
Provable Guarantees for Generative Behavior Cloning: Bridging Low-Level Stability and High-Level Behavior

Adam Block
MIT
Cambridge, MA
ablock@mit.edu

Ali Jadbabaie
MIT
Cambridge, MA
jadbabai@mit.edu

Daniel Pfrommer
MIT
Cambridge, MA
dpfrom@mit.edu

Max Simchowitz
MIT
Cambridge, MA
msimchow@csail.mit.edu

Russ Tedrake
MIT
Cambridge, MA
russt@mit.edu

Abstract

We propose a theoretical framework for studying behavior cloning of complex expert demonstrations using generative modeling. Our framework invokes low-level controllers - either learned or implicit in position-command control - to stabilize imitation around expert demonstrations. We show that with (a) a suitable low-level stability guarantee and (b) a powerful enough generative model as our imitation learner, pure supervised behavior cloning can generate trajectories matching the per-time step distribution of essentially arbitrary expert trajectories in an optimal transport cost. Our analysis relies on a stochastic continuity property of the learned policy we call “total variation continuity” (TVC). We then show that TVC can be ensured with minimal degradation of accuracy by combining a popular data-augmentation regimen with a novel algorithmic trick: adding augmentation noise at execution time. We instantiate our guarantees for policies parameterized by diffusion models and prove that if the learner accurately estimates the score of the (noise-augmented) expert policy, then the distribution of imitator trajectories is close to the demonstrator distribution in a natural optimal transport distance. Our analysis constructs intricate couplings between noise-augmented trajectories, a technique that may be of independent interest. We conclude by empirically validating our algorithmic recommendations, and discussing implications for future research directions for better behavior cloning with generative modeling.

1 Introduction

Training dynamic agents from datasets of expert examples, known as *imitation learning*, promises to take advantage of the plentiful demonstrations available in the modern data environment, in an analogous manner to the recent successes of language models conducting unsupervised learning on enormous corpora of text [67, 70]. Imitation learning is especially exciting in robotics, where mass stores of pre-recorded demonstrations on Youtube [1] or cheaply collected simulated trajectories [42, 20] can be converted into learned robotic policies.

For imitation learning to be a viable path toward generalist robotic behavior, it needs to be able to both represent and *execute* the complex behaviors exhibited in the demonstrated data. An approach that has shown tremendous promise is *generative behavior cloning*: fitting generative models, such

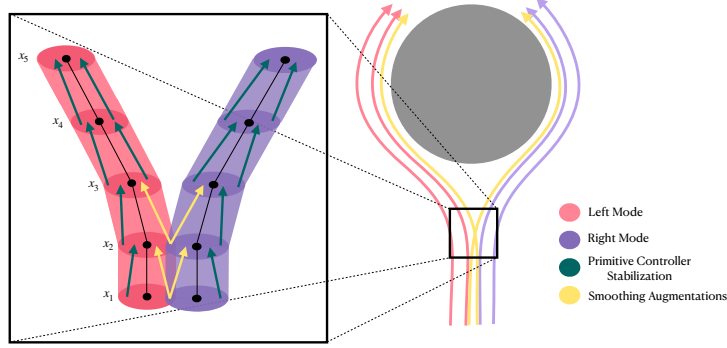


Figure 1: Consider demonstration trajectories exhibiting two modes: a “go left” and “go right” mode around an obstacle depicted in red and purple, respectively. To avoid compounding error, we imitate sequences of simple low-level feedback controllers we call “primitive controllers”, not simply raw actions. Intuitively, primitive controllers provide “tubes” around each demonstration trajectory where the system can be stabilized. Depicted in yellow, our data-noising procedure described below “fills in the gaps” in the demonstration, switching between modes in a well-behaved manner, and whilst allowing the primitive controllers to manage the stabilization.

as diffusion models [2, 19, 33], to expert demonstrations with pure supervised learning. In this paper, we ask: **When can generative behavior cloning imitate arbitrarily complex expert behavior?**

In this paper, we are interested in how *algorithmic choices* interface with the *dynamics of the agent’s environment* to render imitation possible. The key challenge separating imitation learning from vanilla supervised learning is one of *compounding error*: when the learner executes the trained behavior in its environment, small mistakes can accumulate into larger ones; this in turn may bring the agent to regions of state space not seen during training, leading to larger-still deviations from intended trajectories. Without the strong requirement that the learner can interactively query the expert at new states [40, 57], it is well understood that ensuring some form of *stability* in the imitation learning procedure is indispensable [68, 27, 50]. While many natural notions of stability exist for simple behaviors, how to *enforce* stability when imitating more complex behaviors remains an open question. Multi-modal trajectories present a key example of this challenge: consider a robot navigating around an obstacle; because there is no difference between navigating around the object to the right and around to the left, the dataset of expert trajectories may include examples of both options. This bifurcation of good trajectories can make it difficult for the agent to effectively choose which direction to go, possibly even causing the robot to oscillate between directions and run into the object. [19]. Moreover, human demonstrators correlate current actions with the past in order to *commit* to either a right or left path, which makes even formulating the idea of an “expert policy” a conceptually challenging one. Lastly, bifurcations are necessarily *incompatible* with previous notions of stability derived from classical control theory [68, 27, 50]. **In this work, we investigate how these strong and often unrealistic assumptions on the expert policy can be replaced by practical (and often realistic) assumptions on available algorithms.**

1.1 Contributions.

As in previous work, we formalize behavior cloning in two stages: at *train-time*, we learn a map from observations to distributions over actions, supervised by (state, action)-pairs from expert demonstrations coming from N_{exp} independent expert trajectories, while at *test-time*, the learned map, or *policy*, is executed on random initial states (distributed identically to initial training states). Following the schematic of existing theoretical analyses of behavior cloning [68, 50, 27], we demonstrate that a policy trained by minimizing a certain supervised learning objective on expert demonstrations induces trajectories that approximate those of expert demonstrations. Our work considers a significantly more general setting than past theoretical literature, and one which reflects the strength of *generative models* for imitation. One corollary of our key contributions is summarized in the following informal statement. The main technical insights leading to the proof of the theorem are detailed in the bullet points below it, and depicted in Figure 1.

Theorem (informal). Consider a generative behavior cloner $\hat{\pi}$ that learns to predict sequences of expert actions on horizon H , **along with low-level controllers that locally stabilize the trajectories.** Then, with a suitable data noising strategy, for all times $h \leq H$,

$$\begin{aligned} & \mathbb{P}[\text{expert \& imitator trajectories disagree at some time } h \text{ by } \geq \varepsilon] \\ & \leq \mathcal{O}_{\text{Iss}} \left(H\varepsilon + \frac{1}{\varepsilon^2} \sum_h \mathbb{E}_{\text{expert},h} [\mathcal{W}_1(\mathbb{P}_{\text{expert actions}}, \mathbb{P}_{\text{imitator actions}})] \right) \end{aligned}$$

where $\mathbb{E}_{\text{expert},h}[\mathcal{W}_1(\mathbb{P}_{\text{expert actions}}, \mathbb{P}_{\text{imitator actions}})]$ denotes a 1-Wasserstein distance in an appropriate metric between the conditional distribution over expert and imitator actions given the observation at time step h , and where \mathcal{O}_{Iss} hides constants depending polynomially on the stability properties of the low-level controllers, defined formally in [Section 3.1](#).

We now detail the key ingredients of our results.

1. We imitate stochastic demonstrators that may exhibit both complex correlations between actions in their trajectories (e.g. be non-Markovian) and multi-modal behavior. The natural object to imitate in this setting is the conditional probability distribution of expert actions given recent states, but marginalized over past states. We require said conditional action distribution to be learnable by a **generative model**, but otherwise arbitrarily complex: in particular, the conditional distribution of an expert actions given the state can be discontinuous (in any natural distance metric) as a function of state, as in the bifurcation depicted in [Figure 1\(right\)](#).
2. We obtain **rigorous, theoretical guarantees** and without requiring either interactive data collection (e.g. DAGGER [57, 40]), or access to gradients of the expert policy (as in TASIL[50]). Instead, we replace these assumption with an oracle, described below, which **synthesizes stabilizing, low-level policies** along training demonstrations—the green arrows in [Figure 1\(left\)](#). This mirrors recent work on generative behavior cloning that find that providing state-commands through inverse dynamics controllers [33, 2] or position-command controllers of end effectors [19] leads to substantially improved performance.
3. We also apply a subtle-yet-significant modification to a popular **data noising** strategy, which we show yields both theoretical and empirical benefits. Data noising ensures a helpful property we denote *total variation continuity* that interpolates between modes in probability space (without naively averaging their trajectories in world space). This effectively “fills in the missing gaps” in bifurcations, as indicated by yellow arrows in [Figure 1](#).

Our main results, [Theorems 1 and 2](#), are reductions from imitation of complex expert trajectories to supervised generative learning of a specific conditional distribution. For concreteness, [Theorem 3](#) instantiates the generative modeling with Denoising Diffusion Probabilistic Models (DDPMs) of sufficient regularity and expressivity (as investigated empirically in [19, 48, 26]), and establishes end-to-end guarantees for imitation of complex trajectories with sample complexity polynomial in relevant problem parameters. Our analysis framework exposes that any sufficient powerful generative learner obtains similar guarantees. Finally, we empirically validate the benefits of our proposed smoothing strategy in simulated robotic manipulation tasks. We now summarize the algorithmic choices and analytic ideas that facilitate our reduction.

Abridged Related Work. Due to space, we defer a full comparison to past work to [Appendix B](#). DDPMs, proposed in [29, 60], along with their relatives have seen success in image generation [62, 55], along with imitation learning (without data augmentation) [33, 19, 48], which is the starting point of our work. Smoothing data augmentation is ubiquitous in modern imitation learning [40] and our approach corresponds to that of [36] but with noise added at inference time. Despite the benefits of adaptive data collection [58, 40], adaptive demonstrations are more expensive to collect. Previous analyses of imitation learning without adaptive data collection have focused on classical control-theoretic notions of stability, notably incremental stability, [68, 27, 50], which require continuity, Markovianity, and often determinism, and preclude the bifurcations permitted in our setting.

2 Setting

Notation and Preliminaries. [Appendix A](#) gives a full review of notation. Bold lower-case (resp. upper-case) denote vectors (resp. matrices). We abbreviate the concatenation of sequences via $\mathbf{z}_{1:n} = (\mathbf{z}_1, \dots, \mathbf{z}_n)$. Norms $\|\cdot\|$ are Euclidean for vectors and operator norms for matrices unless

otherwise noted. Rigorous probability-theoretic preliminaries are provided in [Appendix F](#). In short, all random variables take values in Polish spaces \mathcal{X} (which include real vector spaces), the space of Borel distributions on \mathcal{X} is denoted $\Delta(\mathcal{X})$. We rely heavily on *couplings* from optimal transport theory: given measures $X \sim P$ and $X' \sim P'$ on \mathcal{X} and \mathcal{X}' respectively, $\mathcal{C}(P, P')$ denotes the space of joint distributions $\mu \in \Delta(\mathcal{X} \times \mathcal{X}')$ called “couplings” such that $(X, X') \sim \mu$ has marginals $X \sim P$ and $X' \sim P'$. $\Delta(\mathcal{X} | \mathcal{Y})$ denotes the space of conditional probability distributions $Q : \mathcal{Y} \rightarrow \Delta(\mathcal{X})$, formally called probability *kernels*; [Appendix F](#) rigorously justifies that, in our setting, all conditional distributions can be expressed as kernels (which we do throughout the paper without comment). Finally $\mathbf{I}_\infty(\mathcal{E})$ denotes the indicator taking value 1 if \mathcal{E} is true and ∞ otherwise.

Dynamics and Demonstrations. We consider a discrete-time, control system with states $\mathbf{x}_t \in \mathcal{X} := \mathbb{R}^{d_x}$, and inputs $\mathbf{u}_t \in \mathcal{U} := \mathbb{R}^{d_u}$, obeying the following nonlinear dynamics

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t), \quad t \geq 1. \quad (2.1)$$

Given length $T \in \mathbb{N}$, we call sequences $\boldsymbol{\rho}_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \in \mathcal{P}_T := \mathcal{X}^{T+1} \times \mathcal{U}^T$ *trajectories*. For simplicity, we assume that (2.1) is deterministic and address stochastic dynamics in [Appendix N](#). Though the dynamics are Markov and deterministic, we consider a stochastic and possibly *non-Markovian* demonstrator, which allows for the multi-modality described in the [Section 1](#).

Definition 2.1 (Expert Distribution). Let $\mathcal{D}_{\text{exp}} \in \Delta(\mathcal{P}_T)$ denote an *expert distribution* over trajectories to be imitated. $\mathcal{D}_{\mathbf{x}_1}$ denotes the distribution of \mathbf{x}_1 under $\boldsymbol{\rho}_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \sim \mathcal{D}_{\text{exp}}$.

Primitive Controllers. Our approach is to imitate not just actions, but simple local *control policies*. In the body of this paper, we consider affine mappings $\mathcal{X} \rightarrow \mathcal{U}$ (redundantly) parameterized as $\mathbf{x} \mapsto \bar{\mathbf{u}} + \bar{\mathbf{K}}(\mathbf{x} - \bar{\mathbf{x}})$; we call these *primitive controllers*, denoted with $\kappa = (\bar{\mathbf{u}}, \bar{\mathbf{x}}, \bar{\mathbf{K}}) \in \mathcal{K}$. We describe the synthesis of these controllers in [Appendix D](#), and extend our results to general families of parameterized controllers in [Appendix E](#). We argue in [Appendix E](#) that primitive controllers are in fact standard practice, and implicit via robotic position control in many applications of diffusion to robotic behavior cloning.

Chunking Policies and Indices. The expert distribution \mathcal{D}_{exp} may involve non-Markovian sequences of actions. We imitate these sequences via *chunking policies*. Fix a *chunk length* $\tau_{\text{chunk}} \in \mathbb{N}$ and *observation length* $\tau_{\text{obs}} \leq \tau_{\text{chunk}}$, and define time indices $t_h = (h-1)\tau_{\text{chunk}} + 1$. For simplicity, we assume τ_{chunk} divides T , and set $H = T/\tau_{\text{chunk}}$. Given a $\boldsymbol{\rho}_T \in \mathcal{P}_T$, define the *trajectory-chunks* and *observation chunks*

$$\begin{aligned} \mathbf{s}_h &:= (\mathbf{x}_{t_{h-1}:t_h}, \mathbf{u}_{t_{h-1}:t_h-1}) \in \mathcal{S} := \mathcal{P}_{\tau_{\text{chunk}}} && \text{(trajectory-chunks)} \\ \mathbf{o}_h &:= (\mathbf{x}_{t_h-\tau_{\text{obs}}+1:t_h}, \mathbf{u}_{t_h-\tau_{\text{obs}}+1:t_h-1}) \in \mathcal{O} := \mathcal{P}_{\tau_{\text{obs}}-1} && \text{(observation-chunks)} \end{aligned}$$

for $h > 1$, and $\mathbf{s}_1 = \mathbf{o}_1 = \mathbf{x}_1$ (for simplicity, we embed \mathbf{o}_1 into $\mathcal{P}_{\tau_{\text{obs}}-1}$ via zero-padding). We call τ_{chunk} -length sequences of primitive controllers *composite actions*

$$\mathbf{a}_h = \kappa_{t_h:t_{h-1}} \in \mathcal{A} := \mathcal{K}^{\tau_{\text{chunk}}}. \quad \text{(composite actions)}$$

A *chunking policy* $\pi = (\pi_h)$ consists of functions π_h mapping observation-chunks \mathbf{o}_h to distributions $\Delta(\mathcal{A})$ over composite actions and interacting with the dynamics (2.1) by $\mathbf{a}_h = \kappa_{t_h:t_{h-1}} \sim \pi_h(\mathbf{o}_h)$, and executing $\mathbf{u}_t = \kappa_t(\mathbf{x}_t)$. We let $d_{\mathcal{A}} = \tau_{\text{chunk}}(d_x + d_u + d_x d_u)$ denote the dimension of the space \mathcal{A} of composite actions. The chunking scheme is represented in [Figure 2](#), demonstrating the rationale for using primitive controllers over open-loop actions. [Remark C.1](#) describes our rationale for studying *states* over generic observations, and considering time-dependent policies.

Desideratum. The quality of imitation of a deterministic policy is naturally measured in terms of step-wise closeness of state and action [68, 50]. With stochastic policies, however, two rollouts of even the same policy can visit different states. We propose measuring *distributional closeness* via *couplings* introduced in the preliminaries above. We define the following losses, focusing on the *marginal distributions* between trajectories.

Definition 2.2. Given $\varepsilon > 0$ and a (chunking) policy π , the (marginal distribution) imitation loss is $\mathcal{L}_{\text{marg}, \varepsilon}(\pi) := \max_{t \in [T]} \inf_{\mu} \mathbb{P}_{\mu} [\max \{ \|\mathbf{x}_{t+1}^{\text{exp}} - \mathbf{x}_{t+1}^{\pi}\|, \|\mathbf{u}_t^{\text{exp}} - \mathbf{u}_t^{\pi}\| \} > \varepsilon]$, where the infimum is over all couplings μ between the distribution of $\boldsymbol{\rho}_T$ under \mathcal{D}_{exp} and that induced by the policy π as described above, such that $\mathbb{P}_{\mu}[\mathbf{x}_1^{\text{exp}} = \mathbf{x}_1^{\pi}] = 1$.

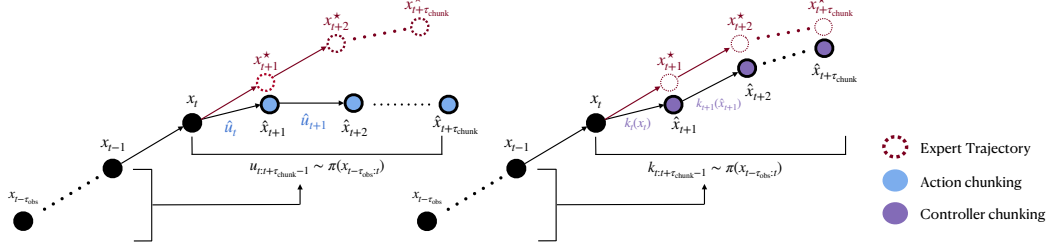


Figure 2: Graphical comparison of an action-chunk based policy (left) as described in [19], versus the primitive-controller chunking policy (right) proposed in this paper. The primitive controller paradigm allows for stabilizing back to the original expert trajectory, whereas using generated actions in an open-loop fashion may cause divergence from the expert in the presence of unstable system dynamics. We refer to **composite actions** as the sequence of primitive controllers given on the right.

Under stronger conditions (whose necessity we establish), we can also imitate joint distributions over actions (Appendix J). Observe that $\mathcal{L}_{\text{fin},\varepsilon} \leq \mathcal{L}_{\text{marg},\varepsilon}$, and that both losses are equivalent to Wasserstein-type metrics on bounded domains. These losses are also equivalent to Lévy-Prokhorov metrics [64] under re-scaling of the Euclidean metric (even for unbounded domains), and also correspond to total variation analogues of shifted Renyi divergences [5, 6]. While empirically evaluating these infima over couplings is challenging, $\mathcal{L}_{\text{marg},\varepsilon}$ upper bounds the difference in expectation between any bounded and Lipschitz control cost decomposing across time steps, states and inputs, and $\mathcal{L}_{\text{fin},\varepsilon}$ upper bounds differences in final-state costs; see Appendix J for discussion.

Diffusion Models. Our analysis provides imitation guarantees when chunking policies π_h select a_h via a sufficiently accurate generative model. Given their recent success, instantiate our analysis for the popular Denoising Diffusion Probabilistic Models (DDPM) framework [18, 41] that allows the learner to sample from a density $q \in \Delta(\mathbb{R}^d)$ assuming that the *score* $\nabla \log q$ is known to the learner. More precisely, suppose the learner is given an observation o_h and wishes to sample $a_h \sim q(\cdot | o_h)$ for some family of probability kernels $q(\cdot | \cdot)$. A DDPM starts with some a_h^0 sampled from a standard Gaussian noise and iteratively “denoises” for each DDPM-time step $0 \leq j < J$:

$$a_h^j = a_h^{j-1} - \alpha \cdot s_{\theta,h}(a_h^{j-1}, o_h, j) + 2 \cdot \mathcal{N}(0, \alpha^2 \mathbf{I}), \quad (2.2)$$

where $s_{\theta,h}(a_h^j, o_h, j)$ estimates the true score $s_{*,h}(a_h, o_h, \alpha j)$, formally defined for any continuous argument $t \leq J\alpha$ to be $s_{*,h}(a, o_h, t) := \nabla_a \log q_{h,[t]}^*(a | o_h)$, where $q_{h,[t]}^*(\cdot | o_h)$ is the distribution of $e^{-t} a_h^{(0)} + \sqrt{1 - e^{-2t}} \gamma$ with $a_h^{(0)}$ is sampled from the target distribution we which to sample from, and $\gamma \sim \mathcal{N}(0, \mathbf{I})$ is a standard Gaussian. We denote by $\text{DDPM}(s_\theta, o_h)$ the law of a_h^J sampled according to the DDPM using $s_\theta(\cdot, o_h, \cdot)$ as a score estimator. Preliminaries on DPPMs are detailed in Appendix L.

3 Conditional sampling with stabilization suffices for behavior cloning

We show that trajectories of the form given in Definition 2.1 can be efficiently imitated if (a) we are given a *synthesis oracle*, described below, that produces low-level control policies that locally stabilize chunks of the trajectory with primitive controllers and (b) we can learn to generate certain appropriate distributions over composite actions, i.e. sequences of primitive controllers. All the following results apply to affine primitive controllers introduced in Section 2 and assume that the system dynamics are second-order smooth and locally stabilizable. In Appendix E, we show that our results still hold with general families of parametric primitive controllers, provided that these controllers induce the same local stability guarantee.

The synthesis oracle. We say primitive controller (cf. Section 2) $\kappa_{1:T} \in \mathcal{K}^T$ is *consistent* with a trajectory $\rho = (x_{1:T+1}, u_{1:T}) \in \mathcal{P}_T$ if $\bar{x}_t = x_t$ and $\bar{u}_t = u_t$ for all $t \in [T]$; note that this implies that $\kappa_t(x_t) = u_t$ for all t . A *synthesis oracle* synth maps $\mathcal{P}_T \rightarrow \mathcal{K}^T$ such that, for all $\rho_T \in \mathcal{P}_T$, $\kappa_{1:T} = \text{synth}(\rho_T)$ is consistent with ρ_T . For our theory, we assume access to a synthesis oracle at training time, and assume the ability to estimate conditional distributions over joint sequences of primitive controllers; Appendix K explains how this can be implemented by solving Ricatti equations if dynamics are known (e.g. in a simulator), smooth, and stabilizable. In our experimental environment, control inputs are desired robot configurations, which the simulated

robot executes by applying feedback gains. As discussed in [Appendix E](#), learned or hand-coded low-level controllers are popular in practical implementations of generative behavior cloning. We discuss the merits of studying imitation learning with a synthesis oracle in depth in [Appendix C.3](#).

Notions of distance. While restricting ourselves to affine primitive controllers, our approximation error of generative behavior cloner is measured in terms of optimal transport distances that use the following “maximum distance.” Given two composite actions $\mathbf{a} = (\bar{\mathbf{u}}_{1:\tau_{\text{chunk}}}, \bar{\mathbf{x}}_{1:\tau_{\text{chunk}}}, \bar{\mathbf{K}}_{1:\tau_{\text{chunk}}})$ and $\mathbf{a}' = (\bar{\mathbf{u}}'_{1:\tau_{\text{chunk}}}, \bar{\mathbf{x}}'_{1:\tau_{\text{chunk}}}, \bar{\mathbf{K}}'_{1:\tau_{\text{chunk}}})$, we define

$$\mathbf{d}_{\max}(\mathbf{a}, \mathbf{a}') := \max_{1 \leq k \leq \tau_{\text{chunk}}} (\|\bar{\mathbf{u}}_k - \bar{\mathbf{u}}'_k\| + \|\bar{\mathbf{x}}_k - \bar{\mathbf{x}}'_k\| + \|\bar{\mathbf{K}}_k - \bar{\mathbf{K}}'_k\|). \quad (3.1)$$

Distances between policies are defined via natural optimal transport costs. Given two policies $\pi = (\pi_h)$, $\pi' = (\pi'_h)$ and observation chunk \mathbf{o}_h , we define an induced optimal transport cost

$$\Delta_\varepsilon(\pi_h(\mathbf{o}_h), \pi'_h(\mathbf{o}_h)) := \inf_{\mu} \mathbb{P}_{(\mathbf{a}_h, \mathbf{a}'_h) \sim \mu} [\mathbf{d}_{\max}(\mathbf{a}_h, \mathbf{a}'_h) > \varepsilon],$$

where the \inf_{μ} denotes the infimum over all couplings between $\mathbf{a}_h \sim \pi_h(\mathbf{o}_h)$ and $\mathbf{a}'_h \sim \pi'_h(\mathbf{o}_h)$. Δ_ε corresponds to a relaxed Lévy-Prokhorov metric [64], and can always be bounded, via Markov’s inequality, by $\Delta_\varepsilon(\pi_h, \pi'_h \mid \mathbf{o}_h) \leq \frac{1}{\varepsilon} \mathcal{W}_{1, \mathbf{d}_{\max}}(\pi_h(\mathbf{o}_h), \pi'_h(\mathbf{o}_h))$, where $\mathcal{W}_{1, \mathbf{d}_{\max}}(\pi_h(\mathbf{o}_h), \pi'_h(\mathbf{o}_h))$ denotes the 1-Wasserstein distance between $\mathbf{a}_h \sim \pi_h(\mathbf{o}_h)$ and $\mathbf{a}'_h \sim \pi'_h(\mathbf{o}_h)$.

3.1 Incremental Stability and the Synthesis Oracle.

We assume that synthesis oracle above produces *incrementally stabilizing* control gains, in the sense first proposed by [9]. Incremental stability has emerged as a natural desirable property for imitation limitation [50, 68, 27], because it forces the expert to be robust to small perturbations of their policy. We now supply a formal definition. Given a primitive controller $\kappa : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_u}$, define the closed loop dynamic map $f_{\text{cl}, \kappa}(\mathbf{x}, \delta \mathbf{u}) := f(\mathbf{x}, \kappa(\mathbf{x}) + \delta \mathbf{u})$. Thus, composite action \mathbf{a} is *consistent* with a trajectory chunk $\mathbf{s} = (\mathbf{x}_{1:\tau_{\text{chunk}}+1}, \mathbf{u}_{1:\tau})$ if $\mathbf{x}_{t+1} = f_{\text{cl}, \kappa_t}(\mathbf{x}_t, \mathbf{0})$ for $1 \leq t \leq \tau_{\text{chunk}}$.¹

Definition 3.1 (Time-Varying Incremental Stability). Let $\gamma(\cdot)$ be a class K function, $\beta(\cdot, \cdot)$ be class KL function, and let $\mathbf{a} = (\kappa_1, \kappa_2, \dots, \kappa_\tau)$ denote a sequence of primitive controllers (i.e. a composite action when $\tau = \tau_{\text{chunk}}$). Given a sequence of input perturbations $\delta \mathbf{u}_{1:\tau} \in (\mathbb{R}^{d_u})^\tau$ and initial condition $\boldsymbol{\xi} \in \mathbb{R}^{d_x}$, let $\mathbf{x}_{i+1}^{\mathbf{a}}(\delta \mathbf{u}_{1:\tau}, \boldsymbol{\xi}) = f_{\text{cl}, \kappa_i}(\mathbf{x}_i^{\mathbf{a}}(\delta \mathbf{u}_{1:\tau}, \boldsymbol{\xi}), \delta \mathbf{u}_i)$, with $\mathbf{x}_1^{\mathbf{a}} = \boldsymbol{\xi}$. We say that composite action \mathbf{a} is time-varying incrementally input-to-state stable (t-ISS) with moduli $\gamma(\cdot), \beta(\cdot, \cdot)$ if for all $\boldsymbol{\xi}, \boldsymbol{\xi}' \in \mathbb{R}^{d_x}$, $0 \leq i \leq \tau$, $\|\mathbf{x}_i^{\mathbf{a}}(\mathbf{0}_{1:\tau}, \boldsymbol{\xi}) - \mathbf{x}_i^{\mathbf{a}}(\delta \mathbf{u}_{1:\tau}, \boldsymbol{\xi}')\| \leq \beta(\|\boldsymbol{\xi} - \boldsymbol{\xi}'\|, \tau) + \gamma(\max_{1 \leq s \leq i-1} \|\delta \mathbf{u}_s\|)$. Given parameters $c_\gamma, c_\xi > 0$ we say that \mathbf{a} is local t-ISS at $\boldsymbol{\xi}_0$ if the above holds only for all $\boldsymbol{\xi}, \boldsymbol{\xi}', \delta \mathbf{u}_{1:\tau}$ such that $\|\boldsymbol{\xi} - \boldsymbol{\xi}_0\|, \|\boldsymbol{\xi}' - \boldsymbol{\xi}_0\| \leq c_\xi$ and $\max_t \|\delta \mathbf{u}_t\| \leq c_\gamma$.

Incremental stability implies that as the initial conditions $\|\boldsymbol{\xi} - \boldsymbol{\xi}'\| \rightarrow 0$ and $\max_{0 \leq s \leq i-1} \|\delta \mathbf{u}_s\| \rightarrow 0$, the trajectories induced by taking rolling out \mathbf{a} from $\boldsymbol{\xi}$, and rolling out \mathbf{a}' with additive input perturbations $\delta \mathbf{u}_{1:\tau}$ tend to zero in norm. This behavior needs only hold for initial conditions in a small neighborhood of a nominal state $\boldsymbol{\xi}_0$. Importantly, the perturbations $\delta \mathbf{u}_{1:\tau}$ are fixed perturbations of inputs, applied to the *closed loop behavior* under the controllers. Our notion of incremental stability are similar too, but subtly different similar notions of past work. We provide an extended comparisons in [Appendix E.2](#). Our main assumption is that the synthesis oracle described above produced primitive controllers which are consistent with, and incrementally stabilizing for, the demonstrated trajectories. [Figure 1](#) demonstrates the effect of stabilizing primitive controllers.

Assumption 3.1. We assume that our synthesis oracle enjoys the following property. Let $\boldsymbol{\rho}_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \sim \mathcal{D}_{\text{exp}}$, and let $\kappa_{1:T} = \text{synth}(\boldsymbol{\rho}_T)$, partitioned into composite actions $\mathbf{a}_{1:H}$, with $\kappa_t(\mathbf{x}) = \bar{\mathbf{K}}_t(\mathbf{x} - \bar{\mathbf{x}}_t) + \bar{\mathbf{u}}_t$. We assume that, with probability one, $\kappa_{1:T}$ is consistent with $\boldsymbol{\rho}_T$ ², and that, for each $1 \leq h \leq H$, $\mathbf{a}_h = (\kappa_{t_h:t_h+\tau_{\text{chunk}}-1})$ is local t-ISS at \mathbf{x}_{t_h} with moduli γ, β and parameters $c_\beta, c_\xi > 0$. We further assume that γ and β take the form

$$\gamma(u) = \bar{c}_\gamma \cdot u, \quad \beta(u, k) = \bar{c}_\beta e^{-(k-1)L_\beta} \cdot u, \quad \bar{c}_\gamma, \bar{c}_\beta > 0, \quad L_\beta \in (0, 1].$$

¹Below, we recall definitions of classes of comparison functions in nonlinear control [38] as follows: we say a univariate function $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is *class K* if it is strictly increasing and satisfies $\gamma(0) = 0$. We say a bivariate function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{Z}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is *class KL* if $x \mapsto \beta(x, t)$ is *class K* for each $t \geq 0$, and $t \mapsto \beta(x, t)$ is nonincreasing in t .

²Note that this implies $\bar{\mathbf{x}}_t = \mathbf{x}_t$ and $\bar{\mathbf{u}}_t = \mathbf{u}_t$.

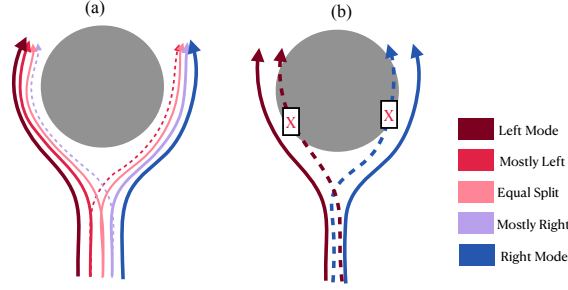


Figure 3: Graphical representation of total variation continuity (TVC) using the running “left mode/right mode” example. Panel (a) depicts a policy π which is TVC, and thus interpolates between left and right modes probabilistically. Importantly, the TVC property applies to the distribution over composite actions, i.e. sequences of *primitive controllers*, as, in Definition 3.3; this ensures, for example, that following the left mode from slightly to the right of the obstacle (purple dotted line) still stabilizes to the idealized left mode trajectory (red). In panel (b), we consider a policy which for TVC applies to the sequences of *raw control inputs* (which is *not* what occurs in Definition 3.3). This can lead to naive mode-switching that collides with the gray obstacle.

Lastly, we assume that for the expert trajectories and the primitive controllers drawn as above, it holds that satisfy $\max\{\|\mathbf{x}_t\|, \|\mathbf{u}_t\|\} \leq R_{\text{dyn}}$ and $\|\mathbf{K}_t\| \leq R_{\mathbf{K}}$ with probability one.

In Appendix K, we show that Assumption 3.1 holds whenever (a) the dynamics of our system are smooth (but not necessarily linear!) (b) the affine gains are chosen to stabilize the Jacobian linearizations of the system around the nominal trajectory.

Definition 3.2 (Problem constants). Throughout, we refer to constants $c_1, c_2, c_3, c_4, c_5 > 0$, which are polynomial in the terms in Assumption 3.1, and which are defined formally in Appendix K.

3.2 Simplified guarantees under total variation continuity

This section presents our main theoretical result: if one learns a chunking policy $\hat{\pi}$ that can compute the conditional distribution of composite actions at time steps given observation-chunks, then a stochastically smoothed version of this policy, $\hat{\pi}_\sigma$, has low imitation error. Define, for any length $\tau \in \mathbb{N}$, the *trajectory distance* between trajectories $\boldsymbol{\rho} = (\mathbf{x}_{1:\tau+1}, \mathbf{u}_{1:\tau})$, $\boldsymbol{\rho}' = (\mathbf{x}'_{1:\tau+1}, \mathbf{u}'_{1:\tau}) \in \mathcal{P}_\tau$

$$d_{\text{traj}}(\boldsymbol{\rho}, \boldsymbol{\rho}') := \max_{1 \leq k \leq \tau+1} \|\mathbf{x}_k - \mathbf{x}'_k\| \vee \max_{1 \leq k \leq \tau} \|\mathbf{u}_k - \mathbf{u}'_k\|. \quad (3.2)$$

In particular, we define $d_{\text{traj}}(\mathbf{o}_h, \mathbf{o}'_h)$ and $d_{\text{traj}}(s_h, s'_h)$ by viewing these as trajectories of length $\tau_{\text{obs}} - 1$ and τ_{chunk} , respectively. Lastly, we define a per-timestep restriction of the expert distribution. In this section, we consider the case where the learner policy satisfies a total variation continuity (TVC) condition, defined below.

Definition 3.3 (TVC of Chunking Policies). We say that a chunking policy $\pi = (\pi_h)$ is total variation continuous with modulus $\gamma_{\text{TVC}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, written $\gamma_{\text{TVC}}\text{-TVC}$, if, for all $h \in [H]$ and any observation-chunks $\mathbf{o}_h, \mathbf{o}'_h \in \mathcal{P}_{\tau_{\text{obs}}-1}$, $\text{TV}(\pi_h(\mathbf{o}_h), \pi_h(\mathbf{o}'_h)) \leq \gamma_{\text{TVC}}(d_{\text{traj}}(\mathbf{o}_h, \mathbf{o}'_h))$.

We depict the TVC property using our running left-right obstacle example in Figure 3. We stress that, in Definition 3.3, the TV bound on $\text{TV}(\pi_h(\mathbf{o}_h), \pi_h(\mathbf{o}'_h))$ applies to the *composite actions* consisting of primitive controllers $\mathbf{a}_h = \kappa_{t_h:t_h+\tau_{\text{chunk}}-1} \sim \pi_h(\mathbf{o}_h)$; it does not upper bound the TV distance between raw control inputs. Indeed, ensuring TVC of the latter can lead to the failure modes depicted in Figure 3(b). Next, we extract an expert “policy” from the expert demonstrations.

Definition 3.4 (Expert “policy” with synthesized controllers). For $h \in [H]$, we let $\mathcal{D}_{\text{exp},h}$ denote the joint distribution of $(\mathbf{a}_h, \mathbf{o}_h)$, induced by drawing a trajectory $\boldsymbol{\rho}_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \sim \mathcal{D}_{\text{exp}}$ from the expert distribution, $\kappa_{1:T} = \text{synth}(\boldsymbol{\rho}_T)$ be the associated primitive controllers, letting $\mathbf{o}_h = (\mathbf{x}_{t_h-\tau_{\text{obs}}+1:t_h}, \mathbf{u}_{t_h-\tau_{\text{obs}}+1:t_h-1})$ be the associated observation-chunk at time h , and $\mathbf{a}_h = \kappa_{t_h:t_h+1-1}$ the associated composite action. We let $\pi_h^*(\cdot) : \mathcal{O} \rightarrow \Delta(\mathcal{A})$ denote the condition distribution of $\mathbf{a}_h \mid \mathbf{o}_h$ under $\mathcal{D}_{\text{exp},h}$.

The conditional distributions $\pi_h^*(\cdot)$ are estimated when training a generative model to predict \mathbf{a}_h from observations \mathbf{o}_h . Note that $\pi_h^*(\cdot)$ (and $\mathcal{D}_{\text{exp},h}$) is defined in terms of *both* expert demonstration

from \mathcal{D}_{exp} and the associated synthesized primitive controllers. In [Lemma J.6](#), we show that when the synthesis oracle $\kappa_{1:T} = \text{synth}(\mathbf{p}_T)$ produces primitive controllers consistent with the trajectories, then $\pi^* = (\pi_h^*)$ produces the same marginals over states as \mathcal{D}_{exp} ; that is, $\mathcal{L}_{\text{marg},\varepsilon}(\pi^*) = 0$.

Theorem 1. Suppose [Assumption 3.1](#) holds, and suppose that $0 \leq \varepsilon < c_2$, and $\tau_{\text{chunk}} \geq c_3$. Then, for any non-decreasing non-negative $\gamma_{\text{TVC}}(\cdot)$ and γ_{TVC} -TVC chunking policy $\hat{\pi}$, it holds that $\mathcal{L}_{\text{marg},\varepsilon}(\hat{\pi}) \leq H\gamma_{\text{TVC}}(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{\mathbf{o}_h \sim \mathcal{D}_{\text{exp},h}} \Delta_{(\varepsilon/c_1)}(\pi_h^*(\mathbf{o}_h), \hat{\pi}_h(\mathbf{o}_h))$, which is at most $H\gamma_{\text{TVC}}(\varepsilon) + \frac{c_1}{\varepsilon} \sum_{h=1}^H \mathbb{E}_{\mathbf{o}_h \sim \mathcal{D}_{\text{exp},h}} [\mathcal{W}_{1,d_{\text{max}}}(\pi_h^*(\mathbf{o}_h), \hat{\pi}_h(\mathbf{o}_h))]$.

The above result reduces the marginal imitation error of $\hat{\pi}$ to the sum over optimal transport errors between $\hat{\pi}$ and a \mathbf{o}_h chosen by the expert demonstrators. Thus, if these are small, the local stabilization properties of the primitive controllers guaranteed by [Assumption 3.1](#) ensure that errors compound at most linearly in problem horizon. The key ideas of the proof are given [Appendix D](#), via a general template for imitation learning of general stochastic policies. This template is instantiated with a details in [Appendix J](#).

3.3 A general guarantee via data noising.

To circumvent assuming that the learner’s policy is TVC, we study estimating the conditionals under a popular data augmentation technique [\[36\]](#), where the learner is trained to imitate the conditional sequence of a $\cdot \mid \tilde{\mathbf{o}}_h$, where $\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$ adds σ^2 -variance Gaussian noise to the true observation-chunk. To understand this better, consider the following *smoothed* policy:

Definition 3.5 (The smoothed policy). Let $\hat{\pi} = (\hat{\pi}_h)$ be a chunking policy. We define the *smoothed policy* $\hat{\pi}_\sigma = (\hat{\pi}_{\sigma,h})$ by letting $\hat{\pi}_{\sigma,h}(\cdot \mid \tilde{\mathbf{o}}_h)$ be distributed as $\hat{\pi}_h(\cdot \mid \tilde{\mathbf{o}}_h)$, where $\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$.

[Appendix J.7.2](#) show’s that Pinsker’s inequality implies noising automatically enforces TVC. This suggests that we can use some form of data noising to enforce the TVC property in [Definition 3.3](#). Let’s now consider a related problem: trying to estimate the optimal distribution over composite actions *conditioned on* a noised observation. This gives rise to a *deconvolution* of the expert policy, which can be thought as an inverse operation of data noising.

Definition 3.6 (Noised Data Distribution and Deconvolution Policy). Let $\mathcal{D}_{\text{exp},h}$ be as in [Definition 3.4](#). Define $\mathcal{D}_{\text{exp},\sigma,h}$ as the distribution over $(\tilde{\mathbf{o}}_h, \mathbf{a}_h)$ generated by $(\mathbf{o}_h, \mathbf{a}_h) \sim \mathcal{D}_{\text{exp},h}$ and $\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$. We define the *deconvolution policy* $\pi_{\text{dec},\sigma,h}^*(\tilde{\mathbf{o}}_h)$ as the conditional distribution of $\mathbf{a}_h \mid \tilde{\mathbf{o}}_h$ under $\mathcal{D}_{\text{exp},\sigma,h}$.

Analogously to π^* , the policy $\pi_{\text{dec},\sigma,h}^*$ is what a generative model trained to generate \mathbf{a}_h from noised observations $\tilde{\mathbf{o}}_h$ of $\mathbf{o}_h \sim \mathcal{D}_{\text{exp}}$ learns to generate. Our next theorem states that, if our $\hat{\pi}$ approximates the idealized conditional distribution of composite actions given noised observations, then $\hat{\pi}_\sigma$, the smoothed policy, imitates the expert distribution with provable bounds on its imitation error:

Theorem 2 (Reduction to conditional sampling under noising). Suppose [Assumption 3.1](#) holds. Let $c_1, \dots, c_5 > 0$, defined in [Definition 3.2](#), and let $\Theta_{\text{Iss}}(x)$ denote a term which is upper and lower bounded by a x times a polynomial in those constants and their inverses. Then, for $\varepsilon \leq \Theta_{\text{Iss}}(1)$, if we choose $\sigma = \varepsilon / \Theta_{\text{Iss}}(\sqrt{d_x} + \log(1/\varepsilon))$ and let $\tau_{\text{chunk}} \leq c_3$ and $\tau_{\text{chunk}} - \tau_{\text{obs}} \geq \frac{1}{L_B} \log(c_1/\varepsilon)$,

$$\mathcal{L}_{\text{marg},\varepsilon}(\hat{\pi}_\sigma) \leq \Theta_{\text{Iss}}(\varepsilon H \sqrt{\tau_{\text{obs}}} \cdot (\sqrt{d_x} + \log(\frac{1}{\varepsilon}))) + \sum_{h=1}^H \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp},\sigma,h}} \left[\Delta_{(\varepsilon^2)}(\pi_{\text{dec},\sigma,h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h)) \right],$$

which is upper bounded by at most $\Theta_{\text{Iss}}(\varepsilon H \sqrt{\tau_{\text{obs}}} \cdot (\sqrt{d_x} + \log(1/\varepsilon))) + \frac{1}{\varepsilon^2} \sum_{h=1}^H \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp},\sigma,h}} \left[\mathcal{W}_{1,d_{\text{max}}}(\pi_{\text{dec},\sigma,h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h)) \right]$.

To reiterate, [Theorem 2](#) guarantees imitation of the distribution of marginals and final states of \mathcal{D}_{exp} by replacing the explicit TVC assumption with noising, and the resulting guarantee applies to the *smoothed policy* $\hat{\pi}_\sigma$ which adds smoothing noise back in. [Appendix J](#) gives a number of additional results. In [Appendix I](#), we show that the proof framework, outlined in [Appendix D](#), which under lies the proofs of [Theorems 1](#) and [2](#), is essentially sharp in the worst case. Moreover, in [Appendix C.3](#), we discuss the merits and drawbacks of our use of the synthesis oracle, and how it circumvents some of the challenges encountered in behavior cloning in past work. The key intuition behind the proof of [Theorem 2](#) is depicted in [Figure 4](#), and full proof sketch is deferred to [Appendix C.2](#)

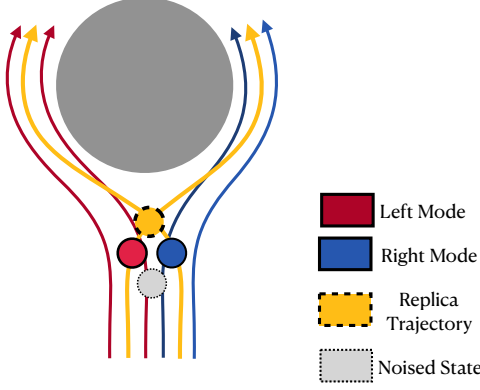


Figure 4: Multi-modal demonstrations traverse an obstacle left or right, exhibiting a pure bifurcation. We consider perturbing expert data on the right mode (blue circle) to a noised datum (gray circle). We show that generative behavior cloners learn to deconvolve this noise, creating a virtual “replica” sample (red circle) following the left mode, such that the replica and original are i.i.d. given the noised one. When the red circle’s primitive controllers are rolled from from the blue circle, this leads to a trajectory (yellow circle) which interpolates across the bifurcations. Marginalizing over this process, the yellow trajectories probabilistically interpolate between red and blue modes, and (approximately) match the per-time-step marginal over expert distributions.

4 HINT: Instantiating Data Noising with DDPMs

We now instantiate [Theorem 2](#) by showing that one can learn a policy $\hat{\pi}$ for which the error terms in [Theorems 1](#) and [2](#) are small by fitting a DDPM to noise-smoothed data. Our proposed algo-

Algorithm 1 Hierarchical Imitation via Noising at Inference Time (HINT)

[h]

- 1: **Initialize** Synthesis oracle `synth`, sample sizes $N_{\text{exp}}, N_{\text{aug}} \in \mathbb{N}$, $\sigma \geq 0$, DDPM step size $\alpha > 0$, DDPM horizon J , function class $\{\mathbf{s}_\theta\}_{\theta \in \Theta}$, gain magnitude $R > 0$, empty data buffer $\mathcal{D} \leftarrow \emptyset$.
`% For no smoothing, set $\sigma = 0$ and $N_{\text{aug}} = 1$`
 - 2: **for** $n = 1, 2, \dots, N_{\text{exp}}$ **do**
 - 3: Sample $\mathbf{p}_T = (x_{1:T+1}, u_{1:T}) \sim \mathcal{D}_{\text{exp}}$ and set $\kappa_{1:T} = \text{synth}(\mathbf{p})$
`% Segment $\mathbf{o}_{1:H}$ from \mathbf{p}_T and $\mathbf{a}_{1:H}$ from $\kappa_{1:T}$`
 - 4: **for** $i = 1, 2, \dots, N_{\text{aug}}$ and $h = 1, 2, \dots, H$ **do**
 - 5: Sample $\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$, $j_h \sim \text{Unif}([J])$ and $\gamma_h \sim \mathcal{N}(0, (j_h \alpha)^2 \mathbf{I})$.
 - 6: $\mathcal{D} \leftarrow \mathcal{D} \cdot \text{append}(\{(\mathbf{a}_h, \tilde{\mathbf{o}}_h, j_h, \gamma_h, h)\})$
 - 7: Fit $\theta \in \arg \min_{\theta \in \Theta} \mathcal{L}_{\text{DDPM}}(\theta, \mathcal{D})$, and let $\hat{\pi} = (\hat{\pi}_\sigma)$ be given by $\hat{\pi}(\cdot \mid \mathbf{o}_h) = \text{DDPM}(\mathbf{s}_{\theta,h}, \mathbf{o}_h)$.
 - 8: **return** $\hat{\pi}_\sigma = (\hat{\pi}_{\sigma,h})$, by smoothing $\hat{\pi}$ as per [Definition 3.5](#).
-

rithm, HINT ([Algorithm 1](#)) combines DDPM-learning of chunked policies as in [\[19\]](#) with a popular form of data-augmentation [\[36\]](#). We collect N_{exp} expert trajectories, synthesize gains, and segment trajectories into observation-chunks \mathbf{o}_h and composite actions \mathbf{a}_h as described in [Section 2](#). We perturb each \mathbf{o}_h to form N_{aug} chunks $\tilde{\mathbf{o}}_h$, as well as horizon indices $j \in [J]$ and inference noises $\gamma \sim \mathcal{N}(0, (\alpha j)^2 \mathbf{I})$, and add these tuples $(\mathbf{a}_h, \tilde{\mathbf{o}}_h, j_h, \gamma_h, h)$ to our data \mathcal{D} . We end the training phase by minimizing the standard DDPM loss [\[62\]](#):

$$\mathcal{L}_{\text{DDPM}}(\theta, \mathcal{D}) = \sum_{(\mathbf{a}_h, \tilde{\mathbf{o}}_h, j_h, \gamma_h, h) \in \mathcal{D}} \left\| \gamma_h - \mathbf{s}_{\theta,h} \left(e^{-\alpha j} \mathbf{a}_h + \sqrt{1 - e^{-2\alpha j}} \gamma_h, \tilde{\mathbf{o}}_h, j_h \right) \right\|^2. \quad (4.1)$$

Our algorithm differs subtly from past work in [Line 8](#): motivated by [Theorem 2](#), we add smoothing noise *back in* at test time. Here, the notation $\text{DDPM}(\mathbf{s}_{\theta,h}, \cdot) \circ \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$ means, given \mathbf{o}_h , we perturb it to $\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$, and sample $\mathbf{a}_h \sim \text{DDPM}(\mathbf{s}_{\theta,h}, \tilde{\mathbf{o}}_h)$. We now state an informal guarantee for HINT, deferring a formal statement to [Appendix C.5](#).

Theorem (Informal Theorem). Suppose that the system dynamics are smooth and that [Assumption 3.1](#) holds for the linearized system. Then there is a choice of the parameters in HINT that is polynomial in all problem parameters such that for N_{exp} , polynomially large in problem parameters, $\mathcal{L}_{\text{marg},\varepsilon}(\hat{\pi}_\sigma) \leq \Theta(\varepsilon H \sqrt{\tau_{\text{obs}}}(\sqrt{d_x} + \log(1/\varepsilon)))$ with high probability.

4.1 Experimental Results

In this section, we demonstrate the benefits of diffusing low level controllers, and of our approach to data noising. We explain the environments in greater detail, along with all training and compu-

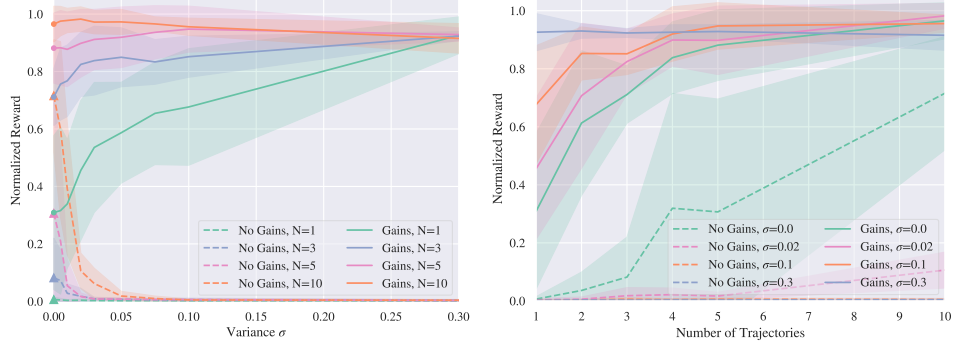


Figure 5: Performance of diffusing chunks of actions $\bar{\mathbf{u}}_{t_{h-1}:t_h-1}$ ("No Gains") versus jointly diffusing actions $\mathbf{u}_{t_{h-1}:t_h-1}$, reference states $\bar{\mathbf{x}}_{t_{h-1}:t_h-1}$ and gains $\mathbf{K}_{t_{h-1}:t_h-1}$ for a 2-D quadrotor system with thrust-and-torque-based control. Different noise levels σ and number of trajectories N are shown. Mean and standard deviation are shown across 5 training seeds.

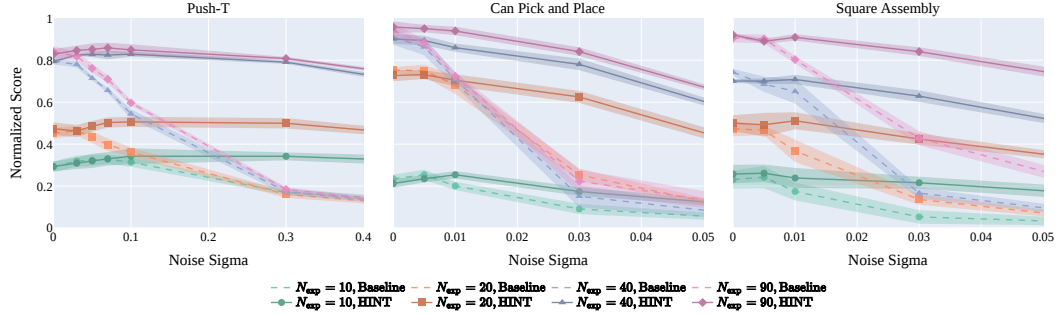


Figure 6: Performance of baseline $\hat{\pi}$ and noise-injected $\hat{\pi} \circ W_\sigma$ HINT policy for different σ . We use 4 training seeds with 50 and 22 test trajectories per seed for PushT and Can and Square Environments respectively. Mean and standard deviation of the test performance on the 3 best checkpoints across the 4 seeds are plotted. The σ values correspond to noise in the normalized $[-1, 1]$ range.

tational details in Appendix O.³ Figure 5 compares the performance of diffusing (chunks of) raw control inputs to diffusing (chunks of) gain matrices for a canonical model of a 2-d quadrotor. We find that diffusing gain matrices yields dramatic improvements in performance, in particular allowing a *single imitated trajectory* to outperform learning raw control inputs from 10 demonstrations.

Next, empirically evaluate the effect on policy performance of our proposal to inject noise back into the dynamics at inference time. We consider three challenging robotic manipulation tasks studied in prior work: PushT block-pushing [19]; Robomimic Can Pick-and-Place and Square Nut Assembly [42] (we direct the reader to Chi et al. [19] for an extensive empirical investigation into the performance of diffusion policies in the un-noised $\sigma = 0$ regime). We display the results of our experiments in Figure 6. Observe that the performance degradation of the replica policy from the unnoised $\sigma = 0$ variant is minimal across all environments and even leads to a slight but noticeable improvement in the small-noise regime for PushT (and low-data Can Pick and Place). In the presence of non-negligible noise HINT significantly outperforms the conventional policy $\hat{\pi}$ (obtained by noising observations at training but not test time), as predicted by our theory.

5 Discussion

This work considerably loosened assumptions placed on the *expert distribution* by introducing a synthesis oracle responsible for stabilization. How best to achieve low-level stabilization remains an open question. We hope that this work encourages further empirical research into improving the stability of imitation learning, either via the hierarchical route proposed in this paper or via new innovations.

³Code for PushT, Robomimic experiments can be found at https://github.com/pfrommerd/diffusion_policy_pt. Quadrotor experiments are in <https://github.com/pfrommerd/stanza/>

Acknowledgments

We would like to thank Cheng Chi for his extensive help in running the code in [19] and Benjamin Burchfiel for numerous insightful discussions and intuitions about DDPM policies. Lastly, we thank Alexandre Amice for his helpful feedback on an earlier draft of this manuscript. AB acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 as well as the support of DOE through award DE-SC0022199. MS and RT acknowledge support from Amazon.com Services LLC grant; PO# 2D-06310236. DP and AJ acknowledge support from the Office of Naval Research under ONR grant N00014-23-1-2299.

References

- [1] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016.
- [2] A. Ajay, Y. Du, A. Gupta, J. Tenenbaum, T. Jaakkola, and P. Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022.
- [3] A. Ajay, A. Gupta, D. Ghosh, S. Levine, and P. Agrawal. Distributionally adaptive meta reinforcement learning. *arXiv preprint arXiv:2210.03104*, 2022.
- [4] A. Ajay, S. Han, Y. Du, S. Li, A. Gupta, T. Jaakkola, J. Tenenbaum, L. Kaelbling, A. Srivastava, and P. Agrawal. Compositional foundation models for hierarchical planning. *arXiv preprint arXiv:2309.08587*, 2023.
- [5] J. Altschuler and K. Talwar. Privacy of noisy stochastic gradient descent: More iterations without more privacy loss. *Advances in Neural Information Processing Systems*, 35:3788–3800, 2022.
- [6] J. M. Altschuler and S. Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *arXiv preprint arXiv:2302.10249*, 2023.
- [7] B. D. Anderson and J. B. Moore. *Optimal control: linear quadratic methods*. Courier Corporation, 2007.
- [8] O. Angel and Y. Spinka. Pairwise optimal coupling of multiple random variables. *arXiv preprint arXiv:1903.00632*, 2019.
- [9] D. Angeli. A lyapunov approach to incremental stability properties. *IEEE Transactions on Automatic Control*, 47(3):410–421, 2002.
- [10] M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.
- [11] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- [12] P. L. Bartlett, A. Montanari, and A. Rakhlin. Deep learning: a statistical viewpoint. *Acta numerica*, 30:87–201, 2021.
- [13] A. Block, Y. Mroueh, and A. Rakhlin. Generative modeling with denoising auto-encoders and langevin sampling. *arXiv preprint arXiv:2002.00107*, 2020.
- [14] A. Block, Y. Mroueh, A. Rakhlin, and J. Ross. Fast mixing of multi-scale langevin dynamics under the manifold hypothesis. *arXiv preprint arXiv:2006.11166*, 2020.
- [15] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [16] J. D. Chang, M. Uehara, D. Sreenivas, R. Kidambi, and W. Sun. Mitigating covariate shift in imitation learning via offline data without great coverage. *arXiv preprint arXiv:2106.03207*, 2021.
- [17] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.

- [18] S. Chen, S. Chewi, J. Li, Y. Li, A. Salim, and A. Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *NeurIPS 2022 Workshop on Score-Based Methods*, 2022.
- [19] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [20] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn. Robonet: Large-scale multi-robot learning. *arXiv preprint arXiv:1910.11215*, 2019.
- [21] P. De Haan, D. Jayaraman, and S. Levine. Causal confusion in imitation learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [22] R. Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.
- [23] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine. One-shot visual imitation learning via meta-learning. In *Conference on robot learning*, pages 357–368. PMLR, 2017.
- [24] D. J. Foster and A. Rakhlin. L-infinity vector contraction for rademacher complexity. *arXiv preprint arXiv:1911.06468*, 6, 2019.
- [25] J. W. Hagoood and B. S. Thomson. Recovering a function from a dini derivative. *The American Mathematical Monthly*, 113(1):34–46, 2006.
- [26] P. Hansen-Estruch, I. Kostrikov, M. Janner, J. G. Kuba, and S. Levine. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*, 2023.
- [27] A. Havens and B. Hu. On imitation learning of linear control policies: Enforcing stability and robustness constraints via lmi conditions. In *2021 American Control Conference (ACC)*, pages 882–887. IEEE, 2021.
- [28] D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dorundo, R. Desai, T. Zhu, S. Parajuli, M. Guo, et al. Jacob steinhardt et justin gilmer. the many faces of robustness: A critical analysis of out-of-distribution generalization. *arXiv preprint arXiv:2006.16241*, 2020.
- [29] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [30] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- [31] A. Hussein, E. Elyan, M. M. Gaber, and C. Jayne. Deep imitation learning for 3d navigation tasks. *Neural computing and applications*, 29:389–404, 2018.
- [32] A. Hyvärinen and P. Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- [33] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [34] C. Jin, P. Netrapalli, R. Ge, S. M. Kakade, and M. I. Jordan. A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*, 2019.
- [35] L. P. Kaelbling and T. Lozano-Pérez. Hierarchical task and motion planning in the now. In *2011 IEEE International Conference on Robotics and Automation*, pages 1470–1477. IEEE, 2011.
- [36] L. Ke, J. Wang, T. Bhattacharjee, B. Boots, and S. Srinivasa. Grasping with chopsticks: Combating covariate shift in model-free imitation learning for fine manipulation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6185–6191. IEEE, 2021.
- [37] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8077–8083. IEEE, 2019.
- [38] H. Khalil. *Nonlinear Systems*. Pearson Education. Prentice Hall, 2002. ISBN 9780130673893. URL https://books.google.com/books?id=t_d1QgAACAAJ.
- [39] I. Kostrikov, D. Yarats, and R. Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. *arXiv preprint arXiv:2004.13649*, 2020.
- [40] M. Laskey, J. Lee, R. Fox, A. Dragan, and K. Goldberg. Dart: Noise injection for robust imitation learning. In *Conference on robot learning*, pages 143–156. PMLR, 2017.

- [41] H. Lee, J. Lu, and Y. Tan. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pages 946–985. PMLR, 2023.
- [42] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. *arXiv preprint arXiv:2108.03298*, 2021.
- [43] T. Marcucci, J. Umenberger, P. A. Parrilo, and R. Tedrake. Shortest paths in graphs of convex sets. *arXiv preprint arXiv:2101.11565*, 2021.
- [44] A. Maurer. A vector-contraction inequality for rademacher complexities. In *Algorithmic Learning Theory: 27th International Conference, ALT 2016, Bari, Italy, October 19-21, 2016, Proceedings 27*, pages 3–17. Springer, 2016.
- [45] M. Mezard and A. Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [46] D. Misra. Mish: A self regularized non-monotonic activation function. *arXiv preprint arXiv:1908.08681*, 2019.
- [47] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- [48] T. Pearce, T. Rashid, A. Kanervisto, D. Bignell, M. Sun, R. Georgescu, S. V. Macua, S. Z. Tan, I. Momennejad, K. Hofmann, et al. Imitating human behaviour with diffusion models. *arXiv preprint arXiv:2301.10677*, 2023.
- [49] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [50] D. Pfrommer, T. Zhang, S. Tu, and N. Matni. Tasil: Taylor series imitation learning. *Advances in Neural Information Processing Systems*, 35:20162–20174, 2022.
- [51] D. Pfrommer, M. Simchowitz, T. Westenbroek, N. Matni, and S. Tu. The power of learned locally linear models for nonlinear policy optimization. *arXiv preprint arXiv:2305.09619*, 2023.
- [52] Y. Polyanskiy and Y. Wu. *Information Theory: From Coding to Learning*. Cambridge University Press, 2022+.
- [53] M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In *Conference on Learning Theory*, pages 1674–1703. PMLR, 2017.
- [54] A. Rakhlin, K. Sridharan, and A. B. Tsybakov. Empirical entropy, minimax regret and minimax risk. *Bernoulli Society for Mathematical Statistics and Probability*, 2017.
- [55] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [56] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [57] S. Ross and D. Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668. JMLR Workshop and Conference Proceedings, 2010.
- [58] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [59] N. M. Shafiullah, Z. Cui, A. A. Altanzaya, and L. Pinto. Behavior transformers: Cloning k modes with one stone. *Advances in neural information processing systems*, 35:22955–22968, 2022.

- [60] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [61] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [62] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [63] Y. Song, S. Garg, J. Shi, and S. Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pages 574–584. PMLR, 2020.
- [64] V. Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.
- [65] X. Sun, S. Yang, and R. Mangharam. Mega-dagger: Imitation learning with multiple imperfect experts. *arXiv preprint arXiv:2303.00638*, 2023.
- [66] R. Tedrake. Lqr-trees: Feedback motion planning on sparse randomized trees. 2009.
- [67] R. Thoppilan, D. De Freitas, J. Hall, N. Shazeer, A. Kulshreshtha, H.-T. Cheng, A. Jin, T. Bos, L. Baker, Y. Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.
- [68] S. Tu, A. Robey, T. Zhang, and N. Matni. On the sample complexity of stability constrained imitation learning. In *Learning for Dynamics and Control Conference*, pages 180–191. PMLR, 2022.
- [69] R. Van Handel. Probability in high dimension. Technical report, PRINCETON UNIV NJ, 2014.
- [70] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [71] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [72] C. Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.
- [73] C. Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [74] P. Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [75] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- [76] T. Westenbroek, M. Simchowitz, M. I. Jordan, and S. S. Sastry. On the stability of nonlinear receding horizon control: a geometric perspective. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 742–749. IEEE, 2021.
- [77] Y. Wu and K. He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [78] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5628–5635. IEEE, 2018.
- [79] S. Zhao and A. Grover. Decision stacks: Flexible reinforcement learning via modular generative models. *arXiv preprint arXiv:2306.06253*, 2023.
- [80] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.

Contents

| | | |
|-----------|--|-----------|
| 1 | Introduction | 1 |
| 1.1 | Contributions. | 2 |
| 2 | Setting | 3 |
| 3 | Conditional sampling with stabilization suffices for behavior cloning | 5 |
| 3.1 | Incremental Stability and the Synthesis Oracle. | 6 |
| 3.2 | Simplified guarantees under total variation continuity | 7 |
| 3.3 | A general guarantee via data noising. | 8 |
| 4 | HINT: Instantiating Data Noising with DDPMs | 9 |
| 4.1 | Experimental Results | 9 |
| 5 | Discussion | 10 |
| I | General Results and Discussion | 18 |
| A | Notation and Organization of Appendix | 18 |
| A.1 | Notation Summary | 18 |
| A.2 | Organization of the Appendix | 19 |
| B | Complete Related Work | 19 |
| C | Further Main Results and Discussion | 21 |
| C.1 | Discussion surround Theorem 1 | 21 |
| C.2 | Proof sketch of Theorem 2 | 22 |
| C.3 | Merits and Drawbacks of the Synthesis Oracle | 22 |
| C.4 | Formal Assumptions for Analysis of HINT | 23 |
| C.5 | Theoretical Guarantee for HINT. | 24 |
| D | Analysis Overview | 24 |
| D.1 | Structure of the proof. | 24 |
| D.2 | Instantiation for control | 26 |
| D.3 | A Guarantee in the Composite MDP, and the derivation of Theorem 3 | 26 |
| D.4 | Proof Overview of Theorem 4 | 27 |
| E | Generalization to Generic Incrementally Stable Primitive Controllers | 29 |
| E.1 | Main Results for General Primitive Controllers | 30 |
| E.2 | Comparison to prior notions of stability. | 31 |
| II | Composite MDP | 34 |

| | | |
|------------|---|-----------|
| F | Measure-Theoretic Background | 34 |
| F.1 | Kernels, Regular Conditional Probabilities and Gluing | 35 |
| F.2 | Optimal Transport and Kernel Couplings | 36 |
| F.3 | Data Processing Inequalities | 37 |
| F.3.1 | Deferred lemmas for the data processing inequalities | 38 |
| F.4 | Proof of Proposition F.3 | 38 |
| F.5 | A simple union-bound recursion. | 40 |
| G | Warmup: Analysis Without Augmentation | 40 |
| G.1 | Relaxing the TVC Condition | 42 |
| H | Imitation in the Composite MDP | 43 |
| H.1 | A generalization of Theorem 4 | 43 |
| H.2 | Proof of Theorem 8 | 45 |
| H.2.1 | Proof Overview and Coupling Construction | 45 |
| H.2.2 | Properties of smoothing, deconvolution, and replicas. | 47 |
| H.2.3 | Bounding the marginal imitation gaps in terms of the teleporting sequence error | 49 |
| H.2.4 | Formal proof of Theorem 8 | 49 |
| H.2.5 | Proof of Lemma H.10 | 53 |
| H.3 | Proof of Theorem 4, and generalization to direct decompositions | 55 |
| I | Lower Bounds | 56 |
| I.1 | Sharpness of Proposition D.2 and Theorem 4 | 57 |
| I.2 | π_{rep}^* and π^* induce the same marginals but different joint distributions, even with memoryless dynamics | 59 |
| I.3 | π_{rep}^* and π^* can have different marginals, implying necessity of $\gamma_{\text{IPS}, \mathcal{S}}$ | 60 |
| I.4 | π_{rep}^* and π_{dec}^* have different marginals, even with memoryless dynamics | 61 |
| III | The Control Setting | 63 |
| J | End-to-end Guarantees and the Proof of Theorem 3 | 63 |
| J.1 | Preliminaries | 63 |
| J.1.1 | Direct Decomposition and Smoothing Kernel. | 64 |
| J.1.2 | Chunking Policies. | 64 |
| J.1.3 | Preliminaries for joint-distribution imitation. | 65 |
| J.2 | Translating Control Imitation Losses to Composite-MDP Imitation Gaps | 65 |
| J.3 | Proofs of main results | 66 |
| J.3.1 | Proof of Theorem 1 | 66 |
| J.3.2 | Proof of Theorem 2 and a more precise statement. | 66 |
| J.3.3 | Proof of Theorem 10 | 67 |

| | | |
|----------|---|------------|
| J.3.4 | Proof of Theorem 3 | 68 |
| J.4 | Imitation of the joint trajectory under total variation continuity of demonstrator policy | 69 |
| J.4.1 | Proof of Theorem 11 | 70 |
| J.4.2 | Proof of Proposition J.2 | 71 |
| J.5 | Imitation in total variation distance | 72 |
| J.6 | Consequence for expected costs | 73 |
| J.7 | Useful Lemmata | 74 |
| J.7.1 | On the trajectories induced by π^* from \mathcal{D}_{exp} | 74 |
| J.7.2 | Concentration and TVC of Gaussian Smoothing. | 75 |
| J.7.3 | Total Variation Telescoping | 76 |
| K | Stability in the Control System | 76 |
| K.1 | Proof of Proposition K.1 | 77 |
| K.1.1 | Proof of Lemma K.3 | 80 |
| K.2 | Proof of Lemma K.2 | 80 |
| K.3 | Synthesized Linear Controllers are Incrementally Stabilizing | 81 |
| K.4 | Proof of Proposition K.6 (incremental stability of affine gains) | 82 |
| K.5 | Proof of Lemma K.7 (state perturbation) | 84 |
| K.6 | Ricatti synthesis of stabilizing gains. | 85 |
| K.6.1 | Proof of Proposition K.10 (Ricatti synthesis of gains) | 86 |
| K.7 | Solutions to recursions | 88 |
| L | Sampling and Score Matching | 90 |
| L.1 | Simplifying Notation | 91 |
| L.2 | Denoising Diffusion Probabilistic Models | 91 |
| L.3 | Score Estimation | 93 |
| M | Proofs for Generic Incrementally Stable Primitive Controllers | 97 |
| M.1 | Generalization of analysis in the composite MDP | 97 |
| M.2 | State-Conditioned Input-Stability and IPS in the Composite MDP via t-ISS | 100 |
| M.3 | Concluding the proof of Theorems 5 and 6 | 101 |
| M.3.1 | Proof of Theorem 6 | 101 |
| M.3.2 | Proof of Theorem 5 | 101 |
| N | Extensions and Further Results | 103 |
| N.1 | Removing the necessity for minimal chunk length via stronger synthesis oracle | 103 |
| N.2 | Noisy Dynamics | 104 |
| N.3 | Robustness to Adversarial Perturbations | 104 |
| N.4 | Deconvolution Policies and Total Variation Continuity | 105 |
| O | Experiment Details | 107 |

| | |
|--|-----|
| O.1 Compute and Codebase Details | 107 |
| O.2 Environment Details | 108 |
| O.3 Gain Synthesis | 109 |
| O.4 DDPM Model and Training Details. | 109 |

Part I

General Results and Discussion

A Notation and Organization of Appendix

In this appendix, we collect the notation we use throughout the paper, as well as providing a high level organization of the appendices.

A.1 Notation Summary

In this section, we summarize some of the notation used throughout the work, divided by subject.

Measure Theory We always let \mathcal{X} denote a Polish space, $\mathcal{B}(\mathcal{X})$ the Borel-algebra on \mathcal{X} , and $\Delta(\mathcal{X})$ the set of borel probability measures on \mathcal{X} . For a random variable X on \mathcal{X} , we let P_X denote the law of X . For random variables X, Y , we let $\mathcal{C}(P_X, P_Y)$ denote the set of couplings of these measures and for laws P_1, P_2 . We write $P_1 \otimes P_2$ for the product measure. We will generally reserve P to denote measure, Q and W for probability kernels, and μ for a joint measure on several random variables.

When $P_1, P_2 \in \Delta(\mathcal{X})$ are laws on the same space, we let $\text{TV}(P_1, P_2)$ denote the total variation distance. We write $P_1 \ll P_2$ if P_1 is absolutely continuous with respect to P_2 . Given a Polish space \mathcal{X} and element $x \in \mathcal{X}$, we let $\delta_x \in \Delta(\mathcal{X})$ denote the dirac-delta measure supported on the set $\{x\} \in \mathcal{B}(\mathcal{X})$ (note that, in a Polish space, the singleton $\{x\}$ set is closed, and therefore Borel).

Norms and linear algebra notation. We use bold lower case vector \mathbf{z} to denote vectors, and bold upper case \mathbf{Z} to denote matrices. We let $\mathbf{z}_{1:K} = (\mathbf{z}_1, \dots, \mathbf{z}_K)$ and $\mathbf{Z}_{1:K} = (\mathbf{Z}_1, \dots, \mathbf{Z}_K)$ denote concatenations. The norms $\|\cdot\|$ denote Euclidean norms on vectors and operator norms on matrices. We identify the spaces \mathcal{P}_k with Euclidean vectors in the standard sense. Given a Euclidean vector $\mathbf{z} \in \mathbb{R}^d$, $\mathcal{N}(\mathbf{z}, \sigma^2 \mathbf{I})$ denote the multivariate normal distribution on \mathbb{R}^d with covariance $\sigma^2 \mathbf{I}$.

Control notation. We let $\mathbf{x}_t \in \mathbb{R}^{d_x}$ denote control states, $\mathbf{u}_t \in \mathbb{R}^{d_u}$ denote control inputs, and $\boldsymbol{\rho}_\tau \in \mathcal{P}_\tau$ denotes trajectories $(\mathbf{x}_{1:\tau+1}, \mathbf{u}_{1:\tau})$. T denotes the time horizon of imitation, so $\boldsymbol{\rho}_T \sim \mathcal{P}_T$. Our dynamics are $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$; for our main results (Section 3), we suppose $f(\mathbf{x}, \mathbf{u}) = \mathbf{x} + \eta f_\eta(\mathbf{x}, \mathbf{u})$, parametrizing dynamics in the form of an Euler discretization with step $\eta > 0$.

Recall that primitive controllers κ take the form $\kappa(\mathbf{x}) = \bar{\mathbf{K}}(\mathbf{x} - \bar{\mathbf{x}}) + \bar{\mathbf{u}}$, where terms with $(\bar{\cdot})$, $\bar{\mathbf{K}}, \bar{\mathbf{x}}, \bar{\mathbf{u}}$, denote parameters of the primitive controller. The space of these is \mathcal{K} .

We also recall the chunk-length τ_{chunk} and observation length τ_{obs} satisfying $0 \leq \tau_{\text{obs}} \leq \tau_{\text{chunk}}$. We recall the definition of the trajectory-chunk s_h and observation-chunk in \mathbf{o}_h in Section 2, which introduced the indexing h , such that $t_h = (h-1)\tau_{\text{chunk}} + 1$. Recall also the composite actions $\mathbf{a}_h = (\kappa_{t_h:t_{h+1}-1}) \in \mathcal{A} = \mathcal{K}^{\tau_{\text{chunk}}}$ as the concatenation of τ_{chunk} primitive controllers.

