sup_{(s,a) \in S \times A} W_p(\mu_1(s,a), \mu_2(s,a)),$$

where  $W_p$  is the p-Wasserstein distance on (X, d) as in Definition 1.

**Lemma 1** (Push-forward law identity). Let Z be a random variable with distribution  $\mu$ , and let f be any measurable function. Then

$$f_{\#}\mu = \operatorname{Law}(f(Z)).$$

Proof. For any Borel set A,

1136 
$$\Pr(f(Z) \in A) = \Pr(Z \in f^{-1}(A)) = \mu(f^{-1}(A)) = f_{\#}\mu(A).$$

Since this holds for all A, we conclude  $f_{\#}\mu = \text{Law}(f(Z))$ .

Lemma 2 (Affine form of the Sobolev Bellman operator). Fix (s,a). Draw independent exogenous noise

$$\varepsilon_r \sim p(\varepsilon_r), \quad \varepsilon_f \sim p(\varepsilon_f), \quad \varepsilon_\pi \sim p(\varepsilon_\pi),$$

and set

$$r = r(s, a; \varepsilon_r), \quad s' = f(s, a; \varepsilon_f), \quad a' = \pi(s'; \varepsilon_\pi), \quad X' \sim \eta^{S_a}(s', a').$$

1144 Define the random affine map

$$\Phi_{s,a}ig(x;\,arepsilon_r,arepsilon_f,arepsilon_\piig) \ = \ big(s,a;arepsilon_rig) \ + \ \mathcal{L}ig(s,a;arepsilon_f,arepsilon_\piig)[\,x\,], \ \ \textit{where}\ b = igg(rac{r}{\partial_a r}igg), \ \ \mathcal{L} = A + N\,\mathcal{D}_s$$

as in Eq. equation 18. Then

$$T_{\pi}^{S_a} \eta^{S_a}(s, a) = \operatorname{Law}(\Phi_{s,a}(X')).$$

In other words, the Sobolev Bellman update is an affine pushforward of the joint return–gradient under exogenous noise.

*Proof.* Fix a Borel set  $A \subset \mathbb{R}^{1+m}$ . Writing out the randomness explicitly in terms of  $\varepsilon_r, \varepsilon_f, \varepsilon_\pi$ , we have

$$\Pr(\Phi_{s,a}(X') \in A) = \mathbb{E}_{\varepsilon_r,\varepsilon_f,\varepsilon_\pi} \Big[ \Pr(\Phi_{s,a}(X') \in A \mid \varepsilon_r,\varepsilon_f,\varepsilon_\pi) \Big].$$

Conditioning on  $(\varepsilon_r, \varepsilon_f, \varepsilon_\pi)$  turns  $\Phi_{s,a}$  into the fixed affine map  $x \mapsto b + \mathcal{L}[x]$ , so

$$\Pr(\Phi_{s,a}(X') \in A \mid \varepsilon_r, \varepsilon_f, \varepsilon_\pi) = \Pr(b + \mathcal{L}[X'] \in A).$$

Since  $X' \sim \eta^{S_a}(s', a')$ , an application of Lemma 1

$$\left[ (b + \mathcal{L})_{\#} \eta^{S_a}(s', a') \right] = \text{Law} \left( b + \mathcal{L}(X') \right)$$

gives

$$\Pr(b + \mathcal{L}[X'] \in A) = [(b + \mathcal{L})_{\#} \eta^{S_a}(s', a')](A).$$

Putting these together,

$$\Pr(\Phi_{s,a}(X') \in A) = \mathbb{E}_{\varepsilon_r, \varepsilon_f, \varepsilon_\pi} \Big[ (b + \mathcal{L})_{\#} \eta^{S_a}(s', a')(A) \Big].$$

Finally, by the explicit-noise integral form equation 20, this expectation is exactly the definition of  $(T_{\pi}^{S_a} \eta^{S_a})(s,a)(A)$ . Hence

$$\Pr(\Phi_{s,a}(X') \in A) = (T_{\pi}^{S_a} \eta^{S_a})(s,a)(A),$$

and since this holds for every Borel A, we conclude  $\text{Law}(\Phi_{s,a}(X')) = (T_{\pi}^{S_a} \eta^{S_a})(s,a)$ .

**Lemma 3** (Affine push-forward contraction in a normed space). Let (X, d) be a normed vector space equipped with the metric d induced by its norm, and let

$$F \colon X \to X, \qquad F(x) = b + L[x],$$

where  $b \in X$  is fixed and  $L \colon X \to X$  is a bounded linear operator. Define its Lipschitz constant

$$||L||_d = \sup_{x \neq y} \frac{d(L[x], L[y])}{d(x, y)} < \infty.$$

Then for any two probability measures  $\alpha, \beta$  on X with finite pth moment,

$$W_p(F_\#\alpha, F_\#\beta) \le ||L||_d W_p(\alpha, \beta)$$

In particular, when  $L = \gamma I$  this recovers Lemma 3 of Zhang et al. (2021).

*Proof.* Fix any  $\varepsilon > 0$ , and choose a coupling

that is  $\varepsilon$ -optimal, i.e.  $\pi$  is within  $\varepsilon$  of the infimum:

$$\left(\int_{Y \cup Y} d(x,y)^p d\pi(x,y)\right)^{1/p} < W_p(\alpha,\beta) + \varepsilon.$$

Push this coupling forward under  $F \times F$  to obtain

$$\pi' = (F \times F)_{\#}\pi \in \Pi(F_{\#}\alpha, F_{\#}\beta).$$

 $\pi \in \Pi(\alpha, \beta)$ 

 Then

$$\begin{split} W_p\big(F_\#\alpha,\,F_\#\beta\big)^p &\leq \int_{X\times X} d(u,v)^p\,d\pi'(u,v) & \text{(by definition of }W_p) \\ &= \int_{X\times X} d\big(F(x),F(y)\big)^p\,d\pi(x,y) & (\pi' = (F\times F)_\#\pi) \\ &= \int_{X\times X} d\big(L[x],L[y]\big)^p\,d\pi(x,y) & (F(x)-F(y)=L[x-y]) \\ &\leq \|L\|_d^p \int_{X\times X} d(x,y)^p\,d\pi(x,y) & \left(d(L[x],L[y])\leq \|L\|_d\,d(x,y)\right) \\ &< \|L\|_d^p \big(W_p(\alpha,\beta)+\varepsilon\big)^p & (\pi \text{ is }\varepsilon\text{-optimal}). \end{split}$$

Taking the pth root and letting  $\varepsilon \to 0$  yields

$$W_p(F_{\#}\alpha, F_{\#}\beta) \leq ||L||_d W_p(\alpha, \beta).$$

**Lemma 4** (Mixture non-expansion, conditional form Zhang et al. (2021) - Lemma 4). Let C be a random variable on  $(\Omega, \mathcal{F}, \rho)$ , and let  $Z_1, Z_2$  be  $\mathbb{R}^d$ -valued random variables. Let's have  $p \geq 1$  and suppose there exists  $\delta \geq 0$  such that for every  $c \in \Omega$ ,

$$W_p(\text{Law}(Z_1 \mid C = c), \text{Law}(Z_2 \mid C = c)) \leq \delta.$$

Then the marginal distributions satisfy

$$W_p(\text{Law}(Z_1), \text{Law}(Z_2)) \leq \delta.$$

Equivalently,

$$\sup_{c \in \Omega} W_p(\operatorname{Law}(Z_1 \mid C = c), \operatorname{Law}(Z_2 \mid C = c)) \leq \delta$$

$$\Longrightarrow$$

$$W_p(\operatorname{Law}(Z_1), \operatorname{Law}(Z_2)) \leq \delta$$

In other words, integrating (averaging) over the conditioning index preserves, and never increases, the worst-case Wasserstein distance.

*Proof.* This proof follows exactly Zhang et al. (2021). Fix any  $\varepsilon > 0$ . By assumption, for every  $c \in \Omega$ ,

$$W_p(\operatorname{Law}(Z_1 \mid C = c), \operatorname{Law}(Z_2 \mid C = c)) \le \delta,$$

so there exists a coupling

$$\pi_c \in \Pi(\text{Law}(Z_1 \mid C = c), \text{Law}(Z_2 \mid C = c))$$

$$\int_{\mathbb{R}^d \times \mathbb{R}^d} d(x, y)^p \, d\pi_c(x, y) \leq (\delta + \varepsilon)^p.$$

By the law of total probability, the marginals are

$$\operatorname{Law}(Z_1)(\cdot) = \int_{\Omega} \operatorname{Law}(Z_1 \mid C = c)(\cdot) \, \rho(dc),$$
$$\operatorname{Law}(Z_2)(\cdot) = \int_{\Omega} \operatorname{Law}(Z_2 \mid C = c)(\cdot) \, \rho(dc).$$

Define a global coupling  $\Pi$  on  $\mathbb{R}^d \times \mathbb{R}^d$  by

$$\Pi(U) = \int_{\Omega} \pi_c(U) \, \rho(dc), \qquad U \subseteq \mathbb{R}^d \times \mathbb{R}^d.$$

Then for any measurable  $A \subset \mathbb{R}^d$ ,

$$\Pi(A \times \mathbb{R}^d) = \int_{\Omega} \pi_c(A \times \mathbb{R}^d) \, \rho(dc) = \int_{\Omega} \operatorname{Law}(Z_1 \mid C = c)(A) \, \rho(dc) = \operatorname{Law}(Z_1)(A),$$

and similarly  $\Pi(\mathbb{R}^d \times A) = \text{Law}(Z_2)(A)$ , so  $\Pi \in \Pi(\text{Law}(Z_1), \text{Law}(Z_2))$ .

$$\begin{split} W_p \big( \mathrm{Law}(Z_1), \ \mathrm{Law}(Z_2) \big)^p & \leq \int_{\mathbb{R}^d \times \mathbb{R}^d} d(x,y)^p \ d\Pi(x,y) & \text{(by definition of } W_p) \\ & = \int_{\Omega} \Big[ \int d(x,y)^p \ d\pi_c(x,y) \Big] \rho(dc) & \text{(by definition of } \Pi) \\ & \leq \int_{\Omega} (\delta + \varepsilon)^p \ \rho(dc) & (\pi_c \text{ is } \varepsilon\text{-optimal}) \\ & = (\delta + \varepsilon)^p. \end{split}$$

Therefore,

$$W_p(\text{Law}(Z_1), \text{Law}(Z_2)) \leq (\delta + \varepsilon),$$

and since  $\varepsilon > 0$  was arbitrary, we conclude

$$W_p(\text{Law}(Z_1), \text{Law}(Z_2)) \leq \delta.$$

 **Lemma 5** (Mixture *p*-convexity for  $W_p$ ). Let (X,d) be a metric space,  $p \in [1,\infty)$ , and let  $(\Omega, \mathcal{F}, \rho)$  be a probability space. Let  $(\mu_c)_{c \in \Omega}, (\nu_c)_{c \in \Omega} \subset \mathcal{P}_p(X)$  be measurable families. Then

 $\boxed{W_p \bigg( \int_{\Omega} \mu_c \, \rho(dc), \, \int_{\Omega} \nu_c \, \rho(dc) \bigg) \, \leq \, \Bigg( \int_{\Omega} W_p(\mu_c, \nu_c)^p \, \rho(dc) \Bigg)^{1/p}}.}$ 

## *Proof.* Step 1: $\varepsilon$ -optimal couplings for each c.

Fix  $\varepsilon > 0$ . For each  $c \in \Omega$ , pick an  $\varepsilon$ -optimal coupling  $\pi_c^{\varepsilon} \in \Pi(\mu_c, \nu_c)$  such that

$$\int_{X\times X} d(x,y)^p \, \pi_c^{\varepsilon}(dx,dy) \leq W_p(\mu_c,\nu_c)^p + \varepsilon.$$

(At this step we only use existence for each fixed c; no measurability in c is needed yet.)

#### Step 2: Measurable selection and mixed coupling.

Assume the family  $(\pi_c^{\varepsilon})_{c\in\Omega}$  can be chosen measurably, so that  $c\mapsto\pi_c^{\varepsilon}$  is a probability kernel. We then define the mixed coupling

$$\Pi^{\varepsilon}(U) \; := \; \int_{\Omega} \pi_c^{\varepsilon}(U) \, \rho(dc), \qquad U \subseteq X \times X \; \text{Borel}.$$

For any measurable  $A \subseteq X$ ,

$$\Pi^{\varepsilon}(A \times X) = \int_{\Omega} \pi_{c}^{\varepsilon}(A \times X) \, \rho(dc) = \int_{\Omega} \mu_{c}(A) \, \rho(dc) = \left(\int_{\Omega} \mu_{c} \, \rho(dc)\right)(A),$$

and similarly

$$\Pi^{\varepsilon}(X \times A) = \int_{\Omega} \pi_{c}^{\varepsilon}(X \times A) \, \rho(dc) = \int_{\Omega} \nu_{c}(A) \, \rho(dc) = \left(\int_{\Omega} \nu_{c} \, \rho(dc)\right)(A).$$

Hence  $\Pi^{\varepsilon}$  has the mixed marginals  $\int_{\Omega} \mu_c \, \rho(dc)$  and  $\int_{\Omega} \nu_c \, \rho(dc)$ , i.e.

$$\Pi^{\varepsilon} \in \Pi \Big( \int_{\Omega} \mu_{c} \, \rho(dc), \, \int_{\Omega} \nu_{c} \, \rho(dc) \Big).$$

#### Step 3: Bound the transport cost of the mixed coupling.

Since  $(c, x, y) \mapsto d(x, y)^p$  is nonnegative and measurable and  $c \mapsto \pi_c^{\varepsilon}$  is a probability kernel, Tonelli's theorem allows us to exchange the order of integration in (c, x, y):

$$\begin{split} \int_{X\times X} d(x,y)^p \, \Pi^\varepsilon(dx,dy) &= \int_{X\times X} d(x,y)^p \, \Big( \int_\Omega \pi_c^\varepsilon(dx,dy) \, \rho(dc) \Big) \\ &= \int_\Omega \Big( \int_{X\times X} d(x,y)^p \, \pi_c^\varepsilon(dx,dy) \Big) \, \rho(dc) \\ &\leq \int_\Omega \left( W_p(\mu_c,\nu_c)^p + \varepsilon \right) \rho(dc) \\ &= \int_\Omega W_p(\mu_c,\nu_c)^p \, \rho(dc) \, + \, \varepsilon. \end{split}$$

## Step 4: Take the infimum over couplings and pass to the limit.

By definition of  $W_p$ ,

$$W_p\Big(\int \mu_c \, d\rho, \, \int \nu_c \, d\rho\Big)^p \, \leq \, \int_{X\times X} d(x,y)^p \, \Pi^\varepsilon(dx,dy) \, \leq \, \int_\Omega W_p(\mu_c,\nu_c)^p \, \rho(dc) + \varepsilon.$$

Taking pth roots and letting  $\varepsilon \downarrow 0$  yields

$$W_p \Big( \int_{\Omega} \mu_c \, \rho(dc), \, \int_{\Omega} \nu_c \, \rho(dc) \Big) \, \leq \, \Bigg( \int_{\Omega} W_p (\mu_c, \nu_c)^p \, \rho(dc) \Bigg)^{1/p}.$$

**Lemma 6** (Spectral norm of a block-triangular matrix). Let  $A \in \mathbb{R}^{m \times m}$ ,  $B \in \mathbb{R}^{n \times m}$ , and  $C \in \mathbb{R}^{n \times n}$ . Then

$$\left\| \begin{pmatrix} 1 & 0 & 0 \\ 0 & A & 0 \\ 0 & B & C \end{pmatrix} \right\|_{2} \le \max\{1, \|A\|_{2}, \|C\|_{2}\} + \|B\|_{2}.$$

Proof. Split

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & A & 0 \\ 0 & B & C \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & C \end{pmatrix}}_{D} + \underbrace{\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & B & 0 \end{pmatrix}}_{E}.$$

By the triangle inequality for the operator-2 norm,

$$||M||_2 \le ||D||_2 + ||E||_2.$$

Since D is block-diagonal,  $||D||_2 = \max\{1, ||A||_2, ||C||_2\}$ , and E has only the single nonzero block B, so  $||E||_2 = ||B||_2$ . Substitution gives the claimed bound.

**Theorem 3** (Supremum-p-Wasserstein contraction of the **complete** Sobolev Bellman operator). Let

$$T^{S_{s,a}}_{\pi}: (\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m+n})) \longrightarrow (\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m+n})),$$

be the complete Sobolev Bellman operator bootstrapping the vector  $(Z, \partial_a Z, \partial_s Z)$ .

The distributional Sobolev Bellman operators are defined by the reparameterized policy  $\pi(s'; \varepsilon_{\pi})$ , the transition  $f(s, a; \varepsilon_f)$  and reward  $r(s, a; \varepsilon_r)$ . By assumptions from Section C.1 we write

$$\|\mathcal{L}\|_{d} = \sup_{s,a} \sup_{\varepsilon_{f},\varepsilon_{\pi}} \left\| \mathcal{L}(s,a;\varepsilon_{f},\varepsilon_{\pi}) \right\|_{2} \leq \gamma \, \max\{1, \ L_{f,a}L_{\pi}, \ L_{f,s}\} + L_{f,s}L_{\pi} \} \ = \ \gamma \, \kappa_{\mathrm{full}},$$

then for any two Sobolev return-distribution functions  $\eta_1, \eta_2$ ,

$$|\bar{W}_p(T_{\pi}^{S_{s,a}}\eta_1, T_{\pi}^{S_{s,a}}\eta_2)| \leq ||\mathcal{L}||_d |\bar{W}_p(\eta_1, \eta_2)| \leq \gamma \kappa_{\text{full}} |\bar{W}_p(\eta_1, \eta_2)|.$$

In particular,  $T_{\pi}^{S_{s,a}}$  is a  $\|\mathcal{L}\|_d$ -contraction whenever  $\|\mathcal{L}\|_d < 1$ , and a sufficient condition for this is  $\gamma \kappa_{\text{full}} < 1$ .

*Proof.* We show that the Sobolev Bellman map  $T_{\pi}^{S_{s,a}}$  is a  $\|\mathcal{L}\|_d$ -contraction in the supremum-p-Wasserstein metric.

By definition have,

$$\bar{W}_p(T_{\pi}^{S_{s,a}}\eta_1, T_{\pi}^{S_{s,a}}\eta_2) = \sup_{s,a} W_p(T_{\pi}^{S_{s,a}}\eta_1(s,a), T_{\pi}^{S_{s,a}}\eta_2(s,a)).$$

Let's fix an arbitrary pair (s, a). Then we draw the *same* exogenous noises for both updates

$$\varepsilon_r \sim p(\varepsilon_r), \quad \varepsilon_f \sim p(\varepsilon_f), \quad \varepsilon_\pi \sim p(\varepsilon_\pi),$$

set

$$r = r(s, a; \varepsilon_r), \quad s' = f(s, a; \varepsilon_f), \quad a' = \pi(s'; \varepsilon_\pi),$$

and sample

$$X_1' \sim \eta_1^{S_{s,a}}(s',a'), \qquad X_2' \sim \eta_2^{S_{s,a}}(s',a').$$

Define the random affine map

$$\Phi_{s,a}(x; \varepsilon_r, \varepsilon_f, \varepsilon_\pi) = b(s, a; \varepsilon_r) + \mathcal{L}(s, a; \varepsilon_f, \varepsilon_\pi)[x],$$

where 
$$b = \begin{pmatrix} r \\ \partial_a r \\ \partial_s r \end{pmatrix}$$
,  $\mathcal{L} = \gamma \begin{pmatrix} 1 & 0 & 0 \\ 0 & f_a^\top \pi_s^\top & 0 \\ 0 & f_s^\top \pi_s^\top & f_s^\top \end{pmatrix}$ 

as in Eq.21.

