



HAL
open science

Bayesian at heart: Towards autonomic outflow estimation via generative state-space modelling of heart rate dynamics

Fernando E Rosas, Diego Candia-Rivera, Andrea I Luppi, Yike Guo, Pedro A.M. Mediano

► To cite this version:

Fernando E Rosas, Diego Candia-Rivera, Andrea I Luppi, Yike Guo, Pedro A.M. Mediano. Bayesian at heart: Towards autonomic outflow estimation via generative state-space modelling of heart rate dynamics. *Computers in Biology and Medicine*, 2023, 170, pp.107857. 10.1016/j.compbiomed.2023.107857 . hal-04855383

HAL Id: hal-04855383

<https://hal.science/hal-04855383v1>

Submitted on 24 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/376791413>

Bayesian at heart: Towards autonomic outflow estimation via generative state-space modelling of heart rate dynamics

Article in *Computers in Biology and Medicine* · December 2023

DOI: 10.1016/j.combiomed.2023.107857

CITATIONS

6

READS

77

5 authors, including:



Fernando E. Rosas
University of Sussex

234 PUBLICATIONS 4,683 CITATIONS

[SEE PROFILE](#)



Diego Candia-Rivera
Paris Brain Institute - ICM

47 PUBLICATIONS 454 CITATIONS

[SEE PROFILE](#)



Andrea Luppi
University of Cambridge

99 PUBLICATIONS 1,703 CITATIONS

[SEE PROFILE](#)



Bayesian at heart: Towards autonomic outflow estimation via generative state-space modelling of heart rate dynamics

Fernando E. Rosas^{a,b,c,d,*}, Diego Candia-Rivera^e, Andrea I. Luppi^{f,g}, Yike Guo^h, Pedro A.M. Mediano^{i,j}

^a School of Engineering and Informatics, University of Sussex, United Kingdom

^b Centre for Psychedelic Research, Department of Brain Science, Imperial College London, United Kingdom

^c Centre for Complexity Science, Imperial College London, London, United Kingdom

^d Centre for Eudaimonia and Human Flourishing, University of Oxford, United Kingdom

^e Sorbonne Université, Paris Brain Institute (ICM), INRIA, CNRS, INSERM, AP-HP, Hôpital Pitié-Salpêtrière, 75013, Paris, France

^f University Division of Anaesthesia, University of Cambridge, Cambridge, United Kingdom

^g Montreal Neurological Institute, McGill University, Montreal, Canada

^h Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong

ⁱ Department of Computing, Imperial College London, South Kensington, London, United Kingdom

^j Department of Psychology, University of Cambridge, Cambridge, United Kingdom

ARTICLE INFO

Keywords:

Autonomic nervous system

Heart rate

Bayesian statistics

State-space modelling

Non-linear methods

ABSTRACT

Recent research is revealing how cognitive processes are supported by a complex interplay between the brain and the rest of the body, which can be investigated by the analysis of physiological features such as breathing rhythms, heart rate, and skin conductance. Heart rate dynamics are of particular interest as they provide a way to track the sympathetic and parasympathetic outflow from the autonomic nervous system, which is known to play a key role in modulating attention, memory, decision-making, and emotional processing. However, extracting useful information from heartbeats about the autonomic outflow is still challenging due to the noisy estimates that result from standard signal-processing methods. To advance this state of affairs, we propose a novel approach in how to conceptualise and model heart rate: instead of being a mere summary of the observed inter-beat intervals, we introduce a modelling framework that views heart rate as a hidden stochastic process that drives the observed heartbeats. Moreover, by leveraging the rich literature of state-space modelling and Bayesian inference, our proposed framework delivers a description of heart rate dynamics that is not a point estimate but a posterior distribution of a generative model. We illustrate the capabilities of our method by showing that it recapitulates linear properties of conventional heart rate estimators, while exhibiting a better discriminative power for metrics of dynamical complexity compared across different physiological states.

1. Introduction

The autonomic nervous system (ANS) is responsible for the physiological adjustments necessary for the regulation of the body and contributes to appropriately responding to environmental demands. The ANS is composed of both afferent and efferent pathways that involve multiple neural structures including the vagus nerve, medulla, brainstem, insula and other cortical and sub-cortical regions [1], and the result of their interactions constitutes what is known as *autonomic outflow* [2,3]. Crucially, the impact of sympathetic and parasympathetic outflow goes far beyond driving homeostatic regulation and “flight or fight” responses, being also involved in cognitive processes such as attention, memory, decision-making, and emotional processing [4]. Thus,

measurement of the autonomic outflow can provide valuable insights into the physiological substrates that support human cognition [5].

A popular approach to investigate autonomic outflow is by analysing properties of how heart rate changes over time. More specifically, a large body of literature has focused on the investigation of *heart rate variability* (HRV) [6] – i.e. the patterns of variation between the time intervals between successive heartbeats – which is thought to reflect the changes in the sympathetic and parasympathetic branches of the autonomic nervous system [2]. The richness of this variability is not fully described merely by the standard deviation of the inter-beat intervals, and different markers have been proposed to capture a range of dynamic patterns indicative of health conditions

* Corresponding author at: School of Engineering and Informatics, University of Sussex, United Kingdom.

E-mail address: f.rosas@sussex.ac.uk (F.E. Rosas).

<https://doi.org/10.1016/j.complbiomed.2023.107857>

Received 19 June 2023; Received in revised form 24 November 2023; Accepted 11 December 2023

Available online 23 December 2023

0010-4825/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

and cognitive states [7]. While conventional metrics of HRV are based on spectral decomposition or other linear methods, these methods are unable to capture non-linear fluctuations in heart rate — which have been shown to convey valuable information [8,9]. Building on this principle, there is a progressively growing literature that provides complementary insights on HRV patterns by investigating fractality and signal complexity via metrics such as permutation entropy [10,11] and detrended fluctuation analysis [12,13], just to name a few. Overall, researchers have a rich set of tools to use cardiac signals to investigate autonomic outflow.

This promising state of affairs is, however, overshadowed by the fact that extracting useful information from cardiac data is highly non-trivial. Heart rate is typically conceived to be a summary statistic (i.e. a simplified description) of the inter-beat interval time series, the estimation of which is based on frequentist principles and results in a point estimate. This approach, however, has critical limitations: (i) it is not able to incorporate relevant prior knowledge (e.g. what type of heart rate values are physiologically plausible), (ii) it is unable to deliver metrics of confidence of the estimation, and (iii) its value relies only on the inter-beat interval of the present moment and hence estimation errors in inter-beat intervals result in important overestimations of heart rate variability. Unfortunately, a poor estimation of heart rate dynamics greatly hinders the capabilities of any downstream analysis. Overall, there is a great need for more powerful and flexible estimation methods that may allow us to better characterise and fully extract the information contained in heart rate dynamics [14].

To address these important limitations, in this paper we propose to rethink the process of estimating heart rate dynamics to better capture the autonomic outflow. For this purpose, we introduce a new framework to model and estimate heart rate dynamics based on state-space modelling [15] and Bayesian statistics [16]. Using state-space modelling principles, our framework conceives the heart rate not as a summary statistic but as a hidden (i.e. not directly measurable) process which drives the observed sequence of heart beats. Moreover, using Bayesian techniques our approach delivers not a point estimate but a posterior distribution, which encodes the likelihood of possible heart rate trajectories given an observed sequence of inter-beat intervals. To illustrate the capabilities of the proposed framework, we use the posterior distribution to build estimates of non-linear properties of heart rate dynamics. We show that these estimates exhibit higher discriminative power to distinguish between different physiological states. Our proposed framework expands researchers' toolkit to study the autonomic outflow, opening new opportunities to develop effective biomarkers for different conditions — e.g. risk stratification of cardiovascular conditions or tracking the underlying physiology of embodied cognitive processing.

