2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 23-27, 2022, Kyoto, Japan

Adaptive Online Sampling of Periodic Processes with Application to Coral Reef Acoustic Abundance Monitoring

Seth McCammon, Nadège Aoki, T. Aran Mooney, Yogesh Girdhar

***Abstract*— In this paper, we present an approach that enables long-term monitoring of biological activity on coral reefs by extending mission time and adaptively focusing sensing resources on high-value periods. Coral reefs are one of the most biodiverse ecosystems on the planet; yet they are also among the most imperiled: facing bleaching, ecological community collapses due to global climate change, and degradation from human activities. Our proposed method improves the ability of scientists to monitor biological activity and abundance using passive acoustic sensors. We accomplish this by extracting periodicities from the observed abundance, and using them to predict future abundance. This predictive model is then used with a Monte Carlo Tree Search planning algorithm to schedule sampling at periods of high biological activity, and power down the sensor during periods of low activity. In simulated experiments using long-term acoustic datasets collected in the US Virgin Islands, our adaptive Online Sensor Scheduling algorithm is able to double the lifetime of a sensor while simultaneously increasing the average observed acoustic activity by 21%.**

I. Introduction

Observing phenomena over long time periods in remote environments, such as in space or underwater, is a challenging problem. The environments' remoteness makes it difficult to deploy robots or to recover them. Consequently, any robot deployed for one of these sensing tasks must maximize its own sensing effectiveness given its limited available resources. The primary resource constraint is typically battery power, which limits the total number of observations a robot can collect over the course of its deployment. Temporarily shutting down non-essential systems, such as sensors, decision-making computers, and actuators allows robots to conserve power [1], [2], prolonging mission duration. However this comes at the cost of missing potentially valuable observations. By intelligently selecting when to emerge from hibernation, a robot can extend its sampling time without sacrificing the quality of the samples obtained. In this paper we focus on this problem of adaptive sampling of a time-varying phenomena, motivated by the challenge of monitoring biological activity on coral reefs.

Coral reefs are one of the most biologically productive areas on the planet, home to over 25% of all marine species [3]. However, they are also among the most fragile. Reef biodiversity and overall health is facing devastating recent degradation from such stressors as ocean warming,

The authors are with the Woods Hole Oceanographic Institution, Woods Hole, MA. {smccammon, naoki, amooney, ygirdhar}@ whoi.edu

This work was supported in part by NSF NRI Grants 1734400 and 2133029, NSF OTIC Award Number 2024077, the Postdoctoral Scholar Program at Woods Hole Oceanographic Institution, with funding provided by the Doherty Foundation, and the Investment in Science Fund at WHOI.

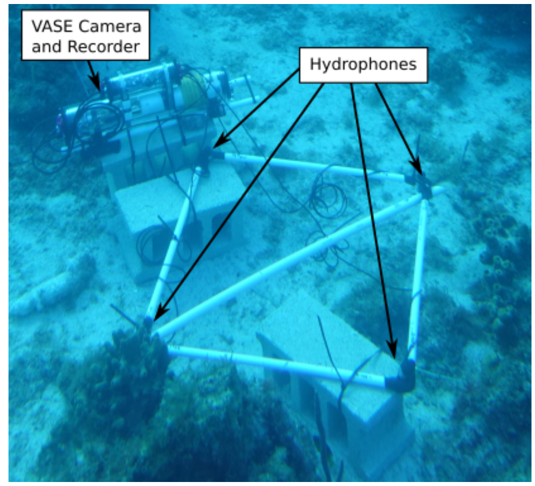

Fig. 1: The Visual-Acoustic Sensor Emplacement (VASE). The camera and acoustic data collected by VASE is processed in realtime to compute bioactivity measures. The system then adaptively make decisions about when to put the sensor in a low-power sleep mode, while optimizing the total observed bioactivity.

overfishing [4] and ocean acidification [5]. Thus it is critical to identify and track reef health and biodiversity in order to understand the efficacy of marine protected areas, help prioritize and manage reefs under threat, and identify the states to which reefs will need to be restored [6]. Tracking coral reef health and biodiversity through time requires measurements of the abundance and biological activity of organisms present on or within a reef [7]. Traditionally, these measures are computed by hand, using human divers swimming along coral reefs. However, this is a costly and labor-intensive approach. One promising alternative is to use sensors deployed on reefs to collect long-term datasets. Since many reef organisms are sonically active, passive acoustics is a potentially powerful tool for reef monitoring, however current analyses of passive acoustic data are limited in their ability to measure biodiversity from acoustic data since large gaps remain in our understanding of which acoustic signals are produced by which species, particularly among fishes [8].

To fill in this gap, larger and richer datasets are required that augment acoustic data with other sensors that can better characterize biodiversity. One sensor capable of producing these datasets is Visual-Acoustic Sensor Emplacement (VASE), shown in Fig. 1. VASE simultaneously records camera imagery and acoustic data, building datasets that can more directly measure biological activity on reefs. Due to its increased sensor payload, VASE consumes more power than traditional hydrophone recorders. Consequently, to improve VASE's ability to monitor coral reefs, new methods are

2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) | 978-1-6654-7927-1/22/$31.00 ©2022 IEEE | DOI: 10.1109/IROS47612.2022.9982217

needed that will enable it to conserve energy and focus observations only during periods of high activity when it is more likely to simultaneously observe fish visually and acoustically.

