An Algorithm for the Automatic Analysis of Signals From an Oyster Heart Rate Sensor

Andrew D. Hellicar, Ashfaqur Rahman, Daniel V. Smith, Greg Smith,
John McCulloch, Sarah Andrewartha, and Andrea Morash

***Abstract*— An *in situ* optical oyster heart rate sensor generates signals requiring frequency estimation with properties different to human ECG and speech signals. We discuss the method of signal generation and highlight a number of these signal properties. An optimal heart rate estimation approach was identified by application of a variety of frequency estimation techniques and comparing results to manually acquired values. Although a machine learning approach achieved the best performance, accurately estimating 96.8% of the heart rates correctly, a median filtered autocorrelation approach achieved 93.7% with significantly less computational requirement. A method for estimating heart rate variation is also presented.**

***Index Terms*— Biomedical signal processing, frequency estimation, machine learning.**

I. Introduction

HEART rate and heart rate variability in intertidal invertebrates are proxies for physiological responses (such as oxygen transport, neural feedback, stress) and tracking these quantities is highly informative. For example they allow the impact of external environmental factors to be quantified and enable optimization of conditions for breeding and consumption stock. Consequently there is a need for heart rate monitoring systems for these animals.

Measurement of heart rate in humans is a well established and mature science [1]; ECG systems detect electrical signals generated by the heart. These signals are amenable to analysis and were first analysed by computer in the 1960s [2]. Real-time heart rate computer analysis systems were in use by the 1980s [3]. Analysis is based on detecting the repetition and shape of the ECG signal structure referred to as the QRS complex. Unfortunately such approaches are not applicable to intertidal invertebrates (such as the oyster) which possess a shell preventing use of externally applied ECG sensors. Early experiments on the bivalve class of intertidal invertebrates observed heart rate by measuring the heart's electrical impedance via implanted electrodes [4]. However, it is critical that the system measuring heart rate does not modify the animal's heart rate by (for example) increasing animal stress. To mitigate the invasive nature and degradation of electrical contacts in the impendence approach an alternative optical coupling method was proposed [5]. This approach illuminates the heart with an optical signal and detects the intensity of reflected light which is modified as a consequence of any motion of the heart's surface. For thin shelled intertidal invertebrates a near infrared source and detector are collocated and positioned on the shell's outer surface. The NIR signal propagates through the shell without need for drilling into the shell. For thick shelled invertebrates such as the oyster, the approach is similar with the exception of drilling a small hole into the shell [6] which would otherwise block the optical signal. The light emitter and detector are mounted and sealed to ensure the integrity of the shell is maintained. Optical techniques avoid direct contact with the animal; however, the approach introduces its own problems in terms of signal analysis due to signal complexity as pointed out in a recent review paper [7]. Therefore there is a need to develop robust signal analysis approaches operating over long installations possible with these sensors. In this paper we analyze the signal from a sensor in situ, and develop algorithms to detect short term average heart rate and variation of heart rate within short (20s) signal sequences.

The oyster biosensor [8] (Fig. 1) provides information on the animal through both shell gape and heart rate detection. Gape refers to the opening of the two shells which occurs when the oyster is submerged and feeding and/or taking in oxygen. The gape sensor (based on a Hall effect sensor) detects open and closed states of the oyster shell by measuring changes in the output voltage of a transducer located on the upper shell in response to a magnetic field induced by a magnet located on the lower shell. The heart rate sensor illuminates the oyster heart with a fixed intensity 950nm (NIR) optical signal generated by an LED source, and measures the reflected optical intensity via an IR diode detector. The LED source and detector are collocated and mounted within 5mm of the heart surface. The sensor responds with a time constant of 20μs to reflections within a cone angle 55 degrees from the sensor axis out to a range of 5mm. The principal of operation of the heart rate sensor is that the heart's surface reflects light which is measured by the IR sensor. Any motion

Manuscript received January 20, 2015; revised March 16, 2015; accepted March 21, 2015. Date of publication April 13, 2015; date of current version June 17, 2015. This work was supported in part by the Tasmanian Government through the Tasmanian Department of Economic Development and in part by the CSIRO Food Futures Flagship. The associate editor coordinating the review of this paper and approving it for publication was Dr. M. R. Yuce.

A. D. Hellicar, A. Rahman, D. V. Smith, G. Smith, and J. McCulloch are with Commonwealth Scientific and Industrial Research Organisation, Battery Point, TAS 7004, Australia (e-mail: andrew.hellicar@csiro.au; ashfaqur.rahman@csiro.au; daniel.v.smith@csiro.au; gregory.smith@csiro.au; john.mcculloch@csiro.au).

S. Andrewartha and A. Morash are with Commonwealth Scientific and Industrial Research Organisation, Battery Point, TAS 7004, Australia, and also with the University of Tasmania, Sandy Bay, TAS 7005, Australia (e-mail: andrew.hellicar@csiro.au; andrea.morash@csiro.au).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSEN.2015.2422375

(a)

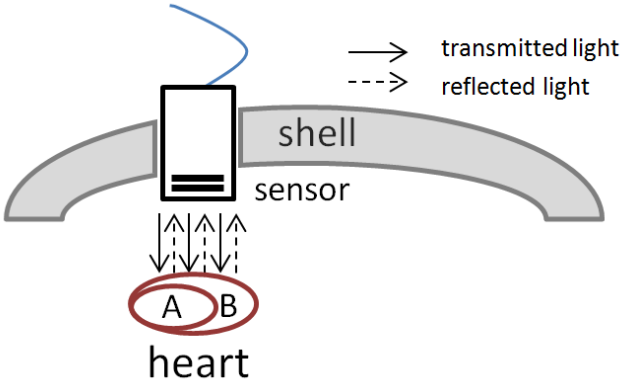

(b)

