

# **Title:** The Cardiac Timing Toolbox (CaTT): Testing for physiologically plausible effects of cardiac timing on behaviour

**Authors:** Maxine T. Sherman<sup>1,2,3</sup>, Hao-Ting Wang<sup>1,2</sup>, Sarah N. Garfinkel<sup>2,4</sup>, Hugo D. Critchley<sup>1,2</sup>

## **Affiliations:**

[1] Sackler Centre for Consciousness Science, University of Sussex, Brighton, UK

[2] Brighton and Sussex Medical School, Brighton, UK

[3] School of Engineering and Informatics, University of Sussex, Brighton, UK

[4] Institute of Cognitive Neuroscience, UCL, London, UK

**Corresponding author:** Maxine Sherman, [maxinesherman@gmail.com](mailto:maxinesherman@gmail.com)

**Data availability:** The toolbox, with associated code and data presented in this manuscript, is freely available for download and non-commercial reuse at <https://github.com/MaxineSherman/CaTT>

**Funding:** MTS, HTW, and HDC are supported by and thank the Mortimer and Theresa Sackler Foundation.

**Acknowledgements:** CaTT is dependent on several other toolboxes and we thank the respective authors for releasing their code on free-to-share licences: M. Sabarimalai Manikandan, K. P. Soman, D. Muir, and F. Cramer. We also thank A. Galvez-Pol, R. McConnell and J. Kilner for sharing their data openly, allowing us to use their valuable dataset here.

## **Highlights:**

- We present an easy-to-use MATLAB toolbox to test for cardiac timing effects.
- We propose a physiologically motivated method of quantifying cardiac time.
- We propose a set of statistical tests with minimal distributional assumptions.
- We demonstrate the utility of our approach by reanalysing open source data.

## **Abstract**

There is a long history of, and renewed interest in, cardiac timing effects on behaviour and cognition. Cardiac timing effects may be identified by expressing events as a function of their location in the cardiac cycle, and applying circular (i.e. directional) statistics to test cardiac time-behaviour associations. Typically this approach ‘stretches’ all points in the cardiac cycle equally, but this is not necessarily physiologically valid. Moreover, many tests impose distributional assumptions that are not met by such data. We present a set of statistical techniques robust to this, instantiated within our new Cardiac Timing Toolbox (CaTT) for MATLAB: A physiologically-motivated method of wrapping behaviour to the cardiac cycle; and a set of non-parametric statistical tests that control for common confounds and distributional characteristics of these data. Using a reanalysis of previously published data, we guide readers through analyses using CaTT, aiding researchers in identifying physiologically plausible associations between heart-timing and cognition.

## Introduction

There is a resurgence of interest in how physiological states of the body influence aspects of cognition and emotion (1,2). From a historical context of the 'peripheral' theory of emotional feelings (3,4), notions of embodiment are increasingly incorporated into contemporary theoretical models in consciousness science (5,6), and in influential computational models of brain function that emphasise the predictive control of the body's changing visceral state (7–9).

Information about the changing visceral state is continuously received by the brain, and one particularly important interoceptive channel carrying information about the dynamic state of cardiac activation is pressure-sensing baroreceptors. These are activated as blood is ejected from the heart into the aorta and carotid arteries with each heart beat (10), and signal, through afferent nerves to the brainstem, the strength and timing of individual heart beats and hence the state of cardiovascular activation. Since these signals peak around the heart beat (cardiac ventricular systole) and are reduced in the interval between heart beats (ventricular diastole), the influence of cardiovascular activation signals on perception, cognition and action may be revealed (even in states of low arousal) by testing for cardiac timing effects. In other words, one can test whether behaviours are facilitated or inhibited at different phases of the beat-to-beat cardiac cycle.

The cardiac cycle describes the repeated sequence of electrical, muscular and haemodynamic changes preceding, during and after ejection of blood from the heart. The key haemodynamic phases of the cardiac cycle are ventricular systole, when the ventricular muscle contracts blood and blood is ejected from the heart, i.e. the 'heart beat', and ventricular diastole - the period between heart beats when ventricles fill. Ventricular contraction is triggered by the propagation of myocardial depolarisation, apparent in the standard electrocardiogram (ECG) as the QRS complex (notably the R wave peak). The subsequent ejection of blood is maximal over the period 250-300ms after the R wave, during which time arterial baroreceptors in aorta and carotid sinus are most active (Fig. 1). The ECG T wave, peaking around 300ms after the R wave, reflects ventricular repolarisation (electrical recovery of the ventricular myocardium). Baroreceptor quiescence follows, in anticipation of the next heart beat.




Figure 1. The cardiac cycle. The QRS complex represents ventricular depolarisation; the T wave represents ventricular repolarisation.

Since Lacey and Lacey's work on the relevance of heart rate to behaviour (11), effects of cardiovascular activation on other aspects of perception and cognition have also been linked to baroreceptor signalling and cardiac cycle effects (2,5). Similarly, neural evoked responses to painful and somatosensory stimuli also vary with cardiac cycle (12,13). The prediction here is that the impact on cognition should be greater when baroreceptor output is maximum (ventricular systole) than when it is not (ventricular diastole). To test for effects of cardiac timing on behaviour, and thus for associations between bodily state and cognition, this branch of cardiac interoception research has taken one of two main approaches.

One involves *a priori* time-locking stimulus presentation to systole (set according to the R wave peak) or diastole (set as  $R + \delta$ , where  $\delta$  is set by the researcher) and comparing behaviour, such as fear intensity ratings and other emotional categories (14,15), pain responses (13,16–19), memory performance (20,21), and motor responses (22–24), across the two timing conditions. Relative to diastole, pain and somatosensory perception are attenuated at systole (12,13,19,25). Another more naturalistic approach is to collect ongoing ECG (or similar data) during a task and test for associations between cardiac timing and behaviour *post-hoc* (25–29).

Although in the first time-locking approach, stimuli can of course be presented at multiple intervals across the cardiac cycle (19), this often leads to unwieldy experimental designs and a potentially disproportionate number of trials time-locked to systole. One advantage with the second approach (which comes at the cost of being unable to make directional inferences, i.e. behaviour may change heart rate) is that the researcher can identify whether particular

points, or phases, of the cardiac cycle other than ventricular systole or diastole facilitate performance. More importantly, the researcher does not have to commit to a fixed latency between systole and diastole. Instead, cardiac timing is expressed relative to the R wave peak.

Nevertheless, testing for such associations between cardiac time and behaviour is non-trivial. We elaborate on these issues later in the manuscript, but briefly, the first issue is that such tests necessitate the use of circular statistics (statistics for directional variables such as compass directions or months of the year), and the set of available statistical tests for circular variables is far smaller than for their linear counterparts. Of course, there will be cases where (linear) time since the last R wave peak is of primary interest (e.g. effects relating to pulse transit time), and in these cases circular statistics will not be necessary. The second issue is that many circular statistical tests make distributional assumptions that will not necessarily be met by the data at hand. Finally, the assumptions underlying standard transformations from milliseconds since the last R wave to cardiac time as expressed as an angle – specifically the assumption that one can uniformly stretch all timepoints - are not necessarily reasonable either.

These complications are important to consider and address to provide valid measures of interoceptive influences on mental processes. In studies focusing on individual differences in the perception of bodily signals, there is considerable ongoing debate about the construct validity of Schandry's heart beat counting task (30) and similar heart beat tracking tasks. Here, performance is confounded by the influence of non-interoceptive information (e.g. exteroceptive information (31) and knowledge of one's average heart rate (32,33)) to the extent that quantifications of interoceptive accuracy are affected by true heart beat timing to only a small degree (34). These discussions continue to motivate improved paradigms for characterising the accuracy of and biases in interoceptive reports (35–37). Nevertheless, there has been considerably less attention given development in measurement and statistical analysis for cardiac timing studies.

In this manuscript, we discuss and develop statistical methods for quantifying cardiac timing effect on behaviour. We propose a novel method for expressing behaviour as a (mathematical) function of the cardiac cycle, and a set of robust statistical analyses to address whether cardiac time differs across response types, correlates with responses, or facilitates/inhibits responses. We further present an easy-to-use MATLAB (The MathWorks Inc. Natick, USA) toolbox – the Cardiac Timing Toolbox (CaTT) - that implements these analyses, and that also can import, pre-process and visualise the user's data. Finally, we

illustrate how our toolbox can be used to probe, sensitively and robustly, associations between cardiac timing and behaviour by reanalysing a previously published, open dataset.

