

417 **Supplementary Materials for: Switching Autoregressive Low-rank**
418 **Tensor Models**

419

420 **Table of Contents**

- 421 – **Appendix A:** SALT Optimization via Tensor Regression.
- 422 – **Appendix B:** SALT approximates a (Switching) Linear Dynamical System.
- 423 – **Appendix C:** Single-subspace SALT.
- 424 – **Appendix D:** Synthetic Data Experiments.
- 425 – **Appendix E:** Modeling Mouse Behavior.
- 426 – **Appendix F:** Modeling *C. elegans* neural data.
- 427 – **Appendix G:** Code Availability.

428 **A SALT Optimization via Tensor Regression**

429 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
 430 tensor $\mathcal{A} \in \mathbb{R}^{N_1 \times N_2 \times N_3}$, which we model via a Tucker decomposition,

$$\mathcal{A} = \sum_{i=1}^{D_1} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} g_{ijk} \mathbf{u}_{:i} \circ \mathbf{v}_{:j} \circ \mathbf{w}_{:k}, \quad (10)$$

431 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 $\mathbf{W} \in$
 432 $\mathbb{R}^{N_3 \times D_3}$, respectively, and g_{ijk} are entries in the core tensor $\mathcal{G} \in \mathbb{R}^{D_1 \times D_2 \times D_3}$.

433 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
 434 decomposition of \mathcal{A} as,

$$\mathcal{A} \times_{2,3} \mathbf{X}_t = \mathcal{A}_{(1)} \text{vec}(\mathbf{X}_t) \quad (11)$$

$$= \mathbf{U} \mathcal{G}_{(1)} (\mathbf{V}^\top \otimes \mathbf{W}^\top) \text{vec}(\mathbf{X}_t) \quad (12)$$

$$= \mathbf{U} \mathcal{G}_{(1)} \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W}) \quad (13)$$

435 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
 436 tensors. *Note that these equations assume that matricization and vectorization are performed in*
 437 *row-major order, as in Python but opposite to what is typically used in Wikipedia articles.*

438 Equation (13) can be written in multiple ways, and these equivalent forms will be useful for deriving
 439 the updates below. We have,

$$\mathcal{A} \times_{2,3} \mathbf{X}_t = \mathbf{U} \mathcal{G}_{(1)} (\mathbf{I}_{D_2} \otimes \mathbf{W}^\top \mathbf{X}_t^\top) \text{vec}(\mathbf{V}^\top) \quad (14)$$

$$= \mathbf{U} \mathcal{G}_{(1)} (\mathbf{V}^\top \mathbf{X}_t \otimes \mathbf{I}_{D_3}) \text{vec}(\mathbf{W}) \quad (15)$$

$$= [\mathbf{U} \otimes \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W})] \text{vec}(\mathcal{G}). \quad (16)$$

440 We minimize the negative log likelihood by coordinate descent.

441 **Optimizing the output factors** Let

$$\tilde{\mathbf{x}}_t = \mathcal{G}_{(1)} \text{vec}(\mathbf{V}^\top \mathbf{X}_t \mathbf{W}) \quad (17)$$

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

$$\mathcal{L}(\mathbf{U}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \mathbf{U} \tilde{\mathbf{x}}_t)^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \mathbf{U} \tilde{\mathbf{x}}_t). \quad (18)$$

443 This is a standard least squares problem with solution

$$\mathbf{U}^* = \left(\sum_t \mathbf{y}_t \tilde{\mathbf{x}}_t^\top \right) \left(\sum_t \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \right)^{-1}. \quad (19)$$

444 **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
 445 on $\text{vec}(\mathcal{G})$ in eq. (16). The NLL as a function of $\mathbf{g} = \text{vec}(\mathcal{G})$ is,

$$\mathcal{L}(\mathbf{g}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{g})^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{g}). \quad (20)$$

446 The minimizer of this quadratic form is,

$$\mathbf{g}^* = \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{X}}_t \right)^{-1} \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \mathbf{y}_t \right) \quad (21)$$

447 **Optimizing the input factors** Let

$$\tilde{\mathbf{X}}_t = \mathbf{U}\mathcal{G}_{(1)}(\mathbf{I}_{D_2} \otimes \mathbf{W}^\top \mathbf{X}_t^\top) \quad (22)$$

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

$$\mathcal{L}(\mathbf{v}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{v})^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{v}). \quad (23)$$

449 The minimizer of this quadratic form is,

$$\mathbf{v}^* = \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{X}}_t \right)^{-1} \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \mathbf{y}_t \right) \quad (24)$$

450 **Optimizing the lag factors** Let

$$\tilde{\mathbf{X}}_t = \mathbf{U}\mathcal{G}_{(1)}(\mathbf{V}^\top \mathbf{X}_t \otimes \mathbf{I}_{D_3}) \quad (25)$$

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

$$\mathcal{L}(\mathbf{w}) = \frac{1}{2} \sum_t (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{w})^\top \mathbf{Q}^{-1} (\mathbf{y}_t - \tilde{\mathbf{X}}_t \mathbf{w}). \quad (26)$$

452 The minimizer of this quadratic form is,

$$\mathbf{w}^* = \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{X}}_t \right)^{-1} \left(\sum_t \tilde{\mathbf{X}}_t^\top \mathbf{Q}^{-1} \mathbf{y}_t \right) \quad (27)$$

453 **Multiple discrete states** If we have discrete states $z_t \in \{1, \dots, H\}$ and each state has its own
 454 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
 455 the E-step, the summations in coordinate updates are weighted by $\omega_t^{(h)}$. For example, the coordinate
 456 update for the core tensors becomes,

$$\mathbf{g}^{(h)*} = \left(\sum_t \omega_t^{(h)} \tilde{\mathbf{X}}_t^{(h)\top} \mathbf{Q}^{(h)-1} \tilde{\mathbf{X}}_t^{(h)} \right)^{-1} \left(\sum_t \omega_t^{(h)} \tilde{\mathbf{X}}_t^{(h)\top} \mathbf{Q}^{(h)-1} \mathbf{y}_t \right) \quad (28)$$

457 **B SALT approximates a (Switching) Linear Dynamical System**

458 We now re-state and provide a full proof for Proposition 1.

459 **Proposition 1** (Low-Rank Tensor Autoregressions Approximate Stable Linear Dynamical Systems).
 460 Consider a stable linear time-invariant Gaussian dynamical system. We define the steady-state
 461 Kalman gain matrix as $\mathbf{K} = \lim_{t \rightarrow \infty} \mathbf{K}_t$, and $\mathbf{\Gamma} = \mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{C})$. The matrix $\mathbf{\Gamma} \in \mathbb{R}^{D \times D}$ has
 462 eigenvalues $\lambda_1, \dots, \lambda_D$. Let $\lambda_{\max} = \max_d |\lambda_d|$; for a stable LDS, $\lambda_{\max} < 1$ [22]. Let n denote the
 463 number of real eigenvalues and m the number of complex conjugate pairs. Let $\hat{\mathbf{y}}_t^{(\text{LDS})} = \mathbb{E}[\mathbf{y}_t | \mathbf{y}_{<t}]$
 464 denote the predictive mean under a steady-state LDS, and $\hat{\mathbf{y}}_t^{(\text{SALT})}$ the predictive mean under a SALT
 465 model. An order- L Tucker-SALT model with rank $n + 2m$, or a CP-SALT model with rank $n + 3m$, can
 466 approximate the predictive mean of the steady-state LDS with error $\|\hat{\mathbf{y}}_t^{(\text{LDS})} - \hat{\mathbf{y}}_t^{(\text{SALT})}\|_{\infty} = \mathcal{O}(\lambda_{\max}^L)$.

