Predictive Sparse Manifold Transform

Yujia Xie*
YXIE@GATECH.EDU
Georgia Institute of Technology, Atlanta, USA

Xinhui Li*
XINHUILI@GATECH.EDU
Georgia Institute of Technology, Atlanta, USA

Vince D. Calhoun
VCALHOUN@GATECH.EDU
Georgia Institute of Technology, Atlanta, USA

Abstract

We present *Predictive Sparse Manifold Transform (PSMT)*, a minimalistic, interpretable and biologically plausible framework for learning and predicting natural dynamics. PSMT incorporates two layers where the first *sparse coding* layer represents the input sequence as sparse coefficients over an overcomplete dictionary and the second *manifold learning* layer learns a geometric embedding space that captures topological similarity and dynamic temporal linearity in sparse coefficients. We apply PSMT on a natural video dataset and evaluate the reconstruction performance with respect to contextual variability, the number of sparse coding basis functions and training samples. We then interpret the dynamic topological organization in the embedding space. We next utilize PSMT to predict future frames compared with two baseline methods with a static embedding space. We demonstrate that PSMT with a dynamic embedding space can achieve better prediction performance compared to static baselines. Our work establishes that PSMT is an efficient unsupervised generative framework for prediction of future visual stimuli.

Keywords: predictive sparse manifold transform, sparse coding, dynamic predictive coding, manifold learning, slow feature analysis, spatiotemporal dynamics

1. Introduction

"Intelligence is the capacity of the brain to predict the future by analogy to the past." – Jeff Hawkins [4]

One fundamental problem of intelligence is how to precisely and efficiently predict future sensory stimuli based on prior knowledge and experience. Our perceived world can be usually viewed as a temporal sequence of natural scenes with smooth variability from one to the next. A sequence of natural images can be represented via sparse coefficients over an overcomplete dictionary, which characterizes spatially localized, oriented, bandpass receptive fields of simple cells in mammalian primary visual cortex [6, 7]. However, the learned sparse codes are not organized to reveal the topological structure of the spatiotemporal dynamics of visual stimuli. Sparse manifold transform (SMT) [1] adds another *manifold learning* layer [10] on top of the *sparse coding* layer [6] to capture topological similarity of the sparse coefficients. It models the continuous temporal transformation of the input sequence via a linear trajectory in the geometric embedding space. Reconstruction of the input sequence from the geometric embedding space can well approximate the original sequence.

---

* Equal contribution.

Can we learn the structured spatiotemporal dynamics from the existing sequence and predict the future sequence leveraging the dynamic embedding space? In this study, we propose *Predictive Sparse Manifold Transform (PSMT)*, a framework inspired by SMT and dynamic predictive coding [1, 5] to learn and predict spatiotemporal dynamics in natural scenes. We process a natural video as a sequence of image frames and apply PSMT on the dataset using different hyperparameters. We interpret the topological organization of the geometric embedding matrix over time and demonstrate the dynamic nature of the embedding space. We next assess the reconstruction error with respect to the contextual variability in natural scenes. We further predict the future frames using PSMT and two other baseline approaches with a static embedding space in two different contextual settings, and evaluate the prediction performance. Our results demonstrate that using only two layers, PSMT with a dynamic embedding space can achieve better prediction performance compared to two static baselines, thus establishing a minimalistic, interpretable and biologically plausible framework for learning and predicting natural dynamics.

2. Methods

2.1. Predictive Sparse Manifold Transform

We outline the Predictive Sparse Manifold Transform (PSMT) framework here and expand its details in Appendix A. Past temporal natural image sequence is processed as $\mathbf{X} = \{\boldsymbol{x}_t\}_{t=T-H+1}^{t=T-1}$ with $\boldsymbol{x}_t \in \mathbb{R}^N$, where T is the index of current frame and H is the number of previous frames. Future inputs $\mathbf{X}^* = \{\boldsymbol{x}_t\}_{t=T}^{t=T+K-1}$, where K is the number of future frames to be predicted, are processed similarly and incorporated in PSMT iteratively as delayed inputs (Step 1).

