
Time-Series Analysis via Low-Rank Matrix Factorization Applied to Infant-Sleep Data

Sheng Liu¹ Mark Cheng² Hayley Brooks^{3,4} Wayne Mackey⁴
David J. Heeger^{4,5} Esteban G. Tabak⁶ Carlos Fernandez-Granda^{1,6}

¹Center for Data Science, NYU ²Technische Universität Berlin

³Department of Psychology, University of Denver ⁴Center for Neural Science, NYU

⁵Department of Psychology, NYU

⁶Courant Institute of Mathematical Sciences, NYU

Abstract

We propose a nonparametric model for time series with missing data based on low-rank matrix factorization. The model expresses each instance in a set of time series as a linear combination of a small number of shared basis functions. Constraining the functions and the corresponding coefficients to be nonnegative yields an interpretable low-dimensional representation of the data. A time-smoothing regularization term ensures that the model captures meaningful trends in the data, instead of overfitting short-term fluctuations. The low-dimensional representation makes it possible to detect outliers and cluster the time series according to the interpretable features extracted by the model, and also to perform forecasting via kernel regression. We apply our methodology to a large real-world dataset of infant-sleep data gathered by caregivers with a mobile-phone app. Our analysis automatically extracts daily-sleep patterns consistent with the existing literature. This allows us to compute sleep-development trends for the cohort, which characterize the emergence of circadian sleep and different napping habits.

1 Introduction

Analysis of time series is a fundamental challenge in modern healthcare applications. The data are often very noisy, incomplete, and high dimensional, which precludes the direct application of standard machine-learning tools. We propose a method to extract meaningful features from such data, which are associated to regular intervals, such as hours, days, weeks, etc. This yields an interpretable representation of the data in terms of a superposition of daily, hourly, or weekly patterns that evolve over time. The representation can be used to perform several tasks, including trend extraction, outlier detection, clustering, and forecasting. We apply the proposed methodology to study infant sleep data gathered by caregivers via a mobile app. Understanding infant sleeping behavior is a problem of great importance to researchers, clinicians and, of course, parents. Traditional studies tend to be limited either to small number of infants from relatively homogeneous backgrounds observed intermittently (e.g. for a few hours, one day per month) or to data obtained from questionnaires, instead of direct observations [1]. Large-scale datasets logged directly by users present an unprecedented opportunity for data-driven analysis of sleep development across large populations.

Technical significance. Our main methodological contribution is a novel approach to analyze a set of multiple time series with missing entries. The goal of the analysis is to extract meaningful patterns associated to regular intervals that partition the time series, such as days, hours, months, etc. To achieve this, we combine ideas from low-rank matrix completion [2], singular spectrum analysis [3], and nonnegative matrix factorization [4]. We represent the data in terms of a small number of basis functions that capture common patterns associated to the intervals of interest. The coefficients of the representation encode the time evolution of these patterns in each of the time series. Constraining the functions and coefficients to be nonnegative avoids cancellations in the model and yields more interpretable basis functions. A smoothing regularization term on the coefficients ensures that they

capture meaningful trends in the data, instead of overfitting short-term fluctuations. The proposed model is validated on a real-world dataset, where we verify that the learned representation can be used to detect outliers and cluster the time series.

Clinical relevance. Recent works suggest that sleep affects memory and language development [5, 6], and may be used to predict attention and behavior problems [7]. The analysis of large-scale infant-sleep data gathered with hand-held devices therefore opens up the door to many promising clinical applications. In this work, we showcase several of them on a real-world dataset. Even though we do not make any prior assumptions about sleeping behavior (beyond nonnegativity and temporal smoothness), our analysis automatically extracts daily-sleep patterns—such as night sleep combined with one or two daily naps—that are consistent with the existing literature.

Related work. There is no previous analysis of infant-sleep data applying low-rank modeling (to the best of our knowledge). Ref. [8] studies sleep development from mobile-phone data using more traditional techniques, such as averaging in predetermined time intervals. Ref. [9] applies a growth mixture model to analyze patterns of developmental change in infants’ sleep. Other relevant studies in infant sleep development include [1, 10–12]. The ingredients in our methodological approach are inspired by several existing techniques, but to the best of our knowledge their combination is novel. Low-rank matrix completion [2] and nonnegative matrix factorization [4] have been widely used to perform imputation, denoising and factor analysis of high-dimensional data in collaborative filtering, topic modeling, and many other domains. Diverse forms of temporal regularization have been proposed in the application of these models to time series in [13–17]. However, these works consider a single multivariate time series, as opposed to a set of multiple one-dimensional time series. The idea of reshaping a single time series as a matrix dates back to at least the work of [3] on singular spectrum analysis. The reshaping procedure differs from ours because the rows overlap (the matrix is Toeplitz), so the learned patterns are not associated to fixed time intervals as in our case. Our approach is closer to very recent work [18], where the authors apply low-rank models to perform denoising, imputation and forecasting on a single time series (as opposed to multiple time series, which is our setting of interest). Code and data are available at <https://cims.nyu.edu/~s15924/infant.html>.

2 Methods

2.1 Low-rank modeling of time-series data

We consider a dataset of N discrete time series $y^{[n]}$, sampled at the same rate. The time series may have missing entries. We partition each time series into intervals of a predefined length ℓ . Our main assumption is that the intervals can be approximated as superpositions of a small number of shared basis functions. In the case of the infant-sleep data the intervals correspond to days, because the aim is to extract daily sleeping patterns. We express each partitioned time series as a matrix $Y^{[n]}$ with ℓ columns, where the rows contain non-overlapping segments of length ℓ ,

$$Y^{[n]} := \begin{bmatrix} y^{[n]}[1] & y^{[n]}[2] & \dots & y^{[n]}[\ell] \\ y^{[n]}[\ell + 1] & y^{[n]}[\ell + 2] & \dots & y^{[n]}[2\ell] \\ y^{[n]}[2\ell + 1] & y^{[n]}[2\ell + 2] & \dots & y^{[n]}[3\ell] \\ \dots & \dots & \dots & \dots \end{bmatrix}. \quad (1)$$

For simplicity, we assume that each row is present or missing, but our methodology is straightforward to extend to the case where single entries are missing. Our goal is to learn a joint low-rank representation for these matrices, such that each $Y^{[n]}$ is well approximated by a rank r decomposition of the form

$$Y^{[n]}(t, i) \approx \sum_{j=1}^r C_j^{[n]}(t) F_j(i), \quad 1 \leq i \leq \ell, \quad (2)$$

for all observed values of t . F_1, \dots, F_r are basis functions shared by all the time series. The model yields a representation of each $Y^{[n]}$ as $r \ll \ell$ coefficient time series $C_1^{[n]}, \dots, C_r^{[n]}$. The dimensionality of the data is consequently reduced by a factor of ℓ/r . Figure 1 shows a diagram of the low-rank model. In the application to infant-sleep data, the basis functions F_1, \dots, F_r represent *daily-sleep patterns* common to all the infants in the population. The coefficients $C_1^{[n]}, \dots, C_r^{[n]}$ used to combine these basis vectors are unique to each infant. They describe the sleep-development dynamics of the infant in terms of the daily-sleep patterns.

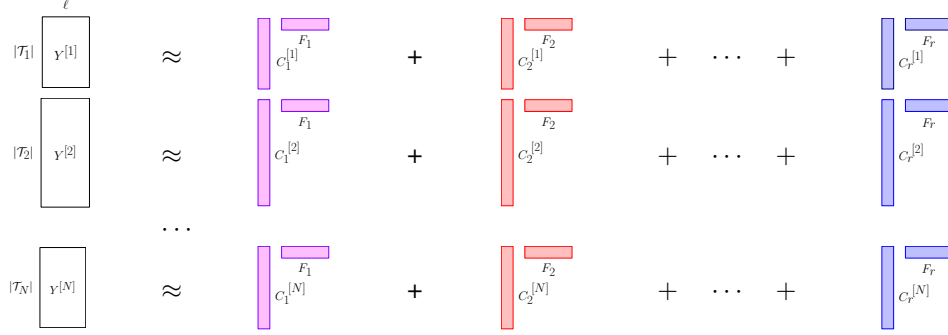


Figure 1: Proposed low-rank model applied on a set of time series. Each time series, represented by a matrix $Y^{[n]}$, $1 \leq n \leq N$, is approximated by a sum of r components corresponding to basis functions F_1, \dots, F_r , shared across the different time series. The coefficients $C_1^{[n]}, \dots, C_r^{[n]}$ provide a representation of reduced dimensionality for each $Y^{[n]}$. \mathcal{T}_n is the set of observed rows in $Y^{[n]}$.