## Methods

### Overview of the Cardiac Timing Toolbox

The toolbox is designed for researchers with little experience using Matlab, to guide the user through the entire analysis pipeline. The toolbox can be downloaded at <https://github.com/MaxineSherman/CaTT>. The functions contained enable the user to preprocess the ECG, conduct statistical analyses on cardiac time-behaviour associations, and perform data visualisation. Detailed descriptions of the functions, shown in Fig.2, can be found by calling, for example, 'help catt\_denoise' in Matlab, and tutorial scripts (one for preprocessing, one for analysis) can be found in the CaTT demo folder. An example pipeline for preprocessing and analysing data using CaTT can be found in Table 1. Demo scripts for preprocessing, analysing and visualising the data can be found in the main folder of the toolbox (catt\_demo\_preprocessing, catt\_demo\_analysis, and catt\_demo\_plotting respectively).

*Table 1. Example pipeline for preprocessing and analysing ECG data with CaTT*

Analysis step	Code
1. Initialize toolbox	catt_init
2. Load data into MATLAB workspace	Custom code
3a. If unknown, estimate sample rate	catt_estimate_srate
3b. If unknown, estimate timestamps	catt_estimate_times
4. Import data into CaTT	catt_import
5. Denoise the ECG data	catt_denoise
6. Identify R wave peaks and the peak and end of the T wave	catt_heartbeat_detection
7. Epoch ECG data into RR intervals	catt_epoch
8. Remove RR intervals with artefacts	catt_manualRejection
9. Estimate IBIs and remove RR intervals with extreme or nonsensical IBIs	catt_IBI
10. Estimate heart rate variability	catt_HRV
11. Estimate QT latency from ECG data, from heart rate, or set according to toolbox parameters	catt_estimate_qt
12. Express onsets as cardiac angles	catt_wrap2heart
13. For each participant, estimate whether onsets are more likely at a particular cardiac angle	catt_bootstrap_clust
14. Pool individual participant z-scores into a group z-score and group p-value	catt_z2p
15. Test for group-level consistency	catt_consistency
16. Test for correlation between cardiac angles and behavior	catt_bootstrap_corr
17. Test for difference in mean/median cardiac angle across conditions	catt_bootstrap_diff
18. Plot the onsets, IBIs and cardiac angles to depict results	catt_plot_circ catt_plot_ibi_dist catt_plot_onset_dist

This manuscript will concentrate on explaining the motivation for and advances of the toolbox, as well as presenting analyses performed with CaTT on real data. To anticipate, we use RR interval to refer to the ECG data between two successive R wave peaks, and interbeat interval (IBI) to refer to the latency (in ms) between the R wave peaks.


### Installing the toolbox

To download the toolbox, click the “Code” button on the Github page ([github.com/MaxineSherman/CaTT](https://github.com/MaxineSherman/CaTT)). The source code will be packaged into a zip file. Move the zip file into a parent folder of choice and unzip it. Then, in Matlab, navigate to the parent folder and use the following command to add the toolbox to path:


```
addpath(genpath('CaTT-main'));
```

We recommend using MATLAB version 2021b, on which this toolbox was developed and tested.

**A**



**B**



*Figure 2.* Outline of the CaTT Toolbox. **(A)** List of Matlab functions included in the toolbox, grouped by their usage. **(B)** Example flow chart showing the order in which to call the functions. The user needs to begin by initialising the toolbox and calling the data. Next, if they wish to preprocess using CaTT then they would denoise, detect the peaks of the R and T waves, epoch the data into heart beats, then reject epochs with artefacts. Participant-level statistics heart beat variability (HRV) and mean interbeat interval (IBI, in ms) are calculated. Next, the user must estimate the QT interval (according to the method they set in the `catt_opts`), onsets are wrapped to the cardiac cycle, and finally, participant- and group-level statistical tests are run.

## Initializing the toolbox

The toolbox is initialized by calling the function `catt_init`, which specifies default settings for subsequent preprocessing, statistical inference and data visualisation steps. Opening the function file reveals a set of parameters, loaded into the structure `catt_opts`, which can be changed by the user according to their needs. Alternatively, parameters can be changed at the command line, for example the sample rate (which defaults at 512Hz) can be changed to 1000Hz as follows:

```
catt_opts.fs = 1000;
```

Scripts that call CaTT should begin by declaring the to-be-created CaTT options structure, `catt_opts`, as a *global* variable and then initialising the options by calling `catt_init`, as follows:

```
global catt_opts;
catt_init;
```

## Importing and cleaning data

To preprocess ECG data in CaTT the toolbox requires, at a minimum, 3 pieces of data: a vector of continuous, (not epoched) ECG data; a vector of timestamps for each ECG sample (in ms); and onsets of events to timelock to the cardiac cycle, for example, stimulus presentation times (in ms). These onsets correspond to the datapoints that will later be converted to cardiac time, expressed in angles. If the researcher does not have timestamps/the sampling rate then, assuming data were sampled continuously, they/it can be constructed from the size of the ECG vector and the sampling rate/timestamps using the function `catt_estimate_times/catt_estimate_srate`. There is also an optional fourth variable one can use: *responses*. This should be numeric, and the same size as `onsets_ms`. It can be a binary variable (e.g. 0 = incorrect; 1 = correct), an ordinal variable (e.g. Likert scale ratings, or indeed any rating scale, for each stimulus presented) or a continuous variable (e.g. reaction times to the stimulus presented). Alternatively, the researcher might choose to store condition-level information here (e.g. 1 = angry face; 2 = neutral face; 3 = emotional face). The purpose of this variable is for use later-on, when comparing associations between cardiac timing and treatment effect across conditions or response types.

In this manuscript (and in the toolbox scripts) we name to the data structure created and used by CaTT as `catt`. This structure is created by passing the three or four aforementioned variables to `catt_import` as follows:



```
catt = catt_import( ECG, times_ms, onsets_ms );
```

or

```
catt = catt_import( ECG, times_ms, onsets_ms, responses );
```

Typically, researchers will have data from many different participants. Each participant's data structure should be stored into a *cell array* to facilitate group-level statistical tests at later stages of analysis. For example, Participant One's data will be stored in `catt{1}`, Participant Two's data in `catt{2}`, etc.

### **Preprocessing the cardiac data**

While users are of course free to import data that are already pre-processed, the toolbox can also pre-process raw ECG data. The pre-processing pipeline is adopted from that proposed in (26). First, continuous ECG data are detrended. This is done by applying median filter 200ms wide to the ECG, followed by another 600ms wide median filter. The result constitutes a baseline that is subtracted from the raw signal. Second, the detrended data are denoised using a three-level Daubechies wavelet filter with one vanishing moment. These steps are performed by the script `catt_denoise`.

The next step is to detect peaks of the R and T waves, and the T wave end, by calling `catt_heartbeat_detection`. For R wave peak detection, CaTT implements the powerful detection method developed by Manikandan & Soman (38), using the authors' code, which is packaged into CaTT (<https://github.com/hongzuL/A-novel-method-for-detecting-R-peaks-in-electrocardiogram-signal>). If the user finds that R wave peak detection is not working well (because their ECG data are particularly noisy) then they can change the parameter `catt_opts.rdetection_thresh`. This parameter determines whether a deflection in the ECG is categorised as an R wave peak or not, and we would therefore recommend increasing its value if the ECG data are particularly noisy.

To detect the peaks and ends of each T wave, we implemented the algorithm by Vásquez-Seisdedos and colleagues (39), which is commonly adopted in the literature (12,25,27): For each R wave peak detected, peaks of the corresponding T wave are found by finding the maximum ECG amplitude in the period between R+200ms and R+500ms. The minimum and maximum values can be changed by changing the settings for `catt_opts.RT_min` and `catt_opts.RT_max` respectively. Detecting the end of each T wave involves searching a constrained region after the T wave peak but before the R wave peak for the point at which the gradient indicates the end of the wave.