The problem of selecting observation periods that maximize the number of fish VASE's camera is expected to capture while respecting the constraints imposed by the limited battery power is an online version of the Sensor Scheduling Problem (OSSP) [9]. We assume that the biological activity evolves as a periodic function of time. There are strong priors on certain periods, such as diel and tidal cycles for coral reef biological activity that could be used to inform the sampling schedule [7]. However, such a fixed approach may limit the ability of VASE to target high-value times, since a fixed sampling schedule may introduce aliasing of other frequencies.

The primary contribution of this paper is a novel framework for informatively sampling a periodic time series. Our approach requires no prior data, instead it learns the periodicities of an intermittently sampled function online, and adapts a sensor's sampling behavior to spend its energy budget more efficiently by focusing on periods of high-value observations. We demonstrate our approach using time series data collected over long-term hydrophone deployments in the US Virgin Islands.

II. Related Work

A. Acoustic Monitoring of Biological Abundance

The widespread distribution and isolation of many habitats found along secluded coastal areas, reefs, seamounts and insular habitats can make traditional approaches to biological monitoring logistically difficult and expensive. Research cruises often result in high ship time costs and typically allow only intermittent and limited opportunities for assessing the conditions at many sites [10]. Moored instruments capable of measuring a wide range of environmental parameters, such as surface and subsurface temperatures, salinity, wave energy, and current flow provide measures of the physical habitat, but do not obtain data directly about the biological activity taking place at a location. As a result, many significant ecological events, such as disease outbreaks, episodic infestations (e.g. harmful algal blooms), reactions to climate change (e.g. massive coral bleaching), and the effect of storms, oil spills and poaching often occur undetected, complicating the interpretation of long-term monitoring data.

An emerging area of study uses hydrophones to listen to the reef "soundscape", as a more direct measure of biological activity on the reef. This is because sounds present in many marine habitats can be an effective indicator of many biological processes, such as spawning events [11], feeding [12], and social communication [13] among many species of fish, invertebrates, and aquatic mammals. Passive acoustics sensors have now been applied on a variety of habitats to track differences in reef health, fish abundances, coral cover, human presence, and related parameters [7], [14]. However, these acoustic recorders, like many robots, rely on a self-contained battery to provide power.

B. Irrevocable Decision Making

Many problems in robotics, particularly ones in time-varying domains, contain elements of irrevocable decisions [15], [16], [17]. An irrevocable decision is one that once made, cannot be revised, through returning and selecting a different alternative. Perhaps the best known irrevocable decision problem is the Cayley-Moser problem [18], often called the hiring problem or secretary problem, where a manager is presented with a sequence of independently and identically distributed (IID) candidates with differing qualifications for a position. The manager's goal is to select the most qualified candidate. In the classical version of the problem, a near-optimal algorithm exists, that selects the best candidate with probability $1/e$ [19]. Many variants of this problem exist, such as selecting the k most qualified candidates [20], [21], selecting candidates with different associated costs [22], or selecting candidates when the candidates do not appear according to a uniform distribution [23]. The problem that we consider differs from the versions of the Cayley-Moser problem in several key ways. The first of these is that we, like Flashpohler et al. [23], do not assume that the samples we may observe appear IID, but instead vary according to a periodic function. Additionally, we go a step beyond the work presented in [23], and do not assume prior knowledge of the period of the function of interest.

C. Spatiotemporal Monitoring

The OSSP can also be thought of as a monitoring problem, similar to the informative path planning problem, where the goal is to plan a set of observation windows that maximizes the total utility of the resultant observations. In the case where the biological utility function is known, this reduces to a sensor placement problem, similar to the optimal stopping problem explored in Best et al. [16]. Since we do not know the value of the biological utility, and instead it is only revealed through observations, we require a different and online planning approach.

A common approach to spatiotemporal monitoring is to use a Gaussian Process to compile the set of observations collected by a robot into a time-varying belief of the state of the world [24], [25]. One of the primary challenges with this approach is that many of the common Gaussian Process kernels (e.g. squared-exponential or Matèrn) struggle to extrapolate the behavior of the latent function beyond the set of observations [25]. One notable exception is the periodic kernel [26] that, with correctly tuned hyperparameters, can achieve good predictive performance while extrapolating beyond observations [27]. However, many of the approaches to time-series modelling which incorporate periodicity rely on prior knowledge of the period, or seasonality, of the signal [28], [29]. For monitoring previously unobserved processes, this prior knowledge cannot always be guaranteed. While the periodicity hyperparameters can be estimated online using the observed data with *maximum a posteriori* methods, this process is computationally expensive ($O(n^3)$), and generally relies on good prior estimates to find the global maximum parameters. Fortunately, the frequency-domain analysis of time series data provides an alternative. The Fourier Transform, and its variants, allow us to decompose a periodic time series signal into a set of component frequencies and intensities with a complexity of $O(n\ log\ n)$ [30]. There are two approaches to including the Fourier transform in developing GP kernels, by either directly transforming the GP hyperparameters into the frequency domain [31] or constructing a composite kernel with automatically detected frequency components [32]. However, the application of these approaches has been limited to time series prediction, and to the best of our knowledge have not been applied to any realtime monitoring task.

III. Problem Formulation

We can formulate the Online Sensor Scheduling problem as a Partially Observable Markov Decision Process (POMDP). A POMDP is fully defined by a set of states, $\mathcal{S}$, set of actions $\mathcal{A}$, a conditional state transition function, $T(s, a, s')$, a set of observations, Ω, and the conditional probabilities of making each observation O, and a reward function $R(\cdot)$ [33].