2.2 Nonnegative matrix factorization with time smoothing

In order to extract meaningful patterns, we incorporate additional structure on the low-rank model defined by Eq. (2). First, we constrain the basis functions and the coefficients to be nonnegative. In applications such as image processing and topic modeling, nonnegative constraints often yield low-rank models that are more interpretable [4]. Second, we favor coefficients that vary smoothly in time. The aim is to capture gradual tendencies in the data, instead of overfitting rapid, local fluctuations. In order to fit a low-rank model with these characteristics we solve the following optimization problem:

$$\begin{aligned} \text{minimize} \quad & \sum_{n=1}^N \sum_{t \in \mathcal{T}_n} \sum_{i=1}^{\ell} \left(Y^{[n]}(t, i) - \sum_{j=1}^r C_j^{[n]}(t) F_j(i) \right)^2 + \lambda \sum_{n=1}^N \sum_{j=1}^r \left\| \mathcal{D}C_j^{[n]} \right\|_2^2 \quad (3) \\ \text{such that} \quad & C_j^{[n]}(t) \geq 0, F_j(i) \geq 0, \|F_j\|_2 = 1, 1 \leq j \leq r, 1 \leq n \leq N, t \in \mathcal{T}_n, 1 \leq i \leq \ell, \end{aligned}$$

where $\lambda > 0$ is a parameter that controls the tradeoff between the fit error and the regularization term, and the rest of the notation is consistent with Eq. (2). \mathcal{D} is a finite-difference operator that computes a discrete approximation to the derivative of the coefficients (we use a second-order operator). Crucially, the presence of the time-smoothing operator \mathcal{D} links the data of each individual time series. Otherwise, the low-rank model would be equivalent to performing principal-component analysis (PCA [19]) or nonnegative matrix factorization (NMF [4]) over the whole time-series dataset, interpreting each vector $Y^{[n]}(\cdot, t)$ as a separate data point. We call our approach nonnegative matrix factorization with time smoothing (NMF-TS). For the infant-sleep data, NMF-TS yields smooth coefficients as shown in Figure 7. In contrast, applying NMF or PCA results in very irregular coefficients, which overfit short-term variations in the data. The resulting low-rank approximations learned by PCA, NMF and NMF-TS for a specific infant are shown in Figure 5. Additional examples of NMF-TS coefficients are shown in Figure 8.

3 Results

In this section we report the results of applying the proposed low-rank model to analyze the infant-sleep dataset described in Section A of the appendix.

The basis functions obtained by applying the proposed low-rank decomposition to the infant sleep data are directly interpretable in terms of sleeping behavior: (1) daytime sleep, (2) nighttime sleep with two naps, (3) nighttime sleep with one nap, (4) nighttime sleep starting and ending early, (5) nighttime sleep starting and ending late. The patterns are shown in the left column of Figure 2 (and also in the rightmost column of Figure 6). The coefficients in the low-rank approximation of each infant indicate the degree to which each of these patterns is present at each day of age. They capture the dynamics of sleep development in terms of the learned daily-sleep patterns. We can therefore visualize and quantify these dynamics using the order statistics of the coefficients $C_1^{[n]}, \dots, C_5^{[n]}$ over

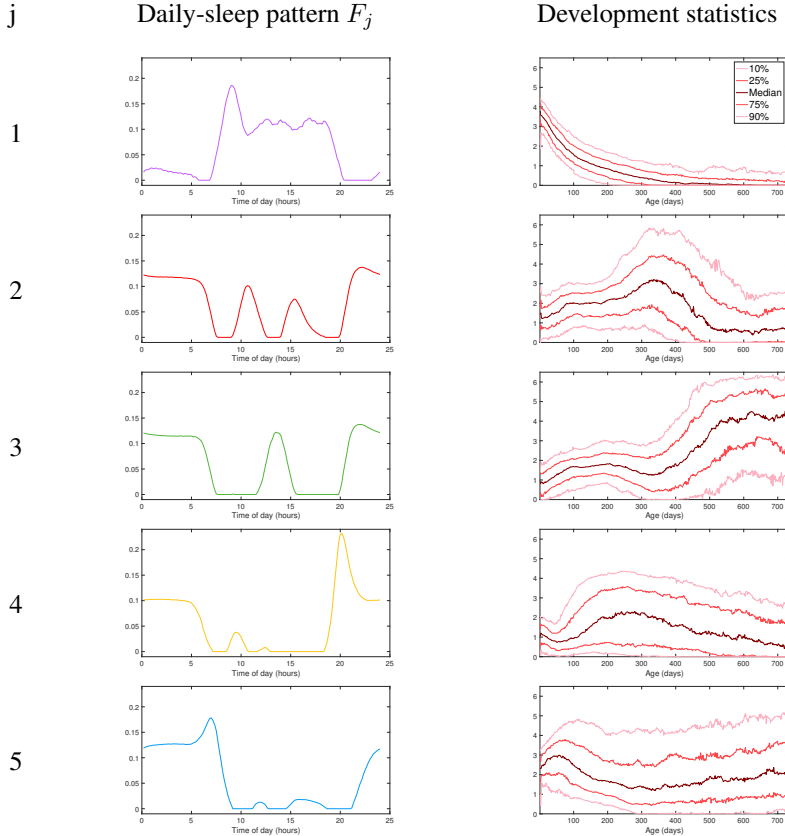


Figure 2: The left column shows the daily-sleep patterns learned by the proposed method. The right column shows the order statistics of the corresponding coefficients in the low-rank model.

the whole dataset. Figure 2 shows the median, the first and third quartiles and the 10th and 90th percentiles at each day of age.

Pattern 1 (daytime sleep) is prominent mainly during the first days of age. It decreases rapidly over the first year, signaling the emergence of circadian sleep. Pattern 2 (nighttime sleep with two naps) is particularly significant between 200 and 500 days of age. Pattern 3 (nighttime sleep with a single nap) emerges in the second year. The evolution of these three factors suggest that there are three stages in early-sleep development: irregular sleep, nighttime sleep with two naps, and nighttime sleep with a single nap. This is consistent with the literature on infant sleep development (see [10–12]). Patterns 4 and 5 account for shifts in the nighttime-sleep schedules of the different infants in the population. Their order statistics remain stable once nighttime sleep consolidates.

In Section C of the appendix, we show the results of applying our methodology to detect anomalous individuals, to cluster the cohort into groups with different sleeping tendencies, and to obtain improved predictions of future sleep behavior.

4 Discussion and Future Directions

In this work we propose a nonparametric model for the analysis of multiple time series based on low-rank matrix estimation. The method is shown to extract meaningful patterns from infant-sleep data, which can be used to characterize general trends in the population, detect anomalous sleeping behavior, cluster the infants according to their sleeping habits, and perform forecasting. A limitation of our approach is that it does not account for shifts in the schedule within the individual time series. Modifying low-rank models to enforce invariance of the basis functions to such shifts is an interesting direction of future research. Other directions include using additional features (such as feeding habits) to improve forecasting, incorporating nonlinearities in the low-rank model, and designing parametric models of infant sleep behavior based on the patterns extracted by our method.

