Biomedical Signal Processing and Control 7 (2012) 333–341

Contents lists available at ScienceDirect

Biomedical Signal Processing and Control

journal homepage: www.elsevier.com/locate/bspc

Real-time detection of workload changes using heart rate variability☆

Adam Hoover[a,*], Anirud Singh[a], Stephanie Fishel-Brown[b], Eric Muth[b]

[a] Electrical and Computer Engineering Department, Clemson University, Clemson, SC 29634-0915, United States
[b] Department of Psychology, Clemson University, Clemson, SC 29634-0915, United States

ARTICLE INFO

Article history:
Received 7 April 2011
Received in revised form 15 June 2011
Accepted 11 July 2011
Available online 15 August 2011

Keywords:
ECG
Change point detection
Heart rate variability
Workload
State change
CUSUM

ABSTRACT

This work presents a novel approach to detecting real-time changes in workload using heart rate variability (HRV). We propose that for a given workload state, the values of HRV vary in a sub-range of a Gaussian distribution. We describe methods to monitor a HRV signal in real-time for change points based upon sub-Gaussian fitting. We tested our method on subjects sitting at a computer performing a low workload surveillance task and a high workload video game task. The proposed algorithm showed superior performance compared to the classic CUSUM method for detecting task changes.

1. Introduction

The purpose of this study is to determine if heart rate variability (HRV) can be used to detect mental workload changes in real-time. We define mental workload as the task demand placed upon a sedentary person. Higher levels of workload have been shown to result in physiological and cognitive changes including decreased attention and concentration [9], increased muscle tension [5], and coordination difficulties [41]. Those changes often negatively impact performance [3,11,17,21].

The autonomic nervous system (ANS) plays a crucial role in regulating physiological changes in response to workload [8,21,53,54]. The parasympathetic nervous system (PNS), the branch of the ANS responsible for relaxing the body, decreases in activity in response to increases in workload. Thus, a variety of physiological indicators of arousal respond to changes in workload [20,37].

In the present study, we consider the effect of changing mental workload upon HRV. Researchers have demonstrated that increased task complexity and attention results in decreased HRV [1,33,38,49]. The high frequency component of HRV between 0.15 and 0.5 Hz, also known as respiratory sinus arrhythmia (RSA), has been widely established as an indirect index of the PNS [10,18,19,24]. We successfully demonstrated that RSA can be used off-line to measure changes in mental workload between a high and low workload task [12,13].

In recent years, however, there has been an increased interest in adaptive automation [35]. Adaptive automation requires the real-time measurement and analysis of physiological signals. Specifically, adaptive systems can modify a task or set of tasks in real-time to help optimize performance or mediate negative mental effects. For example, Katsis et al. [25] studied the use of heartrate, galvanic skin response and respiration for classifying the anxiety state of a subject during therapeutic sessions. Cannon et al. [4] studied the use of EEG and ECG signals for monitoring the cognitive load of a subject. Potential applications of these types of works include aviation [14,36,55], driving [27], and other tasks requiring vigilance [15]. Real-time measurement and analysis of HRV data may provide an unobtrusive and continuous means for adaptive systems to prevent detrimental effects to performance from high levels of workload [3,11,17,21].

1.1. Change point detection

The goal of the current work is to detect change points in HRV in real-time. Fig. 1 illustrates the problem of change point detection. A signal is being monitored over time, and at some point changes from one state (e.g., reasonable workload) to a second state (e.g., detrimental workload). The goal is to detect the point in time at which the change occurs. Several issues make the problem difficult, including noise in the measurements, natural variations in the signal, and ambiguity about the definition of "state". The latter issue can be approached through the use of statistics to model

☆ This work was supported by an Office of Naval Research grant.
* Corresponding author. Tel.: +1 864 656 3377.
E-mail address: ahoover@clemson.edu (A. Hoover).

1746-8094/$ – see front matter
doi:10.1016/j.bspc.2011.07.004

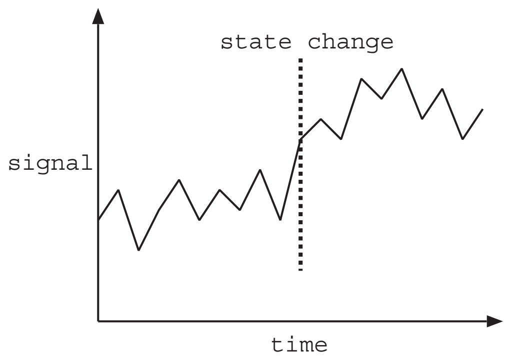

Fig. 1. A change point for a signal.

the properties of the signal. This approach can be traced to work in industrial control in the 1920s [44], where the concern was to monitor manufacturing processes and detect deteriorating conditions or failures. Today, the same approach can be found in a number of fields, including economics [26], weather monitoring [42], biology and genetics [7], signal processing and voice monitoring [46], and computer network security [51]. Reviews of related work in the most recent decades can be found in [2,16,43].

The problem differs depending upon whether or not the changes are being detected off-line versus in real-time. In the off-line case, an entire data set is available for analysis, and classic data grouping and segmentation methods can be applied. In the real-time case, we desire to detect changes as soon as possible while the signal is being monitored. Fig. 2 illustrates the additional facets of the problem. It is expected that some amount of time must elapse after the change has occurred before it can be detected. The delay is the time between the change and its detection. Statistics concerning this delay must be considered in addition to the true positive and false alarm rates, collectively indicating how well the change point detector is performing [2,16]. Real-time change point detection is prevalent in biomedical problems, such as monitoring intensive care patients [6], monitoring anesthesia [56], and monitoring pregnancy contractions [28].

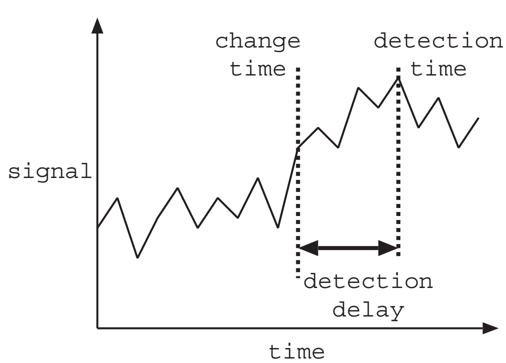

Fig. 2. Detecting a change in real-time.

