
Spectral Temporal Contrastive Learning

Sacha Morin^{†,1,2}, Somjit Nath^{†,2,3}, Samira Ebrahimi Kahou^{2,3,4}, Guy Wolf^{1,2,4}
¹Université de Montréal, ²Mila, ³École de technologie supérieure, ⁴CIFAR AI Chair

Abstract

Learning useful data representations without requiring labels is a cornerstone of modern deep learning. Self-supervised learning methods, particularly contrastive learning (CL), have proven successful by leveraging data augmentations to define positive pairs. This success has prompted a number of theoretical studies to better understand CL and investigate theoretical bounds for downstream linear probing tasks. This work is concerned with the temporal contrastive learning (TCL) setting where the sequential structure of the data is used instead to define positive pairs, which is more commonly used in RL and robotics contexts. In this paper, we adapt recent work on Spectral CL to formulate Spectral Temporal Contrastive Learning (STCL). We discuss a population loss based on a state graph derived from a time-homogeneous reversible Markov chain with uniform stationary distribution. The STCL loss enables to connect the linear probing performance to the spectral properties of the graph, and can be estimated by considering previously observed data sequences as an ensemble of MCMC chains.

1 Introduction

In recent years, Self-Supervised Learning (SSL) has gained popularity in fields with a wealth of unlabeled data such as computer vision and natural language processing. SSL has proven empirically successful: linear probing of ImageNet representations learned by both contrastive and non-contrastive SSL methods has shown surprising performance, being almost on par with supervised models [Bardes et al., 2021]. This success has prompted some authors to investigate SSL objectives from a theoretical perspective. Of particular interest to us is a recent work by [HaoChen et al., 2021] which formulates a learning objective based on a data augmentation graph and applies spectral graph theory tools to study optimal representations and derive error bounds for linear probing.

While data augmentation is perhaps the most common way of defining positive and negative pairs in a contrastive setting, potentially any known structure over the data points can be leveraged for contrastive learning. In particular, temporal contrastive learning (TCL) relies on the temporal structure of the data: data points that are close in time in observed sequences are treated as positives. Representation learning is usually achieved over a set of previously observed sequences. This paradigm finds application for representation learning for videos [Dave et al., 2022], reinforcement learning (RL) [Erraqabi et al., 2022], navigation [Yang et al., 2020, Morin et al., 2023], and potentially any other setting where sequential data is available. Observe that we are interested in **learning an embedding for each observation**, as opposed to learning a single embedding for the entire sequence of observations. This work aims to adapt the theoretical findings of HaoChen et al. [2021] to the temporal setting. Our main contributions are as follows:

1. We derive a TCL objective, which we dub Spectral Temporal Contrastive Learning (STCL);
2. STCL is based on a state graph which is itself derived from a time-homogenous reversible Markov Chain;

[†]Equal Contribution

3. Under a uniform stationary distribution assumption, the STCL minimizers can be directly characterized in terms of the eigenvectors of the state graph Laplacian, which grants access to additional tools to bound the linear probe error.

2 Background

Self-Supervised Learning. Given an input data matrix X , **self-supervised learning** (SSL) can be summarized as learning a “useful” data representation Z in the absence of clearly defined labels, often using a deep network. In practice, usefulness is subsequently measured by using Z in some supervised downstream task to predict some labels y using a linear model. This evaluation procedure is known as **linear probing** and ensures that relevant information (i.e., y) can be recovered from Z using a low capacity model. Naturally, y cannot be used while fitting Z .

Theoretical Works on Contrastive Learning. Arora et al. [2019] provides guarantees on a linear classification downstream task for a representation, Z learned by contrastive learning. One of the important assumptions of this work is that the positive pairs have the same conditional distribution, which basically translates to the positive pairs of augmentations being conditionally independent given the class label of the downstream task. There are some additional works Lee et al. [2020] in this aspect that focus solely on reconstruction-based methods, where learning Z happens by reconstructing the current input X in a lower dimensional space. In such methods, the authors assume conditional independence with respect to latent variables, for example, when we want to classify scenery like forests, seas, or deserts from images, we can use the background color as a latent variable. More interestingly, they weaken this assumption and provide bounds for approximate conditional independence(Lee et al. [2020], Theorem 4.2). Recently, Tosh et al. [2020] finds guarantees for linear probing by assuming there exists a hidden variable such that the positive pairs are conditionally independent given that hidden variable.