Fig. 1. Oyster attached to biosensor. (a) Photo of sensor and oyster. (b) Schematic of sensor inserted into hole in oyster's shell. Heart size depends on whether it is in A: contraction phase (systole) or B: expansion phase (diastole). Reflected light rays shown for expanded heart. For contracted heart fewer rays will strike heart with corresponding change in detected signal level.

in the heart causes perturbation in the surface geometry, thus modifying the intensity of the reflected optical signal. If the heart's surface geometry varies periodically (as is the case for a beating heart) the measured signal exhibits a period equal to that of the heart's motion. Measurement of the signal period is used to infer heartbeat period and consequently heart rate. The sensor is aligned such that signals are detectable both when the oyster is open and closed. Reflected light intensity is sampled at 100Hz over 20 seconds, resulting in sequences containing 2000 samples. This sampling regime is selected to capture multiple heartbeats per sequence, which are expected to have durations within the range of 0.5 – 5 seconds [9]. Currently one sequence is recorded approximately every six minutes for the duration of sensor installation, which has been running for nine months. One gape value is recorded for each sequence. Three classes of sequence occur: saturated (Fig. 2a), aperiodic (Fig. 2b), and periodic (Fig. 3). Any algorithm must classify a sequence as one of these three classes, and determine the heart rate and heart rate variation of sequences classified as periodic.

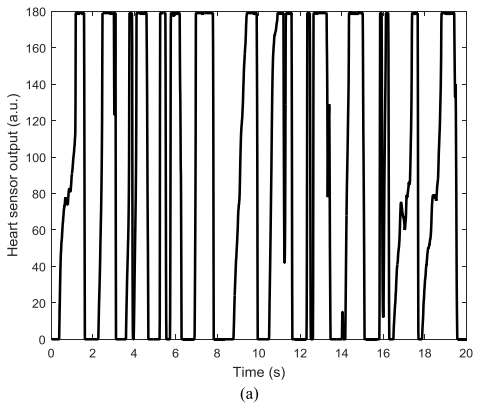

(a)

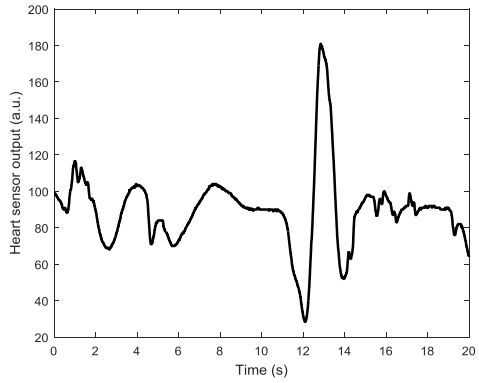

(b)

Fig. 2. Examples of saturated (a) and aperiodic (b) sequences.

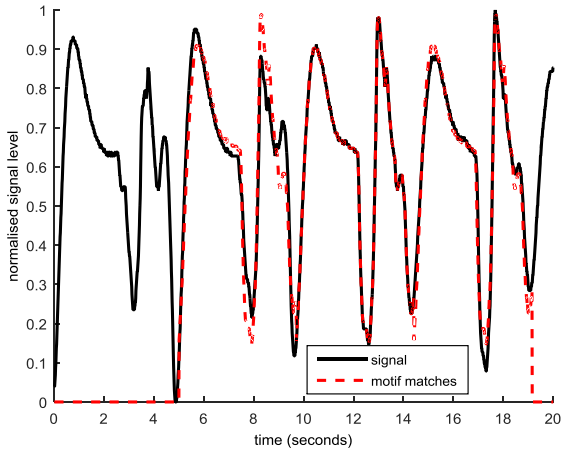

Fig. 3. Example of periodic time domain signals including 4 heart beats. Note peak value within beat occurs for first pulse in two beats up to 10s, and second pulse for last 2 beats. Sequence motif (described in Section II) shown in red/dashed overlays signal from 9.5 to 14.5 seconds. Best shifted motif matches shown from 5.0-9.5 seconds and 14.5-19 seconds. 5-9.5 and 15-20 seconds.

Saturated signals likely occur due to the presence of external light sources, or due to sub-optimal biasing of detector electronics. Aperiodic signals likely occur in the absence of strong signal coupled off the heart surface, or absence of a beating heart. Periodic sequences result as a consequence of the beating heart periodically modifying the intensity of the reflected signal coupled into the detector. The dual chamber heart comprises a smaller ventricle and larger

auricle which each beat once per cycle with a time delay between the two chamber contractions. Captured periodic signals (Fig. 3) can exhibit a strong second-order harmonic due to the combination of two temporally offset pulses of different amplitude. Motion of both chambers is detectable either by direct illumination, or physical coupling through the heart tissue into the illuminated region. The optical reflection sensing modality is sensitive to a combination of specular and multiple reflections, and is therefore highly sensitive to geometry. Furthermore the geometrical configuration evolves over time and any technique must be robust against shape and amplitude modification of the pulse.

Many approaches exist for frequency estimation. One popular method is the autocorrelation method (ACM) [10]. For a signal sequence S containing L samples the sequence mean is subtracted and a signal S^p of length $2L$ is generated by zero padding ($L = 2000$). The period T is estimated by finding m that maximizes the autocorrelation function:

$$A(m) = \frac{1}{L} \sum_{t=0}^{2L-1-m} S^p(t)\, S^p(t+m) \tag{1}$$

such that $m \in [m^{min}, L/2]$, where m^{min} is the lag at which the autocorrelation first becomes negative, which removes the signal coherence time. The sequence is classified as periodic if peak of A is larger than a threshold parameter which depends on signal quality and noise level [11]. The ACM is efficiently implemented by calculating autocorrelation using the Fourier transform approach.