Most methods for change point detection model the signal in terms of its statistics, and look for changes in those statistics. Traditional solutions to the problem have taken one of two approaches: (a) segment the data at the change point into two distributions, each having its own statistics, or (b) detect a substantial change in the distribution statistics as new samples are obtained. Perhaps the most popular example of the latter approach is the cumulative sum (CUSUM) method [34] (see [50] for recent work), monitoring for a change in the mean.

Other approaches have been proposed. Due to their nature, physiological signals may in some cases be modeled as a slowly changing signal plus a noise component [39]; separating these two components allows for simpler change detection. Outliers in the sampled data can be considered a separate problem or incorporated into the change detection problem [48] (we do not consider outliers in this work). In the case of a slowly changing signal, polynomials can be fit to the data, using deviations from the polynomials to detect changes [32]. Wavelets can be used to detect changes in the frequency space of a signal [30].

In this paper, we propose a new approach to the problem of HRV change detection. Our approach was motivated by casual observation of our data. We observed that HRV data during different tasks often accumulated into distributions that looked "sub-Gaussian" [45]. We have developed a method to solve for sub-Gaussian ranges using least-squares fitting, and for monitoring for change points using an overlap statistic. This differs from all previous works in that we do not monitor for changes in the mean and standard deviation of the distribution of the signal. Instead, we consider the case where samples are primarily drawn from one portion of the distribution, and then in a subsequent interval from another portion of the same distribution. This type of change may occur in many biomedical and physiological signals, where it is not the underlying distribution that changes, but rather the range of operation of the system.

In Section 2, we describe a method to monitor a HRV signal in real-time for change points based upon sub-Gaussian fitting. In Section 3, we present experiments on HRV where we use our methods to detect state changes. We compare our approach to the CUSUM method and show that our method provides superior results for this data. While HRV change detection is our primary focus, we believe our methods could be applied to many other problems involving the monitoring of other biomedical and physiological signals.

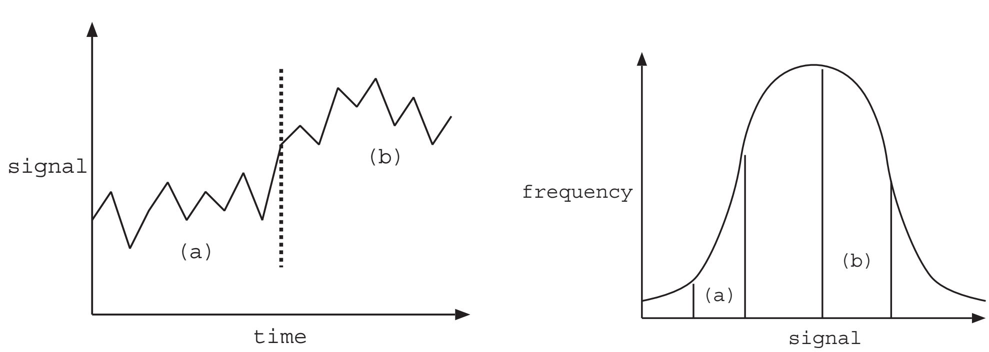

Fig. 3. The state of a signal can be defined as a subrange in its Gaussian distribution.

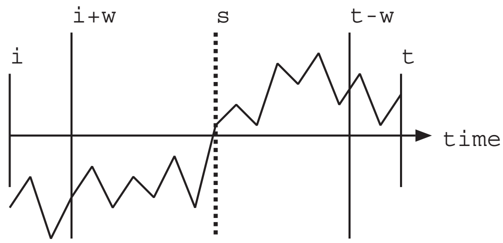

Fig. 4. The indices and variables used in our method.

2. Methods

Fig. 3 illustrates our approach. We assume that during interval (a), data samples are primarily drawn from one subrange (the left side of the distribution in this example); during interval (b), samples are primarily drawn from a different subrange.

Our change point detector requires all the data obtained since the last detected change point to be buffered. Each time a new data point is acquired and added to the buffer, the following steps are taken. A range of possible split points in the buffer is analyzed. Each possible split point breaks the data into two sequences, one prior to the split point and one after the split point. A sub-Gaussian function is fit to the frequency distribution of both sequences. The possible split point with the best sub-Gaussian fits (smallest residuals) is identified as the best split. These two sub-Gaussian fits are then analyzed for overlap. A change point occurs when the sub-Gaussian fits of the two sequences are substantially separate, as determined by an overlap statistic. Once a change point is detected, the buffer is flushed of all data prior to the change point, and the process repeats.

Formally, let a signal be represented as $x_i, x_{i+1}, \ldots, x_t$ where $i-1$ is the index of the last detected change point and t is the index of the most recently acquired data point. Fig. 4 illustrates the timeline of the process. We search for the index s such that the range of data $x_i, \ldots, x_s$ and the range of data $x_{s+1}, \ldots, x_t$ yield the best fitting sub-Gaussian distributions. The variable w is the minimum amount of data required on either side of the change point. Therefore, at least $2w$ data needs to be buffered during the operation of our algorithm. An index j (not pictured in Fig. 4) is varied from $i+w$ to $t-w$ to search for the ideal split point s in the data sequence.

The steps to implement our method may be stated as follows:

1. Let the input signal be represented by $x_i, x_{i+1}, \ldots, x_t$ where each x_i is a discrete measurement, the current time is represented by t, and x_{i-1} is the last detected change point.
2. Let a proposed change point time j iterate from $i+w$ to $t-w$ where w is the minimum window size (see Fig. 4). The minimum window size refers to the minimum amount of data that can be aggregated into a state.
3. Compute the frequency distribution $f_{i,j}$ for data prior to the proposed split point j and the frequency distribution $f_{j+1,t}$ for data subsequent to the proposed split point j.
4. Compute $E(f_{i,j})$ and $E(f_{j+1,t})$ as the residuals of the sub-Gaussian fits to the two frequency distributions.
5. The index j that provides the smallest residual $E(f_{i,j}) + E(f_{j_1,t})$ is chosen as the best split point.
6. An overlap statistic $O([a_1, b_1], [a_2, b_2])$ is computed from the sub-Gaussian fits of the two window sequences using the parameters a_1, b_1, a_2, b_2 calculated during the sub-Gaussian fitting.
7. Compare the calculated overlap statistic O to a threshold P (discussed further in Section 2.3) to decide if a state change has occurred.
8. If a state change is detected, then output the detected change point, flush the buffer of all data prior to j, and let $i = j + 1$.
9. Read the new data point, increment t and continue from step 1.