By Lemma 2, we have

$$T_{\pi}^{S_{s,a}} \eta_i(s,a) = \text{Law}(\Phi_{s,a}(X_i')), \qquad i = 1, 2,$$

so

$$W_p(T_{\pi}^{S_{s,a}}\eta_1(s,a), T_{\pi}^{S_{s,a}}\eta_2(s,a)) = W_p(\text{Law}(\Phi_{s,a}(X_1')), \text{Law}(\Phi_{s,a}(X_2'))).$$

Since Law $(X_i') = \eta_i^{S_{s,a}}(s', a')$ , define

$$C = \sup_{u,v} W_p(\eta_1^{S_{s,a}}(u,v), \eta_2^{S_{s,a}}(u,v)),$$

so that for every (s', a'),

$$W_p\!\!\left(\mathrm{Law}(X_1'),\,\mathrm{Law}(X_2')\right) = W_p\!\!\left(\eta_1^{S_{s,a}}(s',a'),\,\eta_2^{S_{s,a}}(s',a')\right) \ \le \ C.$$

Condition on the full noise triple  $(\varepsilon_r, \varepsilon_f, \varepsilon_\pi)$ , so that  $\Phi_{s,a}$  is a deterministic affine map. By Lemma 3,

$$W_{p}(\operatorname{Law}(\Phi_{s,a}(X_{1}') \mid \varepsilon_{r}, \varepsilon_{f}, \varepsilon_{\pi}), \operatorname{Law}(\Phi_{s,a}(X_{2}') \mid \varepsilon_{r}, \varepsilon_{f}, \varepsilon_{\pi}))$$

$$\leq \|\mathcal{L}(s, a; \varepsilon_{f}, \varepsilon_{\pi})\|_{d} W_{p}(\eta_{1}^{S_{s,a}}(s', a'), \eta_{2}^{S_{s,a}}(s', a'))$$

$$\leq \|\mathcal{L}(s, a; \varepsilon_{f}, \varepsilon_{\pi})\|_{d} C.$$

Taking the supremum over  $\varepsilon_f, \varepsilon_\pi$  and then applying Lemma 4 yields

$$W_p(\text{Law}(\Phi_{s,a}(X_1')), \text{Law}(\Phi_{s,a}(X_2'))) \leq \sup_{\varepsilon_f, \varepsilon_\pi} \|\mathcal{L}(s, a; \varepsilon_f, \varepsilon_\pi)\|_d C.$$

Using  $T_{\pi}^{S_{s,a}} \eta_i(s,a) = \text{Law}(\Phi_{s,a}(X_i'))$ , we conclude

$$W_p(T_{\pi}^{S_{s,a}}\eta_1(s,a), T_{\pi}^{S_{s,a}}\eta_2(s,a)) = W_p(\text{Law}(\Phi_{s,a}(X_1')), \text{Law}(\Phi_{s,a}(X_2')))$$

$$\leq \left(\sup_{\varepsilon, t, \varepsilon_{\pi}} \|\mathcal{L}(s,a;\varepsilon_f,\varepsilon_{\pi})\|_d\right) C.$$

Finally, taking the supremum over (s, a) yields

$$\begin{split} \bar{W}_p \left( T_{\pi}^{S_{s,a}} \eta_1, \, T_{\pi}^{S_{s,a}} \eta_2 \right) &= \sup_{s,a} W_p \left( T_{\pi}^{S_{s,a}} \eta_1(s,a), \, T_{\pi}^{S_{s,a}} \eta_2(s,a) \right) \\ &\leq \sup_{s,a} \left( \sup_{\varepsilon_f, \varepsilon_{\pi}} \| \mathcal{L}(s,a;\varepsilon_f, \varepsilon_{\pi}) \|_d \right) C \\ &= \| \mathcal{L} \|_d \, C \\ &= \| \mathcal{L} \|_d \, \bar{W}_p (\eta_1, \, \eta_2). \end{split}$$

so  $T_{\pi}^{S_{s,a}}$  is a contraction with coefficient  $\|\mathcal{L}\|_d$ .

By Section C.1, we have

$$||f_a(s, a; \varepsilon_f)||_2 \le L_{f,a}, \quad ||f_s(s, a; \varepsilon_f)||_2 \le L_{f,s}, \quad ||\pi_s(s'; \varepsilon_\pi)||_2 \le L_\pi.$$

Hence for each  $(s, a, \varepsilon_f, \varepsilon_\pi)$  and applying Lemma 6 we have

$$\begin{split} \left\| \mathcal{L}(s, a; \varepsilon_{f}, \varepsilon_{\pi}) \right\|_{2} &= \gamma \left\| \begin{pmatrix} 1 & 0 & 0 \\ 0 & f_{a}^{\top} \pi_{s}^{\top} & 0 \\ 0 & f_{s}^{\top} \pi_{s}^{\top} & f_{s}^{\top} \end{pmatrix} \right\|_{2} \\ &\leq \gamma \left( \underbrace{\max\{1, \|f_{a}^{\top} \pi_{s}^{\top}\|_{2}, \|f_{s}^{\top}\|_{2}\}}_{\text{diagonal blocks}} + \underbrace{\|f_{s}^{\top} \pi_{s}^{\top}\|_{2}}_{\text{off-diagonal block } B} \right) \\ &\leq \gamma \left( \max\{1, \|f_{a}\|_{2} \|\pi_{s}\|_{2}, \|f_{s}\|_{2} \right) + \|f_{s}\|_{2} \|\pi_{s}\|_{2} \right) \\ &\leq \gamma \left( \max\{1, L_{f,a} L_{\pi}, L_{f,s}\} + L_{f,s} L_{\pi} \right) = \gamma \kappa_{\text{full}}. \end{split}$$

Since the last inequality holds for every choice of  $(s, a, \varepsilon_f, \varepsilon_\pi)$ , taking the supremum gives

$$\|\mathcal{L}\|_d = \sup_{s,a,\varepsilon_f,\varepsilon_\pi} \|\mathcal{L}(s,a;\varepsilon_f,\varepsilon_\pi)\|_2 \le \gamma \kappa_{\text{full}}.$$

**Lemma 7** (Fixed-point law of the **complete** Sobolev Bellman operator). *Define the infinite-horizon* return and its full action- and state-gradients under policy  $\pi$  by

$$Z(s,a) = \sum_{t=0}^{\infty} \gamma^t r_t, \qquad G^a(s,a) = \partial_a \sum_{t=0}^{\infty} \gamma^t r_t, \qquad G^s(s,a) = \partial_s \sum_{t=0}^{\infty} \gamma^t r_t,$$

and let

$$\eta^{\pi}(s, a) = \text{Law}(Z(s, a), G^{a}(s, a), G^{s}(s, a)).$$

Then  $\eta^{\pi}$  is a fixed point of the complete Sobolev Bellman operator:

 $T_{\pi}^{S_{s,a}} \eta^{\pi} = \eta^{\pi}.$ 

*Proof.* Recalling the one-step affine update

$$\Phi_{s,a}(x;\varepsilon_r,\varepsilon_f,\varepsilon_\pi) = b_{\text{full}}(s,a;\varepsilon_r) + \mathcal{L}_{\text{full}}(s,a;\varepsilon_f,\varepsilon_\pi)[x],$$

the Bellman recursion and its derivatives combine to

$$(Z(s,a), G^a(s,a), G^s(s,a)) = \Phi_{s,a}(Z(s',a'), G^a(s',a'), G^s(s',a')).$$

By Lemma 2 we have

$$\operatorname{Law}(Z(s,a), G^{a}(s,a), G^{s}(s,a)) = \operatorname{Law}(\Phi_{s,a}(Z(s',a'), G^{a}(s',a'), G^{s}(s',a')))$$
$$= (T_{\pi}^{S_{s,a}} \eta^{\pi})(s,a).$$

Since this holds for every (s,a), we conclude  $T_\pi^{S_{s,a}}\eta^\pi=\eta^\pi.$ 

**Theorem 4** (Contraction coefficient of the action–gradient (**incomplete**) Sobolev Bellman operator). *Let* 

$$T_{\pi}^{S_a}: (\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m})) \longrightarrow (\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m}))$$

be the action–gradient Sobolev Bellman operator updating only  $H(s,a) = (Z(s,a), \partial_a Z(s,a))$ .

Let's have  $p \geq 1$ . Fix any two return–gradient laws  $\eta_1, \eta_2 : \mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m})$ . Assume

$$||f_a||_2 \le L_{f,a}, \quad ||\pi_s||_2 \le L_{\pi},$$

and that there is a constant  $L_{\mathcal{D}_s} \geq 0$  such that for every (s', a'), if

$$\zeta^* \in \arg\inf_{\zeta \in \Gamma(\eta_1(s',a'), \eta_2(s',a'))} \mathbb{E}_{(H'_1, H'_2) \sim \zeta} [\|H'_1 - H'_2\|_2^p]$$

is an optimal coupling, then, when  $(Z'_i, \partial_a Z'_i)$  are the first and second components of  $H'_i$ ,

$$\left(\mathbb{E}_{\zeta^*} \|\partial_s Z_1' - \partial_s Z_2'\|_2^p\right)^{1/p} \leq L_{\mathcal{D}_s} \left(\mathbb{E}_{\zeta^*} \|(Z_1', \partial_a Z_1') - (Z_2', \partial_a Z_2')\|_2^p\right)^{1/p}. \tag{23}$$

Here  $\Gamma(\mu, \nu)$  is the set of all couplings of  $\mu$  and  $\nu$ . In practice, neural networks with bounded weights and Lipschitz activations make both outputs and gradients Lipschitz, so  $L_{\mathcal{D}_s} < \infty$ .

Set

$$\kappa_{\text{eff}} = \max\{1, L_{f,a}L_{\pi}\} + L_{f,a}L_{\mathcal{D}_s}.$$

Then

$$\bar{W}_p(T_\pi^{S_a}\eta_1, T_\pi^{S_a}\eta_2) \leq \gamma \kappa_{\text{eff}} \bar{W}_p(\eta_1, \eta_2).$$

If  $\gamma \kappa_{\text{eff}} < 1$ ,  $T_{\pi}^{S_a}$  is a strict contraction.

*Proof.* Here and below, all norms  $\|\cdot\|$  denote the Euclidean norm on vectors, and when applied to any bounded linear operator they denote the corresponding induced operator norm; moreover, every  $W_p$  is taken with respect to the ground metric  $d(x,y) = \|x-y\|$ .

Fix (s, a) and draw one sample  $(\varepsilon_r, \varepsilon_f, \varepsilon_\pi)$ . Write

$$s' = f(s, a; \varepsilon_f), \quad a' = \pi(s'; \varepsilon_\pi),$$

and define

$$\Phi(H') = b(s, a; \varepsilon_r) + A(s, a; \varepsilon_f, \varepsilon_\pi) H'(s', a') + N(s, a; \varepsilon_f) \partial_s Z'(s', a').$$

By the uniform Jacobian bounds of Section C.1,

$$\sup_{s,a,\varepsilon_f} \|f_a(s,a;\varepsilon_f)\| \leq L_{f,a}, \qquad \sup_{s',\varepsilon_\pi} \|\pi_s(s';\varepsilon_\pi)\| \leq L_\pi,$$

$$||A(s, a; \varepsilon_f, \varepsilon_\pi)||_2 = \gamma \max\{1, ||f_a(s, a; \varepsilon_f)||_2 ||\pi_s(s'; \varepsilon_\pi)||_2\} \le \gamma \max\{1, |L_{f,a}L_\pi\},$$

and similarly

$$||N(s, a; \varepsilon_f)||_2 = \gamma ||f_a(s, a; \varepsilon_f)||_2 \le \gamma L_{f, a}.$$

Under an optimal coupling  $\zeta$  of  $\eta_1(s',a')$  and  $\eta_2(s',a')$ , let  $(X_1',X_2') \sim \zeta$  be the paired vectors in  $\mathbb{R}^{1+m}$  and  $(G_1',G_2')$  their state-gradients. Then

$$\begin{split} \|\Phi(X_1') - \Phi(X_2')\| &= \left\|A(X_1' - X_2') + N(G_1' - G_2')\right\| & \text{(definition of } \Phi) \\ &\leq \|A(X_1' - X_2')\| + \|N(G_1' - G_2')\| & \text{(triangle inequality)} \\ &\leq \|A\|_2 \, \|X_1' - X_2'\|_2 + \|N\|_2 \, \|G_1' - G_2'\|_2 & \text{(operator-norm bound)} \end{split}$$

We begin by expressing the expected pth power of the output difference in terms of  $\Phi$ :

$$\mathbb{E}_{\zeta} [\|\Phi(X_1') - \Phi(X_2')\|^p] = \mathbb{E}_{\zeta} [\|b + AX_1' + NG_1' - (b + AX_2' + NG_2')\|^p]$$
  
=  $\mathbb{E}_{\zeta} [\|A(X_1' - X_2') + N(G_1' - G_2')\|^p].$ 

Applying the triangle inequality inside the expectation and using that  $x \mapsto x^p$  is order-preserving yields

$$\mathbb{E}_{\zeta} \left[ \|\Phi(X_1') - \Phi(X_2')\|^p \right] \le \mathbb{E}_{\zeta} \left[ (\|A(X_1' - X_2')\| + \|N(G_1' - G_2')\|)^p \right].$$

Taking the pth root and using Minkowski's inequality gives

$$\left(\mathbb{E}_{\zeta} \|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p} \le \left(\mathbb{E}_{\zeta} \|A(X_1' - X_2')\|^p\right)^{1/p} + \left(\mathbb{E}_{\zeta} \|N(G_1' - G_2')\|^p\right)^{1/p}.$$

We use the following chain,

$$(\mathbb{E}||MX||^p)^{1/p} \le (\mathbb{E}[||M||^p ||X||^p])^{1/p} = ||M|| (\mathbb{E}||X||^p)^{1/p}.$$

to factor out the operator norms of A and N

$$\left(\mathbb{E}_{\zeta} \|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p} \le \|A\| \left(\mathbb{E}_{\zeta} \|X_1' - X_2'\|^p\right)^{1/p} + \|N\| \left(\mathbb{E}_{\zeta} \|G_1' - G_2'\|^p\right)^{1/p}.$$

Invoking the derivative–coupling bound from Equation 23 then gives

$$\left(\mathbb{E}_{\zeta} \|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p} \leq \left(\|A\| + \|N\| L_{\mathcal{D}_s}\right) \left(\mathbb{E}_{\zeta} \|X_1' - X_2'\|^p\right)^{1/p}.$$

Recognizing that, because  $\zeta$  is an optimal coupling, the latter expectation equals the p-Wasserstein distance, we have

$$\left(\mathbb{E}_{\zeta} \|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p} = \left(\|A\| + \|N\| L_{\mathcal{D}_s}\right) W_p(\eta_1(s', a'), \eta_2(s', a')).$$

Define, to induce a valid coupling between the transformed measures,

$$\xi = (\Phi \times \Phi)_{\#} \zeta.$$

Then for any measurable  $B \subset \mathbb{R}^{1+m} \times \mathbb{R}^{1+m}$ ,

$$\xi(B) = \zeta((\Phi \times \Phi)^{-1}(B)).$$

Hence  $\xi$  has marginals  $\Phi_{\#}\eta_1$  and  $\Phi_{\#}\eta_2$ , so  $\xi$  is indeed a coupling of the push-forward measures, i.e.

$$\xi \in \Gamma(\Phi_{\#}\eta_1, \Phi_{\#}\eta_2).$$

By definition of  $W_n$ ,

$$W_p(\Phi_{\#}\eta_1, \Phi_{\#}\eta_2) \leq \left(\mathbb{E}_{(U,V)\sim\xi} \|U - V\|^p\right)^{1/p} = \left(\mathbb{E}_{\zeta} \|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p}.$$

Replacing the left side we have

$$W_p(\Phi_\#\eta_1, \Phi_\#\eta_2) \le (\|A\| + \|N\| L_{\mathcal{D}_s}) W_p(\eta_1(s', a'), \eta_2(s', a')).$$

Substituting the operator-norm bounds  $||A|| \le \gamma \max\{1, L_{f,a}L_{\pi}\}$  and  $||N|| \le \gamma L_{f,a}$  yields

$$W_p(\Phi_{\#}\eta_1, \Phi_{\#}\eta_2) \leq \gamma(\max\{1, L_{f,a}L_{\pi}\} + L_{f,a}L_{\mathcal{D}_s}) W_p(\eta_1(s', a'), \eta_2(s', a')).$$

Noting that for each noise draw  $c=(\varepsilon_r,\varepsilon_f,\varepsilon_\pi)$ , with  $s'=f(s,a;\varepsilon_f)$  and  $a'=\pi(s';\varepsilon_\pi)$ , we have the pointwise bound

$$W_p(\Phi_{c\#}\eta_1(s',a'), \Phi_{c\#}\eta_2(s',a')) \leq \gamma \kappa_{\text{eff}} W_p(\eta_1(s',a'), \eta_2(s',a')).$$

By Lemma 5 (mixture *p*-convexity)

$$W_{p}^{p}(T_{\pi}^{S_{a}}\eta_{1}(s, a), T_{\pi}^{S_{a}}\eta_{2}(s, a)) = W_{p}^{p}(\mathbb{E}_{c}[\Phi_{c\#}\eta_{1}(s', a')], \mathbb{E}_{c}[\Phi_{c\#}\eta_{2}(s', a')])$$

$$\leq \mathbb{E}_{c}[W_{p}^{p}(\Phi_{c\#}\eta_{1}(s', a'), \Phi_{c\#}\eta_{2}(s', a'))]$$

$$\leq (\gamma \kappa_{\text{eff}})^{p} \mathbb{E}_{c}[W_{p}^{p}(\eta_{1}(s', a'), \eta_{2}(s', a'))]$$

$$\leq (\gamma \kappa_{\text{eff}})^{p} \bar{W}_{p}^{p}(\eta_{1}, \eta_{2}).$$

Taking pth roots yields

$$W_p(T_{\pi}^{S_a}\eta_1(s,a), T_{\pi}^{S_a}\eta_2(s,a)) \leq \gamma \, \kappa_{\text{eff}} \, \bar{W}_p(\eta_1,\eta_2).$$

Finally, taking the supremum over (s, a),

$$\bar{W}_p(T_\pi^{S_a}\eta_1, T_\pi^{S_a}\eta_2) \leq \gamma \kappa_{\text{eff}} \bar{W}_p(\eta_1, \eta_2).$$

**Theorem 5** (Contraction coefficient of the action–gradient (**incomplete**) Sobolev Bellman operator — local regularity version). *Let* 

$$T_{\pi}^{S_a}: \left(\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m})\right) \longrightarrow \left(\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m})\right)$$

be the action–gradient Sobolev Bellman operator updating only  $H(s,a) = (Z(s,a), \partial_a Z(s,a))$ . Fix  $p \geq 1$  and two return–gradient laws  $\eta_1, \eta_2 : \mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m})$ . Assume

$$||f_a||_2 \leq L_{f,a}, \qquad ||\pi_s||_2 \leq L_{\pi},$$

and suppose there exists  $L_{D_s} \ge 0$  such that for every (s', a') and every pair of value functions  $Z_1, Z_2$  in the model class, writing  $h := Z_1 - Z_2$ ,

$$\|\partial_s h(s', a')\|_2 \le L_{D_s} \| (h(s', a'), \partial_a h(s', a')) \|_2.$$
 (24)

(In practice, this holds if the architecture or modeling constraints ensure that  $\partial_s Z$  is a bounded linear functional of  $(Z, \partial_a Z)$  at each (s', a').)

