
Contrastive learning of strong-mixing continuous-time stochastic processes

Bingbin Liu

Pradeep Ravikumar
Carnegie Mellon University
{bingbinl,pradeepr,aristeski}@cs.cmu.edu

Andrej Risteski

Abstract

Contrastive learning is a family of self-supervised methods where a model is trained to solve a classification task constructed from unlabeled data. It has recently emerged as one of the leading learning paradigms in the absence of labels across many different domains (e.g. brain imaging, text, images). However, theoretical understanding of many aspects of training, both statistical and algorithmic, remain fairly elusive.

In this work, we study the setting of time series—more precisely, when we get data from a strong-mixing continuous-time stochastic process. We show that a properly constructed contrastive learning task can be used to estimate the transition kernel for small-to-mid-range intervals in the diffusion case. Moreover, we give sample complexity bounds for solving this task and quantitatively characterize what the value of the contrastive loss implies for distributional closeness of the learned kernel. As a byproduct, we illuminate the appropriate settings for the contrastive distribution, as well as other hyperparameters in this setup.

1 Introduction

One of the paradigms of learning from unlabeled data that has seen a lot of recent work in various application domains is “self-supervised learning”. These methods supervise the training process with information inherent to the data without requiring human annotations, and have been applied across computer vision, natural

language processing, reinforcement learning and scientific domains.

Despite the popularity, they are still not very well understood—both on the theoretical and empirical front—often requiring extensive trial and error to find the right pairing of architecture and learning method. In particular, it is often hard to pin down what exactly these methods are trying to learn, and it is even harder to determine what is their statistical and algorithmic complexity.

The specific family of self-supervised approaches we focus on in this work is *contrastive learning*, which constructs different types of tuples by utilizing certain structures in the data and trains the model to identify the types. For an example in vision, [Chen et al. \(2020\)](#) apply two random augmentations (e.g. crops and discolorations) on each training image, and form pairs that are labeled as either positive or negative depending on whether two augmentations are from the same image or not. In NLP, one of the tasks in [Devlin et al. \(2018\)](#); [Tosh et al. \(2020a\)](#) trains the model to predict whether two half-sentences are from the same original sentence.

In this paper, we focus on understanding a natural type of contrastive learning tasks for time series data—a natural structure in NLP ([Devlin et al., 2018](#); [Tosh et al., 2020a](#)), finance ([Ait-Sahalia et al., 2008](#)), and brain imaging research ([Hyvarinen and Morioka, 2016](#)) More precisely, we focus on data coming from a discretization of a *diffusion* process—a common modeling assumption in many of these domains—and show that a natural distinguishing task we set up on pairs of samples from the time series approximately learns the transition kernel of the stochastic process.

Note, a diffusion process is a continuous-time stochastic process and we are interested in learning transition kernels for “mid-range” time intervals, that is, intervals that are potentially too large for the Euler scheme to be accurate. These transition kernels are not easy to learn in general through standard maxi-

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

mum likelihood methods, as closed-form solutions are complicated (Ait-Sahalia et al., 2008) and often do not exist, and empirical estimations can also be challenging (Milstein et al., 2004). To our knowledge, our work is the first one to use contrastive learning to learn such transition kernels. Moreover, we provide a statistical complexity analysis—that is, analyzing the number of samples required to learn a good approximation of the transition kernel. This helps quantify certain aspects of contrastive learning—how should we choose the contrast distribution, and how a small loss on the contrastive task transfers to closeness of the transition kernel estimate.

2 Related Work

There is a large body of recent empirical work on self-supervised learning in general, which we won’t make an effort to survey in full, as it does not directly relate to our results.

There have been some recent works on trying to understand theoretically why and when self-supervised learning works. The closest ones in spirit to our work are Tosh et al. (2020b) and Hyvarinen and Morioka (2016), but there are significant differences with both. Tosh et al. (2020b) focus on a data distribution coming from LDA (topic modeling), and characterize the kinds of downstream classification tasks the learned predictor is useful for. Hyvarinen and Morioka (2016) focus on a time series setting as well but with several differences as highlighted below.

First, they work with a latent-variable model, and show that their method recovers some function of the latent. One example parametrization is an exponential family, and the function of the recovered latent variable depends on the choice of the exponential family.

Second, they assume the data in the time series can be subdivided into “blocks”, such that the distribution remains the same in each block and is sufficiently different from the others. In practice, it is not clear how to choose these “blocks” or how to even verify the assumptions needed on them. We do not need this “blocking”, but our data needs to come from the stationary distribution of the process.

Third, they do not provide an analysis on statistical complexity. In particular, important aspects of how various hyperparameters are chosen and affect the quality of the learned predictor—the size of the blocks, the amount of “difference” between the blocks—are not clear.

For temporally dependent and stationary data, another related work is by Hyvarinen and Morioka (2017). The setup is however different: Hyvarinen and

Morioka (2017) focus on discrete-time data with autocorrelations, whereas we analyze a continuous-time diffusion, leading to different setups and goals for the contrastive task. Moreover, in contrast to our finite sample analysis, their results describe only the asymptotic behaviors, which can hide certain statistical aspects of the algorithm as discussed earlier.

In the simpler iid setting, a classical precursor paper to this is by Gutmann and Hyvärinen (2010), who apply the contrastive learning approach to learning a distribution from iid samples—by setting up a classification task to distinguish between samples from the target distribution and a simple “contrast” distribution. Their analysis is again asymptotic and only provide sample efficiency bounds in the asymptotic limit (i.e. as the number of samples goes to infinity). More recently, such classical approaches have been combined with more modern generative models based approaches to generate better contrast distributions (Gao et al., 2020), and augmented with intermediate tasks to better handle dissimilar target and contrast distributions (Rhodes et al., 2020).

Finally, other papers on empirical and theoretical properties of contrastive learning that are worth mentioning include Purushwalkam and Gupta (2020); Tian et al. (2020) and Saunshi et al. (2019); Wang and Isola (2020)—these are not directly comparable to what we are doing here, as the data models are quite different, as is the flavor of guarantees they show. In particular, these papers work with iid data and focus on learning good representations that can perform well on certain supervised tasks, whereas we use contrastive learning to perform distribution learning, that is, learning the transition kernels.

3 Main Results

This section formally states our results. We will start with specifying the distributional model for the data and the contrastive learning task, and build intuitions on what the task aims to achieve before stating the formal guarantees.