A frequency distribution is computed by tabulating the occurrences of each possible signal value x_n. Formally, the frequency distribution $f_{i,j}$ can be defined as

$$\forall n, \quad freq(x_n) = \int_i^j (x = x_n) \tag{1}$$

In practice, the frequency distribution $f(x_n)$ is calculated as a histogram with appropriate bin sizes. For example, assuming the signal is normalized with a mean equal to zero and a standard deviation equal to one, then the frequency distribution could be calculated using a histogram with bins ranging from −4 to +4 where each bin size is 0.2, providing a 40 point function for sub-Gaussian fitting.

The following subsections describe our methods for sub-Gaussian fitting and computing an overlap statistic.

2.1. Gaussian fitting

The function for a Gaussian distribution may be written as

$$y = ke^{-(x-\mu)^2/2\sigma^2} \tag{2}$$

where μ and σ are the mean and standard deviation of the distribution, and k is a scaling constant. Fitting a Gaussian curve to a set of data generally involves solving for these three parameters. However, we assume that the mean and standard deviation are known a priori, and so are only interested in solving for k. Given a frequency distribution $f(x_n)$, an error function for fitting can be defined as

$$e_n = f(x_n) - ke^{-(x_n-\mu)^2/2\sigma^2} \tag{3}$$

where in practice n is discretized in bin intervals as previously described. The total residual error can be defined as:

$$E = \sum(e_n^2) = \sum_n (f(x_n) - ke^{-(x_n-\mu)^2/2\sigma^2})^2 \tag{4}$$

To solve for the k that provides the best fit, we take the partial derivative of E with respect to k and set this function equal to zero:

$$\frac{\partial E}{\partial k} = 0 \tag{5}$$

Thus we get:

$$2\sum_n (f(x_n) - ke^{-(x_n-\mu)^2/2\sigma^2})(-e^{-(x_n-\mu)^2/2\sigma^2}) = 0 \tag{6}$$

Solving for k gives:

$$k = \frac{\sum_n (f(x_n)e^{-(x_n-\mu)^2/2\sigma^2})}{\sum_n (e^{-(x_n-\mu)^2/2\sigma^2})^2} \tag{7}$$

Fig. 5 shows an example of a Gaussian function fit to a discrete frequency distribution with $\mu = 0$ and $\sigma = 1$.

2.2. Sub-Gaussian fitting

We define a function for a sub-Gaussian distribution as follows:

$$y = \begin{cases} 0 & x < a \\ ke^{\dfrac{-(x-\mu)^2}{2\sigma^2}} & a \le x \le b \\ 0 & x > b \end{cases} \tag{8}$$

The function follows a regular Gaussian distribution within the range a to b, and is zero otherwise. As before, we assume that the

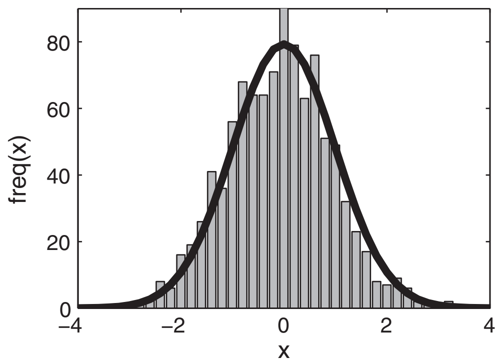

Fig. 5. A Gaussian function fit to a set of data.

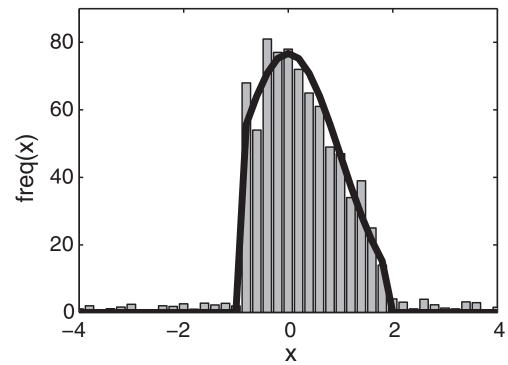

Fig. 6. A sub-Gaussian function fit to a set of data.

parameters μ and σ are known a priori, and so we are only interested in solving for the parameters k, a and b. Given a frequency distribution $f(x_n)$, an error function for fitting can be defined as

$$e_n = \begin{cases} f(x_n) & x < a \\ f(x_n) - ke^{\dfrac{-(x_n-\mu)^2}{2\sigma^2}} & a \le x \le b \\ f(x_n) & x > b \end{cases} \tag{9}$$

The total residual error can be defined as:

$$E = \sum(e_n^2) = \sum_{n<a} f(x_n)^2 + \sum_{n=a}^{b} (f(x_n) - ke^{-(x_n-\mu)^2/2\sigma^2})^2 + \sum_{n>b} f(x_n)^2 \tag{10}$$

Given specific values for a and b, the parameter k can be found using Eq. (5). For the partial derivative of E with respect to k, the parts of the function for $x<a$ and $x>b$ are constants. Thus,

$$\frac{\partial E}{\partial k} = 2\sum_{n=a}^{b} (f(x_n) - ke^{-(x_n-\mu)^2/2\sigma^2})(-e^{-(x_n-\mu)^2/2\sigma^2}) \tag{11}$$

Setting Eq. (11) equal to zero and solving for k gives:

$$k = \frac{\displaystyle\sum_{n=a}^{b} (f(x_n) e^{-(x_n-\mu)^2/2\sigma^2})}{\displaystyle\sum_{n=a}^{b} (e^{-(x_n-\mu)^2/2\sigma^2})^2} \tag{12}$$

The parameters a and b are obtained through a brute force approach. For each possible combination of a and b, the parameter k is found using Eq. (12). The total residual E for each set of values (a, b, k) is found using Eq. (10). Each E is scaled according to its respective k as E/k, in order to compare across sets. The set that produces the minimum E/k is selected as the optimal fit. In practice, the frequency space $f(x_n)$ is discretized, so that the total number of sets (a, b) evaluated is finite. The search space for a and b can be further reduced by requiring a minimum range $b-a$, for example equal to 1σ, a minimum value for a, for example -4σ, and a maximum value for b, for example 4σ. In this case the search range for (a, b) is $([-4\sigma \ldots 2\sigma], [a+\sigma \ldots 4\sigma])$.

