

Estimating Heart Rate and Detecting Feeding Events of Fish Using an Implantable Biologger

Yiran Shen

Data61, CSIRO, Australia
yiran.shen@csiro.au

Reza Arablouei

Data61, CSIRO, Australia
reza.arablouei@csiro.au

Frank de Hoog

Data61, CSIRO, Australia
frank.dehoog@csiro.au

Jaques Malan

Oceans & Atmosphere, CSIRO,
Australia
jaques.malan@csiro.au

James Sharp

Oceans & Atmosphere, CSIRO,
Australia
james.sharp@csiro.au

Sara Shouri

Data61, CSIRO, Australia
sara.shouri@yahoo.com

Timothy D. Clark

School of Life and Environmental
Sciences, Deakin University, Australia
t.clark@deakin.edu.au

Carine Lefevre

Australian Institute of Marine Science
c.lefevre@aims.gov.au

Frederieke Kroon

Australian Institute of Marine Science
f.kroon@aims.gov.au

Andrea Severati

Australian Institute of Marine Science
a.severati@aims.gov.au

Brano Kusy

Data61, CSIRO, Australia
brano.kusy@csiro.au

ABSTRACT

Monitoring of physiology and behavior of marine animals living undisturbed in their natural habitats can provide valuable data on their well-being and response to environmental stressors. We focus on detection of feeding of predatory fish using implantable biologgers that record electrocardiogram (ECG) signals. We propose a novel processing pipeline for resource-constrained embedded systems that can infer higher-level information, such as heart-rate and feeding events, from the ECG signals. Our main contribution is a lightweight change-detection algorithm, that can reliably detect fish feeding in noisy heart-rate data based on unique statistical properties of feeding-induced changes in heart-rate. We evaluate our approach using an in-house biologger that we surgically implant in twelve coral trouts over a period of ten weeks. We show that our signal processing pipeline performs well with noisy ECG signals overall. Specifically, our heart-rate estimation algorithm achieves errors of less than one beat per minute even in scenarios where popular algorithms used by domain scientists perform poorly. Furthermore, our feeding detection algorithm achieves good accuracy and matches the performance of state-of-the-art algorithms while requiring significantly less memory and computational resources. This work is an important first step towards long-term monitoring of high-level condition and health of marine animals in the wild.

KEYWORDS

Animal energetics, biosensor, change detection, heart-rate estimation, resource-constrained system.

1 INTRODUCTION

Industrialization and economic growth have come at the expense of negative impacts on many natural ecosystems. There is a global consensus that a better stewardship of our planet is imperative [20, 35].

However, despite a substantial progress in environmental management, recent news is not good: Australia's Great Barrier Reef has been severely impacted by warm temperatures in an unprecedented scale leading to up to 30% die-off of coral in 2016 [23]; mass extinction of reptiles and amphibians has been observed in Latin America [2]; and the average rate of extinction in vertebrate species over the last century has been estimated to be 100 times higher than the background rate [9].

Environmental management plans and policies are only as good as the quality of the models that they rely on to quantify the effects of various stressors on the environment. Advances in remote and in-situ sensing technologies have improved our capability to collect detailed geomorphological, geophysical, and geochemical data across vast areas. However, many natural habitats harbor intricate ecological interactions between thousands of animal and plant species. This complexity makes it difficult to estimate the accurate physiological state of any ecosystem or to predict undesirable and potentially irreversible changes [6]. It is important to bridge this gap and develop effective technologies for monitoring physiology and behavior of animals in coastal marine ecosystems.

In this paper, we present a signal processing pipeline for reliable detection of fish feeding using heart-rate data (see Figure 1). Feeding is an important component of fish energetics, i.e., the balance sheet of energy intake against energy expenditure. Energetics relates to important health and body condition indicators, including metabolic factors, physical exertions, tissue synthesis, reproductive success, and stress. Frequency and quantity of feeding can also provide valuable insights into food chain composition and predation rates.

On the technical level, we assume that a fish is implanted with a tiny biologger that collects and processes ECG data, and communicates the results back to the user to achieve long-term monitoring.

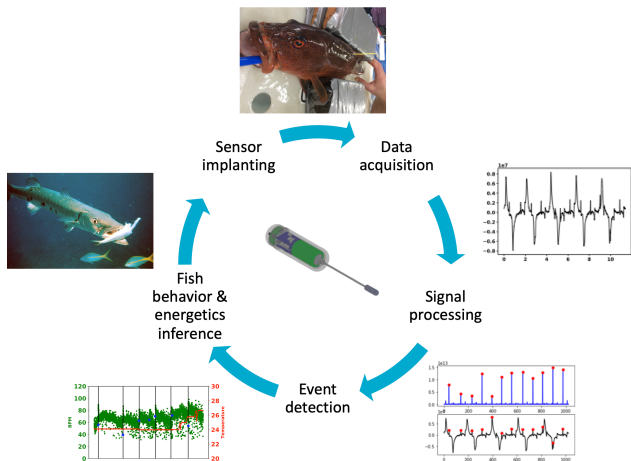


Figure 1: Use of implantable biosensor for monitoring fish energetics in marine ecosystems.

The question is, can we use the ECG data from ectothermic¹ fish to estimate their feeding? It has been observed that the heart rate of fish is elevated while feeding. The heart rate decreases gradually after the feeding over a period whose length depends on the amount of the feed [11]. This can be used as an indicator to estimate the energy intake of fish. However, heart rate of fish can vary due to several factors such as muscle activity, digestion, ambient temperature, and a range of stressors. Our goal is to develop an algorithm that reliably detects feeding events from heart-rate time-series data extracted from noisy ECG signals. A key fact is that food digestion in fish increases the metabolic rate that can be observed as increased heart-rate. The digestion typically takes longer than a typical fight or flight response, so careful statistical evaluation can detect feeding events even if the underlying signal is noisy.

We propose signal processing and analytics algorithms that process noisy ECG data to extract heart-rate and identify feeding events. The noise in the ECG data can be significant due to multiple factors. The main noise source is the activity of other muscles. Fish motion can generate electric signatures or physical movement of the ECG electrodes and introduce peaks that resemble heart beats. Some erroneous peaks may be identified and rejected using accelerometry data. However, other physiological activities, such as respiration, can be more subtle and may not have a strong signature in accelerometry data despite distorting the ECG data severely. We propose several strategies to address the error sources including heartbeat peak augmentation, robust peak detection, and a set of peak validation criteria applied as post-filtering. The preprocessing steps allow us to obtain good quality heart-rate data from noisy ECG signals.

Preprocessed heart-rate data can still combine several signals induced by different physiological activities. Our key contribution is a lightweight Bayesian algorithm capable of recognizing feeding-related signals despite the presence of other physiological signals

¹Ectothermic animals regulate their body temperature by relying on external heat sources from the environment.

and noise. The algorithm works by detecting a change in the underlying statistical properties of heart-rate that is consistent with fish feeding. Recall that food digestion has a signature that is statistically different from other activities. Specifically, a feeding event causes a rapid increase in the heart-rate that is sustained over a substantial period of time. The algorithm looks for significant changes in the statistics of the heart-rate series over a short period of time. The rate and duration of the observed statistical changes are the key parameters that allow us to distinguish fish feeding from other activities.

We pay special attention to computation and memory resources of the proposed algorithms as our goal is to run the algorithms directly on embedded micro-controllers. This will enable long-term operation of biologgers under strict energy and communication budgets that are typical for wild-life monitoring.