An alternative Hanning window (HWE) approach introduced in [12] weights the sequence S with a Hanning window. The method follows the ACM with the modification replacing $A(m)$ by $W(m) \cdot A(m)$ in Equation (1) before calculating the peak and estimating T. $W(m)$ is the reciprocal of the Hanning window's autocorrelation function. Therefore the HWE approach is weighted towards selecting the longer lags. Other lag domain methods include centre clipped autocorrelation (CCA) [13] and normalised autocorrelation (NAC) [14] approaches. The centre-clipped approach sets samples with amplitudes less than a fixed threshold (often 50% of maximum signal amplitude) to zero resulting in a flattening of the spectrum and sharpening of the lag domain. However, since the low amplitude values are removed the method is sensitive to variation in the high amplitude region, (such as may be the case for specular reflections off a surface). The NAC divides the lag domain values by the power in the delayed version of the signal. This weighting offsets the lag domain tapering caused by zero padding of the signal; however, it is also therefore prone to selecting longer period signals. Time domain methods include the average magnitude difference function (AMD) [15] and the maximum likelihood method (MLH) [16]. The AMD and MLH both optimize measures calculated on the time domain signal and delayed version of the signal directly (difference and sum respectively). They are similar to autocorrelation approach without the relatively computationally expensive use of multiplications required in the lag domain approaches. Finally a popular frequency domain method is the cepstrum method (CEP) [17] which is useful for period estimation where a strongly periodic signal is convolved with a second signal. CEP recasts the corresponding product in the frequency domain into a sum via application of a logarithm. Each technique has specific advantages and disadvantages, this motivated a machine learning approach [18] using the output of the six frequency estimation approaches (ACM, CCA, NAC, AMD, MLH, CEP) as features which can exploit the benefits of each method.

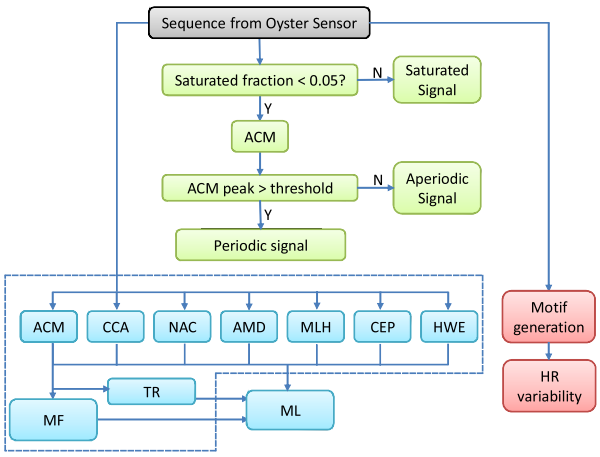

Fig. 4. Flow chart of implemented algorithm including classification phase (top/green), heart rate estimation phase (bottom/blue) and heart rate variability phase (right/red).

Here we extend the preliminary results reported in [18] by adding the HWE estimator (a total of 7 basic estimators), and add two aggregate estimators: tracking (TR) and median filtering (MF) which use information over multiple consecutive sequences. We use these nine estimators as features for a machine learning approach which achieves accuracy superior to that achieved in [18]. We extend the analysis by presenting the performance of all approaches and their sensitivity to selecting incorrect period ranges. We also highlight some of the types of signals and sources of error affecting our approach. Finally we introduce a method of heart rate variance estimation, not previously implemented on this class of sensor. As these sensors become more widely deployed the issues we discuss will need resolving in algorithms embedded in the sensor itself. Therefore computationally efficient and accurate approaches are required.

II. Method

The algorithm is split into three phases: a classification phase identifies the signal type, a heart rate estimation phase determines heart rate for periodic signals, finally heart rate variability is estimated. A diagram of the algorithm is shown in Fig. 4. The classification phase identifies saturated sequences (Fig. 2a) by detecting the proportion of time the signal is saturated. The ACM is used to differentiate between periodic and aperiodic signals. For remaining periodic signals all seven basic estimators named in Section I were implemented generating seven estimates of heart rate. The ACM is biased towards shorter lags due to the presence of zero padding and care is needed to ensure the 2nd harmonic due to the

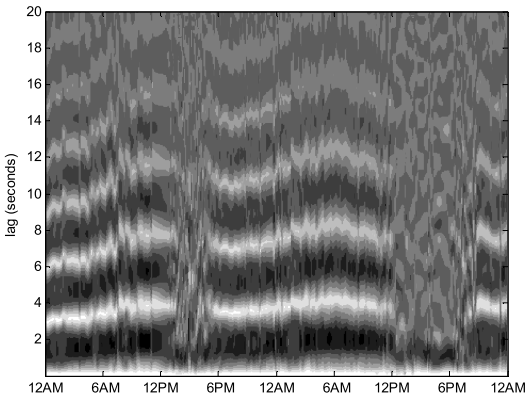

Fig. 5. Lag domain amplitude(white positive, black negative) from sequences gathered over a 48 hour time frame November 3–November 4, 2014.

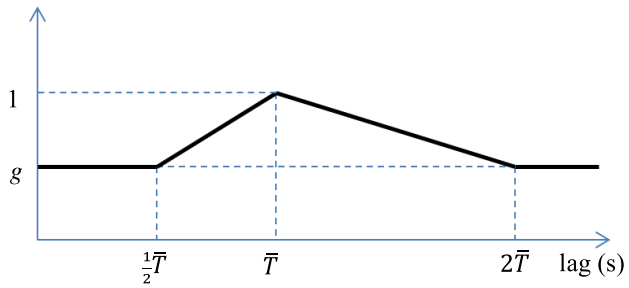

Fig. 6. Tracking estimator's weighting function for scaling autocorrelation approaches' amplitude in the lag domain.

dual chamber beating is not erroneously selected. Alternative methods such as the NAC and HWE exist to compensate for this.