467 *Proof.* A stationary linear dynamical system (LDS) is defined as follows:

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{b} + \boldsymbol{\epsilon}_t \quad (29)$$

$$\mathbf{y}_t = \mathbf{C}\mathbf{x}_t + \mathbf{d} + \boldsymbol{\delta}_t \quad (30)$$

468 where $\mathbf{y}_t \in \mathbb{R}^N$ is the t -th observation, $\mathbf{x}_t \in \mathbb{R}^D$ is the t -th hidden state, $\boldsymbol{\epsilon}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{Q})$, $\boldsymbol{\delta}_t \stackrel{\text{i.i.d.}}{\sim}$
 469 $\mathcal{N}(\mathbf{0}, \mathbf{R})$, and $\boldsymbol{\theta} = (\mathbf{A}, \mathbf{b}, \mathbf{Q}, \mathbf{C}, \mathbf{d}, \mathbf{R})$ are the parameters of the LDS.

470 Following the notation of Murphy [19], the one-step-ahead posterior predictive distribution for the
 471 observations of the LDS defined above can be expressed as:

$$p(\mathbf{y}_t | \mathbf{y}_{1:t-1}) = \mathcal{N}(\mathbf{C}\boldsymbol{\mu}_{t|t-1} + \mathbf{d}, \mathbf{C}\boldsymbol{\Sigma}_{t|t-1}\mathbf{C}^T + \mathbf{R}) \quad (31)$$

472 where

$$\boldsymbol{\mu}_{t|t-1} = \mathbf{A}\boldsymbol{\mu}_{t-1} + \mathbf{b} \quad (32)$$

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t \mathbf{r}_t \quad (33)$$

$$\boldsymbol{\Sigma}_{t|t-1} = \mathbf{A}\boldsymbol{\Sigma}_{t-1}\mathbf{A}^T + \mathbf{Q} \quad (34)$$

$$\boldsymbol{\Sigma}_t = (\mathbf{I} - \mathbf{K}_t\mathbf{C})\boldsymbol{\Sigma}_{t|t-1} \quad (35)$$

$$p(\mathbf{x}_1) = \mathcal{N}(\boldsymbol{\mu}_{1|0}, \boldsymbol{\Sigma}_{1|0}) \quad (36)$$

$$\mathbf{K}_t = (\boldsymbol{\Sigma}_{t|t-1}^{-1} + \mathbf{C}^T\mathbf{R}\mathbf{C})^{-1}\mathbf{C}^T\mathbf{R}^{-1} \quad (37)$$

$$\mathbf{r}_t = \mathbf{y}_t - \mathbf{C}\boldsymbol{\mu}_{t|t-1} - \mathbf{d}. \quad (38)$$

473 We can then expand the mean $\mathbf{C}\boldsymbol{\mu}_{t|t-1} + \mathbf{d}$ as follows:

$$\mathbf{C}\boldsymbol{\mu}_{t|t-1} + \mathbf{d} = \mathbf{C} \sum_{l=1}^{t-1} \boldsymbol{\Gamma}_l \mathbf{A} \mathbf{K}_{t-l} \mathbf{y}_{t-l} + \mathbf{C} \sum_{l=1}^{t-1} \boldsymbol{\Gamma}_l (\mathbf{b} - \mathbf{A} \mathbf{K}_{t-l} \mathbf{d}) + \mathbf{d} \quad (39)$$

474 where

$$\boldsymbol{\Gamma}_l = \prod_{i=1}^{l-1} \mathbf{A}(\mathbf{I} - \mathbf{K}_{t-i}\mathbf{C}) \quad \text{for } l \in \{2, 3, \dots\}, \quad (40)$$

$$\boldsymbol{\Gamma}_1 = \mathbf{I}. \quad (41)$$

475 Theorem 3.3.3 of Davis and Vinter [22] (reproduced with our notation below) states that for a
 476 stabilizable and detectable system, the $\lim_{t \rightarrow \infty} \boldsymbol{\Sigma}_{t|t-1} = \boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma}$ is the unique solution of the
 477 discrete algebraic Riccati equation

$$\boldsymbol{\Sigma} = \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^T - \mathbf{A}\boldsymbol{\Sigma}\mathbf{C}^T(\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}^T + \mathbf{R})^{-1}\mathbf{C}\boldsymbol{\Sigma}\mathbf{A}^T + \mathbf{Q}. \quad (42)$$

478 As we are considering stable autonomous LDSs here, the system is stabilizable and detectable, as all
 479 unobservable states are themselves stable [22, 36]

480 **Theorem 3.3.3** [Reproduced from Davis and Vinter [22], updated to our notation and context].

481 (a) If the pair (\mathbf{A}, \mathbf{C}) is detectable then there exists at least one non-negative solution to the discrete

482 algebraic Riccati equation (42).

483 (b) If further the pair (\mathbf{A}, \mathbf{C}) is stabilizable then this solution Σ is unique, and $\Sigma_{t|t-1} \rightarrow \Sigma$ as
 484 $t \rightarrow \infty$, where $\Sigma_{t|t-1}$ is the sequence generated by (32)-(38) with arbitrary initial covariance Σ_0 .
 485 The matrix $\Gamma = \mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{C})$ is stable, where \mathbf{K} is the Kalman gain corresponding to Σ , i.e.

$$\mathbf{K} = (\Sigma^{-1} + \mathbf{C}^T \mathbf{R} \mathbf{C})^{-1} \mathbf{C}^T \mathbf{R}^{-1} \quad (43)$$

486 **Proof.** See Davis and Vinter [22]. Note that Davis and Vinter [22] define the Kalman gain as $\mathbf{A}\mathbf{K}$.