- **State**: The OSSP state is comprised of four elements: the current time, t, the amount of biological activity in the environment, $f(t)$, the robot's binary state $s_t \in$ [sleep, wake], and the amount of battery remaining B_t. Three of these components: t, S_t, and B_t, are fully observable by the robot, while knowledge about biological activity can only be gleaned by making observations.
- **Actions**: The robot has two actions observing and sleeping. Each action has an associated cost C_{obs} and C_{slp} where $0 < C_{slp} << C_{obs}$.
- **State Transition Function**: The known components of the state, t, S_t, and B_t, transition deterministically with each action. Taking the sleep or wake action adjusts the robot's internal state to the respective sleeping or observing state, while deducting the action's corresponding cost from the available remaining battery power. While the state transition function for $f(t)$ is unknown to the robot, we do assume it takes a periodic form with unknown frequency components.
- **Observations**: The robot is able to make noisy observations of $f(t)$ using its hydrophones. The noise is assumed to be Gaussian, with a mean centered on $f(t)$.
- **Reward**: To enable the scientific objective of measuring acoustic activity at a coral reef, we use a sum-of-activity reward function for an action sequence, $\mathcal{A}$ given by

$$R(\mathcal{A}) = \sum_{a_t \in \mathcal{A}} f(t) \times \mathbb{1}(a_t), \tag{1}$$

  where $\mathbb{1}(\cdot)$ is the indicator function with a value of 1 when a_t is to sense, and 0 when a_t is to wait. This function rewards the robot for observing high levels of biological activity, maximizing the chances of observing rare species on the reef.

The unique challenges of the OSSP are captured by the state transition function, $T(s, a, s')$. As discussed in Section II, the problem that we address in this paper shares the constraint of irrevocable actions with the optimal stopping problem and the hiring problem. Since time is a component of our POMDP state, the robot can never return to a previous state, once the decision to take any action is made, the decision can never be repeated. The second key aspect of $T(\cdot)$ is the assumption that the transition of $f(\cdot)$ is periodic. We will show how this assumption can be exploited in Section IV-B to improve the robot's belief estimation. The final component of $T(\cdot)$ that we would like to highlight is the relationship between the action cost and the overall mission budget. This type of budget constraint is all too common in field robotics applications. Here, it imposes a knapsack constraint on the optimal action sequence. Since both actions have nonzero cost, there is a direct trade-off between maximizing the total number of sensing actions and delaying sensing to higher-value periods of time.

IV. Method

The solution to the POMDP outlined in Section III is the action sequence, $\mathcal{A}$ that maximizes the reward given by Equation 1. Computing the exact solution to any POMDP is PSPACE-hard [33], and so instead of solving it directly, we choose to adopt a receding horizon approach, where at each decision point, we estimate $f(\cdot)$ using the observations in O. Since $f(\cdot)$ is the only unknown component of the POMDP, we can use the estimate, $\hat{f}(\cdot)$ to construct a more tractable MDP, which can then be solved.

There are three main components to our acoustic monitoring method. The first of these is the acoustic observation processing pipeline itself, which allows us to compute a bioactivity measure using the raw acoustic data collected by the sensor's hydrophones. The second component is the belief modelling which uses a periodic Gaussian Process to extrapolate the bioactivity time series into the future. The final component is the temporal planning algorithm that performs the action selection to maximize the POMDP reward function.

A. Biological Activity Estimation from Acoustics

The first step in conducting acoustic monitoring of coral reef activity is finding an activity measure which can be easily calculated from an acoustic time series. Multiple field studies have found that observations of coral reef fish density and diversity correlate with recorded sound levels at low frequencies and that these sound levels vary over daily and seasonal timescales [34]. Higher frequencies tend to be dominated by sounds of certain invertebrates, such as snapping shrimp, and are less useful for fish biodiversity monitoring purposes. Accordingly, in developing our acoustic monitoring approach we choose to use root-mean-square sound pressure levels (SPL$_{rms}$) calculated across a low-frequency (100-2000 Hz) "fish band" as a time-varying metric of reef biological acoustic activity [7]. This metric is useful, but limited by its inability to easily distinguish biological activity from other sources of low-frequency noise, such as boat traffic or rain. In practice this noise would be post-processed and removed manually by trained experts.

