SLEEP, 2022, 1–16

https://doi.org/10.1093/sleep/zsac189
Advance Access Publication Date: 6 August 2022
Original Article

ORIGINAL ARTICLE

Dynamic models of obstructive sleep apnea provide robust prediction of respiratory event timing and a statistical framework for phenotype exploration

Shuqiang Chen[1,], Susan Redline[2,3], Uri T. Eden[4] and Michael J. Prerau[2,3,*,]

[1]Graduate Program for Neuroscience, Boston University, Boston, MA, USA, [2]Division of Sleep and Circadian Disorders, Brigham and Women's Hospital, Boston, MA, USA, [3]Department of Medicine, Harvard Medical School, Boston, MA, USA and [4]Department of Mathematics and Statistics, Boston University, Boston, MA, USA

*Corresponding author. Michael J. Prerau, Division of Sleep and Circadian Disorders, Brigham and Women's Hospital, 221 Longwood Avenue, Boston, MA 02115, USA. Email: mprerau@bwh.harvard.edu.

Abstract

Obstructive sleep apnea (OSA), in which breathing is reduced or ceased during sleep, affects at least 10% of the population and is associated with numerous comorbidities. Current clinical diagnostic approaches characterize severity and treatment eligibility using the average respiratory event rate over total sleep time (apnea-hypopnea index). This approach, however, does not characterize the time-varying and dynamic properties of respiratory events that can change as a function of body position, sleep stage, and previous respiratory event activity. Here, we develop a statistical model framework based on point process theory that characterizes the relative influences of all these factors on the moment-to-moment rate of event occurrence. Our results provide new insights into the temporal dynamics of respiratory events, suggesting that most adults have a characteristic event pattern that involves a period of normal breathing followed by a period of increased probability of respiratory event occurrence, while significant differences in event patterns are observed among gender, age, and race/ethnicity groups. Statistical goodness-of-fit analysis suggests consistent and substantial improvements in our ability to capture the timing of individual respiratory events using our modeling framework. Overall, we demonstrate a more statistically robust approach to characterizing sleep disordered breathing that can also serve as a basis for identifying future patient-specific respiratory phenotypes, providing an improved pathway towards developing individualized treatments.

Statement of Significance

Obstructive sleep apnea (OSA) is a dynamic process, yet the primary diagnostic metric—the apnea-hypopnea index (AHI)—describes only the average respiratory event rate. To reclaim these lost dynamics, we develop a rigorous statistical approach for estimating an "instantaneous AHI", which models the moment-by-moment event rate as a function of body position, sleep stage, and the timing of past events. This model acts as a highly individualized respiratory fingerprint, which we show can accurately predict the precise timing of future events. We also demonstrate robust model differences in age, sex, and race across a large population. Overall, this approach provides a substantial advancement in OSA characterization for individuals and populations, with the potential for improved patient phenotyping and outcome prediction.

Submitted: 18 May, 2022; **Revised:** 25 July, 2022

Graphical Abstract

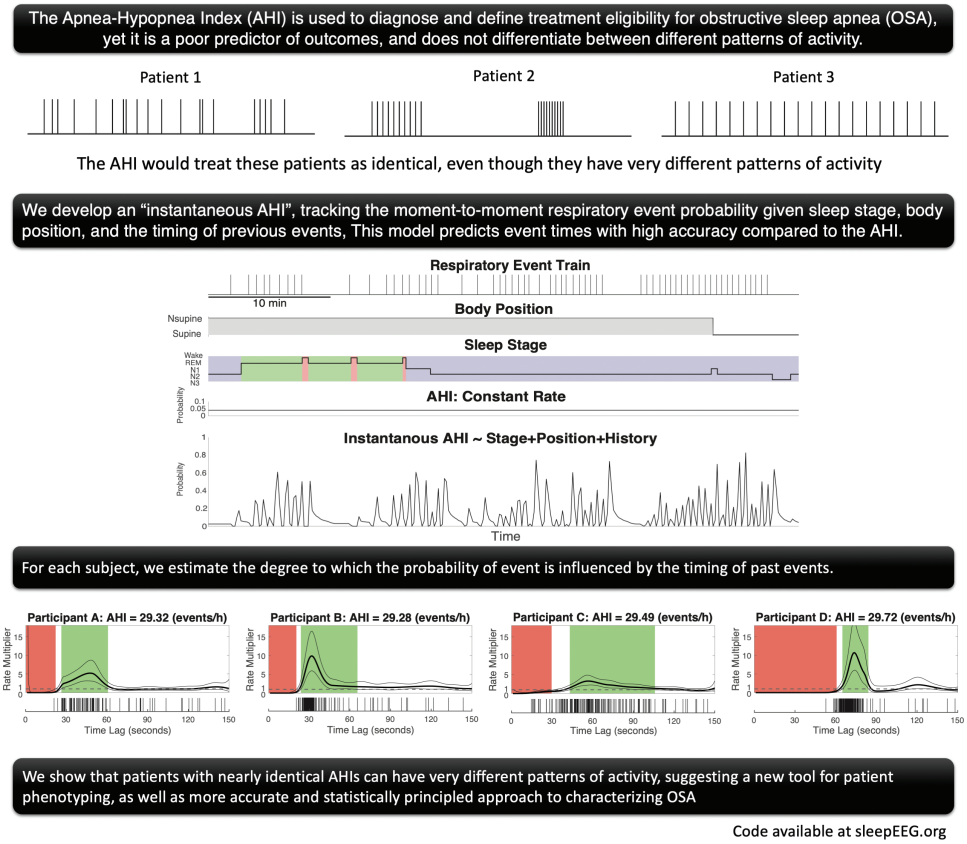

Key words: point processes; sleep apnea; statistical models

Introduction

Obstructive sleep apnea (OSA) is a condition in which there are recurrent periods where breathing ceases or is disrupted during sleep despite continued respiratory effort [1]. OSA affects at least 10% of the population (~30 million people) within the United States, with an increased risk of up to 50% in populations with comorbidities related to cardiovascular disease, obesity, age, and diabetes [2–8]. If left untreated, OSA can increase the risk of numerous health issues including heart failure, stroke, and dementia [4, 9–13]. It is estimated that the economic impact of undiagnosed OSA in the United States is $149.6 billion annually, including $26.2 billion in motor vehicle accidents and $6.5 billion in workplace accidents per year [14] resulting from excessive daytime sleepiness. A major challenge in treating OSA is that it is a complex disorder with variability in its clinical and physiological characteristics with multiple approaches to treatment [15], including surgery, continuous positive airway pressure (CPAP), positional therapy, and oral appliances. Given the highly individualized nature of OSA [16–19], the success of a given treatment and adherence vary widely [20–22], such that patients may need to try several variations of interventions, and may reject treatment altogether before finding a sustainable solution.

Moreover, OSA is a highly dynamic time-varying process, governed by numerous intrinsic and extrinsic factors including sleep architecture, body position, sleep state, sleep stage, time of night effects, and fluid retention [23–27]. The relative influence of these factors is specific to each individual, moderated by that patient's underlying biophysical makeup and environment. Despite these known dynamics, OSA is characterized by a single summary metric, the apnea-hypopnea index (AHI), which is the average rate of respiratory events (apneas plus hypopneas) per hour during sleep [28, 29]. As such, the AHI ignores any temporal patterns in OSA [27]. It is not surprising then that the AHI has been shown to have a high degree of uncertainty and is a poor predictor of clinical outcomes [27, 30–32]. Yet, the AHI is the metric by which patients are diagnosed, treatment eligibility is determined, and federal approval for new devices and treatments is assessed.

Figure 1a illustrates the shortcomings of the AHI, showing a schematic with respiratory event times for three hypothetical patients with same AHI, but with substantially different levels of variability and periodicity of respiratory events. While the temporal patterns differ greatly, these patients are indistinguishable under the AHI and would thus be viewed as clinically similar cases. It is therefore vital to develop metrics of OSA dynamics that reflect the temporal patterns of respiratory events and the factors that influence event rate. In doing so, we can better quantify the mechanisms that influence a patent's respiratory event rate, which could greatly inform treatment decisions and improve outcomes.

Though stochastic modeling of event timing dynamics has had a long history of adoption through numerous fields (including seismology, finance, neuroscience, climatology, sleep architecture, etc.) [33–39], it has not taken hold within clinical sleep medicine for the characterization of OSA events. This is due in part to the fact that standards for OSA diagnosis were