With cardiac events identified, the third step is to epoch the data, not into trials but into discrete RR intervals. This is done by `catt_epoch`. Each epoch begins at an R wave peak and ends at the sample preceding the next, with relevant information stored in the field `catt.RR`. Relevant information here includes: T wave timings; R-T time (the R wave peak to the end of the T wave); the inter-beat interval; if present, onset time expressed relative to the last R wave peak; and, if present, the response associated with the onset. Note that any epoch not containing an onset or response will have these fields populated with NaN.

Manual artefact rejection is performed by the script `catt_manualRejection`. The toolbox will open a mounted figure into the MATLAB GUI and plot, epoch-by-epoch, the processed ECG with detected R wave peaks and T waves overlaid. The user then clicks R wave peaks to reject the RR interval. Rejected intervals are greyed-out and stored (along with associated behaviour, onset and conditions) into a separate field of the data structure called `catt.rej`.

Once pre-processing is complete, the script calls `catt_IBI` to calculate: the vector of all IBIs in a trial or epoch; the IBI of the cardiac cycle (RR interval) in which the onset appeared; and the onset, expressed as time in milliseconds since the last R wave peak.

An important, second data quality check is performed by `catt_IBI`: the function will flag IBIs which are implausibly long or short, perhaps due to human error during manual artefact rejection or due to measurement noise. This is done in 3 ways:

- i. IBIs are z-scored and compared to a threshold called `catt_opts.BPM_extreme_z`. Its default value is 3, but this can be changed in `catt_init`. If  $|z| > \text{catt\_opts.BPM\_extreme\_z}$  then the corresponding RR is excluded from further analysis. To switch off exclusion-by-zscore the user can set `catt_opts.BPM_extreme_z` to `inf` (because  $|z|$  will never be greater than infinity)
- ii. IBIs are converted to beats per minute (bpm). Very short IBIs, corresponding to very high bpm, will be excluded if  $\text{bpm} > \text{catt\_opts.BPM\_max}$ . The default value is 160bpm. To switch off exclusion of very short IBIs, set `catt_opts.BPM_max` to `inf`.

- iii. Very long IBIs, corresponding to very low bpm, will be excluded if  $\text{bpm} < \text{catt\_opts.BPM\_min}$ . The default value is 40bpm. To switch off exclusion of very long IBIs, set  $\text{catt\_opts.BPM\_min}$  to 0.

Excluded RR intervals will be stored in  $\text{catt.rej}$ . A histogram of retained versus removed IBIs can be plotted by calling:

```
plot_on = true
catt = catt_IBI( catt, plot_on )
```

Finally, heart rate variability (HRV) is calculated using the function  $\text{catt\_HRV}$ . By default, the function calculates HRV as RMSSD (the root mean squared of successive differences), but it can also be calculated using other measures as well.

### **Expressing position in the cardiac cycle as angular data**

Cardiac timing – where in the RR interval a response was made – is, like any time (e.g. months of the year), circular. Circular variables are difficult to use for statistical inference, because unlike linear variables (e.g. continuous fear ratings or rating scale responses) where the difference between two values is simply their subtraction, circular variables “wrap”, thereby requiring modular arithmetic. To illustrate, while  $22 + 3 = 25$ , three hours after 22:00 hours is not 25:00 hours, but 01:00 (1am). Similarly, while in linear space 11 is closer to 8 (difference of 3) than 1 (difference of 10), when working on 12-hour clock time the opposite is true: 11am is further from 8am (difference of 3) than 1pm (difference of 2). When stimulus or response onsets are expressed in terms of their position in the cardiac cycle, any statistical inference would be performed on an “inter-beat-interval-hour clock”, that is, using circular arithmetic where the modulus on each trial is the inter-beat interval, which itself is subject to variability over time.

In order to test cardiac timing effects on behaviour, onsets (which refers to either stimulus onsets or response times – whatever you wish to time-lock to your cardiac data) need to be expressed as a function of the cardiac cycle. This is done by expressing time as degrees in radians, denoted by  $\theta$ . Just as a quarter of an hour would be represented as a  $90^\circ$  degree ( $0.5\pi$  rad) angle on a clock face – the 15-minute mark – in cardiac time, a quarter of a cardiac cycle, would be represented as 90 degrees ( $0.5\pi$  rad) as well. There are two methods available for wrapping these times to the ECG:

The first expresses the  $k^{\text{th}}$  onset in milliseconds,  $o_k$ , as a proportion of the latency between the two surrounding R wave peaks, denoted  $IBI_k$ . Cardiac time  $\theta_{R\ peak_k}$  is given in radians by:

$$\theta_{R\ peak_k} = \frac{2\pi o_k}{IBI_k}$$

To wrap a vector of onsets (*onsets*) to R wave peaks via a vector of inter-beat intervals expressed in milliseconds (*IBIs*), first you would set the parameter structure as shown below, then you would call the wrapping function:

```
catt_opts.wrap2 = 'rpeak';
or
catt_opts.wrap2 = 'twav';
catt.wrapped = catt_wrap2heart(catt);
```

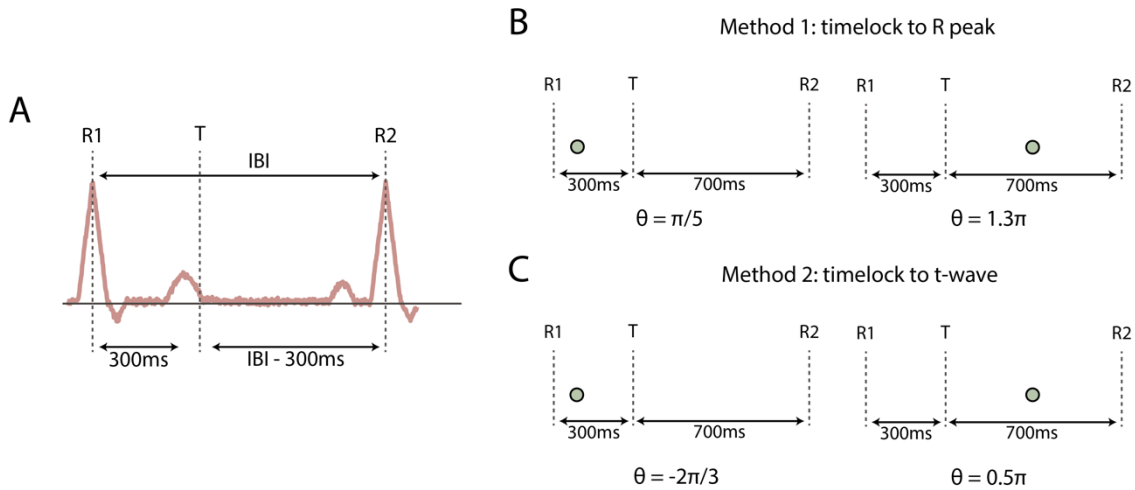
This first method treats all time points in the cardiac cycle equally, such that the relationship between  $\theta_{R\ peak_k}$  and the IBI is the same, regardless of where in the cardiac cycle the onset appeared (e.g. before versus after the T wave). This means that onsets on two different trials can take the same value of  $\theta_{R\ peak_k}$ , despite one falling before the T wave (i.e. prior to ventricular repolarisation) and another falling after the T wave.

The toolbox affords a second, more biologically plausible method for expressing onsets in cardiac time. We know that the latency between the QRS complex and the T wave is largely constant; the interval from R wave peak to T wave peak lasts approximately 300ms (with further small within-participant differences that depend on heart rate). The variability in IBIs is primarily driven by variability in ventricular repolarization and diastole time (i.e. the interval between the T wave and next R wave) under autonomic control (40). Accordingly, the CaTT's second wrapping method does not treat ECG samples across the cardiac cycle equally. Instead, onsets are expressed as a proportion of the cardiac cycle, but relative to the T wave. Onsets occurring between R wave peak and T wave end versus T wave end and R wave peak are transformed separately, and as follows:

$$\theta_{T\ wave_k} = \begin{cases} \frac{\pi(o_k - T - R)}{T - R}, & o_k \text{ before } T \text{ wave} \\ \frac{\pi(o_k - T - R)}{IBI_k - (T - R)}, & o_k \text{ after } T \text{ wave} \end{cases}$$

What this means is that onsets falling before the T wave have negative angles that represent how far away the onset was from the T wave, as a proportion of R-T time. Furthermore, when angles are negative then the onset can be categorised as having occurred during systole.