487 The convergence of the Kalman gain also implies that each term in the sequence Γ_l converges to

$$\Gamma_l = \prod_{i=1}^{l-1} \mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{C}) = (\mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{C}))^{l-1} = \Gamma^{l-1}, \quad (44)$$

488 where, concretely, we define $\Gamma = \mathbf{A}(\mathbf{I} - \mathbf{K}\mathbf{C})$. We can therefore make the following substitution and
 489 approximation

$$\mathbf{C}\mu_{t|t-1} + \mathbf{d} \stackrel{\lim_{t \rightarrow \infty}}{=} \mathbf{C} \sum_{l=1}^{t-1} \Gamma^l \mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{C} \sum_{l=1}^{t-1} \Gamma^l (\mathbf{b} - \mathbf{A}\mathbf{K}\mathbf{d}) + \mathbf{d} \quad (45)$$

$$= \mathbf{C} \sum_{l=1}^L \Gamma^l \mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{C} \sum_{l=1}^L \Gamma^l (\mathbf{b} - \mathbf{A}\mathbf{K}\mathbf{d}) + \mathbf{d} + \sum_{l=L+1}^{\infty} \mathcal{F}(\Gamma^l) \quad (46)$$

$$\approx \mathbf{C} \sum_{l=1}^L \Gamma^l \mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{C} \sum_{l=1}^L \Gamma^l (\mathbf{b} - \mathbf{A}\mathbf{K}\mathbf{d}) + \mathbf{d} \quad (47)$$

490 The approximation is introduced as a result of truncating the sequence to consider just the “first” L
 491 terms, and discarding the higher-order terms (indicated in blue). It is important to note that each term
 492 in (45) is the sum of a geometric sequence multiplied elementwise with \mathbf{y}_t .

493 There are two components we prove from here. First, we derive an element-wise bound on the error
 494 introduced by the truncation, and verify that under the conditions outlined that the bound decays
 495 monotonically in L . We then show that Tucker and CP decompositions can represent the truncated
 496 summations in (47), and derive the minimum rank required for this representation to be exact.

497 **Bounding The Error Term** We first rearrange the truncated terms in (45), where we define
 498 $\mathbf{x}_l \triangleq \mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{b} - \mathbf{A}\mathbf{K}\mathbf{d}$

$$\sum_{l=L+1}^{\infty} \mathcal{F}(\Gamma^l) = \mathbf{C} \sum_{l=L+1}^{\infty} \Gamma^l \mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{C} \sum_{l=L+1}^{\infty} \Gamma^l (\mathbf{b} - \mathbf{A}\mathbf{K}\mathbf{d}) + \mathbf{d}, \quad (48)$$

$$= \sum_{l=L+1}^{\infty} \mathbf{C}\Gamma^l \mathbf{x}_l, \quad (49)$$

$$= \sum_{l=L+1}^{\infty} \mathbf{C}\mathbf{E}\Lambda^{l-1}\mathbf{E}^{-1}\mathbf{x}_l, \quad (50)$$

$$= \sum_{l=L+1}^{\infty} \mathbf{P}\Lambda^{l-1}\mathbf{q}_l, \quad (51)$$

499 where $\mathbf{E}\Lambda\mathbf{E}^{-1}$ is the eigendecomposition of Γ , $\mathbf{P} \triangleq \mathbf{C}\mathbf{E}$, and $\mathbf{q}_l \triangleq \mathbf{E}^{-1}\mathbf{x}_l$. We now consider the
 500 infinity-norm of the error, and apply the triangle and Cauchy-Schwartz inequalities. We can write the
 501 bound on the as

$$\epsilon = \left| \left(\sum_{l=L+1}^{\infty} \mathcal{F}(\Gamma^l) \right)_n \right|, \quad \text{where } n = \arg \max_k \left| \left(\sum_{l=L+1}^{\infty} \mathcal{F}(\Gamma^l) \right)_k \right| \quad (52)$$

$$= \left| \sum_{l=L+1}^{\infty} \sum_{d=1}^D p_{nd} \lambda_d^{l-1} q_{l,d} \right|, \quad (53)$$

$$\leq \sum_{l=L+1}^{\infty} \sum_{d=1}^D |p_{nd}| |\lambda_d^{l-1}| |q_{l,d}|. \quad (54)$$

541 superdiagonal 3-way tensor, where $g_{ijk} = \mathbb{1}_{i=j=k}$. We can then rearrange the mean to:

$$\mathbf{C}\boldsymbol{\mu}_{t|t-1} + \mathbf{d} \approx \mathbf{C} \sum_{l=1}^L \mathbf{E}\boldsymbol{\Lambda}^{l-1} \mathbf{E}^{-1} \mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{C} \sum_{l=1}^L \boldsymbol{\Gamma}^l (\mathbf{b} - \mathbf{A}\mathbf{K}\mathbf{d}) + \mathbf{d} \quad (77)$$

$$= \mathbf{C} \sum_{l=1}^L \mathbf{J}\mathbf{P}_{l-1} \mathbf{S}\mathbf{A}\mathbf{K}\mathbf{y}_{t-l} + \mathbf{m} \quad (78)$$

$$= \mathbf{U} \sum_{l=1}^L \mathbf{P}_{l-1} \mathbf{V}^\top \mathbf{y}_{t-l} + \mathbf{m} \quad (79)$$

$$= \sum_{l=1}^L \sum_i^{n+3m} \sum_j^{n+3m} \sum_k^{n+3m} g_{ijk} \mathbf{u}_{:i} \circ \mathbf{v}_{:j} (p_{l-1,kk} \mathbf{y}_{t-l}) + \mathbf{m} \quad (80)$$

$$= \sum_i^{n+3m} \sum_j^{n+3m} \sum_k^{n+3m} g_{ijk} \mathbf{u}_{:i} \circ \mathbf{v}_{:j} (\mathbf{X}_t \mathbf{w}_{:k}) + \mathbf{m} \quad (81)$$

$$= \left[\sum_{i=1}^{n+3m} \sum_{j=1}^{n+3m} \sum_{k=1}^{n+3m} g_{ijk} \mathbf{u}_{:i} \circ \mathbf{v}_{:j} \circ \mathbf{w}_{:k} \right] \times_{2,3} \mathbf{X}_t + \mathbf{m} \quad (82)$$

542 And so concludes the proof. □

543 **C Single-subspace SALT**

544 Here we explicitly define the generative model of multi-subspace and single-subspace Tucker-SALT
 545 and CP-SALT. Single-subspace SALT is analogous to single-subspace SLDSs (also defined below),
 546 where certain emission parameters (e.g., \mathbf{C} , \mathbf{d} , and \mathbf{R}) are shared across discrete states. This reduces
 547 the expressivity of the model, but also reduces the number of parameters in the model. Note that both
 548 variants of all models have the same structure on the transition dynamics of \mathbf{z}_t .

549 **Multi-subspace SALT** Note that the SALT model defined in (6) and (7) in the main text was a
 550 multi-subspace SALT. We repeat the definition here for ease of comparison

$$\mathbf{y}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N} \left(\left(\sum_{i=1}^{D_1} \sum_{j=1}^{D_2} \sum_{k=1}^{D_3} g_{ijk}^{(z_t)} \mathbf{u}_{:i}^{(z_t)} \circ \mathbf{v}_{:j}^{(z_t)} \circ \mathbf{w}_{:k}^{(z_t)} \right) \times_{2,3} \mathbf{y}_{t-1:t-L} + \mathbf{b}^{(z_t)}, \boldsymbol{\Sigma}^{(z_t)} \right), \quad (83)$$

551 $D_1 = D_2 = D_3 = D$ and \mathcal{G} is diagonal for CP-SALT.