3.1 Setup

We will assume our data comes from continuous time series: namely $\{\mathbf{x}_t\}_{t \geq 0} \subset \mathbb{R}^d$, drawn according to a stochastic process called the *Langevin diffusion*¹, de-

¹The results we state can more generally be stated about an Itô diffusion, namely a stochastic differential equation of the type $d\mathbf{x}_t = -g(\mathbf{x}_t)dt + \sigma(\mathbf{x}_t)dW_t$, $\forall t \geq 0$ satisfying similar regularity conditions to ours. We chose the simplest setting for clarity of exposition.

finied by the stochastic differential equation

$$d\mathbf{x}_t = -\nabla f(\mathbf{x}_t)dt + \sqrt{2}dW_t, \quad \forall t \geq 0, \quad (1)$$

for $f : \mathbb{R}^d \rightarrow \mathbb{R}$ a convex function, and $\{W_t\}_{t \geq 0} \subset \mathbb{R}^d$ a Wiener process, i.e. $W_s - W_t \sim \mathcal{N}(0, (s-t)\mathbf{I}_d)$, $\forall s \geq t \geq 0$. For the reader unfamiliar with diffusions, we can think of a diffusion process as the limit of a discrete sequence of noisy gradient updates with a fresh Gaussian noise: as $\eta \rightarrow 0$, the discrete sequence defined by $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \nabla f(\mathbf{x}_t) + \sqrt{2\eta} \xi_t$ where $\xi_t \sim \mathcal{N}(0, \mathbf{I})$ converges to the continuous time diffusion (Bhattacharya et al., 1978). The simplest instantiation of this, when f is quadratic (i.e. ∇f is linear), gives rise to the Ornstein–Uhlenbeck process, which has broad applications in science and finance modeling.

It is well-known (Bhattacharya et al., 1978) that the stationary distribution of the above process is the distribution $\pi(\mathbf{x}) \propto e^{-f(\mathbf{x})}$, under relatively mild regularity conditions on f . We will assume that x_0 (and hence all subsequent x_t) marginally follow π —i.e. the process is *stationary*.

We will also need several common assumptions on the f in the generative process.

Assumption 1 (Strong convexity). f is ρ -strongly convex.

Assumption 2 (Smoothness of f). f is infinitely differentiable, L_0 -smooth, and ∇f is L_1 -smooth.²

Assumption 3 (Linear growth). There exists a positive constant $K < \infty$, such that $\forall \mathbf{x} \in \mathcal{Z}$, $\|\nabla f(\mathbf{x})\| \leq K(1 + \|\mathbf{x}\|)$.

Assumption 1 ensures the least singular value of the Hessian is lower bounded by ρ —which ensures that $\int_{\mathbf{x}} e^{-f(\mathbf{x})}$ is finite. Assumption 3 ensures the existence of a solution to equation 1, and Assumption 2 ensures the solution of equation 1 is unique.³ We refer the readers to Ait-Sahalia et al. (2008) for formal justifications of these assumptions. We denote with \mathbf{x}_* the minimizer of f , and assume $\mathbf{x}_* = \vec{0}$ for convenience.

Finally, denote $B := \mathbb{E}_\pi \|\mathbf{x}\|$. Note that our assumptions on f guarantee a bounded B : let $Z_\pi := \int_{\mathbf{x}} \exp(-f(\mathbf{x})) d\mathbf{x}$ denote the partition function of the

stationary distribution π , and with $\mathbf{x}_* = \vec{0}$, we have

$$\begin{aligned} B &= \frac{1}{Z_\pi} \int_{\mathbf{x}} \|\mathbf{x}\| \exp(-f(\mathbf{x})) d\mathbf{x} \\ &\leq \frac{1}{Z_\pi} \int_{\mathbf{x}} \|\mathbf{x}\| \exp\left(-f(\mathbf{x}_*) - \frac{\|\mathbf{x}\|^2}{2(1/\rho)}\right) d\mathbf{x} \\ &= \pi(\mathbf{x}_*) \mathbb{E}_{\mathcal{N}(0, \frac{1}{\rho} \mathbf{I}_d)} \|\mathbf{x}\| \leq \pi(\mathbf{x}_*) \sqrt{\mathbb{E}_{\mathcal{N}(0, \frac{1}{\rho} \mathbf{I}_d)} \|\mathbf{x}\|^2} \\ &= \pi(\mathbf{x}_*) \sqrt{\frac{d}{\rho}}. \end{aligned} \quad (2)$$

3.1.1 Contrastive learning task

We choose the contrastive task to be binary classification on observations from the diffusion defined in equation 1. For $\eta = O_{L_1, L_2, \rho}(1)$ —i.e. any η sufficiently small as a function of the regularity parameters of f —we will consider the observations at integer multiples of η , namely $L_{\tilde{\mathcal{X}}} := \{\tilde{\mathbf{x}}_{i\eta}\} \subset \mathbb{R}^n$, and let $T > 0$ be length of the (continuous-time) sequence covered by these observations. Suppose the number of observations in $L_{\tilde{\mathcal{X}}}$ is $2m$ where $2m = \lfloor T/\eta \rfloor$.

The binary classification task is defined on a sequence of pairs of points denoted as $S_X := \{(\mathbf{x}_{2i\eta}, \mathbf{x}'_{2i\eta})\}_{i=0}^{m-1}$, where $\mathbf{x}_{2i\eta} = \tilde{\mathbf{x}}_{2i\eta}$, and $\mathbf{x}'_{2i\eta}$ is chosen in one of the two ways:

- With probability 1/2, we let $\mathbf{x}'_{2i\eta} = \tilde{\mathbf{x}}_{(2i+1)\eta}$ and output $(\mathbf{x}_{2i\eta}, \mathbf{x}'_{2i\eta})$ with label 1. We call these *positive* pairs.
- With probability 1/2, we sample $\mathbf{x}'_{2i\eta} \sim q$ for some *contrast* proposal distribution q and output $(\mathbf{x}_{2i\eta}, \mathbf{x}'_{2i\eta})$ with label 0. (We will specify the restrictions on q momentarily.) We call these *negative* pairs.

Intuitively, the task asks the model to distinguish the noise distribution q from the η -time transition kernel of the process $p_*^\eta : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, which is defined as

$$p_*^\eta(\mathbf{x}, \mathbf{x}') := \Pr(\mathbf{x}_{(t+1)\eta} = \mathbf{x}' | \mathbf{x}_{t\eta} = \mathbf{x}). \quad (3)$$

What we need to assume on the contrast distribution q is that it is sufficiently close to p_*^η (algorithmically, it also needs to have a pdf that is efficient to evaluate). Specifically, define a constant $c_q \geq 1$, such that the ratio between p_*^η and the proposal distribution q is bounded as

$$\frac{1}{c_q} \leq \frac{p_*^\eta(\mathbf{x}, \mathbf{x}')}{q(\mathbf{x}')} \leq c_q, \quad \forall \mathbf{x}, \mathbf{x}'. \quad (4)$$

We will show later that a smaller c_q is more preferable, which amounts to choosing a proposal distribution q that closely tracks the data distribution. This is consistent with observations in previous works on noise

²Recall a function f is L smooth if for any \mathbf{x}, \mathbf{y} in the support, $f(\mathbf{x}) \leq f(\mathbf{y}) + \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + \frac{L}{2} \|\mathbf{x} - \mathbf{y}\|_2^2$.