Abstractions in the composite MDP. The composite MDP is a deterministic MDP with composite-states $\mathbf{s} \in \mathcal{S}$ and composite-actions $\mathbf{a} \in \mathcal{A}$, and (possibly time-varying) deterministic transition dynamics $F_h : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ for $1 \leq h \leq H$. The goal is to imitate a policy $\pi^* = (\pi_h^*)_{1 \leq h \leq H}$, in terms of imitation gaps $\Gamma_{\text{joint}, \varepsilon}$ and $\Gamma_{\text{marg}, \varepsilon}$ defined in Definition D.1. We refer the reader to Appendix D for the relevant terminology, and to Appendix D.2 for its instantiation in our original control setting.

A.2 Organization of the Appendix

We now describe the organization of our many appendices. In [Appendix E](#), we generalize some of our results to accommodate general incrementally stabilizing primitive controllers. In [Appendix E.2](#), we expand on our abbreviated discussion of related work in the body as well as provide a more detailed comparison of our notion of stability [Definition D.5](#) with those found in prior work.

After the preliminaries on organization, notation, and related work, we divide our appendices into three parts.

Part I: General Results and Discussion In this section, we describe major results and discussion omitted from the main body in the interest of space. This present appendix contains notation and organization. Subsequently, [Appendix B](#) provides a comprehensive discussion of related work. [Appendix C](#) provides discussion for [Theorem 1](#), a proof sketch of [Theorem 2](#), and the requisite assumptions and formal statement of our guarantee for HINT. [Appendix D](#) provides a detailed overview of the analysis, and [Appendix E](#) extends our results from affine primitive controllers to general ones.

Part II: The Composite MDP. In the first part of the Appendix, we expand on and provide rigorous proofs of statements and results pertaining to the composite MDP as considered in [Appendix D](#). We begin by providing a detailed background in [Appendix F](#) on the requisite measure theory we use to make our arguments rigorous. In particular, we provide definitions of probability kernels and couplings, as well as measurability properties of optimal transport couplings. In [Appendix G](#), we provide the full proof of [Proposition D.2](#), as warm-up to the proof of [Theorem 4](#). In particular, the argument substantially simplifies if we consider the case of no added augmentation (when $\sigma = 0$ in HINT) and we present a coupling construction that implies the analogous bound in the presence of an additional assumption. The heart of the first part of our appendices is [Appendix H](#), where we rigorously prove a generalization of [Theorem 4](#) by constructing a sophisticated coupling between the imitator and demonstrator trajectories. We conclude the first part of our appendices by proving a number of lower bounds in the composite MDP setting in [Appendix I](#), which demonstrate the tightness of our arguments in [Appendix H](#).

Part III: Instantiations We continue our appendices in the second part, which is concerned with the instantiation of the composite MDP in with incremental stability. The heart of the second part of our appendices is [Appendix J](#), which provides the final, end-to-end guarantees and the proof of [Theorem 1](#) and [Theorems 2](#) and [3](#). We also provide a number of variations on this result, including stronger guarantees on imitation of the joint trajectories ([Appendix J.4](#)), guarantees on HINT under the assumption that sampling is close in total variation ([Appendix J.5](#)). We also show in [Proposition J.5](#) that most natural cost functions have similar expected values on imitator and demonstrator trajectories assuming that the imitation losses are small. In [Appendix K](#), we provide a detailed proof that the control setting considered in [Section 2](#) satisfies the stability properties required by our analysis of the composite MDP and prove [Proposition D.1](#), and [Appendix K.6](#) in particular explains how to synthesize stabilizing gains, as assumed in [Section 2](#). With the stability properties thus proven, we proceed in [Appendix L](#) to instantiate our conditional sampling guarantees with DDPMs. In particular, by applying earlier work, we state and prove [Theorem 13](#), which guarantees that with sufficiently many samples, in our setting we can ensure that the learned DDPM provides samples close in the relevant optimal transport distance to the expert distribution. We also explain in [Remark L.4](#) why stronger total variation guarantees on sampling are unrealistic in our setting.

[Appendix M](#) generalizes the proofs in [Appendix J](#) to the generic primitive controllers considered in [Appendix E](#). We then provide a number of extensions of our main results in [Appendix N](#), including to the setting of noisy dynamics ([Appendix N.2](#)). Finally, in [Appendix O](#), we expand the discussion of our experiments, including training and compute details, environment details, and a link to our code for the purpose of reproducibility.

B Complete Related Work

Imitation Learning. Over the past few years, there has been a significant surge of interest in utilizing machine learning techniques for the execution of exceedingly intricate manipulation and control

tasks. Imitation learning, whereby a policy is trained to mimic expert demonstrations, has emerged as a highly data efficient and effective method in this domain, with application to self-driving vehicles [30, 15, 10], visuomotor policies [23, 78], and navigation tasks [31]. A widely acknowledged challenge of imitation learning is distribution shift: since the training and test time distributions are induced by the expert and trained policies respectively, compounding errors in imitating the expert at test-time can lead the trained policy to explore out-of-distribution states [57]. This distribution shift has been shown to result in the imitator making incorrect judgements regarding observation-action causality, often with catastrophic consequences [21]. Prior work in this domain has predominantly attempted to mitigate this issue in the non-stochastic setting via online data augmentation strategies, sampling new trajectories to mitigate distribution shift [58, 57, 40]. Among this class of methods, the DAgger algorithm in particular has seen widespread adoption [57, 65, 37]. These approaches have the drawback that sampling new trajectories or performing queries on the expert is often expensive or intractable. Due to these limitations, recent developments have focused on novel algorithms and theoretical guarantees for imitation learning in an offline, non-interactive environment [16, 50]. Our work is similarly focused on the offline setting, but is capable of handling stochastic, non-Markovian demonstrators. Unlike [50], we do not require our expert demonstrations to be sampled from a stabilizing expert policy, instead utilizing a synthesis oracle to stabilize around the provided demonstrations. This is a significantly weaker requirement and enables the development of high-probability guarantees for human demonstrators, where sampling new trajectories and reasoning about the stability properties is not possible.

Denosing Diffusion Probabilistic Models and other Generative Approaches. Denosing Diffusion Probabilistic Models (DDPMs) [60, 29] and their variant, Annealed Langevin Sampling [62], have seen enormous empirical success in recent years, especially in state-of-the-art image generation [55, 47, 61]. More relevant to this paper is their application to imitation learning, where they have seen success even without the proposed data augmentation in Janner et al. [33], Chi et al. [19], Pearce et al. [48], Hansen-Estruch et al. [26]. DDPMs rely on learning the score function of the target distribution, which is generally accomplished through some kind of denoised estimation [32, 74, 63]. On the theoretical end, annealed Langevin sampling has been studied with score estimators under a variety of assumptions including the manifold hypothesis and some form of dissipativity [53, 13, 14], although these works have generally suffered from an exponential dependence on ambient dimension, which is unacceptable in our setting. Of greatest relevance to the present paper are the concurrent works of Chen et al. [18], Lee et al. [41] that provide polynomial guarantees on the quality of sampling using a DDPM assuming that the score functions are close in an appropriate mean squared error sense. We take advantage of these latter two works in order to provide concrete end-to-end bounds in our setting of interest. To our knowledge, ours is the first work to consider the application of DDPMs to imitation learning under a rigorous theoretical framework, although we emphasize that this does not constitute a strong technical contribution as opposed to an instantiation of earlier work for the sake of completeness and concreteness.

Recent work has also shown that transformer architectures [80, 17, 59] can also serve as probabilistic models for predicting sequences of robotic actions, and can represent multi-modality to varying degrees. Notably, these approaches also rely heavily on the action-chunking which we consider in this paper [80].

Hierarchical Planning. Hierarchy has long been applied in robotic learning and planning to abstract away low-level primitives. Task-and-motion planning (TAMP) [35] reduces robotic motion represents planning via sequences of discrete primitives — a “mode sequence” — constrained to which the optimization problem is continuous. LQR trees [66] proposes using linear quadratic regulator (LQR) trees to efficiently cover a control state space with local feedback laws so as to compute a motion plan that reaches a desired goal or behavior, subject to stability guarantees. Graph of Convex Sets (GCS), a more recent innovation, decomposes constraint sets into convex regions, each of which represents nodes in a planning graph [43]. More recent work has used hierarchy to leverage the power of large learned models for solving tasks that contain multiple types of data inputs⁴ [4], and as modules for generating multiple forms of supervision [79].

⁴These tasks are also called multimodal, where modes here refer to types of data. In distinction, multi-modality in this paper refers to multiple modes within the distribution over expert demonstrations

Smoothing Augmentations. Data augmentation with smoothing noise has become such common practice, its adoption is essentially folklore. While augmentation of actions which noise is common practice for exploration (see, e.g. [40]), it is widely accepted that noising actions in the learned policy is not best practice, and thus it is more common to add noise to the *states* at training time, preserving target actions as fixed [36]. Our work gives an interpretation of this decision as enforcing that the learned policy obey the distributional continuity property we term TVC (Definition D.3), so that the policy selects similar actions on nearby states. Previous work has interpreted noise augmentation as providing robustness. Data augmentation has been demonstrated to provide more robustness in RL from pixels [39], adaptive meta-learning [3], in more traditional supervised learning as well [28].

C Further Main Results and Discussion

We begin with a general remark.

Remark C.1 (Do we need state observation or time-varying policies?). In practical applications, behavior cloning policies respond not to measurements of physical system state, but rather visual observations and/or tactile feedback. Additionally, learned policies do not explicitly take a time index h as input. This allows these policies to perform flexibly across tasks with varying time horizons, and to automatically reset after encountered obstacles.

In contrast, our formulation requires policies to be (a) functions of system state and (b) vary with the chunk index h . If visual observations or tactile measurements are *sufficient* to recover system state, then we can view these data as redundant states, and thus (a) is not a restriction. Moreover, as described in Remark C.6, a limited portion of theoretical results do hold for policies which do not vary with h . In general, restrictions (a) and (b) are necessary for our analysis because they allow us to analyze the imitator behavior in a Markovian fashion. Without these restrictions, one would have to reason about (a') uncertainty over state given observation, or (b') variation in expert behavior across different time steps h . Removing these restrictions is an exciting direction for future work.

C.1 Discussion surround Theorem 1

Remark C.2 (The $\varepsilon = 0$ case). If we were able to bound the policy error $\Delta_{(\varepsilon)}$ with $\varepsilon = 0$ — which corresponds to estimating $a_h \mid o_h$ in *total variation* distance — the imitation learning problem would be trivialized, and neither the TVC condition above or the noise-injection based smoothing in the section below would not be needed (see Appendix J). Appendix L explains that the needed assumptions for this stronger sense of approximate sampling do not hold in our setting, because expert distributions over actions typically lie on low-dimensional manifolds.

Remark C.3 (On the TVC assumption). It is true that any $\hat{\pi}$ implemented as a DDPM with a Lipschitz activation with bounded-magnitude parameters is indeed TVC. Unfortunately, these Lipschitz constants can be too large to be meaningful in practical scenarios, scaling exponentially with network depth. In addition, the absence of smoothing σ may make the corresponding DDPM learning problem more challenging. Hence, in what follows, we shall require the additional sophistication of smoothing with Gaussian noise of variance $\sigma^2 > 0$ for meaningful guarantees.

Furthermore, we show in Appendix G.1 that the TVC assumption, which measures total variation distance between nearby $\hat{\pi}_h(\cdot)$ at nearby observations, can be relaxed to variant which measures the probability (under a minimal coupling) that actions differ by some tolerance. However, this tolerance has to be quite small, and as we argue, any reasonable notion of Wasserstein continuity is unlikely to suffice.

Remark C.4 (Imitation of the joint distribution). Suppose the expert distribution \mathcal{D}_{exp} has at most τ_{obs} -bounded memory (defined formally in Definition J.5). Then $\mathcal{L}_{\text{joint},\varepsilon}(\hat{\pi})$ satisfies the same upper bound (E.2), where $\mathcal{L}_{\text{joint},\varepsilon}(\hat{\pi})$, formally defined in Definition J.4, measures an optimal transport distance between the *joint distribution* of the expert trajectory and the one induced by $\hat{\pi}$.

Remark C.5 (Is chunking necessary?). In Appendix N.1, we show that we can remove the required lower bound on τ_{chunk} — allowing, in particular, the choice of $\tau_{\text{chunk}} = 1$ — under the slightly stronger condition that our synthesis oracle ensures that the entire sequence of primitive controllers $\kappa_{1:T}$ on the whole horizon T are incrementally stabilizing. However, chunking is known to yield empirical benefits [80], and training models to predict action-chunks of longer duration than the agent acts on is also observed to improve performance [19].

Remark C.6 (Are time-varying policies necessary? (continuing Remark C.1)). In practice, time-invariant policies $\hat{\pi}$ which do depend on the h -index are preferred because they are more resilient to varying-horizon tasks, and can automatically “reset” when they encounter an obstacle. Here, we note that if $\pi_h^*(o_h)$, the conditional distribution of π_h^* given o_h , is independent of h — that is, the expert is Markov and time-invariant given o_h — then the term $\Delta_{(\varepsilon/c_1)}(\pi_h^*(o_h), \hat{\pi}_h(o_h))$ on the right-hand side of (E.2) can be made small by choosing a time-invariant $\hat{\pi}$. Thus, certain expert behavior can indeed be imitated by time-invariant policies. However, we do require time-varying policies to imitate *arbitrary experts*. And, in addition, the data smoothing strategy described below requires a time-varying $\hat{\pi}$. Extending our results to time-invariant $\hat{\pi}$ is an interesting direction for future inquiry, and we suspect that this may require some further notion of cost-to-go to made the formulation feasible.

C.2 Proof sketch of Theorem 2

As with Theorem 1, the key ideas of the proof are given Appendix D, expressed in terms of a general abstraction for behavior cloning we call the “composite MDP”. This template is instantiated with a details in Appendix J. Moreso than Theorem 1, the proof of Theorem 2 requires sophisticated couplings between expert and learner trajectories, and in particular. The intuition is based on the observation that $\hat{\pi}_{\sigma,h}$ mimic $\pi_{\text{rep},\sigma,h}^* := (\pi_{\text{dec},\sigma,h}^*)_{\sigma}$, the smoothing of the deconvolution policy. Inspired by replica pairs in statistical physics, we call $\pi_{\text{rep},\sigma}^*$ the “replica” policy because actions from $\pi_{\text{rep},\sigma,h}^*$ can be thought of as actions from π_h^* that have been noised and deconvolved. This implies:

Fact C.1. Let $o_h \sim \mathcal{D}_{\text{exp},h}$. Then, the distributions of $a_h \sim \pi_h^*(o_h)$, and $a'_h \sim \pi_{\text{rep},\sigma,h}^*(o_h)$, marginalized over o_h , are identical.

This observation can be interpreted as meaning that smoothing and deconvolution are inverse operations at the distributional level. In particular, for a moment, consider an idealized environment where a_h , and not o_h , perfectly determined the dynamics (e.g. by teleportation), then $\pi_{\text{rep},\sigma}^* = (\pi_{\text{rep},\sigma,h}^*)$ and $\pi^* = (\pi_h^*)$ would induce the same dynamics (and, as remarked above, $\mathcal{L}_{\text{marg},\varepsilon}(\pi^*) = 0$).

For non-idealized environments, the argument goes as follows: We couple $a_h \sim \pi_h^*(o_h)$, and $a'_h \sim \pi_{\text{rep},\sigma,h}^*(o_h)$ so that a'_h has the distribution of $a'_h \sim \pi^*(o'_h)$, where o'_h is distributed as o_h and, with high probability, o'_h and o_h are $\tilde{O}(\sigma)$ -close in Euclidean norm. This coupling is depicted Figure 4. We then argue that the dynamics induced by $\pi_{\text{rep},\sigma}^*$ track those of π^* by roughly a similar margin. Moreover, by smoothing $\hat{\pi}_h$ and $\pi_{\text{dec},h}^*$ and applying Jensen’s inequality,

$$\mathbb{E}_{o_h \sim \mathcal{D}_{\text{exp},h}} [\Delta_{(\varepsilon^2)}(\pi_{\text{rep},\sigma,h}^*(o_h), \hat{\pi}_{\sigma,h}(o_h))] \leq \mathbb{E}_{\tilde{o}_h \sim \mathcal{D}_{\text{exp},\sigma,h}} [\Delta_{(\varepsilon^2)}(\pi_{\text{dec},\sigma,h}^*(\tilde{o}_h), \hat{\pi}_h(\tilde{o}_h))].$$

Consequently, when the right hand side of the above inequality is small, the noised policy $\hat{\pi}_{\sigma}$ tracks the replica policy $\pi_{\text{rep},\sigma}^*$, which we have shown to be track π^* (and thus track \mathcal{D}_{exp}). \square

C.3 Merits and Drawbacks of the Synthesis Oracle

The role of a synthesis oracle satisfying Assumption 3.1 is to replace a strong assumption on the stability of an *expert demonstration* with an *algorithmic assumption* that allows post-hoc stabilizing of expert demonstrations. This approach presents a natural question: in what sense is this tradeoff a sensible one? To answer this, consider a paradigmatic case, where the demonstrations solve some complicated task in a smooth, nonlinear control system. Suppose further that the one-step dynamics of the system are known, but that the expert demonstrations come from some optimal control law which is computationally prohibitive to compute, or possibly even some mixture of different, mutually-incompatible trajectories. Assuming the Jacobian linearizations of the nonlinear system are stabilizable (see Appendix K for further details), one can implement a synthesis oracle for affine gains directly by solving a Ricatti recursion on the Jacobian-linearized dynamics around each expert trajectory. . Conceptually, this has the following interpretation: **Our framework reduces the problem of imitating a complex expert trajectory to (i) supervised generative modeling and (ii) solving strictly local control problems.**

That is, we offload complex behavior of the expert being imitated, and reduce the learner’s burden to solving local control problems that are significantly simpler than global planning. For more general systems, Appendix E addresses possibly non-affine stabilizing gains, and discusses how these may

arise from standard practices of using robotic position control or inverse dynamics. [Appendix E.2](#) compares our hierarchical approach to stability to standard formulations that apply to the expert distribution, and we show how the latter rule out the possibility for complex behaviors such as bifurcated trajectories.

Limitations and Future Directions. Our above example required access to differentiable (indeed, smooth) system dynamics. Stabilizing systems with contact dynamics remains an outstanding challenge. More generally, an overtly hierarchical approach may be inefficient for many reasons, notably (1) the dimension of the primitive controller may be much higher than the dimension of raw control inputs; and (2) when the high-level and low-level controllers are parametrized by the neural networks, explicit hierarchy with separate models may preclude shared representation learning. Developing a more comprehensive approach to stability (perhaps one that does not require explicit gain synthesis, and extends to non-smooth systems) is an exciting direction for future work. Nevertheless, we think that **our conceptual contribution of decoupling low-level stability and generative matching of demonstrated behavior will prove useful in future endeavors for reliable and performant behavior cloning.**

C.4 Formal Assumptions for Analysis of HINT

We now state the assumptions required for our theoretical guarantees on HINT. We require access to a class of score functions rich enough to represent the following deconvolution conditionals. To define these, we introduce the following distribution

Definition C.1. Recall the policy $\pi_{\text{dec},\sigma,h}^*$ from [Definition 3.6](#). Given $\mathbf{o}_h \in \mathcal{O}$, let $\pi_{\text{dec},\sigma,h,[t]}^*(\mathbf{o}_h) \in \Delta(\mathcal{A})$ denote the law of $\mathbf{a}_{h,[t]} := e^{-t}\mathbf{a}_h^{(0)} + \sqrt{1 - e^{-2t}}\boldsymbol{\gamma}$, where $\mathbf{a}_h \sim \pi_{\text{dec},\sigma,h}^*(\mathbf{o}_h)$, and $\boldsymbol{\gamma} \sim \mathcal{N}(0, \mathbf{I})$ is independent white Gaussian noise. In words, $\mathbf{a}_{h,[t]}$ is generated from the Ornstein–Uhlenbeck process that interpolates \mathbf{a}_h with white noise. To avoid notational clutter, we suppress the dependence of $\pi_{\text{dec},\sigma,h,[t]}^*$ on σ via $\pi_{\text{dec},h,[t]}^* = \pi_{\text{dec},\sigma,h,[t]}^*$ when σ is clear from context.

Next, we need an assumption of bounded statistical complexity. We opt for the popular *Rademacher complexity* [11]. In defining this quantity, recall that scores are vector-valued, necessitating the vector analogue of Rademacher complexity [44, 24], studied for score matching in [13].

Definition C.2 (Function Class Θ and Rademacher Complexity). Consider a class of score functions of the form $\Theta = \{(\mathbf{s}_{\theta,h})_{1 \leq h \leq H} | \theta \in \Theta\}$, where $\mathbf{s}_{\theta,h}$ maps triples $(\mathbf{a}, \mathbf{o}_h, j)$ of composite actions $\mathbf{a} \in \mathcal{A}$, observation-chunks \mathbf{o}_h , and DDPM-steps $j \in \mathbb{N}$ to vectors in \mathbb{R}^{d_A} . For each chunk $h \in [H]$, DDPM-step $j \in \mathbb{N}$ and discretization size α , define the vector- Rademacher complexity of Θ as

$$\mathcal{R}_{n,h,j}(\Theta; \alpha) := \mathbb{E} \left[\sup_{\theta \in \Theta} \frac{1}{n} \sum_{k=1}^n \left\langle \boldsymbol{\varepsilon}^{(k)}, \mathbf{s}_{\theta,h} \left(\mathbf{a}_{h,[j\alpha]}^{(k)}, \tilde{\mathbf{o}}_h^{(k)}, j \right) \right\rangle \right],$$

where $\boldsymbol{\varepsilon}_k \in \mathbb{R}^d$ are i.i.d. random vectors with Rademacher coordinates, and where $(\tilde{\mathbf{o}}_h^{(k)}, \mathbf{a}_{h,[j\alpha]}^{(k)})$ are i.i.d. samples from $\mathcal{D}_{\sigma,h,[j\alpha]}$.

Assumption C.1. We suppose that, for any $\sigma > 0$, we are given a class of score functions $\Theta = \Theta(\tau_{\text{chunk}}, \tau_{\text{obs}}, \sigma, \alpha)$ of the form in [Definition C.2](#) which satisfies the following conditions:

- (a) *Realizability*: there exists a θ_* such that, for all $h \in [H]$ and $j \in \mathbb{N}$, $\mathbf{s}_{\theta_*,h}(\mathbf{a}, \mathbf{o}_h, j)$ is the score function of $\pi_{\text{dec},h,[j\alpha]}^*(\mathbf{o}_h)$ at $\mathbf{a} \in \mathcal{A}$.
- (b) *The Rademacher complexity of Θ has polynomial decay in n and growth in α* :

$$\sup_{j \in \mathbb{N}} \max_{h \in [H]} \mathcal{R}_{n,h,j}(\Theta; \alpha) \leq C_{\Theta} \alpha^{-1} n^{-\frac{1}{\nu}},$$

where $C_{\Theta} = C_{\Theta}(\sigma, \tau_{\text{chunk}}, \tau_{\text{obs}}) > 0$.

- (c) *The scores have linear growth*; that is, there exists some $C_{\text{grow}} = C_{\text{grow}}(\sigma, \tau_{\text{chunk}}, \tau_{\text{obs}}) > 0$ such that

$$\sup_{j \in \mathbb{N}, h \in [H]} \sup_{\theta \in \Theta} \|\mathbf{s}_{\theta}(\mathbf{a}, \tilde{\mathbf{o}}_h, j)\| \leq C_{\text{grow}} \alpha^{-1} (1 + \|\mathbf{a}\| + \|\tilde{\mathbf{o}}_h\|),$$

As discussed in [Appendix L](#), generalizing to polynomial growth is straightforward.

As justified in [Appendix L](#), our decay condition on the Rademacher complexity is natural for statistical learning, and holds for most common function classes (often with $\nu \leq 2$ and even more benign dependence on J, α); our results can easily extend to approximate realizability as well. Our Rademacher bound depends implicitly on chunk and observation lengths $\tau_{\text{chunk}}, \tau_{\text{obs}} > 0$ and implicitly on dimension $d_{\mathcal{A}}$ via C_{Θ} . Realizability is motivated by the approximation power of deep neural networks [12]. Lastly, we do expect realizability to hold uniformly over $j \geq 0$ because, as $j \rightarrow 0$, the corresponding scores corresponds to a scaled identity function (i.e. the score of a standard Gaussian).

C.5 Theoretical Guarantee for HINT.

We now state our guarantee for HINT, invoking assumptions from the section above. Recall that $d_{\mathcal{A}}$ denotes the dimension of composite actions, and that c_1, \dots, c_5 are as in [Definition 3.2](#).

Theorem 3. Suppose [Assumption 3.1](#) holds. Let $c_1, \dots, c_5 > 0$, defined in [Definition 3.2](#), and let $\Theta_{\text{Iss}}(x)$ denote a term which is upper and lower bounded by a x times a polynomial in those constants and their inverses. Let $\varepsilon \leq \Theta_{\text{Iss}}(1)$, if we choose $\sigma = \varepsilon / \Theta_{\text{Iss}}(\sqrt{d_x} + \log(1/\varepsilon))$ and let $\tau_{\text{chunk}} \leq c_3$ and $\tau_{\text{chunk}} - \tau_{\text{obs}} \geq \frac{1}{L_{\beta}} \log(c_1/\varepsilon)$. Consider running HINT for $\sigma > 0$ with parameters J, α polynomial in the parameters given in [Assumption 3.1](#) specified in [Appendix L](#). Then, Then, if

$$N_{\text{exp}} \geq \text{poly}(C_{\Theta}(\sigma, \tau_{\text{chunk}}, \tau_{\text{obs}}), 1/\varepsilon, R_{\mathbf{K}}, d_{\mathcal{A}}, \log(1/\delta))^{\nu} > 0,$$

then with probability $1 - \delta$, the policy $\hat{\pi}_{\sigma}$ returned by HINT satisfies

$$\mathcal{L}_{\text{marg}, \varepsilon}(\hat{\pi}_{\sigma}) \leq \Theta_{\text{Iss}}\left(\varepsilon H \sqrt{\tau_{\text{obs}}} \cdot (\sqrt{d_x} + \log(1/\varepsilon))\right).$$

In addition, consider running HINT with $\sigma = 0$, and suppose $C_{\Theta}(\sigma, \tau_{\text{chunk}}, \tau_{\text{obs}})|_{\sigma=0}$ is finite. Then, for N_{exp} satisfying the same bound as above, it holds that with probability $1 - \delta$, the policy $\hat{\pi}$ produced by HINT satisfies the guarantees of [Theorem 1](#) up to an additive factor of $H\varepsilon$ on the event that $\hat{\pi}$ happens to be γ -TVC.

[Theorem 3](#) instantiates [Theorem 2](#) by bounding the policy error terms $\Delta_{(\varepsilon^2)}(\pi_{\text{dec}, \sigma, h}^*(\tilde{o}_h), \hat{\pi}_h(\tilde{o}_h))$ in that theorem when $\hat{\pi}$ is the policy learned by fitting the DDPM in [Line 7](#). The formal bound on N_{exp} is given in (J.6) in [Appendix J](#). While making τ_{chunk} larger appears to improve the above bound, it generally increases the statistical and computational challenge of learning the DDPM itself. The guarantees for score matching are derived in [Appendix L](#) by applying [18, 41, 13]; these are applied to [Theorem 2](#) in [Appendix J](#).

D Analysis Overview

Our analysis abstracts away the vector-valued dynamics into a deterministic MDP – the *composite MDP* – with *composite-states* $s \in \mathcal{S}$ and *composite-actions* $a \in \mathcal{A}$, corresponding to trajectory-chunks and composite-actions in [Section 2](#). We abstract away our dynamics as

$$s_{h+1} = F_h(s_h, a_h), \quad h \in \{1, 2, \dots, H\} \quad (\text{D.1})$$

A *composite-policy* π is a sequence of kernels $\pi_1, \pi_2, \dots, \pi_H : \mathcal{S} \rightarrow \Delta(\mathcal{A})$. We let P_{init} denote the distribution of initial state s_1 , and D_{π} denote the distribution of $(s_{1:H+1}, a_{1:H})$ subject to $s_1 \sim P_{\text{init}}$, $a_h \mid s_{1:h}, a_{1:h-1} \sim \pi_h(s_h)$, and the composite-dynamics (D.1). We assume that we have an optimal policy π^* to be imitated, and define P_h^* as the marginal distribution of s_h under D_{π^*} ; ultimately, we shall take π^* to be the policy defined in [Definition 3.4](#).

D.1 Structure of the proof.

We begin by explaining key objects, stability and continuity properties required in the composite MDP. Then, [Appendix D.2](#) relates the composite MDP to our original setting by taking composite-states $s_h = s_h$ as chunks, and taking composite actions as sequences of primitive controllers $a_h = \kappa_{t_h:t_{h+1}-1}$ as in [Section 2](#). We also explain why relevant stability and continuity conditions are met. Finally, we derive [Theorem 2](#) from a generic guarantee for smoothed imitation learning in the composite MDP, [Theorem 4](#), and from sampling guarantees in [Appendix L](#).

We consider two pseudometrics on the space \mathcal{S} : $d_{\mathcal{S}}, d_{\text{TVC}} : \mathcal{S}^2 \rightarrow \mathbb{R}_{\geq 0}$, and a function $d_{\mathcal{A}} : \mathcal{A}^2 \rightarrow \mathbb{R}_{\geq 0}$. For convenience, *do not require* $d_{\mathcal{A}}$ to satisfy the axioms of a pseudometric. We use $d_{\mathcal{S}}$ and $d_{\mathcal{A}}$ to measure error between states and actions, respectively, and $d_{\text{TVC}}(\cdot, \cdot)$ for a probabilistic continuity property described below. In terms of $d_{\mathcal{S}}$ and $d_{\mathcal{A}}$, we consider three measures of imitation error: error on the (i) joint distribution of trajectories ($\Gamma_{\text{joint}, \varepsilon}$) (ii) marginal distribution of trajectories ($\Gamma_{\text{marg}, \varepsilon}$) and (iii) one-step error in actions ($d_{\text{os}, \varepsilon}$). Formally:

Definition D.1 (Imitation Errors). Given an error parameter $\varepsilon > 0$, define the **joint-error** $\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \parallel \pi^*) := \inf_{\mu_1} \mathbb{P}_{\mu_1} [\max_{h \in [H]} \max\{d_{\mathcal{S}}(s_{h+1}^*, \hat{s}_{h+1}), d_{\mathcal{A}}(a_h^*, \hat{a}_h)\} > \varepsilon]$, where the first infimum is over trajectory couplings $((\hat{s}_{1:H+1}, \hat{a}_{1:H}), (s_{1:H+1}^*, a_{1:H}^*)) \sim \mu_1 \in \mathcal{C}(\mathcal{D}_{\hat{\pi}}, \mathcal{D}_{\pi^*})$ satisfying $\mathbb{P}_{\mu_1}[\hat{s}_1 = s_1^*] = 1$. Define the **marginal error** $\Gamma_{\text{marg}, \varepsilon}(\hat{\pi} \parallel \pi^*) := \max_{h \in [H]} \{\inf_{\mu_1} \mathbb{P}_{\mu_1} [d_{\mathcal{S}}(s_{h+1}^*, \hat{s}_{h+1}) > \varepsilon], \inf_{\mu_1} \mathbb{P}_{\mu_1} [d_{\mathcal{A}}(a_h^*, \hat{a}_h) > \varepsilon]\}$ to be the same as the joint-gap, with the “max” outside the probability and inf over couplings. Lastly, define the **one-step error** $d_{\text{os}, \varepsilon}(\hat{\pi}_h(s) \parallel \pi_h^*(s)) := \inf_{\mu_2} \mathbb{P}_{\mu_2} [d_{\mathcal{A}}(\hat{a}_h, a_h^*) \leq \varepsilon]$, where the infimum is over $(a_h^*, \hat{a}_h) \sim \mu_2 \in \mathcal{C}(\hat{\pi}_h(s), \pi_h^*(s))$.

Stability. Our hierarchical approach offloads stability of stochastic π^* onto that of its composite-actions a_h , instantiated as *primitive controllers* (not raw inputs!). This allows us to circumvent more challenging incremental senses of stability (see [Appendix E.2](#) for further discussion).

Definition D.2 (Input-Stability). A trajectory $(s_{1:H+1}, a_{1:H})$ is **input-stable** if all sequences $s'_1 = s_1$ and $s'_{h+1} = F_h(s'_h, a'_h)$ satisfy $d_{\mathcal{S}}(s'_{h+1}, s_{h+1}) \vee d_{\text{TVC}}(s'_{h+1}, s_{h+1}) \leq \max_{1 \leq j \leq h} d_{\mathcal{A}}(a'_j, a_j)$, $\forall h \in [H]$. A policy π is **input-stable** if $(s_{1:H}, a_{1:H}) \sim \mathcal{D}_{\pi}$ is **input-stable** almost surely.

TVC. Continuity of probability kernels and policies in TV distance are measured in terms of d_{TVC} .

Definition D.3. For a measure-space \mathcal{X} and non-decreasing $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, we call a probability kernel $W : \mathcal{S} \rightarrow \Delta(\mathcal{X})$ **γ -total variation continuous (γ -TVC)** if, for all $s, s' \in \mathcal{S}$, $\text{TV}(W(s), W(s')) \leq \gamma(d_{\text{TVC}}(s, s'))$. A policy π is **γ -TVC** if $\pi_h : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ is γ -TVC $\forall h \in [H]$.

Data Noising. In [Appendix G](#), we show that under the strong condition that the learned policy $\hat{\pi}$ is γ -TVC, then HINT with no noise injection ($\sigma = 0$) learns the distribution. Frequently, however, $\hat{\pi}$ may not satisfy this condition, such as when the ground truth π^* is not also TVC. We circumvent this by introducing a *smoothing kernel* $W_{\sigma} : \mathcal{S} \rightarrow \Delta(\mathcal{S})$ that corresponds to the data noising procedure; in HINT we let the kernel be a Gaussian, sending o_h to $\mathcal{N}(o_h, \sigma^2 \mathbf{I}) \in \Delta(\mathcal{P}_{o_h})$. We will thus be able to replace TVC of $\hat{\pi}$ with TVC of W_{σ} . We now introduce a few key objects.

Definition D.4. Given a policy π , we define its **smoothed policy** $\pi \circ W_{\sigma}$ via components $(\pi \circ W_{\sigma})_h = \pi_h \circ W_{\sigma} : \mathcal{S} \rightarrow \Delta(\mathcal{A})$. For π^* fixed, define the **smoothed distribution** $P_{\text{aug}, h}^*$ as the joint distribution over $(s_h^* \sim P_h^*, a_h^* \sim \pi_h^*(s_h^*), \tilde{s}_h^* \sim W_{\sigma}(s_h^*))$, with $a_h^* \perp \tilde{s}_h^* \mid s_h^*$. The **deconvolution policy** π_{dec}^* is defined by letting $\pi_{\text{dec}, h}^*(s)$ denote the distribution of $a_h^* \mid \tilde{s}_h^* = s$, where a_h^*, \tilde{s}_h^* are drawn from $P_{\text{aug}, h}^*$. Finally, the **replica policy** is $\pi_{\text{rep}}^* = \pi_{\text{dec}}^* \circ W_{\sigma}$.

The operator $\pi \circ W_{\sigma}$ composes π with the smoothing kernel. The deconvolution policy π_{dec}^* captures the distribution of actions under π^* given a smoothed state, and corresponds to $\pi_{\text{dec}}^* = (\pi_{\text{dec}, h}^*)_{h=1}^H$. We argue that if a policy $\hat{\pi}$ approximates π_{dec}^* at each step, then $\hat{\pi} \circ W_{\sigma}$ imitates $\pi_{\text{rep}}^* = \pi_{\text{dec}}^* \circ W_{\sigma}$. We explain the “replica policy”, and importance of imitating it, after we state our main theorem. First, we define a notion of stability to smoothing, taking $d_{\text{TVC}}, d_{\mathcal{S}}, d_{\mathcal{A}}$ as given.

Definition D.5. For a non-decreasing maps $\gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, \mathcal{S}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ a pseudometric $d_{\text{IPS}} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ (possibly other than $d_{\mathcal{S}}$ or d_{TVC}), and $r_{\text{IPS}} > 0$, we say a policy π is **$(\gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, \mathcal{S}}, d_{\text{IPS}}, r_{\text{IPS}})$ -input-&-process stable (IPS)** if the following holds for any $r \in [0, r_{\text{IPS}}]$. Consider any sequence of kernels $W_1, \dots, W_H : \mathcal{S} \rightarrow \Delta(\mathcal{S})$ satisfying $\max_{h, s \in \mathcal{S}} \mathbb{P}_{\tilde{s} \sim W_h(s)} [d_{\text{IPS}}(\tilde{s}, s) \leq r] = 1$, and define a process $s_1 \sim P_{\text{init}}, \tilde{s}_h \sim W_h(s_h), a_h \sim \pi_h(\tilde{s}_h)$, and $s_{h+1} := F_h(s_h, a_h)$. Then, almost surely, (a) the sequence $(s_{1:H+1}, a_{1:H})$ is input-stable (b) $\max_{h \in [H]} d_{\text{TVC}}(F_h(\tilde{s}_h, a_h), s_{h+1}) \leq \gamma_{\text{IPS}, \text{TVC}}(r)$ and (c) $\max_{h \in [H]} d_{\mathcal{S}}(F_h(\tilde{s}_h, a_h), s_{h+1}) \leq \gamma_{\text{IPS}, \mathcal{S}}(r)$.

Condition (a) means that the policy $\tilde{\pi}_h$ defined by $\tilde{\pi}_h = \pi_h \circ W_h$ is input-stable. In the appendix, we instantiate $W_{1:H}$ not as W_{σ} , but as (a truncation of) *replica kernels* $W_{\text{rep}, h}^*$ for which $\pi_{\text{rep}, h}^* = \pi_h^* \circ W_{\text{rep}, h}^*$. We show that the replica kernel inherits any concentration satisfied by W_{σ} , ensuring (via truncation) that $\mathbb{P}_{\tilde{s} \sim W_h(s)} [d_{\text{IPS}}(\tilde{s}, s) \leq r] \leq r$. Conditions (b & c) merely require that one-step dynamics are robust to small changes in state, measured in terms of both d_{TVC} and $d_{\mathcal{S}}$.

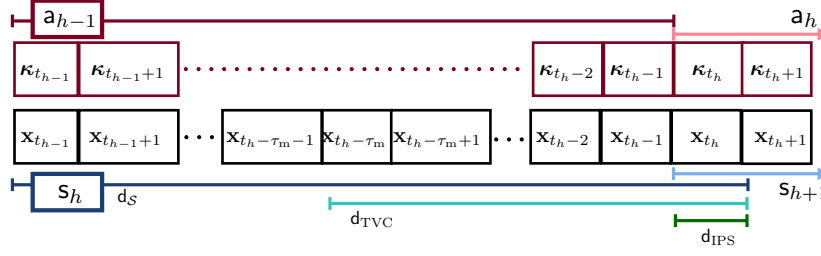


Figure 7: Schematic depicting the composite MDP. States \mathbf{x} and stabilizing gains κ are chunked into composite states \mathbf{s} and composite actions \mathbf{a} (control inputs \mathbf{u} not depicted). The distance labels correspond to the domain over which each distance is evaluated. Note that \mathbf{a}_h begins at the same time that \mathbf{s}_{h+1} does, an indexing convention that we adopt to make the notation in the proofs simpler.

D.2 Instantiation for control

Here we explain the mapping from the control setting of interest to the composite MDP; in so doing we distinguish between the case $h > 1$ and $h = 1$ with reference to composite-states. Recall the definitions of d_{\max} and d_{traj} from (3.1) and (3.2). In the former case, $\mathbf{s}_h = (\mathbf{x}_{t_{h-1}:t_h}, \mathbf{u}_{t_{h-1}:t_h-1}) \in \mathcal{P}_{\tau_{\text{obs}}}$, and $\mathbf{a}_h = \kappa_{t_h:t_{h+1}-1}$ (as in Section 2). Importantly, \mathbf{a}_h *are primitive controllers, which allows us to meet the strong stability condition in Definition D.2*. Figure 7 provides a visual aid for the subtle indexing. For $\mathbf{s}_h = \mathbf{s}_h, \mathbf{s}'_h = \mathbf{s}'_h$, we define $d_S(\mathbf{s}_h, \mathbf{s}'_h) = \max_{t \in [t_{h-1}:t_h]} \|\mathbf{x}_t - \mathbf{x}'_t\| \vee \max_{t \in [t_{h-1}:t_h-1]} \|\mathbf{u}_t - \mathbf{u}'_t\|$, which measures distance on the full subtrajectory, $d_{\text{TVC}}(\mathbf{s}_h, \mathbf{s}'_h) = \max_{t \in [t_h-\tau_{\text{obs}}:t_h]} \|\mathbf{x}_t - \mathbf{x}'_t\| \vee \max_{t \in [t_h-\tau_{\text{obs}}:t_h-1]} \|\mathbf{u}_t - \mathbf{u}'_t\|$, which measures distance on the last τ_{obs} steps, and $d_{\text{IPS}}(\mathbf{s}_h, \mathbf{s}'_h) = \|\mathbf{x}_{t_h} - \mathbf{x}'_{t_h}\|$, which is only on the last step. In the latter case, when $h = 1$, we let $\mathbf{s}_1 = \mathbf{x}_1 \in \mathcal{X}$, and we let $d_S, d_{\text{TVC}}, d_{\text{IPS}}$ all denote the Euclidean distance on \mathcal{X} . In all cases, the transition dynamics F_h are induced by the dynamics (2.1) with $\mathbf{u}_t = \kappa_t(\mathbf{x}_t)$. Finally, for $\mathbf{a} = (\bar{\mathbf{u}}_{1:\tau_{\text{chunk}}}, \bar{\mathbf{x}}_{1:\tau_{\text{chunk}}}, \bar{\mathbf{K}}_{1:\tau_{\text{chunk}}})$ and $\mathbf{a}' = (\bar{\mathbf{u}}'_{1:\tau_{\text{chunk}}}, \bar{\mathbf{x}}'_{1:\tau_{\text{chunk}}}, \bar{\mathbf{K}}'_{1:\tau_{\text{chunk}}})$, we recall from (3.1) $d_{\max}(\mathbf{a}, \mathbf{a}') := \max_{1 \leq k \leq \tau_{\text{chunk}}} (\|\bar{\mathbf{u}}_k - \bar{\mathbf{u}}'_k\| + \|\bar{\mathbf{x}}_k - \bar{\mathbf{x}}'_k\| + \|\bar{\mathbf{K}}_k - \bar{\mathbf{K}}'_k\|)$. We choose a d_A that takes value ∞ when primitive controllers are too far apart

$$d_A(\mathbf{a}, \mathbf{a}') := c_1 d_{\max}(\mathbf{a}, \mathbf{a}') \cdot \mathbf{I}_{\infty}\{d_{\max}(\mathbf{a}, \mathbf{a}') \leq c_2\} \quad (\text{D.2})$$

\mathbf{I}_{∞} is the indicator taking infinite value when the event \mathcal{E} fails to hold and 1 otherwise, and c_1 and c_2 are constants depending polynomially on all problem parameters, given in Appendix K.

We let the expert policy π^* be the concatenation of policies π_h^* , each of which is defined to be the distribution of \mathbf{a}_h conditioned on \mathbf{o}_h under \mathcal{D}_{exp} (see Appendix J for a rigorous definition). As noted above, we take the smoothing kernel W_{σ} to map \mathbf{o}_h to a $\mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I}) \in \Delta(\mathcal{P}_{\mathbf{o}_h})$, which that same appendix shows is $\frac{1}{2\sigma}$ -TVC (w.r.t. d_{TVC} defined above). We note that under these substitutions, the deconvolution policy $\pi_{\text{dec}}^* = (\pi_{\text{dec},h}^*)_{h=1}^H$ is *precisely as defined in Definition C.1*.

Appendix K shows that Assumption 3.1 imply that π^* enjoys the IPS property in the composite MDP thus instantiated, along with many more granular stability guarantees for time-varying affine feedback in nonlinear control systems, which may be of independent interest.

Proposition D.1. Let $c_3, c_4, c_5 > 0$ be as given in Appendix K (and polynomial in the terms in Assumption 3.1). Suppose $\tau_{\text{chunk}} \geq c_3$, and let $r_{\text{IPS}} = c_4$, $\gamma_{\text{IPS}, \text{TVC}}(u) = c_5 u \exp(-L_{\beta}(\tau_{\text{chunk}} - \tau_{\text{obs}}))$, $\gamma_{\text{IPS}, S}(u) = c_5 u$. Then, for $d_S, d_{\text{TVC}}, d_{\text{IPS}}$ as above, we have that π^* is $(\gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, S}, d_{\text{IPS}}, r_{\text{IPS}})$ -IPS.

D.3 A Guarantee in the Composite MDP, and the derivation of Theorem 3

With the substitutions in Appendix D.2, it suffices to prove an imitation guarantee in the composite MDP, assuming π^* is IPS, and $\hat{\pi}$ is close to π_{dec}^* in the appropriate sense.

Theorem 4. Suppose π^* is $(\gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, S}, d_{\text{IPS}}, r_{\text{IPS}})$ -IPS and W_{σ} is γ_{σ} -TVC. Let $\varepsilon > 0$, $r \in (0, \frac{1}{2}r_{\text{IPS}}]$; define $p_r := \sup_{\mathbf{s}} \mathbb{P}_{\mathbf{s}' \sim W_{\sigma}(\mathbf{s})}[d_{\text{IPS}}(\mathbf{s}', \mathbf{s}) > r]$ and $\varepsilon' := \varepsilon + \gamma_{\text{IPS}, S}(2r)$. Then, for any

policy $\hat{\pi}$, both $\Gamma_{\text{joint},\varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$ and $\Gamma_{\text{marg},\varepsilon'}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ are upper bounded by

$$H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS,TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*)). \quad (\text{D.3})$$

Deriving Theorem 2 from Theorem 4. A full proof is given in [Appendix J](#). The key steps are using the stability guarantee of [Proposition D.1](#), the aforementioned TVC-bound on W_σ , and Gaussian concentration to bound p_r with the bound in [Theorem 4](#) to conclude. Moreover, we can show that, with $\tilde{s}_h^* = \tilde{o}_h$,

$$\sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*)) = \sum_{h=1}^H \mathbb{E}_{\tilde{o}_h \sim \mathcal{D}_{\text{exp},\sigma,h}} [\Delta_{(\varepsilon^2)}(\pi_{\text{dec},\sigma,h}^*(\tilde{o}_h), \hat{\pi}_h(\tilde{o}_h))].$$

Thus, [Theorem 4](#) provides the desired guarantee for imitating the policy π^* constructed in [Appendix D.2](#). We conclude with a subtle but powerful observation: that π^* as constructed has trajectories with the same marginals (but possibly different joint distributions) as $\rho_T \sim \mathcal{D}_{\text{exp}}$. \square

A simplified guarantee when $\hat{\pi}$ is TVC. Before we sketch the proof of [Theorem 4](#), we present the simpler guarantee that underpins [Theorem 1](#) in [Section 3](#). This guarantee considers *no smoothing*, but where $\hat{\pi}$ is guaranteed to be γ -TVC. As discussed in [Remark C.3](#), this situation is practically unrealistic when $\hat{\pi}$ is instantiated with a DDPM. Still, the following result is more transparent, and its proof sketch will inform the proof sketch of [Theorem 4](#) following it.

Proposition D.2. Suppose for simplicity that $d_S = d_{\text{TVC}}$. Let π^* be input-stable w.r.t. (d_S, d_A) and let $\hat{\pi}$ be γ -TVC. Then, for all $\varepsilon > 0$,

$$\Gamma_{\text{joint},\varepsilon}(\hat{\pi} \parallel \pi^*) \leq H\gamma(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} d_{\text{os},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi^*(s_h^*)) \quad (\text{D.4})$$

Proof Sketch of Proposition D.2. We couple a trajectory (s_h^*, a_h^*) induced by π^* with a trajectory (\hat{s}_h, \hat{a}_h) induced by $\hat{\pi}$. To do so, consider an inductive event on which all past states and actions are close

$$\mathcal{E}_h = \{\forall j \leq h, d_{\text{TVC}}(\hat{s}_j, s_j^*) \vee d_S(\hat{s}_j, s_j^*) \leq \varepsilon\} \cap \{\forall j \leq h-1, d_A(\hat{a}_j, a_j^*) \leq \varepsilon\},$$

which can be made to hold with probability one for $h = 1$ by ensuring $\hat{s}_1 = s_1^* \sim P_{\text{init}}$. It suffices to find a μ coupling under which $\mathbb{P}_\mu[\mathcal{E}_{H+1}]$ is bounded by the righthand side of [\(D.4\)](#).

When \mathcal{E}_h holds, then as $\hat{\pi}$ is γ -TVC, $\hat{a}_h \sim \hat{\pi}_h(\hat{s}_h)$ is $\gamma(\varepsilon)$ -close in TV distance to an interpolating action $\hat{a}_h^{\text{inter}} \sim \pi_h(s_h^*)$. Thus, there is a coupling under which $\mathbb{P}[\hat{a}_h \neq \hat{a}_h^{\text{inter}}] \leq \gamma(\varepsilon)$. Moreover, by definition of $d_{\text{os},\varepsilon}$, there exists a coupling under which $\mathbb{P}[d_A(\hat{a}_h^{\text{inter}}, a_h^*) > \varepsilon] \leq \mathbb{E}_{s_h^* \sim P_h^*} d_{\text{os},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi^*(s_h^*))$. Thus, by “gluing” the couplings together (an operation rigorously justified in [Appendix F](#)), we have

$$\mathbb{P}[\mathcal{E}_h \cap \{d_A(\hat{a}_h, a_h^*) > \varepsilon\}] \leq \gamma(\varepsilon) + \mathbb{E}_{s_h^* \sim P_h^*} d_{\text{os},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi^*(s_h^*)). \quad (\text{D.5})$$

Invoking the definition of input stability ([Definition D.2](#)), \mathcal{E}_h and $d_A(\hat{a}_h, a_h^*) \leq \varepsilon$ imply \mathcal{E}_{h+1} . Therefore, by selecting couplings appropriately, $\mathbb{P}[\mathcal{E}_h \cap \mathcal{E}_{h+1}^c]$ is also bounded by the righthand side of [\(D.5\)](#). As the events (\mathcal{E}_h) are nested, we can finally telescope to bound $\mathbb{P}[\mathcal{E}_{H+1}]$ by summing up these terms for each h . A full proof is given in [Appendix G](#). \square

D.4 Proof Overview of Theorem 4

In this sketch, we focus on bounding $\Gamma_{\text{joint},\varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$; we address $\Gamma_{\text{marg},\varepsilon'}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ at the end of the section. Our argument constructs a coupling between a *replica trajectory* over $(s_h^{\text{rep}}, a_h^{\text{rep}})$ sampled using the replica policy π_{rep}^* , and an *imitator trajectory* (\hat{s}_h, \hat{a}_h) sampled from $\hat{\pi}_\sigma$. The construction of this coupling draws inspiration from the notion of replica pairs in statistical physics (hence, the name replica) [\[45\]](#). We refer the reader back to [Figure 4](#) above the proof sketch in [Section 3.3](#) for a visual depiction.

The replica kernel. A central object in our proof is the *replica kernel* $W_{\text{rep},h}^* : \mathcal{S} \rightarrow \Delta(\mathcal{S})$, defined so that $\pi_{\text{rep},h}^* = \pi_h^* \circ W_{\text{rep},h}^*$, and constructed by analogy to the replica policy in Definition D.4. The key property of the replica kernel is that it preserves marginals: if $s_h \sim P_h^*$ is drawn from the distribution of states under the expert demonstrations, then $s_h' \sim W_{\text{rep},h}^*(s_h)$ is also distributed as P_h^* ; in other words, P_h^* is a fixed point of $W_{\text{rep},h}^*$:

Fact D.1 (Replica Property). It holds that $P_h^* = W_{\text{rep},h}^* \circ P_h^*$.

A second crucial property is that we can represent the replica kernel as a convolution between W_σ and a deconvolution kernel. Thus, data-processing implies that $W_{\text{rep},h}^*$ inherits TVC from W_σ .

Fact D.2 (Replica Kernel is TVC). If W_σ is γ_{TVC} -TVC, $W_{\text{rep},h}^*$ is as well.

Constructing the replica and teleporting trajectories. Because the replica kernel satisfies $\pi_{\text{rep},h}^* = W_{\text{rep},h}^* \circ \pi_h^*$, we can realize the replica trajectory via

$$\tilde{s}_h^{\text{rep}} \sim W_{\text{rep},h}^*(s_h^{\text{rep}}), \quad a_h^{\text{rep}} \sim \pi_h^*(\tilde{s}_h^{\text{rep}}), \quad s_{h+1}^{\text{rep}} = F_h(s_h^{\text{rep}}, a_h^{\text{rep}}), \quad s_1^{\text{rep}} \sim P_{\text{init}}.$$

We then introduce a *teleporting* trajectory obeying the an almost identical generative process:

$$\tilde{s}_h^{\text{tel}} \sim W_{\text{rep},h}^*(s_h^{\text{tel}}), \quad a_h^{\text{rep}} \sim \pi_h^*(\tilde{s}_h^{\text{tel}}), \quad s_{h+1}^{\text{tel}} = F_h(\tilde{s}_h^{\text{tel}}, a_h^{\text{tel}}), \quad a_1^{\text{tel}} \sim P_{\text{init}}.$$

In words, s_h^{tel} “teleports” to an independent and identically distributed copy conditional on the noise augmentation \tilde{s}_h^{tel} , and draws an action from the expert policy evaluated on the new state. The replica and teleporting sequences differ only in the transitions: whereas $s_{h+1}^{\text{rep}} = F_h(s_h^{\text{rep}}, a_h^{\text{rep}})$ transitions from its *current* state s_h^{rep} , the telporting trajectory transition $s_{h+1}^{\text{tel}} = F_h(\tilde{s}_h^{\text{tel}}, a_h^{\text{tel}})$ from the *replica-drawn*, “teleported” state $\tilde{s}_h^{\text{tel}} \sim W_{\text{rep},h}^*(s_h^{\text{tel}})$. By iteratively applying Fact D.1, and the fact that $s_h^{\text{tel}} \sim \pi_h^*(\tilde{s}_h^{\text{tel}})$, we can make the following claim:

Fact D.3. For each h , s_h^{tel} and \tilde{s}_h^{tel} are distributed according to P_h^* , the marginal distribution of states under the expert policy. Hence, because $s_h^{\text{tel}} \sim P_h^*$, we know that $\hat{\pi}$ and π_h^* are close under the distribution of states induced teleporting sequence.

Constructing the coupling. We now describe how to use the teleporting trajectory to couple the replica and imitator trajectory. We begin by coupling the replica and imitator trajectories. The following diagram explains how we construct the coupling:

$$\underbrace{(\tilde{s}^{\text{rep}} \leftrightarrow \tilde{s}^{\text{tel}}), (a^{\text{rep}} \leftrightarrow a^{\text{tel}})}_{\gamma_{\text{TVC}}, \gamma_{\text{IPS}}, \text{TVC, and induction}} \rightarrow \underbrace{(a^{\text{tel}} \leftrightarrow \hat{a}^{\text{tel,inter}})}_{\text{learning and sampling}} \rightarrow \underbrace{(\hat{a}^{\text{tel,inter}} \leftrightarrow \hat{a}^{\text{rep,inter}})}_{\gamma_{\text{TVC}} \text{ and induction}} \rightarrow \underbrace{(\hat{a}^{\text{rep,inter}} \leftrightarrow \hat{a})}_{\gamma_{\text{TVC}}, \text{input-stability, and induction}}. \quad (\text{D.6})$$

As the dynamics are deterministic, (D.6) determines the coupling of states $s_h^{\text{rep}}, s_h^{\text{tel}}, \hat{s}_h$ as well.

- (a) We begin by arguing that replica and teleporting trajectories are close to one another. The argument is inductive: suppose that $d_{\text{TVC}}(s_h^{\text{rep}}, s_h^{\text{tel}})$ are close. By TVC of the replica kernel, \tilde{s}_h^{rep} and \tilde{s}_h^{tel} are close in TV, and so there is a coupling under which

$$\mathbb{P}[(a_h^{\text{rep}}, \tilde{s}_h^{\text{rep}}) \neq (a_h^{\text{tel}}, \tilde{s}_h^{\text{tel}})] \text{ is small.}$$

Next, recall $p_r := \sup_s \mathbb{P}_{s' \sim W_\sigma(s)}[d_{\text{IPS}}(s', s) > r]$ as defined in Theorem 4, which describes the concentration behavior of W_σ . We use a Bayesian concentration argument (Lemma H.5) to ensure that with probability of failure at most $2p_r$, $d_{\text{IPS}}(\tilde{s}_h^{\text{rep}}, \tilde{s}_h^{\text{rep}}) \leq 2r$. We then use IPS (Definition D.5) to argue that, with the same failure probability,

$$s_{h+1}^{\text{rep}} = F_h(s_h^{\text{rep}}, a_h^{\text{rep}}) \text{ is within } \mathcal{O}(r) \text{ } d_{\text{TVC}}\text{-distance of } F_h(\tilde{s}_h^{\text{rep}}, a_h^{\text{rep}}).$$

Thus, when $(a_h^{\text{rep}}, \tilde{s}_h^{\text{rep}}) = (a_h^{\text{tel}}, \tilde{s}_h^{\text{tel}})$, we obtain that s_{h+1}^{rep} is close to $F_h(a_h^{\text{tel}}, \tilde{s}_h^{\text{tel}}) = s_{h+1}^{\text{tel}}$ in d_{TVC} as well. For more detail, consult Figure 10 in the appendix.

- (b) Because the marginals of \tilde{s}_h^{tel} are distributed according to P_h^* , we can argue that a (fictitious) action $\hat{a}_h^{\text{tel,inter}} \sim (\hat{\pi}_h \circ W_\sigma)(\tilde{s}_h^{\text{tel}})$ is close to a_h^{tel} . Indeed, by the data processing inequality, it is bounded by the closeness of $\hat{\pi}_h$ and $\pi_{\text{dec},h}^*$ on $\tilde{s}_h^{\text{tel}} \sim W_\sigma(s_h^{\text{tel}})$, $s_h^{\text{tel}} \sim P_h^*$; this is controlled by the contribution of $d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*))$ on the right-hand-side of (M.2).

- (c) From part (a) of the coupling, we expect that s_h^{rep} and s_h^{tel} are close. As $\hat{\pi} \circ W_\sigma$ is γ_{TVC} by the data-processing inequality, it follows that $\hat{a}_h^{\text{tel,inter}}$ is close in TV distance to another fictitious action $\hat{a}_h^{\text{rep,inter}} \sim (\hat{\pi}_h \circ W_\sigma)(s_h^{\text{rep}})$.
- (d) Finally, we argue that that trajectory of actions induced by taking the replica-interpolating action $\hat{a}_h^{\text{rep,inter}}$ is close to TV to the imitation trajectories induced by taking \hat{a}_h . The argument is similar in spirit to the proof of [Proposition D.2](#) sketched above, and uses both TVC of the smoothed policy $\hat{\pi} \circ W_\sigma$, and the form of input-stability guaranteed by [Definition D.5](#).

Bounding the marginal gap. Because the teleporting sequence s_h^{tel} has marginals P_h^* , bounding the marginal gap amounts to controlling the distance between s_h^{tel} and \hat{s}_h . This follows from more-or-less the same manipulations.⁵

Measure-theoretic considerations. We construct conditional couplings between pairs of the aforementioned trajectories, each of which corresponds to a certain optimal transport cost. That past trajectories can be associated to optimal couplings measurably is non-trivial, and proven in [Proposition F.3](#). To conclude, we apply a measure theoretic result ([Lemma F.2](#)) to “glue” the pairwise couplings together and establish the main result. The full proof is given in [Appendix H](#), relying on measure-theoretic details in [Appendix F](#). \square

E Generalization to Generic Incrementally Stable Primitive Controllers

In this section, we consider a generalization of the theory to allow for general, nonlinear primitive controllers, as long as they obey the incremental stability considered in Pfrommer et al. [50]. We consider controllers of the form:

$$\kappa(\mathbf{x}) = \kappa(\mathbf{x}; \theta), \quad \theta \in \Theta.$$

We assume that $\Theta \subset \mathbb{R}^{d_\Theta}$ is a measurable subset of a finite dimensional space and (2) that $\kappa(\mathbf{x}; \theta)$ is jointly piecewise-Lipschitz with at most countably many pieces. We define composite actions just as for linear primitive controllers:

$$\mathbf{a} = (\kappa_1, \kappa_2, \dots, \kappa_{\tau_{\text{chunk}}}).$$

We view $\mathcal{K} := \{\kappa(\cdot; \theta) | \theta \in \Theta\}$ and $\mathcal{A} = \mathcal{K}^{\tau_{\text{chunk}}}$ as Polish spaces with the Euclidean metric on the controller parameters θ (resp. sequences of control parameters). Our definition of incremental stability from [Definition 3.1](#) applies verbatim to these more general controllers.

Example E.1 (Approximate Inverse Dynamics & Position Control). A natural example of the above is where θ corresponds to a sequence of position commands supplied to a robotic position controller, as in [19], or where θ is a of state-command given to an inverse dynamics model, as in [3]. In these settings, we can actually regard θ as the “control action,” and envision the closed loop system of {system + position controller/inverse dynamics model} as itself being incrementally stable. However, our framework is considerably more general, and allows us, for example, to diffuse other parameters governing the performance of the low level controllers as well (e.g. joint spring constants in robotic position control).

In this section, we replace [Assumption 3.1](#) with the more general assumption that allows arbitrary forms of the incremental stability moduli:

Assumption 3.1b. Let $\gamma(\cdot)$ be class K, $\beta(\cdot, \cdot)$ be class KL, and let $c_\xi, c_\beta, c_\gamma > 0$ be postive constants. We assume access to a synthesis oracle $\text{synth} : \mathcal{P}_T \rightarrow \mathcal{A}^H$ such that, with probability 1 over $\rho_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \sim \mathcal{D}_{\text{exp}}, \mathbf{a}_{1:H} = \text{synth}(\rho_T)$ satisfies the following properties:

- $\mathbf{a}_h = \kappa_{t_h:t_{h+1}-1}$ is consistent with $\mathbf{s}_{h+1} = (\mathbf{x}_{t_h:t_{h+1}}, \mathbf{u}_{t_h:t_{h+1}-1})$; equivalently,

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \kappa_t(\mathbf{x}_t)), \quad t \in [T].$$

- \mathbf{a}_h is locally t-ISS at \mathbf{x}_{t_h} with moduli $\gamma(\cdot)$, $\beta(\cdot, \cdot)$ and constants $c_\gamma, c_\beta, c_\xi > 0$.

⁵A key difference is that we pick up an additive factor of $\gamma_{\text{IPS}, S}$ (measuring sensitivity of d_S to the smoothing from W_σ) in our tolerance ε' .

E.1 Main Results for General Primitive Controllers

To state our main results, we begin by defining distances on primitive controllers, as well as the induced imitation error of a policy $\hat{\pi}$. Throughout, $\mathcal{D}_{\sigma,h}$ denotes the noise-smoothed expert data distribution over (\tilde{o}_h, a_h) as in Definition 3.6 (with a_h now representing sequence of general primitive controllers, not the linear ones considered in the main body).

Definition E.1. Define the local-distance between composite actions $a = \kappa_{1:\tau_{\text{chunk}}}$, $a' = \kappa'_{1:\tau_{\text{chunk}}} \in \mathcal{A}$ at state \mathbf{x} and scale $\alpha > 0$ as

$$d_{\text{loc},\alpha}(a, a' | \mathbf{x}) := \max_{1 \leq i \leq \tau_{\text{chunk}}} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \|\kappa_i(\mathbf{x}_i + \delta \mathbf{x}) - \kappa'_i(\mathbf{x}_i + \delta \mathbf{x})\|,$$

where above $\mathbf{x}_1 = \mathbf{x}$, $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \kappa_t(\mathbf{x}))$, and $a = \kappa_{1:\tau_{\text{chunk}}}$. Finally, we define

$$\Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha) := \inf_{\mu \in \overline{\mathcal{C}}_{\sigma,h}(\hat{\pi})} \mathbb{P}_{(a_h, a'_h) \sim \mu} [d_{\text{loc},\alpha}(a, a' | \mathbf{x}_{t_h}) > \varepsilon],$$

where $\overline{\mathcal{C}}_{\sigma,h}(\hat{\pi})$ denotes the set of couplings of $(o_h, \tilde{o}_h, a_h, a'_h)$, induced by drawing $(o_h, a_h) \sim \mathcal{D}_{\text{exp},h}$, $\tilde{o}_h \sim \mathcal{N}(o_h, \sigma^2 \mathbf{I})$, and $a'_h \sim \hat{\pi}_h(\tilde{o}_h)$, and where above \mathbf{x}_{t_h} is the last state in o_h . Note that the only degree of freedom for in selecting elements of $\overline{\mathcal{C}}_{\sigma,h}(\hat{\pi})$.

Remark E.1 (On the distance $d_{\text{loc},\alpha}$). In words, $d_{\text{loc},\alpha}(a, a' | \mathbf{x}_{1:\tau_{\text{chunk}}})$ measures the supremal distance between the primitive controllers comprising a, a' , along radius- α neighborhoods of a given sequence $\mathbf{x}_{1:\tau_{\text{chunk}}}$. This supremal distance was studied in Pfrommer et al. [50] and motivates their proposed algorithm TaSIL. Unlike affine primitive controllers, the supremal distance between general primitive controllers may indeed dependence on the localizing sequence $\mathbf{x}_{1:\tau_{\text{chunk}}}$.

Having defined our distance, $\Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha)$ measures the probability that this difference exceeds some threshold $\alpha > 0$, under appropriate couplings where \tilde{o}_h is induced by noising the expert distribution, a is the corresponding action, a' is from the policy $\hat{\pi}$, and the localizing sequence follows from rolling out a on the last state $\tilde{\mathbf{x}}_{t_h}$ of \tilde{o}_h . $\Delta_{\text{ISS},\sigma,h}$ is the natural analogue of the distances consider in Theorem 2. One subtlety of $\Delta_{\text{ISS},\sigma,h}$ is that the couplings are constructed such that a_h is the action associated with o_h . This is conceptually correct because it specifies a localizing-state for $d_{\text{loc},\alpha}(a_h, a'_h | \mathbf{x}_{t_h})$ (which is not an issue for affine primitive controllers). When a_h arises from a synthesis oracle, \mathbf{x}_{t_h} lies on the trajectory from which a_h is synthesized.

Theorem 2 generalizes as follows.

Theorem 5 (Generalization of Theorem 2 to general primitive controllers). Assume Assumption 3.1b, and let $\varepsilon > 0$ satisfy

$$\gamma^{-1}(\beta(2\gamma(\varepsilon), \tau_{\text{chunk}})) \leq \varepsilon \leq \min\{c_\gamma, \gamma^{-1}(c_\varepsilon/4)\} \quad (\text{E.1})$$

Define

$$\omega = 2\sqrt{5d_x + 2\log\left(\frac{2\sigma}{\gamma(\varepsilon)}\right)}, \quad \varepsilon_1 = 2\beta(2\gamma(\varepsilon), 0) + 2\beta(2\sigma\omega, 0), \quad \varepsilon_2 = 2\beta(2\gamma(\varepsilon), 0).$$

Then, if $\sigma \leq c_\varepsilon/4\omega$ and $\gamma(\varepsilon) \leq 2\sigma$, we have

$$\mathcal{L}_{\text{marg},\varepsilon_1}(\hat{\pi}) \leq \frac{3H\sqrt{2\tau_{\text{obs}} - 1}}{2\sigma} (2\varepsilon_2 + \beta(2\sigma\omega, \tau_{\text{chunk}} - \tau_{\text{obs}})) + \sum_{h=1}^H \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \varepsilon_1).$$

Before continuing, let us remark on the parameters and the scaling. Here, ε parametrizes an error scale. ε_1 captures both the imitation error, as well as the radius in which $\Delta_{\text{ISS},\sigma,h}$ is evaluated. ε_2 contributes to the upper bound on the $\mathcal{L}_{\text{marg},\varepsilon_1}$ normalized by $1/\sigma$. We recall that $\mathcal{L}_{\text{marg},\varepsilon_1}$ measures probabilities of deviating from the marginal by a magnitude of at most ε_1 under optimal couplings. To drive the upper bound on $\mathcal{L}_{\text{marg},\varepsilon}$ to zero, we need (a) $\Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \varepsilon_1) \rightarrow 0$, (b) $\varepsilon_2/\sigma \rightarrow 0$ and (c) $\frac{1}{\sigma}\beta(2\sigma\omega, \tau_{\text{chunk}} - \tau_{\text{obs}}) \rightarrow 0$, which requires $\tau_{\text{chunk}} - \tau_{\text{obs}}$ to grow and β to decay in its second argument. To drive the tolerance ε_1 to zero, we require that $2\beta(2\gamma(\varepsilon), 0) + 2\beta(2\sigma\omega, 0)$ to tend to zero, which requires $\sigma \rightarrow 0$ as well.⁶ The following remark examines the typical scalings of these terms and checks that (E.1) is generally easy to satisfy.

⁶Notice that $\varepsilon_1 \rightarrow \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \varepsilon_1)$ is non-increasing, so making ε_1 smaller does not increase this term (though making ε smaller does).

Remark E.2. Suppose that for some $c > 0$, $q_1, q_2 \in (0, 1]$, we have $\gamma(u) = cu^{q_1}$ and $\beta(u, \tau) = u^{q_2} \phi(\tau)$ for some decreasing function ϕ . This is the scaling studied in [50], and indeed for smooth systems with stabilizable systems, our analysis essentially shows that we can take $q_1 = q_2 = 1$ and $\phi(\tau)$ to decay exponentially in τ as shown in Appendix K, and which reflects in Theorem 2. For the more general power scalings, (E.1) reads

$$c^{q_2-1} \phi(\tau_{\text{chunk}})^{1/q_1} \varepsilon^{q_2} \leq \varepsilon \leq \text{constant}.$$

If $q_2 = 1$ (e.g. the stabilizable, smooth case), then this is satisfied whenever ε is sufficiently small and τ_{chunk} is sufficiently large. Otherwise, it one has to take $\varepsilon \geq \frac{1}{c} \phi(\tau_{\text{chunk}})^{1/q_1(1-q_2)}$, which becomes increasingly permissive as τ is enlarged. Moreover, for these scalings, we have $\omega = \mathcal{O}(\log(1 + \sigma/\varepsilon^{q_1}))$, $\varepsilon_1 = \mathcal{O}(\varepsilon^{q_1 q_2} + \sigma \omega) = \varepsilon_1 = \tilde{\mathcal{O}}(\varepsilon^{q_1 q_2} + \sigma)$, and $\varepsilon_2 = \mathcal{O}(\varepsilon^{q_1 q_2})$. In the regime where $q_1 = q_2$, this recovers the scaling observed in Theorem 2.

Similarly, we can generalize Theorem 1 to the general controller setting.

Theorem 6. Suppose Assumption 3.1b holds, and suppose that $\varepsilon > 0$ and $\tau_{\text{chunk}} \in \mathbb{N}$ satisfies (E.1). Then, for any non-decreasing non-negative $\gamma(\cdot)$ and γ -TVC chunking policy $\hat{\pi}$,

$$\mathcal{L}_{\text{marg}, \varepsilon_1}(\hat{\pi}) \leq H\gamma(\varepsilon) + \sum_{h=1}^H \Delta_{\text{Iss}, \sigma, h}(\hat{\pi}; \varepsilon, \varepsilon_1), \quad \varepsilon_1 := 2\beta(2\gamma(\varepsilon), 0) \quad (\text{E.2})$$

In addition, suppose the expert distribution \mathcal{D}_{exp} has at most τ_{obs} -bounded memory (defined formally in Definition J.5). Then $\mathcal{L}_{\text{joint}, \alpha(\varepsilon)}(\hat{\pi})$ satisfies the same upper bound (E.2), where $\mathcal{L}_{\text{joint}, \alpha}(\hat{\pi})$, formally defined in Definition J.4, measures an optimal transport distance between the *joint distribution* of the expert trajectory and the one induced by $\hat{\pi}$.

The proofs of Theorems 5 and 6 are given in the Appendix M, generalizing the proofs of Theorems 1 and 2 in the main text, respectively. As with Theorem 1, Appendix N.1 shows that we can replace the condition on the chunk length τ_{chunk} on and on ε in (K.4) with the condition $\varepsilon \leq c_\gamma$ and the vacuous condition $\tau_{\text{chunk}} \geq 1$, provided that the synthesis oracle produces entire primitive controllers for which the entire sequences $\kappa_{1:T}$ are incrementally stabilizing.

E.2 Comparison to prior notions of stability.

Prior theoretical work in imitation learning focuses either on constraining the learned policy to be stable [27, 68] or assumes the expert policy is suitably stable [50]. The principal notion of stability used in these prior works is *incremental-input-to-state* stability of the closed-loop system under a deterministic, but possibly sophisticated time-independent controller $\pi : \mathcal{X} \rightarrow \mathcal{U}$. Importantly, this work considers the imitation of a *joint distribution* over *sequences* of simple controllers κ we call the “primitive controllers”. These approach necessitate subtle differences in our choice of definitions described below.

In what follows, we let γ be a class K function and β be a class KL function, as described above Definition 3.1.

Definition E.2 (Incremental Input-to-State Stability). We say a policy $\pi : \mathcal{X} \rightarrow \mathcal{U}$ satisfies *Incremental Input-To-State Stability* (δ -ISS) with moduli γ and β if for any two initial conditions $\xi_1, \xi_2 \in \mathcal{X}$, the closed-loop dynamics under policy $\pi : \mathcal{X} \rightarrow \mathcal{U}$ given by $f_{\text{cl}}(\mathbf{x}_t, \Delta_t) = f(\mathbf{x}_t, \pi(\mathbf{u}_t) + \Delta_t)$ satisfies:

$$\|\mathbf{x}_t(\xi_1; \delta \mathbf{u}_{1:t}) - \mathbf{x}_t(\xi_2; \mathbf{0}_{1:t})\| \leq \beta(\|\xi_1 - \xi_2\|) + \gamma\left(\max_{0 \leq s \leq t-1} \|\delta \mathbf{u}_s\|\right),$$

where $\mathbf{x}_t(\xi; \delta \mathbf{u}_{1:t-1})$ is the state at time t under f_{cl} with $x_0 = \xi$ and input perturbations $\delta \mathbf{u}_{1:t-1}$.

We say that it satisfies π satisfies *local* δ -ISS with parameter c if the above holds for all of identical initial conditions $\xi_1 = \xi_2$ (with $\beta(0) = 0$) and for $\delta \mathbf{u}_{1:t-1}$ satisfying $\max_{1 \leq s \leq t} \|\delta \mathbf{u}_s\| \leq c$.

Notice that for $\xi_1 \neq \xi_2$, the β -term necessitates that the dynamics converge irrespective of initial condition. Without time-varying dynamics this can only be achieved by a policy which stabilizes to an equilibrium point, as a policy which tracks a reference trajectory is unable to “forget” the initial condition. Constraining learned policies such that they satisfy this notion of stability is also

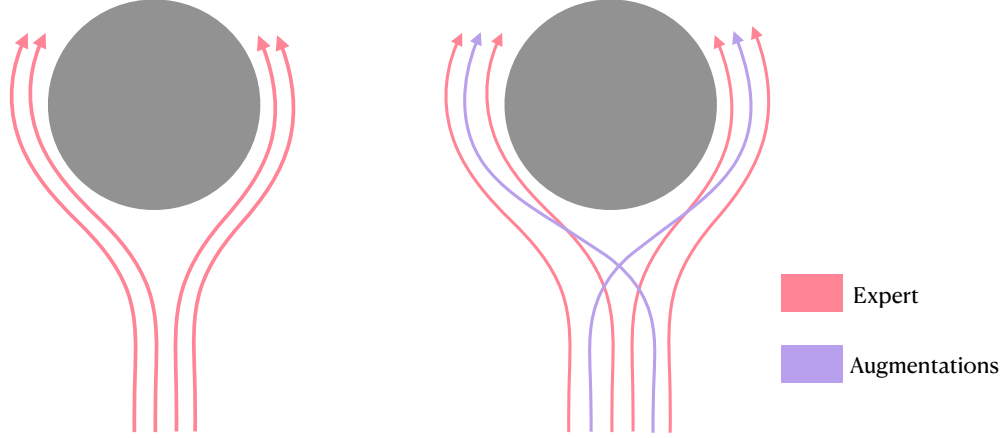


Figure 8: Instance of bifurcation, where augmentation is necessary for stability. The example on the left has an expert demonstrator bifurcating around a circular obstacle. The example on the right demonstrates the utility of augmentations, allowing for trajectories that navigate around the object in the direction farther from their starting point.

challenging. Tu et. al. [68] attempt to do so through regularization while Haven et. al. [27] use matrix inequalities to satisfy this stability property under linear dynamics. Pfrommer et. al. [50] avoid this difficulty only requiring local incremental stability. This weaker notion of incremental stability simply postulates the existence of a (local) input-perturbation to state-perturbation gain function γ . Since this stability property does not necessitate convergence across with different initial conditions and only under input perturbations of magnitude $\leq c$, this only necessitates that the expert policy can correct from small input perturbations.