552 **Single-subspace Tucker-SALT** In single-subspace methods, the output factors are shared across
 553 discrete states

$$\mathbf{y}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N} \left(\mathbf{U} \left(\mathbf{m}^{(z_t)} + \left(\sum_{j=1}^{D_2} \sum_{k=1}^{D_3} \mathbf{g}_{:jk}^{(z_t)} \circ \mathbf{v}_{:j}^{(z_t)} \circ \mathbf{w}_{:k}^{(z_t)} \right) \times_{2,3} \mathbf{y}_{t-1:t-L} \right) + \mathbf{b}, \boldsymbol{\Sigma}^{(z_t)} \right), \quad (84)$$

554 where $\mathbf{m}^{(z_t)} \in \mathbb{R}^{D_1}$.

555 **Single-subspace CP-SALT** Single-subspace CP-SALT requires an extra tensor compared to
 556 Tucker-SALT, as this tensor can no longer be absorbed in to the core tensor.

$$\mathbf{y}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N} \left(\mathbf{U}' \left(\mathbf{m}^{(z_t)} + \mathbf{P}^{(z_t)} \left(\left(\sum_{j=1}^{D_2} \sum_{k=1}^{D_3} \mathbf{g}_{:jk}^{(z_t)} \circ \mathbf{v}_{:j}^{(z_t)} \circ \mathbf{w}_{:k}^{(z_t)} \right) \times_{2,3} \mathbf{y}_{t-1:t-L} \right) \right) + \mathbf{b}, \boldsymbol{\Sigma}^{(z_t)} \right), \quad (85)$$

557 where $\mathbf{U}' \in \mathbb{R}^{N \times D_1'}$, $\mathbf{P}^{(z_t)} \in \mathbb{R}^{D_1' \times D_1}$, $\mathbf{m}^{(z_t)} \in \mathbb{R}^{D_1'}$, $D_1 = D_2 = D_3 = D$, and \mathcal{G} is diago-
 558 nal.

559 **Multi-subspace SLDS** Multi-subspace SLDS is a much harder optimization problem, which we
 560 found was often numerically unstable. We therefore do not consider multi-subspace SLDS in these
 561 experiments, but include its definition here for completeness

$$\mathbf{x}_t \sim \mathcal{N} \left(\mathbf{A}^{(z_t)} \mathbf{x}_{t-1} + \mathbf{b}^{(z_t)}, \mathbf{Q}^{(z_t)} \right), \quad (86)$$

$$\mathbf{y}_t \sim \mathcal{N} \left(\mathbf{C}^{(z_t)} \mathbf{x}_t + \mathbf{d}^{(z_t)}, \mathbf{R}^{(z_t)} \right). \quad (87)$$

562 **Single-subspace SLDS** Single-subspace SLDS was used in all of our experiments, and is typically
 563 used in practice [8, 15]

$$\mathbf{x}_t \sim \mathcal{N} \left(\mathbf{A}^{(z_t)} \mathbf{x}_{t-1} + \mathbf{b}^{(z_t)}, \mathbf{Q}^{(z_t)} \right), \quad (88)$$

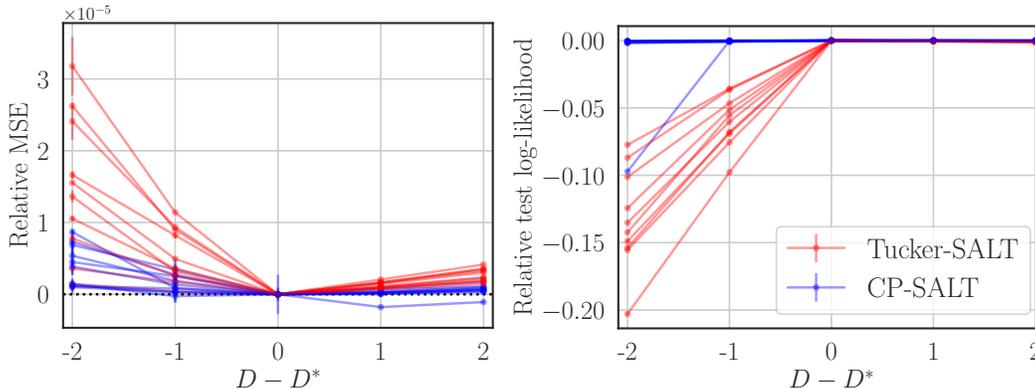
$$\mathbf{y}_t \sim \mathcal{N} (\mathbf{C} \mathbf{x}_t + \mathbf{d}, \mathbf{R}). \quad (89)$$

564 **D Synthetic Data Experiments**

565 **D.1 Extended Experiments for Proposition 1**

566 In Section 5.1 we showed that Proposition 1 can accurately predict the required rank for CP- and
 567 Tucker-SALT models. We showed results for a single LDS for clarity. We now extend this analysis
 568 across multiple random LDS and SALT models. We randomly sampled LDSs with latent dimensions
 569 ranging from 4 to 10, and observation dimensions ranging from 9 to 20. For each LDS, we fit
 570 5 randomly initialized CP-SALT and Tucker-SALT models with $L = 50$ lags. We varied the
 571 rank of our fit SALT models according to the rank predicted by Proposition 1. Specifically, we
 572 computed the estimated number of ranks for a given LDS, denoted D^* , and then fit SALT models
 573 with $\{D^* - 2, D^* - 1, D^*, D^* + 1, D^* + 2\}$ ranks. According to Proposition 1, we would expect
 574 to see the reconstruction error of the autoregressive tensor be minimized, and for prediction accuracy
 575 to saturate, at $D = D^*$.

576 To analyze these model fits, we first computed the average mean squared error of the autoregressive
 577 tensor corresponding to the LDS simulation, as a function of SALT rank relative to the rank required
 578 by Proposition 1. We see, as predicted by Proposition 1, that error in the autoregressive tensor is
 579 nearly always minimized at D^* (Figure 7A). Tucker-SALT was always minimized at D^* . Some
 580 CP-SALT fits have lower MSE at ranks other than predicted by Proposition 1. We believe this is
 581 due to local minima in the optimization. We next investigated the test log-likelihood as a function of
 582 the relative rank (Figure 7B). Interestingly, the test log-likelihood shows that Tucker-SALT strongly
 583 requires the correct number of ranks for accurate prediction, but CP-SALT can often perform well
 584 with fewer ranks than predicted (although still a comparable number of ranks to Tucker-SALT). As
 585 in Figure 2, these analyses empirically confirm Proposition 1.



(a) Normalized MSE of autoregressive tensor. (b) Normalized log-likelihood on held-out test set.