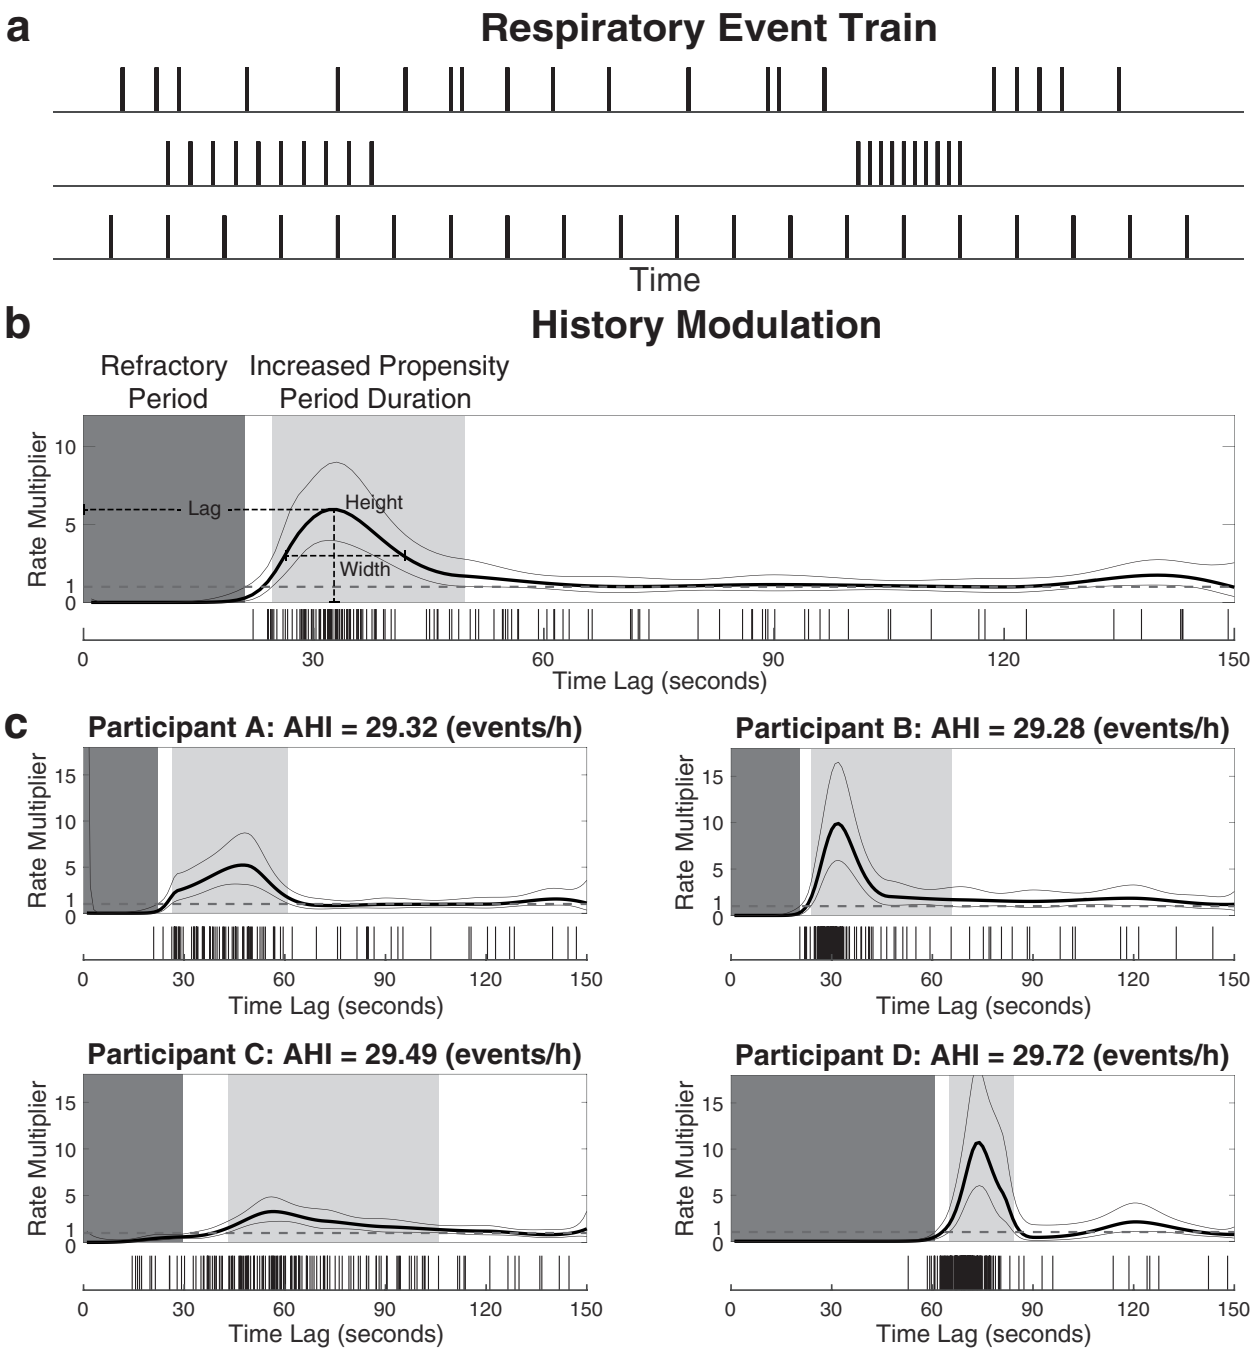

	Refractory Period	Increased Propensity Period Duration	Height	Lag	Width
Participant A	24s	35s	5.22	47s	23s
Participant B	22s	41s	9.90	32s	11s
Participant C	32s	65s	3.30	57s	29s
Participant D	61s	19s	10.71	74s	11s

Figure 1. Differences in history dependence structure provide a basis for characterizing respiratory event phenotypes. (a) A schematic plot shows three patients with identical AHI values even though they exhibit potentially different respiratory event phenotypes. (b) Modeling the history dependence of respiratory events. The history modulation plot (black curve) shows the multiplicative effect on event rate based on the time since the last observed event, with 95% confidence intervals. The aligned inter-event intervals (vertical lines, bottom) reflect the structure of the history modulation curve fit by the PSH model. We can summarize the history modulation curve by defining refractory period (dark gray), increased propensity period duration (light gray), increased propensity peak height, lag, and width. (c) We show another four participants with similar AHIs that possess different history dependence architectures, summary statistics of the history modulation curves for these participants are shown in the table. By characterizing history dependence in terms of interpretable statistics, we can quantify individual differences in respiratory patterns.

Table 1. The clinical report is a useful tool for characterizing the effects of simple conditions on respiratory event occurrence, while parametric models of the respiratory events can provide a statistically principled description of the event dynamics

a Clinical sleep report

	AHI (events/h)	Time (min)	Event count
All sleep	20.44	416.78	142
Supine	54.52	114.45	104
Nonsupine	7.54	302.33	38
REM	53.61	73.87	66
Supine	76.75	41.43	53
Nonsupine	24.05	32.43	13
NREM	13.30	342.92	76
Supine	41.91	73.02	51
Nonsupine	5.56	269.90	25
N1	47.56	25.23	20
Supine	61.59	10.72	11
Nonsupine	37.20	14.52	9
N2	15.68	206.68	54
Supine	51.66	45.30	39
Nonsupine	5.58	161.38	15
N3	1.08	111.00	2
Supine	3.53	17.00	1
Nonsupine	0.64	94.00	1

b Model output

Sleep stage (nonsupine)	Rate (events/h)	95% CI
REM	17.28	[13.49, 22.13]
NREM	7.96	[6.06, 10.46]
N1	18.35	[11.74, 28.68]
N2	8.61	[6.28, 11.79]
N3	0.69	[0.17, 2.79]
Significant REM dominance (χ^2 test, $p < .0001$)		
Body position	Multiplier	95% CI
Supine	4.75	[3.21, 7.02]
Significant supine dominance (χ^2 test, $p < .0001$)		

Panel **a** displays a clinical report that computes the AHI under different conditions, revealing that this experimental participant has supine and REM dominant sleep disordered breathing; panel **b** shows the results of the PS model with the corresponding 95% confidence intervals for the same participant in a, which shows significant REM and supine dominance.

initially developed based on what physicians could easily calculate in the days prior to the incorporation of computers into the clinic [40]. By developing dynamic modeling approaches for OSA, we can provide a much clearer picture of respiratory event dynamics during sleep, leading to a better understanding of the factors and mechanisms underlying patient-to-patient variability, thereby improving clinical assessment of patients and enabling individualized treatment optimization.

In this paper, we aim to improve on the AHI by developing a framework that can describe temporal patterns and capture the moment-by-moment influences of multiple factors on the instantaneous respiratory event rate. To do so, we will build regression models of respiratory event dynamics using a point process framework. Point processes are mathematical models that describe discrete events that are localized in space or time, such as neural spikes, earthquakes, or motor vehicle accidents. Point processes have had long-standing use across a variety of diverse fields, including neuroscience [41], physics [42], finance [43], and earth sciences [44], but have yet, to our knowledge, been applied to explicitly model event dynamics during OSA.

Perhaps the most well-known class of point process is the Poisson process, for which all the events are independent and can be modeled with a single rate parameter. Given that the AHI is a single rate that does not incorporate temporal information, computing it is mathematically equivalent to fitting a Poisson process model to the data. Consequently, our previous work has shown that a Poisson process framework can provide tools for performing statistical tests, computing confidence intervals, and measuring goodness-of-fit, and can be applied to existing AHI methods [27]. A general point process framework is therefore a logical extension of this work, as it is a natural statistical approach for developing models that capture moment-by-moment influences of physiological variables on the rate of respiratory events.

Here, we define a general point process framework for modeling respiratory events that extends the currently used clinical application of AHI to allow for the influence of many factors, including past events, using well-studied tools for model fitting and goodness-of-fit assessment. We apply these methods to sleep data across a large population of adult participants to obtain patient-specific profiles of respiratory event dynamics and the factors influencing them. We further show that, given the resulting model fit for an individual, the timing of individual respiratory events becomes predictable to a large degree within a single night. The goal of this study is to provide a strong methodological foundation for this framework, as well as to provide a strong proof-of-concept for applications to large scale datasets. In doing so, we demonstrate the potential of patient-specific respiratory event patterns as tool for future development of respiratory phenotypes, with the long-term aim of predicting clinical outcomes and informing treatment and clinical decision-making.

Methods

Study overview

The goal of this study is to capture the patterns and influences governing respiratory event activity and thus exploit dynamics (otherwise ignored by the AHI) to provide a more principled basis for characterizing OSA. To do so, we apply a point process framework using a generalized linear model (GLM) [45, 46] which quantifies the relationship between features of sleep (e.g., body position, stage, etc.) and the probability of seeing a respiratory event at a given moment. A general point process can be expressed in terms of a *conditional intensity function,* $\lambda(t \mid H_t)$, which is the "instantaneous rate" at time t, given H_t, the history of past events up to, but not including, time t. The conditional intensity can then be expressed as a function of the variables that influence these events. In this study, we express the conditional intensity of respiratory events as a function of body position (supine vs non-supine) and sleep stage (rapid eye movement [REM], non-REM stages 1–3 [N1–N3]), and previously occurring events. In doing so, we provide a rigorous statistical framework to determine which factors influence the moment-by-moment rate of respiratory events for an individual during OSA, and the relative magnitude of these influences, which provide a basis for phenotyping. See Appendix A and B for the overview of point process modeling and GLM.

We compare three models of respiratory event rate: (1) an "AHI model" with a constant rate, which reflects the current diagnostic standard, (2) a "Position-Stage model" (PS model) in which rate changes dynamically with body position and sleep stage, and (3) a "Position-Stage-History model" (PSH model) in which past events provide an additional influence on rate dynamics (see Appendix C for model specification). We fit these models to scored respiratory event time data from 936 participants from a community based cohort, the Multi-Ethnic Study of Atherosclerosis (MESA) dataset, available from the National Sleep Research Resource [47, 48]. Model performance was assessed, and the influence of each factor was tested for statistical significance. We characterize the results at the level of single participants, at the level of the entire MESA participants, as well at the level of different population groups defined by gender, age, and race/ethnicity.

Data description

To examine the ability of the models to capture OSA temporal dynamics across a large, heterogeneous population of participants, we examine polysomnography data from the Multi-Ethnic Study of Atherosclerosis (MESA), obtained from the National Sleep Research Resource (www.sleepdata.org) [47, 48], which includes manually-scored respiratory events, sleep staging, and body position for a single night of polysomnography. For each individual, we computed the AHI using all apneas and hypopneas associated with a 3% oxygen desaturation. We limited our analyses to participants with full night AHI values above 15, totaling 936 participants (513/423 M/F, age: mean 69.57 ± 8.89, total in-bed time (minutes): mean 482.52 ± 81.26, AHI (events/h): mean 33.40 ± 15.88. The distributions of AHI and total in-bed time are shown in Sup. Figure 1.

Respiratory event termination times were used as the timing for point process events, given event start times are more difficult to precisely define. Event times were then discretized into 1-second intervals and the time series the number of respiratory events (0 or 1) terminating in each of those intervals was computed. The sleep positions were labeled as Supine (laying on the back) and Nonsupine (front or side sleeping). Note, separation of Nonsupine into Front, Left and Right positions, yielded similar results in terms of contribution to the model (Sup. Figure 2). We therefore used Supine and Nonsupine for clinical interest and model simplicity. The sleep stages were labeled as Wake, rapid eye-movement sleep (REM), and non-REM stages 1–3 (N1, N2, and N3).