Despite the gaps between sequences, the slow change in oyster heart rate results in similar periodic structure which is apparent when plotting the autocorrelation function for successive samples (Fig. 5). This long term stability and the HWE lag domain weighting approach motivated a simple tracking (TR) estimator which is biased towards selecting heart rates similar to those estimated from the previous sequences where N_{TR} is the number of sequences used in the calculation. The motivation is to reduce the occurrence of erroneous heart rate estimates at half and double the actual rate (typical errors for estimators on signals with periodic beats containing double peak structure). To achieve this bias the lag space is weighted with a piece-wise linear window. The window was maximal (and unit valued) at the mean detected period ($\bar{T}$) over the previous N_{TR} sequences, and linearly decreases for smaller (and larger) lags until reaching a minimum value $g < 1$ at lags corresponding to half (and double) mean periods. All lags outside this range are also weighted by g, the weighting window is shown in Fig. 6. Optimal values of g and N_{TR} were estimated for the tracking estimator.

We also applied a median filtering (MF) estimator applied to the best performing basic estimator:

1) *The median period was calculated for the N^1_{MF} sequences prior to and including the sequence being analysed.*
2) *Sequences with periods more than 20% different to the median period were removed from the set of N^1_{MF} periods.*
3) *From remaining sequences the median period was calculated for the N^2_{MF} sequences prior to and including the sequence being analsyed.*

The longer time frame corresponding to N^1_{MF} provides robust longer term averages, whereas the potentially shorter time frame N^2_{MF} allow accurate estimations of short term transitions in heart rate. These values were optimally evaluated for the median filter approach.

Finally we implemented a machine learning approach using nine frequency estimators (within dashed block in Fig. 4) as input features. Three WEKA [19] implementations were assessed including a three layer multilayer perceptron (MLP) with six hidden nodes (larger networks showed no improvement in performance). A Linear regression (LR) and a support vector regression (SMO).

Performance of the techniques was quantified by comparing estimated heart rates with values acquired by manually observing sequence class and heart rate within the first 2 months of collected data. After removing saturated sequences from the data set, sequences were randomly selected and classified as either periodic or aperiodic based on visual inspection until 100 sequences were labelled for each class (200 in total). Sixteen regions of consecutive periodic signals were identified and the periods of 9 sequences (evenly spaced throughout each region) were measured by visual inspection resulting in a total of 144 measurements. Ground truth periods were calculated by linearly interpolating the measured sequence periods across each span of sequences resulting in a total of 1450 sequences with ground truth period data. For the machine learning approach the 1450 sequences were split into 10 sets of 145 consecutive sequences and a 10-fold cross-validation used to evaluate performance, training on 9 sets and testing on the remaining set.

In addition to heart rate the variation of heart rate is also of interest as an indicator. Whereas heart rate variation over longer time frames can be calculated as the difference between heart rates estimated from sequences with the desired temporal separation, short term (intra-sequence) variation requires an estimate of the duration of each heartbeat. Because the structure of the heartbeat signal can change within a sequence a motif approach was used to estimate the temporal locations of each beat within the sequence. For each sequence a subsequence of length T was identified as the motif that is most representative of the multiple beat signals within the sequence. Temporal shifts of the motif within the sequence that best reconstructed the entire sequence were used to estimate beat delay and consequently estimate heart rate variance.

A subsequence $S^T_{t_0}(t)$ of sequence $S(t)$ is uniquely defined by length T and starting time t_0 such that $S^T_{t_0}(t) = S(t + t_0)$. We zero pad the sequence to length L such that $S^T_{t_0}(t) = 0$ for t > T. The sequence motif is the subsequence that captures more of the periodic structure in the signal than any other subsequence. To quantify this we calculate the

cross-correlation $C(\tau)$ of $S_{t_0}^T$ with S:

$$C(\tau) = \mathcal{F}^{-1}\left\{\mathcal{F}\left\{S_{t_0}^T(t)\right\} \cdot \mathcal{F}^*\left\{S(t)\right\}\right\} \quad (2)$$

Where $\mathcal{F}$ is the Fourier operator and $*$ the complex conjugate. The cross-correlation requires normalization to eliminate bias towards selecting motifs in the high amplitude regions of the signal. Both the power in the sub-sequence and the distribution of power throughout the sequence contribute to this bias. A normalization factor is calculated from the cross correlation of instantaneous signal power S^2 with a zero padded step function H of duration T:

$$N(\tau) = \mathcal{F}^{-1}\left\{\mathcal{F}\left\{S^2(t)\right\} \cdot \mathcal{F}^*\{H(t)\}\right\}$$

Where H(t) = 1 for t≤T and H(t)=0 for t>T. Consequently $N(t_0)$ is the power in the subsequence and $N(\tau)$ the power in the target sequence starting at time τ of length T.