Comparing local δ -ISS and Definition 3.1. As stated above, past work consider imitation of a fixed, but possibly complex deterministic controller π . In contrast, we imitate joint distributions over *sequences* of primitive controllers $\mathbf{a} = (\kappa_{1:\tau})$. Moreover, our “primitive” controllers are intended to be much simpler than the policy π considered in past work; e.g. the affine controllers considered in the body in this work. Indeed, the real “policy” we try to imitate is potentially very complex expert distribution \mathcal{D}_{exp} , and these primitive controllers serve to stabilize to this distribution. To account for these differences, we modify the local stability considered by Pfrommer et al. [50] in three respects.

- Our notion of stability, Definition 3.1, is applied to fixed-length sequences of controllers $\mathbf{a} = (\kappa_1, \dots, \kappa_\tau)$; past notions of incremental stability are for time-varying controllers and are infinite horizon.
- Definition 3.1 only requires that our notion of incremental stability holds for initial conditions ξ in a radius c_ξ of a nominal initial condition ξ_0 . The reason for this can be seen by considering just the affine primitive controllers studied in the body: time-varying feedback that stabilizes the linearization of a smooth dynamical system is only stabilizing of the actual system in a tube around the nominal trajectory.
- Unlike the local notion of δ -ISS considered in Pfrommer et al. [50], we *do* require considering stability from different initial conditions $\xi_1 \neq \xi_2$. This is because we re-apply incremental stability at each chunk h , and must account for imitation error accumulated up to that point.

The power of a hierarchical approach to stability. Through the introduction of a synthesis oracle which can generate locally stabilizing primitive controllers, we decouple the stability properties of the expert’s behavior from the stabilizability of the underlying dynamical system. This allows for reasoning about generalization in the presence of bifurcations or conflicting demonstrations, which is precluded by local δ -ISS since an expert policy cannot simultaneously stabilize to multiple branches of a bifurcation. For a concrete example, consider Figure 8. Indeed, continuity is the *sine*

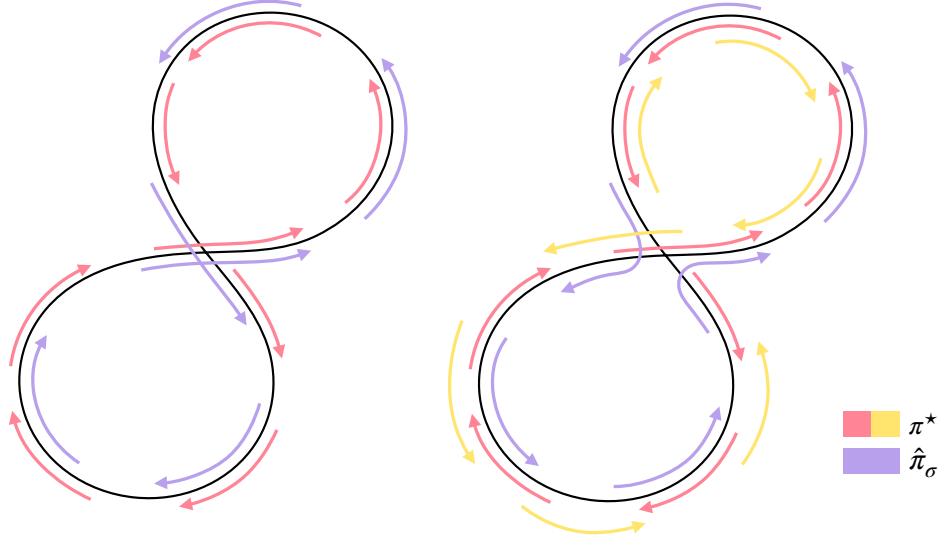


Figure 9: Instance where $\hat{\pi}_\sigma$ and π^* induce the same marginals and joint distributions (left), but in the presence of expert demonstration trajectories that traverse the figure eight both clockwise and counterclockwise directions, $\hat{\pi}_\sigma$ may switch with some probability between demonstrations where they overlap.

qua non of stability and the example given demonstrates the necessity of augmentation to enforce the former. In detail, the figure illustrates an example where an agent is navigating around an obstacle, providing a bifurcation. Without augmentation, the demonstrator trajectories always navigate around the obstacle in the direction closer to their starting point, leading to a sharp discontinuity along a bisector of the obstacle. On the other hand, the data augmentations allow for the policy to have some probability of navigating around the obstacle in the “wrong” direction, which leads to the notion of continuity we consider: total variation continuity.

Because our notion of stability is applied in chunks, our theory is sufficiently flexible so as to allow for the learned policy to switch between expert demonstrations in a manner preserving the marginal distributions but not consistent with the joint distribution across the entire trajectory. This flexibility is illustrated in Figure 9, where we suppose that the demonstrator distribution consists both of trajectories traversing a figure “8” consistently in either a clockwise or counter-clockwise manner, with both orientations represented in the data set. Due to the multi-modality at the critical point in the trajectory, there is ambiguity about which loop to traverse next; specifically, there may exist a policy that randomly select which loop to traverse each time the critical point is visited in such a way that the marginal distributions on states and actions is the same as that induced by the demonstrator. Such a policy will, by definition, preserve the correct *marginal* distributions across states and actions; at the same time, this policy has a different *joint* distribution across all time steps from the demonstrator due to the possibility of traversing the same loop twice in a row.

Part II

Composite MDP

F Measure-Theoretic Background

In this section, we introduce the prerequisite notions from probability theory that we use to formally construct the couplings in [Appendices G and H](#). We begin by introducing general preliminaries, followed by kernels, regular conditional probabilities and a “gluing” lemma in [Appendix F.1](#). We then show that optimal transport costs commute in an appropriate sense with conditional probabilities ([Proposition F.3](#) in [Appendix F.2](#)). We use the preliminaries in the previous sections to derive certain optimal-transport and data processing inequalities in [Appendix F.3](#). We prove [Proposition F.3](#) in [Appendix F.4](#). Finally, we state a simple union bound lemma ([Lemma F.11](#) in [Appendix F.5](#)) of use in later appendices.

General preliminaries. We rely extensively on the exposition in Durrett [22] and refer the reader there for a more thorough introduction. Throughout, we assume there is a Polish space Ω such that all random variables of interest are mappings $X : \Omega \rightarrow \mathcal{X}$, where \mathcal{X} is also Polish. Here, the σ -algebras are always the Borel algebras (the σ -algebra generated by open subsets), denoted $\mathcal{B}(\Omega)$ and $\mathcal{B}(\mathcal{X})$.

The space of (Borel) probability distributions on \mathcal{X} is denoted $\Delta(\mathcal{X})$, and measurability is meant in the Borel sense. Given a measure μ on a space $\mathcal{X} \times \mathcal{Y}$, we say that $X \sim P_X$ under μ if, for all $A \in \mathcal{B}(\mathcal{X})$, $\mu(X \in A) = P_X(A)$.

We adopt standard information theoretic notation to denote joint, marginal, and conditional distributions on vectors of random variables. In particular, if random variables X, Y are distributed according to P , we denote by P_X as the marginal over X , $P_{X|Y}$ as the conditional of $X|Y$ under P , and $P_{X,Y}$ as the joint distribution when this needs to be emphasized.

Definition F.1 (Couplings). Let \mathcal{X}, \mathcal{Y} be Polish spaces and let $P_X \in \Delta(\mathcal{X})$ and $P_Y \in \Delta(\mathcal{Y})$. The set of couplings $\mathcal{C}(P_X, P_Y)$ denotes the set of measure $\mu \in \Delta(\mathcal{X} \times \mathcal{Y})$ such that, $(X, Y) \sim \mu$ has marginals $X \sim P_X$ and $Y \sim P_Y$.⁷ We let $P_X \otimes P_Y \in \mathcal{C}(P_X, P_Y)$ denote the *independent coupling* under which X and Y are independent.

It is standard that $P_X \otimes P_Y$ is always a valid coupling, and hence $\mathcal{C}(P_X, P_Y)$ is nonempty. Couplings have the advantage that they can be used to design many probability-theoretic distances. Through the paper, we use the total variation distance.

Definition F.2 (Total Variation Distance). Let $P_1, P_2 \in \Delta(\mathcal{X})$. We define the total variation distance $\text{TV}(P_1, P_2) := \sup_{A \in \mathcal{B}(\mathcal{X})} |P_1(A) - P_2(A)|$

The total variation distance can be expressed in terms of couplings as follows [52].

Lemma F.1. Let $P_1, P_2 \in \Delta(\mathcal{X})$. Then,

$$\text{TV}(P_1, P_2) = \inf_{\mu \in \mathcal{C}(P_1, P_2)} \mathbb{P}_{(X_1, X_2) \sim \mu} \{X_1 \neq X_2\}.$$

Moreover, there exists a coupling μ_* attaining the infimum.

Support and absolute continuity. We will also require the definition of the support of a measure.

Definition F.3. Given a measure μ on a Borel space (Ω, \mathcal{F}) , we define the *support* $\text{supp}(\mu)$ to be the closure in the topology given by the metric of the set $\{\omega \in \Omega | \mu(U) > 0 \text{ for all open } U \ni \omega\}$.

In addition, we require the definition of absolute continuity.

Definition F.4 (Absolute Continuity). We say that $P \in \Delta(\mathcal{X})$ is absolutely continuous with respect to law $P' \in \Delta(\mathcal{X})$, written $P \ll P'$, if for $A \in \mathcal{B}(\mathcal{X})$, $P'(A) = 0$ implies $P(A) = 0$.

We now go into greater detail on the kinds of couplings that we consider.

⁷More pedantically, for all Borel sets $A_1 \in \mathcal{B}(\mathcal{X})$, $\mu(A_1 \times \mathcal{Y}) = P_X(A_1)$ all Borel sets $A_2 \in \mathcal{B}(\mathcal{X})$, $\mu(\mathcal{X} \times A_2) = P_Y(A_2)$.

F.1 Kernels, Regular Conditional Probabilities and Gluing

One key technical challenge in proving results in the sequel is the fact that we need to “glue” together multiple different couplings. Specifically, while it may be the case that there exist pairwise couplings which satisfy desired properties, there exists a coupling such that the probability of the relevant event is small, it is not obvious that there exists a *single* coupling such that all of these probabilities are small *simultaneously*. There are two natural ways to do this gluing: the first, using regular conditional probabilities we provide here. The second, involving a sophisticated construction of Angel and Spinka [8] requires stronger assumptions on the pseudo-metric, but generalizing beyond Polish spaces, we simply remark can be substituted with a loss of a constant factor.

Kernels. We begin by introducing the notion of a kernel.

Definition F.5 (Kernels). Let (Ω, \mathbb{P}) be a probability space and let X denote a random variable on this space. For a given σ -algebra \mathcal{G} , and map $Q : \Omega \times \mathcal{G} \rightarrow [0, 1]$, we say that Q is a probability kernel if the following two conditions are satisfied:

1. For all measurable events A , the map $\omega \mapsto Q(\omega, A)$ is measurable.
2. For almost every $\omega \in \Omega$, the map $A \mapsto Q(\omega, A)$ is a probability measure.

We can combine a probability kernel with a probability measure on \mathcal{Y} to yield joint distributions over $\mathcal{X} \times \mathcal{Y}$.

Definition F.6. Given an $P_Y \in \Delta(\mathcal{Y})$, we define the probability measure $\text{law}(Q_{X|Y}; P_Y) \in \Delta(\mathcal{X} \times \mathcal{Y})$ such that $\mu = \text{law}(Q_{X|Y}; P_Y)$ satisfies⁸

$$\mu(A \times B) = \mathbb{E}_{Y \sim P_Y} [Q_{X|Y}(A | Y) \mathbf{I}\{Y \in B\}], \quad \forall A \in \mathcal{B}(\mathcal{X}), B \in \mathcal{B}(\mathcal{Y}). \quad (\text{F.1})$$

We let $Q_{X|Y} \circ P_Y \in \Delta(\mathcal{X})$ denote the measure for which $\mu = Q_{X|Y} \circ P_Y$ satisfies

$$\mu(A) = \mathbb{E}_{Y \sim P_Y} [Q_{X|Y}(A | Y)], \quad \forall A \in \mathcal{B}(\mathcal{X})$$

From these, we define the space of conditional couplings as follows.

Definition F.7 (Kernel Couplings). Let $P_Y \in \Delta(\mathcal{Y})$, and $Q_{X_i|Y} \in \Delta(\mathcal{X} | \mathcal{Y})$ for $i \in \{1, 2\}$. We let $\mathcal{C}_{P_Y}(Q_{X_1|Y}, Q_{X_2|Y})$ denote the space of measures $\mu \in \Delta(\mathcal{X}_1 \times \mathcal{X}_2 \times \mathcal{Y})$ over random variables (X_1, X_2, Y) such that $(X_i, Y) \sim \text{law}(Q_{X_i|Y}; P_Y)$ for $i \in \{1, 2\}$.

Note that a similar construction to the independent coupling ensures $\mathcal{C}_{P_Y}(Q_{X_1|Y}, Q_{X_2|Y})$ is nonempty, namely considering the measure $\mu(A_1 \times A_2 \times B_2) = \mathbb{E}_{Y \sim P_Y} [Q_{X_1|Y}(A_1 | Y) Q_{X_2|Y}(A_2 | Y) \mathbf{I}\{Y \in B\}]$.

Regular Conditional Probabilities. We now recall a standard result that conditional probabilities can be expressed through kernels in our setting.

Theorem 7 (Theorem 5.1.9, Durrett [22]). If Ω is a Polish space and \mathbb{P} is a probability measure on the Borel sets of Ω , such that random variables $(X, Y) \sim \mathbb{P}$ in spaces \mathcal{X} and \mathcal{Y} , then there exists a kernel $Q(\cdot | \cdot) \in \Delta(\mathcal{X} | \mathcal{Y})$ such that, for all $A \in \mathcal{B}(\mathcal{X})$ and \mathbb{P} -almost every y , the (standard) conditional probability $\mathbb{P}[X \in A | Y] = Q(A | y)$. We can $Q(\cdot | \cdot)$ the *regular conditional probability measure*.

Regular conditional probabilities allow one to think of conditional probabilities in the most intuitive way, i.e., for two random variables X, Y , the map $Y \mapsto \mathbb{P}(X \in A | Y)$ is a probability kernel. This will be the essential property that we use below.

Gluing. Finally, regular conditional probabilities allow us to “glue together” couplings which share a common random variable.

⁸Recall that $\mathcal{B}(\mathcal{X} \times \mathcal{Y})$ is generated by sets $A \times B \in \mathcal{B}(\mathcal{X}) \times \mathcal{B}(\mathcal{Y})$, so (F.1) defines a unique probability measure

Lemma F.2 (Gluing Lemma). Suppose that X, Y, Z are random variables taking value in Polish spaces $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$. Let $\mu_1 \in \Delta(\mathcal{X} \times \mathcal{Y}), \mu_2 \in \Delta(\mathcal{Y} \times \mathcal{Z})$ be couplings of (X, Y) and (Y, Z) respectively. Then there exists a coupling $\mu \in \Delta(\mathcal{X} \times \mathcal{Y} \times \mathcal{Z})$ on (X, Y, Z) such that under $\mu, (X, Y) \sim \mu_1$ and $(Y, Z) \sim \mu_2$.

Proof. Let $Q(\cdot | Y)$ be a regular conditional probability for Z given Y under μ_2 (whose existence is ensured by [Theorem 7](#)).

We construct μ by first sampling $(X, Y) \sim \mu_1$ and then sampling $Z \sim Q(\cdot | Y)$; observe that by the second property in [Definition F.5](#), this is a valid construction. It is immediate that under μ , we have $(X, Y) \sim \mu_1$ and thus we must only show that $(Y, Z) \sim \mu_2$ to conclude the proof. Let A, B be two measurable sets and we see that

$$\begin{aligned} \mathbb{P}_\mu((Y, Z) \in A \times B) &= \mathbb{E}_{Y \sim \mu} [\mathbb{P}_\mu((Y, Z) \in A \times B | Y)] \\ &= \mathbb{E}_{Y \sim \mu} [\mathbb{E}_{(Y, Z) \sim \mu} [\mathbf{I}[Y \in A] \cdot \mathbf{I}[Z \in B] | Y]] \\ &= \mathbb{E}_{Y \sim \mu} [\mathbf{I}[Y \in A] \cdot \mathbb{E}_\mu [\mathbf{I}[Z \in B] | Y]] \\ &= \mathbb{E}_{Y \sim \mu} [\mathbf{I}[Y \in A] \cdot \mathbb{P}_{\mu_2}(Z \in B | Y)] \\ &= \mu_2((Y, Z) \in A \times B), \end{aligned}$$

where the first equality follows from the tower property of expectations, the second follows by definition of conditional probability, the third follows from the definition of conditional expectation, the fourth follows by the first property from [Definition F.5](#), and the last follows from the fact that the marginals of Y under μ and under μ_2 are the same. The result follows. \square

F.2 Optimal Transport and Kernel Couplings

As shown above for the TV distance, many measures of distributional distance can be quantified in terms of *optimal transport* costs; these are quantities expressed as infima, over all couplings, of the expectation of a certain lower-semicontinuous functions. We show that if the optimal transport costs between two kernels $Y \rightarrow \Delta(\mathcal{X}_i)$ are controlled pointwise, then for any $P_Y \in \Delta(\mathcal{Y})$, there exists a joint distribution over (X_1, X_2, Y) which attains the minimal transport cost.

Proposition F.3. Let $\mathcal{X}_1, \mathcal{X}_2, \mathcal{Y}$ be Polish spaces, and let $P_Y \in \Delta(\mathcal{Y})$, and $Q_i \in \Delta(\mathcal{X}_i | \mathcal{Y})$. for $i \in \{1, 2\}$. Finally, let $\phi : \mathcal{X}_1 \times \mathcal{X}_2 \rightarrow \mathbb{R}$ be lower semicontinuous and bounded below. Then, the following function

$$\psi(y) := \inf_{\mu \in \mathcal{C}(Q_1(y), Q_2(y))} \mathbb{E}_{(X_1, X_2) \sim \mu} [\phi(X_1, X_2)]$$

is a measurable function of y and there exists some $\mu_\star \in \mathcal{C}_{P_Y}(Q_1, Q_2)$ such that

$$\mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} [\phi(X_1, X_2)] = \mathbb{E}_{Y \sim P_Y} \psi(Y).$$

In particular it holds μ_\star -almost surely that

$$\mathbb{E}_{\mu_\star} [\phi(X_1, X_2) | Y] = \psi(Y).$$

We prove the above proposition in [Appendix F.4](#). One useful consequence is the following identity for the total variation distance.

Corollary F.1. Let \mathcal{X}, \mathcal{Y} be Polish spaces, and let $P_Y \in \Delta(\mathcal{Y})$, and $Q_i \in \Delta(\mathcal{X} | \mathcal{Y})$, for $i \in \{1, 2\}$. Then, there exists a coupling $\mu_\star \in \mathcal{C}_{P_Y}(Q_1, Q_2)$ such that

$$\mathbb{P}_{\mu_\star}[X_1 \neq X_2] = \mathbb{E}_{Y \sim P_Y} \text{TV}(Q_1(\cdot | Y), Q_2(\cdot | Y)),$$

with the left-hand side integrand being measurable.

Proof. Using [Lemma F.1](#), we can represent total variation as an optimal transport cost with $\phi(x_1, x_2) = \mathbf{I}\{x_1 \neq x_2\}$. Note that $\phi(x_1, x_2)$ is lower semicontinuous, being the indicator of an open set. Thus, the result follows from [Proposition F.3](#) with $\mathcal{X} = \mathcal{X}_1 = \mathcal{X}_2$, and $\phi(x_1, x_2) = \mathbf{I}\{x_1 \neq x_2\}$. \square

F.3 Data Processing Inequalities

We now derive two *inequalities*. First, we recall the classical version for the total variation distance, and check that a well-known identity holds in our setting.

Lemma F.4 (Data Processing for Total Variation). Let $P_{Y_1}, P_{Y_2} \in \Delta(\mathcal{Y})$ and let $Q_X \in \Delta(\mathcal{X} \mid \mathcal{Y})$. Then,

$$\text{TV}(Q_X \circ P_{Y_1}, Q_X \circ P_{Y_2}) \leq \text{TV}(\text{law}(Q_X; P_{Y_1}), \text{law}(Q_X; P_{Y_2})) = \text{TV}(P_{Y_1}, P_{Y_2}).$$

Proof. The first inequality is just the data processing inequality [52, Theorem 7.7], which also shows that $\text{TV}(\text{law}(Q_X; P_{Y_1}), \text{law}(Q_X; P_{Y_2})) \geq \text{TV}(P_{Y_1}, P_{Y_2})$. To prove the reverse inequality, we use Lemma F.1 to find a coupling μ_Y such that (P_{Y_1}, P_{Y_2}) such that $\mathbb{E}[\mathbf{I}\{Y_1 \neq Y_2\}] = \text{TV}(Y_1, Y_2)$.

Define a probability kernel in $\Delta(\mathcal{X} \times \mathcal{X} \mid \mathcal{Y}_1 \times \mathcal{Y}_2)$ via defining the set $B = \{(x_1, x_2) \in \mathcal{X} \times \mathcal{X} : x_1 = x_2\} \subset \mathcal{X} \times \mathcal{X}$, and define for $A \in \mathcal{B}(\mathcal{X} \times \mathcal{X})$,

$$Q(A \mid y_1, y_2) = \begin{cases} Q_X(\pi_1(A \cap B) \mid y_1) & y_1 = y_2 \\ Q_X(\cdot \mid y_1) \otimes Q_X(\cdot \mid y_2)(A) & \text{otherwise} \end{cases}$$

In a Polish space, Lemmas F.6 and F.7 imply that $A \mapsto Q_X(\pi_1(A \cap B) \mid y_1)$ for each y_1 is a valid measure, and it is standard that the product measures $Q_X(\cdot \mid y_1) \otimes Q_X(\cdot \mid y_2)(A)$ are valid. Moreover, this construction ensures that for $\mu = \text{law}(Q; \mu_Y)$,

$$\mathbb{P}_\mu[\{Y_1 = Y_2\} \text{ and } \{X_1 \neq X_2\}] = 0. \quad (\text{F.2})$$

Lastly, one can check that under $\mu = \text{law}(Q; \mu_Y)$, that $(X_1, Y_1) \sim \text{law}(Q_X; P_{Y_1})$ and $(X_2, Y_2) \sim \text{law}(Q_X; P_{Y_2})$. Thus, μ can be regarded as an element of $\mathcal{C}(\text{law}(Q_X; P_{Y_1}), \text{law}(Q_X; P_{Y_2}))$. Hence, Lemma F.1 implies that

$$\begin{aligned} \text{TV}(\text{law}(Q_X; P_{Y_1}), \text{law}(Q_X; P_{Y_2})) &\leq \text{TV}(\mathbb{P}_\mu[(X_1, Y_1) \neq (X_2, Y_2)]) \\ &= \mathbb{P}_\mu[Y_1 \neq Y_2] + \mathbb{P}_\mu[\{Y_1 = Y_2\} \text{ and } \{X_1 \neq X_2\}] \\ &= \mathbb{P}_{\mu_*}[Y_1 \neq Y_2] \quad (\text{Eq. (F.2)}) \\ &= \mathbb{P}_{(Y_1, Y_2) \sim \mu_Y}[Y_1 \neq Y_2] \\ &= \text{TV}(P_{Y_1}, P_{Y_2}). \quad (\text{construction of } \mu_Y) \end{aligned}$$

□

Next, we derive a general data processing inequality for optimal costs. This result is a corollary of Proposition F.3.

Lemma F.5 (Another Data Processing Inequality for Optimal Transport). Let $\mathcal{X}_1, \mathcal{X}_2, \mathcal{Y}$ be Polish spaces, and let $P_Y \in \Delta(\mathcal{Y})$, and $Q_i \in \Delta(\mathcal{Y} \mid \mathcal{X}_i)$. for $i \in \{1, 2\}$. Denote by $Q_i \circ P_Y$ the marginal of X_i under $(X_i, Y) \sim \text{law}(Q_i; P_Y)$. Then,

$$\inf_{\mu \in \mathcal{C}(Q_1 \circ P_Y, Q_2 \circ P_Y)} \mathbb{E}_{X_1, X_2 \sim \mu} \phi(X_1, X_2) \leq \mathbb{E}_{Y \sim \mu_Y} \left(\inf_{\mu' \in \mathcal{C}(Q_1(Y) \circ Q_2(Y))} \mathbb{E}_{X_1, X_2 \sim \mu'} \phi(X_1, X_2) \right).$$

Proof. One can check that any coupling in $\mu \in \mathcal{C}(Q_1 \circ P_Y, Q_2 \circ P_Y)$ can be obtained by marginalizing Y in a certain coupling of $\mu' \in \mathcal{C}(\text{law}(Q_1; P_Y), \text{law}(Q_2; P_Y))$, and any coupling in the latter can be marginalized to a coupling in the former. Hence,

$$\inf_{\mu \in \mathcal{C}(Q_1 \circ P_Y, Q_2 \circ P_Y)} \mathbb{E}_{X_1, X_2 \sim \mu} \phi(X_1, X_2) = \inf_{\mu \in \mathcal{C}(\text{law}(Q_1; P_Y), \text{law}(Q_2; P_Y))} \mathbb{E}_{X_1, X_2, Y_1, Y_2 \sim \mu} \phi(X_1, X_2)$$

Moreover, to every measure $\mu \in \mu_{P_Y}(Q_1, Q_2)$ over (X_1, X_2, Y) , Lemma F.8 implies that there exists a coupling $\mu' \in \mathcal{C}(\text{law}(Q_1; P_Y), \text{law}(Q_2; P_Y))$ over (X_1, X_2, Y_1, Y_2) such (X_1, X_2) have the same marginals under μ and μ' . Therefore,

$$\inf_{\mu \in \mathcal{C}(\text{law}(Q_1; P_Y), \text{law}(Q_2; P_Y))} \mathbb{E}_{X_1, X_2, Y_1, Y_2 \sim \mu} \phi(X_1, X_2) \leq \inf_{\mu' \in \mathcal{C}_{P_Y}(Q_1, Q_2)} \mathbb{E}_{X_1, X_2, Y \sim \mu} \phi(X_1, X_2).$$

Finally, the right hand side is equal to $\mathbb{E}_{Y \sim \mu_Y} (\inf_{\mu' \in \mathcal{C}(Q_1(Y) \circ Q_2(Y))} \mathbb{E}_{X_1, X_2 \sim \mu'} \phi(X_1, X_2))$ by Proposition F.3. □

F.3.1 Deferred lemmas for the data processing inequalities

Lemma F.6. Let \mathcal{X} be a Polish space. Then, the set $\{(x_1, x_2) \in \mathcal{X} \times \mathcal{X} : x_1 \neq x_2\}$ is open in $\mathcal{X} \times \mathcal{X}$.

Proof. The diagonal is closed in any Polish space by definition of the topology. The result follows. \square

Lemma F.7. Let \mathcal{X} be a Polish space, and let $\pi_1, \pi_2 : \mathcal{X} \times \mathcal{X}$ denote the projection mappings onto each coordinate. Then, for any $A \in \mathcal{B}(\mathcal{X} \times \mathcal{X})$, $\pi_1(A)$ and $\pi_2(A)$ are in $\mathcal{B}(\mathcal{X})$.

Proof. The projection map is open so the result follows immediately by definition of the Borel algebra. \square

Lemma F.8. Let \mathcal{X}, \mathcal{Y} be Polish spaces, and let $\mu \in \Delta(\mathcal{X} \times \mathcal{Y})$. Then, there is a measure $\mu' \in \Delta(\mathcal{X} \times \mathcal{Y} \times \mathcal{Y})$ satisfying

$$\mu'(A \times \mathcal{Y}) = \mu(A), \quad \forall A \in \mathcal{B}(\mathcal{X} \times \mathcal{Y})$$

and

$$\mu'(\mathcal{X} \times \{(y_1, y_2) : y_1 = y_2\}) = 1$$

Proof. Define the set $B_+ = \{(y_1, y_2) : y_1 = y_2\}$. One can check that $\mu'(A \times B) = \mu(A \times \pi_1(B \cap B_+))$, where $\pi_1 : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{Y}$ is the projection onto the first coordinate, extends to a valid measure. \square

F.4 Proof of Proposition F.3

In the case that $\phi(\cdot, \cdot)$ is continuous, the result follows from Villani et al. [73, Corollary 5.22]. For general lower-semicontinuous ϕ , our argument adopts the strategy of “Step 3” of the proof of Villani [72, Theorem 1.3]. This shows that there exists a sequence $\phi_n \uparrow \phi$ pointwise, such that each ϕ_n is uniformly bounded. Define

$$\psi_n(y) := \inf_{\mu \in \mathcal{C}(\mathbf{Q}_1(y), \mathbf{Q}_2(y))} \mathbb{E}_{(X_1, X_2) \sim \mu} [\phi_n(X_1, X_2)].$$

Then, for each n , the continuous case implies that there exists a measure $\mu_{*,n} \in \mathcal{C}_{\nu_Y}(\mathbf{Q}_1, \mathbf{Q}_2)$ such that

$$\mathbb{E}_{Y \sim \nu_Y} \psi_n(Y) = \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{*,n}} [\phi_n(X_1, X_2)] \quad (\text{F.3})$$

Recall now the definition

$$\psi(y) = \inf_{\mu \in \mathcal{C}(\mathbf{Q}_1(y), \mathbf{Q}_2(y))} \mathbb{E}_{(X_1, X_2) \sim \mu} [\phi(X_1, X_2)].$$

Claim F.9. $\psi(y)$ is measurable and satisfies $\psi_n(y) \uparrow \psi(y)$ pointwise.

Proof. We can write

$$\begin{aligned} \sup_{n \geq 0} \psi_n(y) &= \sup_{n \geq 0} \inf_{\mu \in \mathcal{C}(\mathbf{Q}_1(y), \mathbf{Q}_2(y))} \mathbb{E}_{(X_1, X_2) \sim \mu} [\phi_n(X_1, X_2)] \\ &\stackrel{(i)}{=} \inf_{\mu \in \mathcal{C}(\mathbf{Q}_1(y), \mathbf{Q}_2(y))} \mathbb{E}_{(X_1, X_2) \sim \mu} [\phi(X_1, X_2)] = \psi(y). \end{aligned}$$

Here, (i) follows from the “Step 3” in the proof of Villani [72, Theorem 1.3], which shows that any optimal transport cost C of a lowersemicontinuous ϕ is equal to a limit of the costs C_n of any bounded continuous $\phi_n \uparrow \phi$. In our case, we fix each y , so $C = \psi(y)$ and $C_n = \psi_n(y)$. It is clear that $\psi_n(y)$ is increasing, so for each y , $\psi_n(y) \uparrow \psi(y)$. As ψ is the pointwise monotone limit of ψ_n , it is measurable. \square

Claim F.10. The set of couplings of $\mathcal{C}_{\mathbf{P}_Y}(X_1, X_2)$ is compact in the weak topology.

Proof. Recall that $\Delta(\mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2)$ denote the set of Borel measures on $\mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2$. This set is also a Polish space in the weak topology. The subset $\mathcal{C}_{P_Y}(X_1, X_2) \subset \Delta(\mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2)$ is compact if and only if it is relatively compact and closed.

To show relative compactness, Prokhorov's theorem means that it suffices to show that $\mu_{P_Y}(Q_1, Q_2)$ is tight, i.e. for all $\varepsilon > 0$, there exists a compact $\mathcal{K}_\varepsilon \subset \mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2$ such that for any $\mu \in \mathcal{C}_{P_Y}(X_1, X_2)$, $\mathbb{P}_\mu[(Y, X_1, X_2) \in \mathcal{K}_\varepsilon] \geq 1 - \varepsilon$. This follows by setting $\mathcal{K} = \mathcal{K}_{Y,\varepsilon} \times \mathcal{K}_{X_1,\varepsilon} \times \mathcal{K}_{X_2,\varepsilon}$, where the sets are such that $\mathbb{P}_{P_Y}[Y \notin \mathcal{K}_{Y,\varepsilon}] \geq 1 - \varepsilon/3$ and $\mathbb{P}_{Q_i}[X_i \notin \mathcal{K}_{X_i,\varepsilon}] \geq 1 - \varepsilon/3$, where Q_i is the marginal of X_i given by $Y \sim P_Y$, $X_i \sim P_i(\cdot | Y)$ (such sets exist because $\mathcal{X}_1, \mathcal{X}_2, \mathcal{Y}$ are Polish).

To check that $\mathcal{C}_{P_Y}(Q_1, Q_2) \subset \Delta(\mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2)$ is closed, it suffices to show that it is sequentially closed (as $\Delta(\mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2)$ is Polish). To this end, consider any sequence $\mu_n \in \mathcal{C}_{P_Y}(Q_1, Q_2)$ such that $\mu_n \xrightarrow{\text{weak}} \mu \in \Delta(\mathcal{Y} \times \mathcal{X}_1 \times \mathcal{X}_2)$ in the weak topology. By definition, this means that for any $i \in \{1, 2\}$ and any continuous and bounded $f_i : \mathcal{Y} \times \mathcal{X}_i \rightarrow \mathbb{R}$,

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mu_n} f_i(Y, X_i) = \mathbb{E}_\mu f_i(Y, X_i).$$

For all $\mu_n \in \mathcal{C}_{P_Y}(Q_1, Q_2)$, $\mathbb{E}_{\mu_n} f_i(Y, X_i) = \mathbb{E}_{Y \sim \nu_Y} \mathbb{E}_{X_i \sim \nu_i(\cdot | Y_i)} f_i(Y, X_i)$. Thus,

$$\mathbb{E}_\mu f_i(Y, X_i) = \mathbb{E}_{Y \sim \nu_Y} \mathbb{E}_{X_i \sim \nu_i(\cdot | Y_i)} f_i(Y, X_i), \quad \text{for all continuous, bounded } f_i : \mathcal{Y} \times \mathcal{X}_i \rightarrow \mathbb{R}.$$

Hence, the marginal distribution of (Y, X_i) under μ must be equal to that of $(Y \sim P_Y, X_i \sim Q_i(\cdot | Y))$ for $i \in \{1, 2\}$, which means $\mu \in \mathcal{C}_{P_Y}(Q_1, Q_2)$. \square

By compactness, there exists (passing to a subsequence if necessary) a $\mu_\star \in \mathcal{C}_{P_Y}(Q_1, Q_2)$ such that $\mu_{\star,n} \xrightarrow{\text{weak}} \mu_\star$ in the weak topology. Then, as ϕ_m is continuous and bounded, it follows that for all m ,

$$\begin{aligned} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} [\phi_m(X_1, X_2)] &= \limsup_{n \rightarrow \infty} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,n}} [\phi_m(X_1, X_2)] && (\mu_{\star,n} \xrightarrow{\text{weak}} \mu_\star) \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,n}} [\phi_n(X_1, X_2)] && (\phi_m \leq \phi_n \text{ for } n \geq m) \\ &= \limsup_{n \rightarrow \infty} \mathbb{E}_Y \psi_n(Y) && ((F.3)) \\ &= \mathbb{E}_Y \lim_{n \rightarrow \infty} \psi_n(Y) && (\text{Monotone Convergence}) \\ &= \mathbb{E}_Y \psi(Y). && (\text{Claim F.9}) \end{aligned}$$

Thus, by the monotone convergence theorem,

$$\begin{aligned} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} [\phi(X_1, X_2)] &= \mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} \left[\lim_{m \rightarrow \infty} \phi_m(X_1, X_2) \right] \\ &= \lim_{m \rightarrow \infty} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} [\phi_m(X_1, X_2)] \\ &\leq \lim_{m \rightarrow \infty} \mathbb{E}_Y \psi(Y) = \mathbb{E}_Y \psi(Y). \end{aligned}$$

Similarly, repeating some of the above steps,

$$\begin{aligned} \mathbb{E}_Y \psi(Y) &= \limsup_{n \rightarrow \infty} \mathbb{E}_Y \psi_n(Y) \\ &= \limsup_{n \rightarrow \infty} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,n}} [\phi_n(X_1, X_2)] \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,m}} [\phi_n(X_1, X_2)] && (\mu_{\star,n} \text{ is the optimal coupling for } \phi_n) \\ &\leq \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,m}} \left[\lim_{n \rightarrow \infty} \phi_n(X_1, X_2) \right] && (\text{monotone convergence}) \\ &\leq \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,m}} [\phi(X_1, X_2)]. \end{aligned}$$

Hence, $\mathbb{E}_Y \psi(Y) \leq \liminf_{m \geq 1} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,m}} [\phi(X_1, X_2)]$. By assumption, $\phi(X_1, X_2)$ is lower semicontinuous and bounded from below. Thus, the Portmanteau theorem [22] implies that, as $\mu_{\star,m} \xrightarrow{\text{weak}} \mu_\star$,

$$\liminf_{m \geq 1} \mathbb{E}_{(X_1, X_2, Y) \sim \mu_{\star,m}} [\phi(X_1, X_2)] = \mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} [\phi(X_1, X_2)].$$

Hence, $\mathbb{E}_Y \psi(Y) \leq \mathbb{E}_{(X_1, X_2, Y) \sim \mu_\star} [\phi(X_1, X_2)]$, proving the reverse inequality.

Proof of the last statement. To prove the last statement, we observe that if $\mu_* \in \mathcal{C}_{P_Y}(Q_1, Q_2)$ then there exists a version of $(\mu_*)_{X, X'|Y}$ that is a regular conditional probability and such that for almost every y it holds that $(\mu_*)_{X, X'|y} \in \mathcal{C}(Q_1(y), Q_2(y))$. Indeed, the existence of a version that is a regular conditional probability is immediate by [Theorem 7](#). To see that this version is a valid coupling of $Q_1(y)$ and $Q_2(y)$, observe that under μ_* , the joint law of $(X, Y) \sim Q_1$ and thus the conditional distribution under μ_* of $X|Y$ is determined up to sets of Q_1 -measure 0. In particular, again by [Theorem 7](#), there exists a regular conditional probability that is a version of $(\mu_*)_{X|y}$ and this must agree almost everywhere with $(Q_1)_{X|y} = Q_1(y)$. The same argument holds for X' and thus $(\mu_*)_{X, X'|y} \in \mathcal{C}(Q_1(y), Q_2(y))$ for almost every y . Thus, by definition of ψ as an infimum, it holds for almost every y that

$$\psi(y) \leq \mathbb{E}_{(X, X') \sim (\mu_*)|Y} [\phi(X, X')].$$

By the second claim of the proposition, we also have that

$$\mathbb{E}_{\mu_*} [\phi(X_1, X_2)] = \mathbb{E}_{\mu_*} [\psi(Y)].$$

Because the expectations are equal and one function is pointwise almost everywhere dominated by the other function, the two functions must be equal almost everywhere, concluding the proof. \square

F.5 A simple union-bound recursion.

Finally, we also use the following version of the union bound extensively in our recursion proofs.

Lemma F.11. For any event \mathcal{E} and events $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_H$, it holds that

$$\mathbb{P}[(\mathcal{Q} \cap \bigcap_{h=1}^H \mathcal{B}_h)^c] \leq \mathbb{P}[\mathcal{Q}^c] + \mathbb{P} \left[\exists h \in [H] \text{ s.t. } \left(\mathcal{Q} \cap \bigcap_{j=1}^{h-1} \mathcal{B}_j \cap \mathcal{B}_h^c \right) \text{ holds} \right]$$

Proof. Note that

$$\left(\mathcal{Q} \cap \bigcap_{h=1}^H \mathcal{B}_h \right)^c = \mathcal{Q}^c \cup \left(\mathcal{Q} \cap \left(\bigcap_{h=1}^H \mathcal{B}_h \right)^c \right) = \mathcal{Q}^c \cup \bigcup_{h=1}^H \mathcal{Q} \cap \mathcal{B}_h \cap \bigcap_{j=1}^{h-1} \mathcal{B}_j^c.$$

The result follows by a union bound. \square

G Warmup: Analysis Without Augmentation

In this section, we give a simplified analysis that replaces the smoothing kernels W_σ with the assumption that the learner policy $\hat{\pi}$ is already total variation continuous. The removal of the coupling kernel makes the coupling construction considerably simpler while still communicating some intuition for the full proof in [Appendix H](#).

Throughout this section, we make the following assumptions on the state and action spaces, along with their associated metrics:

Assumption G.1. We assume that \mathcal{S} and \mathcal{A} are Polish spaces. This means they are metrizable, but we do not annotate their metrics because, e.g. the metric on \mathcal{S} may be other than d_S . We further assume that

- d_S, d_{TVC} are pseudometrics and Borel measurable function from $\mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$
- For any $\varepsilon \geq 0$, the set $\{(a, a') \in \mathcal{A} \times \mathcal{A} : d_{\mathcal{A}}(a, a') > \varepsilon\}$ is an open subset of $\mathcal{A} \times \mathcal{A}$; i.e. $d_{\mathcal{A}}(\cdot, \cdot)$ is lower semicontinuous. In particular, this means $d_{\mathcal{A}}$ is a Borel measurable function.

Recall the definitions of total variation continuity (TVC) and input-stability in [Appendix D](#).

Proof. The key to the proof is to construct an appropriate “interpolating sequence” of actions $\hat{a}_{1:H}^{\text{inter}}$ to which we couple both $(s_{1:H+1}^*, a_{1:H}^*)$ and $(\hat{s}_{1:H+1}, \hat{a}_{1:H})$. This technique will be used in a significantly more sophisticated manner in the sequel to prove the analogous result with smoothing.

Let \mathcal{F}_h denote the σ -algebra generated by $(s_{1:h}^*, a_{1:h}^*)$, $(\hat{s}_{1:h}, \hat{a}_{1:h})$, and $\hat{a}_{1:h}^{\text{inter}}$, and let \mathcal{F}_0 denote the σ -algebra generated by s_1^*, \hat{s}_1 . We construct couplings of the following form:

- The initial states are generated as $\mathbf{s}_1^* = \hat{\mathbf{s}}_1 \sim P_{\text{init}}$.
- The dynamics are determined by F_h :

$$\mathbf{s}_{h+1}^* = F_h(\mathbf{s}_h^*, \mathbf{a}_h^*), \quad \hat{\mathbf{s}}_{h+1} = F_h(\hat{\mathbf{s}}_h, \hat{\mathbf{a}}_h) \quad (\text{G.1})$$

In particular, $\mathbf{s}_{h+1}^*, \hat{\mathbf{s}}_{1:h+1}$ are \mathcal{F}_h measurable.

- The conditional distributions of the primitive controllers satisfy the following

$$\mathbf{a}_h^* \mid \mathcal{F}_{h-1} \sim \pi_h^*(\mathbf{s}_h^*), \quad \hat{\mathbf{a}}_{h-1} \mid \mathcal{F}_{h-1} \sim \hat{\pi}_h(\hat{\mathbf{s}}_h), \quad \hat{\mathbf{a}}_h^{\text{inter}} \mid \mathcal{F}_h \sim \hat{\pi}_h(\mathbf{s}_h^*). \quad (\text{G.2})$$

Note that if μ satisfies the above construction, then $(\mathbf{s}_{1:H+1}^*, \mathbf{s}_{1:H}^*) \sim D_{\pi^*}$ and $(\hat{\mathbf{s}}_{1:H+1}, \hat{\mathbf{a}}_{1:H}) \sim D_{\hat{\pi}}$.

Specifying the rest of the coupling. It remains to specify the coupling of the terms in (G.2). We establish our coupling sequentially. Let $\mu^{(0)}$ denote the coupling of $\hat{\mathbf{s}}_1 = \mathbf{s}_1^* \sim P_{\text{init}}$.

Assume we have constructed the coupling up to state $h-1$. For ease, let Y_{h-1} denote the random variable corresponding to $(\mathbf{s}_{1:h}^*, \hat{\mathbf{s}}_{1:h}, \mathbf{a}_{1:h-1}^*, \hat{\mathbf{a}}_{1:h-1}, \hat{\mathbf{a}}_{1:h-1}^{\text{inter}})$; note that Y_{h-1} is \mathcal{F}_{h-1} -measurable (as $\hat{\mathbf{s}}_h, \mathbf{s}_h^*$ are determined by the dynamics (G.1)). Observe that, by the assumption of $\hat{\pi}_h$ being TVC, it holds that

$$\text{TV}(\mathbb{P}_{\hat{\mathbf{a}}_h \mid Y_{h-1}}, \mathbb{P}_{\hat{\mathbf{a}}_h^{\text{inter}} \mid Y_{h-1}}) \leq \gamma(\text{d}_{\text{TVC}}(\hat{\mathbf{s}}_h, \mathbf{s}_h^*)).$$

Thus by Lemma F.1, there exists a coupling $\mu_1^{(h)}$ between $Y_{h-1}, \hat{\mathbf{a}}_h, \hat{\mathbf{a}}_h^{\text{inter}}$, with $Y_{h-1} \sim \mu^{(h-1)}$ such that it holds that

$$\mathbb{P}[\hat{\mathbf{a}}_h \neq \hat{\mathbf{a}}_h^{\text{inter}}] \leq \mathbb{E}_{\mu^{(h-1)}}[\gamma(\text{d}_{\text{TVC}}(\hat{\mathbf{s}}_h, \mathbf{s}_h^*))].$$

Similarly by Proposition F.3, there is a coupling $\mu_2^{(h)}$ of $Y_{h-1}, \hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*$ such that

$$\mathbb{P}_{\mu_2^{(h)}}[\text{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*) > \varepsilon] \leq \mathbb{E}_{\mathbf{s}_h^* \sim \mu^{(h-1)}}[\text{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\mathbf{s}_h^*), \pi_h^*(\mathbf{s}_h^*))].$$

By the gluing lemma Lemma F.2 and a union bound, we may construct a coupling $\mu^{(h)}$ of $Y_h, \hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*, \hat{\mathbf{a}}_h$ such that (almost surely),

$$\begin{aligned} & \mathbb{P}_{\mu^{(h)}}[\{\text{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*) > \varepsilon\} \cup \{\hat{\mathbf{a}}_h \neq \hat{\mathbf{a}}_h^{\text{inter}}\} \mid \mathcal{F}_{h-1}] \\ &= \mathbb{P}_{\mu^{(h)}}[\{\text{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*) > \varepsilon\} \cup \{\hat{\mathbf{a}}_h \neq \hat{\mathbf{a}}_h^{\text{inter}}\} \mid Y_{h-1}] \\ &\leq \gamma(\text{d}_{\text{TVC}}(\hat{\mathbf{s}}_h, \mathbf{s}_h^*)) + \text{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\mathbf{s}_h^*), \pi_h^*(\mathbf{s}_h^*)) \end{aligned} \quad (\text{G.3})$$

Thus inductively, we may continue this construction for $h \leq H$ and let $\mu = \mu^{(H)}$.

Concluding the proof. Define the event $\mathcal{B}_h := \{\text{d}_{\mathcal{A}}(\mathbf{a}_h, \hat{\mathbf{a}}_h^{\text{inter}}) \leq \varepsilon\}$ and $\mathcal{C}_h = \{\hat{\mathbf{a}}_h^{\text{inter}} = \hat{\mathbf{a}}_h\}$. Then, by Lemma F.11

$$\mathbb{P}_{\mu} \left[\left(\bigcap_{h=1}^H \mathcal{B}_h \cap \mathcal{C}_h \right)^c \right] \leq \sum_{h=1}^H \mathbb{P}_{\mu} \left[\left(\bigcap_{j=1}^{h-1} \mathcal{B}_j \cap \mathcal{C}_j \right) \cap (\mathcal{B}_h^c \cup \mathcal{C}_h^c) \right]. \quad (\text{G.4})$$

Note first that $(\bigcap_{j=1}^{h-1} \mathcal{B}_j \cap \mathcal{C}_j)$ is \mathcal{F}_{h-1} measurable. On this event, input stability at $\hat{\mathbf{a}}_j^{\text{inter}} = \hat{\mathbf{a}}_j$, $1 \leq j \leq h-1$, implies that

$$\text{d}_{\mathcal{S}}(\mathbf{s}_h^*, \hat{\mathbf{s}}_h) \leq \varepsilon.$$

Thus, (G.3) implies that

$$\begin{aligned} \mathbb{P}_{\mu} \left[\left(\bigcap_{j=1}^{h-1} \mathcal{B}_j \cap \mathcal{C}_j \right) \cap (\mathcal{B}_h^c \cup \mathcal{C}_h^c) \right] &\leq \mathbb{E}_{\mu}[\gamma(\text{d}_{\text{TVC}}(\hat{\mathbf{s}}_h, \mathbf{s}_h^*)) \mathbf{I}\{\text{d}_{\text{TVC}}(\hat{\mathbf{s}}_h, \mathbf{s}_h^*) \leq \varepsilon\} + \text{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\mathbf{s}_h^*), \pi_h^*(\mathbf{s}_h^*)) \mid \mathcal{F}_{h-1}] \\ &\leq \gamma(\varepsilon) + \mathbb{E}_{\mu}[\mathbb{E}_{\mu}[\text{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\mathbf{s}_h^*), \pi_h^*(\mathbf{s}_h^*)) \mid \mathcal{F}_{h-1}]] \\ &= \gamma(\varepsilon) + \mathbb{E}_{\mu}[\text{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\mathbf{s}_h^*), \pi_h^*(\mathbf{s}_h^*))] \\ &= \gamma(\varepsilon) + \mathbb{E}_{\mathbf{s}_h^* \sim P_h^*} \mathbb{E}_{\mu}[\text{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\mathbf{s}_h^*), \pi_h^*(\mathbf{s}_h^*))], \end{aligned}$$

where the first equality follows from the tower rule for conditional expectations and the second follows because $s_h^* \sim P_h^*$ under μ . Summing and applying (G.4) implies that

$$\mathbb{P}_\mu \left[\left(\bigcap_{h=1}^H \mathcal{B}_h \cap \mathcal{C}_h \right)^c \right] \leq H\gamma(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} [\mathbf{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(s_h^*), \pi_h^*(s_h^*))].$$

Again, invoking input stability and the definitions $\mathcal{B}_h := \{\mathbf{d}_{\mathcal{A}}(\mathbf{a}_h, \hat{\mathbf{a}}_h^{\text{inter}}) \leq \varepsilon\}$ and $\mathcal{C}_h = \{\hat{\mathbf{a}}_h^{\text{inter}} = \hat{\mathbf{a}}_h\}$, $(\bigcap_{h=1}^H \mathcal{B}_h \cap \mathcal{C}_h)^c$ implies that

$$\max_{1 \leq h \leq H} \max \{\mathbf{d}_{\mathcal{S}}(s_{h+1}^*, \hat{s}_{h+1}), \mathbf{d}_{\mathcal{A}}(\mathbf{a}_h^*, \hat{\mathbf{a}}_h)\} \leq \varepsilon.$$

This concludes the proof. \square

G.1 Relaxing the TVC Condition

In this section, we show our results hold with the following generalization of TVC,

Definition G.1 (Relaxed TVC). We say $\hat{\pi}$ satisfies (γ, ε') -relaxed TVC if, for all h , and $s, s' \in \mathcal{S}$,

$$\inf_{\mu} \mathbb{P}_{(\mathbf{a}, \mathbf{a}')} [\mathbf{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') > \varepsilon'] \leq \gamma(\mathbf{d}_{\mathcal{S}}(s, s')),$$

where \inf_{μ} is the infimum over all couplings μ with $\mathbf{a} \sim \hat{\pi}_h(s)$ and $\mathbf{a}' \sim \hat{\pi}_h(s')$.

The main result of this section is as follows.

Proposition G.1 (Generalization of Proposition D.2). Let $\varepsilon > \varepsilon_1, \varepsilon_2 > 0$. Let π^* be input-stable w.r.t. $(\mathbf{d}_{\mathcal{S}}, \mathbf{d}_{\mathcal{A}})$ and let $\hat{\pi}$ be (γ, ε_1) -relaxed TVC. Further, suppose that $\mathbf{d}_{\mathcal{A}}$ (which need not satisfy the triangle inequality), satisfies

$$\{\mathbf{d}_{\mathcal{A}}(\mathbf{a}', \mathbf{a}) \leq \varepsilon_2\} \text{ and } \{\mathbf{d}_{\mathcal{A}}(\mathbf{a}'', \mathbf{a}') \leq \varepsilon_1\} \text{ implies } \{\mathbf{d}_{\mathcal{A}}(\mathbf{a}'', \mathbf{a}) \leq \varepsilon\}, \quad (\text{G.5})$$

for all $\mathbf{a}, \mathbf{a}', \mathbf{a}'' \in \mathcal{A}$ and all $\varepsilon_1, \varepsilon_2, \varepsilon$ given above. Then,

$$\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \parallel \pi^*) \leq H\gamma(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbf{d}_{\text{os}, \varepsilon_2}(\hat{\pi}_h(s_h^*) \parallel \pi^*(s_h^*)).$$

We remark that (G.5) holds the distance $\mathbf{d}_{\mathcal{A}}$ defined in Appendix D.2 whenever ε is sufficiently small, and $\varepsilon_1 + \varepsilon_2 = \varepsilon$.

Proof Sketch of Proposition G.1. The proof is nearly identical to the standard proof under (non-relaxed) TVC given above. The only difference is we replace the events $\{\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*) > \varepsilon\}$ and $\{\hat{\mathbf{a}}_h \neq \hat{\mathbf{a}}_h^{\text{inter}}\}$ with the events $\{\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*) > \varepsilon_1\}$ and $\{\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h, \hat{\mathbf{a}}_h^{\text{inter}}) > \varepsilon_2\}$. By (G.5), the intersection of the complement these events implies

$$\{\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \mathbf{a}_h^*) \leq \varepsilon_2\} \cap \{\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h, \hat{\mathbf{a}}_h^{\text{inter}}) \leq \varepsilon_1\} \text{ implies } \{\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h, \mathbf{a}_h^*) \leq \varepsilon\},$$

which allows the same argument to be followed. \square

Remark G.1 (Why Wasserstein continuity is not enough). Given that relaxed TVC is enough, we may be tempted to believe that $\hat{\pi}$ need only satisfy a Wasserstein continuity condition with linear $\gamma(\cdot)$. However, this is not quick strong enough to imply our argument.

For simplicity, assume $\mathbf{d}_{\mathcal{A}}$ is a metric and uppose that we have $\hat{\pi}$ is continuous with respect to the p -Wasserstein distance with metric $\mathbf{d}_{\mathcal{A}}$. This means that, when $\mathbf{d}_{\mathcal{S}}(s, s') \leq \varepsilon_0$, we can hope by Markov's inequality that

$$\inf_{\mu} \mathbb{E}_{\mu} [\mathbf{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}')^p]^{1/p} \leq L \mathbf{d}_{\mathcal{S}}(s, s'),$$

where μ are the couplings of $\mathbf{a} \sim \hat{\pi}_h(s)$ and $\mathbf{a}' \sim \hat{\pi}_h(s')$. Then, by Markov's inequality, the Most we have hope for is that

$$\inf_{\mu} \mathbb{P}_{\mu} [\mathbf{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') \geq tL \mathbf{d}_{\mathcal{S}}(s, s')] \leq \frac{1}{t^p}.$$

In particular, using the argument from the proof of our proposition,

$$\inf_{\mu} \mathbb{P}_{\mu}[\mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{inter}}, \hat{\mathbf{a}}_h) \geq tL\mathbf{d}_{\mathcal{S}}(\mathbf{s}_h^*, \hat{\mathbf{s}}_h)] \leq \frac{1}{t^p}.$$

On the other hand, our stability assumption can only bound $\mathbf{d}_{\mathcal{S}}(\mathbf{s}_{h+1}^*, \hat{\mathbf{s}}_{h+1}) \leq \max_{1 \leq i \leq h} \mathbf{d}_{\mathcal{A}}(\hat{\mathbf{a}}_i^{\text{inter}}, \hat{\mathbf{a}}_i)$. Thus, for $L > 1$, an any target ε , we can at most hope for bounds which scale as L^H in error. This is essentially the same issue tackled in Pfrommer et al. [50] by matching higher-order derivatives. What matching these derivatives in Wasserstein space would mean is still unclear.

H Imitation in the Composite MDP

In this section, we prove our imitation guarantees in the composite MDP under the full generality of data augmentation. The majority of this section is devoted to proving a more general version of [Theorem 4](#) that applies to vectorized notions of distance and helps tighten our bounds when instantiated in the control setting. In [Appendix H.1](#), we introduce some notation and state our most general result, [Theorem 8](#). We then proceed to show that [Theorem 4](#) follows from [Theorem 8](#) and in [Appendix H.2](#), we provide a detailed and rigorous proof of the main result. In [Appendix H.3](#), we show that the more general [Theorem 8](#) implies [Theorem 4](#) from the text.

Throughout, we also assume \mathcal{S} admits a direct decomposition. This is useful to capture the fact that we only apply smoothing on the \mathbf{o}_h coordinates (observation chunk), not the full trajectory chunk \mathbf{s}_h .

Definition H.1 (Direct Decomposition). Let $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ is a direct decomposition. We let $\phi_{\mathbf{o}}$ and $\phi_{/\mathbf{o}}$ denote projections onto the \mathcal{O} and $\mathcal{S}_{/\mathcal{O}}$ components, respectively. We say that the $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ is *compatible* with the dynamics if $F_h((\mathbf{o}, \mathbf{v}), \mathbf{a}) = F_h((\mathbf{o}, \mathbf{v}'), \mathbf{a})$ for all $\mathbf{v}, \mathbf{v}' \in \mathcal{S}_{/\mathcal{O}}$ and $\mathbf{o} \in \mathcal{O}$, and *compatible* with policy π if $\pi_h((\mathbf{o}, \mathbf{v}), \mathbf{a}) = \pi_h((\mathbf{o}, \mathbf{v}'), \mathbf{a})$; we define compatibility of a kernel W and of a pseudometric $\mathbf{d}(\cdot, \cdot) : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ with $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ similarly.

We emphasize that compatibility of dynamics with a direct decomposition does not make \mathbf{v} irrelevant because $\mathbf{d}_{\mathcal{S}}$ still depends on \mathbf{v} . For the purposes of the instantiation for control in the following appendix, we wish to control the imitation gaps on distances that do depend on \mathbf{v}_h , even though \mathbf{v}_h does not figure directly into the dynamics. Note that as defined, \mathbf{v}_h does depend on the dynamics up until time $h - 1$ and thus it is necessary to deal with this component in order to provide guarantees in $\mathbf{d}_{\mathcal{S}}$.

H.1 A generalization of Theorem 4

We now state a generalization of [Theorem 4](#), which replaces a single distance by a vector of distances of dimension K ; this will be useful for our instantiation of the composite MDP as a chunked control system in our final application (in particular, for deriving a bound on $\mathcal{L}_{\text{fin}, \varepsilon}$). It also showcases the most general structure accommodated by our proof technique.

We begin by defining some notation:

- Let $K \in \mathbb{N}$ denote a dimension
- Let $\vec{\varepsilon} \in \mathbb{R}_{\geq 0}^K$ denote a vector of tolerances
- Let $\vec{\mathbf{d}}_{\mathcal{S}}(\cdot, \cdot)$ denote a vector of pseudometrics $\mathbf{d}_{\mathcal{S}, i}$ on \mathcal{S}
- Let $\vec{\mathbf{d}}_{\mathcal{A}}$ denote a vector of non-negative functions $\mathbf{d}_{\mathcal{A}, i} : \mathcal{A}^2 \rightarrow \mathbb{R}_{\geq 0}$, not necessarily pseudometrics.
- Let \preceq denote vector wise inequality, and let the symbols \wedge and \vee be generalized to denote entrywise minima and maxima. Similarly, addition of vectors is coordinate wise with scalars assumed to be broadcast appropriately.
- We let $\mathbf{d}_{\mathcal{S}, 1} = \mathbf{d}_{\text{TV}}$ denote the metric we consider for evaluating total variation distance.

We generalize We assume the following measure-theoretic regularity conditions, generalizing [Assumption G.1](#) as follows.

Assumption H.1. We assume that \mathcal{S} and \mathcal{A} are Polish spaces. This means they are metrizable, but we do not annotate their metrics because, e.g. the metric on \mathcal{S} may be other than $d_{\mathcal{S}}$. We further assume that

- $d_{\mathcal{S},i}$ is a pseudometric and Borel measurable function from $\mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$.
- For any $\varepsilon \geq 0$, the set $\{(a, a') \in \mathcal{A} \times \mathcal{A} : d_{\mathcal{A},i}(a, a') > \varepsilon\}$ is an open subset of $\mathcal{A} \times \mathcal{A}$; i.e. $d_{\mathcal{A},i}(\cdot, \cdot)$ is lower semicontinuous. In particular, this means $d_{\mathcal{A},i}$ is a Borel measurable function. Note that this implies that the

$$\{(a, a') \in \mathcal{A} \times \mathcal{A} : \vec{d}_{\mathcal{A}}(a, a') \not\leq \vec{\varepsilon}\}.$$

is open and thus measurable.

Note that the above assumption is the natural vectorized generalization of [Assumption G.1](#). Next, we define vector versions of our imitation errors.

Definition H.2 (Imitation Errors, vector version). Given error parameter $\vec{\varepsilon} \in \mathbb{R}_{\geq 0}^K$, define

- The **vector joint-error**

$$\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi} \parallel \pi^*) := \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\exists h \in [H] : \vec{d}_{\mathcal{S}}(\hat{s}_{h+1}, s_{h+1}^*) \vee \vec{d}_{\mathcal{A}}(a_h^*, \hat{a}_h) \not\leq \vec{\varepsilon} \right],$$

where the infimum is over trajectory couplings $((\hat{s}_{1:H+1}, \hat{a}_{1:H}), (s_{1:H+1}^*, a_{1:H}^*)) \sim \mu_1 \in \mathcal{C}(\mathcal{D}_{\hat{\pi}}, \mathcal{D}_{\pi^*})$ satisfying $\mathbb{P}_{\mu_1}[\hat{s}_1 = s_1^*] = 1$.

- The **vector marginal error**

$$\vec{\Gamma}_{\text{marg}, \vec{\varepsilon}}(\hat{\pi} \parallel \pi^*) := \max_{h \in [H]} \max \left\{ \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\vec{d}_{\mathcal{S}}(\hat{s}_{h+1}, s_{h+1}^*) \not\leq \vec{\varepsilon} \right], \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\vec{d}_{\mathcal{A}}(a_h^*, \hat{a}_h) \not\leq \vec{\varepsilon} \right] \right\}$$

the same as the to joint-gap, with the “max” outside the probability and infimum over couplings.

- The **vector-wise one-step error**

$$\vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}_h(s) \parallel \pi_h^*(s)) := \inf_{\mu_2} \mathbb{P}_{\mu_2} \left[\vec{d}_{\mathcal{A}}(\hat{a}_h, a_h^*) \not\leq \vec{\varepsilon} \right],$$

where the infimum is over $(a_h^*, \hat{a}_h) \sim \mu_2 \in \mathcal{C}(\hat{\pi}_h(s), \pi_h^*(s))$.

We now describe input stability.

Definition H.3 (Input-Stability, vector version). A trajectory $(s_{1:H+1}, a_{1:H})$ is **input-stable** w.r.t. $(\vec{d}_{\mathcal{S}}, \vec{d}_{\mathcal{A}})$ if all sequences $s'_1 = s_1$ and $s'_{h+1} = F_h(s'_h, a'_h)$ satisfy

$$d_{\mathcal{S},i}(s'_{h+1}, s_{h+1}) \leq \max_{1 \leq j \leq h} d_{\mathcal{A},i}(a'_j, a_j), \quad \forall h \in [H], i \in [K]$$

Finally, define input process stability. A slight technicality is that, in our instantiation, π^* is taken to be a suitable regular condition probability of the joint distribution \mathcal{D}_{exp} of expert trajectories. This means that π^* can only really satisfy desired regularity conditions on states visited with positive probability by \mathcal{D}_{exp} . We address this subtlety by considering the following definition generalizing [Definition D.5](#) in the body. We also restrict the kernels under consideration to those which produce distributions *absolutely continuous* ([Definition F.4](#)) with respect to P_h^* , and denoted with the \ll comparator. More specifically, we only care about absolute continuity under the projections onto the \mathcal{O} component of \mathcal{S} .

Definition H.4 (Input & Process Stability, vector version). Let $p_{\text{IPS}} \in (0, 1)$, $\vec{\gamma}_{\text{IPS}} = (\gamma_{\text{IPS},i})_{1 \leq i \leq K}$ be a collection non-decreasing maps $\gamma_{\text{IPS},i} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, let $d_{\text{IPS}} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ be a pseudometric (possibly other than any of the $d_{\mathcal{S},i}$), and $r_{\text{IPS}} > 0$. We say a policy π^* is $(\vec{\gamma}_{\text{IPS}}, d_{\text{IPS}}, r_{\text{IPS}}, p_{\text{IPS}})$ -**(vectorwise-input- & process stable (vIPS))** if the following holds for any $r \in [0, r_{\text{IPS}}]$:

Consider any sequence of kernels $W_h : \mathcal{S} \rightarrow \Delta(\mathcal{S})$, $1 \leq h \leq H$, satisfying

$$\forall h, s \in \mathcal{S} : \mathbb{P}_{\tilde{s} \sim W_h(s)}[d_{\text{IPS}}(\tilde{s}, s) \leq r] = 1, \quad \phi_{\mathcal{O}} \circ W_h(s) \ll \phi_{\mathcal{O}} \circ P_h^*. \quad (\text{H.1})$$

Define a process $s_1 \sim P_{\text{init}}$, $\tilde{s}_h \sim W_h(s_h)$, $a_h \sim \pi_h(\tilde{s}_h)$, and $s_{h+1} := F_h(s_h, a_h)$. Then, with probability at least $1 - p_{\text{IPS}}$,

- (a) the sequence $(s_{1:H+1}, a_{1:H})$ is input-stable w.r.t (\vec{d}_S, \vec{d}_A) (as defined by [Definition H.3](#)).
- (b) $\max_{h \in [H]} d_{S,i}(F_h(\tilde{s}_h, a_h), s_{h+1}) \leq \gamma_{IPS,i}(r)$.

We can now state our desired generalization.

Theorem 8. Suppose that there

- (a) π^* is $(\vec{\gamma}_{IPS}, d_{IPS}, r_{IPS}, p_{IPS})$ -vector IPS in the sense of [Definition H.4](#).
- (b) There is a direct decomposition of $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$, which associated projection maps ϕ_o and $\phi_{/o}$, and which is compatible with the dynamics, and policies π^* , $\hat{\pi}$, and smoothing kernel W_σ , and d_{IPS} .
- (c) $\phi_o \circ W_\sigma$ is γ_σ -TVC with respect to the pseudometric $d_{TVC} = d_{S,1}$.

Let $\hat{\pi}_\sigma$ be any policy which is $\hat{\gamma}$ -TVC, also w.r.t. $d_{TVC} = d_{S,1}$. Finally, let $\vec{\varepsilon} \in \mathbb{R}_{\geq 0}^K$, $r \in (0, \frac{1}{2}r_{IPS}]$, and define

$$p_r := \sup_s \mathbb{P}_{s' \sim W_\sigma(s)}[d_{IPS}(s', s) > r], \quad \vec{\varepsilon}_{\text{marg}} := \vec{\varepsilon} + \vec{\gamma}_{IPS}(2r).$$

Then,

- For any policy $\hat{\pi}$, both $\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*)$ and $\vec{\Gamma}_{\text{marg}, \vec{\varepsilon}_{\text{marg}}}(\hat{\pi}_\sigma \parallel \pi^*)$ are upper bounded by

$$p_{IPS} + H(2p_r + \hat{\gamma}(\vec{\varepsilon}_1) + (\hat{\gamma} + \gamma_\sigma) \circ \gamma_{IPS, \text{TVC}}(2r)) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}_{\sigma, h}(s_h^{\text{tel}}) \parallel \pi_{\text{rep}, h}^*(s_h^{\text{tel}})) \quad (\text{H.2})$$

- In the special case where $\hat{\pi}_\sigma = \hat{\pi} \circ W_\sigma$, we can take $\hat{\gamma} = \gamma_\sigma$, and obtain that $\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*)$ and $\vec{\Gamma}_{\text{marg}, \vec{\varepsilon}_{\text{marg}}}(\hat{\pi}_\sigma \parallel \pi^*)$ are upper bounded by

$$p_{IPS} + H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{IPS, \text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} \vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}, h}^*(\tilde{s}_h^*)). \quad (\text{H.3})$$

We note that [Theorem 4](#) is a special case of [Theorem 8](#) and prove the former assuming the latter here at the end of the section.

H.2 Proof of Theorem 8

H.2.1 Proof Overview and Coupling Construction

We begin with an intuitive overview of the proof and partially construct the relevant intermediate trajectories used to define our coupling, after which we sketch the organization of the rest of [Appendix H.2](#).

The proof proceeds by constructing a sophisticated coupling between the law of a trajectory evolving according to $\hat{\pi}$ and a trajectory evolving according to π_{rep}^* by introducing several intermediate sequences of composite states and composite actions.

We partially specify this coupling below and formally construct it in [Appendix H.2.4](#). Our construction is recursive and relies on the input and process stability as well as total variation continuity to show that if the trajectories generated by π_{rep}^* and $\hat{\pi}$ are close in $\vec{d}_{\text{os}, \vec{\varepsilon}}$ evaluated on states at step h , then they will remain close at step $h+1$. There are a number of technical subtleties involved, especially those of a measure-theoretic nature, but much of the intuition can be gleaned from the following partial specification of the coupling μ over composite-state $(\hat{s}_{1:H}, s_{1:H}^{\text{rep}}, s_{1:H}^{\text{tel}}, \tilde{s}_{1:H}^{\text{tel}}) \subset \mathcal{S}$, composite-actions $(\hat{a}_{1:H}^{\text{rep}}, \hat{a}_{1:h}, a_{1:H}^{\text{tel}}) \subset \mathcal{K}$ and interpolating composite-actions, $(\hat{a}_{1:H}^{\text{rep}, \text{inter}}, \hat{a}_{1:H}^{\text{tel}, \text{inter}}) \subset \mathcal{A}$.

To define the construction, we define the probability kernels corresponding to the replica and deconvolution policies. Note that these are slightly different from the definitions in the body due to the use of the direct decomposition; the intuition is the same, however.

Definition H.5 (Replica and Deconvolution Kernels). Let $P_{\text{aug},h}^{\text{proj}}$ denote the joint distribution over $(o_h^*, s_h^*, \tilde{o}_h^*, a_h^*)$ under the generative process

$$s_h^* \sim P_h^*, \quad a_h^* \sim \pi_h^*(s_h^*), \quad o_h^* = \phi_o(s_h^*), \quad \tilde{o}_h^* \sim \phi_o \circ W_\sigma(s_h^*)$$

For $o \in \mathcal{O}$, let $W_{\text{dec},\mathcal{O},h}^*(o)$ denote the distribution of o_h^* conditioned on $\tilde{o}_h^* = o$, under $P_{\text{aug},h}^{\text{proj}}$. Given $s = (o, v)$, define

$$\begin{aligned} W_{\text{dec},h}^*(s) &= W_{\text{dec},\mathcal{O},h}^*(\phi_o(s)) \otimes \delta_{\phi_{/o}(s)}, \\ W_{\text{rep},h}^*(s) &= W_{\text{dec},h}^* \circ (W_\sigma(\phi_o(s)) \otimes \delta_{\phi_{/o}(s)}) = (W_{\text{dec},\mathcal{O},h}^* \circ W_\sigma(\phi_o(s))) \otimes \delta_{\phi_{/o}(s)}. \end{aligned}$$

where we recall the dirac-delta δ . Equivalently, $W_{\text{dec},h}^*(s)$ denotes the conditional sequence of (\tilde{o}, v) , where $v = \phi_{/o}(s)$, and $\tilde{o} \sim W_{\text{dec},\mathcal{O},h}^*(s)$; $W_{\text{rep},h}^*$ can be expressed similarly.

We remark that $W_{\text{dec},h}^*$ and $W_{\text{rep},h}^*$ are both kernels and by [Theorem 7](#), we may assume that the joint distribution over $(s_h^*, \tilde{s}_h^{\text{tel}})$ admits a regular conditional probability and thus these constructions are well-defined.

Remark H.1. Note that the kernels $W_{\text{dec},h}^*$ and $W_{\text{rep},h}^*$ are compatible with the decomposition $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ by construction. Moreover, note that if $s = (o, v)$, $\phi_{/o} \circ W_{\text{dec},h}^*(s) = \phi_{/o} \circ W_{\text{rep},h}^*(s)$ is the dirac-delta distribution supported on v .

Lemma H.1. Under our the assumption that π^* and W_σ are compatible with the direct decomposition,

$$\pi_{\text{dec},h}^*(s) = \pi^* \circ W_{\text{dec},h}^*, \quad \pi_{\text{rep},h}^*(s) = \pi^* \circ W_{\text{rep},h}^*$$

Proof. This follows immediately because π^* and W_σ are compatible with the direct decomposition, and by the definition of [Definition D.4](#). \square

A template for the coupling. Our couplings are partially specified by the following generative process, and what remains unspecified are couplings between random variables at each step h . In what follows, let \mathcal{F}_0 denote the σ -algebra generatively by $\hat{s}_1 = s_1^{\text{rep}} = s_1^{\text{tel}}$. Let \mathcal{F}_h denote the sigma-algebra generated by $(\hat{s}_{1:h}, s_{1:h}^{\text{rep}}, s_{1:h}^{\text{tel}})$, $(a_{1:h}^{\text{rep}}, \tilde{s}_{1:h}^{\text{rep}}, \tilde{s}_{1:h}^{\text{tel}}, a_{1:h}^{\text{tel}}, \hat{a}_{1:h})$, and $(\hat{a}_{1:h}^{\text{rep,inter}}, \hat{a}_{1:h}^{\text{tel,inter}})$.

- The initial states are drawn as

$$\hat{s}_1 = s_1^{\text{rep}} = s_1^{\text{tel}} \sim P_{\text{init}}.$$

- The dynamics satisfy

$$\hat{s}_{h+1} = F_h(\hat{s}_h, \hat{a}_h), \quad s_{h+1}^{\text{rep}} = F_h(s_h^{\text{rep}}, a_h^{\text{rep}}), \quad s_{h+1}^{\text{tel}} = F_h(\tilde{s}_h^{\text{tel}}, a_h^{\text{tel}})$$

Note that determinism of the dynamics implies that s_{h+1}^{tel} , s_{h+1}^{rep} and \hat{s}_{h+1} are \mathcal{F}_h -measurable.

- We generate

$$\begin{aligned} \tilde{s}_h^{\text{rep}} \mid \mathcal{F}_{h-1} &\sim W_{\text{rep},h}^*(s_h^{\text{rep}}), \quad a_h^{\text{rep}} \mid \mathcal{F}_{h-1}, \tilde{s}_h^{\text{rep}} \sim \pi_h^*(\tilde{s}_h^{\text{rep}}), \\ \tilde{s}_h^{\text{tel}} \mid \mathcal{F}_{h-1} &\sim W_{\text{rep},h}^*(s_h^{\text{tel}}), \quad a_h^{\text{tel}} \mid \mathcal{F}_{h-1}, \tilde{s}_h^{\text{tel}} \sim \pi_h^*(\tilde{s}_h^{\text{tel}}). \\ \hat{a}_h \mid \mathcal{F}_{h-1} &\sim \hat{\pi}_{\sigma,h}(\hat{s}_h) \end{aligned}$$

Importantly, we note that, marginalizing over \tilde{s}_h^{tel} and \tilde{s}_h^{rep} , respectively, $a_h^{\text{tel}} \mid \mathcal{F}_{h-1} \sim \pi_{\text{rep},h}^*(s_h^{\text{tel}})$ and $a_h^{\text{rep}} \mid \mathcal{F}_{h-1} \sim \pi_{\text{rep},h}^*(s_h^{\text{rep}})$.