Set

$$\kappa_{\text{eff}} = \max\{1, L_{f,a}L_{\pi}\} + L_{f,a}L_{D_s}.$$

Then

$$\bar{W}_p\left(T_{\pi}^{S_a}\eta_1, T_{\pi}^{S_a}\eta_2\right) \leq \gamma \, \kappa_{\text{eff}} \, \bar{W}_p(\eta_1, \eta_2)$$

In particular, if  $\gamma \kappa_{\text{eff}} < 1$ , then  $T_{\pi}^{S_a}$  is a strict contraction on  $(\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m}), \bar{W}_p)$ .

*Proof.* All vector norms  $\|\cdot\|$  are Euclidean norms; operator norms are the corresponding induced norms. All  $W_p$  distances use the ground metric  $d(x,y) = \|x-y\|$ .

Fix (s, a) and a realization  $c = (\varepsilon_r, \varepsilon_f, \varepsilon_\pi)$  of the exogenous noise. Let

$$s' = f(s, a; \varepsilon_f), \qquad a' = \pi(s'; \varepsilon_\pi).$$

The one–step backup of H can be written in the affine form

$$\Phi(H') = b(s, a; \varepsilon_r) + A(s, a; \varepsilon_f, \varepsilon_\pi) H'(s', a') + N(s, a; \varepsilon_f) \partial_s Z'(s', a'),$$

where  $H' = (Z', \partial_a Z')$ . By the uniform Jacobian bounds (Section C.1),

$$||A(s, a; \varepsilon_f, \varepsilon_\pi)|| = \gamma \max\{1, ||f_a(s, a; \varepsilon_f)|| ||\pi_s(s'; \varepsilon_\pi)||\} \le \gamma \max\{1, L_{f,a}L_\pi\},$$

and

$$||N(s, a; \varepsilon_f)|| = \gamma ||f_a(s, a; \varepsilon_f)|| \le \gamma L_{f,a}.$$

Let  $\zeta^*$  be an optimal coupling between  $\eta_1(s',a')$  and  $\eta_2(s',a')$ . Draw  $(X_1',X_2') \sim \zeta^*$ , where  $X_i' = (Z_i', \partial_a Z_i') \in \mathbb{R}^{1+m}$ , and write  $G_i' = \partial_s Z_i'(s',a')$ . Then

$$\begin{split} \|\Phi(X_1') - \Phi(X_2')\| &= \left\|A(X_1' - X_2') + N(G_1' - G_2')\right\| & \text{(definition of } \Phi) \\ &\leq \|A(X_1' - X_2')\| + \|N(G_1' - G_2')\| & \text{(triangle inequality)} \\ &\leq \|A\| \, \|X_1' - X_2'\| + \|N\| \, \|G_1' - G_2'\| & \text{(operator-norm bound)}. \end{split}$$

Applying the local regularity bound equation 24 pointwise to  $h := Z'_1 - Z'_2$  at (s', a') gives

$$||G_1' - G_2'|| = ||\partial_s h(s', a')|| \le L_{D_s} ||(h(s', a'), \partial_a h(s', a'))|| = L_{D_s} ||X_1' - X_2'||.$$

Hence, for every paired draw  $(X'_1, X'_2)$ ,

$$\|\Phi(X_1') - \Phi(X_2')\| \le (\|A\| + \|N\| L_{D_s}) \|X_1' - X_2'\|.$$

Raising to the pth power, taking expectations under  $\zeta^*$ , and then taking pth roots yields

$$\left(\mathbb{E}_{\zeta^*} \|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p} \leq \left(\|A\| + \|N\| L_{D_s}\right) \left(\mathbb{E}_{\zeta^*} \|X_1' - X_2'\|^p\right)^{1/p}.$$

By optimality of  $\zeta^*$ ,

$$\left(\mathbb{E}_{\zeta^*} \|X_1' - X_2'\|^p\right)^{1/p} = W_p(\eta_1(s', a'), \eta_2(s', a')).$$

Define, to induce a valid coupling between the transformed measures,

$$\xi = (\Phi \times \Phi)_{\#} \zeta^*.$$

Then for any Borel  $B \subset \mathbb{R}^{1+m} \times \mathbb{R}^{1+m}$ ,

$$\xi(B) = \zeta^* ((\Phi \times \Phi)^{-1}(B)).$$

Hence  $\xi$  has marginals  $\Phi_{\#}\eta_1(s',a')$  and  $\Phi_{\#}\eta_2(s',a')$ , i.e.  $\xi \in \Gamma(\Phi_{\#}\eta_1(s',a'),\Phi_{\#}\eta_2(s',a'))$ . By the definition of  $W_p$ ,

$$W_p\left(\Phi_{\#}\eta_1(s',a'), \Phi_{\#}\eta_2(s',a')\right) \leq \left(\mathbb{E}_{(U,V)\sim\xi}\|U-V\|^p\right)^{1/p} = \left(\mathbb{E}_{\zeta^*}\|\Phi(X_1') - \Phi(X_2')\|^p\right)^{1/p}.$$

Combining the previous two inequalities and inserting the bounds on ||A|| and ||N|| gives, for this fixed noise realization,

$$W_p(\Phi_{\#}\eta_1(s',a'),\Phi_{\#}\eta_2(s',a')) \leq \gamma(\max\{1,L_{f,a}L_{\pi}\} + L_{f,a}L_{D_s})W_p(\eta_1(s',a'),\eta_2(s',a')).$$

Finally,  $T_{\pi}^{S_a}\eta_i(s,a)$  is the mixture (over  $\varepsilon_r, \varepsilon_f, \varepsilon_\pi$ ) of the pushforwards  $\Phi_{c\#}\eta_i(s',a')$  with  $(s',a')=(f(s,a;\varepsilon_f),\pi(s';\varepsilon_\pi))$ . By Lemma 5 (mixture p-convexity) and the fixed–noise bound above,

$$W_{p}^{p}(T_{\pi}^{S_{a}}\eta_{1}(s, a), T_{\pi}^{S_{a}}\eta_{2}(s, a)) = W_{p}^{p}(\mathbb{E}_{c}[\Phi_{c\#}\eta_{1}(s', a')], \mathbb{E}_{c}[\Phi_{c\#}\eta_{2}(s', a')])$$

$$\leq \mathbb{E}_{c}[W_{p}^{p}(\Phi_{c\#}\eta_{1}(s', a'), \Phi_{c\#}\eta_{2}(s', a'))]$$

$$\leq (\gamma \kappa_{\text{eff}})^{p} \mathbb{E}_{c}[W_{p}^{p}(\eta_{1}(s', a'), \eta_{2}(s', a'))]$$

$$\leq (\gamma \kappa_{\text{eff}})^{p} \bar{W}_{p}^{p}(\eta_{1}, \eta_{2}).$$

Taking pth roots yields

$$W_p(T_{\pi}^{S_a}\eta_1(s, a), T_{\pi}^{S_a}\eta_2(s, a)) \leq \gamma \kappa_{\text{eff}} \bar{W}_p(\eta_1, \eta_2).$$

Taking the supremum over (s, a) completes the proof:

$$\bar{W}_p(T_{\pi}^{S_a}\eta_1, T_{\pi}^{S_a}\eta_2) \leq \gamma \kappa_{\text{eff}} \bar{W}_p(\eta_1, \eta_2).$$

**Lemma 8** (Fixed-point law of the **incomplete** Sobolev Bellman operator). *Define the infinite-horizon return and its full action-gradient under policy*  $\pi$  *by* 

 $Z(s,a) = \sum_{t=0}^{\infty} \gamma^t r_t, \qquad G(s,a) = \partial_a \sum_{t=0}^{\infty} \gamma^t r_t,$ 

and let  $\eta^{\pi}(s, a) = \text{Law}(Z(s, a), G(s, a))$ . Then  $\eta^{\pi}$  is a fixed point of the Sobolev Bellman operator:

$$T_{\pi}^{S_a} \eta^{\pi} = \eta^{\pi}.$$

*Proof.* Recalling the one-step affine update

$$\Phi_{s,a}(x;\,\varepsilon_r,\varepsilon_f,\varepsilon_\pi) = b(s,a;\varepsilon_r) + \mathcal{L}(s,a;\varepsilon_f,\varepsilon_\pi)[\,x\,],$$

the Bellman recursion and its derivative combine to

$$(Z(s,a), G(s,a)) = \Phi_{s,a}(Z(s',a'), G(s',a')).$$

By Lemma 2 we have

$$\text{Law}(Z(s, a), G(s, a)) = \text{Law}(\Phi_{s, a}(Z(s', a'), G(s', a'))) = (T_{\pi}^{S_a} \eta^{\pi})(s, a).$$

Since this holds for every (s, a), we conclude  $T^{S_a}_{\pi} \eta^{\pi} = \eta^{\pi}$ .

Corollary 5.1 (Convergence of Sobolev evaluation iterates). Under the conditions of Theorem 3, let  $\kappa$  denote the contraction constant with  $\kappa = \kappa_{\rm eff}$  for the incomplete operator or  $\kappa = \kappa_{\rm full}$  for the complete one and suppose  $\gamma \kappa < 1$ . For any initial Sobolev return–distribution function  $\eta_0$ , define the sequence

$$\eta_{n+1} = T_{\pi}^S \, \eta_n,$$

where  $T_{\pi}^{S}$  may be either the incomplete  $(S_{a})$  or complete  $(S_{s,a})$  Sobolev operator. Then by Banach's fixed-point theorem the iterates converge to the unique fixed point  $\eta^{\pi}$  (cf. Lemmas 7, 8):

$$\bar{W}_p(\eta_n, \eta^\pi) \leq (\gamma \kappa)^n \bar{W}_p(\eta_0, \eta^\pi) \xrightarrow[n \to \infty]{} 0.$$

In particular,  $\eta_n \to \eta^{\pi}$  in the supremum-p-Wasserstein metric.

#### D.2 PROOF OF PROPOSITION 3.1

**Lemma 9** (Mean-difference bound via  $W_1$ ). Let X, Y be  $\mathbb{R}^d$ -valued random variables with distributions  $\mu = \operatorname{Law}(X)$  and  $\nu = \operatorname{Law}(Y)$ , and assume  $\mathbb{E}||X|| < \infty$ ,  $\mathbb{E}||Y|| < \infty$ . Then

$$\left\| \mathbb{E}[X] - \mathbb{E}[Y] \right\| \leq W_1(\mu, \nu),$$

Here,  $\|\cdot\|$  denotes the Euclidean norm on  $\mathbb{R}^d$ , and  $W_1$  is taken with respect to the ground metric  $d(x,y) = \|x-y\|$ .

Proof. Let

$$m_X = \mathbb{E}[X], \quad m_Y = \mathbb{E}[Y].$$

If  $m_X \neq m_Y$ , define the unit vector

$$u = \frac{m_X - m_Y}{\|m_X - m_Y\|}.$$

Then the scalar function  $f(x) = u^{\top}x$  satisfies

$$|f(x) - f(y)| = |u^{\top}(x - y)| \le ||u|| ||x - y|| = ||x - y||,$$

so  $||f||_{\text{Lip}} \le 1$ . By Kantorovich–Rubinstein duality Villani et al. (2009),

$$W_1(\mu,\nu) = \sup_{\|g\|_{\mathrm{Lip}} \le 1} \left| \mathbb{E}[g(X)] - \mathbb{E}[g(Y)] \right| \ge \left| \mathbb{E}[f(X)] - \mathbb{E}[f(Y)] \right|.$$

But  $\mathbb{E}[f(X)] - \mathbb{E}[f(Y)] = u^{\top}(m_X - m_Y) = \|m_X - m_Y\|$ , hence  $\|m_X - m_Y\| \le W_1(\mu, \nu)$ . If  $m_X = m_Y$ , the inequality is trivial.

**Proposition 2.** Let  $\pi$  be an  $L_{\pi}$ -Lipschitz continuous policy, and let  $G(s) = \operatorname{Law}(\nabla_a Z^{\pi}(s,a) \mid_{a=\pi(s)})$  and  $\hat{G}(s) = \operatorname{Law}(\nabla_a \hat{Z}(s,a) \mid_{a=\pi(s)})$  denote the true and estimated distributions of the action-gradients at  $a=\pi(s)$ , respectively. Then the error between the true and estimated policy gradients is bounded by

$$\left\| \nabla_{\theta} J(\theta) - \nabla_{\theta} \hat{J}(\theta) \right\| \leq \frac{L_{\pi}}{1 - \gamma} \mathbb{E}_{s \sim d_{\pi}^{\mu}} \left[ W_{1} \left( \nabla_{a} Z^{\pi}(s, a) \big|_{a = \pi(s)}, \nabla_{a} \hat{Z}(s, a) \big|_{a = \pi(s)} \right) \right].$$

This proposition is a distributional extension of Proposition 3.1 from D'Oro & Jaskowski (2020).

Proof.

#### **Step 1: True and Estimated Policy Gradients**

The true policy gradient is given by:

$$\nabla_{\theta}J(\theta) = \frac{1}{1-\gamma}\mathbb{E}_{s\sim d^{\mu}_{\pi}}\left[\mathbb{E}\left[\nabla_{a}Z^{\pi}(s,a)\big|_{a=\pi(s)}\right]\nabla_{\theta}\pi(s)\right],$$

The estimated policy gradient is:

$$\nabla_{\theta} \hat{J}(\theta) = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\pi}^{\mu}} \left[ \mathbb{E} \left[ \nabla_{a} \hat{Z}(s,a) \big|_{a=\pi(s)} \right] \nabla_{\theta} \pi(s) \right].$$

## **Step 3: Policy Gradient Error**

The  $L^2$  norm of the difference between the true and estimated policy gradients is:

$$\begin{split} \left\| \nabla_{\theta} J(\theta) - \nabla_{\theta} \hat{J}(\theta) \right\| &= \left\| \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d_{\pi}^{\mu}} \left[ \left( \mathbb{E} \left[ \nabla_{a} Z^{\pi}(s, a) \big|_{a = \pi(s)} \right] - \mathbb{E} \left[ \nabla_{a} \hat{Z}(s, a) \big|_{a = \pi(s)} \right] \right) \right. \\ &\times \left. \nabla_{\theta} \pi(s) \right] \right\|. \end{split}$$

#### Step 4: Applying the Triangle Inequality and Lipschitz Continuity

Using the triangle inequality and the Lipschitz continuity of the policy  $(\|\nabla_{\theta}\pi(s)\| \leq L_{\pi})$ , we have:

$$\begin{split} \left\| \nabla_{\theta} J(\theta) - \nabla_{\theta} \hat{J}(\theta) \right\| &\leq \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d_{\pi}^{\mu}} \Big[ \left\| \mathbb{E} \left[ \nabla_{a} Z^{\pi}(s, a) \big|_{a = \pi(s)} \right] - \mathbb{E} \left[ \nabla_{a} \hat{Z}(s, a) \big|_{a = \pi(s)} \right] \right\| \\ &\times \left\| \nabla_{\theta} \pi(s) \right\| \Big] \\ &\leq \frac{L_{\pi}}{1 - \gamma} \mathbb{E}_{s \sim d_{\pi}^{\mu}} \left[ \left\| \mathbb{E} \left[ \nabla_{a} Z^{\pi}(s, a) \big|_{a = \pi(s)} \right] - \mathbb{E} \left[ \nabla_{a} \hat{Z}(s, a) \big|_{a = \pi(s)} \right] \right\| \Big] \,. \end{split}$$

#### Step 5: Bounding the mean-difference by $W_1$

Let

$$X = \nabla_a Z^{\pi}(s, a) \big|_{a=\pi(s)}, \qquad Y = \nabla_a \hat{Z}(s, a) \big|_{a=\pi(s)}.$$

By Lemma 9, we immediately have

$$\left\| \mathbb{E}[X] - \mathbb{E}[Y] \right\| \leq W_1(X, Y) = W_1(\nabla_a Z^{\pi}(s, a) \big|_{a = \pi(s)}, \nabla_a \hat{Z}(s, a) \big|_{a = \pi(s)}).$$

## **Step 6: Conclusion**

Combining the results from the previous steps, we established that the  $L^2$  norm of the difference between the true and estimated policy gradients can be bounded as follows:

$$\left\| \nabla_{\theta} J(\theta) - \nabla_{\theta} \hat{J}(\theta) \right\| \leq \frac{L_{\pi}}{1 - \gamma} \left\| \mathbb{E}_{s \sim d_{\pi}^{\mu}} \left[ W_{1} \left( \nabla_{a} Z^{\pi}(s, a) \big|_{a = \pi(s)}, \nabla_{a} \hat{Z}(s, a) \big|_{a = \pi(s)} \right) \right].$$

## E PRACTICAL DIFFICULTIES WITH WASSERSTEIN FOR TRAINING

Adversarial  $W_1$  is a loose proxy for true Wasserstein. Some works have directly applied WGAN training to distributional RL, for example Freirich et al. (2019), who cast return distributions as targets for adversarial matching. In practice, WGAN training replaces the exact Kantorovich–Rubinstein dual (Arjovsky et al., 2017; Gulrajani et al., 2017) with a *parametric* discriminator plus approximate Lipschitz control (weight clipping or gradient penalties). This induces three sources of deviation from the true distance: finite-capacity approximation error, imperfect Lipschitz enforcement, and optimization error. Systematic analyses show that the resulting WGAN losses can correlate poorly with the actual Wasserstein metric and need not be meaningful approximations of it, undermining proofs that presume access to the exact  $W_1$ . See both the empirical and theoretical analysis of Mallasto et al. (2019) and the critique by Stanczuk et al. (2021).

Computational cost of exact OT in multiple dimensions. Even ignoring estimator issues, computing multivariate  $W_p$  exactly on mini-batches is costly: building the pairwise cost matrix requires  $O(m^2)$  memory, and solving the discrete OT problem takes  $O(m^3 \log m)$  time (Genevay et al., 2019). This makes per-update calls prohibitive in distributional RL.

**Sample complexity.** Beyond runtime, OT also suffers from poor statistical efficiency: the empirical Wasserstein distance converges to its population value at rate  $O(n^{-1/d})$  in dimension d, compared to  $O(n^{-1/2})$  for MMD (Genevay et al., 2019). This slow convergence further limits its practicality in high-dimensional RL.

## F BACKGROUND ON MMD

**Definition 3** (Maximum Mean Discrepancy as an IPM). Let  $k: X \times X \to \mathbb{R}$  be a symmetric, positive-semi-definite reproducing kernel with RKHS  $\mathcal{H}$  and feature map

$$\phi \colon X \to \mathcal{H}, \qquad \phi(x) = k(x, \cdot).$$

For probability measures P, Q on X, define their mean embeddings

$$\mu_P = \int_X \phi(x) dP(x), \quad \mu_Q = \int_X \phi(x) dQ(x).$$

Then the Maximum Mean Discrepancy is

$$MMD_k(P,Q) := \|\mu_P - \mu_Q\|_{\mathcal{H}},$$

whose square admits the kernel expansion

$$\begin{aligned} \text{MMD}_{k}^{2}(P,Q) &= \|\mu_{P} - \mu_{Q}\|_{\mathcal{H}}^{2} \\ &= \iint k(x,x') \, dP(x) \, dP(x') \, + \, \iint k(y,y') \, dQ(y) \, dQ(y') \\ &- \, 2 \iint k(x,y) \, dP(x) \, dQ(y) \\ &= \iint k(x,y) \, d(P-Q)(x) \, d(P-Q)(y). \end{aligned}$$

Moreover,  $\mathrm{MMD}_k$  coincides with the integral probability metric (IPM) over the unit ball of  $\mathcal{H}$ , namely

$$\mathrm{MMD}_{k}(P,Q) = \sup_{\substack{f \in \mathcal{H} \\ \|f\|_{\mathcal{H}} \le 1}} \left\{ \mathbb{E}_{x \sim P}[f(x)] - \mathbb{E}_{y \sim Q}[f(y)] \right\} = \|\mu_{P} - \mu_{Q}\|_{\mathcal{H}},$$

as shown in Gretton et al. (2012).