In PSMT, we aim to predict future inputs $\boldsymbol{x}_{t=T-1+k}$ for $k = 1, ..., K$ given all previous inputs at $\boldsymbol{x}_{t=T-1+k-1}$. For the ease of presentation, we use footnotes *prev* to represent $t = T - 1 + k - 1$, *current* to represent $t = T - 1 + k$, and *future* to represent $t = T + k$.

Algorithm 1: Predictive Sparse Manifold Transform

Input: $T, H, K, M, N, f, \mathbf{X} = \{\boldsymbol{x}_t\}_{t=T-H+1}^{t=T-1}, \mathbf{X}^* = \{\boldsymbol{x}_t\}_{t=T}^{t=T+K-1}$;   // $\mathbf{X}^*$: delayed input

Output: $\hat{\boldsymbol{x}}_{t=T+k}$ for $k = 1, ..., K$

Sparse coding: $(\boldsymbol{\Phi}_M, \mathbf{A}) \leftarrow (\boldsymbol{\Phi}_M, \mathbf{A} : \mathbf{X} \approx \boldsymbol{\Phi}_M \mathbf{A})$;   // $\boldsymbol{\Phi}_M$: dictionary with M items

for *k = 1, ..., K* **do**

  Step 1: $\boldsymbol{x}_{\text{current}} := \boldsymbol{x}_{t=T-1+k} \leftarrow \mathbf{X}^*_k$;   // receive most-recent input at $t = T-1+k$

  Step 2: $\boldsymbol{\alpha}_{\text{current}} := \boldsymbol{\alpha}_{t=T-1+k}$ s.t. $\boldsymbol{x}_{T-1+k} \approx \boldsymbol{\Phi}_M \boldsymbol{\alpha}_{T-1+k}$; // get most-recent sparse code from $\boldsymbol{\Phi}_M$

  Step 3: $\mathbf{A} \leftarrow [\mathbf{A}; \boldsymbol{\alpha}_{\text{current}}]$;   // update $\mathbf{A}$ to include most-recent sparse code

  Step 4: $\mathbf{A}^* \leftarrow \mathbf{A}_{[t=T-H+k-1:t=T+k-1]}$;   // include most recent H entries in A

  Step 5: $\mathbf{P}^* \leftarrow \text{argmin}_{\mathbf{P}}\{\|\mathbf{P}\mathbf{A}^*\mathbf{D}\|_F^2 : \mathbf{P}\mathbf{V}\mathbf{P}^\top = \mathbf{I}\}$;   // compute embedding matrix $\mathbf{P}^*$

  Step 6: $\boldsymbol{\beta}^* \leftarrow \mathbf{P}^*\boldsymbol{\alpha}^*_{\text{future}} \approx 2\mathbf{P}^*\boldsymbol{\alpha}_{\text{current}} - \mathbf{P}^*\boldsymbol{\alpha}_{\text{prev}}$; // predict next embedding at $t = T+k$

  Step 7: $\boldsymbol{\alpha}^{\text{REC}}_{\text{future}} \leftarrow \text{argmin}_{\boldsymbol{\alpha}}\{\|\mathbf{P}^*\boldsymbol{\alpha} - \boldsymbol{\beta}^*\|_F^2 + \lambda \mathbf{z}^\top \boldsymbol{\alpha} : \boldsymbol{\alpha} \succeq 0\}$;   // reconstruct next sparse code