Fig. 6 shows an example of a sub-Gaussian function fit to a discrete frequency distribution with $\mu=0$, $\sigma=1$, $a=-0.8$ and $b=2.2$.

2.3. Overlap statistic

Given two sequences of data with sub-Gaussian fits, we seek to determine the amount of overlap of the sub-Gaussians. We do this by determining the amount of overlap of the range $[a_1, b_1]$ of the first sub-Gaussian to the range $[a_2, b_2]$ of the second sub-Gaussian. Because these ranges may or may not overlap, we first compute the intersection $[a', b']$ of the two ranges as follows:

$$a\prime = \begin{cases} a_1 & a_2 \le a_1 \text{ and } a_1 \le b_2 \\ a_2 & a_1 < a_2 \text{ and } a_2 \le b_1 \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

$$b\prime = \begin{cases} b_1 & a_2 \le b_1 \text{ and } b_1 \le b_2 \\ b_2 & a_1 \le b_2 \text{ and } b_2 < b_1 \\ 0 & \text{otherwise} \end{cases}$$

Using these values, the overlap with respect to $[a_1, b_1]$ is

$$P_1 = \frac{b\prime - a\prime}{b_1 - a_1} \tag{14}$$

The overlap with respect to $[a_2, b_2]$ is

$$P_2 = \frac{b\prime - a\prime}{b_2 - a_2} \tag{15}$$

The total overlap is computed as the average of the two:

$$Overlap = \frac{P_1 + P_2}{2} \tag{16}$$

The range of values that the overlap statistic can take is [0, 1].

2.4. Pseudocode

Our algorithm can be implemented by the following pseudo code:

Read data point x_t
while (x_t)
  Update buffer: $x_i, x_{i+1}, x_{i+2}, \ldots, x_t$
  for $j = i + w$ to $t - w$
    $E(j) = subGaussianFits(i, j, t)$
  $min(E(j)) \to s, [a_1, b_1], [a_2, b_2]$
  $O = Overlap(a_1, b_1, a_2, b_2)$
  if $O < P$ then state change detected at s, $i = s + 1$

3. Experimental results

We tested our methods on data obtained from the monitoring of HRV. We first describe our measure of HRV. We then describe our data set and the experiments performed on this data.

3.1. Measure of HRV

The electrical activity of the heart, as measured by the electrocardiogram (ECG), can be used to construct an event series that indicates the time between individual heartbeats. HRV describes cyclical variations in an inter-beat interval (IBI) series related to autonomic nervous system activity [47]. Roughly speaking, this analysis can provide a measure of the restfulness of the participant. Our measure of HRV is described in detail in [23,40]. Briefly, a 64 second window of the most recent IBI values is analyzed using a fast Fourier transform. The magnitude of the power in the frequency range 9–30 cycles per minute is found. This magnitude is log-normalized and taken as our measure of HRV. In general, larger values indicate the subject is more rested, while smaller values indicate the subject is more aroused. The window of IBI data is continuously updated as new heartbeats are detected.

3.2. Tasks and workload changes

We created a custom program that engages the participant in two different tasks, a "shooting task" and a "surveillance task". The shooting task was intended to increase mental workload, causing a decrease in our HRV measure. The surveillance task was intended to decrease mental workload, causing an increase in our HRV measure. Fig. 7 shows a screenshot of the shooting task. The participant controls the spaceship at the bottom of the screen in action that resembles a video game. Waves of enemy space ships appear in sequence, containing anywhere from 5 to 10 ships. At random times and at random places, ammunition and health icons appear at the top of the screen and fall towards the bottom. The goals of the shooting task are to shoot the enemy space ships using the left mouse button, avoid getting hit with enemy fire, and replenish ammunition and life by moving over the randomly falling ammunition and life icons.

The surveillance task used the same general graphics. The main difference was that all firing was disabled; the enemies did not fire at the participant, and the participant could not fire at the enemies. Deaths could not occur, and the amount of health and ammunition of the participant could not be changed. Instead, at random times, a special red-colored enemy ship would appear among the normally grey-colored enemy ships. At such times, the participant was instructed to press the spacebar.

Our program was designed to randomly switch between these two tasks at paired time intervals of 30 s, 1 min, 2 min, 4 min, and 8 min. Participant performance on the combined tasks was used to calculate monetary compensation, on the order of \$4–\$70 (US). Simultaneous to both tasks, the program also posed mental arithmetic questions as a secondary task [13], to simulate realistic multi-tasking conditions.

3.3. Data

Forty-five participants participated in the dual-task paradigm. Each trial lasted 31 min ($2 \times$ each interval of 30 s, 1 min, 2 min, 4 min and 8 min). The order of the sequence of task pairs for each participant was determined using a Latin square counterbalancing method. The NASA-TLX (Task Load Index) [22], a subjective workload questionnaire, confirmed that the tasks had the desired effects of evoking more or less mental workload [12].

In off-line analysis, it was found that the minimum amount of time per task for which a difference in the HRV measure could be found was 2 min [12]. This reflects the nature of the measure as a minute-to-minute indicator of the restfulness of the participant. For this experiment, we considered only those portions of the data for each participant consisting of the 4 min and 8 min task pairs; the HRV data for each of these was manually segmented out of the 31 min recording, giving us 90 total test segments. Fig. 8 presents some examples of these segments. For both these plots, and for every recording in our data set, the task was changed exactly in the middle of the recording. Thus, in this experiment we define our goal as the automatic detection of a change point occurring at the middle of each recorded segment. Due to the nature of our HRV measure, and some small variability in the start and switch times for tasks, we define a correct detection as any within ± 30 s of the task change. Any other change points detected are considered false positives.

These examples demonstrate the difficulty of our change point detection problem. In the left plot, a relative change in the signal is visible to the naked eye. The data in the first half of the recording appears on average higher than the data in the second half of the recording. However, simple thresholding techniques would fail badly at identifying the change because of the oscillatory nature of the signal. In the right plot, there is very little discernible difference between the signal in each half of the recording. In this case, most of the visually apparent difference is due to the increased oscillatory behavior in the latter half of the recording.

