
Appendix: Switching Autoregressive Low-rank Tensor Models

Anonymous Author(s)

Affiliation

Address

email

1 A SALT Optimization via Tensor Regression

2 Let $\mathbf{y}_t \in \mathbb{R}^{N_1}$ be the t -th outputs and $\mathbf{X}_t \in \mathbb{R}^{N_2 \times N_3}$ be the t -th inputs. The regression weights are a
3 tensor $\mathcal{A} \in \mathbb{R}^{N_1 \times N_2 \times N_3}$, which we model via a Tucker decomposition,

$$\mathcal{A} = \sum_{i=1}^{D_1} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} g_{ijk} \mathbf{u}_{:i} \circ \mathbf{v}_{:j} \circ \mathbf{w}_{:k}, \quad (1)$$

4 where \mathbf{u}_i , \mathbf{v}_j , and \mathbf{w}_k are columns of the factor matrices $\mathbf{U} \in \mathbb{R}^{N_1 \times D_1}$, $\mathbf{V} \in \mathbb{R}^{N_2 \times D_2}$, and $\mathbf{W} \in \mathbb{R}^{N_3 \times D_3}$,
5 respectively, and g_{ijk} are entries in the core tensor $\mathcal{G} \in \mathbb{R}^{D_1 \times D_2 \times D_3}$.

6 Consider the linear model, $\mathbf{y}_t \sim \mathcal{N}(\mathcal{A} \times_{2,3} \mathbf{X}_t, \mathbf{Q})$ where $\mathcal{A} \times_{2,3} \mathbf{X}_t$ is defined using the Tucker decom-
7 position of \mathcal{A} as,

$$\mathcal{A} \times_{2,3} \mathbf{X}_t = \mathcal{A}_{(1)} \text{vec}(\mathbf{X}_t) \quad (2)$$

$$= \mathbf{U} \mathcal{G}_{(1)} (\mathbf{V}^\top \otimes \mathbf{W}^\top) \text{vec}(\mathbf{X}_t) \quad (3)$$

$$= \mathbf{U} \mathcal{G}_{(1)} \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W}) \quad (4)$$

8 where $\mathcal{A}_{(1)} \in \mathbb{R}^{N_1 \times N_2 N_3}$ and $\mathcal{G}_{(1)} \in \mathbb{R}^{D_1 \times D_2 D_3}$ are mode-1 matricizations of the corresponding tensors.
9 *Note that these equations assume that matricization and vectorization are performed in row-major*
10 *order, as in Python but opposite to what is typically used in Wikipedia articles.*

11 Equation (4) can be written in multiple ways, and these equivalent forms will be useful for deriving
12 the updates below. We have,

$$\mathcal{A} \times_{2,3} \mathbf{X}_t = \mathbf{U} \mathcal{G}_{(1)} (\mathbf{I}_{D_2} \otimes \mathbf{W}^\top \mathbf{X}_t^\top) \text{vec}(\mathbf{V}^\top) \quad (5)$$

$$= \mathbf{U} \mathcal{G}_{(1)} (\mathbf{V}^\top \mathbf{X}_t \otimes \mathbf{I}_{D_3}) \text{vec}(\mathbf{W}) \quad (6)$$

$$= [\mathbf{U} \otimes \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W})] \text{vec}(\mathcal{G}). \quad (7)$$

13 We minimize the negative log likelihood by coordinate descent.

14 **Optimizing the output factors** Let

$$\tilde{\mathbf{x}}_t = \mathcal{G}_{(1)} \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W}) \quad (8)$$

15 for fixed \mathbf{V} , \mathbf{W} , and \mathcal{G} . The NLL as a function of \mathbf{U} is,

$$\mathcal{L}(\mathbf{U}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \mathbf{U} \tilde{\mathbf{x}}_t)^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \mathbf{U} \tilde{\mathbf{x}}_t). \quad (9)$$

16 This is a standard least squares problem with solution

$$\mathbf{U}^* = \left(\sum_t \mathbf{y}_t \tilde{\mathbf{x}}_t^\top \right) \left(\sum_t \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \right)^{-1}. \quad (10)$$

17 **Optimizing the core tensors** Let $\tilde{\mathbf{X}}_t = \mathbf{U} \otimes \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W}) \in \mathbb{R}^{N_1 \times D_1 D_2 D_3}$ denote the coefficient on
 18 $\text{vec}(\mathcal{G})$ in eq. (7). The NLL as a function of $\mathbf{g} = \text{vec}(\mathcal{G})$ is,

$$\mathcal{L}(\mathbf{g}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{g})^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{g}). \quad (11)$$

19 The minimizer of this quadratic form is,

$$\mathbf{g}^\star = \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{X}}_t \right)^{-1} \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \mathbf{y}_t \right) \quad (12)$$