It should be noted that the original MESA study coded "race" as either "White, Caucasian", "Chinese American", "Black, African-American", or "Hispanic", which are overly reductive characterizations. Herein, we replace "Chinese American" with "Asian", to reflect more accurately the Asian, Pacific Islander, and other participants within the underlying population represented under this category.

Statistical tests

Chi-square tests were performed to compute position and stage dominance (Table 1), Kolmogorov-Smirnov (KS) tests were used to evaluate the goodness-of-fit of the models (Figure 2b), as well as to evaluate the difference of the distributions (Sup. Figure 2). In Figure 3a, permutation tests with global bounds were performed to compare history curves across all groups [49], and *t*-tests were conducted to compare the rates (∗, ∗∗, ∗∗∗ denote the *p*-value <.05, <.01, <.001 separately). Chi-square tests were also used to compare deviance reduction for different models (Figure 4b). All statistical analyses were performed in MATLAB_R2021a. Significance levels of 0.05 were used, if not otherwise specified.

Code toolbox

We have created a MATLAB code toolbox for analyzing sleep apnea dynamics, which is publicly available at http://sleepEEG.org.

Results

Characterizing respiratory event dynamics in individual participants

OSA is predominantly mediated by biophysical obstructions to the airway, thus factors such as body position and sleep stage are known to influence the prevalence of respiratory events [50, 51]. By identifying the conditions under which events are most likely to occur for a specific patient, clinicians may be guided in the selection of an appropriate individualized treatment. For example, if a patient has supine dominant events, a clinician may initially suggest a positional therapy that keeps them off their back during sleep. In practice, clinicians are aided in this analysis by reports of condition-dependent AHIs. Table 1a illustrates a typical section from a clinical sleep report, using an example participant from the MESA dataset. The overall AHI across the total sleep time (All Sleep) of 20.44 indicates moderate–severe apnea, which would qualify the participant for treatment. The conditional AHIs show that respiratory events are supine and REM dominant, with an AHI of 76.75 during REM while supine.

If we instead fit the data using the *Position-Stage* (PS) model, which describes event rate dynamically in terms of body position and sleep stage, we can expand upon the clinical report and provide a compact, statistically principled characterization of the effect of condition on the underlying respiratory event rate, as well as the uncertainty about the effect size. Table 1b shows the PS model parameters expressed as rates for each stage during nonsupine sleep, along with a multiplier that is applied when the participant is supine. Given our past work demonstrating the extreme variability associated with AHI estimates [27], we also report the 95% confidence interval for each parameter and explicitly test for stage or position sensitivity. At each point in time, the fitted model along with the current stage and position defines the time-varying AHI. For example, to compute the supine REM event rate, we multiply the REM nonsupine rate (17.28 events/h) by the supine multiplier (4.75) to get $17.28 \times 4.75 = 82.08$ (events/h).

We observed a marked difference in event rate as a function of position (475% increase in supine vs nonsupine), and stage (217% increase in REM vs NREM). Given the clinical importance of stage and positional dominance in characterizing and treating OSA, we performed hypothesis tests to assess the statistical significance of these influences. In this participant, we find significant REM (χ^2 test, $p < .0001$) and supine (χ^2 test, $p < .0001$) dominance. This is also evidenced by the fact that the REM and NREM 95% confidence intervals are non-overlapping, and that the 95% confidence interval for the supine multiplier does not overlap with 1. Overall, this modeling approach provides a highly compact and informative representation of the

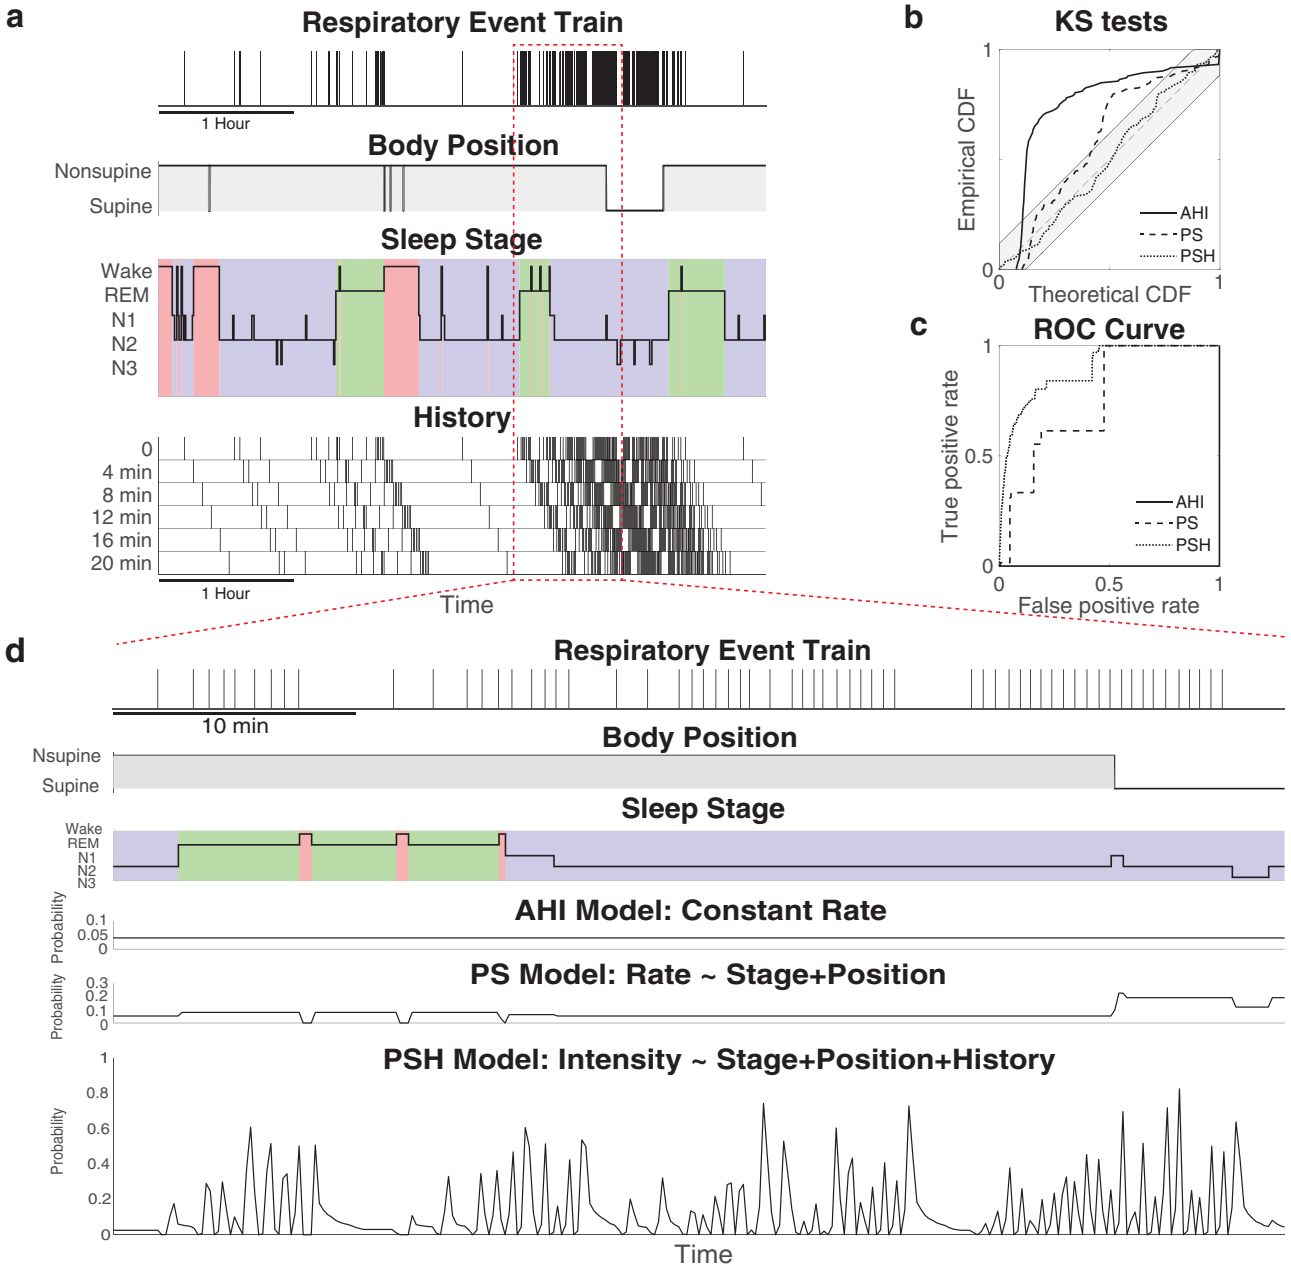

Figure 2. The addition of history dependence greatly improves the goodness-of-fit over AHI and the predictability of individual respiratory event times. (a) Modeling components and goodness-of-fit for an experimental participant. From top to bottom, respiratory event train, sleep position, hypnogram, and history. (b) The KS plot shows goodness-of-fit for event timing given the model estimates. We show the KS plots for the following models: model with constant rate (AHI, solid), model with stage and position (PS, dashed), model with stage, position, and history dependence (PSH, dotted). (c) The ROC curve for the AHI model (solid), the PS model (dashed) and the PSH model (dotted), which computed based on the probability of seeing events in 10-second intervals from (d). (d) In a short time segment (~50 min), the probability of observing an event in 10-second intervals is shown for AHI, PS, and PSH. The PSH model uses temporal information from the history dependence structure to provide much stronger predictions for individual events. Note the order of magnitude differences in the scales of probability on the y-axes. Overall, these results show model improvement with stage and position, however, the addition of history dependence provides the greatest amount of information and increases the level of predictability.

data traditionally contained in clinical reports while providing simple measures of statistical uncertainty.