A subsequence $S_{t_0}^T$ can repeat L/T times throughout the sequence of length L. We evaluate the peaks of the cross-correlation at temporal lag steps of size T corresponding to $\tau_i = iT + t_o \ modulo \ T$, where $i \in \{1, floor(L/T)\}$. To calculate heart rate variation we allow for period drift within a sequence and therefore look for maximum values within a range $\{\tau_i - kT, \tau_i + kT\}$ where k depends on the expected maximum heart rate variation within a sequence. Finally given the estimated heart rate L the sequence motif is defined by delay t_m which corresponds to the subsequence delay t_0 which maximizes the following expression:

$$t_m = argmax_{t_0}\left\{\sum_{i=1}^{floor\left(\frac{L}{T}\right)} \max_{\tau\in\{\tau_i-KT,\tau_i+KT\}} \left[\frac{C(\tau)}{N(t_0)\,N(\tau)}\right]\right\} \quad (3)$$

The expected maximum heart rate variation within a sequence was set to $k = 0.1$. After detecting the best motif the temporal shifts of the motif τ that resulted in maximizing expression (3) are used for estimating beat to beat variation estimation within a sequence.

III. Results

The first 2 months of available data was assessed. Of the 9000 sequences, approximately 14% were saturated, 61% periodic, and the remainder aperiodic. Of the approaches the ACM generated the best classification of the aperiodic/periodic signals achieving 89% accuracy with a detection threshold of 1/3 with respect to the 200 sequences for which periodic/aperiodic ground truth was available. Therefore ACM was used for the classification stage in Fig. 4.

In terms of the saturated signals, a histogram of saturated signals against time of day is shown in Fig. 7 and indicates this detection approach is heavily influenced by the presence of sunlight. An explanation for the generation of saturated sequences (Fig 2a) in the presence of sunlight is that the IR sensor is detecting bursts of sunlight. The presence of waves on the water surface directly above the oyster results in an instantaneous distribution of sunlight below the water surface. This distribution contains regions of increased sunlight due to a focusing effect via refraction through the air/water interface. Wave motion then results in temporal variation of sunlight incident onto the oyster and consequently coupled into the detector.

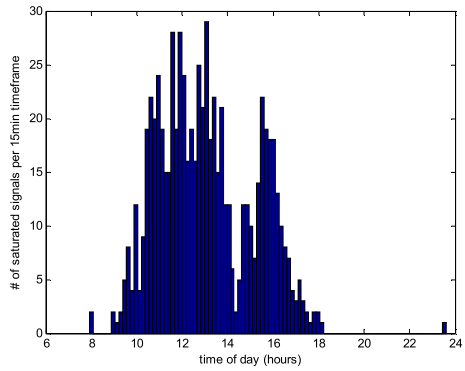

Fig. 7. Histogram of saturated sequence class against time of day.

TABLE I
Accuracy Results for Implemented Methods. *TR Uses ACM, g = 0.6, N_{TR} Corresponds to 4.5 Hours. **MF Uses ACM, N_{MF}^1 and N_{MF}^2 Correspond to 4.5 Hours and 1 Hour, Respectively

Approach		HR estimates % accurate	Heart beat period estimates % split into low (L), correct (C), and high (H) ranges		
			L	C	H
Lag	ACM	83.2	0.3	89.7	10.0
	CCA	82.7	0.1	88.8	11.1
	NAC	11.6	2.6	11.9	85.5
	HWW	71.5	1.9	77.9	20.1
Time	AMD	13.5	0.0	14.0	86.0
	MLH	7.0	0.0	7.0	93.0
Frequency	CEP	33.0	3.7	35.7	60.6
Aggregated	TR*	85.2	0.3	91.3	8.4
	MF**	93.7	0.0	99.9	0.1
Machine Learning	MLP	96.3	0.7	99.2	0.1
	LR	96.7	0.7	99.2	0.1
	SMO	96.8	0.7	99.2	0.1

The ultimate aim of heart rate estimation is to find associations with physiological response of the oysters. For analysis by biological scientists a 90% accuracy was targeted. Furthermore due to the interpolated nature of the generated ground truth values, the ground truths were estimated to be within 4% of the actual periods. This was based on comparing accuracy from a coarser interpolation. Consequently a heart rate estimate was deemed accurate if it was within 10% of the manually acquired ground truth. The percentages of heart rates estimated accurately for each approach are shown in Table I. The most accurate methods are the machine learning approaches followed by the aggregated approaches. The ACM achieves the highest accuracy of the single (lag, time and frequency) estimators (83.2%). Multiple ACM heart rate estimates were then aggregated to calculate

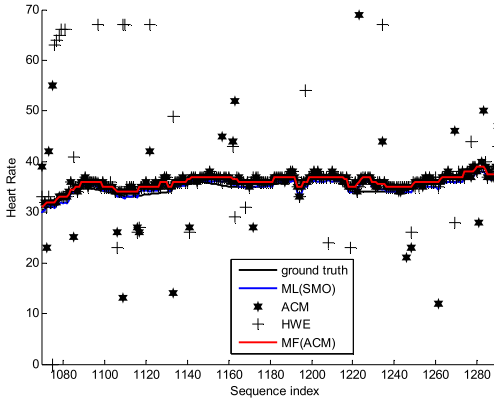

Fig. 8. Plot of estimated heart rates for a selection of approaches.

the Tracking and Median Filtered estimators. Optimal values of N_{TR} corresponded to time frames of 4-5 hours with optimal $g = 0.6$ giving an accuracy of 85.2%. The median filter approach achieved 93.7% with values of N_{TR} and N_{TR} corresponding to 4.5 hours and 1 hour respectively. Note these optimal values are sensitive to the level of heart rate variation over time frames corresponding to hours. Values of N_{TR} and N_{TR} corresponding to 2 hours and 20 minutes respectively still achieved 92% accuracy, but may be more accurate for estimating heart rates where variation is more rapid than the current installation. Finally the Machine learning approach achieved superior results, the best approach (based on a support vector regression) achieving 96.8% accuracy.