We conducted a controlled experiment to collect data from live fish. We used commercial and in-house biologgers that collect electrocardiogram (ECG), tri-axial acceleration, ambient temperature, and atmospheric pressure/depth data. While the loggers collect data related to both energy intake and expenditure, we focus our analysis on ECG data only. We implanted biologgers in coral trouts in a facility with six fish tanks and conducted feeding trials over a period of ten weeks. Water temperature in the tanks was configured to simulate daily cycles during different weather seasons. During the experiment, we fed the fish with several pieces of food and manually recorded the feeding event data including date, time, and size of the feed. At the end of the experiment, the data was retrieved from the biologgers by removing them from the fish as the devices do not currently have any through-water communication capability.

Our analysis of the ECG data in this paper can be divided into two parts. First, we estimate heart-rate from noisy ECG signals. We have asked a domain expert to manually label ECG peaks in a subset of the fish data, so that we can quantify performance of our heart-rate estimation algorithm. The algorithm achieves high performance with an average error of 0.14 beats per minute (bpm), substantially outperforming three existing heart-rate estimation algorithms in scenarios where significant noise is present.

Second, we evaluate our Bayesian change-detection algorithm in simulation and on the empirical data traces from the fish. We pose the feeding detection as a classification problem, where the goal is to determine whether a fish feeds during a 24-hour period. Our algorithm achieves an equal-error rate² of 0.15 on the empirical data. The algorithm performs similarly to the state-of-the-art algorithms for change detection. However, it consumes less data and requires considerably less memory and computations facilitating its operation on embedded micro-controllers.

2 BACKGROUND AND RELATED WORK

The area of implantable and wearable sensors for free-living animals has seen a rapid progress in the last two decades driven mainly by the advances in miniaturization and bio-compatibility of electronics [32]. Sensors have been deployed to measure motion, physiological, behavioral, and environmental animal data. They

²The error rate that balances false-positive and false-negative error rates.

have contributed to improved understanding of migratory pathways, foraging patterns, habitat use, and responses of animals to environmental changes.

In this section, we review the relevant literature and discuss key motivations and developments in biosensor research.

2.1 Sensing Animal Physiology

Physiological sensors have revolutionized our understanding of animal species in the wild [37]. Animals can react to human presence through fight or flight responses and may trigger release of stress hormones leading to increased heart-rate, respiration, and energy availability. The long-term effects of stress due to frequent disturbance has been shown to result in lower reproductive success and higher disease susceptibility [18]. Heart-rate biosensors have been specifically used in a wide range of studies measuring the energetic response of animals to human activities such as aircraft noise [7], ecotourism at breeding sites [18], and UAV noise [17]. It has been shown that physiological responses of animals may not always translate into observable behavioral responses. Thus, it is important to consider the implications of long-term stress on animal well-being even when short-term negative effects are not conspicuous.

2.2 Bioenergetics

Organisms and their cells must perform work to stay alive, grow, and reproduce using energy harvested from various sources. Bioenergetics studies the energy flows in living cells of individual animals [27]. The underlying concept of energetics is relatively simple. Organisms consume the chemical energy stored in food to synthesize complex molecules and generate motion, heat, and metabolism byproducts. Energetic status of an animal has implications on both individual animals as well as population-level processes. Bioenergetics has been used to explain animal behavior, selection of migration routes, habitat selection, and behavioral response to environmental changes [16, 24]. Characterization of feed intake that we study in this paper is an important component of bioenergetics models.

2.3 Fish Biosensors

A range of commercial sensors (e.g., Star-Oddi [26]) and research prototypes have been used to collect information about fish behavior in general and fish energetics in particular [16, 32]. Our focus here is on the detection and characterization of fish feeding.

2.3.1 Movement. Accelerometer and electromyogram (EMG) sensors can provide data on fish activity levels and rates of movement. They can be calibrated to provide information about energy expenditure or to classify specific fish behavior [13, 22, 38]. While accelerometer data can be used to classify fish feeding, the method suffers from false positives as it is difficult to distinguish successful and failed predation events, using accelerometer data alone.

2.3.2 Temperature. While primarily used to provide insights into habitat use [34], implanted temperature sensors can yield information about feeding for species of fish that are not completely ectothermic. For example, some tuna fish and lamnid sharks possess vascular heat exchangers that function to retain metabolically-derived heat in specific regions of the body. This enables some

species to exhibit a heat increment following feeding. This so-called “heat increment of feeding” has been used to quantify individual meal sizes [10, 11, 36]. However, temperature-based feeding detection only applies to a handful of fish species.

2.3.3 ECG/Heart-Rate. For most fish, the heat increment of feeding is not visible because they rapidly lose any metabolically-derived heat across the gills and surface of the body. Nevertheless, meal sizes can still be estimated in these species by exploiting the link between digestion and increased blood flow. The increase in blood flow during the digestive period is typically mediated by an increase in heart-rate [4, 13, 14]. Several implantable sensors have been proposed in the literature and are even commercially available from Star-Oddi [34]. Our algorithms are independent of the underlying hardware. However, it is worth noting that unique design of our in-house heart-rate logger allows it to minimize fish discomfort while recording ECG data at a superior quality, when compared to commercial sensors.

2.4 ECG Algorithms

2.4.1 Heart-rate. Heart-rate can be obtained by analyzing frequency components of ECG time-series data. Common preprocessing steps include band-pass and low-pass filters to remove noise outside of the target frequency range [21]. As the heart-rate frequency may vary for different species, preprocessing algorithms require careful tuning. Popular techniques for ECG data analysis include frequency domain transformations such as Fourier [15] and wavelet transformation [33]. The final step is to apply peak detectors, to either find the dominant peak in the frequency domain, or to directly find peak-to-peak distances within time domain. Post-processing can further improve heart-rate estimation accuracy, by enforcing custom parameters, e.g., minimum peak-to-peak distance, or peak amplitude thresholds. Peak-to-peak distance is trivially related to heart-rate through the known sampling rate.

2.4.2 Feeding Detection. There exists limited prior work on detection of fish feeding in the lab or in the wild. Vision-based approach [31] works well when the water is clear and fish is within the camera field of view. However, its usability is limited by the sensing coverage. Acceleration-based detection [5] using data from implantable biologgers solves the coverage problem, but suffers from false detections caused by unsuccessful prey motions. In this paper, we rely on the heart-rate time-series extracted from ECG signals. We reformulate the feeding detection problem as the change detection problem which has been well-studied in the literature [3]. For example, Bayesian inference algorithms for change detection have been proposed, capable of operating both in offline [19] and online [1] modes. These algorithms estimate the probability of a significant change of statistical properties of a time-series signal and label each such occasion as the change point. Our contribution extends this work and proposes a lightweight change detection algorithm that has the potential to be executed in real time on resource-constraint micro-processors.

3 FISH FEEDING DETECTION AND CHARACTERIZATION

In this section, we present the algorithms for estimation of fish heart-rate from noisy ECG data and detection of feeding events.

3.1 Processing Pipeline Overview

We summarize our processing pipeline in Figure 2. The process starts with signal acquisition by biollogger that lasts for the duration of the experiment. Data is then downloaded to the computer and preprocessed to filter out errors. This process generates *samples* defined as 11.38-second-long snapshots of ECG, accelerometer, and other sensor data. Samples are separated by a period of 3.5 minutes during which no data is recorded. We apply our heart-rate estimation algorithm to each ECG sample to compute the corresponding median heart-rate. A post-filtering step discards low-quality samples and interpolates data to infer the missing and discarded values. Finally, we propose a new lightweight change detection algorithm to determine feeding events based on the heart-rate values.