³Milder conditions on f ensure uniqueness/existence of solutions and not essential for our proofs—we assume this for simplicity of exposition.

contrastive learning (NCE) that a closer q makes the contrastive task harder and hence tends to work better in practice (Gutmann and Hyvärinen, 2010; Gao et al., 2020). Formally, a larger c_q will give a looser bound on the KL divergence between p^η and p_*^η , as we will see in theorem 2 and its proof.

The model we use for the supervised task is denoted as $h : \mathbb{R}^{2n} \rightarrow \mathbb{R}$, which takes in a $(\mathbf{x}, \mathbf{x}')$ pair and predicts the probability of the pair being positive. We assume the output of h to be bounded in $[0, 1]$.⁴ We denote the function class h belongs to as \mathcal{H} , and train h with the ℓ_2 loss:

$$\ell(h, \{(\mathbf{x}, \mathbf{x}'), y\}) = (h(\mathbf{x}, \mathbf{x}') - y)^2 \quad (5)$$

Let the *empirical risk* \hat{R} of a model h and loss ℓ associated with a training set S_X be defined as usual, and taking the expectation over S_X gives the *population risk* R :

$$\begin{aligned} \hat{R}_{S_X}(\ell \circ h) &:= \frac{1}{|S_X|} \sum_{i=1}^{|S_X|} \ell(h, \{(\mathbf{x}_{2i\eta}, \mathbf{x}'_{2i\eta}), y\}), \\ R(\ell \circ h) &:= \mathbb{E}_{S_X} \hat{R}_{S_X}(\ell \circ h). \end{aligned} \quad (6)$$

The *generalization gap* is defined as the maximum difference between the above two in the class of classifiers that we consider:

$$\Phi(S_X) := \sup_{h \in \mathcal{H}} \left[\left| \hat{R}_{S_X}(\ell \circ h) - R(\ell \circ h) \right| \right] \quad (7)$$

By way of remarks: the ℓ_2 loss is chosen since it is bounded, Lipschitz and strongly-convex, which makes the generalization bound calculations more manageable. It would be interesting to also provide bounds for cross-entropy or other losses.

We also need concepts capturing the complexity of the function class: the *empirical Rademacher complexity* $\hat{\mathfrak{R}}$ of a function class \mathcal{H} is defined with a given dataset S of size m as

$$\hat{\mathfrak{R}}_S(\mathcal{H}) := \frac{1}{m} \mathbb{E}_\varepsilon \left[\sup_{h \in \mathcal{H}} \left| \sum_{i \in [m]} \varepsilon_i h(x_i) \right| \middle| S = (x_1, \dots, x_m) \right] \quad (8)$$

The *Rademacher complexity* is defined by taking the expectation over the dataset S as

$$\mathfrak{R}_m(\mathcal{H}) := \mathbb{E}_{S:|S|=m} \hat{\mathfrak{R}}_S(\mathcal{H}). \quad (9)$$

Assumption 4 (Rademacher Complexity). *We will assume the Rademacher complexity of \mathcal{H} satisfies $\mathfrak{R}_m(\mathcal{H}) = O(K_\eta(\mathcal{H})\sqrt{\log m/m})$, where $K_\eta(\mathcal{H})$ depends on both \mathcal{H} and the task setup η .*

⁴This can easily be enforced, for example, by having a sigmoid layer at the end of a neural network.

The expression for $\mathfrak{R}_m(\mathcal{H})$ is common in standard generalization bounds. For instance, such dependency is achieved when the square root of the VC dimension of \mathcal{H} is bounded by $K_\eta(\mathcal{H})$ (Mohri et al., 2012).

3.2 Characterizing the Optimum given Infinite Data

To gain intuition on what the contrastive task does, we first characterize the optimum of the contrastive learning objective in the limit of infinite data. We note that similar analyses have appeared in other works on variants of contrastive learning, e.g. Hyvärinen and Morioka (2016); Tosh et al. (2020a,b). We show:

Lemma 1 (Population optimum). *The optimum of the contrastive learning objective*

$$\arg \min_h \mathbb{E}_{((x, x'), y)} \ell(h, \{(\mathbf{x}, \mathbf{x}'), y\})$$

satisfies

$$h^*(\mathbf{x}, \mathbf{x}') = \frac{p_*^\eta(\mathbf{x}, \mathbf{x}')}{q(\mathbf{x}') + p_*^\eta(\mathbf{x}, \mathbf{x}')$$

Proof. The proof proceeds by expanding the expectation in question, and a variance-bias like calculation. Namely, for a fixed (x, x') , taking the expectation over y gives:

$$\begin{aligned} \mathbb{E}_y \ell(h, \{(\mathbf{x}, \mathbf{x}'), y\}) &= (h(\mathbf{x}, \mathbf{x}') - \Pr(y = 1 | \mathbf{x}, \mathbf{x}'))^2 \\ &\quad + \Pr(y = 1 | \mathbf{x}, \mathbf{x}') (1 - \Pr(y = 1 | \mathbf{x}, \mathbf{x}')) \end{aligned} \quad (10)$$

The last term of equation 10 does not depend on h , so the minimum is achieved when $h^*(\mathbf{x}, \mathbf{x}') = \Pr(y = 1 | \mathbf{x}, \mathbf{x}')$. Expanding $\Pr(y = 1 | \mathbf{x}, \mathbf{x}')$ by the Bayes rule, we get

$$\begin{aligned} h^*(\mathbf{x}, \mathbf{x}') &= \Pr(y = 1 | \mathbf{x}, \mathbf{x}') \\ &= \frac{\Pr(\mathbf{x}, \mathbf{x}' | y = 1) \Pr(y = 1)}{\Pr(\mathbf{x}, \mathbf{x}' | y = 0) \Pr(y = 0) + \Pr(\mathbf{x}, \mathbf{x}' | y = 1) \Pr(y = 1)} \\ &= \frac{\Pr(\mathbf{x}, \mathbf{x}' | y = 1)}{\Pr(\mathbf{x}, \mathbf{x}' | y = 0) + \Pr(\mathbf{x}, \mathbf{x}' | y = 1)} \\ &= \frac{\pi(\mathbf{x}) p_*^\eta(\mathbf{x}, \mathbf{x}')}{\pi(\mathbf{x}) q(\mathbf{x}') + \pi(\mathbf{x}) p_*^\eta(\mathbf{x}, \mathbf{x}')} = \frac{p_*^\eta(\mathbf{x}, \mathbf{x}')}{q(\mathbf{x}') + p_*^\eta(\mathbf{x}, \mathbf{x}')} \end{aligned} \quad (11)$$

□

Note that the above proof uses essentially nothing about q other than that it is known: this is why population level analyses of contrastive objectives (e.g. like Hyvärinen and Morioka (2016)) may fail to capture many non-asymptotic aspects of the contrastive task.

