Published in final edited form as:

Comput Stat Data Anal. 2015 September ; 89: 1–11. doi:10.1016/j.csda.2015.03.001.

Modeling sleep fragmentation in sleep hypnograms: An instance of fast, scalable discrete-state, discrete-time analyses

Bruce J. Swihart^{a,*}, Naresh M. Punjabi^b, and Ciprian M. Crainiceanu^a

^a Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, Baltimore, MD, 21205, United States

^b Johns Hopkins University, Baltimore, MD, United States

Abstract

Methods are introduced for the analysis of large sets of sleep study data (hypnograms) using a 5-state 20-transition-type structure defined by the American Academy of Sleep Medicine. Application of these methods to the hypnograms of 5598 subjects from the Sleep Heart Health Study provide: the first analysis of sleep hypnogram data of such size and complexity in a community cohort with a range of sleep-disordered breathing severity; introduce a novel approach to compare 5-state (20-transition-type) to 3-state (6-transition-type) sleep structures to assess information loss from combining sleep state categories; extend current approaches of multivariate survival data analysis to clustered, recurrent event discrete-state discrete-time processes; and provide scalable solutions for data analyses required by the case study. The analysis provides detailed new insights into the association between sleep-disordered breathing and sleep architecture. The example data and both R and SAS code are included in online supplementary materials.

Keywords

Competing risks; Multi-state; Poisson regression; Recurrent event; Sleep-disordered breathing; Stratified

1. Introduction

HIGHLIGHTS

- We explore associations of sleep-disordered breathing and sleep structure.
- 5-state hypnograms are systematically compared to 3-state.
- We analyze a community cohort of 5598 subjects (2.7 million rows total).

* Corresponding author. Tel.: +1 443 287 8782; fax: +1 410 955 0958. bruce.swihart@gmail.com (B.J. Swihart). URL:<http://www.biostat.jhsph.edu/~bswihart/> (B.J. Swihart)..

Supplementary materials.

Supporting information for this article is available online: datasets, R and SAS code in a 27 MB.zip file (<http://www.biostat.jhsph.edu/%7EBswihart/Publications/pophyp.zip>).

- We reduce analysis time from 8 hours to 30 s.

An individual's sleep is conceptualized as a hypnogram, a discrete-state discrete-time stochastic process (Fig. 1). Currently, the field of sleep science broadly generalizes a typical sleep progression as Wake (W, on the hypnogram axis) to Stage 1 (1) to Stage 2 (2) to Stage Slow-wave (S) back to Stage 2 and then Rapid Eye Movement (R). The progression makes up one sleep cycle, which lasts approximately 60–90 min and repeats through the night.

Judging from the three examples of Fig. 1, sleep is often more complex than the purportedly ‘typical’ pattern shown of the top panel. The top panel fits the generalization well, but the other two do not, with many detours and interruptions to the charted course of a “typical” sleep: the middle panel has more alternations between Stage 2 and Stage Slow-wave; while the bottom panel has a duration in Wake before leaving Stage 1 and a much more fragmented Stage Slow-wave portion. Given that every 30 s the trajectory can change to any other of the four states or remain in the current state makes for a diverse functional space for one hour snippets of three individuals, let alone for an overall sleep time of typically 7 h for the thousands of individuals in sleep epidemiology investigations.

1.1. Previous methods used in the analysis of hypnogram data

To remove potential confusion about novelty, we provide a brief review of related publications with the problems they addressed, along with a summarizing graphic (Fig. 2). Hypnogram data have been the topic of previous analytic frameworks and data applications demonstrated on modest sample sizes involving one homogeneous group or for comparing no more than two groups. For example, previous statistical methodology for sleep focused on an important clinical goal of relating time-varying hormone levels to the sleep process modeled as a reduced set of transition-types for the number of states considered (see Fahrmeir and Klinger, 1998; Yassouridis et al., 1999; Aalen et al., 2004; Kneib and Hennerfeind, 2008; Kalus et al., 2009). Norman et al. (2006) established differences in the stability of sleep and severity of sleep-disordered breathing in a sleep-runs analysis (which is a 2-state 1-transition-formulation with no intra-subject repeated events clustering). For S states, a full and flat model paradigm would include all $S \times (S - 1)$ pairwise transition-types. Swihart et al. (2008) analyzed the full 3-state 6 transition-type paradigm with Poisson regression for relative transition counts and multi-state survival models for relative transition rates to study sleep stability. Swihart et al. (in press) implemented a random effects Bayesian Poisson regression.