- Lastly, we select interpolating actions via

$$\hat{a}_h^{\text{rep,inter}} \mid \mathcal{F}_{h-1} \sim \hat{\pi}_{\sigma,h}(s_h^{\text{rep}}), \quad \hat{a}_h^{\text{tel,inter}} \mid \mathcal{F}_{h-1} \sim \hat{\pi}_{\sigma,h}(s_h^{\text{tel}})$$

We will say μ is “respects the construction” as shorthand to mean that μ obeys the above equations. The coupling is illustrated graphically in [Figure 10](#). We now establish several key properties of the above constructions, separated into a subsection for the sake of clarity.

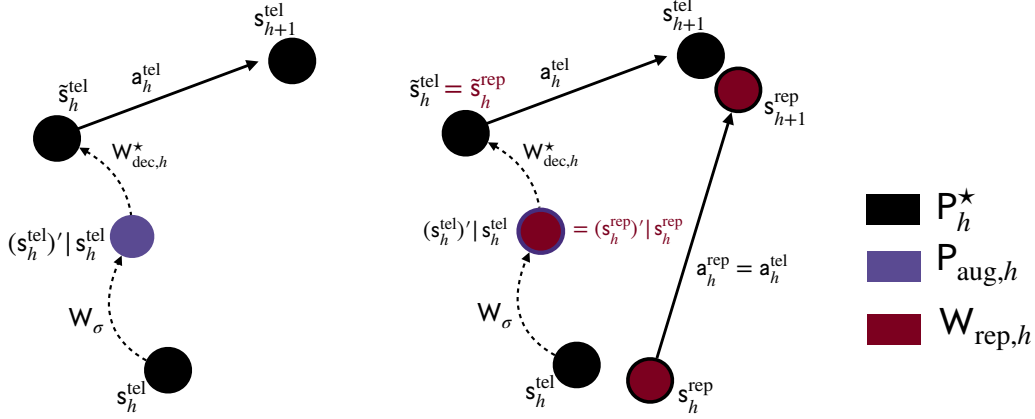


Figure 10: Graphical illustration of the coupling, in the special case where $\mathcal{O} = \mathcal{S}$ for simplicity. **On the left** is the teleporting sequence, with $\tilde{s}_h^{\text{tel}} \sim W_{\text{rep},h}^*(s_h^{\text{tel}}) = W_{\text{dec},h}^* \circ W_\sigma(s_h^{\text{tel}})$. We represent the teleporting explicitly by noising s_h^{tel} to become $(s_h^{\text{tel}})'$ by applying W_σ and then applying $W_{\text{dec},h}^*$ to complete the “teleporting” to \tilde{s}_h^{tel} . We then apply $a_h^{\text{tel}} \sim \pi_h^*(\tilde{s}_h^{\text{tel}})$, and continue onto s_{h+1}^{tel} from the teleported state $\tilde{s}_{h+1}^{\text{tel}}$. **On the right**, we illustrate the replica sequence next to the teleporting sequence. We start with s_h^{rep} , which is close to s_h^{tel} (a consequence of our proof). We then apply the replica kernel to achieve \tilde{s}_h^{rep} . Our argument uses that $W_{\text{rep},h}^* = W_{\text{dec},h}^* \circ W_\sigma$ is TVC (a consequence of TVC of W_σ as shown in Lemma H.2). We depict this property pictorially: since W_σ is TVC and s_h^{tel} and s_h^{rep} are close, we can couple things in such a way that, with good probability, $(s_h^{\text{tel}})' \sim W_\sigma(s_h^{\text{tel}})$ and $(s_h^{\text{rep}})' \sim W_\sigma(s_h^{\text{rep}})$ are equal. We then extend the coupling to that $\tilde{s}_h^{\text{rep}} = \tilde{s}_h^{\text{tel}}$ on the event $\{(s_h^{\text{tel}})' = (s_h^{\text{rep}})'\}$, both being drawn by applying $W_{\text{dec},h}^*$ to both of $(s_h^{\text{tel}})' = (s_h^{\text{rep}})'$. We extend the coupling once more so that $a_h^{\text{tel}} \sim \pi_h^*(\tilde{s}_h^{\text{tel}})$ and $a_h^{\text{rep}} \sim \pi_h^*(\tilde{s}_h^{\text{rep}})$ are equal on this good probability event. Using our notion of stability, IPS, and the fact that s_h^{rep} and s_h^{tel} are close, the good probability event on which a_h^{tel} and a_h^{rep} are equal implies that s_{h+1}^{rep} remains close to s_{h+1}^{tel} . We remark that our actual analysis never explicitly computes the $(\cdot)'$ -terms drawn from W_σ ; rather, these terms appear implicitly in our definitions of $W_{\text{rep},h}^*$ and the verification of its TVC property.

Organization of the remaining parts of Appendix H.2. In Appendix H.2.2, we prove several prerequisite properties of the construction given above, including concentration of the smoothing kernel, and key properties of the replica distribution. Next, Appendix H.2.3 shows that, due to these properties of the replica distribution, we can bound the marginal imitation gap by controlling the tracking of the teleporting sequence constructed above. Finally, in Appendix H.2.4 we formally construct the coupling and rigorously prove Theorem 8.

H.2.2 Properties of smoothing, deconvolution, and replicas.

In this section, we establish several useful properties of smoothed and replica policies. We begin by showing that smoothed policies are TVC.

Lemma H.2. The following hold

- For any h , $\phi_o \circ W_{\text{rep},h}^*$ and $\pi_{\text{rep},h}^*$ are γ_σ TVC.
- If π is any policy compatible with the direct decomposition $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ (in the sense of Definition H.1), then $\pi \circ W_\sigma$ is γ_σ -TVC.

Proof. We observe that $\phi_o \circ W_{\text{rep},h}^* = \phi_o \circ W_{\text{dec},h}^* \circ W_\sigma(s)$. Moreover, we observe $W_{\text{dec},h}^*$ satisfies $\phi_o \circ W_{\text{dec},h}^*(s) = W_{\text{dec},\mathcal{O},h}^* \circ \phi_o$, so that $\phi_o \circ W_{\text{rep},h}^* = W_{\text{dec},\mathcal{O},h}^* \circ \phi_o \circ W_\sigma(s)$. As $\phi_o \circ W_\sigma$ is TVC, the first claim is a consequence of the data-processing inequality Lemma F.4. The second uses the fact that all listed objects involve composition of kernels with W_σ . \square

Next, we show that the replica construction preserves marginals.

Lemma H.3 (Marginal-Preservation). There exists a coupling \mathbb{P} of $\mathbf{o}_h \sim \phi_o \circ P_h^*$, $\mathbf{o}'_h \sim \phi_o \circ W_\sigma(\mathbf{o}_h, \cdot)$ (where (\cdot) denotes an irrelevant argument due to compatibility of W_σ with the direct decomposition), and $\tilde{\mathbf{o}}_h \sim \phi_o \circ W_{\text{rep},h}^*(\mathbf{o}_h, \cdot)$ (again, (\cdot) denotes an irrelevant argument) such that

$$(\mathbf{o}_h, \mathbf{o}'_h) \stackrel{d}{=} (\tilde{\mathbf{o}}_h, \mathbf{o}'_h).$$

In particular, for $\mathbf{s}_h^{\text{tel}}$ and $\tilde{\mathbf{s}}_h^{\text{tel}}$ as in our construction, the marginal distributions of $\phi_o(\mathbf{s}_h^{\text{tel}})$ and $\phi_o(\tilde{\mathbf{s}}_h^{\text{tel}})$ are the same, where $\mathbf{s}_h^{\text{tel}} \sim P_h^*$ and $\tilde{\mathbf{s}}_h^{\text{tel}} \mid \mathbf{s}_h^{\text{tel}} \sim W_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}})$.

Proof. By [Assumption G.1](#) and [Theorem 7](#), we may assume that all joint distributions' conditional probabilities are regular conditional probabilities and thus almost surely equal to a kernel. Moreover, since all kernels are compatible with the direct decomposition, it suffices to prove the special case of the trivial direct-decomposition where $\mathcal{O} = \mathcal{S}$. Fix a common measure \mathbb{P} over which $\mathbf{s}_h^{\text{tel}}$, $\tilde{\mathbf{s}}_h^{\text{tel}}$, and \mathbf{s}'_h are defined such that $\mathbf{s}_h^{\text{tel}} \sim P_h^*$, $\mathbf{s}'_h \sim W_\sigma(\mathbf{s}_h^{\text{tel}})$, and $\tilde{\mathbf{s}}_h^{\text{tel}} \sim W_{\text{dec},h}(\mathbf{s}'_h)$. Then for any measurable sets A, B , we have

$$\begin{aligned} \mathbb{P}(\mathbf{s}_h^{\text{tel}} \in A, \mathbf{s}'_h \in B) &= \mathbb{P}(\mathbf{s}'_h \in B) \cdot \mathbb{E}_{\mathbf{s}'_h} [\mathbf{I}[\mathbf{s}'_h \in B] \cdot \mathbb{P}(\mathbf{s}_h^{\text{tel}} \in A \mid \mathbf{s}'_h)] \\ &= \mathbb{P}(\mathbf{s}'_h \in B) \cdot \mathbb{E}_{\mathbf{s}'_h} [\mathbf{I}[\mathbf{s}'_h \in B] \cdot \mathbb{P}(\tilde{\mathbf{s}}_h^{\text{tel}} \in A \mid \mathbf{s}'_h)] \\ &= \mathbb{P}(\tilde{\mathbf{s}}_h^{\text{tel}} \in A, \mathbf{s}'_h \in B), \end{aligned}$$

where the first equality holds by the fact that we are working with regular conditional probabilities and Bayes' rule, the second equality holds by the definition of the deconvolution kernel above, and the last equality holds again by Bayes' rule and the tower rule for conditional expectations.

To prove the second statement, we apply induction, again assuming that $\mathcal{O} = \mathcal{S}$ as in the proof of the first statement. Note that $\mathbf{s}_1^{\text{tel}} \sim P_1^* = P_{\text{init}}$, and $\tilde{\mathbf{s}}_1^{\text{tel}} \sim W_{\text{rep},1}^* \circ P_1^*$. Thus, from the first part of the lemma, $\phi_o(\mathbf{s}_1^{\text{tel}}) \sim \phi_o \circ P_1^*$. Now, suppose the induction holds up to step h . Then, $\tilde{\mathbf{s}}_h^{\text{tel}} \sim P_h^*$, as $\mathbf{a}_h^{\text{tel}} \sim \pi_h^*(\mathbf{a}_h^{\text{tel}})$, then $\mathbf{s}_{h+1}^{\text{tel}} = F_h(\tilde{\mathbf{s}}_h^{\text{tel}}, \mathbf{a}_h^{\text{tel}}) \sim P_{h+1}^*$. Again $\tilde{\mathbf{s}}_{h+1}^{\text{tel}} \sim W_{\text{rep},h+1}^*(\mathbf{s}_{h+1}^{\text{tel}})$, so that $\tilde{\mathbf{s}}_{h+1}^{\text{tel}}$ has marginal $W_{\text{rep},h+1}^* \circ P_{h+1}^* = P_{h+1}^*$, as needed. \square

We further show that $W_{\text{rep},h}$ can be defined to be absolutely continuous with respect to P_h^* .

Lemma H.4. The kernel $W_{\text{rep},h}$ satisfies that $\phi_o \circ W_{\text{rep},h} \ll \phi_o \circ P_h^*$ as laws, validating the second condition in [\(H.1\)](#). It further holds that $\phi_o \circ W_{\text{dec},h} \ll \phi_o \circ P_h^*$.

Proof. The first statement follows immediately from [Lemma H.3](#) because these distributions are the same. The second statement follows immediately from the tower law of conditional expectation and the definition of $W_{\text{dec},h}$. \square

Lastly, we establish that the replica kernel inherits all concentration properties from the smoothing kernel.

Lemma H.5 (Replica Concentration). Recall that

$$p_r := \sup_{\mathbf{s}} \mathbb{P}_{\mathbf{s}' \sim W_\sigma(\mathbf{s})} [\mathbf{d}_{\text{IPS}}(\mathbf{s}', \mathbf{s}) > r].$$

We then have

$$\mathbb{P}_{\mathbf{s}_h \sim P_h^*, \tilde{\mathbf{s}}_h \sim W_{\text{rep},h}^*(\mathbf{s}_h)} [\mathbf{d}_{\text{IPS}}(\tilde{\mathbf{s}}_h, \mathbf{s}_h) > 2r] \leq 2p_r.$$

Proof. Again, all terms – W_σ , $W_{\text{rep},h}^*$, $W_{\text{dec},h}^*$ and \mathbf{d}_{IPS} – are compatible with the direct decomposition, it suffices to consider the case of the trivial direct decomposition under which $\mathcal{O} = \mathcal{S}$.

Let \mathbb{P} denote a distribution over $\mathbf{s}_h \sim P_h^*$, $\mathbf{s}'_h \sim W_\sigma(\mathbf{s}_h)$, and $\tilde{\mathbf{s}}_h \sim W_{\text{dec},h}^*(\mathbf{s}'_h)$. In this special case, we see that $\tilde{\mathbf{s}}_h \mid \mathbf{s}_h \sim W_{\text{rep},h}^*(\mathbf{s}_h)$ ⁹. By a union bound,

$$\begin{aligned} \mathbb{P}_{\mathbf{s}_h \sim P_h^*, \tilde{\mathbf{s}}_h \sim W_{\text{rep},h}^*(\mathbf{s}_h)} [\mathbf{d}_{\text{IPS}}(\mathbf{s}_h, \tilde{\mathbf{s}}_h) > 2r] &\leq \mathbb{P}[\mathbf{d}_{\text{IPS}}(\tilde{\mathbf{s}}_h, \mathbf{s}'_h) > r] + \mathbb{P}[\mathbf{d}_{\text{IPS}}(\mathbf{s}_h, \mathbf{s}'_h) > r] \\ &= 2\mathbb{P}[\mathbf{d}_{\text{IPS}}(\mathbf{s}_h, \mathbf{s}'_h) > r] \leq 2p_r, \end{aligned}$$

where the equality follows from the first statment of [Lemma H.3](#). \square

⁹Notice that, for general $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$, this condition would become $\phi_o(\tilde{\mathbf{s}}_h) \mid \phi_o(\mathbf{s}_h) \sim \phi_o \circ W_{\text{rep},h}^*(\phi_o(\mathbf{s}_h), \cdot)$, where the \cdot argument is irrelevant.

Remark H.2. Note that, in the previous lemma, it suffices that the following weaker condition holds: $\mathbb{P}_{s \sim P_h^*, s' \sim W_\sigma(s)}[\mathbf{d}_{\text{IPS}}(s', s) > r] \leq p_r$, i.e. for concentration to hold only in distribution over $s \sim P_h^*$, instead of *uniformly* over states.

H.2.3 Bounding the marginal imitation gaps in terms of the teleporting sequence error

Before turning to the proof of [Theorem 8](#), we verify that closeness to the *teleporting sequences* suffices to control error in marginal gap to π^* . The key property here is that the teleporting sequence, as shown in [Lemma H.3](#), has the same marginal distribution over states as does π^* .

Lemma H.6. Let μ be any coupling obeying the construction of the couplings above. Then,

$$\vec{\Gamma}_{\text{marg}, \vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi^*) \leq \mathbb{P}_\mu \left[\exists h \in [H] : \left\{ \vec{\mathbf{d}}_S(s_{h+1}^{\text{tel}}, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right\} \cup \left\{ \vec{\mathbf{d}}_{\mathcal{A}}(a_h^{\text{tel}}, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right\} \right]$$

Proof. We begin with a (reverse) union bound.

$$\begin{aligned} & \mathbb{P}_\mu \left[\exists h \in [H] : \left\{ \vec{\mathbf{d}}_S(s_{h+1}^{\text{tel}}, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right\} \cup \left\{ \vec{\mathbf{d}}_{\mathcal{A}}(a_h^{\text{tel}}, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right\} \right] \\ & \geq \max_h \max \left\{ \mathbb{P}_\mu \left[\vec{\mathbf{d}}_S(s_{h+1}^{\text{tel}}, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right], \mathbb{P}_\mu \left[\vec{\mathbf{d}}_{\mathcal{A}}(a_h^{\text{tel}}, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right] \right\}. \end{aligned}$$

By [Lemma H.3](#) implies that s_h^{tel} has the marginal distribution of $s_h^* \sim P_h^*$. Moreover, by construction, for each h , $a_h^{\text{tel}} \mid \mathcal{F}_h \sim \pi_{\text{rep}, h}^*(s_h^{\text{tel}})$. Thus, for each h , s_{h+1}^{tel} and a_h^{tel} have the same *marginals* as the marginals as s_{h+1}^* and a_h^* under the distribution D_{π^*} induced by π^* . Hence,

$$\begin{aligned} \mathbb{P}_\mu \left[\vec{\mathbf{d}}_S(s_{h+1}^{\text{tel}}, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right] & \geq \inf_{\mu_1} \mathbb{P} \left[\vec{\mathbf{d}}_S(s_{h+1}^*, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right] \\ \mathbb{P}_\mu \left[\vec{\mathbf{d}}_{\mathcal{A}}(a_h^{\text{tel}}, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right] & \geq \inf_{\mu_1} \mathbb{P} \left[\vec{\mathbf{d}}_{\mathcal{A}}(a_h^*, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right], \end{aligned}$$

where the \inf_{μ_1} is, as in [Definitions D.1](#) and [H.2](#), the infimum over couplings between D_{π^*} and $D_{\hat{\pi}}$. Thus,

$$\begin{aligned} & \mathbb{P}_\mu \left[\exists h \in [H] : \left\{ \vec{\mathbf{d}}_S(s_{h+1}^{\text{tel}}, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right\} \cup \left\{ \vec{\mathbf{d}}_{\mathcal{A}}(a_h^{\text{tel}}, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right\} \right] \\ & \geq \max_h \max \left\{ \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\vec{\mathbf{d}}_S(s_{h+1}^*, \hat{s}_{h+1}) \not\leq \vec{\varepsilon}_{\text{marg}} \right], \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\vec{\mathbf{d}}_{\mathcal{A}}(a_h^*, \hat{a}_h) \not\leq \vec{\varepsilon}_{\text{marg}} \right] \right\} \\ & := \vec{\Gamma}_{\text{marg}, \vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi^*). \end{aligned}$$

□

H.2.4 Formal proof of Theorem 8

We now proceed to formally prove [Theorem 8](#)

Key Events. For the random variables defined above, we define three groups of events.

- The *coupling events*, denoted by \mathcal{B} , which are controlled by carefully selecting a coupling.
- The *inductive events*, denoted by \mathcal{C} , which we condition on when bounding the probability of the coupling events.
- The *stability events*, denoted by \mathcal{Q} , which take advantage of the stability properties of the imitation policy.

Definition H.6 (Coupling Events). Define the events

$$\begin{aligned}
\mathcal{B}_{\text{tel},h} &= \{ \mathbf{a}_h^{\text{rep}} = \mathbf{a}_h^{\text{tel}}, \phi_o(\tilde{\mathbf{s}}_h^{\text{rep}}) = \phi_o(\tilde{\mathbf{s}}_h^{\text{tel}}) \} \\
\mathcal{B}_{\text{est},h} &= \{ \vec{d}_{\mathcal{A}}(\hat{\mathbf{a}}_h^{\text{tel,inter}}, \mathbf{a}_h^{\text{tel}}) \not\leq \vec{\varepsilon} \} \\
\mathcal{B}_{\text{inter},h} &= \{ \hat{\mathbf{a}}_h^{\text{tel,inter}} = \hat{\mathbf{a}}_h^{\text{rep,inter}} \} \\
\mathcal{B}_{\hat{\mathbf{a}},h} &= \{ \hat{\mathbf{a}}_h^{\text{rep,inter}} = \hat{\mathbf{a}}_h \} \\
\mathcal{B}_{\text{all},h} &= \mathcal{B}_{\text{inter},h} \cap \mathcal{B}_{\text{tel},h} \cap \mathcal{B}_{\text{est},h} \cap \mathcal{B}_{\hat{\mathbf{a}},h} \\
\bar{\mathcal{B}}_{\text{all},h} &= \bigcap_{j=1}^h \mathcal{B}_{\text{all},h}
\end{aligned}$$

Notice that each of the events above are \mathcal{F}_h -measurable. Moreover, note that on $\bar{\mathcal{B}}_{\text{all},h}$, $\max_{1 \leq j \leq h} \phi_{\text{IS}}(\hat{\mathbf{a}}_j, \mathbf{a}_j^{\text{rep}}) \leq \varepsilon$.

Definition H.7 (Inductive Event). Define the events

$$\begin{aligned}
\mathcal{C}_{\hat{\mathbf{s}},h} &= \{ \vec{d}_{\mathcal{S}}(\mathbf{s}_h^{\text{rep}}, \hat{\mathbf{s}}_h) \leq \vec{\varepsilon} \}, \\
\mathcal{C}_{\text{tel},h} &= \{ \vec{d}_{\mathcal{S}}(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) \leq \vec{\gamma}_{\text{IPS}}(2r) \} \\
\mathcal{C}_{\text{all},h} &:= \mathcal{C}_{\hat{\mathbf{s}},h} \cap \mathcal{C}_{\text{tel},h} \\
\bar{\mathcal{C}}_{\text{all},h} &= \bigcap_{j=1}^h \mathcal{C}_{\text{all},j}
\end{aligned}$$

Notice that all the above events are \mathcal{F}_{h-1} -measurable, due to determinism of the dynamics. Note that also $\mathbb{P}_{\mu}[\bar{\mathcal{C}}_{\text{all},1}] = 1$ for any μ that respects the construction (as $\mathbf{s}_1^{\text{rep}} = \mathbf{s}_1^{\text{tel}} = \hat{\mathbf{s}}_1$).

Definition H.8 (Stability Events). Define the events

$$\begin{aligned}
\mathcal{Q}_{\text{close}} &:= \{ \forall h \in [H] : d_{\text{IPS}}(\mathbf{s}_h^{\text{rep}}, \tilde{\mathbf{s}}_h^{\text{rep}}) \leq 2r \} \\
\mathcal{Q}_{\text{IS}} &:= \{ (\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}}) \text{ is input-stable w.r.t. } (\vec{d}_{\mathcal{S}}, \vec{d}_{\mathcal{A}}) \} \\
\mathcal{Q}_{\text{IPS}} &:= \{ \vec{d}_{\mathcal{S}}(F_h(\tilde{\mathbf{s}}_h^{\text{rep}}, \mathbf{a}_h^{\text{rep}}), \mathbf{s}_{h+1}^{\text{rep}}) \leq \vec{\gamma}_{\text{IPS}} \circ d_{\text{IPS}}(\tilde{\mathbf{s}}_h^{\text{rep}}, \mathbf{s}_h^{\text{rep}}), \quad 1 \leq j \leq H \} \\
\mathcal{Q}_{\text{all}} &:= \mathcal{Q}_{\text{IPS}} \cap \mathcal{Q}_{\text{close}}.
\end{aligned}$$

In words, $\mathcal{Q}_{\text{close}}$ the event on which $\mathbf{s}_h^{\text{rep}}$ and $\tilde{\mathbf{s}}_h^{\text{rep}} \sim W_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}})$ are close, and \mathcal{Q}_{IS} and \mathcal{Q}_{IPS} ensure consequences of (vector) input-stability and (vector) input process stability holds.

Steps of the proof. First, we use stability to reduce the event $\bar{\mathcal{C}}_{\text{all},h+1}$ to $\bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h}$:

Lemma H.7 (Stability Claim). By construction,

$$\bar{\mathcal{C}}_{\text{all},h+1} \subset \mathcal{Q}_{\text{all}} \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h}.$$

Proof. It suffices to show that on $\mathcal{Q}_{\text{all}} \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h}$, $\vec{d}_{\mathcal{S}}(\mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}) \leq \vec{\varepsilon}$ and $\vec{d}_{\mathcal{S}}(\mathbf{s}_{h+1}^{\text{rep}}, \mathbf{s}_{h+1}^{\text{tel}}) \leq \vec{\gamma}_{\text{IPS}}(2r)$. By applying the event \mathcal{Q}_{IS} to the sequence $\mathbf{a}'_h = \hat{\mathbf{a}}_h$ and $\mathbf{s}'_h = \hat{\mathbf{s}}_h$, we have that on $\mathcal{Q}_{\text{all}} \subset \mathcal{Q}_{\text{IS}}$ that

$$\forall h \in [H], i \in [K], \quad d_{\mathcal{S},i}(\mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}) \leq \max_{1 \leq j \leq h} d_{\mathcal{A},i}(\mathbf{a}_j^{\text{rep}}, \hat{\mathbf{a}}_j)$$

For the next point, note that the compatibility of the dynamics with the direct decomposition $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{\mathcal{O}}$ implies that there exists a dynamics map $F_{\mathcal{O},h}$ for which

$$F_h(\mathbf{s}, \mathbf{a}) = F_{\mathcal{O},h}(\phi_o(\mathbf{s}), \mathbf{a}).$$

Similarly, by applying \mathcal{Q}_{IPS} and $\mathcal{Q}_{\text{close}}$ and the event $\{\phi_o(\tilde{s}_h^{\text{rep}}) = \phi_o(\tilde{s}_h^{\text{tel}}), \mathbf{a}_h^{\text{tel}} = \mathbf{a}_h^{\text{rep}}\}$ on $\mathcal{B}_{\text{tel},h}$, it holds that on $\mathcal{Q}_{\text{all}} \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h}$ that, for all $h \in [H]$,

$$\begin{aligned}
\vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, F_h(\tilde{s}_h^{\text{rep}}, \mathbf{a}_h^{\text{rep}})) &= \vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, F_{o,h}(\phi_o(\tilde{s}_h^{\text{rep}}), \mathbf{a}_h^{\text{rep}})) \\
&= \vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, F_{o,h}(\phi_o(\tilde{s}_h^{\text{tel}}), \mathbf{a}_h^{\text{tel}})) & (\mathcal{B}_{\text{tel},h}) \\
&= \vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, F_h(\tilde{s}_h^{\text{tel}}, \mathbf{a}_h^{\text{tel}})) \\
&= \vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, \mathbf{s}_{h+1}^{\text{tel}}) \\
&\leq \tilde{\gamma}_{\text{IPS}} \circ \mathbf{d}_{\text{IPS}}(\mathbf{s}_j^{\text{tel}}, \mathbf{s}_j^{\text{tel}}) & (\mathcal{Q}_{\text{IPS}}) \\
&\leq \tilde{\gamma}_{\text{IPS}} \circ \mathbf{d}_{\text{IPS}}(2r). & (\mathcal{Q}_{\text{close}})
\end{aligned}$$

□

From Lemma H.7, we decompose our error probability as follows:

Lemma H.8 (Key Error Decomposition). Let μ respect the construction (in the sense of Appendix H.2.1). Then, for any coupling μ which respects the construction,

$$\vec{\Gamma}_{\text{joint},\vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \vee \vec{\Gamma}_{\text{marg},\vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi^*) \leq \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\bar{\mathcal{B}}_{\text{all},h}^c \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h-1}] \quad (\text{H.4})$$

Proof. In what follows, we use $\vec{v} \vee \vec{w}$ to denote the entrywise maximum of two vectors of the same dimension. Define the events $\mathcal{E}_h := \bar{\mathcal{C}}_{\text{all},h+1} \cap \bar{\mathcal{B}}_{\text{all},h}$. Observe that the events are nested: $\mathcal{E}_h \supset \mathcal{E}_{h+1}$, and that on \mathcal{E}_H , we have that for all $h \in [H]$

$$\begin{aligned}
\vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}) \vee \vec{d}_A(\mathbf{a}_h^{\text{rep}}, \hat{\mathbf{a}}_h) &\preceq \vec{\varepsilon} \vee \vec{d}_A(\mathbf{a}_h^{\text{rep}}, \hat{\mathbf{a}}_h) & (\mathcal{C}_{\hat{\mathbf{s}},h+1} \supset \bar{\mathcal{C}}_{\text{all},h+1} \supset \mathcal{E}_h) \\
&\preceq \vec{\varepsilon}. & (\bar{\mathcal{B}}_{\text{all},h} \supset \mathcal{E}_h)
\end{aligned}$$

On $\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H$, we have that

$$\max_h \vec{d}_S(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) \leq \tilde{\gamma}_{\text{IPS}}(2r), \quad \text{and} \quad \mathbf{a}_h^{\text{tel}} = \mathbf{a}_h^{\text{rep}}$$

Thus, by the triangle inequality and $\vec{\varepsilon}_{\text{marg}} = \vec{\varepsilon} + \tilde{\gamma}_{\text{IPS}}(2r)$, on $\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H$,

$$\max_h \vec{d}_S(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) \leq \vec{\varepsilon}_{\text{marg}}, \quad \text{and} \quad \vec{d}_A(\mathbf{a}_h^{\text{tel}}, \hat{\mathbf{a}}_h) = \vec{d}_A(\mathbf{a}_h^{\text{rep}}, \hat{\mathbf{a}}_h) \leq \vec{\varepsilon} \leq \vec{\varepsilon}_{\text{marg}}.$$

Thus,

$$\begin{aligned}
&\mathbb{P}_\mu \left[\exists h \in [H] : \left\{ \vec{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}) \vee \vec{d}_A(\mathbf{a}_h^{\text{rep}}, \hat{\mathbf{a}}_h) \not\preceq \vec{\varepsilon} \right\} \cup \left\{ \vec{d}_S(\mathbf{s}_{h+1}^{\text{tel}}, \hat{\mathbf{s}}_{h+1}) \vee \vec{d}_A(\mathbf{a}_h^{\text{tel}}, \hat{\mathbf{a}}_h) \not\preceq \vec{\varepsilon}_{\text{marg}} \right\} \right] \\
&\leq \mathbb{P}_\mu[(\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H)^c] & (\text{H.5})
\end{aligned}$$

In particular, this shows that

$$\vec{\Gamma}_{\text{joint},\vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \leq \mathbb{P}_\mu[(\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H)^c],$$

and similarly, by Lemma H.6,

$$\vec{\Gamma}_{\text{marg},\vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi^*) \leq \mathbb{P}_\mu[(\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H)^c]$$

As $(\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}}) \sim \mathbf{D}_{\pi_{\text{rep}}^*}$, (H.5) shows that

$$\vec{\Gamma}_{\text{joint},\vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \vee \vec{\Gamma}_{\text{marg},\vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \leq \mathbb{P}_\mu[(\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H)^c].$$

Let us conclude by bounding $\mathbb{P}_\mu[(\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H)^c]$. Using the nesting structure $\mathcal{E}_h = \bigcap_{j=1}^h \mathcal{E}_j$, the peeling lemma, [Lemma F.11](#), and a union bound, it holds that

$$\begin{aligned}
\mathbb{P}_\mu[(\mathcal{Q}_{\text{all}} \cap \mathcal{E}_H)^c] &\leq \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \mathbb{P}[\exists h \in [H] \text{ s.t. } (\mathcal{Q}_{\text{all}} \cap \mathcal{E}_{h-1} \cap \mathcal{E}_h^c) \text{ holds}] \\
&\leq \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\mathcal{Q}_{\text{all}} \cap \mathcal{E}_{h-1} \cap \mathcal{E}_h^c \text{ holds}] \\
&= \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\mathcal{Q}_{\text{all}} \cap \bar{\mathcal{B}}_{\text{all},h-1} \cap \bar{\mathcal{C}}_{\text{all},h} \cap (\bar{\mathcal{B}}_{\text{all},h} \cap \bar{\mathcal{C}}_{\text{all},h+1})^c \text{ holds}] \\
&= \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\mathcal{Q}_{\text{all}} \cap \bar{\mathcal{B}}_{\text{all},h-1} \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h}^c] \\
&= \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\mathcal{Q}_{\text{all}} \cap \bar{\mathcal{B}}_{\text{all},h-1} \cap \bar{\mathcal{C}}_{\text{all},h} \cap \mathcal{B}_{\text{all},h}^c],
\end{aligned}$$

where the last step invokes [Lemma H.7](#). \square

Next, we bound the contribution of $\mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c]$ in (H.4), uniformly over all couplings.

Lemma H.9. For all μ which respect the construction,

$$\mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] \leq p_{\text{IPS}} + 2Hp_r.$$

Proof. $\mathbb{P}_\mu[\mathcal{Q}_{\text{close}}^c] = \mathbb{P}_\mu[\exists h : d_{\text{IPS}}(\mathbf{s}_h^{\text{tel}}, \tilde{\mathbf{s}}_h^{\text{tel}}) > 2r] \leq 2Hp_r$ by [Lemma H.5](#) and a union bound.

Let us now bound $\mathbb{P}_\mu[\mathcal{Q}_{\text{close}} \cap \mathcal{Q}_{\text{IPS}}^c] \leq \mathbb{P}_\mu[\mathcal{Q}_{\text{IPS}}^c \mid \mathcal{Q}_{\text{close}}]$. Define the kernels $W_h(\mathbf{s})$ to be equal to the kernel $W_{\text{rep},h}(\mathbf{s})$ conditioned on the event $\mathbf{s}' \sim W_{\text{rep},h}(\mathbf{s})$ satisfies $d_{\text{IPS}}(\mathbf{s}', \mathbf{s}) \leq 2r$. Then, conditional on $\mathcal{Q}_{\text{close}}$, we see that the sequence $(\mathbf{s}_{1:H+1}^{\text{rep}}, \tilde{\mathbf{s}}_{1:H}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}})$ obeys the generative process $\tilde{\mathbf{s}}_h^{\text{rep}} \mid \tilde{\mathbf{s}}_{1:h-1}^{\text{rep}}, \mathbf{s}_{1:h}^{\text{rep}}, \mathbf{a}_{1:h-1}^{\text{rep}} \sim W_h(\mathbf{s})$, $\mathbf{a}_h^{\text{rep}} \mid \tilde{\mathbf{s}}_{1:h}^{\text{rep}}, \mathbf{s}_{1:h}^{\text{rep}}, \mathbf{a}_{1:h-1}^{\text{rep}} \sim \pi_h^*(\tilde{\mathbf{s}}_h^{\text{rep}})$, $\mathbf{s}_{h+1}^{\text{rep}} = F_h(\mathbf{s}_h^{\text{rep}}, \mathbf{a}_h^{\text{rep}})$. By construction, for each h , $\mathbb{P}_{\mathbf{s}' \sim W_{\text{rep},h}(\mathbf{s})}[d_{\text{IPS}}(\mathbf{s}', \mathbf{s}) > 2r] = 0$. Thus, the definition of (vector) input process stability ([Definition B.4](#)) and assumption $r \leq \frac{1}{2}r_{\text{IPS}}$ implies that $\mathbb{P}_\mu[\mathcal{Q}_{\text{IPS}}^c \mid \mathcal{Q}_{\text{close}}] \leq p_{\text{IPS}}$. \square

The remaining step of the proof is therefore to bound the second term in (H.4).

Lemma H.10. There exists a coupling μ which respects the construction and satisfies the following for any $h \in [H]$

$$\begin{aligned}
&\mathbb{P}_\mu[\mathcal{B}_{\text{all},h}^c \mid \mathcal{F}_{h-1}] \\
&\leq \hat{\gamma} \circ d_{\text{TVC}}(\mathbf{s}_h^{\text{rep}}, \hat{\mathbf{s}}_h) + (\hat{\gamma} + \gamma_\sigma) \circ d_{\text{TVC}}(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) + \vec{d}_{\text{os},\varepsilon}(\hat{\pi}_{\sigma,h}(\mathbf{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}})), \mu\text{-almost surely}
\end{aligned}$$

Consequently, for all $h \in [H]$,

$$\begin{aligned}
&\mathbb{P}_\mu[\mathcal{B}_{\text{all},h}^c \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h-1}] \\
&\leq \hat{\gamma}(\varepsilon_1) + (\hat{\gamma} + \gamma_\sigma) \circ \gamma_{\text{IPS},\text{TVC}}(2r) + \mathbb{E}_\mu[\vec{d}_{\text{os},\varepsilon}(\hat{\pi}_{\sigma,h}(\mathbf{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}}))]
\end{aligned}$$

Moreover, $\mathbf{s} \mapsto \vec{d}_{\text{os},\varepsilon}(\hat{\pi}_{\sigma,h}(\mathbf{s}) \parallel \pi_{\text{rep},h}^*(\mathbf{s}))$ is measurable.

Proof Sketch. We begin by giving a high level overview of the construction, which is done recursively. The key technical tool is [Lemma F.2](#) above, which allows us to transform any coupling μ between random variables (X, Y) into a probability kernel $\mu(\cdot \mid X)$ mapping instances of X to probability distributions on Y such that $(X, Y) \sim \mu$ has the same law as $(X, Y \sim \mu(\cdot \mid X))$. For each h , we then show that, assuming the coupling has kept the states and controls close together until time $h-1$, this will imply the following chain:

$$\underbrace{(\mathbf{a}^{\text{rep}} \leftrightarrow \mathbf{a}^{\text{tel}})}_{\gamma_{\text{TVC}} \text{ and induction}} \rightarrow \underbrace{(\mathbf{a}^{\text{tel}} \leftrightarrow \hat{\mathbf{a}}^{\text{tel,inter}})}_{\text{learning and sampling}} \rightarrow \underbrace{(\hat{\mathbf{a}}^{\text{tel,inter}} \leftrightarrow \hat{\mathbf{a}}^{\text{rep,inter}})}_{\gamma_{\text{TVC}} \text{ and induction}} \rightarrow \underbrace{(\hat{\mathbf{a}}^{\text{rep,inter}} \leftrightarrow \hat{\mathbf{a}})}_{\gamma_{\text{TVC}} \text{ and induction}},$$

where the bidirectional arrows indicate individual couplings between the laws of the random variables that are constructed by the method outlined in text below and the single directional arrows denote the probability kernels described above. The full proof of the lemma is given in [Appendix H.2.5](#). \square

Concluding the proof. Here, we finish the proof of [Theorem 8](#). Recall that we wish to bound $\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \vee \vec{\Gamma}_{\text{marg}, \vec{\varepsilon}_{\text{marg}}}(\hat{\pi}_\sigma \parallel \pi^*)$. We begin by bounding $\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \vee \vec{\Gamma}_{\text{marg}, \vec{\varepsilon}_{\text{marg}}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*)$. In light of [Lemma H.8](#), it suffices to bound

$$\mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\bar{\mathcal{B}}_{\text{all}, h}^c \cap \bar{\mathcal{C}}_{\text{all}, h} \cap \bar{\mathcal{B}}_{\text{all}, h-1}],$$

where μ is the coupling in [Lemma H.10](#). Applying [Lemma H.9](#) and [Lemma H.10](#),

$$\begin{aligned} & \mathbb{P}_\mu[\mathcal{Q}_{\text{all}}^c] + \sum_{h=1}^H \mathbb{P}_\mu[\bar{\mathcal{B}}_{\text{all}, h}^c \cap \bar{\mathcal{C}}_{\text{all}, h} \cap \bar{\mathcal{B}}_{\text{all}, h-1}] \\ & \leq p_{\text{IPS}} + 2Hp_r + \sum_{h=1}^H \mathbb{P}_\mu[\bar{\mathcal{B}}_{\text{all}, h}^c \cap \bar{\mathcal{C}}_{\text{all}, h} \cap \bar{\mathcal{B}}_{\text{all}, h-1}] \\ & \leq p_{\text{IPS}} + H(2p_r + \hat{\gamma}(\vec{\varepsilon}_1) + (\hat{\gamma} + \gamma_\sigma) \circ \gamma_{\text{IPS}, \text{Tvc}}(2r)) + \sum_{h=1}^H \mathbb{E}_{\mathbf{s}_h^{\text{tel}} \sim \mu} \vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}_{\sigma, h}(\mathbf{s}_h^{\text{tel}}) \parallel \pi_{\text{rep}, h}^*(\mathbf{s}_h^{\text{tel}})) \end{aligned}$$

To conclude, we note that for any μ which respects the construction, [Lemma H.3](#) ensures that $\mathbf{s}_h^{\text{tel}}$ as the marginal distribution of $\mathbf{s}_h^* \sim \pi_h^*$. Thus, the above is at most

$$p_{\text{IPS}} + H(2p_r + \hat{\gamma}(\vec{\varepsilon}_1) + (\hat{\gamma} + \gamma_\sigma) \circ \gamma_{\text{IPS}, \text{Tvc}}(2r)) + \sum_{h=1}^H \mathbb{E}_{\mathbf{s}_h^* \sim \mathbf{P}_h^*} \vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}_{\sigma, h}(\mathbf{s}_h^*) \parallel \pi_{\text{rep}, h}^*(\mathbf{s}_h^*)) \quad (\text{H.6})$$

which concludes the proof of (H.2) for $\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi} \parallel \pi_{\text{rep}}^*)$.

To prove (H.3) for $\vec{\Gamma}_{\text{joint}, \vec{\varepsilon}}(\hat{\pi} \parallel \pi_{\text{rep}}^*)$, we consider the special case that $\hat{\pi}_\sigma = \hat{\pi} \circ W_\sigma$. By definition, $\hat{\pi}_{\sigma, h} = \hat{\pi} \circ W_\sigma$. Thus, the data-processing inequality for optimal transport ([Lemma F.5](#))

$$\vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}_{\sigma, h}(\mathbf{s}_h^*) \parallel \pi_{\text{rep}, h}^*(\mathbf{s}_h^*)) \leq \mathbb{E}_{\mathbf{s}_h' \sim W_\sigma(\mathbf{s}_h^*)} \vec{d}_{\text{os}, \vec{\varepsilon}}(\hat{\pi}(\mathbf{s}_h') \parallel \pi_{\text{dec}, h}^*(\mathbf{s}_h')),$$

for all \mathbf{s}_h^* . Substituting this into (H.6), and setting $\hat{\gamma} = \gamma_\sigma$ (in view of [Lemma H.2](#)), finishes the argument.

H.2.5 Proof of Lemma H.10

Recall that [Assumption H.1](#) ensures all of the general measure-theoretic guarantees of [Appendix F](#) hold true in our setting. Notably we need the gluing lemma ([Lemma F.2](#)) and the commuting of optimal transport metrics and conditional probabilities ([Proposition F.3](#)).

Proof strategy. Our proof follows along similar lines as that of [Proposition D.2](#), although with the added complication of including the smoothing. We will inductively construct μ . A useful schematic for the construction at each step is the following diagram:

$$\underbrace{(\tilde{\mathbf{s}}^{\text{rep}} \leftrightarrow \tilde{\mathbf{s}}^{\text{tel}}), (\mathbf{a}^{\text{rep}} \leftrightarrow \mathbf{a}^{\text{tel}})}_{\mathcal{B}_{\text{tel}, h}} \rightarrow \underbrace{(\mathbf{a}^{\text{tel}} \leftrightarrow \hat{\mathbf{a}}^{\text{tel}, \text{inter}})}_{\mathcal{B}_{\text{est}, h}} \rightarrow \underbrace{(\hat{\mathbf{a}}^{\text{tel}, \text{inter}} \leftrightarrow \hat{\mathbf{a}}^{\text{rep}, \text{inter}})}_{\mathcal{B}_{\text{inter}, h}} \rightarrow \underbrace{(\hat{\mathbf{a}}^{\text{rep}, \text{inter}} \leftrightarrow \hat{\mathbf{a}})}_{\mathcal{B}_{\hat{\mathbf{a}}, h}},$$

where the events under each bidirectional arrow refer to the event such ensuring that there exists a coupling such that the objects are close. We then will apply [Lemma F.2](#) to glue the individual couplings together. We will then use [Lemma F.11](#) and a union bound to control the probability under our constructed coupling that any of the relevant events fail to hold, concluding the proof.

Recursive construction of μ . Let $h \geq 1$, and suppose that we have constructed the coupling $\mu^{(1:h-1)}$ for steps $1, \dots, h-1$ which respects the construction. Recall that \mathcal{F}_h denotes the sigma-algebra generated by $(\hat{\mathbf{s}}_{1:h}, \mathbf{s}_{1:h}^{\text{rep}}, \mathbf{s}_{1:h}^{\text{tel}})$, $(\mathbf{a}_{1:h}^{\text{rep}}, \tilde{\mathbf{s}}_{1:h}^{\text{rep}}, \tilde{\mathbf{s}}_{1:h}^{\text{tel}}, \mathbf{a}_{1:h}^{\text{tel}}, \hat{\mathbf{a}}_{1:h})$, and $(\hat{\mathbf{a}}_{1:h}^{\text{rep}, \text{inter}}, \hat{\mathbf{a}}_{1:h}^{\text{tel}, \text{inter}})$. Notice that $\mathbf{s}_{h+1}^{\text{tel}}, \mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}$ are determined by \mathcal{F}_h as well. Similarly, it can be seen from [Definition H.5](#) that $\phi_{/o}(\tilde{\mathbf{s}}_{h+1}^{\text{tel}})$ and $\phi_{/o}(\tilde{\mathbf{s}}_{h+1}^{\text{rep}})$ are also determined by \mathcal{F}_h (since the replica kernel preserves the

$\mathcal{S}_{/\mathcal{O}}$ -components). We summarize all these aforementioned variables in a random variable Y_h . Let \mathcal{F}_0 denote the filtration generated by $s_1^{\text{rep}} = s_1^{\text{tel}} = \hat{s}_1$. We let $Y_0 = (s_1^{\text{rep}}, s_1^{\text{tel}}, \hat{s}_1)$.

Correspondingly, let Z_h denote the random variables $(a_h^{\text{rep}}, \phi_o(\tilde{s}_h^{\text{rep}}), \phi_o(\tilde{s}_h^{\text{tel}}), a_h^{\text{tel}}, \hat{a}_h)$, and $(\hat{a}_h^{\text{rep,inter}}, \hat{a}_h^{\text{tel,inter}})$ such that the joint law of these random variables respects the construction. Our goal is then to specify, for each $h \in [H]$, a joint distribution of (Y_{h-1}, Z_h) . Note that Z_h, Y_{h-1} determines Y_h , and we call this induced law $\mu^{(h)}$.

We begin by specifying joint distributions conditional on Y_{h-1} and subsets of Z_h , then glue them together by the gluing lemma. Below, we use information-theoretic notation.

- By total variation continuity of $\phi_o \circ W_{\text{rep},h}^*$ (Lemma H.2),

$$\text{TV}(\mathbb{P}_{\phi_o(\tilde{s}_h^{\text{rep}})|Y_{h-1}}, \mathbb{P}_{\phi_o(\tilde{s}_h^{\text{tel}})|Y_{h-1}}) \leq \gamma_\sigma \circ d_{\text{TVC}}(s_h^{\text{rep}}, s_h^{\text{tel}}).$$

Because $a_h^{\text{rep}} \sim \pi_h^*(\tilde{s}_{h+1}^{\text{rep}})$ and $a_h^{\text{tel}} \sim \pi_h^*(\tilde{s}_h^{\text{tel}})$, and π^* is compatible with the decomposition $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ (i.e. $\pi_h^*(s)$ is a function of $\phi_o(s)$) Lemma F.4 implies that (almost surely)

$$\text{TV}(\mathbb{P}_{(a_h^{\text{rep}}, \phi_o(\tilde{s}_h^{\text{rep}}))|Y_{h-1}}, \mathbb{P}_{(a_h^{\text{tel}}, \phi_o(\tilde{s}_h^{\text{tel}}))|Y_{h-1}}) \leq \gamma_\sigma \circ d_{\text{TVC}}(s_h^{\text{rep}}, s_h^{\text{tel}}).$$

Hence, Corollary F.1 implies that there exists a coupling $\mu_{\text{tel}}^{(h)}$ over $Y_{h-1}, (\phi_o(\tilde{s}_h^{\text{rep}}), a_h^{\text{rep}}), (\phi_o(\tilde{s}_h^{\text{tel}}), a_h^{\text{tel}})$ respecting the construction such that $Y_h \sim \mu^{(h-1)}$ and such that (almost surely)

$$\mathbb{E}_{\mu_{\text{tel}}^{(h)}}[\mathcal{B}_{\text{tel},h} | Y_{h-1}] = \mathbb{P}_{\mu_{\text{tel}}^{(h)}}[(\phi_o(\tilde{s}_h^{\text{rep}}), a_h^{\text{rep}}) \neq (\phi_o(\tilde{s}_h^{\text{tel}}), a_h^{\text{tel}}) | Y_{h-1}] \leq d_{\text{TVC}}(s_h^{\text{rep}}, s_h^{\text{tel}}).$$

- In our construction, $a_h^{\text{tel}} | Y_{h-1} \sim \pi_{\text{rep},h}^*(s_h^{\text{tel}})$, and $\hat{a}_h^{\text{tel,inter}} | Y_{h-1} \sim \hat{\pi}_{\sigma,h}(s_h^{\text{tel}})$. Thus, by definition of $\vec{d}_{\text{os},\vec{\varepsilon}}$, and the assumption $\mathbf{I}\{\vec{d}_{\mathcal{A}}(\cdot, \cdot) \not\leq \vec{\varepsilon}\}$ is lower semicontinuous, Proposition F.3 implies that we may find a coupling $\mu_{\text{est}}^{(h)}$ of $(a_h^{\text{tel}}, \hat{a}_h^{\text{tel,inter}}, Y_{h-1})$ respecting the construction such that, almost surely,

$$\begin{aligned} \mathbb{P}_{\mu_{\text{est}}^{(h)}}[\mathcal{B}_{\text{est},h}^c | Y_{h-1}] &= \mathbb{P}_{\mu_{\text{est}}^{(h)}}[\vec{d}_{\mathcal{A}}(\hat{a}_h^{\text{tel,inter}}, a_h^{\text{tel}}) \not\leq \vec{\varepsilon} | Y_{h-1}] \\ &= \vec{d}_{\text{os},\vec{\varepsilon}}(\hat{\pi}_{\sigma,h}(s_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(s_h^{\text{tel}})). \end{aligned}$$

Moreover, that same proposition ensures measurability of $s \rightarrow \vec{d}_{\text{os},\vec{\varepsilon}}(\hat{\pi}_{\sigma,h}(s) \parallel \pi_{\text{rep},h}^*(s))$.

- Since $\hat{a}_h^{\text{tel,inter}} | \mathcal{F}_h \sim \hat{\pi}_{\sigma,h}(s_h^{\text{tel}})$ and $\hat{a}_{h+1}^{\text{rep,inter}} | \mathcal{F}_h \sim \hat{\pi}_{\sigma,h}(s_h^{\text{rep}})$, and since $\hat{\pi}_{\sigma,h}(\cdot)$ is $\hat{\gamma}$ -TVC by assumption,

$$\text{TV}(\mathbb{P}_{\hat{a}_h^{\text{tel,inter}}|Y_{h-1}}, \mathbb{P}_{\hat{a}_h^{\text{rep,inter}}|Y_{h-1}}) \leq \hat{\gamma} \circ d_{\text{TVC}}(s_h^{\text{rep}}, s_h^{\text{tel}}).$$

Corollary F.1 implies that there is a coupling $\mu_{\text{inter}}^{(h)}$ between $(\hat{a}_h^{\text{tel,inter}}, \hat{a}_h^{\text{rep,inter}}, Y_{h-1})$ such that

$$\mathbb{P}_{\mu_{\text{inter}}^{(h)}}[\mathcal{B}_{\text{inter},h}^c | Y_{h-1}] = \mathbb{P}_{\mu_{\text{inter}}^{(h)}}[\hat{a}_h^{\text{tel,inter}} \neq \hat{a}_h^{\text{rep,inter}} | Y_{h-1}] \leq \hat{\gamma} \circ d_{\text{TVC}}(s_h^{\text{tel}}, s_h^{\text{rep}})$$

- Similarly, since $\hat{a}_h^{\text{rep,inter}} | \mathcal{F}_{h-1} \sim \hat{\pi}_h(s_h^{\text{rep}})$ and $\hat{a}_{h+1} | \mathcal{F}_{h-1} \sim \hat{\pi}_h(\hat{s}_h)$, $\hat{\pi}_h(\cdot)$ is $\hat{\gamma}$ -TVC, Corollary F.1 implies that there is a coupling $\mu_{\hat{a}}^{(h)}$ between $(\hat{a}_h^{\text{rep,inter}}, \hat{a}_h, Y_{h-1})$ such that

$$\mathbb{P}_{\mu_{\hat{a}}^{(h)}}[\mathcal{B}_{\hat{a},h}^c | Y_{h-1}] = \mathbb{P}_{\mu_{\hat{a}}^{(h)}}[\hat{a}_h \neq \hat{a}_h^{\text{rep,inter}} | Y_{h-1}] \leq \hat{\gamma} \circ d_{\text{TVC}}(s_h^{\text{rep}}, \hat{s}_h)$$

We can then apply the gluing lemma (Lemma F.2) to

$$\begin{aligned} X_{h,1} &= (\phi_o(\tilde{s}_h^{\text{tel}}), a_h^{\text{tel}}, Y_{h-1}) \\ X_{h,2} &= (\phi_o(\tilde{s}_h^{\text{rep}}), a_h^{\text{rep}}, Y_{h-1}) \\ X_{h,3} &= (a_h^{\text{tel}}, \hat{a}_h^{\text{tel,inter}}, Y_{h-1}) \\ X_{h,4} &= (\hat{a}_h^{\text{tel,inter}}, \hat{a}_h^{\text{rep,inter}}, Y_{h-1}) \\ X_{h,5} &= (\hat{a}_h^{\text{rep,inter}}, \hat{a}_h, Y_{h-1}) \end{aligned}$$

with

$$(X_{h,1}, X_{h,2}) \sim \mu_{\text{tel}}^{(h)}, \quad (X_{h,2}, X_{h,3}) \sim \mu_{\text{est}}^{(h)}, \quad (X_{h,3}, X_{h,4}) \sim \mu_{\text{inter}}^{(h)}, \quad (X_{h,4}, X_{h,5}) \sim \mu_{\hat{\mathbf{a}}}^{(h)}.$$

Lemma F.2 guarantees the existence of a coupling $\mu^{(h)}$ consistent with all sub-couplings $\mu_{\text{tel}}^{(h)}$, $\mu_{\text{est}}^{(h)}$, $\mu_{\text{inter}}^{(h)}$, $\mu_{\hat{\mathbf{a}}}^{(h)}$. Then, $\mu^{(h)}$ -almost surely (and using that \mathcal{F}_{h-1} is precisely the σ -algebra generated by Y_{h-1})

$$\begin{aligned} & \mathbb{P}_{\mu^{(h)}}[\mathcal{B}_{\text{all},h}^c \mid \mathcal{F}_{h-1}] \\ & \leq \mathbb{P}_{\mu^{(h)}}[\mathcal{B}_{\text{tel},h}^c \mid \mathcal{F}_{h-1}] + \mathbb{P}_{\mu^{(h)}}[\mathcal{B}_{\text{est},h}^c \mid \mathcal{F}_{h-1}] + \mathbb{P}_{\mu^{(h)}}[\mathcal{B}_{\text{inter},h}^c \mid \mathcal{F}_{h-1}] + \mathbb{P}_{\mu^{(h)}}[\mathcal{B}_{\hat{\mathbf{a}},h}^c \mid \mathcal{F}_{h-1}] \\ & \leq \hat{\gamma} \circ \mathbf{d}_{\text{TVC}}(\mathbf{s}_h^{\text{rep}}, \hat{\mathbf{s}}_h) + (\hat{\gamma} + \gamma_\sigma) \circ \mathbf{d}_{\text{TVC}}(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) + \vec{\mathbf{d}}_{\text{os},\vec{\varepsilon}}(\hat{\pi}_{\sigma,h}(\mathbf{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}})) \\ & = \hat{\gamma} \circ \mathbf{d}_{\text{TVC}}(\mathbf{s}_h^{\text{rep}}, \hat{\mathbf{s}}_h) + (\hat{\gamma} + \gamma_\sigma) \circ \mathbf{d}_{\text{TVC}}(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) + \vec{\mathbf{d}}_{\text{os},\vec{\varepsilon}}(\hat{\pi}_{\sigma,h}(\mathbf{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}})) \end{aligned}$$

This concludes the inductive construction.

For the second statement, notice that the events $\bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h-1}$ are \mathcal{F}_h measurable (thus determined by $\mu^{(h-1)}$) and, when they hold, $\vec{\mathbf{d}}_{\mathcal{S}}(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) \leq \vec{\gamma}_{\text{IPS}}(2r)$ and $\mathbf{d}_{\mathcal{S}}(\mathbf{s}_h^{\text{rep}}, \hat{\mathbf{s}}_h) \leq \vec{\varepsilon}$. For our purposes, we use $\mathbf{d}_{\text{TVC}} = \mathbf{d}_{\mathcal{S},1}(\mathbf{s}_h^{\text{rep}}, \mathbf{s}_h^{\text{tel}}) \leq \gamma_{\text{IPS},\text{TVC}}(2r)$ and $\mathbf{d}_{\mathcal{S}}(\mathbf{s}_h^{\text{rep}}, \hat{\mathbf{s}}_h) \leq \vec{\varepsilon}_1$. Hence,

$$\begin{aligned} \max_{h \in [H]} \mathbb{P}_{\mu}[\mathcal{B}_{\text{all},h}^c \cap \bar{\mathcal{C}}_{\text{all},h} \cap \bar{\mathcal{B}}_{\text{all},h-1}] & \leq \hat{\gamma}(\vec{\varepsilon}_1) + (\hat{\gamma} + \gamma_\sigma) \circ \gamma_{\text{IPS},\text{TVC}}(2r) \\ & \quad + \vec{\mathbf{d}}_{\text{os},\vec{\varepsilon}}(\hat{\pi}_{\sigma,h}(\mathbf{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\mathbf{s}_h^{\text{tel}})). \end{aligned}$$

The result follows.

H.3 Proof of Theorem 4, and generalization to direct decompositions

In this subsection, we consider the special case dealt with in Theorem 4. Note that there always exists a trivial direct decomposition that is compatible with all policies and dynamics simply by letting $\mathcal{S}_{/\mathcal{O}} = \emptyset$ and $\mathcal{S} = \mathcal{O}$. We prove here the version of the result that involves a possibly nontrivial direct decomposition, as we will instantiate this in our control setting by letting $\mathcal{O} = \{\mathbf{o}_h\}$ and $\mathcal{S} = \{\mathbf{s}_h\}$, i.e., projecting \mathbf{s}_h onto the last τ_{obs} coordinates gives \mathbf{o}_h . We further consider a restriction of IPS to consider kernels absolutely continuous with respect to \mathbf{P}_h^* in their \mathcal{O} component.

Definition H.9 (Restricted IPS). For a non-decreasing maps $\gamma_{\text{IPS},\text{TVC}}, \gamma_{\text{IPS},\mathcal{S}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ a pseudometric $\mathbf{d}_{\text{IPS}} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ (possibly other than $\mathbf{d}_{\mathcal{S}}$ or \mathbf{d}_{TVC}), and $r_{\text{IPS}} > 0$, we say a policy π is $(\gamma_{\text{IPS},\text{TVC}}, \gamma_{\text{IPS},\mathcal{S}}, \mathbf{d}_{\text{IPS}}, r_{\text{IPS}})$ -restricted IPS if the following holds for any $r \in [0, r_{\text{IPS}}]$. Consider any sequence of kernels $W_1, \dots, W_H : \mathcal{S} \rightarrow \Delta(\mathcal{S})$ satisfying

$$\max_{h, \mathbf{s} \in \mathcal{S}} \mathbb{P}_{\tilde{\mathbf{s}} \sim W_h(\mathbf{s})}[\mathbf{d}_{\text{IPS}}(\tilde{\mathbf{s}}, \mathbf{s}) \leq r] = 1, \quad \forall \mathbf{s}, \quad \phi_{\mathcal{O}} \circ W_h(\mathbf{s}) \ll \phi_{\mathcal{O}} \circ \mathbf{P}_h^*.$$

and define a process $\mathbf{s}_1 \sim \mathbf{P}_{\text{init}}$, $\tilde{\mathbf{s}}_h \sim W_h(\mathbf{s}_h)$, $\mathbf{a}_h \sim \pi_h(\tilde{\mathbf{s}}_h)$, and $\mathbf{s}_{h+1} := F_h(\mathbf{s}_h, \mathbf{a}_h)$. Then, almost surely, (a) the sequence $(\mathbf{s}_{1:H+1}, \mathbf{a}_{1:H})$ is input-stable (in the sense of Definition D.2) (b) $\max_{h \in [H]} \mathbf{d}_{\text{TVC}}(F_h(\tilde{\mathbf{s}}_h, \mathbf{a}_h), \mathbf{s}_{h+1}) \leq \gamma_{\text{IPS},\text{TVC}}(r)$ and (c) $\max_{h \in [H]} \mathbf{d}_{\mathcal{S}}(F_h(\tilde{\mathbf{s}}_h, \mathbf{a}_h), \mathbf{s}_{h+1}) \leq \gamma_{\text{IPS},\mathcal{S}}(r)$.

Note that the above is a slightly weaker condition than the one in Definition D.5 in the main text and consequently, the following theorem which uses it as an assumption implies Theorem 4 in the body.

Theorem 9. Suppose that

- $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$ is in Definition B.1 and projections $\phi_{\mathcal{O}}, \phi_{\mathcal{S}_{/\mathcal{O}}}$, which is compatible with the dynamics and with given policies $\hat{\pi}, \pi^*$, smoothing kernel W_σ , and pseudometric \mathbf{d}_{IPS} .
- π^* satisfies $(\gamma_{\text{IPS},\text{TVC}}, \gamma_{\text{IPS},\mathcal{S}}, \mathbf{d}_{\text{IPS}}, r_{\text{IPS}})$ -restricted IPS (Definition H.9) and $\phi_{\mathcal{O}} \circ W_\sigma$ is $\gamma_{\sigma\text{-TVC}}$.

Given $\varepsilon > 0$ and $r \in (0, \frac{1}{2}r_{\text{IPS}}]$, define

$$p_r := \sup_{\mathbf{s}} \mathbb{P}_{\mathbf{s}' \sim W_\sigma(\mathbf{s})}[\mathbf{d}_{\text{IPS}}(\mathbf{s}', \mathbf{s}) > r], \quad \varepsilon' := \varepsilon + \gamma_{\text{IPS},\mathcal{S}}(2r)$$

Then, for any policy $\hat{\pi}$, both $\Gamma_{\text{joint},\varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$ and $\Gamma_{\text{marg},\varepsilon'}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ are upper bounded by

$$H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS},\text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec},h}^*(\tilde{s}_h^*)). \quad (\text{H.7})$$

Proof. Consider the special case $K = 2$ with $d_{S,1} = d_{\text{TVC}}$, $d_{S,2} = d_S$, $d_{A,1} = d_{A,2} = d_A$, $p_{\text{IPS}} = 0$ and $\tilde{\varepsilon} = (\varepsilon, \varepsilon)$. In this case, applying (H.3) in Theorem 8, we see that

$$\begin{aligned} & \vec{\Gamma}_{\text{joint},\tilde{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \vee \vec{\Gamma}_{\text{marg},\tilde{\varepsilon}_{\text{marg}}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \\ & \leq p_{\text{IPS}} + H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS},\text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} \vec{d}_{\text{os},\tilde{\varepsilon}}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*)) \end{aligned}$$

We now observe that under this convention,

$$\begin{aligned} \Gamma_{\text{joint},\varepsilon}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) &= \inf_{\mu_1} \mathbb{P}_{\mu_1}[\max_{h \in [H]} d_S(\hat{s}_{h+1}, s_{h+1}^*) \vee d_A(\hat{a}_h, a_h^*) > \varepsilon] \\ &\leq \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\max_{h \in [H]} (d_{\text{TVC}}(\hat{s}_{h+1}, s_{h+1}^*), d_S(\hat{s}_{h+1}, s_{h+1}^*)) \vee (d_A(\hat{a}_h, a_h^*), d_A(\hat{a}_h, a_h^*)) \not\leq \tilde{\varepsilon} \right] \\ &= \vec{\Gamma}_{\text{joint},\tilde{\varepsilon}}(\hat{\pi}_\sigma \parallel \pi_{\text{rep}}^*) \end{aligned}$$

and similarly $\Gamma_{\text{marg},\varepsilon'}(\hat{\pi}_\sigma \parallel \pi^*) \leq \vec{\Gamma}_{\text{marg},\tilde{\varepsilon} + \gamma_{\text{IPS}}(2r)}(\hat{\pi}_\sigma \parallel \pi^*)$. From the construction of \vec{d}_A , however, we see that $\{\vec{d}_A(a, a') \not\leq \tilde{\varepsilon}\} = \{d_A(a, a') > \varepsilon\}$ for all a, a' and thus for all $h \in [H]$,

$$\begin{aligned} \vec{d}_{\text{os},\tilde{\varepsilon}}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_h^*(\tilde{s}_h^*)) &= \inf_{\mu_2} \mathbb{P}_{\mu_2} [\vec{d}_A(\hat{a}_h, a_h^*) \not\leq \tilde{\varepsilon}] \\ &= \inf_{\mu_2} \mathbb{P}_{\mu_2} [d_A(\hat{a}_h, a_h^*) \geq \varepsilon] \\ &= d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_h^*(\tilde{s}_h^*)). \end{aligned}$$

Plugging in to (H.3) concludes the proof. \square

I Lower Bounds

In this section, we establish lower bounds against the imitation results in the composite MDP. Specifically, we show that

- In Appendix I.1 we show that Theorem 4 and Proposition D.2 are sharp in the regime where $\gamma_{\text{IPS},\text{TVC}} = \gamma_{\text{IPS},S} = 0$.
- In Appendix I.2, we show that the marginals of an expert policy π^* and replica policy π_{rep}^* can coincide, but their joint distributions can be different. By considering $\hat{\pi} = \pi_{\text{dec}}^*$ in Theorem 4, this establishes the necessity of considering the marginal imitation gap with respect to π^* .
- In Appendix I.3, we lower bound the distance between *marginal distributions* over states under π^* and π_{rep}^* in the regime where $\gamma_{\text{IPS},S} \neq 0$. This example demonstrates that the dependence of $\gamma_{\text{IPS},S}$ in Theorem 4 is essentially sharp.
- In Appendix I.4, we show that for an expert policy π^* and smoothing kernel W_σ , the state distributions under π_{rep}^* and π_{dec}^* can have different marginals (and thus different joint distributions). By considering $\hat{\pi} = \pi_{\text{dec}}^*$ in Theorem 4, this explains why it is necessary to smooth $\hat{\pi}$ to $\hat{\pi} \circ W_\sigma$.

Taken together, the above counterexamples show that our distinctions between joint and marginal distributions, decision to add noise at inference time, and dependence on almost all problem quantities in Appendix D are sharp. We do not, however, establish necessity of $\gamma_{\text{IPS},\text{TVC}}$ in the interest of brevity; we believe this quantity is necessary. Still, the $\gamma_{\text{IPS},\text{TVC}}$ term contributes a factor exponentially small in τ_{chunk} in Theorem 3, so we deem lower bounds establishing its necessity of lesser importance.

Commonalities of construction. In all but [Appendix I.3](#), we take the action and state spaces to be

$$\mathcal{S} = \mathcal{A} = \mathbb{R},$$

which is the archetypal Polish space [22]. Throughout, we use δ_x to denote the dirac-delta distribution on $x \in \mathbb{R}$. We let $d_S(s', s) = d_{\text{TVC}}(s', s) = |s' - s|$ and $d_A(a', a) = |a' - a|$ all be the Euclidean distance.

I.1 Sharpness of Proposition D.2 and Theorem 4

Here, we demonstrate that [Proposition D.2](#) is tight up to constant factors, and that [Theorem 4](#) is tight up to the terms $\gamma_{\text{IPS}, \text{TVC}}$, $\gamma_{\text{IPS}, S}$ and concentration probability p_r . Consider the simple dynamics

$$F_h(s, a) = a.$$

Note that, as the dynamics are state-independent, we have $\gamma_{\text{IPS}, \text{TVC}}(\cdot) = \gamma_{\text{IPS}, S}(\cdot) \equiv 0$. Furthermore, let us assume policies do not depend on time index h . Let $\pi^* : s \rightarrow \delta_0$ be deterministic, and let $P_{\text{init}} = \delta_0$ be an initial state distribution concentrated on 0. Then, D_{π^*} is the dirac distribution on the all-zero trajectory.

Fix parameters $0 < \varepsilon < \sigma$, and $p \in (0, 1)$. We consider the following smoothing-kernel

$$W_{\varepsilon, \sigma} = \begin{cases} \delta_0 & s \leq 0 \\ (1 - \frac{s}{\sigma})\delta_0 + \frac{s}{\sigma}\delta_\sigma & s \in [0, \sigma] \\ \delta_\sigma & s > \sigma, \end{cases}$$

Define the candidate policy

$$\hat{\pi}_{\varepsilon, p, \sigma}(s) := \begin{cases} (1 - p)\delta_\varepsilon + p\delta_\sigma & s \leq \frac{\varepsilon}{2} \\ \delta_\sigma & s > \frac{\varepsilon}{2} \end{cases}$$

Proposition I.1. For any $p \in (0, 1)$, $0 < \varepsilon < \sigma$, set $\bar{\pi} = \hat{\pi}_{\varepsilon, p, \sigma} \circ W_{\sigma, \varepsilon}$. Then,

- (a) π^* , π_{rep}^* and π_{dec}^* all map $s \rightarrow \delta_0$, $P_h^* = \delta_0$, and thus for any $\tilde{\pi} \in \{\pi^*, \pi_{\text{rep}}^*, \pi_{\text{dec}}^*\}$,

$$\mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{s'_h \sim W_\sigma(s_h^*)} [d_{\text{os}, \varepsilon}(\hat{\pi}_{\varepsilon, p, \sigma}(s'_h) \parallel \tilde{\pi}(s'_h))] = \mathbb{E}_{s_h^* \sim P_h^*} [d_{\text{os}, \varepsilon}(\bar{\pi}(s_h^*) \parallel \tilde{\pi}(s_h^*))] = p.$$

- (b) The kernel $W_{\sigma, \varepsilon}$ is γ_σ -TVC, where $\gamma_\sigma(u) = u/\sigma$.

- (c) For a universal constant $c > 0$,

$$\Gamma_{\text{joint}, \varepsilon}(\bar{\pi} \parallel \pi^*) = \Gamma_{\text{marg}, \varepsilon}(\bar{\pi} \parallel \pi^*) \geq c \min\{1, H(p + \varepsilon/\sigma)\},$$

and the same holds with π^* replaced by π_{rep}^* or π_{dec}^* .

In particular, the above proposition shows that

$$\Gamma_{\text{joint}, \varepsilon}(\bar{\pi} \parallel \pi^*) = \Gamma_{\text{marg}, \varepsilon}(\bar{\pi} \parallel \pi^*) \gtrsim H\gamma_\sigma(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} [d_{\text{os}, \varepsilon}(\bar{\pi}(s_h^*) \parallel \pi^*(s_h^*))],$$

verifying the sharpness of [Proposition D.2](#) (note that $\bar{\pi} = \hat{\pi}_{\varepsilon, p, \sigma} \circ W_\sigma$ is γ_σ TVC). Similarly, our above proposition shows that,

$$\Gamma_{\text{joint}, \varepsilon}(\bar{\pi} \parallel \pi_{\text{rep}}^*) = \Gamma_{\text{marg}, \varepsilon}(\bar{\pi} \parallel \pi^*) \gtrsim H\gamma_\sigma(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} [d_{\text{os}, \varepsilon}(\hat{\pi}_{\varepsilon, p, \sigma}(s_h^*) \parallel \pi_{\text{dec}, h}^*(s_h^*))],$$

verifying that [Theorem 4](#) is sharp up to the additional stability terms $\gamma_{\text{IPS}, \text{TVC}}$, $\gamma_{\text{IPS}, S}$.

Proof. We begin with a computation. Define

$$\eta(s) = 1 - (1 - p)(1 - \frac{s}{\sigma}) = p + (1 - p)\frac{s}{\sigma}$$

We compute

$$\begin{aligned}\bar{\pi} &= \hat{\pi}_{\varepsilon,p,\sigma} \circ W_{\sigma,\varepsilon} = \begin{cases} (1-p)\delta_\varepsilon + p\delta_\sigma & s \leq \frac{\varepsilon}{2} \\ \delta_\sigma & s > \frac{\varepsilon}{2} \end{cases} \circ \begin{cases} \delta_0 & s \leq 0 \\ (1 - \frac{s}{\sigma})\delta_0 + \frac{s}{\sigma}\delta_\sigma & s \in [0, \sigma] \\ \delta_\sigma & s \geq \sigma. \end{cases} \\ &= \begin{cases} (1-p)\delta_\varepsilon + p\delta_\sigma & s \leq 0 \\ (1 - \eta(s))\delta_\varepsilon + \eta(s)\delta_\sigma & 0 \leq s \leq \sigma \\ \delta_\sigma & s > \sigma. \end{cases} \end{aligned} \quad (\text{I.1})$$

In particular,

$$\hat{\pi}(0) = \pi_{\varepsilon,p,\sigma}(0) = (1-p)\delta_\varepsilon + p\delta_\sigma$$

Part (a). Notice that the support of the deconvolution and replica distributions are always in the support of P_h^* , which is always $s = 0$ under π^* . Thus, $\pi^* = \pi_{\text{rep}}^* = \pi_{\text{dec}}^*$. By the same token, for any policy π ,

$$\mathbb{E}_{s_h^* \sim P_h^*} [\mathbf{d}_{\text{os},\varepsilon}(\pi(s_h^*) \parallel \tilde{\pi}_*(s_h^*))] = \mathbb{P}[|\pi(0)| > \varepsilon].$$

Hence, as $\bar{\pi}(0) = \hat{\pi}_{\varepsilon,p,\sigma}(0) = (1-p)\delta_\varepsilon + p\delta_\sigma$, and as $\sigma > \varepsilon$, part (a) follows.

Part (b). Consider $s, s' \in S$. We can assume, from the functional form of $W_{\varepsilon,\sigma}(\cdot)$, that $0 \leq s \leq s' \leq \sigma$. Then,

$$\text{TV}(W_{\varepsilon,\sigma}(s), W_{\varepsilon,\sigma}(s')) = \text{TV}(\delta_0(1 - \frac{s}{\sigma}) + (\frac{s}{\sigma})\delta_\sigma, \delta_0(1 - \frac{s'}{\sigma}) + (\frac{s'}{\sigma})\delta_\sigma) = \frac{|s' - s|}{\sigma},$$

establishing total variation continuity.