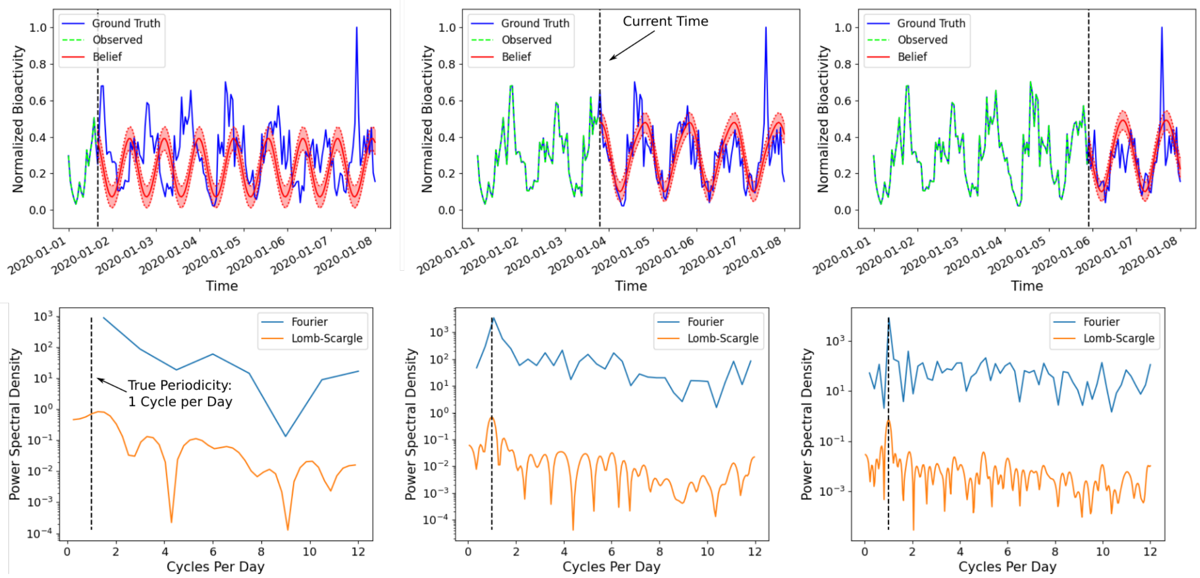

Fig. 2: The proposed Auto-Periodic Gaussian Process predictions with 10, 40, and 70% of the data in the Jan. 2020 Tektite Reef dataset. As the APGP gains more observations, it is able to more accurately predict the period and phase of the periodic acoustic activity. The lower figures show the corresponding estimations of the Power Spectral Density. With fewer observations, the Lomb-Scargle method is still able to identify the true periodicity of 1 daily cycle.

However, since our system is utilizing this data in real-time without surface communications, relying on human annotation is impossible. While it may be possible to automate the detection of anomalous events, it is beyond the scope of the work presented here.

B. *Predicting Future Activity*

To reduce the POMDP outlined in Section III to a tractable MDP, we use a Gaussian Process (GP) to predict the future behavior of the acoustic activity time series. GPs are a widely-used tool for estimating the value of environmental processes [24], [25]. A GP is defined by a mean function, typically assumed to be zero, and a kernel function $K(x, x')$ which defines the covariance of any two points, x and x', in the GP's domain. Many commonly used kernels, such as the Matérn or Squared-Exponential kernels [26] use a length-scale parameter to define the correlation between two points. However, these kernels have poor extrapolative performance beyond the set of observations, since they regress to the mean function in the absence of data [25]. Instead, we choose to model the time series using the periodic kernel,

$$K(x, x') = \sigma^2 \exp\left(\frac{-2\sin^2(T \times (\frac{x-x'}{2}))}{\ell^2}\right), \quad (2)$$

with periodicity controlled by the parameter T. As Tompkins and Ramos [32] showed, setting this parameter correctly is of critical importance to achieving good accuracy in time-series predictions. However, the method they propose to allow these parameters to be learned online relies on a Fourier periodogram. One of the major assumptions of all Fourier-based frequency analysis techniques such as FFT is that $\mathcal{T}$ consists of samples spaced uniformly in time. However, one of the main goals of the VASE system is to intermittently sample the signal of interest, preserving battery power by only sampling when we expect to be able to collect high-value observations. For many applications where the skipped sampling is a relatively small portion of the total number of samples, approximation methods such as interpolation or averaging may be sufficient to restore the even-spacing assumption. Since our approach involves skipping a significant portion of the time series, we require an alternative approach.

To overcome the problem of estimating periodicity from irregularly sampled data, we propose the use of Lomb-Scargle Power Spectral Density [35], [36]. The Lomb-Scargle method uses a least-squares approach to compute the fit of individual sinusoids across a set of candidate frequencies $\mathcal{F} = \{f_1, f_2, ..., f_n\}$,

$$y(t; f, A_f, \phi_f) = A_f \sin(2\pi f(t - \phi_f)), \quad (3)$$

For each candidate frequency, the amplitude, A_f, and phase, ϕ_f are chosen based on a least-squares fit to the observed data minimizing

$$\hat{\chi}^2(f) \equiv \sum_n (y_n - y(t_n; f, A_f, \phi_f))^2 \quad (4)$$

Finally the periodogram is computed by

$$P(f) = \frac{1}{2}(\hat{\chi}_0^2 - \hat{\chi}^2(f)). \quad (5)$$

To choose the range of frequencies, $\mathcal{F}$, over which to evaluate the Lomb-Scargle periodogram, we use the set of positive frequencies over which the standard Fourier periodogram could be computed,

$$\mathcal{F} \triangleq \left[0, \frac{1}{d \times n}, \frac{2}{d \times n}, ..., \frac{n}{2d}\right], \quad (6)$$

where d is the minimum spacing of samples in $\mathcal{T}$, and $n = \lfloor(\max(\mathcal{T}) - \min(\mathcal{T}))/d\rfloor$.

To select T for the GP kernel, we rank the peaks of the Lomb-Scargle periodogram by topographic prominence, the difference in height between each peak and the highest saddle point on each side [37]. We select the k most prominent peaks, and the resultant Auto-Periodic kernel is defined as

$$K_{AP}(x, x') = \sum_{i=0}^{k} \sigma_i^2 \exp\left(\frac{-2\sin^2(T_i \times (\frac{x-x'}{2}))}{\ell_i^2}\right). \quad (7)$$

An example of the Auto-Periodic kernel Gaussian Process (APGP) is shown in Fig. 2.

