A Real-Time Index of Vagal Activity

Adam Hoover, Ph.D.
Electrical and Computer Engineering Department
Clemson University
Clemson, SC 29634-0915
ahoover@clemson.edu

Eric Muth, Ph.D.
Department of Psychology
Clemson University
Clemson, SC 29634-1355
muth@clemson.edu

This work was funded by the Defense Advanced Research Projects Agency's Augmented Cognition Program managed by LCDR Dylan Schmorrow. The grant was administered through the Office of Naval Research as Award No. N000140210347, titled "Enhancement of Training and Performance Through Man-Machine Interactions Sensitive to Human Arousal and Task Difficulty".

Submitted May 2003 to the International Journal for Human-Computer Interaction.

Abstract

A computer system today receives no data regarding the physiological state of the user, but there are many cases where these data could be useful. For example, as the user becomes bored or lethargic, the system could raise the workload or audio-visual feedback to stimulate arousal. As the user becomes tense or strained, the system could lighten the workload or simplify the feedback to lessen arousal. This type of physiological-based closed-loop feedback could be applied in a number of scenarios, such as training, stressful repetitive work (e.g. air-traffic control), and military operations. In this work we describe a computing system that produces a real-time cardiac-based measure of arousal. The measure is based upon changes in respiratory sinus arrhythmia, an established measure of vagal activity. We describe the measure, the measure's potential and limits, the computing system, and the new directions for physiological monitoring research made possible by such a system.

Key words: augmented cognition, physiological monitoring, wearable computing

1 Introduction

A computing system today collects little if any data regarding the state of its human user. In many applications, performance could be improved if the machine had access to data on the physiological and cognitive state of the person. For example, as the user becomes bored or lethargic, the system could raise the workload or audio-visual feedback to stimulate arousal. As the user becomes tense or strained, the system could lighten the workload or simplify the feedback to lessen arousal. These examples of closed-loop feedback could be applied in a number of scenarios, such as training, stressful repetitive work (e.g. air traffic control), sustained equipment operation (e.g. long-haul driving), and military operations.

The Defense Advanced Research Projects Agency's Augmented Cognition Program has led an effort to explore sensors that index an individual's cognitive state. The goal of the program is to develop a closed-loop system that senses the human's state and reacts to it, allowing the human to accomplish more work as part of the human-machine system as outlined in the example applications described in the above paragraph. In order for an individual sensor to be useful in these applications, it must be capable of operating in a generic environment, provide meaningful feedback in continuous real time, and be cost-effective. To our knowledge, the sensing system we describe in this paper is the first to meet all these challenges. Numerous other human physiological measures have been investigated, including but not limited to electroencephalography, pupilography and dynamic position tracking. While these measures show promise for measuring an individual's brain activity, changes in pupil size and posture respectively, and indirectly indexing workload and or cognitive state, they currently either require equipment that is difficult to integrate in an applied environment, are not robust enough to be used in an applied environment, and/or require further basic research before meaningful real time output can be cost-effectively obtained from the sensor.

Our research explores using physiological monitoring and its relation to cognitive state in order to close the human-machine loop. Specifically we have focused on using cardiovascular indices of autonomic nervous system (ANS) activity or arousal. The ANS regulates internal states by acting as a feed-forward and feed-back system from the central nervous system (CNS) to the periphery and from the periphery to the CNS. The parasympathetic nervous system (PNS) is the branch of the ANS that is often associated with homeostatic functions, i.e. returning the organism to "rest". The sympathetic nervous system (SNS) is often associated with "fight or flight", arousing the organism. PNS and SNS interactions are complex in that the PNS and SNS can act independently, coactively or reciprocally [1]. Internal organs such as the heart and stomach respond to PNS and SNS interactions. During arousal, there is typically PNS withdrawal and SNS activation, while heart rate increases and stomach activity decreases.

We use a well-established, non-invasive, cardiovascular index, to objectively assess the parasympathetic activity of the autonomic nervous system, namely, respiratory sinus arrhythmia (RSA) [5]. When cardiac variability, the majority of which occurs between 0.0-1.0 Hz [7], is plotted as a continuous function against time, three periodic fluctuations can be observed. These include a low frequency peak between 0.04-0.08 Hz, a mid frequency peak centered around 0.1 Hz, and a high frequency peak between 0.15-0.5 Hz. The physiological basis for the high frequency component, or RSA, is well known and many studies have validated the use of various RSA measures as indices of PNS activity [2, 4, 6].

In this paper we describe a novel system we have developed that can track ANS arousal based upon tracking RSA. We call this system an *arousal meter.* The system operates in real-time, updating the arousal measure at 4 Hz. The system uses modest computing resources, enabling operation as a background process on a standard desktop or laptop computer. In the rest of this paper we describe in detail the theory and operation of the arousal meter. We then discuss some proof-of-concept experiments demonstrating its potential.