History dependence as the basis for respiratory event phenotyping

While the PS model describes respiratory event rate as a function of discrete brain or physiological states, it is unable to capture temporal patterns of events that are present within each state. To better capture these patterns, we model the inherent temporal dependence structure of respiratory events by adding *history dependence*, which describes the effect of previous event timing on the current rate. We fit the *Position-Stage-History* (PSH) model to the data, which estimates a multiplicative modulation of the event rate due to a prior event at any given time lag. This effect is visualized with a *history modulation plot*, which shows

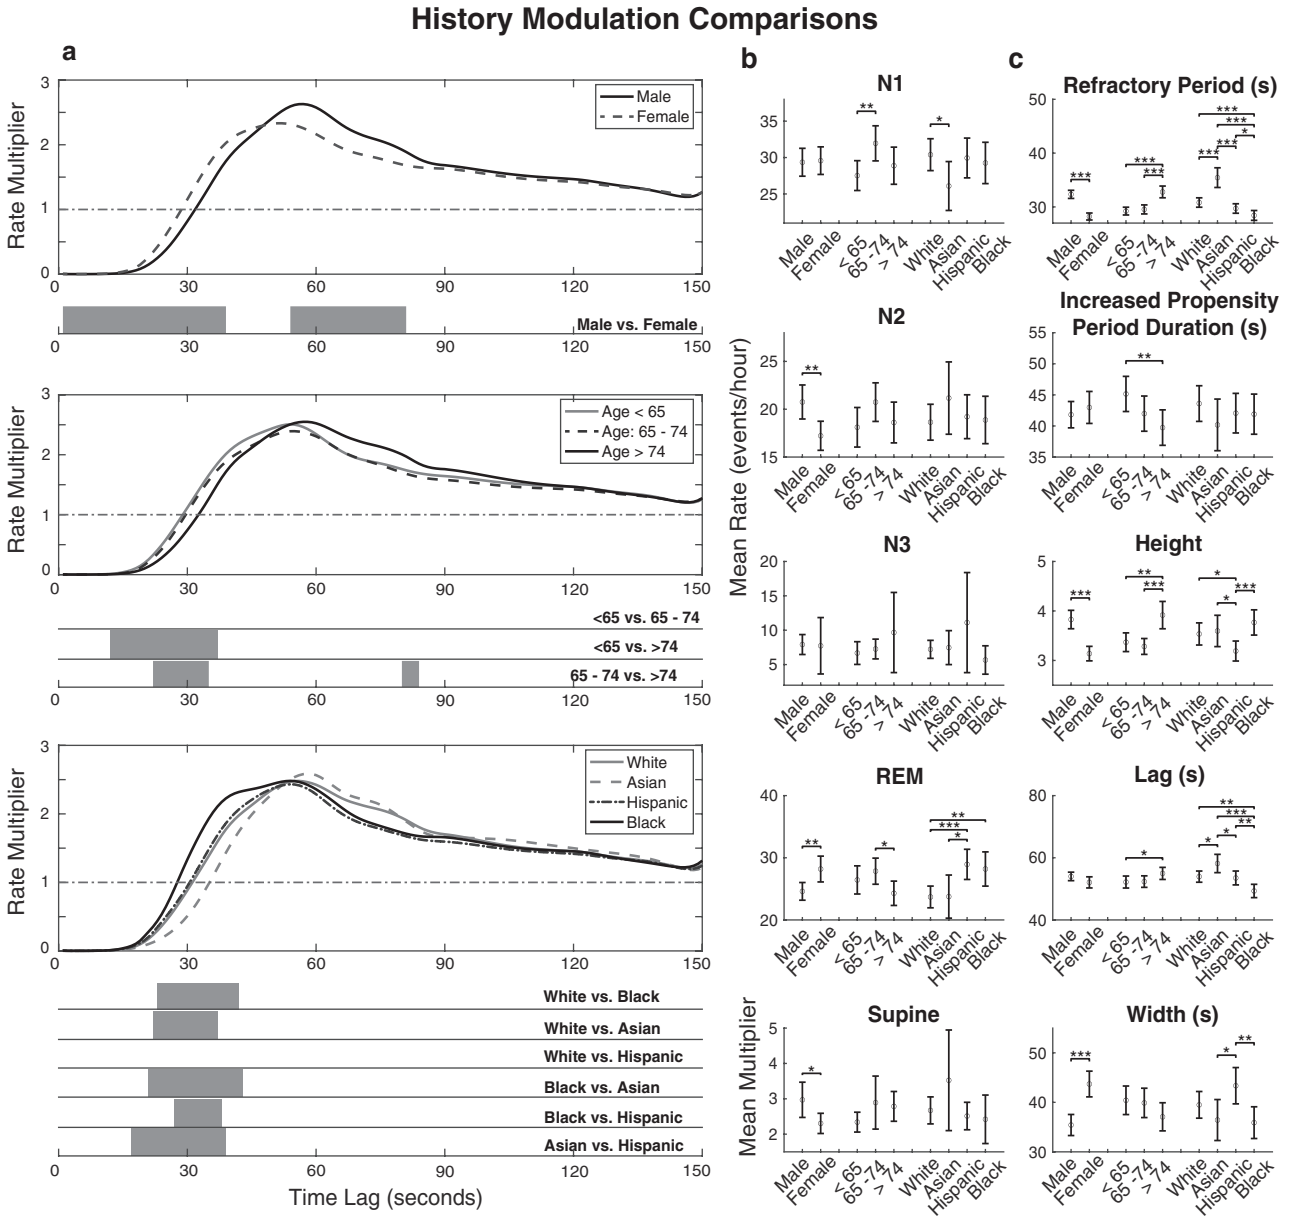

Figure 3. Statistical tests reveal significant differences in history modulation among different groups. (a) The mean history modulation curves with significant regions (gray) using global permutation tests (Significance level: .05). (b) The mean event rate in different sleep stages, as well as the mean supine multiplier for all groups. *t*-test for significant difference on the mean values is conducted between all levels of each factor (*, **, *** denote the *p*-value <.05, <.01, <.001), error bars represent the 95% confidence intervals on the means. (c) The comparisons of all summary statistics of history modulation in all groups.

the value of the rate multiplier as a function of the time since each previous event. This plot allows us to answer the question: How much more likely is there to be a respiratory event now, given that an event was observed X seconds ago?

Figure 1b shows a history modulation plot from another participant from the MESA dataset. The curve (solid black) shows the multiplicative effect and 95% confidence interval for each time lag (time since the last event), which reflects the structure of the inter-event intervals shown (bottom panel). Portions of the modulation curve significantly below 1 indicate a decreased rate, or refractory period during which few or no events occur, and portions above 1 indicate a period of increased propensity of events. For this participant, the curve suggests a significant reduction in rate between 0 s and ~20 s after each event (dark gray region), followed by a period between ~25 s and 55 s after each event during which the activity is significantly enhanced (light gray region). The modulation is maximal at ~35 s with a modulation value close to 6. The modulation curve then falls back towards 1, suggesting no significant history modulation effect at this point after a previous event. Thus, we can say that this participant is prone to bursts of respiratory events, with an event most likely to occur 35 s after the previous event at ~6 times the baseline rate. The addition of history dependence is exceptionally powerful as it enables us to identify specific patterns of events that are more or less likely to occur and predict times at which events become very likely.

The analysis of respiratory events in a dynamic context can also allow us to better characterize heterogeneity across

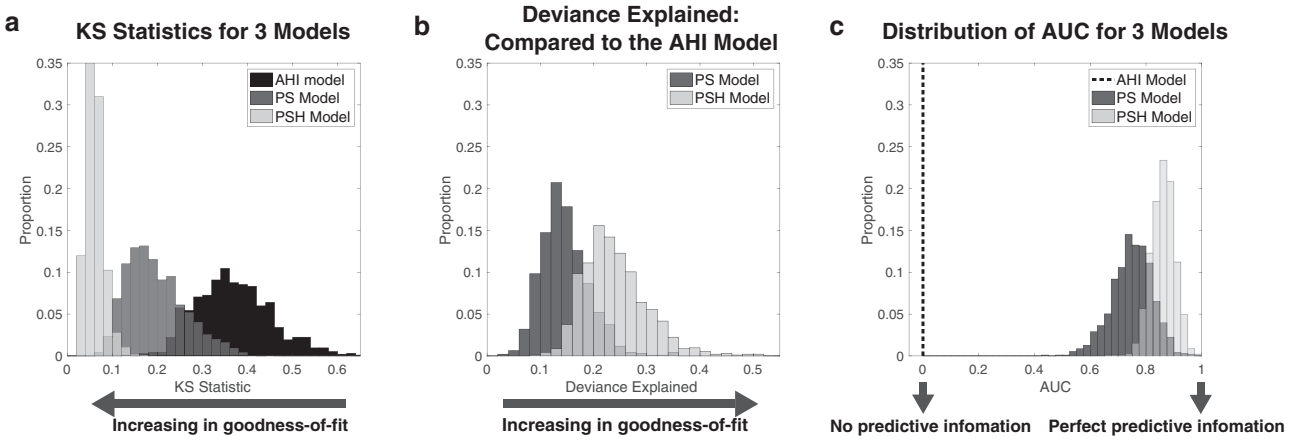

Figure 4. Goodness-of-fit analysis across the population shows PS model with position and stage explains more than AHI model, history components explain most and greatly increase the predictability. The result is conveyed by three ways: (a) shows the distribution of KS statistics for the three models, where AHI model has the largest rejection rate of KS tests across the population, reaching 94.98%, followed by PS model with rejection rate at 84.08%, while the PSH model passes over 97% KS tests; (b) gives the distribution of deviance explained compared to the AHI models across the population, where PSH model explains on average 24% more deviance than the AHI model, Chi-square tests at 1% significance level show that for 99.47% of the patients, adding history component statistically improves the PS model, and the percentage approaches to 100% when compared to the AHI model; (c) compares the distributions of the AUC for the three models, where the PSH model shows the best prediction ability with mean AUC equals to 0.86, while the PS model centers at 0.75 and AHI model always stands at 0. Black, dark gray and light gray refer to the AHI model, PS model, and PSH model, respectively.

participants. Figure 1c shows four additional MESA participants who have nearly identical AHIs of ~30 but different history dependence structures. While the history modulation curves for all participants share the same general form (a refractory period followed by a peak of increased propensity of events, which decreases to 1), the timing and degree of modulation differ for each individual. This variation in history dependence structure reflects highly heterogeneous temporal patterns in the event timing, which are distinctive for each individual despite sharing the same AHI.

To explicitly characterize features of individualized timing properties, we can define summary statistics on the history modulation curve. We quantify the length of the refractory period, as well as parameterize features in the increased propensity period, such as the lag, height at maximum, and width of this period (Figure 1b, see Appendix D for details). The table in Figure 1c shows the values of summary statistics for the four corresponding participants, which reflect the variation seen in the modulation curves. While participants A, B, and C have similar 20–30 s refractory periods, the refractory period for participant D is longer (61 s), suggesting an increased latency between consecutive events. The tall, narrow increased propensity peaks of participants B and D (reflected in the height and width statistics) suggest highly predictable trains of events, whereas the short and broad peaks of participants A and C indicate increased randomness of event timing. Overall, by characterizing history dependence in terms of interpretable parameters, we can quickly form a quantitative basis for the future development of respiratory phenotypes that reflect the temporal patterns of event activity.