3.4. Results

Fig. 9 demonstrates the key point of our methods. It shows sub-Gaussian fits to the data for each task for the recordings shown in Fig. 8. The top row shows the fits for the left recording, the bottom row shows the fits for the right recording. In both cases, the split point was fixed to the exact middle of the recording, so that the result of an ideal detection could be considered. The plot on the left shows the sub-Gaussian fit to the data recorded during the first task, the plot in the middle shows the sub-Gaussian fit to the data recorded during the second task, and the plot on the right shows the overlap of the two fits.

In the case of the upper row, the raw data overlap on either side of the change point. However, the strongest fitting sub-Gaussians to each side do not overlap at all, indicating a strong likelihood of a state change. For the lower row, there is an even larger amount of overlap, so there is less likelihood of a state change. Even so, the sub-Gaussian fits show a visually apparent difference, especially when compared to the minimal visual difference in the raw signal.

We coded our methods in C, compiled using the gnu gcc compiler, and executed on a standard desktop computer running the Ubuntu linux operating system. For all sub-Gaussian fits, we computed the frequency space using a 40 point histogram and searched the $[a, b]$ space as described in Section 2. The parameter w was set to 2 min, reflecting the minimum amount of data to aggregate into a state. The overlap statistic threshold was chosen as 0.25. Processing each recording took only a few seconds; our methods could easily run in real-time.

Fig. 10 shows some examples where our method found a change point corresponding to the task change, with no false positives. The dashed line indicates the change point (s in our algorithm), the dotted line indicates when the change was found (t in our algorithm). The example on the left demonstrates that our approach can detect a change that is quite subtle, and occurs somewhat gradually. The example on the right shows a signal with a visually more apparent change point. However, the signal also shows oscillatory behavior that would cause problems for any simple thresholding approach. Our method successfully identified the change point without triggering false positives during the oscillations.

Fig. 11 shows some examples where our method detected a change point at the task change, but also detected a change point at another time. Visually, the false detections do not appear to truly be false; they occur at places where the signal appears to actually be changing. This is discussed more below. Fig. 12 shows some examples where our method did not find any change points at all.

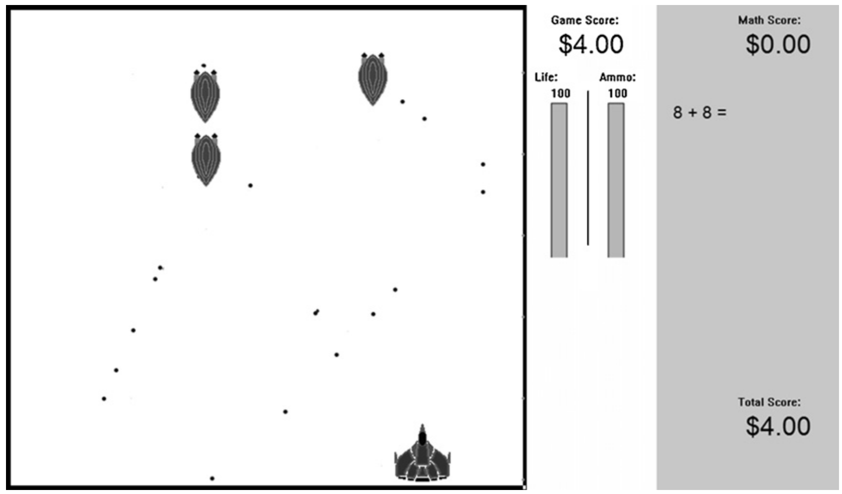

Fig. 7. Screenshot of the higher workload "shooting task".

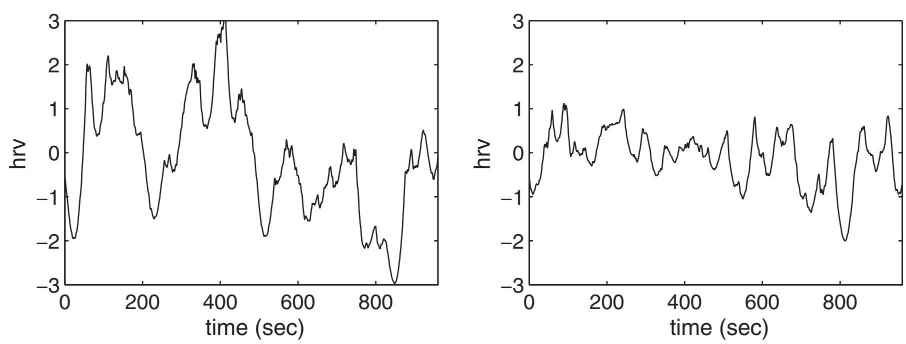

Fig. 8. Two examples from our data set.