2. Results

2.1. A Bayesian approach to model heart rate dynamics

The conventional method to calculate heart rate involves inferring how many beats one would expect per minute on average given the observation of N_b beats over a period of time of T seconds, which leads to the estimate $HR = 60N_b/T$. If one is interested in a dynamical description of how the heart rate fluctuates over time, one can follow the same rationale and reduce the time period to the limit where $N_b \rightarrow 1$ and T becomes equal to the inter-beat interval I_b , leading to the following estimate of the “instantaneous” heart rate:

$$HR_{\text{freq}}(t) = \frac{60}{I_b(t)}. \quad (1)$$

From a statistical perspective, this expression can be understood as the outcome of an elementary frequentist method of inference that delivers a point estimate for the average number of beats per minute — in fact, it is the number of beats one would see if all beats were separated




Fig. 1. Proposed heart rate state-space modelling approach. The observable data (the heart beats) are assumed to be driven by the dynamics of a hidden stochastic process (the heart rate), which cannot be directly measured but can be inferred from the data.

by the same inter-beat interval I_b . As such, it has the strengths and weaknesses of frequentist approaches: it is conceptually simple and computationally lightweight, although it cannot estimate its own uncertainty or incorporate prior knowledge on plausible heart rate values. Furthermore, as $HR_{\text{freq}}(t)$ ignores previous inter-beat interval values, errors in the estimation of $I_b(t)$ inevitably lead to overestimations of heart rate fluctuations.

In contrast, our proposed approach conceives the heart rate as a hidden process that drives the actual observed heart beats, the statistical properties of which can be estimated via generative modelling. Our model involves two time series corresponding to the values of a dynamical process sampled with sampling frequency $f_s = 1/\Delta t$: x_t , which counts the number of heart beats that take place during a temporal bin of length Δt , and z_t , which is the heart rate that drives the corresponding heart beats (Fig. 1). Our framework comprises a generative statistical model (see *Methods*), the key component of which is a probability distribution p that describes the likelihood of observing a given sequence of heart beats x_1, \dots, x_N together with a heart rate time series z_1, \dots, z_N .

Through this model, heart rate dynamics are now described not by a point estimate (i.e. as a single, most likely trajectory) but as obeying the following conditional distribution:

$$HR_{\text{bayes}} : z_1, \dots, z_N \sim p(z_1, \dots, z_N | x_1, \dots, x_N). \quad (2)$$

This posterior distribution describes the most likely heart rate trajectories z_1, \dots, z_N given the observed data x_1, \dots, x_N . Crucially, by following methods proposed in Ref. [17] we avoid the computationally intensive task of computing the explicit posterior distribution, and instead use a Gibbs sampler [18] to efficiently obtain sample trajectories. This allows not only to find the most likely trajectory, but also to estimate uncertainty (e.g. via the posterior variance). These sampled trajectories also allow us to build accurate estimators of non-linear properties of heart rate dynamics — as developed in later sections. The sampling procedure described in Ref. [17], which we follow here, employs two hyperparameters θ and τ , which are related with the connectivity strength between successive samples. In our experiments, sampled trajectories were seen to be fairly insensitive to the choice of τ , while their smoothness strongly depended on θ — low values of θ induce a strong constraint between successive samples, making more unlikely abrupt changes of values. For simplicity, all our calculations use $\tau = 1$ and consider two values of θ : 0.01 for allowing more variation between successive heart rate values, and 1 to make their connectivity stronger.

2.2. Method validation

We tested our proposed method by analysing the autonomic outflow of healthy subjects going through a standard tilt-table protocol. This protocol places participants at various angles to monitor changes in cardiovascular activity in order to assess the balance between the sympathetic and parasympathetic branches of the autonomic nervous system [19]. In particular, when the head is tilted upwards then blood flow to the brain decreases, which triggers activation of the sympathetic nervous system and suppression of the parasympathetic nervous system [20,21].

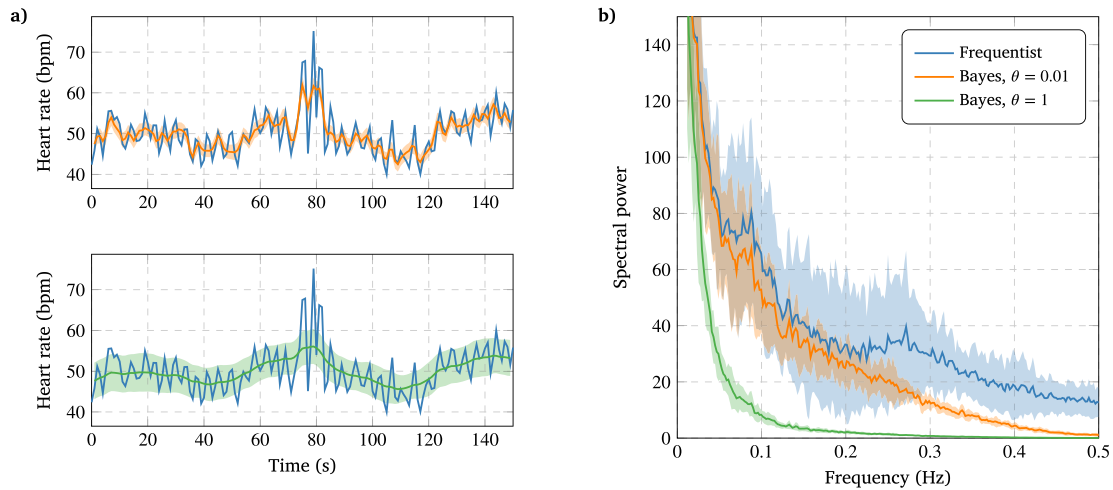


Fig. 2. Illustration of the proposed method in example time series and power spectrum. (a) Comparison between frequentist (blue) and Bayesian (orange and green) estimation of heart rate, of which we display the mean and standard deviation obtained from 100 trajectories sampled from models with a weaker (top, $\theta = 0.01$) and stronger (bottom, $\theta = 1$) coupling strength between successive samples (see Section 4.2.2). (b) Power spectrum of heart dynamics as obtained from frequentist and Bayesian approaches.

We used a public dataset¹ comprising 10 healthy subjects (5 males, average age 28.7 ± 1.2 years) undergoing the tilt-table test [22–24]. Each subject was measured under four conditions (self stand up, slow tilt, fast tilt, and resting baseline) with a one-lead ECG. The resulting time series were used to calculate inter-beat intervals via template-based matching of QRS waves [25]. The sequence of inter-beat intervals was then used to construct two generative models for each subject, using two different values for a free parameter denoted by θ that regulates the variability between successive values of heart rate (see *Methods*) – which is closely related to the bandwidth of the resulting trajectories – and sampled 100 heart rate trajectories from the posterior. Our approach worked as expected, generating smoother and more plausible trajectories than the frequentist approach. Furthermore, results also show how a higher value of the hyperparameter θ filters out fluctuations, resulting on trajectories that keep the main trends while removing abrupt fluctuations. The resulting trajectories and corresponding spectrum are shown in Fig. 2.

To further test our proposed method, we used computational modelling to assess the consistency of the estimated autonomic outflow. Computational modelling has been employed to generate synthetic series of heartbeat dynamics, which help us understand autonomic dynamics in different conditions [26]. Among the different models that have been proposed, the integral pulse frequency modulation (IPFM) models describe the physiological transduction from the autonomic outflow to heartbeat generation [27]. These models are developed under the hypothesis of the existence of a heartbeat function as a real-time modulation function representing stimulation of the sinoatrial node, which is directly related to heartbeat generation [28]. The modulation function in IPFM models considers the combination of heart rate components of sympathetic-vagal control as the inputs to an integrator that generates the heartbeats [29–31].

Accordingly, we tested the trajectories generated by our Bayesian framework using a IPFM model. Specifically, we applied the average of the resulting trajectories as an input, and as an output the model generates synthetic heartbeats that takes the autonomic outflow estimation as an input for an integrate-and-fire process [30], whose details are described in *Methods*. We used the obtained sequence of “synthetic” heartbeats to calculate heart rate using the simple frequentist approach given in Eq. (1), and studied its suitability by comparing it with the frequentist estimation of heart rate obtained from the original heartbeats (Fig. 3b).

Results indicate that the synthetic heart rate matches the measured heart rate both when in supine position (z-score, measured: 1.12 ± 0.11 ; synthetic: 1.12 ± 0.10) and standing up (z-score, measured: 1.29 ± 0.08 ; synthetic: 1.30 ± 0.08). Moreover, the resulting values of mean heart rate were significantly correlated among subjects (Spearman $R = 0.8667$, $p = 0.0027$, Fig. 3c). Furthermore, the observed changes between conditions in the synthetic heart rate recapitulate the ones exhibited in the real data, as observed via a paired Wilcoxon signed-rank test ($p = 0.0051$).