Of particular interest to our paper is the recent trend of analyzing contrastive learning through the lens of spectral graph theory [HaoChen et al., 2021, Balestrierio and LeCun, 2022]. A natural way of manipulating positive relationships over a dataset is to treat them as edges in a graph. Let \mathcal{G} be a graph with the rows of X (data points) as vertices and where edges represent a positive relationship.

HaoChen et al. [2021] leverages the graph framework to propose a Spectral Contrastive Learning objective with provable guarantees for a downstream classification task. They first introduce the notion of **population augmentation graph** where X in fact consists of all possible data points and their augmented views¹. They then formalize the following intuition: it is substantially easier to go from one image to the other using augmentations if both images belong to the same class according to the downstream labels y . This implies that the different subgraphs in \mathcal{G} corresponding to the different classes in y will have strong intra-connectivity but weak inter-connectivity, which in turn ensures that some eigenvectors of the Laplacian of \mathcal{G} will align with the class structure [Ng et al., 2001].

The authors then introduce their SSL objective and show that any minimizer Z^* will correspond to the bottom eigenvectors of the Laplacian of \mathcal{G} , up to a left scaling transform and a right linear transform. They then derive a bound on the linear probe error in the population setting that essentially depends on how well the classes y turn out to be clustered in \mathcal{G} (Theorem 3.8). They further consider additional finite-sample generalization bounds.

The Graph Laplacian in RL. A number of papers formulate a state graph similar to ours and rely on some elements of spectral graph theory to design learning objectives in the context of RL Machado et al. [2017], Erraqabi et al. [2022]. LapRep [Wu et al., 2018] is particularly relevant to our work, as it formulates a similar state graph based on a Markov chain and derives a contrastive learning objective using the classical *graph drawing* objective of [Koren, 2003] as a starting point, which should also be minimized by the Laplacian eigenvectors. Our final objective differs substantially from the LapRep one as we rely on the matrix factorization perspective of [HaoChen et al., 2021] to derive **both** the positive and the negative terms in the contrastive loss, while the LapRep negative term is obtained by relaxing the original orthogonality constraint of graph drawing. Nevertheless, the LapRep method would be an interesting benchmark in future empirical comparisons of STCL.

¹They consider an exponentially large, but finite data space, e.g., all vectors in a subset of \mathbb{R}^d with finite precision representations.

3 Spectral Temporal Contrastive Learning

This section will adapt the theory of HaoChen et al. [2021] to the temporal setting and derive the Spectral Temporal Contrastive Loss (STCL). We will begin by defining a Markov chain and an associated state graph \mathcal{G} in Subsection 3.1. Using this state graph as a temporal analog of the data augmentation graph, we will introduce the population matrix factorization loss from HaoChen et al. [2021] in Subsection 3.2 and characterize its minimizers in terms of the eigenvectors of the normalized graph Laplacian. Finally, Subsection 3.3 shows how a contrastive loss can estimate the population loss without observing \mathcal{G} using samples from the Markov Chain.

3.1 Markov Chain & State Graph Definitions

Consider a time-homogeneous Markov chain over discrete finite states $S = \{i\}_{i=1}^N$ with transition probabilities q_{ij} and some associated observations $x_i \in \mathcal{X}$. We further assume a unique stationary distribution π and let $\rho_{ij} := \pi_i q_{ij}$. We will now build a weighted state graph \mathcal{G} under the assumption that the chain is reversible: $\rho_{ij} = \rho_{ji}$. Using S as vertices we define an adjacency matrix \mathbf{A} such that $\mathbf{A}_{ij} = \rho_{ij}$. The reversibility property ensures that \mathbf{A} is symmetric. Moreover, if the first state in the chain is distributed according to π , then ρ_{ij} can also be interpreted as a joint distribution over transitions, with π acting as a marginal over states and q as a conditional. Computing the elements of the diagonal degree matrix \mathbf{D} , therefore, corresponds to marginalizing states, i.e.

$$D_{ii} = \sum_j \rho_{ij} = \pi_i \implies \mathbf{D} = \text{diag}(\pi). \quad (1)$$