  Step 8: $\hat{\boldsymbol{x}}_{\text{future}} = \boldsymbol{\Phi}_M \boldsymbol{\alpha}^{\text{REC}}_{\text{future}}$;   // estimate 1-step future input
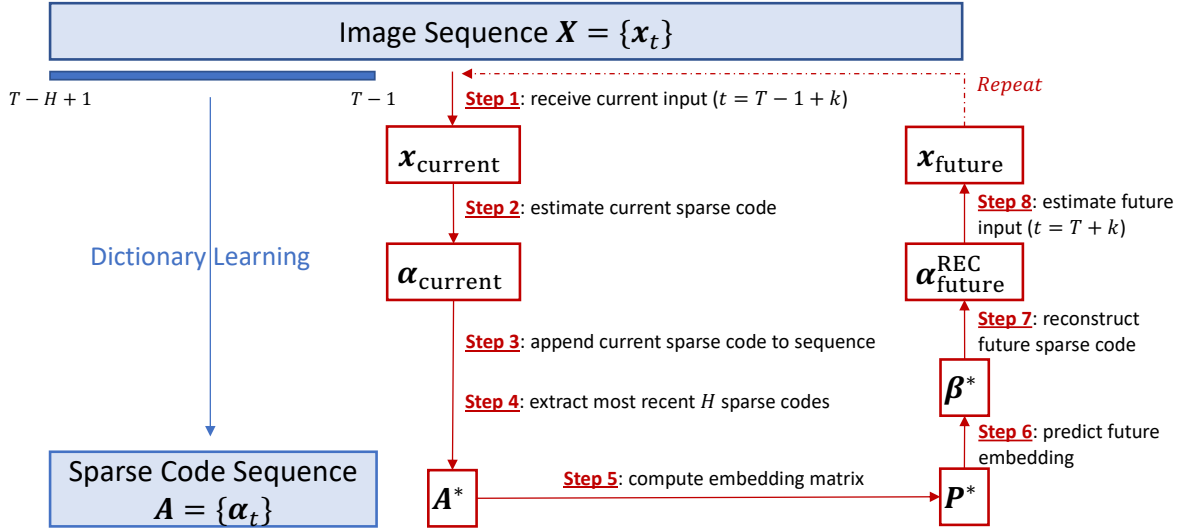
end

Figure 1: **PSMT algorithm illustration.**

2.2. Experiment

Dataset and Preprocessing We record the film *Russian Ark* [9] which was shot in one take and convert the video to a sequence of 10000 image frames. Each frame includes 800×800 pixels, which is then downsampled to 200×200 pixels. We next extract a patch of 16×16 pixels at the center of the downsampled frame and perform standardization on each patch.

Hyperparameter Search To investigate the optimal hyperparameters, we test 9 combinations of hyperparameters including the number of sparse coding basis functions M and the number of training frames H. We set $M = 512, 1024, 2048$ to learn a $2\times$, $4\times$, $8\times$ overcomplete dictionary, respectively. For each M, we evaluate a training size of $2\times, 2.5\times, 3 \times M$, each with a $\frac{3}{4}$ overlapping window. We compute the reconstruction mean squared error (MSE) for each combination.

Dynamic Topological Analysis We analyze the topology of $\boldsymbol{\Phi}_M$ by looking at the embedding matrix $\mathbf{P}^*$ dynamically for each training frame H. Each column in $\mathbf{P}^*$ corresponds to each column in $\boldsymbol{\Phi}_M$ (i.e., each dictionary item). We compute the pairwise cosine-similarity among columns in $\mathbf{P}^*$ and construct a graph $\mathcal{G}$ with each node representing each dictionary item, and each edge representing whether the cosine-similarity between the nodes is greater than a threshold (0.7). Using the optimal hyperparameter combination from the hyperparameter search above, we have $\mathbf{P}^*{}_t$ centered at different time t and consequently $\mathcal{G}_t$. We select a subset of nodes in $\mathcal{G}$ to visualize across time and analyze their connectivity behavior (i.e., neighbor count). We further investigate the implication of similarity in $\mathbf{P}^*$ to topological similarity between dictionary items.

Reconstruction Unlike SMT using all frames to reconstruct, PSMT relies only on H frames prior to time T to make future predictions. By analyzing the reconstruction performance with respect to T and H, we can develop insights into tuning these two parameters for prediction. We implement the original SMT with all frames included and measure the reconstruction MSE at each frame.

Prediction We evaluate the prediction performance with respect to two hyperparameters – start time T and future prediction sequence length K – estimated from reconstruction result. We compare PSMT with two baseline methods. Baseline 1 uses a static embedding matrix $\mathbf{P}^*$ and a static sparse

Table 1: **Hyperparameter search.** The mean $\pm$ standard deviation of the mean squared error ($\downarrow$) after trimming 25% tails for different sparse coding basis functions M and training frames H. More basis functions and training samples lead to better reconstruction performance.