Figure 1: The EZ-IBI heart rate sensor, manufactured by UFI Corporation.

2 Methods

The arousal meter consists of two hardware components, a heart rate sensor and a computational platform (desktop or laptop computer). The heart rate sensor detects heart inter-beat-intervals (IBIs) and forwards them continuously to the computer for arousal analysis. All this hardware is commercially available off-the-shelf. Figure 1 shows an EZ-IBI heart rate sensor manufactured by UFI Corporation (Morro Bay, California). The EZ-IBI measures approximately 13 cm $\times$ 7 cm $\times$ 3 cm and is powered by a 9 V battery. For a computational platform we tried several different desktop and laptop computers, ranging from a 450 MHz Pentium 2 desktop computer to a 1.7 GHz Celeron laptop computer. All these platforms are sufficiently powerful to operate the arousal analysis as a background process while operating additional software. Section 3 examines the issue of the minimum required computational hardware in more detail.

The person to be monitored is connected to the EZ-IBI unit via three electrode leads as seen in Figure 2. Two active recording leads (black) are connected, one on the right side of the person just below the collar bone and one on the left side of the person just below the left breast. These two electrodes are connected to fetrodes (field effect transistors) that serve as amplifiers and increase the signal to noise ratio. These leads are positioned to minimize electrode movement and be in line with the major vector of depolarization of the heart. The third lead (green with grey) serves as a reference for signal noise reduction.

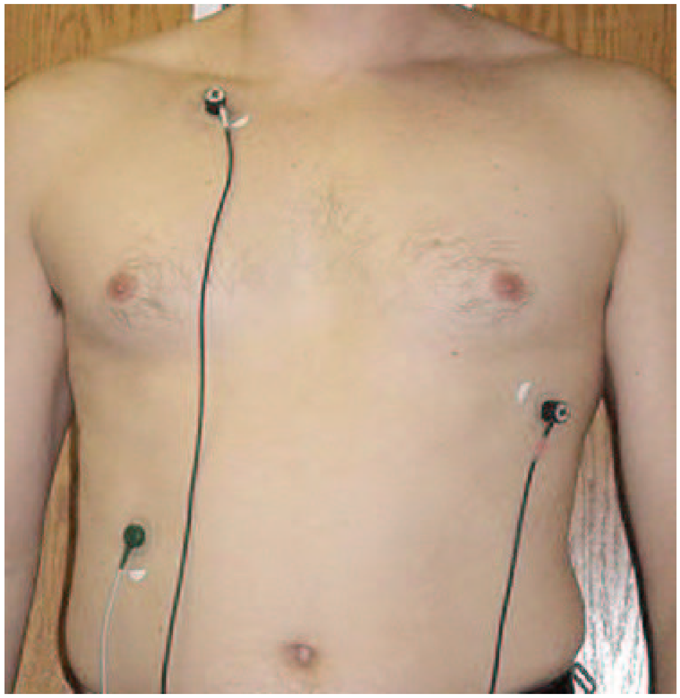

Figure 2: The arousal meter is connected to a person using three electrodes.

The EZ-IBI unit derives IBIs by monitoring the raw electrocardiogram (ECG) and recording the time interval between successive R-spikes of the cardiac QRS interval (see Figure 3). R-spikes are detected using a fixed threshold. Our experience with the EZ-IBI unit is that it is robust to motion generated artifacts, and generally shows less than 0.5% artifactual IBIs over an hour recording. IBIs typically range between 500 and 2000 ms. When IBIs are plotted against time, a waveform is evident in which the three periodic fluctuations described in Section 1 can be observed (called low, mid and high frequency). An arousal gauge reading is derived based upon analysis of the high frequency component.

Each IBI is transmitted to the computer via a serial cable immediately after its detection. In order to facilitate frequency analysis, IBIs must be synchronously resampled over time and must also be buffered over a sufficient period of time. The primary reason we sample is to create time-regularized data (fixed amount of time between consecutive samples), so that Fourier-based frequency analysis is valid. Porges [8] describes a method to resample by fitting a moving polynomial function to the asynchrnous IBIs and then solving for synchronous IBIs

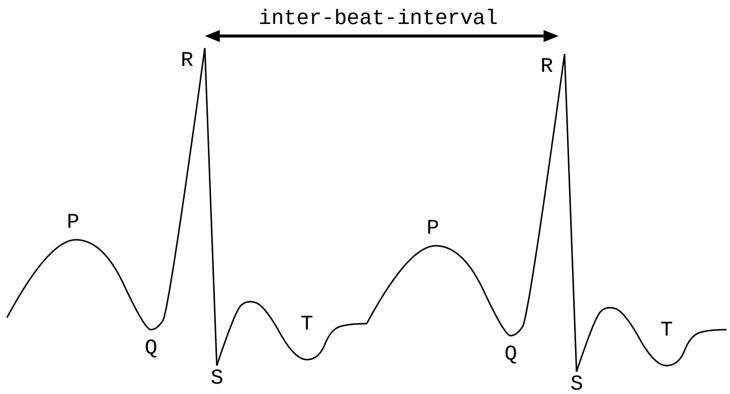

Figure 3: Measurement of a heart interbeat interval (IBI).

