HHS Public Access

Author manuscript

IEEE Trans Biomed Eng. Author manuscript; available in PMC 2021 September 08.

Published in final edited form as:

IEEE Trans Biomed Eng. 2021 September ; 68(9): 2833–2845. doi:10.1109/TBME.2021.3071366.

A Model-Based Framework for Assessing the Physiologic Structure of Electrodermal Activity

Sandya Subramanian [Student Member, IEEE],
Harvard-MIT Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Patrick L. Purdon [Member, IEEE],
Massachusetts General Hospital Department of Anesthesia, Critical Care, and Pain Medicine and Harvard Medical School, USA.

Riccardo Barbieri [Senior Member, IEEE],
Politecnico di Milano Department of Electronics, Informatics and Engineering, Italy and also with the Massachusetts General Hospital Department of Anesthesia, Critical Care, and Pain Medicine, USA.

Emery N. Brown [Fellow, IEEE]
Massachusetts Institute of Technology Picower Institute for Learning and Memory, USA and the Institute for Medical Engineering and Science, Massachusetts General Hospital Department of Anesthesia, Critical Care, and Pain Medicine, USA, and also with Harvard Medical School, USA.

Abstract

Objective: We present a statistical model for extracting physiologic characteristics from electrodermal activity (EDA) data in observational settings.

Methods: We based our model on the integrate-and-fire physiology of sweat gland bursts, which predicts inverse Gaussian (IG) inter-pulse interval structure. At the core of our model-based paradigm is a subject-specific amplitude threshold selection process for EDA pulses based on the statistical properties of four right-skewed models including the IG. By performing a sensitivity analysis across thresholds and fitting all four models, we selected for IG-like structure and verified the pulse selection with a goodness-of-fit analysis, maximizing capture of physiology at the time scale of EDA responses.

Results: We tested the model-based paradigm on simulated EDA time series and data from two different experimental cohorts recorded during different experimental conditions, using different equipment. In both the simulated and experimental data, our model-based method robustly recovered pulses that captured the IG-like structure predicted by physiology, despite large differences in noise level. In contrast, established EDA analysis tools, which attempted to

Corresponding author: Sandya Subramanian. sandya@mit.edu.

This article has supplementary downloadable material available at https://doi.org/10.1109/TBME.2021.3071366, provided by the authors.

estimate neural activity from slower EDA responses, did not provide physiological validation and were susceptible to noise.

Conclusion: We present a computationally efficient, statistically rigorous, and physiology-informed paradigm for pulse selection from EDA data that is robust across individuals and experimental conditions, yet adaptable to varying noise level.

Significance: The robustness of the model-based paradigm and its physiological basis provide empirical support for the use of EDA as a clinical marker for sympathetic activity in conditions such as pain, anxiety, depression, and sleep states.

Keywords

Electrodermal activity; inverse Gaussian; point processes; signal processing; statistics

I. INTRODUCTION

SWEAT gland activity is a fundamental component of our “fight or flight” response. Therefore, it is used as a way to assess sympathetic nervous system activity [1]. Sympathetic activation is induced by emotional and physiologic states such as stress, anxiety, and pain. Electrodermal activity (EDA) measures the electrical conductance of the skin as a proxy for the activity of sweat glands [1]. As sweat glands are more active, due to physiologic or emotional stimulation, the electrical conductance of the skin increases since sweat conducts electricity [1]. The continuous measurement of EDA can be performed inexpensively and non-invasively [1]. For this reason, it has been employed in lie detector tests and as a neuromarketing tool. There is growing interest in the development of real-time algorithms to accurately characterize EDA to track changes in emotional and physiologic states.

EDA consists of two simultaneous levels of activity. The baseline or tonic component represents background or ambient conditions and drifts gradually over minutes [1]. On top of that, pulsatile sweat release events (pulses) make up what is known as the phasic component of EDA, which has a timescale of a few seconds and is thought to correspond to sympathetic nervous system activity [1].

EDA has most commonly been studied in controlled experiments, in which subjects are presented with known stimuli at specific times in an environment with no distractions [2]-[6]. In these circumstances, the times of stimulus presentation act as ‘ground truth’ for sympathetic activation and EDA pulses. However, in many other circumstances, including clinical settings and during longitudinal monitoring, there is no known ground truth for EDA pulses. These are the circumstances under which we would also expect high inter-subject variability, differences due to different recording equipment, and changing situational conditions, since the environment cannot be controlled. Therefore, it is especially important to have a framework in place that establishes the physiological validity of the results of EDA analysis from individual EDA datasets.

Existing EDA analysis tools such as Ledalab [2], [3] and cvxEDA [4] assume that EDA obeys a convolution model: neural point process events are convolved with activity in the sweat gland. This is a challenging approach since neural activity occurs on a time scale

(milliseconds) much faster than EDA pulses (seconds). Therefore, estimating neural activity is an under-constrained optimization problem which to solve requires making assumptions about the size and shape of pulsatile sweat responses in EDA across all subjects, equipment, and conditions. In addition, these existing tools have been validated primarily in controlled experimental datasets. Therefore, existing EDA analysis tools do not offer validation grounded in physiology for EDA data when there is no known ground truth.

Such a validation framework can be developed from our previous work demonstrating how the physiology of sweat glands manifests as statistical structure in observational EDA data. We have previously shown that the inter-pulse interval distribution in EDA data follows an inverse Gaussian distribution [7][8], which agrees with a model of the rise of sweat through the gland to the skin surface as an integrate-and-fire process [9][10]. This process is similar to the mechanism that underlies other point process events such as neuronal and cardiac action potentials, earthquakes, geysers, and volcanoes [11][12]. We showed that the temporal structure in EDA favors right-skewed heavy tailed distributions, including the inverse Gaussian and even heavier tailed distributions such as the lognormal, due to the presence of sparse regions of EDA with low activity and long inter-pulse intervals [7]. This result makes possible the use of low-order models in EDA analyses and increases the signal-to-noise ratio. Most importantly, it provides physiological insight at the time scale of EDA responses to validate results of EDA analysis even without an explicit ground truth.

Our previous analyses suggest that the relationship between the properties of several right-skewed distributions can help identify the presence of noise in any subject cohort [7]. While the heavier-tailed inverse Gaussian and lognormal distributions are indicators of desired statistical structure, the flexibility of the gamma distribution in capturing the light tails likely to occur with random noise in a subject cohort allows it to function as a noise indicator [7]. By examining the patterns of goodness-of-fit of four distributions (inverse Gaussian, lognormal, gamma, exponential) across subjects in a cohort, we can gauge the level of noise in the subject cohort and the set of pulses to extract. We can also identify individual subjects who are clear outliers from the rest of the cohort.

In this work, we define a robust method for analyzing EDA data from any observational EDA subject cohort. The crux of our model-based approach is to extract the important physiological characteristics relevant to dynamic sympathetic activity by focusing on the EDA pulses themselves, without making additional assumptions necessary to make inferences at a much faster time scale. We show that this method is successful in capturing the underlying statistical structure in simulated data, as well as in data from two subject cohorts collected from different populations, under different experimental conditions, and using different equipment. We also show that existing deconvolution-based EDA tools cannot be validated against known physiology in these same subject cohorts. Our method combines the steps of extracting pulses and identifying the relevant statistical structure into an iterative rather than sequential process. This is analogous to what has been done in the case of ‘clusterless decoding’ for spike sorting and decoding in neural spiking data with improved results [13]. In the case of EDA, this iterative process is successful because the underlying model is consistent with physiology.

This paper is organized as follows. In Data, we describe: 1) how we generated literature-standard simulated EDA data; and 2) the two experimental subject cohorts used in this work. In Methods, we outline: 1) how we used our insights about sweat gland physiology to design and test a robust pulse selection paradigm; and 2) how we applied existing EDA analysis tools to two subject cohorts of observational EDA data. In Results, we illustrate the use of this paradigm on simulated data and data from two different subject cohorts. We compare our model-based approach to existing EDA analysis tools in terms of agreement with known physiology. Finally, in the Discussion and Conclusion, we summarize our model-based framework and the implications. We have made the code open source, enabling others to use this methodology and further refine the technique. A preliminary version of this model-based methodology was published in [14].

II. Data

A. Simulated EDA Data

We generated simulated EDA based on the literature standard used by the cvxEDA algorithm [4]. However, we modified the placement of pulses by sampling inter-pulse intervals from an inverse Gaussian distribution rather than random selection. This agrees with our previous discovery based on sweat gland physiology [7]. We also modified the noise level in the simulated EDA data to represent varying degrees of sensor noise. Each simulated EDA dataset was 1-hour in duration.