References

- [1] J. M. Henderson, K. G. France, and N. M. Blampied, “The consolidation of infants’ nocturnal sleep across the first year of life,” *Sleep medicine reviews*, vol. 15, no. 4, pp. 211–220, 2011.
- [2] E. J. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational mathematics*, vol. 9, no. 6, p. 717, 2009.
- [3] R. Vautard, P. Yiou, and M. Ghil, “Singular-spectrum analysis: A toolkit for short, noisy chaotic signals,” *Physica D: Nonlinear Phenomena*, vol. 58, no. 1-4, pp. 95–126, 1992.
- [4] D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, p. 788, 1999.
- [5] S. Seehagen, C. Konrad, J. S. Herbert, and S. Schneider, “Timely sleep facilitates declarative memory consolidation in infants,” *Proceedings of the National Academy of Sciences*, vol. 112, no. 5, pp. 1625–1629, 2015.
- [6] G. Dionne, E. Touchette, N. Forget-Dubois, D. Petit, R. E. Tremblay, J. Y. Montplaisir, and M. Boivin, “Associations between sleep-wake consolidation and language development in early childhood: a longitudinal twin study,” *Sleep*, vol. 34, no. 8, pp. 987–995, 2011.
- [7] A. Sadeh, G. De Marcas, Y. Guri, A. Berger, L. Tikotzky, and Y. Bar-Haim, “Infant sleep predicts attention regulation and behavior problems at 3–4 years of age,” *Developmental neuropsychology*, vol. 40, no. 3, pp. 122–137, 2015.
- [8] J. A. Mindell, E. S. Leichman, J. Composto, C. Lee, B. Bhullar, and R. M. Walters, “Development of infant and toddler sleep patterns: real-world data from a mobile application,” *Journal of sleep research*, vol. 25, no. 5, pp. 508–516, 2016.
- [9] M. Weinraub, R. H. Bender, S. L. Friedman, E. J. Susman, B. Knoke, R. Bradley, R. Houts, and J. Williams, “Patterns of developmental change in infants’ nighttime sleep awakenings from 6 through 36 months of age,” *Developmental Psychology*, vol. 48, no. 6, p. 1511, 2012.
- [10] J. M. Henderson, K. G. France, J. L. Owens, and N. M. Blampied, “Sleeping through the night: the consolidation of self-regulated sleep across the first year of life,” *Pediatrics*, vol. 126, no. 5, pp. e1081–e1087, 2010.
- [11] M. Weissbluth, “Naps in children: 6 months–7 years,” *Sleep*, vol. 18, no. 2, pp. 82–87, 1995.
- [12] I. Iglowstein, O. G. Jenni, L. Molinari, and R. H. Largo, “Sleep duration from infancy to adolescence: reference values and generational trends,” *Pediatrics*, vol. 111, no. 2, pp. 302–307, 2003.
- [13] Z. Chen and A. Cichocki, “Nonnegative matrix factorization with temporal smoothness and/or spatial decorrelation constraints,” in *Laboratory for Advanced Brain Signal Processing, RIKEN, Tech. Rep.*, Citeseer, 2005.
- [14] S. Rallapalli, L. Qiu, Y. Zhang, and Y.-C. Chen, “Exploiting temporal stability and low-rank structure for localization in mobile networks,” in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, pp. 161–172, ACM, 2010.
- [15] V. C. Cheung, K. Devarajan, G. Severini, A. Turolla, and P. Bonato, “Decomposing time series data by a non-negative matrix factorization algorithm with temporally constrained coefficients,” in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3496–3499, IEEE, 2015.
- [16] H.-F. Yu, N. Rao, and I. S. Dhillon, “Temporal regularized matrix factorization for high-dimensional time series prediction,” in *Advances in neural information processing systems*, pp. 847–855, 2016.
- [17] E. G. Tabak and G. Trigila, “Conditional expectation estimation through attributable components,” *Information and Inference*, 2018.
- [18] A. Agarwal, M. J. Amjad, D. Shah, and D. Shen, “Model-agnostic time series analysis via matrix estimation,” *arXiv preprint arXiv:1802.09064*, 2018.
- [19] I. T. Jolliffe, “Principal component analysis and factor analysis,” in *Principal component analysis*, pp. 115–128, Springer, 1986.
- [20] S. Lloyd, “Least squares quantization in pcm,” *IEEE transactions on information theory*, vol. 28, no. 2, pp. 129–137, 1982.

- [21] J. MacQueen *et al.*, “Some methods for classification and analysis of multivariate observations,” in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, pp. 281–297, Oakland, CA, USA, 1967.
- [22] E. A. Nadaraya, “On estimating regression,” *Theory of Probability & Its Applications*, vol. 9, no. 1, pp. 141–142, 1964.
- [23] G. S. Watson, “Smooth regression analysis,” *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 359–372, 1964.

A Cohort and Data Preprocessing

The data used in this work were shared voluntarily by parents using a mobile-phone app, which allows multiple caregivers to track infant sleep. The raw measurements consist of a list of times when the caregiver reported the beginning or end of a sleeping period during the first two years of the infant’s life.

For each infant, days in the data that seem biologically implausible are discarded, since they are likely due to input errors. The criteria to discard days include sleeping periods longer than 16 hours, awake periods longer than 20 hours, no sleep during the night (9pm to 7am), and being surrounded by 5 days or more without data. More than half of the infants are missing at least 50% of the days. Infants for which more than 90 % of the days are missing are not included in the dataset. The resulting number of infants in the dataset is 701 (reduced from original 837 infants).

We represent the time series of sleep data for each infant n as a matrix $Y^{[n]}$. Each row corresponds to a day (row 1 corresponds to the day the infant was born). Columns correspond to the time of the day, discretized into 144 10-minute intervals. Each entry $Y^{[n]}(t, i)$ contains the fraction of time the infant was reported to be asleep on day t during the i th 10-minute interval (1 indicates that the infant slept during the whole period, 0 that they were awake). Figure 8 in the appendix shows several examples.

B Implementation details and parameter selection

We implement our methodology in Matlab, using the function *nmf* to obtain an initialization for the basis functions based on nonnegative matrix factorization, and the function *lsqlin* to solve the nonnegative least squares subproblems defined in Section B.1. We fix a rank of $r := 5$ in our analysis. As shown in Figure 9 of the appendix, the singular values computed over the whole dataset suggest that the data is well approximated by a dimension-5 subspace (there is a jump between the 5th and 6th singular value). Figure 10 in the appendix shows basis functions learned using other ranks. The value of the smoothness regularization parameter is fixed to $\lambda := 10^5$. The results are very robust to changes in this parameter, as illustrated in Figure 11 of the appendix. The basis functions change only slightly when λ is large enough, whereas the coefficient vectors become smoother gradually as we increase λ .

B.1 Optimization algorithm

To minimize the cost function (3) we apply alternating minimization. The coefficients are updated by solving N separate nonnegative least-squares problem, one for each time series,

$$\widehat{C}_1^{[n]}, \dots, \widehat{C}_r^{[n]} := \arg \min_{\substack{C_1, \dots, C_r \in \mathbb{R}^{|\mathcal{T}_n|} \\ C_j(t) \geq 0, \forall j, t}} \sum_{t \in \mathcal{T}_n} \sum_{i=1}^{\ell} \left(Y^{[n]}(t, i) - \sum_{j=1}^r C_j(t) F_j(i) \right)^2 + \lambda \sum_{j=1}^r \| \mathcal{D} C_j \|^2$$

for each n . These N subproblems are completely decoupled and can consequently be solved in parallel. The basis functions are updated by solving the nonnegative least-squares problem,

$$\widehat{F}_1, \dots, \widehat{F}_r := \arg \min_{\substack{F_1, \dots, F_r \in \mathbb{R}^{\ell} \\ F_j(i) \geq 0, \forall j, i}} \sum_{n=1}^N \sum_{t \in \mathcal{T}_n} \sum_{i=1}^{\ell} \left(Y^{[n]}(t, i) - \sum_{j=1}^r C_j^{[n]}(t) F_j(i) \right)^2, \quad (4)$$

and then normalizing each function to have unit ℓ_2 norm.

C Other applications

C.1 Outlier detection

Identifying anomalous sleeping patterns in infants has great potential clinical value. Unfortunately, detecting outliers in a set of time series is very challenging because the data are usually very high dimensional. Points in high-dimensional spaces tend to all be at a similar distance from each other, so that comparing these distances is not meaningful, a phenomenon commonly known as the *curse of dimensionality*. The histogram at the top of Figure 3 illustrates this for our infant-sleep data. We tackle this issue by using the coefficients in our low-rank model as a representation of reduced dimensionality. This has the added benefit that the outliers can be interpreted in terms of the sleeping patterns described in Section 3. We are particularly interested in the first three patterns (irregular sleep, nighttime sleep with two naps, and nighttime sleep with a single nap), since the remaining two are associated to shifts in waking hours, which are not as meaningful from a clinical point of view. The second row of Figure 3 displays the histogram of distances to the medians of each component (shown in Figure 2). In contrast to the histogram for the raw data, the distribution of distances is heavily skewed towards smaller distances. This makes it possible to characterize outliers as infants that are on the right tail of the histogram.

Figure 3 also shows several outliers associated to each of our patterns of interest. By comparing to the general trends in Figure 2, we can analyze how they deviate from the normal behavior in the population. For the daytime-sleep pattern, the first and second outliers develop circadian sleep appears later than usual, whereas the third develops structured naps very early. For the two-nap pattern, the first outlier develops three structured naps exceptionally early, the second has a very fuzzy two-nap schedule, and the third one has two naps until an abnormally late age. For the one-nap pattern, the first outlier never develops a one-nap habit, whereas the second and third have a fuzzy one-nap habit that starts unusually early.