Figure 7: Extended results examining Proposition 1. Results are shown for the ability of SALT to estimate ten randomly generated LDSs, using five SALT repeats for each LDS. MSEs (in panel A) and log-likelihoods (in panel B) are normalized by the mean MSE and mean log-likelihood of SALT models trained with $D = D^*$. D is the rank of the fit SALT model, and D^* is the necessary rank predicted by Proposition 1.

586 **D.2 Quantitative Performance: Synthetic Switching LDS Experiments**

587 We include further results and analysis for the NASCAR[®] and Lorenz attractor experiments presented
 588 in Section 5.2. We compare the marginal likelihood achieved by single-subspace SALT models of
 589 different sizes. We see that SALT outperforms ARHMMs, and can fit larger models (more lags)
 590 without overfitting (Figure 8). Note that the SLDS does not admit exact inference, and so we cannot
 591 readily compute the exact marginal likelihood for the SLDS.

592 **D.3 TVART versus SALT in recovering the parameters of SLDSs**

593 We compared SALT to TVART [24], another tensor-based method for modeling autoregressive
 594 processes. We modified TVART (as briefly described in the original paper) so that it can handle
 595 AR(p) processes, as opposed to only AR(1) processes. TVART is also not a probabilistic model (i.e.,

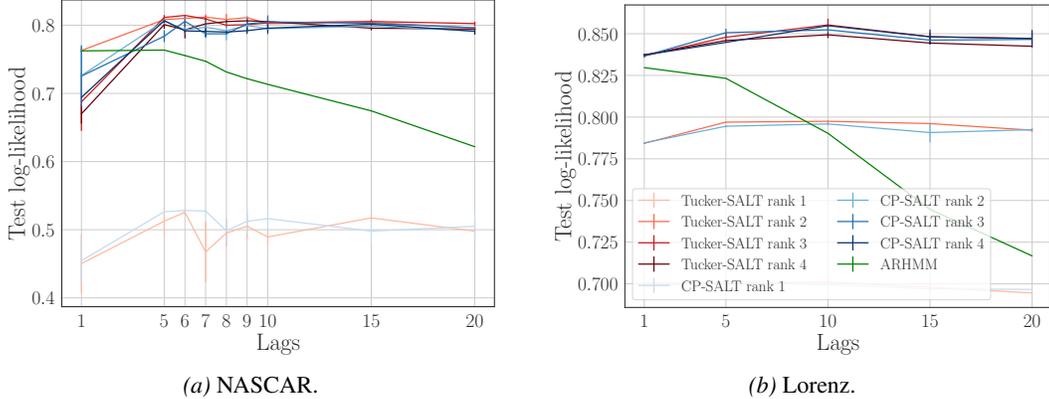


Figure 8: Quantitative performance of different SALT models and ARHMMs (averaged over 3 different runs) on the synthetic experiments presented in Section 5.2. The test-set log likelihood is shown as a function of lags in the SALT model, for both (A) the NASCAR[®] and (B) Lorenz synthetic datasets.

596 cannot compute log-likelihoods), and so we focus our comparison on how well these methods recover
 597 the parameters of a ground-truth SLDS.

598 We used the same SLDS that we used to generate the NASCAR[®] dataset in Section 5.2. We then
 599 used $L = 7$ CP-SALT and Tucker-SALT with ranks 3 and 2, respectively, and computed the MSE
 600 between the ground truth tensor and SALT tensors. For TVART, we used $L = 7$, bin size of 10,
 601 and ranks 2 and 3 to fit the model to the data. We then clustered the inferred dynamics parameters
 602 to assign discrete states. To get the TVART parameter estimation, we computed the mean of the
 603 dynamics parameters for each discrete state and computed the MSE against the ground truth tensor.
 604 The MSE results are as follows:

Table 2: Results comparing SALT and TVART [24] on the NASCAR example.

Model	Rank	Tensor Reconstruction MSE ($\times 10^{-3}$)	Number of parameters
TVART	2	0.423	1.4K
TVART	3	0.488	2.0K
Tucker-SALT	2	0.294	0.6K
CP-SALT	3	0.297	0.7K

605 Table 2 shows that SALT models recover the dynamics parameters of the ground truth SLDS more
 606 accurately. Furthermore, we see that SALT models use fewer parameters than TVART models for
 607 the dataset (as the number of parameters in TVART scales linearly with the number of windows). We
 608 also note that TVART cannot be applied to held-out data, and, without post-hoc analysis, does not
 609 readily have a notion of re-usable dynamics or syllables.

610 D.4 The effect of the number of switches on the recovery of the parameters of the 611 autoregressive dynamic tensors

612 We asked how the frequency of discrete state switches affected SALT’s ability to recover the au-
 613 toregressive tensors. We trained CP-SALT, Tucker-SALT, the ARHMM, all with $L = 5$ lags, and
 614 the SLDS on data sampled from an SLDS with varying number of discrete state switches. The
 615 ground-truth SLDS model had $H = 2$ discrete states, $N = 20$ observations and $D = 7$ dimensional
 616 continuous latent states. The matrix $\mathbf{A}^{(h)}(\mathbf{I} - \mathbf{K}^{(h)}\mathbf{C}^{(h)})$ of each discrete state of the ground-
 617 truth SLDS had 1 real eigenvalue and 3 pairs of complex eigenvalues. We sampled 5 batches of
 618 $T = 15,000$ timesteps of data from the ground-truth SLDS, with $s_n \in \{1, 10, 25, 75, 125\}$ num-
 619 ber of discrete state switches that were evenly spaced out across the data. We then computed the
 620 mean squared error (MSE) between the SLDS tensors and the tensors reconstructed by SALT, the
 621 ARHMM, and the SLDS. (Figure 9). More precisely, we combined the 3rd order autoregressive
 622 tensors from each discrete state into a 4th order tensor, and calculated the MSE based on these 4th

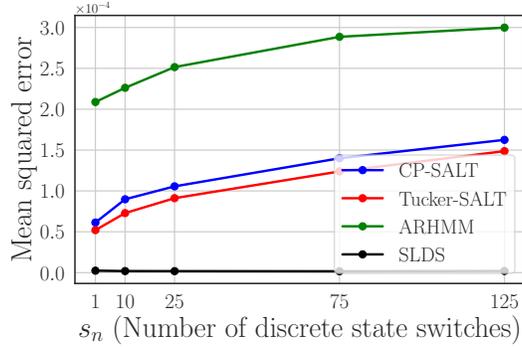


Figure 9: The quality of SALT approximation of SLDSs decreases as the number of discrete state switches increases: The data comes from an SLDS with $H = 2$, $N = 20$, and $D = 7$. 15,000 timesteps were generated, with varying numbers of evenly spaced discrete state switches (x-axis). The mean squared error of reconstructing the autoregressive tensors increased as a function of the number of discrete state switches. Note that we combined the 3rd order autoregressive tensors from each discrete state into a 4th order tensor, and calculated the MSE based on these 4th order tensors.