2.3. Comparable discriminative power of spectral properties

The power of generative modelling can be harnessed to build Bayesian estimators of properties of heart rate dynamics. In particular, for a given property F of a heart rate trajectory z_1, \dots, z_N (e.g. entropy or spectral power), a Bayesian estimator can be built as follows:

$$\hat{F} = \sum_{z_1, \dots, z_N} F(z_1, \dots, z_N) p(z_1, \dots, z_N | x_1, \dots, x_N), \quad (3)$$

where the value of the property F for each possible trajectory is weighted by the likelihood of such trajectory given the observed data. Furthermore, if F is a linear property, then Eq. (3) accepts a shortcut: it reduces to $\hat{F} = F(\bar{z}_1, \dots, \bar{z}_N)$, where $\bar{z}_1, \dots, \bar{z}_N$ is the average trajectory under the posterior $p(z_1, \dots, z_N | x_1, \dots, x_N)$.

We sought to study spectral properties of heart rate dynamics via these Bayesian estimators, and investigated their capability to discriminate between the different physiological states in which the balance between sympathetic and parasympathetic activities is re-arranged. We compared the discriminative power of these Bayesian estimations against the obtained via the standard frequentist approach.

For this purpose, we calculated Bayesian estimators for some linear features of heart rate dynamics, including the mean heart rate and spectral components of the heart rate variance, for each subject on each condition. To calculate spectral components, we first filtered the heart rate trajectories using a bandpass Butterworth filter of order 2 over standard HR bands (0.04–0.15 Hz for low-frequency and 0.15–4 Hz for high-frequency components), and then calculated the variance of the resulting signal. We refer to the variance of the low and high frequency signals as “low-frequency HRV” and “high-frequency HRV”, respectively.

We analysed the obtained values of mean heart rate, low-frequency HRV, and high-frequency HRV obtained for each stage and each subject via linear mixed-effect modelling. Specifically, we constructed models using the heart rate feature as target variable, stage as fixed effect, and

¹ <https://physionet.org/content/prcp/1.0.0/>.




Fig. 3. Method validation. (a) Block diagram of the integral pulse frequency modulation (IPFM) model, which is used to generate synthetic heartbeats (see *Methods*). (b) Group-wise measured and synthetic heart rate series for the change from supine to upright position. The displayed signals correspond to the group median and shaded areas to the median absolute deviation. (c) Scatter plot showing the correlation between measured and synthetic heart rate, in supine and standing up conditions. Dashed black lines are least squares regression lines. All the standing up postural changes were averaged per subject, centred at the time of postural change. Z-scores were computed per subject in the -180 to 180 s interval with respect to the postural change timing.

modelled the effect of subject identity using a random intercept. We compared the effects observed with our Bayesian procedure using two different values for hyperparameter θ : 0.01 (more loose) and 1 (more constrained). We contrasted these results with the ones obtained from a standard frequentist estimation. Results are shown in [Fig. 2](#) and the resulting statistics are presented in [Appendix](#).

Overall, the effects observed for mean heart rate across conditions are very similar for all estimation approaches, while for HRV the model

with $\theta = 0.01$ presents small reductions and the one with $\theta = 1$ presents stronger reductions — consistent with the fact that θ is related to the bandwidth of the model (see *Methods* and also [Fig. 2](#)). Correspondingly, the model with $\theta = 0.01$ achieves much better performance distinguishing between different states than the model with $\theta = 1$, particularly with high-frequency HRV. Moreover, it is worth noticing that the discriminative power of the model with $\theta = 0.01$ is comparable — and sometimes better — than the one from the frequentist method.

2.4. Better performance for measures of dynamical complexity

It is generally recognised that analytic techniques inspired by principles from complexity science [32] are able to describe biological processes— such as the ones that drive heart dynamics — in illuminating ways [2]. When employing such techniques, it is worth noting that if a dynamical property of heart rate dynamics F is non-linear, then its expected value \hat{F} is in general different from the property evaluated on the average trajectory, $F(\bar{z}_1, \dots, \bar{z}_N)$.² The fact that Bayesian estimators can assess this difference is one of their most powerful features. To illustrate how the proposed framework can be used to estimate metrics of dynamical complexity, in the following we consider two of them: the *Hurst exponent* and *sample entropy*.

First, we consider the Hurst exponent of the heart rate time series as estimated via detrended fluctuation analysis (DFA) [33] — a method to determine the statistical self-affinity of a signal which is particularly useful for analysing time series with long-memory processes or $1/f$ noise [34] (see Ref. [35] for a tutorial). While DFA can be used just as a pragmatic way to obtain useful biomarkers, some researchers interpret its effectiveness as indicative of a fractal nature in heart rate fluctuations [36].³ As a simple proof of concept, here we calculated the Hurst exponent via a simple application of DFA using the open-source package *Fathon* [41].

As a second complexity measure, we consider the well-known permutation entropy [42], which analyses the patterns in the heart rate dynamics by classifying them into discrete classes and then studying their frequencies — being an example of a more general approach known as “symbolic dynamics” [10]. Permutation entropy is a popular technique for measuring the pattern complexity in time series analysis in general [43] and heart rate dynamics in particular [44], with extensions that can assess multiple temporal scales [11]. For the sake of simplicity, we performed an elementary calculation of permutation entropy using the open source package *Ordpy* [45].

Both Hurst exponent and permutation entropy were calculated on each of the 100 sampled trajectories per subject, and then averaged to provide the outcome of the proposed Bayesian estimator. For comparison, we also evaluated both properties in the point estimate trajectory provided by the frequentist approach. To contrast the discriminative power of Bayesian and frequentist estimators, we constructed linear-mixed models using the various HR markers as dependent variables, the physiological stage as independent variable, and a random intercept per subject (see *Methods*). Results show that neither the frequentist method or the Bayesian model with $\theta = 0.01$ are able to find significant differences in either metric across conditions. Crucially, however, the Bayesian model with $\theta = 1$ does reveal strong effects in both metrics, particularly in the self stand up and fast tilt conditions, showing that our approach yields more suitable heart rate time series for downstream complexity analyses.

² For example, the well-known Jensen inequality states that, if F is convex, then the mean value of F is usually greater than F evaluated in the mean value.

³ Some argue that this would be related to the fractal structure of blood vessels and of nervous and humoral pathways connecting and hierarchically regulating local blood flows among several vascular beds [37]. However, the precise physiological bases of DFA are a matter of debate, and efforts to explain it have been based on relationships with spectral properties [38,39] and via computational modelling [40].