Part (c) In view of part (a), it suffices to bound gaps relative to π^* . Let \mathbb{P} denote probabilities over $s_{1:H+1}, a_h$ under $\bar{\pi}$. Let $\mathcal{A}_{1,h}$ denote the event that at step h , $a_h = \varepsilon$, and let $\mathcal{A}_{2,h}$ denote the event that $a_h = \sigma$. As the state s_0 is absorbing and as $F_h(s, a) = a_h$, the following events are equal

$$\{\exists h : |a_h| \vee |s_{h+1}| > \varepsilon\} = \mathcal{A}_{2,H}.$$

Hence,

$$\Gamma_{\text{joint},\varepsilon}(\bar{\pi} \parallel \pi^*) = \mathbb{P}[\mathcal{A}_{2,H}].$$

Moreover, as $\mathcal{A}_{2,H}$ is measurable with respect to the marginal of a_H , we also have that

$$\Gamma_{\text{marg},\varepsilon}(\bar{\pi} \parallel \pi^*) = \mathbb{P}[\mathcal{A}_{2,H}].$$

It thus suffices to lower bound $\mathbb{P}[\mathcal{A}_{2,H}]$. By definition of $\bar{\pi}$, the events $\mathcal{A}_{1,h}, \mathcal{A}_{2,h}$ are exhaustive: $\mathcal{A}_{1,h}^c = \mathcal{A}_{2,h}$. Moreover, from (I.1),

$$\mathbb{P}[\mathcal{A}_{2,h+1} \mid \mathcal{A}_{2,h}] = 1, \quad \mathbb{P}[\mathcal{A}_{2,h+1} \mid \mathcal{A}_{1,h}] = \eta(\varepsilon), \quad \mathbb{P}[\mathcal{A}_{1,1}] = 1 - \eta(0) \geq 1 - \eta(\varepsilon).$$

Thus,

$$\begin{aligned}\mathbb{P}[\mathcal{A}_{2,H}] &= \mathbb{P}[\mathcal{A}_{2,H} \mid \mathcal{A}_{2,H-1}] \mathbb{P}[\mathcal{A}_{2,H-1}] + \mathbb{P}[\mathcal{A}_{2,H} \mid \mathcal{A}_{1,H-1}] \mathbb{P}[\mathcal{A}_{1,H-1}] \\ &= \mathbb{P}[\mathcal{A}_{2,H-1}] + \eta(\varepsilon) \mathbb{P}[\mathcal{A}_{1,H-1}] \\ &= \mathbb{P}[\mathcal{A}_{2,H-2}] + \eta(\varepsilon) (\mathbb{P}[\mathcal{A}_{1,H-1}] + \mathbb{P}[\mathcal{A}_{1,H-2}]) \\ &= \eta(\varepsilon) \left(\sum_{h=1}^{H-1} \mathbb{P}[\mathcal{A}_{1,h}] \right) + \mathbb{P}[\mathcal{A}_{2,1}] \\ &\geq \eta(\varepsilon) \left(\sum_{h=1}^{H-1} \mathbb{P}[\mathcal{A}_{1,h}] \right)\end{aligned}$$

Moreover, as s_0 is absorbing,

$$\mathbb{P}[\mathcal{A}_{1,h}] = \mathbb{P}[\mathcal{A}_{1,h} \mid \mathcal{A}_{1,h-1}] \mathbb{P}[\mathcal{A}_{1,h-1}] = (1 - \eta(\varepsilon)) \mathbb{P}[\mathcal{A}_{1,h-1}].$$

Combining with $\mathbb{P}[\mathcal{A}_{1,1}] = (1 - p) \geq (1 - \eta(0)) \geq 1 - \eta(\varepsilon)$, we have $\mathbb{P}[\mathcal{A}_{1,h}] \geq (1 - \eta(\varepsilon))^h$. Hence,

$$\begin{aligned}\mathbb{P}[\mathcal{A}_{2,H+1}] &\geq \eta(\varepsilon) \left(\sum_{h=1}^{H-1} (1 - \eta(\varepsilon))^h \right) \\ &= \eta(\varepsilon) \frac{1 - \eta(\varepsilon) - (1 - \eta(\varepsilon))^H}{1 - (1 - \eta(\varepsilon))} \\ &= 1 - \eta(\varepsilon) - (1 - \eta(\varepsilon))^H \\ &= \Omega(\min\{1, H(\eta(\varepsilon))\})\end{aligned}$$

as $\eta(\varepsilon) \downarrow 0$. Substituting in $\eta(\varepsilon) = p + (1 - p)\varepsilon/\sigma = \Omega(p + \varepsilon/\sigma)$ concludes. \square

I.2 π_{rep}^* and π^* induce the same marginals but different joint distributions, even with memoryless dynamics

We give a simple example where π_{rep}^* and π^* induce the same marginal distributions over trajectories, but different joints. As we show, this example demonstrates the necessity of measuring the marginal imitation error of a smoothed policy, $\Gamma_{\text{marg},\varepsilon}$, over the joint error, $\Gamma_{\text{joint},\varepsilon}$. A graphical (but nonrigorous) demonstration of this issue can be seen in [Figure 9](#) in [Appendix E.2](#).

Again, let $\mathcal{S} = \mathcal{A} = \mathbb{R}$, and $F_h(s, a) = a$. We let

$$W_\sigma(\cdot) = \mathcal{N}(\cdot, \sigma^2)$$

denote Gaussian smoothing. Fix some $\varepsilon > 0$. Define

$$P_{\text{init}} = \frac{1}{2}(\delta_{-\varepsilon} + \delta_{+\varepsilon}), \quad \pi^*(s) = \begin{cases} \delta_{-\varepsilon} & s \leq 0 \\ \delta_{\varepsilon} & s > 0 \end{cases}.$$

Thus, D_{π^*} is supported on the trajectories with $(s_{1:H+1}, a_{1:H})$ being either all ε or all $-\varepsilon$, and

$$P_h^* = P_{\text{init}} = \frac{1}{2}(\delta_{-\varepsilon} + \delta_{+\varepsilon}).$$

Hence, the replica and deconvolution map to distributions supported on $\{\varepsilon, -\varepsilon\}$. Let $\phi_\sigma(\cdot)$ denote the Gaussian PDF with variance σ . Then,

$$W_{\text{dec},h}^*(s) = \frac{\delta_\varepsilon \phi_\sigma(s - \varepsilon) + \delta_{-\varepsilon} \phi_\sigma(s + \varepsilon)}{\phi_\sigma(s - \varepsilon) + \phi_\sigma(s + \varepsilon)}.$$

Moreover,

$$W_{\text{rep},h}^*(s) = \mathbb{E}_{Z \sim \mathcal{N}(0, \sigma^2)} \left[\frac{\delta_\varepsilon \phi_\sigma(s - \varepsilon + Z) + \delta_{-\varepsilon} \phi_\sigma(s + \varepsilon + Z)}{\phi_\sigma(s - \varepsilon + Z) + \phi_\sigma(s + \varepsilon + Z)} \right]. \quad (\text{I.2})$$

One can check that for $\varepsilon \leq \sigma$,

$$W_{\text{rep},h}^*(u\varepsilon) = \Theta \left(\frac{(1 + \frac{c\varepsilon}{\sigma})\delta_{u\varepsilon} + (1 - \frac{c\varepsilon}{\sigma})\delta_{-u\varepsilon}}{2} \right), \quad u \in \{-1, 1\}$$

for $\varepsilon \ll 1$. In particular, for $s \in \{-\varepsilon, \varepsilon\}$

$$\mathbb{P}_{a \sim \pi_{\text{rep},h}^*}(s) [a = -s] \geq \Omega(1). \quad (\text{I.3})$$

In particular, if $(s_{1:H+1}^{\text{rep}}, a_{1:H}^{\text{rep}}) \sim D_{\pi_{\text{rep}}^*}$, then

$$\begin{aligned}\mathbb{P}[\exists h : d(s_h^{\text{rep}}, s_{h+1}^{\text{rep}}) > \varepsilon] &\leq \mathbb{P}[\exists h : s_h^{\text{rep}} = -s_{h+1}^{\text{rep}}] \\ &\leq \mathbb{P}[\exists h : s_h^{\text{rep}} = -a_h^{\text{rep}}] = 1 - \exp(-\Omega(H)),\end{aligned}$$

where in the last step we used (I.3) and the fact that the π_{rep}^* uses fresh randomness at each round. Moreover, as π^* always commits to either an all- ε or all- $(-\varepsilon)$ -trajectory, we see that for any $\mu \in \mathcal{C}(D_{\pi^*}, D_{\pi_{\text{rep}}^*})$ over $(s_{1:H+1}^*, a_{1:H}^*) \sim D_{\pi^*}$ and $(s_{1:H+1}^{\text{rep}}, a_{1:H}^{\text{rep}}) \sim D_{\pi_{\text{rep}}^*}$,

$$\Gamma_{\text{joint},\varepsilon}(\pi_{\text{rep}}^*, \pi^*) \geq \mathbb{P}_\mu[\exists 1 \leq h \leq H : d(s_{h+1}^*, s_{h+1}^{\text{rep}}) > \varepsilon] \geq 1 - \exp(-\Omega(H)),$$

That is, the replica and expert policies have different joint state distribution.

Remark I.1. The above result demonstrates the necessity of measuring the marginal error between $\hat{\pi} \circ W_\sigma$ and π^* in [Theorem 4](#): if we apply that proposition with $\hat{\pi} = \pi_{\text{dec}}^*$, then for all ε , $\mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os}, \varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*)) = 0$. But then $\hat{\pi} \circ W_\sigma = \pi_{\text{rep}}^*$, and we know that $\Gamma_{\text{joint}, \varepsilon}(\pi_{\text{rep}}^*, \pi^*) \geq \mathbb{P}_\mu[\exists 1 \leq h \leq H : d(s_{h+1}^*, s_{h+1}^{\text{rep}}) > \varepsilon] \geq 1 - \exp(-\Omega(H))$. Thus, we cannot hope for smoothed policies to imitate expert demonstrations in joint state distributions without additional assumptions.

Remark I.2 (Importance of chunking). Above we have shown that π_{rep}^* oscillates between ε and $-\varepsilon$ (for actions and subsequent states). We remark that these oscillations can have very deleterious effects on performance on real control systems. This is why it is beneficial to predict entire sequences of trajectories. Indeed, consider a modified construction such that $\mathcal{S} = \mathcal{A} = \mathbb{R}^K$, and $F_h(s, a) = a$. Here, we interpret \mathcal{S} as a sequence of K -control states in \mathbb{R} , and a as sequence of K -actions, denoting the i -th coordinate of s via $s[i]$,

$$\pi^*(s) = \begin{cases} \delta_{-\varepsilon \mathbf{1}} & s[1] \leq 0 \\ \delta_{\varepsilon \mathbf{1}} & s[1] > 0, \end{cases}$$

Then, we can view the oscillations in π_{rep}^* as oscillations between length K trajectories, which is essentially what happens in our analysis for $K = \tau_{\text{chunk}}$.

I.3 π_{rep}^* and π^* can have different marginals, implying necessity of $\gamma_{\text{IPS}, \mathcal{S}}$

Our construction lifts the construction in [Appendix I.2](#) to a two-dimensional state space $\mathcal{S} = \mathbb{R}^2$, keeping one dimensional actions $\mathcal{A} = \mathbb{R}$. Let $s = (s[1], s[2])$ denote coordinate of $s \in \mathcal{S}$. For some parameter ν , the dynamics are

$$s_{h+1} = F_h(s_h, a_h) = (a_h, \nu \cdot (s_h[1] - a_h))$$

We let $d_{\mathcal{S}} = d_{\text{TVC}} = d_{\text{IPS}}$ denote the ℓ_1 norm on $\mathcal{S} = \mathbb{R}^2$. Our initial state distribution is

$$P_{\text{init}} = \frac{1}{2} (\delta_{(\varepsilon, 0)} + \delta_{(-\varepsilon, 0)})$$

We let

$$\pi^*(s) = \begin{cases} \delta_{(-\varepsilon, 0)} & s \leq 0 \\ \delta_{(\varepsilon, 0)} & s > 0. \end{cases}$$

Thus, π^* induces trajectories which either stay on $\delta_{(\varepsilon, 0)}$ or $\delta_{(-\varepsilon, 0)}$.

$$P_h^* = \frac{1}{2} (\delta_{(\varepsilon, 0)} + \delta_{(-\varepsilon, 0)}), \quad \forall h \geq 1.$$

Let

$$W_\sigma(s) = \mathcal{N}(s', \sigma^2)$$

Proposition I.2. In the above construction, we can take $\gamma_{\text{IPS}, \mathcal{S}}(u) \leq \nu \cdot u$ in [Definition D.5](#), and p_r satisfies the conditions in [Theorem 4](#) for $r = 2\sigma\sqrt{\log(1/p_r)}$. Moreover, for any $\varepsilon \leq \sigma$,

$$\Gamma_{\text{marg}, \varepsilon'}(\pi_{\text{rep}}^* \parallel \pi^*) \geq \Omega(1), \quad \varepsilon' = \nu\varepsilon$$

Remark I.3 (Sharpness of $\gamma_{\text{IPS}, \mathcal{S}}$). Before proving this proposition, we note that if we take $\varepsilon = \sigma$ and $r = 2\sigma\sqrt{\log(1/p_r)}$, then $\nu\varepsilon = \tilde{\Omega}(\gamma_{\text{IPS}}(2r))$, showing that our dependence on $\gamma_{\text{IPS}, \mathcal{S}}$ is sharp up to logarithmic factors. Moreover, the looseness up to logarithmic factors in the above point is an artifact of using the Gaussian smoothing W_σ , and can be removed by replacing W_σ with a truncated-Gaussian kernel.

Proof of Proposition I.2. To see $\gamma_{\text{IPS}, \mathcal{S}}(u) \leq \nu \cdot u$, we have $\|F_h(s, a) - F_h(s', a)\| = \|(a, \nu \cdot (s[1] - a)) - (a, \nu \cdot (s'[1] - a))\| = \nu|s[1] - s'[1]| \leq \nu d_{\text{TVC}}(s, s')$. That we can take $r = 2\sigma\sqrt{\log(1/p_r)}$ follows from Gaussian concentration.

To prove the final claim, one can directly generalize (I.2) to find that, for any $b \in \mathbb{R}$,

$$W_{\text{rep},h}^*(s) = \mathbb{E}_{Z \sim \mathcal{N}(0, \sigma^2)} \left[\frac{\delta_{(\varepsilon,0)} \phi_\sigma(s[1] - \varepsilon + Z) + \delta_{(-\varepsilon,0)} \phi_\sigma(s[1] + \varepsilon + Z)}{\phi_\sigma(s[1] - \varepsilon + Z) + \phi_\sigma(s[1] + \varepsilon + Z)} \right].$$

This follows from the observation that $W_{\text{rep},h}^*$ and P_h^* have the same support, and as P_h^* always is support on vectors with second coordinate zero, that the second coordinate of s in $W_{\text{rep},h}^*(s)$ is uninformative. For $\varepsilon \leq \sigma$, we find that

$$W_{\text{rep},h}^*((\varepsilon, b)) = c\delta_{(\varepsilon,0)} + (1-c)\delta_{(-\varepsilon,0)}, c = \Omega(1), b \in \mathbb{R}.$$

and $W_{\text{rep},h}^*((-\varepsilon, b))$ is defined symmetrically, Hence, under $(s_{1:H+1}^{\text{rep}}, a_{1:H}^{\text{rep}}) \sim \pi_{\text{rep}}^*$,

$$\mathbb{P}[s_1^{\text{rep}} \neq a_1^{\text{rep}}] \geq \Omega(1)$$

Moreover, when $s_2^{\text{rep}} \neq a_{h^{\text{rep}}}^{\text{rep}}$, we have that $|s_2^{\text{rep}}[2]| = \nu |s_1^{\text{rep}} - a_1^{\text{rep}}|$, which as π^* is supported on $\{\delta_{(\varepsilon,0)}, \delta_{(-\varepsilon,0)}\}$, means, $|s_2^{\text{rep}}(2)| \geq 2\nu\varepsilon$. Thus,

$$\mathbb{P}[|s_2^{\text{rep}}[2]| \geq 2\nu\varepsilon] \geq \Omega(1)$$

On the other hand, $s_2^* \sim P_h^*$ has $s_2^*[2] = 0$ with probability one. Thus, for any coupling μ between $D_{\pi^*}, D_{\pi_{\text{rep}}^*}$,

$$\mathbb{P}_\mu[|d_S(s_2^{\text{rep}}, s_2^*)| \geq 2\nu\varepsilon] \geq \Omega(1)$$

Thus,

$$\Gamma_{\text{marg}, \nu\varepsilon}(\pi_{\text{rep}}^* \parallel \pi^*) \geq \Omega(1).$$

□

I.4 π_{rep}^* and π_{dec}^* have different marginals, even with memoryless dynamics

Here, we show how π_{rep}^* and π_{dec}^* have different marginals even if the dynamics are memoryless. By considering $\hat{\pi} = \pi_{\text{dec}}^*$ in Theorem 4, the discussion below demonstrates why one needs to consider $\hat{\pi}_\sigma = \hat{\pi} \circ W_\sigma$ in order to obtain small imitation gap.

For simplicity, we use a discrete smoothing kernel W_σ , though the example extends to the Gaussian smoothing kernel in the previous counter example. Again, let $\mathcal{S} = \mathcal{A} = \mathbb{R}$, and $F_h(s, a) = a$. Take

$$\pi^*(s) = \begin{cases} \delta_{-\sigma} & s \leq 0 \\ \delta_\sigma & s > 0 \end{cases}$$

Let us consider an asymmetric initial state distribution

$$P_{\text{init}} = \frac{1}{4}\delta_{-\sigma} + \frac{3}{4}\delta_{+\sigma}.$$

Note then that

$$\forall h, \quad P_h^* = P_{\text{init}} = \frac{1}{4}\delta_{-\sigma} + \frac{3}{4}\delta_\sigma, \quad (\text{I.4})$$

We consider a smoothing kernel,

$$W_\sigma(s) = \begin{cases} (\frac{1}{2} + \frac{s}{4\sigma})\delta_\sigma + (\frac{1}{2} - \frac{s}{4\sigma})\delta_{-\sigma} & -2\sigma \leq s \leq 2\sigma \\ \delta_\sigma & s \geq 2\sigma \\ \delta_{-\sigma} & s \leq -2\sigma \end{cases}$$

The salient part of our construction of W_σ is that

$$W_\sigma(\sigma) = \frac{1}{4}\delta_{-\sigma} + \frac{3}{4}\delta_\sigma, \quad W_\sigma(-\sigma) = \frac{1}{4}\delta_\sigma + \frac{3}{4}\delta_{-\sigma}.$$

Denote the marginals of π_{rep}^* and π_{dec}^* with $P_{\odot,h}^*$ and $P_{\text{dec},h}^*$. One can show via the lack of memory in the dynamics and the structure of π^* that

$$P_{\odot,h+1}^* = W_{\text{rep},h}^* \circ P_{\odot,h}^*, \quad W_{\text{dec},h+1}^* = W_{\text{dec},h}^* \circ P_{\text{dec},h}^*, \quad (\text{I.5})$$

By the replica property (Lemma H.3), $W_{\text{rep},h}^* \circ P_h^* = P_h^*$ for all h . Thus, for all h , (I.4) and (I.5) imply

$$P_{\odot,h}^* = P_h^* = \frac{1}{4}\delta_{-\sigma} + \frac{3}{4}\delta_{+\sigma}. \quad (\text{I.6})$$

The following claim computes $P_{\text{dec},h}^*$.

Claim I.3. Consider any distribution of the form $P = (1 - p)\delta_\sigma + p\delta_{-\sigma}$. Then

$$W_{\text{dec},h}^* \circ P = \left(\frac{9}{10} - \frac{p}{5}\right)\delta_\sigma + \left(\frac{1}{10} + \frac{p}{5}\right)\delta_{-\sigma}.$$

Thus,

$$P_{\text{dec},h+1}^*[-\sigma] = \frac{1}{10} \left(\sum_{i=0}^{h-1} 5^{-i} \right) + \frac{1}{4} 5^{1-h}.$$

Before proving the claim, let us remark on its implications. As $h \rightarrow \infty$,

$$P_{\text{dec},h}^*[-\sigma] \rightarrow \frac{1}{10} \left(\frac{1}{1 - 1/5} \right) = \frac{1}{10} \cdot \frac{5}{4} = \frac{1}{8}.$$

Thus,

$$\lim_{h \rightarrow \infty} P_{\text{dec},h}^* = \frac{7}{8}\delta_\sigma + \frac{1}{8}\delta_{-\sigma},$$

achieving a different stationary distribution that $P_h^* = P_{\odot,h}^*$. This shows that

$$\lim_{H \rightarrow \infty} \Gamma_{\text{marg},\sigma}(\pi_{\text{rep}}^*, \pi_{\text{dec}}^*) \geq \text{TV}\left(\frac{7}{8}\delta_\sigma + \frac{1}{8}\delta_{-\sigma}, \frac{3}{4}\delta_\sigma + \frac{1}{4}\delta_{-\sigma}\right) = \frac{1}{8},$$

which implies that the deconvolution policy π_{dec}^* does approximate π_{rep}^* . From (I.6), it also follows that π_{rep}^* and π^* have identical marginals, so

$$\lim_{H \rightarrow \infty} \Gamma_{\text{marg},\sigma}(\pi^*, \pi_{\text{dec}}^*) \geq \text{TV}\left(\frac{7}{8}\delta_\sigma + \frac{1}{8}\delta_{-\sigma}, \frac{3}{4}\delta_\sigma + \frac{1}{4}\delta_{-\sigma}\right) = \frac{1}{8}$$

as well. In particular, if we take $\hat{\pi} = \pi_{\text{dec}}^*$ in Theorem 4, we see that there is no hope to for bounding $\Gamma_{\text{marg},\varepsilon}(\pi^*, \hat{\pi})$; we must bound $\Gamma_{\text{marg},\varepsilon}(\pi^*, \hat{\pi} \circ W_\sigma)$ (again noting that if $\hat{\pi} = \pi_{\text{dec}}^*$, $\hat{\pi} \circ W_\sigma = \pi_{\text{rep}}^*$).

Proof of Claim I.3. We have that for $s' \in \{-\sigma, \sigma\}$,

$$W_{\text{dec},s'|s}^* = \frac{W_\sigma(s')[s] \cdot P_h^*(s')}{W_\sigma(s')[s] \cdot P_h^*(s') + W_\sigma(-s')[s] \cdot P_h^*(-s')}$$

With $s = s' = \sigma$, the above is

$$W_{\text{dec},h}^*(s' = \sigma \mid s = \sigma) = \frac{\frac{3}{4} \cdot \frac{3}{4}}{\frac{3}{4} \cdot \frac{3}{4} + \frac{1}{4} \cdot \frac{1}{4}} = \frac{9}{10}.$$

And

$$W_{\text{dec},h}^*(s' = \sigma \mid s = -\sigma) = \frac{\frac{1}{4} \cdot \frac{3}{4}}{\frac{1}{4} \cdot \frac{3}{4} + \frac{3}{4} \cdot \frac{1}{4}} = \frac{1}{2}.$$

Hence, for any $p \in [0, 1]$,

$$\begin{aligned} W_{\text{dec},h}^*(s' = \sigma \mid s = -\sigma)((1 - p)\delta_\sigma + p\delta_{-\sigma}) &= ((1 - p)\frac{9}{10} + \frac{p}{2})\delta_\sigma + (1 - ((1 - p)\frac{9}{10} + \frac{p}{2}))\delta_{-\sigma} \\ &= \left(\frac{9}{10} - \frac{p}{5}\right)\delta_\sigma + \left(\frac{1}{10} + \frac{p}{5}\right)\delta_{-\sigma}. \end{aligned}$$

Consequently, by (I.5), we can unfold a recursion to compute

$$\begin{aligned} P_{\text{dec},h+1}^*[-\sigma] &= W_{\text{dec},h}^*(s' = \sigma \mid s = -\sigma)P_{\text{dec},h}^* \\ &= \left(\frac{1}{10} + \frac{P_{\text{dec},h}^*[\sigma]}{5}\right) \\ &= \frac{1}{10} \sum_{i=0}^{h-1} 5^{-i} + P_{\text{dec},1}^*[\sigma] \cdot 5^{1-h} \\ &= \frac{1}{10} \sum_{i=0}^{h-1} 5^{-i} + P_1^*[\sigma] \cdot 5^{1-h} \\ &= \frac{1}{10} \left(\sum_{i=0}^{h-1} 5^{-i} \right) + \frac{1}{4} 5^{1-h}. \end{aligned}$$

□

Part III

The Control Setting

J End-to-end Guarantees and the Proof of Theorem 3

In this section, we provide a number of end-to-end guarantees for the learned imitation policy under various assumptions. The core of the section is the proof of [Theorem 2](#) which provides the basis for the final proof of [Theorem 3](#) in the body by uniting the analysis in the composite MDP from [Appendix H](#), the control theory from [Appendix K](#), and the sampling guarantees from [Appendix L](#). We now summarize the organisation of the appendix:

- In [Appendix J.1](#), we recall the association between the control setting and the composite MDP presented in [Appendix D](#), as well as rigorously instantiating the direct decomposition and the expert policy.
- In [Appendix J.2](#), we establish the correspondence between the imitation losses studied for the composite MDP in [Appendix D](#) with the desiderata in [Section 2](#).
- In [Appendix J.3](#), we provide the proof of [Proposition D.2](#) and [Theorems 2](#) and [3](#). We prove [Theorem 2](#) as a consequence of a more granular guarantee, [Theorem 10](#), which exposes the various tradeoffs in problem parameters.
- In [Appendix J.4](#), we demonstrate that if the demonstrator policy is assumed to be TVC, then we can recover stronger guarantees than those provided in [Theorem 3](#) without this assumption; in particular, we show that we can bound the *joint* imitation loss as well as the marginal and final versions.
- In [Appendix J.5](#), we show that if we were able to produce samples from a distribution close in *total variation* to the expert policy distribution, as opposed to the weaker optimal transport metric that we consider in the rest of the paper, then without any further assumptions, imitation learning is easily achievable.
- In [Appendix J.6](#), we demonstrate the utility of our imitation losses, showing that for Lipschitz cost functions decomposing in natural ways, our imitation losses as defined in [Definition 2.2](#) provide control over the difference in expected cost under expert and imitated distributions.
- Finally, in [Appendix J.7](#), we collect a number of useful lemmata that we use throughout the appendix.

J.1 Preliminaries

Here, we state various preliminaries to the end-to-end theorems. Recall that c_1, \dots, c_5 are constants which are polynomial in the parameters in [Assumption 3.1](#), and are spelled out explicitly in [Appendix K](#). For simplicity, to avoid complications with the boundary effects at $h = 1$, we re-define $h = 1$ -observation chunks \mathbf{o}_1 as elements $\mathcal{O} = \mathcal{P}_{\tau_{\text{obs}}-1}$ by prepending the necessary zeros – i.e. $\mathbf{o}_1 = (0, 0, \dots, 0, \mathbf{x}_1)$ – and similarly modifying $\mathbf{s}_1 \in \mathcal{S} = \mathcal{P}_{\tau_{\text{chunk}}}$ by prepending zeros. We first recall the definitions of the composite-states and -actions from [Appendix D](#). The prepending of zeros in the $h = 1$ case is mentioned above. For $h > 1$, recall that $\mathbf{s}_h = (\mathbf{x}_{t_{h-1}:t_h}, \mathbf{u}_{t_{h-1}:t_h-1})$ and that $\mathbf{a}_h = \kappa_{t_h:t_{h+1}-1}$, where we again emphasize that \mathbf{a}_h begins at the same t that \mathbf{s}_{h+1} does. We further recall the distances

$$d_{\text{traj}}(\boldsymbol{\rho}, \boldsymbol{\rho}') := \max_{1 \leq k \leq \tau+1} \|\mathbf{x}_k - \mathbf{x}'_k\| \vee \max_{1 \leq k \leq \tau} \|\mathbf{u}_k - \mathbf{u}'_k\|$$

defined for trajectories of arbitrary length. Given $\mathbf{s}_h, \mathbf{s}'_h$ with observation-(sub)chunks $\mathbf{o}_h, \mathbf{o}'_h$, we define

$$\begin{aligned} d_{\mathcal{S}}(\mathbf{s}_h, \mathbf{s}'_h) &= d_{\text{traj}}(\mathbf{s}_h, \mathbf{s}'_h) = \max_{t \in [t_{h-1}:t_h]} \|\mathbf{x}_t - \mathbf{x}'_t\| \vee \max_{t \in [t_{h-1}:t_h-1]} \|\mathbf{u}_t - \mathbf{u}'_t\|, \\ d_{\text{TVC}}(\mathbf{s}_h, \mathbf{s}'_h) &= d_{\text{traj}}(\mathbf{o}_h, \mathbf{o}'_h) = \max_{t \in [t_h-\tau_{\text{obs}}:t_h]} \|\mathbf{x}_t - \mathbf{x}'_t\| \vee \max_{t \in [t_h-\tau_{\text{obs}}:t_h-1]} \|\mathbf{u}_t - \mathbf{u}'_t\|, \\ d_{\text{IPS}}(\mathbf{s}_h, \mathbf{s}'_h) &= \|\mathbf{x}_{t_h} - \mathbf{x}'_{t_h}\|. \end{aligned}$$

Finally, for $\mathbf{a} = (\bar{\mathbf{u}}_{1:\tau_{\text{chunk}}}, \bar{\mathbf{x}}_{1:\tau_{\text{chunk}}}, \bar{\mathbf{K}}_{1:\tau_{\text{chunk}}})$ and $\mathbf{a}' = (\bar{\mathbf{u}}'_{1:\tau_{\text{chunk}}}, \bar{\mathbf{x}}'_{1:\tau_{\text{chunk}}}, \bar{\mathbf{K}}'_{1:\tau_{\text{chunk}}})$, recall from (3.1) and (D.2) that

$$\begin{aligned} d_{\max}(\mathbf{a}, \mathbf{a}') &= \max_{1 \leq k \leq \tau_{\text{chunk}}} \|\bar{\mathbf{u}}_k - \bar{\mathbf{u}}'_k\| + \|\bar{\mathbf{x}}_k - \bar{\mathbf{x}}'_k\| + \|\bar{\mathbf{K}}_k - \bar{\mathbf{K}}'_k\| \\ d_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') &:= c_1 d_{\max}(\mathbf{a}, \mathbf{a}') \cdot \mathbf{I}_{\infty}\{d_{\max}(\mathbf{a}, \mathbf{a}') > c_2\} \end{aligned}$$

We note the following fact.

Fact J.1. Suppose that $\varepsilon \leq c_2$. Then $d_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') \leq \varepsilon$ whenever $d_{\max}(\mathbf{a}, \mathbf{a}') \leq \varepsilon/c_1$.

J.1.1 Direct Decomposition and Smoothing Kernel.

This section will invoke the generalizations [Theorem 4](#) which requires TVC only subspace of the state space. This invokes the direct decomposition explained in [Appendix H](#).

Definition J.1 (Direct Decomposition and Smoothing Kernel). We consider the decomposition of $\mathcal{S} = \mathcal{O} \oplus \mathcal{S}_{/\mathcal{O}}$, where $\mathcal{O} = \mathcal{P}_{\tau_{\text{obs}}-1}$ are the coordinates of \mathbf{s}_h corresponding to the observation chunk \mathbf{o}_h , and $\mathcal{S}_{/\mathcal{O}}$ are all remaining coordinates. We let $\phi_{\mathcal{O}} : \mathcal{S} \rightarrow \mathcal{O}$ denote the projection onto the coordinates in \mathcal{O} . We instantiate the smoothing kernel W_{σ} as follows: For $\mathbf{s} = \mathbf{s}_h \in \mathcal{S} = \mathcal{P}_{\tau_{\text{chunk}}}$, we let

$$W_{\sigma}(\mathbf{s}) = \mathcal{N}\left(\mathbf{s}_h, \begin{bmatrix} \sigma^2 \mathbf{I}_{\mathcal{O}} & 0 \\ 0 & 0 \end{bmatrix}\right), \quad (\text{J.1})$$

where $\mathbf{I}_{\mathcal{O}}$ denotes the identity supported on the coordinates in \mathcal{O} as described above.

We note that the above direct decomposition satisfies the requisite compatibility assumptions explained in [Appendix H](#). Note also that d_{IPS} and W_{σ} are compatible with the above direct decomposition.

J.1.2 Chunking Policies.

We continue by centralizing a definition of chunking policies.

Definition J.2 (Policy and Initial-State Distributions). Given an *chunking policy* $\pi = (\pi_h)_{h=1}^H$ with $\pi_h : \mathcal{O} = \mathcal{P}_{\tau_{\text{obs}}-1} \rightarrow \Delta(\mathcal{A})$, we let \mathcal{D}_{π} denote the distribution over $\boldsymbol{\rho}_T$ and $\mathbf{a}_{1:H}$ induced by selecting $\mathbf{a}_h \sim \pi_h(\mathbf{o}_h)$, and rolling out the dynamics as described in [Section 2](#). We extend chunking policies to maps $\pi_h : \mathcal{S} = \mathcal{P}_{\tau_{\text{chunk}}} \rightarrow \Delta(\mathcal{A})$ by expressing $\pi_h = \pi_h \circ \phi_{\mathcal{O}}$ (i.e., projection \mathbf{s}_h onto its \mathbf{o}_h -components). Further, we let \mathbf{P}_{init} denote the distribution of \mathbf{x}_1 under $\boldsymbol{\rho}_T \sim \mathcal{D}_{\text{exp}}$.

Remark J.1. The notation \mathcal{D}_{π} denotes the special case of chunking policies in the control setting of [Section 2](#), whereas we reserve the serif font \mathcal{D}_{π} for the distribution induced by policies in the composite MDP. For composite MDPs instantiated as in [Appendix D.2](#), the two exactly coincide.

Construction of π^* for composite MDP. We now recall the policies π^* and $\pi_{\text{dec},\sigma}^*$, defined in [Definitions 3.4](#) and [3.6](#), respectively.

Definition J.3 (Policies corresponding to \mathcal{D}_{exp}). Define the following sequence kernels $\pi^* = (\pi_h^*)_{h=1}^H$ and $\pi_{\text{dec}}^* = (\pi_{\text{dec},h}^*)_{h=1}^H$ via the following process. Let $\boldsymbol{\rho}_T \sim \mathcal{D}_{\text{exp}}$, and let $\mathbf{a}_{1:H} = \text{synth}(\boldsymbol{\rho}_T)$; further, let $\mathbf{o}_{1:H}$ be the corresponding observation-chunks from $\boldsymbol{\rho}_T$. Let

- $\pi_h^*(\cdot) : \mathcal{O} = \mathcal{P}_{\tau_{\text{obs}}-1} \rightarrow \Delta(\mathcal{A})$ denote a regular conditional probability corresponding to the distribution over \mathbf{a}_h given \mathbf{o}_h in the above construction. In other words, the distribution of $\mathbf{a}_h \mid \mathbf{o}_h$ under $\mathcal{D}_{\sigma=0,h}$.
- Let $\pi_{\text{dec},\sigma,h}^*(\cdot) : \mathcal{O} = \mathcal{P}_{\tau_{\text{obs}}-1} \rightarrow \Delta(\mathcal{A})$ denote a regular conditional probability corresponding to the distribution over \mathbf{a}_h given an augmented $\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$. In other words, the distribution of $\mathbf{a}_h \mid \tilde{\mathbf{o}}_h$ under $\mathcal{D}_{\sigma,h}$.

When instantiating the composite MDP, π^* corresponds to its namesake, and $\pi_{\text{dec},\sigma,h}^*$ to $\pi_{\text{dec},h}^*$. Moreover, π^* as constructed above, \mathbf{P}_h^* denotes the distribution over \mathbf{s}_h under \mathcal{D}_{π^*} . By [Lemma J.6](#), this is in fact equal to the distribution over \mathbf{s}_h under \mathcal{D}_{exp} . Notice further, therefore, that $\phi_{\mathcal{O}} \circ \mathbf{P}_h^*$ is precisely the distribution of \mathbf{o}_h under \mathcal{D}_{exp} .

Remark J.2. We remark that by [Theorem 7](#), π_h^* is unique up to a measure zero set of \mathbf{o}_h as distributed as above, and $\pi_{\text{dec},h}^*$ is unique almost surely for $\tilde{\mathbf{o}}_h$ distributed as above. In particular, since the latter has density with respect to the Lebesgue measure and infinite support, $\pi_{\text{dec},h}^*$ is unique in a Lebesgue almost everywhere sense.

J.1.3 Preliminaries for joint-distribution imitation.

This section introduces a further *joint imitation gap*, which we can make small under a stronger bounded-memory assumption on \mathcal{D}_{exp} stated below.

Definition J.4 (Joint and Final Imitation Gap). Given a chunking policy π , we let

$$\mathcal{L}_{\text{joint},\varepsilon}(\pi) := \inf_{\mu} \mathbb{P}_{\mu} \left[\max_{t \in [T]} \max \{ \|\mathbf{x}_{t+1}^{\text{exp}} - \mathbf{x}_{t+1}^{\pi}\|, \|\mathbf{u}_t^{\text{exp}} - \mathbf{u}_t^{\pi}\| \} > \varepsilon \right],$$

where the infimum is over all couplings between the distribution of ρ_T under \mathcal{D}_{exp} and that induced by the policy π . We also define

$$\mathcal{L}_{\text{fin},\varepsilon}(\pi) := \inf_{\mu} \mathbb{P}_{\mu} [\|\mathbf{x}_{T+1}^{\text{exp}} - \mathbf{x}_{T+1}^{\pi}\| > \varepsilon],$$

the loss restricted to the final states under each distribution.

Controlling $\mathcal{L}_{\text{joint},\varepsilon}(\pi)$ requires various additional stronger assumptions (*which we do not require in [Theorem 3](#)*), one of which is that the demonstrator has bounded memory:

Definition J.5. We say that the demonstration distribution, synthesis oracle pair $(\mathcal{D}_{\text{exp}}, \text{synth})$ have τ -bounded memory if under $\rho_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \sim \mathcal{D}_{\text{exp}}$ and $\mathbf{a}_{1:H} = \text{synth}(\rho_T)$, the conditional distribution of \mathbf{a}_h and $\mathbf{x}_{1:t_h-\tau}, \mathbf{u}_{1:t_h-\tau}$ are conditionally independence given $(\mathbf{x}_{t_h-\tau+1:t_h}, \mathbf{u}_{t_h-\tau+1:t_h-1})$.

We note that enforcing [Definition J.5](#) can be relaxed to a mixing time assumption (see [Remark J.5](#)). Moreover, we stress that we *do not* need the condition in [Definition J.5](#) if we only seek imitation of marginal distributions (as captured by $\mathcal{L}_{\text{marg},\varepsilon}$ and $\mathcal{L}_{\text{fin},\varepsilon}$), as in [Theorem 3](#).

J.2 Translating Control Imitation Losses to Composite-MDP Imitation Gaps

Lemma J.1. Recall the imitation losses [Definitions 2.2](#) and [J.4](#), and the composite-MDP imitation gaps [Definition D.1](#). Further consider, the substitutions defined in [Appendix D.2](#), with π^* instantiated as in [Definition J.3](#). Given policies $\pi = (\pi_h)$ with $\pi_h : \mathcal{O} = \mathcal{P}_{\tau_{\text{obs}}-1} \rightarrow \mathcal{A}$, we can extend $\pi_h : \mathcal{S} = \mathcal{P}_{\tau_{\text{chunk}}} \rightarrow \mathcal{A}$ by the natural embedding of $\mathcal{P}_{\tau_{\text{obs}}-1}$ into $\mathcal{P}_{\tau_{\text{chunk}}}$. Then, for any $\varepsilon > 0$,

$$\mathcal{L}_{\text{marg},\varepsilon}(\pi) \leq \Gamma_{\text{marg},\varepsilon}(\pi \parallel \pi^*).$$

If we instead consider the the substitutions defined in [Appendix D.2](#), but set $d_{\mathcal{S}}$ to equal d_{IPS} , which only measures distance in the final coordinate of each trajectory chunk \mathbf{s}_h ,

$$\mathcal{L}_{\text{fin},\varepsilon}(\pi) \leq \Gamma_{\text{marg},\varepsilon}(\pi \parallel \pi^*), \quad d_{\mathcal{S}}(\cdot, \cdot) \leftarrow d_{\text{IPS}}(\cdot, \cdot) \quad (\text{J.2})$$

Finally, if \mathcal{D}_{exp} has $\tau \leq \tau_{\text{obs}}$ -bounded memory,

$$\mathcal{L}_{\text{joint},\varepsilon}(\pi) \leq \Gamma_{\text{joint},\varepsilon}(\pi \parallel \pi^*).$$

Proof. Let's start with the first bound, let superscript exp denote objects from \mathcal{D}_{exp} and superscript π from \mathcal{D}_{π} , the distribution induced by chunking policy π . Letting \inf_{μ} denote infima over couplings between the two, we have

$$\begin{aligned} \mathcal{L}_{\text{marg},\varepsilon}(\pi) &:= \max_{t \in [T]} \inf_{\mu} \{ \mathbb{P}_{\mu} [\|\mathbf{x}_{t+1}^{\text{exp}} - \mathbf{x}_{t+1}^{\pi}\| > \varepsilon], \mathbb{P}_{\mu} [\|\mathbf{u}_t^{\text{exp}} - \mathbf{u}_t^{\pi}\| > \varepsilon] \} \\ &:= \max_{t \in [T]} \inf_{\mu} \{ \mathbb{P}_{\mu} [\|\mathbf{x}_{t+1}^{\text{exp}} - \mathbf{x}_{t+1}^{\pi}\| \vee \|\mathbf{u}_t^{\text{exp}} - \mathbf{u}_t^{\pi}\| > \varepsilon] \} \\ &\stackrel{(a)}{\leq} \max_{h \in [H]} \inf_{\mu} \left\{ \mathbb{P}_{\mu} \left[\max_{0 \leq i \leq \tau_{\text{chunk}}} \|\mathbf{x}_{t_{h+1}-i}^{\text{exp}} - \mathbf{x}_{t_{h+1}-i}^{\pi}\| \vee \max_{1 \leq i \leq \tau_{\text{chunk}}} \|\mathbf{u}_{t_{h+1}-i}^{\text{exp}} - \mathbf{u}_{t_{h+1}-i}^{\pi}\| \right] \right\} \\ &\leq \max_{h \in [H]} \inf_{\mu} \{ \mathbb{P}_{\mu} [d_{\mathcal{S}}(\mathbf{s}_{h+1}^{\text{exp}}, \mathbf{s}_{h+1}^{\pi})] \}, \end{aligned}$$

where step (a) uses that for any $t \in [H]$, we can find some h such that $t + 1 \in t_{h+1} - \{0, 1, \dots, \tau_{\text{chunk}}\}$ and $t \in t_{h+1} - \{1, 2, \dots, \tau_{\text{chunk}}\}$.¹⁰

From Lemma J.6, s_h^{exp} has the same marginal distribution as $s_h^{\pi^*}$, the distribution induced by π^* in Definition J.3.¹¹ Still, letting $\inf_{\mu'}$ denote infimum over couplings between \mathcal{D}_π and \mathcal{D}_{exp} , equality of these marginals suffices to ensure

$$\mathcal{L}_{\text{marg},\varepsilon}(\pi) \leq \max_{h \in [H]} \inf_{\mu} \left\{ \mathbb{P}_{\mu'} \left[d_S(s_{h+1}^{\pi^*}, s_{h+1}^\pi) \right] \right\},$$

As s_h corresponds to a composite state s_h in the composite MDP, the above is at most $\Gamma_{\text{marg},\varepsilon}(\pi \parallel \pi^*)$ as in definition Definition D.1. For the final-state imitation loss,

$$\begin{aligned} \mathcal{L}_{\text{fin},\varepsilon}(\pi) &:= \inf_{\mu} \mathbb{P}_{\mu} \left[\|\mathbf{x}_{T+1}^{\text{exp}} - \mathbf{x}_{T+1}^\pi\| > \varepsilon \right] \\ &\leq \max_{h \in [H]} \inf_{\mu} \left\{ \mathbb{P}_{\mu} \left[d_{\text{IPS}}(s_h^{\text{exp}}, s_h^\pi) \right] \right\}, \end{aligned}$$

where again d_{IPS} only measures error in the final state of s_h . The corresponding bound in (J.2) follows similarly.

Finally, we have

$$\mathcal{L}_{\text{joint},\varepsilon}(\pi) := \inf_{\mu} \mathbb{P}_{\mu} \left[\max_{t \in [T]} \max \left\{ \|\mathbf{x}_{t+1}^{\text{exp}} - \mathbf{x}_{t+1}^\pi\|, \|\mathbf{u}_t^{\text{exp}} - \mathbf{u}_t^\pi\| \right\} > \varepsilon \right],$$

When \mathcal{D}_{exp} has $\tau \leq \tau_{\text{obs}}$ -bounded memory, then, the expert and π^* -induced trajectories are identically distributed. Therefore, directly from this observation and Definition D.1,

$$\mathcal{L}_{\text{joint},\varepsilon}(\pi) = \inf_{\mu} \mathbb{P}_{\mu} \left[\max_{t \in [T]} \max \left\{ \|\mathbf{x}_{t+1}^{\pi^*} - \mathbf{x}_{t+1}^\pi\|, \|\mathbf{u}_t^{\pi^*} - \mathbf{u}_t^\pi\| \right\} > \varepsilon \right] \leq \Gamma_{\text{joint},\varepsilon}(\pi \parallel \pi^*).$$

□

J.3 Proofs of main results

J.3.1 Proof of Theorem 1

The result is a direct consequence of the following points. First, with our instantiation of the composite MDP, we can bound $\mathcal{L}_{\text{marg},\varepsilon}(\hat{\pi}) \leq \Gamma_{\text{marg},\varepsilon}(\hat{\pi} \parallel \pi^*) \leq \Gamma_{\text{joint},\varepsilon}(\hat{\pi} \parallel \pi^*)$ due to Lemma J.1; and moreover, we have $\mathcal{L}_{\text{joint},\varepsilon}(\hat{\pi}) \leq \Gamma_{\text{joint},\varepsilon}(\hat{\pi} \parallel \pi^*)$ when \mathcal{D}_{exp} has $\tau \leq \tau_{\text{obs}}$ -bounded memory. The bound now follows from Proposition D.2, and the fact that Proposition D.1 verifies the input-stability property, and Fact J.1. □

J.3.2 Proof of Theorem 2 and a more precise statement.

In this section, we derive Theorem 2 from a more precise guarantee that exposes the various algorithmic knobs in a more explicit manner.

Theorem 10. Let Assumption 3.1 hold, and let $c_1, \dots, c_5 > 0$ be as in Definition 3.2. Suppose that the $\varepsilon, \sigma, \tau_{\text{chunk}} > 0$ satisfy $\varepsilon < c_2$, and $\tau_{\text{chunk}} \geq c_3$, and $5d_x + \log\left(\frac{4\sigma}{c_1\varepsilon}\right) \leq c_4^2/(16\sigma^2)$. Then the marginal imitation loss (Definition 2.2) and final-state imitation loss (Definition J.4) of the *smoothed* $\hat{\pi}_\sigma$ are bounded by

$$\begin{aligned} \max \left\{ \mathcal{L}_{\text{marg},\varepsilon_1}(\hat{\pi}_\sigma), \mathcal{L}_{\text{fin},\varepsilon_2}(\hat{\pi}_\sigma) \right\} &\leq H\sqrt{2\tau_{\text{obs}} - 1} \left(\frac{2\varepsilon}{\sigma} + \iota_\sigma(\varepsilon) e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})} \right) \\ &\quad + \sum_{h=1}^H \mathbb{E}_{\tilde{o}_h \sim \mathcal{D}_{\text{exp},\sigma,h}} \Delta_{(\varepsilon/c_1)}(\pi_{\text{dec},\sigma,h}^*(\tilde{o}_h), \hat{\pi}_h(\tilde{o}_h)). \end{aligned} \tag{J.3}$$

and where $\iota_\sigma(\varepsilon) = 6c_5 \sqrt{5d_x + 2 \log\left(\frac{4\sigma}{c_1\varepsilon}\right)}$ is logarithmic in $1/\varepsilon$, and where

$$\varepsilon_1 = \varepsilon + \sigma \iota_\sigma(\varepsilon), \quad \varepsilon_2 = \varepsilon + \sigma e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})} \iota_\sigma(\varepsilon).$$

¹⁰Recall $t_h := (h-1)\tau_{\text{chunk}} + 1$.

¹¹Note the subtlety that the joint distribution of these may defer because π^* has limited memory.

Remark J.3 (Parameter Dependencies). Each term in (J.3) can be made small by decreasing the amount of noise σ in the smoothing, increasing the number of trajectories, and increasing the chunk length τ_{chunk} ; indeed, these are the levers by means of which we derive Theorem 2 just below. Increasing τ_{chunk} comes at the (implicit) expense of increasing the length of composite actions, thereby inducing a more challenging conditional generative modeling problem. Decreasing σ increases sensitivity to the tolerance ε , and, as discussed in Appendix L, may make the underlying generative modeling problem more challenging. Note that the contribution of the additive σ -term in ε_2 , used for the final-state loss $\mathcal{L}_{\text{fin},\varepsilon}$, is exponentially-in- τ_{chunk} smaller than that in ε_1 . Interestingly, our theory suggest no benefit to increasing τ_{obs} (corroborated empirically in [19]).

We now turn the proof of proof of Theorem 2.

Deduction. Theorem 2 from Theorem 10. Fix a desired ε_0 for which

$$\varepsilon_0 < \min\{1/2c_1, 1/\sqrt{c_1/c_2}, 3c_4c_5\}.$$

Then, taking $\varepsilon = c_1\varepsilon_0^2$, we have $\varepsilon < c_2$ and $\varepsilon \leq \varepsilon_0/2$. Select $\sigma = \frac{1}{2}\varepsilon_0/\iota_\star(\varepsilon_0)$. For such an σ , we have that as $\iota_\star(\varepsilon_0) \geq 6c_5$, $\sigma \leq \frac{1}{12c_5}\varepsilon_0$, and thus

$$\iota_\sigma(\varepsilon) = 6c_5\sqrt{5d_x + 2\log\left(\frac{4\sigma}{c_1\varepsilon}\right)} = 6c_5\sqrt{5d_x + 2\log\left(\frac{8\sigma}{c_1^2\varepsilon_0^2}\right)} \leq 6c_5\sqrt{5d_x + 2\log\left(\frac{1}{c_5c_1^2\varepsilon_0}\right)} =: \iota_\star(\varepsilon_0)$$

and therefore, with our vari $\varepsilon_1 := \varepsilon + \sigma\iota_\sigma(\varepsilon) \leq \varepsilon_0/2 + \varepsilon_0/2 \leq \varepsilon_0$. Next, we verify the condition

$$c_4^2/(16\sigma^2) \geq 5d_x + \log\left(\frac{4c_1\sigma}{\varepsilon}\right) = (\iota_\sigma(\varepsilon)^2)/(6c_5)^2$$

Rearranging, we need $\frac{36c_5^2c_4^2}{16} \geq (\sigma\iota_\sigma(\varepsilon))^2$, and as $(\sigma\iota_\sigma(\varepsilon))^2 \leq \varepsilon_0/2$, it then suffices that $\varepsilon_0 \leq 3c_4c_5$, which holds. Therefore, Theorem 2 implies

$$\mathcal{L}_{\text{marg},\varepsilon_0}(\hat{\pi}_\sigma) \leq H\iota_\star(\varepsilon_0)\sqrt{2\tau_{\text{obs}} - 1} \left(2c_1\varepsilon_0 + e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})}\right) + \sum_{h=1}^H \mathbb{E}_{\tilde{o}_h \sim \mathcal{D}_{\text{exp},\sigma,h}} \Delta_{(\varepsilon/c_1)}(\pi_{\text{dec},\sigma,h}^\star(\tilde{o}_h), \hat{\pi}_h(\tilde{o}_h)).$$

The first result follows by relabeling $\varepsilon \leftarrow \varepsilon_0$ and taking $\tau_{\text{chunk}} - \tau_{\text{obs}} \geq \frac{1}{L_\beta} \log(c_1/\varepsilon)$. The second result is a consequence of Markov's inequality, and the well behaved-ness of conditional distributions established throughout Appendix F. \square

J.3.3 Proof of Theorem 10

Proof of Theorem 10. Lets begin by bounding $\mathcal{L}_{\text{marg},\varepsilon}(\pi)$. Recall the definitions of $d_S, d_{\text{TVC}}, d_{\text{IPS}}$ in Appendix D, and let $\mathbf{s}_{1:H+1}^\star$ and $\mathbf{s}_{1:H+1}$ denote the composite states corresponding to a trajectory $(\mathbf{x}_{1:T+1}^\star, \mathbf{u}_{1:T}^\star)$ under π^\star and $(\mathbf{x}_{1:T+1}^\pi, \mathbf{u}_{1:T}^\pi)$, respectively, under the instantiation of the composite MDP in Appendix D.2. We can view π^\star and π (which depend only on observation chunks \mathbf{o}_h) as policies in the composite MDP which are compatible with the decomposition Definition H.1. We make the following points:

- In light of Lemma J.1,

$$\mathcal{L}_{\text{marg},\varepsilon_1}(\pi \parallel \pi^\star) \leq \Gamma_{\text{marg},\varepsilon_1}(\pi \parallel \pi^\star).$$

- By Lemma J.8, a consequence of Pinsker's inequality, it holds that the Gaussian kernel W_σ used in HINT is γ_σ -TVC (w.r.t. d_{TVC}) with

$$\gamma_\sigma(u) = \frac{u\sqrt{2\tau_{\text{obs}} - 1}}{2\sigma} \quad (\text{J.4})$$

- Note that $d_{\text{IPS}}(\mathbf{s}_h, \mathbf{s}'_h) = \|\mathbf{x}_{t_h} - \mathbf{x}'_{t_h}\|$ measures Euclidean distance between the last \mathbf{x} -coordinates of $\mathbf{s}_h, \mathbf{s}'_h$. Moreover, if $\mathbf{s}'_h \sim W_\sigma(\mathbf{s}_h)$ the last coordinate \mathbf{x}'_{t_h} of \mathbf{s}' is distributed as $\mathcal{N}(\mathbf{x}_{t_h}, \sigma^2 I)$. By Lemma J.7 with $d = d_x$, that for $r = 2\sigma \cdot \sqrt{5d_x + 2\log\left(\frac{1}{p}\right)}$

$$p_r = \mathbb{P}_{\mathbf{s}' \sim W_\sigma(\mathbf{s})}[d_{\text{IPS}}(\mathbf{s}, \mathbf{s}') > r] \leq p. \quad (\text{J.5})$$

- As (a) \mathbf{s}_h^* corresponds to \mathbf{s}_h from $\mathbf{p}_T \sim \mathcal{D}_{\text{exp}}$, (b) as $\hat{\pi}, \pi_{\text{dec}}^*$ are functions of \mathbf{o}_h , and (c) by recalling the definition of $\mathbf{d}_{\text{os}, \varepsilon}$ in [Definition D.1](#), $\varepsilon \leq c_2$ ensures

$$\begin{aligned}
& \mathbb{E}_{\mathbf{s}_h^* \sim \mathbf{P}_h^*} \mathbb{E}_{\tilde{\mathbf{s}}_h^* \sim \mathcal{W}_\sigma(\mathbf{s}_h^*)} \mathbf{d}_{\text{os}, \varepsilon}(\pi_h(\tilde{\mathbf{s}}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{\mathbf{s}}_h^*)) \\
&= \mathbb{E}_{\mathbf{o}_h \sim \mathcal{D}_{\text{exp}, h}} \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})} \inf_{\mu \in \mathcal{C}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h))} \mathbb{P}_{(\mathbf{a}, \mathbf{a}') \sim \mu}[\mathbf{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') \geq \varepsilon] \\
&\leq \mathbb{E}_{\mathbf{o}_h \sim \mathcal{D}_{\text{exp}, h}} \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})} \inf_{\mu \in \mathcal{C}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h))} \mathbb{P}_{(\mathbf{a}, \mathbf{a}') \sim \mu}[\mathbf{d}_{\text{max}}(\mathbf{a}, \mathbf{a}') \geq \varepsilon / c_1] \\
&\quad \text{(Fact J.1)} \\
&= \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp}, \sigma, h}} \Delta_{(\varepsilon / c_1)}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h)).
\end{aligned}$$

- Finally, [Proposition K.1](#) (formalizing [Proposition D.1](#)) ensures that under our assumption $\tau_{\text{chunk}} \geq c_3 /$, and let $r_{\text{IPS}} = c_4$, $\gamma_{\text{IPS}, \text{TVC}}(u) = c_5 u \exp(-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}}))$, $\gamma_{\text{IPS}, \mathcal{S}}(u) = c_5 u$ for c_3, c_4, c_5 given in [Appendix K](#). Then, for $\mathbf{d}_{\mathcal{S}}, \mathbf{d}_{\text{TVC}}, \mathbf{d}_{\text{IPS}}$ as above, we have that π^* is $(\gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, \mathcal{S}}, \mathbf{d}_{\text{IPS}}, r_{\text{IPS}})$ -IPS.

Consequently, for $r = 2\sigma \cdot \sqrt{5d_x + 2 \log \left(\frac{4\sigma}{c_1 \varepsilon} \right)} \in (0, \frac{1}{2} r_{\text{IPS}})$, [Theorem 9](#) (which, we recall, generalizes [Theorem 4](#) to account for the direct decomposition structure) implies

$$\begin{aligned}
\mathcal{L}_{\text{marg}, \varepsilon + 2rc_5}(\hat{\pi}_\sigma) &= \mathcal{L}_{\text{marg}, \varepsilon + 2rc_5}(\hat{\pi}_\sigma \parallel \pi^*) \leq \Gamma_{\text{marg}, \varepsilon + 2rc_5}(\hat{\pi}_\sigma \parallel \pi^*) \\
&\leq H \left(\frac{\varepsilon}{2\sigma} + \frac{3}{2\sigma} \sqrt{2\tau_{\text{obs}} - 1} \left(\max \left\{ \varepsilon, 2rc_5 e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})} \right\} \right) \right) \\
&\quad + \sum_{h=1}^H \mathbb{E}_{\mathbf{s}_h^* \sim \mathbf{P}_h^*} \mathbb{E}_{\tilde{\mathbf{s}}_h^* \sim \mathcal{W}_\sigma(\mathbf{s}_h^*)} \mathbf{d}_{\text{os}, \varepsilon}(\pi_h(\tilde{\mathbf{s}}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{\mathbf{s}}_h^*)) \\
&\leq H \sqrt{2\tau_{\text{obs}} - 1} \left(\frac{2\varepsilon}{\sigma} + 6\sigma c_5 \sqrt{5d_x + 2 \log \left(\frac{4\sigma}{c_1 \varepsilon} \right)} e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})} \right) \\
&\quad + \sum_{h=1}^H \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp}, \sigma, h}} \Delta_{(\varepsilon / c_1)}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h)). \\
&\leq H \sqrt{2\tau_{\text{obs}} - 1} \left(\frac{2\varepsilon}{\sigma} + \sigma \iota(\varepsilon) \right) + \sum_{h=1}^H \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp}, \sigma, h}} \Delta_{(\varepsilon / c_1)}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h)).
\end{aligned}$$

Substituting in $\varepsilon_1 = \varepsilon + 2rc_5 = c_1 \varepsilon + 4c_5 \sigma \cdot \sqrt{5d + 2 \log \left(\frac{4\sigma}{c_1 \varepsilon} \right)} \leq \varepsilon + \sigma \iota(\varepsilon)$, the bound on $\mathcal{L}_{\text{marg}, \varepsilon_1}$ is proved.

To show $\mathcal{L}_{\text{fin}, \varepsilon_2}(\hat{\pi}_\sigma)$ satisfies the same bound, we replace $\mathbf{d}_{\mathcal{S}}$ in the above argument (as defined in [Appendix D.2](#)) with $\mathbf{d}_{\mathcal{S}}(\cdot, \cdot) \leftarrow \mathbf{d}_{\text{IPS}}(\cdot, \cdot)$, where again we recall that $\mathbf{d}_{\text{IPS}}(\mathbf{s}_s, \mathbf{s}'_s) = \|\mathbf{x}_{t_h} - \mathbf{x}'_{t_h}\|$ measures differences in the final associated control state. From [Proposition K.1](#), it follows we can take $\gamma_{\text{IPS}, \mathcal{S}}(u) = c_5 u e^{-L_\beta \tau_{\text{chunk}}}$. Thus, we can replace ε_1 above with $\varepsilon_2 := \varepsilon + 4c_5 e^{-L_\beta \tau_{\text{chunk}}} \sigma \cdot (5d_x + 2 \log \left(\frac{1}{\varepsilon} \right))^{1/2}$. This concludes the proof that

$$\mathcal{L}_{\text{marg}, \varepsilon_2}(\hat{\pi}_\sigma) \leq H \sqrt{2\tau_{\text{obs}} - 1} \left(6c_5 \sqrt{5d_x + 2 \log \left(\frac{4\sigma}{c_1 \varepsilon} \right)} e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})} + \frac{2\varepsilon}{\sigma} \right),$$

which can be simplified as needed. \square

J.3.4 Proof of [Theorem 3](#)

Adopt the shorthand $\Delta_h = \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp}, \sigma, h}} \Delta_{(\varepsilon / c_1)}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h))$. From [Theorem 2](#), it suffices to show that with probability at least $1 - \delta$, it holds that $\Delta_h \leq \varepsilon^2$ for all $h \in [H]$. For $d_{\mathcal{A}} = \tau_{\text{chunk}}(d_x + d_u + d_x d_u)$, we have that $\mathbf{a} \in \mathbb{R}^{d_{\mathcal{A}}}$. Note that by [Assumption 3.1](#) it holds \mathcal{D}_{exp} -almost surely that we can crudely bound $\|\mathbf{a}_h\| \leq \sqrt{d_{\mathcal{A}}} \max\{R_{\mathbf{K}}, R_{\text{dyn}}\}$ and thus the condition on

q in [Theorem 13](#) holds for

$$R = \sqrt{d_{\mathcal{A}}} \max\{R_{\mathbf{K}}, R_{\text{dyn}}\}.$$

By [Assumption C.1](#), the conditions on the score class \mathbf{s}_{θ} hold for us to apply [Theorem 13](#). Note that by assumption,

$$N_{\text{exp}} \geq c \left(\frac{C_{\Theta} d R (R \vee \sqrt{d_{\mathcal{A}}}) \log(dn)}{\varepsilon^8} \right)^{4\nu} \vee \left(\frac{d_{\mathcal{A}}^6 (R^4 \vee d_{\mathcal{A}}^2 \log^3(\frac{HndR\sigma}{\delta\varepsilon}))}{\varepsilon^{48}} d_{\mathcal{A}}^2 \right)^{4\nu}, \quad (\text{J.6})$$

where we note that the right hand side is poly $(C_{\Theta}, 1/\varepsilon, R_{\mathbf{K}}, d_{\mathcal{A}}, \log(H/\delta))^{\nu}$, and J and α are set as in [\(L.1\)](#). Taking a union bound over $h \in [H]$ and applying [Theorem 13](#) tells us that with probability at least $1 - \delta$, for all $h \in [H]$, it holds that

$$\mathbb{E}_{\tilde{\mathbf{o}}_h \sim q_{\tilde{\mathbf{o}}_h}} \left[\inf_{\mu \in \mathcal{C}(\mathcal{DDPM}(\mathbf{s}_{\tilde{\theta}}, \tilde{\mathbf{o}}_h), q(\cdot|\tilde{\mathbf{o}}_h))} \mathbb{P}_{(\hat{\mathbf{a}}, \mathbf{a}^*) \sim \mu} (\|\hat{\mathbf{a}} - \mathbf{a}^*\| \geq \varepsilon) \right] \leq \varepsilon.$$

Thus, it holds that with probability at least $1 - \delta$,

$$\sum_{h=1}^H \Delta_h \leq H\varepsilon.$$

Plugging this in to [Theorem 2](#) concludes the proof of the first statement. The proof of the second statement is analogous. \square

J.4 Imitation of the joint trajectory under total variation continuity of demonstrator policy

Here, we show that if the demonstrator policy has (a) bounded memory and (b) satisfies a certain continuity property in total variation distance, then we can imitate the *joint distribution* over trajectories, not just marginals. Recall the joint imitation loss from $\mathcal{L}_{\text{joint}, \varepsilon}$ from [Definition J.4](#).

Theorem 11. Consider the setting [Theorem 10](#), and define as shorthand

$$\Delta_{h, \varepsilon} := \mathbb{E}_{\tilde{\mathbf{o}}_h \sim \mathcal{D}_{\text{exp}, \sigma, h}} \Delta_{(\varepsilon/c_1)}(\pi_{\text{dec}, \sigma, h}^*(\tilde{\mathbf{o}}_h), \hat{\pi}_h(\tilde{\mathbf{o}}_h)).$$

Suppose that, in addition, there is a strictly increasing function $\gamma(\cdot)$ such that for all $\mathbf{o}_h, \mathbf{o}'_h \in \mathcal{O}$,

$$\text{TV}(\pi^*(\mathbf{o}_h), \pi^*(\mathbf{o}'_h)) \leq \gamma(\|\mathbf{o}_h - \mathbf{o}'_h\|),$$

where π^* is defined is the conditional in [Definition J.3](#). Further, suppose that \mathcal{D}_{exp} has $\tau \leq \tau_{\text{obs}}$ bounded memory ([Definition J.5](#)). Then, with $\varepsilon_1 := \varepsilon + \sigma\iota(\varepsilon)$ as in [Theorem 2](#),

$$\begin{aligned} \mathcal{L}_{\text{joint}, \varepsilon_1}(\hat{\pi}_{\sigma}) &\leq H \cdot \text{ERRTV}(\sigma, \gamma) \\ &\quad + H \sqrt{2\tau_{\text{obs}} - 1} \left(\frac{2\varepsilon}{\sigma} + 6c_5 \sqrt{5d_x + 2 \log\left(\frac{4\sigma}{\varepsilon}\right)} e^{-L_{\beta}(\tau_{\text{chunk}} - \tau_{\text{obs}})} \right) + \sum_{h=1}^H \Delta_{h, \varepsilon}. \end{aligned}$$

where we define $d_0 = \tau_{\text{obs}} d_x + (\tau_{\text{obs}} - 1) d_u$ and $u_0 = \gamma(8\sigma \sqrt{d_0 \log(9)})$, and

$$\text{ERRTV}(\sigma, \gamma) = \begin{cases} 2c\sigma\sqrt{d_0} & \text{linear } \gamma(u) = c \cdot u, c > 0 \\ u_0 + \int_{u_0}^{\infty} e^{-\frac{\gamma^{-1}(u)^2}{64\sigma^2}} du & \text{general } \gamma(\cdot) \end{cases}. \quad (\text{J.7})$$

In particular, under [Assumption C.1](#), if

$$N_{\text{exp}} \geq c \left(\frac{C_{\Theta} d_{\mathcal{A}} R (R \vee \sqrt{d_{\mathcal{A}}}) \log(dn)}{(\varepsilon/\sigma)^4} \right)^{4\nu} \vee \left(\frac{d_{\mathcal{A}}^6 (R^4 \vee d_{\mathcal{A}}^2 \log^3(\frac{HndR\sigma}{\delta\varepsilon}))}{(\varepsilon/\sigma)^{24}} d_{\mathcal{A}}^2 \right)^{4\nu},$$

then with probability at least $1 - \delta$, it holds that

$$\mathcal{L}_{\text{joint}, \varepsilon_1}(\hat{\pi}_{\sigma}) \leq H \cdot \text{ERRTV}(\sigma, \gamma) + H \sqrt{2\tau_{\text{obs}} - 1} \left(\frac{3\varepsilon}{\sigma} + 6c_5 \sqrt{5d_x + 2 \log\left(\frac{4\sigma}{\varepsilon}\right)} e^{-L_{\beta}(\tau_{\text{chunk}} - \tau_{\text{obs}})} \right).$$

Remark J.4. The second term in our bound on $\mathcal{L}_{\text{joint},\varepsilon}(\pi)$ is identical to the bound in [Theorem 2](#). The term ERRTVC captures the additional penalty we pay to strengthen for imitation of marginals to imitation of joint distributions. Notice that if $\lim_{u \rightarrow 0} \gamma(u) \rightarrow 0$ and $\gamma(u)$ is sufficiently integrable, then, $\lim_{\sigma \rightarrow 0} \text{Err}(\sigma, \gamma) = 0$. This is most clear in the linear $\gamma(\cdot)$ case, where $\text{Err}(\sigma, \gamma) = \mathcal{O}(\sigma)$.

The proof is given in [Appendix J.4.1](#); it mirrors that of [Theorem 2](#), but replaces [Theorem 4](#) with the following imitation guarantee in the composite MDP abstraction of [Appendix D](#), which bounds the joint imitation gap relative to π^* if π^* is TVC.

Proposition J.2. Consider the set-up of [Appendix D](#), and suppose that the assumptions of [Theorem 9](#), but that, in addition, the expert policy π^* is $\tilde{\gamma}(\cdot)$ -TVC with respect to the pseudometric d_{TVC} , where $\tilde{\gamma} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is strictly increasing. Then, for all parameters as in [Theorem 4](#), and any $\tilde{r} > 0$,

$$\begin{aligned} \Gamma_{\text{joint},\varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi^*) &\leq H \int_0^\infty \max_s \mathbb{P}_{s' \sim W_\sigma(s)}[d_{\text{TVC}}(s, s') > \tilde{\gamma}^{-1}(u)/2] du \\ &\quad + H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS},\text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*)), \end{aligned}$$

where the [term in color](#) on the first line is the only term that differs from the bound in [Theorem 4](#).

Moreover, in the special case where all of the distributions of $d_{\text{TVC}}(s, s') \mid s' \sim W_\sigma(s)$ are stochastically dominated by a common random variable Z , and further more $\tilde{\gamma}(u) = \tilde{c} \cdot u$ for some constant \tilde{c} , then our bound may be simplified to

$$\begin{aligned} \Gamma_{\text{joint},\varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi^*) &\leq 2\tilde{c}H\mathbb{E}[Z] \\ &\quad + H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS},\text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},\varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{s}_h^*)). \end{aligned}$$

Proof Sketch. [Proposition J.2](#) is derived below in [Appendix J.4.2](#). It is corollary of [Theorem 4](#), combined with adjoining the coupling constructed therein to a TV distance coupling between π_{rep}^* (whose joints we *can always* imitate) and π^* . Coupling trajectories induced by π_{rep}^* and π^* relies on the TVC of π^* , as well as concentration of W_σ . \square

Using the above proposition, we can derive the following consequences for imitation of the joint distribution.

J.4.1 Proof of Theorem 11

The proof is nearly identical to that of [Theorem 2](#), with the modifications that we replace our use of [Theorem 4](#) with [Proposition J.2](#) taking $\tilde{\gamma} \leftarrow \gamma$. By [Lemma J.1](#) and the assumption that \mathcal{D}_{exp} has $\tau \leq \tau_{\text{obs}}$ -bounded memory, it suffices to bound the joint-gap in the composite MDP:

$$\mathcal{L}_{\text{joint},\varepsilon}(\pi) \leq \Gamma_{\text{joint},\varepsilon}(\pi \parallel \pi^*).$$

We bound this directly from [Proposition J.2](#). The final statement follows from [Theorem 13](#) in the same way that it does in the proof of [Theorem 3](#).

The only remaining modification, then, is to evaluate the additional additive terms colored in purple in [Proposition J.2](#); we will show that ERRTVC as defined in [\(J.7\)](#) suffices as an upper bound. We have two cases. In both, let $d_0 = \tau_{\text{obs}}d_x + (\tau_{\text{obs}} - 1)d_u$. As d_{TVC} measures the distance between the chunks $\phi_h = \phi_o(s_h)$, $\tilde{\phi}_h = \phi_o(s'_h)$, which have dimension d_0 , and since we $\phi_o \circ W_\sigma(\cdot) = \mathcal{N}(\cdot, \sigma^2 \mathbf{I}_{d_0})$, we have

$$d_{\text{TVC}}(\phi_o \circ s, \phi_o \circ s') \mid s' \sim W_\sigma(s) \stackrel{\text{dist}}{=} \|\gamma\|, \quad \gamma \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{d_0}) \quad (\text{J.8})$$

General $\gamma(\cdot)$. Eq. [\(J.8\)](#) and [Lemma J.7](#) imply that

$$\mathbb{P}_{s' \sim W_\sigma(s)}[d_{\text{TVC}}(s, s')] \leq \exp(-r^2/16\sigma^2), \quad r \geq 4\sigma d_0 \log(9).$$

Hence, if $u_0 = \gamma(8\sigma d_0 \log(9))$, then

$$\mathbb{P}[d_{\text{TVC}}(s, s') > \gamma^{-1}(u)/2] \leq \exp(-\gamma^{-1}(u)^2/64\sigma^2), \quad u \geq u_0.$$

Thus, as probabilities are at most one,

$$\int_0^\infty \max_s \mathbb{P}_{s' \sim W_\sigma(s)}[d_{\text{TVC}}(s, s') > \gamma^{-1}(u)/2] du \leq u_0 + \int_{u_0}^\infty e^{-\frac{\gamma^{-1}(u)^2}{64\sigma^2}} du,$$

as needed.

Linear $\gamma(\cdot)$. In the special case where $\gamma(u) = c(u)$, [Eq. \(J.8\)](#) implies that we can take $Z = \|\gamma\|$ where $\gamma \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{d_0})$ in the second part of [Proposition J.2](#). The corresponding additive term is then $2Hc\mathbb{E}[\|\gamma\|]$. By Jensen's inequality, $\mathbb{E}[\|\gamma\|] \leq \sqrt{\mathbb{E}[\|\gamma\|^2]} = \sqrt{\sigma^2 d_0} = \sigma\sqrt{d_0}$, as needed. \square

J.4.2 Proof of [Proposition J.2](#)

Define the shorthand

$$B := H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS}, \text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{\mathbf{s}_h^* \sim \mathcal{P}_h^*} \mathbb{E}_{\tilde{\mathbf{s}}_h^* \sim W_\sigma(\mathbf{s}_h^*)} \mathbf{d}_{\text{os}, \varepsilon}(\hat{\pi}_h(\tilde{\mathbf{s}}_h^*) \parallel \pi_{\text{dec}}^*(\tilde{\mathbf{s}}_h^*)),$$

and recall that [Theorem 4](#) ensures $\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*) \leq B$. Further, recall from [Definition D.1](#) that

$$\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*) = \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\max_{h \in [H]} \max\{\mathbf{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}), \mathbf{d}_A(\mathbf{a}_h^{\text{rep}}, \hat{\mathbf{a}}_h)\} > \varepsilon \right],$$

where the infimum is over all couplings μ_1 of $(\hat{\mathbf{s}}_{1:H+1}, \hat{\mathbf{a}}_{1:H}) \sim D_{\hat{\pi} \circ W_\sigma}$ and $(\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}}) \sim D_{\pi_{\text{rep}}^*}$ with $\mathbb{P}_{\mu_1}[\hat{\mathbf{s}}_1 = \mathbf{s}_1^{\text{rep}}] = 1$. For any coupling μ_1 , we can consider another coupling μ_2 of $(\mathbf{s}_{1:H+1}^*, \mathbf{a}_{1:H}^*) \sim D_{\pi^*}$ and $(\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}}) \sim D_{\pi_{\text{rep}}^*}$ with $\mathbb{P}_{\mu_2}[\mathbf{s}_1^* = \mathbf{s}_1^{\text{rep}}] = 1$. By the ‘‘gluing lemma’’ ([Lemma F.2](#)), we can construct a combined coupling μ which respects the marginals of μ_1 and μ_2 . This combined coupling induces a joint coupling $\tilde{\mu}_1$ of $D_{\hat{\pi} \circ W_\sigma}$ and D_{π^*} which, by a union bound, satisfies $\mathbb{P}_{\tilde{\mu}_1}[\hat{\mathbf{s}}_1 = \mathbf{s}_1^*] = 1$. Thus, by a union bound, we can bound

$$\begin{aligned} \Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi^*) &\leq \mathbb{P}_{\tilde{\mu}_1} \left[\max_{h \in [H]} \max\{\mathbf{d}_S(\mathbf{s}_{h+1}^*, \hat{\mathbf{s}}_{h+1}), \mathbf{d}_A(\mathbf{a}_h^*, \hat{\mathbf{a}}_h)\} > \varepsilon \right] \\ &\leq \mathbb{P}_{\mu_1} \left[\max_{h \in [H]} \max\{\mathbf{d}_S(\mathbf{s}_{h+1}^{\text{rep}}, \hat{\mathbf{s}}_{h+1}), \mathbf{d}_A(\mathbf{a}_h^{\text{rep}}, \hat{\mathbf{a}}_h)\} > \varepsilon \right] \\ &\quad + \mathbb{P}_{\mu_2} [(\mathbf{s}_{1:H+1}^*, \mathbf{a}_{1:H}^*) \neq (\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}})]. \end{aligned}$$

Passing to the infimum over μ_1, μ_2 ,

$$\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi^*) \leq \underbrace{\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)}_{\leq B} + \inf_{\mu_2} \mathbb{P}_{\mu_2} [(\mathbf{s}_{1:H+1}^*, \mathbf{a}_{1:H}^*) \neq (\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}})],$$

where again μ_2 quantify couplings of $(\mathbf{s}_{1:H+1}^*, \mathbf{a}_{1:H}^*) \sim D_{\pi^*}$ and $(\mathbf{s}_{1:H+1}^{\text{rep}}, \mathbf{a}_{1:H}^{\text{rep}}) \sim D_{\pi_{\text{rep}}^*}$ with $\mathbb{P}_{\mu_2}[\mathbf{s}_1^* = \mathbf{s}_1^{\text{rep}}] = 1$. Bounding the infimum over μ_2 with [Proposition J.4](#), we have

$$\Gamma_{\text{joint}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi^*) \leq B + \sum_{h=1}^H \mathbb{E}_{\mathbf{s}_h^*} \text{TV}(\pi_h^*(\mathbf{s}_h^*), \pi_{\text{rep}, h}^*(\mathbf{s}_h^*))$$

To conclude, it suffices to show the following bound:

Claim J.3. For any $\mathbf{s} \in \mathcal{S}$, $h \in [H]$, and $\tilde{r} \geq 0$, $\text{TV}(\pi_h^*(\mathbf{s}), \pi_{\text{rep}, h}^*(\mathbf{s})) \leq \int_0^\infty \max_{\mathbf{s}} \max_{\mathbf{s}'} \mathbb{P}_{\mathbf{s}' \sim W_\sigma(\mathbf{s})} [\mathbf{d}_{\text{TVC}}(\mathbf{s}, \mathbf{s}') > \tilde{\gamma}^{-1}(u)/2] du$.