C.2 Clustering

In this section we consider the problem of automatically clustering infants according to their sleeping behavior. From a clinical point of view, the goal is to determine whether there are subpopulations of infants with distinct tendencies, as a preliminary step in the study of factors that affect sleep development. Figures 4 and 12 shows the result of separating the data into two clusters using a standard algorithm by [20] to minimize the k-means cost function [21]. Unfortunately, the clustering procedure just captures differences in nighttime sleep schedules that are not clinically meaningful: cluster 1 contains infants that wake up and go to bed earlier than those in cluster 2.

As described in Section 3, three of the daily-sleep patterns extracted by NMF-TS are associated to meaningful sleeping habits, whereas the remaining two just capture shifts in nighttime sleep. We build upon this insight to design a metric that is approximately invariant to shifts. Let $Y^{[n]}$ and $Y^{[n']}$ be the data corresponding to two different infants. The metric is equal to the ℓ_2 norm of the difference between the coefficients associated to the first three patterns for each of the infants,

$$d_{\text{NMF-TS}}(Y^{[n]}, Y^{[n']}) := \sqrt{\frac{1}{|\mathcal{T}_n \cap \mathcal{T}_{n'}|} \sum_{j=1}^3 \sum_{t \in \mathcal{T}_n \cap \mathcal{T}_{n'}} (C_j^{[n]}(t) - C_j^{[n']}(t))^2}. \quad (5)$$

The metric only takes into account entries in $\mathcal{T}_n \cap \mathcal{T}_{n'}$ (i.e. observations present in the two time series). It is normalized by $|\mathcal{T}_n \cap \mathcal{T}_{n'}|$ to render it as invariant as possible to the number of missing entries. Figure 4 and 13 shows the results of applying k-means clustering based on $d_{\text{NMF-TS}}$. The clusters separate the infants according to sleep behavior: cluster 1 contains infants that have a more regular napping behavior, a more consolidated circadian rhythm, and develop a one-nap schedule earlier compared to the infants in cluster 2. Our methodology also makes it possible to cluster the population according to more specific behaviors. This is achieved by only including the coefficients corresponding to a single basis function in Eq (5). Figure 14 shows the cluster centers obtained by applying k means with the metric associated to each of the first three basis functions (instead of all three). The clustering clearly captures differences associated to the corresponding sleeping pattern.

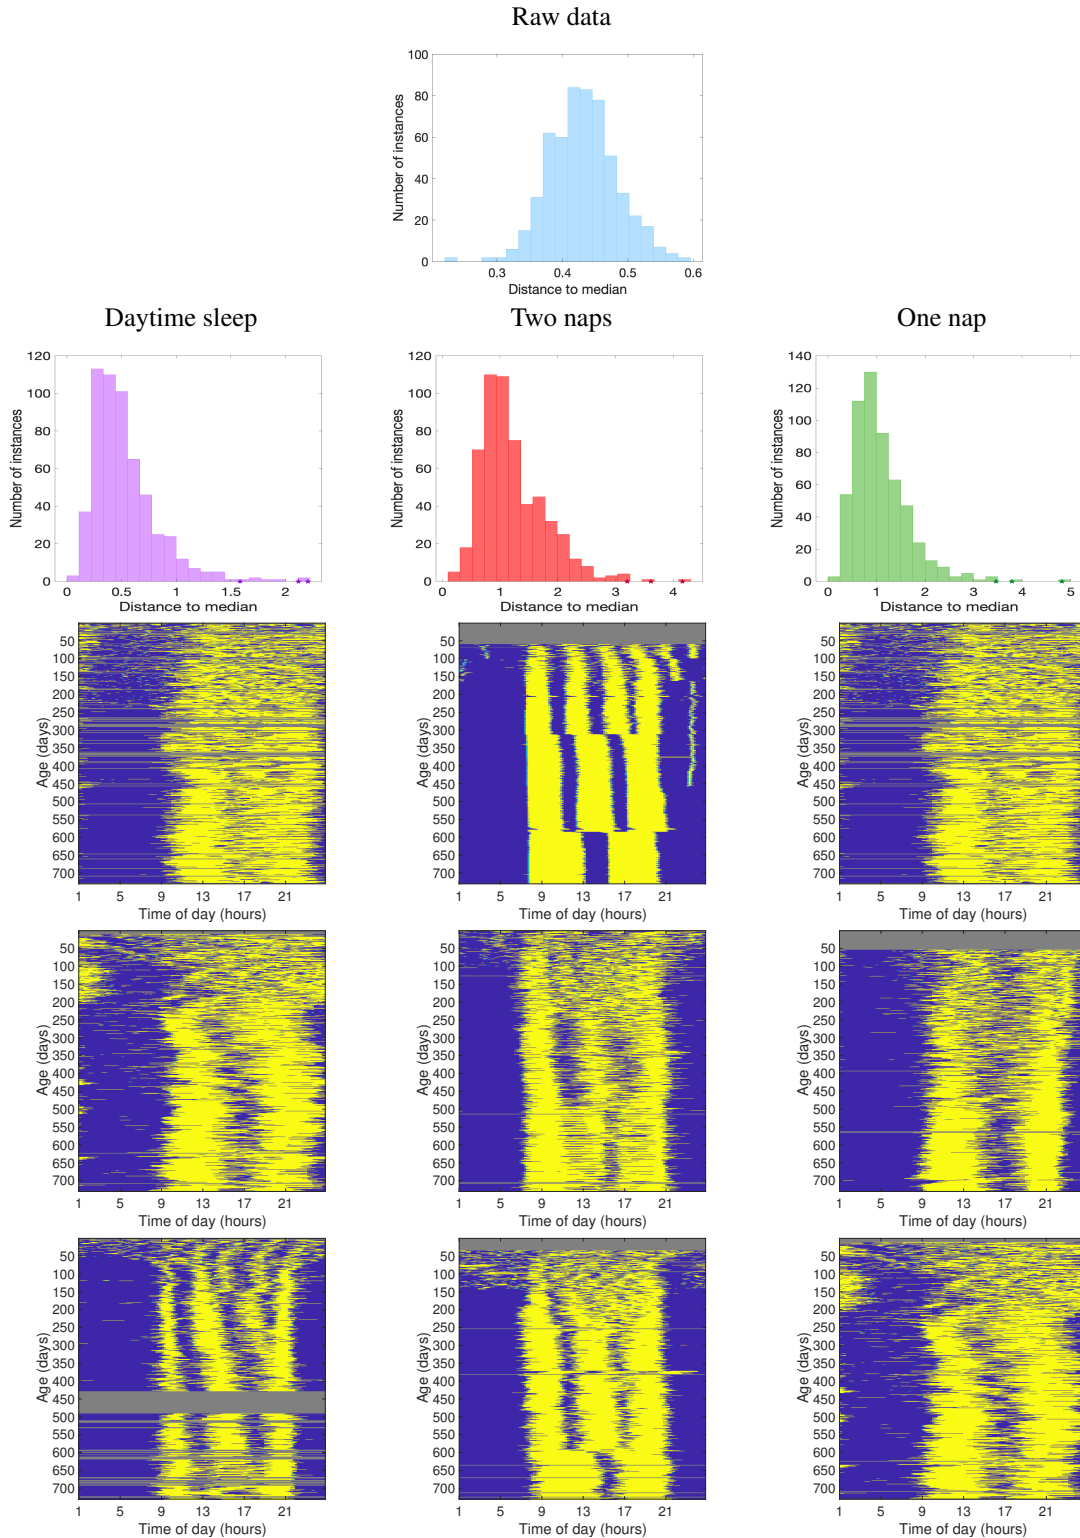


Figure 3: The top graph shows the histogram of the ℓ_2 -norm distances between the time series in the infant-sleep dataset and their median. The second row shows histograms of distances between the low-rank components extracted by the proposed methodology and the corresponding median (shown in Figure 2). For each component we show three examples of outliers marked with stars on the histogram. Yellow indicates that the infant is awake, and blue that they are asleep. Gray values indicate missing data.