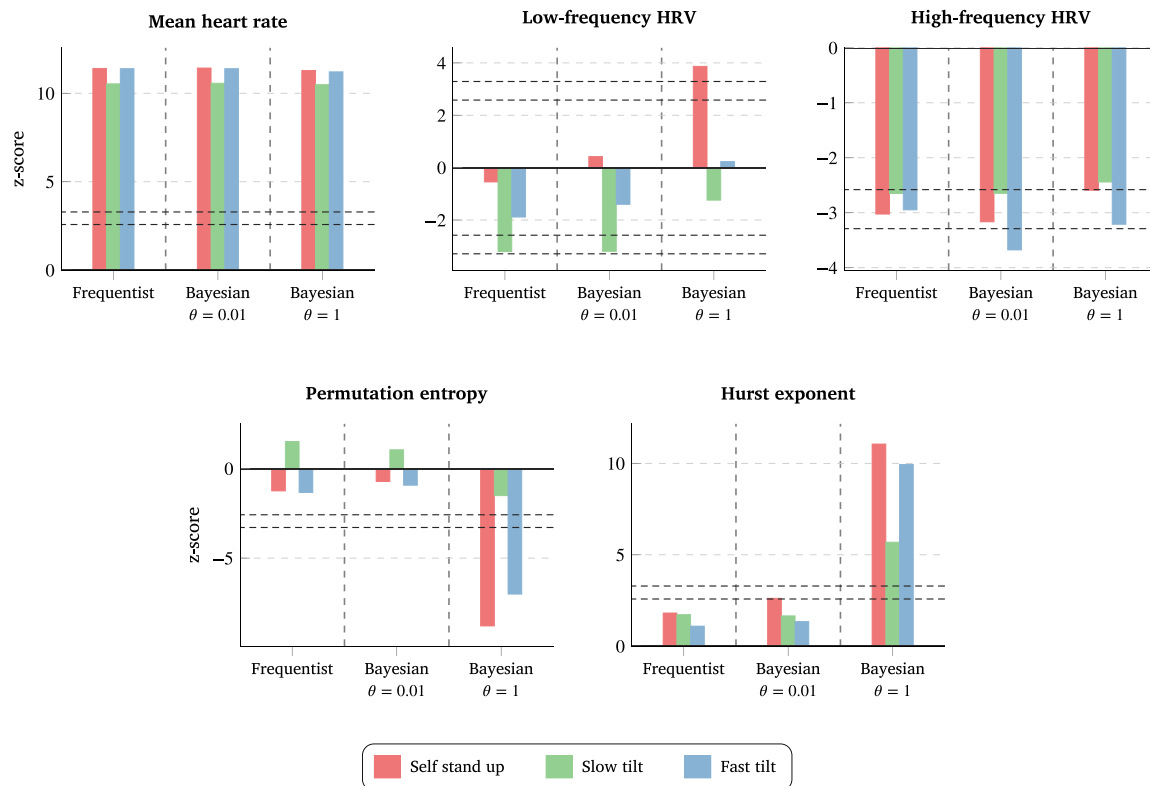


Fig. 4. Effects (in the form of z-scores) obtained from linear mixed-effect modelling of the effect of various physiological states on properties of heart rate dynamics. Top row contains mean heart rate and spectral measures, and bottom row contains measures of dynamical complexity. While the Bayesian model with $\theta = 0.01$ exhibits a similar discriminative power than the frequentist estimator for mean heart rate and spectral HRV, the Bayesian model with $\theta = 1$ performs better for permutation entropy and Hurst exponent. Black horizontal lines represent p -value thresholds corresponding to $p = 0.01$ and $p = 0.001$.

3. Discussion

This paper introduces a new conceptual and modelling paradigm for the assessment of autonomic processes, in which the heart rate is construed as a hidden stochastic process that drives the generation of observed heart beats. Our framework instantiates this idea using the rich toolkit of state-space modelling and Bayesian statistics, with efficient algorithms to sample the resulting model and generate corresponding heart rate trajectories. This framework has great potential for clinical applications monitoring the effect of medical interventions on autonomic dynamics — a potential that our ongoing work is starting to explore for the case of pharmacological interventions [46].

Our results show how sampled trajectories can be used to estimate the posterior distribution of various properties of heart rate dynamics, leading to Bayesian estimators of spectral power or dynamical complexity. In agreement with the literature [19], our results on the autonomic effects in a tilt-table test revealed a transient increase in heart rate and decrease in HRV in response to standing upright that is observed in both Bayesian and standard (frequentist) estimates of autonomic outflow. Importantly, results further show that the proposed Bayesian estimators of dynamical complexity measures such as permutation entropy and the Hurst exponent exhibit more discriminative power than the corresponding estimates obtained from the frequentist approach. It should be noted that while the results presented here highlight the potential of this approach for delivering sensitive autonomic markers, additional work is needed to interpret these effects in terms of precise biological mechanisms.

The estimation framework followed in this work, which the more general framework for modelling point processes introduced in Ref. [17], depends on one effective hyperparameter θ that controls the smoothness of the resulting autonomic trajectories. In this work we have contrasted the effects of two choices of values to this hyperparameter, which delineate the boundaries of realistic trajectories that are

not too smooth (if $\theta \ll 0.01$) or display numerical artefacts (if $\theta \gg 1$). While intermediate values can also be evaluated, at this stage we do not have a general procedure to judge if some values should be preferred over others. Future work focus on biological mechanics may investigate if certain values of θ could provide a more clear lens to investigate specific autonomic phenomena of interest.

Overall, our approach for estimating autonomic outflow ought to be seen as a first instantiation of a powerful and flexible paradigm, which can be significantly extended in the future. For example, the proposed model consists of a hidden process composed of a single time series, which aims to capture the aggregated autonomic outflow. Future work could use this framework to model sympathetic and parasympathetic activities separately, as well as to model the interactions between multiple physiological signals. This could help identify the neural substrates of underlying brain-heart interactions [47,48], and capture the specific nodes from the central autonomic network that are involved [3]. Another limitation of the presented approach is that it does not account for the separate impact of blood pressure and respiration on heart rate [49], which could introduce biases in the results.

It is worth noting that our modelling approach focuses on modelling heart rate dynamics and not heart beats, which are taken as input data for estimating the former. Other approaches are more appropriate for modelling and predicting inter-beat intervals or heartbeat events, such as the methods developed in Refs. [50–52]. In contrast, our approach focuses on the posterior distribution of the heart rate dynamics given the observed heart beats.

Another relevant feature of the proposed modelling approach is that it is well-suited to the analysis of non-stationary autonomic data. Heart beat dynamics can be highly non-stationary in certain conditions, e.g. under postural changes (as in the tilt-table protocol considered in our analyses). Accordingly, our proposed method builds a model that is




Fig. 5. Diagram of our approach's pipeline. For an individual time series of inter-beat intervals, our framework delivers the posterior distribution of a feature of interest (e.g. mean heart rate, or sample entropy). This is done via two intermediate steps: first a number of plausible heart rate trajectories are generated, and then the feature of interest is calculated over each of these.

non-stationary by construction, and is hence fully capable of accounting for non-stationary features of heart rate dynamics. The exploitation of this powerful feature is the subject of ongoing work.

Overall, the present paper puts forward a new conceptual and practical approach to model autonomic dynamics which offers several new research opportunities and potential applications. While here we present simple validation and proof-of-concept, further work is required to delineate the precise physiological implications of the proposed construction. The proposed framework may lead to practically useful biomarkers to diagnose and monitor clinical conditions, while providing a basis for future modelling and simulation studies of the autonomic system and the involved neuronal mechanisms underlying physiological regulation.

4. Methods

4.1. Overall approach

At a high level, our modelling strategy takes a sequence of inter-beat intervals as inputs, and delivers the posterior distribution of an autonomic observable of interest as output. This is done via two intermediate steps, as illustrated graphically in Fig. 5. First, the sequence of inter-beat intervals is used to generate a number of heart rate trajectories compatible with it. Second, an observable of interest is calculated over each of the generated trajectories. This last step generates a number of estimates of the property in question, which together constitute a sample of the posterior distribution of that property. While the analyses presented in Section are based on the mean values of such posterior distributions, they can be harnessed in other ways — e.g. calculating their median, variance, or other moments.

4.2. Modelling of heart rate dynamics

Our model of heart rate dynamics is based on two time series that correspond to the values of dynamical processes with sampling frequency $f_s = 1/\Delta t$: x_t , which counts the number of heart beats in a temporal bin of length Δt (i.e. between the present moment and the previous sample), and z_t , which represents the heart rate that drives the corresponding heart beats (Fig. 1). In general, we assume that the heart beats x_t are observed, and that sequence of heart rates z_t follows a hidden (i.e. unmeasured) stochastic process whose statistical properties can nevertheless be inferred.

In order to build a joint probability distribution over heart beats x_1, \dots, x_N and heart rates z_1, \dots, z_N , we follow the state-space literature [15] in adopting the following assumptions that make inference of the hidden process tractable:

- (i) the dynamics of the hidden process are Markovian; and
- (ii) given the value of z_t , the observable x_t is conditionally independent of all other observations and values of the hidden process.

Thanks to these assumptions, the joint distribution of heart rate and heart beat counts can be expressed as:

$$p(x_1, z_1, \dots, x_T, z_T) = p(z_1)p(x_1|z_1) \prod_{t=2}^T p(z_t|z_{t-1})p(x_t|z_t) \quad (4)$$




Fig. 6. Graphical model representation of the proposed state-space model for heart rate dynamics. The actual heart rate is tracked by z_1, \dots, z_T , while the variables x_1, \dots, x_T correspond to the count of heart beats over a given time interval, and y_1, \dots, y_T are auxiliary variables.

Therefore, the specification of the full model requires only three ingredients: the heart rate dynamics in the form of the conditional probability $p(z_t|z_{t-1})$, the link between heart rate and heart beats $p(x_t|z_t)$, and the distribution of the initial condition $p(z_1)$.

4.2.1. Heart rate dynamics via gamma Markov chains

Here we specify a generative model for heart rate dynamics, which follows the approach introduced in Ref. [17] for general point processes.