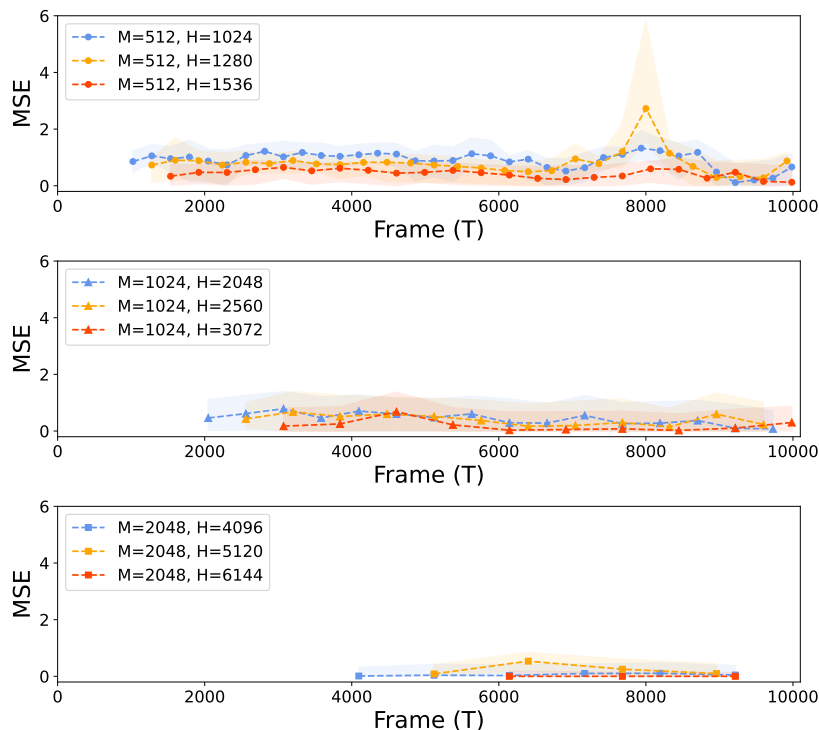| M / H	$2 \times M$	$2.5 \times M$	$3 \times M$
512	0.925 ± 0.129	0.734 ± 0.125	$\mathbf{0.425 \pm 0.106}$
1024	0.484 ± 0.202	0.474 ± 0.220	$\mathbf{0.316 \pm 0.178}$
2048	0.195 ± 0.174	0.309 ± 0.238	$\mathbf{0.190 \pm 0.228}$

Figure 2: **Hyperparameter search.** The median and the interquartile range of the mean squared error ($\downarrow$) at each segment for different sparse coding basis functions M and training frames H. More basis functions and training samples lead to better reconstruction performance.