**Remark (Euclidean densities).** Working in  $\mathbb{R}^d$ , if P and Q admit densities p(x) and q(x) with respect to Lebesgue measure dx, then

$$dP(x) = p(x) dx$$
,  $dQ(x) = q(x) dx$ ,  $d(P - Q)(x) = (p(x) - q(x)) dx$ ,

and each of the above integrals becomes an ordinary Lebesgue integral in x, y.

**Definition 4** (Conditionally positive definite (CPD) kernel — integral form). Let X be a measurable space and let  $k: X \times X \to \mathbb{R}$  be symmetric. We say that k is conditionally positive definite (CPD) if

$$\iint_{X\times X} k(x,x')\,d\mu(x)\,d\mu(x')\ \geq\ 0\qquad \text{for all finite signed measures $\mu$ on $X$ with $\mu(X)=0$}.$$

If the inequality is strict for every nonzero such  $\mu$ , then k is conditionally strictly positive definite (CSPD).

**Proposition 3** (Equivalence of  $\gamma_k$  and RKHS–MMD for CPD kernels). Let  $k: X \times X \to \mathbb{R}$  be conditionally positive definite (CPD) and define

$$\rho_k(x,y) := k(x,x) + k(y,y) - 2k(x,y).$$

Fix  $z_0 \in X$  and set the distance-induced (one-point centered) kernel

$$k^{\circ}(x,y) \ := \ \tfrac{1}{2} \big[ \rho_k(x,z_0) + \rho_k(y,z_0) - \rho_k(x,y) \big] \ = \ k(x,y) - k(x,z_0) - k(z_0,y) + k(z_0,z_0).$$

Then  $k^{\circ}$  is positive definite and admits an RKHS  $\mathcal{H}_{k^{\circ}}$ . For any P, Q with finite integrals,

$$\gamma_k(P,Q)^2 := \iint k(x,y) d(P-Q)(x) d(P-Q)(y)$$

$$= \iint k^\circ(x,y) d(P-Q)(x) d(P-Q)(y)$$

$$= \|\mu_{k^\circ}(P) - \mu_{k^\circ}(Q)\|_{\mathcal{H}_{k^\circ}}^2$$

$$= \text{MMD}_{k^\circ}(P,Q)^2.$$

Justification. This follows from the distance-induced kernel construction and equivalence results in Sejdinovic et al. (2013).

**Proposition 4** (MMD as a Metric on  $\mathcal{P}(X)$ ). Let  $k \colon X \times X \to \mathbb{R}$  be a symmetric kernel. We say that MMD<sub>k</sub> defines a metric on  $\mathcal{P}(X)$  iff k is conditionally strictly positive definite (CSPD), i.e., for every nonzero finite signed Borel measure  $\nu$  with  $\nu(X) = 0$ ,

$$\iint_{X\times X} k(x,y) \, d\nu(x) \, d\nu(y) > 0.$$

Then  $\mathrm{MMD}_k$  satisfies the metric axioms on  $\mathcal{P}(X)$ :

- 1. Nonnegativity:  $\mathrm{MMD}_k(P,Q) \geq 0$ .
- 2. Symmetry:  $MMD_k(P,Q) = MMD_k(Q,P)$ .
- 3. Identity of indiscernibles:  $MMD_k(P,Q) = 0 \Rightarrow P = Q$ .
- 4. Triangle inequality: for any  $P, Q, R \in \mathcal{P}(X)$ ,  $\mathrm{MMD}_k(P,Q) \leq \mathrm{MMD}_k(P,R) + \mathrm{MMD}_k(R,Q)$ .

Justification. This is the standard correspondence between negative-type distances, distance-induced kernels, and RKHS MMD metrics as outlined in Sejdinovic et al. (2013).

#### Examples of kernels inducing a metric.

• Gaussian RBF kernel: for any bandwidth  $\sigma > 0$ ,

$$k_{\sigma}^{\text{RBF}}(x,y) = \exp(-\|x-y\|_{2}^{2}/(2\sigma^{2})),$$

which is characteristic on  $\mathbb{R}^d$  and hence induces a metric on  $\mathcal{P}(\mathbb{R}^d)$  via MMD<sub>kRBF</sub>.

• Multiquadric kernel (Killingberg & Langseth (2023)):

$$k_h^{\text{MQ}}(x,y) = -\sqrt{1 + h^2 \|x - y\|_2^2}, \quad h > 0,$$

which is conditionally strictly positive-definite and thus induces a metric on distributions via  $\mathrm{MMD}_{k_r^{\mathrm{MQ}}}$ .

#### F.1 CONTRACTION UNDER MMD

Contraction guarantees under MMD can be established in much the same way as in Appendix D. This first requires defining the notion of the supremum-MMD, which, as with the supremum-Wasserstein distance, is a worst-case bound over the entire state–action space:

$$\mathrm{MMD}_{\infty}(\eta,\nu) := \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathrm{MMD} \big( \eta(s,a), \, \nu(s,a) \big).$$

As Nguyen et al. (2020) first introduced MMD-based distributional reinforcement learning, they provided criteria under which a kernel induces a contraction in this sup-MMD metric with the standard distributional Bellman operator. We first recall the univariate distributional Bellman operator Bellemare et al. (2017):

$$\big(\mathcal{T}_{\pi}^{\mathrm{Dist}}\eta\big)(s,a) = \mathrm{Law}\big[R(s,a) + \gamma\,Z(s',\pi(s'))\big], \quad s' \sim P(\cdot\mid s,a).$$

**Sufficient conditions.** Let  $k(x,y) = \sum_{i \in I} c_i k_i(x,y)$  be a positive-definite kernel on  $\mathbb{R}$ . If each component  $k_i$  satisfies:

- 1. Shift-invariance:  $k_i(x+c,y+c)=k_i(x,y)$  for all  $c \in \mathbb{R}$ ,
- 2. Scale-sensitivity of order  $\alpha_i$ :  $k_i(cx, cy) = |c|^{\alpha_i} k_i(x, y)$  for all  $c \in \mathbb{R}$ ,

then for any policy  $\pi$ ,

$$\mathrm{MMD}_{\infty}(\mathcal{T}_{\pi}^{\mathrm{Dist}}\eta, \ \mathcal{T}_{\pi}^{\mathrm{Dist}}\nu) \leq \gamma^{\alpha_*/2} \, \mathrm{MMD}_{\infty}(\eta, \nu),$$

where  $\alpha_* = \min_{i \in I} \alpha_i$ .

Contraction under the multiquadric kernel. Another work, whose kernel we use, Killingberg & Langseth (2023) proposed the multiquadric kernel discussed above. They showed a contraction in MMD<sup>2</sup> for that specific kernel. More precisely, this holds for any pair of distributions  $\mu, \nu \in \mathcal{P}(\mathbb{R})$ , letting  $(f_{r,\gamma})_{\#}\mu$  denote the pushforward of  $\mu$  by  $z \mapsto r + \gamma z$ , they show

$$\mathrm{MMD}^{2}\big((f_{r,\gamma})_{\#}\mu,\,(f_{r,\gamma})_{\#}\nu;\,k_{h}^{\mathrm{MQ}}\big) \leq \gamma\,\mathrm{MMD}^{2}\big(\mu,\nu;\,k_{h}^{\mathrm{MQ}}\big).$$

Taking square-roots on both sides gives the pointwise MMD bound

$$\mathrm{MMD}((f_{r,\gamma})_{\#}\mu, (f_{r,\gamma})_{\#}\nu; k_h^{\mathrm{MQ}}) \leq \sqrt{\gamma} \, \mathrm{MMD}(\mu, \nu; k_h^{\mathrm{MQ}}).$$

Define the supremum-MMD over state-action pairs by

$$\mathrm{MMD}_{\infty}(\eta, \nu) := \sup_{(s, a) \in \mathcal{S} \times \mathcal{A}} \mathrm{MMD} \big( \eta(s, a), \, \nu(s, a) \big).$$

Then for any two return–distribution mappings  $\eta, \nu$ ,

$$\begin{aligned} \mathrm{MMD}_{\infty} \big( \mathcal{T}_{\pi}^{\mathrm{Dist}} \eta, \ \mathcal{T}_{\pi}^{\mathrm{Dist}} \nu \big) &= \sup_{(s,a)} \mathrm{MMD} \big( (f_{R(s,a),\gamma})_{\#} \eta(s,a), \ (f_{R(s,a),\gamma})_{\#} \nu(s,a) \big) \\ &\leq \sup_{(s,a)} \sqrt{\gamma} \, \mathrm{MMD} \big( \eta(s,a), \ \nu(s,a) \big) \\ &= \sqrt{\gamma} \, \mathrm{MMD}_{\infty} (\eta, \nu). \end{aligned}$$

Thus the distributional Bellman operator is a  $\sqrt{\gamma}$ -contraction in MMD $_{\infty}$ .

Contraction in the multivariate setting. To the best of our knowledge, the only MMD-contraction result in a multivariate setting is from Wiltzer et al. (2024). They show that when you push each sampled return vector  $Z \in \mathbb{R}^d$  through the affine map  $z \mapsto \mathbf{R}(s,a) + \gamma z$ , the same two requirements, shift-invariance and homogeneity of the kernel, guarantee a  $\gamma^{\alpha/2}$  contraction in the supremum-MMD metric. In other words, by treating each component of the return vector uniformly and applying the identical homogeneity-based argument from the univariate case, one obtains exactly the same geometric shrinkage factor.

This, however, falls short of the setting we require for the distributional Sobolev Bellman operator in Appendix C, where the pushforward is the more general (pseudo)-affine map

$$x \mapsto \Phi_{s,a}(x) = b(s,a) + \mathcal{L}(s,a)[x],$$

and  $\mathcal{L}(s,a)$  need not be a simple diagonal scaling. Characterizing the precise conditions on both the kernel and the (pseudo-)linear operator  $\mathcal{L}(s,a)$  under which this general  $\Phi_{s,a}$  yields a contraction in supremum-MMD remains an open problem.

#### F.2 EMPIRICAL ESTIMATORS OF MMD

In practice, the expectations in Eq. 16 cannot be computed exactly and must be approximated from samples. Given two sets of m samples  $\{x_i\}_{i=1}^m \sim P$  and  $\{y_i\}_{i=1}^m \sim Q$ , one commonly used estimator is the *biased* form (Gretton et al., 2012):

$$\widehat{\text{MMD}}_b^2 = \frac{1}{m^2} \sum_{i,j=1}^m k(x_i, x_j) + \frac{1}{m^2} \sum_{i,j=1}^m k(y_i, y_j) - \frac{2}{m^2} \sum_{i,j=1}^m k(x_i, y_j).$$
(25)

An alternative is the *unbiased* estimator proposed in (Gretton et al., 2012):

$$\widehat{\mathbf{MMD}}_{u}^{2} = \frac{1}{m(m-1)} \sum_{\substack{i,j=1\\i\neq j}}^{m} k(x_{i}, x_{j}) + \frac{1}{m(m-1)} \sum_{\substack{i,j=1\\i\neq j}}^{m} k(y_{i}, y_{j}) - \frac{2}{m^{2}} \sum_{i,j=1}^{m} k(x_{i}, y_{j}).$$
(26)

The biased estimator given above is commonly preferred over the unbiased one in works involving MMD for distributional RL (Nguyen et al., 2020; Killingberg & Langseth, 2023).

## G RESULTS WITH MAX SLICED MMD

## H RESULTS WITH MAX-SLICED MMD (MSMMD)

Corollary 6.2 shows that the *complete* Sobolev Bellman operator  $T_{\pi}^{S_{s,a}}$  contracts by  $c_1(\|\mathcal{L}\|_d)$  under the supremum–max–sliced MQ–MMD metric.

Theorem 7 establishes contraction for the *incomplete* action–gradient operator  $T_{\pi}^{S_a}$ , with coefficient  $c_k(\bar{L})$  where  $\bar{L}=\gamma\,\kappa_{\rm eff}$ .

**Lemma 10** (Affine Bellman update = affine pushforward). Fix (s,a). Let C collect all environment/policy randomness, and let (S',A')=g(s,a;C). Let  $\eta$  map each (x,u) to a law  $\eta(x,u)$  on  $\mathbb{R}^d$ , and let  $X' \sim \eta(S',A')$  (conditionally on C). Given an offset  $b_{s,a}: \operatorname{supp}(C) \to \mathbb{R}^d$  and a measurable matrix map  $L_{s,a}: \operatorname{supp}(C) \to \mathbb{R}^{d \times d}$ , define

$$\Phi_{s,a}(x;C) = b_{s,a}(C) + L_{s,a}(C) x.$$

Then

$$(T^{\pi}\eta)(s,a) = \text{Law}(\Phi_{s,a}(X';C)).$$

*Proof.* Fix a Borel set  $A \subset \mathbb{R}^d$ . Using the definition of pushforward laws,

$$\Pr(\Phi_{s,a}(X';C) \in A) = \mathbb{E}_C \left[ \Pr(b_{s,a}(C) + L_{s,a}(C)X' \in A \mid C) \right]$$
$$= \mathbb{E}_C \left[ \left( x \mapsto b_{s,a}(C) + L_{s,a}(C)x \right)_{\#} \eta(S',A')(A) \right].$$

By definition of the distributional Bellman operator with affine update  $z \mapsto b_{s,a}(C) + L_{s,a}(C)z$  and next index (S', A'), the right-hand side equals  $(T^{\pi}\eta)(s,a)(A)$ . Since this holds for all Borel A, the laws coincide.

**Lemma 11** (Scale–Lipschitz property of (squared) MMD with MQ kernel). Let  $k_h(x,y) = -\sqrt{1+h^2\|x-y\|^2}$  with h>0.

For probability measures  $\mu, \nu$  on  $\mathbb{R}^d$  with finite second moments, define the population MMD<sup>2</sup> by

$$\mathrm{MMD}_{k_h}^2(\mu,\nu) = \mathbb{E} k_h(X,X') + \mathbb{E} k_h(Y,Y') - 2 \mathbb{E} k_h(X,Y),$$

for  $X, X' \sim \mu$  i.i.d. and  $Y, Y' \sim \nu$  i.i.d.

For the scaling map  $S_s: x \mapsto sx$  with  $s \ge 0$ , we have

$$\mathrm{MMD}_{k_h}^2 \big( (S_s)_{\#} \mu, (S_s)_{\#} \nu \big) \leq c_2(s) \, \mathrm{MMD}_{k_h}^2 (\mu, \nu), \qquad c_2(s) := \max\{ s, \ s^2 \}.$$

Consequently, the (unsquared) MMD satisfies

$$\mathrm{MMD}_{k_h} \big( (S_s)_{\#} \mu, (S_s)_{\#} \nu \big) \leq c_1(s) \, \mathrm{MMD}_{k_h} (\mu, \nu), \qquad c_1(s) := \max \{ \sqrt{s}, \ s \}.$$

In particular, for s < 1 the map  $S_s$  is a contraction for both  $MMD_{k_h}^2$  and  $MMD_{k_h}$ .

Proof. Set

$$\phi(r) = \sqrt{1 + h^2 r^2}.$$

With this notation,

$$\mathrm{MMD}_{k_{b}}^{2}(\mu,\nu) = 2 \mathbb{E} \phi(\|X - Y\|) - \mathbb{E} \phi(\|X - X'\|) - \mathbb{E} \phi(\|Y - Y'\|).$$

When  $0 \le s \le 1$ , note that  $\phi(0) = 1$  and  $\phi$  is convex, as we have

$$\phi'(r) = \frac{h^2 r}{\sqrt{1 + h^2 r^2}}, \qquad \phi''(r) = \frac{h^2}{(1 + h^2 r^2)^{3/2}} \ge 0.$$

By convexity, for any  $a, b \in \mathbb{R}$  and  $s \in [0, 1]$ ,

$$\phi((1-s)a+sb) \le (1-s)\phi(a)+s\phi(b).$$

Taking a = 0, b = r, and recalling  $\phi(0) = 1$ , this gives

$$\phi(sr) \le (1-s) + s \phi(r), \qquad 0 \le s \le 1.$$

Applying this inequality inside each expectation, the constants cancel in the linear combination since (2-1-1)(1-s)=0. Therefore

$$\mathrm{MMD}_{k_b}^2((S_s)_{\#}\mu, (S_s)_{\#}\nu) \leq s \, \mathrm{MMD}_{k_b}^2(\mu, \nu).$$

When  $s \ge 1$ , consider  $f(u) = \sqrt{1 + h^2 u}$  for  $u \ge 0$ ; it is concave as

$$f'(u) = \frac{h^2}{2\sqrt{1+h^2u}}, \qquad f''(u) = -\frac{h^4}{4(1+h^2u)^{3/2}} \le 0.$$

By definition,  $\phi(r) = f(r^2)$ . For any  $u \ge 0$  and  $\lambda \ge 1$ , concavity gives

$$f(u) = f\left(\left(1 - \frac{1}{\lambda}\right) \cdot 0 + \frac{1}{\lambda} \cdot (\lambda u)\right) \ge \left(1 - \frac{1}{\lambda}\right) f(0) + \frac{1}{\lambda} f(\lambda u),$$

hence

$$f(\lambda u) \le \lambda f(u) - (\lambda - 1)f(0).$$

Taking  $u = r^2$ ,  $\lambda = s^2$ , and recalling that f(0) = 1, we obtain

$$\phi(sr) = \sqrt{1 + h^2 s^2 r^2} \le s^2 \phi(r) - (s^2 - 1).$$

Again inserting this inequality into the definition of MMD<sup>2</sup>, the constants cancel as before, and we obtain

$$\mathrm{MMD}_{k_h}^2((S_s)_{\#}\mu, (S_s)_{\#}\nu) \leq s^2 \, \mathrm{MMD}_{k_h}^2(\mu, \nu).$$

Combining both cases, the multiplicative factor is s for  $0 \le s \le 1$  and  $s^2$  for  $s \ge 1$ . Hence

$$c_2(s) = \max\{s, s^2\}.$$

Taking square roots gives the corresponding bound for the unsquared MMD,

$$c_1(s) = \max\{\sqrt{s}, s\}.$$

**Lemma 12** (Mixture p-convexity of  $\mathrm{MMD}_k$  in an RKHS). Let  $k: X \times X \to \mathbb{R}$  be a symmetric positive-semidefinite reproducing kernel with RKHS  $(\mathcal{H}, \langle \cdot, \cdot \rangle)$  and feature map  $\phi(x) = k(x, \cdot)$ . Let  $(\Omega, \mathcal{F}, \rho)$  be a probability space, and let  $(\mu_c)_{c \in \Omega}$  and  $(\nu_c)_{c \in \Omega}$  be measurable families of probability measures on X for which the mean embeddings  $\mu_{\mu_c} := \int_X \phi \, d\mu_c$  and  $\mu_{\nu_c} := \int_X \phi \, d\nu_c$  exist in  $\mathcal{H}$ . Define the mixtures  $\bar{\mu} := \int_{\Omega} \mu_c \, \rho(dc)$  and  $\bar{\nu} := \int_{\Omega} \nu_c \, \rho(dc)$ . Assume all mean embeddings and integrals below are well defined. Then for every  $p \geq 1$ ,

$$\mathrm{MMD}_{k}(\bar{\mu}, \bar{\nu}) \leq \left( \int_{\Omega} \mathrm{MMD}_{k}(\mu_{c}, \nu_{c})^{p} \, \rho(dc) \right)^{1/p}$$

Proof. By linearity of mean embeddings,

$$\mu_{\bar{\mu}} = \int_{\Omega} \mu_{\mu_c} \, \rho(dc), \qquad \mu_{\bar{\nu}} = \int_{\Omega} \mu_{\nu_c} \, \rho(dc),$$

where  $\mu_{\mu_c} = \int_X \phi(x) d\mu_c(x)$  and  $\mu_{\nu_c} = \int_X \phi(x) d\nu_c(x)$  are elements of  $\mathcal{H}$ . Thus,

$$\mu_{\bar{\mu}} - \mu_{\bar{\nu}} = \int_{\Omega} v(c) \, \rho(dc), \qquad v(c) := \mu_{\mu_c} - \mu_{\nu_c} \in \mathcal{H}.$$

Hence

$$\begin{aligned} \mathrm{MMD}_k(\bar{\mu}, \bar{\nu}) &= \|\mu_{\bar{\mu}} - \mu_{\bar{\nu}}\|_{\mathcal{H}} = \left\| \int_{\Omega} v(c) \, \rho(dc) \right\|_{\mathcal{H}} \\ &\leq \int_{\Omega} \|v(c)\|_{\mathcal{H}} \, \rho(dc) \qquad \qquad \text{(triangle inequality in } \mathcal{H}) \\ &\leq \left( \int_{\Omega} \|v(c)\|_{\mathcal{H}}^p \, \rho(dc) \right)^{1/p} \qquad \qquad (L^1 \leq L^p \text{ on a probability space)}. \end{aligned}$$

Finally,  $||v(c)||_{\mathcal{H}} = ||\mu_{\mu_c} - \mu_{\nu_c}||_{\mathcal{H}} = \mathrm{MMD}_k(\mu_c, \nu_c)$ , which gives the claim.