B. Experimental Data

For this study, we used EDA recordings from two experiments. The first is EDA we previously collected from 12 healthy volunteers between the ages of 22 and 34 while awake and at rest. The study was approved by the Massachusetts Institute of Technology (MIT) Committee on the Use of Humans as Experimental Subjects (approved 05/09/2018, protocol #1804343699A001). Approximately one hour of EDA data was collected at 256 Hz from electrodes connected to the second and fourth digits of each subject's non-dominant hand. Subjects were seated upright and instructed to remain awake. They were allowed to read, meditate, or watch something on a laptop or tablet, but not to write with the instrumented hand. One subject's data were not included in the analysis because we learned after completing the experiment that the subject occasionally experienced a Raynaud's type phenomenon, which would affect the quality of the EDA data. Data from the remaining 11 subjects were analyzed.

The second subject cohort consists of EDA recorded during a study of propofol-induced unconsciousness from eleven healthy volunteers while under propofol sedation [15]. The protocol was approved by the Massachusetts General Hospital (MGH) Human Research Committee. For all subjects, approximately 3 hours of data were recorded while using target-controlled infusion protocol. The data collection is described in detail in [15]. The infusion rate was increased and then decreased in a total of ten stages of roughly equal lengths to achieve target propofol concentrations of: 0 mg/kg/hr, 1, 2, 3, 4, 5, 3.75, 2.5, 1.25, 0. The two subject cohorts were collected in different years by different study teams and using different recording equipment systems.

III. Methods

A. The Model-Based Framework: Preprocessing and Pulse Extraction

All analyses were performed using Matlab 2019a. Preprocessing consisted of two major steps, 1) detecting and removing artifacts and 2) isolating the phasic component. Both have been described in previously in [7] and [14]. Because of the level of high frequency noise seen in the recording equipment used for the propofol data, those data were additionally low-pass filtered with a cutoff of 3 Hz after artifact removal.

We used a measure of locally-adjusted peak amplitude called prominence to account for the changing background filling level of the sweat glands. To compute prominence for each peak, we used the *findpeaks* algorithm in Matlab. This algorithm adjusts the amplitude of each peak in the signal as the height above the highest of neighboring "valleys" on either side. The valleys are chosen based on the lowest point in the signal between the peak and the next intersection with the signal of equal height on either side.

B. Motivation for Physiology-Informed Pulse Selection

The goal of pulse selection from EDA data is to extract the set of pulses as close as possible to the ground truth of true EDA phasic activity without including sensor noise. Therefore, relying on known properties of sweat gland physiology is key to distinguishing between pulses and noise. We previously showed that the bursting of sweat glands can be modeled as an integrate-and-fire process, which manifests as temporal structure that looks inverse Gaussian or like mixtures of inverse Gaussians in the inter-pulse intervals [7][8]. One way to characterize this is using tail behavior, which captures the heaviness of the tail of a right-skewed distribution using the ultimate settling rate. The ultimate settling rate is computed as the limit of the hazard function as x tends to infinity [16]. Different right-skewed distributions, such as the inverse Gaussian, lognormal, gamma, and exponential, differ in their tail behavior properties [17]. A slower settling rate indicates a heavier tail. Fitting a variety of models allows us to sample a range of tail behaviors (both lighter and heavier) which may represent mixtures of inverse Gaussians. Our previous work showed that EDA data favors heavier tailed models such as the lognormal and inverse Gaussian over lighter tailed models such as the gamma and exponential, likely due to the presence of regions of low phasic activity with long inter-pulse intervals [7]. Using these insights, we proceeded with the analysis in three phases for each subject cohort, detailed in the following sections. Overall, our model-based approach aims to define a novel way of assessing to what degree pulses extracted from EDA data are consistent with known physiology.

C. Phase I: Individual Subject Trends

We screened a range of prominence thresholds from 0.001 to 0.02 in increments of 0.001 and 0.02 to 0.08 in increments of 0.005 for each subject in the cohort. These limits were chosen to span from the smallest detectable deflection to the largest pulses in EDA data. At each prominence threshold, a set of pulses was extracted, from which four inter-pulse interval models were fit (the inverse Gaussian, lognormal, gamma, and exponential) by maximum likelihood [18]. For each, the Kolmogorov-Smirnov (KS) distance was computed as a measure of goodness-of-fit. The KS distance measures the maximum distance between

the theoretical and empirical inter-pulse intervals after they have been rescaled using the time-rescaling theorem [19]. A smaller KS distance indicates a better fit. We also computed a 95% significance cutoff for the KS distance based on the number of pulses extracted [20]. A KS distance greater than this cutoff indicates that the data differ significantly from the model. This process resulted in four model KS-distances and one significance cut-off at each prominence threshold for each subject.

D. Phase II: Subject cohort As a Whole

We plotted the median KS-distance and significance cut-off per prominence threshold across subjects for the entire cohort. The goal of Phase II was to understand the effects of data collection settings, such as equipment used and experimental condition studied, on the EDA recordings from the subject cohort overall. We were specifically interested in the role of each of the four models to identify four specific trends which relate to the tail behavior of the four models:

1. The inverse Gaussian and/or lognormal models reach a “sweet spot” of being the best fitting distributions for some region of prominence thresholds. These models generally do not fit well at the lowest thresholds due to the presence of noise. The size of this “sweet spot” can vary from subject to subject but indicates that the pulse extraction is preserving the desired statistical structure without too much noise. This is the ideal region from which to select the optimal prominence threshold.

2. The exponential model is generally the worst fitting model for the majority of thresholds, since it does not capture the physiology of sweat glands. However, with increasing threshold, the number of pulses decreases and therefore the significance cutoff becomes more generous. If the KS distance of the exponential model is above the significance cutoff at lower thresholds, the point at which it crosses under the significance cutoff marks the threshold at which there are now too few pulses to draw any conclusions. If the exponential model is under the significance cutoff even at lower thresholds (yet it is still the worst fitting model), this rule cannot be used to assess whether the number of pulses is sufficient.

3. The gamma model can be used to identify noise. The gamma model generally fits best in the presence of strong noise, usually at lower thresholds. This is because the presence of noise generally makes inter-pulse intervals short, which favors light-tailed inter-pulse interval distributions. The gamma is the model with the most flexibility in the tail and is the only one of the four models able to capture these very light tails. However, as noise is removed with increasing threshold, the KS-distance of the gamma model increases and “crosses” the lognormal and/or inverse Gaussian, as they fit the data better. We hereby refer to this trend (fitting well at lower thresholds and then slowly “rising out” with increasing threshold) as the “gamma signature”.

4. The gamma signature may be shifted left or right in a subject cohort, indicating the presence of lower or higher levels of noise overall, usually due to recording equipment or experimental conditions. For example, the gamma model may start

with a relatively poor fit at lower thresholds, then improve to being the best model before slowly rising out. This is simply the gamma signature shifted to the right, which indicates that the subject cohort has higher levels of noise. The "sweet spot" for optimal thresholds will also be at higher thresholds for the subjects in the cohort.

We illustrate the methods with an analysis of a single subject for three different thresholds and the four candidate models (Fig. 2). At the lowest threshold (Fig. 2(a), there are many short inter-pulse intervals due to the presence of noise in pulse selection, and therefore the inter-pulse interval distribution has almost no tail. The gamma model is best at accommodating this property. At moderate thresholds with less noise (Fig. 2(b)), the inverse Gaussian and lognormal are the best fit models since the inter-pulse interval distribution has a heavier tail. And finally, at the highest threshold (Fig. 2(c)), pulse extraction clearly excludes several pulses, and there are too few pulses to accurately characterize the inter-pulse interval distribution. The significance cutoffs are very generous, making all of the models appear to fit well. For each subject cohort, once each of these 4 trends had been identified, we proceeded to Phase III.

E. Phase III: Optimal Threshold Selection Per Subject

The KS-distances and significance cut-offs across prominence thresholds computed in Phase I were plotted per subject. The trends were then compared to the overall cohort-level trends observed in Phase II. For a given subject, if all models behaved similarly across prominence thresholds to the overall cohort in terms of goodness-of-fit with respect to the significance cutoff, the optimal prominence threshold was chosen from the "sweet spot" identified in Phase II, in which the inverse Gaussian and lognormal models are the best fitting.