at the desired points in the function. We have avoided using this method because it was patented by Porges, and we are releasing both freeware and commercial versions of our system. In addition, we have found that simple oversampling of the asynchronous values is sufficient for our analysis, although one should be careful to remember that it does not increase the effective sampling rate.

Formally, we are resampling as follows. Let h_{a_i} represent a raw asynchronous IBI value (units are ms) received from the EZ-IBI unit, and h_{s_j} represent a synchronously resampled IBI value (units are ms), where $i = 0, 1, ...$ and $j = 0, 1, ...$ are integer indices of the data. Let Δt (units are ms) represent the synchronous resampling rate. The synchronous IBIs are oversampled from the asynchronous IBIs as

$$h_{s_j} = h_{a_i} \quad \forall j \text{ such that } \sum_{x=0}^{i-1} h_{a_x} \leq j\Delta t < \sum_{x=0}^{i} h_{a_x} \tag{1}$$

where $h_{s_0} = h_{a_0}$. Using Equation 1, an asynchronous IBI is copied at the syncrhonous sampling rate Δt until the next asyncrhonous IBI is detected. In accordance with the Nyquist Sampling Theorem, the sampling rate should be at least twice as great as the maximum expected frequency in the data. In our case, the smallest expected asyncronous IBI is approximately 500 ms, so we set $\Delta t = 250$ ms.

The periodic fluctuations observable in a plot of IBIs over time generally occur in the following ranges: low $\in$ [2...5] cycles per minute (cpm), mid $\in$ [6...9] cpm, and high $\in$ [9...30] cpm (these are the same ranges discussed in Section 1 transformed from Hz to cpm). In order to analyze power in these frequencies it is necessary to buffer at least 15 seconds of data. To provide some robustness we buffer approximately four times that amount of data. We perform power analysis using the Fast Fourier Transform (FFT) as detailed in [3]. This method requires a number of data points equal to a power of two. We therefore buffer the most recent 64 seconds of data (256 data points) for analysis.

Prior to transforming the data, we mean-center the data at zero

$$h_{m_k} = h_{s_j} - \frac{1}{256} \sum_{x=j-255}^{j} h_{s_x} \tag{2}$$

where the index $k = [0...255]$ while the index j refers to the most recently sampled (synchronous) 256 data points. We then apply a Hamming smoothing function

$$h_{h_k} = 0.54 + 0.46 \cos\left(\frac{2\pi w_k}{255}\right) h_{m_k} \tag{3}$$

where $w_k = k \ \forall \ k \in [0...127]$ and $w_k = w_{255-k} \ \forall \ k \in [128...255]$. These two operations are intended to reduce the aliasing caused by windowed Fourier analysis of the data. The FFT returns a number of power measures equal to the number of input data points. In our case, the returned power range is [0...240] cpm in 0.9375 cpm increments (4 samples per second times 60 seconds divided by 256 total data points).

The sequence of operations (synchronously sample new IBI, mean center most recent 256 data points, Hamming smooth those data, then FFT) is performed at the sampling rate Δt. In order to provide additional robustness, the power measures are averaged across the most recent 60 FFT analyses. We then determine the magnitude of the maximum power in the desired high range [9...30] cpm. Let this maximum average power be denoted p_j, where the index j indicates the maximum average power reading computed at time j associated with the most recently sampled IBI h_{s_j}. This power term is what drives the arousal meter gauge reading. However, we assume that different people have different populations of maximum power readings over time. In order to normalize the power readings to each individual, we compute the Z score of the power reading as

$$Z_j = \frac{p_j - \mu_p}{\sigma_p} \tag{4}$$

where μ_p and σ_p are the mean and standard deviation of the power readings for the given individual. We update these statistics continuously as the arousal meter runs using

$$\begin{aligned} s_j &= ({\sigma_p}^2 + {\mu_p}^2)J \\ \mu_p &= \frac{\mu_p J + p_j}{J+1} \\ \sigma_p &= \sqrt{\frac{s_j + {p_j}^2}{J+1} - {\mu_p}^2} \end{aligned} \tag{5}$$

where J is the total number of readings over which the statistics are being computed. Assuming that the distribution of power readings for an individual is close to a normal distribution, we can expect over 86% of the Z score readings to lie in the range [-1.5 ... 1.5]. We therefore cap the readings to lie within this range and drive the arousal meter gauge in the range [0...1] as

$$a_j = 1 - \frac{z_j + 1.5}{3} \tag{6}$$

where a higher a_j reading indicates a higher arousal.