623 order tensors. As expected, the MSE increased with the number of switches in the data, indicating that
 624 the quality of SALT approximation of SLDSs decreases as the frequency of discrete state switches
 625 increases.

626 **E Modeling Mouse Behavior**

627 We include further details for the mouse experiments in Section 5.3.

628 **E.1 Training Details**

629 We used the first 35,949 timesteps of data from each of the three mice, which were collected at
630 30Hz resolution. We used $H = 50$ discrete states and fitted ARHMMs and CP-SALT models with
631 varying lags and ranks. Similar to Wiltchko et al. [2], we imposed stickiness on the discrete state
632 transition matrix via a Dirichlet prior with concentration of 1.1 on non-diagonals and 6×10^4 on
633 the diagonals. These prior hyperparameters were empirically chosen such that the durations of the
634 inferred discrete states and the given labels were comparable. We trained each model 5 times with
635 random initialization for each hyperparameter, using 100 iterations of EM on a single NVIDIA Tesla
636 P100 GPU.

637 **E.2 Video Generation**

638 Here we describe how the mouse behavioral videos were generated. We first determined the CP-
639 SALT hyperparameters as those which led to the highest log-likelihood on the validation dataset.
640 Then, using that CP-SALT model, we computed the most likely discrete states on the train and test
641 data. Given a discrete state h , we extracted slices of the data whose most likely discrete state was
642 h . We padded the data by 30 frames (i.e. 1 second) both at the beginning and the end of each slice
643 for the movie. A red dot appears on each mouse for the duration of discrete state h . We generated
644 such videos for all 50 discrete states (as long as there existed at least one slice for each discrete
645 state) on the train and test data. For a given discrete state, the mice in each video behaved very
646 similarly (e.g., the mice in the video for state 18 "pause" when the red dots appear, and those in
647 the video for state 32 "walk" forward), suggesting that CP-SALT is capable of segmenting the data
648 into useful behavioral syllables. See "MoSeq_salt_videos_train" and "MoSeq_salt_videos_test"
649 in the supplementary material for the videos generated from the train and test data, respectively.
650 "salt_crowd_ i .mp4" refers to the crowd video for state i . We show the principal components for states
651 1, 2, 13, 32, 33, 47 in Figure 10.

652 **E.3 Modeling Mouse Behavior: Additional Analyses**

653 We also investigated whether SALT qualitatively led to a good segmentation of the behavioral data
654 into discrete states, shown in Figure 10. Figure 10A shows a 30 second example snippet of the test
655 data from one mouse colored by the discrete states inferred by CP-SALT. CP-SALT used fewer
656 discrete states to model the data than the ARHMM (Figure 10B). Coupled with the finding that
657 CP-SALT improves test-set likelihoods, this suggests that the ARHMM may have oversegmented
658 the data and CP-SALT may be better able to capture the number of behavioral syllables. Figure 10C
659 shows average test data (with two standard deviations) for a short time window around the onset
660 of a discrete state (we also include mouse videos corresponding to that state in the supplementary
661 materials). The shrinking gray area around the time of state onset, along with the similar behaviors
662 of the mice in the video, suggests that CP-SALT is capable of segmenting the data into consistent
663 behavioral syllables.

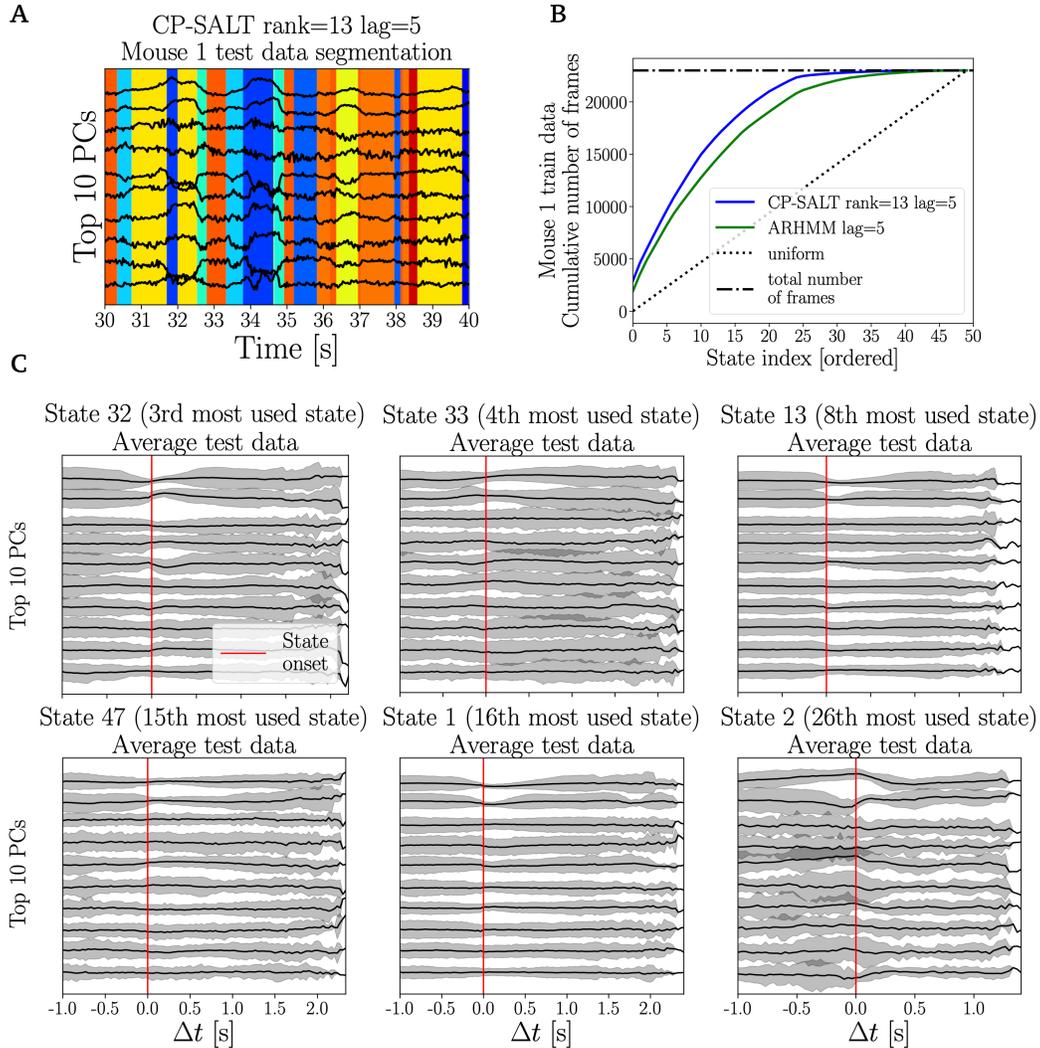


Figure 10: CP-SALT leads to qualitatively good segmentation of the mouse behavioral data into distinct syllables. (A) 30 seconds of test data (Mouse 1) with the discrete states inferred by CP-SALT as the background color. (B) For one mouse, the cumulative number of frames that are captured by each discrete state, where the discrete states are ordered according to how frequently they occur. (C) The average test data, with two standard deviations, for six states of CP-SALT, aligned to the time of state onset. The shrinkage of the gray region around the state onset tells us that CP-SALT segments the test data consistently.