If the models did not behave similarly to the cohort-level trends identified in Phase II, the subject was considered a potential outlier. In some cases, the same trends were seen, but shifted to the right (or left), which was an indication of higher (or lower) levels of noise in that subject's data compared to the rest of the cohort in general. In those cases, the optimal threshold for that subject is likely to be much higher (or lower) than for the rest of the cohort. Finally, after selecting the optimal threshold for each subject, we made a full KS plot with 95% confidence bounds to verify that our pulse extraction and choice of threshold did indeed capture the statistical structure as intended. If the KS plot of a model follows the 45-degree diagonal and stays fully within the confidence bounds, then that model fits the data well. We computed the AIC and KS-distance for all models for each subject as well as the tail settling rate [16].

F. Testing the Model-Based Framework on Simulated Data

We performed 50 test runs on simulated EDA data. We used 10 distinct combinations of inverse Gaussian distribution parameters and noise levels. For the inverse Gaussian models, the range of μ was between 15 and 30 while the range of λ was between 5 and 15. Three distinct noise levels were used, with standard deviations of 0.005, 0.01, and 0.03. Thirty of the runs were at the middle noise level, and the remaining 20 runs were split evenly between the lowest and highest noise levels. The different inverse Gaussian models were included to test whether the model-based framework is robust across physiological and

environment-dependent variation. The different noise levels were included to test whether the model-based framework is robust against varying levels of sensor noise. For each of the 10 distinct combinations, we ran 5 test runs. Since these were simulated data, no cohort-level analyses were done. The final prominence threshold was chosen for each test run according to the same principles discussed for the behavior of each of the four models (inverse Gaussian, lognormal, gamma, exponential). For each test run, we recorded the true number of EDA pulses, the final prominence threshold selected within the model-based framework, the number of pulses extracted by the model-based framework, the number of pulses matched in time between the true and extracted pulses, recall (proportion of true EDA pulses extracted by the model-based framework), precision (proportion of pulses extracted by the model-based framework which were true pulses), and maximum likelihood estimates of inverse Gaussian model parameters with 99% confidence intervals [18].

G. Comparison to Existing EDA Analysis Tools

We extracted pulses from the awake and at rest cohort using an existing EDA analysis pipeline from the Leda Lab [2][3]. This algorithm is based on a deconvolution approach, in which EDA is assumed to be the result of the convolution of a neural input signal with a single impulse response function for each subject, which represents that subject's sweat gland response to a single neural impulse. By recommendation of the authors, we used a continuous deconvolution analysis with the default parameters provided, which includes a peak threshold of 0.001 for all data. For each subject, after pulse extraction, we made a KS plot with 95% confidence bounds to assess whether the extracted pulses recovered statistical structure representative of known physiology.

We performed a qualitative comparison with the cvxEDA algorithm for both subject cohorts, since this algorithm does not provide guidance for pulse extraction from the phasic component [4]. This algorithm also follows a deconvolution approach to compute the inferred neural input from the phasic component. We used all default parameter values provided.

IV. Results

A. Simulated EDA Data

The results of each of the 50 runs of simulated EDA data are detailed in Table S-I, while the summary, including by noise level and overall is included in Table S-II. The overall results showed an average threshold of [0.033, 0.047], an average recall of [0.940, 0.962], and average precision of [0.981 ± 0.006]. This means that across runs, 94% of true pulses were recovered and 98% of recovered pulses were true pulses. However, looking at the breakdown by noise level, it is clear that the final threshold adjusted to the noise level to maintain precision and recall. The average threshold at the highest noise level was significantly higher than at either lower noise level. In terms of the estimates of the inverse Gaussian model parameters, the 99% confidence intervals captured the true value of μ in 96% of runs and the true value of λ in 74% of runs.

B. Awake and at Rest Subject Cohort Phases I and II

The overall goodness-of-fit curves for the four models for the awake and at rest cohort (Fig. 3) show four trends previously mentioned. The lognormal becomes the best-fitting model around a threshold of 0.004, whereas the inverse Gaussian crosses the gamma around a threshold of 0.008. After this point, the lognormal and inverse Gaussians remain the best fitting models. The exponential starts out as and remains the worst-fitting model throughout, except for a very small region near a threshold of 0.017. In a small number of cases, the exponential can appear to perform better than the gamma even though it is a subclass of the gamma. This is likely due to the fact that maximizing the likelihood does not necessarily minimize the KS distance, as reflected by the occasional disagreement between AIC and KS distance, as discussed in [7]. The exponential crosses under the significance cutoff around a threshold of 0.01. The gamma distribution starts as the best-fitting model and crosses both the lognormal and inverse Gaussian. The gamma signature does not appear shifted to the right. Based on all four of these results, it seems that for the majority of subjects in this cohort, the optimal prominence threshold should be in the range of 0.004 to 0.01.

C. Awake and at Rest Subject Cohort Phase III

Subject S6 shows similar goodness-of-fit curves to the overall awake and at rest cohort (Fig. 4(b)). The gamma starts as the best-fitting model and then crosses the lognormal and inverse Gaussian models. The exponential is always the worst-fitting model and crosses below the significance cut-off around a prominence threshold of 0.02. The gamma signature is not shifted to the right. We chose an optimal prominence threshold of 0.004. At that threshold, the pulses selected include all larger pulses as well as some smaller pulses in regions of sparse activity (Fig. 4(c)). The final KS-plot (Fig. 4(d)) shows that the lognormal and inverse Gaussian models remain close to the 45-degree line and within 95% confidence bounds throughout, but neither the exponential nor the gamma model does.

Across the 11 subjects in the awake and at rest cohort, the final prominence thresholds ranged between 0.0025 and 0.027. Nine of these 11 subjects had optimal thresholds between 0.004 and 0.01, as suggested from the full cohort analysis in Phase I. The total number of pulses in the one-hour time window ranged between 97 and 324, including the distantly spaced smaller pulses (Table I). The KS-distance and AIC results largely agreed with each other across subjects. At the respective optimal prominence thresholds for each subject, the lognormal or inverse Gaussian was the best fitting model for all 11 subjects according to KS-distance and 10 out of the 11 subjects according to AIC (Tables I and II). The exponential was the worst of the four models tested for 10 of the 11 subjects according to AIC and KS-distance. Specifically, with respect to KS-distance, the inverse Gaussian and lognormal models were within the significance cutoff for all subjects, the gamma for 8, but the exponential was only within the cutoff for 3 out of 11 subjects. These results suggest that our method was able to extract pulses while preserving the desired statistical structure.

Looking at the tail behavior analysis (Table III), the settling rate of the inverse Gaussian model is always less than that of the exponential and gamma, predicting a heavier tail than the exponential and gamma, even though all three are commonly classified as medium-tailed distributions. The lognormal is commonly classified as a heavy-tailed distribution.

Therefore, our method selects for a heavier tail in pulse extraction, representing heavy-tailed inverse Gaussians or mixtures of them which can be approximated by the lognormal.

Subject S5 presented an interesting case in which the goodness-of-fit curves deviate noticeably from the overall cohort trends (Fig. 5(b)). There are two clear deviations. First, the inverse Gaussian never crosses below the gamma model in goodness-of-fit, so the 'sweet spot' is determined by the lognormal alone. Second, the gamma signature appears shifted to the right, since it starts with a poor fit, then becomes the best fit model before crossing the lognormal. This also shifts the sweet spot and choice of optimal prominence threshold to the right. We accordingly chose an optimal threshold of 0.027. These trends would indicate that this subject's data are characterized by a much higher level of noise than the rest of the cohort. This is also clear when looking at the data themselves (Fig. 5(c)). Interestingly, even though the KS distance selects for the lognormal over the gamma, the gamma is the best fit model according to AIC (Tables I and II). The tail behavior analysis indicates that the gamma has a much higher settling rate for this subject than any of the other models, suggesting that it is able to accommodate the higher degree of noise and therefore shorter inter-pulse intervals with a light tail.

D. Propofol Sedation Subject Cohort Phases I and II

The overall goodness-of-fit curves for all four models for the propofol sedation cohort (Fig. 6) show some similarities to the awake and at rest cohort and some key differences with respect to the four trends previously discussed. Here, the lognormal and inverse Gaussian are always the best fitting models, largely remaining under the significance cut-off throughout. The exponential is almost always the worst-fitting model, although it never crosses below the significance cutoff. The behavior of the gamma is the trend most distinctly different from that of the other cohort. The gamma is never the best-fitting model for this cohort nor is it ever under the significance cutoff. This reflects more stringent significance cutoffs and narrower 95% confidence bounds due to a longer duration of data (and therefore more data points) for each subject compared to the awake and at rest cohort. However, the shape of the gamma signature is still visible, just shifted to the right. The gamma starts with poor fit, then improves slightly to a local minimum around a threshold of 0.004 before slowly rising out. This suggests that the propofol sedation cohort has a much higher level of noise compared to the awake and at rest cohort, and therefore the region of optimal prominence thresholds would also be higher. Based on the Phase II analysis, the optimal prominence thresholds for most subjects in this cohort should be in the range of 0.015 to 0.055, where the gamma has risen out of its local minimum sufficiently and the exponential is still the worst fitting model. Because of the nature of the trends in this cohort, this range is very broad.