Let us first specify the dynamics of the hidden process that models the heart rate, $p(z_t|z_{t-1})$ and $p(z_1)$. The most common instantiation of state-space models considers Gaussian variables [15], however we avoid this approach as we heart rate is a non-negative quantity. Instead, our model considers the hidden dynamics as driven by a Gamma Markov chain (GMC) [53,54], which builds the dynamics of the hidden process z_1, \dots, z_T with the addition of auxiliary variables y_2, \dots, y_T whose joint distribution is a Markov chain of the form (Fig. 6):

$$p(z_1, y_2, \dots, z_T, y_T) = p(z_1) \prod_{j=2}^T p(y_j|z_{j-1})p(z_j|y_j), \quad (5)$$

with the corresponding distributions defined as follows:

$$\begin{aligned} z_1 &\sim G(z_1; \alpha_1, \beta_1), \\ y_t|z_{t-1} &\sim \text{IG}(y_t; \gamma, \gamma z_{t-1}), \\ z_t|y_t &\sim G\left(z_t; \gamma, \frac{\gamma}{y_t}\right), \end{aligned} \quad (6)$$

where α_1, β_1 and γ are hyperparameters, and G and IG denote the Gamma and Inverse Gamma distributions, which are given by

$$\begin{aligned} G(z; \alpha', \beta') &= \frac{\beta'^{\alpha'}}{\Gamma(\alpha')} z^{\alpha'-1} e^{-\beta' z}, \\ \text{IG}(y; \tilde{\alpha}, \tilde{\beta}) &= \frac{\tilde{\beta}^{\tilde{\alpha}}}{\Gamma(\tilde{\alpha})} z^{-\tilde{\alpha}-1} e^{-\tilde{\beta}/z}, \end{aligned} \quad (7)$$

where $\alpha', \tilde{\alpha}$ and are shape parameters and $\beta', \tilde{\beta}$ are scale parameters. Crucially, this choice of auxiliary variables makes the full conditionals to have a simple form:

$$\begin{aligned} y_t|z_{t-1}, z_t &\sim \text{IG}\left(y_t; 2\gamma, \gamma x_{t-1} + \gamma x_t\right), \\ z_t|y_t, y_{t+1} &\sim G\left(z_t; 2\gamma, \frac{\gamma}{y_t} + \frac{\gamma}{y_{t+1}}\right), \\ z_1|y_2 &\sim G\left(z_1; \alpha + \gamma, \beta + \frac{\gamma}{y_{t+1}}\right), \\ z_T|y_T &\sim G\left(z_T; \gamma, \frac{\gamma}{y_T}\right), \end{aligned} \quad (8)$$

which enables an efficient sampling of heart rate trajectories, as explained in the next section.

Let us now specify the distribution that links heart rate and heart beats, $p(x_t|z_t)$. To leverage the powerful methods of Bayesian statistics,

our natural choice is to consider a Poisson distribution — the conjugate likelihood of the Gamma distribution. Hence,

$$x_t | z_t \sim P(x_t; z_t), \quad (9)$$

where $P(x; \lambda) = \lambda^x e^{-\lambda} / x!$ is the Poisson distribution for rate λ . Please note that while the Poisson distribution generally yields a suboptimal fit to empirical distributions of the number of heart beats per bin in general,⁴ our usage of Poisson distributions avoids this issue — as explained in the next section.

Finally, let us discuss the parameters of the model. From Eqs. (8) and (9), it can be seen that the model parameters are only α_1 and β_1 for specifying $p(z_1)$, and γ for determining the connectivity strength between successive samples of the hidden process. Our approach is to choose α_1 and β_1 by performing a maximum likelihood estimation of the parameters of a Gamma distribution that would most likely have generated heart rate estimates as calculated by the simple method given in Eq. (1). For γ , we follow Ref. [17] and consider a prior distribution over it — in this work we use $\gamma \sim \text{Exp}(\gamma; \theta)$ with Exp denoting an exponential distribution, and θ is the hyperparameter of this prior.

4.2.2. Sampling the generative model

A useful feature of Bayesian statistics is that they enable efficient methods to sample the posterior distribution $p(z_1, \dots, z_T | x_1, \dots, x_T)$. This can be done via a *Gibbs sampler* [55,56], a Markov chain Monte Carlo (MCMC) [57] algorithm that can be used to obtain sequences of observations that follow a given distribution when direct sampling is difficult. The Gibbs sampler provides an efficient procedure to extract trajectories from the posterior distribution, as we explain below.

Following Ref. [17], one can show that the posterior distributions have — thanks to the modelling choices — the following simple expressions:

$$\begin{aligned} y_t | z_{t-1}, z_t &\sim \text{IG}(y_t; 2\gamma, \gamma x_{t-1} + \gamma x_t), \\ z_t | y_t, y_{t+1}, x_t &\sim \text{G}\left(z_t; 2\gamma + x_t, \frac{\gamma}{y_t} + \frac{\gamma}{y_{t+1}} + n\Delta t\right), \\ z_1 | y_2, x_1 &\sim \text{G}\left(z_1; \alpha + \gamma + x_1, \beta + \frac{\gamma}{y_{t+1}} + n\Delta t\right), \\ z_T | y_T, x_T &\sim \text{G}\left(z_T; \gamma + x_T, \frac{\gamma}{y_T} + n\Delta t\right). \end{aligned} \quad (10)$$

Furthermore, the (unnormalised) posterior distribution for parameter γ given z_t, y_t can be shown to be the following non-standard distribution [17]:

$$\begin{aligned} \gamma | z_t, y_t &\sim \text{Exp}(\gamma; \theta) \times \left(\frac{\gamma^\gamma}{\Gamma(\gamma)}\right)^{2(T-1)} \prod_{j=2}^T \left(\frac{z_j z_{j-1}}{y_j^2}\right)^\gamma \\ &\times \exp\left(-\gamma \sum_{j=2}^T \frac{z_j + z_{j-1}}{y_j}\right). \end{aligned} \quad (11)$$

Using the above equations, samples of trajectories can be computed as follows. First, the sampler is initialised using independent samples for z_1, \dots, z_T drawn from a Gamma distribution with parameters that best match a simple frequentist estimation of the heart rate (i.e. using the same values of α and β obtained for $p(z_1)$, as specified in the previous section). Also, γ is initialised by the mean value of its prior, i.e. the value of its hyperparameter θ . Using the obtained initial values of z_t and γ , the sampler then iterates as follows:

1. Estimate y_t given x_t and z_t
2. Estimate z_t given y_t and z_{t-1}
3. Estimate γ given z_t and y_t
4. Go to 1

⁴ For example, one can check that while the number of heart beats in a given time window grows linearly with bin length, the variance does not.

In this procedure, the estimation of z_t and y_t is done directly using Eqs. (10). For the estimation of γ we follow Ref. [17] in using a Metropolis–Hastings MCMC using a zero-mean Gaussian with standard deviation τ as the proposal function. The value of τ was selected to achieve an intermediate number of accepted and rejected steps in the Markov chain, and results were empirically insensitive to many choices of τ . Also, please note that when sampling heart rate trajectories, there is no need to keep record of the values of the auxiliary variables y_1, \dots, y_T .

4.2.3. Choosing parameter values

The sampling procedure described in previous subsection has only two hyperparameters, θ and τ , which are related with the prior and estimation of the connectivity strength γ . In our experiments samples were seen to be fairly insensitive to τ , while their smoothness strongly depended on θ . For this reason, we set $\tau = 1$ and consider two values of θ : 0.01 for allowing more variation between successive heart rate values, and 1 to make their connectivity stronger (Fig. 1).

Another parameter to be chosen is the sampling frequency f_s . To determine this value, we followed two criteria: we would like to have a relatively high sampling frequency to be able to assess heart dynamics with good temporal resolution, and we would like to make the Poisson statistics of $p(x_t | z_t)$ fit the distribution of empirical data as well as possible. Luckily, both criteria can be satisfied by adopting a relatively high frequency; in this paper we use $f_s = 3$ Hz. Indeed, the resulting small temporal bins means x_t is effectively a binary variable, such that $x_t = 1$ if there was a heart beat during that bin and $x_t = 0$ otherwise. Furthermore, as the frequencies of events $x_t = 1$ is relatively low (approximately 1/3 for a heart rate near 60 BPM), such statistics can be well fit by a Poisson distribution with low rate.