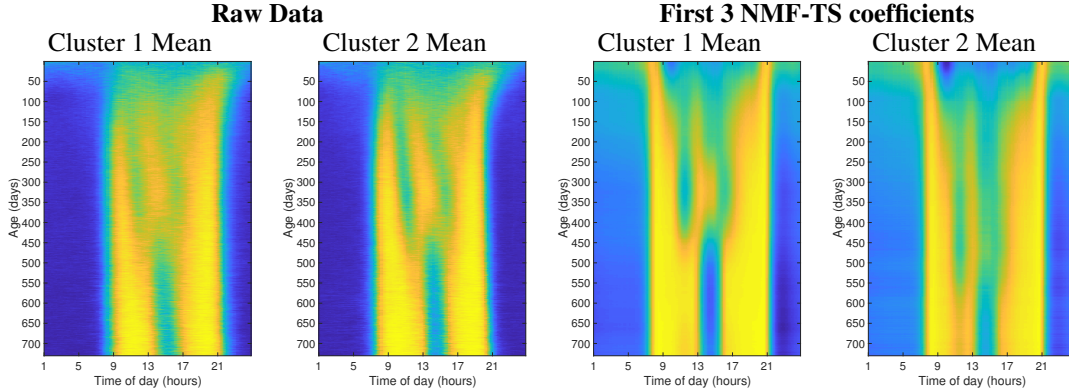


Figure 4: Cluster centers obtained by applying Lloyd’s algorithm to perform k-means clustering with $k := 2$ on the raw data (left) and on the first 3 NMF-TS coefficients (right). Clustering the raw data just captures differences in nighttime sleep schedule, whereas clustering the coefficients uncovers differences in nap behavior.

	Forecasting period (days of age)		
	100-200 from 1-100	300-400 from 200-300	600-700 from 500-600
Mean	0.387 ($3.66 \cdot 10^{-2}$)	0.337 ($3.68 \cdot 10^{-2}$)	0.292 ($3.78 \cdot 10^{-2}$)
KR (raw data)	0.382 ($9.76 \cdot 10^{-2}$)	0.326 ($4.04 \cdot 10^{-2}$)	0.273 ($3.65 \cdot 10^{-2}$)
KR (NMF-TS)	0.374 ($2.81 \cdot 10^{-2}$)	0.318 ($3.27 \cdot 10^{-2}$)	0.270 ($3.54 \cdot 10^{-2}$)
Rank-5 ground truth	0.359 ($3.11 \cdot 10^{-2}$)	0.305 ($3.35 \cdot 10^{-2}$)	0.257 ($3.28 \cdot 10^{-2}$)

Table 1: Average forecasting error per 10-min interval for the infant-sleep dataset on held-out data. The standard deviation of the error over the test set is given in brackets.

C.3 Forecasting sleeping behavior

In this section we consider the problem of forecasting the sleep tendencies of infants based on their past sleeping behavior using kernel regression, as described in Section C.4. In order to produce forecasts that are robust to noise and spurious short-term fluctuations, we propose to estimate the coefficients in the rank-5 NMF-TS decomposition, instead of the raw data. To evaluate our approach we separate the infant-sleep dataset into a training set with 561 time series, and a test set with 140 examples. The low-rank model is fit to the training data. The test data are only used to compute forecasting error. Table 1 reports the performance of the method over three different 100-day periods using sleeping data from the previous 100 days. For each time period, we only take into account time series that have at least 70% observed data in that interval. We compare the result of (1) prediction based on the mean of the training data, (2) kernel regression (KR) using the raw data, and (3) kernel regression using the NMF-TS coefficients. We also compare to the ground-truth rank-5 coefficients computed from the true future data (this is the best-case scenario for a rank-5 model). The error is high for all methods, in the range 0.25-0.4 for true values that are between 0 (awake) and 1 (asleep). This is expected: the dimension of the estimated signals is $100\ell = 14,400$, an order of magnitude larger than the number of examples! We can only hope to predict general trends in the future sleep behavior, which motivates our approach of predicting smooth coefficients associated to a small number of basis functions. The results show that this provides an improvement over using the raw data for all three time intervals.

C.4 Kernel Regression

First, we fit the rank-5 NMF-TS model on the training data to obtain the basis functions F_1, \dots, F_5 and the corresponding coefficients. We then extract the coefficients corresponding to the *past* interval

and to the *future* interval that we want to predict,

$$\left(\left\{ C_{1,\text{past}}^{[n]}, \dots, C_{5,\text{past}}^{[n]} \right\}, \left\{ C_{1,\text{future}}^{[n]}, \dots, C_{5,\text{future}}^{[n]} \right\} \right), \quad 1 \leq n \leq N_{\text{train}}. \quad (6)$$

Given a test example for which we only have past data, we compute its rank-5 NMF-TS coefficients $\{C_{1,\text{past}}^{\text{test}}, \dots, C_{5,\text{past}}^{\text{test}}\}$ using the basis functions obtained from the training set. Then we apply kernel regression, a nonparametric method proposed by [22] and [23], to estimate its future coefficients:

$$\hat{C}_{i,\text{future}}^{\text{test}} := \frac{\sum_{n=1}^{N_{\text{train}}} K(C_{i,\text{past}}^{[n]}, C_{i,\text{past}}^{\text{test}}) C_{i,\text{future}}^{[n]}}{\sum_{n=1}^{N_{\text{train}}} K(C_{i,\text{past}}^{[n]}, C_{i,\text{past}}^{\text{test}})}, \quad 1 \leq i \leq r, \quad (7)$$

The kernel K is Gaussian with standard deviation σ , which is set via cross-validation on the training set. Intuitively, the estimate is a weighted average of instances that have similar past data to the test example. The estimated coefficients are then combined using the basis functions from the training set, to produce a low-rank estimate of future sleep.

D Additional Figures

In this section we provide additional figures to illustrate our methods and analysis. Figure 5 shows Low-rank approximation obtained from different methods. Figure 6 shows basis functions extracted from the infant-sleep dataset using PCA, NMF and NMF-TS. Figure 7 shows coefficients corresponding to the basis functions for the infant whose data is shown in Figure 5. Figure 8 shows raw data from infants randomly selected from the dataset. Each data matrix is shown accompanied by the coefficients of its NMF-TS low-rank representation. Figure 9 shows the singular values computed over the whole dataset. The jump between the 5th and 6th singular value suggests using a rank-5 low-rank representation. Figure 10 shows the basis functions extracted by PCA (top), NMF (center) and the proposed approach (bottom) for different values of the rank r . The value $r = 5$, which we use for our analysis, yields the most interpretable model. Figure 11 illustrates the effect of varying the parameter λ on the NMF-TS basis functions and coefficients. The daily-sleep patterns change only slightly, whereas the age-development become smoother as we increase λ . We set $\lambda = 10^5$ for our analysis. Figure 12 and Figure 13 show the results of applying Lloyd’s algorithm to perform k-means clustering on the raw infant-sleep data and on the three first NMF-TS coefficients respectively. Figures 15, 16 and 17 show examples of forecasts computed as described in Section C.3. The low-rank prediction obtained by applying kernel regression on the NMF-TS coefficients is visually similar to the true rank-5 representation of the future data.

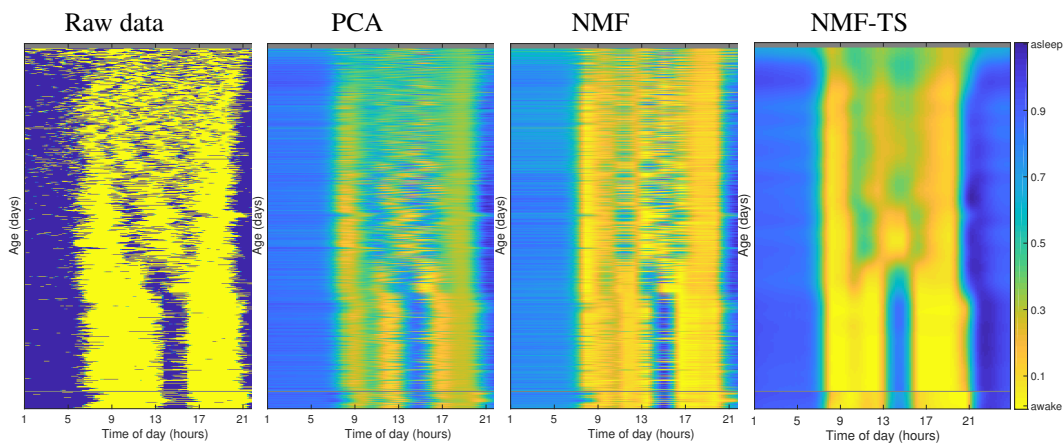


Figure 5: Low-rank approximation (given by equation (2)) of the data corresponding to a specific infant obtained combining the daily-sleep patterns and the age development patterns extracted by PCA (2nd column), NMF (3rd column) and the proposed method (4th column). The raw data is shown in the first column. Yellow indicates that the infant is awake, whereas dark blue indicates that they are asleep. The average fitting error per entry for this infant is equal to 0.091 for PCA, 0.104 for NMF and 0.116 for NMF-TS (as expected, since additional constraints reduce overfitting). Gray values indicate missing data.