664 **F Modeling *C. elegans* Neural Data**

665 We include further details and results for the *C. elegans* example presented in Section 5.4. This
 666 example highlights how SALT can be used to gain scientific insight in to the system.

667 **F.1 Training Details**

668 We used ~ 3200 timesteps of data (recorded at 3Hz) from one worm, for which 48 neurons were
 669 confidently identified. The data were manually segmented in to seven labels (reverse sustained, slow,
 670 forward, ventral turn, dorsal turn, reversal (type 1) and reversal (type 2). We therefore used $H = 7$
 671 discrete states in all models (apart from the GLM). After testing multiple lag values, we selected
 672 $L = 9$ for all models, as these longer lags allow us to examine longer-timescale interactions and
 673 produced better segmentations across models, with only a small reduction in variance explained.
 674 We trained each model 5 times with KMeans initialization, using 100 iterations of EM on a single
 675 NVIDIA Tesla V100 GPU. Models that achieved 90% explained variance on a held-out test set were
 676 then selected and analyzed (similar to Linderman et al. [9]).

677 **F.2 Additional Quantitative Results**

678 Figure 11 shows additional results for training different models. In Figure 11A we see that models
 679 with larger ranks (or latent dimension) achieve higher explained variance. Interestingly, longer lags
 680 can lead to a slight reduction in the explained variance, likely due to overfitting. This effect is less
 681 pronounced in the more constrained single-subspace SALT, but, these models achieve lower explained
 682 variance ratios throughout. Longer lag models allow us to inspect longer-timescale dependencies,
 683 and so are more experimentally insightful. Figure 11B shows the confusion matrix for discrete states
 684 between learned models and the given labels. The segmentations were similar across all models that
 685 achieved 90% explained variance.

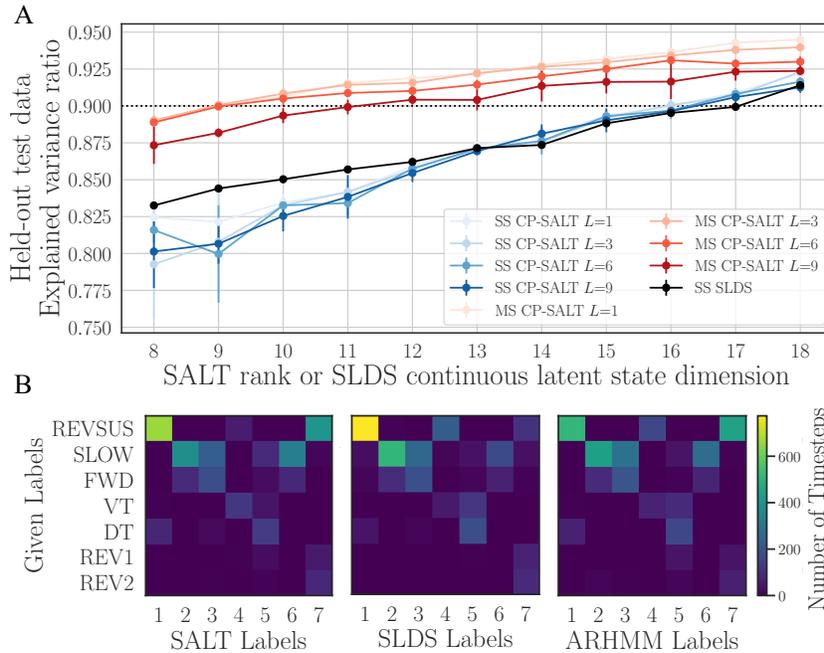


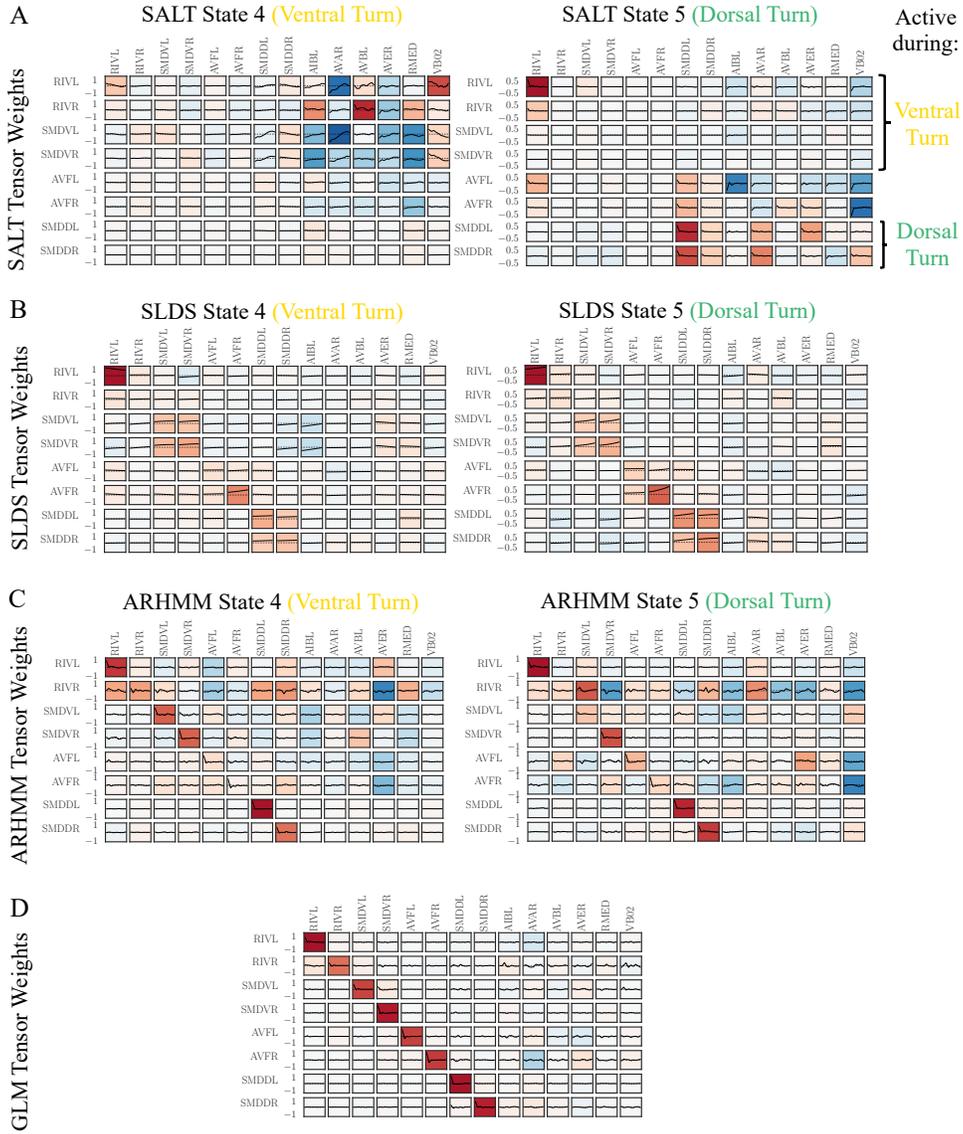
Figure 11: **SALT and SLDS perform comparably on held-out data:** (A): Explained variance on a held-out sequence. Single-subspace (SS) SALT and SLDS perform comparably. Multi-subspace (MS) SALT achieves a higher explained variance with fewer ranks. Multi-subspace SLDS was numerically unstable. (B): Confusion matrices between given labels and predicted labels. All methods produce similar quality segmentations.