3.2 Data Acquisition and Preprocessing

We briefly describe the data acquisition process (more details are in Section 4.1.2). The biollogger periodically records 1024 values of ECG in on-board storage, sampled at 90Hz. Therefore, it records a 11.38-second snapshot of ECG and acceleration data, i.e., a *sample*. Due to memory and energy constraints, the logger then turns off for 3.5 minutes.

The aim of the preprocessing step is to clean the ECG data and convert it to a format suitable for heart-rate estimation. We split the entire data collected by the biollogger into individual 11.38-second-long samples and discard samples that have missing or corrupt values.

3.3 Heart-Rate Estimation

Heart-rate estimation from ECG time-series data is an important step towards detection of feeding events. We show an example of ECG data in Figure 3. Note that the ECG data is noisy and affected by other physiological phenomena such as respiration and body movement.

3.3.1 Sources of error. We first characterize different sources of error in ECG time-series data. Figure 3 clearly shows that other factors are at play that introduce significant periodic signals in our data. We asked a fish physiologist in our team to mark the true ECG peaks corresponding to heart beats and correlate the other spurious ECG peaks with underwater video footage of the fish. We confirmed that the slower fluctuations in the ECG signal (cf. Figure 3) correspond to the movements due to respiration. Similarly, other abrupt body movements, such as rubbing or fighting, introduce strong irregular peaks in the signal. The key observation is that despite having smaller amplitude, heartbeat related ECG peaks are sharper compared to other peaks, indicating they have a unique signature in the frequency domain.

3.3.2 Heartbeat peak augmentation. We propose the following light-weight algorithm to identify heartbeat peaks. Assume we

are working with the l th sample denoted by

$$X^l = \{x_1^l, x_2^l, x_3^l, \dots, x_n^l\}$$

where n equals 1024 in our case. We first calculate the second-order central differences of the values in the sample as

$$x_{g,i}^l = \frac{x_{i-2}^l - 2x_i^l + x_{i+2}^l}{4}. \quad (1)$$

For the boundary points, the difference between the first (last) and second (second last) values are computed. We discard $x_{g,i}^l$ with positive values and square the rest to obtain the set of values

$$X_s^l = \left\{ \left(x_{g,i}^l \right)^2 \text{ if } x_{g,i}^l < 0 \text{ for } i = 1, \dots, n \right\}. \quad (2)$$

This procedure augments the peaks associated with heartbeats, see Figure 3 for an example. The figure demonstrates that smaller but narrower peaks induced by heartbeats are augmented and the wider peaks are suppressed.

3.3.3 Instantaneous heart-rate calculation. After the augmentation step, we apply a peak detection algorithm. A peak detector finds the local maxima of X_s^l that satisfy certain fish biology and signal amplitude conditions, i.e., 1) the distance between any two consecutive peaks should be at least 45, corresponding to the maximum instantaneous heart-rate of 120 beats per minute (bpm) and 2) we only consider peaks whose amplitudes are in top-50 per sample.

Having detected the peaks, we can then use the inverse of the intervals between consecutive peaks as estimates of the instantaneous heart-rate. Specifically, assume that the peak detector finds the location of m peaks, i.e., indexes of m peaks in the sample, as

$$\{l_1^l, l_2^l, l_3^l, \dots, l_m^l\}.$$

The instantaneous heart-rate in bpm between peaks at l_j^l and l_{j+1}^l is calculated as

$$h_j^l = \frac{60}{(l_{j+1}^l - l_j^l) / 90} \quad (3)$$

where 60 is the number of seconds in a minute and 90 is the sampling rate of the ECG signal.

3.3.4 Median heart-rate estimation. We use the median value of the instantaneous heart-rates h_j^l to find the sample heart-rate h^l .

As noted previously, ECG signals are noisy and our filtering algorithm cannot resolve all problems, i.e., it does not remove all non-heart-related ECG peaks and it cannot reconstruct missing heartbeat peaks. We show two typical examples of erroneous and missing heartbeat peaks in Figure 4. Our algorithm has missed a peak between peaks 7 and 8 in the top plot and misidentifies peak 6 as a heartbeat in the bottom plot. These errors can lead to inaccurate instantaneous heart-rate estimation and can impact accuracy of the median estimator, if a significant number of them are present in the sample.

Therefore, we introduce an additional criterion that enforces consistency of the instantaneous heart-rates within each sample. We opted against discarding the partially corrupt samples as it would reduce our dataset size. Instead, we propose a statistics-based




Figure 2: The processing pipeline for heart-rate estimation and feeding event detection. The first and second subfigures from the left present examples of the ECG samples and detected heartbeat peaks (marked in red), respectively. The series of detected heart-rate values of a fish over two weeks are shown in the third and fourth subfigures as green dots. The ground-truth and detected feeding events are marked with black lines and blue dots in the fourth subfigure.




Figure 3: Example of ECG time series (upper) and augmented heartbeat peaks (lower)




Figure 4: Examples of partly corrupt samples with missed or misidentified peaks.

approach to deal with such samples. Assume a set of instantaneous heart-rates denoted by

$$h^I = \{h_1^I, h_2^I, \dots, h_{m-1}^I\}$$

with mean $\mu(h^I)$ and standard deviation $\delta(h^I)$. We consider any sample h_j^I to be a dispensable outlier if it fulfills

$$h_j^I > \mu(h^I) + c\delta(h^I), \quad j = 1, \dots, m \quad (4)$$

where c is a tunable design parameter.

3.3.5 Post-filtering. Pre-filtering removes samples that are severely corrupt or heart-rate estimates that are inconsistent with other data points. We use two additional simple post-filtering criteria to improve the quality of data: we remove instantaneous heart-rate estimates lower than 20 bpm and higher than 120 bpm, which is the expected range for our target fish species. We also discard heart-rate estimates with less than three instantaneous samples. Finally, we apply cubic spline interpolation to fill small gaps in the heart-rate data.

The proportion of the discarded samples for 12 fish from our dataset are shown in Table 1. For majority of the fish, low percentage of data was discarded, indicating that the quality of our data collected in the experiment is generally good.

fish ID	Y1	G1	Y2	R2	Y3	R3
outliers	21.4%	1.42%	0.64%	3.13%	2.02%	0.73%
Fish ID	G4	G5	YB1	YB3	YR1	RG5
outliers	9.79%	8.14%	23.35%	0.74%	5.18%	5.27%

Table 1: Percentage of the data samples discarded as outliers for 12 fish.

3.4 Detecting Feeding Events

We first briefly discuss the key assumptions we make about the relationship between feeding and heart-rate in fish [11]. Fish have a natural baseline heart-rate that relates to the ambient temperature, genotype, health status, and other physiological factors. However, a fish can rapidly change its heart-rate in response to various stimuli including feeding, predation, reproduction, and disease. In particular, feeding has a distinguishing feature as it commonly results in a rapid increase in the heart-rate followed by a gradual decrease over a period of several hours (up to tens of hours depending on size of the feed). Figure 5 shows an example of heart-rate estimates with manually-labeled feeding events over a period of about two weeks.

We observe that each feeding event has a significant impact on the estimated heart-rate trace. This impact can be modeled as a change in the statistical properties of the heart-rate estimates and located using a change-point detection algorithm [3]. Change-point detection algorithms can be roughly categorized into offline and online approaches. Due to high computational complexity and latency, offline approaches are generally not suitable for implementation on resource-constrained devices or long-term deployments.