coefficient $\boldsymbol{\alpha}^*$ obtained from H training frames for all future predictions. Baseline 2 also uses the same static $\mathbf{P}^*$ but updates the sparse coefficient $\boldsymbol{\alpha}_t$ across time. In contrast, PSMT utilizes a dynamic embedding matrix $\mathbf{P}^*$ as well as a dynamic sparse coefficient $\boldsymbol{\alpha}_t$, which is updated at each future frame k. We measure the log MSE between the reconstructed image patch from the sparse coefficient $\boldsymbol{\alpha}$ learned directly from inputs and the predicted image patch from Step 8 in PSMT (see Algorithm 1).
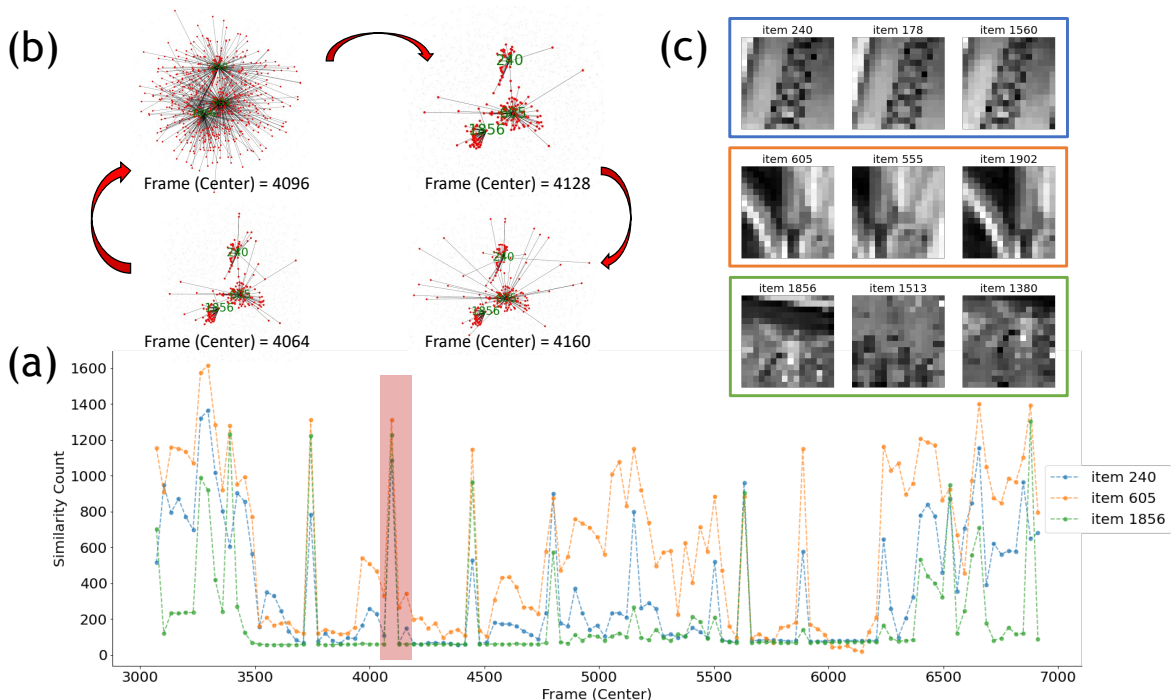
Figure 3: **Dynamic topological analysis of the PSMT embedding space.** (a) Similarity count of three selected dictionary items (240, 605 and 1856) in graph $\mathcal{G}_t$ constructed from $\mathbf{P}^*_t$ center at frame t using $M = 2048$ and $H = 6144$ with $\frac{63}{64}$ overlap. Similarity count is computed as the number of dictionary items in the strongly-connected-component containing selected items. (b) Sequence of (shaded in red) induced subgraph containing selected items. (c) Selected dictionary items with their two most-similar items in $\mathbf{\Phi}_M$.

3. Results

More basis functions and training samples lead to better reconstruction performance. As shown in Table 1 and Figure 2, when using the same number of sparse coding basis functions, the more frames used for training, the lower the average MSEs are generally. Additionally, when using more basis functions, the reconstruction performance is better because the learned dictionary is more expressive.