This choice, however, has two downsides — which luckily do not matter for the current application. First, the relatively regular appearance of one beat every three bins generates a small artefactual fluctuation. We deal with this in an heuristic manner: we do a smoothing step (a rolling mean with a triangle window of 9 samples), and then a subsequent subsampling step keeping one every three samples. With this, our resulting sampling frequency is $f_s = 1$ Hz, which is still high enough for the purpose of studying heart rate dynamics. The second downside is that the model cannot be used to sample meaningful sequences of heart beats, as the conditional independence of the corresponding variables is violated due to the small bin size. However, our goal is not to sample heart beats but heart rate trajectories, and hence this limitation is not important. Importantly, violations to this statistical condition do not imply that the modelling becomes inaccurate — in fact, the efficacy of the posterior of the proposed model beyond this condition is equivalent to the well-known broad efficacy of the Naive Bayes classifier in a wide range of scenarios that do not satisfy its statistical assumptions [58].

4.3. Building Bayesian estimators of properties of heart dynamics

Finally, we can leverage this trajectory sampling technology to build Bayesian estimators of properties of heart rate dynamics, as follows.

Let us use the shorthand notation $\mathbf{z} = (z_1, \dots, z_T)$ for a sampled trajectory of the heart rate dynamics, and consider $F(\mathbf{z})$ to be a scalar function of this trajectory. For example, this could be an estimator of the trajectory's entropy or of its Hurst exponent. Then, using our Gibbs sampler we could extract samples of the posterior distribution $p(F(\mathbf{z}) | x_1, \dots, x_T)$. One can then use these samples to estimate various useful features of this distribution, such as its mean or standard deviation. In particular, the presented analyses estimate various properties of heart rate dynamics as the empirical mean value of this posterior.

Importantly, our experiments suggest that the proposed model for heart rate dynamics is generally non-ergodic. Therefore, our method to sample the posterior is as follows. To generate one trajectory, we run the Gibbs sampler N_t iterations and discard the first N_d . We calculate

the mean value of the remaining $N_r - N_d$ runs and take the result as a single trajectory.⁵ To generate N_s trajectories, we run this procedure N_s times, initialising the Gibbs sampler every time with new random initial conditions. For the presented results, for each time series of heartbeats we run the Gibbs sampler for $N_r = 2 \times 10^4$, and discard the first $N_d = 5 \times 10^3$.

4.4. ECG data and preprocessing

To validate our framework, we analysed the estimation of heart rate on a data set of postural changes. For this, we used a public dataset⁶ comprising one-lead ECG series from 10 healthy subjects (5 males and 5 females, average age 28.7 ± 1.2 years) while undergoing the tilt-table test [22–24]. The subjects were initially asked to remain in a horizontal supine position and to move to a vertical position with the help of either the tilt-table or by self-stand up. The subjects were part of six sessions that were sorted randomly between resting periods: two stand up, two slow tilts (50 s from 0 to 70°), and two fast tilts (2 s from 0 to 70°), while remaining in each condition for approximately 3 min. The entire protocol lasted between 55 and 75 min. This experimental protocol was approved by MIT's Committee On the Use of Humans as Experimental Subjects and the Advisory Board of the MIT-MGH General Clinical Research Centre, as reported in the original publication [59].

For simplicity of the analysis, we aligned all trials by subject between –180 to 180 s with respect to the change from supine position to standing up. ECG preprocessing involved frequency filtering, R-peak detection, and correction of misdetections. ECG data were bandpass-filtered using a Butterworth filter of order 4, between 0.5 and 45 Hz. Heartbeats from QRS waves were identified in an automated process based on a template-based method for detecting R-peaks [25].

4.5. Validation of generated heart rate via an integral pulse frequency modulation (IPFM) model

Mathematically, the heartbeat generation model is based on the temporal integration of the autonomic outflow, denoted by $z(t)$,⁷ which contains the modulations from sympathetic and parasympathetic nervous systems. Given the previous beat occurred at time t_k , the model considers the following integral function:

$$\Phi_k(t) = \int_{t_k}^t z(t) dt . \quad (12)$$

When Φ_k reaches a threshold R (which for simplicity we fix at $R = 1$), then the $k + 1$ -st heartbeat is generated and the integral is reset. Finally, the heartbeat function $x(t)$ is modelled as a sum of Dirac functions $\delta(t)$ positioned at each heart beat timing t_k :

$$x(t) = \sum_k \delta(t - t_k) . \quad (13)$$

Note that, in principle, one could formulate a probabilistic version of IPFM and invert it using Bayes' rule to obtain another method to generate plausible time series of autonomic outflow from observed heart beats. However, here we focused on the GMC model for three reasons: (i) the IPFM model has discontinuous jumps, which pose problems for Bayesian inference; (ii) it cannot properly model ectopic beats; and (iii) the GMC model supports a computationally efficient Gibbs sampler that allows for faster inference. Nonetheless, developing a Bayesian inversion of IPFM is an interesting avenue for future work.

⁵ Technically, we are sampling the posterior distribution of the *mean* heart rate trajectory. Given the non-ergodicity of the sampler, this approach was more empirically powerful than analysing non-averaged sampled trajectories directly.

⁶ <https://physionet.org/content/prcp/1.0.0/>.

⁷ Note that we have changed notation from discrete-time z_i in Section 4.2.1 to continuous-time $z(t)$.

4.6. Linear mixed-effect modelling

The discriminative power of the Bayesian estimates obtained in Section 2.4 were compared with the one of frequentist estimates via linear mixed-effect modelling, with results illustrated in Fig. 4. For clarity: the Bayesian and frequentist methods are used to obtain values, which are then entered into statistical comparison carried out with LME modelling. The linear mixed models were constructed using one HR marker (e.g. mean HR, low-frequency HRV, etc.) as dependent variable and the physiological state (i.e. self stand up, slow tilt, or fast tilt), as a categorical dependent variable, and a random intercept per subject to account for between-subject variability. For making the values of the estimates of the regressions comparable, each of the considered dependent variables were turn into z-scores. Please note that we use linear mixed-modelling merely as a simple way to assess the differences between Bayesian and frequentist estimations, but this is not an integral stage in our proposed new framework: rather, we use this as a way of comparing our new framework against a more traditional approach.

CRediT authorship contribution statement

Fernando E. Rosas: Conceptualization, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Diego Candia-Rivera:** Conceptualization, Formal analysis, Investigation, Methodology, Writing – review & editing. **Andrea I. Luppi:** Conceptualization, Formal analysis, Investigation, Methodology, Writing – review & editing. **Yike Guo:** Conceptualization, Formal analysis, Investigation, Methodology, Writing – review & editing. **Pedro A.M. Mediano:** Conceptualization, Formal analysis, Investigation, Methodology, Writing – review & editing.

Declaration of competing interest

None Declared

Acknowledgements

We thank Daniel Bor, Sarah Garfinkel, and Andreas Roepstorff for inspiring conversations and helpful feedback. F.R. was supported by the Fellowship Programme of the Institute of Cultural and Creative Industries of the University of Kent. A.I.L. is supported by the Molson Neuro-Engineering Fellowship.