E. Propofol Sedation Subject Cohort Phase III

Subject P09 shows similar goodness-of-fit curves to the overall propofol sedation cohort (Fig. 7(b)). The lognormal and inverse Gaussian models are the best fitting throughout. The gamma model signature is shifed to the right slightly, with it reaching a local minimum around a threshold of 0.01 and then rising out, even crossing above the exponential model, which is overall the worst fitting model throughout. We chose an optimal prominence threshold of 0.04. At that threshold, the extracted pulses include all pulse-like activity by

visual inspection. Because of the level of noise inherent in the subject cohort overall, it is likely that the extracted pulses may include some noise. However, there are clear regions with no pulses that are still excluded correctly, such as between 20 and 60 minutes (Fig. 7(c)). In addition, the small zoom-in shows that the extracted pulses do actually correspond to pulse-like activity and not noise (Fig. 7(c)). The final KS-plot (Fig. 7(d)) shows that the lognormal and inverse Gaussian models stay close to the 45-degree line throughout while the gamma and exponential models do not.

Across the 11 subjects in the propofol sedation cohort, the final prominence thresholds used ranged between 0.02 and 0.055 (Table IV), which is fully within the range suggested from the cohort analysis in Phase I. The total number of pulses in the 3 to 4-hour time window ranged between 383 and 1250, which also reflects the higher degree of noise in the cohort. At the respective optimal prominence thresholds for each subject, either the lognormal or inverse Gaussian was the best fitting model for all 11 subjects according to both AIC and KS-distance (Tables IV and S-III). The exponential was the worst of the four models tested for 8 of the 11 subjects according to AIC and KS-distance, with the gamma performing slightly worse for the other 3 subjects. Specifically, with respect to KS-distance, the inverse Gaussian and lognormal models were within the significance cutoff for all subjects, the gamma for only one, the exponential for none of the subjects. Based on the tail behavior analysis (Table S-IV), similar to the awake and at rest cohort, the settling rate of the inverse Gaussian model is always markedly less than that of the exponential and gamma. These results verify that our method was able to extract pulses while capturing the desired inverse Gaussian-like structure in the propofol sedation cohort, a second cohort with very different properties from the awake and at rest cohort.

F. Comparison to Existing EDA Analysis Tools

The Ledalab algorithm extracted almost an order of magnitude more pulses than the model-based framework for each subject in the awake and at rest cohort (Table S-III), likely due to the low default threshold value of 0.001 for pulses. In comparison, with the model-based framework, we chose a threshold at least 2.5 times greater for all subjects. Since this method does not adapt any parameters to the properties of the data at hand, there is no guarantee that pulse extraction will be consistent across subjects or cohorts. This was reflected in the statistical properties of the pulses extracted. The extracted pulses largely look like noise (Fig. 8(a)), and the goodness-of-fit analysis (Fig. 8(b)) indicates that for the majority of subjects (Table S-V, Figs. S20-S29), none of the four models offered an accurate statistical description of the data.

For the cvxEDA method, we could not compare extracted EDA pulses, since this algorithm does not provide guidelines for extracting pulses from the phasic component. Therefore, we performed qualitative comparison, shown in Fig. 9 for three representative subjects. Fig. 9A is an example subject for whom the model-based paradigm and cvxEDA yielded similar qualitative results, showing the estimated neural activity from cvxEDA aligned with extracted pulses from the model-based paradigm. Fig. 9(b) shows an example subject for whom cvxEDA yielded a qualitatively similar phasic component to the model-based paradigm. However, the estimate of neural activity, the ultimate output of cvxEDA, is much

noisier and masks the dynamics seen clearly in the extracted pulses from our model-based paradigm. Finally, Fig. 9(c) shows an example subject for whom both phasic activity and estimated neural activity were fully corrupted by noise with cvxEDA compared to the model-based paradigm. In all cases (Figs. S30-S48), there was no way to validate the estimate of neural activity from cvxEDA against known physiology.

V. Discussion

In this work, we defined a systematic and robust method to extract pulses from EDA data that preserved the statistical structure in the data derived from physiology while excluding noise. This method was used for EDA data in observational settings, without known, controlled experimental stimuli. This method also allowed for an assessment of the signal to noise profile across a whole cohort of data and identification of individual subjects whose data do not behave like the rest of the cohort. We tested this method on simulated EDA data as well as two cohorts, each collected using different equipment and under different experimental conditions. We showed that the method captures the statistical structure in all cases, including with a recall greater than 94% and precision greater than 98% on simulated EDA data (Tables II, IV and S-II). In each subject cohort, the goodness-of-fit properties inherently reflected the signal to noise ratio in that cohort (Figs 3 and 6). Finally, we compared our method to pulse selection using the Ledalab algorithm for EDA and qualitatively to the cvxEDA algorithm and showed that unlike our method, these did not perform consistently on new subject cohorts on which they were not built (Figs 8 and 9).

Our model-based framework for characterizing and analyzing EDA data provided a structured and computationally efficient way to extract the relevant physiological characteristics from the data. We used a total of seven parameters across four models (inverse Gaussian, lognormal, gamma, and exponential) to determine how to extract pulses. By doing so, we reduced the arbitrariness of assigning a threshold to the data simply based on visual inspection or assigning a single threshold across all datasets. The idea of a sensitivity analysis across thresholds reduced the arbitrariness of imposing assumptions about noise, pulse amplitude, or pulse shape across the cohort. Results from simulated EDA data showed that using statistical structure to inform threshold selection allowed the threshold to increase with noise level and preserve performance (Table S-II).

The continuous deconvolution-based Ledalab algorithm, on the other hand, assumed a single unchanging pulse amplitude threshold across all subjects and cohorts. This resulted in too generous of pulse selection, with an average of one pulse every two seconds across all subjects (mean of 1880 pulses in one hour) (Fig. 8(a), Table S-V). This was far too frequent given known sweat gland physiology and other studies of EDA [21]-[27]. In addition, the goodness-of-fit results indicated that none of the models was under the significance cut-off for the majority of subjects at this threshold (Fig. 8(b), Table S-V). This was similar to the threshold screening results of our method at low thresholds in high-noise data. Even if one were to have attempted to threshold the Ledalab pulse selection with a higher threshold, it would have been arbitrarily based on visual inspection alone. CvxEDA, like Ledalab, is a deconvolution-based algorithm. While we could not compare pulse selection with cvxEDA, since it does not have a step of selecting pulses, qualitative comparison of results still

showed inconsistencies across subjects and cohorts, with some clearly affected by noise (Figs 9, S30-S48).

The ultimate goal of these methodologies was to estimate the underlying neural input as a measure of sympathetic activity. However, these methodologies created a more challenging approach. By attempting to estimate a faster time resolution phenomenon (neural activity) that was not directly measured using a slower time resolution phenomenon (sweat gland activity), these methodologies were attempting to solve an under-constrained optimization problem. This could not be done without imposing a number of assumptions about pulse shape, pulse amplitude, and noise level across all datasets. These assumptions ignored variation between subjects, experimental conditions, and recording equipment. In addition, the estimated neural activity was highly sensitive to noise. In contrast, the model-based approach took advantage of the physiology of sweat glands themselves to uncover the necessary physiological characteristics of sympathetic activity without making such assumptions. By allowing the structure in the data to inform pre-processing and analysis, instead of passing the data through a fixed analysis pipeline, we could extract physiological characteristics at the time scale of the EDA response rather than the more challenging question of estimating faster neural activity. The fits we report in this work may be improved upon by use of dynamics models, which we will explore in our future work.

By developing an analysis pipeline that is computationally efficient, statistically rigorous, and physiology-based, we have enabled robust capture of relevant physiological characteristics from EDA data, even in the absence of a known ground truth. We can also track the goodness-of-fit of our models in any setting. This robustness is a significant step forward in allowing EDA to be a clinical marker for sympathetic activity in diverse conditions such as pain, anxiety, depression, and sleep. Our future work will investigate models for capturing the relevant information in the amplitudes of the pulses as well and applying dynamic models for both types of information. Data and code are available on the PhysioNet database [28]. We hope to further refine this methodology and release further versions of the code that incorporate pulse amplitude analysis and other modifications based on the feedback we receive from those who use it on their own data.