Care must be taken in the implementation of timing control for these methods. We assume that all the above calculations can be completed within each period Δt of sampling, and that the arousal meter process will then "sleep" some amount of time before beginning the next iteration. Let t_j (units are ms) represent the actual amount of time taken to complete the calculations during the current iteration. The simple approach to maintain regularity of sampling would be to compute t_j independently each iteration and then sleep $\Delta t - t_j$. However, the inherent uncertainty in precision of measuring time within a computing process introduces a small error each iteration. Over time, this error accumulates. We have observed timing drift on the order of five seconds per minute with this approach. The appropriate alternative we employ is to track the timing on J the total number of sampling periods executed, sleeping $J\Delta t - t_e$ at the end of each iteration, where t_e is the total elapsed time since the arousal meter process started. In the other extreme, if the arousal meter process does not finish its calculations within the required time (for example, when another process running on the computer swaps in for a period greater than 250 ms) then one or more sampling intervals are missed. In this case, the arousal meter process copies the last read IBI value filling in the gap. In our experience, such gaps are small (less than one second) and occur rarely with normal usage (less than once per ten minutes).

Figure 4 shows a snapshot of the graphical user interface of the arousal meter. The most recent synchronous IBI is displayed top-left, and the parameters of the arousal meter are listed on the left side. Part A in Figure 4 shows a plot of the most recent IBIs over time. Part B shows the power readings for the current iteration. Part C1 shows the maximum average power plotted over the last two hours, while part C2 shows the same statistic over only the last minute. This statistic is shown over these two ranges to enable visualization of arousal patterns over varying scales of time. Part D shows a gauge-like reading of the arousal meter score (Equation 6) plotted in a semicircle from $\pi ... 0$ degrees, where 0 degrees (due East) represents the highest arousal. The arousal meter score is also continuously pushed into shared memory where it can be accessed by any other process running on the same computer.

3 Experiments

In this section we describe two proof-of-concept experiments. These experiments demonstrate how the arousal meter works. We also briefly describe two additional studies in progress. These studies show how the arousal meter can close the human-machine loop in new ways, and consequently how it enables new research.