C. Receding Horizon Sensing Scheduling

The final component of our Online Sensor Scheduling algorithm is the planner that takes a belief at the current time, t, and the robot's state, including the remaining battery, and plans an optimal action sequence that maximizes the reward function given in Equation 1. This problem can be formulated as a Markov Decision Process with Time-Varying Rewards (T-MDP). The state space, action space, and state transition function are identical to those of the POMDP outlined in Section III. The key difference is in the T-MDP the reward function is known; it is the APGP belief constructed from the robot's observations. To solve the T-MDP, we chose to use Monte Carlo Tree Search (MCTS) [38]. The MCTS algorithm explores a robot's action space through biased random sampling, trading off between exploring new action sequences and exploiting previously-sampled high-quality action sequences. At each iteration of the MCTS planner, the MCTS tree is expanded by sampling action sequences using the Upper Confidence Bound for Trees metric [38], and a rollout function to expand the action sequence to the full available budget by means of a heuristic rollout. The full rolled-out sequence is then evaluated according to Equation 1, and the reward is backpropogated up the tree, adjusting the expected reward of each node in the action sequence according to the reward of the sampled sequence and number of times each respective node in the tree has been sampled. This process is repeated for a fixed number of iterations. In our implementation we used $n = 10,000$. MCTS has several properties which make it desirable for our realtime field application. First, MCTS is an anytime algorithm, which enables us to directly trade off the amount of computation available to the planner with the quality of the resultant plan. Second, through the rollout function, MCTS only evaluates full action sequences. For nonmyopic planning problems, like the OSSP, this helps MCTS more fully explore the action space of the robot, leading more quickly to high-quality plans.

Pseudocode for the full Online Sensor Scheduling algorithm is shown in Algorithm 1. At each loop through the algorithm, if the robot is awake, it observes the world, constructs its belief, then plans an action sequence using that belief. It then executes the next step in that plan. Since no replanning occurs while the robot is asleep, the robot will sleep for the full duration planned by the MCTS planner before awakening, observing, and continuing with new information. While a plan with no wake actions would result in the robot completely draining its battery without waking, Equation 1 would give such a plan would have a reward of zero. Therefore it would not be chosen over any other plan with nonzero reward. An example of the behavior of Online Sensor Scheduling algorithm is shown in Fig. 3. Initially, the algorithm has only one observation to work with, and so it plans to sample densely based on its uniform belief. However, as more observations are collected, the APGP belief becomes more accurate, and the planner is able to create a sampling schedule that focuses only on the high-value peaks of bioactivity.

Algorithm 1 Online Sensor Scheduling

function OSS(B_{max}, k)
    $t \leftarrow 0, S_0 \leftarrow$ Sense, $B_0 \leftarrow B_{max}, \mathcal{T} \leftarrow \emptyset$
    while $B_t > 0$ **do**
        if $S_t ==$ sleep **then**
            $B_{t+1} \leftarrow B_t - C_{slp}$
        else if $S_t ==$ wake **then**
            $B_{t+1} \leftarrow B_t - C_{obs}$
            $\Psi_t \leftarrow$ **ObserveAcousticPower**()
            $\mathcal{T} \leftarrow \mathcal{T} \cup (t, \Psi_t)$
            $K_{AP} \leftarrow$ **Lomb-ScargleKernel**($\mathcal{T}, \mathcal{F}(\mathcal{T}), k$)
            $bel \leftarrow$ **GP**($\mathcal{T}, K_{AP}(\cdot,\cdot)$)
            $\hat{\mathcal{A}} \leftarrow$ **MCTS**(S_t, bel, B_{t+1})
        $S_{t+1} \leftarrow S_t + \hat{\mathcal{A}}_t$
        $t \leftarrow t + 1$

V. Results

A. Datasets

We tested our approach using datasets collected from three coral reef sites on the island of St. John, U.S. Virgin Islands. The sites are referred to as Tektite, Yawzi, and Cocoloba reefs. Local biological activity among the three sites is known to vary: Tektite reef has the highest levels of hard coral cover (25%) and the highest fish abundances determined from manual surveys, Yawzi has the second highest (15% coral cover) and Cocoloba the lowest (9% coral cover). Time-series acoustic data were collected using SoundTrap ST-300 single-channel acoustic recorders that sampled the reefs on a 10% duty cycle, with one minute of audio recorded every 10 minutes, between November 2019 - August 2020. Recordings from the months of January, March, May, and July 2020 were calibrated post-hoc according to the manufacturer's specifications with hydrophone sensitivities ranging from 176.4 – 178 dB re 1 V/μPa. The calibrated data were subsequently filtered to 100-200 Hz using an 8th order Butterworth bandpass filter and SPL$_{rms}$ values were calculated across each 1 min file. Acoustic activity at all three sites was variable over diel timescales. Dawn and dusk peaks in low-frequency sound were evident and thought to be reflective of increased crepuscular calling activity by soniferous fish. Each reef time series has been divided into four one-month datasets with activity averaged over one-hour intervals, used for statistical analysis.