PSMT embedding space is topologically organized over time. As shown in Figure 3, we focus on three dictionary items (240, 605 and 1856), and each of them is not connected in every $\mathcal{G}_t$ from $\mathbf{P}^*_t$. From Figure 3 (a), we note the actual clusters change over time, which motivates us to iteratively update $\mathbf{P}^*$ for each future frame prediction. From Figure 3 (b), we observe the induced subgraphs from the embedding space $\mathbf{P}^*$ displaying cluster-like behavior and such clusters persist across time. Therefore, each $\mathbf{P}^*_t$ can be topologically organized, so do items in $\mathbf{\Phi}_M$. In Figure 3 (c), we visualize the corresponding items in $\mathbf{\Phi}_M$ and observe the similarity across the rows and dis-similarity across columns, confirming the topology in $\mathbf{P}^*$ corresponds to that in $\mathbf{\Phi}_M$.
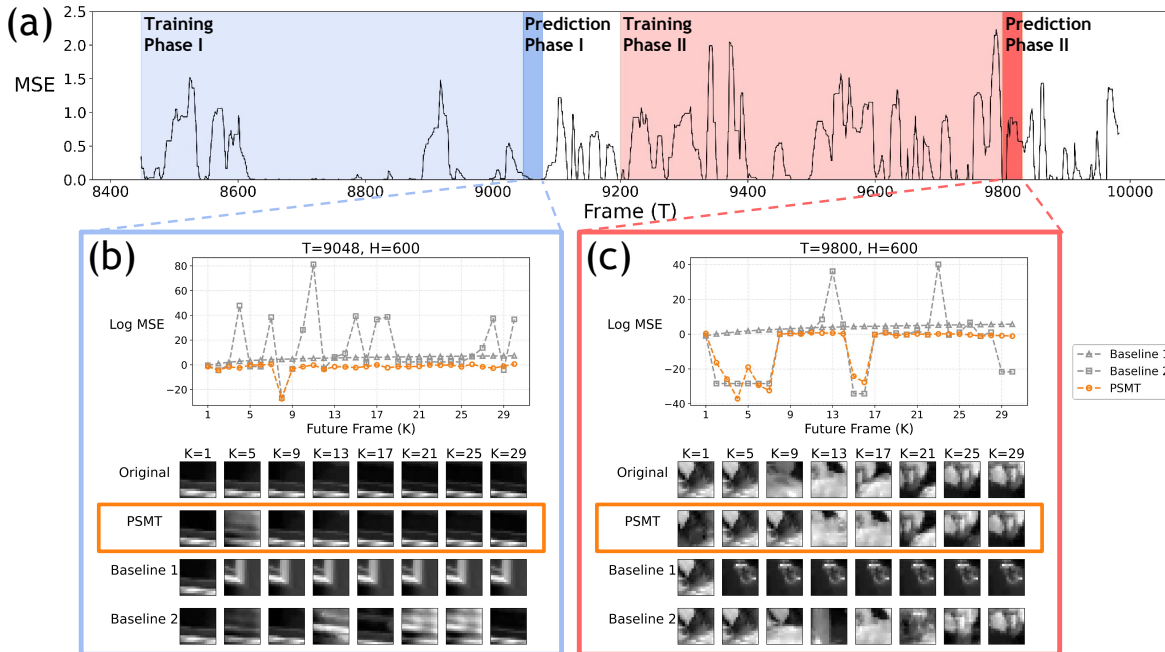
5

Figure 4: **Prediction performance.** (a) MSEs between the reconstructed patches and the original ones (frames $8448 - 9984$) using 512 basis functions. We select two representative phases for prediction, one with relatively static scenes and low MSEs (phase I) and the other with relatively dynamic scenes and high MSEs (phase II). (b-c) Log MSEs between the predicted and original patches ones (top) and examples of predicted patches (bottom) from prediction phase I ($T = 9048$) and II ($T = 9800$). PSMT outperforms two baselines with lower average log MSEs across 30 future frames.

Dynamic PSMT achieves better prediction performance compared to static baselines. We perform prediction on the segment (frames $8448 - 9984$) with minimum average MSE from 512 basis functions for the sake of computational efficiency. As shown in Figure 4 (a), we select two representative prediction phases, one with relatively static scenes and low MSEs (phase I) and the other with relatively dynamic scenes and high MSEs (phase II). We predict 30 future frames using PSMT and two other static baseline approaches in these two phases, respectively. According to Figures 4 (b) and (c), PSMT with a dynamic embedding space outperforms two baseline approaches with a static embedding space. Specifically, PSMT achieves lower average log MSEs ($-2.18, -7.12$) across 30 future frames in both cases. The predicted patches from PSMT better approximate the estimated patches compared to two baselines.

4. Discussion

Contributions Our contributions in this work are multi-fold. First, we propose the PSMT framework for predicting a future temporal sequence of natural scenes. Second, we evaluate the PSMT reconstruction performance with respect to two key hyperparameters and show that more basis func-

tions and training samples usually lead to better performance. Third, we perform dynamic topological analysis of the geometric embedding matrix over time and demonstrate its dynamic nature. Forth, we demonstrate the superiority of PSMT with a dynamic embedding space ($\mathbf{P}^*$ is updated and time-variant) for future temporal sequence prediction compared to two baseline methods with a static embedding space ($\mathbf{P}^*$ is fixed and time-invariant).

Limitations One important limitation of PSMT is the assumption of temporal linearity (imposed by $\mathbf{D}$) in the geometric embedding space, thus it might not be sufficient to capture sharp or large variations in natural scenes such as a flash of light. Another limitation originates from two conditions in sparse coding: (1) the number of basis functions is larger than the image dimension such that the learned dictionary $\mathbf{\Phi}_M$ is overcomplete; (2) the number of frames is larger than the number of basis functions such that the sparse coefficient matrix $\mathbf{A}^*$ is not singular. Thus, small image patches are used to train the model instead of full images. Additionally, we use the analytic solution [10] to derive the embedding matrix $\mathbf{P}^*$ but the solution might be too global to capture local dynamics, especially for a large dictionary with many items. Alternatively, stochastic gradient descent with mini-batches can be used to solve for $\mathbf{P}^*$, which better characterizes locality [1].