Appendix. Statistical results

Average heart rate			
	Estimate	SE	p-value
<i>Frequentist</i>			
Self stand up	10.61	0.93	<2e-16 ***
Slow tilt	9.27	0.88	<2e-16 ***
Fast tilt	10.14	0.89	<2e-16 ***
<i>Bayes $\theta = 0.01$</i>			
Self stand up	10.63	0.93	<2e-16 ***
Slow tilt	9.30	0.88	<2e-16 ***
Fast tilt	10.14	0.89	<2e-16 ***
<i>Bayes $\theta = 1$</i>			
Self stand up	10.50	0.93	<2e-16 ***
Slow tilt	9.23	0.88	<2e-16 ***
Fast tilt	9.98	0.89	<2e-16 ***

Low frequency HRV			
	Estimate	SE	p-value
<i>Frequentist</i>			
Self stand up	-0.90×10^{-4}	1.65×10^{-4}	0.589
Slow tilt	-4.97×10^{-4}	1.55×10^{-4}	0.00188 **
Fast tilt	-2.98×10^{-4}	1.58×10^{-4}	0.06238.
<i>Bayes $\theta = 0.01$</i>			
Self stand up	0.36×10^{-4}	8.53×10^{-5}	0.67212
Slow tilt	-2.56×10^{-4}	8.01×10^{-5}	0.00186 **
Fast tilt	-1.14×10^{-4}	8.14×10^{-5}	0.16308
<i>Bayes $\theta = 1$</i>			
Self stand up	2.29×10^{-5}	5.92×10^{-6}	0.000196 ***
Slow tilt	-0.69×10^{-5}	5.56×10^{-6}	0.218629
Fast tilt	0.13×10^{-5}	5.65×10^{-6}	0.821652
High frequency HRV			
	Estimate	SE	p-value
<i>Frequentist</i>			
Self stand up	-8.59×10^{-4}	2.84×10^{-4}	0.00313 **
Slow tilt	-7.08×10^{-4}	2.67×10^{-4}	0.00922 **
Fast tilt	-7.99×10^{-4}	2.71×10^{-4}	0.00399 **
<i>Bayes $\theta = 0.01$</i>			
Self stand up	-5.35×10^{-5}	1.69×10^{-5}	0.002061 **
Slow tilt	-4.21×10^{-5}	1.59×10^{-5}	0.009295 **
Fast tilt	-5.92×10^{-5}	1.61×10^{-5}	0.000396 ***
<i>Bayes $\theta = 1$</i>			
Self stand up	-2.35×10^{-7}	9.07×10^{-8}	0.522
Slow tilt	-2.08×10^{-7}	8.52×10^{-8}	0.532
Fast tilt	-2.78×10^{-7}	8.66×10^{-8}	0.490
Permutation entropy			
	Estimate	SE	p-value
<i>Frequentist</i>			
Self stand up	-0.009443	0.007571	0.215
Slow tilt	0.010748	0.007114	0.134
Fast tilt	-0.009659	0.007232	0.185
<i>Bayes, $\theta = 0.01$</i>			
Self stand up	-0.005489	0.007582	0.471
Slow tilt	0.007487	0.007125	0.296
Fast tilt	-0.006723	0.007247	0.356
<i>Bayes, $\theta = 1$</i>			
Self stand up	-0.015936	0.001813	4.61e-14 ***
Slow tilt	-0.002576	0.001703	0.134
Fast tilt	-0.012141	0.001732	2.92e-10 ***
Hurst exponent			
	Estimate	SE	p-value
<i>Frequentist</i>			
Self stand up	0.078	0.043	0.0729.
Slow tilt	0.069	0.040	0.0897.
Fast tilt	0.045	0.041	0.2743
<i>Bayes, $\theta = 0.01$</i>			
Self stand up	0.089	0.034	0.00965 **
Slow tilt	0.053	0.032	0.10012
Fast tilt	0.043	0.032	0.18427
<i>Bayes, $\theta = 1$</i>			
Self stand up	0.188	0.017	<2e-16 ***
Slow tilt	0.091	0.016	1.04e-07 ***
Fast tilt	0.159	0.016	<2e-16 ***

References

[1] G.G. Berntson, P.J. Gianaros, M. Tsakiris, Interoception and the autonomic nervous system: Bottom-up meets top-down, Oxford University Press, 2019.
 [2] U.R. Acharya, K. Paul Joseph, N. Kannathal, C.M. Lim, J.S. Suri, Heart rate variability: A review, *Med. Biol. Eng. Comput.* 44 (12) (2006) 1031–1051.

[3] F. Beissner, K. Meissner, K.-J. Bär, V. Napadow, The autonomic brain: an activation likelihood estimation meta-analysis for central processing of autonomic function, *J. Neurosci.: Off. J. Soc. Neurosci.* 33 (25) (2013) 10503–10511.
 [4] J.F. Thayer, E. Sternberg, Beyond heart rate variability, *Ann. New York Acad. Sci.* 1088 (1) (2006) 361–372.
 [5] D. Candia-Rivera, Brain-heart interactions in the neurobiology of consciousness, *Curr. Res. Neurobiol.* 3 (2022) 100050.
 [6] G.D. Clifford, Signal processing methods for heart rate variability (Ph.D. thesis), Oxford University, UK, 2002.
 [7] R. Sassi, S. Cerutti, F. Lombardi, M. Malik, H.V. Huikuri, C.-K. Peng, G. Schmidt, Y. Yamamoto, D. Reviewers, B. Gorenek, G.Y.H. Lip, G. Grassi, G. Kudaiberdieva, J.P. Fisher, M. Zabel, R. Macfadyen, Advances in heart rate variability signal analysis: joint position statement by the e-cardiology ESC working group and the European heart rhythm association co-endorsed by the Asia Pacific heart rhythm society, *EP Europace* 17 (9) (2015) 1341–1353, Publisher: Oxford Academic.
 [8] R. Acharya U, O.W. Sing, L.Y. Ping, T. Chua, Heart rate analysis in normal subjects of various age groups, *Biomed. Eng. Online* 3 (2004) 1–8.
 [9] F. Shaffer, J.P. Ginsberg, An overview of heart rate variability metrics and norms, *Front. Public Health* (2017) 258.
 [10] D. Cysarz, A. Porta, N. Montano, P. Leeuwen, J. Kurths, N. Wessel, Quantifying heart rate dynamics using different approaches of symbolic dynamics, *Eur. Phys. J. Spec. Top.* 222 (2013) 487–500.
 [11] W. Aziz, M. Arif, Multiscale permutation entropy of physiological time series, in: 2005 Pakistan Section Multitopic Conference, IEEE, 2005, pp. 1–6.
 [12] T. Penzel, J.W. Kantelhardt, L. Grote, J.-H. Peter, A. Bunde, Comparison of detrended fluctuation analysis and spectral analysis for heart rate variability in sleep and sleep apnea, *IEEE Trans. Biomed. Eng.* 50 (10) (2003) 1143–1151.
 [13] P. Castiglioni, G. Parati, A. Civijian, L. Quintin, M. Di Rienzo, Local scale exponents of blood pressure and heart rate variability by detrended fluctuation analysis: Effects of posture, exercise, and aging, *IEEE Trans. Biomed. Eng.* 56 (3) (2008) 675–684.
 [14] A.J. Camm, M. Malik, J.T. Bigger, G. Breithardt, S. Cerutti, R.J. Cohen, P. Coumel, E.L. Fallen, H.L. Kennedy, R.E. Kleiger, et al., Heart rate variability: Standards of measurement, physiological interpretation and clinical use. Task force of the European society of cardiology and the North American society of pacing and electrophysiology, *Circulation* 93 (5) (1996) 1043–1065.
 [15] J. Durbin, S.J. Koopman, Time Series Analysis By State Space Methods, vol. 38, OUP Oxford, 2012.
 [16] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Rubin, Bayesian data analysis, Chapman and Hall/CRC, 1995.
 [17] S. Gugushvili, F. van der Meulen, M. Schauer, P. Spreij, Fast and scalable non-parametric Bayesian inference for Poisson point processes, 2018, arXiv preprint arXiv:1804.03616.
 [18] C.M. Bishop, Pattern Recognition and Machine Learning, Springer, 2006.
 [19] A. Porta, E. Tobaldini, S. Guzzetti, R. Furlan, N. Montano, T. Gnecci-Ruscone, Assessment of cardiac autonomic modulation during graded head-up tilt by symbolic analysis of heart rate variability, *Am. J. Physiol. Heart Circ. Physiol.* 293 (1) (2007) H702–708.
 [20] W.H. Cooke, J.B. Hoag, A.A. Crossman, T.A. Kuusela, K.U. Tahvanainen, D.L. Eckberg, Human responses to upright tilt: A window on central autonomic integration, *J. Physiol.* 517 (Pt 2) (1999) 617–628.
 [21] N. Montano, T.G. Ruscone, A. Porta, F. Lombardi, M. Pagani, A. Malliani, Power spectrum analysis of heart rate variability to assess the changes in sympathovagal balance during graded orthostatic tilt, *Circulation* 90 (4) (1994) 1826–1831.
 [22] A.L. Goldberger, L.A. Amaral, L. Glass, J.M. Hausdorff, P.C. Ivanov, R.G. Mark, J.E. Mietus, G.B. Moody, C.K. Peng, H.E. Stanley, PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals, *Circulation* 101 (23) (2000).
 [23] T. Heldt, M. Oefinger, M. Hoshiyama, R. Mark, Circulatory response to passive and active changes in posture, in: *Computers in Cardiology*, 2003, 2003, pp. 263–266, ISSN: 0276-6547.
 [24] T. Heldt, E.B. Shim, R.D. Kamm, R.G. Mark, Computational modeling of cardiovascular response to orthostatic stress, *J. Appl. Physiol.* 92 (3) (2002) 1239–1254, Publisher: American Physiological Society.
 [25] D. Candia-Rivera, V. Catrambone, G. Valenza, The role of electroencephalography electrical reference in the assessment of functional brain–heart interplay: From methodology to user guidelines, *J. Neurosci. Methods* 360 (2021) 109269.
 [26] P. McSharry, G. Clifford, L. Tarassenko, L. Smith, A dynamical model for generating synthetic electrocardiogram signals, *IEEE Trans. Biomed. Eng.* 50 (3) (2003) 289–294.
 [27] R.K. Mohn, Modelling the natural pacemaker of the heart as a pulse-frequency modulator, *Med. Biol. Eng. Comput.* 16 (1) (1978) 90–97.
 [28] M.N. Levy, P.J. Martin, Neural regulation of the heart beat, *Annu. Rev. Physiol.* 43 (1) (1981) 443–453.
 [29] R. Bailón, G. Laouini, C. Grao, M. Orini, P. Laguna, O. Meste, The integral pulse frequency modulation model with time-varying threshold: application to heart rate variability analysis during exercise stress testing, *IEEE Trans. Bio-Med. Eng.* 58 (3) (2011) 642–652.
 [30] D. Candia-Rivera, V. Catrambone, R. Barbieri, G. Valenza, Integral pulse frequency modulation model driven by sympathovagal dynamics: Synthetic vs. real heart rate variability, *Biomed. Signal Process. Control* 68 (2021) 102736.