VI. Conclusion

We have developed a novel paradigm for analyzing EDA data that extracts pulses based on the statistical characteristics of the EDA time series. This model-based paradigm can be used for EDA data in observational settings, without a known ground truth or controlled stimuli. The model-based paradigm performs consistently on simulated EDA data and data from different subject cohorts. It verifies that the selected pulses reflect known sweat gland physiology, a process unique to the model-based paradigm. Unlike existing EDA analysis tools, which have been largely developed for controlled experimental data with ground truth, the model-based paradigm does not make assumptions about pulse shapes, amplitudes, or noise levels across data sets. It also does not seek to estimate the underlying neural activity, which occurs at a faster time scale than the recorded EDA response, and is therefore challenging to infer with accuracy or precision. Accurate pulse selection from EDA is key to inferring underlying sympathetic information at the time scale relevant to EDA responses.

Sympathetic activation is implicated in many physiologic and psychological states, including stress, anxiety, depression, sleep, and pain. Our physiology-based paradigm for robust pulse selection from EDA data opens the door for use of EDA as part of standard clinical measures to track these states.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported in part by the Picower Institute for Learning and Memory, the National Science Foundation Graduate Research Fellowship Program, and in part by the NIH Award P01-GM118629 (to E.N.B.). We would like to thank the staff of the Massachusetts Institute of Technology Clinical Research Center.

References

[1]. Boucsein W, “Principles of electrodermal phenomena,” in Electrodermal Activity. New York, NY, USA: Springer, 2012, pp. 20–105.

[2]. Benedek M, and Kaernbach C, “Decomposition of skin conductance data by means of nonnegative deconvolution,” Psychophysiology, vol. 47, pp. 647–658, 2010, doi: 10.1111/j.1469-8986.2009.00972.x. [PubMed: 20230512]

[3]. Benedek M, and Kaernbach C, “A continuous measure of phasic electrodermal activity,” J. Neurosci. Methods, vol. 190, pp. 80–91, 2010, doi: 10.1016/j.jneumeth.2010.04.028. [PubMed: 20451556]

[4]. Greco Aet al., “A convex optimization approach to electrodermal activity processing,” IEEE Trans. Biomed. Eng, vol. 63, pp. 797–804, 4. 2016, doi: 10.1109/TBME.2015.2474131. [PubMed: 26336110]

[5]. Faghih Ret al., “Characterization of fear conditioning and fear extinction by analysis of electrodermal activity,” in Proc. 37th IEEE Int. Conf. Eng. Med. Biol.. 2015, pp. 7814–7818.

[6]. Alexander Det al.., “Separating individual skin conductance responses in a short interstimulus-interval paradigm,” J. Neurosci. Methods, vol. 146, pp. 116–123, 2005, doi: 10.1016/j.jneumeth.2005.02.001. [PubMed: 15935228]

[7]. Subramanian Set al., “Point Process Temporal Structure Characterizes Electrodermal Activity,” PNAS, vol. 117, no. 42, pp. 26422–26428, 10. 2020. [PubMed: 33008878]

[8]. Subramanian Set al.., “A point process characterization of electrodermal activity,” in Proc. 40th IEEE Int. Conf. Eng. Med. Biol., 2018, pp. 37–40.

[9]. Schrodinger E, “Zur theorie der fall-und steigversuche an teilchen mit brownscher bewegung,” Physikalische Zeitschrift, vol. 16, pp. 289–295, 1915.

[10]. Gerstein G, and Mandelbrot B, “Random walk models for the spike activity of a single neuron,” Biophys. J, vol. 4, pp. 41–68, 1964, doi: 10.1016/s0006-3495(64)86768-0. [PubMed: 14104072]

[11]. Chhikara R, and Folks J, “The Inverse Gaussian Distribution: Theory, Methodology, and Applications,” New York, NY, USA: Marcel Dekker, 1989.

[12]. Barbieri Ret al., “A point-process model of human heartbeat intervals: New definitions of heart rate and heart rate variability,” Amer. J. Physiol. Hear. Circ. Physiol, vol. 288, pp. H424–H435, 2005, doi: 10.1152/ajpheart.00482.2003.

[13]. Deng Xet al., “Clusterless decoding of position from multiunit activity using a marked point process filter,” Neural Comput., vol. 27, no. 7, pp. 1438–1460, 2015. [PubMed: 25973549]

[14]. Subramanian Set al., “A systematic method for preprocessing and analyzing electrodermal activity,” in Proc. IEEE Int. Conf. Eng. Med. Biol., 2019, pp. 6902–6905.

[15]. Purdon PLet al., “Electroencephalogram Signatures of Loss and Recovery of Consciousness from Propofol,” PNAS, vol. 110, pp. E1142–E1151, 2013. [PubMed: 23487781]

[16]. Halliwell L, “Classifying the tails of loss distributions,” Casualty Actuarial Soc. E-Forum, vol. 2, 2013. [Online]. Available: https://www.casact.org/sites/default/files/database/forum_13spforumv2_haliwell.pdf

[17]. Folks J, and Chhikara R, “The inverse Gaussian distribution and its statistical application – A review,” J. Roy. Statist. Soc. Ser. B (Methodological), vol. 40, pp. 263–289, 1978, doi: 10.1111/j.2517-6161.1978.tb01039.x.

[18]. Pawitan Y, “Frequentist properities,” in In All Likelihood. Oxford, U.K: Clarendon Press, 2013, pp. 117–145.

[19]. Brown Eet al., “The time-rescaling theorem and its application to neural spike train data analysis,” Neural Comput., vol. 14, pp. 325–346, 2001, doi: 10.1162/08997660252741149.

[20]. Daley D, and Vere-Jones D, “Basic theory of random measures and point processes,” in An Int. Theory Point Processes: Vol. II: General Theory Structure, New York, NY, USA: Springer, 2007, pp. 2–76.

[21]. Storm Het al., “Palmar skin conductance compared to a developed stress score and to noxious and awakening stimuli on patients in anaesthesia,” Acta Anaesthesiol. Scand, vol. 49, pp. 798–803, 2005, doi: 10.1111/j.1399-6576.2005.00665.x. [PubMed: 15954962]

[22]. Hallin RG, and Torebjork HE, “Afferent and efferent C units recorded from human skin nerves in situ,” Acta Soc. Medicorum, vol. 75, pp. 277–281, 1970.

[23]. Hagbarth Ket al., “General characteristics of sympathetic activity in human skin nerves,” Acta Physiol. Scand, vol. 84, pp. 164–176, 1972, doi: 10.1111/j.1748-1716.1972.tb05167.x. [PubMed: 5015184]

[24]. Hallin RG, and Torebjork HE, “Single unit sympathetic activity in human skin nerves during rest and various manoeuvres,” Acta Physiol. Scand, vol. 92, pp. 303–317, 1974, doi: 10.1111/j.1748-1716.1974.tb05749.x. [PubMed: 4454989]

[25]. Wallin BG, “Sympathetic nerve activity underlying electrodermal and cardiovascular reactions in man,” Psychophysiology, vol. 18, pp. 470–476, 1981, doi: 10.1111/j.1469-8986.1981.tb02483.x. [PubMed: 7267931]

[26]. Wallin BG, “Peripheral sympathetic neural activity in conscious humans,” Annu. Rev. Physiol, vol. 50, pp. 565–576, 1988, doi: 10.1146/annurev.ph.50.030188.003025. [PubMed: 3288106]

[27]. Uncini Aet al.., “The sympathetic skin response: Normal values, elucidation of afferent components and application limits,” J. Neurol. Sci, vol. 87, pp. 299–306, 1988, doi: 10.1016/0022-510x(88)90254-7. [PubMed: 2850351]

[28]. Subramanian S, Barbieri R, and Brown E, “Electrodermal activity of healthy volunteers while awake and at rest (version 1.0),” PhysioNet, Deposited 4. Aug. 2020 10.13026/arty-2540.