Analogously, onsets falling after the end of the T wave take positive angles which represent how far away from T the onset was as a proportion of T-R time. In this way, whatever the IBI is, the angle  $\theta$  will always tell you whether the response came before or after the T wave,



**Figure 3. Wrapping stimulus onsets or response times to the cardiac cycle.** **A** Schematic of one RR interval, from the R wave peak to the T wave end up to the second R wave peak. **B** Expressing onsets (green dot) in cardiac time, R wave peak method. Left panel depicts an onset at 100ms post R1, where the IBI is 1000ms, meaning  $\theta_{R\ peak} = 2\pi(100/1000) = \pi/5$ . Right panel depicts an onset at 650ms post R1, giving  $\theta_{R\ peak} = 2\pi(700/1000) = 1.3\pi$ . **C**. Wrapping onsets to the T wave. For the same values as in panel B, when the onset arrives before the T wave then  $\theta_{T\ wave}$  is negative, and takes the value  $\pi(100-300)/300 = -3\pi/4$  and  $\pi(650-300)/(1000-300) = 0.5\pi$  respectively.

and will respect the latency of the refractory period. Onsets taking positive angles can be classified as occurring during diastole.

To illustrate, suppose we assume the time between the R wave peak and end of the T wave to be fixed at 300ms, and large deviations from 300ms to be attributable to measurement noise. If a stimulus is presented between an R and T wave (in the refractory period), then that stimulus onset time is expressed as a (negative) proportion of that 300ms period. Here, the most extreme values are highly negative, i.e., very far away from T and thus close to R, and the least extreme values (those closest to 0) are very close to T and thus far from R. If instead the stimulus is presented between the T and R wave, then the calculation of  $\theta_{T\ wave_k}$  is different. In this case, we take the onset as a (positive) proportion of the time between T and R wave, i.e., the inter-beat interval minus 300ms. As before, extreme positive values mean that the stimulus was presented just before the next R wave peak and thus far from

the T wave; values close to 0 indicate that the stimulus appeared just after the T wave. This is illustrated in Figure 3.

In cases where the CaTT user has estimated T wave time from their data, these timepoints can be used in the calculation of  $\theta_{T\ wave_k}$ . Where the user is unable to estimate these timepoints (e.g. because they only R wave peak times and not the raw ECG) then  $\theta_{T\ wave_k}$  can be estimated as follows:

One needs to estimate when T wave occurs from the R wave, the participant's heart rate, and from an assumption on average QR-T time. The toolbox assumes the average QT time to be 400ms and the QR time to be 50ms. Therefore, the default R-T time is 350ms. However, there are also three implemented methods for estimating a corrected QR-T interval from heart rate:

1. Bazett's formula (41):  $QT_c = \frac{QT}{\sqrt{bpm/60}}$
2. Fridericia's formula (42):  $QT_c = \frac{QT}{\sqrt[3]{bpm/60}}$
3. Sagie's formula (43):  $QT_c = 1000 \times \left( \frac{QT}{1000} + 0.154 \left( 1 - \frac{bpm}{60} \right) \right)$

where QT is the QT interval (again, here assumed to be 400ms = 350ms + 50ms QR time) and  $QT_c$  is the QT interval corrected for heart rate. Across all of these corrections there is an assumption of a stable heart rate. This is reasonable for many experimental designs, though manipulations or stimuli likely to substantially change heart rate (e.g. delivery of electric shocks) then this assumption may be violated. In these cases we would recommend estimating the end of the T wave from the ECG data. Of course, regardless of the method chosen, expressing stimulus or response times as a function of their time from the previous R wave peak or T wave does not mean that they were *caused* by that latency.

### Statistical inference on cardiac phase

The CaTT statistical inference functions are designed to address three types of analysis question, and in each case the toolbox calls function from the *circstat* toolbox (44) to implement the appropriate tests. For ease of use, this toolbox is packaged into CaTT, and for details on how the statistical tests work, we refer the reader to Berens' manuscript. Furthermore, in each case CaTT uses permutation testing to calculate z-scores and p-values. The analysis types are as follows:

1. Is there an **association** between the phase of the cardiac cycle and some behaviour?
2. Is there a **difference** in cardiac phase between two conditions?

### 3. Are response times **clustered** around some point in the cardiac cycle?

To test for associations between phase and behaviour, we can use circular correlations. For example, one may wish to test whether there is a correlation between how far along the cardiac cycle a face is presented is, and how intensely that face is rated. Alternatively, one could test for a correlation between two cardiac times, for example, the (cardiac) time when a face was presented and the (cardiac) time when a response was made.

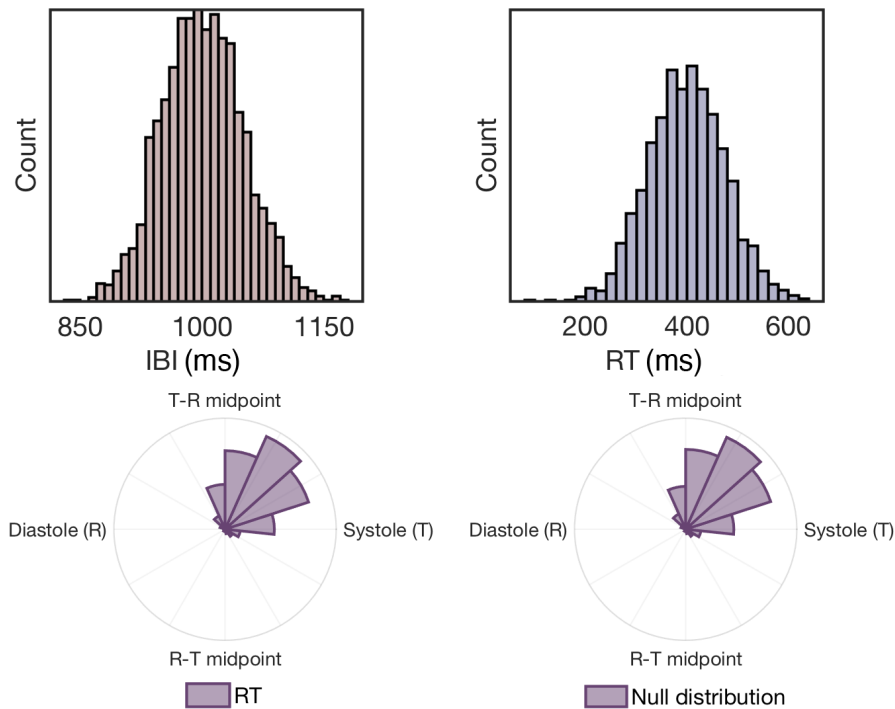
In order to test for such correlations, you can call the function `catt_bootstrap_corr`, passing the function your variables and also whether they are linear (e.g. intensity ratings) or angular (i.e. cardiac time). Correlations are calculated using the *circstat* toolbox function `circ_corrcl` (for correlation between circular and linear variables) or `circ_corrcc` (for correlations between two circular variables) as appropriate. The function calculates p-values via permutation testing, shuffling one of the variables to break the association with the other and recomputing the correlation over many interactions. The p-value is then the proportion of permutations (shuffles) that resulted in a correlation greater than or equal to the empirical correlation.

Similarly, if one wishes to test for a difference in cardiac phase between conditions, for example to see whether a decision to bet versus pass in a game is associated with different portions of the cardiac cycle, the function `catt_bootstrap_diff` can be called. This also uses permutation testing, randomly allocating each datapoint to one or another label over multiple iterations to obtain a p-value. Here, CaTT offers several options: the function can be used on linear (e.g. ratings) or circular (e.g. cardiac timing) data; condition differences can be calculated using the mean (for e.g. Gaussian or Von Mises data) or the median (for e.g. skewed data); and for paired or independent designs.