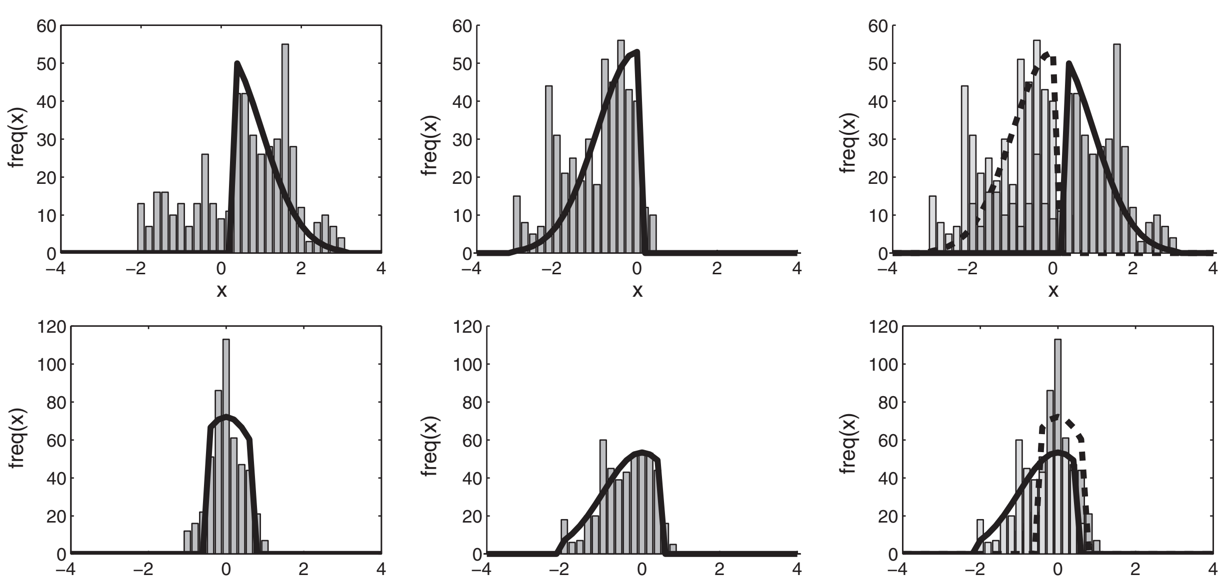

Fig. 9. Sub-Gaussian fits for the two recordings shown in Fig. 8. The plot on the left is for the first half of the data, the plot in the middle is for the second half of the data, and the plot on the right shows the overlap.

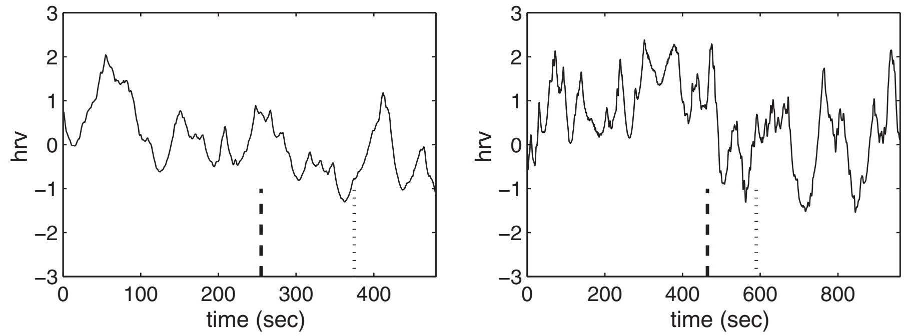

Fig. 10. Two examples where a change point was detected at the task change, with no false positives (dashed line = change point; dotted line = time of detection).

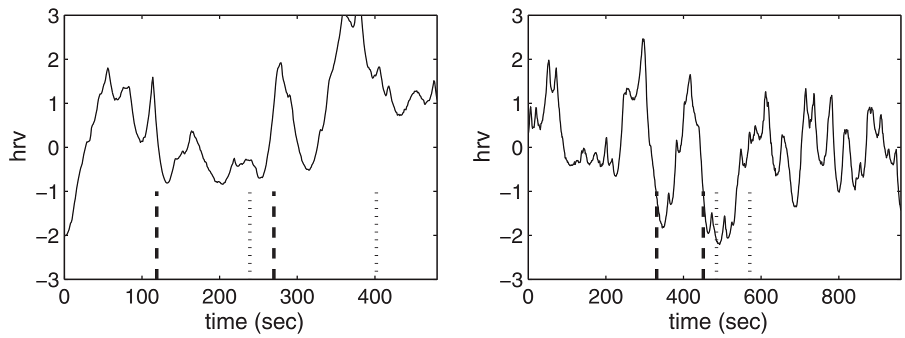

Fig. 11. Two examples where a superfluous change point was detected in addition to the change point at the task change (dashed line = change point; dotted line = detection).

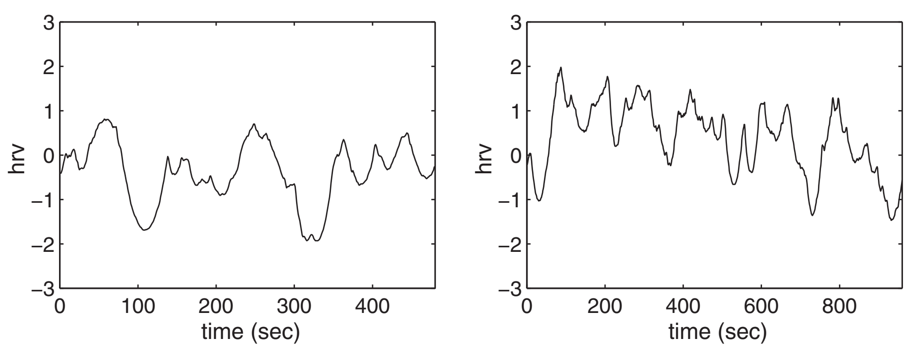

Fig. 12. Two examples where no change points were detected.

In order to evaluate our approach, we compared it against the classic CUSUM method [2,34]. The CUSUM statistic was calculated on the running sample mean.[1] We ran our entire data set of 90 recordings several times, varying the threshold for change point detection from 0.005 to 0.02. For our sub-Gaussian method, we similarly ran our entire data several times, varying the overlap threshold from 0.05 to 0.5. Fig. 13 shows the ROC plots comparing the performance of both methods, in terms of total true positives versus total false positives, across the 90 recordings. Our sub-Gaussian method clearly performed better than the CUSUM method.

As can be seen in these examples, the type of change we are seeking to detect is subtle, if it is there at all. The plot on the right of Fig. 12 shows a slight downward trend in value over time, while the plot on the left does not appear to show any change. This

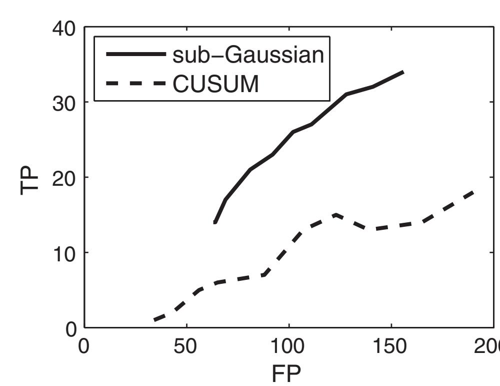

Fig. 13. Comparison of our sub-Gaussian method to the classic CUSUM method on our data. The units are total true positives versus total false positives across 90 recordings, varying the detection threshold for each method.