20 **Optimizing the input factors** Let

$$\tilde{\mathbf{X}}_t = \mathbf{U} \mathcal{G}_{(1)} (\mathbf{I}_{D_2} \otimes \mathbf{W}^\top \mathbf{X}_t^\top) \quad (13)$$

21 for fixed \mathbf{U} , \mathbf{W} , and \mathcal{G} . The NLL as a function of $\mathbf{v} = \text{vec}(\mathbf{V}^\top)$ is,

$$\mathcal{L}(\mathbf{v}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{v})^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{v}). \quad (14)$$

22 The minimizer of this quadratic form is,

$$\mathbf{v}^\star = \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{X}}_t \right)^{-1} \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \mathbf{y}_t \right) \quad (15)$$

23 **Optimizing the lag factors** Let

$$\tilde{\mathbf{X}}_t = \mathbf{U} \mathcal{G}_{(1)} (\mathbf{V}^\top \mathbf{X}_t \otimes \mathbf{I}_{D_3}) \quad (16)$$

24 for fixed \mathbf{U} , \mathbf{V} , and \mathcal{G} . The NLL as a function of $\mathbf{w} = \text{vec}(\mathbf{W})$ is,

$$\mathcal{L}(\mathbf{w}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{w})^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{w}). \quad (17)$$

25 The minimizer of this quadratic form is,

$$\mathbf{w}^\star = \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{X}}_t \right)^{-1} \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \mathbf{y}_t \right) \quad (18)$$

26 **Multiple discrete states** If we have discrete states $z_t \in \{1, \dots, H\}$ and each state has its own
 27 parameters $(\mathcal{G}^{(h)}, \mathbf{U}^{(h)}, \mathbf{V}^{(h)}, \mathbf{W}^{(h)}, \mathbf{Q}^{(h)})$, then letting $\omega_t^{(h)} = \mathbb{E}[z_t = h]$ denote the weights from the
 28 E-step, the summations in coordinate updates are weighted by $\omega_t^{(h)}$. For example, the coordinate
 29 update for the core tensors becomes,

$$\mathbf{g}^{(h)\star} = \left(\sum_t \omega_t^{(h)} \tilde{\mathbf{X}}_t^{(h)\top} \mathbf{Q}^{(h)-1} \tilde{\mathbf{X}}_t^{(h)} \right)^{-1} \left(\sum_t \omega_t^{(h)} \tilde{\mathbf{X}}_t^{(h)\top} \mathbf{Q}^{(h)-1} \mathbf{y}_t \right) \quad (19)$$

30 B SALT approximates a Linear Dynamical System

31 A stationary linear dynamical system (LDS) is defined as follows:

$$\begin{aligned} \mathbf{x}_t &= \mathbf{A} \mathbf{x}_{t-1} + \mathbf{b} + \boldsymbol{\epsilon}_t \\ \mathbf{y}_t &= \mathbf{C} \mathbf{x}_t + \mathbf{d} + \boldsymbol{\delta}_t \end{aligned}$$

32 where $\mathbf{y}_t \in \mathbb{R}^{N_1}$ is the t -th observation, $\mathbf{x}_t \in \mathbb{R}^{D_1}$ is the t -th hidden state, $\boldsymbol{\epsilon}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{Q})$, $\boldsymbol{\delta}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{R})$,
 33 and $\boldsymbol{\theta} = (\mathbf{A}, \mathbf{b}, \mathbf{Q}, \mathbf{C}, \mathbf{d}, \mathbf{R})$ are the parameters of the LDS.

34 Following the notation of Murphy [1], the one-step-ahead posterior predictive distribution for the
 35 observations of the LDS defined above can be expressed as:

$$p(\mathbf{y}_t | \mathbf{y}_{1:t-1}) = \mathcal{N}(\mathbf{C} \boldsymbol{\mu}_{t|t-1} + \mathbf{d}, \mathbf{C} \boldsymbol{\Sigma}_{t|t-1} \mathbf{C}^\top + \mathbf{R})$$