The development of a new approach was spurred by epidemiologic sleep studies, which often have thousands of subjects and research goals involving comparisons of multiple subgroups. The methods described above, save for a modification to the Poisson regression in Swihart et al. (2008), were unable to scale to 5598 subjects under the 5-state 20-transition-type paradigm with the focus of analyzing group differences in the transition process.

Thus, the goal is to provide statistical models that: (1) are not more complex than necessary for comparing (more than 2) groups; (2) characterize transition-type-specific features of the

frequency and rate behaviors observed in Fig. 1; and (3) do not over-simplify the sleep state transition process, as done by the currently accepted 3-state characterization of sleep cycles.

2. Data description and methods

The hypnogram data is supplied by the Sleep Heart Health Study (Quan et al., 1997). Discrete-time discrete-state hypnograms are ASCII files of one line, displaying a symbol to represent the occupied state for sequential and mutually exclusive 30-s epochs. For example, RRRR22121W could be the 10 epoch tail-end of a string, where one R is 30-ss of REM sleep, 2 is Stage 2, 1 is Stage 1, and W is wake (Fig. 3, top left). A transition occurs whenever there is a change of symbol adjacent to one another. There are 5 transitions in this string, chronologically: a REM-Stage 2 (labeled R2) transition, Stage 2–Stage 1 (labeled 21) transition, Stage 1–Stage 2 (labeled 12) transition, Stage 2–Stage 1 (labeled 21) transition, and a Stage 1–Wake (labeled 1W) transition. A corresponding time-at-risk (interchangeably, duration in state, transition time, survival time, sojourn time, failure time, or time-to-event) can be assigned to each transition. The time-at-risk is state-specific and is assigned to all possible transitions (observed and censored) that have the same starting state. In the current example, the observed transitions had the following times at risk: R2 was 2 min, transition-type 21 had 1 min, transition-type 12 had 0.5 min, the second occurrence of transition-type 21 had 0.5 min, 1W had 0.5 min, and no transition out was recorded of the final Wake epoch.

The stages of sleep are collapsable, yielding hypnograms of fewer states. The collapsibility is biologically motivated. For instance, to go from a 5-state to 3-state hypnogram, Stage 1, Stage 2 and Stage Slow-wave are combined into Non-REM (NREM) sleep stage. In our example, RRRR22121W becomes RRRRNNNNNW (Fig. 3, middle left). Continuing in this vein, both NREM and REM sleep stages of 3-state sleep can be collapsed into an “Asleep” stage, making sleep a 2-state process, where RRRRNNNNNW becomes AAAAAAAAAW (Fig. 3, bottom left).

Two methods of analysis are proposed to characterize the sleep transitions and assess covariate effects on these transition measures. A well-known equivalence exists between the likelihoods of a Poisson regression via the generalized linear model (GLM) framework and a parametric survival model with exponentially distributed survival times which feature (piecewise) constant hazards (Holford, 1976, 1980; Laird and Olivier, 1981; Swihart et al., in press). The aforementioned equivalence prompts the investigation of a Poisson regression via a generalized estimating equations (GEE) log-linear model as an approximation to a multi-state proportional hazards model, which are each fitted to a 5-state and 3-state representation of the sleep hypnogram. Evaluating the presupposed, approximate equivalence of these two models is needed, as each serves as a near but distinct approach to those known in the exact equivalence. Herein, “Poisson regression” will be used for the GEE log-linear model, and “multi-state survival model” will represent the stratified multi-state proportional hazards model for competing risks and recurrent events.

For each hypnogram resolution (5-, 3-, or 2-stage), the two models described above each require a different data format. A multi-state survival model requires a format detailing

chronologically each possible transition in a separate row, whether the transition was observed or censored, and the time to transition (Fig. 3). Poisson regression requires transition-type specific total counts of occurrence and total time at risk (tar) for those counts, which is straightforwardly borne of summing the observed (obs) and time-to-event (tte) variables of the properly constructed multi-state survival model format, respectively, by transition-type (shift).