- [31] M. Brennan, M. Palaniswami, P. Kamen, Poincaré plot interpretation using a physiological model of HRV based on a network of oscillators, *Am. J. Physiol. Heart Circ. Physiol.* 283 (5) (2002) H1873–1886.
- [32] M.M. Waldrop, *Complexity: The Emerging Science At the Edge of Order and Chaos*, Simon and Schuster, 1993.
- [33] C.-K. Peng, S.V. Buldyrev, S. Havlin, M. Simons, H.E. Stanley, A.L. Goldberger, Mosaic organization of DNA nucleotides, *Phys. Rev. E* 49 (2) (1994) 1685.
- [34] J.W. Kantelhardt, E. Koscielny-Bunde, H.H. Rego, S. Havlin, A. Bunde, Detecting long-range correlations with detrended fluctuation analysis, *Physica A* 295 (3–4) (2001) 441–454.
- [35] E.A. Ihlen, Introduction to multifractal detrended fluctuation analysis in matlab, *Front. Physiol.* 3 (2012) 141.
- [36] A. Eke, P. Herman, L. Kocsis, L. Kozak, Fractal characterization of complexity in temporal physiological signals, *Physiol. Meas.* 23 (1) (2002) R1.
- [37] P. Castiglioni, D. Lazzeroni, P. Coruzzi, A. Faini, Multifractal-multiscale analysis of cardiovascular signals: A DFA-based characterization of blood pressure and heart-rate complexity by gender, *Complexity* 2018 (2018) 1–14.
- [38] D.P. Francis, K. Willson, P. Georgiadou, R. Wensel, L.C. Davies, A. Coats, M. Piepoli, Physiological basis of fractal complexity properties of heart rate variability in man, *J. Physiol.* 542 (2) (2002) 619–629.
- [39] K. Willson, D.P. Francis, R. Wensel, A.J. Coats, K.H. Parker, Relationship between detrended fluctuation analysis and spectral analysis of heart-rate variability, *Physiol. Meas.* 23 (2) (2002) 385.
- [40] J.L. Rojo-Álvarez, A. Sanchez-Sanchez, O. Barquero-Perez, R. Goya-Esteban, E. Everss, I. Mora-Jimenez, A. Garcia-Alberola, Analysis of physiological meaning of detrended fluctuation analysis in heart rate variability using a lumped parameter model, in: *2007 Computers in Cardiology, IEEE, 2007*, pp. 25–28.
- [41] S. Bianchi, Fathon: A python package for a fast computation of detrended fluctuation analysis and related algorithms, *J. Open Sour. Softw.* 5 (45) (2020) 1828.
- [42] C. Bandt, B. Pompe, Permutation entropy: A natural complexity measure for time series, *Phys. Rev. Lett.* 88 (17) (2002) 174102.
- [43] M. Zanin, L. Zunino, O.A. Rosso, D. Papo, Permutation entropy and its main biomedical and econophysics applications: A review, *Entropy* 14 (8) (2012) 1553–1577.
- [44] C. Bian, C. Qin, Q.D. Ma, Q. Shen, Modified permutation-entropy analysis of heartbeat dynamics, *Phys. Rev. E* 85 (2) (2012) 021906.
- [45] A.A. Pessa, H.V. Ribeiro, Ordpy: A python package for data analysis with permutation entropy and ordinal network methods, *Chaos* 31 (6) (2021) 063110.
- [46] F.E. Rosas, P.A. Mediano, C. Timmermann, A.I. Luppi, D. Candia-Rivera, R. Abbasi-Asl, A. Gazzaley, M.L. Kringelbach, S.D. Muthukumaraswamy, D. Bor, et al., The entropic heart: Tracking the psychedelic state via heart rate dynamics, 2023, *bioRxiv*, 1–11.
- [47] D. Candia-Rivera, V. Catrambone, R. Barbieri, G. Valenza, Functional assessment of bidirectional cortical and peripheral neural control on heartbeat dynamics: A brain-heart study on thermal stress, *NeuroImage* 251 (2022) 119023.
- [48] D. Candia-Rivera, Modeling brain-heart interactions from poincaré plot-derived measures of sympathetic-vagal activity, *MethodsX* 10 (2023) 102116.
- [49] J.-L. Elghozi, D. Laude, A. Girard, Effects of respiration on blood pressure and heart rate variability in humans, *Clin. Exp. Pharmacol. Physiol.* 18 (11) (1991) 735–742.
- [50] R. Barbieri, E.C. Matten, A.A. Alabi, E.N. Brown, A point-process model of human heartbeat intervals: New definitions of heart rate and heart rate variability, *Am. J. Physiol.-Heart Circ. Physiol.* 288 (1) (2005) H424–H435, Publisher: American Physiological Society.
- [51] G. Valenza, L. Citi, J.P. Saul, R. Barbieri, Measures of sympathetic and parasympathetic autonomic outflow from heartbeat dynamics, *J. Appl. Physiol.* 125 (1) (2018) 19–39.
- [52] M.T. Sherman, H.-T. Wang, S.N. Garfinkel, H.D. Critchley, The cardiac timing toolbox (CaTT): Testing for physiologically plausible effects of cardiac timing on behaviour, *Biol. Psychol.* 170 (2022) 108291.
- [53] A.T. Cemgil, O. Dikmen, Conjugate gamma Markov random fields for modelling nonstationary sources, in: *International Conference on Independent Component Analysis and Signal Separation, Springer, 2007*, pp. 697–705.
- [54] O. Dikmen, A.T. Cemgil, Unsupervised single-channel source separation using Bayesian NMF, in: *2009 IEEE Workshop on Applications of Signal Processing To Audio and Acoustics, IEEE, 2009*, pp. 93–96.
- [55] S. Geman, D. Geman, Stochastic relaxation, gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.* (1984) 721–741.
- [56] A.E. Gelfand, A.F. Smith, Sampling-based approaches to calculating marginal densities, *J. Am. Stat. Assoc.* 85 (410) (1990) 398–409.
- [57] S. Brooks, A. Gelman, G. Jones, X.-L. Meng, *Handbook of Markov Chain Monte Carlo*, CRC Press, 2011.
- [58] P. Domingos, M. Pazzani, On the optimality of the simple Bayesian classifier under zero-one loss, *Mach. Learn.* 29 (1997) 103–130.
- [59] T. Heldt, M. Oefinger, M. Hoshiyama, R. Mark, Circulatory response to passive and active changes in posture, in: *Computers in Cardiology, 2003, IEEE, 2003*, pp. 263–266.