The accuracy of any approach is dependent on two factors: the ability to estimate the period within the sequence accurately, and the ability to avoid erroneously selecting double and half periods. Each estimate of heartbeat period was compared to the ground truth. If the estimated period was within the range 2/3 to 3/2 multiples of the ground truth period the estimate was categorized as being in the correct range. For period estimates below and above this range, the estimate was categorized as being in the low or high range respectively. The percentages of estimated periods categorized into these ranges by each approach are shown in Table I. For lag based approaches the accuracy is strongly related to the erroneous selection of longer periods. The bias of ACM and CCA to lower periods via lag domain weighting resulting from sequence zero padding is evident, whereas the normalised approach which removes this weighting suffered the most. The two pulse structure of the heart rate signals has the potential to generate erroneous half period peaks in the lag domain. Despite the lower lag bias of the ACM and CCA the ACM and CCA were not susceptible to the half period error. The time domain methods exhibited similar performance to the NAC. The CEP had the highest of estimates in the low period range, as a frequency domain method this corresponds to the CEP selecting the second harmonic, resulting in half period estimates. The MF approach was highly successful at avoiding half and double period estimates; the averaging effect over multiple samples did not dramatically reduce accuracy. Information in the multiple estimators was utilized by the machine learning approach to improve the accuracy beyond the median filtered results.

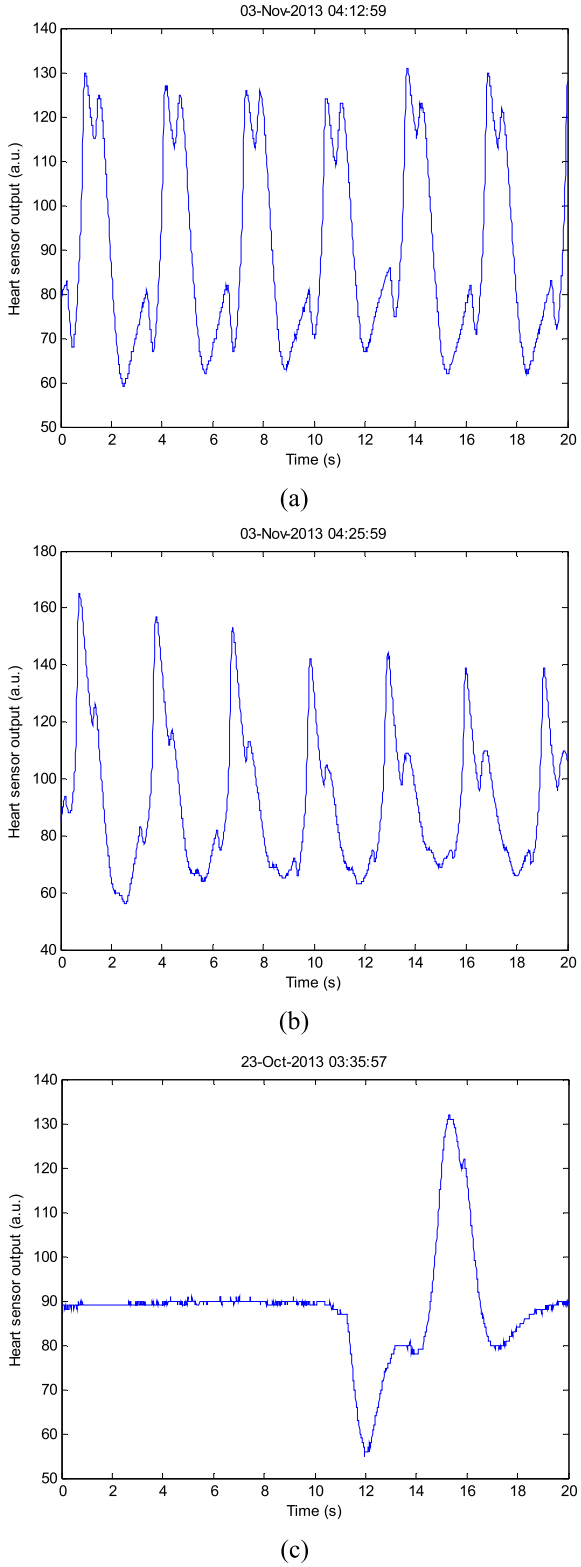

Fig. 9. (a)-(b) Two sequences of 20 s duration from early region of Fig. 3 timeframe at ∼4:13AM and 4:25AM (13 minutes apart). (c) Example sequence including single "beat."

The results show a wide variation in the ability of the methods to determine heart rate which is strongly influenced by the bias of the various methods towards detection of heart rates at double and half the correct heart rate. An example range of sequences is shown in Fig. 8 along with estimated

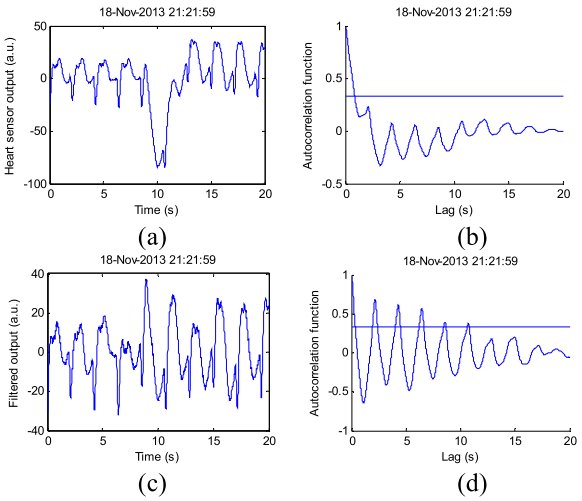

Fig. 10. (a) Original sequence (DC removed). (b) Lag domain of sequence. (c) 2 Hz high-pass filtered sequence. (d) Lag domain of filtered sequence.

heart rates; the ML and MF closely follow the ground truth, whereas ACM and HWE exhibit erroneous readings on some sequences.