2.1. Multi-state survival model

Consider group g and transition-type h : for a survival analysis yielding a transition-type specific log-hazard $\alpha_h(t)$ and effect $\beta_{g,h}$, the multi-state survival model is stratified on transition type h and regressed upon the interaction term involving h :

$$\alpha_h(t) = \alpha_{h0}(t) + g:h + x,$$

where $g:h$ represents the interaction terms sans the main effects, as discussed in Therneau and Grambsch (2000). The t is the associated time-to-event (tte) of the transition. The event may be censored (obs = 0) or observed (obs = 1, see Fig. 3). The x represents demographic covariates as typically included in a multi-state survival model. The coefficients $\beta_{g,h}$ are the focus of the analysis—the demographic covariates are included for adjustment only. Directly exponentiating $\beta_{g,h}$ gives the hazard ratio of group g to the reference group for transition-type h .

To fit the 3-stage, 6 transition-type competing risks model for many subjects among the 4 SDB groups (see Fig. 3 for variable descriptions), the following could be coded in SAS and R:

```
/*In SAS*/
PROC PHREG data = SAfull COVS(AGGREGATE);
CLASS race smokstatus group shift / REF=FIRST PARAM=REFERENCE;
MODEL tte*obs(0) =
  group*shift
  age sex race smokstatus / covb ties =efron RL=WALD;
STRATA shift;
ID pptid;
RUN;

## In R:
coxph(Surv(tte, obs) ~ group:type +
  + age + sex + I(factor(race)) + smokstatus
  + strata(shift)
  + cluster(pptid),
  data=d.5598.36)
## See supplementary materials for data and executable code examples.
```

Ultimately, the proportional hazards assumption is vital and the group effect and other covariates assumed to follow the proportional hazards property can be tested for each of the stratified hazards. A violation of this assumption can be tested by testing for inclusion of the covariate interacted with $\log(t)$ or testing if the corresponding Schoenfeld residuals and $\log(t)$ have a correlation of zero. If the test is significant, then inclusion of the $\log(t)$ interaction in the model is the remedy for the violation. In addition to the stratified model considering all transitions types, an analysis of 20 separate survival models of one transition-type each will be considered.

2.2. Poisson regression

The mean of a Poisson process can be modeled log-linearly with a log-offset of the total time at risk (tar),

$$\log \lambda^{(gh)} = g + h + g:h + \text{offset} \{ \log(tar_h) \} + x.$$

The quantity $\lambda^{(gh)}$ is the rate for group g and type h and exponentiating linear combinations of coefficients will yield relative *rates* of the overall counts between group g and the reference group for transition-type h . The following could be coded in SAS and R:

```
/*In SAS*/
PROC GENMOD data = LLfull ;
CLASS race smokstatus type group pptid /
REF=FIRST PARAM=REFERENCE ;
MODEL counts = type group type*group age race sex smokstatus /
D=POISSON LINK=LOG OFFSET=logtar;
REPEATED SUBJECT=pptid /
WITHINSUBJECT=TYPE CORR=IND SORTED;
RUN;

## In R:
geeglm(counts~ shift*grouplabel + offset(I(log(tar))) +
I(factor(race)) + sex + I(smokstatus) + age,
id=pptid,
data=d.5598,
family="Poisson",
corstr="independence",
scale.fix=TRUE,
wave=type)
## See supplementary materials for data and executable code examples.
```

Noting that each individual has the same number of rows and that each (ordered) row within individual corresponds to the same attribute (transition-type), these log-linear models are well-suited for fitting with GEE. A hallmark of the GEE is the sandwich estimator, which is also used in the multi-state survival model. The method of GEE modeling is widely used,

computationally fast, and can potentially model correlation structure. The $H \times H$ correlation matrix conveys the correlation of the time-adjusted frequencies of a transition-type occurring. Intuitively, negative correlations could be expected due to the competing risks nature of transition-types sharing the same starting state (e.g. transition-type 1W is negatively correlated with 12 because the greater instances of 1W implies fewer instances of 12); whereas positive correlation could be anticipated for transition-types that share the same state as their ending and starting state (e.g. 1W could be positively correlated with W2 because the higher counts of 1W means more occasions of entering wake, from which W2 can occur). Common parametric structures (“exchangeable” or “AR-1”) do not admit both negative and positive correlations, and the unstructured specification is computationally difficult. The realization that correlation is a nuisance in the GEE framework and that the analytic goals are point estimates and confidence intervals gives motivation to initially take a working independent structure. Consistent estimates are produced regardless of correlation structure, and the possibility of bootstrapping subjects to correct confidence intervals is explored. Further discussion on this approach is found in Section 5.

