

A APPENDIX

In this supplementary material, we provide the details of the example derivative computations from Section 2.2, give a justification of the NCPD expansion formula exploited in Section 2.1, and provide further experimental results that we were not able to include in Section 3.

EXAMPLE DERIVATIONS

Here, we justify the derivations provided in the example in Section 2.2. We note that Anaisi et al. (2020); Kolda & Hong (2019) provide similar derivations for the CP tensor decomposition, but their decompositions do not attempt to further decompose the CP factor matrices, and thus, their results are not sufficient for providing derivatives with respect to the \mathbf{A} and \mathbf{S} matrices. Consider the full reconstruction loss function for the order- k tensor \mathbf{X} ,

$$C = \|\mathbf{X} - \llbracket \mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_k \rrbracket\|_F^2,$$

where for some fixed $1 \leq i \leq k$, $\mathbf{X}_i = \mathbf{ABC}$ and consider the gradient $\frac{\partial C}{\partial \mathbf{B}}$. Let

$$\mathbf{H}_i = \mathbf{X}_k \circ \dots \circ \mathbf{X}_{i+1} \circ \mathbf{X}_{i-1} \circ \dots \circ \mathbf{X}_1.$$

Now, we let $\widetilde{\mathbf{X}} = \llbracket \mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_k \rrbracket$. Then, if $\widetilde{\mathbf{X}}_{(i)}$ is the mode- i matricization of $\widetilde{\mathbf{X}}$ (see e.g., Kolda & Bader, 2009), we have

$$\widetilde{\mathbf{X}}_{(i)} = \mathbf{X}_i (\mathbf{X}_k \circ \dots \circ \mathbf{X}_{i+1} \circ \mathbf{X}_{i-1} \circ \dots \circ \mathbf{X}_1)^\top = \mathbf{X}_i \mathbf{H}_i^\top.$$

Thus, if we let $\mathbf{X}_{(i)}$ be the mode- i matricization of \mathbf{X} , we have that

$$C = \|\mathbf{X}_{(i)} - \widetilde{\mathbf{X}}_{(i)}\|_F^2 = \|\mathbf{X}_{(i)} - (\mathbf{ABC})(\mathbf{X}_k \circ \dots \circ \mathbf{X}_{i+1} \circ \mathbf{X}_{i-1} \circ \dots \circ \mathbf{X}_1)^\top\|_F^2.$$

Now, we compute the desired gradient through a series of applications of the chain rule. We then see that

$$\begin{aligned} \frac{\partial C}{\partial \mathbf{B}} &= \frac{\partial C}{\partial \mathbf{X}_i \mathbf{H}_i^\top} \frac{\partial \mathbf{X}_i \mathbf{H}_i^\top}{\partial \mathbf{ABC}} \frac{\partial \mathbf{ABC}}{\partial \mathbf{B}} \\ &= \mathbf{A}^\top \left(\frac{\partial C}{\partial \mathbf{X}_i \mathbf{H}_i^\top} \frac{\partial \mathbf{X}_i \mathbf{H}_i^\top}{\partial \mathbf{X}_i} \right) \mathbf{C}^\top \\ &= 2\mathbf{A}^\top (\mathbf{X}_{(i)} - \mathbf{X}_i \mathbf{H}_i^\top) \mathbf{H}_i \mathbf{C}^\top. \end{aligned}$$

Now, using the calculations above we can proceed in calculating $\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_j)}}$. Gao et al. (2019) show that

if $\left(\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_j)}}\right)^{\mathbf{S}}$ denotes derivative of C with respect to $\mathbf{A}_i^{(\ell_j)}$, holding the \mathbf{S} matrices constant, then we have

$$\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_j)}} = \left(\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_1)}}\right)^{\mathbf{S}} + \sum_{\substack{\ell_1 \leq \ell_2 \leq \mathcal{L}-2 \\ 1 \leq j \leq r}} \mathbf{U}_i^{(\ell_1, \ell_2), j}$$