Figure 6: Basis functions F_1, \dots, F_7 extracted from the infant-sleep dataset using PCA (left), NMF (center) and NMF-TS (right). The rank of the representation is set to $r := 5$. Both NMF and NMF-TS recover patterns that are nonnegative and therefore more interpretable in terms of sleeping behavior. Large values are associated to sleep, whereas values close to zero indicate wakefulness.




Figure 7: Coefficients $C_1^{[r]}, \dots, C_r^{[r]}$ corresponding to the basis functions in Figure 6 for the infant whose data is shown in Figure 5. The patterns obtained from NMF are more interpretable than those obtained from PCA, but are also extremely noisy. In contrast, the patterns recovered by NMF-TS are smooth and can be interpreted in terms of sleeping behavior that are consistent with the infant's raw data in Figure 5.

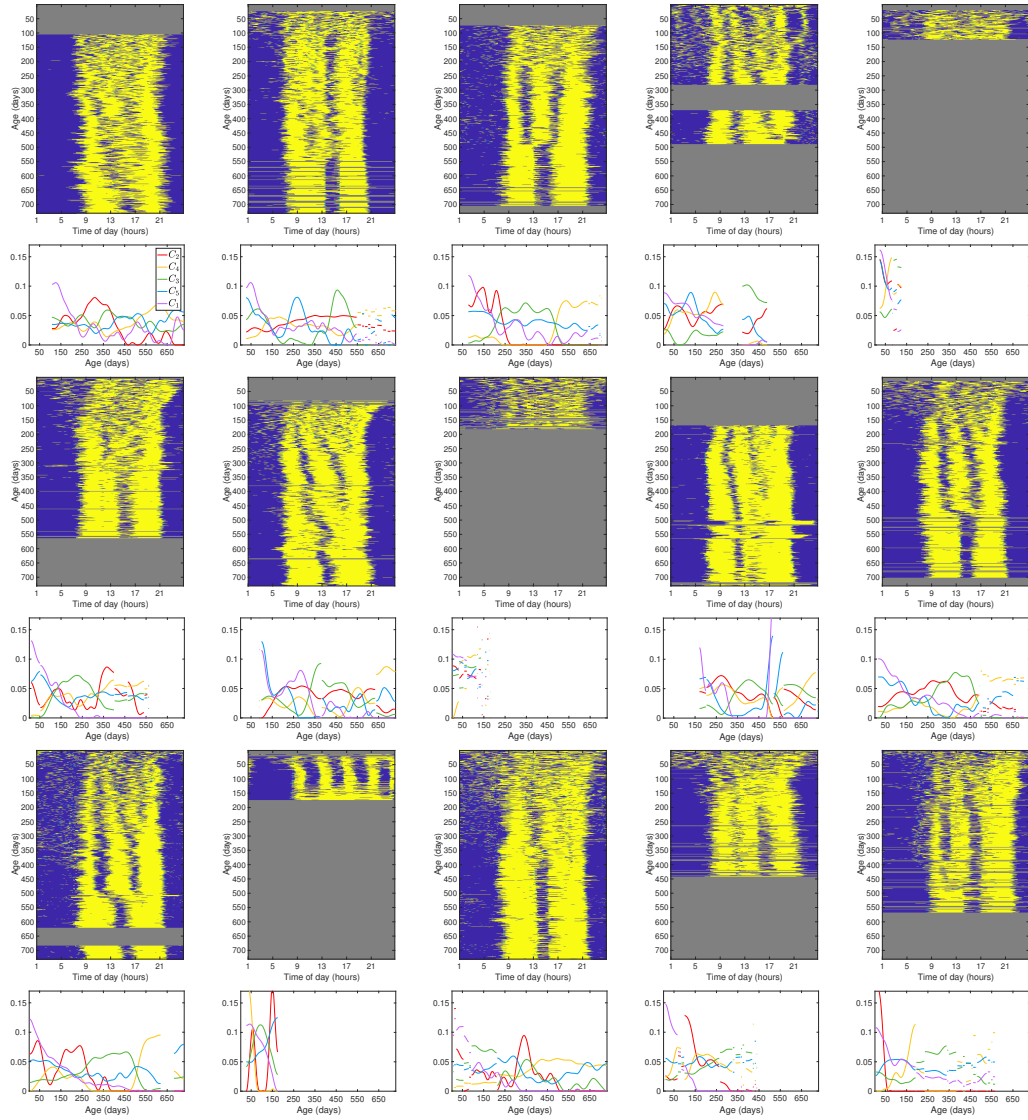


Figure 8: Raw data from infants randomly selected from the dataset. Yellow indicates that the infant is awake, and blue that they are asleep. Gray values indicate missing data. Below each data matrix we show the coefficients from the NMF-TS representation.




Figure 9: First 50 singular values obtained by computing the singular-value decomposition of the infant-sleep dataset.

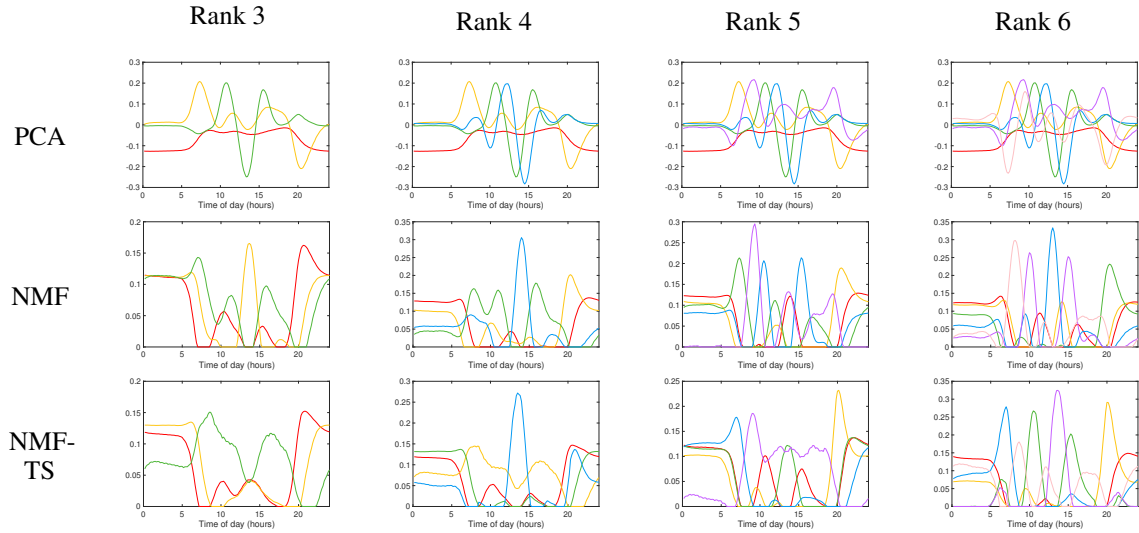


Figure 10: Basis functions extracted from the whole dataset using PCA (top), NMF (center) and NMF-TS (bottom) for different values of the rank r . For NMF-TS we set $\lambda := 10^5$.