3.3 Statement of main results

We claim that a low loss on the contrastive task implies closeness in a learned η -time transition kernel and the ground truth one. We will state the main results here and defer the proofs to Section 5.

Sample complexity bounds We first present the sample complexity for controlling the generalization gap defined in equation 7:

Theorem 1. *If $T = \Omega\left(\frac{B^2 K_\eta(\mathcal{H})^3}{\delta^2 \Delta_{gen}^3} \left(\log \frac{1}{\delta}\right)^{\frac{3}{2}}\right)$, then with probability $1 - \delta$, the generalization gap is bounded by Δ_{gen} .*

Note that the dependency of T on η comes only through the complexity measure $K_\eta(\mathcal{H})$. The reason there isn't additional dependence on η (e.g. the reader might imagine the number of "samples" effectively depends on T/η) is that though decreasing η gives more samples, the samples will be more dependent and hence less useful for generalization. The proof in section 5.1 will formally justify this intuition.

Distribution estimator from classifier for contrastive task Second, we show how to prove guarantees on an estimator for the transition kernel, given a classifier with a small contrastive task loss.

In light of Lemma 1, given a classifier h , define the transition kernel implied by h as

$$p^\eta(\mathbf{x}, \mathbf{x}') := \frac{h(\mathbf{x}, \mathbf{x}')q(\mathbf{x}')}{1 - h(\mathbf{x}, \mathbf{x}')}. \quad (12)$$

We wish to show that if h achieves a small loss, the p^η defined above is in fact close to p^* in some distributional sense.

We will show two types of guarantees, one under the assumption that p^η is somewhat close to p^* , and one for arbitrary p^η . In the first case, we will in fact show that a small loss implies that the learned p^η is close to p^* in a KL divergence sense (more precisely, $\mathbb{E}_x KL(p_*^\eta(\cdot|x)||p^\eta(\cdot|x))$ is small); in the second case, we will show that $\mathbb{E}_{x,x'} |p_*^\eta(x, x') - p^\eta(x, x')|$ is small.

The reason we can extract a stronger result in the first case is that we can leverage the strong convexity of the contrastive loss near the global optimum in an appropriate sense. Intuitively, in a strongly convex loss, a small loss implies closeness of the parameter to the global optimum. Such a property will not hold globally, as the loss may be arbitrarily non-convex as a function of p^η . Still, we will be able to extract a weaker guarantee (and with a less standard notion of distance).

Guarantees for p^η close to p_*^η Let constants $\Delta_{\min}, \Delta_{\max}$ be defined such that

$$\frac{p^\eta(\mathbf{x}, \mathbf{x}')}{p_*^\eta(\mathbf{x}, \mathbf{x}')} \in [\Delta_{\min}, \Delta_{\max}], \quad \forall \mathbf{x}, \mathbf{x}'. \quad (13)$$

and note that $0 < \Delta_{\min} \leq 1 \leq \Delta_{\max}$.

$\Delta_{\min}, \Delta_{\max}$ can be considered as a notion of closeness between p_*^η and p^η . When $\Delta_{\min}, \Delta_{\max}$ are close to 1, that is, when p^η lies in a small neighborhood of p_*^η , we can show the contrastive loss is locally strongly convex with respect to the KL divergence. This allows us to relate the loss to the KL divergence between p_*^η and p^η . Formally, we state the following result:

Theorem 2. *Suppose assumption 1-3 are satisfied and that $\Delta_{\max} \leq \frac{7}{6}$. Suppose the training error of h is $\epsilon_{tr} + \epsilon_*$, where $\epsilon_* := \mathbb{E}_{\{(x, x'), y\}} \left(\frac{p_*^\eta(\mathbf{x}, \mathbf{x}')}{p_*^\eta(\mathbf{x}, \mathbf{x}') + q(\mathbf{x}')} - y \right)^2$ is the optimal error achieved by p_*^η . If the generalization gap is bounded by ϵ_{tr} , then the average KL divergence between the ground truth and learned transition kernel is bounded by the contrastive loss as*

$$\mathbb{E}_{\mathbf{x} \sim \pi} KL(p_*^\eta(\cdot|\mathbf{x})||p^\eta(\cdot|\mathbf{x})) \leq \frac{2(1 + c_q)^5 \epsilon_{tr}}{\Delta_{\min}^2}. \quad (14)$$

Recall that c_q (defined in equation 4) represents how close the contrast distribution q is to p_*^η . Theorem 2 hence explains why a closer q is more preferable, as has been suggested by empirical evidence (Gao et al., 2020). In addition, as mentioned earlier, the bound $\Delta_{\max} \leq \frac{7}{6}$ is required to reason about the convexity of the contrastive loss in the neighborhood of p_*^η . Globally, the loss need not be convex, so it is entirely possible for a p^η faraway from p_*^η in the KL sense to have a small contrastive loss. Nevertheless, we can prove something weaker in this case.

Guarantees for arbitrary p^η In the case of an arbitrary p^η , we prove the following bound on the closeness to p_*^η :

Theorem 3. *Under assumptions 1-3, and let ϵ_*, T, p^η be the same as in Theorem 2. Let $\eta = O_{\rho, L_0, L_1}(1)$, then p^η satisfies*

$$\begin{aligned} & \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\cdot|\mathbf{x})} |p^\eta(\mathbf{x}, \mathbf{x}') - p_*^\eta(\mathbf{x}, \mathbf{x}')| \\ & \leq \sqrt{2\epsilon_{tr}} O\left(\sqrt{\pi(\mathbf{x}_*)} \left(\frac{1}{\rho\eta^2}\right)^{d/4}\right) \end{aligned} \quad (15)$$

We make two remarks about the theorem 3. First, the value of η cannot be too large for the RHS of equation 15 to obtain (i.e. $\eta = O_{\rho, L_0, L_1}(1)$). Analyzing merely the optimum of the contrastive objective would not reveal this.