686 **F.3 Additional Autoregressive Filters**

687 Figures 12 and 13 show extended versions of the autoregressive filters included in Section 5.4. Figure
 688 12 shows the filters learned for ventral and dorsal turns (for which panel A was included in Figure 5),

689 while Figure 13 shows the filters for forward and backward locomotion. Note that the GLM does not
690 have multiple discrete states, and hence the same filters are used across states. We see for ARHMM
691 and SALT that known-behavior tuned neurons have higher magnitude filters (determined by area
692 under curve), whereas the SLDS and GLM do not recover such strong state-specific tuning. Since
693 the learned SLDS did not have stable within-state dynamics, the autoregressive filters could not be
694 computed using Equation (47). We thus show $CA^{(h)l}C^+$ for lag l as a proxy for the autoregressive
695 filters of discrete state h of the SLDS. Note that this is a post-hoc method and does not capture the
696 true dependencies in the observation space.

697 We see that SALT consistently assigns high autoregressive weight to neurons known to be involved
698 in certain behaviors (see Figures 12 and 13). In contrast, the ARHMM identifies these relationships
699 less reliably, and the estimate of the SLDS autoregressive filters identifies few strong relationships.
700 As the GLM only have one “state”, the autoregressive filters are averaged across state, and so few
701 strong relationships are found. This highlights how the low-rank and switching properties of SALT
702 can be leveraged to glean insight into the system.



Lag=9, equivalent to 3 seconds; from left to right

Figure 12: Autoregressive tensors learned by different models (Ventral and Dorsal Turns): (A-C) One-dimensional autoregressive filters learned in two states by SALT, SLDS, ARHMM (identified as ventral and dorsal turns), and **(D)** by a GLM. RIV and SMDV are known to mediate ventral turns, while SMDD mediate dorsal turns [31, 32, 34].

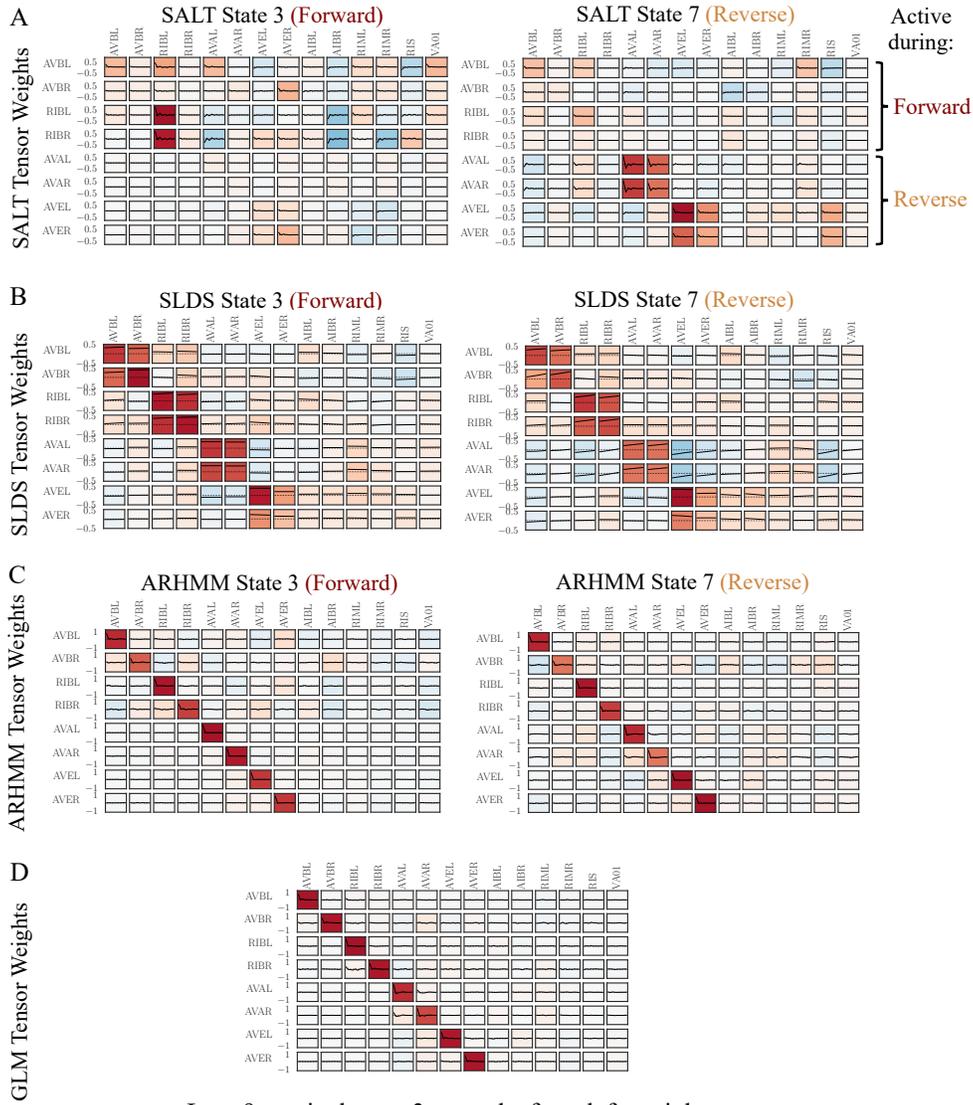


Figure 13: Autoregressive tensors learned by different models (Forward Locomotion and Reversal): (A-C) One-dimensional autoregressive filters learned in two states by SALT, SLDS, ARHMM (identified as forward and reverse), and (D) by a GLM. AVB and RIB are known to mediate forward locomotion, while AVA and AVE are involved in initiating reversals [31, 32, 37, 38].

703 **G Code Availability**

704 As part of Supplementary materials, we include the source code for SALT (in Python) and an example
705 Jupyter Notebook. This notebook shows how to sampling data from SALT, and then re-fit a SALT
706 model to the data with an EM algorithm. We include just the core code here for ease of inspection,
707 and will release the full source code after the review process.