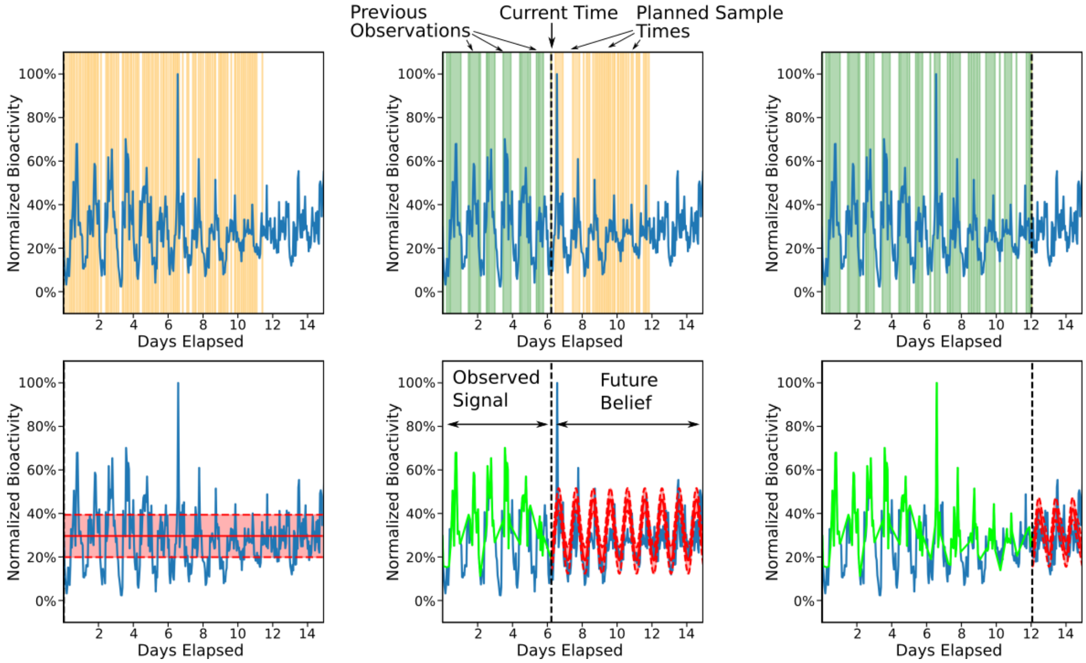

Fig. 3: Demonstration of the OSS planner on the Jan 2020 Tektite Dataset. APGP model uses 1 frequency component. Times shown are the initial belief and plan T = 0, the plan at 50% of budget expended, and the final sensing schedule. The upper images show the belief over time. Initially, with only one datapoint, we just predict a constant function, however we quickly learn the periodicity and use it to produce plans that target high-value peaks.

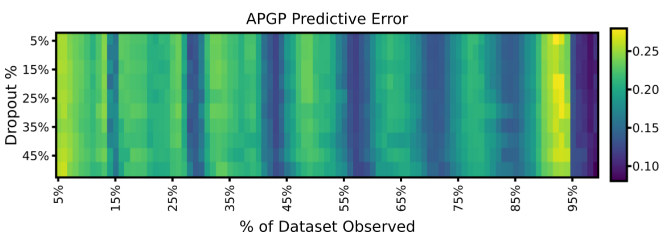

Fig. 4: Root Mean Square Error (RMSE) in units of normalized bioactivity for the APGP prediction on the January 2020 Tektite Reef dataset. As additional data is used to train the APGP (horizontal axis), the prediction accuracy decreases but with a periodic trend. The peak in RMSE between 90 and 95% data used is due to the anomalous spike in activity observed on January 7th, which can be seen in Fig. 2.

B. APGP Predictor

The first set of experiments we performed focused on evaluating the performance of the Auto-Periodic kernel Gaussian Process (APGP). Particularly, we were interested in quantifying the predictive accuracy of the APGP with varying amounts of training data. The results from this experiment are shown in Fig. 4. The APGP lengthscale and process variance hyperparameters, σ^2, and ℓ in Equation 7 were both set to 1, and $k = 1$ frequency component was used. The evaluation metric is Root Mean Square Error (RMSE) on the portion of the dataset not used for training. Intuitively, the overall error decreases as more data is used for training, particularly within as $\mathcal{T}$ grows to contain the first period of the bioacoustic signal. However, one unexpected result from this experiment is a periodic component to this decrease. This trend is a result of the periodogram analysis. Since spectral analysis assumes an infinitely repeating periodic signal, we observe the local minima in error when the observed portion of the dataset is an integer multiple of the dataset's period.

In the same experiments, we also tried randomly removing some portion of the dataset, varying from 0% to 50% of the observations. The intent of this was to simulate the loss of data due to the intermittent sampling in the OSSP planner. However, we found that this random dropout had little to no effect on the overall predictive accuracy. We attribute this to the fact that even with 50% of the observations removed, the signal is still relatively densely sampled compared to the true period. This fact allows the Lomb-Scargle periodogram to correctly identify the true periodicities in the signal, and maintain high predictive accuracy.

C. OSSP Performance

The second set of experiments we performed were to evaluate the performance of the entire Online Sensor Scheduling framework. For these experiments, results are aggregated over all 12 of the calibrated USVI acoustic datasets. We compare the performance of our approach to the standard approach of continuous sampling on two metrics: the percentage of total acoustic activity observed by the robot and the total deployment time before the battery is run dry. Unless otherwise noted, we simulate a battery with enough capacity for 144 hours (6 days) of continuous sampling. We use a sleep cost ratio of $C_{slp}/C_{obs} = 2\%$, and stop the MCTS planner after 10,000 iterations.

We compared the performance of our Auto Periodic Kernel using the Lomb-Scargle periodogram, to a kernel constructed with the Fourier periodogram using the FFT algorithm as described in [32]. On average, the planner using the Lomb-Scargle belief outperformed the Fourier kernel and the baseline approach on the acoustic activity metric with an average improvement of 21% over the baseline and 19% over the Fourier Kernel. This is a result of the emergent interaction between the reward maximization sampling policy given

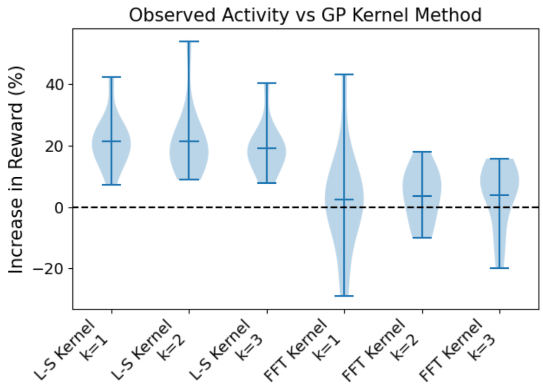

Fig. 5: Performance increase over the baseline method for our Online Sensor Scheduling Algorithm for varying numbers of frequency components (k). Due to the difficulty of accurately identifying the periodic components with missing observations, the Fourier Periodogram APGP does not produce a significant increase in observation quality.