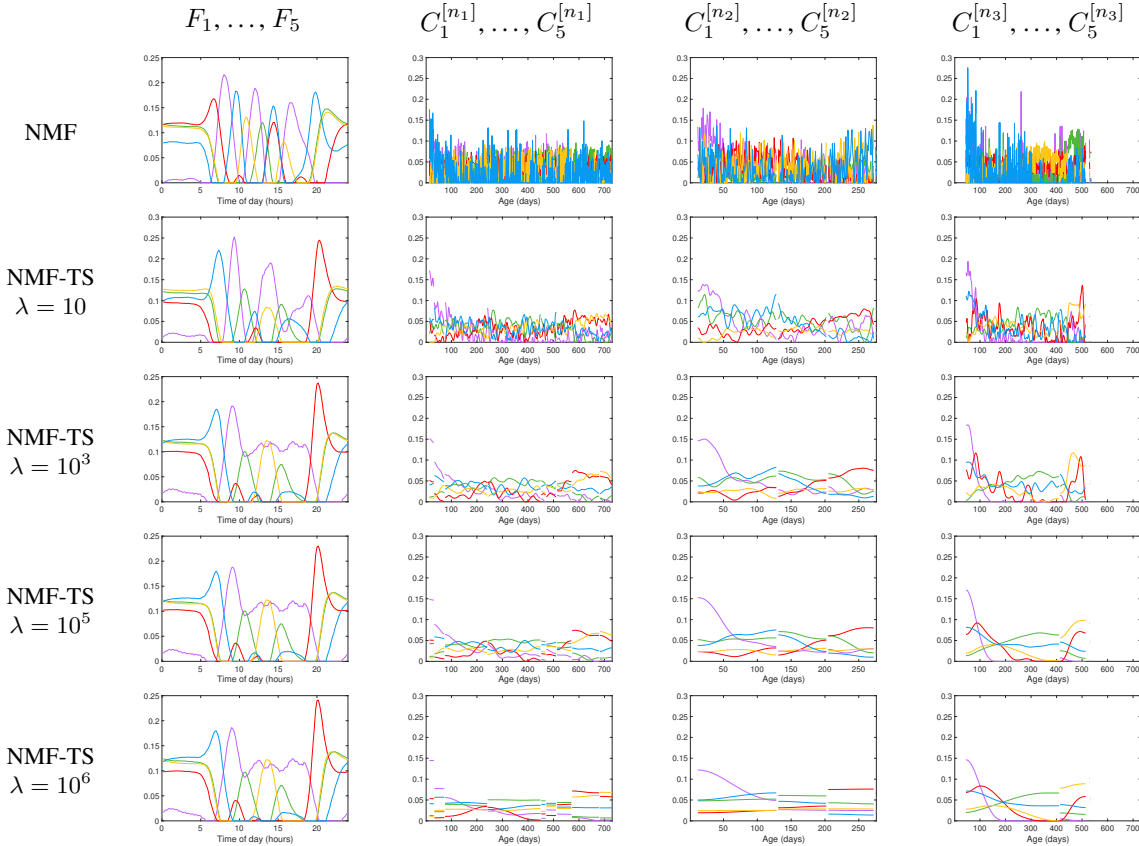


Figure 11: Effect of varying the parameter λ on the basis functions F_1, \dots, F_r and the coefficients $C_1^{[n_1]}, \dots, C_r^{[n_1]}$ for three different time series n_1, n_2, n_3 . The basis functions change only slightly once λ is large enough. The coefficients become smoother as we increase λ .

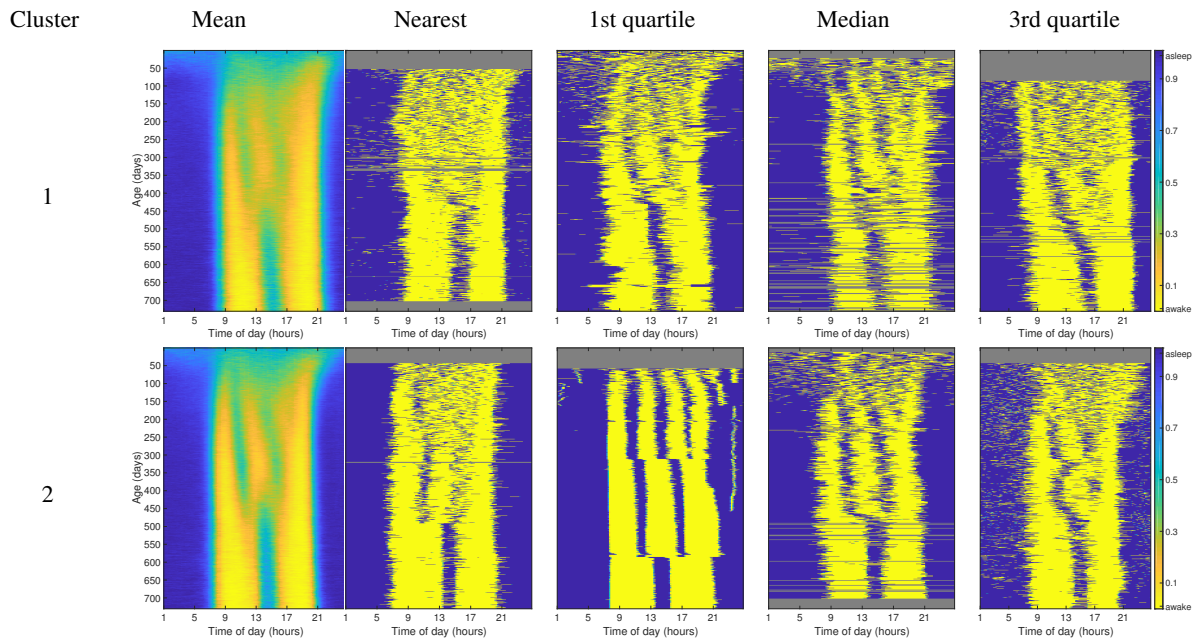


Figure 12: Results of applying Lloyd’s algorithm to perform k-means clustering with $k := 2$ to the infant-sleep data. For the two clusters we show the cluster mean (1st column), the closest infant to the mean (2nd column), and the first quartile, median and third quartile infants in terms of distance to the mean out of infants with at least 80% available data (3rd, 4th and 5th columns respectively). The first cluster contains 308 infants. The second cluster contains 393. Gray values indicate missing data.




Figure 13: Results of applying Lloyd's algorithm to perform k-means clustering with $k := 2$ based on the metric $d_{\text{NMF-TS}}$ defined in Eq. (5). For the two clusters we show the low-rank cluster mean (1st column), the closest infant to the mean (2nd column), and the first quartile, median and third quartile infants in terms of distance to the mean out of infants with at least 80% available data (3rd, 4th and 5th columns respectively). Below each image we also show the coefficients corresponding to the three sleeping patterns included in the metric. The first cluster contains 359 infants. The second cluster contains 342. Gray values indicate missing data.

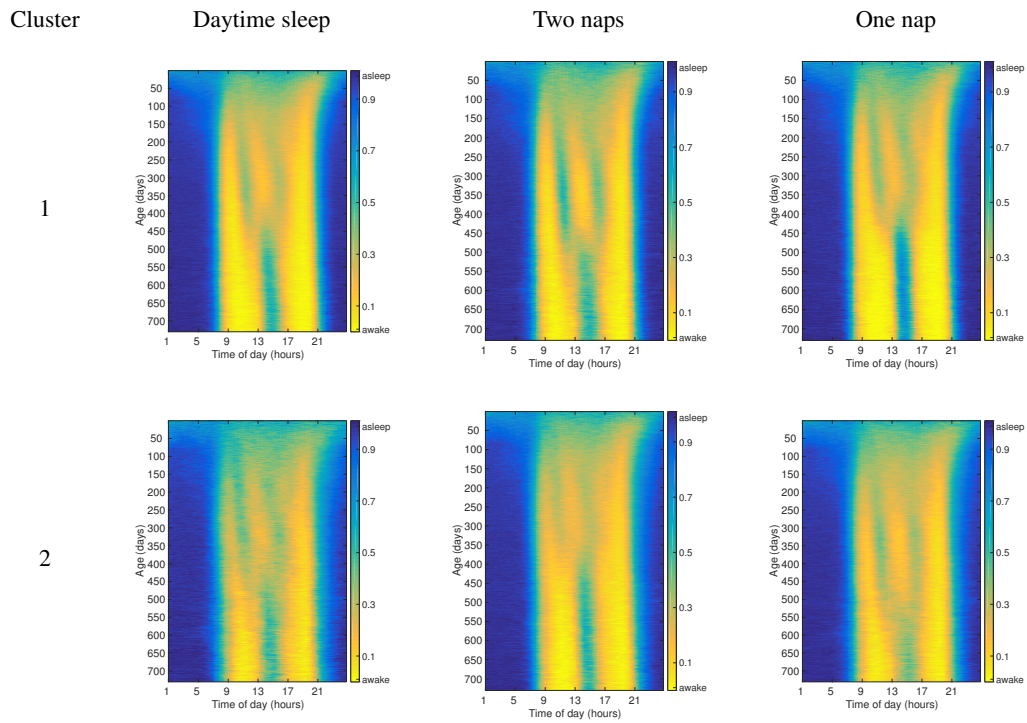


Figure 14: Cluster centers obtained by performing k-means clustering with $k := 2$ using a distance associated to each of the first three basis functions in Figure 2. The difference between the centers depends on the corresponding sleeping pattern. In particular, the centers in the top row show a more clearly marked emergence of circadian sleep, two-nap schedule, and one-nap schedule respectively.