2.3. Notes on parameterization

The transition rates for each group are of interest and are estimated as relative transition-type specific transition rates among the groups. Such estimation necessitates the inclusion of interactions of binary indicator variables for the transition-types and non-reference groups. In each of the previous two models discussed, main effects play a different role. In the multi-state survival model (stratified on transition-type), the interaction between group and transition-type essentially becomes the group indicator in that stratum, rendering the inclusion of the group main effect unnecessary, as well of course the main effect for transition-type. The baseline hazard acts as the referent group. However, in the Poisson regression, the model is not stratified and thus the main effects of group and transition-type are included to provide backing for the interaction terms. For instance, if the main-effects were omitted in the log-linear analysis, the effect for the design variable for $g : h$ would be the transition rate for group g of transition-type h compared to all other transition-types $\{1, \dots, h-1, h+1, \dots, H\}$ for the reference group. That is, without stratification to restrict the comparison, the interaction effect alone is not fully transition-type specific.

3. Application to the SHHS data

The analysis concerns modeling the association between sleep structure and sleep-disordered breathing (SDB) severity. SDB is a condition where the throat fully or partially collapses during sleep, causing a desaturation in blood oxygen levels. The desaturation is corrected through a sympathetic nervous system response dubbed “arousal” to re-open the airway. The arousal affects many systems, including the cardiovascular system, thus the naming of the Sleep Heart Health Study. SDB typically is categorized into 4 levels of severity: Absent, Mild, Moderate, and Severe corresponding to ranges of the respiratory disturbance index at 4% oxygen desaturation (rdi4p) of average events/hr: [0,5), [5,15), [15,30), [30, ∞). The SDB-Absent group will serve as the reference group.

In the Sleep Heart Health Study, 6369 subjects had polysomnograms, and 5639 had polysomnograms of high enough quality to be reliably processed into 5-state hypnograms. Of the 5639, 5598 had complete demographic and covariate information (age, race, rdi4p, sex, and smoking status). For the 5-state resolution, all 20 possible transition-types were formulated into a Poisson regression format dataset (111,960 rows) and a multi-state survival model format (2,716,188 rows). For the 3-state resolution, all 6 possible transition-types were formulated into a Poisson regression format (33,588 rows) and a multi-state survival model format (728,966 rows).

Fig. 3 depicts three hypnogram resolutions for one subject's 10-epoch (5 min) portion of sleep, visualized with spaghetti plots. For populations of sleep hypnograms, visualizing several hypnogram trajectories in a spaghetti plot is prone to over-plotting. A lasagna plot, by contrast, is a heat map of a matrix where element S_{ij} is the state occupied by the i th subject at the j th epoch (Swihart et al., 2010). Therefore, a lasagna plot is a heatmap that displays clustered longitudinal data, with clusters in the rows and time in the columns and eliminates the overlapping of trajectories that plagues spaghetti plots. In addition, lasagna plots are capable of dynamic sorting. Fig. 4 displays three lasagna plots for each of three different state resolutions for 5598 subjects over 1218 epochs (10 h, 9 min). The top panel for a given resolution is unsorted with respect to subjects. The middle panel shows the same lasagna plot where subjects are organized into the four SDB groups (in descending order of severity, for ease of interpretation) and within SDB group by total sleep time. The bottom panel is a within-column within-SDB-group sorting of the lasagna plot in the middle panel, which shows group-level temporal behavior. Note, as the legends of Fig. 4 collapse (from left to right) how much information is lost: Stage Slow-wave has well-defined peaks that alternate with REM across disease severity, and the prevalence of each group being in Stage 1 in the first epoch of sleep onset is decidedly over 50%, decreasing drastically and then stabilizing over the night.