**Lemma 13** (Mixture p-convexity for CPD kernels via the distance-induced RKHS). Let  $k: X \times X \to \mathbb{R}$  be conditionally positive definite (CPD) and let  $k^{\circ}$  be the associated distance-induced (one-point centered) kernel from Proposition 3, so that for all probabilities P, Q with finite integrals,

$$\gamma_k(P,Q) = \text{MMD}_{k^{\circ}}(P,Q).$$

Let  $(\Omega, \mathcal{F}, \rho)$  be a probability space, and let  $(\mu_c)_{c \in \Omega}$  and  $(\nu_c)_{c \in \Omega}$  be measurable families of probability measures on X with finite embeddings for  $k^{\circ}$ . Define the mixtures  $\bar{\mu} := \int_{\Omega} \mu_c \, \rho(dc)$  and  $\bar{\nu} := \int_{\Omega} \nu_c \, \rho(dc)$ . Then for every  $p \geq 1$ ,

$$\gamma_k(\bar{\mu},\bar{\nu}) \leq \left(\int_{\Omega} \gamma_k(\mu_c,\nu_c)^p \, \rho(dc)\right)^{1/p}$$

*Proof.* By Proposition 3,  $\gamma_k = \text{MMD}_{k^{\circ}}$ . Applying Lemma 12 to the PSD kernel  $k^{\circ}$  and the families  $(\mu_c), (\nu_c)$  yields

$$\mathrm{MMD}_{k^{\circ}}(\bar{\mu}, \bar{\nu}) \leq \left(\int_{\Omega} \mathrm{MMD}_{k^{\circ}}(\mu_{c}, \nu_{c})^{p} \, \rho(dc)\right)^{1/p}.$$

Replacing MMD<sub> $k^{\circ}$ </sub> by  $\gamma_k$  via Proposition 3 gives the claim.

**Lemma 14** (Max–sliced affine push-forward contraction — anisotropic linear case). Let  $\Delta$  be a divergence on  $\mathcal{P}(\mathbb{R})$ . Assume that for all  $\mu, \nu \in \mathcal{P}(\mathbb{R})$ :

**(T)** Translation invariance: for every  $t \in \mathbb{R}$ ,

$$\Delta((x \mapsto x + t)_{\#}\mu, (x \mapsto x + t)_{\#}\nu) = \Delta(\mu, \nu).$$

(S) Scale-Lipschitz: there exists a nondecreasing  $c:[0,\infty)\to[0,\infty)$  such that for every  $s\geq 0$ ,

$$\Delta((x \mapsto sx)_{\#}\mu, (x \mapsto sx)_{\#}\nu) \le c(s) \Delta(\mu, \nu).$$

Define the max-sliced lift of  $\Delta$  by

$$\mathbf{MS}\Delta(\mu,\nu) := \sup_{\theta \in \mathbb{S}^{d-1}} \Delta((P_{\theta})_{\#}\mu, (P_{\theta})_{\#}\nu), \qquad P_{\theta}(x) = \langle \theta, x \rangle.$$

Let F(x) = Ax + b with an arbitrary matrix  $A \in \mathbb{R}^{d \times d}$  and  $b \in \mathbb{R}^d$ , and denote  $L := ||A||_{\text{op}} = \sup_{\|v\|=1} ||Av||$ . Then, for all  $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ ,

$$\mathbf{MS}\Delta\big(F_{\#}\mu,F_{\#}\nu\big) \ \leq \ c(L)\,\mathbf{MS}\Delta(\mu,\nu).$$

*Proof.* Fix  $\theta \in \mathbb{S}^{d-1}$  and set  $w_{\theta} := A^{\top} \theta$ .

Case 1:  $w_{\theta} = 0$ . Then  $(P_{\theta} \circ F)(x) = \langle \theta, b \rangle$  is constant, hence

$$\Delta((P_{\theta})_{\#}F_{\#}\mu, (P_{\theta})_{\#}F_{\#}\nu) = 0 \tag{27}$$

$$\leq c(0) \Delta((P_{\phi})_{\#}\mu, (P_{\phi})_{\#}\nu)$$
 for any unit  $\phi$ , (28)

so the desired bound holds trivially.

Case 2:  $||w_{\theta}|| > 0$ . Write  $r_{\theta} := ||w_{\theta}||$  and  $\phi_{\theta} := w_{\theta}/r_{\theta} \in \mathbb{S}^{d-1}$ . For any  $X \sim \mu$  and  $Y \sim \nu$ ,

$$(P_{\theta} \circ F)(X) = \langle \theta, AX + b \rangle = \langle \theta, b \rangle + r_{\theta} \langle \phi_{\theta}, X \rangle,$$

and similarly for Y. By (T) and (S) we obtain

$$\Delta((P_{\theta})_{\#}F_{\#}\mu, (P_{\theta})_{\#}F_{\#}\nu) = \Delta(\operatorname{Law}(r_{\theta}\langle\phi_{\theta}, X\rangle), \operatorname{Law}(r_{\theta}\langle\phi_{\theta}, Y\rangle))$$
(29)

$$\leq c(r_{\theta}) \Delta((P_{\phi_{\theta}})_{\#}\mu, (P_{\phi_{\theta}})_{\#}\nu). \tag{30}$$

**Taking the supremum.** Now take the supremum over  $\theta \in \mathbb{S}^{d-1}$ :

$$\sup_{\theta} \Delta ((P_{\theta})_{\#} F_{\#} \mu, (P_{\theta})_{\#} F_{\#} \nu) \leq \sup_{\theta} c(r_{\theta}) \sup_{\phi} \Delta ((P_{\phi})_{\#} \mu, (P_{\phi})_{\#} \nu). \tag{31}$$

Since  $r_{\theta} = \|A^{\top}\theta\| \le \|A^{\top}\|_{\text{op}} = \|A\|_{\text{op}} = L$  and c is nondecreasing,

$$\sup_{\theta} \Delta ((P_{\theta})_{\#} F_{\#} \mu, (P_{\theta})_{\#} F_{\#} \nu) \le c(L) \operatorname{\mathbf{MS}} \Delta(\mu, \nu).$$
(32)

The left-hand side is exactly  $MS\Delta(F_{\#}\mu, F_{\#}\nu)$ , which proves the claim.

**Lemma 15** (Max-sliced mixture p-convexity). Let  $\Delta$  be a divergence on  $\mathcal{P}(\mathbb{R})$  that is mixture p-convex for some  $p \in [1, \infty)$ : for every probability space  $(\Omega, \mathcal{F}, \rho)$  and measurable families  $(\mu_c), (\nu_c) \subset \mathcal{P}(\mathbb{R})$ ,

$$\Delta \left( \int_{\Omega} \mu_c \, \rho(dc), \, \int_{\Omega} \nu_c \, \rho(dc) \right) \, \leq \, \left( \int_{\Omega} \Delta(\mu_c, \nu_c)^p \, \rho(dc) \right)^{1/p}.$$

Define the max–sliced lift on  $\mathcal{P}(\mathbb{R}^d)$  by

$$\mathbf{MS}\Delta(\mu,\nu) := \sup_{\theta \in \mathbb{S}^{d-1}} \Delta((P_{\theta})_{\#}\mu, (P_{\theta})_{\#}\nu), \qquad P_{\theta}(x) = \langle \theta, x \rangle.$$

Then  $MS\Delta$  is also mixture p-convex:

$$\mathbf{MS}\Delta\bigg(\int_{\Omega}\mu_{c}\,\rho(dc),\,\int_{\Omega}\nu_{c}\,\rho(dc)\bigg)\,\,\leq\,\, \Bigg(\int_{\Omega}\mathbf{MS}\Delta(\mu_{c},\nu_{c})^{p}\,\rho(dc)\Bigg)^{1/p}.$$

*Proof.* Fix  $\theta \in \mathbb{S}^{d-1}$  and set

$$\mu_c^{\theta} := (P_{\theta})_{\#} \mu_c, \qquad \nu_c^{\theta} := (P_{\theta})_{\#} \nu_c \in \mathcal{P}(\mathbb{R}).$$

Pushforward commutes with mixtures:

$$(P_{\theta})_{\#} \Big( \int \mu_c \, d\rho \Big) = \int \mu_c^{\theta} \, d\rho, \qquad (P_{\theta})_{\#} \Big( \int \nu_c \, d\rho \Big) = \int \nu_c^{\theta} \, d\rho.$$

By mixture p-convexity of  $\Delta$  in one dimension,

$$\Delta((P_{\theta})_{\#} \int \mu_c \, d\rho, \, (P_{\theta})_{\#} \int \nu_c \, d\rho) \leq \left(\int \Delta(\mu_c^{\theta}, \nu_c^{\theta})^p \, d\rho\right)^{1/p}. \tag{33}$$

Taking the supremum over  $\theta$  on the left-hand side of equation 33 gives

$$\sup_{\theta} \Delta \left( (P_{\theta})_{\#} \int \mu_c \, d\rho, \, (P_{\theta})_{\#} \int \nu_c \, d\rho \right) \leq \sup_{\theta} \left( \int \Delta (\mu_c^{\theta}, \nu_c^{\theta})^p \, d\rho \right)^{1/p}. \tag{34}$$

Define  $f(\theta,c) := \Delta(\mu_c^{\theta},\nu_c^{\theta})$  and  $h(c) := \sup_{\phi} f(\phi,c) = \mathbf{MS}\Delta(\mu_c,\nu_c)$ . Since  $f(\theta,c) \leq h(c)$  pointwise in c, we obtain for every  $\theta$ ,

$$\left(\int f(\theta,c)^p \, d\rho(c)\right)^{1/p} \, \leq \, \left(\int h(c)^p \, d\rho(c)\right)^{1/p}.$$

Taking  $\sup_{\theta}$  yields

$$\sup_{\theta} \left( \int \Delta(\mu_c^{\theta}, \nu_c^{\theta})^p \, d\rho \right)^{1/p} \leq \left( \int \mathbf{MS} \Delta(\mu_c, \nu_c)^p \, d\rho \right)^{1/p}. \tag{35}$$

Combining equation 34 and equation 35 shows

$$\mathbf{MS}\Delta\bigg(\int \mu_c \, d\rho, \, \int \nu_c \, d\rho\bigg) \, \leq \, \Bigg(\int \mathbf{MS}\Delta(\mu_c, \nu_c)^p \, \rho(dc)\Bigg)^{1/p},$$

as claimed.

Lemma 16 (Supremum–max–sliced contraction of the multivariate distributional Bellman operator (anisotropic linear map)). Let  $\Delta$  be a divergence on  $\mathcal{P}(\mathbb{R})$  and define the max–sliced lift on  $\mathcal{P}(\mathbb{R}^d)$  by

$$\mathbf{MS}\Delta(\mu,\nu) := \sup_{\theta \in \mathbb{S}^{d-1}} \Delta\big( (P_{\theta})_{\#}\mu, \, (P_{\theta})_{\#}\nu \big), \qquad P_{\theta}(x) = \langle \theta, x \rangle.$$

Assume  $\Delta$  satisfies:

- (T) Translation nonexpansion:  $\Delta((x \mapsto x + t)_{\#}\mu, (x \mapsto x + t)_{\#}\nu) \leq \Delta(\mu, \nu)$  for all  $t \in \mathbb{R}$ .
- (S) Scale-Lipschitz: there exists a nondecreasing  $c:[0,\infty)\to [0,\infty)$  such that, for all  $s\geq 0$ ,  $\Delta\big((x\mapsto sx)_\#\mu,\,(x\mapsto sx)_\#\nu\big)\leq c(s)\,\Delta(\mu,\nu).$
- (M<sub>p</sub>) Mixture p-convexity: for every probability space  $(\Omega, \mathcal{F}, \rho_0)$  and measurable families  $(\mu_c), (\nu_c) \subset \mathcal{P}(\mathbb{R}),$

$$\Delta \left( \int \mu_c \, \rho_0(dc), \, \int \nu_c \, \rho_0(dc) \right) \, \leq \, \left( \int \Delta(\mu_c, \nu_c)^p \, \rho_0(dc) \right)^{1/p}, \qquad p \in [1, \infty).$$

**Bellman update (anisotropic linear map).** Fix (s, a). Gather all environment/policy randomness into a single random element C, which determines the successor index through a measurable mapping g:

$$(S', A') := g(s, a; C).$$

At (s,a), apply an affine transformation with a C-dependent translation and an arbitrary C-dependent linear map:

$$b_{s,a}(C) \in \mathbb{R}^d$$
,  $A_{s,a}(C) \in \mathbb{R}^{d \times d}$ .

Conditioned on C, the next sample is drawn from the law at the successor index,

$$X' \mid C \sim \eta(S', A'),$$

and the Bellman update is the push-forward of X' by this affine map:

$$(T^{\pi}\eta)(s,a) := \text{Law}(b_{s,a}(C) + A_{s,a}(C) X').$$

Define, for each C,

$$L(C) := \|A_{s,a}(C)\|_{\text{op}},$$

and the global envelope

$$\bar{L} := \sup_{(s,a)} \sup_{C} L(C).$$

Also define the supremum metric

$$\overline{\mathbf{MS}\Delta}(\eta_1, \eta_2) := \sup_{(s,a)} \mathbf{MS}\Delta(\eta_1(s,a), \eta_2(s,a)).$$

Then, for all  $\eta_1, \eta_2$ ,

$$\overline{\mathbf{MS}\Delta}(T^{\pi}\eta_1, T^{\pi}\eta_2) \leq c(\bar{L}) \overline{\mathbf{MS}\Delta}(\eta_1, \eta_2).$$

*Proof.* Fix (s, a) and condition on C. Set

$$\Phi_{s,a}(x;C) := b_{s,a}(C) + A_{s,a}(C) x, \qquad Z_i := \Phi_{s,a}(X_i';C),$$

with  $X_i' \mid C \sim \eta_i(S', A')$ . By Lemma 10,

$$(T^{\pi}\eta_i)(s,a) = \operatorname{Law}(\Phi_{s,a}(X_i';C)) = \operatorname{Law}(Z_i).$$

**Affine push-forward at fixed** C**.** Applying Lemma 14, which relies on (T) and (S), to the conditional laws  $\text{Law}(X'_i \mid C)$  gives

$$\mathbf{MS}\Delta(\operatorname{Law}(Z_1 \mid C), \operatorname{Law}(Z_2 \mid C)) \leq c(L(C)) \,\mathbf{MS}\Delta(\operatorname{Law}(X_1' \mid C), \operatorname{Law}(X_2' \mid C))$$

$$= c(L(C)) \,\mathbf{MS}\Delta(\eta_1(S', A'), \eta_2(S', A')).$$
(36)

Averaging over C. Lemma 15, which relies on  $(M_p)$ , together with equation 36 yields

$$\mathbf{MS}\Delta(\mathrm{Law}(Z_{1}), \, \mathrm{Law}(Z_{2})) \leq \left(\int \mathbf{MS}\Delta(\mathrm{Law}(Z_{1} \mid C), \, \mathrm{Law}(Z_{2} \mid C))^{p} \, \rho(dC)\right)^{1/p}$$

$$\leq \left(\int \left(c(L(C)) \, \mathbf{MS}\Delta(\eta_{1}(S', A'), \, \eta_{2}(S', A'))\right)^{p} \, \rho(dC)\right)^{1/p}$$
(39)

$$\leq c(\bar{L}) \left( \int \mathbf{MS} \Delta \left( \eta_1(S', A'), \eta_2(S', A') \right)^p \rho(dC) \right)^{1/p}, \tag{40}$$

since c is nondecreasing and  $L(C) \leq \bar{L}$  for all C.

**Supremum bound.** For any realization of C, by definition of the supremum metric,

$$\mathbf{MS}\Delta\big(\eta_1(S',A'),\eta_2(S',A')\big) \leq \sup_{(u,v)} \mathbf{MS}\Delta\big(\eta_1(u,v),\eta_2(u,v)\big) = \overline{\mathbf{MS}\Delta}(\eta_1,\eta_2).$$

Combining this with the previous inequality,

$$\mathbf{MS}\Delta((T^{\pi}\eta_1)(s,a), (T^{\pi}\eta_2)(s,a)) = \mathbf{MS}\Delta(\mathrm{Law}(Z_1), \mathrm{Law}(Z_2))$$
(41)

$$\leq c(\bar{L}) \left( \int \overline{\mathbf{MS}\Delta}(\eta_1, \eta_2)^p \, \rho(dC) \right)^{1/p}$$
 (42)

$$= c(\bar{L}) \, \overline{\mathbf{MS}\Delta}(\eta_1, \eta_2). \tag{43}$$

Taking the supremum over (s, a) completes the proof:

$$\overline{\mathbf{MS}\Delta}(T^{\pi}\eta_1, T^{\pi}\eta_2) \leq c(\bar{L}) \overline{\mathbf{MS}\Delta}(\eta_1, \eta_2).$$

**Lemma 17** (Fixed-point law of the distributional Bellman operator (general linear discount)). *Define the infinite–horizon return under policy*  $\pi$  *recursively by* 

$$Z(s,a) \stackrel{d}{=} \Phi_{s,a}(Z(S',A');C),$$

where C collects the one-step randomness, (S',A')=g(s,a;C) is the successor pair, and

$$\Phi_{s,a}(x;C) := r(s,a;C) + \Gamma(s,a;C) x, \qquad r(s,a;C) \in \mathbb{R}^d, \ \Gamma(s,a;C) \in \mathbb{R}^{d \times d}.$$

Equivalently, along a trajectory  $(S_t, A_t)$  with one-step randomness  $(C_t)_{t>0}$ , set

$$r_t := r(S_t, A_t; C_t), \qquad \Gamma_t := \Gamma(S_t, A_t; C_t), \qquad \Pi_{0:t-1} := \Gamma_0 \Gamma_1 \cdots \Gamma_{t-1} \ (\Pi_{0:t-1} := I_d),$$

and, whenever the series converges,

$$Z(s,a) = \sum_{t=0}^{\infty} \prod_{0:t-1} r_t.$$

Set

$$\eta^{\pi}(s, a) := \text{Law}(Z(s, a)) \in \mathcal{P}(\mathbb{R}^d).$$

$$T_{\pi} \eta^{\pi} = \eta^{\pi}.$$

Proof. By definition,

$$Z(s,a) \stackrel{d}{=} \Phi_{s,a}(Z(S',A');C), \qquad (S',A') = g(s,a;C).$$

Conditioning on C gives

$$Z(S', A') | C \sim \eta^{\pi}(S', A').$$

By the push–forward law (Lemma 1),

$$\operatorname{Law}(Z(s,a)) = \operatorname{Law}(\Phi_{s,a}(X';C)), \qquad X' \mid C \sim \eta^{\pi}(S',A').$$

By definition of the distributional Bellman operator,  $(T_{\pi}\eta^{\pi})(s,a) = \text{Law}(\Phi_{s,a}(X';C))$ , hence  $(T_{\pi}\eta^{\pi})(s,a) = \eta^{\pi}(s,a)$  for all (s,a).