Moreover, though the exponential dependency in η may appear pessimistic, it is in fact the right one. A closer inspection of the left hand side of equation 15 shows that its scaling in η is also $(1/\eta)^{d/2}$ (by Lemma 6)—so the only “extra” exponential factors are the η -independent exponential terms. It is not clear if this can be removed or is essential—or if possibly other losses can remove this kind of dependence.

4 Generalization Machinery for Non-iid Data

At the core of our analysis is a set of tools for non-iid data, which we first build up before discussing the proof. We will use generalization results for data coming from *strong mixing* stochastic processes: namely, the samples are not independent; but, intuitively, after a short amount of time, the samples are “almost independent”. Precisely, we use the notion of β -mixing:

Definition 1 (β -mixing). *For a stationary Markov process, the β -mixing coefficient is defined as the average TV distance between the distribution after running the process for t time with a given starting point, and the stationary distribution π :*

$$\beta(t) = \mathbb{E}_{\mathbf{x}} \text{TV}(P^t(\cdot|x_0 = \mathbf{x}), \pi) \quad (16)$$

A process is said to be β -mixing if $\lim_{t \rightarrow \infty} \beta(t) = 0$.

The β -mixing coefficient of a discrete-time sequence is defined similarly, with the conditional distribution defined between points in the sequence.

We note that β -mixing can be defined more generally on processes that may are not necessarily stationary or Markov. The above definition is the cleanest version that suffices for our setting.

The reason this will be useful for us is that when our data is β -mixing, we will be able to use generalization bounds similar to those we have for iid data. More precisely, we will leverage the following result by [Mohri and Rostamizadeh \(2009\)](#), which when applied to our setting becomes:

Lemma 2 (Rademacher complexity bound, [Mohri and Rostamizadeh \(2009\)](#), Theorem 1). *Let S_X form a β -mixing sequence with stationary distribution π . Then, for some $\delta \in (0, 1)$, for every μ such that $\delta > 2(\mu - 1)\beta(T/2\mu)$, with probability at least $1 - \delta$, the generalization gap $\Phi(S_X)$ is bounded by*

$$\begin{aligned} \Phi(S_X) &:= \sup_{h \in \mathcal{H}} \left[R(\ell \circ h) - \hat{R}_{S_X}(\ell \circ h) \right] \\ &\leq \left\{ \mathfrak{R}_\mu(\mathcal{H}) + \sqrt{\frac{\log(2/(\delta - \Delta_{appr}^\mu))}{2\mu}} \right\} \quad (17) \end{aligned}$$

where $\Delta_{appr}^\mu := 2(\mu - 1)\beta_{S_X}(T/2\mu)$, and $\mathfrak{R}_\mu(\mathcal{H})$ is the Rademacher complexity of \mathcal{H} over samples of size μ drawn iid from π .

The result is proved using a technique called *blocking* from [Yu \(1994\)](#). The idea is to divide a dependent sequence of samples into 2μ blocks of consecutive points, such that when the block size $\frac{T}{2\mu}$ is sufficiently large, every other block would be approximately independent because of the fast mixing. The generalization analysis can hence be divided into two steps, one for applying standard generalization bound on i.i.d. data (i.e. the blocks), and the other for bounding the approximation error of treating dependent blocks as independent ones. The term Δ_{appr}^μ is a consequence of the derivation in [Yu \(1994\)](#) and accounts for errors of approximating non-iid data with iid ones.

We will proceed by first showing fast mixing, then applying the generalization bounds above.

4.1 Proving β -mixing

We will first show β -mixing of the sequence S_X of pairs $(\mathbf{x}, \mathbf{x}')$ as constructed in Section 3; that is, by choosing \mathbf{x} from the diffusion process, and then choosing \mathbf{x}' to be η -time after in the process or from a proposal distribution with equal probability. Intuitively, this would suggest that once two points are sufficiently apart, they will be approximately independent, on which standard generalization bounds apply. Formally, we have the following result:

Lemma 3. *The β -mixing coefficients for the sequence S_X defined in Section 3 is $\beta_{S_X}(t) = O\left(\frac{B}{\sqrt{t}}\right)$.*

Proof. We will prove this by showing the sequence of pairs S_X shares the same β coefficients as the sequence of points $L_{\bar{X}}$ (Lemma 4). Then, since $L_{\bar{X}}$ is itself β -mixing (Lemma 5), the claim follows. \square

Having the same β coefficients between S_X and $L_{\bar{X}}$ makes intuitive sense, since the sequence of pairs can be considered as a mixture of a dependent sequence and an independent sequence, and adding the independent one should not worsen the mixing coefficient.

Lemma 4. $\beta_{S_X}(t) = \beta_{L_{\bar{X}}}(t)$.

Proof. First note that S_X is Markov and stationary, since the temporal dependency only comes from the first elements in the pairs, which are points in $L_{\bar{X}}$ that

is itself Markov and stationary:

$$\begin{aligned}
 & \Pr\left((\mathbf{x}_{2(i+1)\eta}, \mathbf{x}'_{2(i+1)\eta}) | (\mathbf{x}_0, \mathbf{x}'_0), \dots, (\mathbf{x}_{2i\eta}, \mathbf{x}'_{2i\eta})\right) \\
 &= \Pr\left(\mathbf{x}'_{2(i+1)\eta} | \mathbf{x}_{2(i+1)\eta}\right) \Pr\left(\mathbf{x}_{2(i+1)\eta} | \mathbf{x}_0, \dots, \mathbf{x}_{2i\eta}\right) \\
 &= \Pr\left(\mathbf{x}'_{2(i+1)\eta} | \mathbf{x}_{2(i+1)\eta}\right) \Pr\left(\tilde{\mathbf{x}}_{2(i+1)\eta} | \tilde{\mathbf{x}}_{2i\eta}\right) \\
 &= \Pr\left(\mathbf{x}'_{2(i+1)\eta} | \mathbf{x}_{2(i+1)\eta}\right) \Pr\left(\mathbf{x}_{2(i+1)\eta} | \mathbf{x}_{2i\eta}\right)
 \end{aligned} \tag{18}$$

The mixing coefficient of S_X can then be calculated explicitly, leading to $\beta_{S_X}(2i\eta) = \beta_{L_{\tilde{X}}}(2i\eta)$. The details are deferred to the appendix. \square

Next, we bound the TV distance, as a function of t , between the stationary distribution π and the distribution after running the diffusion for time t given any starting point:

Lemma 5 (Bubeck et al. (2018), Proposition 4). *Let $B := \mathbb{E}_\pi \|x\|$. For any $t > 0$, $\forall \mathbf{x} \in \mathcal{X}$,*

$$TV(\mathbb{P}(\mathbf{x}_t | \mathbf{x}_0 = \mathbf{x}), \pi) \leq \frac{B}{\sqrt{2\pi t}},$$

where $\mathbb{P}(\mathbf{x}_t | \mathbf{x}_0 = \mathbf{x})$ denotes the distribution after running the diffusion for time t conditioned on being at \mathbf{x} at time 0.