Finally, one may seek to test if responses, for example button-presses, are more likely to occur at particular portions of the cardiac cycle. This would be a test of circular uniformity, asking whether cardiac times are clustered about some point. To test this, the toolbox function `catt_bootstrap_clust` would be called. The function offers two choices of test: The Rayleigh test (45), which assumes a Von Mises distribution (“circular Normality”) and that, if non-uniformity (“clustering”) is present, that it is unimodal; and Rao’s spacing test (46), which does not have distributional assumptions (akin to a “non-parametric” test). Both are implemented by *circstat* code (47). This third analysis type – testing for (lack of) circular uniformity – is arguably more complicated than the former two, for reasons that will be laid out in the following section.

## The importance of permutation testing

Frequentist statistical tests (e.g. t-test, ANOVA) are typically employed in experimental research in cognitive neuroscience and cognitive psychology to test for effects of group on



**Figure 4.** The importance of permutation testing. **A-B.** Histograms depicting the distribution of two simulated, independent variables: inter-beat interval (left) and reaction time (right). **C.** Circular histogram of simulated reaction times, wrapped to the (simulated) cardiac cycle. Despite the variables being entirely independent and so no true association between cardiac timing and RT, strong non-uniformity is present. This is depicted by the clustering of angles. **D.** The null distribution, obtained by wrapping RTs to shuffled IBIs. The “observed” phase concentration seen in panel C is no different from the null.

outcome, or to test for associations (e.g. Pearson correlation). These tests make Gaussian assumptions and set the null hypothesis at zero; no correlation, no group difference, etc.

The problem with testing for associations between cardiac (or indeed any) timing on behaviour is twofold: First, Normality (of errors) is not a reasonable assumption for circular variables because Gaussian support is over a linear rather than circular space. This issue is straightforward to address: Analogous statistical tests (such as the circular t-test) are available for circular variables if the relevant assumptions are met.



The second problem, which is more complicated to address, is the assumption that the effect under the null hypothesis is zero. This second point is the most crucial, because even if one uses circular analogues of t-tests or ANOVAs there still may be a problem with statistical inference.

Why might the effect under the null take some non-zero value? One such case would be when we test for an association between two periodic variables. Suppose, for example, that a researcher collects ECG while a participant presses a button whenever they feel like making a voluntary action. The researcher wants to know if participants are more likely to press a button at particular points in the cardiac cycle. If an individual participant's heart rate is 60 bpm, then an R wave peak comes every ~1000ms (Fig. 4A). If there is also rhythmicity in the participant's button presses (perhaps they press every 400ms, Fig 4B), then the researcher may well find a correlation between cardiac timing and behaviour, even if the two variables are entirely independent. This is illustrated in Fig. 4C (visualisation by the function `catt_plot_circ`) where it can be seen that there is a concentration of button presses about a particular portion of the cardiac cycle, just after the T wave. This indicates that, despite there being no causal relationship between cardiac time and button press, the null hypothesis is not zero: some clustering (non-uniformity) is expected.

### Using permutation testing for non-uniformity

To control for the above, the CaTT implements permutation testing for tests of non-uniformity (as well as for tests of difference and association), producing surrogate data from which  $H_0$  can be estimated. How is the permutation test for non-uniformity performed? Under the null hypothesis, there is no relationship between cardiac time and behaviour. Therefore, we reason that wrapping behaviour to IBIs from other trials should not reduce non-uniformity if the null hypothesis is true. This is analogous to permutation tests for correlations, where over many iterations the one variable is shuffled while the other is held constant and the correlation recomputed. Here, behaviour (onsets) is wrapped to shuffled IBIs to break the association between the behaviour on some trial and the cardiac cycle it appeared in. A clustering test – either the “parametric” Rayleigh test (which assumes a Von Mises distribution; the circular analogue of Gaussianity) or the “non-parametric” Rao's Spacing test with no distributional assumptions– is performed on each iteration. In both cases the tests are implemented by the `circstat` toolbox (47). The null distribution is shown in Fig. 3D. Here it is made apparent that, despite the strong concentration of angles prior to shuffling, this strong concentration is entirely expected under the null.

Once all permutation iterations are complete, one has a simulated null distribution (48), from which the mean and standard deviation of the null can be calculated. From here, it is simple to express the circular non-uniformity (the test statistic) as a z-score, that can later be combined with the z-scores of other participants as necessary:

$$Z = \frac{U - \bar{U}^*}{SD(U^*)}$$

where  $U$  is the observed test statistic,  $U^*$  is the distribution of test statistics under the null and  $SD$  is standard deviation. The permutation test also gives a p-value, which is simply the proportion of permutations in which the test statistic was greater or equal to  $U$ . For the simulated data in Fig. 3C, the z-score is -0.4931, giving a permutation p-value of .622.

If the user is running analyses on multiple participants (as will often be the case), then a p-value for the group is needed. To obtain this, the toolbox uses Stouffer's method (32) to obtain a pooled z-score. Here, for  $k$  participants, the set of  $k$  individual z-scores is summed and divided by  $\sqrt{k}$ . The result of this function is a z-score, meaning that the two-tailed p-value is  $2(1 - \Phi(|z|))$ , where  $\Phi$  is the Normal cumulative distribution function. This conversion from individual z-scores to a p-value for the group is implemented by the function `catt_z2p`. This approach of estimating participant-level effect sizes and combining them using Stouffer's method has been recommended for cases where  $H_0$  is unknown (49).

### **Using permutation testing to test for consistent non-uniformity across participants**

This above permutation test can tell us whether non-uniformity is present at the group level. However, more than that, one would also want to know that the non-uniformity is such that there is some group-level cardiac time around which angles are clustered. If most participants exhibit non-uniformity but each is skewed towards a different cardiac time, then we cannot infer an effect of cardiac timing on behaviour. There are several reasons why this could be the case, including distributional properties of the onsets or of the inter-beat intervals. For example, if they are highly skewed then certain cardiac phase angles would be less likely to occur than others.

If participants' distributions of cardiac times (expressed in angles) are Von Mises distributed then one could just take the circular mean. However, this is not necessarily the case and an unreasonable assumption to make without examination of the data. Accordingly, we propose

the following test (with no distributional assumptions), implemented by the function `catt_consistency`:

If the user wraps to the R wave peak (systole), then for each participant, the angular data are first binned into  $k$  equally-spaced bins, e.g. 8 bins of  $\pi/4$  ( $45^\circ$ ; the default). One's choice of bin number will depend on how many trials each participant completes, balanced against the required sensitivity of the test. Next, the proportion of data in each bin is calculated.

If the data are uniformly distributed, then for each participant each bin should contain  $1/k$  of the data. To test for consistency across participants, `catt_consistency` runs  $k$  one-sample tests of difference against this expected value of  $1/k$  and corrects for multiple comparisons using the FDR method. Because FDR is more lenient than other corrections (e.g. familywise error) it is important to be cautious and conservative when interpreting results. If indeed, a participant exhibits non-uniformity (as determined by `catt_bootstrap_clust`) and consistency in where that non-uniformity lies (as determined by `catt_consistency`) then the researcher can conclude that some behaviour is more/less likely to occur at particular phases of the cardiac cycle. Effect sizes for each phase bin (in `output.diff_perc`) are given as the percentage difference between the observed and expected proportion of the data in each bin. Furthermore, these preferred (and/or un-preferred) phases can be specified as occurring somewhere between the limits of the bins the researcher chose.

The consistency test requires an adjustment if the user has wrapped onsets to the T wave (diastole): the period between diastole and the subsequent R wave peak is typically longer than the RT interval, and therefore, all else being equal, one would expect a greater proportion of onsets to fall after diastole than before. Accordingly, even in the absence of any consistency one would still expect unequal proportions of onsets falling before versus after diastole. To resolve this issue, onsets falling before diastole (which have negative angles) are binned into four (or  $k/2$ ) equally-spaced bins between  $-\pi$  and  $0$ , and then bin counts converted to proportions by considering only the proportion of negative angles. This process is repeated for positive angles, binning them into four ( $k/2$ ) equally-spaced bins between  $0$  and  $\pi$ . There are therefore two different bin sizes used: one (which will be smaller) for  $-\pi$ rad to  $0$ rad and another (larger) size for  $0$ rad to  $\pi$ rad. Once binning is complete, the consistency test continues as described in the previous paragraph, comparing proportions to those expected under circular uniformity:  $1/k$ .