**Theorem 6** (Supremum–max–sliced MMD contraction via the generic lemma). Let  $k : \mathbb{R} \to \mathbb{R}$  be a 1D kernel and let  $\mathrm{MMD}_k$  be its 1D maximum mean discrepancy. Assume:

(T) Translation nonexpansion (1D). For all  $t \in \mathbb{R}$  and  $\alpha, \beta \in \mathcal{P}(\mathbb{R})$ ,  $\mathrm{MMD}_k((x \mapsto x + t)_{\#}\alpha, (x \mapsto x + t)_{\#}\beta) \leq \mathrm{MMD}_k(\alpha, \beta)$ .

**Scale–Lipschitz (1D).** There exists a nondecreasing  $c_k : [0, \infty) \to [0, \infty)$  such that for all  $s \ge 0$  and  $\alpha, \beta$ ,  $\mathrm{MMD}_k((x \mapsto sx)_\# \alpha, (x \mapsto sx)_\# \beta) \le c_k(s) \, \mathrm{MMD}_k(\alpha, \beta)$ .

 $(\mathbf{M}_p)$  Mixture p-convexity.  $\mathrm{MMD}_k$  is mixture p-convex on  $\mathcal{P}(\mathbb{R})$ .

Define the max-sliced divergence  $\mathbf{MSMMD}_k(\mu, \nu) := \sup_{\theta \in \mathbb{S}^{d-1}} \mathrm{MMD}_k((P_\theta)_\# \mu, (P_\theta)_\# \nu)$  and the supremum metric  $\overline{\mathbf{MSMMD}}_k(\eta_1, \eta_2) := \sup_{(s,a)} \mathbf{MSMMD}_k(\eta_1(s,a), \eta_2(s,a))$ .

Consider the anisotropic affine Bellman update  $(T^{\pi}\eta)(s,a) = \text{Law}(b_{s,a}(C) + A_{s,a}(C) X')$  with  $X' \mid C \sim \eta(S',A')$ ,  $L(C) := ||A_{s,a}(C)||_{\text{op}}$ , and  $\bar{L} := \sup_{(s,a) \in C} L(C)$ . Then, for all  $\eta_1, \eta_2$ ,

$$\overline{\mathbf{MSMMD}}_k (T^{\pi} \eta_1, T^{\pi} \eta_2) \leq c_k(\overline{L}) \overline{\mathbf{MSMMD}}_k (\eta_1, \eta_2).$$

In particular, if  $c_k(\bar{L}) < 1$ , then  $T^{\pi}$  is a contraction on  $(S \times A \to \mathcal{P}(\mathbb{R}^d), \overline{\mathbf{MSMMD}}_k)$ .

*Proof.* Assumption (T) is exactly the 1D translation nonexpansion for  $MMD_k$ . gives the 1D scale-Lipschitz control with  $c = c_k$ . (M<sub>p</sub>) supplies the mixture p-convexity.

Applying Lemma 16 with  $\Delta = \text{MMD}_k$  then yields

$$\overline{\mathbf{MSMMD}}_k(T^{\pi}\eta_1, T^{\pi}\eta_2) \leq c_k(\bar{L}) \overline{\mathbf{MSMMD}}_k(\eta_1, \eta_2).$$

**Corollary 6.1** (Supremum-max-sliced contraction for MQ-MMD). Let  $k_h(x,y) = -\sqrt{1+h^2\|x-y\|^2}$  with h>0 (the multiquadric kernel). Assume the Bellman update is  $X'\mid C\sim \eta(S',A')$  and  $(T^\pi\eta)(s,a)=\mathrm{Law}\big(b_{s,a}(C)+A_{s,a}(C)\,X'\big)$ , with  $L(C):=\|A_{s,a}(C)\|_{\mathrm{op}}$  and  $\bar{L}:=\sup_{(s,a)}\sup_{C}L(C)$ . Then, for all return-law maps  $\eta_1,\eta_2$ ,

$$\overline{\mathbf{MSMMD}}_{k_h} (T^{\pi} \eta_1, T^{\pi} \eta_2) \leq c_1(\overline{L}) \overline{\mathbf{MSMMD}}_{k_h} (\eta_1, \eta_2),$$
$$c_1(s) := \max\{\sqrt{s}, s\}.$$

*Proof.* Since  $k_h$  is radial it depends only on  $\|x-y\|$ , hence  $\mathrm{MMD}_{k_h}$  is translation–nonexpansive (T). By Lemma 11,  $\mathrm{MMD}_{k_h}$  satisfies the 1D scale–Lipschitz property with factor  $c_1(s) = \max\{\sqrt{s}, s\}$ . Moreover,  $k_h$  is conditionally positive definite (CPD) by Theorem 3.1 of Killingberg & Langseth (2023), so mixture p–convexity ( $\mathbf{M}_p$ ) for  $\mathrm{MMD}_{k_h}$  follows from Lemma 13. Applying Lemma 16 with  $\Delta = \mathrm{MMD}_{k_h}$  yields the stated contraction with constant  $c_1(\bar{L})$ .

**Corollary 6.2** (Supremum-max-sliced contraction of the (**complete**) Sobolev Bellman operator). Let  $T_{\pi}^{S_{s,a}}$  be the complete Sobolev Bellman operator defined in Theorem 3, bootstrapping the vector  $(Z, \partial_a Z, \partial_s Z)$ . Equip  $\mathcal{P}(\mathbb{R}^{1+m+n})$  with the max-sliced MMD metric based on the MQ kernel  $k_h(x,y) = -\sqrt{1+h^2\|x-y\|^2}$ . Then, for all Sobolev return-law maps  $\eta_1, \eta_2$ ,

$$\overline{\mathbf{MSMMD}}_{k_h} (T_{\pi}^{S_{s,a}} \eta_1, T_{\pi}^{S_{s,a}} \eta_2) \leq c_1(\|\mathcal{L}\|_d) \overline{\mathbf{MSMMD}}_{k_h} (\eta_1, \eta_2),$$
$$c_1(s) = \max\{\sqrt{s}, s\}.$$

*Proof.* The proof follows the same structure as Corollary 6.1, except that the affine map now arises from the Sobolev Bellman operator rather than the standard distributional Bellman update. The only difference is that the linear component of the affine map is now  $\mathcal{L}(s, a; \varepsilon_f, \varepsilon_\pi)$  defined in Theorem 3. By the argument there, its operator norm is bounded by  $\|\mathcal{L}\|_d$ , hence the contraction factor becomes  $c_1(\|\mathcal{L}\|_d)$ .

**Lemma 18** (Triangle inequality for MMD (PSD) and  $\gamma_k$  (CPD)). Let X be a measurable space. For probability measures P, Q, R on X, the following triangle inequalities hold:

$$\mathrm{MMD}_k(P,Q) \leq \mathrm{MMD}_k(P,R) + \mathrm{MMD}_k(R,Q)$$
 if k is PSD,

and

$$\gamma_k(P,Q) \leq \gamma_k(P,R) + \gamma_k(R,Q)$$
 if k is CPD.

Proof. When k is positive semi-definite with RKHS  $\mathcal{H}$ , we have

$$\begin{aligned} \text{MMD}_{k}(P, Q) &= \|\mu_{P} - \mu_{Q}\|_{\mathcal{H}} \\ &\leq \|\mu_{P} - \mu_{R}\|_{\mathcal{H}} + \|\mu_{R} - \mu_{Q}\|_{\mathcal{H}} \\ &= \text{MMD}_{k}(P, R) + \text{MMD}_{k}(R, Q), \end{aligned}$$

by the triangle inequality for the norm in H.

When k is conditionally positive definite, one may fix  $z_0 \in X$  and form the distance-induced kernel

$$k^{\circ}(x,y) = k(x,y) - k(x,z_0) - k(z_0,y) + k(z_0,z_0),$$

which is positive semi-definite. Proposition 3 ensures that  $\gamma_k(P,Q) = \mathrm{MMD}_{k^{\circ}}(P,Q)$ , so applying the previous argument to  $k^{\circ}$  yields the desired inequality.

**Theorem 7** (Supremum–max–sliced MMD contraction of the **incomplete** Sobolev Bellman operator). Let  $T_{\pi}^{S_a}: (\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m})) \to (\mathcal{S} \times \mathcal{A} \to \mathcal{P}(\mathbb{R}^{1+m}))$  be the action–gradient (incomplete) Sobolev Bellman operator that updates  $H = (Z, \partial_a Z)$  as in Eq. equation 17.

Assumptions on the 1D base divergence. Let the one-dimensional base divergence be  $\Delta = \text{MMD}_k$  for a CPD, shift-invariant kernel k on  $\mathbb{R}$ . Assume:

(T) Translation nonexpansion:  $\mathrm{MMD}_k((x \mapsto x + t)_{\#}\mu, (x \mapsto x + t)_{\#}\nu) \leq \mathrm{MMD}_k(\mu, \nu)$  for all  $t \in \mathbb{R}$ .

**Scale–Lipschitz:** There exists a nondecreasing  $c_k:[0,\infty)\to[0,\infty)$  such that  $\mathrm{MMD}_k((x\mapsto sx)_\#\mu,(x\mapsto sx)_\#\nu)\leq c_k(s)\,\mathrm{MMD}_k(\mu,\nu)$  for all  $s\geq 0$ .

 $(\mathbf{M}_p)$  Mixture p-convexity (with p=1): Lemma 13 holds, and its max-sliced lift Lemma 15 applies.

Jacobian bounds. Assume

$$||f_a||_2 \le L_{f,a}, \qquad ||\pi_s||_2 \le L_{\pi}.$$

For the linear blocks in equation 17,

$$A(s,a;\varepsilon_f,\varepsilon_\pi) = \gamma \begin{pmatrix} 1 & 0 \\ 0 & f_a^\top \, \pi_s^\top \end{pmatrix}, \qquad N(s,a;\varepsilon_f) = \gamma \begin{pmatrix} 0 \\ f_a^\top \end{pmatrix}.$$

Set

$$||A||_{\text{op}} \leq \gamma \max\{1, L_{f,a}L_{\pi}\}, \qquad ||N||_{\text{op}} \leq \gamma L_{f,a}, \qquad \bar{L} := \sup_{s,a,\varepsilon_f,\varepsilon_{\pi}} ||[A\ N]||_{\text{op}}.$$

**Distributional lifting assumption.** For each (s', a'), there exists a measurable lifting map

$$\mathcal{J}_{s',a'}: \mathcal{P}(\mathbb{R}^{1+m}) \longrightarrow \mathcal{P}(\mathbb{R}^{1+m+n}),$$

sending a law  $\mu$  of  $H'=(Z',\partial_a Z')$  to a joint law  $\mathcal{J}_{s',a'}(\mu)$  of  $(H',\partial_s Z'(s',a'))$ , whose first marginal is  $\mu$  and whose second marginal matches the true  $\partial_s Z'$ -law under  $\mu$ . Assume there is a constant  $L_{\rm aug} \geq 1$  such that

$$\mathbf{MSMMD}_k(\mathcal{J}_{s',a'}(\mu), \mathcal{J}_{s',a'}(\nu)) \leq L_{\mathrm{aug}} \mathbf{MSMMD}_k(\mu,\nu), \quad \forall \mu, \nu.$$

**Conclusion.** Define  $\overline{\mathbf{MSMMD}}_k(\eta_1, \eta_2) = \sup_{(s,a)} \mathbf{MSMMD}_k(\eta_1(s,a), \eta_2(s,a))$ . Then, for all return–gradient maps  $\eta_1, \eta_2$ ,

$$\overline{\mathbf{MSMMD}}_{k}(T_{\pi}^{S_{a}}\eta_{1}, T_{\pi}^{S_{a}}\eta_{2}) \leq L_{\mathrm{aug}} c_{k}(\bar{L}) \overline{\mathbf{MSMMD}}_{k}(\eta_{1}, \eta_{2}).$$

In particular, if  $L_{\rm aug} \ c_k(\bar{L}) < 1$ , then  $T_\pi^{S_a}$  is a strict contraction.

**Remark** (MQ kernel). For the multiquadric kernel  $k_h(x,y) = -\sqrt{1 + h^2(x-y)^2}$ , Lemma 11 gives  $c_k(s) = \max\{\sqrt{s}, s\}$ . Hence the sufficient condition is  $L_{\rm aug}\sqrt{\bar{L}} < 1$  when  $\bar{L} \le 1$ , and  $L_{\rm aug}\bar{L} < 1$  when  $\bar{L} \ge 1$ .

*Proof.* Fix (s, a) and a noise draw  $C = (\varepsilon_r, \varepsilon_f, \varepsilon_\pi)$ . Set

$$s' = f(s, a; \varepsilon_f), \qquad a' = \pi(s'; \varepsilon_\pi).$$

By Lemma 10, the update is an affine pushforward:

$$(T_{\pi}^{S_a} \eta_i)(s, a) = \text{Law}(\Phi_{s,a}(X_i'; C)), \qquad \Phi_{s,a}(x) = b + Ax + Nu,$$

with  $X_i' \sim \eta_i(s', a')$  and u the corresponding  $\partial_s Z'$ .

**Augmented variable.** Let  $\tilde{X}'_i = (H'_i, U'_i) \in \mathbb{R}^{(1+m)+n}$  with law  $\tilde{\eta}_i = \mathcal{J}_{s',a'}(\eta_i(s',a'))$ . Then, conditionally on C,

$$\langle \theta, \Phi_{s,a}(H_i') \rangle = \langle \theta, b \rangle + \langle [A \ N]^\top \theta, \tilde{X}_i' \rangle.$$

**Apply affine–anisotropic lemma.** By Lemma 14, with (T) and,

$$\mathbf{MSMMD}_{k}\Big(\mathrm{Law}(\Phi_{s,a}(H_{1}^{\prime})\mid C),\ \mathrm{Law}(\Phi_{s,a}(H_{2}^{\prime})\mid C)\Big) \leq c_{k}(\|[A\mid N]\|_{\mathrm{op}})\ \mathbf{MSMMD}_{k}(\tilde{\eta}_{1},\tilde{\eta}_{2}).$$

**Lift back to**  $(Z, \partial_a Z)$ . By the lifting assumption,

$$\mathbf{MSMMD}_k(\tilde{\eta}_1, \tilde{\eta}_2) \leq L_{\text{aug}} \mathbf{MSMMD}_k(\eta_1(s', a'), \eta_2(s', a')).$$

Average over noise. Using mixture convexity for max-slice,

$$\mathbf{MSMMD}_{k}((T_{\pi}^{S_{a}}\eta_{1})(s, a), (T_{\pi}^{S_{a}}\eta_{2})(s, a)) \leq L_{\text{aug}} c_{k}(\sup_{C} ||[A \ N]||_{\text{op}}) \overline{\mathbf{MSMMD}}_{k}(\eta_{1}, \eta_{2}).$$

Taking the supremum over (s, a) yields the claim.

## I BACKGROUND ON THE WORLD-MODEL

#### I.1 CONDITIONAL VARIATIONAL AUTO-ENCODERS

A principled invertible generative model can be obtained from a Variational Auto-Encoder (VAE) (Kingma, 2013). More interestingly for us are conditional VAE Sohn et al. (2015) which we briefly introduce.

A Conditional Variational Autoencoder (cVAE) is a generative model that learns to generate new samples from a distribution conditioned on given input information. In our case, the cVAE models the distribution of next states and rewards conditioned on current states and actions.

Formally, the cVAE consists of two components:

• **Encoder**: The encoder  $q_{\zeta}(\varepsilon \mid s', r; s, a)$  maps the observed next state s' and reward r, conditioned on the current state-action pair (s, a), to a latent variable  $\varepsilon$ , typically modeled as a Gaussian distribution with diagonal covariance matrix:

$$q_{\zeta}(\varepsilon \mid s', r; s, a) = \mathcal{N}(\varepsilon; \mu_{\zeta}(s', r, s, a), \sigma_{\zeta}^{2}(s', r, s, a) \odot I). \tag{44}$$

- **Decoder**: The decoder  $p_{\psi}(s', r \mid \varepsilon; s, a)$  reconstructs the next state s' and reward r from the latent variable  $\varepsilon$ , conditioned on the current state-action pair (s, a).
- **Prior**: The prior  $p_v(s', r \mid \varepsilon; s, a)$  allows more flexibility in the latent space than a simple standard Gaussian  $\mathcal{N}(0; I)$  It is also parametrized as a multivariate diagonal Gaussian:

$$q_{\nu}(\varepsilon \mid s, a) = \mathcal{N}(\varepsilon; \mu_{\nu}(s, a), \sigma_{\nu}^{2}(s, a) \odot I). \tag{45}$$

The objective of a cVAE is to maximize the Evidence Lower Bound (ELBO), which balances accurate reconstruction of the input with a regularization term that ensures the learned posterior distribution remains close to the prior distribution. The objective is as follows

$$\mathcal{L}_{\text{cVAE}}(\zeta, \psi) = \mathbb{E}_{q_{\zeta}(\varepsilon \mid s', r; s, a)} \left[ \log p_{\psi}(s', r \mid \varepsilon; s, a) \right] -\lambda_{\text{KL}} \times D_{\text{KL}} \left( q_{\zeta}(\varepsilon \mid s', r; s, a) \parallel p_{v}(\varepsilon \mid s, a) \right).$$
(46)

The first term encourages faithful reconstruction of the next state and reward, while the second term regularizes the posterior distribution to remain close to a standard Gaussian prior. Assuming the decoder  $p_{\psi}(s',r\mid\varepsilon;s,a)$  is Gaussian with a fixed variance, the reconstruction term reduces to an L2 loss, which can be estimated using the difference between the reconstructed samples and the true samples.