With the definition of β -mixing, Lemma 5 shows that $L_{\tilde{X}}$ itself is β -mixing, as long as $B < \infty$, $\beta_{L_{\tilde{X}}}(t\eta) = TV(\mathbb{P}(\mathbf{x}_{t\eta} | \mathbf{x}_0 = \mathbf{x}), \pi) = O(\frac{1}{\sqrt{t\eta}}) \rightarrow 0$ as $t \rightarrow \infty$.

5 Proofs of Main Results

We are now ready to prove the main results in section 3.3. We will start with the finite sample generalization bound, and map the loss on the contrastive task to the KL divergence between the learned and true transition kernels, assuming the former lies in a neighborhood of the latter. We will finish with the proof for theorem 3 where the closeness assumption is lifted.

5.1 Proof of the Generalization Bound

Let's first prove the sample complexity bound for generalization, where we use results in Mohri and Ros-tamizadeh (2009) to choose the optimal μ to bound the generalization gap.

Proof of Theorem 1. Following notations in lemma 2, let μ denote the number of “effective” training samples. Substituting in the choice of $T = \Omega\left(\frac{B^2 K_\eta(\mathcal{H})^3}{\delta^2 \Delta_{gen}^3} \left(\log \frac{1}{\delta}\right)^{\frac{3}{2}}\right)$, we have $\Delta_{appr}^\mu = O\left(\frac{B}{\sqrt{T}} \mu^{\frac{3}{2}}\right) \leq \delta$. Recall the empirical Rademacher complexity is $\mathfrak{R}_\mu = O\left(K_\eta(\mathcal{H}) \sqrt{\log \mu / \mu}\right)$. Then, choosing μ

such that $\mu = \Omega\left(\frac{K_\eta(\mathcal{H}) \sqrt{\log(1/(\delta - \Delta_{appr}))}}{\Delta_{gen}^2}\right)$, it can be checked that the following is satisfied:

$$C \sqrt{\frac{1}{\mu}} \left(K_\eta(\mathcal{H}) \sqrt{\log \mu} + \sqrt{-\log(\delta - \Delta_{appr})} \right) \leq \Delta_{gen} \tag{19}$$

The calculation details can be found in the appendix. \square

5.2 Proof of Theorem 2

Theorem 2 states that when p^η is close to p_*^η and the population contrastive loss is not much worse than the optimal value, the KL divergence between p_*^η and p^η is also small. At a high level, with p^η close to p_*^η , we can do a “multiplicative” Taylor expansion of p^η around p_*^η . Then, it can be shown that the second derivative is strictly positive with a proper choice of Δ_{max} . This is similar in spirit to the notion of strong convexity with respect to KL, from the difference in losses.

Proof of Theorem 2. Recall that in section 3.3 we defined constants Δ_{min} , Δ_{max} such that $\forall \mathbf{x}, \mathbf{x}', \frac{p^\eta(\mathbf{x}, \mathbf{x}')}{p_*^\eta(\mathbf{x}, \mathbf{x}')} \in [\Delta_{min}, \Delta_{max}]$. We can equivalently write this relation as $p^\eta = p_*^\eta(1 + \delta)$ with $\delta \in [\Delta_{min} - 1, \Delta_{max} - 1]$. By the mean value theorem, $\exists \xi \in [\Delta_{min} - 1, \Delta_{max} - 1]$, such that

$$\begin{aligned}
 & \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta} \text{KL}(p_*^\eta \| p^\eta) = \mathbb{E}_\pi \int p_*^\eta \log \frac{p_*^\eta}{p_*^\eta(1 + \delta)} \\
 &= - \mathbb{E}_\pi \int p_*^\eta \log(1 + \delta) \\
 &= - \mathbb{E}_\pi \int p_*^\eta \left(\delta - \int_{s=0}^{\delta} \frac{1}{2(1+s)^2} s \, ds \right) \\
 &\stackrel{(i)}{=} - \mathbb{E}_\pi \int p_*^\eta \left(\delta - \frac{1}{2} \frac{1}{(1+\xi)^2} \delta^2 \right) \\
 &\stackrel{(ii)}{\leq} \frac{1}{2\Delta_{min}^2} \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta} \delta^2.
 \end{aligned} \tag{20}$$

where (i) applies the mean value theorem to the second order Taylor expansion around 0, and (ii) uses $\int_{\mathbf{x}} p_*^\eta \delta = 0$ since p^η, p_*^η both integrates to 1.

Rewriting the gap between the population loss between $p^\eta(\mathbf{x}, \mathbf{x}')$ and $p_*^\eta(\mathbf{x}, \mathbf{x}')$ as a function of $\delta(\mathbf{x}, \mathbf{x}')$, we get:

$$r(\delta) := \left(\frac{p_*^\eta q \delta}{(p_*^\eta(1 + \delta) + q)(p_*^\eta + q)} \right)^2 \tag{21}$$

where the dependence on \mathbf{x}, \mathbf{x}' is omitted for clarity.

We would like to lower bound $r''(\delta)$. By the mean

value theorem, for any $\mathbf{x}, \mathbf{x}', \exists \xi' \in [\Delta_{\min} - 1, \Delta_{\max} - 1]$,

$$\begin{aligned} r''(\delta) &= r''(0) + r'''(0)\delta + \frac{1}{2}r''''(\xi')\delta^2 \\ &\geq \frac{2(p_*^\eta)^2 q^2}{(p_*^\eta + q)^5} ((7 - 6\Delta_{\max})p_*^\eta + q) \\ &\quad + \frac{12(p_*^\eta)^5 q^2}{p_*^\eta + q} \cdot \frac{5p_*^\eta + 3q - 2p_*^\eta \Delta_{\max}}{(\Delta_{\max} p_*^\eta + q)^6} \end{aligned} \quad (22)$$