## Results

## Applying CaTT to real data

To illustrate how CaTT can be used to perform analyses on behavioural and ECG data, we present a re-analysis of a freely available, published dataset (3). In this study, Galvez-Pol and colleagues collected continuous ECG while participants performed a visual search task. They asked whether blinks, saccades or fixations were more likely to occur at particular phases of the cardiac cycle, and answering this question requires performing a test of non-uniformity. However, because tests of uniformity cannot accommodate data where trials are nested in participants, the authors took the circular mean cardiac angle for each participant and tested for non-uniformity of these means. Here, we exploit the potential of permutation testing – the ability to run tests within-participants and combine estimated z-scores into a group-level p-value – applying our proposed approach to Galvez-Pol and colleagues' data. The code for this analysis can be found in the toolbox ([https://github.com/MaxineSherman/CaTT/blob/main/catt\\_demo\\_analysis.m](https://github.com/MaxineSherman/CaTT/blob/main/catt_demo_analysis.m)). If it is not already present, the script will download the authors' data from OSF.

In this open dataset, the available features include a time index, the IBI of each cardiac cycle in which an oculomotor event was made, and the time of the oculomotor event expressed in milliseconds since the last R wave peak.

After initialising the toolbox by calling

```
global catt_opts
catt_init;
```

we set the sampling frequency to 1000Hz, the method of estimating QT intervals to 'fixed' and the default QT time to 400ms. We did this because, without the raw ECG data, QT times could not be estimated:

```
catt_opts.fs = 1000;
catt_opts.qt_method = 'fixed';
catt_opts.qt_default = 400;
```

To run our analysis, for each participant  $i$  and each type of oculomotor event (fixations, saccades and blinks) we created a field called RR (as would be created by the function `catt_epoch`) and loaded in the relevant data: the time since the last R wave peak for each fixation was loaded it into a structure called `group{i}.RR.catt.onset`; dummy time-points and indices were loaded into `group{i}.RR.times` and

`group{i}.RR.catt.idx_RR` respectively, and all responses were left blank by setting `group{i}.RR.catt.response` to `nan`.




Figure 5. Histograms depicting the percentage of (from left) fixations, saccades and blinks removed for each participant. Note that the histograms for fixations and saccades are identical. This is because, in this dataset, fixations always followed saccades and they both fell in the same RR interval.

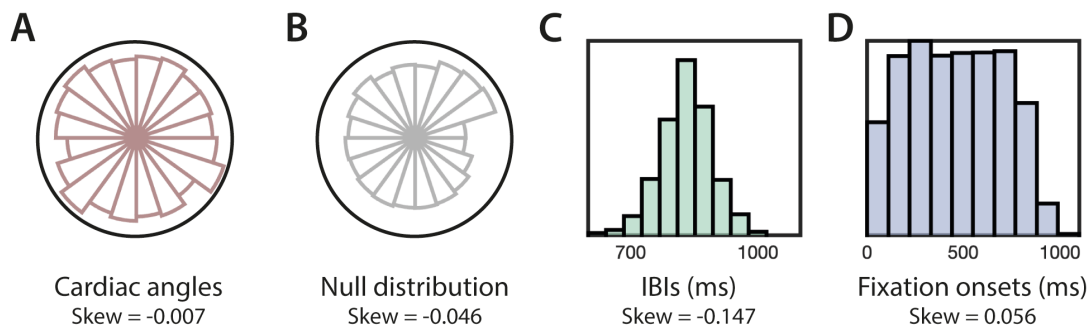
While we could not run manual artefact rejection, we could screen trials with extreme IBIs. This was done by running `catt_IBI` using the default parameters (excluding RR intervals where the IBI was more than 3SD from the mean, or which implied a heart rate  $> 160$ bpm or  $< 40$ bpm). No participant had more than 3.5% of events excluded (Fig. 5).

Then, we ran participant-by-participant non-parametric permutation tests, using the Rao test, and loaded the output back into the `catt` structure. This was done by running, for each participant  $i$ :

```
[~,~,group{i}.catt] = catt_bootstrap_clust(group{i}.catt, 'rao');
```

With this complete, we obtained the group z-value by calling `catt_z2p(catt)`. Individual participant z-values ranged from -2.76 to -7.03, and so the group z-value was -28.62 and the corresponding p-value  $< .001$ . Note that the combined group-level z is substantially larger than individuals' scores because of the cumulative effect of all z-scores having the same sign. The z-scores here were negative, meaning that the data were more, not less, uniform than expected under the null. The reason for this is as follows: participants exhibited a skewed circular distribution of cardiac angles, such that certain angles were *less* likely than others (see Fig. 6A for a representative example). This skew was present under the null as well (Fig. 6B), and to a far greater degree, hence the negative z-score. Upon inspection of the behavioural data, it can be seen that this was due to a skewed distribution of IBIs and

uniform distribution of onsets (Fig 6C-D), leading to certain angles being necessarily less likely a priori. Accordingly, participants were more uniform than expected because at those



*Figure 6.* Data from a representative participant. **A.** Distribution of participant’s fixations, expressed as degrees in the cardiac cycle **B.** The null distribution, constructed by permutation testing. **C-D.** Histograms depicting the participant’s IBIs and onsets respectively. Note that the IBIs are skewed and the onsets broadly uniformly distributed. Skewness values for each distribution are given under the x-axis labels

underrepresented phases, there were *more* fixations occurring, leading to a “rounder” (more uniform) distribution in Fig. 6A than 6B. This issue highlights the importance of using statistical tests with minimal distributional assumptions when looking for cardiac-behaviour relationships.

Having established the presence of deviation from uniformity (here, more uniformity than expected), the second question is whether the location of non-uniformity is consistent across participants (because fixations are more likely at particular phases of the cardiac cycle), or inconsistent (because the non-uniformity is driven by the distribution of onsets and IBIs). To address this question, we ran `catt_consistency(catt)`. As can be seen in Table 2, bin proportions did not significantly deviate from their expected values after correction for multiple comparisons, though a small difference of 1.35% that did not survive correction was found between  $1.75\pi$  and  $2\pi$  rad (just before the R wave peak). Accordingly, we cannot convincingly conclude that fixations are more likely at particular phases of the cardiac cycle: It is possible that the observed clustering was an artefact of the skewed data distributions.

We repeated this analysis for the timing of saccades (see Table 2). Here, results showed that, again, fixation times deviated from uniformity (group  $Z = -33.52$ , group  $p < .001$ ) but now there was also significant group-level consistency in where these deviations occurred: fixations were over-represented by 1.75% at  $0.5\pi$ - $0.75\pi$  ( $p = .01$ ). This cardiac phase is relatively early in the cycle, and broadly consistent with the location originally reported. Fixations were also disproportionately less likely to occur, by 2.38%, just before the R wave peak ( $p < .001$ ). Finally, we performed the analysis for blinks. Again, blink times were more uniform than expected (group  $Z = -8.08$ ,  $p < .001$ ). However, no consistency was found,

meaning that no phase of the cardiac cycle was associated with more or fewer blinks at the group-level.

*Table 2.* Results for the consistency tests on fixation, saccade and blink data, wrapping to the R wave peak. Results are percentage deviation from expected proportions,  $p$  values, and whether the test was significant following FDR-correction

Bin	Fixations			Saccades			Blinks		
	Difference (%)	$p$	Sig. (FDR)	Difference (%)	$p$	Sig. (FDR)	Difference (%)	$p$	Sig. (FDR)
0 (systole) to 0.25 $\pi$	0.17	.797	N	1.32	.075	N	-15.97	.012	N
0.25 $\pi$ to 0.5 $\pi$	1.15	.067	N	0.83	.246	N	-16.92	.019	N
0.5 $\pi$ to 0.75 $\pi$	0.98	.154	N	1.75	.009	Y	4.11	.728	N
0.75 $\pi$ to $\pi$	-0.81	.292	N	-0.85	.184	N	-10.15	.111	N
$\pi$ to 1.25 $\pi$	0.32	.640	N	-0.65	.370	N	-5.94	.395	N
1.25 $\pi$ to 1.5 $\pi$	-0.83	.209	N	-0.25	.681	N	4.86	.522	N
1.5 $\pi$ to 1.75 $\pi$	0.37	.560	N	0.24	.696	N	-0.90	.912	N
1.75 $\pi$ to 2 $\pi$	-1.35	.043	N	-2.38	< .001	Y	-3.82	.665	N

The previous analyses found that saccades are more likely to occur approximately 25%-37.5% of the way in to the cardiac cycle and less likely just before the R wave peak, but this cannot itself tell us where that period falls in systole or diastole. To address this question, the analyses were repeated, timelocking oculomotor events to the T wave instead: our novel method of examining cardiac timing effects.. Oculomotor events were time-locked such that diastole was set at a fixed latency of R + 350. We had to estimate the R wave peak in this way because we did not have the raw ECG data.