Another way to look at the sleep data is through empirical transition probabilities from one epoch to the next averaged every 10 min for 8 h (Fig. 5). Fig. 5(A) and (C) shows that the transition probabilities over the night look very similar across SDB severities. The plots on the diagonal in Fig. 5(B) compared to those off diagonal indicate that staying in a state once in a state are the higher probability occurrences. Also note that some transition-types occur rarely, such as 1S, RS, and S1 (Fig. 5(D)).

A two-model analysis is conducted each on 3-state and 5-state resolution data. The two models are competitors, in a sense: the (stratified, recurrent event, competing risks) multi-state survival model stands to honor the sleep process better, but may be more computationally intensive than the Poisson regression via GEE. The application will show the Poisson regression gives computationally faster yet similar results. The two resolutions are competitors as well. If the 3-state analogues of the 5-state estimates reflect direction, magnitude, and significance of results then consideration could be given to using the 3-state resolution of sleep, as collapsing states did not obscure finer-level effects. In addition to comparing the analogues, the 5-state resolution transition-types with no 3-state analogue must also be analyzed for effects when deciding to use exclusively a lower resolution.

For each transition-type, the 95% confidence intervals and point estimates of the relative rate ratios (RR) for the Poisson regression and hazard ratios (HR) for the multi-state survival model can be clustered as three vertical lines, left-to-right increasing in terms of SDB severity. The three lines represent the three comparisons of the risk of transition for increasingly severe SDB versus the referent SDB-absent group. Those clusters of three lines by transition-type can then be organized in plots with other clusters to give view of the modeled relationship of SDB on sleep itself.

4. Results

As given by the two models, the SDB dose response clusters are visualized in Fig. 6. The first transition-type is 1W on the horizontal axis in the top left panel in Fig. 6. Three confidence intervals for the relative rate ratios of 1W are displayed (in red, from left to right): mild SDB to absent, moderate SDB to absent, and severe SDB to absent. As the relative SDB severity increases, so does the point estimate, indicating that 1W happens at faster rates for more severe levels of SDB. A similar pattern is displayed in green for 2W in the adjacent spot on the horizontal axis. The next two clusters are SW and RW, making the first four transition-types in the plot all the types that enter Wake. The remaining four spots on the horizontal axis are the transition-types that *exit* Wake. The plot below the Wake entering–exiting plot is for Stage 1, and below that is the Stage 2 entering–exiting plot (left middle panel). The four transition-types that exit Stage 2 are in green and show an interesting divergence (left middle panel): with increasing SDB severity, transition rates exiting to wake and the lightest stage of sleep (2W and 21) increase and transition rates that exit into the deepest sleep and REM (2S and 2R) decrease. Both deep sleep and REM sleep are thought to be very important in physical and mental recuperative processes. This analysis indicates that as SDB severity increases, waking up from Stage 1 or 2 and going from sleep to a lighter sleep (Stage 2 to Stage 1) happen at increasingly quicker rates while the rates of getting deep sleep and REM from the most common Stage 2 decrease. These findings imply that SDB severity affects sleep structure in a way that favors light sleep and waking stages at the expense of REM and deep sleep.

Compare the top left panel with the top right panel of Fig. 6 to compare the results of a Poisson regression to a multi-state survival model. All of the estimates, trends, and inferences are similar. Each pair of plots in each row can be compared for how the two modeling approaches depict the association of SDB severity and sleep structure. Overall the Poisson regression seems to serve as a good approximation. This might indicate that the survival times are approximately exponentially distributed.

In considering the two resolutions, a 3-state resolution would completely omit the relationship seen for intra-NREM transitions of type 21 and 2S, because no transition would be recorded—just contiguous time in NREM. A 5-state analysis deserves serious consideration given the strong trend of intra-NREM transitions and increasing SDB severity (significantly more type 21 and fewer type 2S transitions) would be obscured with a 3-state hypnogram analysis.