Figure 5: Median heart-rate values together with manually recorded feeding times.

3.4.1 Lightweight Bayesian Online Change-Point Detection. We propose an algorithm for detecting feeding events that is based on existing Bayesian online change detection algorithm [1]. We modify the algorithm to suit the resource constraints of our implantable biollogger. The proposed algorithm can process long-term heart-rate time-series data and relies on the fact that the heart-rate of fish remains elevated after feeding resulting in statistical properties that are distinct from the baseline.

Bayesian online change detection and growing resource consumption. The algorithm proposed in [1] estimates the probability of change for each point of a time series in a sequential manner using a Bayesian inference approach. Let the sequence of values in a time series be $\{x_1, x_2, x_3, \dots, x_t, x_{t+1}, \dots, x_T\}$ with M change points at indexes $\{c_1, c_2, \dots, c_m, c_{m+1}, \dots, c_M\}$, which separate the sequence into $M - 1$ segments. The Bayesian approach assumes the values within each segment are independent and identically distributed (i.i.d.) and are generated by a random process with probability density function (PDF) $p(x_t|\theta_\rho)$ where θ_ρ is a parameter for the ρ th segment. The underlying random processes of different segments are assumed to be *independent*. The Bayesian approach estimates the probability of change P^c for each data point based on the probabilities at the previous data points and the prior probability for the current point. Then, the change points can be determined by thresholding the probability of change.

The run-length variable r_t is defined for x_t to provide more robust estimation of the probability of change at time t . Intuitively, the run-length represents how many of the future points are needed to compute P_t^c at current time t and these future points are treated as being from the same distribution as the current point. The run-length process introduces system delay. For example, if $r_t = k$ and the time interval between two consecutive data points is d minutes, the inherent delay of the Bayesian-based online algorithm is kd

minutes. However, for monitoring fish feeding, e.g., in aquaculture, the primary concern is whether the fish is regularly fed and real-time response is not necessary.

Bayesian online change detection is a continuous chain-like processing algorithm. The probability P_t^c for current instance t is dependent on all the previous data points and their probability distributions. Therefore, the probabilities and parameters of the past distribution should be recorded. This means the required memory and computations increase over time and may exceed the capacity of our low-cost biollogger in a long-term deployment.

Chain-breaking points. The resources required for a long-term deployment are primarily determined by the length of the chain used in Bayesian inference, which utilizes all past information. We propose a lightweight version of Bayesian online change-point detection by inserting chain-breaking points at appropriate positions to restrict the memory usage while incurring minimal degradation in performance, i.e., the accuracy of estimating P_t^c .

We insert two types of chain-breaking points to break the chain and discard the obsolete prior information. The first is when the algorithm finds a point with high probability of change (P_t^c) at time t . As discussed above, the data points in each segment are assumed to be i.i.d. and the PDFs of different segments are independent. Therefore, the inaccuracies introduced by breaking the chain at the points with high P_t^c will be negligible. The second type of chain-breaking points corresponds to the case when the length of the chain exceeds a pre-defined length dictated by the resource constraints of the embedded system. However, we need to take care of the corner case when a new change point appears right after we break the old chain due to exceeding the maximum length. Specifically, we need to ensure that a sufficient number of points from the old distribution is recorded, for the purposes of the new change point calculation.

3.4.2 Resource Consumption Analysis. Our fish feeding detection algorithm is intended for in-situ processing on low-cost biosensors. Therefore, resource consumption is a key factor in the algorithm design. As mentioned before, the resource consumption of the original Bayesian online change-point detection algorithm [1] increases over time as all the previous information must be kept in the memory to infer the probability of change for the latest point. Hence, the growing memory requirement limits its practical usage on resource-constrained devices or long-term deployments. Our extension of the algorithm with the chain-breaking mechanisms limits the maximum memory and computations required by the algorithm. We introduce parameter L , the maximum length of the data-point sequence stored in the memory, to control the resource usage of the change-point detection algorithm.

Let the maximum length be L and the we need maintain a generalized Student's t -distribution for Bayesian-based inference. The algorithm keeps five variables for each data point: one variable to hold the probability value and four variables to update the parameters of generalized Student's t -distribution [25]. Therefore, the maximum number of parameters needed to be kept in memory for our proposed algorithm is $5L$ as compared to the original algorithm [1] that stores $5N$ parameters for a data trace of length N . As the computation of the probability of change is proportional to the number of previous points considered, our lightweight approach




Figure 6: Our implantable biosensor compared with coin (a) and its inner layout (b).

saves significant amounts of computation when N is large. For example, when $N = 10,000$ and $L = 1000$, the proposed algorithm performed around eight times faster than the original algorithm on our dataset.

4 EXPERIMENTS

We evaluate the heart-rate estimation and feeding detection algorithms in experiments with live fish using in-house and commercial Star-Oddi biologgers in the National Sea Simulator (SeaSim). SeaSim is a premiere marine research facility located at the Australian Institute for Marine Sciences (AIMS) in Townsville, Australia. The facility provides laboratory space that allows researchers to verify complex scientific hypothesis in long-term experiments. We first present the in-house biologger in more detail and then outline the experimental design.

4.1 Biologger Design

It has been shown, that externally attached devices have negative impacts on the swimming performance of fish, compromising both well-being of the animal as well as the quality of the data [8]. Therefore, we opted for an implantable sensor design.

While commercial sensors exist, we designed our own biologger as a platform for running proprietary signal processing and analytics algorithms as well as to experiment with the physical sensor dimensions to enable minimally-invasive long-term data collection from live fish. Specifically, we positioned one of the electrodes at the end of a flexible shielded wire (see Figure 6), rather than on the main sensor body as is typical in commercial sensors. This allows us to position the electrode close to fish heart whereas the main biosensor body can be positioned in the fish’s body cavity to reduce its adverse impact on the fish. As we will show later in the paper, this novel design allows our sensor to collect higher quality ECG data compared to commercial single-body sensors.

4.1.1 Hardware Design. The key objective for our hardware design was small physical size. We use a half-AA-size Li-Po primary cell with capacity of 1.2 Ah at 3.6V. The battery can be replaced when depleted. The enclosure is made of 3D-printed acrylic with a stainless-steel end cap on one side that acts as an ECG electrode and

Sensor	Sample rate	Information
ECG	90 Hz	Bioelectrical signals generated by heart muscle depolarizations.
Accelerometer	90 Hz	Acceleration along 3-axis, representative of activity levels and fish body orientation.
Temperature	one	Internal body temperature, representative of ambient temperature for ectotherms.
Pressure	one	Internal body pressure, representative of depth.

Table 2: Sensor components of our implantable biosensor.

a flexible shielded wire lead on the other side as the second ECG electrode. The spatial separation of the two electrodes is the key feature of our sensor design and helps to simultaneously improve signal fidelity and fish well-being as discussed above. The biologger can record high-frequency tri-axial acceleration, ECG, ambient temperature, and water pressure sensor data in an on-board memory (see Table 2). We did not include any communication interface on the logger and the data needs to be downloaded using wired connection after the experiment.

The biologger is controlled with a 32-bit Atmel AVR micro-controller to acquire data from all sensors. Hardware constraints are typical for the CPU family: the maximum speed of 84MHz, 128KB of SRAM, and 256KB of program memory. We use a 256 MB on-board NAND flash memory for data logging.