[29]. Sano Aet al., “Quantitative analysis of wrist electrodermal activity during sleep,” Int. J. Psychophysiol, vol. 94, pp. 382–389, 2014, doi: 10.1016/j.ijpsycho.2014.09.011. [PubMed: 25286449]

[30]. Storm Het al., “Skin conductance correlates with perioperative stress,” Acta Anaesthesiol. Scand, vol. 46, pp. 887–895, 2002, doi: 10.1034/j.1399-6576.2002.460721.x. [PubMed: 12139547]

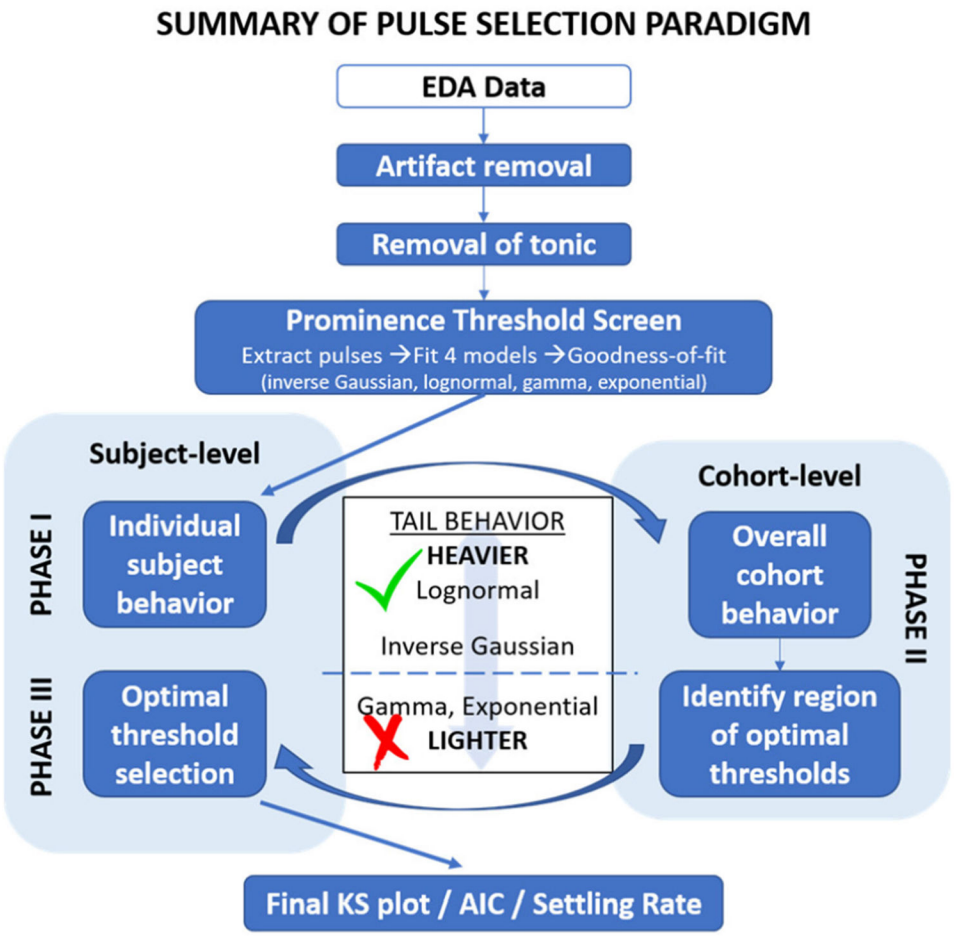

Fig. 1.

A schematic of the methods, starting from raw EDA data to verifying the goodness-of-fit of the chosen prominence threshold. In the center, the tail behavior properties of the 4 models are summarized.

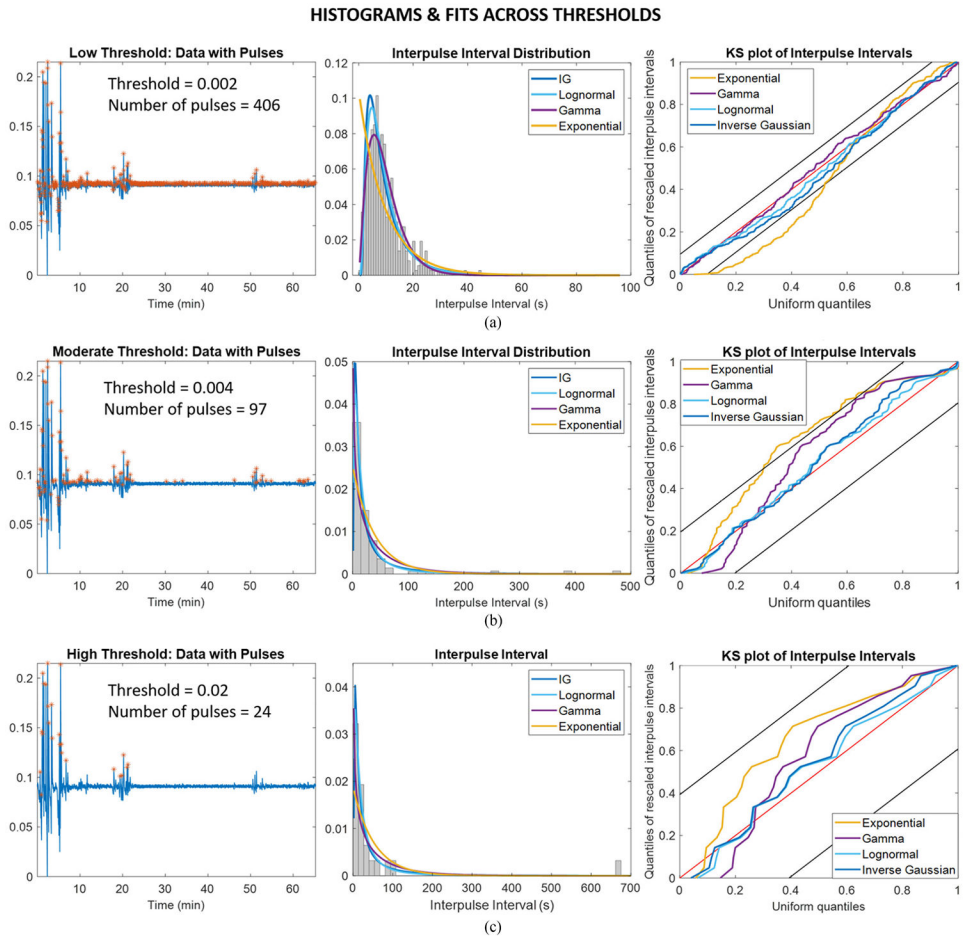

Fig. 2.

For a single subject's data (S6), the threshold for pulse selection was gradually increased. At each of three representative thresholds, (a) 0.002 (low), (b) 0.004 (moderate), and (c) 0.02 (high), the pulses extracted, shape of the inter-pulse interval distribution, and the goodness-of-fit of the 4 distributions are shown.

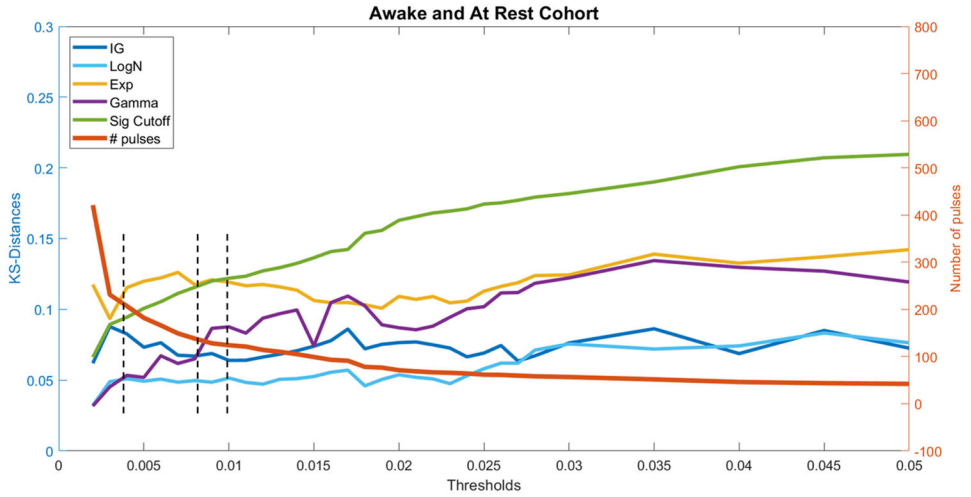

Fig. 3.

Overall goodness-of-fit curves for all 4 distributions (left y-axis) and number of pulses extracted (right y-axis) across the awake and at rest cohort, dashed vertical lines indicate, from left to right, (1) where the lognormal crosses under the gamma, (2) where the inverse Gaussian crosses under the gamma, and (3) where the exponential crosses under the significance cutoff.