Equipped with the matrix \mathbf{D} , we can define the normalized adjacency matrix as $\bar{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$ and the normalized Laplacian as $\bar{\mathbf{L}} = \mathbf{I} - \bar{\mathbf{A}}$. Observe that $\bar{\mathbf{A}}$ and $\bar{\mathbf{L}}$ share the same eigenvectors, and if λ is an eigenvalue of $\bar{\mathbf{A}}$ then $1 - \lambda$ is an eigenvalue of $\bar{\mathbf{L}}$. This fact will allow to interchangeably use the top eigenvectors of $\bar{\mathbf{A}}$ or the bottom eigenvectors $\bar{\mathbf{L}}$ in theoretical arguments.

3.2 STCL in the Population Setting

Population Objective. We consider the problem of learning vertex Euclidean representations from the high dimensional observations \mathbf{X} using a deep encoder $f : \mathcal{X} \rightarrow \mathbb{R}^k$. In this section, we assume knowledge of \mathcal{G} and $\bar{\mathbf{A}}$. We further assume in the population setting that we have exactly one observation per state, which allows to build a population representation matrix $\mathbf{Z} \in \mathbb{R}^{N \times k}$ by stacking $f(x_i)$ as rows. To learn our feature extractor f , we will adapt the objective from HaoChen et al. [2021] to our Markov chain setting. Consider the low-rank matrix factorization objective

$$\mathcal{L}_{mf}(\mathbf{Z}) = \|\bar{\mathbf{A}} - \mathbf{D}^{\frac{1}{2}} \mathbf{Z} \mathbf{Z}^\top \mathbf{D}^{\frac{1}{2}}\|_F^2. \quad (2)$$

and the minimizers $\mathbf{Z}^* = \text{argmin}_{\mathbf{Z} \in \mathbb{R}^{N \times k}} \mathcal{L}_{mf}(\mathbf{Z})$. The rescaling by $\mathbf{D}^{\frac{1}{2}}$ effectively means that we consider rescaled representations $\sqrt{\pi_i} f(x_i)$ in the learning objective. As we will see, this rescaling will prove critical in the derivation of a sampling-based learning objective (Section 3.3), but will complicate the analysis of \mathbf{Z}^* .

The matrix factorization approach allows us to characterize the minimizers \mathbf{Z}^* in terms of the eigenvectors of $\bar{\mathbf{A}}$, which in turn provides rich information on the geometry of \mathcal{G} . Indeed, the Eckart–Young–Mirsky Theorem [Eckart and Young, 1936] ensures that any rank k minimizer \mathbf{Z}^* of Equation 2 can be expressed as $\mathbf{Z}^* = \mathbf{D}^{-\frac{1}{2}} \mathbf{U}_k \mathbf{\Lambda}_k^{\frac{1}{2}} \mathbf{Q}$ where, $\mathbf{U}_k \in \mathbb{R}^{N \times k}$ denotes the top k eigenvectors of $\bar{\mathbf{A}}$, $\mathbf{\Lambda}_k \in \mathbb{R}^{k \times k}$ is the truncated diagonal matrix of the top eigenvalues, and $\mathbf{Q} \in \mathbb{R}^{k \times k}$ is an orthogonal matrix. \mathbf{U}_k can be equivalently understood as the bottom eigenvectors of $\bar{\mathbf{L}}$.

Linear Probing. Following the SSL setting, we then freeze the representations \mathbf{Z} and perform linear probing by minimizing a regression loss with respect to some target vector $\mathbf{y} \in \mathbb{R}^N$:

$$\begin{aligned} \mathcal{C}(\mathbf{w}) &= \min_{\mathbf{w} \in \mathbb{R}^k} \|\mathbf{y} - \mathbf{Z}^* \mathbf{w}\|^2 \\ &= \min_{\mathbf{w} \in \mathbb{R}^k} \|\mathbf{y} - \mathbf{D}^{-\frac{1}{2}} \mathbf{U}_k \mathbf{w}\|^2 \end{aligned}$$

where the second line results from the fact that both $\mathbf{\Lambda}_k^{\frac{1}{2}}$ and \mathbf{Q} are invertible, which allows to freely reparametrize \mathbf{w} . Linear probing can in principle reach zero error whenever $\mathbf{y} \in \text{span}(\mathbf{D}^{-\frac{1}{2}} \mathbf{U}_k)$.