History dependence significantly improves respiratory event time predictions

Figure 2a and b illustrate the impact of sleep position, stage, and history dependence on the rate of respiratory events using another participant from the MESA dataset. For this participant, we fit the AHI, PS, and PSH models and computed the goodness-of-fit of each model to the data using Kolmogorov-Smirnov (KS) plots [41]. Figure 2a shows the respiratory event times of a 4.5-hour segment along with the factors used by each model to predict these events (body position, stage, and history dependence). Note, the visualization of history component here uses the simplest history basis which is a set of predictors that include respiratory event trains lagged by time orders. In practice, we used 150-second history and employed the cardinal spline basis which acts as a smoothed version of the indicator basis. Additionally, the splines provide a more intuitive sense of what history components look like and how the historical events affect the model for estimates of current probability of event occurrence. Details of the spline functions are provided in Appendix C. Figure 2b shows the KS plots for the AHI (solid), PS (dashed), and PSH (dotted) model fits. The KS statistic measures the largest deviation between the KS plot and the y = x line, with significant lack of fit indicated for any KS plot that leaves the gray region. While the AHI and PS models fail the KS test (KS statistics = 0.50 and 0.28, respectively), the PSH model, which adds history dependence, passes the KS test (KS statistic = 0.07), with the curve lying completely within the gray region. This suggests that the temporal dynamics captured by the history dependent model are critical for characterizing the temporal structure of the events for this participant.

Furthermore, the addition of different model components can dramatically improve the ability to predict the precise timing of respiratory events. This is illustrated in Figure 2d, which compares the probability of seeing respiratory events in 10-second intervals for each of the different models in ~50 minute time window. The AHI model assumes a constant rate of events, which means that the probability of observing an event in any 10-second bin is quite small, below 0.5%. For the PS model, the estimated probability is a step function that depends on the current position/stage combination. During supine N1 sleep the estimated probability of an event in any 10-second interval is around 2%, while during nonsupine N3 sleep it falls well below 0.5%. Respiratory events are therefore more predictable under

this model than under the AHI model. Finally, in the PSH model, we can see that after each event occurs, the probability in subsequent 10-second bins that the next event will occur in that bin undergoes a consistent dynamic pattern that starts low and increases rapidly before the event actually occurs. When multiple events arrive in sequence, the probability at event times rises further, reaching values nearing 80%. At such times, based on the preceding pattern of events, we can be quite confident that about 30 s after a previous event, another event will occur in the next 10 s or shortly thereafter. At the event times, the PSH model predicts events with ~16 times higher probability than the AHI model. Although more difficult to see, the predicted probability during non-event periods is substantially lower under the PSH model than the AHI model. Thus, the PSH model predicts both event times and non-event times with much higher accuracy than do the AHI or PS models.

Figure 2c depicts a receiver operating characteristic (ROC) curve showing the ability of each model to predict whether or not an event will occur in the next 10 s [52]. Each model provides a probability of an event in each 10-second time bin, and the ROC curve shows how the rate of false positives and true positives change as a function of different probability thresholds for predicting an event. This curve is constructed so that the area under the curve (AUC) is equal to the probability that the corresponding model will give a higher probability of an event in bins where events occur than in those where events do not occur [53]. For the constant AHI model, there is no difference in the probability of an event between different intervals, and so the AUC is 0. For the PS model, the AUC increases substantially to a value of 0.754, based on the fact that events are more likely to occur during intervals with specific sleep stage and position combinations. For the PSH model, the AUC again increases sizably to a value of 0.890, reflecting the large increase in the predictability of respiratory events at specific periods after previous events.

In summary, as we extend our model to include additional predictors, we move from a static descriptive statistic of respiratory event occurrence over the night to a dynamic process that makes the respiratory events more predictable. These modeling results also provide evidence that history dependence is a highly-informative feature for the prediction of respiratory event times. Moreover, these results provide insight into underlying mechanisms and information for use for targeted time-dependent interventions such as positional devices or PAP levels which can be programmed to change over the need to address dynamic changes in airway collapsibility.

Population analysis of sleep disordered breathing

Given the ability to quantify the patient-specific profiles of respiratory event dynamics, we can characterize the variability in these factors over populations. The modeling results across the entire population of 936 participants from the MESA dataset are summarized in Sup. Figure 3. Figure 3 summarizes the population results as a function of self-reported gender (male, female), age category (<65, 65–74, >74), and race/ethnicity (Black, Asian, White, and Hispanic).

Figure 3a shows average history modulation curves and regions of significant differences between demographic groups (gray regions). We see significant differences in modulation within the first 30–40 s after each event, suggesting that refractory time may be an important distinguishing factor between groups. When comparing history dependence as a function of gender, we see a significant region between 55 s and 80 s, with males having a higher average modulation during those times, suggesting a longer interval between apnea events in males versus females. This is potentially due to the greater average height in males leading to longer circulation time.

Figure 3b compares the model parameters representing mean rate per sleep stage as well as the multiplier for the supine position. In general, these results corroborate previous analyses of the MESA dataset [48, 54]. The inferences from this model related to the prevalence of REM and supine dominance also corroborate findings described in the literature [55, 56].

Figure 3c compares derived statistics from the history modulation curves as a function of demographics. In particular, the refractory period and increased propensity peak height tend to differ significantly across gender, age, and race/ethnicity. The increased propensity peak lag shows significant age and race/ethnicity differences as well. Since the history modulation features are estimated while simultaneously accounting for AHI, correlation analysis confirms the weak correlations between history modulation features and AHI (Sup. Figure 4).

Additional analysis (Sup. Figure 3c) shows that across nearly all participants, significant history modulation structure (i.e. the influence of the previous event on rate) lasts between 60 s and 180 s following each event. Although this general structure is common among all participants, the specific shape of the influence of past spiking as a function of time, and therefore the patterns of events that emerge, can vary substantially from individual to individual. Notably, all history modulation statistics demonstrate strong heterogeneity, as evidenced by the long tails of the distributions (Sup. Figure 3d).

Overall, these results show wide variation in respiratory patterns beyond the effect of stage and position, with significant differences between key demographic factors. This highlights the viability of this approach as new and potentially valuable basis for patient phenotyping.

We next perform standard goodness-of-fit and classification performance metrics on the population data to quantify the improvement in model fit and predictive power as a function of model selection, the results of which are summarized in Figure 4. For the entire population, we look at the distribution of KS statistics, deviance explained, and area under the ROC curve.

In Figure 4a, we see the distribution of KS statistics for each model, with a smaller KS statistic values reflecting better goodness-of-fit. The AHI model has the poorest fit, with the largest mean KS statistic of 0.70 and ~95% of models rejected by a KS test at the 0.05 significance level. The PS model improves over the AHI with a mean KS statistic of 0.20, but still ~84% of models are rejected by the KS test. Finally, the PSH model has the smallest mean KS statistic value of 0.02, and only ~2.5% of the models are rejected by the KS test. Figure 4b shows the distribution of deviance explained relative to the AHI, for PS and PSH models. The PS model on average explains 14.62% more deviance than the AHI model, while PSH model explains 23.98%. Chi-square tests show that the addition of history dependence in the PSH model provides a significant improvement for 99.47% of the participants over the PS model and for 100% of the participants over the AHI model. Additional analysis of the relationship between model errors and AHI are shown in Sup. Figure 5. In fact, the overall mean squared deviance residual increases with increasing AHI. A more detailed analysis separating out

the event and no-event intervals shows that this increase is dominated by a very small increase in the predicted rate during no-event intervals, however, there are substantial improvements of model fitting in event times with increasing AHI, and PSH model shows remarkable reduction in mean squared deviance residual compared to AHI model and PS model.

Figure 4c shows the distribution of AUC for ROC curves using the models to predict whether or not a respiratory event will occur in any 10-second interval. The AUC for the constant AHI model is always 0, since the model does not differentiate between times where events are more or less likely. The distribution of the AUC for the PS model is centered around 0.75, which reflects the fact that events are more likely to occur at specific combinations of sleep stage and position than others. The AUC for the PSH model has a distribution with a mean that increases further to about 0.86, reflecting the increase in predictability related to the history of past events, even when accounting for differences in event rates due to sleep stage and position. Thus, history dependence provides consistent and significant improvements in event predictability. Overall, history dependence acts as a critical component in modeling, explaining, and predicting OSA temporal dynamics.

Discussion

Sleep-disordered breathing is a dynamic process in which the rate of respiratory events is influenced by multiple factors. Despite its dynamic nature, clinical diagnosis tends to collapse the complex processes underlying OSA to a single metric measuring the average rate of respiratory event occurrence, the AHI. Thus, potentially valuable information is being lost by ignoring OSA temporal dynamics and variability. Here, we develop a general point process framework that enables us to look at the instantaneous rate of respiratory events as a function of any combination of the factors influencing it, including the timing of past events. This approach can provide insight in mechanistic heterogeneity, which can be informative in therapeutic decision making.

The comparisons between the AHI model, PS model, and PSH model highlight the importance of each of the model factors. Sleep position, stage, and the history of past respiratory events consistently have significant and substantial effect of the instantaneous event intensity. For the participant shown in Figure 2, a shift from N3 to N1 corresponds to a 3-fold increase in the instantaneous event rate, and a shift from nonsupine to supine sleep corresponds to a 1.6-fold increase. Similarly, specific patterns of past events make future events substantially more likely. In the example in Figure 2d, the PSH model predicts the probability of an event occurring to be nearly 16-fold higher than the AHI model at those times when events actually occur. Such increases in predictability are common when history dependence is included into point process models. Adding history dependence also has the effect of drastically improving the goodness-of-fit of these models as expressed through the deviance and KS statistics. KS tests showed significant model mis-fit for nearly all participants using the AHI model, and were still present in the overwhelming majority of participants when accounting for the effects of sleep stage and position; including history in these models consistently improved the quality of the fit so that nearly no participants showed significant lack of fit.

Population analysis using the PSH model corroborated previous clinical findings while also revealing features of respiratory event patterns that are consistent across individuals and features that are patient specific. Across the entire population, the model fits confirmed previous findings that respiratory event occurrence is reduced in deeper sleep stages when arousal threshold is elevated and upper airway collapsibility less than light sleep. The fits also corroborate prior work showing that more respiratory events occur in supine than other positions [57], consistent with mechanisms of apnea attributed to unfavorable airway geometry and reduced lung volume during supine sleep [58, 59]; the model fits also revealed significant differences in different population cohorts, such as the level of REM and supine dominant activity in males and females, as well as their refractory period and increased propensity peak height. Our result aligns with the previous findings showing females have more respiratory events in REM and less events in supine position compared to males, reflecting the lower loop gain and less airway collapsibility in females [54, 60]. These findings show the potential for this framework to address specific epidemiological questions and mechanistic hypotheses in future applications.

The model fits also revealed that most participants share a common history modulation structure following each respiratory event: a consistent refractory period followed by an increased propensity period when events are more common than during periods not immediately preceded by an event. Yet, we also observed a broad range of different history modulation structures across the population. Given that different history modulation structure will lead to distinct patters of events from individual to individual, further understanding the statistical structure of these patterns could assist clinicians with phenotyping a patient's disease, allowing them to personalize therapeutic strategies. A key factor in understanding the clinical applicability of history modulation will be the characterization of long-term stability in modulation structure. While these data had one study night available, future work will focus on understanding the temporal stability of history modulation curves over time using data sets with repeated studies per participant. Additionally, the availability of large sleep databases paired with genetic data will enable future studies to characterize heritability and other genetic factors involved in history modulation.