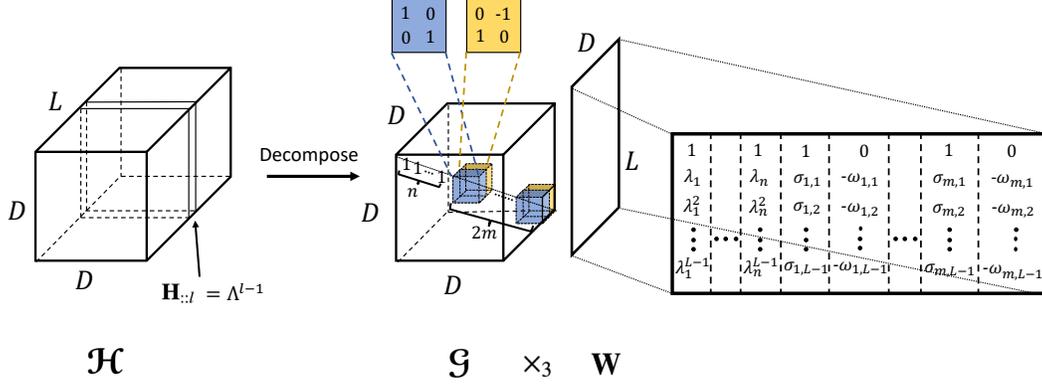


Figure 1: **Decomposition of \mathcal{H} into \mathcal{G} and \mathcal{W} such that $\mathcal{H} = \mathcal{G} \times_3 \mathcal{W}$:** Given an LDS whose $\mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{C})$ has n real eigenvalues and m pairs of complex eigenvalues, this decomposition illustrates how Tucker-SALT can approximate the LDS well with rank $n + 2m$.

64 With $\mathbf{V} = (\mathbf{E}^{-1} \mathbf{A} \mathbf{K})^T$, $\mathbf{U} = \mathbf{C} \mathbf{E}$, $\mathbf{m} = \mathbf{C} \sum_{l=1}^{N_3} \Gamma_{l,\text{stable}}(\mathbf{b} - \mathbf{A} \mathbf{K} \mathbf{d}) + \mathbf{d}$, and $\mathbf{X}_t = \mathbf{y}_{t-1:t-N_3}$, we can rearrange
 65 the mean to:

$$\begin{aligned}
 \mathbf{C} \boldsymbol{\mu}_{t|t-1} + \mathbf{d} &\approx \mathbf{C} \sum_{l=1}^{N_3} \mathbf{E} \boldsymbol{\Lambda}^{l-1} \mathbf{E}^{-1} \mathbf{A} \mathbf{K} \mathbf{y}_{t-l} + \mathbf{C} \sum_{l=1}^{N_3} \Gamma_{l,\text{stable}}(\mathbf{b} - \mathbf{A} \mathbf{K} \mathbf{d}) + \mathbf{d} \\
 &= \mathbf{U} \sum_{l=1}^{N_3} \mathbf{H}_{::l} \mathbf{V}^T \mathbf{y}_{t-l} + \mathbf{m} \\
 &= \mathbf{U} \sum_{l=1}^{N_3} (\mathcal{G} \bar{\times}_3 \mathbf{w}_l) \mathbf{V}^T \mathbf{y}_{t-l} + \mathbf{m} \\
 &= \mathbf{U} \sum_{l=1}^{N_3} ((\mathcal{G} \times_2 \mathbf{V}) \bar{\times}_3 \mathbf{w}_l) \mathbf{y}_{t-l} + \mathbf{m} \\
 &= \mathbf{U} \sum_{l=1}^{N_3} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} \mathbf{g}_{:jk} \circ \mathbf{v}_{:j} (\mathbf{w}_{lk} \mathbf{y}_{t-l}) + \mathbf{m} \\
 &= \mathbf{U} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} \mathbf{g}_{:jk} (\mathbf{v}_{:j}^\top \mathbf{X}_t \mathbf{w}_{:k}) + \mathbf{m} \\
 &= \sum_{i=1}^{D_1} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} \mathbf{u}_{:i} g_{ijk} (\mathbf{v}_{:j}^\top \mathbf{X}_t \mathbf{w}_{:k}) + \mathbf{m} \\
 &= \left[\sum_{i=1}^{D_1} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} g_{ijk} \mathbf{u}_{:i} \circ \mathbf{v}_{:j} \circ \mathbf{w}_{:k} \right] \times_{2,3} \mathbf{X}_t + \mathbf{m}
 \end{aligned}$$

66 **CP Tensor Regression** By rearranging $\mathbf{E}, \boldsymbol{\Lambda}^l$, and \mathbf{E}^{-1} into \mathbf{J}, \mathbf{P}_l , and \mathbf{S} respectively as fol-
 67 lows:

$$\mathbf{J} = \left[\mathbf{a}_1 \dots \mathbf{a}_n \mathbf{b}_1 + \mathbf{c}_1 \mathbf{b}_1 \mathbf{c}_1 \dots \mathbf{b}_m + \mathbf{c}_m \mathbf{b}_m \mathbf{c}_m \right]$$