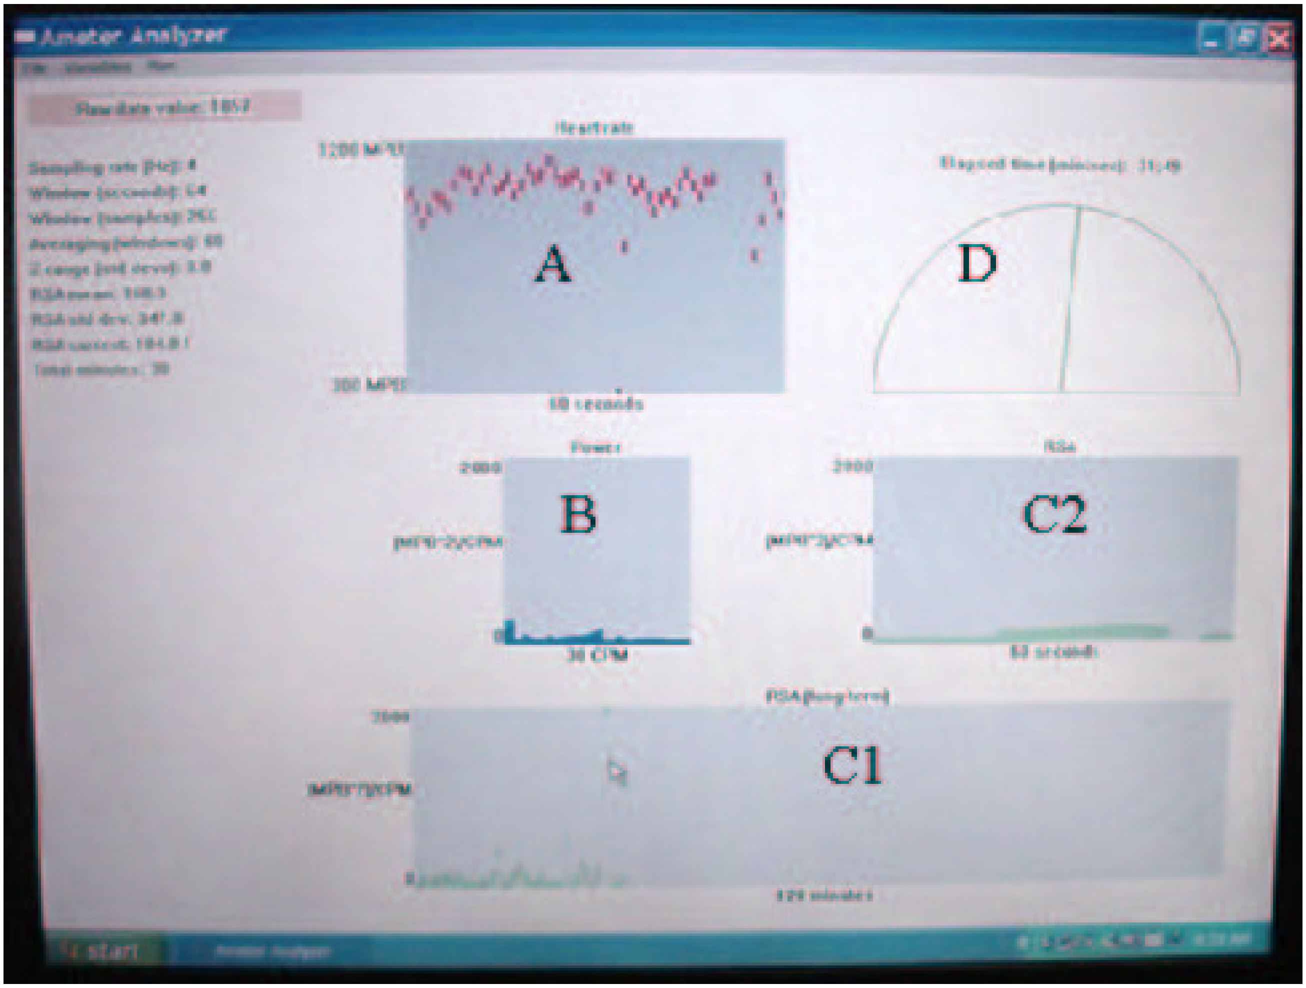

Figure 4: Screen capture of the arousal meter interface (see text).

3.1 Gauges task experiment

In this experiment, volunteers completed a set of six different tasks of varying physical and cognitive difficulty. The arousal meter was run continuously to determine the average arousal measure of the volunteers during each task. It was expected that the tasks of varying physical difficulty would produce significantly different arousal measures. The main hypothesis to be tested was whether or not the tasks of varying cognitive difficulty would produce significantly different arousal measures.

The first three tasks were variations of the Gauges Task, a dual attention task developed by Select International Corporation (San Diego, California) and used with permission for

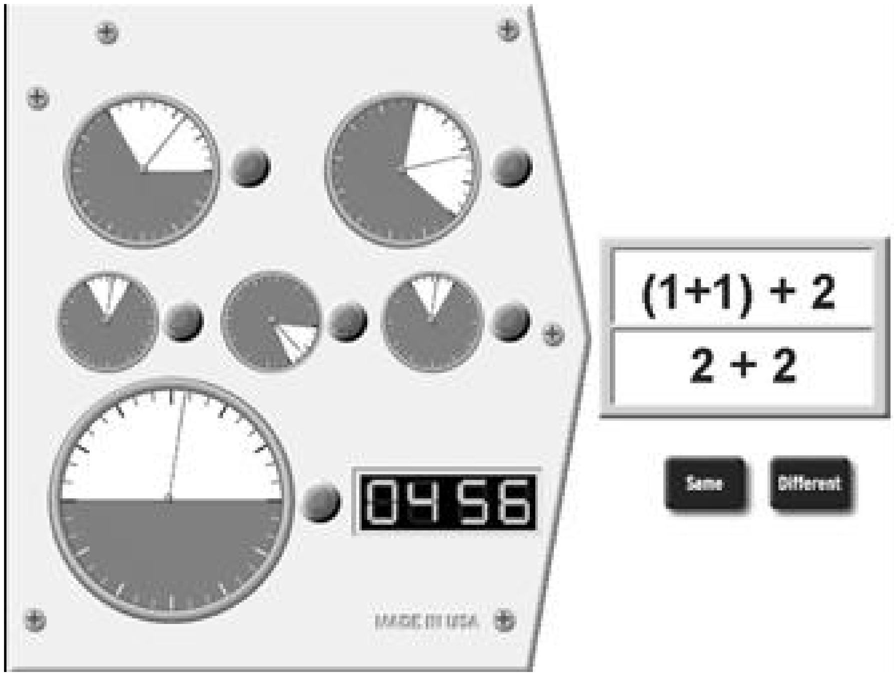

Figure 5: Screen capture of the interface for the gauges task.

this research. The basic screen displayed in this task is shown in Figure 5. The goal in this task is to calculate the values of the two arithmetic expressions on the right and to determine if they are identical, responding with the "same" and "different" buttons under the expressions. At the same time, it is necessary to monitor the six gauges on the left of the display and to reverse the direction of the needle as it approaches the red area. The needle of each gauge moves slowly and steadily in either direction. The small buttons to the right of each gauge reverse the direction. The participant is penalized for the number of incorrect responses to the math items and for the number of seconds that the gauge needles fall in the red zone. The three versions of this task varied in difficulty: slow gauges and easy math, fast gauges and easy math, fast gauges and hard math. The order of presentation of the levels of difficulty was counterbalanced so that each one appeared equally often as first, second, or third.