To further explore the connection between resolutions for 3-state transition-types that combine 5-state transition-types, analogue plots can be made. Fig. 7 shows the 1 member set to 3 member set mappings for the 3-state and 5-state models. In the top left panel, we see that for most severe SDB groups, there is a discordance of significant findings in 5-state sleep (1R and SR are significantly higher rates in severe SDB relative to SDB-absent; 2R is significantly lower) and they “cancel” out in 3-state sleep as seen by the NR insignificance across SDB severities. As for the other plots, we see similar shapes between the resolutions, possibly indicating that 1W and 2W are drivers of NW; R1 of RN; and W1 and W2 of NW.

5. Modeling insights

To fit the multi-state survival model, prudent software and hardware choices are suggested. Regardless of OS platform, 64-bit SAS and R are recommended. If encountered, remedying a violation of proportional hazards with $\log(t)$ interactions requires R version 2.13 or later of `coxph()` for its implementation of `tt()`. Running 64-bit SAS in Windows utilizing a 3.40 GHz quad-core processor with 16 GB of RAM, the 5-state resolution Poisson regression with independent working correlation structure of the previous section took 13 s compared to 8.5 h for the multi-state survival model (13 h for $\log(t)$ interactions proportional hazards correction). The long computational times for the survival models are due to the required accounting for recurrent events within subject for a considerable number of subjects with SAS option `COVS(aggregate)` and R option `cluster(id)`: eliminating these options reduces computational times to mere minutes. Given the Poisson regression gave similar results, the Poisson regression can be used quickly and repeatedly as an exploratory tool for an investigation, and when a final model is suspected then fit the corresponding survival model or bootstrap subjects for corrected Poisson regression intervals (bootstrapping 1000 times was on par computationally with fitting the survival model) (Sherman and le Cessie, 1997).

Modeling of the correlation can be attempted, but proves challenging. The unstructured working correlation structure has many parameters and the specifiable common structures struggle to reflect the competing risks nature of the process. If estimating the unstructured correlation specification is prohibitive, a sample correlation matrix (or, an appropriately found nearest positive definite matrix to that data-based calculation) can be “user-specified”.

Producing the RRs of Poisson regression requires linear combinations of the group main effect and interaction terms (see Online Supporting Information). The multi-state survival model parameterization, due to no main effects and only interaction variables, requires no linear combinations to produce the HRs. Handling multiple groups in the group-defining condition and several transition-types when comparing resolutions takes organizational care. We advocate keeping rows of different resolutions analogous to one another for ease of comparison among the resolutions as well as making entering–exiting state plots for learning the “story” of the data analysis. While PROC PHREG does not require that interaction variables be manually coded, doing so may be beneficial for organizing the results and output.

6. Discussion

One could consider parallelizing the multi-state survival model into 20 separate Cox regressions modeling only the recurrent events of one particular type of transition and no competing risk information. The point estimates (not shown) are similar to those displayed in Fig. 6. While there is a time savings (the quickest ran at 3 min, the longest in 75) relative to the multi-state survival model, the time savings does not best the Poisson regression.

The Fine-Gray competing risks regression model for cumulative incidence may be of use for smaller datasets (Fine and Gray, 1999; Zhou et al., 2011, 2012). The function `crrSC:::crrs()` in R was executed on the 5598 subjects and ran for 21 days on a dedicated node before terminating. The current code implementation on CRAN allows for stratified or recurrent event analyses, not both (Zhou and Latouche, 2013). A reviewer suggested computing cumulative incidence functions for each starting state might be computationally feasible. See Latouche et al. (2013) for a case study for analyzing hazards and cumulative incidence side-by-side.

The equivalence between a Poisson regression via a log-linear GLMM with a log (time at risk) offset and parametric multi-state survival modeling assuming exponential survival times and piecewise constant hazards is well known. Thus, middle ground exists between the Poisson regression via GEE and the multi-state survival models put forth, but implementation is not as straightforward, requiring further data manipulation and customized programming code (Swihart et al., in press).