Results were broadly consistent with those from the R wave peak analysis (see Table 3). For all oculomotor events types, distributions were more uniform than under the null ( $p < .001$ ). As before, we found no group-level consistency for fixations or blinks. Saccades were 1.67% less likely at diastole (-0.5 $\pi$ -0),  $p = .002$ , with no difference found in any other bin. This analysis found an underrepresentation of saccades, whereas versus when locking to the R wave peak we found an overrepresentation of saccades, however note that the locations in the cardiac cycle that these bins refer to are different.

Next, using CaTT we can ask whether there are differences in the ‘preferred’ cardiac phases for saccades, fixations and blinks. To do this, we ran, for each participant, independent-samples bootstrapped tests of difference using the function `catt_bootstrap_diff`. These permutation tests gave us three z-scores per participant, each pertaining to one of the three contrasts. We used the function `catt_z2p` to merge the per-participant z-scores into group-level z-scores and group-level p-values for each contrast.

Table 3. Results for the consistency tests on fixation, saccade and blink data, wrapping to the T wave. Results are percentage deviation from expected proportions,  $p$  values, and whether the test was significant following FDR-correction

Bin	Fixations			Saccades			Blinks		
	Difference (%)	$p$	Sig. (FDR)	Difference (%)	$p$	Sig. (FDR)	Difference (%)	$p$	Sig. (FDR)
$-\pi$ (systole) to $-0.75\pi$	1.30	.020	N	0.40	.443	N	-4.97	.422	N
$-0.75\pi$ to $-0.5\pi$	-0.95	.123	N	0.33	.535	N	0.66	.917	N
$-0.5\pi$ to $-0.25\pi$	-0.09	.884	N	0.94	.159	N	-2.52	.671	N
$-0.25\pi$ to 0 (diastole)	-0.26	.649	N	-1.67	.002	Y	6.82	.485	N
0 (diastole) to $0.25\pi$	-1.20	.209	N	1.17	.110	N	-2.75	.723	N
$0.25\pi$ to $0.5\pi$	1.08	.301	N	-0.41	.623	N	-3.41	.680	N
$0.5\pi$ to $0.75\pi$	0.36	.681	N	0.20	.882	N	4.57	.679	N
$0.75\pi$ to $\pi$ (systole)	-0.23	.789	N	-0.96	.194	N	1.59	.926	N

Results showed no significant difference at the group-level between mean cardiac time of saccades versus fixations (wrapping to R wave:  $z = -0.82$ ,  $p = .415$ , wrapping to T wave:  $z = -0.42$ ,  $p = .674$ ), saccades versus blinks (wrapping to R:  $z = 0.52$ ,  $p = .601$ , wrapping to T:  $z = 0.83$ ,  $p = .408$ ), or fixations versus blinks (wrapping to R wave:  $z = 1.83$ ,  $p = .068$ , wrapping to T wave:  $z = 1.03$ ,  $p = .303$ ). This means that we found no evidence for any of the oculomotor events being executed earlier or later in the cardiac cycle than any other.

Finally, we can ask whether the “preferred phase” for each type of oculomotor event (fixations, saccades and blinks) are correlated. Here, a significant (positive) correlation would suggest that there is some reliable lag between the two events. One would expect an association between fixations and saccades (because they are related oculomotor events), but not necessarily an association between blinks and fixations or blinks and saccades – we ran this analysis to illustrate how CaTT can be used for correlational analyses.

For each participant and each event type, we calculated the (circular) mean cardiac angle for that event (i.e. the average location of that event in the cardiac cycle), relative to the R wave peak. Then using the function `catt_bootstrap_corr` we ran a circular-circular correlation between each event type. There was an almost perfect correlation between preferred phases for fixations and saccades,  $\rho(30) = .95$ ,  $p < .001$ , consistent with those events occurring very close together in time. There was no significant correlation between saccades and blinks,  $\rho(30) = .20$ ,  $p = .265$ , nor between fixations and blinks,  $\rho(30) = .12$ ,  $p = .471$ .



## Discussion

In this manuscript we have presented the open-source Cardiac Timing Toolbox for MATLAB, which can preprocess raw ECG data, detect R wave peaks and the T wave, perform data quality checks, and run trial- and group-level statistical analyses to test for associative relationships between human behaviour and cardiac timing. We are not aware of comparable toolboxes, and hope that this software can facilitate statistically rigorous work on interoception even with minimal coding experience.

We presented a reanalysis of freely available, previously published data (26), which showed that the toolbox is able to decisively and robustly detect associations between cardiac timing and behaviour, exploiting trial-level data rather than having to average within participants, which is problematic when the onsets are non-Von Mises (an issue which is of course not unique to the dataset we used). Furthermore, CaTT is able to test for relationships between behaviour and the cardiac cycle in ways that respect several important features of these unique datasets; namely, the broadly fixed nature of the Q-T interval; the potential for rhythmicity in both the cardiac and behavioural data, which would lead to artefactual non-uniformity; and the skewed distributions of inter-beat intervals. The statistical functions implemented use permutation testing and correction for multiple comparisons exclusively, alongside minimal distributional assumptions, to maximise statistical rigour. In future versions of the toolbox we hope to expand the set of statistical tests to include more complex group-level models, e.g. circular analogues of ANOVA and Bayesian models.

We hope that these analysis approaches, made straightforward to implement by our code, will enhance and facilitate research on relationships between cardiac activation and behaviour. However, we fully acknowledge that statistical analysis in this area is an ongoing project, and that further development is needed to map better behaviour to the physiological changes occurring over the cardiac cycle. One example here is that when the location of the T wave is inferred from participants' heart rate, results are vulnerable to small changes in how the RT latency is set. Our implementation of T wave detection here overcomes this problem, however this is only possible to use when the researcher has access to the raw ECG data (and the signal to noise ratio is high enough to permit T wave detection). Another avenue for development is extending the toolbox to accommodate pulse-oximetry data. While one can import R wave peaks detected from pulse oximetry data into CaTT, we hope to implement pulse-oximetry R and T detection in CaTT's preprocessing architecture to aid ease-of-use. Finally, we hope to incorporate a tool for performing both a priori power analysis (to determine the number of participants and trials required), and post-hoc power analysis (to estimate the power of a given test).

In CaTT, we have both implemented traditional circular analyses (e.g. the Rayleigh test), and also proposed our own approaches (e.g. for testing consistency without making distributional assumptions). We fully expect that there are superior approaches to those we have selected, and we hope that the pipeline we propose can act as a benchmark for future work.