Our biologger currently does not feature any wireless underwater communication capability, although we have a miniature acoustic (sonar) communication module under development. Considering the cost and range of underwater sonar communication and the scarcity of available resources on the device, transmitting raw data wirelessly is impractical. However, by taking advantage of the lightweight signal-processing algorithms presented in the previous section, the device will require to transmit only higher level information, such as median heart-rates or detected feeding events, a few times per day allowing continuous operation of the sensor for months or years, depending on the application.

4.1.2 Biosignal Acquisition. The ECG sensor is an analogue component that translates bio-electrical signals generated by the heart to numerical values using a 24-bit analog-to-digital converter (ADC). The small potential difference between the two electrodes is amplified by 1700 times and a basic 1.6 Hz to 100 Hz bandpass filter is applied before digitization at 90 samples per second. The ADC offers further analog amplification if required.

In order to minimize drain on the battery, the biosensor is completely powered down except for the real-time clock (RTC). Upon RTC alarm, power is applied and the sensor records one sample sequence, before returning to power-off state. The sample sequence consists of a single measurement of pressure and temperature and continuous measurement of acceleration and ECG for a duration of about 11.4 seconds at 90 samples per second. This continuous sample period is long enough to capture a number of heart beats in order to calculate beats per minute as well as inter-beat variance.

The sampling interval can be configured to suit the deployment length and battery/memory capacity.

4.1.3 Time Synchronization. Time-synchronization of the data is achieved by using the on-board RTC. Synchronization proceeds in two steps. First, we synchronize the RTC chip with universal time coordinated (UTC) time when the sensor is prepared for deployment and record the time and date in a log file. Second, we synchronize the RTC with UTC time again when the sensor is retrieved and record the time and date in a log file. The difference between the elapsed on-board and UTC times allows us to calculate clock drift in the RTC. After estimating the clock drift, we retrospectively correct timestamps of all samples in our data trace. This simple clock offset compensation mechanism only works if the RTC clock drifts linearly over time. This assumption is not true in real-world conditions, as changes in temperature and battery voltage will cause the RTC clock to drift non-linearly. However, we have measured the overall clock drift to be on the order of approximately 10 seconds over a 10-week period. Hence the non-linear drift introduces negligible errors for the purposes of the analysis in this paper.

4.1.4 Power Consumption. Given the energy resources at our disposal are limited, low power consumption was a key objective in the sensor design. Our CPU had several low-power modes to reduce its idle power consumption. However, we opted for a configuration with a switch that cuts power to all circuits other than the RTC during the sleep mode. This helped us to achieve the current draw of $1.76 \mu A$ in sleep mode. In the active mode, we had reduced the clock frequency of CPU to 12 MHz which was sufficient for our computations, resulting in $6.24 mA$ current draw³. Our average power consumption during the experiment was approximately $70 \mu A$, due to the 3.5 minute sampling duty cycle. Given the capacity of our battery, we could have significantly increased the sampling frequency, however, we were limited to the 3.5 minute interval by our storage constraints.

4.2 National Sea Simulator (SeaSim)

The SeaSim facility enables researchers to examine the impact of complex environmental changes in tropical marine environments through large and long-term experiments. Specifically, scientists can manipulate key environmental factors such as light, temperature, acidity/pCO₂, salinity, sedimentation and contaminants using fine-scale control reflecting current and projected future environmental conditions. The facility provides full service that includes design and assembly of experimental equipment and staff to maintain the species in a good condition and health over long periods of time.

4.3 Experimental Design

We have designed our experiment to validate the hypothesis that feeding can be detected from heart-rate signals of coral reef fish species. While we ran the experiments in fish tanks, we have spent a considerable effort to ensure that the conditions were realistic. For example, the temperature profile of the water was set to follow seasonal and daily patterns corresponding to different weather

³This includes the current drawn by CPU, sensors, and NAND flash memory.

seasons in Townsville, Australia, which is located on the Great Barrier Reef. The fish interacted with each other and the equipment during the day, resulting in a wide range of daily activities, such as fighting, swimming, rubbing, and feeding. ECG and heart-rate data were impacted by all of these factors.

4.3.1 Fish species. Coral trout (*Plectropomus leopardus*) were obtained from the Australian Reef Fish Trading Company and were wild-caught from reefs around Cairns and Mackay, Australia. Fish were kept in quarantine tanks with flow-through seawater for a two-week period during which they were fed two or three times weekly. Following quarantine and acclimatization, fish were implanted with biosensors and transferred into six fish tanks for the experiment. Each tank was fitted with a lid to prevent fish from jumping out of the tank and large PVC pipes were provided as structure and environmental enrichment to the tank.



Figure 7: Left: fish tank with six coral trouts. Right: coral trout surgery in progress.

4.3.2 Experimental tanks. For the main experiment, we designed six fit-for-purpose fish tanks (see Figure 7) in close collaboration with SeaSim. Each experimental tank held six coral trouts and all six tanks were located in a temperature controlled room without daylight. Tanks were supplied by a flow-through of temperature controlled seawater that was set to maintain desired dissolved oxygen and ammonia levels in the tanks. The temperature as well as the artificial lighting intensity were managed by the general Control System of the SeaSim facility, to replicate the photo-thermo period seasonal baseline of local reefs. Each of the 6 systems was equipped with air supply and air stones and, to assist with the control of ammonia, the tank water was recirculated through a custom built bio-filter.

Daily fish health checks and key maintenance tasks were performed during all stages of the experiment, including cleaning of tanks, water changes and visual inspection of each experimental tank system. In addition, key water quality parameters such as dissolved oxygen (DO) and ammonia (NH₃) were measured and corrective action was taken immediately if required to ameliorate the water quality parameters.

4.3.3 Biosensor implantation and fish recovery. Following a two-week quarantine and acclimatization period, fish were implanted with our biosensors and 5 commercial Star-Oddi sensors and given a colored dorsal tag for individual identification. Surgical approaches followed general practice in literature ([12, 13]) and in accordance

with our animal ethics approval ⁴. After surgery, fish were given a three-week recovery period and placed into the six experimental tanks (six fish per tank) at a temperature of 24.5°C. Fish were not fed during the first week of the recovery period to support healing of the abdominal incision. During the following two weeks of the recovery period, fish were fed a maintenance ration every two to three days. This maintenance ration consisted of approximately 20 to 30 gram of thawed pilchards per coral trout.

4.3.4 Experimental treatment. After the three-week recovery period, all six tanks were progressively changed to their target treatment temperature (0.5°C/day). The six tanks were divided into three different temperature treatments:

- (1) tanks 1 and 4 were set at 28.6°C (high-temperature);
- (2) tanks 2 and 3 were set at 26.1°C (mid-temperature);
- (3) tanks 5 and 6 were set at 23.6°C (low-temperature).

These temperature treatments were subsequently maintained for a period of six weeks. Afterwards, the fish were individually netted and euthanized for the removal of the biosensor and subsequent download of the data.

4.3.5 Feeding trials. The feeding occurred every 2-3 days to give fish an opportunity to feed up to 11 times. The feed was thawed pinkies or pilchards most of the time but we also used freshly killed damselfish once a fortnight (4 adult *Acanthochromis* per coral trout). The meal size was randomized for each coral trout, although it proved difficult to follow the feeding schedule due to some fish feeding aggressively and other fish refusing to take food. Observations of feeding were manually recorded in log sheets for each feeding event and each fish (identified using external tags). Additionally, 2 random tanks were instrumented with GoPro cameras that recorded the whole feeding process.