Future Work We plan to extend PSMT to learn the dynamic topological organization of brain networks from neuroimaging data, bridging the gap between visual perception and mental representation.

5. Conclusion

We present Predictive Sparse Manifold Transform (PSMT), a two-layer unsupervised generative model to learn and predict natural dynamics. Our results show that PSMT embedding space is topologically organized over time and PSMT using a dynamic embedding space can achieve better prediction performance compared to static baselines, thus establishing an efficient prediction framework.

6. Acknowledgments

References

[1] Yubei Chen, Dylan Paiton, and Bruno Olshausen. The sparse manifold transform. *Advances in neural information processing systems*, 31, 2018.

[2] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. 2004.

[3] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. URL https://www.gurobi.com.

[4] Jeff Hawkins and Sandra Blakeslee. *On intelligence*. Macmillan, 2004.

[5] Linxing Preston Jiang and Rajesh PN Rao. Dynamic predictive coding: A new model of hierarchical sequence learning and prediction in the cortex. *bioRxiv*, pages 2022–06, 2022.

[6] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.

[7] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.

[8] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.

[9] Alexander Sokurov. Russian ark. URL https://youtu.be/PECz8C7m_Yo.

[10] Max Vladymyrov and Miguel Á Carreira-Perpinán. Locally linear landmarks for large-scale manifold learning. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13*, pages 256–271. Springer, 2013.

Appendix A. PSMT Algorithm Details

Here, we describe the Predictive Sparse Manifold Transform (PSMT) framework in detail. PSMT is an iterative procedure that uses the past and current image inputs to predict the next input in a sliding fashion. The numerical inputs to PSMT include

1. T: the starting point for prediction;
2. H: the window size (i.e., number of most-recent consecutive inputs from the past) used for prediction;
3. K: number of future inputs to be predicted;
4. M: number of items in the overcomplete sparse-coding dictionary;
5. N: number of pixels in the input;
6. f: dimension in the embedded space.

Additionally, we have the past image sequence as $\mathbf{X} = \{\boldsymbol{x}_t\}_{t=T-H+1}^{t=T-1}$ and incorporate the future input $\mathbf{X}^* = \{\boldsymbol{x}_t\}_{t=T}^{t=T+K-1}$ iteratively in PSMT.

We initialize PSMT by computing overcomplete sparse-coding dictionary (item size M) $\boldsymbol{\Phi}_M \in \mathbb{R}^{N \times M}$ with $N < M$ together with the corresponding non-negative sparse codes $\mathbf{A} = \{\boldsymbol{\alpha}_t\}_{t=T-H+1}^{T-1}$. That is, $\mathbf{A} \in \mathbb{R}^{M \times (H-1)}$ with $M < H$ subject to $\boldsymbol{\alpha}_t \succeq 0$, to approximate the input $\mathbf{X}$. Here, the Least Angle Regression (LARS) algorithm [2] from sklearn [8] is used to solve this estimation:

$$\mathbf{X} \approx \boldsymbol{\Phi}_M \mathbf{A}. \tag{1}$$

At each time $t = T + k - 1$ for $k = 1, ..., K$, we receive the input $\boldsymbol{x}_{t=T+k-1}$ and want to predict the next input $\hat{\boldsymbol{x}}_{t=T+k}$ using PSMT. For the ease of presentation, we use footnotes *prev* to represent $t = T - 1 + k - 1$, *current* to represent $t = T - 1 + k$, and *future* to represent $t = T + k$.