[1] Several other statistics were tested, such as the actual mean (computed a priori); the reported results are for the statistic that showed the best results in the CUSUM method.

can be explained by considering the nature of the scenario under which our data was gathered. The original intent was to determine if a change in task *caused* a change in the HRV measure. In this experiment we are addressing the opposite question; namely, can we automatically detect a task change by monitoring HRV. Our goal is hampered by the fact that the task may not have caused the participant to change his or her HRV, or that changes in the HRV measure may have occurred at times other than when the task changed. For example, a participant may rouse him or herself during the surveillance task, at times causing a fluctuation in our measure that is independent of the task change. It has been demonstrated previously that vigilance tasks can raise workload, for example depending on the stressfulness of the task [31,52]. We conducted a two-tail independent samples *t*-test on each of our 90 test segments, comparing the data prior to the task change to the data after the task change. For 27 of the segments, the t statistic ($\alpha = 0.05$) showed either no difference or a reverse (from the expected) change in workload. Thus, our results must be viewed in the context of what our change point detector is being challenged to detect. We believe that the type of result shown in Fig. 10 demonstrates the potential of our approach.

4. Conclusions

In this paper, we presented a novel approach to detect a change in mental workload based upon the real-time monitoring of HRV. Our methods are based upon fitting a sub-Gaussian function to the sequences of data preceding and succeeding a suspected change point. An overlap statistic evaluates the overlap of the two fits, which can be thresholded to determine whether or not a change has occurred. We described an algorithm to implement our method, and demonstrated its use on detecting task changes between a computer shooting game and a computer vigilance game for 45 participants. For this data, our method showed better performance compared to the classic CUSUM method for detecting task changes.

As can be seen in the examples from Section 3, we feel that our methods can successfully detect changes that are quite subtle. In future work we would like to apply our methods to other types of data, particularly where the ground truth could be more easily defined. We believe our methods could be applied to many other problems, such as those involving the monitoring of biomedical and physiological signals. In these cases and perhaps others, it is viable to consider a system where it is the range of operation that changes rather than the distribution statistics.

Concerning the detection of mental workload from heart monitoring, it may be that other measures besides HRV (or besides our particular measure of HRV) may provide a better indicator for detecting state change. The method described in this paper could be applied to any measure, and allows for the comparison of different measures, which is an interesting direction for future work. In addition, while we assumed our measure follows a Gaussian distribution over the long term, it may be that for other measures or other data sets another distribution is more appropriate. Frisen [16] discusses this issue in the context of the CUSUM method. In future work it would be interesting to adapt our method to alternative distributions.

References

[1] M. Althaus, L. Mulder, G. Mulder, A. van Roon, R. Minderaa, Influence of respiratory activity on the cardiac response pattern to mental effort, Psychophysiology 35 (1998) 420–430.
[2] M. Basseville, I. Nikiforov, in: Detection of Abrupt Changes: Theory and Application, Prentice-Hall Publishing, New York, 1993.
[3] B. Beith, Subjective workload under individual and team performance conditions, in: Proceedings of the Human Factors Society 31st Annual Meeting, (1987) pp. 67–71.
[4] J. Cannon, P. Krokhmal, R. Lenth, R. Murphey, An algorithm for online detection of temporal changes in operator cognitive state using real-time psychophysiological data, Biomedical Signal Processing and Control 5 (2010) 229–236.
[5] S. Cathcart, D. Pritchard, Relationships between arousal-related moods and episodic tension-type headaches: a biophysiological study, Headache: Journal of Head and Face Pain 38 (3) (2004) 214–221.
[6] S. Charbonnier, G. Becq, L. Biot, On-line segmentation algorithm for continuously monitored data in intensive care units, IEEE Transactions on Biomedical Engineering 51 (3) (2004) 484–491.
[7] J. Chen, Y.-P. Wang, A statistical change point model approach for the detection of DNA copy number variations in array CGH data, IEEE/ACM Transactions on Computational Biology and Bioinformatics 6 (4) (2009) 529–541.
[8] P. Comens, D. Reed, M. Mette, Physiologic responses of pilots flying high performance aircraft, Aviation, Space, and Environmental Medicine 58 (1987) 205–210.
[9] S. Dickman, Dimensions of arousal: wakefulness and vigor, Human Factors 44 (3) (2002) 429–442.
[10] D. Eckberg, Human sinus arrhythmia as an index of vagal cardiac outflow, Journal of Applied Physiology 54 (1983) 961–966.
[11] F. Eggemeier, Properties of workload assessment techniques, in: P.A. Hancock, M. Meshkati (Eds.), Human Mental Workload, 1988, pp. 41–62.
[12] S. Fishel, E. Muth, A. Hoover, L. Gugerty, Determining the resolution of a real-time arousal gauge, in: In the Proceedings of SPIE, vol. 6218, Chemical and Biological Sensing VII, 2006.
[13] S. Fishel, E. Muth, A. Hoover, Establishing appropriate physiological baseline procedures for real-time physiological measurement, Journal of Cognitive Engineering and Decision Making 1 (3) (2007) 286–308.
[14] F. Freeman, P. Mikulka, L. Prinzel, M. Scerbo, Evaluation of an adaptive automation system using three EEG indices with a visual tracking system, Biological Psychology 50 (1999) 61–76.
[15] F. Freeman, P. Mikulka, M. Scerbo, L. Scott, An evaluation of an adaptive automation system using a cognitive vigilance task, Biological Psychology 67 (2004) 283–297.
[16] M. Frisen, Properties and use of the Shewhart method and its followers, Sequential Analysis 26 (2007) 171–193.
[17] D. Gopher, E. Donchin, in: K.R. Boff, L. Kaufman, J.P. Thomas (Eds.), Handbook of Perception and Human Performance, 1986, pp. 41–49.
[18] P. Grossman, J. Karemaker, W. Wieling, Prediction of tonic parasympathetic cardiac control using respiratory sinus arrhythmia: the need for respiratory control, Psychophysiology 28 (1991) 201–216.
[19] P. Grossman, Respiratory and cardiac rhythms as windows to central and autonomic biobehavioral regulation: selection of window frames, keeping the panes clean, and viewing neural topography, Biological Psychology 34 (1992) 131–161.
[20] P. Grossman, G. Stemmler, E. Meinhardt, Paced respiratory sinus arrhythmia as an index of cardiac parasympathetic tone during varying behavioral tasks, Psychophysiology 27 (1990) 404–416.
[21] S. Hart, J. Hauser, Inflight application of three pilot workload measurement techniques, Aviation, Space, and Environmental Medicine 58 (1987) 402–410.
[22] S. Hart, L. Staveland, in: P.A. Hancock, N. Meshkati (Eds.), Human Mental Workload, 1988, pp. 239–250.
[23] A. Hoover, E. Muth, A real-time index of vagal activity, International Journal of Human–Computer Interaction 17 (2) (2004) 197–209.
[24] P. Katona, F. Jih, Respiratory sinus arrhythmia: noninvasive measure of parasympathetic cardiac control, Journal of Applied Physiology 39 (1975) 801–805.
[25] C. Katsis, D. Katertsidis, C. Fotiadis, An integrated system based on physiological signals for the assessment of affective states in patients with anxiety disorders, Biomedical Signal Processing and Control 6 (3) (2011) 261–268.
[26] J.-Y. Kim, Detection of change in persistence of a linear time series, Journal of Econometrics 95 (2000) 97–116.
[27] S. Lal, A. Craig, A critical review of the psychophysiology of driver fatigue, Biological Psychology 55 (2001) 173–194.
[28] P. La Rosa, A. Nehorai, H. Eswaran, C. Lowery, H. Preissl, Detection of uterine MMG contractions using a multiple change point estimator and the K-means cluster algorithm, IEEE Transactions on Biomedical Engineering 55 (2) (2008) 453–467.
[30] S. MacDougall, A. Nandi, R. Chapman, Multiresolution and hybrid Bayesian algorithms for automatic detection of change points, IEEE Proceedings Vision, Image, and Signal Processing 154 (4) (1998) 280–286.
[31] G. Matthews, et al., Task engagement, cerebral blood flow velocity, and diagnostic monitoring for sustained attention, Journal of Experimental Psychology: Applied 16 (2) (2010) 187–203.
[32] M. Mboup, C. Join, M. Fliess, An on-line change-point detection method, in: In the Proceedings of 16th Mediterranean Conference on Control and Automation, 2008, pp. 1290–1295.
[33] H. Middleton, A. Sharma, D. Agouzoul, B. Sahakian, T. Robbins, Contrasts between the cardiovascular concomitants of tests of planning and attention, Psychophysiology 36 (1999) 610–618.
[34] E. Page, Continuous inspection schemes, Biometrika 41 (1/2) (1954) 100–115.
[35] R. Parasuraman, K. Cosenzo, E. De Visser, Adaptive automation for human supervision of multiple uninhabited vehicle: effects of change detection, situation awareness, and mental workload, Military Psychology 21 (2) (2009) 270–297.
[36] R. Parasuraman, M. Mouloua, R. Molley, Effects of adaptive task allocation on monitoring of automated systems, Human Factors 38 (1996) 665–679.