4.3.6 Discussion on deployment in the wild. As described above, currently, the fish implanted with the biologgers need to be recaptured to retrieve the biologgers and the collected data. This is one of the reasons for our choice of coral trout as the subject species for the experiment since coral trout stay around the same reef for their entire lives and are highly likely to be recaptured. However, there are several challenges that should be addressed to enable more effective deployment of our biologgers in the wild. First, the device has to be equipped with underwater communication capability to stream the data or high-level information inferred from the data back to a gateway and eliminate the need for the retrieval of the biologger. Second, the size of the biologger ought to be minimized through optimizing the hardware design. We note that in our experience, our biologger is one of the best existing solutions and has shown promise for future deployment in the wild once the underwater communication component is added. It minimizes communication bandwidth by running sophisticated algorithms on-board, its sensed physiological signals are accurate, and its size is similar to the state-of-the-art commercialized implantable biologgers such as Star-Oddi's products [30].

⁴Animal ethics approval was obtained from the James Cook University's Animal Ethics Committee with approval number A2314.

5 PERFORMANCE EVALUATION

We have selected twelve fish from the trials for the analysis in this section, based on the biosensor performance (some sensors failed due to hardware/software problems) and based on the feeding performance of the fish (some fish ate substantially less food than expected). All of the data comes from our biologgers as none of the commercial Star-Oddi sensors returned good quality ECG or heart-rate data. Our conjecture is that despite following the manufacturer instructions on installing the sensors in the fish, the sensors were simply too far from the heart, resulting in low signal-to-noise ratio of the heart-rate signals.

5.1 Heart-rate Estimation

We evaluate the accuracy of our heart-rate estimation algorithm in comparison with other existing approaches. We randomly selected 100 ECG samples from Fish G1 and manually labeled the locations of the peaks corresponding to heartbeats. We performed the labeling in consultation with our fish physiologist to ensure having high-quality ground truth.

The evaluation metric that we use is the mean absolute error (MAE) calculated as the average of the absolute differences between the ground truth and the estimated median heart-rate values. We evaluate the performance of our algorithm by comparing it with three competing methods, which are the wavelet-based algorithm of [33], filtering-based algorithm of [21] and Labchart [29] software. The wavelet-based method applies wavelet transformation and peak detection to estimate the heart-rate. The filtering-based method filters the ECG signal using band-pass and low-pass filters with carefully-tuned bandwidths then detects the peaks of the filtered signal. Labchart is a popular software tool for life science data analysis used by fish biologists to analyze fish ECG time-series.

methods	proposed	filtering-based	wavelet-based	Labchart
MAE (bpm)	0.14	2.2	30.8	20.9

Table 3: MAE of different heart-rate estimation methods.

Table 3 presents the evaluation results for the four methods. The results show that the proposed algorithm achieves the lowest MAE at 0.14 bpm. Due to the noisy nature of the ECG signals collected, wavelet-based method and Labchart performed poorly. The MAE for Labchart and Wavelet-based are 20.9 bpm and 30.8 bpm, respectively. Filtering-based method achieves reasonable accuracy if we tune the filter bandwidths carefully. Nonetheless, our method achieves 15 times better performance in terms of MAE (0.14 bpm vs. 2.2 bpm).

5.2 Feeding Event Detection

5.2.1 Evaluation Goal and Metrics. The goal of the evaluation in this section is to demonstrate that our proposed online change detection algorithm, which is built on the Bayesian change-point detection algorithm proposed in [1] and devised to cope with constrained resources, achieves an accuracy comparable to that of the original algorithm on both synthetic and real-world datasets. We are not aware of any existing online change-point detection algorithm suitable for resource-constrained embedded systems or long-term deployment in real-world scenarios.




Figure 8: Evaluation results on synthetic datasets: the two online algorithms perform similarly for (a) different maximum chain lengths L and (b) for a wide range of future points. Low approximation error (a) shows that our online algorithm closely tracks the original one. Finally, while offline algorithm provides better results (c), the performance gap is not significant.

We use a number of different evaluation metrics in this section to demonstrate the performance of different change detection approaches, namely:

- the false positive rate (FPR): the probability of negative events being falsely detected as positive events;
- the false negative rate (FNR): the probability of positive events being falsely detected as negative events;
- the approximation error: the average of the absolute differences of the entries of two vectors containing probabilities of change;
- the overall error rate (OER): the proportion of false detections over all positive and negative events;
- the equal error rate: the value of FPR (or FNR) when FPR and FNR are balanced (become equal);
- the event localization precision: the distance between the locations of correctly detected events and their corresponding ground truth.

5.2.2 Evaluations on Synthetic Datasets. The aim of evaluation on synthetic datasets is to compare our lightweight change-point detection algorithm to prior art, specifically, with the offline Bayesian approach and the original Bayesian online approach. We call these algorithms *Lightweight Online Bayes*, *Offline Bayes*, and *Original Online Bayes*, respectively.

We follow the same approach as in the related literature [28] and generate the synthetic data by making changes at certain points in the statistical properties of simulated i.i.d. time-series. Each time-series consists of 6000 data points drawn from a Gaussian distribution whose statistical properties, i.e., mean and standard deviation, change every 1000 data points. Therefore, each time-series data sequence has 5 change points. We independently generate 30 time-series each with 5 change points at random locations and report the evaluation results that are averaged over all 30 simulation trials.

During the evaluation, we observed that FPR is zero in most cases as the statistic properties remain the same within a segment until a change point is encountered. Therefore, we use only FNR to demonstrate the performance of the change detection approaches with different parameters. In addition, we use OER to evaluate the performance of the considered algorithms with real-world datasets

as OER combines the false and negative errors and is more suitable for complex and noisy real-world data.

Figure 8 compares the performance of the algorithms for different parameter values, which include a) the maximum length of the chain L , which is related to the memory required to store the relevant parameters, b) the number of future points used to estimate the probability of change point, and c) the threshold for detecting change points. Note that parameters a) and b) are not relevant for the Offline Bayes algorithm.

Figure 8(a) shows the results for FNR. We can observe that our lightweight algorithm performs better with increasing L and achieves the same accuracy as the original online Bayes for $L > 300$. Moreover, the approximation error between the outputs of the lightweight and original approaches shows decreasing trend with the growth of the maximum length indicating that the proposed lightweight approach can better approximate the original approach. In real world deployments, we can set the maximum length according to the available resource.

The impact of the number of *future points* required to estimate the current probability of change is shown in Figure 8(b). The *maximum length* L is set as 300 and the threshold for determining the change point is set as 0.5. The results show that FNRs of both lightweight and original algorithms are around 0.08 when the number of *future points* is equal to or more than 10. Note that the detection latency increases when more *future points* are needed.

Finally, we evaluate FNR of all three approaches with *maximum length* and *future points* parameters set to 300 and 30, respectively. The results show that the lightweight algorithm achieves similar FNR as the original online algorithm. Although the offline approach produced the lowest FNR, the performance gap between these three approaches is insignificant.

Recall that we do not include FPR in these results as the metric is very small when using synthetic data. However, the choice of the threshold entails a trade-off between FNR and FPR, which will be discussed in the real-world evaluation.