The methods put forth stand to aid the investigation of sleep itself with sleep-related and non-sleep-related health outcomes. In the application we analyzed SDB predicting changes in sleep stage structure using fast, scalable commonplace routines to enable analyses of other datasets of larger sizes by those with basic computing proficiency and resources. Future work would be to continue down hypothesized causal pathways and connect the transition-type-specific count, time at risk, and rate features of sleep and predict a non-sleep related outcome, say, heart rate variability. Another direction of research would be to account for the longitudinal aspects of SHHS, as the sleep-EEG feature extraction work has (Crainiceanu et al., 2009). Doing so may ultimately provide better diagnostic tools and further understanding of how sleep interacts with health.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The work of Swihart and Crainiceanu was supported by Award Number R01NS060910 from the National Institute Of Neurological Disorders And Stroke. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute Of Neurological Disorders And Stroke or the National Institutes of Health.

References

- Aalen O, Fosen J, Weedon-Fekjær H, Borgan Ø, Husebye E. Dynamic analysis of multivariate failure time data. *Biometrics*. 2004; 60(3):764–773. [PubMed: 15339300]
- Crainiceanu C, Caffo B, Di C, Punjabi N. Nonparametric signal extraction and measurement error in the analysis of electroencephalographic activity during sleep. *J. Amer. Statist. Assoc.* 2009; 104(486):541–555.
- Fahrmeir L, Klinger A. A nonparametric multiplicative hazard model for event history analysis. *Biometrika*. 1998; 85(3):581.
- Fine JP, Gray RJ. A proportional hazards model for the subdistribution of a competing risk. *J. Amer. Statist. Assoc.* 1999; 94(446):496–509.
- Holford T. Life tables with concomitant information. *Biometrics*. 1976; 32(3):587–597. [PubMed: 963172]
- Holford T. The analysis of rates and of survivorship using log-linear models. *Biometrics*. 1980; 36(2): 299–305. [PubMed: 7407317]
- Kalus S, Kneib T, Steiger A, Holsboer F, Yassouridis A. A new strategy to analyze possible association structures between dynamic nocturnal hormone activities and sleep alterations in humans. *Am. J. Physiol.-Regul. Integr. Comp. Physiol.* 2009; 296(4):R1216–R1227. [PubMed: 19144755]
- Kneib T, Hennerfeind A. Bayesian semi parametric multi-state models. *Stat. Model.* 2008; 8(2):169.
- Laird N, Olivier D. Covariance analysis of censored survival data using log-linear analysis techniques. *J. Amer. Statist. Assoc.* 1981; 76(374):231–240.
- Latouche A, Allignol A, Beyersmann J, Labopin M, Fine JP. A competing risks analysis should report results on all cause-specific hazards and cumulative incidence functions. *J. Clin. Epidemiol.* 2013; 66(6):648–653. [PubMed: 23415868]
- Norman R, Scott M, Ayappa I, Walsleben J, Rapoport D. Sleep continuity measured by survival curve analysis. *Sleep*. 2006; 29(12):1625–1631. [PubMed: 17252894]
- Quan S, Howard T, Iber C, Kiley J, Nieto F, O'Connor G, Rapoport D, Redline S, Robbins J, Samet J, et al. The sleep heart health study: design, rationale, and methods. *Sleep (New York, NY)*. 1997; 20(12):1077–1085.
- Sherman M, le Cessie S. A comparison between bootstrap methods and generalized estimating equations for correlated outcomes in generalized linear models. *Comm. Statist. Simulation Comput.* 1997; 26(3):901–925.
- Swihart B, Caffo B, Bandeen-Roche K, Punjabi N. Characterizing sleep structure using the hypnogram. *J. Clin. Sleep Med.: JCSM: Off. Publ. Am. Acad. Sleep Med.* 2008; 4(4):349–355.
- Swihart B, Caffo B, Crainiceanu C, Punjabi N. Mixed effect Poisson log-linear models for clinical and epidemiological sleep hypnogram data. *Stat. Med.* 2012; 31(9) in press.
- Swihart B, Caffo B, James B, Strand M, Schwartz B, Punjabi N. Lasagna plots: a saucy alternative to spaghetti plots. *Epidemiology (Cambridge, Mass.)*. 2010; 21(5):621–625.
- Therneau, T.; Grambsch, P. *Modeling Survival Data: Extending the Cox Model*. Springer; 2000.
- Yassouridis A, Steiger A, Klinger A, Fahrmeir L. Modelling and exploring human sleep with event history analysis. *J. Sleep Res.* 1999; 8(1):25–36. [PubMed: 10188133]
- Zhou B, Fine J, Latouche A, Labopin M. Competing risks regression for clustered data. *Biostatistics*. 2012; 13(3):371–383. [PubMed: 22045910]
- Zhou, B.; Latouche, A. *crrSC: competing risks regression for Stratified and Clustered data.. R package version 1.1*. 2013. URL: <http://CRAN.R-project.org/package=crrSC>
- Zhou B, Latouche A, Rocha V, Fine J. Competing risks regression for stratified data. *Biometrics*. 2011; 67(2):661–670. [PubMed: 21155744]