The matrix $D^{-\frac{1}{2}}$ prevents a direct connection to the spectral properties of \mathcal{G} , since the normalized adjacency matrix eigenvectors U_k do not directly span the space of interest. In HaoChen et al. [2021], the authors overcome this limitation by considering a classification task. The weights in $D^{-\frac{1}{2}}$ are positive and only act as temperature parameters in the softmax used by the classifier probe, with no impact on the logit rankings. For the purpose of training a linear classifier, the authors claim that the representation $D^{-\frac{1}{2}}U_k$ is as useful as U_k .

Uniform Stationary Distribution. In the regression setting, $D^{-\frac{1}{2}}$ matters more. Observe that in the specific case where π is uniform over S , $D^{-\frac{1}{2}}$ effectively acts as a scalar scaling of all dimensions by $\frac{1}{\sqrt{N}}$. In this case, $\text{span}(U_k)$ properly describes the set of linear tasks achievable with the learned representation and allows to directly apply spectral graph theory. We further observe that under the uniform π assumption, $\bar{A} = \frac{1}{N}A$ and the normalization step does not affect the eigenvectors of the original adjacency and Laplacian matrices.

Tasks. Describing the set of ideal linear tasks as $\mathbf{y} \in \text{span}(U_k)$ is an appealing proposition since U_k is known to provide a Euclidean embedding that captures the geometry of \mathcal{G} [Spielman, 2019]. To illustrate this point, we show the top eigenvectors of \bar{A} for a ring graph in Figure 1 and a 2D grid graph in Figure 2. Pairs of state with high transition probabilities are mapped closely in terms of Euclidean distance. As we will further explore in our experiments in Appendix 5.1, this property can be leveraged for state/pose prediction tasks that strongly correlate with the underlying state graph connectivity. U_k is also known to encode the cluster structure of the graph when such clusters exist [Ng et al., 2001]. In navigation tasks, this fact can be exploited to predict different "rooms" in the environment and learn RL options [Machado et al., 2017].

3.3 STCL Loss Estimator

Contrastive loss. In practice, we do not have access to the graph \mathcal{G} or all the states in S to train f . However, it turns out that \mathcal{L}_{m_f} (Equation 2) can be expanded as

$$\mathcal{L}_{m_f}(\mathbf{Z}) \propto -2 \sum_{ij} \rho_{ij} f(x_i)^\top f(x_j) + \sum_{ij} \pi_i \pi_j (f(x_i)^\top f(x_j))^2 \quad (3)$$

$$\mathcal{L}(f) := -2 \mathbb{E}_{(i,j) \sim \rho} [f(x_i)^\top f(x_j)] + \mathbb{E}_{i \sim \pi, j \sim \pi} [(f(x_i)^\top f(x_j))^2]. \quad (4)$$

We can see a contrastive loss emerging from Equation 4: the first term pushes positive pairs sampled from the transition joint ρ to be aligned. The second one samples "negative pairs" independently from the stationary distribution π and is minimized whenever their representations are orthogonal.

Sampling. We cannot directly sample from π and ρ to estimate \mathcal{L} . However, it is common to assume in TCL settings that we have access to a previously collected set of sequences $\{(x_t^m)_{t=1}^M\}_{m=1}^M$ which we can use to train f . Examples of this setting includes exploring an environment with a robot [Morin et al., 2023] or pretraining on human videos for offline RL [Ma et al., 2022].

We consider the observed sequences as chains sampled from the Markov chain defined in Section 3.1. MCMC theory suggests that a very long chain or an ensemble of chains, ideally from different starting states, should eventually mix and allow to approximately sample from π (and from ρ by first sampling x_t from π and then sampling x_{t+1} from observed transitions from x_t). For our experiments, we simply treat our observations as samples from π , even if mixing is not guaranteed. We provide early experimental results in Section 5.1.

4 Conclusion and Future Works

In this work, we extended the results of HaoChen et al. [2021] to the temporal contrastive learning setting by considering the state graph induced by an underlying Markov chain. In future work, we hope to leverage Lemma B.6 of HaoChen et al. [2021] —and potentially additional assumptions on the task \mathbf{y} and the graph \mathcal{G} — to obtain more explicit and interpretable guarantees. Moreover, the uniform stationary assumption is somewhat restrictive. We hope that a more general argument can handle the left scaling by $D^{-\frac{1}{2}}$ to better characterize the span of the STCL minimizers. We also aim to relax the reversibility assumption and extend the analysis to continuous state spaces.