where $\mathbf{U}_i^{(\ell_1, \ell_2), j}$ relates C to $\mathbf{A}_i^{(\ell_1)}$ through $\mathbf{S}_i^{(\ell_2)}$ and $\mathbf{S}_i^{(\ell_1)}$, is defined column-wise (j), and depends upon $\left(\frac{\partial C}{\partial \mathbf{S}_i^{(\ell_2)}}\right)^*$, the derivative of C with respect to $\mathbf{S}_i^{(\ell_2)}$ holding $\mathbf{S}_i^{(\ell_2+1)}, \dots, \mathbf{S}_i^{(\mathcal{L}-2)}$ constant. Thus, $\left(\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_j)}}\right)^{\mathbf{S}}$ and $\left(\frac{\partial C}{\partial \mathbf{S}_i^{(\ell_j)}}\right)^*$ are sufficient to calculate $\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_j)}}$. We calculate the gradient $\left(\frac{\partial C}{\partial \mathbf{A}_i^{(\ell_j)}}\right)^{\mathbf{S}}$ where

$$C = \|\mathbf{X} - \llbracket \widetilde{\mathbf{X}}_1, \widetilde{\mathbf{X}}_2, \dots, \widetilde{\mathbf{X}}_k \rrbracket\|_F^2$$

and $\widetilde{\mathbf{X}}_i = \mathbf{A}_i^{(0)} \mathbf{A}_i^{(1)} \dots \mathbf{A}_i^{(\mathcal{L}-2)} \mathbf{S}_i^{(\mathcal{L}-2)}$. Since we can assume that $\mathbf{A}_i^{(\ell_1)}$ is independent of all other \mathbf{A} 's and \mathbf{S} 's, we have that

$$\left(\frac{\partial \mathcal{C}}{\partial \mathbf{A}_i^{(\ell_j)}} \right)^{\mathbf{S}} = 2 \left(\mathbf{A}_i^{(0)} \mathbf{A}_i^{(1)} \dots \mathbf{A}_i^{(\ell_j-1)} \right)^{\top} \left(\mathbf{X}_{(i)} - \mathbf{X}_i \mathbf{H}_i^{\top} \right) \mathbf{H}_i \left(\mathbf{A}_i^{(\ell_{j+1})} \dots \mathbf{A}_i^{(\mathcal{L}-2)} \mathbf{S}_i^{(\mathcal{L}-2)} \right)^{\top}. \quad (17)$$

Now, we calculate $\left(\frac{\partial \mathcal{C}}{\partial \mathbf{S}_i^{(\ell_j)}} \right)^*$. Since we can assume that $\mathbf{S}_i^{(\ell_j)}$ is independent of all other \mathbf{A} 's and \mathbf{S} 's, we have that

$$\left(\frac{\partial \mathcal{C}}{\partial \mathbf{S}_i^{(\ell_j)}} \right)^* = 2 \left(\mathbf{A}_i^{(0)} \mathbf{A}_i^{(1)} \dots \mathbf{A}_i^{(\ell_j)} \right)^{\top} \left(\mathbf{X}_{(i)} - \mathbf{X}_i \mathbf{H}_i^{\top} \right) \mathbf{H}_i. \quad (18)$$

Thus, we have the required derivatives to evaluate $\frac{\partial \mathcal{C}}{\partial \mathbf{A}_i^{(\ell_j)}}$.

HNCPD EXPANSION

We now provide brief justification of the expansion of the NCPD in terms of later factorizations used in Section 2.1 that is,

$$\llbracket \widetilde{\mathbf{X}}_1, \widetilde{\mathbf{X}}_2, \dots, \widetilde{\mathbf{X}}_k \rrbracket = \sum_{1 \leq j_1, j_2, \dots, j_k \leq r^{(0)}} \alpha_{j_1, j_2, \dots, j_k} \left((\mathbf{A}_1^{(0)})_{:,j_1} \otimes (\mathbf{A}_2^{(0)})_{:,j_2} \otimes \dots \otimes (\mathbf{A}_k^{(0)})_{:,j_k} \right)$$

where $\alpha_{j_1, j_2, \dots, j_k} = \sum_{p=1}^r (\mathbf{S}_1^{(0)})_{j_1, p} (\mathbf{S}_2^{(0)})_{j_2, p} \dots (\mathbf{S}_k^{(0)})_{j_k, p}$.