Proof. To show this claim, we note that we can represent (via the notation in [Appendix H.3](#)) $\pi_{\text{rep}, h}^*(\mathbf{s}) = \pi_h^* \circ W_{\text{rep}, h}^*(\mathbf{s})$, where $W_{\text{rep}, h}^*$ is the replica-kernel defined in [Definition H.5](#). Thus, we can construct a coupling of $\mathbf{a}^* \sim \pi_h^*(\mathbf{s})$ and $\mathbf{a}^{\text{rep}} \sim \pi_{\text{rep}, h}^*(\mathbf{s})$ by introducing an intermediate state $\mathbf{s}' \sim W_{\text{rep}, h}^*(\mathbf{s})$ and $\mathbf{a}^{\text{rep}} \sim \pi_h^*(\mathbf{s}')$. By [Lemma F.4](#), the fact that TV distance is bounded by one, and the assumption that π^* is $\tilde{\gamma}$ -TVC, we then have

$$\text{TV}(\pi_h^*(\mathbf{s}), \pi_{\text{rep}, h}^*(\mathbf{s})) \leq \mathbb{E}_{\mathbf{s}' \sim W_{\text{rep}, h}^*(\mathbf{s})} \text{TV}(\pi_h^*(\mathbf{s}), \pi_h^*(\mathbf{s}')).$$

Recall the well-known formula that, for a non-negative random variable X , $\mathbb{E}[X] = \int_0^\infty \mathbb{P}[X > u] du$ [\[22\]](#). From this formula, we find

$$\begin{aligned} \text{TV}(\pi_h^*(\mathbf{s}), \pi_{\text{rep}, h}^*(\mathbf{s})) &\leq \int_0^\infty \mathbb{P}[\text{TV}(\pi_h^*(\mathbf{s}), \pi_h^*(\mathbf{s}')) > u] du \\ &\stackrel{(i)}{\leq} \int_0^\infty \mathbb{P}[\mathbf{d}_{\text{TVC}}(\mathbf{s}, \mathbf{s}') > \tilde{\gamma}^{-1}(u)] du \end{aligned}$$

where in (i) we used that $\text{TV}(\pi_h^*(s), \pi_h^*(s')) \leq \tilde{\gamma}(\text{d}_{\text{TV}}(s, s'))$ and that, as $\tilde{\gamma}(\cdot)$ is strictly increasing, we have the equality of events $\{\text{TV}(\pi_h^*(s), \pi_h^*(s')) > u\} = \{\text{d}_{\text{TV}}(s, s') > \tilde{\gamma}^{-1}(u)\}$. Arguing as in the proof of [Lemma H.5](#), we have that $\mathbb{P}_{s' \sim W_\sigma(s)}[\text{d}_{\text{TV}}(s, s') > \tilde{\gamma}^{-1}(u)] \leq \max_s \mathbb{P}_{s' \sim W_\sigma(s)}[\text{d}_{\text{TV}}(s, s') > \tilde{\gamma}^{-1}(u)/2]$. Hence, we conclude

$$\text{TV}(\pi_h^*(s), \pi_{\text{rep}, h}^*(s)) \leq \int_0^\infty \max_s \mathbb{P}_{s' \sim W_\sigma(s)}[\text{d}_{\text{TV}}(s, s') > \tilde{\gamma}^{-1}(u)/2] du$$

which proves the first guarantee. \square

With the above claim proven, we conclude the proof of the first statement of [Proposition J.2](#). For the second statement, we observe that under the stated stochastic domination assumption by Z , and if $\tilde{\gamma}(u) = \tilde{c} \cdot u$, then $\max_s \mathbb{P}_{s' \sim W_\sigma(s)}[\text{d}_{\text{TV}}(s, s') > \tilde{\gamma}^{-1}(u)/2] \leq \mathbb{P}[Z > \frac{u}{2\tilde{c}}]$. Hence, by a change of variables $u = \frac{t}{2\tilde{c}}$,

$$\int_0^\infty \max_s \mathbb{P}_{s' \sim W_\sigma(s)}[\text{d}_{\text{TV}}(s, s') > \tilde{\gamma}^{-1}(u)/2] du \leq \int_0^\infty \mathbb{P}[Z > \frac{u}{2\tilde{c}}] = 2\tilde{c} \int_0^\infty \mathbb{P}[Z > u] = 2\tilde{c}\mathbb{E}[Z],$$

where again we invoke that Z must be nonnegative (to stochastically dominate non-negative random variables), and thus used the expectation formula referenced above. \square

J.5 Imitation in total variation distance

Here, we notice that estimating the score in TV distance fascilliates estimation in the composite MDP, with no smoothing:

Theorem 12. Define the error term

$$\Delta_{\text{TV}, h}(\hat{\pi}) := \sum_{h=1}^H \mathbb{E}_{\mathbf{o}_h \sim \mathcal{D}_{\text{exp}, h}} \Delta_{(\varepsilon)}(\pi_h^*(\mathbf{o}_h), \hat{\pi}_h(\mathbf{o}_h)) \Big|_{\varepsilon=0}$$

We have the characterization $\Delta_{\text{TV}, h}(\hat{\pi}) = \mathbb{E}_{\mathbf{o}_h} \text{TV}(\pi_h^*(\mathbf{o}_h), \hat{\pi}_h(\mathbf{o}_h))$ where \mathbf{o}_h has the distribution induced by $\hat{\pi}$, and $\pi^*(\mathbf{o}_h)$ denotes the distribution of $\mathbf{a}_h \mid \mathbf{o}_h$ under \mathcal{D}_{exp} as in [Definition J.3](#). Then, under no additional assumption (not even those in [Section 3](#)),

$$\mathcal{L}_{\text{fin}, \varepsilon=0}(\hat{\pi}_\sigma) \leq \mathcal{L}_{\text{marg}, \varepsilon=0}(\hat{\pi}_\sigma) \leq \sum_{h=1}^H \Delta_{\text{TV}, h}(\hat{\pi})$$

In addition π^* has τ -bounded memory ([Definition J.5](#)) for $\tau \leq \tau_{\text{obs}}$, then for $\mathcal{L}_{\text{joint}, \varepsilon}$ as in [Definition J.4](#),

$$\mathcal{L}_{\text{joint}, \varepsilon=0}(\hat{\pi}) \leq \sum_{h=1}^H \Delta_{\text{TV}, h}(\hat{\pi})$$

The first part of the above theorem is a direct consequence [Proposition F.3](#). Then, the second part of the theorem follows by combining the proposition below in the composite MDP, together with the correct instantiations for control, and [Lemma J.1](#) to convert $\mathcal{L}_{\text{marg}, \varepsilon}$ and $\mathcal{L}_{\text{fin}, \varepsilon}$ into $\Gamma_{\text{marg}, \varepsilon} \leq \Gamma_{\text{joint}, \varepsilon}$, and $\Gamma_{\text{joint}, \varepsilon}$, respectively.

Proposition J.4. Consider the composite MDP setting of [Appendix D](#). Then, there exists a coupling

$$\text{TV}(\mathcal{D}_{\hat{\pi}}, \mathcal{D}_{\pi^*}) \leq \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \text{TV}(\pi_h^*(s_h^*), \hat{\pi}_h(s_h^*))$$

Thus, there exists a couple $\mu \in \mathcal{C}(\mathcal{D}_{\pi^*}, \mathcal{D}_{\hat{\pi}})$ of $(s_{1:H+1}^*, a_{1:H}^*) \sim \mathcal{D}_{\pi^*}$ and $(\hat{s}_{1:H+1}, \hat{a}_{1:H}) \sim \mathcal{D}_{\hat{\pi}}$ such that $\mathbb{P}_\mu[(s_{1:H+1}^*, a_{1:H}^*) \neq (\hat{s}_{1:H+1}, \hat{a}_{1:H})]$ is bounded by the right-hand side of the above display. Moreover, this coupling can be constructed such that $\mathbb{P}_\mu[s_1^* = \hat{s}_1]$.

Proof of Proposition J.4. This is a direct consequence of [Lemma J.9](#), with $P_1 \leftarrow P_{\text{init}}$, and Q_{h+1} corresponding to the kernel for sampling $a_h^* \sim \pi^*(s_h^*)$ and incrementing the dynamics $s_{h+1}^* = F_h(s_h^*, a_h^*)$, and Q'_h the same for $\hat{a}_h \sim \hat{\pi}_h(\hat{s}_h)$, and similar incrementing of the dynamics. \square

J.6 Consequence for expected costs

Finally, we prove [Proposition J.5](#), which shows that it is sufficient to control the imitation losses in [Definition 2.2](#) if we wish to control the difference of a Lipschitz cost function between the learned policy and the expert distribution:

Proposition J.5. Recall the marginal and final imitation losses in [Definition 2.2](#), and also the joint imitation loss in [Definition J.4](#). Consider a cost function $\mathfrak{J} : \mathcal{P}_T \rightarrow \mathbb{R}$ on trajectories $\rho_T \in \mathcal{P}_T$. Finally, let $\rho_T \sim \mathcal{D}_{\text{exp}}$, and let $\rho'_T \sim \mathcal{D}_\pi$ be under the distribution induced by π . Then,

- (a) If $\max_{\rho_T} |\mathfrak{J}(\rho_T)| \leq B$, and ρ_T is L Lipschitz in the Euclidean norm¹² (treating ρ_T as Euclidean vector in $\mathbb{R}^{(T+1)d_x + Td_u}$), then

$$|\mathbb{E}_{\mathcal{D}_{\text{exp}}}[\mathfrak{J}(\rho_T)] - \mathbb{E}_{\mathcal{D}_\pi}[\mathfrak{J}(\rho'_T)]| \leq \sqrt{2T}L\varepsilon + 2B\mathcal{L}_{\text{joint},\varepsilon}(\pi).$$

- (b) If \mathfrak{J} decomposes into a sum of costs, $\mathfrak{J}(\rho) = \ell_{T+1,1}(\mathbf{x}_{1:T}) + \sum_{t=1}^T \ell_{t,1}(\mathbf{x}_t) + \ell_{t,2}(\mathbf{u}_t)$, where $\ell_{t,1}(\cdot), \ell_{t,2}(\cdot)$ are L -Lipschitz and bounded in magnitude in B . Then,

$$|\mathbb{E}_{\mathcal{D}_{\text{exp}}}[\mathfrak{J}(\rho_T)] - \mathbb{E}_{\mathcal{D}_\pi}[\mathfrak{J}(\rho'_T)]| \leq 4TB\mathcal{L}_{\text{marg},\varepsilon}(\pi) + 2TL\varepsilon.$$

- (c) $\mathfrak{J}(\rho) = \ell_{T+1,1}(\mathbf{x}_{T+1})$ depends only on \mathbf{x}_{T+1} , then

$$|\mathbb{E}_{\mathcal{D}_{\text{exp}}}[\mathfrak{J}(\rho_T)] - \mathbb{E}_{\mathcal{D}_\pi}[\mathfrak{J}(\rho'_T)]| \leq 2B\mathcal{L}_{\text{fin},\varepsilon}(\pi) + L\varepsilon$$

Thus, for our imitation guarantees to apply to most natural cost functions used in practice, it suffices to control the imitation losses defined above.

Proof of Proposition J.5. Let $\rho_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T}) \sim \mathcal{D}_{\text{exp}}$, and let $\rho'_T = (\mathbf{x}'_{1:T+1}, \mathbf{u}'_{1:T})$ be under the distribution induced by π .

Part (a). For any coupling μ between the two under which $\mathbf{x}_1 = \mathbf{x}'_1$, and let $\mathcal{E}_\varepsilon := \{\max_t \|\mathbf{x}_{t+1} - \mathbf{x}'_{t+1}\| \vee \|\mathbf{u}_t - \mathbf{u}'_t\| \leq \varepsilon\}$.

$$\begin{aligned} |\mathbb{E}[\mathfrak{J}(\rho_T)] - \mathbb{E}[\mathfrak{J}(\rho'_T)]| &= |\mathbb{E}_\mu[\mathfrak{J}(\rho_T) - \mathfrak{J}(\rho'_T)]| \\ &\leq \mathbb{E}_\mu[|\mathfrak{J}(\rho_T) - \mathfrak{J}(\rho'_T)|] \\ &\leq 2B\mathbb{P}_\mu[\mathcal{E}_\varepsilon^c] + \mathbb{E}_\mu[|\mathfrak{J}(\rho_T) - \mathfrak{J}(\rho'_T)|\mathbf{I}\{\mathcal{E}_\varepsilon\}] \end{aligned}$$

By passing to an infimum over couplings, $\inf_\mu \mathbb{P}_\mu[\mathcal{E}_\varepsilon^c] \leq \mathcal{L}_{\text{joint},\varepsilon}(\pi)$. Moreover, we observe that under μ , $\mathbf{x}_1 = \mathbf{x}'_1$, and the remaining coordinates, $(\mathbf{x}_{2:T+1}, \mathbf{u}_{1:T})$ and $(\mathbf{x}'_{2:T+1}, \mathbf{u}'_{1:T})$ are the concatenation of $2T$ vectors, so the Euclidean norm of the concatenations $\|\rho_T - \rho'_T\|$ is at most $\sqrt{2T} \max_t \|\mathbf{x}_{t+1} - \mathbf{x}'_{t+1}\| \vee \|\mathbf{u}_t - \mathbf{u}'_t\|$, which on \mathcal{E}_ε is at most $\sqrt{2T}\varepsilon$. Using Lipschitz-ness of \mathfrak{J} concludes.

Part (b) Using the adaptive decomposition of the cost and the fact that \mathbf{x}_1 and \mathbf{x}'_1 have the same distributions,

$$\begin{aligned} |\mathbb{E}[\mathfrak{J}(\rho_T)] - \mathbb{E}[\mathfrak{J}(\rho'_T)]| &= \left| \sum_{t=1}^T (\mathbb{E}[\ell_{t,1}(\mathbf{x}_{t+1})] - \mathbb{E}[\ell_{t,1}(\mathbf{x}'_{t+1})]) + (\mathbb{E}[\ell_{t,2}(\mathbf{u}_t)] - \mathbb{E}[\ell_{t,2}(\mathbf{u}'_t)]) \right| \\ &\leq \sum_{t=1}^T |\mathbb{E}[\ell_{t,1}(\mathbf{x}_{t+1})] - \mathbb{E}[\ell_{t,1}(\mathbf{x}'_{t+1})]| + |\mathbb{E}[\ell_{t,2}(\mathbf{u}_t)] - \mathbb{E}[\ell_{t,2}(\mathbf{u}'_t)]| \end{aligned}$$

Applying similar arguments as in part (a) to each term, we can bound

$$\max \{ |\mathbb{E}[\ell_{t,1}(\mathbf{x}_{t+1})] - \mathbb{E}[\ell_{t,1}(\mathbf{x}'_{t+1})]|, |\mathbb{E}[\ell_{t,2}(\mathbf{u}_t)] - \mathbb{E}[\ell_{t,2}(\mathbf{u}'_t)]| \} \leq 2B\mathcal{L}_{\text{marg},\varepsilon}(\pi) + L\varepsilon.$$

Summing over the $2T$ terms concludes.

Part (c). Follows similar to part (b). □

¹²Of course, Lipschitzness in other norms can be derived, albeit with different T dependence

J.7 Useful Lemmata

J.7.1 On the trajectories induced by π^* from \mathcal{D}_{exp}

The key step in all of our proofs is to relate the expert distribution over trajectories $\rho_T \sim \mathcal{D}_{\text{exp}}$ to the distribution induced by the chunking policy π^* in [Definition J.3](#), which induces distribution \mathcal{D}_{π^*} .

Lemma J.6. There exists a sequence of probability kernels π_h^* mapping $\mathbf{o}_h \rightarrow \Delta(\mathcal{A})$ such that the chunking policy $\pi^* = (\pi_h^*)_{1 \leq h \leq H}$ satisfies the following:

- (a) $\pi_h^*(\mathbf{o}_h)$ is equal to the almost-sure conditional probability of \mathbf{a}_h conditioned on \mathbf{o}_h under $\rho_T \sim \mathcal{D}_{\text{exp}}$ and $\mathbf{a}_{1:H} = \text{synth}(\rho_T)$.
- (b) The marginal distribution over each $\mathbf{s}_h \sim P_h^*$ (as defined in [Definition J.3](#)) is the same as the marginals of each \mathbf{s}_h under $\rho_T \sim \mathcal{D}_{\text{exp}}$, and hence, $(\mathbf{s}_h, \mathbf{a}_h) \sim \mathcal{D}_{\text{exp}}$ has the same distribution as $(\mathbf{s}_h, \mathbf{a}_h)$ where $\mathbf{s}_h \sim P_h^*, \mathbf{a}_h \mid \mathbf{s}_h \sim \pi_h^*(\mathbf{s}_h)$.
- (c) If \mathcal{D}_{exp} has τ -bounded memory ([Definition J.5](#)) and if $\tau \leq \tau_{\text{obs}}$, then the joint distribution of ρ_T induced by π^* is equal to the joint distribution over ρ_T under \mathcal{D}_{exp} .
- (d) Again, let P_h^* be as defined in [Definition J.3](#). Consider any sequence of kernels $W_{1:H}$ satisfying the constraint that $\phi_o \circ W_h(\mathbf{s}) \ll \phi_o \circ P_h^*$ for all \mathbf{s}, h . Consider a sequence of states $\mathbf{s}_{1:H+1}$ drawn as in [Definition H.9](#) by $\mathbf{s}_1 \sim P_{\text{init}}, \tilde{\mathbf{s}}_h \sim W_h(\mathbf{s}_h)$, and $\mathbf{s}_a \sim \pi_h^*(\tilde{\mathbf{s}}_h)$, $\mathbf{s}_{s+1} = F_h(\mathbf{s}_h, \mathbf{a}_h)$. Then, let $\tilde{\mathbf{o}}_h = \phi_o(\tilde{\mathbf{s}}_h)$ and $\mathbf{a}_h \sim \pi_h^*(\mathbf{a}_h)$. Then the distribution of $(\tilde{\mathbf{o}}_h, \mathbf{a}_h)$ is absolutely continuous with respect to the distribution of $(\mathbf{o}_h, \mathbf{a}_h) \sim \mathcal{D}_{\text{exp}}$.

Remark J.5 (Replacing τ -bounded memory with mixing). We can replace that τ -bounded memory condition to the following mixing assumption. Define the chunk $\rho_{i \leq j} = (\mathbf{x}_{i:j}, \mathbf{u}_{i:j-1})$. Define the measures

$$\begin{aligned} Q_h(\mathbf{o}_h) &= \mathbb{P}_{\mathbf{a}_{1:h-1}, \rho_{1:t_h-\tau_{\text{obs}}-1}, \mathbf{a}_{h:H}, \rho_{t_h:T+1} | \mathbf{o}_h} \\ Q_h^\otimes(\mathbf{o}_h) &= \mathbb{P}_{\mathbf{a}_{1:h-1}, \rho_{1:t_h-\tau_{\text{obs}}-1} | \mathbf{o}_h} \otimes \mathbb{P}_{\mathbf{a}_{h:H}, \rho_{t_h:T+1} | \mathbf{o}_h} \end{aligned}$$

which describes the conditional distribution of the whole trajectory without \mathbf{o}_h and the product-distribution of the conditional distributions of the before- \mathbf{o}_h part of the trajectory, and after \mathbf{o}_h -part. Under the condition

$$\mathbb{E}_{\mathbf{o}_h \text{ from } \rho_T \sim \mathcal{D}_{\text{exp}}} \text{TV}(Q_h(\mathbf{o}_h), Q_h^\otimes(\mathbf{o}_h)) \leq \varepsilon_{\text{mix}}(\tau_{\text{obs}}),$$

which measures how close the before- and after- \mathbf{o}_h parts of the trajectory are to being conditionally independent, one can leverage [Lemma J.9](#) to show that

$$\text{TV}(\mathcal{D}_{\pi^*}, \mathcal{D}_{\text{exp}}) \leq H \varepsilon_{\text{mix}}(\tau_{\text{obs}})$$

[Lemma J.6](#) corresponds to the special when $\varepsilon_{\text{mix}} = 0$.

Proof of [Lemma J.6](#). We prove each part in sequence

Part (a). follows from the fact that all random variables are in real vector spaces, and thus Polish spaces. Hence, we can invoke the existence of regular conditional probabilities by [Theorem 7](#).

Part (b). This follows by marginalization and Markovianity of the dynamics. Specifically, let $(\rho_T^*, \mathbf{a}_{1:H}^*)$ be a trajectory and composite actions induced by the chunking policy π^* , and let $(\rho_T, \mathbf{a}_{1:H})$ be the same induced by \mathcal{D}_{exp} . Let \mathbf{o}_h^* denote observation chunks of ρ_T^* , and let \mathbf{o}_h observation chunks of ρ_T (length $\tau_{\text{obs}} - 1$); similarly, denote by \mathbf{s}_h^* and \mathbf{s}_h the respective trajectory chunks (length $\tau_{\text{chunk}} \geq \tau_{\text{obs}}$).

We argue inductively that the trajectory chunks \mathbf{s}_h^* and \mathbf{s}_h are identically distributed for each h . For $h = 1$, \mathbf{s}_1^* and \mathbf{s}_1 are identically distributed according to $\mathcal{D}_{\mathbf{x}_1}$. Now assume we have show that \mathbf{s}_h^* and \mathbf{s}_h are identically distributed. As observation chunks are sub-chunks of trajectory chunks, this means that \mathbf{o}_h^* and \mathbf{o}_h are identically distributed. By part (a), it follows that $(\mathbf{o}_h^*, \mathbf{a}_h^*)$ and $(\mathbf{o}_h, \mathbf{a}_h)$ are identically distributed. In particular, $(\mathbf{x}_{t_h}^*, \mathbf{a}_h^*)$ and $(\mathbf{x}_{t_h}, \mathbf{a}_h)$ are identically distributed, where $\mathbf{x}_{t_h}^*$ (resp \mathbf{x}_{t_h}) these denote the t_h -th control state under π^* (resp. \mathcal{D}_{exp}). By Markovianity of the dynamics, \mathbf{s}_{h+1}^* and \mathbf{s}_{h+1} are functions of $(\mathbf{x}_{t_h}^*, \mathbf{a}_h^*)$ and $(\mathbf{x}_{t_h}, \mathbf{a}_h)$, respectively, \mathbf{s}_{h+1}^* and \mathbf{s}_{h+1} are identically distributed, as needed.

Part (c). When \mathcal{D}_{exp} has τ -bounded memory and $\tau \leq \tau_{\text{obs}}$, then we have the almost-sure equality

$$\mathbb{P}_{\mathcal{D}_{\text{exp}}}[\mathbf{a}_h \in \cdot \mid \mathbf{x}_{1:t_h}, \mathbf{u}_{1:t_h}] = \mathbb{P}_{\mathcal{D}_{\text{exp}}}[\mathbf{a}_h \in \cdot \mid \mathbf{o}_h] = \pi_h^*(\mathbf{o}_h)[\mathbf{a}_h \in \cdot].$$

Finally, $\mathbf{x}_{t_h+1:t_{h+1}}, \mathbf{u}_{t_h:t_{h+1}-1}$ are determined by \mathbf{x}_{t_h} and \mathbf{a}_h , this inductively establishes equality of the joint-trajectory distributions.

Part (d) That the distributions of $\tilde{\mathbf{o}}_h = \phi_o(\tilde{\mathbf{s}}_h)$ under the construction in part (d) is absolutely continuous with respect to $\tilde{\mathbf{o}}_h$ under \mathcal{D}_{exp} coincide is immediate from the condition of W_h . The second part follows from part (a). \square

J.7.2 Concentration and TVC of Gaussian Smoothing.

We now include two easy lemmata necessary for the proof. The first shows that p_r is small when r is $\Theta(\sigma)$ by elementary Gaussian concentration:

Lemma J.7. Suppose that $\gamma \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ is a centred Gaussian vector with covariance $\sigma^2 \mathbf{I}$ in \mathbb{R}^d for some $\sigma > 0$. Then for all $p > 0$, it holds with probability at least $1 - p$ that

$$\|\gamma\| \leq 2\sigma \cdot \sqrt{2d \log(9) + 2 \log\left(\frac{1}{p}\right)} \leq 2\sigma \cdot \sqrt{5d + 2 \log\left(\frac{1}{p}\right)}$$

Moreover, for $r \geq 4\sigma \sqrt{d \log(9)}$, $\mathbb{P}[\|\gamma\| \geq r] \leq \exp(-r^2/16\sigma^2)$.

Proof. We apply the standard covering based argument as in, e.g., Vershynin [71, Section 4.2]. Note that

$$\|\gamma\| = \sup_{\mathbf{w} \in \mathcal{S}^{d-1}} \langle \gamma, \mathbf{w} \rangle,$$

where \mathcal{S}^{d-1} is the unit sphere in \mathbb{R}^d . Let \mathcal{U} denote a minimal $(1/4)$ -net on \mathcal{S}^{d-1} and observe that a simple computation tells us that

$$\sup_{\mathbf{w} \in \mathcal{S}^{d-1}} \langle \gamma, \mathbf{w} \rangle \leq 2 \cdot \max_{\mathbf{w} \in \mathcal{U}} \langle \mathbf{w}, \gamma \rangle.$$

A classical volume argument (see for example, Vershynin [71, Section 4.2]) tells us that $|\mathcal{U}| \leq 9^d$. A classical Gaussian tail bound tells us that for any $\mathbf{w} \in \mathcal{S}^{d-1}$, it holds that for any $r > 0$,

$$\mathbb{P}(\langle \mathbf{w}, \gamma \rangle > r) \leq e^{-\frac{r^2}{2\sigma^2}}.$$

Thus by a union bound, we have

$$\mathbb{P}(\|\gamma\| > r) \leq |\mathcal{U}| \cdot \max_{\mathbf{w} \in \mathcal{U}} \mathbb{P}\left(\|\gamma\| > \frac{r}{2}\right) \leq 9^d \cdot e^{-\frac{r^2}{8\sigma^2}}.$$

Inverting concludes the proof. \square

The second lemma shows that the relevant smoothing kernel is TVC:

Lemma J.8. For any $\sigma > 0$, let ϕ_o and W_σ be as in Definition J.1 kernel, then W_σ is γ_{TVC} -TVC for with respect to \mathbf{d}_{TVC} (as defined in Appendix D.2)

$$\gamma_{\text{TVC}}(u) = \frac{u\sqrt{2\tau_{\text{obs}} - 1}}{2\sigma}.$$

Proof. Recall that ϕ_o denotes projection onto the \mathcal{O} -component of the direct decomposition in Definition H.1, i.e. projects onto the observation chunk \mathbf{o}_h . We apply Pinsker's inequality [52]: Then, for $\mathbf{s}, \mathbf{s}' \in \mathbb{R}^p$, we have

$$\text{TV}(\phi_o \circ W_\sigma(\mathbf{s}), \phi_o \circ W_\sigma(\mathbf{s}')) \leq \sqrt{\frac{1}{2} \cdot \text{D}_{\text{KL}}(\phi_o \circ W_\sigma(\mathbf{s}) \parallel \phi_o \circ W_\sigma(\mathbf{s}'))}.$$

Note that for $s = s_h$ with corresponding observation chunk \mathbf{o}_h $\phi_o \circ W_\sigma(s) \sim \mathcal{N}(\mathbf{o}_h, \sigma^2 \mathbf{I})$. Similarly, for \mathbf{o}'_h corresponding to s' , $\phi_o \circ W_\sigma(s') \sim \mathcal{N}(\mathbf{o}'_h, \sigma^2 \mathbf{I})$. Hence,

$$D_{\text{KL}}(\phi_o \circ W_\sigma(s) \parallel \phi_o \circ W_\sigma(s')) \leq \frac{\|\mathbf{o}_h - \mathbf{o}'_h\|^2}{2\sigma^2}.$$

Thus, we conclude $\text{TV}(\phi_o \circ W_\sigma(s), \phi_o \circ W_\sigma(s')) \leq \frac{\|\mathbf{o}_h - \mathbf{o}'_h\|}{2\sigma}$. Finally, we upper bound the Euclidean norm $\|\mathbf{o}_h - \mathbf{o}'_h\|$ of vectors consisting of $2\tau_{\text{obs}} - 1$ sub-vectors via d_{TVc} (which is the maximum Euclidean norm of these subvectors) via $\|\mathbf{o}_h - \mathbf{o}'_h\| \leq \sqrt{2\tau_{\text{obs}} - 1} d_{\text{TVc}}(s, s')$. \square

J.7.3 Total Variation Telescoping

Lemma J.9 (Total Variation Telescoping). Let $\mathcal{Y}_1, \dots, \mathcal{Y}_H, \mathcal{Y}_{H+1}$ be Polish spaces. Let $P_1 \in \Delta(\mathcal{Y}_1)$, and let $Q_h, Q'_h \in \Delta(\mathcal{Y}_h \mid \mathcal{X}, \mathcal{Y}_{1:h-1})$, $h > 1$. Define $P'_1 = P_1$, and recursively define

$$P_h = \text{law}(Q_h; P_{h-1}), \quad P'_h = \text{law}(Q'_h; P'_{h-1}), \quad h > 1.$$

Then,

$$\text{TV}(P_{H+1}, P'_{H+1}) \leq \sum_{h=1}^H \mathbb{E}_{Y_{1:h} \sim P_h} \text{TV}(Q_{h+1}(\cdot \mid Y_{1:h}), Q'_{h+1}(\cdot \mid Y_{1:h}))$$

Moreover, there exists a coupling of $\mu \in \mathcal{C}(P_{H+1}, P'_{H+1})$ over $Y_{1:H+1} \sim P_{H+1}$ and $Y'_{1:H+1} \sim P'_{H+1}$ such that

$$\mathbb{P}_\mu[Y_1 = Y'_1] = 1, \quad \mathbb{P}_\mu[Y_{1:H+1} \neq Y'_{1:H+1}] \leq \sum_{h=1}^H \mathbb{E}_{Y_{1:h} \sim P_h} \text{TV}(Q_{h+1}(\cdot \mid Y_{1:h}), Q'_{h+1}(\cdot \mid Y_{1:h})).$$

Proof. To prove the first part of the lemma, define $Q'_{i,j}$ for $2 \leq i \leq j \leq H+1$ by $Q'_{i,i} = Q_i$ define $Q'_{i,j}$ by appending $Q'_{i,j}$ to $Q'_{i,j-1}$. and $\text{law}(Q'_{i,j}; (\cdot)) = \text{law}(Q'_j; \text{law}(Q_{i,j-1}; (\cdot)))'$. We now define

$$P^{(i)} = \text{law}(Q'_{i+1,H+1}; P_i),$$

with the convention $\text{law}(Q'_{H+2,H+1}; P_{H+1}) = P_{H+1}$. Note that $P^{(H+1)} = P_{H+1}$, and $P^{(1)} = P'_{H+1}$. Then, because TV distance is a metric,

$$\text{TV}(P_{H+1}, P'_{H+1}) \leq \sum_{h=1}^H \text{TV}(P^{(h)}, P^{(h+1)})$$

Moreover, we can write $P^{(i)} = \text{law}(Q'_{i+2,H+1}; \text{law}(Q'_{i+1}; P_i))$ and $P_{i+1} = \text{law}(Q_{i+1}; P_i)$. Thus,

$$\begin{aligned} \text{TV}(P^{(i)}, P^{(i+1)}) &= \text{TV}(\text{law}(Q'_{i+2,H+1}; \text{law}(Q'_{i+1}; P_i)), \text{law}(Q'_{i+2,H+1}; \text{law}(Q_{i+1}; P_i))) \\ &\quad \text{(Lemma F.4)} \\ &= \text{TV}(\text{law}(Q'_{i+1}; P_i), \text{law}(Q_{i+1}; P_i)) \\ &= \mathbb{E}_{Y_{1:i} \sim P_i} \text{TV}(Q'_i(Y_{1:i}), Q_i(Y_{1:i})). \end{aligned} \quad \text{(Corollary F.1)}$$

This completes the first part of the demonstration (noting symmetry of TV). The second part follows from [Corollary F.1](#), by letting $Y \leftarrow Y_1$, and $X \leftarrow Y_{2:H+1}$ in that lemma. \square

K Stability in the Control System

This section proves our various stability conditions. Precisely, we establish the following guarantee:

Proposition K.1. Let $c_\gamma, c_\xi, \bar{c}_\beta, \bar{c}_\gamma, L_\beta$ be the constants defined in [Assumption 3.1](#). In terms of these, define

$$\begin{aligned}\alpha &= \bar{c}_\beta (4\bar{c}_\gamma \min\{c_\gamma, c_\xi/4\bar{c}_\gamma\} + c_\xi) \\ c_1 &:= 4\bar{c}_\gamma \bar{c}_\beta (2 + \alpha L_{\text{stab}} + 2R_{\text{dyn}}) \\ c_2 &:= \max\{1, c_1\}^{-1} \min\{c_\gamma, c_\xi/2\bar{c}_\gamma\} \\ c_3 &:= \frac{1}{L_\beta} \log(2e\bar{c}_\beta) \\ c_4 &:= c_\xi/2 \\ c_5 &:= 2\bar{c}_\beta\end{aligned}$$

For actions $\mathbf{a} = (\kappa_k)_{1 \leq k \leq \tau_{\text{chunk}}}$ where $\kappa_k(\mathbf{x}) = \bar{\mathbf{u}}_k + \bar{\mathbf{K}}_k(\mathbf{x} - \bar{\mathbf{x}}_k)$ are affine primitive controllers, define $\mathbf{d}_{\max}(\mathbf{a}, \mathbf{a}') := \max_{1 \leq k \leq \tau_{\text{chunk}}} (\|\bar{\mathbf{u}}_k - \bar{\mathbf{u}}'_k\| + \|\bar{\mathbf{x}}_k - \bar{\mathbf{x}}'_k\| + \|\bar{\mathbf{K}}_k - \bar{\mathbf{K}}'_k\|)$, and let

$$\mathbf{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') := c_1 \mathbf{d}_{\max}(\mathbf{a}, \mathbf{a}') \cdot \mathbf{I}_\infty \{\mathbf{d}_{\max}(\mathbf{a}, \mathbf{a}') \leq c_2\}$$

Then, if $\tau_{\text{chunk}} \geq c_3/\eta$, the policy π^* as defined in [Definition J.3](#) satisfies $(r_{\text{IPS}}, \gamma_{\text{IPS,TVc}}, \gamma_{\text{IPS,S}}, \mathbf{d}_{\text{IPS}})$ -restricted IPS ([Definition H.9](#)) with $\mathbf{d}_{\mathcal{A}}$ as above, and with

$$r_{\text{IPS}} = c_4, \quad \gamma_{\text{IPS,TVc}}(u) = c_5 u \exp(-\eta(\tau_{\text{chunk}} - \tau_{\text{obs}})/L_{\text{stab}}), \quad \gamma_{\text{IPS,S}}(u) = c_5 u.$$

[Appendix K.1](#) proves [Proposition K.1](#), based on a lemma whose proof is given in [Appendix K.2](#).

In what follows, we justify our assumption of a stabilizing synthesis oracle, [Assumption 3.1](#). First, [Appendix K.4](#) shows that if the system dynamics are *smooth*, then time-varying affine controllers whose gains stabilize the Jacobian linearization of the given system satisfy [Definition 3.1](#). This result is stated in [Appendix K.3](#), along with the requisite assumptions, and proven in [Appendix K.4](#), based on a lemma whose proof is given in [Appendix K.5](#).

Finally, [Appendix K.6](#) shows how a synthesis oracle can produce gains which stabilize the linearized dynamics can obtained by solving the Riccati equation, assuming sufficient dynamical regularity. Finally, [Appendix K.7](#) gives the solutions to various scalar recursions used in the proofs throughout.

K.1 Proof of [Proposition K.1](#)

We now translate the incremental stability guarantee about into the IPS guarantee needed by [Proposition K.1](#). The core technical ingredient is the following lemma, whose proof we defer to [Appendix K.2](#).

Lemma K.2 (Trajectory-tracking via t-ISS). Consider a given sequence $(\tilde{\mathbf{x}}_{t_h})$, and suppose that $\mathbf{a}_h = \kappa_{t_h:t_{h+1}-1}$ is local-t-ISS at $\tilde{\mathbf{x}}_{t_h}$ for each $1 \leq h \leq H$ (with parameters as in [Definition 3.1](#)). Consider consistent trajectories $(\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T})$ satisfying

$$\mathbf{u}_t = \kappa_t(\mathbf{x}_t), \quad \mathbf{u}'_{t+1} = \kappa'_t(\mathbf{x}'_t), \quad \mathbf{x}_1 = \mathbf{x}'_1, \quad \max_h \|\tilde{\mathbf{x}}_{t_h} - \mathbf{x}_{t_h}\| \leq r \leq c_\xi/2$$

Further, define the sequence $(\hat{\mathbf{x}}_t)$ by setting, for each h ,

$$\hat{\mathbf{x}}_{t_h} := \tilde{\mathbf{x}}_{t_h}, \quad \hat{\mathbf{x}}_{t_h+i} := f(\hat{\mathbf{x}}_{t_h+i}, \kappa_t(\hat{\mathbf{x}}_{t_h+i-1})), \quad i \in \{1, 2, \dots, \tau_{\text{chunk}} - 1\} \quad (\text{K.1})$$

Then, the following guarantees hold

- (a) $\|\mathbf{x}_{t_h+i} - \hat{\mathbf{x}}_{t_h+i}\| \leq \beta(r, i)$ for $i \in \{0, 1, 2, \dots, \tau_{\text{chunk}} - 1\}$ and $h \in [H]$.
- (b) Suppose that $\varepsilon > 0$ satisfies

$$\gamma^{-1}(\beta(2\gamma(\varepsilon), \tau_{\text{chunk}})) \leq \varepsilon \leq \min\{c_\gamma, \gamma^{-1}(c_\xi/4)\}$$

and that one of the following hold

$$\max_{1 \leq t \leq T} \sup_{\|\delta \mathbf{x}\| \leq \alpha(\varepsilon)} \|\kappa_t(\mathbf{x}_t + \delta \mathbf{x}) - \kappa'_t(\mathbf{x}_t + \delta \mathbf{x})\| \leq \varepsilon, \quad \alpha(\varepsilon) := 2\beta(2\gamma(\varepsilon), 0), \quad (\text{K.2})$$

$$\max_{1 \leq t \leq T} \sup_{\|\delta \mathbf{x}\| \leq \alpha(\varepsilon) + \beta(r, 0)} \|\kappa_t(\hat{\mathbf{x}}_t + \delta \mathbf{x}) - \kappa'_t(\hat{\mathbf{x}}_t + \delta \mathbf{x})\| \leq \varepsilon, \quad (\text{K.3})$$

Then for all $h \in [H]$, $i \in \{0, 1, \dots, \tau_{\text{chunk}}\}$, and $t \in [T]$,

$$\|\mathbf{u}_t - \mathbf{u}'_t\| \leq \varepsilon \leq \alpha \quad \|\mathbf{x}_{t_h+i} - \mathbf{x}'_{t_h+i}\| \leq \beta(2\gamma(\varepsilon), i) + \gamma(\varepsilon) \leq \alpha$$

As a consequence, we derive the following reduction from IPS and input-stability in the composite MDP to t-ISS.

Definition K.1 (Instantiation of the composite MDP for general primitive controllers). In this section, we summarize the instantiation of the MDP in [Appendix J](#):

- States $\mathbf{s}_h = \mathbf{s}_h$ and $\mathbf{d}_S, \mathbf{d}_{\text{TVC}}, \mathbf{d}_{\text{IPS}}$ are just as in [Appendix D](#). Moreover, $\mathbf{o}_h = \phi_o \circ \mathbf{s}_h$.
- The kernel $W_\sigma(\cdot)$ is the same as (J.1) in [Appendix J](#), applying $\mathcal{N}(0, \sigma^2 I)$ noise in the coordinates in \mathbf{o}_h .
- Actions \mathbf{a}_h are sequences of affine primitive controllers $\kappa_{1:\tau_{\text{chunk}}}$.
- $\pi^* = (\pi_h^*)$ be the policy induced by the conditional distribution of $\mathbf{a} \mid \mathbf{o}_h$ as constructed in [Definition J.3](#) in [Appendix J](#).

Lemma K.3. Instantiate the composite MDP as in [Definition K.1](#), with π^* as in [Definition J.3](#). Furthermore, suppose that under $(\rho_T, \mathbf{a}_{1:H}) \sim \mathcal{D}_{\text{exp}}$ with $\rho_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_T)$, the following both hold with probability one:

- Each action \mathbf{a}_h satisfies our notion of incremental stability ([Definition 3.1](#)) with moduli $\gamma(\cdot), \beta(\cdot, \cdot)$, constants c_γ, c_ξ
- $\mathbf{x}_t \in \mathcal{X}_0$ for some set $\mathcal{X}_0 \subset \mathbb{R}^{d_x}$, and $\kappa_t \in \mathcal{K}_0$ for some set of primitive controllers $\mathcal{K}_0 \subset \mathcal{K}$.¹³

Finally, let $\varepsilon_0 > 0$ satisfy (E.1), that is:

$$\gamma^{-1}(\beta(2\gamma(\varepsilon_0), \tau_{\text{chunk}})) \leq \varepsilon_0 \leq \min\{c_\gamma, \gamma^{-1}(c_\xi/4)\}, \quad (\text{K.4})$$

For given $\alpha > 0$, let $D_\alpha(\mathbf{a}, \mathbf{a}')$ be a function which, for all composite actions $\mathbf{a} = \kappa_{1:\tau_{\text{chunk}}}$ satisfying $\kappa_i \in \mathcal{K}_0$ all arbitrary composite actions $\mathbf{a}' = \kappa'_{1:\tau_{\text{chunk}}} \in \mathcal{K}^{\tau_{\text{chunk}}}$, satisfies

$$D_\alpha(\mathbf{a}, \mathbf{a}') \geq \sup_{\mathbf{x} \in \mathcal{X}_0} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \max_{1 \leq i \leq \tau_{\text{chunk}}} \|\kappa_i(\mathbf{x}_i + \delta \mathbf{x}) - \kappa'_i(\mathbf{x}_i + \delta \mathbf{x})\|.$$

and let

$$\bar{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}'; \alpha) := \psi(D_\alpha(\mathbf{a}, \mathbf{a}')) \cdot \mathcal{I}_\infty \{D_\alpha(\mathbf{a}, \mathbf{a}') \leq \varepsilon_0\}, \quad \psi(u) := 2\beta(2\gamma(u), 0).$$

Then, the following hold:

- (a) π^* is input-stable with respect to $\mathbf{d}_S, \mathbf{d}_{\text{TVC}}$ as defined in [Appendix D](#)

$$\mathbf{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}') = \bar{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}'; \psi(\varepsilon)),$$

- (b) For any $r_{\text{IPS}} \leq c_\xi/2$, π^* is $(r_{\text{IPS}}, \gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, S}, \mathbf{d}_{\text{IPS}})$ - restricted-IPS ([Definition H.9](#)) with

$$\mathbf{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}') = \bar{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}'; \psi(\varepsilon) + \beta(r_{\text{IPS}}, 0)), \quad \gamma_{\text{IPS}, \text{TVC}}(r) = \beta(r, \tau_{\text{chunk}} - \tau_{\text{obs}}), \quad \gamma_{\text{IPS}, S}(r) = \beta(r, 0),$$

Note that the above lemma holds for general forms of incremental stability. Let us now instantiate it for the form of incremental stability of the form established in [Proposition K.6](#).

Corollary K.1. Suppose that $\gamma(\varepsilon) = \bar{c}_\gamma \cdot \varepsilon$ and $\beta(\varepsilon, k) = \bar{c}_\gamma \phi(k) \cdot \varepsilon$. Then, as long as we take

$$2\phi(\tau_{\text{chunk}})\bar{c}_\beta \leq 1, \quad \varepsilon_0 := \min\{c_\gamma, c_\xi/4\bar{c}_\gamma\},$$

and setting $\psi(\varepsilon) = \varepsilon$, we have that

- (a) π^* is input-stable with $\mathbf{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}') = \bar{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}'; 4\bar{c}_\gamma \bar{c}_\beta \varepsilon_0)$.

- (b) For any $r_{\text{IPS}} = c_\xi$, π^* is $(r_{\text{IPS}}, \gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, S}, \mathbf{d}_{\text{IPS}})$ - restricted-IPS ([Definition H.9](#)) with

$$\mathbf{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}') = 4\bar{c}_\beta \bar{c}_\gamma \bar{d}_\mathcal{A}(\mathbf{a}, \mathbf{a}'; \bar{c}_\beta(4\bar{c}_\gamma \varepsilon_0 + r)), \quad \gamma_{\text{IPS}, \text{TVC}}(r) = \bar{c}_\beta \phi(\tau_{\text{chunk}} - \tau_{\text{obs}})r, \quad \gamma_{\text{IPS}, S}(r) = \bar{c}_\beta r$$

We are now ready to prove the main result of this appendix.

¹³This can be directly generalized to a constraint on the composite states \mathbf{s}_h and composite actions \mathbf{a}_h .

Proof of Proposition K.1. Note that, by assumption, we are in the regime of Corollary K.1, with $\phi(k) = e^{-L_\beta(k-1)}$ and $\varepsilon_0 := \min\{c_\gamma, c_\xi/4\bar{c}_\gamma\}$. We note that, under our assumption $L_\beta \leq 1$,

$$\phi(k) = e^{L_\beta} e^{-L_\beta(k-1)} \leq e \cdot e^{-L_\beta(k-1)}. \quad (\text{K.5})$$

Hence, $2\phi(\tau_{\text{chunk}})\bar{c}_\beta \leq 1$ for $\tau_{\text{chunk}} \geq c_3 = \log(2e\bar{c}_\beta)/L_\beta$.

Next, we develop D_α . Express the primitive controllers $\kappa_i = (\bar{\mathbf{u}}_i, \bar{\mathbf{x}}_i, \bar{\mathbf{K}}_i)$ and $\kappa'_i = (\bar{\mathbf{u}}'_i, \bar{\mathbf{x}}'_i, \bar{\mathbf{K}}'_i)$. Recall

$$d_{\max}(\mathbf{a}, \mathbf{a}') = \max_{1 \leq i \leq \tau_{\text{chunk}}} \max\{\|\bar{\mathbf{u}}_i - \bar{\mathbf{u}}'_i\| + \|\bar{\mathbf{x}}_i - \bar{\mathbf{x}}'_i\| + \|\bar{\mathbf{K}}_i - \bar{\mathbf{K}}'_i\|\}.$$

By assumption, the expert distribution \mathcal{D}_{exp} ensures that $\|\mathbf{x}_t\| \leq R_{\text{dyn}}$ and that $\|\bar{\mathbf{K}}_t\| \leq R_{\mathbf{K}}$. Moreover, it also ensures $\|\bar{\mathbf{x}}_t\| \leq R_{\text{dyn}}$, since under the expert distribution, $\bar{\mathbf{x}}_t = \mathbf{x}_t$. Thus, to find an upper bound on the distance $D_\alpha(\mathbf{a}, \mathbf{a}')$, it suffices to take $\mathcal{X}_0 = \{\mathbf{x} : \|\mathbf{x}\| \leq R_{\text{dyn}}\}$ and bound the following quantity for all $\mathbf{a} = \kappa_{1:\tau_{\text{chunk}}}$ and $\mathbf{a}' = \kappa'_{1:\tau_{\text{chunk}}}$ for which $\kappa_i = (\bar{\mathbf{u}}_i, \bar{\mathbf{x}}_i, \bar{\mathbf{K}}_i)$ satisfies $\|\bar{\mathbf{x}}_i\| \leq R_{\text{dyn}}$ and $\|\bar{\mathbf{K}}_i\| \leq R_{\mathbf{K}}$:

$$\begin{aligned} & \sup_{\mathbf{x}: \|\mathbf{x}\| \leq R_{\text{dyn}}} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \max_{1 \leq t \leq \tau_{\text{chunk}}} \|\kappa_t(\mathbf{x}_t + \delta \mathbf{x}) - \kappa'_t(\mathbf{x}_t + \delta \mathbf{x})\| \\ &= \sup_{\mathbf{x}: \|\mathbf{x}\| \leq R_{\text{dyn}}} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \max_{1 \leq t \leq \tau_{\text{chunk}}} \|\bar{\mathbf{u}}_t - \bar{\mathbf{u}}'_t + (\bar{\mathbf{K}}_t - \bar{\mathbf{K}}'_t)(\mathbf{x}_t + \delta \mathbf{x}) + \bar{\mathbf{K}}_t \bar{\mathbf{x}}_t - \bar{\mathbf{K}}'_t \bar{\mathbf{x}}'_t\| \\ &\leq \max_{1 \leq t \leq \tau_{\text{chunk}}} \|\bar{\mathbf{u}}_t - \bar{\mathbf{u}}'_t\| + \|\bar{\mathbf{K}}_t - \bar{\mathbf{K}}'_t\|(\alpha + \sup_{\mathbf{x}: \|\mathbf{x}\| \leq R_{\text{dyn}}} \|\mathbf{x}\|) + \|\bar{\mathbf{K}}_t(\bar{\mathbf{x}}_t - \bar{\mathbf{x}}'_t)\| + \|(\bar{\mathbf{K}}_t - \bar{\mathbf{K}}'_t)\bar{\mathbf{x}}'_t\| \\ &\quad \max_{1 \leq t \leq \tau_{\text{chunk}}} \|\bar{\mathbf{u}}_t - \bar{\mathbf{u}}'_t\| + \|\bar{\mathbf{K}}_t - \bar{\mathbf{K}}'_t\|(\alpha + R_{\text{dyn}}) + R_{\mathbf{K}}\|\bar{\mathbf{x}}_t - \bar{\mathbf{x}}'_t\| + \|\bar{\mathbf{K}}_t - \bar{\mathbf{K}}'_t\|\|\bar{\mathbf{x}}'_t\| \\ &\quad (\|\bar{\mathbf{K}}_t\| \leq R_{\mathbf{K}}) \\ &\leq \max_{1 \leq t \leq \tau_{\text{chunk}}} \|\bar{\mathbf{u}}_t - \bar{\mathbf{u}}'_t\| + \|\bar{\mathbf{K}}_t - \bar{\mathbf{K}}'_t\|(\alpha + 2R_{\text{dyn}} + \|\bar{\mathbf{x}}_t - \bar{\mathbf{x}}'_t\|) + R_{\mathbf{K}}\|\bar{\mathbf{x}}_t - \bar{\mathbf{x}}'_t\| \\ &\quad (\|\bar{\mathbf{x}}_t\| \leq R_{\text{dyn}}) \\ &\leq d_{\max}(\mathbf{a}, \mathbf{a}')(1 + R_{\mathbf{K}}\alpha + 2R_{\text{dyn}} + d_{\max}(\mathbf{a}, \mathbf{a}')); \end{aligned}$$

that is, we can take

$$D_\alpha(\mathbf{a}, \mathbf{a}') = d_{\max}(\mathbf{a}, \mathbf{a}')(1 + R_{\mathbf{K}}\alpha + 2R_{\text{dyn}} + d_{\max}(\mathbf{a}, \mathbf{a}')).$$

Now, set $\alpha = 4\bar{c}_\beta \bar{c}_\gamma \varepsilon_0 + \bar{c}_\beta r_{\text{IPS}} = \bar{c}_\beta (4\bar{c}_\gamma \min\{c_\gamma, c_\xi/4\bar{c}_\gamma\} + c_\xi)$. For $c_1 = 4\bar{c}_\gamma \bar{c}_\beta (2 + \alpha R_{\mathbf{K}} + 2R_{\text{dyn}})$ and $c_2 = \max\{1, c_1\}^{-1} \min\{c_\gamma, c_\xi/4\bar{c}_\gamma\}$. Then if $d_{\max}(\mathbf{a}, \mathbf{a}') \leq c_2$, then,

$$D_\alpha(\mathbf{a}, \mathbf{a}') \leq d_{\max}(\mathbf{a}, \mathbf{a}')(2 + R_{\mathbf{K}}\alpha + 2R_{\text{dyn}}) \leq \min\{c_\gamma, c_\xi/4\bar{c}_\gamma\}.$$

and, in particular,

$$\bar{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}' \mid \alpha) \leq 4\bar{c}_\beta \bar{c}_\gamma ((2 + R_{\mathbf{K}}\alpha + 2R_{\text{dyn}})d_{\max}(\mathbf{a}, \mathbf{a}')) = c_1 d_{\max}(\mathbf{a}, \mathbf{a}')$$

Hence, unconditionally,

$$\bar{d}_{\mathcal{A}}(\mathbf{a}, \mathbf{a}' \mid \alpha) \leq c_1 d_{\max}(\mathbf{a}, \mathbf{a}') \mathbf{I}_{\infty}\{d_{\max}(\mathbf{a}, \mathbf{a}') \leq c_2\}$$

Thus, π^* satisfies $(r_{\text{IPS}}, \gamma_{\text{IPS, TVC}}, \gamma_{\text{IPS, S}}, d_{\text{IPS}})$ -restricted-IPS (Definition H.9) with $r_{\text{IPS}} = c_\xi/2 = c_4$

$$\begin{aligned} d_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') &= c_1 d_{\max}(\mathbf{a}, \mathbf{a}') \mathbf{I}_{\infty}\{d_{\max}(\mathbf{a}, \mathbf{a}') \leq c_2\} \\ \gamma_{\text{IPS, TVC}}(r) &= \bar{c}_\beta \cdot \phi(k), \quad \gamma_{\text{IPS, S}}(r) = \bar{c}_\beta r \end{aligned}$$

Using (K.5) and recalling $c_5 = e\bar{c}_\beta$, we conclude $(r_{\text{IPS}}, \gamma_{\text{IPS, TVC}}, \gamma_{\text{IPS, S}}, d_{\text{IPS}})$ -restricted-IPS (Definition H.9) with $r_{\text{IPS}} = c_\xi/2 = c_4$ and

$$\begin{aligned} d_{\mathcal{A}}(\mathbf{a}, \mathbf{a}') &= c_1 d_{\max}(\mathbf{a}, \mathbf{a}') \mathbf{I}_{\infty}\{d_{\max}(\mathbf{a}, \mathbf{a}') \leq c_2\} \\ \gamma_{\text{IPS, TVC}}(r) &= c_5 e^{-L_\beta(\tau_{\text{chunk}} - \tau_{\text{obs}})}, \quad \gamma_{\text{IPS, S}}(r) = c_5 r. \end{aligned}$$

This concludes the proof. \square

K.1.1 Proof of Lemma K.3

Let's prove Lemma K.3(a). Let $(s_{1:H+1}, a_{1:H})$ be drawn from the distribution induces by π^* , and let $a'_{1:H}$ be some other sequences of actions. The primitive controllers and states under the instantiation of the composite MDP for $a_{1:H}, a'_{1:H}$ respectively be $\kappa_{1:T}, \kappa'_{1:T}$ and $\mathbf{x}_{1:T+1}$. Note that, by Lemma J.6(b), each \mathbf{x}_t has the same marginals as under the expert distribution \mathcal{D}_{exp} and similarly so does a_h , so by the assumption of the lemma, $\mathbf{x}_t \in \mathcal{X}_0$ and $\kappa_t \in \mathcal{K}_0$ with probability one. Thus,

$$\sup_{\mathbf{x} \in \mathcal{X}_0} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \max_{t_h \leq t \leq t_{h+1}-1} \|\kappa_t(\mathbf{x}_t + \delta \mathbf{x}) - \kappa'_t(\mathbf{x}_t + \delta \mathbf{x})\| \leq D_\alpha(a, a').$$

In particular if $\varepsilon \leq \varepsilon_0$ and $D_\alpha(a, a') \leq \varepsilon$, then

$$\sup_{\mathbf{x} \in \mathcal{X}_0} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \max_{t_h \leq t \leq t_{h+1}-1} \|\kappa_t(\mathbf{x}_t + \delta \mathbf{x}) - \kappa'_t(\mathbf{x}_t + \delta \mathbf{x})\| \leq \varepsilon \leq \varepsilon_0,$$

By Lemma K.2, and the fact that $\beta(\varepsilon, i)$ is non-increasing in i , we find $\max_h d_S(s_h, s'_h) = \max_t \{\|\mathbf{x}_t - \mathbf{x}'_t\|, \|\mathbf{u}_t - \mathbf{u}'_t\|\} \leq \psi(\varepsilon)$, as needed.

To prove Lemma K.3(b), let $(s_{1:H+1}, \tilde{s}_{1:H+1}, a_{1:H})$ be as in the definition of restricted IPS (Definition H.9), let $a'_{1:H}$ be an alternative sequence of composite actions, and unpack these into $(\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T})$, $(\tilde{\mathbf{x}}_{1:T+1}, \tilde{\mathbf{u}}_{1:T})$, $\kappa_{1:T}$ and $\kappa'_{1:T}$ as above. We let $\tilde{o}_h = (\tilde{\mathbf{x}}_{t_h - \tau_{\text{obs}} + 1:t_h}, \tilde{\mathbf{u}}_{t_h - \tau_{\text{obs}} + 1:t_h}) = \phi_o \circ \tilde{s}_h$ denote the observation-chunk associated with \tilde{s}_h . It follows from Lemma J.6(d) and the construction in (Definition H.9) that the distribution (\tilde{o}_h, a_h) under this construction is absolutely continuous w.r.t. the distribution of (o_h, a_h) under \mathcal{D}_{exp} . In particular, this implies that $a_h = \kappa_{t_h:t_{h+1}-1}$ satisfies the incremental stability condition on $\tilde{\mathbf{x}}_{t_h}$, as well as the following property: let $\hat{s}_{h+1} = F_h(\tilde{s}_h, a_h)$, which concretely are states $(\hat{\mathbf{x}}_{t_h:t_{h+1}}, \hat{\mathbf{u}}_{t_h:t_{h+1}-1})$ corresponding to the dynamics induced by rolling out $a_h = \kappa_{t_h:t_{h+1}-1}$ from $\hat{\mathbf{x}}_{t_h}$, depicted in (K.1). Then, absolute continuity of (\tilde{o}_h, a_h) with respect to its analogues under \mathcal{D}_{exp} implies that $\hat{\mathbf{x}}_{t_h:t_{h+1}}$ is absolutely continuous w.r.t. the distribution of $\mathbf{x}_{t_h:t_{h+1}}$ under \mathcal{D}_{exp} . Hence, $\hat{\mathbf{x}}_t \in \mathcal{X}_0$ for $t_h \leq t \leq t_{h+1}$. By a similarly argument, we also have $\kappa_t \in \mathcal{K}_0$ with probability one. Thus, we have

$$\sup_{\mathbf{x} \in \mathcal{X}_0} \sup_{\|\delta \mathbf{x}\| \leq \alpha} \max_{t_h \leq t \leq t_{h+1}-1} \|\kappa_t(\mathbf{x} + \delta \mathbf{x}) - \kappa'_t(\mathbf{x} + \delta \mathbf{x})\| \leq D_\alpha(a_h, a'_h),$$

Hence, whenever $D_\alpha(a_h, a'_h) \leq \varepsilon$ for $\alpha = \psi(\varepsilon) + \beta(r, 0)$, then

$$\max_{1 \leq t \leq T} \sup_{\mathbf{x} \in \mathcal{X}_0} \sup_{\|\delta \mathbf{x}\| \leq \psi(\varepsilon) + \beta(r, 0)} \|\kappa_t(\mathbf{x} + \delta \mathbf{x}) - \kappa'_t(\mathbf{x} + \delta \mathbf{x})\| \leq \varepsilon \leq \varepsilon_0, \quad (\text{K.6})$$

then we find (using $\hat{\mathbf{x}}_t \in \mathcal{X}_0$)

$$\max_{1 \leq t \leq T} \sup_{\|\delta \mathbf{x}\| \leq \psi(\varepsilon) + \beta(r, 0)} \|\kappa_t(\hat{\mathbf{x}}_t + \delta \mathbf{x}) - \kappa'_t(\hat{\mathbf{x}}_t + \delta \mathbf{x})\| \leq \varepsilon \leq \varepsilon_0$$

Thus when (K.6) is true for all h , Lemma K.2 again implies $\max_h d_S(s_h, s'_h) = \max_t \{\|\mathbf{x}_t - \mathbf{x}'_t\|, \|\mathbf{u}_t - \mathbf{u}'_t\|\} \leq \psi(\varepsilon)$ (again, using $\beta(\cdot, i)$ being non-increasing in i). This concludes the proof. \square

K.2 Proof of Lemma K.2

We begin with the following simplifying observation, which follows from considering the definition of local t-ISS with $\delta \mathbf{u}_t \equiv 0$ at time $t = 0$:

Observation K.4. $\beta(m, u) \geq u$ for any $u \in [0, c_\xi]$.

The inequality $\|\mathbf{x}_{t_h+i} - \hat{\mathbf{x}}'_{t_h,i}\| \leq \beta(r, i)$ is an immediate consequence of local-t-ISS of a_h at $\hat{\mathbf{x}}_{t_h,0}$. Note further that this means that

$$\|\mathbf{x}_{t_h+i} - \hat{\mathbf{x}}'_{t_h,i}\| \leq \beta(r, i) \leq \beta(r, 0) \leq r \leq c_\xi/2. \quad (\text{K.7})$$

Let us prove $\|\mathbf{x}_{t_h+i} - \mathbf{x}'_{t_h+i}\| \leq \beta(2\gamma(\varepsilon), i) + \gamma(\varepsilon)$. Next, define $\delta \mathbf{u}_t = \kappa'_t(\mathbf{x}'_t) - \kappa_t(\mathbf{x}'_t)$ and $\delta \mathbf{x}_t = \mathbf{x}'_t - \mathbf{x}_t$. We begin by fixing a chunk h and arguing along the lines of Pfrommer et al. [50, Proposition 3.1]. In what follows, we assume either (K.2) or (K.3), restated here for convenience:

$$\max_{1 \leq t \leq T} \sup_{\|\delta \mathbf{x}\| \leq \alpha(\varepsilon)} \|\kappa_t(\mathbf{x}_t + \delta \mathbf{x}) - \kappa'_t(\mathbf{x}_t + \delta \mathbf{x})\| \leq \varepsilon, \quad \alpha(\varepsilon) := 2\beta(2\gamma(\varepsilon), 0), \quad \text{or} \quad (\text{K.8})$$

$$\max_{1 \leq t \leq T} \sup_{\|\delta \mathbf{x}\| \leq \alpha(\varepsilon) + \beta(r, 0)} \|\kappa_t(\hat{\mathbf{x}}_t + \delta \mathbf{x}) - \kappa'_t(\hat{\mathbf{x}}_t + \delta \mathbf{x})\| \leq \varepsilon, \quad (\text{K.9})$$

Claim K.5. Fix $c_0 > 0$. Suppose that, at a given step h , $\|\delta \mathbf{x}_{t_h}\| \leq c_0 \leq c_\xi/2$, and that $\beta(c_0, 0) + \gamma(\varepsilon) \leq \alpha$. Then, for all $0 \leq i \leq \tau_{\text{chunk}} - 1$, $\|\delta \mathbf{u}_{t_h+i}\| \leq \varepsilon \leq \alpha$ and

$$\forall 0 \leq i \leq \tau_{\text{chunk}}, \quad \|\delta \mathbf{x}_{t_h+i}\| \leq \beta(c_0, i) + \gamma(\varepsilon) \leq \alpha$$

Proof. We perform induction over $t \geq t_h$. Assume inductively that $\|\delta \mathbf{x}_t\| \leq \beta(c_0, t - t_h) + \gamma(\varepsilon) \leq \alpha$ and $\max_{1 \leq s \leq t-1} \|\delta \mathbf{u}_s\| \leq \varepsilon$; note that this base case $t = t_h$ holds as $\beta(c_0, 0) \leq \alpha$ by [Observation K.4](#) and our assumption on c_0 . From the inductive hypothesis and the condition [\(K.8\)](#),

$$\|\delta \mathbf{u}_t\| = \|\kappa'_t(\mathbf{x}'_t) - \kappa_t(\mathbf{x}_t)\| \leq \max_{t_h \leq t \leq t_{h+1}-1} \sup_{\|\delta \mathbf{x}\| \leq \alpha} \|\kappa_t(\mathbf{x}_t + \delta \mathbf{x}) - \kappa'_t(\mathbf{x}_t + \delta \mathbf{x})\| \leq \varepsilon.$$

Note that by [\(K.7\)](#), [\(K.9\)](#) also suffices for the above to hold. Hence, in either case $\max_{1 \leq s \leq t} \|\delta \mathbf{u}_s\| \leq \Delta_h$. As $\|\mathbf{x}_{t_h} - \mathbf{x}'_{t_h}\| \leq c_\xi/2$, the triangle inequality and [\(K.7\)](#) imply $\|\mathbf{x}'_{t_h} - \hat{\mathbf{x}}_{t_h}\| \leq c_\xi$. This, and the fact that $\varepsilon \leq c_\gamma$, allows us to invoke our definition of incremental stability in [Definition 3.1](#), implying

$$\|\delta \mathbf{x}_{t+1}\| \leq \beta(c_0, t + 1 - t_h) + \gamma(\varepsilon),$$

as needed. \square

To conclude, we argue inductively on h that we can take $c_0 = 2\gamma(\varepsilon)$ in the above claim. First note that $2\gamma(\varepsilon) \leq c_\xi/2$ by assumption. Thus, from [Observation K.4](#), $\gamma(\varepsilon) = \frac{1}{2} \cdot 2\gamma(\varepsilon) \leq \frac{1}{2} \beta(2\gamma(\varepsilon), 0)$. Hence, for $c_0 = 2\gamma(\varepsilon)$ $\beta(c_0, 0) + \gamma(\varepsilon) \leq \frac{3}{2} \beta(2\gamma(\varepsilon), 0) \leq \alpha$. Moreover, by assumption $\delta \mathbf{x}_1 = 0$, the bound $\|\delta \mathbf{x}_{t_h}\| \leq 2\gamma(\varepsilon)$ holds trivially for step $h = 1$. Assuming it holds for h , [Claim K.5](#) yields

$$\forall 0 \leq i \leq \tau_{\text{chunk}}, \quad \|\delta \mathbf{x}_{t_h+i}\| \leq \beta(c_0, i) + \gamma(\varepsilon) \leq \beta(c_0, 0) + \gamma(\varepsilon) \leq \alpha,$$

where the final inequality follows from the computation above. Moreover, by taking $i = \tau_{\text{chunk}}$,

$$\|\delta \mathbf{x}_{t_{h+1}}\| = \|\delta \mathbf{x}_{t_h+\tau_{\text{chunk}}}\| \leq \beta(2\gamma(\varepsilon), \tau_{\text{chunk}}) + \gamma(\varepsilon) \leq 2\gamma(\varepsilon),$$

where the last inequality is by the assumption of the lemma. \square

K.3 Synthesized Linear Controllers are Incrementally Stabilizing

In this section, we give a sufficient condition for incremental stability of affine primitive controllers. Recall our notation of a length- K control trajectory is denoted $\boldsymbol{\rho} = (x_{1:K+1}, u_{1:K}) \in \mathcal{P}_K = (\mathbb{R}^{d_x})^{K+1} \times (\mathbb{R}^{d_u})^K$. Given such a trajectory, the *Jacobian linearizations* are denoted

$$\mathbf{A}_k(\boldsymbol{\rho}) := \frac{\partial}{\partial \mathbf{x}} f_\eta(\mathbf{x}_k, \mathbf{u}_k), \quad \mathbf{B}_k(\boldsymbol{\rho}) := \frac{\partial}{\partial \mathbf{u}} f_\eta(\mathbf{x}_k, \mathbf{u}_k)$$

for $k \in [K]$. Recalling our dynamics map $f(\cdot, \cdot)$, and step size $\eta > 0$, we say that $\boldsymbol{\rho}$ is *feasible* if, for all $k \in [K]$,

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \quad \text{where } f(\mathbf{x}, \mathbf{u}) = \mathbf{x} + \eta f_\eta(\mathbf{x}, \mathbf{u}).$$

We now introduce a notion of *regularity* on the dynamics, which essentially enforces boundedness and smoothness.

Definition K.2 (Trajectory Regularity). A control trajectory $\boldsymbol{\rho} = (\mathbf{x}_{1:K+1}, \mathbf{u}_{1:K})$ is $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ -regular if for all $k \in [K]$ and all $(\mathbf{x}'_k, \mathbf{u}'_k) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_u}$ such that $\|\mathbf{x}'_k - \mathbf{x}_k\| \vee \|\mathbf{u}_k - \mathbf{u}'_k\| \leq R_{\text{dyn}}$,¹⁴

$$\|\nabla f_\eta(\mathbf{x}'_k, \mathbf{u}'_k)\|_{\text{op}} \leq L_{\text{dyn}}, \quad \|\nabla^2 f_\eta(\mathbf{x}'_k, \mathbf{u}'_k)\|_{\text{op}} \leq M_{\text{dyn}}.$$

We also recall the definitions around Jacobian stabilization. We start with a definition of Jacobian stabilization for feedback gains, from which we then recover the definition of Jacobian stabilization for primitive controllers given in the body.

¹⁴Here, $\|\nabla^2 f_\eta(\mathbf{x}'_t, \mathbf{u}'_t)\|_{\text{op}}$ denotes the operator-norm of a three-tensor.

Definition K.3 (Jacobian Stability). Consider $R_{\mathbf{K}}, L_{\text{stab}}, B_{\text{stab}} \geq 1$. Consider sequence of gains $\mathbf{K}_{1:K} \in (\mathbb{R}^{d_u \times d_u})^K$ and trajectory $\boldsymbol{\rho} = (\mathbf{x}_{1:K+1}, \mathbf{u}_{1:K}) \in \mathcal{P}_K$. We say that $(\boldsymbol{\rho}, \mathbf{K}_{1:K})$ -is $(R_{\mathbf{K}}, B_{\text{stab}}, L_{\text{stab}})$ -Jacobian Stable if $\max_k \|\mathbf{K}_k\|_{\text{op}} \leq B_{\text{stab}}$, and if the closed-loop transition operators defined by

$$\Phi_{\text{cl},k,j} := (\mathbf{I} + \eta \mathbf{A}_{\text{cl},k-1}) \cdot (\mathbf{I} + \eta \mathbf{A}_{\text{cl},k-2}) \cdot (\dots) \cdot (\mathbf{I} + \eta \mathbf{A}_{\text{cl},j})$$

with $\mathbf{A}_{\text{cl},k} = \mathbf{A}_k(\boldsymbol{\rho}) + \mathbf{B}_{k-1}(\boldsymbol{\rho})\mathbf{K}_{k-1}$ satisfies the following inequality

$$\|\Phi_{\text{cl},k,j}\|_{\text{op}} \leq B_{\text{stab}} \left(1 - \frac{\eta}{L_{\text{stab}}}\right)^{k-j}.$$

The following proposition is proven in [Appendix K.4](#), establishing incremental stability of affine gains.

Proposition K.6 (Incremental Stability of Affine Primitive Controller). Suppose that $\bar{\boldsymbol{\rho}} = (\bar{\mathbf{x}}_{1:K+1}, \bar{\mathbf{u}}_{1:K})$ is $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ regular, and suppose $(\bar{\boldsymbol{\rho}}, \bar{\mathbf{K}}_{1:K})$ is $(R_{\mathbf{K}}, B_{\text{stab}}, L_{\text{stab}})$ stable. Suppose that $\eta \leq L_{\text{stab}}/2$, that $R_{\mathbf{K}} \geq 1$, define the constants

$$\begin{aligned} c_{\xi,1} &= \frac{1}{4R_{\mathbf{K}}B_{\text{stab}}} \min \left\{ 1, \frac{1}{4L_{\text{stab}}M_{\text{dyn}}R_{\mathbf{K}}B_{\text{stab}}} \right\} \\ c_{\xi,2} &= \min \left\{ \frac{1}{96B_{\text{stab}}M_{\text{dyn}}R_{\mathbf{K}}^2}, \frac{R_{\text{dyn}}}{32R_{\mathbf{K}}} \right\} \\ c_{\xi} &= \min\{c_{\xi,1}, c_{\xi,2}/2\} \\ c_{\gamma} &= \min \left\{ \frac{1}{48B_{\text{stab}}M_{\text{dyn}}R_{\mathbf{K}}^2}, \frac{R_{\text{dyn}}}{16L_{\text{stab}}R_{\mathbf{K}}} \right\} \\ \bar{c}_{\beta} &:= 16B_{\text{stab}} \\ \bar{c}_{\gamma} &:= 8L_{\text{stab}}B_{\text{stab}}L_{\text{dyn}} \end{aligned}$$

and set

$$\beta(u, k) = \bar{c}_{\beta} \left(1 - \frac{\eta}{L_{\text{stab}}}\right)^{k-1} \cdot u, \quad \gamma(u) := \bar{c}_{\gamma} \cdot u$$

Then, the controllers $\kappa_k(\mathbf{x}) = \bar{\mathbf{K}}_k(\mathbf{x} - \bar{\mathbf{x}}_k) + \bar{\mathbf{u}}_k$ are incrementally stabilizing in the sense of [Definition 3.1](#) with moduli $\gamma(\cdot)$ and $\beta(\cdot, \cdot)$ and constants c_{ξ}, c_{γ} as above.

K.4 Proof of [Proposition K.6](#) (incremental stability of affine gains)

We require the following lemma, proven in the section below.

Lemma K.7 (Stability to State Perturbation). Let $\bar{\boldsymbol{\rho}} = (\bar{\mathbf{x}}_{1:K+1}, \bar{\mathbf{u}}_{1:K}) \in \mathcal{P}_K$ be an $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ -regular and feasible path, and let $\mathbf{K}_{1:K}$ be gains such that $(\bar{\boldsymbol{\rho}}, \mathbf{K}_{1:K})$ is $(R_{\mathbf{K}}, B_{\text{stab}}, L_{\text{stab}})$ -stable. Assume that $R_{\mathbf{K}} \geq 1$, $L_{\text{stab}} \geq 2\eta$. Fix another \mathbf{x}_1 and define another trajectory $\boldsymbol{\rho}$ via

$$\mathbf{u}_k = \bar{\mathbf{u}}_k + \mathbf{K}_k(\mathbf{x}_k - \bar{\mathbf{x}}_k), \quad \mathbf{x}_{k+1} = \bar{\mathbf{x}}_k + \eta f_{\eta}(\mathbf{x}_k, \mathbf{u}_k)$$

Then, if

$$\|\mathbf{x}_1 - \bar{\mathbf{x}}_1\| \leq c_{\xi,1} := \frac{1}{4R_{\mathbf{K}}B_{\text{stab}}} \min \left\{ 1, \frac{1}{4L_{\text{stab}}M_{\text{dyn}}R_{\mathbf{K}}B_{\text{stab}}} \right\},$$

then

- $\|\mathbf{x}_{k+1} - \bar{\mathbf{x}}_{k+1}\| \leq 2B_{\text{stab}}\|\mathbf{x}_1 - \bar{\mathbf{x}}_1\|\beta_{\text{stab}}^k$.
- $(\boldsymbol{\rho}, \mathbf{K}_{1:K})$ is $(R_{\mathbf{K}}, 2B_{\text{stab}}, L_{\text{stab}})$ -stable.
- $\|\mathbf{B}_k(\boldsymbol{\rho})\| \leq L_{\text{dyn}}$, and in addition, the trajectory $\boldsymbol{\rho}$ is $(R_{\text{dyn}}/2, L_{\text{dyn}}, M_{\text{dyn}})$ -regular.