References

- Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- Jeff Z. HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *CoRR*, abs/2106.04156, 2021. URL <https://arxiv.org/abs/2106.04156>.
- Ishan Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. Tclr: Temporal contrastive learning for video representation. *Computer Vision and Image Understanding*, 219:103406, 2022.
- Akram Erraqabi, Marlos C. Machado., Mingde Zhao, Sainbayar Sukhbaatar, Alessandro Lazaric, Denoyer Ludovic, and Yoshua Bengio. Temporal abstractions-augmented temporally contrastive learning: An alternative to the laplacian in rl. In James Cussens and Kun Zhang, editors, *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pages 641–651. PMLR, 01–05 Aug 2022. URL <https://proceedings.mlr.press/v180/erraqabi22a.html>.
- Ge Yang, Amy Zhang, Ari Morcos, Joelle Pineau, Pieter Abbeel, and Roberto Calandra. Plan2vec: Unsupervised representation learning by latent plans. In *Learning for Dynamics and Control*, pages 935–946. PMLR, 2020.
- Sacha Morin, Miguel Saavedra-Ruiz, and Liam Paull. One-4-all: Neural potential fields for embodied navigation. *arXiv preprint arXiv:2303.04011*, 2023.
- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *CoRR*, abs/1902.09229, 2019. URL <http://arxiv.org/abs/1902.09229>.
- Jason D. Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *CoRR*, abs/2008.01064, 2020. URL <https://arxiv.org/abs/2008.01064>.
- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. *CoRR*, abs/2008.10150, 2020. URL <https://arxiv.org/abs/2008.10150>.
- Randall Balestriero and Yann LeCun. Contrastive and non-contrastive self-supervised learning recover global and local spectral embedding methods, 2022.
- Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 14, 2001.
- Marlos C Machado, Marc G Bellemare, and Michael Bowling. A laplacian framework for option discovery in reinforcement learning. In *International Conference on Machine Learning*, pages 2295–2304. PMLR, 2017.
- Yifan Wu, George Tucker, and Ofir Nachum. The laplacian in RL: learning representations with efficient approximations. *CoRR*, abs/1810.04586, 2018. URL <http://arxiv.org/abs/1810.04586>.
- Yehuda Koren. On spectral graph drawing. In Tandy Warnow and Binhai Zhu, editors, *Computing and Combinatorics*, pages 496–508, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-45071-9.
- C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. doi: 10.1007/BF02288367.
- Daniel Spielman. Spectral and algebraic graph theory. *Yale lecture notes, draft of December*, 4:47, 2019.
- Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. *arXiv preprint arXiv:2210.00030*, 2022.

5 Appendix




Figure 1: Top Eigenvectors of \bar{A} for a Markov Chain over a **ring graph** with 100 states.




Figure 2: Top Eigenvectors of \bar{A} for a Markov Chain over a **11x10 grid graph**.

5.1 Experiments

We consider 2 toy experiments. The first dataset consists of images of a 3D Teapot model rotated over 400 different angles Weinberger et al. [2004], and the downstream task is to predict the Teapot angle pose $(\cos \theta, \sin \theta)$. The second dataset is based on images generated by a Grid World environment, where the probe task will be to predict the (x, y) coordinates of the agent. We collect random trajectories with random uniform policies in both environments and train STCL with a Resnet-18 backbone and representations of dimension $k = 8$. As a baseline, we use PCA to extract the top 8 principal components and run the same linear probe. We visualize the outputs of the probes in Figure 3 and their R^2 metrics in Table 1.

Table 1: Coefficient of determination (R^2) of a supervised linear probe for PCA, Spectral Temporal Contrastive learning (STCL) and Ground Truth (GT). Please refer to main text for details.

	PCA Probe \uparrow	STCL Probe \uparrow	GT \uparrow
Teapot	0.80	0.98	1.00
Grid World	0.19	0.96	1.00




Figure 3: Outputs of the linear probe for the PCA baseline (second column) and STCL (third column). Ground truth targets are shown in the last column.