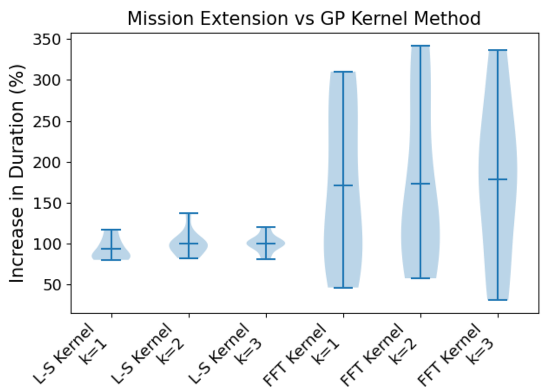

Fig. 6: Increase in observation time using our Online Sensor Scheduling algorithm over the baseline. The Fourier Periodogram APGP tended to over-sparsely sample the bioactivity, leading to much longer mission times, but at the cost of reduced reward.

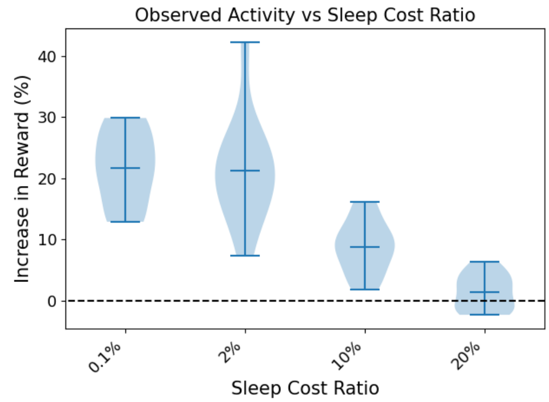

Fig. 7: Effects of varying the sleep cost ratio on observed biological activity. As the ratio decreases, there is an increase in the average amount of biological activity observed.

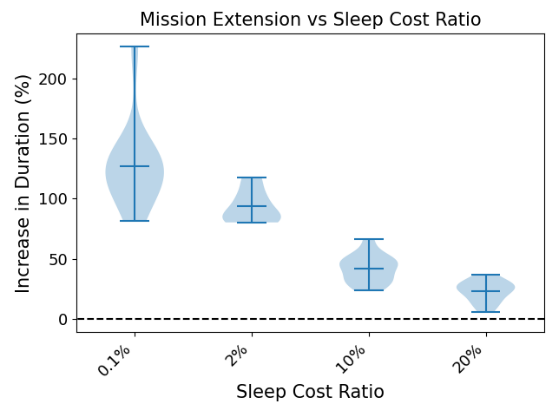

Fig. 8: Effects of varying the sleep cost ratio on mission duration. As the ratio decreases, there is an increase in the average mission length.

perfect knowledge of the signal, and the optimal sampling policy for resolving the signal for Fourier analysis. To resolve a signal with period T, we need to sample above the Nyquist Frequency of $2/T$. However, to maximize the POMDP reward, we want to only sample at the peaks of the signal, a period of T. This makes it impossible to actually resolve the true periodicity using Fourier analysis, leading to poor belief estimates. While the Lomb-Scargle approach does also suffer from these competing needs, its least-squares approach is better able to leverage a small set of dense sampling to fit the overall periodic signal. Both signals achieved intermittent sampling behavior, with the Lomb-Scargle kernel doubling the sampling mission duration, and the Fourier kernel increasing it by 175%. However, we found that the Fourier approach samples over-sparsely, which leads to longer mission durations, but no corresponding increase in observed activity over the baseline.

We also compared the Lomb-Scargle Auto Periodic Kernel with the Fourier kernel across different numbers of frequency components. As the results in Fig. 5 and Fig. 6 show, there is no significant difference in performance on either metric with increasing number of frequency components. There are several possible explanations for this. First, the signals that we consider contain a single main frequency component. The periodograms from Fig. 2 show that the daily cycle contains over an order of magnitude more power than the other frequency components. The smoothing of the predicted signal by the GP lengthscale and variance parameters is sufficient to cover the observation noise in the true signal, resulting in high-quality predictions. The second possible explanation is that any improvements in predictive accuracy are lost due to overfitting to a noisy training signal.

In the experiments discussed so far, we have assumed a sleep cost ratio of 2% relative to the cost of observing. However different sensors are optimized differently for low power states. To evaluate the impact of this ratio on the ability of the OSS planning algorithm to select high-value observation windows, we evaluated our sensor framework at different ratios. As the results in Fig. 7 and Fig. 8 show, as this ratio decreases (i.e. as sleeping becomes cheaper relative to sensing), we can spend more time sleeping to delay observations to future high-value times, leading to longer sampling missions and overall higher observed biological activity. This result is a consequence of the direct tradeoff between sensing and observing enforced by the knapsack constraint of the robot's battery.

VI. Conclusion and Future Work

In this paper we have presented a method for extending the lifespan of a sensor while increasing the quality of its observations. We apply this approach to improve the monitoring of a diverse and vital, yet imperiled ecosystem: coral reefs. Our approach identifies frequency components of a time series, with no requirements for prior data, to predict the series' future states. This prediction is used by a non-myopic planning algorithm to solve the optimal stopping problem, scheduling sensor observations in high-value time periods, and sending the robot to sleep in low-value periods. We demonstrate our method using acoustic datasets collected in the US Virgin Islands. Future work involves adding spatial dimensions to our approach. This extension would allow us to predict the time and place of biological activity, enabling a mobile sensor platform, like an AUV to pre-position itself to observe activity, minimizing its impact on the reef. Additionally, we plan to deploy this algorithm on VASE during upcoming field experiments.