1. Step 1 to 4 in Algorithm 1: upon receiving the most recent input $\boldsymbol{x}_{\text{current}}$, we compute its sparse code $\boldsymbol{\alpha}_{\text{current}}$ using the initialized dictionary $\boldsymbol{\Phi}_M$. We incorporate it into matrix $\mathbf{A}$. Now, $\mathbf{A} \in \mathbb{R}^{M \times (H+k-1)}$. We only include that latest H sparse codes in PSMT by extracting $t = T - H + k - 1$ to $t = T - 1 + k$ from $\mathbf{A}$ into $\mathbf{A}^*$, making $\mathbf{A}^* \in \mathbb{R}^{M \times (H-1)}$:

$$\mathbf{A}^* \leftarrow \mathbf{A}_{[t=T-H+k-1:t=T-1+k]}. \tag{2}$$

2. Step 5 in Algorithm 1: we solve for the embedding matrix $\mathbf{P}^*$ assuming that the temporal trajectory of $\mathbf{A}$ is linear in the embedding space:

$$\mathbf{P}^* := \text{argmin}_{\mathbf{P}}\{\|\mathbf{P}\mathbf{A}^*\mathbf{D}\|_F^2 : \mathbf{P}\mathbf{V}\mathbf{P}^\top = \mathbf{I}\}, \tag{3}$$

where $\mathbf{D} \in \mathbb{R}^{H \times H}$ is a second-order derivative operator where $\mathbf{D}_{t-1,t} = -0.5$, $\mathbf{D}_{t,t} = 1$, $\mathbf{D}_{t,t+1} = -0.5$ and other entries are zeros. The closed form solution is $\mathbf{P}^* = \mathbf{U}^\top \mathbf{V}^{-\frac{1}{2}}$ with $\mathbf{P}^* \in \mathbb{R}^{f \times M}$ given by [10] where $\mathbf{V}$ is the covariance matrix of $\mathbf{A}^*$ and $\mathbf{U}$ is a matrix of f trailing eigenvectors of the matrix $\mathbf{Q} = \mathbf{V}^{-\frac{1}{2}} \mathbf{A}^* \mathbf{D}\mathbf{D}^\top \mathbf{A}^{*\top} \mathbf{V}^{-\frac{1}{2}}$.

3. Step 6 in Algorithm 1: assume continuous smooth transformation of inputs in the embedding space, adapted from SMT (i.e., $\mathbf{P}^*\boldsymbol{\alpha}_t^* \approx \frac{1}{2}\mathbf{P}^*\boldsymbol{\alpha}_{t-1} + \frac{1}{2}\mathbf{P}^*\boldsymbol{\alpha}_{t+1}$), the future frame $t + 1$ can be approximated as:

$$\mathbf{P}^*\boldsymbol{\alpha}_{t+1}^* \approx 2\mathbf{P}^*\boldsymbol{\alpha}_t - \mathbf{P}^*\boldsymbol{\alpha}_{t-1}. \tag{4}$$

That is, at k, we predict the embedded vector $\boldsymbol{\beta}^*$ at $t = T + k$ as

$$\boldsymbol{\beta}^* := \mathbf{P}^*\boldsymbol{\alpha}_{\text{future}}^* \approx 2\mathbf{P}^*\boldsymbol{\alpha}_{\text{current}} - \mathbf{P}^*\boldsymbol{\alpha}_{\text{prev}}. \tag{5}$$

4. Step 7 in Algorithm 1: we recover the future sparse code at $t = T + k$ from $\boldsymbol{\beta}^* \in \mathbb{R}^f$ by solving the following optimization using Gurobi [3] with $\mathbf{z} = [\|\boldsymbol{p}_1^*\|_2, \ldots, \|\boldsymbol{p}_N^*\|_2]^\top$ (i.e., columns in $\mathbf{P}^*$).

$$\boldsymbol{\alpha}_{\text{future}}^{\text{REC}} := \operatorname{argmin}_{\boldsymbol{\alpha}} \{ \|\mathbf{P}^* \boldsymbol{\alpha} - \boldsymbol{\beta}^*\|_F^2 + \lambda \mathbf{z}^\top \boldsymbol{\alpha} : \boldsymbol{\alpha} \succeq 0 \}. \tag{6}$$

5. Step 8 in Algorithm 1: we predict the future input at $t = T + k$ by:

$$\boldsymbol{x}_{\text{future}} = \boldsymbol{\Phi}_M \boldsymbol{\alpha}_{\text{future}}^{\text{REC}}. \tag{7}$$