This lemma is proven in [Appendix K.5](#) just below. Now, set $\beta_{\text{stab}} = (1 - \eta/L_{\text{stab}})$, and define $\delta\mathbf{x}_k = \mathbf{x}'_k - \mathbf{x}_k$. Let $\mathbf{A}_t = \frac{\partial}{\partial x} f_{\eta}(x, u)|_{(x,u)=(\mathbf{x}_k, \mathbf{u}_k)}$

$$\mathbf{x}'_{t+1} = \mathbf{x}_k + \eta f_{\eta}(\mathbf{x}'_k, \kappa_k(\mathbf{x}'_k) + \delta\mathbf{u}_k), \quad \mathbf{x}_{t+1} = \mathbf{x}_k + \eta f_{\eta}(\mathbf{x}_k, \kappa_k(\mathbf{x}_k))$$

This means that

$$\delta \mathbf{x}_{t+1} = \underbrace{(\mathbf{I} + \eta(\mathbf{A}_k + \mathbf{B}_k \bar{\mathbf{K}}_k))}_{=\mathbf{A}_{\text{cl},k}} \delta \mathbf{x}_k + \eta \mathbf{B}_k \delta \mathbf{u}_k + \eta \text{rem}_k,$$

where $\text{rem}_k = f_\eta(\mathbf{x}'_k, \kappa_k(\mathbf{x}'_k) + \delta \mathbf{u}_k) - f_\eta(\mathbf{x}_k, \kappa_k(\mathbf{x}_k)) - (\mathbf{A}_k + \mathbf{B}_k \bar{\mathbf{K}}_k) \delta \mathbf{x}_k - \mathbf{B}_k \delta \mathbf{u}_k$. Defining $\Phi_{\text{cl},k,j} := \mathbf{A}_{\text{cl},t-1} \mathbf{A}_{\text{cl},t-2} \dots \mathbf{A}_{\text{cl},s}$ and unfolding the recursion,

$$\delta \mathbf{x}_{t+1} = \eta \sum_{s=1}^t \Phi_{\text{cl},k+1,j+1} (\mathbf{B}_s \delta \mathbf{u}_k + \text{rem}_k) + \Phi_{\text{cl},k+1,1} \delta \mathbf{x}_1$$

Define $\varepsilon_k = \|\delta \mathbf{x}_k\|$ and $\varepsilon_{\mathbf{u}} := \max_{1 \leq t \leq T} \|\delta \mathbf{u}_k\|$. Then, we have

$$\begin{aligned} \varepsilon_{t+1} &\leq \eta \sum_{j=1}^k \|\Phi_{\text{cl},k+1,j+1}\| (L_{\text{dyn}} \varepsilon_{\mathbf{u}} + \|\text{rem}_k\|) + \|\Phi_{\text{cl},k+1,1}\| \varepsilon_1 \\ &\stackrel{(i)}{\leq} 2B_{\text{stab}} \left(\eta \sum_{j=1}^k \beta_{\text{stab}}^{k-j} (L_{\text{dyn}} \varepsilon_{\mathbf{u}} + \|\text{rem}_k\|) + \beta_{\text{stab}}^k \varepsilon_1 \right) \\ &\stackrel{(ii)}{\leq} 2B_{\text{stab}} \left(\eta \sum_{j=1}^k \beta_{\text{stab}}^{k-j} (L_{\text{dyn}} \varepsilon_{\mathbf{u}} + M_{\text{dyn}}((1 + 2R_{\mathbf{K}}^2) \varepsilon_k^2 + 2\varepsilon_{\mathbf{u}}^2)) + \beta_{\text{stab}}^t \varepsilon_1 \right) \\ &\stackrel{(iii)}{\leq} 2B_{\text{stab}} \left(\eta \sum_{j=1}^k \beta_{\text{stab}}^{k-j} (2L_{\text{dyn}} \varepsilon_{\mathbf{u}} + M_{\text{dyn}}(1 + 2R_{\mathbf{K}}^2) \varepsilon_k^2) + \beta_{\text{stab}}^t \varepsilon_1 \right) \end{aligned}$$

where we in (i) $\|\Phi_{\text{cl},k,j}\| \leq 2B_{\text{stab}} \beta_{\text{stab}}^{t-s}$, and (ii) follows by [Claim K.8](#), stated and proven below, and the following inductive hypothesis

$$\max_{1 \leq j \leq k} \varepsilon_k \leq C' = \frac{R_{\text{dyn}}}{4R_{\mathbf{K}}}, \quad (\text{Inductive Hypothesis})$$

and (ii) uses the assumption $\varepsilon_{\mathbf{u}} \leq \frac{L_{\text{dyn}}}{2M_{\text{dyn}}}$. Setting $\Delta_1 = 2B_{\text{stab}} \varepsilon_1$, $\Delta_2 = 4B_{\text{stab}} L_{\text{dyn}} \varepsilon_{\mathbf{u}}$ and $C = 2B_{\text{stab}} M_{\text{dyn}} (1 + 2R_{\mathbf{K}}^2) \leq 6B_{\text{stab}} M_{\text{dyn}} R_{\mathbf{K}}^2$, [Lemma K.17](#) implies

$$\varepsilon_k \leq 4\Delta_1 \beta_{\text{stab}}^{k-1} + 2L_{\text{stab}} \Delta_2 = \underbrace{8B_{\text{stab}} \varepsilon_1 \beta_{\text{stab}}^{k-1}}_{\beta(\varepsilon_1, k)} + \underbrace{8L_{\text{stab}} B_{\text{stab}} L_{\text{dyn}} \varepsilon_{\mathbf{u}}}_{\gamma(\varepsilon_{\mathbf{u}})}$$

provided that that $\Delta_2 \leq \min \left\{ \frac{1}{8CL}, \frac{C'}{4L} \right\}$ and $\Delta_1 \leq \min \left\{ \frac{1}{16CL}, \frac{C'}{8} \right\}$ for $L = L_{\text{stab}}$. Substituting in relevant quantities and keeping the shorthand $L = L_{\text{stab}}$, it suffices that

$$\begin{aligned} \min \left\{ \frac{1}{16CL}, \frac{C'}{8} \right\} &\geq \min \left\{ \frac{1}{96B_{\text{stab}} M_{\text{dyn}} R_{\mathbf{K}}^2}, \frac{R_{\text{dyn}}}{32R_{\mathbf{K}}} \right\} \geq \varepsilon_1 \\ &\quad \underbrace{\hspace{10em}}_{=c_{\xi,2}/2} \\ \min \left\{ \frac{1}{8CL}, \frac{C'}{4L} \right\} &\geq \min \left\{ \frac{1}{48B_{\text{stab}} M_{\text{dyn}} R_{\mathbf{K}}^2}, \frac{R_{\text{dyn}}}{16L_{\text{stab}} R_{\mathbf{K}}} \right\} \geq \varepsilon_{\mathbf{u}}. \\ &\quad \underbrace{\hspace{10em}}_{=c_{\gamma}} \end{aligned}$$

□

Claim K.8. Suppose that $\varepsilon_{\mathbf{u}} \leq \frac{R_{\text{dyn}}}{4}$ and $\varepsilon_k \leq \frac{R_{\text{dyn}}}{4R_{\mathbf{K}}}$. Then, $\|\text{rem}_k\| \leq M_{\text{dyn}}((1 + 2R_{\mathbf{K}}^2) \varepsilon_k^2 + 2\varepsilon_{\mathbf{u}}^2)$.

Proof. Define $\mathbf{u}_k = \kappa_k(\mathbf{x}_k)$ and $\delta \mathbf{u}'_k = \kappa_k(\mathbf{x}'_k) \delta \mathbf{u}_k - \mathbf{u}_k$. We have that $\delta \mathbf{u}'_k = \delta \mathbf{u}_k + \kappa_k(\mathbf{x}'_k) - \kappa_k(\mathbf{x}_k) = \delta \mathbf{u}_k + \bar{\mathbf{K}}_k(\mathbf{x}'_k - \mathbf{x}_k)$. We bound $\|\delta \mathbf{u}'_k\| \leq \|\delta \mathbf{u}_k\| + \|\bar{\mathbf{K}}_k(\mathbf{x}'_k - \mathbf{x}_k)\| \leq \|\delta \mathbf{u}_k\| + R_{\mathbf{K}} \|\delta \mathbf{x}_k\|$, where we recall $\|\bar{\mathbf{K}}_k\| \leq R_{\mathbf{K}}$ and $\delta \mathbf{x}_k = \mathbf{x}'_k - \mathbf{x}$. By assumption and definition, $\|\delta \mathbf{u}_k\| \leq \varepsilon_{\mathbf{u}}$ and by definition of ε_k we conclude that

$$\|\delta \mathbf{x}_k\| \leq \varepsilon_k, \quad \|\delta \mathbf{u}'_k\| \leq (1 + R_{\mathbf{K}}) \varepsilon_k + \varepsilon_{\mathbf{u}} \quad (\text{K.10})$$

Consider a curve $\mathbf{x}_k(s) = \mathbf{x}_k + s\delta\mathbf{x}_k$ and $\mathbf{u}_k(s) = \delta\mathbf{u}'_k + \mathbf{u}_k$. With these definition

$$\begin{aligned} \text{rem}_k &= f_\eta(\mathbf{x}_k(1), \mathbf{u}_k(1)) - f_\eta(\mathbf{x}_k(0), \mathbf{u}_k(0)) - (\mathbf{A}_k + \mathbf{B}_k \bar{\mathbf{K}}_k) \delta\mathbf{x}_k - \mathbf{B}_k \delta\mathbf{u}_k \\ &= f_\eta(\mathbf{x}_k(1), \mathbf{u}_k(1)) - f_\eta(\mathbf{x}_k(0), \mathbf{u}_k(0)) - \mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k \delta\mathbf{u}'_k \\ &= \underbrace{\frac{\partial}{\partial s}(f_\eta(\mathbf{x}_k(s), \mathbf{u}_k(s))) - \mathbf{A}_k \delta\mathbf{x}_k + \mathbf{B}_k \delta\mathbf{u}'_k}_{=0} \\ &\quad + \int_0^s (1-s)^2 \frac{\partial^2}{\partial s^2}(f_\eta(\mathbf{x}_k(s), \mathbf{u}_k(s)))^\top \Big|_{s=0} (\delta\mathbf{x}_k, \delta\mathbf{u}'_k) ds \end{aligned}$$

Thus,

$$\begin{aligned} \|\text{rem}_k\| &\leq \frac{1}{2} \sup_{s \in [0,1]} \left\| \frac{\partial^2}{\partial s^2}(f_\eta(\mathbf{x}_k(s), \mathbf{u}_k(s))) \right\| \leq M_{\text{dyn}} \sup_s \left\| \frac{\partial}{\partial s}(\mathbf{x}_k(s), \mathbf{u}_k(s)) \right\|^2 \\ &\stackrel{(i)}{\leq} M_{\text{dyn}} (\|\delta\mathbf{x}_k\|^2 + \|\delta\mathbf{u}'_k\|^2) \\ &\leq M_{\text{dyn}} ((1 + 2R_{\mathbf{K}}^2)\varepsilon_k^2 + 2\varepsilon_{\mathbf{u}}^2), \end{aligned} \tag{K.10} \text{ and Am-GM}$$

To justify inequality (i), we observe that $\boldsymbol{\rho} = (\mathbf{x}_{1:K+1}, \mathbf{u}_{1:K})$ is $(R_{\text{dyn}}/2, L_{\text{dyn}}, M_{\text{dyn}})$ regular. Note that $\sup_s \|\mathbf{x}_k(s) - \mathbf{x}_k\| = \|\delta\mathbf{x}_k\|$ and $\sup_s \|\mathbf{u}_k(s) - \mathbf{u}_k\| = \|\delta\mathbf{u}'_k\|$. Hence, by the definition of trajectory regularity (Definition K.2), (i) holds as long as we check that $\|\delta\mathbf{x}_k\| \vee \|\delta\mathbf{u}'_k\| \leq R_{\text{dyn}}/4$. As $\|\delta\mathbf{x}_k\| \vee \|\delta\mathbf{u}_k\| \leq \max\{R_{\mathbf{K}}\varepsilon_k + \varepsilon_{\mathbf{u}}, \varepsilon_k\}$ and as we take $R_{\mathbf{K}} \geq 1$, it suffices that $\varepsilon_{\mathbf{u}} \leq \frac{R_{\text{dyn}}}{4}$ and $\varepsilon_k \leq \frac{R_{\text{dyn}}}{4R_{\mathbf{K}}}$, which is ensured by the claim. \square

K.5 Proof of Lemma K.7 (state perturbation)

Define $\bar{\Delta}_{\mathbf{x},k} = \mathbf{x}_k - \bar{\mathbf{x}}_k$. Then

$$\begin{aligned} \bar{\Delta}_{\mathbf{x},k+1} &= \bar{\Delta}_{\mathbf{x},k} + \eta (f_\eta(\mathbf{x}_k, \bar{\mathbf{u}}_k + \bar{\mathbf{K}}_k(\mathbf{x}_k - \bar{\mathbf{x}}_k)) - f_\eta(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k)) \\ &= \bar{\Delta}_{\mathbf{x},k} + \eta(\mathbf{A}_k(\bar{\boldsymbol{\rho}}) + \mathbf{B}_k(\bar{\boldsymbol{\rho}})\mathbf{K}_k)\bar{\Delta}_{\mathbf{x},k} + \text{rem}_k, \end{aligned} \tag{K.11}$$

where

$$\text{rem}_k = f_\eta(\mathbf{x}_k, \bar{\mathbf{u}}_k + \mathbf{K}_k(\mathbf{x}_k - \bar{\mathbf{x}}_k)) - f_\eta(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k) - (\mathbf{A}_k(\bar{\boldsymbol{\rho}}) + \mathbf{B}_k(\bar{\boldsymbol{\rho}})\mathbf{K}_k)\bar{\Delta}_{\mathbf{x},k}.$$

Claim K.9. Take $R_{\mathbf{K}} \geq 1$, and suppose $\|\bar{\Delta}_{\mathbf{x},k}\| \leq R_{\text{dyn}}/2R_{\mathbf{K}}$. Then,

$$\|\bar{\mathbf{x}}_k - \mathbf{x}_k\| \vee \|\bar{\mathbf{u}}_k - \mathbf{u}_k\| \leq R_{\text{dyn}}/2, \tag{K.12}$$

and $\|\text{rem}_k\| \leq M_{\text{dyn}} R_{\mathbf{K}}^2 \|\bar{\Delta}_{\mathbf{x},k}\|^2$.

Proof. Let $\mathbf{u}_k = \bar{\mathbf{u}}_k + \mathbf{K}_k(\mathbf{x}_k - \bar{\mathbf{x}}_k)$. The conditions of the claim imply $\|\mathbf{u}_k - \bar{\mathbf{u}}_k\| \vee \|\mathbf{x}_k - \bar{\mathbf{x}}_k\| \leq R_{\text{dyn}}/2$. From Taylor's theorem and the fact that $\bar{\boldsymbol{\rho}}$ is $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ -regular imply that

$$\begin{aligned} \|f_\eta(\mathbf{x}_k, \mathbf{u}_k) - f_\eta(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k)\| &\leq \frac{1}{2} M_{\text{dyn}} (\|\mathbf{x}_k - \bar{\mathbf{x}}_k\|^2 + \|\mathbf{u}_k - \bar{\mathbf{u}}_k\|^2) \\ &\leq \frac{1}{2} (1 + R_{\mathbf{K}}^2) M_{\text{dyn}} \|\mathbf{x}_k - \bar{\mathbf{x}}_k\|^2 \leq R_{\mathbf{K}}^2 M_{\text{dyn}} \|\bar{\Delta}_{\mathbf{x},k}\|^2, \end{aligned}$$

where again use $R_{\mathbf{K}} \geq 1$ above. \square

Solving the recursion from (K.11), we have

$$\bar{\Delta}_{\mathbf{x},k+1} = \eta \sum_{j=1}^k \boldsymbol{\Phi}_{\text{cl},k+1,j+1} \text{rem}_j + \boldsymbol{\Phi}_{\text{cl},k+1,1} \bar{\Delta}_{\mathbf{x},1}.$$

Set $\beta_{\text{stab}} := (1 - \frac{\eta}{L_{\text{stab}}})$, so that $M := \frac{\eta}{\beta_{\text{stab}}^{-1} - 1} = L_{\text{stab}}$. By assumption, $\|\Phi_{\text{cl},k,j}\| \leq B_{\text{stab}}\beta_{\text{stab}}^{k-j}$, so using [Claim K.9](#) implies that, if $\max_{j \in [k]} \|\bar{\Delta}_{\mathbf{x},j}\| \leq R_{\text{dyn}}/2R_{\mathbf{K}}$ for all $j \in [k]$,

$$\|\bar{\Delta}_{\mathbf{x},k+1}\| \leq \eta \sum_{j=1}^k B_{\text{stab}} M_{\text{dyn}} R_{\mathbf{K}}^2 \beta_{\text{stab}}^{k-j} \|\bar{\Delta}_{\mathbf{x},j}\|^2 + B_{\text{stab}} \beta_{\text{stab}}^k \|\bar{\Delta}_{\mathbf{x},1}\|.$$

Applying [Lemma K.15](#) with $\alpha = 0$, $C_1 = B_{\text{stab}} M_{\text{dyn}} R_{\mathbf{K}}^2$, and $C_2 = B_{\text{stab}} \geq 1$ and $M = L_{\text{stab}}$ (noting $\beta_{\text{stab}} \geq 1/2$), it holds that for $\|\bar{\Delta}_{\mathbf{x},1}\| = \varepsilon_1 \leq 1/4MC_1C_3 = 1/4L_{\text{stab}}M_{\text{dyn}}R_{\mathbf{K}}^2B_{\text{stab}}^2$,

$$\|\bar{\Delta}_{\mathbf{x},k+1}\| \leq 2B_{\text{stab}}\|\bar{\Delta}_{\mathbf{x},1}\|(1 - \frac{\eta}{L_{\text{stab}}})^k.$$

To ensure the inductive hypothesis that $\max_{j \in [k]} \|\bar{\Delta}_{\mathbf{x},j}\| \leq R_{\text{dyn}}R_{\mathbf{K}}$, it suffices to ensure that $2B_{\text{stab}}\|\bar{\Delta}_{\mathbf{x},1}\| \leq R_{\text{dyn}}/2R_{\mathbf{K}}$, which is assumed by the lemma. Thus, we have shown that, if

$$\|\bar{\Delta}_{\mathbf{x},1}\| \leq \min \left\{ \frac{R_{\text{dyn}}}{2R_{\mathbf{K}}B_{\text{stab}}}, \frac{1}{8L_{\text{stab}}M_{\text{dyn}}R_{\mathbf{K}}^2B_{\text{stab}}^2} \right\},$$

it holds that $\|\bar{\Delta}_{\mathbf{x},k+1}\| \leq 2B_{\text{stab}}\|\bar{\Delta}_{\mathbf{x},1}\|(1 - \frac{\eta}{L_{\text{stab}}})^k \leq R_0$ for all k .

Next, we address the stability of the gains for the perturbed trajectory ρ . Using $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ -regularity of $\bar{\rho}$ and [\(K.12\)](#),

$$\begin{aligned} & \|\mathbf{A}_k(\rho) + \mathbf{B}_k(\rho)\mathbf{K}_k - \mathbf{A}_k(\bar{\rho}) + \mathbf{B}_k(\bar{\rho})\mathbf{K}_k\| \\ &= \left\| \begin{bmatrix} \mathbf{A}_k(\rho) - \mathbf{A}_k(\bar{\rho}) & \hat{\mathbf{B}}_k(\rho) - \mathbf{B}_k(\bar{\rho}) \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \mathbf{K}_k \end{bmatrix} \right\| \\ &= \left\| (\nabla f_{\eta}(\hat{\mathbf{x}}_k, \mathbf{u}_k) - \nabla f_{\eta}(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k)) \begin{bmatrix} \mathbf{I} \\ \mathbf{K}_k \end{bmatrix} \right\| \\ &\leq M_{\text{dyn}} \|(\mathbf{x}_k - \bar{\mathbf{x}}_k, \mathbf{K}_k(\mathbf{x}_k - \bar{\mathbf{x}}_k))\| \left\| \begin{bmatrix} \mathbf{I} \\ \mathbf{K}_k \end{bmatrix} \right\| \\ &= M_{\text{dyn}} \|\mathbf{x}_k - \bar{\mathbf{x}}_k\| \left\| \begin{bmatrix} \mathbf{I} \\ \mathbf{K}_k \end{bmatrix} \right\|^2 \leq M_{\text{dyn}} \|\mathbf{x}_k - \bar{\mathbf{x}}_k\| (1 + \|\mathbf{K}_k\|_{\text{op}}^2) \\ &= M_{\text{dyn}} \|\mathbf{x}_k - \bar{\mathbf{x}}_k\| \left\| \begin{bmatrix} \mathbf{I} \\ \mathbf{K}_k \end{bmatrix} \right\|^2 \leq M_{\text{dyn}} \|\mathbf{x}_k - \bar{\mathbf{x}}_k\| (1 + \|\mathbf{K}_k\|_{\text{op}}^2) \\ &\leq 2R_{\mathbf{K}}^2 M_{\text{dyn}} \|\mathbf{x}_k - \bar{\mathbf{x}}_k\| \\ &\leq 4B_{\text{stab}} R_{\mathbf{K}}^2 M_{\text{dyn}} \|\mathbf{x}_1 - \bar{\mathbf{x}}_1\| \beta_{\text{stab}}^{k-1}, \quad \beta_{\text{stab}} = (1 - \frac{\eta}{L_{\text{stab}}}). \end{aligned}$$

Invoking [Lemma K.18](#) with $\beta_{\text{stab}} \geq 1/2$, $\|\hat{\Phi}_{\text{cl},k,j}\| \leq 2B_{\text{stab}}\beta_{\text{stab}}^{k-j}$ for all j, k provided that $4B_{\text{stab}}R_{\mathbf{K}}^2M_{\text{dyn}}\|\mathbf{x}_1 - \bar{\mathbf{x}}_1\| \leq 1/4B_{\text{stab}}L_{\text{stab}}$, which requires $\|\mathbf{x}_1 - \bar{\mathbf{x}}_1\| \leq 1/16B_{\text{stab}}^2R_{\mathbf{K}}^2L_{\text{stab}}M_{\text{dyn}}$.

The last part of the lemma uses $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ -regularity of $\bar{\rho}$ and [\(K.12\)](#).

K.6 Ricatti synthesis of stabilizing gains.

In this section, we show that under a certain *stabilizability* condition, it is always possible to synthesize primitive controllers satisfying Jacobian stability, [Definition K.3](#), with reasonable constants. We begin by defining our notion of stabilizability; we adopt the formulation based on Jacobian linearizations of non-linear systems the discrete analogue of the senses proposed in which is consistent with [\[51, 76\]](#).

Definition K.4 (Stabilizability). A control trajectory $\rho = (\mathbf{x}_{1:K+1}, \mathbf{u}_{1:K}) \in \mathcal{P}_K$ is $L_{\mathcal{S}/\mathcal{O}}$ -Jacobian-Stabilizable if $\max_k \mathcal{V}_k(\rho) \leq L_{\mathcal{S}/\mathcal{O}}$, where for $k \in [K+1]$, $\mathcal{V}_k(\rho)$ is defined by

$$\begin{aligned} \mathcal{V}_k(\rho) &:= \sup_{\xi: \|\xi\| \leq 1} \left(\inf_{\tilde{\mathbf{u}}_{1:s}} \|\tilde{\mathbf{x}}_{K+1}\|^2 + \eta \sum_{j=k}^K \|\tilde{\mathbf{x}}_j\|^2 + \|\tilde{\mathbf{u}}_j\|^2 \right) \\ &\text{s.t. } \tilde{\mathbf{x}}_k = \xi, \quad \tilde{\mathbf{x}}_{j+1} = \tilde{\mathbf{x}}_j + \eta(\mathbf{A}_j(\rho)\tilde{\mathbf{x}}_j + \mathbf{B}_j(\rho)\tilde{\mathbf{u}}_j), \end{aligned}$$

Here, for simplicity, we use Euclidean-norm costs, though any Mahalanobis-norm cost induced by a positive definite matrix would suffice. We propose to synthesize gain matrices by performing a standard Ricatti update, normalized appropriately to take account of the step size $\eta > 0$ (see, e.g. Appendix F in [51]).

Definition K.5 (Ricatti update). Given a path $\rho \in \mathcal{P}_k$ with $\mathbf{A}_k = \mathbf{A}_k(\rho)$, $\mathbf{B}_k = \mathbf{B}_k(\rho)$ we define

$$\begin{aligned}\mathbf{P}_{K+1}^{\text{ric}}(\rho) &= \mathbf{I}, \quad \mathbf{P}_k^{\text{ric}}(\rho) = (\mathbf{I} + \eta \mathbf{A}_{\text{cl},k}(\rho))^\top \mathbf{P}_{k+1}^{\text{ric}}(\rho) (\mathbf{I} + \eta \mathbf{A}_{\text{cl},k}(\rho)) + \eta (\mathbf{I} + \mathbf{K}_k(\rho) \mathbf{K}_k(\rho)^\top) \\ \mathbf{K}_k^{\text{ric}}(\rho) &= (\mathbf{I} + \eta \mathbf{B}_k^\top \mathbf{P}_{k+1}^{\text{ric}}(\rho) \mathbf{B}_k)^{-1} (\mathbf{B}_k^\top \mathbf{P}_{k+1}^{\text{ric}}(\rho)) (\mathbf{I} + \eta \mathbf{A}_k) \\ \mathbf{A}_{\text{cl},k}^{\text{ric}}(\rho) &= \mathbf{A}_k + \mathbf{B}_k \mathbf{K}_k(\rho).\end{aligned}$$

The main result of this section is that the parameters $(R_{\mathbf{K}}, B_{\text{stab}}, L_{\text{stab}})$ in Definition K.3 can be bounded in terms of L_{dyn} in Definition K.2, and the bound $L_{\mathcal{S}/\mathcal{O}}$ defined above.

Proposition K.10 (Instantiating the Lyapunov Lemma). Let $L_{\text{dyn}}, L_{\mathcal{S}/\mathcal{O}} \geq 1$, and let $\rho = (\mathbf{x}_{1:K+1}, \mathbf{u}_{1:K})$ be $(R_{\text{dyn}}, L_{\text{dyn}}, M_{\text{dyn}})$ -regular and $L_{\mathcal{S}/\mathcal{O}}$ -Jacobian Stabilizable. Suppose further that $\eta \leq 1/5 L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}$. Then, $(\rho, \mathbf{K}_{1:K}^{\text{ric}})$ -is $(R_{\mathbf{K}}, B_{\text{stab}}, L_{\text{stab}})$ -Jacobian Stable, where

$$R_{\mathbf{K}} = \frac{4}{3} L_{\mathcal{S}/\mathcal{O}} L_{\text{dyn}}, \quad B_{\text{stab}} = \sqrt{5} L_{\text{dyn}} L_{\mathcal{S}/\mathcal{O}}, \quad L_{\text{stab}} = 2 L_{\mathcal{S}/\mathcal{O}}$$

Proposition K.10 is proven in Appendix K.6.1 below. A consequence of the above proposition is that, given access to a smooth local model of dynamics, one can implement the synthesis oracle by computing linearizations around demonstrated trajectories, and solving the corresponding Ricatti equations as per the above discussions to synthesize the correct gains.

K.6.1 Proof of Proposition K.10 (Ricatti synthesis of gains)

Throughout, we use the shorthand $\mathbf{A}_k = \mathbf{A}_k(\rho)$ and $\mathbf{B}_k = \mathbf{B}_k(\rho)$, recall that $\|\cdot\|$ denotes the operator norm when applied to matrices. We also recall our assumptions that $L_{\text{dyn}}, L_{\mathcal{S}/\mathcal{O}} \geq 1$. We begin by translating our stabilizability assumption (Definition K.4) into the the \mathbf{P} -matrices in Definition K.5. The following statement recalls Lemma F.1 in [51], an instantiation of well-known solutions to linear-quadratic dynamic programming (e.g. [7]).

Lemma K.11 (Equivalence of stabilizability and Ricatti matrices). Consider a trajectory $(\mathbf{x}_{1:K}, \mathbf{u}_{1:K})$, and define the parameter $\Theta := (\mathbf{A}_{\text{jac}}(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k), \mathbf{B}_{\text{jac}}(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k))_{k \in [K]}$. Then, for all $k \in [K]$,

$$\forall k \in [K], \quad \mathcal{V}_k(\rho) = \|\mathbf{P}_k(\Theta)\|_{\text{op}}$$

Hence, if ρ is $L_{\mathcal{S}/\mathcal{O}}$ -stabilizable,

$$\max_{k \in [K+1]} \|\mathbf{P}_k(\Theta)\|_{\text{op}} \leq L_{\mathcal{S}/\mathcal{O}}.$$

Lemma K.12 (Lyapunov Lemma, Lemma F.10 in [51]). Let $\mathbf{X}_{1:K}, \mathbf{Y}_{1:K}$ be matrices of appropriate dimension, and let $\mathbf{Q} \succeq \mathbf{I}$ and $\mathbf{Y}_k \succeq 0$. Define $\Lambda_{1:K+1}$ as the solution of the recursion

$$\Lambda_{K+1} = \mathbf{Q}, \quad \Lambda_k = \mathbf{X}_k^\top \Lambda_{k+1} \mathbf{X}_k + \eta \mathbf{Q} + \mathbf{Y}_k$$

Define the operator $\Phi_{j+1,k} = \mathbf{X}_j \cdot \mathbf{X}_{j-1} \cdot \dots \cdot \mathbf{X}_k$, with the convention $\Phi_{k,k} = \mathbf{I}$. Then, if $\max_k \|\mathbf{I} - \mathbf{X}_k\|_{\text{op}} \leq \kappa \eta$ for some $\kappa \leq 1/2\eta$,

$$\|\Phi_{j,k}\|^2 \leq \max\{1, 2\kappa\} \max_{k \in [K+1]} \|\Lambda_k\| (1 - \eta \alpha)^{j-k}, \quad \alpha := \frac{1}{\max_{k \in [K+1]} \|\Lambda_{1:K+1}\|}.$$

Claim K.13. If ρ is $(0, L_{\text{dyn}}, \infty)$ -regular, then for all k , $\mathbf{A}_k = \mathbf{A}_k(\rho)$ and $\mathbf{B}_k = \mathbf{B}_k(\rho)$ satisfy $\max_{k \in [K]} \max\{\|\mathbf{A}_k\|, \|\mathbf{B}_k\|\} \leq L_{\text{dyn}}$.

Proof. For any $k \in [K]$,

$$\max\{\|\mathbf{A}_k\|, \|\mathbf{B}_k\|\} = \max \left\{ \left\| \frac{\partial}{\partial x} f(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k) \right\|, \left\| \frac{\partial}{\partial u} f(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k) \right\| \right\} \leq \|\nabla f(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k)\| \leq L_{\text{dyn}},$$

where the last inequality follows by regularity. \square

Claim K.14. Recall $\mathbf{K}_k^{\text{ric}}(\boldsymbol{\rho}) = (\mathbf{I} + \eta \mathbf{B}_k^\top \mathbf{P}_{k+1}^{\text{ric}}(\boldsymbol{\rho}) \mathbf{B}_k)^{-1} (\mathbf{B}_k^\top \mathbf{P}_{k+1}^{\text{ric}}(\boldsymbol{\rho})) (\mathbf{I} + \eta \mathbf{A}_k)$. Then, if $\boldsymbol{\rho}$ is $L_{\mathcal{S}/\mathcal{O}}$ -stabilizable and $(0, L_{\text{dyn}}, \infty)$ -regular, and if $\eta \leq 1/3L_{\text{dyn}}$,

$$\|\mathbf{K}_k^{\text{ric}}(\boldsymbol{\rho})\| \leq \frac{4}{3} L_{\mathcal{S}/\mathcal{O}} L_{\text{dyn}}$$

Proof. We bound

$$\begin{aligned} \|\mathbf{K}_k^{\text{ric}}(\boldsymbol{\rho})\| &\leq \|\mathbf{B}_k\| \|\mathbf{P}_{k+1}^{\text{ric}}(\boldsymbol{\rho})\| (1 + \eta \|\mathbf{A}_k\|) \\ &\leq L_{\text{dyn}} (1 + \eta L_{\text{dyn}}) \|\mathbf{P}_{k+1}^{\text{ric}}(\boldsymbol{\rho})\| && \text{(Claim K.13)} \\ &\leq L_{\mathcal{S}/\mathcal{O}} L_{\text{dyn}} (1 + \eta L_{\text{dyn}}) && \text{(Lemma K.11, } L_{\mathcal{S}/\mathcal{O}} \geq 1) \\ &\leq \frac{4}{3} L_{\mathcal{S}/\mathcal{O}} L_{\text{dyn}} && (\eta \leq 1/3L_{\text{dyn}}) \end{aligned}$$

□

Proof of Proposition K.10. We want to show that $\mathbf{K}_{1:K}^{\text{ric}}(\boldsymbol{\rho})$ is $(R_{\mathbf{K}}, B_{\text{stab}}, L_{\text{stab}})$ -stabilizing. Claim K.14 has already established that $\max_{k \in [K]} \|\mathbf{K}_k^{\text{ric}}(\boldsymbol{\rho})\| \leq R_{\mathbf{K}} = \frac{4}{3} L_{\mathcal{S}/\mathcal{O}} L_{\text{dyn}}$.

To prove the other conditions, we apply Lemma K.12 with $\mathbf{Y}_k = \mathbf{K}_k(\boldsymbol{\Theta}) \mathbf{K}_k(\boldsymbol{\Theta})$, $\mathbf{Q} = \mathbf{I}$, and $\mathbf{X}_k = \mathbf{I} + \eta \mathbf{A}_{\text{cl},k}(\boldsymbol{\Theta})$. From Definition K.5, let have that the term Λ_k in Lemma K.12 is precise equal to $\mathbf{P}_k(\boldsymbol{\Theta})$. From Lemma K.11,

$$\max_{k \in [K+1]} \|\mathbf{P}_k(\boldsymbol{\Theta})\|_{\text{op}} = \max_{k \in [K+1]} \mathcal{V}_k(\boldsymbol{\rho}) \leq L_{\mathcal{S}/\mathcal{O}}.$$

This implies that if $\max_k \|\mathbf{X}_k - \mathbf{I}\| \leq \kappa \eta \leq 1/2$, we have

$$\|\Phi_{\text{cl},j,k}(\boldsymbol{\Theta})\|^2 = \|(\mathbf{X}_j \cdot \mathbf{X}_{j-1} \cdots \mathbf{X}_k)\| \leq \max\{1, 2\kappa\} L_{\mathcal{S}/\mathcal{O}} \left(1 - \frac{\eta}{L_{\mathcal{S}/\mathcal{O}}}\right)^{j-k}.$$

It suffices to find an appropriate upper bound κ . We have

$$\begin{aligned} \|\mathbf{X}_k - \mathbf{I}\| &= \|\eta \mathbf{A}_{\text{cl},k}(\boldsymbol{\Theta})\| \leq \eta (\|\mathbf{A}_k\| + \|\mathbf{B}_k\| \|\mathbf{K}_k(\boldsymbol{\Theta})\|) \\ &\leq \eta L_{\text{dyn}} (1 + \|\mathbf{K}_k(\boldsymbol{\Theta})\|) \\ &\leq \eta L_{\text{dyn}} (1 + \frac{4}{3} L_{\text{dyn}} L_{\mathcal{S}/\mathcal{O}}) && \text{(Claim K.14)} \\ &\leq \frac{7}{3} \eta L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}} && (L_{\mathcal{S}/\mathcal{O}}, L_{\text{dyn}} \geq 1) \end{aligned}$$

Setting $\kappa = \frac{7}{3} L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}$, we have that as $\eta \leq \frac{1}{5L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}} \leq \min\{\frac{3}{14L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}}, \frac{1}{3L_{\text{dyn}}}\}$ (recall $L_{\text{dyn}}, L_{\mathcal{S}/\mathcal{O}} \geq 1$), we can bound

$$\max\{1, 2\kappa\} \leq \max\left\{1, \frac{14}{3} L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}\right\} \leq \max\{1, 5L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}\} = 5L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}},$$

where again recall $L_{\mathcal{S}/\mathcal{O}}, L_{\text{dyn}} \geq 1$. In sum, for $\eta \leq \frac{1}{5L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}}$, we have

$$\|\Phi_{\text{cl},j,k}\|^2 \leq 5L_{\text{dyn}}^2 L_{\mathcal{S}/\mathcal{O}}^2 \left(1 - \frac{\eta}{L_{\mathcal{S}/\mathcal{O}}}\right)^{j-k}.$$

Hence, using the elementary inequality $\sqrt{1-a} \leq (1-a/2)$,

$$\|\Phi_{\text{cl},j,k}\| \leq \sqrt{5} L_{\text{dyn}} L_{\mathcal{S}/\mathcal{O}} \left(1 - \frac{\eta}{L_{\mathcal{S}/\mathcal{O}}}\right)^{(j-k)/2} \leq \sqrt{5} L_{\text{dyn}} L_{\mathcal{S}/\mathcal{O}} \left(1 - \frac{\eta}{2L_{\mathcal{S}/\mathcal{O}}}\right)^{j-k},$$

which shows that we can select $B_{\text{stab}} = \sqrt{5} L_{\text{dyn}} L_{\mathcal{S}/\mathcal{O}}$ and $L_{\text{stab}} = 2L_{\mathcal{S}/\mathcal{O}}$. □

K.7 Solutions to recursions

This section contains the solutions to various recursions.

Lemma K.15 (First Key Recursion). Let $C_1 > 0$, $C_2 \geq 1/2$, $\beta_{\text{stab}} \in (0, 1)$, and suppose $\varepsilon_1, \varepsilon_2, \dots$ is a sequence satisfying $\varepsilon_1 \leq \bar{\varepsilon}_1$, and

$$\varepsilon_{k+1} \leq C_2 \beta_{\text{stab}}^k \bar{\varepsilon}_1 + C_1 \eta \sum_{j=1}^k \beta_{\text{stab}}^{k-j} \varepsilon_j^2$$

Then, as long as $C_1 \leq \beta(1 - \beta)/2\eta$, it holds that $\varepsilon_k \leq 2C_2 \beta_{\text{stab}}^{k-1} \bar{\varepsilon}_1$ for all k .

Proof. Consider the sequence $\nu_k = 2C_2 \beta_{\text{stab}}^{k-1} \bar{\varepsilon}_1$. Since $C_2 \geq 1/2$, we have $\nu_1 \geq \bar{\varepsilon}_1 \geq \varepsilon_1$. Moreover,

$$\begin{aligned} C_2 \beta_{\text{stab}}^k \bar{\varepsilon}_1 + C_1 \sum_{j=1}^k \beta_{\text{stab}}^{k-j} \nu_j &= C_2 \beta_{\text{stab}}^k \bar{\varepsilon}_1 + 2C_1 C_2 \sum_{j=1}^k \beta_{\text{stab}}^{k+j-2} \bar{\varepsilon}_1 \\ &= C_2 \beta_{\text{stab}}^k \bar{\varepsilon}_1 \left(1 + \frac{2C_1}{\beta} \sum_{j=0}^{k-1} \beta_{\text{stab}}^j \right) \\ &\leq C_2 \beta_{\text{stab}}^k \bar{\varepsilon}_1 \left(1 + \frac{2C_1 \eta}{\beta(1 - \beta)} \right) \end{aligned}$$

Hence, for $C_1 \leq \beta(1 - \beta)/2\eta$, we have $C_2 \beta_{\text{stab}}^k \bar{\varepsilon}_1 + C_1 \sum_{j=1}^k \beta_{\text{stab}}^{k-j} \nu_j \leq 2C_2 \bar{\varepsilon}_1 \beta_{\text{stab}}^k \leq \nu_{k+1}$. This shows that the (ν_k) sequence dominates the (ε_k) sequence, as needed. \square

Lemma K.16 (Second Key Recursion). Let $c, \Delta, \eta > 0$, $\beta_{\text{stab}} \in (0, 1)$ and let $\varepsilon_1, \varepsilon_2, \dots$ satisfy $\varepsilon_1 \leq c$ and

$$\varepsilon_{k+1} \leq c \beta_{\text{stab}}^k + c \eta \Delta \beta_{\text{stab}}^{k-1} \sum_{j=1}^k \varepsilon_j.$$

Then, if $\Delta \leq \frac{\beta(1-\beta)}{2c\eta}$, $\varepsilon_{k+1} \leq 2c \beta_{\text{stab}}^k$ for all k .

Proof. Consider the sequence $\nu_k = 2c \beta_{\text{stab}}^{k-1}$. Since $\varepsilon_1 \leq c$, $\nu_1 \geq \varepsilon_1$. Moreover,

$$\begin{aligned} c \beta_{\text{stab}}^k + c \eta \Delta \beta_{\text{stab}}^{k-1} \sum_{j=1}^k \nu_j &\leq c \beta_{\text{stab}}^k + 2c^2 \eta \Delta \beta_{\text{stab}}^{k-1} \sum_{j=1}^k \beta_{\text{stab}}^{j-1} \\ &\leq c \beta_{\text{stab}}^k + 2c^2 \eta \Delta \beta_{\text{stab}}^{k-1} \frac{1}{1 - \beta} \\ &\leq c \beta_{\text{stab}}^k \left(1 + 2c \Delta \frac{\eta}{\beta(1 - \beta)} \right). \end{aligned}$$

Hence, for $\Delta \leq \frac{\beta(1-\beta)}{2c\eta}$, the above is at most $2c \beta_{\text{stab}}^k \leq \nu_{k+1}$. This shows that the (ν_k) sequence dominates the (ε_k) sequence, as needed. \square

Lemma K.17 (Third Key Recursion). Let $\eta > 0$, $\beta = (1 - \frac{\eta}{L})$, $L \geq 2\eta$, and let $C, C', \Delta_2, \Delta_1 > 0$. Suppose that $\varepsilon_1, \varepsilon_2, \dots$ satisfies,

$$\varepsilon_{t+1} \leq \eta \sum_{s=1}^t \beta^{t-s} (\Delta_2 + C \varepsilon_s^2) + \beta^t \Delta_1$$

whenever $\max_{1 \leq s \leq t} \varepsilon_s \leq C'$. Suppose that

$$\Delta_2 \leq \frac{1}{\max\{8CL, 4LC'\}}, \quad \Delta_1 \leq \frac{1}{\max\{16CL, 8C'\}}$$

Then, for all t ,

$$\varepsilon_t \leq 4\Delta_1 \beta^{t-1} + 2L\Delta_2.$$

Proof. Consider $\bar{\varepsilon}_t = \alpha_1 \beta^{t-1} + \alpha_2$, with $\alpha_1 \geq \Delta_1$ and $\alpha_2 > 0$. As long as $\alpha_1 + \alpha_2 \leq C'$, we have $\bar{\varepsilon}_t \leq C'$ for all t . To show $\bar{\varepsilon}_t \geq \varepsilon_t$, it suffices that $\bar{\varepsilon}_t \geq \eta \sum_{s=1}^t \beta^{t-s} (\Delta_2 + C\bar{\varepsilon}_s^2) + \beta^t \Delta_1$. To have this occur, we need

$$\begin{aligned}
& \eta \sum_{s=1}^t \beta^{t-s} (\Delta_2 + C\bar{\varepsilon}_s^2) + \beta^t \Delta_1 \\
& \leq \eta \sum_{s=1}^t \beta^{t-s} (\Delta_2 + 2C\alpha_1^2 \beta^{2(s-1)} + 2C\alpha_2^2) + \beta^t \Delta_1 \\
& \leq \eta \sum_{s=1}^t \beta^{t-s} (\Delta_2 + 2C\alpha_1^2 \beta^{2(s-1)} + 2C\alpha_2^2) + \beta^t \Delta_1 \\
& = (\Delta_2 + 2C\alpha_2^2) \cdot (\eta \sum_{s=1}^t \beta^{t-s}) + 2C\alpha_1^2 \cdot (\eta \sum_{s=1}^t \beta^{t-s} \beta^{2(s-2)}) + \beta^t \Delta_1 \\
& = (\Delta_2 + 2C\alpha_2^2) \cdot (\eta \sum_{s=1}^t \beta^{t-s}) + \beta^{t-1} (2C\alpha_1^2 \cdot (\eta \sum_{s=1}^t \beta^{s-1}) + \Delta_1) \\
& \leq L(\Delta_2 + 2C\alpha_2^2) + \beta^{t-1} (2C\alpha_1^2 L + \Delta_1),
\end{aligned}$$

where the last step upper bounds the geometric series with $\eta = (1 - \eta/L)$. Assuming $\eta \leq L/2$, the above is at most

$$L(\Delta_2 + 2C\alpha_2^2) + 2\beta^t (2C\alpha_1^2 L + \Delta_1).$$

Matching terms, it is enough that

$$\alpha_2 \geq L(\Delta_2 + 2C\alpha_2^2), \quad \alpha_1 \geq 2(2C\alpha_1^2 L + \Delta_1), \quad \alpha_1 + \alpha_2 \leq C'$$

Let's choose $\alpha_2 = 2L\Delta_2$ and $\alpha_1 = 4\Delta_1$. Then, it is enough that

$$L\Delta_2 \geq 8CL^2\Delta_2^2, \quad 2\Delta_1 \geq 32C\Delta_1^2 L, \quad (2L\Delta_2 + 4\Delta_1) \leq \frac{1}{C'}$$

For this, it suffices that $\Delta_2 \leq \frac{1}{\max\{8CL, 4LC'\}}$ and $\Delta_1 \leq \frac{1}{\max\{16CL, 8C'\}}$. \square

Lemma K.18 (Matrix Product Perturbation). Define matrix products

$$\Phi_{k,j} = \mathbf{X}_{k-1} \cdot \mathbf{X}_{k-2} \cdots \mathbf{X}_j, \quad \Phi'_{k,j} = \mathbf{X}'_{k-1} \cdot \mathbf{X}'_{k-2} \cdots \mathbf{X}'_j.$$

Let $\eta, \Delta, c > 0$ and $\beta_{\text{stab}} \in (0, 1)$. If (a) $\Phi_{k,j} \leq \beta_{\text{stab}}^{k-j}$ for all $j \leq k$, (b) $\|\mathbf{X}_j - \mathbf{X}'_j\| \leq \eta \Delta \beta_{\text{stab}}^{j-1}$ for all $j \geq 1$ and (c) $\Delta \leq \frac{\beta(1-\beta)}{2c\eta}$, then, for all $j \leq k$, $\|\Phi'_{k,j}\| \leq 2c\beta_{\text{stab}}^{k-j}$.

Proof. Without loss of generality, take $j = 1$. Then, letting $\Delta_k = (\mathbf{X}'_k - \mathbf{X}_k)$,

$$\begin{aligned}
\Phi'_{k+1,1} &= \mathbf{X}'_k \cdot \mathbf{X}'_{k-2} \cdots \mathbf{X}'_1 \\
&= \mathbf{X}'_k \cdot \Phi'_{k,1} \\
&= \Delta_k \Phi'_{k,1} + \mathbf{X}_k \Phi'_{k,1} \\
&= \Delta_k \Phi'_{k,1} + \mathbf{X}_k \Delta_{k-1} \Phi'_{k-2,1} + \mathbf{X}_k \mathbf{X}_{k-1} \Phi'_{k-2,1} \\
&= \Phi_{k+1,k+1} \Delta_k \Phi'_{k,1} + \Phi_{k+1,k} \Delta_{k-1} \Phi'_{k-2,1} + \Phi_{k+1,k} \Phi'_{k-2,1} \\
&= \sum_{j=1}^k \Phi_{k+1,j+1} \Delta_j \Phi'_{j,1} + \Phi_{k+1,1}.
\end{aligned}$$

Thus,

$$\begin{aligned}
\|\Phi'_{k+1,1}\|_{\text{op}} &\leq c\eta \sum_{j=1}^k \beta_{\text{stab}}^{k-j} \|\mathbf{X}_j - \mathbf{X}'_j\| \|\Phi'_{j,1}\| + c\beta_{\text{stab}}^k \\
&\leq c\eta \beta_{\text{stab}}^{k-1} \Delta \sum_{j=1}^k \|\Phi'_{j,1}\| + c\beta_{\text{stab}}^k. \quad (\|\mathbf{X}_j - \mathbf{X}'_j\| \leq \eta \Delta \beta_{\text{stab}}^{j-1})
\end{aligned}$$

Define $\varepsilon_j = \|\Phi'_{j,1}\|$. Then, $\varepsilon_1 = 1 \leq c$, so Lemma K.16 implies that for $\Delta \leq \frac{(1-\beta)\beta}{2\eta}$, $\|\Phi'_{k,1}\| := \varepsilon_k \leq 2c\beta_{\text{stab}}^k$ for all k . \square

L Sampling and Score Matching

In this section, we provide a rigorous guarantee on the quality of sampling from the learned DDPM under [Assumption C.1](#). We begin by recalling the basic motivation for Denoising Diffusion Probabilistic Models (DDPMs) and explain how we train them. We then apply results from Chen et al. [18] to show that if we have learned the conditional score function, then sampling can be done efficiently. While Block et al. [13] demonstrated that unconditional score learning can be learned through standard statistical learning techniques, we generalize these results to the case of conditional score learning and conclude the section by proving that with sufficiently many samples, we can efficiently sample from a distribution close to our target.

We organize the section as follows:

- We then state the main result of the section, [Theorem 13](#), which provides a high probability upper bound on the number of samples n required in order to sample from DDPM trained on a given score estimate such that the sample is close in our optimal transport metric to the target distribution.
- In particular, in [\(L.1\)](#), we give the exact polynomial dependence of the sampling parameters α and J on the parameters of the problem.
- Before embarking on the proof, [Appendix L.1](#) introduces simplifying notation; notably, dropping the dependence on subscript h , replacing the score dependence on j with a class Θ_j , and denoted $\mathcal{D}_{\sigma,h,[t]}$ as simply $q_{[t]}$.
- We break the proof of [Theorem 13](#) into two sections. First, in [Appendix L.2](#), we recall a result of Chen et al. [18], Lee et al. [41] that shows that it suffices to accurately learn the score in the sense that if the score estimate is accurate in the appropriate sense, then the DDPM will adequately sample from a distribution close to the target.
- In [Remark L.4](#), we emphasize the conditions that would be required to sample in total variation and explain why they do not hold in our setting.
- Then, in [Appendix L.3](#), we apply statistical learning techniques, similar to those in Block et al. [13], to show that with sufficiently many samples, we can effectively learn the score. We detail in [Remark L.8](#) how the realizability part of [Assumption C.1](#) can be relaxed.
- Finally, we conclude the proof of [Theorem 13](#) by combining the two intermediate results detailed above.

To begin, we define our notion of statistical complexity:

We now state the main result of this section.

Theorem 13. Fix $1 \leq h \leq H$, let q denote $\mathcal{D}_{\sigma,h}$, d denote $d_{\mathcal{A}}$, and suppose that $(\mathbf{a}_i, \mathbf{o}_{h,i}) \sim q$ are independent for $1 \leq i \leq n$. Suppose that the projection of q onto the first coordinate has support (as defined in [Definition F.3](#)) contained in the euclidean ball of radius $R \geq 1$ in \mathbb{R}^d . For $\varepsilon > 0$, set

$$J = c \frac{d^3 R^4 (R + \sqrt{d})^4 \log\left(\frac{dR}{\varepsilon}\right)}{\varepsilon^{20}}, \quad \alpha = c \frac{\varepsilon^8}{d^2 R^2 (R + \sqrt{d})^2}. \quad (\text{L.1})$$

for some universal constant $c > 0$. Suppose that for all $1 \leq j \leq J$, the following hold:

- There exists a function class Θ_j containing some θ_j^* such that $\mathbf{s}_*(\cdot, \cdot, j\alpha) = \mathbf{s}_{\theta_j^*}(\cdot, \cdot, j\alpha) = \nabla \log q_{[j\alpha]}(\cdot|\cdot)$, where $q_{[\cdot]}$ is defined in [Section 2](#).
- The following holds for some $\delta > 0$:

$$\sup_{\substack{\theta, \theta' \in \Theta_j \\ \|\mathbf{a}\| \vee \|\mathbf{a}'\| \leq R + \sqrt{d \log\left(\frac{2nd}{\delta}\right)} \\ \|\mathbf{o}_h\| \leq R}} \|\mathbf{s}_\theta(\mathbf{a}, \mathbf{o}_h, t) - \mathbf{s}_{\theta'}(\mathbf{a}', \mathbf{o}_h, t)\| \leq c \frac{d^2 (R + \sqrt{d \log\left(\frac{2nd}{\delta}\right)})^2}{\varepsilon^8}.$$

- [Assumption C.1](#) holds and thus, for all $j \in [J]$, it holds that $\mathcal{R}_n(\Theta_j) \leq C_\Theta \alpha^{-1} n^{-1/\nu}$ for some $\nu \geq 2$ and all $n \in \mathbb{N}$ and, moreover, the linear growth condition is satisfied.

- The parameter $\hat{\theta} = \hat{\theta}_{1:J}$ is defined to be the empirical minimizer of $\mathcal{L}_{\text{DDPM}}$ from [Section 3](#).

If

$$n \geq c \left(\frac{C_{\Theta} \alpha^{-1} d R (R \vee \sqrt{d}) \log(dn)}{\varepsilon^4} \right)^{4\nu} \vee \left(\frac{d^6 (R^4 \vee d^2 \log^3(\frac{ndR}{\varepsilon\delta}))}{\varepsilon^{24}} d^2 \right)^{4\nu},$$

then with probability at least $1 - \delta$, it holds that

$$\mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\inf_{\mu \in \mathcal{E}(\text{DDPM}(\mathbf{s}_{\hat{\theta}}, \mathbf{o}_h), q(\cdot | \mathbf{o}_h))} \mathbb{P}_{(\hat{\mathbf{a}}, \mathbf{a}^*) \sim \mu} (\|\hat{\mathbf{a}} - \mathbf{a}^*\| \geq \varepsilon) \right] \leq 3\varepsilon.$$

Remark L.1. We emphasize that the exact value of the polynomial dependence (and in particular its pessimism) stem from the guarantees of Chen et al. [18], Lee et al. [41] regarding the quality of sampling with DDPMs. We remark below that the learning process itself does not incur such poor polynomial dependence except via these guarantees. Furthermore, we do not expect the sampling guarantees of those two works to be tight in any sense and such a poor polynomial dependence is not observed in practice. Rather, we include the bounds of Chen et al. [18], Lee et al. [41] so as to provide a fully rigorous end-to-end guarantee showing that polynomially many samples suffice to do imitation learning under our assumptions.

Remark L.2. A subtle difference between the presentation in the body and that here is the dependence of the complexity of Θ on the parameter α . We phrase the complexity guarantee as we did in the body in order to emphasize the dependence on the algorithmic parameter. If we let C'_{Θ} denote a constant such that $\mathcal{R}_n(\Theta) \leq C'_{\Theta} (\alpha/n)^{-1/\nu}$, then the sample complexity above becomes

$$n \geq c \left(\frac{C'_{\Theta} \log(dn)}{\alpha} \right)^{4\nu} \vee \left(\frac{d^2 (R^2 \vee d^2 \log^3(\frac{ndR}{\varepsilon\delta}))}{\alpha^2 \varepsilon^{16}} \right)^{4\nu}.$$

Critically, the guarantee of the quality of our DDPM is not in TV, but rather an optimal transport distance tailored to the problem at hand. As discussed in [Remark C.2](#), it is precisely this weaker guarantee that makes the problem challenging.

L.1 Simplifying Notation

We substantially simplify the notation in this appendix to suppress all dependence on h . In particular, we fix some $h \in [H]$ and consider $\mathbf{o}_h \sim \mathcal{D}_{\sigma, h, [0]}$. We let q denote $\mathcal{D}_{\sigma, h}$ and d denote $d_{\mathcal{A}}$. We further fix some $\sigma > 0$ and let $q_{[t]}$ denote the law of $\mathbf{a} \mid \mathbf{o}_h$ according to $\mathcal{D}_{\sigma, h, [t]}$ for the sake of notational simplicity. Furthermore, to emphasize that our analysis of the statistical learning theory decomposes accross DDPM time steps, we denote by Θ_j the function class $\mathbf{s}_{\theta}(\cdot, \cdot, \alpha j)$. We (redundantly) keep the dependence on t in the function evaluation for the sake of clarity. All other notation is defined *in situ*.

We emphasize that while our theoretical analysis treats each $\mathbf{s}_{\theta, h}$ separately, empirically one sees better success in training the score estimates jointly; on the other hand, the focus of this paper is not on sampling and score estimation and so we make the simplifying assumption for the sake of convenience.

L.2 Denoising Diffusion Probabilistic Models

We begin by motivating the sampling procedure described in (2.2), which is derived by fixing a horizon T and considering the continuum limit as $\alpha \downarrow 0$ and $J = \frac{T}{\alpha}$. More precisely, consider a forward process satisfying the stochastic differential equation (SDE) for $0 \leq t \leq T$:

$$d\mathbf{a}^t = -\mathbf{a}^t dt + \sqrt{2} dB_t, \quad \mathbf{a}^0 \sim q,$$

where B_t is a Brownian motion on \mathbb{R}^d and \mathbf{a}^0 is sampled from the target density. Applying the standard time reversal to this process results in the following SDE:

$$d\mathbf{a}_{\leftarrow}^{T-t} = (\mathbf{a}_{\leftarrow}^t + 2\nabla \log q_{T-t}(\mathbf{a}_{\leftarrow}^t)) dt + \sqrt{2} dB_t, \quad \mathbf{a}_{\leftarrow}^0 \sim q_T,$$

where q_t is the law of \mathbf{a}^t . Because the forward process mixes exponentially quickly to a standard Gaussian, in order to approximately sample from q , the learner may sample $\hat{\mathbf{a}}_{\leftarrow}^0 \sim \mathcal{N}(0, \mathbf{I})$ and evolving $\hat{\mathbf{a}}_{\leftarrow}^t$ according to the SDE above. Note that the classical Euler-Maruyama discretization of the above procedure is exactly (2.2), but with the true score $\nabla \log q_{T-t}$ replaced by score estimates $\mathbf{s}_\theta(\cdot, T-t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$; we may hope that if $\mathbf{s}_\theta(\cdot, T-t) \approx \nabla \log q_{T-t}$ as functions, then the procedure in (2.2) produces a sample close in law to q . Indeed, the following result provides a quantitative bound:

Theorem 14 (Corollary 4, Chen et al. [18]). Suppose that a distribution q on \mathbb{R}^d is supported on some ball of radius $R \geq 1$. Let C be a universal constant, fix $\varepsilon > 0$, and let α, J be set as in (L.1). If we have a score estimator $\mathbf{s}_\theta : \mathbb{R}^d \times [\tau] \rightarrow \mathbb{R}^d$ such that

$$\max_{j \in [J]} \mathbb{E}_{\mathbf{a} \sim q_{[\alpha j]}} \left[\left\| \mathbf{s}_\theta(\mathbf{a}, j) - \nabla \log q_{[\alpha j]}(\mathbf{a}) \right\|^2 \right] \leq \varepsilon^4,$$

then

$$\sup_{f: \|f\|_\infty \vee \|\nabla f\|_\infty \leq 1} \mathbb{E}_{\hat{\mathbf{a}} \sim \text{Law}(\mathbf{a}^J)} [f(\hat{\mathbf{a}})] - \mathbb{E}_{\mathbf{a}^* \sim q} [f(\mathbf{a}^*)] \leq \varepsilon^2,$$

where \mathbf{a}^J is sampled as in (2.2).

Remark L.3. As a technical aside, we note that Chen et al. [18, Corollary 4] applies to an “early stopped” DDPM, in the sense that the denoising is stopped in slightly fewer than J steps. On the other hand, for the choice of α given above, Chen et al. [18, Lemma 20 (a)] demonstrates that this distribution is ε^2 -close in Wasserstein distance to the sample produced by using all J steps and so by multiplying C above by a factor of 2 the guarantee is preserved. Because in practice we do not stop the DDPM early, we phrase Theorem 14 in the way above as opposed to the more complicated version with the early stopping.

Remark L.4. While [18, 41] show that if \mathbf{s}_θ is close to the $\mathbf{s}_{*,h}$ in $L^2(q_{[\alpha j]})$ and q satisfies mild regularity properties, then the law of \mathbf{a}_h^J will be close in total variation to q . Unfortunately, the required regularity of q , that the score is Lipschitz, is too strong to hold in many of our applications, such as when the data lie close to a low-dimensional manifold. In such cases, Chen et al. [18] provided guarantees in a weaker metric on distributions. We emphasize that even with full dimensional support, the Lipschitz constant of $\nabla \log q$ is likely large and thus the dependence on this constant appearing in Chen et al. [18, Theorem 2] is unacceptable. In particular, this subtle point is what necessitates the intricate construction of our paper; as remarked in Section 3, if we could expect the score to be sufficiently regular and producing a sample close in total variation to the target distribution were feasible, the problem would be trivial.

While Theorem 14 applies to unconditional sampling, it is easy to derive conditional sampling guarantees as a corollary.

Corollary L.1. Suppose that q is a joint distribution on actions \mathbf{a} and observations $\mathbf{o}_h \in \mathbb{R}^{d'}$. Further assume that the marginals over \mathbb{R}^d are fully supported in a ball of radius $R \geq 1$. Then there exists a universal constant C such that for all small $\varepsilon > 0$, if J and α are set as in (L.1) and

$$\mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\max_{j \in [J]} \mathbb{E}_{\mathbf{a} \sim q_{[\alpha j]}(\cdot | \mathbf{o}_h)} \left[\left\| \mathbf{s}_\theta(\mathbf{a}, j, \mathbf{o}_h) - \nabla \log q_{[\alpha j]}(\mathbf{a} | \mathbf{o}_h) \right\|^2 \right] \right] \leq \varepsilon^4, \quad (\text{L.2})$$

then

$$\mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\inf_{\mu \in \mathcal{C}(\text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h), q(\cdot | \mathbf{o}_h))} \mathbb{P}_{(\hat{\mathbf{a}}, \mathbf{a}^*) \sim \mu} (|\hat{\mathbf{a}} - \mathbf{a}^*| \geq \varepsilon) \right] \leq 3\varepsilon$$

Proof. We begin by showing an intermediate result,

$$\mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\sup_{f: \|f\|_\infty \vee \|\nabla f\|_\infty \leq 1} \mathbb{E}_{\hat{\mathbf{a}} \sim \text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h)} [f(\hat{\mathbf{a}})] - \mathbb{E}_{\mathbf{a}^* \sim q(\cdot | \mathbf{o}_h)} [f(\mathbf{a}^*)] \right] \leq 3\varepsilon^2. \quad (\text{L.3})$$

using Theorem 14. Let

$$\mathcal{A} = \left\{ \max_{j \in [J]} \mathbb{E}_{\mathbf{a} \sim q_{[\alpha j]}(\cdot | \mathbf{o}_h)} \left[\left\| \mathbf{s}_\theta(\mathbf{a}, j, \mathbf{o}_h) - \nabla \log q_{[\alpha j]}(\mathbf{a} | \mathbf{o}_h) \right\|^2 \right] \leq \varepsilon^2 \right\}.$$

By Markov's inequality and (L.2), it holds that

$$\mathbb{P}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}}(\mathcal{A}^c) \leq \frac{\varepsilon^4}{\varepsilon^2} = \varepsilon^2$$

and thus

$$\begin{aligned} & \mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\sup_{f: \|f\|_\infty \vee \|\nabla f\|_\infty \leq 1} \mathbb{E}_{\hat{\mathbf{a}} \sim \text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h)} [f(\hat{\mathbf{a}})] - \mathbb{E}_{\mathbf{a}^* \sim q(\cdot|\mathbf{o}_h)} [f(\mathbf{a}^*)] \right] \\ &= \mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\mathbf{I}[\mathcal{A}] \sup_{f: \|f\|_\infty \vee \|\nabla f\|_\infty \leq 1} \mathbb{E}_{\hat{\mathbf{a}} \sim \text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h)} [f(\hat{\mathbf{a}})] - \mathbb{E}_{\mathbf{a}^* \sim q(\cdot|\mathbf{o}_h)} [f(\mathbf{a}^*)] \right] \\ &+ \mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\mathbf{I}[\mathcal{A}^c] \sup_{f: \|f\|_\infty \vee \|\nabla f\|_\infty \leq 1} \mathbb{E}_{\hat{\mathbf{a}} \sim \text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h)} [f(\hat{\mathbf{a}})] - \mathbb{E}_{\mathbf{a}^* \sim q(\cdot|\mathbf{o}_h)} [f(\mathbf{a}^*)] \right] \\ &\leq \mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\mathbf{I}[\mathcal{A}] \inf_{q' \in \Delta(\mathbb{R}^d)} W_2(q(\cdot|\mathbf{o}_h), q') + \text{TV}(q', \text{Law}(\pi^\tau)) \right] + 2\varepsilon^2. \end{aligned}$$

For each \mathbf{o}_h , we may apply Theorem 14 and observe that for $\mathbf{o}_h \in \mathcal{A}$,

$$\sup_{f: \|f\|_\infty \vee \|\nabla f\|_\infty \leq 1} \mathbb{E}_{\hat{\mathbf{a}} \sim \text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h)} [f(\hat{\mathbf{a}})] - \mathbb{E}_{\mathbf{a}^* \sim q(\cdot|\mathbf{o}_h)} [f(\mathbf{a}^*)] \leq \varepsilon^2,$$

which proves (L.3). Now, for any fixed \mathbf{o}_h , by Markov's inequality and the definition of Wasserstein distance,

$$\inf_{\mu \in \mathcal{C}(\text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h), q(\cdot|\mathbf{o}_h))} \mathbb{P}_{(\hat{\mathbf{a}}, \mathbf{a}^*) \sim \mu} (\|\hat{\mathbf{a}} - \mathbf{a}^*\| \geq \varepsilon) \leq \frac{W_1(\text{DDPM}(\mathbf{s}_\theta, \mathbf{o}_h), q(\cdot|\mathbf{o}_h))}{\varepsilon}.$$

The result follows. \square

Note that the guarantee in Corollary L.1 is precisely what we need to control the one step imitation error in Theorem 4; thus, the problem of conditional sampling has been reduced to estimating the score. In the subsequent section, we will apply standard statistical learning techniques to provide a nonasymptotic bound on the quality of a score estimator.

L.3 Score Estimation

In the previous section we have shown that conditional sampling can be reduced to the problem of learning the conditional score. While there exist non-asymptotic bounds for learning the unconditional score [13], they apply to a slightly different score estimator than is typically used in practice. Here we upper bound the estimation error in terms of the complexity of the space of parameters Θ .

Observe that in order to apply Corollary L.1, we need a guarantee on the error of our score estimate in $L^2(q_{[\alpha j]})$ for all $j \in [J]$. Ideally, then, for fixed \mathbf{o}_h and $t = \alpha j$, we would like to minimize $\mathbb{E}_{\mathbf{a} \sim q_{[t]}} \left[\left\| \mathbf{s}_\theta(\mathbf{a}, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}|\mathbf{o}_h) \right\|^2 \right]$, where the inner norm is the Euclidean norm on \mathbb{R}^d . Unfortunately, because $q_{[t]}$ itself is unknown, we cannot even take an empirical version of this loss. Instead, through a now classical integration by parts [32, 74, 62], this objective can be shown to be equivalent to minimizing

$$\mathcal{L}_{\text{DDPM}}(\theta, \mathbf{a}, \mathbf{o}, t) = \mathbb{E}_{\mathbf{a} \sim q_{[t]}} \left[\left\| \mathbf{s}_\theta \left(e^{-t} \mathbf{a} + \sqrt{1 - e^{-2t}} \boldsymbol{\gamma}, \mathbf{o}_h, t \right) + \frac{1}{\sqrt{1 - e^{-2t}}} \boldsymbol{\gamma} \right\|^2 \right].$$

Because we are really interested in the expectation over the joint distribution $(\mathbf{a}, \mathbf{o}_h)$, we may take the expectation over \mathbf{o}_h and recover (4.1) as the empirical approximation. We now prove the following result for a single time step t :

Proposition L.1. Suppose that q is a distribution such that $q(\cdot|\mathbf{o}_i)$ is supported on a ball of radius R for q -almost every \mathbf{o}_h . For fixed $j \in [J]$ and α from (L.1), let $t = j\alpha$ and suppose that there is some

$\theta^* \in \Theta_j$ such that $\mathbf{s}_*(\cdot, \cdot, t) = \mathbf{s}_{\theta^*}(\cdot, \cdot, t) = \nabla \log q_{[t]}(\cdot|\cdot)$, i.e., \mathbf{s}_θ is rich enough to represent the true score at time t . Suppose further that the class of functions $\{\mathbf{s}_\theta | \theta \in \Theta_j\}$ satisfies for all $\theta \in \Theta_j$,

$$\sup_{\substack{\theta, \theta' \in \Theta_j \\ \|\mathbf{a}\| \vee \|\mathbf{a}'\| \leq R \\ \|\mathbf{o}_h\| \leq R}} \|\mathbf{s}_\theta(\mathbf{a}, \mathbf{o}_h, t) - \mathbf{s}_{\theta'}(\mathbf{a}', \mathbf{o}_h, t)\| \leq c \frac{d^2(R + \sqrt{d \log(\frac{2nd}{\delta})})^2}{\varepsilon^8}$$

for some universal constant $c > 0$. Recall the Rademacher term $\mathcal{R}_n(\Theta_j)$ defined in [Definition C.2](#), and let

$$\hat{\theta} \in \arg \min_{\theta \in \Theta} \sum_{i=1}^n \mathcal{L}_{\text{DDPM}}(\theta, \mathbf{a}_i, \mathbf{o}_i, t)$$

for independent and identically distributed $(\mathbf{a}_i, \mathbf{o}_i) \sim q$. Then it holds with probability at least $1 - \delta$ over the data that

$$\begin{aligned} & \mathbb{E}_{(\mathbf{a}_t, \mathbf{o}_h) \sim q_{[t]}} \left[\|\mathbf{s}_{\hat{\theta}}(\mathbf{a}_t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}_t | \mathbf{o}_h)\|^2 \right] \\ & \leq c \cdot \sqrt{\frac{\log(dn)}{1 - e^{-2t}}} \left(\mathcal{R}_n(\Theta) + \frac{d^2(R + \sqrt{d \log(\frac{2nd}{\delta})})^2}{\varepsilon^8} \cdot \sqrt{\frac{d \log(\frac{4dn}{\delta})}{n}} \right). \end{aligned}$$

Remark L.5. We note that while we assume a linearly growing score function for the sake of simplicity, our analysis easily handles any polynomial growth with a mild resulting change in the constants, omitted for the sake of simplicity.

Before we provide a proof, we recall the following result:

Lemma L.2. Suppose that q is supported in a ball of radius R and let $t \geq \alpha$ for α as in [\(L.1\)](#). Then $\nabla \log q_{[t]}(\cdot|\cdot)$ is L -Lipschitz with respect to the first parameter for

$$L = \frac{dR^2(R \vee \sqrt{d})^2}{\varepsilon^8}.$$

In particular,

$$\sup_{\substack{\|\mathbf{a}\| \vee \|\mathbf{a}'\| \leq R \\ \mathbf{o}_h}} \|\nabla \log q_{[t]}(\mathbf{a} | \mathbf{o}_h) - \nabla \log q_{[t]}(\mathbf{a}' | \mathbf{o}_h)\| \leq 2LR$$

and there exists some assignment of Θ and \mathbf{s}_θ that satisfies the boundedness condition in [Proposition L.1](#).

Proof. The first statement follows from replacing the ε in Chen et al. [[18](#), Lemma 20 (c)] by ε^2 . The second statement follows immediately from the first. \square

Remark L.6. Note that a slight variation of this result is what leads to the dependence on α in the growth parameter in [Assumption C.1](#) allowing for realizability. Indeed, by Chen et al. [[18](#), Lemma 20], it holds that the true score of $q_{[\alpha]}$ is

$$L = \frac{1}{1 - e^{-2\alpha}} \vee \frac{|1 - e^{-2\alpha}(1 + R^2)|}{(1 - e^{-2\alpha})^2}$$

Lipschitz, which is $O(\alpha^{-1})$ for $\alpha \ll 1$.