The fourth task for all subjects was a 5 minute rest period. This task is of low physiological difficulty and allows arousal comparison to tasks of higher physiological difficulty.

The fifth task for all subjects was a 2 minute effort to maintain pressure on a hand grip device. Each subject was first asked to squeeze the device as hard as he or she could. Once the maximum effort was determined, the subject was then instructed to maintain 25% effort of that maximum for 2 minutes. A researcher monitored the level and verbally instructed the subject to increase or decrease effort as needed.

The sixth and final task for all subjects was a dual attention task in which the subject was required to fix attention on a spatial shape that moved about the screen and at the same time to count down from 200 to 0 by 3's. The screen at all times showed a blue circle and a red square. Every 3 seconds, these two shapes changed location. The subject was instructed to keep his or her eyes on the blue circle while counting. The resuults of this task are not reported here but data from this task were used in the normalization procedure described below.

Twenty students were recruited as test subjects, of which five were excluded for various reasons (one served as the pilot, two were unable to complete the study, and two had problems during data collection that voided the data).

Ideally, the arousal meter would be calibrated to each individual by having the individual wear the arousal meter for a given amount of time and complete a representative set of tasks of varying physiological difficulty. Thus when the individual started the actual experimental tasks the arousal meter would have sufficient data to calculate a stable mean and standard deviation of arousal (the μ_p, σ_p and J terms described in Equations 4-5). However, in this study, there were time limitations on data collection and pre-experiment calibration was not possible. Instead, we analyzed the data from this study utilizing a post-study calibration procedure. First, within an individual, the average power μ_p was calculated for each task. Second, a mean and standard deviation of power was calculated for each individual across the tasks. Each individual's average power for each task was then converted to a z-score using the individual's mean and standard deviation of power.

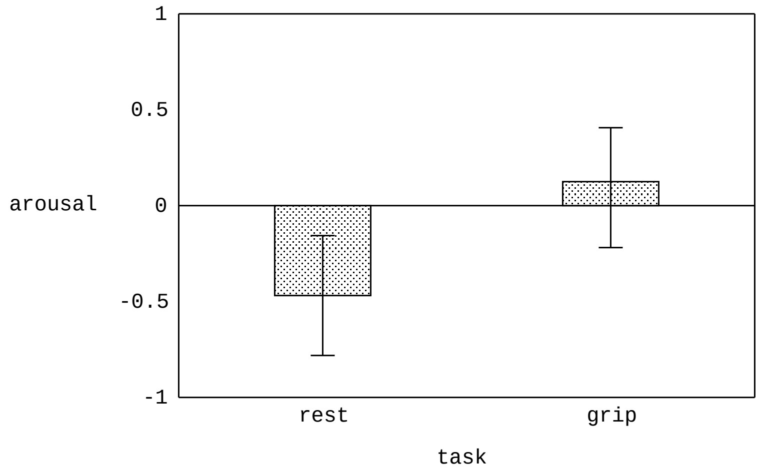

Figure 6: Average subject arousal on physical tasks.

The goal of this study was to see if we could detect a "state shift" in arousal, a change from relaxed to aroused or vice versa (operationally defined as crossing zero in the normalized scale). This was accomplished by comparing a resting period to handgrip, a known physiological challenge. It was expected that this would result in a state shift from relaxed (negative in the normalized scale) to aroused (positive in the normalized scale). Other tasks with unknown arousal levels (gauges 1, 2 and 3) were also compared to resting to see if any of these tasks caused a state shift. Note that it was not the goal of this study to compare all of the tasks to each other (hence why an omnibus ANOVA was not employed), in order to test all relative arousal changes, but only to determine if any of the tasks causes a "state shift" from rest.

Figure 6 shows the average arousal level for all participants on the two physical tasks (rest and hand grip). As expected, the hand grip task was significantly more arousing than the rest task ($t[14] = 1.810, p < 0.05$). These data confirm that the arousal measure described in this paper is sensitive to manipulations that cause physiological (sympathetic) arousal.