5.2.3 Feeding Event Detection on Fish Datasets. We now evaluate the performance of our lightweight algorithm on the heart-rate traces obtained from the ECG signals of fish in the real-world experiments (see Section 4). We aim to determine whether a fish feeds during any given day, i.e., a period of 24 hours. Before the

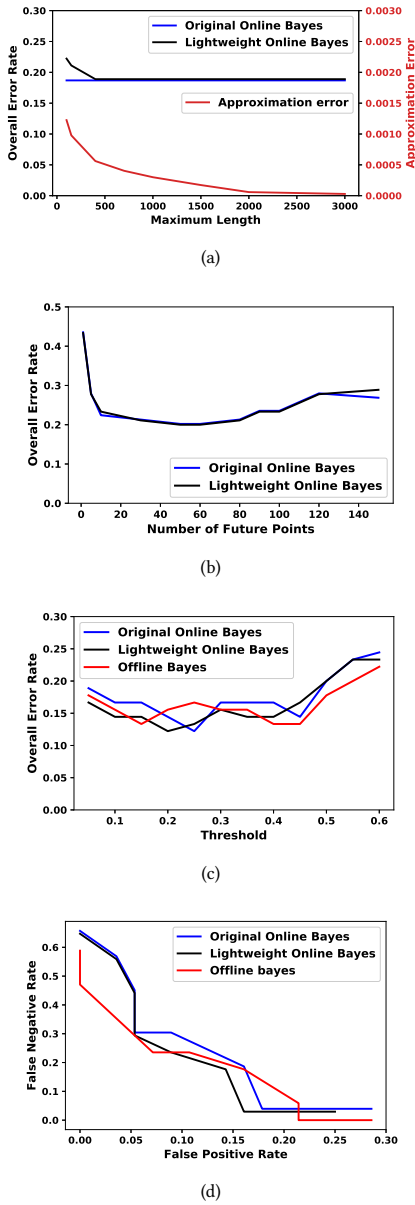


Figure 9: Evaluation results on real-world datasets: the algorithms perform similarly as we vary the key parameters.

evaluation, we first divide the heart-rate estimates into multiple 24-hour segments with 12 hours overlap to increase the number of samples for feeding detection. Then, according to the recorded feeding time-sheet, we mark each sample as positive (containing feeding events) or negative (with no feeding event). The event detection algorithms are applied on each segment to yield positive or negative predictions and the evaluation results are obtained by comparing the predictions with the ground truth.

We first evaluate the performance of the fish feeding detection algorithm using the OER metric, subject to different values of *maximum length*, number of *future points*, and *threshold* parameters.

Figure 9(a) shows that OER decreases with the increase of maximum length and levels at around 400. However, the approximation error keeps decreasing, which is intuitive as the lightweight algorithm can better approximate the original online Bayes with more data stored in the memory.

Figure 9(b) shows that as the number of *future points* increases up to 20, OER drops significantly. OER then levels off and starts gradually increasing after we reach 80 future points, indicating that there is an optimal parameter setting that yields the best accuracy. Intuitively, the algorithm will stop benefiting from the increased number of future points at a point where the future data stops correlating strongly with food digestion and heart-rate starts getting impacted by the daily temperature cycle.

We next evaluate OER of all three approaches with different values of the threshold, see Figure 9(c). We can observe that the three approaches achieve comparable low error rates and none shows any superior performance over the others.

The trade-off between FPR and FNR is important in event detection systems. We use FPR and FNR as the metrics to compare the performance of the three approaches. In a detection system, both low FPR and FNR are desired. The trade-off between FPR and FNR should be carefully tuned according to specific requirements of the target application. The final evaluation results are shown as the FPR-FNR curve in Figure 9(d). Specifically, for our proposed lightweight Bayes approach, the equal error rate is 0.15 when FPR and FNR are balanced, which indicates the algorithm misses the feeding events with a probability of 15% when 15% of non-feeding segments are falsely detected as containing feeding events.

Finally, we estimate how close temporally are the detected feeding points to the manually-labeled ground truth. We compare the locations of the correctly detected feeding events and their corresponding ground truth in the heart-trace traces. The average event localization precision is about 11 minutes.

6 CONCLUSION AND FUTURE WORK

We address the problem of long-term monitoring of condition and health of wild fish free-living in coastal ecosystems. We implement and validate two key algorithms in this paper and present an implantable biologger for in-situ monitoring of fish physiology. The biologger collects multi-modal signals including ECG, acceleration, temperature, and pressure, which can be used to estimate heart-rate, respiration rate, and movement of the fish.

We implanted a number of biologgers in coral trouts in an experiment that ran over a period of ten weeks and collected feeding-related data in different environmental conditions. We evaluate our heart-rate estimation and feeding event detection algorithms using the empirical data from fish trials. The results show that our algorithms are robust to noise and achieve heart-rate estimation errors of 0.14 bpm and feeding detection overall error rate of 0.15. Our novel lightweight change detection algorithm achieves similar performance to the state-of-the-art change detection algorithms while significantly reducing their computation and memory requirements. This facilitates implementing the algorithms on embedded

hardware, an important step towards long-term in-situ monitoring systems that rely on bandwidth-constrained underwater communications.

In future work, we will extend the biollogger with an acoustic communications interface and implement feeding detection algorithms directly on the device. We aim to validate the approach in long-term experiments in the wild. Although we have devised the proposed algorithms specifically to be run on resource-constrained embedded systems, the implementation is challenging as it has to be carried out in an optimal fashion to minimize the resource consumption. Depending on the availability of the memory and computational capability as well as the available power budget, certain hyperparameters may require careful tuning to yield a desirable trade-off between the competing aspects of resource consumption and performance. We will also investigate more sophisticated signal separation and data fusion algorithms to utilize the available multi-modal data more effectively. For example, one of our goals is to estimate the respiration rate of the fish from the ECG signal, which may provide additional information for the study of fish energetics.