Assuming the encoder parametrizes a Gaussian with diagonal covariance and that the prior is also Gaussian with identity covariance and zero mean, the KL divergence can be estimated from encoded input samples as

$$D_{KL}(\zeta) = \mathbb{E}_{(s,a)} \left[ \frac{1}{2} \sum_{j=1}^{d} \left( 1 + \log(\sigma_{\zeta,j}^{2}(s,a)) - \mu_{\zeta,j}^{2}(s,a) - \sigma_{\zeta,j}^{2}(s,a) \right) \right]. \tag{47}$$

## J ALGORITHM DESIGN, PSEUDO-CODE AND BASELINE

In this section we describe the motivation for some design choices of our method. Section J.1 sets the design for the toy reinforcement learning experiment of Section 6.1 and the Mujoco experiments from Section 6.2. We follow by providing the pseudo-code in Section J.2. Finally in Section J.3 we discuss our implementation of our primary baseline MAGE D'Oro & Jaskowski (2020).

#### J.1 DESIGN CHOICES

Firstly, most experiments are conducted in the *data-efficient* setting, where the number of updates per interaction with the environment (UTD ratio) is larger than one, making stability and overestimation bias critical concerns. MAGE adds gradient regularization on top of TD3 (Fujimoto et al., 2018), which itself incorporates several modifications to DDPG (Lillicrap et al., 2016). Below, we describe each modification and whether it was retained in our implementation.

**Target policy smoothing:** this technique applies random noise to the policy's action when estimating the TD learning target. It is the only TD3 (and MAGE) modification we retained, as other changes cluttered the implementation and introduced noise that could interact poorly with distributional modeling. It is worth noting that omitting target policy smoothing in similar settings is not uncommon (Singh et al., 2022; Kuznetsov et al., 2020).

**Delayed policy update:** the policy is updated once for every two critic updates, which helps stabilize learning by preventing premature policy shifts.

**Double estimation:** an ensemble of two critics,  $Q_1, Q_2$  or  $Z_1, Z_2$ , is used, with the bootstrapped target taken as the minimum of the two estimates at the next state (including policy smoothing noise). This modification introduces an underestimation bias to counteract the well-known overestimation

issue in value-based methods (Van Hasselt et al., 2016). We observed double estimation to be critical for stable performance in the data-efficient setting: without it, average Q-values quickly diverged. Overestimation bias can also be addressed in the distributional setting, as demonstrated by TQC (Kuznetsov et al., 2020). A straightforward approach to induce underestimation is to truncate the top p% of values in the critic's target distribution; in our setup, truncating as few as 25% of values proved highly effective. Following both TQC (Chen et al., 2021), we employ an ensemble of two distributional critics whose samples are concatenated.

### J.2 PSEUDO-CODES

In this subsection, we present the pseudo-code for the MSMMD estimation procedure, the MMD estimation procedure, and the full algorithm for policy evaluation and improvement.

- Algorithm 1 (Estimation of MSMMD): approximates the supremum over projection directions by gradient ascent on the unit sphere, yielding an empirical max–sliced MMD between two sets of samples.
- Algorithm 2 (MMD Estimation of Sobolev samples): leverages the world model to generate differentiable Sobolev-return samples, and applies truncation to mitigate overestimation bias as in Kuznetsov et al. (2020).
- Algorithm 3 (Full DSDPG algorithm): integrates our Sobolev-MMD components into the DDPG framework Lillicrap et al. (2016).

# Algorithm 1 Estimation of MSMMD from empirical samples

```
Input: Empirical samples X = \{x_i\}_{i=1}^N \subset \mathbb{R}^d, Y = \{y_i\}_{i=1}^N \subset \mathbb{R}^d
Input: Kernel k defining the MMD; gradient steps T; step size \eta
Initialize a unit direction: w \sim \mathcal{N}(0, I_d); \quad \theta \leftarrow w/||w||
                                                                                           // random unit direction
Project-optimize over directions: for t = 1, ..., T do
     u_i \leftarrow \langle \theta, x_i \rangle, \quad v_i \leftarrow \langle \theta, y_i \rangle \quad \text{for } i = 1, \dots, N
                                                                                            // project to 1D along 	heta
     J(\theta) \leftarrow \text{MMD}_k(\{u_i\}_{i=1}^N, \{v_i\}_{i=1}^N)
                                                                                                // maximize MMD over 	heta
     g \leftarrow \nabla_{\theta} J(\theta)
                                                                                   // gradient w.r.t. direction
     w \leftarrow w + \eta g
                                                               // ascent step in unconstrained space
     \theta \leftarrow w/||w||
                                                                 // re-normalize onto the unit sphere
\bar{\theta} \leftarrow \text{stop\_grad}(\theta)
                                                           // stop gradient on the final direction
Output: \widehat{MSMMD}(X,Y) \leftarrow \widehat{MMD}_k(\{\langle \bar{\theta}, x_i \rangle\}_{i=1}^N, \{\langle \bar{\theta}, y_i \rangle\}_{i=1}^N)
```

```
2700
             Algorithm 2 Estimation of MMD<sup>2</sup> loss via imagination of transition samples with Sobolev samples
2701
               1: Require: Number of samples M, kernel k, discount factor \gamma \in (0,1)
2702
              2: Require: Truncation percentage p \in [0, 100]
2703
              3: Require: Distributional critic Z_{\phi}(s, a, \varepsilon)
2704
              4: Require: Policy network \pi_{\theta}(s)
2705
              5: Require: Conditional VAE (cVAE) with prior p_v(\varepsilon) and decoder p_{\psi}(s', r \mid s, a, \varepsilon)
2706
              6: Input: Transition sample (s, a)
2707
              7: Input: Online critic parameter \phi, target critic parameters \phi'_1 and \phi'_2
2708
              8: Input: Target policy parameter \theta'
2709
              9: Input: Boolean flag use_action_gradient
             10: Output: Gradient estimation of MMD with respect to \phi
2710
                                                                                                       {Sample latent variable from the prior}
             11: Draw \varepsilon \sim p_{\upsilon}(\varepsilon)
2711
             12: (\hat{s}', \hat{r}) \sim p_{\psi}(\hat{s}', r \mid s, a, \varepsilon)
13: a' \leftarrow \pi_{\theta'}(\hat{s}')
                                                                                                      {Generate transition using the decoder}
2712
                                                                                                                 {Action from target policy on \hat{s}'}
2713
            13. a \leftarrow h_{\theta}(s)

14. Sample Z_{1:M} \stackrel{i.i.d.}{\sim} Z_{\phi}(s, a)

15. Sample Z_{\text{next},1:M}^{(1)} \stackrel{i.i.d.}{\sim} Z_{\phi'_{1}}(\hat{s}', a'),

16. Sample Z_{\text{next},1:M}^{(2)} \stackrel{i.i.d.}{\sim} Z_{\phi'_{2}}(\hat{s}', a')
2714
                                                                                                                       {Samples from online critic}
2715
2716
                                                                                                               {Samples from each target critic}
2717
             17: Z_{\text{next},1:2M} \leftarrow \text{concat}(Z_{\text{next}}^{(1)}, Z_{\text{next}}^{(2)})
18: if use_action_gradient then
2718
                                                                                                             {Concatenate target-critic samples}
2719
                      for each 1 \le i \le 2M do Y_i \leftarrow \mathbf{f}_{\hat{r},\hat{s}',\gamma}^S(Z_{\text{next},i})
             19:
2720
             20:
                                                                                                              {Bellman target Sobolev samples}
2721
                          \nabla_a Y_i \leftarrow \nabla_a \mathbf{f}_{\hat{r},\hat{s}',\gamma}^S(Z_{\text{next},i})
             21:
                                                                                                            {Action gradient of Bellman target}
2722
                          Y_i \leftarrow \operatorname{concat}(Y_i, \nabla_a Y_i)
2723
             22:
             23:
                       end for
2724
             24:
                       for each 1 \le i \le M do
2725
                           \nabla_a Z_i \leftarrow \nabla_a Z_\phi(s, a)
             25:
                                                                                                                {Action gradient of online critic}
2726
                           Z_i \leftarrow \operatorname{concat}(Z_i, \nabla_a Z_i)
             26:
2727
             27:
2728
             28: else
2729
                       for each 1 < i < 2M do
             29:
2730
                          Y_i \leftarrow \mathbf{f}_{\hat{r},\hat{s}',\gamma}^S(Z_{\text{next},i})
             30:
2731
                       end for
             31:
2732
             32: end if
2733
             33: Prune top p\% of \{Y_i\}_{i=1}^{2M} by return magnitude
                                                                                                                                  {Low-bias target set}
2734
             34: Compute MMD<sup>2</sup>:
             35: \text{MMD}^2 \leftarrow \sum_{i \neq j} \left[ k(Z_i, Z_j) - 2k(Z_i, Y_j) + k(Y_i, Y_j) \right] 36: Return: \text{MMD}^2
2735
2736
2737
```

```
2754
            Algorithm 3 Distributional Sobolev Deterministic Policy Gradient (DSDPG)
2755
              1: Require: Number of samples M, number of policy samples M_{\text{policy}}, kernel k, discount factor
2756
                  \gamma \in (0,1), learning rates \alpha_{\theta}, \alpha_{\phi}
2757
              2: Require: Two distributional critics Z_{\phi_1}(s,a) and Z_{\phi_2}(s,a) with targets Z_{\phi'_1}(s,a), Z_{\phi'_2}(s,a)
2758
              3: Require: Actor network \pi_{\theta}(s) with target \pi_{\theta'}(s)
2759
              4: Require: Replay buffer \mathcal{D}
2760
              5: Require: Boolean flag use_action_gradient
2761
              6: Require: Parameters of cVAE: Encoder (\zeta), Decoder (\psi), Prior (v)
2762
              7: Input: Initial parameters \theta, \phi_1, \phi_2 and target parameters \theta' \leftarrow \theta, \phi'_1 \leftarrow \phi_1, \phi'_2 \leftarrow \phi_2
2763
              8: Input: Policy update frequency d
              9: for each episode do
2764
            10:
                      Initialize a random process \mathcal{N} for action exploration
2765
                      Observe initial state s_0
            11:
2766
            12:
                      for each step in the episode do
2767
            13:
                         a_t \leftarrow \pi_{\theta}(s_t) + \mathcal{N}_t
                                                                                                  {Select action with exploration noise}
2768
            14:
                         Execute a_t, observe (s_{t+1}, r_t)
2769
            15:
                         Store (s_t, a_t, r_t, s_{t+1}) in replay buffer \mathcal{D}
2770
                         \zeta, \psi, \upsilon \leftarrow \text{train\_world\_model}(s_t, a_t, s_{t+1}, r_t)
                                                                                                    {Train cVAE world model and update
            16:
2771
                         parameters}
2772
            17:
                      end for
2773
            18:
                      for each gradient step do
                         Sample mini-batch \{(s_i, a_i)\}_{i=1}^N \sim \mathcal{D}
2774
            19:
                                                                                                                    {Replay buffer sampling}
                         Compute Distributional Loss via MMD (cf. Alg. 2):
2775
            20:
            21:
                         NB. Samples are concatenated and top p\% are removed as in TQC (Kuznetsov et al., 2020)
2776
2777
            22:
                         L_{Z_1} \leftarrow \text{MMD\_Sobolev}(Z_{\phi_1}, (Z_{\phi_1'}, Z_{\phi_2'}), s_i, a_i, \pi_{\theta'}, \gamma, \texttt{use\_action\_gradient})
2778
                         L_{Z_2} \leftarrow \text{MMD\_Sobolev}(Z_{\phi_2}, (Z_{\phi_1'}, Z_{\phi_2'}), s_i, a_i, \pi_{\theta'}, \gamma, \text{use\_action\_gradient})
            23:
2779
                         Update critics:
            24:
2780
                         \phi_1 \leftarrow \phi_1 - \alpha_\phi \nabla_{\phi_1} L_{Z_1}
            25:
2781
                         \phi_2 \leftarrow \phi_2 - \alpha_\phi \nabla_{\phi_2} L_{Z_2}
            26:
2782
                         if gradient step mod d = 0 then
            27:
2783
                             Sample M_{\text{policy}} values from critic for each state in the batch:
            28:
                             Z_{\text{policy},1:M_{\text{policy}}}^{(i)} \overset{i.i.d.}{\sim} Z_{\phi_1}(s_i,\pi_{\theta}(s_i)) \quad \text{ for each } i \in \{1,\dots,N\}
2784
            29:
2785
                             Compute actor loss using double summation:
            30:
2786
                             \begin{array}{l} L_{\pi} = -\frac{1}{N} \sum_{i=1}^{N} \frac{1}{M_{\text{policy}}} \sum_{j=1}^{M_{\text{policy}}} Z_{\text{policy},j}^{(i)} \\ \text{Update actor: } \theta \leftarrow \theta - \alpha_{\theta} \nabla_{\theta} L_{\pi} \end{array}
            31:
2787
2788
            32:
            33:
                             Update target networks:
2789
                             \phi_1' \leftarrow \tau \phi_1 + (1 - \tau)\phi_1'
            34:
                             \phi_2' \leftarrow \tau \phi_2 + (1 - \tau) \phi_2'
\theta' \leftarrow \tau \theta + (1 - \tau) \theta'
            35:
2791
            36:
2792
            37:
                         end if
2793
            38:
                      end for
2794
            39: end for
2795
```

## J.3 BASELINE - MAGE

2796279727982799

2805

2806

2807

The primary baseline of this work is MAGE D'Oro & Jaskowski (2020), as they were the first to propose to use the action gradient to steer policy evaluation. The primary difference with our work is they did so deterministically. Hence our MMD setup collapses to a simple regression setting.

As a primary baseline, we consider the deterministic counterpart of the Sobolev-distributional backup. Define

$$\begin{split} f_{\text{det}}(s, a) &= \begin{bmatrix} f_{\text{det}}^{\text{ret}}(s, a) \\ f_{\text{det}}^{\text{act}}(s, a) \end{bmatrix}, \\ f_{\text{det}}^{\text{ret}}(s, a) &= r(s, a) + \gamma \, Q_{\phi}(s', \pi(s')), \\ f_{\text{det}}^{\text{act}}(s, a) &= \nabla_{a} r(s, a) + \gamma \, \nabla_{a} Q_{\phi}(s', \pi(s')). \\ \mathbf{L}^{S_{a}}(\phi; s, a) &= \left| f_{\text{det}}^{\text{ret}}(s, a) - Q_{\phi}(s, a) \right|^{2} + \lambda_{S} \left\| f_{\text{det}}^{\text{act}}(s, a) - \nabla_{a} Q_{\phi}(s, a) \right\|^{2}. \end{split}$$

In practice, s' and r are sampled from the stochastic world model (cVAE) as in DSDPG. Furthermore, the above L2 terms are replaced by Huber losses, with the gradient-term weight set to 5 just as in D'Oro & Jaskowski (2020).

### K TOY SUPERVISED LEARNING

To motivate our algorithm, we demonstrate its ability to learn the joint distribution over both the output and gradient of a random function in a supervised setup. We compare deterministic Sobolev training Czarnecki et al. (2017) against our Distributional Sobolev training and show that only the latter can capture the full variability of both outputs and gradients.

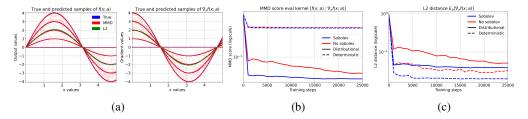


Figure 4: (a) Samples of the marginals of the full Sobolev distribution  $[f(x;a); \nabla_x f(x;a)]$ : output (left) and gradient (right). Blue: samples from the true distribution; red: samples from the Distributional Sobolev model trained via MMD; green: samples from the deterministic Sobolev baseline Czarnecki et al. (2017). (b) Biased MMD score (lower is better) on the joint variable. (c)  $L_2$  error between predicted sample mean and true mean. NB. Predicted (red and green) and true (blue) samples are highly overlapping.

The task involves learning a one-dimensional conditional distribution p(y|x), defined as a mixture of sinusoids  $f(x;a) = a \times \sin(x)$ , where the latent variable a is uniformly drawn from  $\{0,1,2,3,4\}$ . The distributions over outputs f(x;a) and their gradient  $\nabla_x f(x;a)$  are depicted in Figure 4a. It compares an MMD-based model and a regression-based model, trained using stochastic gradient descent with identical architectures. In an unlimited data regime, new pairs of x and a were drawn for each batch, and for each x, four a values were sampled with replacement, yielding samples  $(x,y_{1:4})$ .

As expected, the MMD-based model captures the full joint distribution  $[f(x;a); \nabla_x f(x;a)]$ , whereas the regression baseline collapses to the conditional mean  $\mathbb{E}_a[f(x;a); \nabla_x f(x;a)]$ . Figure 4b plots the MMD score (lower is better) on this joint variable using an evaluation kernel, and Figure 4c shows the  $L_2$  error between the regression model and the sample mean of the Sobolev generator. As a result, the MMD model better matches the entire distribution, the regression model slightly outperforms on the mean, and both methods effectively exploit the gradient signal (blue curves).

The distributional variant via MMD and the regression one via L2 used the same architecture except for some noise of dimension 10 drawn from  $\mathcal{N}(0;I)$  concatenated to the input for the distributional generator. For each pair  $(x,y_{1:4})$ , four samples were drawn from the generator. Both were trained using Rectified Adam Liu et al. (2019) optimizer with a learning rate of  $1 \times 10^{-3}$  and  $(\beta_0,\beta_1) = (0.5,0.9)$ . Neural network is a simple MLP with 2 hidden layers of 256 neurons and Swish nonlinearities (Ramachandran et al., 2017).

Maximum Mean Discrepancy (MMD) was estimated using a mixture of RBF kernels with bandwidths  $\sigma_i$  from the set  $\{\sigma_1, \sigma_2, \dots, \sigma_7\} = \{0.01, 0.05, 0.1, 0.5, 1, 10, 100\}$ . We used the biased estimator from Eq. 25.

The equation for a mixture of RBF kernels is given by:

$$k^{\text{mix}}(x,y) = \sum_{i=1}^{7} \exp\left(-\frac{\|x-y\|^2}{2\sigma_i^2}\right). \tag{48}$$

The evaluation kernel we used was the Rational Quadratic  $k_{\alpha}^{RQ}$  with  $\alpha = 1$  with

$$k_{\alpha}^{\text{RQ}}(x,y) = \left(1 + \frac{\|x - y\|^2}{2\alpha}\right)^{-\alpha} \tag{49}$$

Regarding the dataset, the  $(x,y_{1:4})$  pairs were drawn with  $x \sim \mathcal{U}[0;5]$  and a was draw from  $\{0,1,2,3,4\}$  with replacement. In the limited data regime, the pairs  $(x,y_{1:4})$  were drawn once and stayed fix. The batch size was thus equal to the number of points in the dataset. In the unlimited data regime 256 new pairs were drawn for each batch. Every experiment was ran for 25 000 batch sampled and thus the same number of SGD steps.

Limited data regime In both supervised and reinforcement learning tasks, the assumption of unlimited data is unrealistic. Here, we explore how the performance of the two methods, MMD-based and the regression-based model, diverges when the amount of available data is restricted. We use the same setup as before, but with a fixed number of  $(x, y_{1:4})$  pairs. Several aspects of the learnt models can be inspected. In order to assess stability, we report the average norm of the second-order derivative over the input space. For accuracy, we measure the average L2 losses between the true expected gradient and the predicted gradient. Results are shown in Figure 5. As can be seen, the deterministic model tends to overfit rapidly, while the distributional variant (MMD) proves more robust, maintaining better performance even with constrained data. Notably, the second-order derivative for the deterministic model escalates sharply as data becomes constrained, indicating instability in its approximation. In Appendix K.1, we discuss different common tricks to mitigate overfitting and related issues in RL.