For heart rate sensors deployed on cheap, lightweight hardware the MF(ACM) approach is ideal as the more accurate machine learning approach requires labelled training data and additional computational complexity. However; for offline analysis with labelled data the machine learning approach provides the most accurate results.

The signals exhibit a number of unique properties. For example the signals can switch between regions of periodic and aperiodic signals (observable in lag domain Fig. 5). Although the heart rate varies slowly the signal shape can vary rapidly. This is demonstrated by the sequences in Fig. 9a, b which are separated by approximately 13 minutes. These sequences are periodic with identical heart rates but exhibit significantly different shapes. Sudden signal deviations often occur either in isolation (Fig. 9c) or within a periodic sequence (Fig. 10a). Isolated deviations likely indicate the presence of a short term geometry variation caused either by a single heartbeat or the shell opening and closing. Higher frequency sampling of gape would identify the cause. Deviations within periodic sequences are often ∼2s in duration, with a corresponding change in the periodic structure either side of the deviation. An example of this is seen by comparing the signal waveform in Fig. 10a 0s-5s versus the 15s-20s signal waveform. We believe these dips occur when the oyster rapidly closes and opens the shell, thus changing the coupling geometry. Furthermore these dips cause problems for the ACM due to the first peak (2s in Fig. 10b) becoming obscured by the spread of the self coherence peak (0s in Fig. 10b). The spread is introduced as a consequence of the sequence mean value being significantly different to the periodic structure local mean value resulting in temporarily localized DC offsets which extend the coherence time. The two sequences recorded adjacent to Fig. 9a were correctly identified with 2s period. One solution is to high pass filter the sequence which removes the localized DC offset and exposes the lag peaks (Fig. 10d). However this lowers the sensitivity of the method to heart rate detection in other sequences. As these errors are present within single sequences, the median filtering approach eliminates these errors.

IV. Conclusion

We discussed a range of properties of signals generated by a novel oyster heart rate monitor and consequences on signal analysis. These properties include sudden changes in signal level, transition from periodic to aperiodic signals, large saturated signal interference, and evolution of the periodic structure within the sequences. We implemented a range of existing methods for estimating heart rate, along with some tailored approaches aggregating multiple sequences, and a machine learning approach. The approaches were all able to estimate heart rate. The lag domain approaches biased to shorter lags were the best single sequence simple estimators achieving 83.2% accuracy. However none of the existing methods achieved better than 90% accuracy. This accuracy required a machine learning approach which achieved the best results (96.8%) using all available estimators as input features. Alternatively an approach utilizing multiple consecutive sequences (the median filtered autocorrelation method) was required which achieved acceptable accuracy (93.7%). The ultimate choice is determined by available computational resources. Future work would require implementation of the algorithms within the sensor platform and require computationally efficient means in calculating heart rate variation.

Acknowledgment

The authors gratefully acknowledge B. Taylor from the University of Tasmania whose oyster tag design is the basis of the *in situ* oyster tag design used in this paper, and without which this work would not be possible.

References

[1] E. Braunwald, "Cardiology: The past, the present, and the future," *J. Amer. College Cardiol.*, vol. 42, no. 12, pp. 2031–2041, 2003.

[2] R. A. Bruce, J. A. Mazzarella, J. W. Jordan, Jr., and E. Green, "Quantitation of QRS and ST segment responses to exercise," *Amer. Heart J.*, vol. 71, no. 4, pp. 455–466, Apr. 1966.

[3] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. BME-32, no. 3, pp. 230–236, Mar. 1985.

[4] E. R. Trueman, "Activity and heart rate of bivalve molluscs in their natural habitat," *Nature*, vol. 214, pp. 832–833, May 1967.

[5] M. H. Depledge and B. B. Andersen, "A computer-aided physiological monitoring system for continuous, long-term recording of cardiac activity in selected invertebrates," *Comparative Biochem. Physiol. A, Physiol.*, vol. 96, no. 4, pp. 473–477, 1990.

[6] P. A. Ritto, J. G. Contreras, and J. J. Alvarado-Gil, "Monitoring of heartbeat by laser beam reflection," *Meas. Sci. Technol.*, vol. 14, no. 3, p. 317, 2003.

[7] N. P. Burnett *et al.*, "An improved noninvasive method for measuring heartbeat of intertidal animals," *Limnol. Oceanography, Methods*, vol. 11, no. 2, pp. 91–100, Feb. 2013.

[8] CSIRO. *Biosensors for Oysters*. [Online]. Available: http://www.csiro.au/Organisation-Structure/Flagships/Food-Futures-Flagship/Breed-Engineering-Theme/Aquaculture-Biosensors-Oysters.aspx, accessed Jan. 2014.

[9] K. H. Park, Y.-S. Kim, E.-Y. Chung, S.-N. Choe, and J.-J. Choo, "Cardiac responses of Pacific oyster crassostrea gigas to agents modulating cholinergic function," *Comparative Biochem. Physiol. C*, vol. 139, no. 4, pp. 303–308, 2004.