Figure 7 shows the average arousal level for all participants on the three variations of the gauges task. The easy and moderate levels of the gauges task did not show significant differences in arousal from the baseline (the rest task). However, the difficult level of the gauges task produced an average arousal significantly higher than the average arousal during

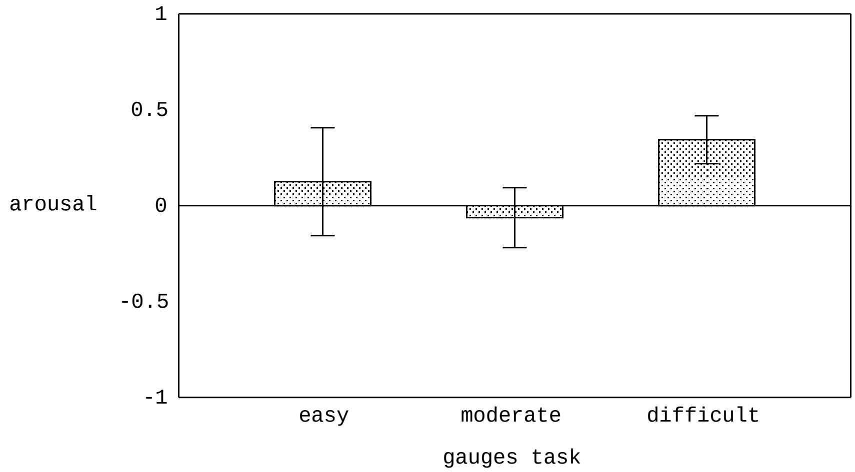

Figure 7: Average subject arousal on the varying levels of a cognitive task.

the rest task ($t[14] = 2.354, p < 0.05$). This provides some evidence that cognitive taskload has an effect on the arousal measure described in this work.

3.2 Day monitoring

Changes in arousal can occur over relatively short (e.g seconds) or long (e.g. an hour) periods of time. The arousal meter produces a continuous reading at 4 Hz, so it is possible to monitor events of relatively short duration. The gauges task study just described is an example where the events of interest each occured for a few minutes. By compressing the arousal meter readings over time (for example, taking the average reading over a range of time) it is also possible to monitor events of longer duration. Figure 8 shows an arousal meter log for a person wearing the arousal meter during a nine hour period. The data shows a low average arousal early in the morning (waking up), a little relaxation around lunch time (parasympathetic response to eating), and then a peak mid afternoon (around the time the subject was giving a public presentation). These data were not part of a formal study, so testing conditions were not controlled. We present them only to show how the arousal meter is also capable of providing data useful for monitoring events of relatively longer duration.

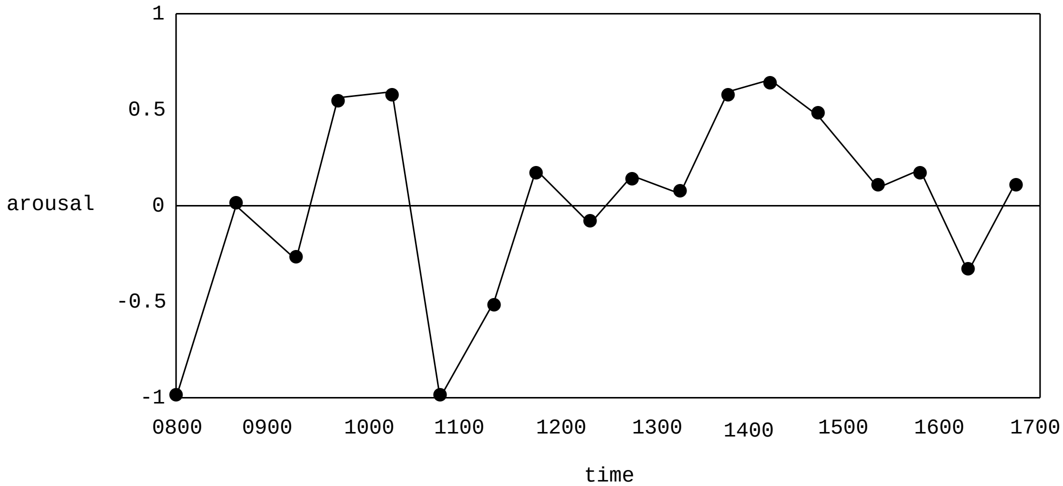

Figure 8: Arousal meter log for one subject during nine hour period.

3.3 Ongoing studies

We currently have two simultaneous research tracks underway. One track is a basic science track that is investigating the current arousal meter analysis approach and how it relates to other subjective and objective measures as well as looking at alternative heart-rate variability analyses. An example of a study in this track is a sleep deprivation experiment for which data have been collected and data analysis is underway. In this sleep experiment participants were instructed to go to sleep between 11 PM and 1 AM the night before the study and to get out of bed between 7 and 9 AM the morning of the study and not to nap during the day. Participants were asked to report to the sleep research laboratory at 5 PM on the day of the study and remained in the lab until approximately 11 AM the next morning, being kept awake (sleep deprived) for approximately 24 hrs. During the sleep deprivation, heart rate was monitored continuously and the participants completed a number of low, moderate and high workload tasks with task performance being recorded. Other physiological measures such as body temperature and eye closure were also monitored and subjective measures obtained at certain times during the study. This study will allow for an evaluation of the arousal meter under varying task workload states as well as arousal states and may lead to better arousal meter algorithms.