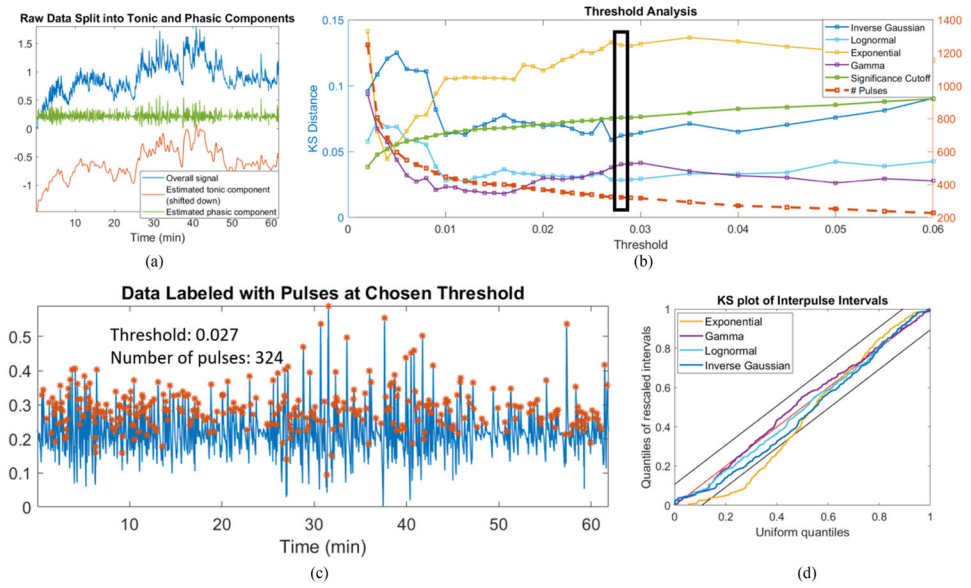

Fig. 4.

Results for Subject S5 from the awake and at rest cohort, showing noticeable differences from the trends of the cohort as a whole. (a) Preprocessing of data by splitting into tonic and phasic components, (b) Screening of thresholds with chosen threshold marked with bolded rectangle, (c) Pulses extracted at chosen threshold, (d) Full KS-plot showing goodness-of-fit at chosen threshold.

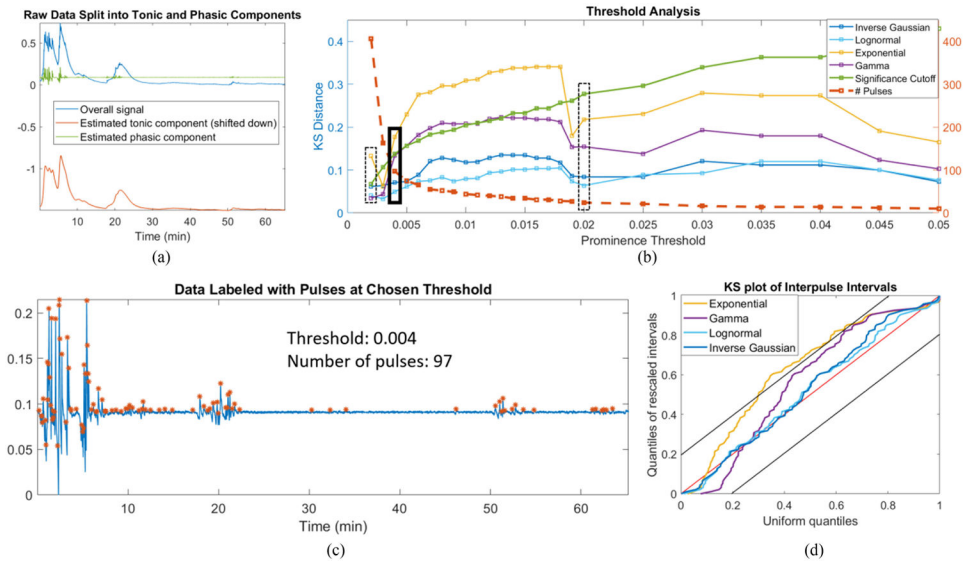

Fig. 5.

Results for Subject S6 from the awake and at rest cohort, showing agreement with the trends of the cohort as a whole. (a) Preprocessing of data by splitting into tonic and phasic components, (b) Screening of thresholds with chosen threshold marked with bolded rectangle, dashed rectangles indicate other thresholds (low and high) also shown in Fig. 1, (c) Pulses extracted at chosen threshold, (d) Full KS-plot showing goodness-of-fit at chosen threshold.

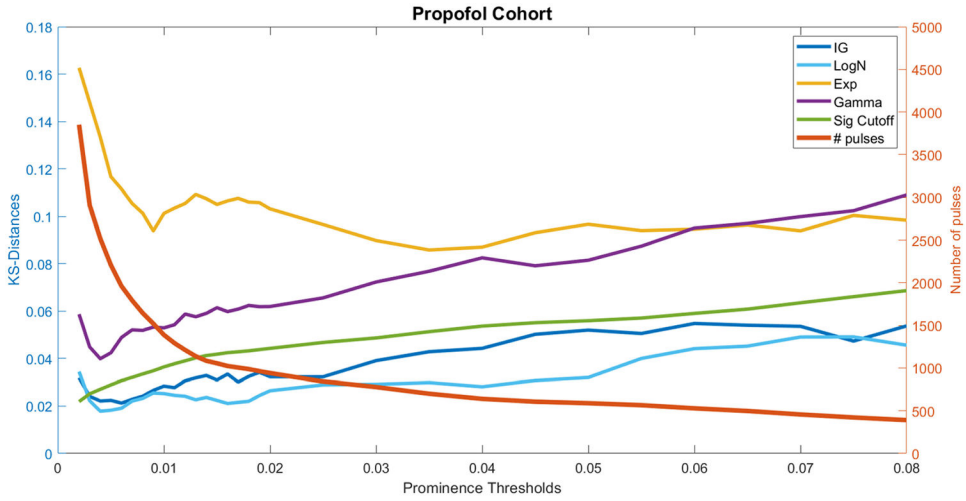

Fig. 6.

Overall goodness-of-fit curves for all 4 distributions (left y-axis) and number of pulses extracted (right y-axis) across the propofol sedation cohort.

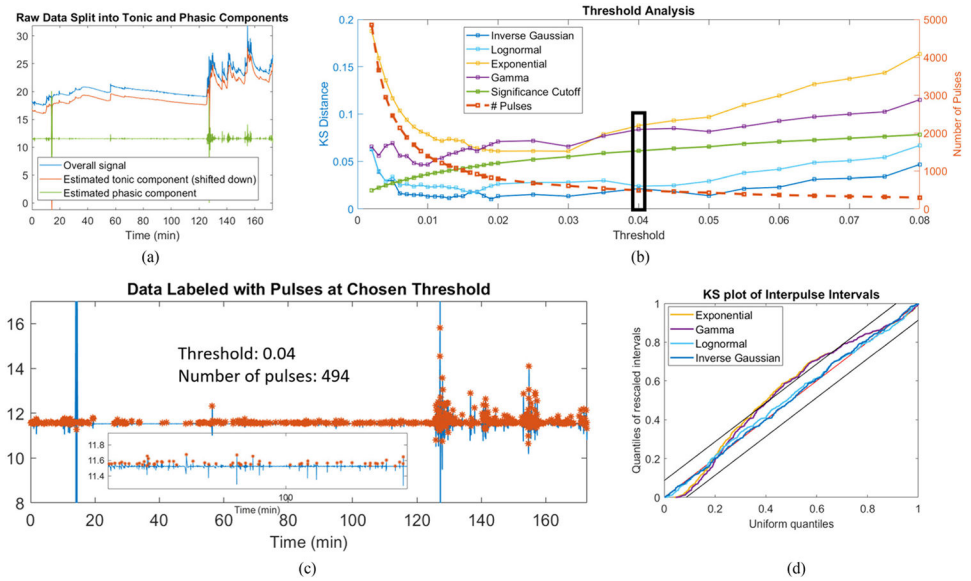

Fig. 7.

Results for Subject P8 from the propofol cohort, showing agreement with the trends of the cohort as a whole. (a) Preprocessing of data by splitting into tonic and phasic components, (b) Screening of thresholds with chosen threshold marked with bolded rectangle, (c) Pulses extracted at chosen threshold, (d) Full KS-plot showing goodness-of-fit at chosen threshold.