## References

1. Azzalini D, Rebollo I, Tallon-Baudry C. Visceral signals shape brain dynamics and cognition. *Trends Cogn Sci.* 2019;23(6):488–509.
2. Critchley HD, Garfinkel SN. The influence of physiological signals on cognition. *Curr Opin Behav Sci.* 2018;19:13–8.
3. James W. What is an Emotion? *Mind.* 1884 Jan 26;9(34):188–205.
4. Lange CG. Om Sindsbevaegelser: Et Psyko-fysiologisk Studie. In: Lange CG, James W, editors. *The Emotions.* Baltimore: Williams and Wilkins; 1922.
5. Seth AK, Suzuki K, Critchley HD. An interoceptive predictive coding model of conscious presence. *Front Psychol.* 2011 Jan;2(January):395.
6. Gu X, FitzGerald THB, Friston KJ. Modeling subjective belief states in computational psychiatry: interoceptive inference as a candidate framework. *Psychopharmacology (Berl).* 2019;236(8):2405–12.
7. Petzschner FH, Garfinkel SN, Paulus MP, Koch C, Khalsa SS. Computational models of interoception and body regulation. *Trends Neurosci.* 2021;44(1):63–76.
8. Allen M, Friston KJ. From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese.* 2018;195(6):2459–82.
9. Allen M, Levy A, Parr T, Friston KJ. In the Body's Eye: The Computational Anatomy of Interoceptive Inference [Internet]. 2019 Apr [cited 2021 Dec 20] p. 603928. Available from: <https://www.biorxiv.org/content/10.1101/603928v1>
10. Mancina G, Mark AL. Arterial baroreflexes in humans. In: Shepherd JT, Abboud FM, editors. *Handbook of Physiology Section 2: The cardiovascular system, peripheral circulation and organ blood flow.* Washington: American Physiological Society; 1983.
11. Lacey BC, Lacey JI. Studies of heart rate and other bodily processes in sensorimotor behavior. In: Obrist PA, Black AH, Brener J V., DiCara L, editors. *Cardiovascular psychophysiology: Current issues in response mechanisms, biofeedback and methodology.* 1974. p. 538–64.
12. Al E, Iliopoulos F, Forschack N, Nierhaus T, Grund M, Motyka P, et al. Heart–brain interactions shape somatosensory perception and evoked potentials. *Proc Natl Acad Sci.* 2020 May 12;117(19):10575–84.
13. Gray MA, Minati L, Paoletti G, Critchley HD. Baroreceptor activation attenuates attentional effects on pain-evoked potentials. *PAIN®.* 2010;151(3):853–61.

14. Leganes-Fonteneau M, Buckman JF, Suzuki K, Pawlak A, Bates ME. More than meets the heart: systolic amplification of different emotional faces is task dependent. *Cogn Emot.* 2020;
15. Gray MA, Beacher FD, Minati L, Nagai Y, Kemp AH, Harrison NA, et al. Emotional appraisal is influenced by cardiac afferent information. *Emotion.* 2012;12(1):180.
16. Elbert T, Rockstroh B, Lutzenberger W, Kessler M, Pietrowsky R, Birbaumer N. Baroreceptor stimulation alters pain sensation depending on tonic blood pressure. *Psychophysiology.* 1988;25(1):25–9.
17. Dworkin BR, Elbert T, Rau H, Birbaumer N, Pauli P, Droste C, et al. Central effects of baroreceptor activation in humans: attenuation of skeletal reflexes and pain perception. *Proc Natl Acad Sci.* 1994;91(14):6329–33.
18. Droste C, Kardos A, Brody S, Greenlee MW, Roskamm H, Rau H. Baroreceptor stimulation: pain perception and sensory thresholds. *Biol Psychol.* 1994;37(2):101–13.
19. Wilkinson M, McIntyre D, Edwards L. Electrocutaneous pain thresholds are higher during systole than diastole. *Biol Psychol.* 2013;94(1):71–3.
20. Li X, Swallow K, Chiu M, De Rosa E, Anderson AK. Does the body give the brain an attentional boost? Examining the relationship between attentional and cardiac gating. *Biol Psychol.* 2018;
21. Garfinkel SN, Barrett AB, Minati L, Dolan RJ, Seth AK, Critchley HD. What the heart forgets: Cardiac timing influences memory for words and is modulated by metacognition and interoceptive sensitivity. *Psychophysiology.* 2013;
22. Larra MF, Finke JB, Wascher E, Schächinger H. Disentangling sensorimotor and cognitive cardioafferent effects: A cardiac-cycle-time study on spatial stimulus-response compatibility. *Sci Rep.* 2020;
23. Makowski D, Sperduti M, Blondé P, Nicolas S, Piolino P. The heart of cognitive control: Cardiac phase modulates processing speed and inhibition. *Psychophysiology.* 2020;
24. Rae CL, Botan VE, Gould Van Praag CD, Herman AM, Nyssönen JAK, Watson DR, et al. Response inhibition on the stop signal task improves during cardiac contraction. *Sci Rep.* 2018;
25. Motyka P, Grund M, Forschack N, Al E, Villringer A, Gaebler M. Interactions between cardiac activity and conscious somatosensory perception. *Psychophysiology.* 2019;
26. Galvez-Pol A, McConnell R, Kilner JM. Active sampling in visual search is coupled to the cardiac cycle. *Cognition.* 2020;
27. Kundendorf S, Klotzsche F, Akbal M, Villringer A, Ohl S, Gaebler M. Active information sampling varies across the cardiac cycle. *Psychophysiology.* 2019;
28. Ohl S, Wohltat C, Kliegl R, Pollatos O, Engbert R. Microsaccades are coupled to heartbeat. *J Neurosci.* 2016;
29. Martins Q, McIntyre D, Ring C. Effects of baroreceptor stimulation on performance of the Sternberg task: A cardiac cycle time study of working memory. *Biol Psychol.* 2014;

30. Schandry R. Heart Beat Perception and Emotional Experience. *Psychophysiology*. 1981;
31. Desmedt O, Luminet O, Corneille O. The heartbeat counting task largely involves non-interoceptive processes: Evidence from both the original and an adapted counting task. *Biol Psychol*. 2018;
32. Desmedt O, Van Den Houte M, Walentynowicz M, Dekeyser S, Luminet O, Corneille O. A Systematic Review and Meta-analysis on the Association Between Heartbeat Counting Task Performance and Mental Disorders and Their Risk Factors Among Adults. 2020;
33. Desmedt O, Corneille O, Luminet O, Murphy J, Bird G, Maurage P. Contribution of time estimation and knowledge to heartbeat counting task performance under original and adapted instructions. *Biol Psychol*. 2020;154:107904.
34. Desmedt O, Luminet O, Maurage P, Corneille O. What If the Heartbeat Counting Task Required No Measure of Cardiac Activity? 2020;
35. Harrison OK, Garfinkel SN, Marlow L, Finnegan S, Marino S, Nanz L, et al. The Filter Detection Task for measurement of breathing-related interoception and metacognition. *BioRxiv*. 2020;
36. Larsson D, Esposito G, Critchley Hu, Dienes Z, Garfinkel S. Sensitivity to changes in rate of heartbeats as a measure of interoceptive ability. 2020;
37. Legrand N, Nikolova N, Correa C, Brændholt M, Stuckert A, Kildahl N, et al. The heart rate discrimination task: a psychophysical method to estimate the accuracy and precision of interoceptive beliefs. *Biol Psychol*. 2021 Dec 11;108239.
38. Manikandan MS, Soman KP. A novel method for detecting R-peaks in electrocardiogram (ECG) signal. *Biomed Signal Process Control*. 2012;
39. Vázquez-Seisdedos CR, Neto JE, Marañón Reyes EJ, Klautau A, Limão de Oliveira RC. New approach for T-wave end detection on electrocardiogram: Performance in noisy conditions. *Biomed Eng OnLine*. 2011 Sep 9;10(1):77.
40. Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology. Heart Rate Variability. *Circulation*. 1996 Mar 1;93(5):1043–65.
41. BAZETT HC. AN ANALYSIS OF THE TIME-RELATIONS OF ELECTROCARDIOGRAMS. *Ann Noninvasive Electrocardiol*. 1997;
42. Fridericia LS. The duration of systole in an electrocardiogram in normal humans and in patients with heart disease. *Annals of Noninvasive Electrocardiology*. 2003.
43. Sagie A, Larson MG, Goldberg RJ, Bengtson JR, Levy D. An improved method for adjusting the QT interval for heart rate (the Framingham Heart Study). *Am J Cardiol*. 1992;
44. Berens P. *CircStat: A MATLAB toolbox for circular statistics*. *J Stat Softw*. 2009;31(10).
45. Fisher NI. *Statistical analysis of circular data*. Cambridge university press; 1995.

46. Batschelet E. Circular statistics in biology. Acad PRESS 111 FIFTH AVE N Y NY 10003 1981 388. 1981;
47. Berens P. CircStat : A MATLAB Toolbox for Circular Statistics . J Stat Softw. 2009;
48. Theiler J, Prichard D. Using 'surrogate surrogate data'to calibrate the actual rate of false positives in tests for nonlinearity in time series. Fields Inst Comm. 1997;11:99.
49. Dowding I, Haufe S. Powerful statistical inference for nested data using sufficient summary statistics. Front Hum Neurosci. 2018;12:103.