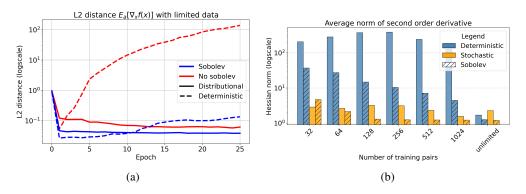


Figure 5: Toy supervised learning problem. Comparison between MMD-based or L2 based modeling. Left panel: training curve of L2 loss (logscale) on gradient between true conditional expectation with regression prediction and with empirical mean from MMD-based model. Sobolev (blue) used gradient information to train either using MMD (full line) or L2 regression (dashed line). Right panel: average over the input space of the second order derivative (logscale) of predicted gradient from deterministic model (blue), MMD / stochastic (yellow) and with gradient information / Sobolev (dashed). Metrics averaged over 5 seeds.

## K.1 COMMON TRICKS

Overfitting with limited data is a common issue in regression tasks. Early stopping seems an obvious solution in this case but we emphasize that it requires an evaluation criterion that is not always

available (i.e in policy evaluation). Other solutions include weight regularization Krogh & Hertz (1991), dropout Srivastava et al. (2014), Bayesian neural networks Blundell et al. (2015), ensembling Chua et al. (2018), and spectral normalization Zheng et al. (2023), all of which often reduce network capacity. To address similar issues, Fujimoto et al. (2018) proposed adding noise to the target to match, effectively smoothing the critic. As argued by Ball & Roberts (2021), this method can be seen as indirectly acting like spectral normalization, encouraging smoother gradients and effectively reducing the magnitude of the second derivative. Appendix K.2 shows how noise scale impacts overfitting by inducing bias. On the other hand, we propose avoiding such assumptions by using generative modeling to add latent freedom.

### K.2 Adding noise

Inspired by Fujimoto et al. (2018), we added some independent noise on x for each  $(x, y_{1:4})$ . Noise scale  $\sigma$  was in  $\{0.01, 0.1, 0.5\}$ . For each new batch sampled it was sampled from a standard Gaussian  $\eta \sim \mathcal{N}(0; \sigma^2)$  and added as  $\tilde{x} = x + \eta$ 

In Figure 6-7, we can see the impact of the various noise scales on the predictions of the deterministic regression. As can be seen, adding noise on x as a positive effect in terms of stabilizing the gradient but it induces an bias that grows with the scale of the noise. Moreover, this noise depend on the application and makes strong assumption about the function to learn. The stabilizing effect of additive noise can further be seen in Figure 8-9 where both the L2 loss and average second order derivative are displayed as function on the number of sampled locations.

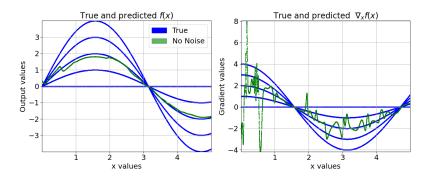


Figure 6: Toy supervised learning problem. Comparison of samples from the true five-mode distribution with predictions made by a deterministic model trained with L2 loss (green). The output space is shown on the left, and the gradient space on the right. Results obtained after 25,000 training steps.

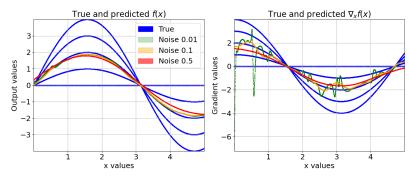


Figure 7: Toy supervised learning problem. Comparison between true samples from the distribution of five modes and deterministic models trained with varying levels of additive noise on their input data. Low level of noise (green), medium level of noise (orange), high level of noise (red). Results obtained after 25,000 training steps.

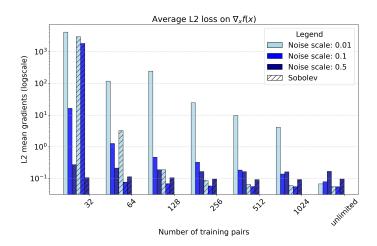


Figure 8: Toy supervised learning problem. Comparison of the L2 loss between the predicted gradient and the conditional expectation of the true distribution. Different scales of additive noise on the input are compared: low noise (light blue), medium noise (medium blue), and high noise (dark blue), alongside Sobolev training (dashed). Results are shown after 25,000 training steps.

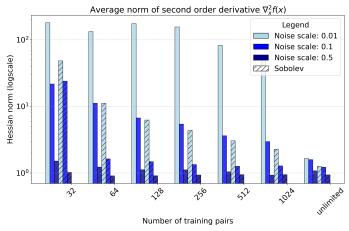


Figure 9: Toy supervised learning problem. Comparison of the average second order derivative norm over the input space. Different scales of additive noise on the input are compared: low noise (light blue), medium noise (medium blue), and high noise (dark blue), alongside Sobolev training (dashed).

## L TOY REINFORCEMENT LEARNING

### L.1 ENVIRONMENT

In this appendix we provide full details of our custom toy environment, summarized in Table 1.

**Environment Dynamics.** A point mass moves in a two-dimensional continuous state space, and the agent controls its acceleration in [-1,1] per axis (Acceleration range). At each step the mass's velocity is updated via Euler integration with friction (Friction coefficient) and then its position is advanced by the new velocity (Integration time step). The mass has unit mass (Mass of agent) and radius 0.5 (Agent radius), and if its center leaves the square of half-width 3 (Bounding box half-width) the episode ends with no reward (Reward for leaving bounding box).

**Partial Observability via Memory.** At the start of each episode the mass is placed uniformly in a square of half-width 1 (Initialization area half-width). We then sample one of  $M \in \{3,4,5,6\}$  hidden bonus locations arranged radially around the center. Reaching the correct location of radius 0.5 (Bonus radius) yields a terminal reward (Reward for reaching bonus). The agent does not know which location is active, but each visit to a location sets a binary memory flag, making the MDP partially observable and forcing exploration.

Controlling Distributional Modes via M. By sweeping  $M \in \{3,4,5,6\}$  we directly tune the number of modes in the return distribution across episodes, from concentrated when M is small to highly multimodal when M is large.

Parameter	Value
Maximum episode length	25 steps
Reward for reaching bonus	10
Reward for leaving bounding box	0
Acceleration range	[-1,1] per axis
Mass of agent	1
Agent radius	0.5
Bonus radius	0.5
Bounding box half-width	3
Initialization area half-width	1
Friction coefficient	0.1
Integration time step	0.5

Table 1: Summary of the toy environment's key parameters.

#### L.2 EXPERIMENTAL DETAILS

The settings were similar to Section M appart from the TQC Kuznetsov et al. (2020) truncation parameter p which was set to 5%.

### M REINFORCEMENT LEARNING EXPERIMENTS

Here we describe the architectures, optimizers and other hyperparameters of the full Distributional Sobolev Deterministic Policy Gradient algorithm.

**Policy network** Policy network is a MLP with 2 hidden layers of 400 neurons. The non-linearity was Swish (Ramachandran et al., 2017). Final activations are mapped to the output space using a linear transformation followed by a  $\tanh$  non-linearity. The policy network was optimized using the Rectified Adam Liu et al. (2019) with a learning rate of  $1 \times 10^{-4}$ .

**Crtic network** Critic network architecture is almost the same as for the policy network. Importantly, it is kept constant for experiments using normal DDPG and DSDPG apart from noise concatenated on the input [s; a]. The network is a MLP with 400 neurons, skip connections from the input and Swish activations. No non-linearity was applied on the output after the last linear layer.

Table 2: Hyperparameters for the DDPG and DSDPG experiments on Mujoco environments

Item	Value
Discount $\gamma$	0.99
Polyak averaging $ au$	0.005
Buffer size	$10^{6}$
Batch size	256
Exploration noise scale	0.1
Critic learning rate	$1 \times 10^{-4}$
Policy learning rate	$1 \times 10^{-4}$
cVAE learning rate	$3 \times 10^{-4}$
cVAE KL weight	0.1
cVAE latent dim	$ \mathcal{S} +1$
Critic input noise dim	64
Number of samples dist.	10
%p truncation (TQC)	25%
maximum slicing optimization steps	100
maximum slicing LR	1e-4
maximum slicing optimizer	Adam (Kingma, 2014)

The critic network is optimized using Rectified Adam Liu et al. (2019) with  $(\beta_1, \beta_2) = (0.9, 0.999)$  and a learning rate of  $1 \times 10^{-4}$ .

**Conditional VAE world model** The encoder, decoder and prior networks are MLPs with 3 hidden layers, each containing 1024 neurons. Skip connections are applied from the input to each hidden layer. The cVAE was optimized using Adam, as we observed using RAdam to systematically diverge,  $(\beta_1, \beta_2) = (0.9, 0.999)$  and a learning rate of  $3 \times 10^{-4}$ .

The prior is a learnt diagonal multivariate Gaussian  $\mathcal{N}(\varepsilon; \mu_v(s, a), \sigma_v^2(s, a) \odot I)$  with a latent dimension equal to the size of the random variable being modeled, which is  $|\mathcal{S}| + 1$  for (s', r). Following D'Oro & Jaskowski (2020); Zhu et al. (2024), the cVAE predicts the difference between the current and next observation,  $\delta_s = s' - s$  which is then added back to s, along with the reward r.

Conditional Generative Moment Matching For the distributional critic, noise vectors were concatenated with the state-action pairs (s,a) and passed through the same architecture as the deterministic critic. The noise dimension was set to 64. Noise was transformed by an independant 2 layers small MLP of width 64 with Swish activation before being passed to the critic. For each state-action pair, 10 samples were drawn to update both the critic and the policy. The multiquadratic kernel  $k_h^{\rm MQ}(x,y) = -\sqrt{1+h^2\|x-y\|_2^2}$  Killingberg & Langseth (2023) was used, with the kernel parameter h set to 100.

### M.1 FULL CURVES AND WALL-CLOCK TIME

The full evaluation curves on the 6 environments are reported in Figure 10. The wall-clock time of the different methods is displayed in Table 3.

### N ADDITIONAL EXPERIMENTS - RL

In this section we provide futher experiments along three axis: kernel bandwidth (Section N.1), noise scale (Section N.2), number of samples in the distributional methods (Section N.3) and capacity of the world model (Section N.4).

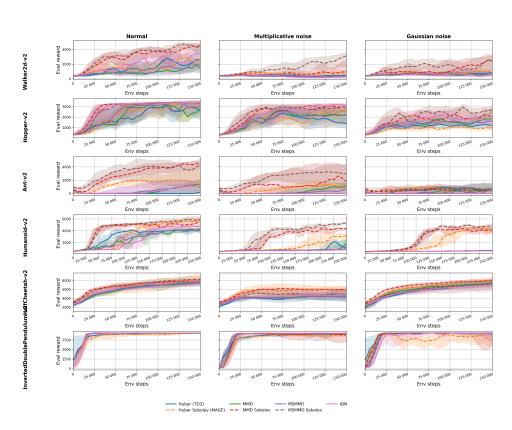


Figure 10: Evaluation of DSDPG (MMD Sobolev), deterministic Sobolev/MAGE D'Oro & Jaskowski (2020), TD3-Huber Fujimoto et al. (2018), IQN Dabney et al. (2018a), and standard MMD Nguyen et al. (2020); Killingberg & Langseth (2023) on six MuJoCo tasks. Results are reported over 10 random seeds. The median is displayed with 25%-75% IQR. Window smoothing with window size 3.

Method	Time for 1000 iterations (s)
MSMMD Sobolev	62.5
MSMMD	45.5
MMD Sobolev	40.0
MMD	35.7
IQN	35.7
Huber Sobolev	31.3
Huber	30.3

Table 3: Wall-clock time to perform 1000 iterations of each method under Humanoid-v2. Experiments were run on a single Nvidia H100 GPU.

### N.1 MULTIQUADRIC KERNEL BANDWIDTH

We recall the multiquadric kernel

$$k_h^{\text{MQ}}(x,y) = -\sqrt{1 + h^2 \|x - y\|_2^2}.$$

As shown in Killingberg & Langseth (2023), selecting a proper value for h is critical as it affects both the expressiveness of the kernel and the numerical stability of its MMD estimator: a too-small h yields a nearly constant kernel, reducing its discrimination power, while a too-large h increases the gradient magnitude, scaling on the order of h, which can lead to exploding gradients and training oscillations unless mitigated by techniques such as kernel rescaling or gradient clipping.

In Section 6.2, the default value for h was 100 which worked reliably. Here we test the values h=10 and h=250. The results are depicted in Figure 12 and 11 where we display the final evaluation performance and the evaluation curves. The comparison is made on the environments with and without multiplicative noise (as in Section 6.2). As can be seen, overall, the performance of MMD and DSDPG (MMD Sobolev) seems to be quite insensitive to the kernel bandwidth. Our choice of h=100 seems to be robust.

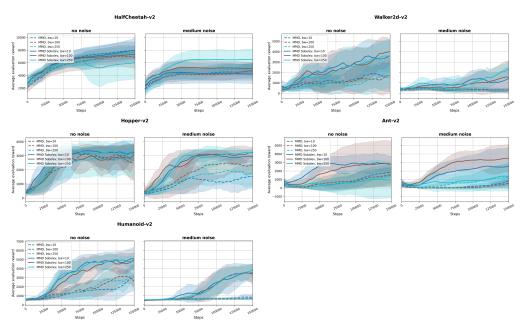


Figure 11: Full curves of DSDPG against baselines on 5 Mujoco environments.

### N.2 Noise scale

As the noise scale we exposed in the Section 6.2, which we called *medium*, already allowed to observe a gap between DSDPG and the baseline, we tried to increase the noise scale. We moved

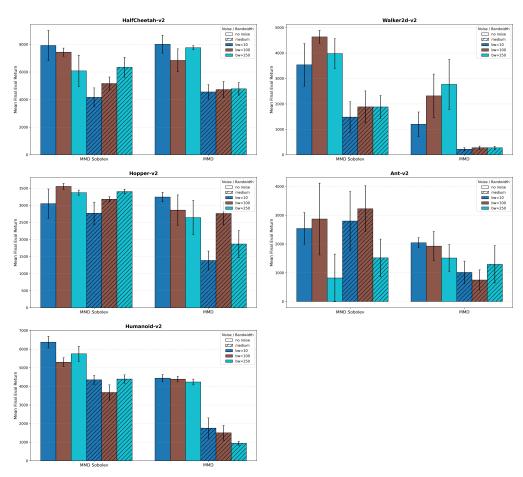


Figure 12: Comparison of bandwidths under the multiquadric kernel for MMD distributional RL and Sobolev MMD distributional RL. Bar plots of the final averaged evaluation reward. Mean over 5 seeds.

from  $n \sim \mathcal{U}[0.8, 1.2]$  to  $n \sim \mathcal{U}[0.7, 1.3]$ . This changes already made the tasks harder. The results are displayed in Figure 13. As can be seen the good performance of DSDPG (MMD Sobolev) depicted in Section 6.2 can be extended to the larger noise scale, as our method maintains a consistent gap against the baselines (especially MAGE).

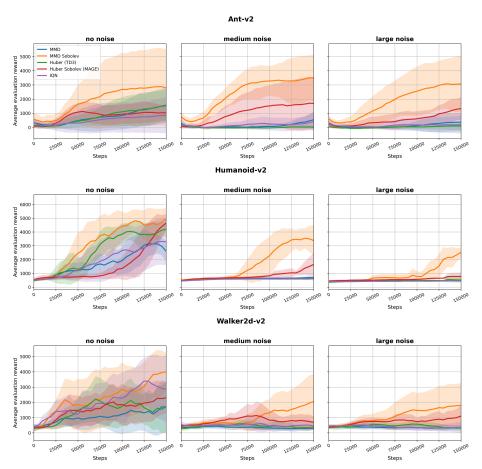


Figure 13: Comparison of multiplicative noise scales. Averaged evaluation sum of rewards over 5 seeds.

### N.3 NUMBER OF SAMPLES

Sensitivity to the number of samples is an important design question as the cost of generating samples will scale linearly and the cost of estimating MMD will scale quadratically. It is also intersting to verify how sensitive our method is to that parameter. The results displayed in Section 6.2 used 10 samples to model the Sobolev distributions. Here we additionally test with 5 and 25 samples. In Figure 15 and 14 we display the final evaluation returns and evaluation curves while varying the number of samples. As can be seen, there is no clear trend to be found. The Ant-v2 environment seems to be particularly sensitive to this parameter. Apart from this example we observe our method is robust to the number of samples as at least two different values maintained performance.

## N.4 WORLD MODEL CAPACITY

The world model size is, along the number of samples, one of the primary parameter driving the overall cost of our method. As for MAGE D'Oro & Jaskowski (2020), we backpropagate through the world model to infer the gradients of the target. Here we change the width of each neural network in the cVAE world model from 1024 to 256. We observe in Figure 16, on the single Walker2d-v2 environment we tested, that more than insensitive to model size, it seems the choice of 1024 was sub-

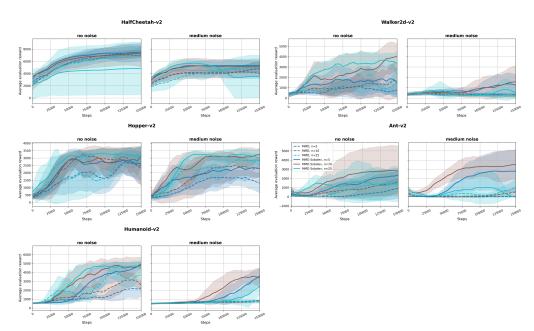


Figure 14: Comparison of number of samples used in the MMD methods (usual MMD an DSDPG denoted as MMD Sobolev). Evaluation curves are mean over 5 seeds.

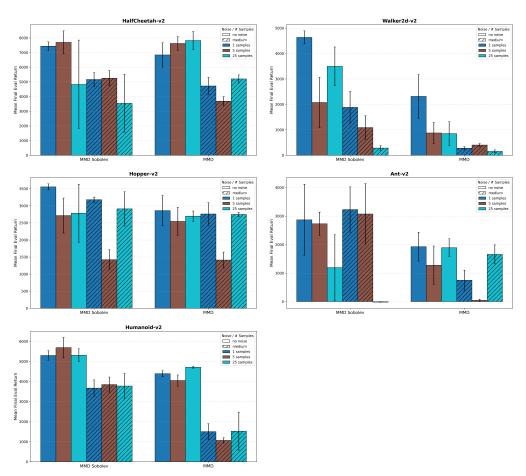


Figure 15: Comparison of number of samples used in the MMD methods (usual MMD an DSDPG denoted as MMD Sobolev). Final evaluation performance as mean over 5 seeds.

optimal for this environment. This suggests our method is robust to this design choice. This should be evaluated on more demanding (higher dimensional) environments (Ant-v2, Humanoid-v2).

### Walker2d-v2 - medium noise

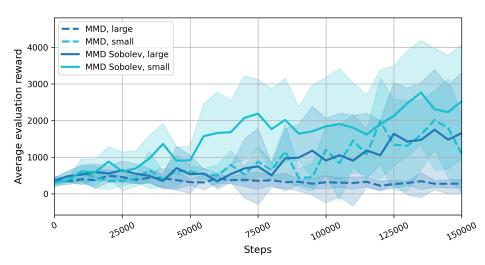


Figure 16: Comparison of world model size for MMD distributional RL and Sobolev MMD distributional RL. Bar plots of the final averaged evaluation reward. Mean over 5 seeds.

### O LLM USAGE

We used an LLM-based assistant to support the preparation of this paper. In particular, it was employed to (i) rephrase draft paragraphs for clarity and suggest alternative framings of related work, (ii) format proofs, explore directions, and verify intermediate steps, (iii) assist in debugging code, (iv) suggest LaTeX equation formatting, and (v) help identify relevant theoretical results in preceding works. All core research contributions, including the development of theoretical results, algorithms, and experiments, were carried out by the authors.