The Taylor series converges when $|\delta| \leq \frac{p_*^\eta + q}{p_*^\eta} = 1 + \frac{q}{p_*^\eta}$, which always holds under our assumption on δ . Moreover, we require $\Delta_{\max} \leq \frac{7}{6}$ to ensure $(7 - 6\Delta_{\max})p_*^\eta + q > 0$. Then,

$$\begin{aligned} \frac{5p_*^\eta + 3q - 2p_*^\eta \Delta_{\max}}{(\Delta_{\max} p_*^\eta + q)^6} &\geq \frac{1}{3} \frac{8p_*^\eta + 9q}{(\frac{1}{6})^6 (7p + 6q)^6} \\ &\geq \frac{6^6}{3} \frac{8p_*^\eta + 9q}{(7p_*^\eta + 7q)^6} = \frac{8 \cdot 6^6}{3 \cdot 7^6} \frac{1}{(p_*^\eta + q)^5} \geq \frac{1}{(p_*^\eta + q)^5}. \end{aligned} \quad (23)$$

Substituting this back to equation 22 gives

$$\begin{aligned} r''(\delta) &\geq \frac{2(p_*^\eta)^2 q^3}{(p_*^\eta + q)^5} + \frac{24(p_*^\eta)^5 q^2}{(p_*^\eta + q)^6} \\ &\geq 2 \cdot \frac{1}{(1 + \frac{q}{p_*^\eta})^2} \cdot \frac{1}{(1 + \frac{p_*^\eta}{q})^3} \geq \frac{2}{(1 + c_q)^5}. \end{aligned} \quad (24)$$

We can then derive an upper bound on $\mathbb{E}\delta^2$. Recall that in theorem 2 the gap between the population loss of the learned h and that of h_* is set to be $2\epsilon_{tr}$:

$$\begin{aligned} 2\epsilon_{tr} &= \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim \frac{p_*^\eta + q}{2}} [r(\delta(\mathbf{x}, \mathbf{x}')) - r(0)] \\ &= \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim \frac{p_*^\eta + q}{2}} \int_0^\delta (\delta - t) r''(t) dt \\ &\geq \frac{r_{\min}}{2} \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim \frac{p_*^\eta + q}{2}} \delta^2 \geq \frac{r_{\min}}{4} \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta} \delta^2 \end{aligned} \quad (25)$$

where we denote $r_{\min} := \frac{2}{(1 + c_q)^5}$.

This means $\mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta} \delta^2 \leq \frac{8\epsilon_{tr}}{r_{\min}}$. Together with equation 20, we can bound the average KL as

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim \pi} \text{KL}(p_*^\eta(\cdot | \mathbf{x}) \| p^\eta(\cdot | \mathbf{x})) &\leq \frac{1}{2\Delta_{\min}^2} \mathbb{E}_{\pi, p_*^\eta} \delta^2 \\ &\leq \frac{4\epsilon_{tr}}{r_{\min} \Delta_{\min}^2} = \frac{2(1 + c_q)^5 \epsilon_{tr}}{\Delta_{\min}^2}. \end{aligned} \quad (26)$$

□

5.3 Proof of Theorem 3

Up to this point, we have reasoned about the generalization gap and the relation between the loss and distributional closeness when p^η is in the proximity of p_*^η . There is one last piece missing: we need to characterize what the value of the loss implies for p^η when its relation to p_*^η is unknown.

This is not an obvious task because the loss guarantees that the squared difference in equation 10 is small *on average* over \mathbf{x}, \mathbf{x}' according to our data distribution. This does not necessarily imply the squared difference in its numerator, ⁵ i.e. $(p^\eta(\mathbf{x}, \mathbf{x}') - p_*^\eta(\mathbf{x}, \mathbf{x}'))^2$, is small. For example, if $q(\mathbf{x}') \ll \min\{p^\eta(\mathbf{x}, \mathbf{x}'), p_*^\eta(\mathbf{x}, \mathbf{x}')\}$, then the above difference would be small regardless of the values of $p^\eta(\mathbf{x}, \mathbf{x}'), p_*^\eta(\mathbf{x}, \mathbf{x}')$.

We will leverage the following estimates on the transition kernel of the Langevin diffusion:

Lemma 6 (Gobet (2002), Proposition 1.2). *Under assumptions 1-3, $\exists c, C > 1$, such that*

$$\begin{aligned} p^\eta(\mathbf{x}, \mathbf{x}') &\geq \frac{1}{c} \frac{1}{\eta^{d/2}} e^{-C \frac{\|\mathbf{x} - \mathbf{x}'\|^2}{\eta}} e^{-C\eta \|\mathbf{x}\|^2}, \\ p_*^\eta(\mathbf{x}, \mathbf{x}') &\leq c \frac{1}{\eta^{d/2}} e^{-\frac{1}{c} \frac{\|\mathbf{x} - \mathbf{x}'\|^2}{\eta}} e^{C\eta \|\mathbf{x}\|^2}. \end{aligned} \quad (27)$$

The theorem in Gobet (2002) holds actually in a substantially more general setting than ours: it only requires that the drift (in our setting ∇f) and diffusion coefficient are in $C^{1+\gamma}, \gamma > 0$.

With this result in mind, as well as the previous lemmas, we are ready to prove theorem 3:

Proof of Theorem 3. Recall that the optimal solution of the contrastive task satisfies $h^* = \frac{p_*^\eta(\mathbf{x}, \mathbf{x}')}{q(\mathbf{x}') + p_*^\eta(\mathbf{x}, \mathbf{x}')}$ by Lemma 1, and that the population contrastive loss is no more than $2\epsilon_{tr}$ over the optimal ϵ_* achieved by h^* . This gives

$$\begin{aligned} 2\epsilon_{tr} &\geq \mathbb{E}_{\mathbf{x} \sim \pi} \mathbb{E}_{\mathbf{x}' \sim \frac{1}{2}(p_*^\eta(\mathbf{x}, \cdot) + q)} \ell(h) \\ &\geq \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim \pi} \mathbb{E}_{\mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} \ell(\hat{h}) \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim \pi} \mathbb{E}_{\mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} \left(\frac{q(p^\eta - p_*^\eta)}{(q + p^\eta)(q + p_*^\eta)} \right)^2 \end{aligned} \quad (28)$$

We now use the above loss bound to upper bound $\mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} |p^\eta(\mathbf{x}, \mathbf{x}') - p_*^\eta(\mathbf{x}, \mathbf{x}')|$. For notational convenience, we will drop \mathbf{x}, \mathbf{x}' when it is clear from the context.