REFERENCES

- [1] Ryan Prescott Adams and David JC MacKay. 2007. Bayesian online changepoint detection. *arXiv preprint arXiv:0710.3742* (2007).
- [2] John Alroy. 2015. Current extinction rates of reptiles and amphibians. *Proceedings of the National Academy of Sciences* (2015).
- [3] Samaneh Aminikhanghahi and Diane J Cook. 2017. A survey of methods for time series change point detection. *Knowledge and information systems* 51, 2 (2017), 339–367.
- [4] JD Armstrong. 1986. Heart rate as an indicator of activity, metabolic rate, food intake and digestion in pike, *Esox lucius*. *Journal of Fish Biology* 29 (1986), 207–221.
- [5] Therese Arvén Norling. 2017. Remotely monitoring heart-rate and feeding behaviour of fish by using electronic sensor-tags. (2017).
- [6] BE Beisner, DT Haydon, and K. Cuddington. 2003. Alternative stable states in ecology. *Frontiers in Ecology and the Environment* 1, 7 (2003), 376–382.
- [7] I.-A. Bisson, L. K. Butler, T. J. Hayden, P. Kelley, J. S. Adelman, L. M. Romero, and M. C. Wikelski. 2011. Energetic response to human disturbance in an endangered songbird. *Animal Conservation* 14, 5 (2011), 484–491.
- [8] C. J. Bridger and R. K. Booth. 2003. The Effects of Biotelemetry Transmitter Presence and Attachment Procedures on Fish Physiology and Behavior. *Reviews in Fisheries Science* 11(1) (2003).
- [9] Gerardo Ceballos, Paul R. Ehrlich, Anthony D. Barnosky, Andrés García, Robert M. Pringle, and Todd M. Palmer. 2015. Accelerated modern human-induced species losses: Entering the sixth mass extinction. *Science Advances* 1, 5 (2015).
- [10] TD Clark, WT Brandt, J Nogueira, LE Rodriguez, M Price, CJ Farwell, and BA Block. 2010. Postprandial metabolism of Pacific bluefin tuna (*Thunnus orientalis*). *Journal of Experimental Biology* 213, 14 (2010), 2379–2385.
- [11] TD Clark, BD Taylor, RS Seymour, D Ellis, J Buchanan, QP Fitzgibbon, and PB Frappell. 2008. Moving with the beat: heart rate and visceral temperature of free-swimming and feeding bluefin tuna. *Proceedings of the Royal Society B: Biological Sciences* 275, 1653 (2008), 2841–2850.
- [12] Timothy Darren Clark, SG Hinch, BD Taylor, PB Frappell, and AP Farrell. 2009. Sex differences in circulatory oxygen transport parameters of sockeye salmon (*Oncorhynchus nerka*) on the spawning ground. *Journal of Comparative Physiology B* 179, 5 (2009), 663–671.
- [13] Timothy Darren Clark, E Sandblom, SG Hinch, DA Patterson, PB Frappell, and AP Farrell. 2010. Simultaneous biologging of heart rate and acceleration, and their relationships with energy expenditure in free-swimming sockeye salmon (*Oncorhynchus nerka*). *Journal of Comparative Physiology B* 180, 5 (2010), 673–684.
- [14] Timothy Darren Clark, E. Sandblom, S. G. Hinch, D. A. Patterson, P. B. Frappell, and A. P. Farrell. 2010. Simultaneous biologging of heart rate and acceleration, and their relationships with energy expenditure in free-swimming sockeye salmon (*Oncorhynchus nerka*). *Journal of Comparative Physiology B* 180, 5 (01 Jun 2010), 673–684. <https://doi.org/10.1007/s00360-009-0442-5>
- [15] Gari D Clifford. 2002. *Signal processing methods for heart rate variability*. Ph.D. Dissertation. Oxford University, UK.
- [16] Steven J. Cooke, Jacob W. Brownscombe, Graham D. Raby, Franziska Broell, Scott G. Hinch, Timothy D. Clark, and Jayson M. Semmens. 2016. Remote bioenergetics measurements in wild fish: Opportunities and challenges. *Comparative Biochemistry and Physiology Part A: Molecular and Integrative Physiology* 202 (2016), 23 – 37.
- [17] Mark A. Dittner, John B. Vincent, Leland K. Werden, Jessie C. Tanner, Timothy G. Laske, Paul A. Iaizzo, David L. Garshelis, and John R. Fieberg. 2015. Bears Show a Physiological but Limited Behavioral Response to Unmanned Aerial Vehicles. *Current Biology* 25, 17 (2015), 2278 – 2283.
- [18] Ursula Ellenberg. 2017. *Impacts of Penguin Tourism*. Springer International Publishing, Cham, 117–132.
- [19] Paul Fearnhead. 2006. Exact and efficient Bayesian inference for multiple changepoint problems. *Statistics and computing* 16, 2 (2006), 203–213.
- [20] David Griggs, Mark Stafford-Smith, Owen Gaffney, Johan Rockström, Marcus C. Öhman, Priya Shyamsundar, Will Steffen, Gisbert Glaser, Norichika Kanie, and Ian Noble. 2013. Sustainable development goals for people and planet. *Nature* 495 (20 03 2013), 305 EP –.
- [21] Patrick S Hamilton. 2002. Open source ECG analysis software documentation. *Computers in cardiology 2002* (2002), 101–104.
- [22] Scott G Hinch, Emily M Standen, Michael C Healey, and Anthony P Farrell. 2002. Swimming patterns and behaviour of upriver-migrating adult pink (*Oncorhynchus gorbuscha*) and sockeye (*O. nerka*) salmon as assessed by EMG telemetry in the Fraser River, British Columbia, Canada. In *Aquatic Telemetry*. Springer, 147–160.
- [23] Terry P. Hughes, James T. Kerry, Andrew H. Baird, Sean R. Connolly, Andreas Dietzel, C. Mark Eakin, Scott F. Heron, Andrew S. Hoey, Mia O. Hoogenboom, Gang Liu, Michael J. McWilliam, Rachel J. Pears, Morgan S. Pratchett, William J. Skirving, Jessica S. Stella, and Gergely Torda. 2018. Global warming transforms coral reef assemblages. *Nature* 556, 7702 (2018), 492–496.
- [24] Matthew L. Keefer and Christopher C. Caudill. 2016. Estimating thermal exposure of adult summer steelhead and fall Chinook salmon migrating in a warm impounded river. *Ecology of Freshwater Fish* 25, 4 (2016), 599–611.
- [25] Johannes Kulick. 2016. Bayesian Changepoint Detection Implementation. https://github.com/hildensia/bayesian_changepoint_detection.
- [26] Miles D Lamare, Tracey Channon, Chris Cornelisen, and Murray Clarke. 2009. Archival electronic tagging of a predatory sea star—testing a new technique to study movement at the individual level. *Journal of Experimental Marine Biology and Ecology* 373, 1 (2009), 1–10.
- [27] Albert Lehninger, David L. Nelson, and Michael M. Cox. 2008. *Lehninger Principles of Biochemistry*. W. H. Freeman.
- [28] Rakesh Malladi, Giridhar P Kalamangalam, and Behnaam Aazhang. 2013. Online Bayesian change point detection algorithms for segmentation of epileptic activity. In *2013 Asilomar Conference on Signals, Systems and Computers*. IEEE, 1833–1837.
- [29] Youzou Miyadera, Shoichi Nakamura, Taisuke Nanashima, and Setsuo Yokoyama. 2008. LabChart: A Support System for Collaborative Research Activities in University Laboratories and its Practical Evaluations. In *2008 12th International Conference Information Visualisation*. IEEE, 169–178.
- [30] Star Oddi. 2019. <https://www.star-oddi.com/>.
- [31] Eyal Shamur, Miri Zilka, Tal Hassner, Victor China, Alex Liberzon, and Roi Holzman. 2016. Automated detection of feeding strikes by larval fish using continuous high-speed digital video: a novel method to extract quantitative data from fast, sparse kinematic events. *Journal of Experimental Biology* 219, 11 (2016), 1608–1617.
- [32] Bograd SJ, Block BA, Costa DP, and Godley BJ. 2010. Biologging technologies: new tools for conservation. Introduction. *Endangered Species Research* 10 (2010), 1–7.
- [33] Andrew Tan. 2016. Electrocardiograms: QRS detection using wavelet analysis. (2016).
- [34] Vinay Udyawer, Colin A Simpfendorfer, Michelle R Heupel, and Timothy D Clark. 2017. Temporal and spatial activity-associated energy partitioning in free-swimming sea snakes. *Functional ecology* 31, 9 (2017), 1739–1749.
- [35] Economist Intelligence Unit. 2015. *The Blue Economy: Growth, Opportunity and a Sustainable Ocean Economy*. (2015).
- [36] RE Whitlock, A Walli, P Cermeño, LE Rodriguez, C Farwell, and BA Block. 2013. Quantifying energy intake in Pacific bluefin tuna (*Thunnus orientalis*) using the heat increment of feeding. *Journal of Experimental Biology* 216, 21 (2013), 4109–4123.
- [37] A.D.M. Wilson, M. Wikelski, R.P. Wilson, and S.J. Cooke. [n.d.]. Utility of biological sensor tags in animal conservation. *Conservation Biology* 29, 4 ([n. d.]), 1065–1075.
- [38] S.M. Wilson, S.G. Hinch, E.J. Eliason, A.P. Farrell, and S.J. Cooke. 2013. Calibrating acoustic acceleration transmitters for estimating energy use by wild adult Pacific salmon. *Comparative Biochemistry and Physiology Part A: Molecular and Integrative Physiology* 164, 3 (2013), 491 – 498.