There are a variety of ways to extend the modeling methods and data analyses explored in this work. For one, the models we presented included only a small set of factors that can impact sleep disordered breathing. It would be natural to explore the influence of additional physiological and clinical covariates by adding them to the model. For example, changes in neural electrophysiological activity have been associated with apnea severity [61–63], which might be captured by adding to the model variables related to the power in specific frequency bands of the EEG. One advantage of our model-based framework is that it allows for disambiguation of the influences of confounded covariates. Sleep stage is, by definition, correlated with features of the EEG, but a point process model including both could help researchers identify additional information in the EEG not captured by sleep staging. Our model did not explicitly include the information about the duration of the respiratory events, which could be another important factor to include.

Another way to extend these models is to include interactions between the factors influencing the respiratory event rate. All of the models we explored assumed that the influences

of each factor were multiplicatively separable. Thus, in the PSH model, each participant has a typical apnea pattern determined by the history modulation, and only the frequency of that pattern is modulated by changes in sleep stage and position. By adding terms that include interactions between history, stage, and position, we would generate a model where distinct patterns of respiratory events could be fit for different sleep stage and position combinations.

Future work applying these methods to larger datasets that include measures of disease severity and therapeutic outcomes will help determine the value of inferences from these models. We posited above that history modulation curves could be viewed as patient specific signatures of respiratory event dynamics. Analyses over larger datasets could help determine whether these signatures exist in a continuum along one or more dimensions or whether they cluster into a small number of phenotypic groups. In either case, statistical models can be constructed relating features of the history dependence curves to clinical disease measures. If respiratory event patterns are predictive of patient prognoses or responses to various therapies, clinicians can use model fits for individual patient sleep data to guide clinical decision making.

Additionally, the striking increase in the predictability of individual respiratory events using these models might help guide novel therapeutic approaches. For example, this approach could be used as an improved control signal for therapeutics to apply positive airway pressure, neural stimulation, or positional interventions precisely at times when respiratory events are most likely to occur. Because of the dynamic nature of sleep disordered breathing, it may take interrupting only a few events to dramatically decrease the total number of events over a full night's sleep.

Statistical modeling provides well-studied tools for identifying predictable structure in dynamic data, for determining the factors that influence that structure, and for quantifying uncertainty. Point process models therefore offer a powerful tool for sleep scientists and clinicians to move beyond simple descriptive statistics to better understand the dynamics of sleep disordered breathing and to provide more effective treatment options.

Funding

This work was supported by the National Institute of Neurological Disorders and Stroke (NINDS), United States, under Grant R01 NS-096177 (M.J.P.). S.R. was partly supported by HL R35 1358181.

Acknowledgments

The authors would like to acknowledge early implementation work by Karoline Weber [75]. The Multi-Ethnic Study of Atherosclerosis (MESA) Sleep Ancillary study was funded by NIH-NHLBI Association of Sleep Disorders with Cardiovascular Health Across Ethnic Groups (RO1 HL098433). MESA is supported by NHLBI funded contracts HHSN268201500003I, N01-HC-95159, N01-HC-95160, N01-HC-95161, N01-HC-95162, N01-HC-95163, N01-HC-95164, N01-HC-95165, N01-HC-95166, N01-HC-95167, N01-HC-95168 and N01-HC-95169 from the National Heart, Lung, and Blood Institute, and by cooperative agreements UL1-TR-000040, UL1-TR-001079, and UL1-TR-001420 funded by NCATS. The National Sleep Research Resource was supported by the National Heart, Lung, and Blood Institute (R24 HL114473, 75N92019R002).

Disclosure Statement

Financial Disclosure. SR reports consultant fees for work unrelated to this manuscript from Jazz Pharma, Esiai Inc, Lilly Inc and Apnimed Inc.

Non-financial Disclosure. None.

Data Availability

The datasets generated during and/or analyzed during the current study are available in the Multi-Ethnic Study of Atherosclerosis (MESA), from the National Sleep Research Resource (www.sleepdata.org) [47, 48]. Code is available at http://sleepEEG.org/.

Appendix A: Point Process Framework

To develop a point process framework, we first define a function called the counting process, $N(t)$, which denotes the total number of respiratory events that occur up to and including time t, for $t \in (0, \mathrm{TST}]$, where TST represents a participant's total sleep time (the amount of non-Wake time during a recording). The AHI is therefore defined as the total event count divided by the total sleep time, $\mathrm{AHI} = N(\mathrm{TST})/\mathrm{TST}$, which expresses the average rate of events over the entire night's sleep.

It is clear, however, that respiratory events are not independent. We thus need to move towards a framework that allows us to capture the temporal patterns in the data. A general point process can be expressed in terms of a conditional intensity function [45], which is defined as

$$\lambda(t|H_t) = \lim_{\Delta \to 0} \frac{P[N(t+\Delta) - N(t) = 1|H_t]}{\Delta}, \quad (1)$$

for which $\lambda(t \mid H_t)$ is the "instantaneous rate" at time t, given H_t, the past history of events up to, but not including time t. Whereas the AHI defines the average rate of events over the entire night, the conditional intensity defines an instantaneous rate of events that can be influenced by variables such as sleep position and stage as well as by previously occurring events.

For any time interval of sufficiently small length, Δ, equation (1) indicates that the probability of observing an respiratory event in that interval is approximately proportional to this conditional intensity:

$$P[\text{One respiratory event in } [t, t+\Delta) \mid H_t] \approx \lambda(t \mid H_t)\Delta.$$

For a specific pattern of respiratory events observed at times $s_1 < s_2 < \ldots < s_n$ over the interval $[0, T]$, the likelihood of observing those event times is given by [45]

$$\prod_{i=1}^{n} [\lambda(s_i \mid H_{s_i})] e^{-\int_0^T \lambda(t \mid H_t) dt} \quad (2)$$

Thus, the conditional intensity defines the probability of observing a respiratory event moment by moment, which can be used to compute the likelihood of any pattern of events over the course of the night.

Appendix B: Model Inference Methods

Defining a point process model of respiratory events involves expressing the conditional intensity, $\lambda(t \mid H_t)$, as a function of the variables that influence these events. Here we focus on variables of known clinical interest including sleep stage and position, as well as variables that can be used to capture temporal patterns of respiratory events, including the history of prior events. One well-studied class of point process models can be expressed under the broader statistical framework of generalized linear models (GLMs) [41]. Most commonly, point process GLMs are constructed by letting the log of the conditional intensity be a linear function of the model parameters and functions of the variables that influence the rate of events,

$$\log(\lambda(t \mid H_t)) = \sum_{j=1}^{p} \beta_j g_j(x_t),$$

where p is number of model parameters, β_j is the jth model parameter, and $g_j(x_t)$ is some function of the variables, x_t, that influence the intensity of events.

Point process GLMs have several advantageous properties. They can flexibly capture the influences of multiple variables, even when the influences are non-linear [41]; the model parameters often have simple interpretations [64]; and they have convex likelihoods allowing for computationally efficient estimation [41, 65, 66]. There exist simple, computationally efficient algorithms to compute the maximum likelihood estimators for GLM parameters, such as iteratively reweighted least squares (IRLS) [46]. Robust software packages implementing these algorithms are available in most statistical analysis platforms and coding languages (e.g. R, MATLAB, and Python). The resulting maximum likelihood estimators are asymptotically unbiased and asymptotically have minimum mean-squared error. The IRLS algorithm also computes the Fisher information, which determines the standard errors of the model parameters and is used to compute test statistics to determine whether any collection of variables significantly influence the intensity of events [66].

Point process GLMs also provide well-studied approaches for assessing the goodness-of-fit of the model to the observed data. The model deviance generalizes the notion of residual model variance from linear regression and can be used to compare models or to quantify the amount of variability explained by the model. Deviance residuals can be computed and plotted to determine whether there remains unmodeled structure after fitting the GLM [46, 67]. If the model is correct, the time-rescaling theorem [68] can be used to remap the event times into a homogenous Poisson process. After rescaling, Kolmogorov-Smirnov (KS) plots can be used to compare the distribution of inter-event times to those predicted by the model. A well-fit model will produce a KS plot that closely follows a 45-degree line and stays within its significance bounds. KS plots that are not contained in these bounds suggest lack-of-fit in the model.

Appendix C: Modeling Respiratory Event Dynamics

In this section we describe several candidate models for characterizing respiratory event dynamics. We start with a constant rate model, which mimics the AHI, and progressively build to a full model that incorporates body position, stage, and past event history in the estimation of an "instantaneous AHI".

Constant Rate (AHI) Model

As stated, the AHI expresses the sample rate of events over the observation period. A point process where all of the events are independent of each other and are characterized by a single rate parameter is called a Poisson process. Therefore, we construct a point process GLM that assumes that the rate of the respiratory events is constant throughout the whole night to explore the statistical properties of a model based on a single AHI statistic:

$$\log(\lambda(t \mid H_t)) = \beta_0 \quad (3)$$

Here, β_0 is a model parameter that can be interpreted as the log of the event rate. It is simple to show that the maximum likelihood estimate of β_0 is exactly the log of the AHI statistic. This model posits that respiratory events occur as a *homogeneous Poisson process* over the entire night [45].

We refer to this constant rate model as the AHI model because it returns the AHI statistic. Additionally, it allows us to compute confidence bounds about the AHI [27], to evaluate the sufficiency of this statistic to characterize the observed patterns of respiratory events, and to assess whether adding variables to the model provides a statistically significant improvement in fitting these data.

Note that the parameter β_0 is interpreted as the log rate, and that we obtain the model estimate for the rate by exponentiating it. We can multiply this rate by the length of any sleep period to compute the expected number of events in that period, based on the model fit.

Position and Stage (PS) Dependent Model

Clinicians are often interested in how respiratory events depend on clinical factors such as sleep position and stage. Some patients have predominantly supine OSA and others predominantly REM. Typically, a night's sleep will be partitioned based on these factors, and separate AHIs computed for each sleep position (e.g. AHI_{supine}) or sleep stage (e.g. AHI_{REM}). Guided by this clinical practice, we explore point process GLMs including sleep position and stage (which we call the PS model) with the following form

$$\log(\lambda(t \mid H_t)) = \overbrace{\beta_p I_P(t)}^{\text{Position}} + \overbrace{\Sigma_{s \in S} \beta_s I_s(t, s)}^{\text{Stage}} \quad (4)$$

Where S represents the set of sleep stages: {N1, N2, N3, REM, Wake}, β_p and β_s refer to the model

parameters reflecting the effects of body position and sleep stage, respectively. The indicator functions $I_P(t)$ and $I_s(t,s)$, express whether a participant is in the corresponding position or stage at time t. Specifically,

$$I_P(t) = \begin{cases} 1 & \text{if Nonsupine at } t \\ 0 & \text{if Supine at } t \end{cases}, \quad I_s(t,s) = \begin{cases} 1 & \text{if stage} = s \text{ at } t \\ 0 & \text{otherwise} \end{cases}$$

The PS model posits that respiratory events occur as an inhomogeneous Poisson process with a rate that depends on sleep stage and position. Comparing the PS to the AHI model allows us to determine the influence of these variables on the rate of respiratory events. The values of $\exp(\beta_s)$ are event rates for each corresponding sleep stage when the participant is supine, and the value of $\exp(\beta_p)$ represents a multiplier to the event rate when the participant in non-supine. If the value of β_p is significantly different from zero, it suggests that there is strong evidence for an effect of sleep position on the event rate. If the values of β_s are significantly different from each other, it suggests that there is strong evidence for an effect of sleep stage on the event rate. By comparing the model deviance of the PS model to that of the AHI model, we can determine the amount of variability in the data that is explained by differences in the event rate due to sleep stage and position. Here, we assume that sleep stage and body position have multiplicatively separable effects, but the modeling framework allows us easily to extend the PS model to include interaction terms.