We have that by definition,

$$\llbracket \widetilde{\mathbf{X}}_1, \widetilde{\mathbf{X}}_2, \dots, \widetilde{\mathbf{X}}_k \rrbracket = \sum_{p=1}^r \left((\widetilde{\mathbf{X}}_1)_{:,p} \otimes (\widetilde{\mathbf{X}}_2)_{:,p} \otimes \dots \otimes (\widetilde{\mathbf{X}}_k)_{:,p} \right).$$

We also have that $\widetilde{\mathbf{X}}_i = \mathbf{A}_i^{(0)} \mathbf{S}_i^{(0)}$ for $1 \leq i \leq k$, so we have that for each column p , $1 \leq p \leq r$ of $\widetilde{\mathbf{X}}_i$,

$$(\widetilde{\mathbf{X}}_i)_{:,p} = \sum_{j=1}^{r^{(0)}} (\mathbf{S}_i^{(0)})_{j,p} (\mathbf{A}_i^{(0)})_{:,j}.$$

Thus, by the linearity of the outer product we have that

$$\left((\widetilde{\mathbf{X}}_1)_{:,p} \otimes (\widetilde{\mathbf{X}}_2)_{:,p} \otimes \dots \otimes (\widetilde{\mathbf{X}}_k)_{:,p} \right) = \sum_{1 \leq j_1, j_2, \dots, j_k \leq r^{(0)}} \alpha_{p, j_1, j_2, \dots, j_k} \left((\mathbf{A}_1^{(0)})_{:,j_1} \otimes (\mathbf{A}_2^{(0)})_{:,j_2} \otimes \dots \otimes (\mathbf{A}_k^{(0)})_{:,j_k} \right)$$

where $\alpha_{p, j_1, j_2, \dots, j_k} = (\mathbf{S}_1^{(0)})_{j_1, p} (\mathbf{S}_2^{(0)})_{j_2, p} \dots (\mathbf{S}_k^{(0)})_{j_k, p}$. Now, by noting that

$$\alpha_{j_1, j_2, \dots, j_k} = \sum_{p=1}^r \alpha_{p, j_1, j_2, \dots, j_k}$$

we arrive at the original statement.

SYNTHETIC EXPERIMENT

In this section, we provide the additional views of the synthetic tensor and computed approximations from Section 3.1. In Figure 2 in the main text, for visualization we displayed the projection of each tensor onto the third mode. In Figure 9, we display the projections of these tensors onto all three modes. We see that due to the simple block structure used to produce the synthetic data tensor, the three modes all tell a similar story; that is, Neural NCPD is able to recover meaningful structure along all three modes.