The second track is taking a leap forward in an applied approach to investigate the utility of closed-loop systems that use physiological feedback. Data are currently being collected on the effectiveness of a prototype closed-loop system. The system embeds the arousal meter into a computer task. The computer task is a game with primary and secondary task components. There is an assistance mode for the secondary task that can be turned on and off manually or automatically based on the arousal meter score. The goal is to manipulate the level of computer assistance in response to changes in arousal level with the expectation of keeping arousal at a level that optimizes performance. There is one independent variable, control of computer assistance, with three between-subjects conditions: manual user control, automated control based on heart rate variability, and a control condition where players are yoked to the heart rate information of the previous automated-condition participant. The dependent variable is the score on the game. We hypothesize that computer-controlled assistance manipulation based on arousal meter data will result in the highest scores on the game.

4 Conclusion

The physical size of all the required components, including computing resources, enables us to hypothesize the construction of a wearable arousal meter in the near future. The arousal meter process uses less than 10% of the processor time on a 450 MHz Pentium 2 desktop computer, and requires less than 100 KB of memory (for both program and data). It should be feasible to use a relatively generic embedded processor to run the arousal meter process. This processor could be placed inside the EZ-IBI heartrate sensor, along with circuitry to wirelessly transmit the arousal meter reading. We are currently working with UFI Corporation on this design with plans to build a prototype wearable arousal meter.

In addition to the ongoing studies described in Section 3.3, we are working to improve the arousal measure. We are studying the wavelet transform and its applicability to IBI analysis. While IBI fluctuations are periodic, they are not necessarily sinusoidal. It may be possible to apply a basis function that more exactly matches the patterns being sought. The wavelet transform also does not suffer from aliasing due to windowed analysis. For these two reasons, we believe that using the wavelet transform will increase the fidelity of our arousal measure.

Although our IBI detector produces few errors, we have found that even a single IBI error causes the arousal measure to spike for a short period surrounding the bad IBI. According to a recent report [10] from the Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology, and a comprehensive review paper by Porges and Byrne [9], there are as yet no known methods for automated IBI artifact detection and correction. However, in order for a closed-loop system to operate using any sort of IBI analysis, automated error detection and correction is essential. This is a topic of our current research.

The arousal meter described in this work is based only on measurement of the PNS subcomponent of ANS activity. This results in limited sensitivity to maximally arousing situations which are driven by the SNS. We are examining measures that provide more comprehensive indication of ANS activity. We believe the data collected utilizing the sleep deprivation paradigm described above and the applications of tools such as the wavelet transform will lead to more complete heart-rate variability derived ANS measures of arousal. This will in turn lead to the development of a better arousal meter.

References

[1] G. Berntson, J. Cacioppo and K. Quigly, "Autonomic determinism: The modes of autonomic control, the doctrine of autonomic space, and the laws of autonomic constraint", in *Psychological Review*, vol. 98, pp. 459-487, 1991.

[2] D. Eckberg, "Human sinus arrhythmia as an index of vagal cardiac outflow", in *Journal of Applied Physiology*, vol. 54, pp. 961-966, 1983.

[3] R. Gonzalez and R. Woods, *Digital Image Processing*, Addison-Wesley Publishing, Reading MA, 1992.

[4] P. Grossman, J. Karemaker and W. Wieling, “Prediction of tonic parasympathetic cardiac control using respiratiory sinus arrhythmia: The need for respiratory control”, in *Psychophysiology*, vol. 28, pp. 201-216, 1991.

[5] P. Grossman, “Respiratory and cardiac rhythmis as windows to central and autonomic biobehavioral regulation: Selection of window frames, keeping the panes clean and viewing the neural topography”, in *Biological Psychology*, vol. 34, pp. 131-161, 1992.

[6] P. Katona and F. Jih, “Respiratory sinus arrhythmia: noninvasive measure of parasympathetic cardiac control”, in *Journal of Applied Physiology*, vol. 39, pp. 801-805, 1975.

[7] E. Mezzacappa, D. Kindlon and F. Earls, “The utility of spectral analytic techniques in the study of the autonomic regulation of beat-to-beat heart rate variability”, in *International Journal of Methods in Psychiatric Research*, vol. 4, pp. 29-44, 1994.

[8] S. Porges, “Method and apparatus for evaluating rhythmic oscillations in aperiodic physiological response systems”, U.S. patent #4,510,944, April 16, 1985.

[9] S. Porges and E. Byrne, “Research methods for measurement of heart rate and respiration”, in *Biological Psychology*, vol. 34, pp. 93-130, 1992.

[10] 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”, in *Circulation*, vol. 93, pp. 1043-1065, 1996.