Fig. 1. The first hour of sleep for three individuals, visualized by discrete-time discrete-state spaghetti plots known as sleep hypnograms. The vertical axis displays the five stages of sleep: Wake, Rapid-Eye Movement (REM), Stage 1, Stage 2, and Stage Slow-wave labeled “W”, “R”, “1”, “2”, “S”, respectively. The horizontal axis denotes time from sleep onset in 30-s epochs for 60 min. Note: these hypnograms ignore the first transition from Wake.

Fig. 2.

A summarization of statistical methodology literature with sleep applications. Four facets are displayed: Number of Subjects, Number of Groups, Number of States, and Number of Transition-Types. On the vertical axis are the labels for the cited paper, and in each facet they are ordered based on the attribute of the facet. The proposed approach, Swihart 2015, is handling more subjects, more groups, more states and more transition-types. Green points represent a full and flat paradigm, where all $S(S-1)$ transition-types are modeled for the S states considered. This figure appears in color in the electronic version of this article. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 3.

Sleep hypnograms (left) of the same 10 epoch (5 min) sleep trajectory as represented in 5-stage, 3-stage, and 2-stage sleep resolution (horizontal axis in epochs, with the 1st, 5th, and 9th epoch labeled) and accompanying Poisson regression and multi-state survival model data formats (right). Column-name descriptions of the variables: *order* is the order of occurrence of the transitions; *shift* is the transition-type label; *obs* is the binary indicator of whether the transition was observed (1) or censored (0); *tte* is the time to event in minutes; *type* is a numerical variable in 1-to-1 correspondence with *shift*; *count* is the number of times the transition-type *type* was observed during the total time at risk (*tar*) in minutes.

Fig. 4.

Lasagna plots for 5-state, 3-state, and 2-state sleep. Each lasagna plot has 5598 rows (subjects) and 1218 columns (epochs). The top row of lasagna plots displays subjects in no particular order. The second row shows subjects grouped into SDB severity group (absent to severe; top to bottom within the plot) and ranked by total sleep time within severity group. The bottom row of lasagna plots is those of the middle row sorted within columns within severity group, highlighting the group-level temporal dynamics of Stage Slow-wave. The bottom _{left} panel has no bounding box so as to show the Stage 1 (S1) dynamic in the first epochs more clearly. This figure appears in color in the electronic version of this article. (For

interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 5.

Top row: Transition probabilities on the untransformed scale; Bottom panels: transition probabilities on the log-scale; Left column: faceted by SDB severity; Right column: faceted by next state. (C) is the log-transform of (A)—likewise for (D) and (B). This figure appears in color in the electronic version of this article. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 6.

Entering–Exiting plots for 5-state resolution. On the left, Relative Rates as a function of sleep disordered breathing. On the right, Hazard Ratios. Each of the 5 plots in a column is made displaying the 8 transition-types involving the entering and exiting of a state (top to bottom: Wake, Stage 1, Stage 2, Slow-wave, REM). The dots are the point estimates for increasing SDB severity versus the reference SDB-Absent group and the whiskers are the 95% confidence intervals. Comparing plots within rows demonstrates how similar the estimates are between the modeling approaches. This figure appears in color in the electronic version of this article. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 7.

Compare the estimated trends of 3-state resolution (clockwise from upper-left: NR, NW, WN, RN) to the corresponding 5-state analogues ($\{1R, 2R, SR\}$, $\{1W, 2W, SW\}$, $\{W1, W2, WS\}$, $\{R1, R2, RS\}$, respectively). The vertical dashed line separates 3-state (black) from 5-state (colors from Fig. 6) estimates in each panel. This figure appears in color in the electronic version of this article. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)