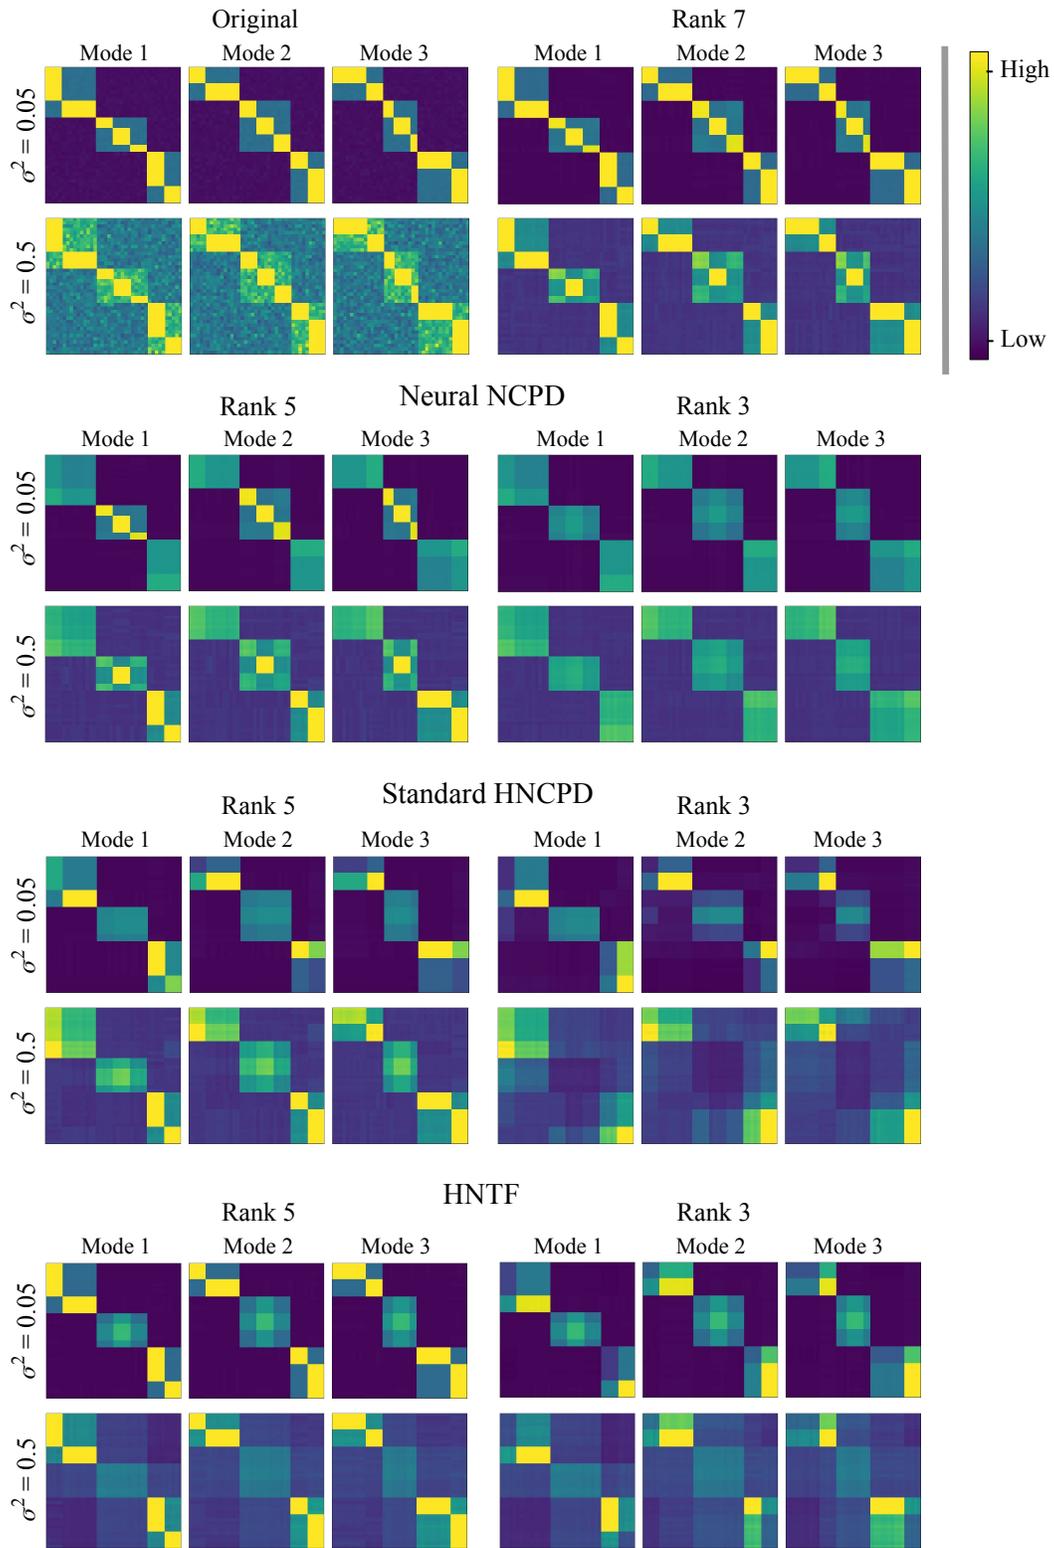


Figure 9: Here we display the projections onto all three modes for the original data tensor \mathbf{X} and approximations of \mathbf{X} at ranks $r = 7$, $r^{(0)} = 5$, and $r^{(1)} = 3$ produced by Neural NCPD, Standard HNCPD, and HNTF at two levels of noise.