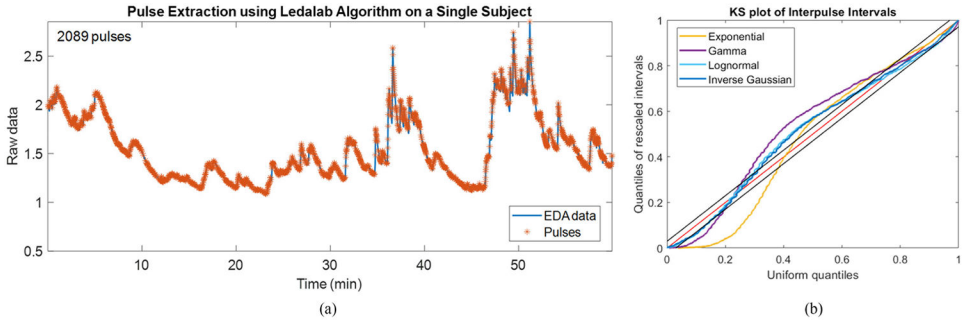

Fig. 8.

(a) Pulse selection and (b) goodness-of-fit results using the Ledalab algorithm for Subject S8 from the awake and at rest cohort.

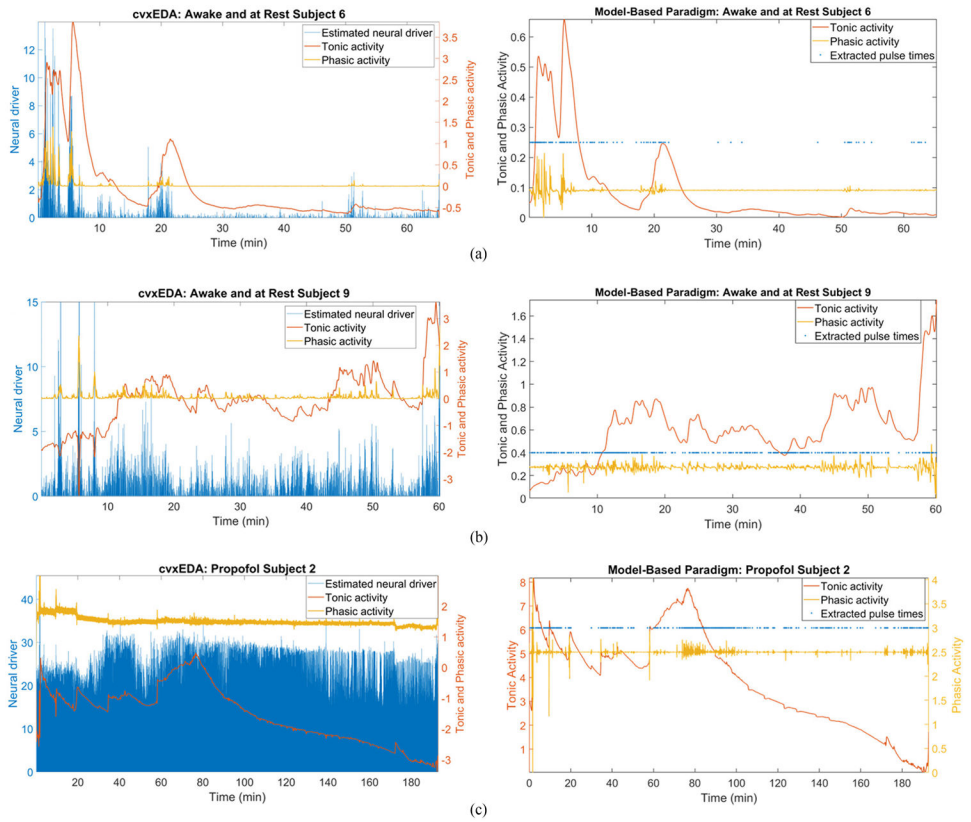

Fig. 9.

Results from the cvxEDA algorithm for (a) Subject S6 from the awake and at rest cohort, appearing similar qualitatively to the model-based framework, (b) Subject S9 from the awake and at rest cohort, appearing similar with respect to phasic activity to the model-based framework but yielding a very noisy estimate of neural activity compared to extracted pulses, and (c) Subject P2 from the propofol sedation cohort, appearing drastically different from the model-based framework and clearly corrupted by noise.

TABLE I

AIC Results for the Awake and At Rest Cohort.

	Thresh	Num. Dulses	IG	LogN	Gamma	Exp
S1	0.004	149	1156.494	**1148.853**	1222.312	1222.398
S2	0.004	283	**1904.077**	1909.798	1996.249	1996.701
S3	0.005	112	914.059	**896.049**	995.094	995.910
S4	0.008	170	1329.109	**1319.333**	1343.492	1370.524
S5	0.027	324	2183.330	2148.833	**2140.502**	2218.675
S6	0.004	97	**830.576**	832.572	888.536	900.063
S7	0.01	124	1050.422	**1046.177**	1068.990	1086.153
S8	0.008	269	**1846.806**	1849.936	1871.999	1928.629
S9	0.01	299	2100.344	**2033.972**	2041.055	2083.859
S10	0.0025	111	973.622	**973.547**	977.151	997.975
S11	0.006	112	**975.447**	977.170	1002.139	1002.841

The best model for each subject is indicated in bold.

Thresh = threshold, Num. pulses = number of pulses, IG = inverse Gaussian, LogN = lognormal, Exp = exponential.

TABLE II

KS-Distance Results for the Awake and At Rest Cohort.

	Num. pulses	Sig. Cutoff	IG	LogN	Gamma	Exp
S1	149	0.111	0.083*	**0.051***	0.121	0.115
S2	283	0.081	**0.055***	0.062*	0.0899	0.0902
S3	112	0.128	0.121*	**0.054***	0.159	0.169
S4	170	0.104	0.050*	**0.041***	0.063*	0.132
S5	324	0.075	0.059*	**0.029***	0.038*	0.133
S6	97	0.138	0.071*	**0.050***	0.134*	0.177
S7	124	0.122	0.064*	**0.053***	0.092*	0.121*
S8	269	0.083	**0.021***	0.025*	0.061*	0.117
S9	299	0.079	0.078*	**0.028***	0.051*	0.100
S10	111	0.129	0.049*	**0.042***	0.055*	0.109*
S11	112	0.128	**0.026***	0.031*	0.079*	0.088*

The best model for each subject is indicated in bold. Asterisks indicate that the model was under the significance cutoff.

Num. pulses = number of pulses, Sig. cutoff = significance cutoff, IG = inverse Gaussian, LogN = lognormal, Exp = exponential.

TABLE III

Settling Rate Results for the Awake and At Rest Cohort

	IG	LogN	Gamma	Exp
S1	0.019	0	0.051	0.044
S2	0.028	0	0.089	0.079
S3	0.009	0	0.026	0.031
S4	0.029	0	0.085	0.047
S5	0.052	0	0.181	0.088
S6	0.004	0	0.017	0.025
S7	0.020	0	0.058	0.033
S8	0.051	0	0.146	0.075
S9	0.035	0	0.142	0.083
S10	0.019	0	0.056	0.029
S11	0.011	0	0.037	0.030

IG = inverse Gaussian, LogN = lognormal, Exp = exponential.

TABLE IV

KS-Distance Results for the Propofol Sedation Cohort

	Thresh	Num. pulses	Sig. Cutoff	IG	LogN	Gamma	Exp
P1	0.035	727	0.0504	0.049*	**0.032***	0.078	0.107
P2	0.020	383	0.0694	**0.037***	0.039*	0.101	0.110
P3	0.035	762	0.0492	0.030*	**0.019***	0.069	0.068
P4	0.025	1010	0.0427	0.038*	**0.021***	0.080	0.077
P5	0.055	566	0.0571	0.053*	**0.044***	0.091	0.084
P6	0.020	838	0.0469	0.038*	**0.023***	0.059	0.089
P7	0.035	1250	0.0384	**0.025***	0.028*	0 070	0.086
P8	0.040	494	0.0611	**0.021***	0.024*	0.084	0.088
P9	0.055	575	0.0566	0.044*	**0.024***	0.063	0.083
P10	0.021	627	0.0542	**0.021***	0.038*	0.085	0.124
P11	0.030	778	0.0487	0.031*	**0.015***	0.036*	0.092

The best model for each subject is indicated in bold. Asterisks indicate that the model was under the significance cutoff.

Num. pulses = number of pulses, Sig. cutoff = significance cutoff, IG = inverse Gaussian, LogN = lognormal, Exp = exponential.