[37] S. Porges, Vagal tone: a physiologic marker of stress vulnerability, Pediatrics 90 (3) (1992) 498–504.
[38] S. Porges, D. Raskin, Respiratory and heart rate components of attention, Journal of Experimental Psychology 81 (1969) 497–503.
[39] M. Raifel, S. Ron, Estimation of slowly changing components of physiological signals, IEEE Transactions on Biomedical Engineering 44 (3) (1997) 215–220.
[40] J. Rand, A. Hoover, S. Fishel, J. Moss, J. Pappas, E. Muth, Real-time correction of heart interbeat intervals, IEEE Transactions on Biomedical Engineering 54 (2007) 946–950.
[41] C. Robazza, L. Bortoli, V. Nougler, Physiological arousal and performance in elite archers: a field study, European Psychologist 3 (4) (1998) 263–270.
[42] P. Romilly, Time series modelling of global mean temperature for managerial decision-making, Journal of Environmental Management 76 (2005) 61–70.
[43] S. Shaban, Change point problem and two-phase regression: an annotated bibliography, International Statistical Review 48 (1) (1980) 83–93.
[44] W. Shewhart, in: Economic Control of Quality of Manufactured Product, Macmillan Publishing, London, 1931.
[45] A. Singh, Automated state change detection of a Gaussian distributed signal, Master's thesis, Electrical & Computer Engineering Dept., Clemson University, 2008.
[46] R. Tahmasbi, S. Rezaei, Change point detection in GARCH models for voice activity detection, IEEE Transactions on Audio, Speech, and Language Processing 16 (5) (2008) 1038–1046.
[47] Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology, Heart rate variability: standards of measurement, physiological interpretation, and clinical use, Circulation 93 (March (5)) (1996) 1043–1065.
[48] J. Takeuchi, K. Yamanishi, A unifying framework for detecting outliers and change points from time series, IEEE Transactions on Knowledge and Data Engineering 18 (4) (2006) 482–492.
[49] J. Veltman, A. Gaillard, Physiological indices of workload in a simulated flight task, Biological Psychology 42 (1996) 323–342.
[50] G. Verdier, N. Hilgert, J.-P. Vila, Optimality of CUSUM rule approximations in change-point detection problems: application of nonlinear state-space systems, IEEE Transactions on Information Theory 54 (11) (2008) 5102–5112.
[51] H. Wang, D. Zhang, K. Shin, Change-point monitoring for the detection of DoS attacks, IEEE Transactions on Dependable and Secure Computing 1 (4) (2004) 193–208.
[52] J. Warm, G. Matthews, J. Finomore, Vigilance, workload and stress, in: P. Hancock, J. Szalma (Eds.), Performance Under Stress, Ashgate Publishing (2008).
[53] G. Wilson, Air-to-ground training missions: a psychophysiological workload analysis, Ergonomics 36 (36) (1993) 1071–1087.
[54] G. Wilson, An analysis of mental workload in pilots during flight using multiple psychophysiological measures, International Journal of Aviation Psychology 12 (1) (2002) 3–18.
[55] G. Wilson, C. Russell, Performance enhancement in an uninhabited air vehicle task using psychophysiologically determined adaptive aiding, Human Factors 49 (6) (2007) 1005–1018.
[56] P. Yang, G. Dumont, J. Ansermino, Adaptive change detection in heart rate trend monitoring in anesthetized children, IEEE Transactions on Biomedical Engineering 53 (11) (2006) 2211–2219.