TEMPORAL DOCUMENT ANALYSIS EXPERIMENT

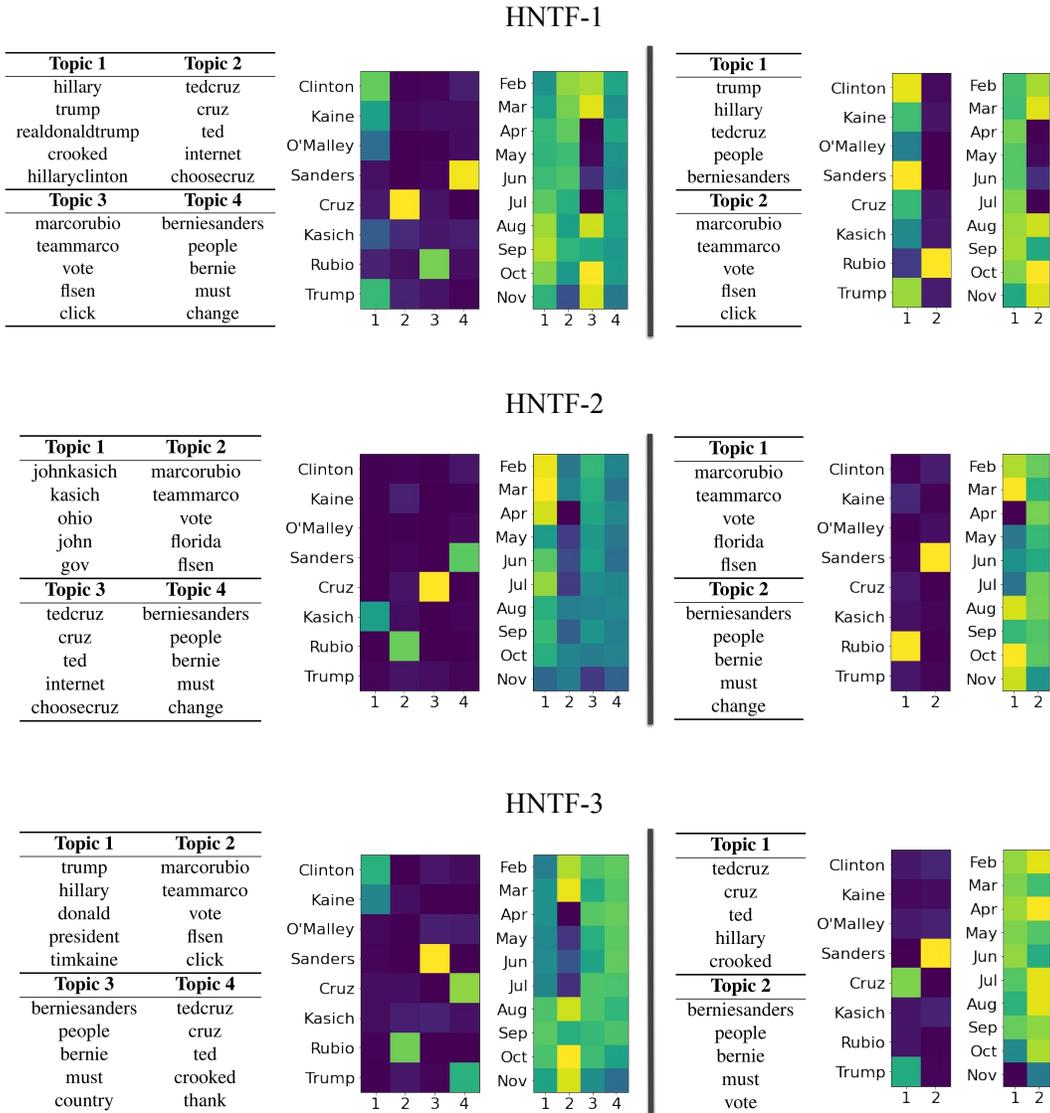


Figure 10: Here we display a three-layer HNTF on the Twitter dataset from Section 3.2 at ranks $r = 8$, $r^{(0)} = 4$, and $r^{(1)} = 2$, run separately for each of the possible ordering on the data tensor. We display the top keywords and heatmaps of topics in the candidate and temporal modes at ranks 4 (left) and 2 (right). We note that the rank 8 factorization is identical to that of Neural NCPD, so we do not re-display it here (see Section 3.2).

In Figure 10, we display the results from running HNTF on the Twitter dataset in Section 3.2 excluding the topics at rank 8 because they are identical to those learned by Neural NCPD (see Section 3.2). We see that while the factorization for the first possible ordering is similar to that of Neural NCPD and contains significant meaningful topic modeling information, the other two orderings lose significant information by the last layer and, and have topic presence and from only 2 or 3 of the eight candidates.