[10] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. Englewood Cliffs, NJ, USA: Prentice-Hall, 2009
[11] M. G. Elfeky, W. G. Aref, and A. K. Elmagarmid, "Using convolution to mine obscure periodic patterns in one pass," in *Proc. EDBT*, 2004, pp. 605–620.
[12] P. Boersma, "Accurate short-term analysis of the fundamental frequency and the harmonics-to-noise ratio of a sampled sound," in *Proc. IFA*, vol. 17. 1993 pp. 97–110.
[13] J. J. Dubnowski, R. W. Schafer, and L. Rabiner, "Real-time digital hardware pitch detector," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-24, no. 1, pp. 2–8, Feb. 1976.
[14] D. Talkin, "A robust algorithm for pitch tracking (RAPT)," in *Speech Coding and Synthesis*. Amsterdam, The Netherlands: Elsevier, 1995, ch. 14, pp. 495–518,
[15] M. J. Ross, H. Shaffer, A. Cohen, R. Freudberg, and H. J. Manley, "Average magnitude difference function pitch extractor," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 22, no. 5, pp. 353–362, Oct. 1974.
[16] A. M. Noll, "Pitch determination of human speech by the harmonic product spectrum, the harmonic sum spectrum and a maximum likelihood estimate," in *Proc. Symp. Comput. Process. Commun.*, vol. 19, 1970, pp. 779–797.
[17] A. M. Noll, "Cepstrum pitch determination," *Acoust. Soc. Amer.*, vol. 41, no. 2, pp. 293–309, Feb. 1967.
[18] A. D. Hellicar, A. Rahman, D. Smith, G. Smith, and J. McCulloch, "A neural network and SOM based approach to analyse periodic signals: Application to oyster heart-rate data," in *Proc. IJCNN*, Beijing, China, Jul. 2014, pp. 2211–2217.
[19] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The WEKA data mining software: An update," *SIGKDD Explorations*, vol. 11, no. 1, pp. 10–18, Jun. 2009.

Andrew D. Hellicar received the bachelor's (Hons.) degree in engineering from the University of Tasmania, in 1997, and the Ph.D. degree from Monash University, in 2007. In 2012, he returned to Tasmania, where he conducts research on machine learning techniques applied to the agriculture domain with the ICT Centre, Commonwealth Scientific and Industrial Research Organisation (CSIRO). Since 2002, he has been with the Radio Physics Laboratory in Computational Electromagnetics and Distributed Computing, CSIRO, before leading projects in terahertz imaging and terrestrial communication systems and undertaking research on antennas for radio astronomy and satellite communication.

Ashfaqur Rahman (SM'12) received the Ph.D. degree in information technology from Monash University, Gippsland, VIC, Australia. He worked on specific machine learning problems, including ensemble learning and fusion, feature selection/weighting methods, genetic algorithm-based optimization, and image segmentation and classification. He is currently a Machine Learning Researcher for more than ten years. He is a Research Scientist with Commonwealth Scientific and Industrial Research Organisation (CSIRO), Hobart, TAS, Australia. He is the Leader of the Computational Intelligence Team with CSIRO. He has authored around 70 peer-reviewed journal articles, book chapters, and conference papers. He serves as a Reviewer of prestigious conferences and journals. He was the Program Committee Chair of the 2013 International Conference on Digital Image Computing: Techniques and Applications.

Daniel V. Smith received the Ph.D. degree in telecommunication engineering from the University of Wollongong, in 2007. From 2008 to 2009, he was a Post-Doctoral Fellow with the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Hobart, Australia. He was then appointed as a Research Scientist with the Intelligent Sensing and Systems Laboratory, CSIRO, where he currently works. The focus of his work has been upon the research and development of analytical models for environmental and agricultural applications using sensors and sensor networks. His research interests include adaptive filtering, time series modeling, and data uncertainty representations.

Greg Smith received the B.E. (Hons.) degree and the M.Sc. degree in computer science from the University of Technology, Sydney, in 1987 and 1998, respectively, and the Ph.D. degree from the University of Sydney, in 2004. He is currently with the Commonwealth Scientific and Industrial Research Organisation, Hobart, TAS. His current interests are in data mining, software for big data analytics, and agent-based modeling.

John McCulloch received the Degree from the University of Tasmania (UTas), in 1993. He was a Computer Systems Engineer with UTas. He lectured on experimental design and mechatronics before joining Commonwealth Scientific and Industrial Research Organisation (CSIRO), where he has been managing biosensor and decision support projects in the marine and pond aquaculture space. He is currently a Research Engineer with the Autonomous Systems Program, Digital Productivity Flagship, CSIRO. His example projects are a data management and visualization system for the Tasmanian Shellfish Quality Assurance Program, and oyster biosensor development such that real time physiological data (heart rate and gape) is now being gathered from animals on farms in Southern Tasmania.

Sarah Andrewartha received the Ph.D. (Hons.) degree from Melbourne's La Trobe University. She focused on understanding how the environment influenced the development and physiology of a broad range of animals. She was a Post-Doctoral Fellow with the University of North Texas, where she was involved in research on the development of acid-base regulation in chicken embryos, and epigenetic inheritance patterns in response to changes in the environment of Daphnia (tiny freshwater crustaceans with a transparent body, often used in biological research). She is currently a Comparative Physiologist and Post-Doctoral Fellow with the Commonwealth Scientific and Industrial Research Organisation, Hobart, Australia. She uses biosensors to monitor the effect of environmental changes on shellfish and fish physiology. In the future, this information will be used by aquaculture farmers to assist farm management, based upon measurements of animal fitness. She is examining body temperature influences on metabolism, ventilation, and blood acid-base regulation in developing Tamar wallaby joeys, and the consequences of the nest environment for young reptiles and amphibians.

Andrea Morash received the B.Sc. degree in biology from Mount Allison University, NB, Canada, in 2005, and the Ph.D. degree in biology from McMaster University, ON, Canada, in 2010. From 2010 to 2012, she completed her first post-doctoral fellowship with the University of Cambridge funded by an NSERC Post-Doctoral Fellowship Award. She joined the University of Tasmania and the Commonwealth Scientific and Industrial Research Organisation as a Post-Doctoral Researcher in 2013. Her research interests focus on the effects of environmental stress on animal energy production and utilization and how this impacts their physiology and ecology.