Position, Stage, and History Dependent (PSH) Model

Both the AHI and PS models assume that the probability distribution of future events is independent of past events. Thus, these models treat OSA as a Poisson process, and cannot be used to distinguish between patients with similar rates but different temporal patterns of events. However, it is clear that respiratory events have some impact on future events for at least some period of time due to the physiological mechanisms of sleep apnea. To account for the possibility of dependence between respiratory events, we extend the PS model by adding functions of the history of previous respiratory events to the GLM (which we call the PSH model) as follows:

$$\log(\lambda(t \mid H_t)) = \overbrace{\beta_p I_P(t)}^{\text{Position}} + \overbrace{\sum_{s\in S}\beta_s I_s(t,s)}^{\text{Stage}} + \overbrace{\sum_{k=1}^{K} h_k g_k(H_t)}^{\text{History}} \quad (5)$$

Where $g_k(H_t)$ represents a collection of basis functions that smooth the influence of the event history H_t, K is the number of basis function used to capture history dependence, and h_k are the model parameters whose values determine which patterns of events are more or less likely. Here, we used cardinal spline basis functions for $g_k(H_t)$, details are provided below.

Spline Functions

There are many ways to include the past history of respiratory events in a point process model. Here we chose to use spline functions to provide a parsimonious representation of the past history of respiratory events. Splines are piecewise, low-degree polynomial curves that are continuous and smooth interpolations between control points, called knots. Splines are flexible functions, and thus able to capture a wide range of modulation structure. One commonly used spline is the cardinal spline, which uses third-order polynomials to interpolate between control points. Cardinal splines have been successfully used to model patterns of neural spiking events in electrophysiology data [69–74].

Let (x_i, p_i) represent the lag and associated magnitude of the i^{th} knot. The equation for the cubic cardinal spline at a time point x between control points x_i and x_{i+1} is defined as:

$$Sp(u) = \left(2u^3 - 3u^2 + 1\right)p_i + \left(u^3 - 2u^2 + u\right)p_i' + (-2u^3 + 3u^2)p_{i+1} + (u^3 - u^2)p_{i+1}' \quad (6)$$

Here, $u = \frac{x - x_i}{x_{i+1} - x_i}$, and p_i' is the derivative at the ith control point, which is determined by a tension parameter s and its neighboring control points:

- $p_i' = \frac{s}{l_1}(p_{i+1} - p_{i-1})$, where $l_1 = \frac{x_{i+1} - x_{i-1}}{x_{i+1} - x_i}$
- $p_{i+1}' = \frac{s}{l_2}(p_{i+2} - p_i)$, where $l_2 = \frac{x_{i+2} - x_i}{x_{i+1} - x_i}$

The cardinal spline can therefore be expressed in matrix form as:

$$Sp(u) = \begin{bmatrix} u^3 & u^2 & u & 1 \end{bmatrix} \begin{bmatrix} -\frac{s}{l_1} & \left(2 - \frac{s}{l_2}\right) & \left(\frac{s}{l_1} - 2\right) & \frac{s}{l_2} \\ \frac{2s}{l_1} & \left(\frac{s}{l_2} - 3\right) & \left(3 - \frac{2s}{l_1}\right) & -\frac{s}{l_2} \\ -\frac{s}{l_1} & 0 & \frac{s}{l_1} & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} p_{i-1} \\ p_i \\ p_{i+1} \\ p_{i+2} \end{bmatrix} \quad (7)$$

Equation (7) shows that each cardinal spline between two knots is determined by the four adjacent knots and the shape is controlled by the tension parameter s. For a detailed overview of incorporating spline models into a GLM framework and other spline functions, see Sarmashghi et al. [69].

Implementation Details

All the models were fit using the GLM package in MATLAB. Since there is no event history at the beginning of each night, we used the first 150 s of each night's sleep to construct an initial respiratory event history. History dependence was then fit using a cardinal spline basis function with a tension parameter of 0.5. With end points at 0 s and 150 s, four knots were placed evenly between the 10th percentile of inter-event intervals and 90 s, with another knot at 120 s. Two additional knots placed at –10 s and 160 s were used to determine the derivatives of the spline function at the end points. Code is available at http://sleepEEG.org/.

Appendix D: Modulation Curve Summary Statistics

Given the wide range of history dependence structure observed across the population, we explored a set of summary statistics that captured salient features of these curves. From the model parameters, we get a 95% upper confidence bound (UCB) and lower confidence bound (LCB) for the modulation curve at each time lag. Using these, we define the following five model summary statistics:

1. Refractory period—The first lag at which the UCB is greater than or equal to 1.

2. Increased propensity period duration—The length of the period from when the LCB is first greater than or equal to 1 to when the LCB returns to be less than or equal to 1.
3. Increased propensity peak height—The maximum value of the history modulation.
4. Increased propensity peak lag—The lag corresponding to peak height.
5. Increased propensity peak width—The half peak bandwidth of the history modulation.

These summary statistics were computed for all participants in the MESA cohort.

Supplementary Material

Supplementary material is available at *SLEEP* online.

References

1. Greenberg H, *et al*. Obstructive sleep apnea. In: *Principles and Practice of Sleep Medicine*. Elsevier; 2017:1110–1124.e6. doi:10.1016/B978-0-323-24288-2.00114-8
2. Kwon Y, *et al*. Sleep, sleep apnea and atrial fibrillation: questions and answers. *Sleep Med Rev*. 2018;**39**:134–142. doi:10.1016/j.smrv.2017.08.005.
3. Muraki I, *et al*. Sleep apnea and type 2 diabetes. *J Diabetes Investig*. 2018;**9**(5):991–997. doi:10.1111/jdi.12823.
4. Tsai M, *et al*. Sleep apnea in heart failure. *Curr Treat Options Cardiovasc Med* 2018;**20**(4):1–11. doi:10.1007/s11936-018-0624-0.
5. Anzai T, *et al*. Association between central sleep apnea and atrial fibrillation/flutter in Japanese-American men: The Kuakini Honolulu Heart Program (HHP) and Honolulu-Asia Aging Study (HAAS). *J Electrocardiol*. 2020;**61**:10–17. doi:10.1016/j.jelectrocard.2020.05.005.
6. Peppard PE, *et al*. Increased prevalence of sleep-disordered breathing in adults. *Am J Epidemiol*. 2013;**177**(9):1006–1014. doi:10.1093/aje/kws342.
7. Kohler M. Risk factors and treatment for obstructive sleep apnea amongst obese children and adults. *Curr Opin Allergy Clin Immunol*. 2009;**9**(1):4–9. doi:10.1097/ACI.0b013e32831d8184.
8. Zhang J, *et al*. The association of neck circumference with incident congestive heart failure and coronary heart disease mortality in a community-based population with or without sleep-disordered breathing. *BMC Cardiovasc Disord*. 2018;**18**(1):108. doi:10.1186/s12872-018-0846-9.
9. Baril AA, *et al*. Biomarkers of dementia in obstructive sleep apnea. *Sleep Med Rev*. 2018;**42**:139–148. doi:10.1016/j.smrv.2018.08.001.
10. Daulatzai MA. Evidence of neurodegeneration in obstructive sleep apnea: Relationship between obstructive sleep apnea and cognitive dysfunction in the elderly: OSA and Neuropathogenesis of Late-Onset Alzheimer's Disease. *J Neurosci Res*. 2015;**93**(12):1778–1794. doi:10.1002/jnr.23634.
11. Azarbarzin A, *et al*. The sleep apnea-specific hypoxic burden predicts incident heart failure. *Chest*. 2020;**158**(2):739–750. doi:10.1016/j.chest.2020.03.053.
12. Mazzotti DR, *et al*. Symptom subtypes of obstructive sleep apnea predict incidence of cardiovascular outcomes. *Am J Respir Crit Care Med*. 2019;**200**(4):493–506. doi:10.1164/rccm.201808-1509OC.
13. Ogilvie RP, *et al*. Joint effects of OSA and self-reported sleepiness on incident CHD and stroke. *Sleep Med*. 2018;**44**:32–37. doi:10.1016/j.sleep.2018.01.004.
14. Watson NF. Health care savings: the economic value of diagnostic and therapeutic care for obstructive sleep apnea. *J Clin Sleep Med*. 2016;**12**(08):1075–1077. doi:10.5664/jcsm.6034.
15. Osman AM, *et al*. Obstructive sleep apnea: current perspectives. *NSS*. 2018;**10**:21–34. doi:10.2147/nss.s124657.
16. Edwards BA, *et al*. More than the sum of the respiratory events: personalized medicine approaches for obstructive sleep apnea. *Am J Respir Crit Care Med*. 2019;**200**(6):691–703. doi:10.1164/rccm.201901-0014TR.
17. Limoges E, *et al*. Atypical sleep architecture and the autism phenotype. *Brain*. 2005;**128**(Pt 5):1049–1061. doi:10.1093/brain/awh425.
18. Coughlin K, *et al*. Phenotypes of obstructive sleep apnea. *Otolaryngol Clin North Am*. 2020;**53**(3):329–338. doi:10.1016/j.otc.2020.02.010.
19. Zinchuk AV, *et al*. Phenotypes in obstructive sleep apnea: a definition, examples and evolution of approaches. *Sleep Med Rev*. 2017;**35**:113–123. doi:10.1016/j.smrv.2016.10.002.
20. Sawyer AM, *et al*. A systematic review of CPAP adherence across age groups: clinical and empiric insights for developing CPAP adherence interventions. *Sleep Med Rev*. 2011;**15**(6):343–356. doi:10.1016/j.smrv.2011.01.003.
21. Aloia MS. Understanding the problem of poor CPAP adherence. *Sleep Med Rev*. 2011;**15**(6):341–342. doi:10.1016/j.smrv.2011.04.002.
22. Simon SL, *et al*. Barriers to treatment of paediatric obstructive sleep apnoea: development of the adherence barriers to continuous positive airway pressure (CPAP) questionnaire. *Sleep Med*. 2012;**13**(2):172–177. doi:10.1016/j.sleep.2011.10.026.
23. White LH, *et al*. Night-to-night variability in obstructive sleep apnea severity: relationship to overnight rostral fluid shift. *J Clin Sleep Med*. 2015;**11**(02):149–156. doi:10.5664/jcsm.4462.
24. Tschopp S, *et al*. Night-to-night variability in obstructive sleep apnea using peripheral arterial tonometry: a case for multiple night testing. *J Clin Sleep Med*. 2021;**17**(9):1751–1758. doi:10.5664/jcsm.9300.
25. Stöberl AS, *et al*. Night-to-night variability of obstructive sleep apnea. *J Sleep Res*. 2017;**26**(6):782–788. doi:10.1111/jsr.12558.
26. Roeder M, *et al*. Night-to-night variability of respiratory events in obstructive sleep apnoea: a systematic review and meta-analysis. *Thorax*. 2020;**75**(12):1095–1102. doi:10.1136/thoraxjnl-2020-214544.
27. Thomas RJ, *et al*. Quantifying statistical uncertainty in metrics of sleep disordered breathing. *Sleep Med*. 2020;**65**:161–169. doi:10.1016/j.sleep.2019.06.003.
28. Iber. The AASM Manual for the Scoring of Sleep and Associated Events: Rules, Terminology and Technical Specifications, 1st ed. Westchester, IL: American Academy of Sleep Medicine, 2007. https://scholar.google.com/citations?view_op=view_citation&hl=en&user=mLTec7cAAAAJ&citation_for_view=mLTec7cAAAAJ:tOudhMTPpwUC.
29. Berry RB, *et al*. Rules for scoring respiratory events in sleep: update of the 2007 AASM manual for the scoring of sleep and associated events: deliberations of the sleep apnea definitions task force of the American Academy of Sleep Medicine. *J Clin Sleep Med*. 2012;**08**(05):597–619. doi:10.5664/jcsm.2172.
30. Joosten SA, *et al*. Statistical uncertainty of the apnea-hypopnea index is another reason to question the utility of this metric. *Sleep Med*. 2020;**65**:159–160. doi:10.1016/j.sleep.2019.07.009.