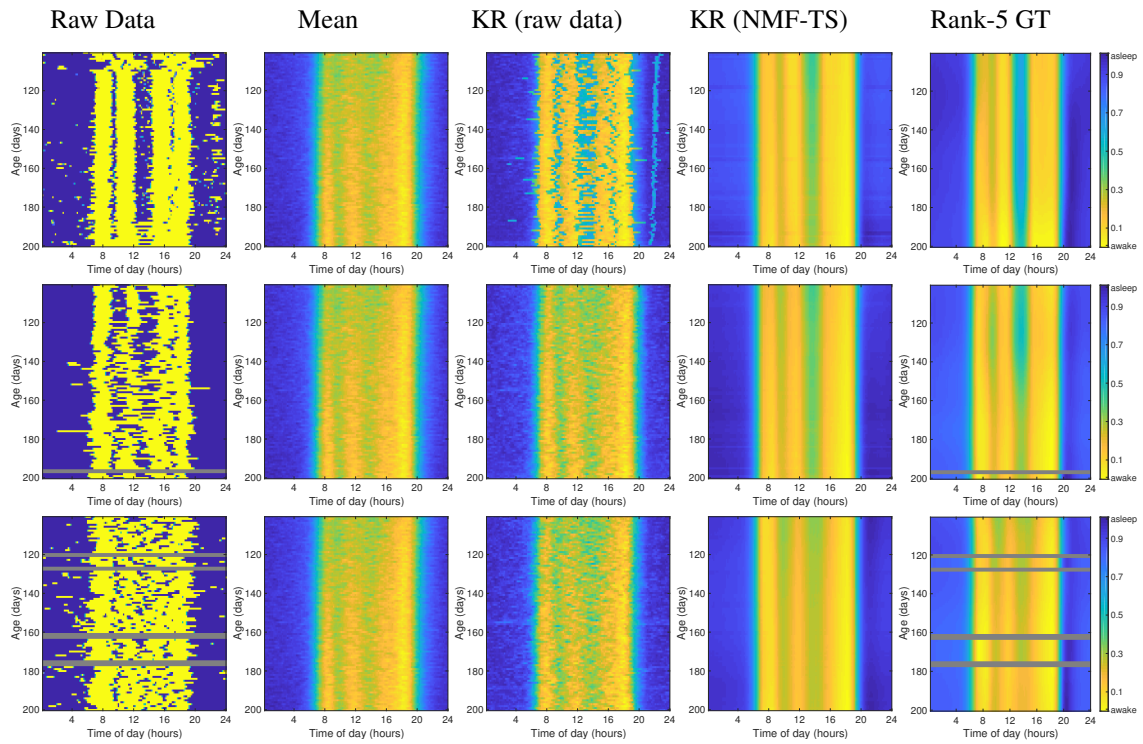


Figure 15: Examples of sleep forecasts from days 100 to 200 of age, computed from the data of days 0 to 100. For each example, we show the future data, the mean of the training set, the kernel-regression (KR) using the raw training data, the kernel-regression prediction based on NMF-TS coefficients, and the rank-5 representation of the future data.

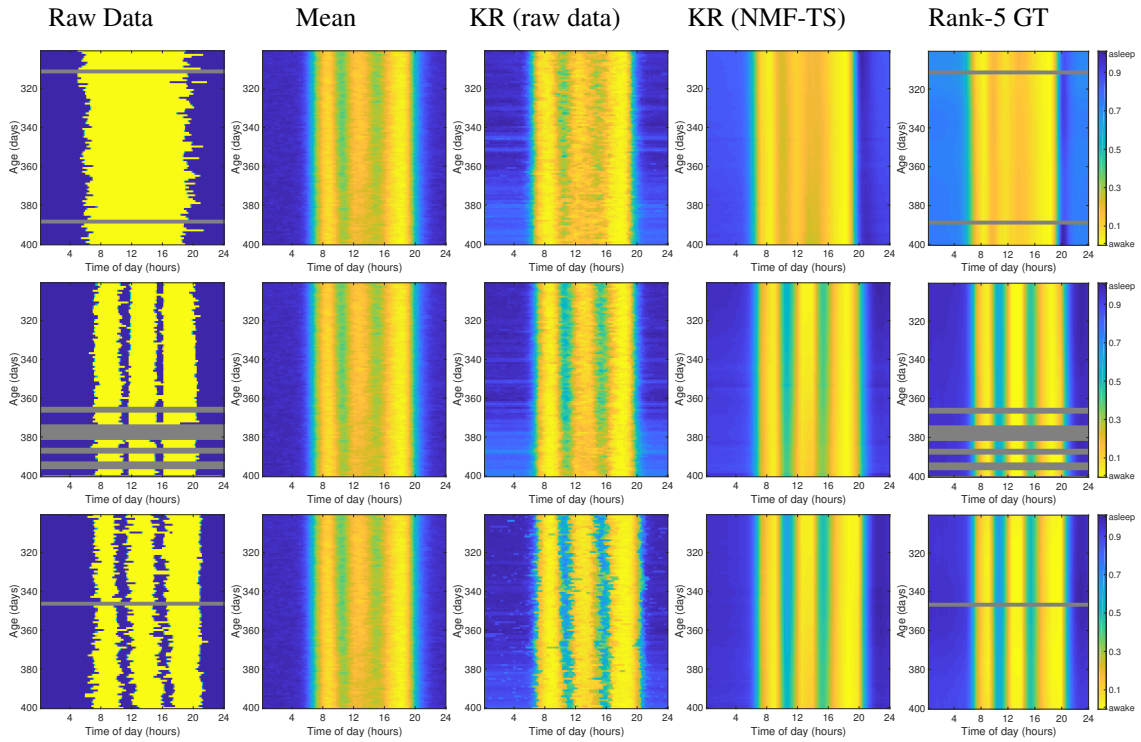


Figure 16: Examples of sleep forecasts from days 300 to 400 of age, computed from the data of days 200 to 300. For each example, we show the future data, the mean of the training set, the kernel-regression (KR) using the raw training data, the kernel-regression prediction based on NMF-TS coefficients, and the rank-5 representation of the future data.

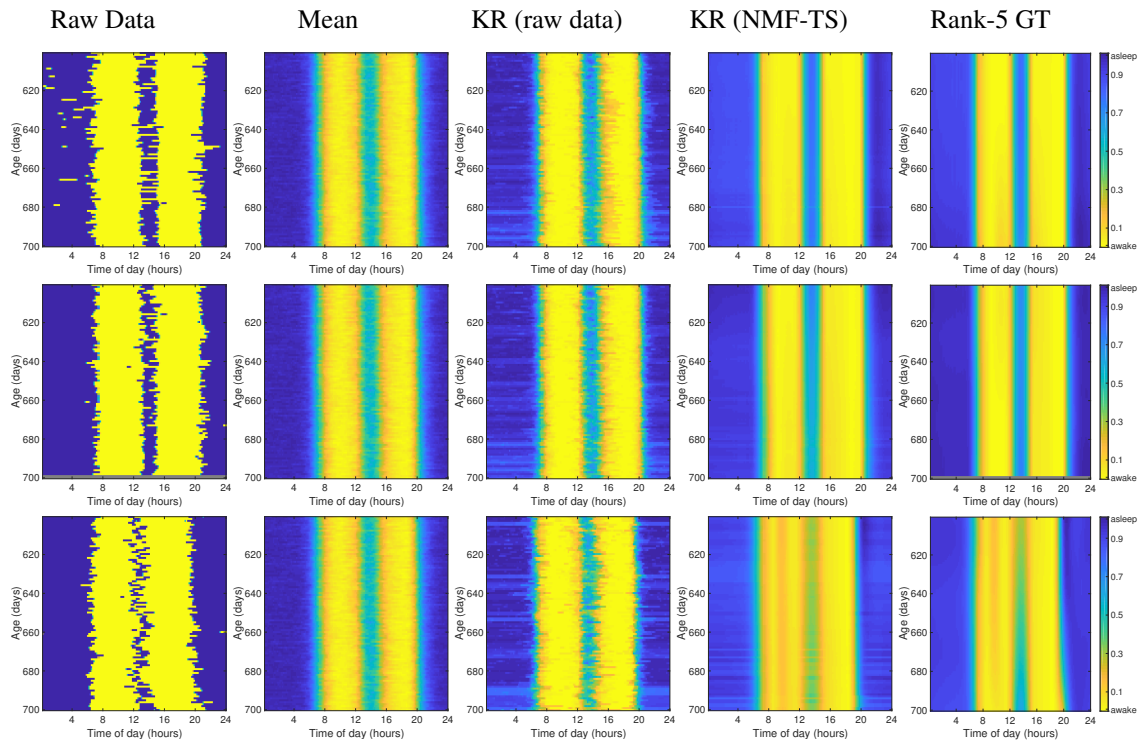


Figure 17: Examples of sleep forecasts from days 600 to 700 of age, computed from the data of days 500 to 600. For each example, we show the future data, the mean of the training set, the kernel-regression (KR) using the raw training data, the kernel-regression prediction based on NMF-TS coefficients, and the rank-5 representation of the future data.