Define a function r_q with $r_q(p) = \frac{p}{p+q}$. The population risk can now be written as $\mathbb{E}_{\mathbf{x}, \mathbf{x}'} (r_q(p^\eta) - r_q(p_*^\eta))^2$. Note that $r'_q(p) = \frac{q}{(p+q)^2}$ and r_q is concave in p , hence

$$r'_q(\max\{p^\eta, p_*^\eta\}) \cdot (p^\eta - p_*^\eta) \leq r_q(p^\eta) - r_q(p_*^\eta) \quad (29)$$

Using equation 29 and Cauchy-Schwarz, we have

$$\begin{aligned} &\mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} |p^\eta - p_*^\eta| \\ &= \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} \left[\frac{|p^\eta - p_*^\eta| r'_q(\max\{p^\eta, p_*^\eta\})}{r'_q(\max\{p^\eta, p_*^\eta\})} \right] \end{aligned} \quad (30)$$

⁵Recall that the squared term in equation 10 can be expanded as $(h(\mathbf{x}, \mathbf{x}') - \Pr(y = 1 | \mathbf{x}, \mathbf{x}'))^2 = \left(\frac{p}{p+q} - \frac{p_*}{p_*+q} \right)^2 = \frac{(p-p_*)^2 q^2}{(p+q)^2 (p_*+q)^2}$.

$$\begin{aligned} &\leq \mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} \left[(r_q(p^\eta) - r_q(p_*^\eta)) \cdot \frac{1}{r'_q(\max\{p^\eta, p_*^\eta\})} \right] \\ &\leq \sqrt{\mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} (r_q(p^\eta) - r_q(p_*^\eta))^2} \\ &\quad \cdot \sqrt{\mathbb{E}_{\mathbf{x} \sim \pi, \mathbf{x}' \sim p_*^\eta(\mathbf{x}, \cdot)} (\max\{p^\eta, p_*^\eta\} + q)^4 / q^2} \end{aligned}$$

where the first term is the population risk on the contrastive task, which is bounded by $2\epsilon_{tr}$.

We will proceed to bound the second term. Since p^η, p_*^η are both assumed to satisfy Assumption 1-3, lemma 6 allows us to bound the quantity of interest in equation 30: let c_*, C_* and \hat{c}, \hat{C} be the constants in lemma 6 for p_*^η and p^η respectively. Denote $C_u = \max\{C_*, \hat{C}\}$, $C_l = \min\{c_*, \hat{c}\}$. Recall that ρ is the strong convexity constant of f . It can be shown that if η is sufficiently small as a function of these constants (e.g. $\eta = \frac{\rho}{10C_u}$), let $\sigma^2 = \frac{C_l \eta}{2}$, then

$$\mathbb{E}_{\mathbf{x}, \mathbf{x}'} \frac{(\max\{p^\eta, p_*^\eta\} + q)^4}{q^2} \leq O\left(\pi(\mathbf{x}_*) \left(\frac{1}{\rho \eta^2}\right)^{d/2}\right) \quad (31)$$

The proof applies lemma 6 and the strong convexity of f to simplify the expression with a Gaussian-integral like calculation; the details are deferred to the appendix.

Plugging this inequality back in equation 30 gives the statement of the theorem. \square

The proof of theorem 3 can be adapted straightforwardly to accommodate the boundedness assumptions in theorem 2, namely, when $\frac{p^\eta}{p_*^\eta}$ and $\frac{p_*^\eta}{q}$ are bounded by $(\Delta_{\min}, \Delta_{\max})$ and $[\frac{1}{c_q}, c_q]$ respectively. In this case, the right hand side of equation 31 will be updated to $(2c_q + 1)^4 \mathbb{E}_{\mathbf{x}, \mathbf{x}'} (p_*^\eta)^2$. The exponential dependency on η is however still present, a consequence of lemma 6.

6 Conclusion

We study contrastive learning objectives in time-series settings—particularly, when the data comes from a strong-mixing diffusion process. We provide both sample complexity bounds and quantitative results on the proximity of the learned transition kernel, given a good classifier for a judiciously chosen contrastive task.

This is a first-cut work, and many natural open problems remain. For instance, how do other objectives (e.g. cross-entropy loss) perform? Are there better contrastive objectives than the proposed one that have a better scaling with dimension? Can we analyze the algorithmic effects of different choices of the contrast distribution q ? Finally, can we analyze more complicated (e.g. latent-variable) diffusion processes using similar methods?

Acknowledgements

We acknowledge the support of NSF via IIS-1909816 and OAC-1934584.

References

- Y. Ait-Sahalia et al. Closed-form likelihood expansions for multivariate diffusions. *The Annals of Statistics*, 36(2):906–937, 2008.
- R. Bhattacharya et al. Criteria for recurrence and existence of invariant measures for multidimensional diffusions. *The Annals of Probability*, 6(4):541–553, 1978.
- S. Bubeck, R. Eldan, and J. Lehec. Sampling from a log-concave distribution with projected langevin monte carlo. *Discrete & Computational Geometry*, 59(4):757–783, 2018.
- T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- R. Gao, E. Nijkamp, D. P. Kingma, Z. Xu, A. M. Dai, and Y. N. Wu. Flow contrastive estimation of energy-based models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7518–7528, 2020.
- E. Gobet. Lan property for ergodic diffusions with discrete observations. In *Annales de l’IHP Probabilités et statistiques*, volume 38, pages 711–737, 2002.
- M. Gutmann and A. Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304, 2010.
- A. Hyvarinen and H. Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ica. In *Advances in Neural Information Processing Systems*, pages 3765–3773, 2016.
- A. Hyvarinen and H. Morioka. Nonlinear ica of temporally dependent stationary sources. In *Artificial Intelligence and Statistics*, pages 460–469. PMLR, 2017.
- G. N. Milstein, J. G. Schoenmakers, V. Spokoiny, et al. Transition density estimation for stochastic differential equations via forward-reverse representations. *Bernoulli*, 10(2):281–312, 2004.
- M. Mohri and A. Rostamizadeh. Rademacher complexity bounds for non-iid processes. In *Advances*

- in *Neural Information Processing Systems*, pages 1097–1104, 2009.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT Press, 2012.
- S. Purushwalkam and A. Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. *arXiv preprint arXiv:2007.13916*, 2020.
- B. Rhodes, K. Xu, and M. U. Gutmann. Telescoping density-ratio estimation. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4905–4916. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/33d3b157ddc0896addfb22fa2a519097-Paper.pdf>.
- N. Saunshi, O. Plevrakis, S. Arora, M. Khodak, and H. Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pages 5628–5637, 2019.
- Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, and P. Isola. What makes for good views for contrastive learning. *arXiv preprint arXiv:2005.10243*, 2020.
- C. Tosh, A. Krishnamurthy, and D. Hsu. Contrastive estimation reveals topic posterior information to linear models. *arXiv preprint arXiv:2003.02234*, 2020a.
- C. Tosh, A. Krishnamurthy, and D. Hsu. Contrastive learning, multi-view redundancy, and linear models, 2020b.
- T. Wang and P. Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. *arXiv preprint arXiv:2005.10242*, 2020.
- B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, pages 94–116, 1994.