31. Punjabi NM. COUNTERPOINT: is the apnea-hypopnea index the best way to quantify the severity of sleep-disordered breathing? No. *Chest*. 2016;**149**(1):16–19. doi:10.1378/chest.14-2261.
32. Vgontzas AN. Excessive daytime sleepiness in sleep apnea: it is not just apnea hypopnea index. *Sleep Med*. 2008;**9**(7):712–714. doi:10.1016/j.sleep.2008.05.001.
33. Brockwell PJ, *et al*. *Time Series: Theory and Methods (Springer Series in Statistics)*. 2nd ed. 1991. 2nd printing 2009. Softcover reprint of the original 2nd ed. 1991. Springer; 2009. New York. http://books.google.com/books?id=_DcYu_EhVzUC&printsec=frontcover&dq=TimeSeries+Theory+and+Methods+brockwell&hl=&cd=1&source=gbs_api. Accessed March 5, 2021.
34. Durbin J, *et al*. *Time Series Analysis by State Space Methods*. New York: Oxford University Press; 2001. http://scholar.google.com/scholar?q=related:axQQYgu8FIMJ:scholar.google.com/&hl=en&num=20&as_sdt=0,5&as_ylo=2001&as_yhi=2001. Accessed March 5, 2021.
35. Kim CJ, *et al*. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. **Vol. 1**. Cambridge, Mass: The MIT Press; 1999. http://ideas.repec.org/b/mtp/titles/0262112388.html. Accessed March 5, 2021.
36. Kitagawa G, *et al*. *Smoothness Priors Analysis of Time Series*. New York, NY: Springer-Verlag New York; 1996. http://lccn.loc.gov/96022800. Accessed March 5, 2021.
37. Snyder DL, *et al*. *Random Point Processes in Time and Space*. Springer Science & Business Media; 2012. New York. http://www.worldcat.org/title/random-point-processes-in-time-and-space/oclc/958524536. Accessed March 5, 2021.
38. Lo CC, *et al*. Dynamics of sleep-wake transitions during sleep. *EPL*. 2002;**57**(5):625. doi:10.1209/epl/i2002-00508-7.
39. Lo CC, *et al*. Common scale-invariant patterns of sleep–wake transitions across mammalian species. *Proc Natl Acad Sci USA*. 2004;**101**(50):17545–17548. doi:10.1073/pnas.0408242101.
40. Guilleminault C, *et al*. Sleep apnea syndromes and related sleep disorders. In R. L. Williams, I. Karacan, & C. A. Moore (Eds.), Sleep disorders: Diagnosis and treatment (2nd ed., pp. 47–71). New York: Wiley.
41. Truccolo W, *et al*. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol*. 2005;**93**(2):1074–1089. doi:10.1152/jn.00697.2004.
42. Tempel E, *et al*. Galaxy spin alignment in filaments and sheets: observational evidence. *ApJ*. 2013;**775**(2):L42. doi:10.1088/2041-8205/775/2/l42.
43. Engle RF, *et al*. Trades and quotes: a bivariate point process. *J Financial Econ*. 2003;**1**(2):159–188. doi:10.1093/jjfinec/nbg011.
44. Ogata Y. Space-time point-process models for earthquake occurrences. *Ann Inst Stat Math*. 1998;**50**(2):379–402. doi:10.1023/a:1003403601725.
45. Daley DJ, *et al*. *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods*. Springer: New York; 2002.
46. McCullagh P, *et al*. *Generalized Linear Models*. 2nd ed. Boca Raton: Routledge; 2019. doi:10.1201/9780203753736.
47. Zhang GQ, *et al*. The national sleep research resource: towards a sleep data commons. *J Am Med Inform Assoc*. 2018;**25**(10):1351–1358. doi:10.1093/jamia/ocy064.
48. Chen X, *et al*. Racial/ethnic differences in sleep disturbances: the Multi-Ethnic Study of Atherosclerosis (MESA). *Sleep*. 2015;**38**(6):877–888. doi:10.5665/sleep.4732.
49. Fujisawa S, *et al*. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci*. 2008;**11**(7):823–833. doi:10.1038/nn.2134.
50. Yalciner G, *et al*. Association of sleep time in supine position with apnea-hypopnea index as evidenced by successive polysomnography. *Sleep Breath*. 2017;**21**(2):289–294. doi:10.1007/s11325-016-1401-5.
51. Siddiqui F, *et al*. Half of patients with obstructive sleep apnea have a higher NREM AHI than REM AHI. *Sleep Med*. 2006;**7**(3):281–285. doi:10.1016/j.sleep.2005.10.006.
52. Fawcett T. An introduction to ROC analysis. *Pattern Recognit Lett*. 2006;**27**(8):861–874. doi:10.1016/j.patrec.2005.10.010.
53. Muschelli J. ROC and AUC with a binary predictor: a potentially misleading metric. *J Classif*. 2020;**37**(3):696–708. doi:10.1007/s00357-019-09345-1.
54. Won CHJ, *et al*. Sex differences in obstructive sleep apnea phenotypes, the multi-ethnic study of atherosclerosis. *Sleep*. 2020;**43**(5). doi:10.1093/sleep/zsz274.
55. Eiseman NA, *et al*. The impact of body posture and sleep stages on sleep apnea severity in adults. *J Clin Sleep Med*. 2012;**8**(6):655–66A. doi:10.5664/jcsm.2258.
56. Menon A, *et al*. Influence of body position on severity of obstructive sleep apnea: a systematic review. *ISRN Otolaryngol*. 2013;**2013**:670381. doi:10.1155/2013/670381.
57. Shahveisi K, *et al*. Sleep architecture in patients with primary snoring and obstructive sleep apnea. *Basic Clin Neurosci*. 2018;**9**(2):147–156. doi:10.29252/NIRP.BCN.9.2.147.
58. McSharry DG, *et al*. Physiological mechanisms of upper airway hypotonia during REM sleep. *Sleep*. 2014;**37**(3):561–569. doi:10.5665/sleep.3498.
59. Joosten SA, *et al*. Supine position related obstructive sleep apnea in adults: pathogenesis and treatment. *Sleep Med Rev*. 2014;**18**(1):7–17. doi:10.1016/j.smrv.2013.01.005.
60. O'Connor C, *et al*. Gender differences in the polysomnographic features of obstructive sleep apnea. *Am J Respir Crit Care Med*. 2000;**161**(5):1465–1472. doi:10.1164/ajrccm.161.5.9904121.
61. Ratnavadivel R, *et al*. Marked reduction in obstructive sleep apnea severity in slow wave sleep. *J Clin Sleep Med*. 2009;**5**(6):519–524.
62. Svanborg E, *et al*. EEG frequency changes during sleep apneas. *Sleep*. 1996;**19**(3):248–254. doi: 10.1093/sleep/19.3.248
63. Walsleben JA, *et al*. The utility of topographic EEG mapping in obstructive sleep apnea syndrome. *Sleep*. 1993;**16**(8 Suppl):S76–S78. doi:10.1093/sleep/16.suppl_8.S76
64. Kramer MA, *et al*. *Case Studies in Neural Data Analysis: A Guide for the Practicing Neuroscientist*. Cambridge, MA, USA: MIT Press; 2016.
65. Santner TJ, *et al*. *The Statistical Analysis of Discrete Data*. New York: Springer Science & Business Media; 2012.
66. Pawitan Y. *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. Oxford: OUP; 2001.
67. Kass RE, *et al*. *Analysis of Neural Data*. New York, NY: Springer New York; 2014. doi:10.1007/978-1-4614-9602-1
68. Brown EN, *et al*. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput*. 2002;**14**(2):325–346. doi:10.1162/08997660252741149.
69. Sarmashghi M, *et al*. Efficient spline regression for neural spiking data. *PLoS One*. 2021;**16**(10):1–19. doi:10.1371/journal.pone.0258321.
70. Spencer E, *et al*. A procedure to increase the power of Granger-causal analysis through temporal smoothing. *J Neurosci Methods*. 2018;**308**:48–61. doi:10.1016/j.jneumeth.2018.07.010.

71. Farhoodi S, *et al*. The problem of perfect predictors in statistical spike train models. *Neurons Behav Data Anal Theory*. 2021;**5**(3):1–20. doi:10.51628/001c.27667.
72. Frank LM, *et al*. Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: an adaptive filtering approach. *J Neurosci*. 2002;**22**(9):3817–3830. doi:10.1523/JNEUROSCI.22-09-03817.2002.
73. Eden UT, *et al*. Using point process models to describe rhythmic spiking in the subthalamic nucleus of Parkinson's patients. In: 2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society; 2011:757–760. doi:10.1109/IEMBS.2011.6090173
74. Meng L, *et al*. A unified approach to linking experimental, statistical and computational analysis of spike train data. *PLoS One*. 2014;**9**(1):1–10. doi:10.1371/journal.pone.0085269.
75. Weber K. *Advances in Point Process Modeling: Feature Selection, Goodness-of-Fit and Novel Applications* [Unpublished doctoral dissertation]. Boston, MA: Boston University; 2018.