We also require the following standard result:

Lemma L.3. If $\mathcal{R}_n(\Theta_j)$ is defined as in [Definition C.2](#), then

$$\mathbb{E}_{\gamma_1, \dots, \gamma_n} \left[\sup_{\substack{\theta \in \Theta_j \\ 1 \leq j \leq J}} \frac{1}{n} \cdot \sum_{i=1}^n \langle \mathbf{s}_\theta(\mathbf{a}, \mathbf{o}_i, j), \gamma_i \rangle \right] \leq \sqrt{\pi \log(dn)} \cdot \mathcal{R}_n(\Theta_j)$$

Proof. This statement is classical and follows immediately from the fact that the norm of a Gaussian is independent from its sign as well as the fact that $\mathbb{E}[\max_{i,j}(\gamma_i)_j] \leq \sqrt{\pi \log(dn)}$ by classical Gaussian concentration. See Van Handel [69] for more details. \square

Proof of Proposition L.1. Let P_n denote the empirical measure on n independent samples $\{(\mathbf{a}_i, \mathbf{o}_i, \gamma_i)\}$ and let $\mathbf{a}_i^t = e^{-t}\mathbf{a}_i + \sqrt{1 - e^{-2t}}\gamma_i$. Let $C_t = \sqrt{1 - e^{-2t}}$ and observe that by definition and realizability,

$$P_n \left(\|C_t \mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \gamma\|^2 \right) \leq P_n \left(\|C_t \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h) - \gamma\|^2 \right). \quad (\text{L.4})$$

We emphasize that by Lemma L.2, realizability does not make the result vacuous. Adding and subtracting $C_t \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h)$ from the left hand inequality, expanding and rearranging, we see that

$$\begin{aligned} C_t^2 P_n \left(\|\mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h)\|^2 \right) &\leq 2C_t \cdot P_n \left(\langle \mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h), \gamma \rangle \right) \\ &\leq 2C_t \cdot P_n \left(\sup_{\theta \in \Theta_j} \langle \mathbf{s}_{\theta}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h), \gamma \rangle \right). \end{aligned}$$

We now claim that with probability at least $1 - \delta$, it holds that

$$\begin{aligned} P_n \left(\sup_{\theta \in \Theta} \langle \mathbf{s}_{\theta}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h), \gamma \rangle \right) &\leq \mathbb{E} \left[P_n \left(\sup_{\theta \in \Theta_j} \langle \mathbf{s}_{\theta}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h), \gamma \rangle \right) \right] \\ &\quad + B \cdot \sqrt{\frac{d \log \left(\frac{2d}{\delta} \right)}{n}}, \end{aligned}$$

where

$$B = c \frac{d^2 (R + \sqrt{d \log \left(\frac{2nd}{\delta} \right)})^2}{\varepsilon^8} \quad (\text{L.5})$$

for some universal constant $c > 0$. To see this, we claim that with probability at least $1 - \frac{\delta}{2}$, it holds that $\|\mathbf{a}_i^t\| \leq c \left(R + \sqrt{d \log \left(\frac{2nd}{\delta} \right)} \right)$ for all $1 \leq i \leq n$. Indeed, this follows by Gaussian concentration in Jin et al. [34, Lemmata 1 & 2]. We may now apply Lemma L.2 to a bound on the osculation of $\mathbf{s}_{\theta} - \nabla \log q_{[t]}$ in the ball of the above radius. Conditioning on the event that $\|\mathbf{a}_i^t\|$ is bounded by the above, we may argue as in Wainwright [75, Theorem 4.10] that if we let the function

$$G = G(\mathbf{a}_1, \mathbf{o}_1, \dots, \mathbf{a}_n, \mathbf{o}_n) = P_n \left(\sup_{\theta \in \Theta_j} \langle \mathbf{s}_{\theta}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h), \gamma \rangle \right),$$

then for any i , on the event of bounded norm, replacing $(\mathbf{a}_i, \mathbf{o}_i)$ with $(\mathbf{a}'_i, \mathbf{o}'_i)$ and leaving other terms unchanged changes ensures that $|G - G'| \leq \frac{2B}{n} \gamma_i$. Thus by Jin et al. [34, Corollary 7] and a union bound, the claim holds. Because γ is mean zero, we have

$$\begin{aligned} \mathbb{E} \left[P_n \left(\sup_{\theta \in \Theta} \langle \mathbf{s}_{\theta}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h), \gamma \rangle \right) \right] &\leq \mathbb{E} \left[P_n \left(\sup_{\theta \in \Theta} \langle \mathbf{s}_{\theta}(\mathbf{a}^t, \mathbf{o}_h, t), \gamma \rangle \right) \right] \\ &\leq \sqrt{\pi \log(dn)} \cdot \mathcal{R}_n(\Theta_j), \end{aligned}$$

where the last inequality follows by Lemma L.3 and the fact that $t = jJ$. Summing up the argument until this point and rearranging tells us that with probability at least $1 - \delta$, it holds that

$$P_n \left(\|\mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h)\|^2 \right) \leq \frac{2}{C_t} \sqrt{\pi \log(nd)} \cdot \mathcal{R}_n(\Theta) + \frac{B}{C_t} \cdot \sqrt{\frac{d \log \left(\frac{2nd}{\delta} \right)}{n}},$$

with B given in (L.5). We now use a uniform norm comparison between population and empirical norms to conclude the proof. Indeed, it holds by Rakhlin et al. [54, Lemma 8.i & 9] that there exists a critical radius

$$r_n \leq cB \log^3(n) \mathcal{R}_n(\Theta_j)^2$$

such that with probability at least $1 - \delta$,

$$\begin{aligned} \mathbb{E}_{(\mathbf{a}^t, \mathbf{o}_h) \sim q_{[t]}} \left[\left\| \mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h) \right\|^2 \right] \\ \leq 2 \cdot P_n \left(\left\| \mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h) \right\|^2 \right) + cr_n + c \frac{\log\left(\frac{1}{\delta}\right) + \log \log n}{n}, \end{aligned}$$

where again c is some universal constant. Combining this with our earlier bound on the empirical distance and a union bound, after rescaling δ , we have that

$$\begin{aligned} \mathbb{E}_{(\mathbf{a}^t, \mathbf{o}_h) \sim q_{[t]}} \left[\left\| \mathbf{s}_{\hat{\theta}}(\mathbf{a}^t, \mathbf{o}_h, t) - \nabla \log q_{[t]}(\mathbf{a}^t | \mathbf{o}_h) \right\|^2 \right] &\leq \frac{4}{C_t} \sqrt{\pi \log(nd)} \cdot \mathcal{R}_n(\Theta_j) + \frac{2B}{C_t} \cdot \sqrt{\frac{d \log\left(\frac{4nd}{\delta}\right)}{n}} \\ &\quad + cB \log^3(n) \cdot \mathcal{R}_n^2(\Theta_j) + c \frac{\log\left(\frac{2}{\delta}\right) + \log \log(n)}{n} \end{aligned}$$

with probability at least $1 - \delta$. This concludes the proof. \square

Remark L.7. For the sake of simplicity, in the proof of [Proposition L.1](#) we applied uniform deviations and recovered the “slow rate” of $\mathcal{R}_n(\Theta)$, up to logarithmic factors. Indeed, if we were to further assume that the score function class is star-shaped around the true score, we could recover a faster rate, as was done in the case of unconditional sampling in Block et al. [13] with a slightly different loss. While in our proof the appeal to Rakhlin et al. [54] to control the population norm by the empirical norm could be replaced with a simpler uniform deviations argument because we have already given up on the fast rate, such an argument is necessary in the more refined analysis. As the focus of this paper is not on the sampling portion of the end-to-end analysis, we do not include a rigorous proof of the case of fast rates for the sake of simplicity and space.

Remark L.8. While we assumed for simplicity that the score was realizable with respect to our function class for every time $t = \alpha j$, this condition can be relaxed to approximate realizability in a standard way. In particular, if the score is ε -far away from some function representable by our class in a pointwise sense, then we can add an ε to the right hand side of (L.4) with minimal modification to the proof.

With [Proposition L.1](#), and a union bound, we recover the following result:

Proposition L.4. Suppose that the conditions on \mathbf{s}_{θ} in [Proposition L.1](#) continue to hold. Suppose further that $\|\mathbf{s}_{\theta}(\mathbf{a}, \mathbf{o}_h, t)\| \leq C_{\text{grow}}(1 + \|\mathbf{a}\| + \|\mathbf{o}_h\|)$ for all \mathbf{a} and some universal constant $C_{\text{grow}} > 0$. Let J and α be as in (L.1) and suppose that $\alpha \leq \frac{1}{2}$. Then, with probability at least $1 - \delta$ over \mathcal{D}' , it holds that

$$\begin{aligned} \mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\max_{j \in [J]} \mathbb{E}_{\mathbf{a} \sim q_{[\alpha j]}(\cdot | \mathbf{o}_h)} \left[\left\| \mathbf{s}_{\theta}(\mathbf{a}, j, \mathbf{o}_h) - \nabla \log q_{[\alpha j]}(\mathbf{a} | \mathbf{o}_h) \right\|^2 \right] \right] \\ \leq c \frac{d(R \vee \sqrt{d})^2 \log(dn)}{\varepsilon^4} \mathcal{R}_n(\Theta) + c \frac{d^3 \left((R \vee \sqrt{d})^2 + d \log\left(\frac{ndR}{\delta\varepsilon}\right) \right)}{\varepsilon^{12}} \sqrt{\frac{d \log\left(\frac{4dnR}{\delta\varepsilon}\right)}{n}} \end{aligned}$$

In particular if

$$\mathcal{R}_n(\Theta_j) \leq C_{\Theta} n^{-1/\nu}$$

for some $\nu \geq 2$ and all $j \in [J]$, then for

$$n \geq cC_{\text{grow}} \left(\frac{C_{\Theta} \alpha^{-1} d(R \vee \sqrt{d})^2 \log(dn)}{\varepsilon^4} \right)^{4\nu} \vee \left(\frac{d^6 (R^4 \vee d^3 \log^3\left(\frac{ndR}{\delta\varepsilon}\right))}{\varepsilon^{24}} d^2 \right)^{4\nu}$$

it holds that with probability at least $1 - \delta$,

$$\mathbb{E}_{\mathbf{o}_h \sim q_{\mathbf{o}_h}} \left[\max_{j \in [J]} \mathbb{E}_{\mathbf{a} \sim q_{[\alpha j]}(\cdot | \mathbf{o}_h)} \left[\left\| \mathbf{s}_{\theta}(\mathbf{a}, j, \mathbf{o}_h) - \nabla \log q_{[\alpha j]}(\mathbf{a} | \mathbf{o}_h) \right\|^2 \right] \right] \leq \varepsilon^4.$$

Proof. We begin by proving the result on the event that $\|\mathbf{a}\| \vee \|\mathbf{o}_h\| \leq C(R \vee \sqrt{d}) \log\left(\frac{Jn(R \vee \sqrt{d})}{\delta\varepsilon}\right)$.

Note that

$$1 - e^{-2t} \geq 1 - e^{-2\alpha} \geq \alpha$$

because $2\alpha \leq 1$. We now apply [Proposition L.1](#) and take a union bound over $j \in [J]$. All that remains is to demonstrate that the contribution of the event that \mathbf{a}^j is outside the above defined ball is negligible. To do this, observe that by Lee et al. [41, Lemma 4.15], there is some $C > 0$ such that \mathbf{a}^j is $C(R \vee \sqrt{d})$ -subGaussian. By the sublinearity of the growth of \mathbf{s}_θ in \mathbf{a} , as well as the Lipschitzness of $q_{[\alpha j]}$ following from Chen et al. [18, Lemma 20], bounding a maximum by a sum, and the elementary computation in [Lemma L.5](#), we have that the expectation of this term on this event is bounded by $\frac{CC_{\text{grow}}}{n}$. The result follows. \square

We note that in our simplified analysis, we have assumed that $N_{\text{aug}} = 1$, i.e., for each sample, we take a single noise level from the path. In practice, we use many augmentations per sample. Again, as the focus of our paper is not on score estimation and sampling, we treat this as a simple convenience and leave open to future work the problem of rigorously demonstrating that multiple augmentations indeed help with learning. Finally, for a discussion on bounding $\mathcal{R}_n(\Theta)$, see Wainwright [75].

Proof of Theorem 13. We note that the proof follows immediately from combining [Corollary L.1](#) with [Proposition L.4](#). \square

We conclude the section with the following elementary computation used above:

Lemma L.5. Suppose that X is a σ -subGaussian random variable on \mathbb{R} . Then for any $r \geq 1$,

$$\mathbb{E}[|X| \cdot \mathbf{I}[|X| > r]] \leq C \frac{\sigma}{r} \cdot e^{-\frac{r^2}{2\sigma^2}}$$

Proof. This is an elementary computation. Indeed,

$$\begin{aligned} \mathbb{E}[|X| \cdot \mathbf{I}[|X| > r]] &= \int_r^\infty \mathbb{P}(|X| > t) dt \leq C \int_r^\infty e^{-\frac{t^2}{2\sigma^2}} dt \\ &\leq C \cdot \int_r^\infty \frac{t}{r} e^{-\frac{t^2}{2\sigma^2}} dt \\ &\leq C \frac{\sigma}{r} \cdot e^{-\frac{r^2}{2\sigma^2}}. \end{aligned}$$

The result follows. \square

M Proofs for Generic Incrementally Stable Primitive Controllers

This section proves [Theorems 5](#) and [6](#), generalizing our guarantees to general primitive controllers. Note that, in this more general setting, we can no longer expect to bound the norm of the difference between two controllers evaluated at some point \mathbf{x} $\kappa_t(\mathbf{x}) - \kappa_t(\mathbf{x}')$ by differences in their parameter values. Instead, we opt for the more local notion of distance considered in [Theorems 5](#) and [6](#), via the localized distance $d_{\text{loc},\alpha}$ considered in [Definition E.1](#). To this end, [Appendix M.1](#) begins by generalizing the analysis of the composite MDP to allow the distance $d_{\mathcal{A}}$ take an additional state-argument (in order to capture the localization of the distance in $d_{\text{loc},\alpha}$). [Appendix M.2](#) then converts our assumption of incremental stability, [Assumption 3.1b](#), into the IPS stability conditioned required in the composite MDP. Finally, we conclude of our intended results in [Appendix M.3](#), following the same arguments as for affine primitive controllers in [Appendix J](#).

M.1 Generalization of analysis in the composite MDP

Here, we consider a generalization of the analysis of the composite MDP where we allow $d_{\mathcal{A}}$ to depend on state. Our analysis follows [Appendix H.3](#) and the proof of [Theorem 9](#). All notation here borrows from that section. Formally, we consider

$$d_{\mathcal{A};\mathcal{S}}(\cdot, \cdot \mid \cdot) : (\mathcal{A} \times \mathcal{A}) \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}.$$

We recall the direct decomposition in [Definition H.1](#) of $\mathcal{S} = \mathcal{Z} \oplus \mathcal{S}_{/\mathcal{O}}$, where we recall that \mathcal{Z} is the component that coincides with the ‘ \mathbf{o}_h ’ component of the state in our instantiation. Further, recall that $\phi_{\mathcal{O}}$ is the projection onto the \mathcal{Z} component.

Condition M.1 (Measurability). We require that $d_{\mathcal{A};\mathcal{S}}$ is measurable, and that, for all s , the set $\{(a', a) : d_{\mathcal{A};\mathcal{S}}(a', a; s) > \varepsilon\}$ is open. and that $(a', a, s) \mapsto d_{\mathcal{A};\mathcal{S}}(a', a; s)$ is measurable. We also assume that $d_{\mathcal{A};\mathcal{S}}(a', a; s)$ only depends on s through $\phi_o(s)$.

We re-define a state-conditioned input stability as follows

Definition M.1. We say that a sequence $(s_{1:H+1}, a_{1:H})$ is state-conditioned input-stable with respect to an auxilliary sequence $\tilde{s}_{1:H+1}$ if

$$d_{\mathcal{S}}(s'_{h+1}, s_{h+1}) \vee d_{\text{TVC}}(s'_{h+1}, s_{h+1}) \leq \max_{1 \leq j \leq h} d_{\mathcal{A};\mathcal{S}}(a'_j, a_j \mid \tilde{s}_j), \quad \forall h \in [H]$$

We now define a one-step error which is state dependent (allowing for $d_{\mathcal{A};\mathcal{S}}$). To simplify the exposition, we define marginal gaps which ignore the now-state-dependent $d_{\mathcal{A};\mathcal{S}}$.

Definition M.2 (Modified Imitation Gaps). Define the state

$$d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_h(s), \pi_h^*(s) \mid s') := \inf_{\mu_2} \mathbb{P}_{\mu_2} [d_{\mathcal{A};\mathcal{S}}(\hat{a}_h, a_h^* \mid s') \leq \varepsilon],$$

where the infimum is over couplings $(a_h^*, \hat{a}_h) \sim \mu_2 \in \mathcal{C}(\hat{\pi}_h(s), \pi_h^*(s))$. Further define

$$\Gamma_{\text{marg},\mathcal{S},\varepsilon} := \inf_{\mu_1} \mathbb{P}_{\mu_1} \left[\max_{h \in [H]} \max \{d_{\mathcal{S}}(s_{h+1}^*, \hat{s}_{h+1}) > \varepsilon\} \right], \quad \Gamma_{\text{joint},\mathcal{S},\varepsilon} := \max_{h \in [H]} \inf_{\mu_1} \mathbb{P}_{\mu_1} [d_{\mathcal{S}}(s_{h+1}^*, \hat{s}_{h+1}) > \varepsilon]$$

where above μ_1 ranges over the same couplings as in [Definition D.1](#).

Guarantees under TVC of $\hat{\pi}$. We now generalize [Proposition D.2](#) under the assumption that $\hat{\pi}$ is TVC.

Proposition M.1 (Generalization of [Proposition D.2](#)). Let π^* be state-conditioned input-stable w.r.t. $(d_{\mathcal{S}}, d_{\mathcal{A};\mathcal{S}})$ and let $\hat{\pi}$ be γ -TVC. Then, for all $\varepsilon > 0$,

$$\Gamma_{\text{joint},\mathcal{S},\varepsilon}(\hat{\pi} \parallel \pi^*) \leq H\gamma(\varepsilon) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi_h^*(s_h^*) \mid s_h^*).$$

Proof Sketch. The proof is nearly identical to the proof of [Proposition D.2](#) in [Appendix G](#). The only difference is that, when we measure the distance between $a_h^* \sim \pi_h^*(s_h^*)$ and $\hat{a}_h^{\text{inter}} \sim \hat{\pi}_h(s_h^*)$, this distance is specified at s_h^* . Hence, we replace $d_{\mathcal{A}}(\hat{a}_h^{\text{inter}}, a_h^*)$ with $d_{\mathcal{A}}(\hat{a}_h^{\text{inter}}, a_h^* \mid s_h^*)$. This leads to use replacing $d_{\text{os},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi_h^*(s_h^*))$ with $d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi_h^*(s_h^*) \mid s_h^*)$ in the final bound. \square

Guarantees with smoothing kernel. Next, we turn to the generalization of [Theorems 4](#) and [9](#) to allow for state-conditioned action distances.

Definition M.3. Given non-decreasing maps $\gamma_{\text{IPS},\text{TVC}}, \gamma_{\text{IPS},\mathcal{S}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ a pseudometric $d_{\text{IPS}} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ (possibly other than $d_{\mathcal{S}}$ or d_{TVC}), and $r_{\text{IPS}} > 0$, we say a policy π is $(\gamma_{\text{IPS},\text{TVC}}, \gamma_{\text{IPS},\mathcal{S}}, d_{\text{IPS}}, r_{\text{IPS}})$ -state-conditioned-restricted IPS if it satisfies the conditions of [Definition H.9](#), with the only modification that for the constructed $s_{1:H+1}, a_{1:H}, \tilde{s}_{1:H}$, the condition that $s_{1:H+1}, a_{1:H}$ is input-stable is replaced with “state-conditioned input stable with respect to the sequence $\tilde{s}_{1:H}$.” More precisely, the condition is met if the following holds for any $r \in [0, r_{\text{IPS}}]$. Consider any sequence of kernels $W_1, \dots, W_H : \mathcal{S} \rightarrow \Delta(\mathcal{S})$ satisfying

$$\max_{h,s \in \mathcal{S}} \mathbb{P}_{\tilde{s} \sim W_h(s)} [d_{\text{IPS}}(\tilde{s}, s) \leq r] = 1, \quad \forall s, \quad \phi_o \circ W_h(s_h) \ll \phi_o \circ P_h^*.$$

and define a process $s_1 \sim P_{\text{init}}, \tilde{s}_h \sim W_h(s_h), a_h \sim \pi_h(\tilde{s}_h)$, and $s_{h+1} := F_h(s_h, a_h)$. Then, almost surely,

- (a) the sequence $(s_{1:H+1}, a_{1:H})$ is state-conditioned input stable with respect to the sequence $\tilde{s}_{1:H}$
- (b) $\max_{h \in [H]} d_{\text{TVC}}(F_h(\tilde{s}_h, a_h), s_{h+1}) \leq \gamma_{\text{IPS},\text{TVC}}(r)$ and (c) $\max_{h \in [H]} d_{\mathcal{S}}(F_h(\tilde{s}_h, a_h), s_{h+1}) \leq \gamma_{\text{IPS},\mathcal{S}}(r)$.

Theorem 15. Consider the setting of [Theorem 9](#), but with π^* satisfies $(\gamma_{\text{IPS}, \text{TVC}}, \gamma_{\text{IPS}, \mathcal{S}}, d_{\text{IPS}}, r_{\text{IPS}})$ -state-conditioned-restricted IPS ([Definition M.3](#)) rather than (standard) restricted IPS ([Definition H.9](#)). Again, let $\varepsilon > 0$ and $r \in (0, \frac{1}{2}r_{\text{IPS}}]$, and efine

$$p_r := \sup_s \mathbb{P}_{s' \sim W_\sigma(s)}[d_{\text{IPS}}(s', s) > r], \quad \varepsilon' := \varepsilon + \gamma_{\text{IPS}, \mathcal{S}}(2r)$$

Then, for any policy $\hat{\pi}$, both $\Gamma_{\text{joint}, \mathcal{S}, \varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$ and $\Gamma_{\text{marg}, \mathcal{S}, \varepsilon'}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ are upper bounded by

$$H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS}, \text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os}, \mathcal{S}, \varepsilon}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec}, h}^*(\tilde{s}_h^*) \mid s_h^*).$$

Proof. The proof follows by modifying [Theorem 8](#), and hence [Theorem 9](#) as a consequence. The key change is that we replace the event $\mathcal{B}_{\text{est}, h} = \left\{ d_{\mathcal{A}}(\hat{a}_h^{\text{tel, inter}}, a_h^{\text{tel}}) > \varepsilon \right\}$ ¹⁵ in [Definition H.6](#) with

$$\mathcal{B}_{\text{est}, h} = \left\{ d_{\mathcal{A}}(\hat{a}_h^{\text{tel, inter}}, a_h^{\text{tel}} \mid \tilde{s}_h^{\text{tel}}) > \varepsilon \right\}, \quad (\text{M.1})$$

and replace the event \mathcal{Q}_{IS} in [Definition H.8](#) with

$$\mathcal{Q}_{\text{IS}} := \{s_{1:H+1}^{\text{rep}}, a_{1:H}^{\text{rep}} \text{ is state-conditioned input-stable w.r.t. } \tilde{s}_{1:H}^{\text{rep}}\}$$

We also define the following event

$$\mathcal{Q}'_{\text{IS}, h} := \{s_{1:h+1}^{\text{rep}}, a_{1:h}^{\text{rep}} \text{ is state-conditioned input-stable w.r.t. } \tilde{s}_{1:h}^{\text{tel}}\},$$

which considers input stability for $h \leq H$ steps and shifts the reference sequence from $\tilde{s}_{1:h}^{\text{rep}}$ to $\tilde{s}_{1:h}^{\text{tel}}$. What changes as a result of these argument is as follows:

- We check that [Lemma H.7](#) goes through:

$$\bar{\mathcal{C}}_{\text{all}, h+1} \subset \mathcal{Q}_{\text{all}} \cap \bar{\mathcal{C}}_{\text{all}, h} \cap \bar{\mathcal{B}}_{\text{all}, h}.$$

The first modification here is that, when $s_{1:H+1}^{\text{rep}}, a_{1:H}^{\text{rep}}$ is input stable with respect to $\tilde{s}_{1:H}^{\text{rep}}$,

$$d_{\mathcal{S}}(\hat{s}_{h+1}, s_{h+1}^{\text{rep}}) \vee d_{\text{TVC}}(\hat{s}_{h+1}, s_{h+1}^{\text{rep}}) \leq \max_{1 \leq j \leq h} d_{\mathcal{A}; \mathcal{S}}(\hat{a}_j, a_j^{\text{rep}} \mid \tilde{s}_j^{\text{rep}}), \quad \forall h \in [H]$$

Since $d_{\mathcal{A}; \mathcal{S}}(a, a' \mid s)$ depends only on s through $\phi_o(s)$, then when

$$\bigcap_{1 \leq j \leq h} \mathcal{B}_{\text{tel}, j} := \{a_j^{\text{rep}} = a_j^{\text{tel}}, \phi_o(s_j^{\text{rep}}) = \phi_o(\tilde{s}_j^{\text{tel}})\} \subset \bar{\mathcal{B}}_{\text{all}, h},$$

holds, we have

$$\max_{1 \leq j \leq h} d_{\mathcal{A}; \mathcal{S}}(\hat{a}_j, a_j^{\text{rep}} \mid \tilde{s}_j^{\text{rep}}) = \max_{1 \leq j \leq h} d_{\mathcal{A}; \mathcal{S}}(\hat{a}_j, a_j^{\text{tel}} \mid \tilde{s}_j^{\text{tel}}), \quad \forall h \in [H]$$

Finally, on $\bar{\mathcal{B}}_{\text{all}, h}$, we have $\hat{a} = \hat{a}^{\text{tel, inter}}$, so we get

$$\max_{1 \leq j \leq h} d_{\mathcal{A}; \mathcal{S}}(\hat{a}_j, a_j^{\text{rep}} \mid \tilde{s}_j^{\text{rep}}) = \max_{1 \leq j \leq h} d_{\mathcal{A}; \mathcal{S}}(\hat{a}_j^{\text{tel, inter}}, a_j^{\text{tel}} \mid \tilde{s}_j^{\text{tel}}), \quad \forall h \in [H],$$

which is at most ε under our second definition of $\mathcal{B}_{\text{est}, h}$.

- [Lemma H.8](#) goes unchanged
- We check that [Lemma H.9](#) goes through. This follows from the definition of state-conditioned input-stable, using $s_{1:H+1}^{\text{rep}}, a_{1:H}^{\text{rep}}, \tilde{s}_{1:H}^{\text{rep}}$ as $s_{1:H+1}, a_{1:H}, \tilde{s}_{1:H}$, so that when $\mathbb{P}_\mu[\mathcal{Q}_{\text{IS}}^c \cap \mathcal{Q}_{\text{close}}] = 0$.¹⁶

¹⁵This is the special case of $\mathcal{B}_{\text{est}, h} = \left\{ \vec{d}_{\mathcal{A}}(\hat{a}_h^{\text{tel, inter}}, a_h^{\text{tel}}) \not\leq \varepsilon \right\}$ with $\vec{d}_{\mathcal{A}}$ being scalar valued (i.e. all coordinate identical).

¹⁶Here, we replace p_{IPS} in that lemma with failure probability 0.

- Recall that $d_{\mathcal{A};\mathcal{S}}(a, a'; s)$ depends only on s through $\phi_o(s)$. Hence, under the event

$$\bigcap_{1 \leq j \leq h} \mathcal{B}_{\text{tel},j} := \{a_j^{\text{rep}} = a_j^{\text{tel}}, \phi_o(\tilde{s}_j^{\text{rep}}) = \phi_o(\tilde{s}_j^{\text{tel}})\}, \subset \bar{\mathcal{B}}_{\text{all},h}$$

we have that \mathcal{Q}_{IS} implies $\mathcal{Q}'_{\text{IS},h}$.

- Lemma H.10 replaces $d_{\text{os},\varepsilon}(\hat{\pi}_{\sigma,h}(\tilde{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\tilde{s}_h^{\text{tel}}))$ with $d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_{\sigma,h}(\tilde{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\tilde{s}_h^{\text{tel}}) \mid \tilde{s}_h^{\text{tel}})$, i.e. the state-conditioned one-step error (that we condition on \tilde{s}_h^{tel} comes from our re-definition of $\mathcal{B}_{\text{est},h}$ in (M.1). Thus, we get $\Gamma_{\text{joint},\mathcal{S},\varepsilon}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$ and $\Gamma_{\text{marg},\mathcal{S},\varepsilon'}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ are upper bounded by

$$H(2p_r + 3\gamma_\sigma(\max\{\varepsilon, \gamma_{\text{IPS},\text{TVC}}(2r)\})) + \sum_{h=1}^H \mathbb{E}_{\tilde{s}_h^{\text{tel}}} d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_{\sigma,h}(\tilde{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\tilde{s}_h^{\text{tel}}) \mid \tilde{s}_h^{\text{tel}}). \quad (\text{M.2})$$

- Because \tilde{s}_h^{tel} has marginal $s_h^* \sim P_h^*$, we can replace the terms $\mathbb{E}_{\tilde{s}_h^{\text{tel}}} d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_{\sigma,h}(\tilde{s}_h^{\text{tel}}) \parallel \pi_{\text{rep},h}^*(\tilde{s}_h^{\text{tel}}) \mid \tilde{s}_h^{\text{tel}})$ with $\mathbb{E}_{s_h^* \sim P_h^*} d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_{\sigma,h}(s_h^*) \parallel \pi_{\text{rep},h}^*(s_h^*) \mid s_h^*)$.
- Using the same data-processing argument as in the proof as in Theorem 8, we can bound $\mathbb{E}_{s_h^* \sim P_h^*} d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_{\sigma,h}(s_h^*) \parallel \pi_{\text{rep},h}^*(s_h^*) \mid s_h^*) \leq \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{s_h^* \sim W_\sigma(s_h^*)} d_{\text{os},\mathcal{S},\varepsilon}(\hat{\pi}_h(s_h^*) \parallel \pi_{\text{dec},h}^*(s_h^*) \mid s_h^*)$.

□

M.2 State-Conditioned Input-Stability and IPS in the Composite MDP via t-ISS

Lemma K.3 reduced IPS in the composite MDP to incremental stability in a form that applies primarily to affine primitive controllers. In this section, we generalize the lemma further to depend on a more localized distance reflecting the state-conditioned distance $d_{\mathcal{A};\mathcal{S}}(\cdot, \cdot; \cdot)$ in the composite MDP. Recall the local-distance between composite actions $a = \kappa_{1:\tau_{\text{chunk}}}$, $a' = \kappa'_{1:\tau_{\text{chunk}}} \in \mathcal{A}$ at state \mathbf{x} and scale $\alpha > 0$, defined in Definition E.1 as

$$d_{\text{loc},\alpha}(a, a' \mid \mathbf{x}) := \max_{1 \leq i \leq \tau_{\text{chunk}}} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha} \|\kappa_i(\mathbf{x}_i + \delta \mathbf{x}) - \kappa'_i(\mathbf{x}_i + \delta \mathbf{x})\|,$$

where above $\mathbf{x}_1 = \mathbf{x}$, $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \kappa_t(\mathbf{x}))$ with $a = \kappa_{1:\tau_{\text{chunk}}}$.

Lemma M.2. Instantiate the composite MDP as in Definition K.1, with π^* as in Definition J.3. Furthermore, suppose that under $(\rho_T, a_{1:H}) \sim \mathcal{D}_{\text{exp}}$ with $\rho_T = (\mathbf{x}_{1:T+1}, \mathbf{u}_T)$, the following both hold with probability one:

- Each action a_h satisfies our notion of incremental stability (Definition 3.1) with moduli $\gamma(\cdot)$, $\beta(\cdot, \cdot)$, constants $c_\gamma, c_\xi > 0$ (i.e. Assumption 3.1b holds)

Finally, let $\varepsilon_0 > 0$ satisfy (E.1), that is:

$$\gamma^{-1}(\beta(2\gamma(\varepsilon_0), \tau_{\text{chunk}})) \leq \varepsilon_0 \leq \min\{c_\gamma, \gamma^{-1}(c_\xi/4)\},$$

Further, given $\tilde{s} \in \mathcal{S} = \mathcal{P}_{\tau_{\text{chunk}}}$ with last-step $\tilde{x}_{\tau_{\text{chunk}}}$, consider the distance-like function

$$\bar{d}_{\mathcal{A}}(a, a'; \alpha, \tilde{s}) := \psi(d_{\text{loc},\alpha}(a, a' \mid \tilde{x}_{\tau_{\text{chunk}}})) \cdot \mathcal{I}_\infty \{d_{\text{loc},\alpha}(a, a' \mid \tilde{x}_{\tau_{\text{chunk}}}) \leq \varepsilon_0\}, \quad \psi(u) := 2\beta(2\gamma(u), 0).$$

Then, the following hold:

- π^* is state-conditioned input-stable (Definition M.1) with respect to $d_{\mathcal{S}}, d_{\text{TVC}}$ as defined in Appendix D

$$d_{\mathcal{A}}(a, a'; s) = \bar{d}_{\mathcal{A}}(a, a'; \psi(\varepsilon), s),$$

- For any $r_{\text{IPS}} \leq c_\xi/2$, π^* is $(r_{\text{IPS}}, \gamma_{\text{IPS},\text{TVC}}, \gamma_{\text{IPS},\mathcal{S}}, d_{\text{IPS}})$ -state-conditioned restricted-IPS (Definition M.3) with

$$d_{\mathcal{A}}(a, a'; s) = \bar{d}_{\mathcal{A}}(a, a'; \psi(\varepsilon) + \beta(r_{\text{IPS}}, 0), s), \quad \gamma_{\text{IPS},\text{TVC}}(r) = \beta(r, \tau_{\text{chunk}} - \tau_{\text{obs}}), \quad \gamma_{\text{IPS},\mathcal{S}}(r) = \beta(r, 0),$$

Proof of Lemma M.2. The proof is nearly identical to that of Lemma K.3, based on Lemma K.2. The only difference is that, rather than using the worst-case bound $\mathbf{x}_t \in \mathcal{X}_0$, we condition on the relevant states. For part (a), we consider $(s_{1:H+1}, a_{1:H})$ be drawn from the distribution induces by π^* , and let $a'_{1:H}$ be some other sequences of actions, and measure $d_{A,S}(a_h, a_h; s_h)$. Thus the relevant control-state to condition on is \mathbf{x}_{t_h} in the construction of Lemma K.2. For verifying (b), we instead condition on $\tilde{\mathbf{x}}_{t_h}$ because, as in Definition M.3, we measure the input-state stability condition for restricted-state-conditioned-IPS with sequence to the states $\tilde{s}_{1:H}$. \square

M.3 Concluding the proof of Theorems 5 and 6

M.3.1 Proof of Theorem 6

The result is a direct consequence of the following points. First, with our instantiation of the composite MDP, we can bound $\mathcal{L}_{\text{marg},\varepsilon}(\hat{\pi}) \leq \Gamma_{\text{marg},S,\varepsilon}(\hat{\pi} \parallel \pi^*) \leq \Gamma_{\text{joint},S,\varepsilon}(\hat{\pi} \parallel \pi^*)$ by the same argument in Lemma J.1¹⁷; by a similar argument, we have $\mathcal{L}_{\text{joint},\varepsilon}(\hat{\pi}) \leq \Gamma_{\text{joint},S,\varepsilon}(\hat{\pi} \parallel \pi^*)$ when \mathcal{D}_{exp} has $\tau \leq \tau_{\text{obs}}$ -bounded memory. The bound now follows from Proposition M.1, the fact that Lemma M.2 verifies the input-stability property (with $\varepsilon, \tau_{\text{chunk}}$ satisfies (E.1)). \square

M.3.2 Proof of Theorem 5

We begin with a lemma that upper bounds the imitation gaps by $\Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha(\varepsilon) + 2\beta(2r, 0))$ and other relevant terms. Essentially, the following lemma combines the general imitation guarantee in Theorem 15 with the incremental stability analysis in Lemmas K.2 and M.2.

Lemma M.3. Consider the instantiation of the composite MDP as in Definition K.1, let $r \leq c_\xi/4$, and recall $\alpha(\varepsilon) = 2\beta(2\gamma(\varepsilon), 0)$. Further suppose that ε satisfy Eq. (E.1). Then, the the modified imitation gaps (whose definition we recall Definition M.2) $\Gamma_{\text{joint},S,\alpha(\varepsilon)}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$ and $\Gamma_{\text{marg},S,\alpha(\varepsilon)+\beta(2r,0)}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ are bounded above by

$$H(4p_r + 3\gamma_\sigma(\max\{\alpha(\varepsilon), \beta(2r, \tau_{\text{chunk}} - \tau_{\text{obs}})\})) + \sum_{h=1}^H \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha(\varepsilon) + 2\beta(2r, 0)).$$

Before proving the lemma, lets quickly show how it implies the desired theorem. We bound

$$\begin{aligned} \mathcal{L}_{\text{marg},2\beta(2\gamma(\varepsilon),0)+2\beta(2r,0)} &\leq \mathcal{L}_{\text{marg},2\beta(2\gamma(\varepsilon),0)+\beta(2r,0)} \\ &= \mathcal{L}_{\text{marg},\alpha(\varepsilon)+\beta(2r,0)}(\hat{\pi}) \leq \Gamma_{\text{marg},S,\alpha(\varepsilon)+\beta(2r,0)}(\hat{\pi} \circ W_\sigma \parallel \pi^*), \end{aligned}$$

where the last inequality is due to as in the proof of Theorem 6, the intermediate inequality uses the definition of $\alpha(\varepsilon)$, and the first inequality uses anti-monotonicity of $\mathcal{L}_{\text{marg},\varepsilon}$ in ε . Moreover, as shown in the proof of Theorem 2 in (J.4) and (J.5), we can take $\gamma_\sigma(u) = \frac{u\sqrt{2\tau_{\text{obs}}-1}}{2\sigma}$ and for $p_r \leq p$

when $r = \sigma\omega_p$, $\omega_p := 2\sqrt{5d_x + 2\log\left(\frac{1}{p}\right)}$ and $W_\sigma(\cdot)$ is the Gaussian Kernel in (J.1). Hence, we conclude that if $\sigma \leq c_\xi/4\omega_p$,

$$\mathcal{L}_{\text{marg},\varepsilon_1(p)} \leq H\left(4p + \frac{3\sqrt{2\tau_{\text{obs}}-1}}{2\sigma}(\max\{\varepsilon_2, \beta(2\sigma\omega_p, \tau_{\text{chunk}} - \tau_{\text{obs}})\})\right) + \sum_{h=1}^H \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \varepsilon_1(p)),$$

where above $\varepsilon_1(p) = 2\beta(2\gamma(\varepsilon), 0) + 2\beta(2\sigma\omega_p, 0)$ and $\varepsilon_2 = 2\beta(2\gamma(\varepsilon), 0)$, as needed. Since $\gamma(\varepsilon) \leq 2\sigma$, in we can choose $p = \frac{\gamma(\varepsilon)}{\sigma^2}$ and upper bound $\omega_p \leq \omega(\varepsilon) := 2\sqrt{5d_x + 2\log\left(\frac{2\sigma}{\gamma(\varepsilon)}\right)}$.

The bound now follows from this upper bound and the bound $4p = 4\frac{2\gamma(\varepsilon)}{\sigma^8} \leq 4\frac{2\beta(2\gamma(\varepsilon), 0)}{\sigma^8} \leq \frac{\varepsilon_2}{2\sigma}$, the first inequality follows from Observation K.4. \square

¹⁷Here, $\Gamma_{\text{marg},S,\varepsilon}, \Gamma_{\text{joint},S,\varepsilon}$ are defined in in Definition M.2. The only difference between these the standard gaps $\Gamma_{\text{marg},\varepsilon}, \Gamma_{\text{joint},\varepsilon}$ consider in Definition M.2 is that they drop the closeness on composite actions, which is immaterial for $\mathcal{L}_{\text{marg},\varepsilon}(\hat{\pi})$.

Proof of Lemma M.3. Recall the replica and deconvolution kernels $W_{\text{dec},h}^*(\cdot), W_{\text{rep},h}(\cdot)$ defined in Definition H.5. We have that

$$\mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},S,\alpha}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec},h}^*(\tilde{s}_h^*) \mid s_h^*) = \mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} \inf_{\mu} \mathbb{P}_{\mu}[d_{\mathcal{A};S}(a', a \mid s_h^*) > \alpha] \quad (\text{M.3})$$

where \inf_{μ} is over all couplings $a_h \sim \pi_{\text{dec},h}^*(\tilde{s}_h^*), a'_h \sim \hat{\pi}(\tilde{s}_h^*)$. By the gluing lemma (Lemma F.2), each coupling in μ is equivalent to a coupling $\hat{\mu}$ over $(s_h^*, \tilde{s}_h^*, \hat{s}_h^*, a, a')$ where

- $s_h^* \sim P_h^*, \tilde{s}_h^* \sim W_\sigma(s_h^*)$
- $\hat{s}_h^* \mid \tilde{s}_h^* \sim W_{\text{dec},h}^*(\tilde{s}_h^*)$ for the deconvolution kernel defined in
- $a_h \sim \pi_h^*(\hat{s}_h^*)$ and $a'_h \sim \hat{\pi}(\tilde{s}_h^*)$

For couplings $\hat{\mu}$ of this form, and for $r > 0$, then, we can bound (M.3) via

$$\begin{aligned} & \inf_{\hat{\mu}} \mathbb{P}_{\hat{\mu}}[d_{\mathcal{A};S}(a'_h, a_h \mid s_h^*) \leq \alpha] \\ &= \inf_{\hat{\mu}} \mathbb{E}_{\hat{\mu}} \mathbf{I}\{d_{\mathcal{A};S}(a'_h, a_h \mid s_h^*) \leq \alpha\} \\ &= \inf_{\hat{\mu}} \mathbb{E}_{\hat{\mu}} \left[\mathbf{I} \left\{ \sup_{\hat{s}: d_{\text{IPS}}(s, \hat{s}_h^*) \leq 2r} d_{\mathcal{A};S}(a'_h, a_h \mid \hat{s}) > \alpha \right\} + \mathbf{I}\{d_{\text{IPS}}(s_h^*, \hat{s}_h^*) > 2r\} \right]. \end{aligned}$$

Because for any $\hat{\mu}$, $\hat{s}_h^* \mid \tilde{s}_h^* \sim W_{\text{dec},h}^*(\tilde{s}_h^*)$, we see that the joint distribution s_h^*, \hat{s}_h^* is independent of the coupling $\hat{\mu}$ and follows the replica distribution: $\hat{s}_h^* \mid s_h^* \sim W_{\text{rep},h}(s_h^*)$. Consequently, by the Bayesian concentration lemma Lemma H.5, the expected value of the term $\mathbf{I}\{d_{\text{IPS}}(s_h^*, \hat{s}_h^*) > 2r\}$ is at most $2p_r$. Hence,

$$\mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},S,\alpha}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec},h}^*(\tilde{s}_h^*) \mid s_h^*) \leq 2p_r + \inf_{\hat{\mu}} \mathbb{E}_{\hat{\mu}} \left[\mathbf{I} \left\{ \sup_{\hat{s}: d_{\text{IPS}}(s, \hat{s}_h^*) \leq 2r} d_{\mathcal{A};S}(a'_h, a_h \mid \hat{s}) > \alpha \right\} \right]$$

Again, marginalizing over s_h^* and using the form of the conditions, the right hand side of the above

$$\mathbb{E}_{s_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(s_h^*)} d_{\text{os},S,\alpha}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec},h}^*(\tilde{s}_h^*) \mid s_h^*) \leq 2p_r + \inf_{\hat{\mu}} \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{s}: d_{\text{IPS}}(s, \hat{s}_h^*) \leq 2r} d_{\mathcal{A};S}(a'_h, a_h \mid \hat{s}) > \alpha \right] \quad (\text{M.4})$$

Next, we recall the function $\psi(u) := 2\beta(2\gamma(u), 0)$, instantiate $\alpha = \psi(\varepsilon)$, $\alpha' = \alpha + \beta(2r, 0)$, $r_{\text{IPS}} = 2r$, and set $d_{\mathcal{A};S}(a'_h, a_h \mid s)$ to be

$$d_{\mathcal{A};S}(a'_h, a_h \mid s) = \psi(d_{\text{loc},\alpha}(a, a' \mid \tilde{\mathbf{x}}_{\tau_{\text{chunk}}})) \cdot \mathcal{I}_{\infty} \{d_{\text{loc},\alpha}(a, a' \mid \tilde{\mathbf{x}}_{\tau_{\text{chunk}}}) \leq \varepsilon_0\},$$

Using that d_{IPS} measures the Euclidean distance between the last control state of composite-state, we

$$\begin{aligned} \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{s}: d_{\text{IPS}}(s, \hat{s}_h^*) \leq 2r} d_{\mathcal{A};S}(a'_h, a_h \mid \hat{s}) > \alpha \right] &= \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{s}: d_{\text{IPS}}(\hat{s}, \hat{s}_h^*) \leq 2r} d_{\text{loc},\alpha'}(a_h, a'_h \mid \hat{\mathbf{x}}_{t_h}) > \varepsilon \right] \\ &= \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{\mathbf{x}}_{t_h}: \|\mathbf{x}_{t_h} - \hat{\mathbf{x}}_{t_h}\| \leq 2r} d_{\text{loc},\alpha'}(a_h, a'_h \mid \hat{\mathbf{x}}_{t_h}) > \varepsilon \right] \\ &= \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{\mathbf{x}}_{t_h}: \|\mathbf{x}_{t_h} - \hat{\mathbf{x}}_{t_h}\| \leq 2r} \max_{0 \leq i < \tau_{\text{chunk}}} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha'} \|(\kappa_{t_h+i} - \kappa'_{t_h+i})(\hat{\mathbf{x}}_{t_h+i})\| > \varepsilon \right] \\ &\leq \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{\mathbf{x}}_{t_h}: \|\mathbf{x}_{t_h} - \hat{\mathbf{x}}_{t_h}\| \leq 2r} \max_{0 \leq i < \tau_{\text{chunk}}} \sup_{\delta \mathbf{x}: \|\delta \mathbf{x}\| \leq \alpha''} \|(\kappa_{t_h+i} - \kappa'_{t_h+i})(\mathbf{x}_{t_h+i})\| > \varepsilon \right] \\ &\quad \text{(the inequality just replaces } \alpha' \text{ with } \alpha'') \\ &= \mathbb{P}_{\hat{\mu}} \left[\sup_{\hat{\mathbf{x}}_{t_h}: \|\mathbf{x}_{t_h} - \hat{\mathbf{x}}_{t_h}\| \leq 2r} d_{\text{loc},\alpha''}(a_h, a'_h \mid \mathbf{x}_{t_h}) > \varepsilon \right], \quad (\text{M.5}) \end{aligned}$$

$\hat{\mathbf{x}}_{t_h}$ is the first state in \hat{s} , and \mathbf{x}_{t_h} the first state in \hat{s}_h^* , $\hat{\mathbf{x}}_{t_h:t_h+\tau_{\text{chunk}}-1} = \text{rollout}(\mathbf{a}_h; \hat{\mathbf{x}}_{t_h})$, $\hat{\mathbf{x}}_{t_h:t_h+\tau_{\text{chunk}}-1}^* = \text{rollout}(\mathbf{a}_h; \hat{\mathbf{x}}_{t_h}^*)$, and finally,

$$\alpha'' := \underbrace{\alpha'}_{=\alpha+\beta(r_{\text{ips}},0)} + \Delta, \quad \Delta := \sup_{\hat{\mathbf{x}}_{t_h} : \|\mathbf{x}_{t_h} - \hat{\mathbf{x}}_{t_h}\| \leq 2r} \sup_{0 \leq i < \tau_{\text{chunk}}} \|\hat{\mathbf{x}}_{t_h:t_h+\tau_{\text{chunk}}-1} - \mathbf{x}_{t_h:t_h+\tau_{\text{chunk}}-1}\|. \quad (\text{M.6})$$

Now, we see that $\hat{\mu}$ ranges all couplings of the form

- $\hat{s}_h^* \sim P_h^*$ and $\tilde{s}_h^* \sim W_\sigma(\hat{s}_h^*)$ (by inverting the deconvolution)
- $\mathbf{a}_h' \sim \tilde{s}_h^*$ and $\mathbf{a}_h \sim \pi_h^*(\hat{s}_h^*)$,

which we can see (under our instantiation of the composite MDP under [Appendices D and J](#)) is equivalently to $\hat{\mu}$ ranging over all couplings in $\bar{\mathcal{C}}_{\sigma,h}(\hat{\pi})$. Hence, by [Assumption 3.1b](#) (i.e. t-ISS of \mathbf{a}_h at \mathbf{a}_h), we can bound Δ in (M.6) (using $r \leq c_\xi/4$) by $\Delta \leq \beta(2r, 0)$. Hence, we can bound $\alpha'' \leq \alpha + 2\beta(2r, 0)$, and thus we conclude (from (M.5) and (M.4)) that

$$\begin{aligned} \mathbb{E}_{\hat{s}_h^* \sim P_h^*} \mathbb{E}_{\tilde{s}_h^* \sim W_\sigma(\hat{s}_h^*)} \mathbf{d}_{\text{os},S,\alpha}(\hat{\pi}_h(\tilde{s}_h^*) \parallel \pi_{\text{dec},h}^*(\tilde{s}_h^*) \mid \mathbf{s}_h^*) &\leq 2p_r + \inf_{\hat{\mu} \in \bar{\mathcal{C}}_{\sigma,h}(\hat{\pi})} \mathbb{P}_{\hat{\mu}}[\mathbf{d}_{\text{loc},\alpha''}(\mathbf{a}_h, \mathbf{a}_h' \mid \text{rollout}(\mathbf{a}_h; \mathbf{x}_{t_h})) > \varepsilon] \\ &= 2p_r + \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha + 2\beta(2r, 0)), \end{aligned}$$

where the last equality is by definition of $\Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha + 2\beta(2r, 0))$. Consequently, from [Theorem 15](#), for any policy $\hat{\pi}$, both $\Gamma_{\text{joint},S,\alpha}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*)$ and $\Gamma_{\text{marg},S,\alpha+\gamma_{\text{ips},S}(2r)}(\hat{\pi} \circ W_\sigma \parallel \pi^*)$ are upper bounded by

$$H(4p_r + 3\gamma_\sigma(\max\{\alpha, \gamma_{\text{ips},\text{TVC}}(2r)\})) + \sum_{h=1}^H \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha + 2\beta(2r, 0)).$$

Substituting in $\gamma_{\text{ips},\text{TVC}}(r) = \beta(r, \tau_{\text{chunk}} - \tau_{\text{obs}})$, $\gamma_{\text{ips},S}(r) = \beta(r, 0)$, we conclude that

$$\begin{aligned} &\Gamma_{\text{joint},S,\alpha}(\hat{\pi} \circ W_\sigma \parallel \pi_{\text{rep}}^*) \vee \Gamma_{\text{marg},S,\alpha+\beta(2r,\tau_{\text{chunk}}-\tau_{\text{obs}})}(\hat{\pi} \circ W_\sigma \parallel \pi^*) \\ &\leq H(4p_r + 3\gamma_\sigma(\max\{\alpha, \beta(2r, 0)\})) + \sum_{h=1}^H \Delta_{\text{ISS},\sigma,h}(\hat{\pi}; \varepsilon, \alpha + 2\beta(2r, 0)). \end{aligned}$$

Substituting in $\gamma_{\text{ips},\text{TVC}}(r) \leq \beta(r, \tau_{\text{chunk}} - \tau_{\text{obs}})$, $\gamma_{\text{ips},S}(r) \leq \beta(r, 0)$ from [Lemma M.2](#), as well as $\alpha = \psi(\varepsilon) = 2\beta(2\gamma(\varepsilon), 0)$ concludes. \square

N Extensions and Further Results

N.1 Removing the necessity for minimal chunk length via stronger synthesis oracle

Theorem 16. Support we replace [Assumption 3.1](#) in [Theorem 1](#) with the assumption that our trajectory oracle produces *entire sequences of gains* $\kappa_{1:T}$ which satisfy time-varying incremental stability ([Definition 3.1](#)) on the whole trajectory. Then,

- The conclusion of [Theorem 1](#) holds
- we no longer need the condition $\tau_{\text{chunk}} \geq c_3$; taking $\tau_{\text{chunk}} = 1$ suffices.
- The constants c_1, c_2 depend only on c_γ and \bar{c}_γ . That is, c_ξ and terms associated with β can be vacuously large.

Analogous, if we replace [Assumption 3.1b](#) in [Theorem 6](#) with the assumption that $\kappa_{1:T}$ satisfies the time-varying incremental stability condition, then

- The conclusion of [Theorem 6](#) holds
- We no longer need the condition $\tau_{\text{chunk}} \geq c_3$; taking $\tau_{\text{chunk}} = 1$ suffices. Moreover, we can replace the condition (E.1) of ε with the simpler condition $\varepsilon \leq c_\gamma$.

- Lastly, one can replace $\varepsilon_1 = 2\beta(2\gamma(\varepsilon), 0)$ in (E.2) with the term $\varepsilon_1 = \gamma(\varepsilon)$.

The proof of [Theorem 16](#) follows by replacing [Lemma K.2](#) with the following simpler lemma that recapitulates Pfrommer et al. [50, Proposition 3.1], and propagating the argument through the proof.

Lemma N.1. Consider two consistent trajectories $(\mathbf{x}_{1:T+1}, \mathbf{u}_{1:T})$ and $(\mathbf{x}'_{1:T+1}, \mathbf{u}'_{1:T})$, as well as sequences of primitive controller $\kappa_{1:T}, \kappa'_{1:T}$, such that $\mathbf{x}_1 = \mathbf{x}'_1$, and $\mathbf{u}_t = \kappa_t(\mathbf{x}_t)$, $\mathbf{u}'_t = \kappa'_t(\mathbf{x}'_t)$. Suppose that

$$\max_t \sup_{\mathbf{x}: \|\mathbf{x} - \mathbf{x}_t\| \leq \gamma(\varepsilon)} \|\kappa_t(\mathbf{x}) - \kappa'_t(\mathbf{x}_t)\| \leq \varepsilon.$$

Then, $\max_t \|\mathbf{u}_t - \mathbf{u}'_t\| \leq \varepsilon$ and $\max_t \|\mathbf{x}_t - \mathbf{x}'_t\| \leq \gamma(\varepsilon)$

N.2 Noisy Dynamics

We can directly extend our imitation guarantees in the composite MDP to settings with noise:

$$\mathbf{s}_{h+1} \sim F_h^{\text{noise}}(\mathbf{s}_h, \mathbf{a}_h, \mathbf{w}_h), \quad \mathbf{w}_h \sim P_{\text{noise},h}, \quad (\text{N.1})$$

where the noises are independent of states and of each other. Indeed, (N.1) can be directly reduced to the no-noise setting by lifting “actions” to pairs $(\mathbf{a}_h, \mathbf{w}_h)$, and policies π to encompass their distribution of actions, and over noise.

Another approach is instead to condition on the noises $\mathbf{w}_{1:H}$ first, and treat the noise-conditioned dynamics as deterministic. Then one can take expectation over the noises and conclude. The advantage of this approach is that the couplings constructed thereby is that the trajectories experience identical sequences of noise with probability one.

Extending the control setting to incorporate noise is doable but requires more effort:

- If the *demonstrations are noiseless*, then one can still appeal to the synthesis oracle to synthesis stabilizing gains. However, one needs to (ever so slightly) generalize the proofs of the various stability properties (e.g. IPS in [Proposition D.1](#)) to accomodate system noise.
- If the demonstrations themselves have noise, one may need to modify the synthesis oracle setup somewhat. This is because the synthesis oracle, if it synthesizes stabilizing gains, will attempt to get the learner to stabilize to a noise-perturbed trajectory. This can perhaps be modified by synthesizing controllers which stabilize to smoothed trajectories, or by collecting demonstrations of desired trajectories (e.g. position control), and stabilizing to the these states than than to actual states visited in demonstrations.

N.3 Robustness to Adversarial Perturbations

Our results can accomodate an even more general framework where there are both noises as well adversarial perturbations. We explain this generalization in the composite MDP.

Specifically, consider a space \mathcal{E} of adversarial perturbations, as well as \mathcal{W} of noises as above. We may posit a dynamics function $F^{\text{adv}} : \mathcal{S} \times \mathcal{A} \times \mathcal{W} \times \mathcal{A} \rightarrow \mathcal{S}$, and consider the evolution of an imitator policy $\hat{\pi}$ under the adversary

$$\begin{aligned} \hat{\mathbf{s}}_{h+1} &= F_h^{\text{adv}}(\hat{\mathbf{s}}_h, \hat{\mathbf{a}}_h, \mathbf{w}_h, \mathbf{e}_h), \quad \mathbf{w}_h \sim P_{\text{noise},h} \\ \hat{\mathbf{a}}_h &\sim \hat{\pi}_h(\mathbf{s}_h) \\ \mathbf{e}_h &\sim \pi_h^{\text{adv}}(\hat{\mathbf{s}}_{1:h}, \mathbf{a}_{1:h}, \mathbf{w}_{1:h}, \mathbf{e}_{1:h-1}), \\ \hat{\mathbf{s}}_1 &\sim \pi_0^{\text{adv}}(\mathbf{s}_1), \quad \mathbf{s}_1 \sim P_{\text{init}}. \end{aligned}$$

By contrast, we can model the demonstrator trajectory as arising from noisy, but otherwise unperturbed trajectories:

$$\mathbf{s}_{h+1}^* \sim F_h^{\text{adv}}(\mathbf{s}_h^*, \mathbf{a}_h^*, \mathbf{w}_h, 0), \quad \mathbf{w}_h \sim P_{\text{noise},h}, \quad \mathbf{a}_h^* \sim \pi_h^*(\mathbf{s}_h^*), \quad \mathbf{s}_1^* \sim P_{\text{init}}.$$

To reduce the composite-MDP in [Appendix D](#), we can view the combination of adversary π^{adv} and imitator $\hat{\pi}$ as a combined policy, and the π^* with zero augmentation as another policy; here, we

would then treat actions as $\tilde{a} = (a, e)$. Then, one can consider modified senses of stability which preserve trajectory tracking, as well as a modification of d_A to a function measuring distances between $\tilde{a} = (a, e)$ and $\tilde{a}' = (a', e')$. The extension is rather mechanical, and we fit details. Note further that, by including a $\pi_0^{\text{adv}}(s_1)$, we can modify the analysis to allow for subtle differences in initial state distribution. This would in turn require strengthening our stability assumptions to allow stability to initial state (e.g., the definition of incremental stability as exposted by [50]).

N.4 Deconvolution Policies and Total Variation Continuity

While our strongest guarantees hold for the replica policies, where we add noise both as a data augmentation at training time *and* at test time, many practitioners have seen some success with the deconvolution policies where noise is only added at training time. We note that [Proposition D.2](#) holds when the learned policy is TVC; without noise at training time this certainly will not hold when the expert policy is not TVC. We show here that the deconvolution expert policy is TVC under mild assumptions, which lends some credence to the empirical success of deconvolution policies.

Precisely, we show that, under reasonable conditions, deconvolution is total variation continuous. In particular, suppose that $\mu \in \Delta(\mathbb{R}^d)$ is a Borel probability measure and p is a density with respect to μ . Further suppose that Q is a density with respect to the Lebesgue measure on \mathbb{R}^d . Suppose that $\mathbf{x} \sim p$, $\mathbf{w} \sim Q$, and let $\tilde{\mathbf{x}} = \mathbf{x} + \mathbf{w}$. Denote the deconvolution measure of \mathbf{x} given $\tilde{\mathbf{x}}$ as $p(\cdot|\tilde{\mathbf{x}})$. We show that this measure is continuous in TV.

Proposition N.2. Let $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ be fixed, let $p : \mathbb{R}^d \rightarrow \mathbb{R}$ denote a probability density, and let $Q : \mathbb{R}^d \rightarrow \mathbb{R}$ denote a function such that $\nabla^2 Q$ and $\nabla \log Q$ exist and are continuous on the set

$$\mathcal{X} = \{(1-t)\tilde{\mathbf{x}} + t\tilde{\mathbf{x}}' - \mathbf{x} | \mathbf{x} \in \text{supp } p \text{ and } t \in [0, 1]\}$$

Then it holds that

$$\text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}')) \leq \|\tilde{\mathbf{x}} - \tilde{\mathbf{x}}'\| \cdot \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla \log Q(\mathbf{x})\|.$$

By [Lemma F.4](#), any policy composed with the total variation kernel is thus total variation continuous with a linear γ_{TVC} ; moreover, the Lipschitz constant is given by the maximal norm of the score of the noise distribution. For example, if Q is the density of a Gaussian with variance σ^2 , then $\gamma_{\text{TVC}}(u) \leq \frac{\sup_{\mathbf{x}} \|\mathbf{x}\|}{\sigma^2}$ is dimension independent.

Remark N.1. Note that our notation is intentionally different from that in the body to emphasize that this is a general fact about abstract probability measures. We may instantiate the guarantee in the control setting of interest by letting $\mathbf{x} = \mathbf{o}_h$ and consider Q to be a Gaussian (for example) kernel. In this case, we see that the deconvolution policy of [Definition C.1](#) is automatically TVC.

To prove [Proposition N.2](#), we begin with the following lemma:

Lemma N.3. Let $\tilde{\mathbf{x}} \in \mathbb{R}^d$ be fixed and suppose that $\nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x})$ exists for all $\mathbf{x} \in \text{supp } p$. Then, for all $\mathbf{x} \in \text{supp } p$, it holds that $\nabla_{\tilde{\mathbf{x}}} p(\mathbf{x}|\tilde{\mathbf{x}})$ exists. Furthermore,

$$\int \|\nabla p(\mathbf{x}|\tilde{\mathbf{x}})\| d\mu(\mathbf{x}) \leq 2 \sup_{\mathbf{x} \in \text{supp } p} \|\nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x})\|,$$

where the gradient above is with respect to $\tilde{\mathbf{x}}$.

Proof. We begin by noting that if $\nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x})$ exists, then so does $\nabla Q(\tilde{\mathbf{x}} - \mathbf{x})$. By Bayes' rule,

$$p(\mathbf{x}|\tilde{\mathbf{x}}) = \frac{p(\mathbf{x})Q(\tilde{\mathbf{x}} - \mathbf{x})}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')}.$$

We can then compute directly that

$$\nabla p(\mathbf{x}|\tilde{\mathbf{x}}) = \frac{p(\mathbf{x})\nabla Q(\tilde{\mathbf{x}} - \mathbf{x})}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} - \frac{p(\mathbf{x})Q(\tilde{\mathbf{x}} - \mathbf{x}) \cdot \int \nabla Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')}{\left(\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')\right)^2},$$

where the exchange of the gradient and the integral is justified by Lebesgue dominated convergence and the assumption of differentiability of Q and thus existence is ensured. We have now that

$$\begin{aligned}
\|\nabla p(\mathbf{x}|\tilde{\mathbf{x}})\| &= \frac{p(\mathbf{x})Q(\tilde{\mathbf{x}} - \mathbf{x})}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} \cdot \left\| \nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x}) - \frac{\int \nabla Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} \right\| \\
&= \frac{p(\mathbf{x})Q(\tilde{\mathbf{x}} - \mathbf{x})}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} \cdot \left\| \nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x}) - \frac{\int (\nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x}')) \cdot Q(\tilde{\mathbf{x}} - \mathbf{x})p(\mathbf{x}')d\mu(\mathbf{x}')}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} \right\| \\
&\leq \left(\sup_{\mathbf{x} \in \text{supp } p} \|\nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x})\| \right) \cdot \frac{p(\mathbf{x})Q(\tilde{\mathbf{x}} - \mathbf{x})}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} \cdot \left(1 + \frac{\int Q(\tilde{\mathbf{x}} - \mathbf{x})p(\mathbf{x}')d\mu(\mathbf{x}')}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} \right) \\
&= \left(2 \sup_{\mathbf{x} \in \text{supp } p} \|\nabla \log Q(\tilde{\mathbf{x}} - \mathbf{x})\| \right) \cdot \frac{p(\mathbf{x})Q(\tilde{\mathbf{x}} - \mathbf{x})}{\int Q(\tilde{\mathbf{x}} - \mathbf{x}')p(\mathbf{x}')d\mu(\mathbf{x}')} .
\end{aligned}$$

Now, integrating over \mathbf{x} makes the second factor 1, concluding the proof. \square

We will now make use of the theory of Dini derivatives ([25]) to prove a bound on total variation.

Lemma N.4. For fixed $\tilde{\mathbf{x}}, \tilde{\mathbf{x}}'$ and $0 \leq t \leq 1$, let the upper Dini derivative

$$D^+ \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t)) = \limsup_{h \downarrow 0} \frac{\text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_{t+h})) - \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t))}{h},$$

where

$$\tilde{\mathbf{x}}_t = (1-t)\tilde{\mathbf{x}} + t\tilde{\mathbf{x}}'.$$

If $\nabla \log Q(\tilde{\mathbf{x}}_t - \mathbf{x})$ exists and is finite for all $\mathbf{x} \in \text{supp } p$ and $t \in [0, 1]$, then

$$\text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}')) \leq \int_0^1 D^+ \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t)) dt. \quad (\text{N.2})$$

Proof. We compute:

$$\begin{aligned}
2 |\text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_{t+h})) - \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t))| &= \left| \int |p(\mathbf{x}|\tilde{\mathbf{x}}) - p(\mathbf{x}|\tilde{\mathbf{x}}_{t+h})| - |p(\mathbf{x}|\tilde{\mathbf{x}}) - p(\mathbf{x}|\tilde{\mathbf{x}}_t)| d\mu(\mathbf{x}) \right| \\
&\leq \int |p(\mathbf{x}|\tilde{\mathbf{x}}_{t+h}) - p(\mathbf{x}|\tilde{\mathbf{x}}_t)| d\mu(\mathbf{x}). \quad (\text{N.3})
\end{aligned}$$

Observe that by the assumption on Q and Lemma N.3, $p(\mathbf{x}|\tilde{\mathbf{x}}_t)$ is differentiable and thus continuous in $\tilde{\mathbf{x}}_t$. We therefor see that the function

$$t \mapsto \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t))$$

is continuous as $\tilde{\mathbf{x}}_t$ is linear in t . By Hagood and Thomson [25, Theorem 10], (N.2) holds. \square

We now bound the Dini derivatives:

Lemma N.5. Let $\tilde{\mathbf{x}}, \tilde{\mathbf{x}}' \in \mathbb{R}^d$ such that for all $t \in [0, 1]$ it holds that

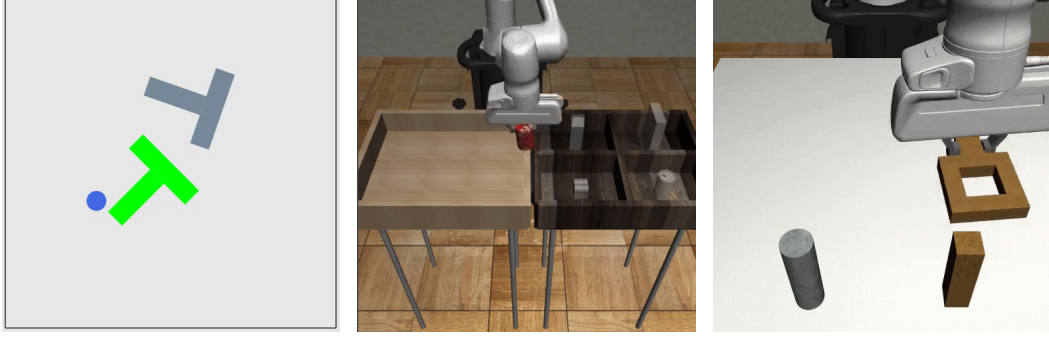
$$\sup_{\mathbf{x} \in \text{supp } p} \left| \frac{d^2}{dt^2} (p(\mathbf{x}|\tilde{\mathbf{x}}_t)) \right| = C < \infty,$$

where the derivative is applied on $\tilde{\mathbf{x}}_t$. If the assumptions of Lemmas N.3 and N.5 hold, then

$$D^+ \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t)) \leq \|\tilde{\mathbf{x}} - \tilde{\mathbf{x}}'\| \cdot \sup_{\substack{\mathbf{x} \in \text{supp } p \\ t \in [0, 1]}} \|\nabla \log Q(\tilde{\mathbf{x}}_t - \mathbf{x})\|.$$

Proof. By definition,

$$D^+ \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t)) = \limsup_{h \downarrow 0} \frac{\text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_{t+h})) - \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t))}{h}.$$



(a) PushT Environment [19]. The blue circle is the manipulation agent, while the green area is the target position which the agent must push the blue T block into. (b) Can Pick-and-Place Environment [42]. The grasper must pick up a can from the left bin and place it into the correct bin on the right side. (c) Square Nut Assembly Environment [42]. The grasper must pick up the square nut (the position of which is randomized) and place it over the square peg.

Figure 11: Environment Visualizations.

Fix some t and some small h . By (N.3), it holds that

$$|\text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_{t+h})) - \text{TV}(p(\cdot|\tilde{\mathbf{x}}), p(\cdot|\tilde{\mathbf{x}}_t))| \leq \frac{1}{2} \cdot \int |p(\mathbf{x}|\tilde{\mathbf{x}}_{t+h}) - p(\mathbf{x}|\tilde{\mathbf{x}}_t)| d\mu(\mathbf{x}).$$

By Taylor’s theorem, it holds that

$$p(\mathbf{x}|\tilde{\mathbf{x}}_{t+h}) - p(\mathbf{x}|\tilde{\mathbf{x}}_t) = h \cdot \frac{d}{dt} (p(\mathbf{x}|\tilde{\mathbf{x}}_t)) + h^2 \cdot \frac{d^2}{dt^2} (p(\mathbf{x}|\tilde{\mathbf{x}}_{t'}))$$

for some $t' \in [0, 1]$. By the chain rule, we have

$$\frac{d}{dt} (p(\mathbf{x}|\tilde{\mathbf{x}}_t)) = \langle \tilde{\mathbf{x}}' - \tilde{\mathbf{x}}, \nabla p(\mathbf{x}|\tilde{\mathbf{x}}_t) \rangle,$$

and thus,

$$|p(\mathbf{x}|\tilde{\mathbf{x}}_{t+h}) - p(\mathbf{x}|\tilde{\mathbf{x}}_t)| \leq h \cdot \|\tilde{\mathbf{x}} - \tilde{\mathbf{x}}'\| \cdot \|\nabla p(\mathbf{x}|\tilde{\mathbf{x}}_t)\| + h^2 C$$

Now, applying Lemma N.3 and plugging into the previous computation concludes the proof. \square

We are finally ready to state and prove our main result:

Proof of Proposition N.2. Note that

$$\frac{d^2}{dt^2} (p(\mathbf{x}|\tilde{\mathbf{x}}_t)) = (\tilde{\mathbf{x}} - \tilde{\mathbf{x}}')^T \nabla^2 p(\mathbf{x}|\tilde{\mathbf{x}}_t) (\tilde{\mathbf{x}} - \tilde{\mathbf{x}}')$$

and thus is bounded if and only if $\nabla^2 p(\mathbf{x}|\tilde{\mathbf{x}}_t)$ is bounded. An elementary computation shows that if $\nabla^2 Q$ exists and is continuous on \mathcal{X} , then $\nabla^2 p(\mathbf{x}|\tilde{\mathbf{x}}_t)$ is bounded in operator norm on \mathcal{X} . Thus the assumption in Lemma N.5 holds. Applying Lemma N.4 then concludes the proof. \square

O Experiment Details

O.1 Compute and Codebase Details

Code. For our experiments we build on the existing PyTorch-based codebase and standard environment set provided by Chi et al. [19] as well as the robomimic demonstration dataset Mandelkar et al. [42].¹⁸

¹⁸The modified codebase with instructions for running the experiments is available at the following anonymous link: https://www.dropbox.com/s/vzw0gvk1fd3yadw/diffusion_policy.zip?dl=0. We will provide a public github repository for the final release.

Compute. We ran all experiments using 4 Nvidia V100 GPUs on an internal cluster node. For each environment running all experiments depicted in Figure 6 took 12 hours to complete with 20 workers running simultaneously for a total of approximately 10 days worth of compute-hours. Between all 20 workers, peak system RAM consumption totaled about 500 GB.

O.2 Environment Details

For simplicity the stabilization oracle `synth` is built into the environment so that the diffusion policy effectively only performs positional control. See Appendix O for visualizations of the environments.

PushT. The PushT environment introduced in [19] is a 2D manipulation problem simulated using the PyMunk physics engine. It consists of pushing a T-shaped block from a randomized start position into a target position using a controllable circular agent. The synthesis oracle runs a low-level feedback controller at a 10 times higher to stabilize the agent’s position towards a desired target position at each point in time via acceleration control. Similar to Chi et al. [19], we use a position-error gain of $k_p = 100$ and velocity-error gain of $k_v = 20$. The observation provided to the DDPM model consists of the x,y coordinates of 9 keypoints on the T block in addition to the x,y coordinates of the manipulation agent, for a total observation dimensionality of 20.

For rollouts on this environment we used trajectories of length $T = 300$. Policies were scored based on the maximum coverage between the goal area and the current block position, with > 95 percent coverage considered an “successful” (score = 1) demonstration and the score linearly interpolating between 0 and 1 for less coverage. A total of 206 human demonstrations were collected, out of which we use a subset of 90 for training.

Can Pick-and-Place. This environment is based on the Robomimic [42] project, which in turn uses the MuJoCo physics simulator. For the low-level control synthesis we use the feedback controller provided by the Robomimic package. The position-control action space is 7 dimensional, including the desired end manipulator position, rotation, and gripper position, while the observation space includes the object pose, rotation in addition to position and rotation of all linkages for a total of 23 dimensions. Demonstrations are given a score of 1 if they successfully complete the pick-and-place task and a score of 0 otherwise. We roll out 400 timesteps during evaluation and for training use a subset of up to 90 of the 200 “proficient human” demonstrations provided.

Square Nut Assembly. For Square Nut Assembly, which is also Robomimic-based [42], we use the same setup as the Can Pick and Place task in terms of training data, demonstration scoring, and low-level positional controller. The observation, action spaces are also equivalent to the Can Pick-and-Place task with 23 and 7 dimensions respectively.

2D Quadcopter. The 2D quadcopter system is described by the state vector: $(x, z, \phi, \dot{x}, \dot{z}, \dot{\phi})$, with input $u = (u_1, u_2)$, and dynamics:

$$\begin{aligned}\ddot{x} &= -u_1 \sin(\phi)/m, \\ \ddot{z} &= u_1 \cos(\phi)/m - g, \\ \ddot{\phi} &= u_2/I_{xx}.\end{aligned}$$

The specific constants we use are $m = 0.8$, $g = 9.8$, and $I_{xx} = 0.5$. We integrate these dynamics using forward Euler with step size $\tau = 0.01$. The task is to move the quadcopter to the origin state. The cost function we used for the MPC expert is:

$$c((x, z, \phi, \dot{x}, \dot{z}, \dot{\phi}), (u_1, u_2)) = x^2 + z^2 + \phi^2 + \dot{x}^2 + \dot{z}^2 + \dot{\phi}^2 + 0.5(u_1 - mg)^2 + 0.1u_2^2.$$

We constructed a per-timestep reward function using this cost function:

$$r((x, z, \phi, \dot{x}, \dot{z}, \dot{\phi}), (u_1, u_2)) = \exp(-c((x, z, \phi, \dot{x}, \dot{z}, \dot{\phi}), (u_1, u_2))),$$

such that the MPC cost minimization corresponds to maximizing the reward used to benchmark the trained models.

O.3 Gain Synthesis

For the quadcopter gain-diffusion experiments, we synthesize stabilizing gains for each (\bar{x}_t, \bar{u}_t) pair in our training data by analytically differentiating the dynamics $x_{t+1} = f(x_t, u_t)$ given in O.2 at \bar{x}_t, \bar{u}_t and applying infinite-horizon LQR to the linearized system $x_{t+1} = A(x_t - \bar{x}_t) + B(u_t - \bar{u}_t)$ where $A = \partial_x f(\bar{x}_t, \bar{u}_t)$, $B = \partial_u f(\bar{x}_t, \bar{u}_t)$. In particular, we solve the discrete time algebraic Ricatti equation:

$$P = A^\top P A - (A^\top P B)(R + B^\top P B)^{-1} + Q,$$

where for simplicity we used identity matrices for R, Q . Using P we computed the gains:

$$K = -(R + B^\top P B)^{-1} B^\top P A.$$

Since the timesteps of the simulator are small, we experimentally find that this is sufficient in order to stabilize the system over the diffused chunks and produces significantly less variance in the gains than performing time-varying discrete LQR over the chunks to synthesize the gains.

O.4 DDPM Model and Training Details.

PushT, Can-Pick-and-Place, Square Nut Assembly. For these experiments we use the same 1-D convolutional UNet-style [56] architecture employed by [19], which is in turn adapted from Janner et al. [33]. This principally consists of 3 sets of downsampling 1-dimensional convolution operations using Mish activation functions [46], Group Normalization (with 8 groups) [77], and skip connections with 64, 128, and 256 channels followed by transposed convolutions and activations in the reversed order. The observation and timestep were provided to the model with Feature-wise Linear Modulation (FiLM) [49], with the timestep encoded using sin-positional encoding into a 64 dimensional vector.

During training and evaluation we utilize a squared cosine noise schedule [47] with 100 timesteps across all experiments. For training we use the AdamW optimizer with linear warmup of 500 steps, followed by an initial learning rate of 1×10^{-4} combined with cosine learning rate decay over the rest of the training horizon. For PushT models we train for 800 epochs and evaluate test trajectories every 200 epochs while for Can Pick-and-Place and Square Nut Assembly we evaluate performance every 250 epochs and train for a total of 1500 epochs.

The diffusion models are conditioned on the previous two observations trained to predict a sequence of 16 target manipulator positions, starting at the first timestep in the conditional observation sequence. The 2nd (corresponding to the target position for the current timestep) through 9th generated actions are emitted as the $\tau_c = 8$ length action sequence and the rest is discarded. Extracting a subsequence of a longer prediction horizon in this manner has been shown to improve performance over just predicting the $H = 8$ action sequence directly [19].

2D Quadcopter. For the 2D quadcopter experiments, we used a 5 layer MLP with hidden feature dimensions of 128, 128, 64, and 64 for all experiments, with sine-positional encoding of dimension 64 and FiLM to condition on the diffusion timestep and observation chunk. We use the same optimizer setup as the PushT, Pick-and-Place, and Square Assembly experiments with a batch size of 64 and a total of 200 epochs or 20,000 training iterations, whichever is larger.

We predict sequences of 8 control inputs, conditioned on two previous observations, where the 2nd control input in this sequence corresponds to the current timestep. For gain diffusion experiments, this includes a sequence of 8 control inputs, reference states, and gains. Similar to our other experiments, the 2nd through 5th generated actions of this sequence are emitted.

Augmentation Procedure. For $\sigma > 0$ we generate new perturbed observations per training iteration, effectively using $N_{\text{aug}} = N_{\text{epoch}}$ augmentations. We find this to be easier than generating and storing N_{aug} augmentations with little impact on the training and validation error. Noise is injected after the observations have been normalized such that all components lie within $[-1, 1]$ range. Performing noise injection post normalization ensures that the magnitude of noise injected is not affected by different units or magnitudes.