VIDEO DATA EXPERIMENT

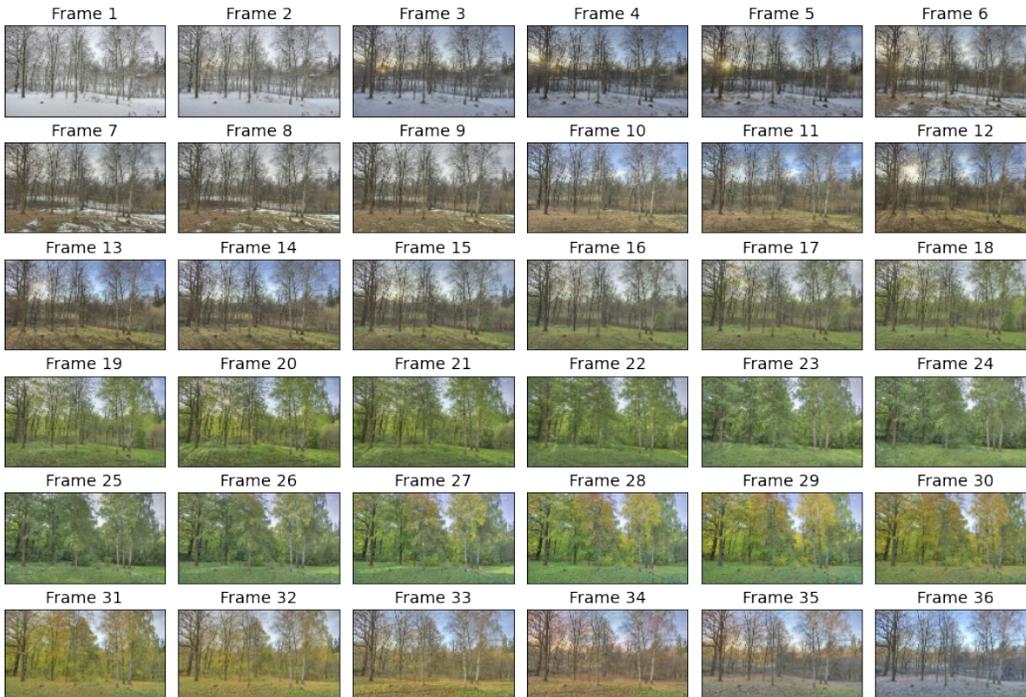


Figure 11: Here we display the first 36 of 37 frames of the time lapse video dataset from Section 3.3 (The 37th frames is included in Figure 6)

In Figure 11, we display the first 36 of 37 frames of the time lapse video dataset from Section 3.3 (the 37th frame is included in Figure 6) in order to make it clear how seasons progress throughout the frames. We see that the video begins in the white winter months, transitions to spring at around frame 16, and stays green until it transitions to fall around frame 28.

In Figure 12, we display the $\mathcal{S}_3^{(0)}$ matrix (top) and $\mathcal{S}_3^{(1)}$ matrix (bottom) produced by Neural NCPD on the time-lapse video tensor described in Section 3.3. By examining the \mathcal{S} matrices from our Neural NCPD algorithm, we are also able to see the hierarchical relationship between the topics from different ranks. In the $\mathcal{S}_3^{(0)}$ matrix, we see the hierarchical relationship between the rank 6 and rank 8 topics. In the $\mathcal{S}_3^{(1)}$ matrix, we see the hierarchical relationship between the rank 3 and rank 8 topics. We note that the $\mathcal{S}_3^{(0)}$ matrix (top) illustrates that topic one of rank 6 NCPD is closely related to topic eight of rank 8 NCPD, and $\mathcal{S}_3^{(1)}$ (bottom) similarly illustrates that topic two of rank 3 NCPD is closely related to topic eight in rank 8 NCPD; these relationships are unsurprising because, as seen in Figure 7 in the main text, these topics are present temporally during winter and fall and spatially in the sky behind the trees.

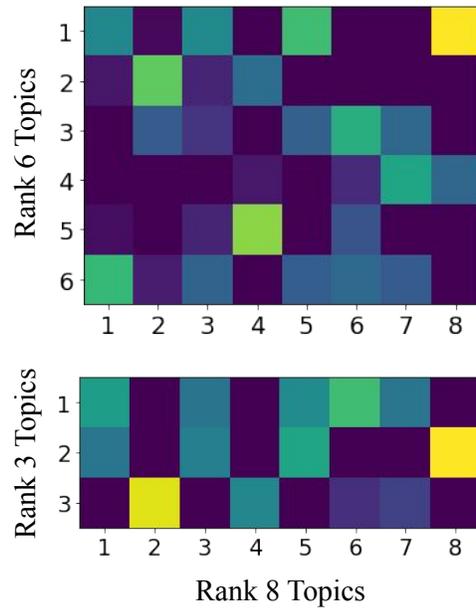


Figure 12: The $S_3^{(0)}$ matrix (top) and $S_3^{(1)}$ matrix (bottom) produced by Neural NCPD on the time-lapse video tensor described in Section [3.3](#)