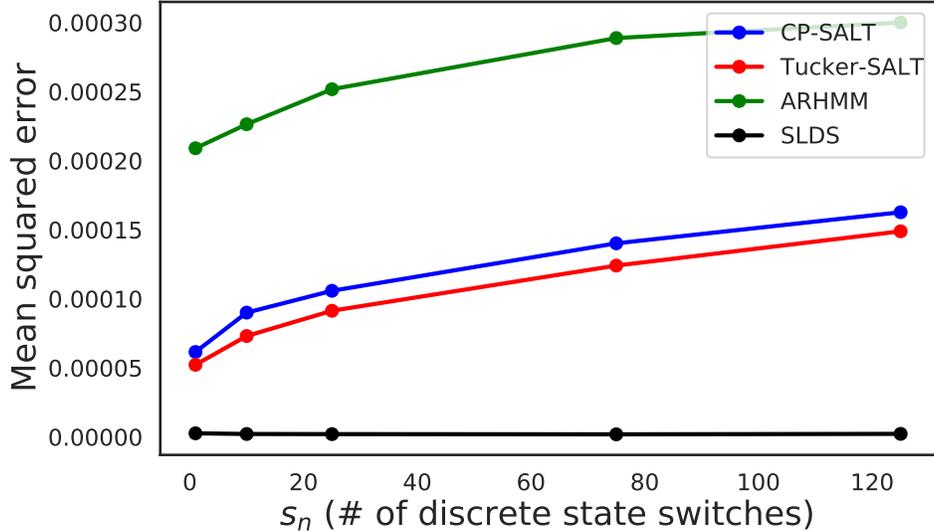


Figure 2: The quality of SALT approximation of SLDSs decreases as the number of discrete state switches increases: The data comes from an SLDS with $H = 2$, $N = 20$, and $D = 7$. 15,000 timesteps were generated, with varying numbers of evenly spaced out discrete state switches (x-axis). The mean squared error of reconstructing the autoregressive tensors increased with the number of discrete state switches.

80 This is because in SALT, the autoregressive dynamics only depend on the current discrete state.
 81 However, in an SLDS, when we marginalize out the continuous latent states, the autoregressive
 82 dynamics in the observation space will depend on previous discrete states as well.

83 C Modeling Mouse Behavior: Videos

84 Here we describe how the mouse behavioral videos were generated. We first determined the CP-
 85 SALT hyperparameters as those which led to the highest log-likelihood on the validation dataset.
 86 Then, using that CP-SALT model, we computed the most likely discrete states on the train and test
 87 data. Given a discrete state h , we extracted slices of the data whose most likely discrete state was
 88 h . We padded the data by 30 frames (i.e. 1 second) both at the beginning and the end of each slice
 89 for the movie. A red dot appears on each mouse for the duration of discrete state h . We generated
 90 such videos for all 50 discrete states (as long as there existed at least one slice for each discrete state)
 91 on the train and test data. For a given discrete state, the mice in each video behaved very similarly
 92 (e.g., the mice in the video for state 18 “pause” when the red dots appear, and those in the video for
 93 state 32 “walk” forward), suggesting that CP-SALT is capable of segmenting the data into useful
 94 behavioral syllables.

95 D Additional Synthetic Data Experiments

96 D.1 The effect of the number of switches on the recovery of the parameters of the 97 autoregressive dynamic tensors

98 We asked how the number of discrete state switches affected SALT’s ability to recover the autoregres-
 99 sive tensors. We trained CP-SALT, Tucker-SALT, the ARHMM, all with $L = 5$ lags, and the SLDS
 100 on data sampled from an SLDS with varying number of discrete state switches. The ground-truth
 101 SLDS model had $H = 2$ discrete states, $N = 20$ observations and $D = 7$ dimensional continuous
 102 latent states. The matrix $\mathbf{A}^{(h)}(\mathbf{I} - \mathbf{K}^{(h)}\mathbf{C}^{(h)})$ of each discrete state of the ground-truth SLDS had 1 real
 103 eigenvalue and 3 pairs of complex eigenvalues. We sampled 5 batches of $T = 15,000$ timesteps of
 104 data from the ground-truth SLDS, with $s_n \in \{1, 10, 25, 75, 125\}$ number of discrete state switches that
 105 were evenly spaced out across the data. We then computed the mean squared error (MSE) between
 106 the SLDS tensors and the tensors reconstructed by SALT, the ARHMM, and the SLDS. (Figure 2).
 107 More precisely, we combined the 3rd order autoregressive tensors from each discrete into a 4th order
 108 tensor, and calculated the MSE based on these 4th order tensors. As expected, the MSE increased

109 with the number of switches in the data, indicating that the quality of SALT approximation of SLDSs
110 decreases as the number of discrete state switches increases.

111 **E Code**

112 As part of Supplementary materials, we include the source code for SALT (in Python) and an example
113 Jupyter Notebook for sampling data from SALT and fitting SALT to the data with an EM algorithm.
114 SALT code requires the [ssm-jax](#) package.

115 **References**

- 116 [1] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- 117 [2] Panqanamala Ramana Kumar and Pravin Varaiya. *Stochastic systems: Estimation, identification,*
118 *and adaptive control*. SIAM, 2015.
- 119 [3] Tohru Katayama et al. *Subspace methods for system identification*, volume 1. Springer, 2005.