References

[1] H. C. Woithe *et al.*, "Slocum glider energy measurement and simulation infrastructure," in *proc. IEEE/MTS OCEANS Conference, Sydney*, 2010, pp. 1–8.

[2] K. Smith Jr, A. Sherman, P. McGill, R. Henthorn, J. Ferreira, T. Connolly, and C. Huffard, "Abyssal benthic rover, an autonomous vehicle for long-term monitoring of deep-ocean processes," *Science Robotics*, vol. 6, no. 60, 2021.

[3] N. Knowlton, R. E. Brainard, R. Fisher, M. Moews, L. Plaisance, and M. J. Caley, "Coral reef biodiversity," *Life in the world's oceans: diversity distribution and abundance*, pp. 65–74, 2010.

[4] T. P. Hughes *et al.*, "Climate change, human impacts, and the resilience of coral reefs," *Science*, vol. 301, no. 5635, pp. 929–933, 2003.

[5] O. Hoegh-Guldberg *et al.*, "Coral reefs under rapid climate change and ocean acidification," *Science*, vol. 318, no. 5857, pp. 1737–1742, 2007.

[6] M. Adjeroud *et al.*, "Recurrent disturbances, recovery trajectories, and resilience of coral assemblages on a south central pacific reef," *Coral reefs*, vol. 28, no. 3, pp. 775–780, 2009.

[7] M. B. Kaplan, T. A. Mooney, J. Partan, and A. R. Solow, "Coral reef species assemblages are associated with ambient soundscapes," *Marine Ecology Progress Series*, vol. 533, pp. 93–107, 2015.

[8] T. A. Mooney *et al.*, "Listening forward: approaching marine biodiversity assessments using acoustic methods," *Royal Society open science*, vol. 7, no. 8, p. 201287, 2020.

[9] L. Zhao, W. Zhang, J. Hu, A. Abate, and C. J. Tomlin, "On the optimal solutions of the infinite-horizon linear sensor scheduling problem," *IEEE Transactions on Automatic Control*, vol. 59, no. 10, pp. 2825–2830, 2014.

[10] National Research Council *et al.*, *Science at Sea: Meeting Future Oceanographic Goals with a Robust Academic Research Fleet*. National Academy Press, Washington, DC, 2009.

[11] P. S. Lobel, "Sounds produced by spawning fishes," *Environmental Biology of Fishes*, vol. 33, no. 4, pp. 351–358, 1992.

[12] T. C. Tricas and K. S. Boyle, "Parrotfish soundscapes: implications for coral reef management," *Marine Ecology Progress Series*, vol. 666, pp. 149–169, 2021.

[13] L. Van Oosterom, J. Montgomery, A. Jeffs, and C. Radford, "Evidence for contact calls in fish: conspecific vocalisations and ambient soundscape influence group cohesion in a nocturnal species," *Scientific Reports*, vol. 6, no. 1, pp. 1–8, 2016.

[14] S. L. Nedelec, S. D. Simpson, M. Holderied, A. N. Radford, G. Lecellier, C. Radford, and D. Lecchini, "Soundscapes and living communities in coral reefs: temporal and spatial variation," *Marine Ecology Progress Series*, vol. 524, pp. 125–135, 2015.

[15] J. Das *et al.*, "Data-driven robotic sampling for marine ecosystem monitoring," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1435–1452, 2015.

[16] G. Best, W. Martens, and R. Fitch, "Path planning with spatiotemporal optimal stopping for stochastic mission monitoring," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 629–646, 2017.

[17] C. Y. H. Lee, G. Best, and G. A. Hollinger, "Optimal sequential stochastic deployment of multiple passenger robots," in *proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 8934–8940.

[18] L. Moser, "On a problem of cayley," *Scripta Math*, vol. 22, pp. 289–292, 1956.

[19] P. Freeman, "The secretary problem and its extensions: A review," *International Statistical Review/Revue Internationale de Statistique*, pp. 189–206, 1983.

[20] Y. Girdhar and G. Dudek, "Optimal online data sampling or how to hire the best secretaries," in *proc. Canadian Conference on Computer and Robot Vision*. IEEE, 2009, pp. 292–298.

[21] M. Bateni, M. Hajiaghayi, and M. Zadimoghaddam, "Submodular secretary problem and extensions," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 2010, pp. 39–52.

[22] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg, "A knapsack secretary problem with applications," in *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*. Springer, 2007, pp. 16–28.

[23] G. Flaspohler, N. Roy, and Y. Girdhar, "Near-optimal irrevocable sample selection for periodic data streams with applications to marine robotics," in *proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 5691–5698.

[24] K.-C. Ma, L. Liu, H. K. Heidarsson, and G. S. Sukhatme, "Data-driven learning and planning for environmental sampling," *Journal of Field Robotics*, vol. 35, no. 5, pp. 643–661, 2018.

[25] S. McCammon *et al.*, "Ocean front detection and tracking using a team of heterogeneous marine vehicles," *Journal of Field Robotics*, vol. 38, no. 6, pp. 854–881, 2021.

[26] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer school on machine learning*. Springer, 2003, pp. 63–71.

[27] N. Durrande, J. Hensman, M. Rattray, and N. D. Lawrence, "Detecting periodicities with gaussian processes," *PeerJ Computer Science*, vol. 2, p. 50, 2016.

[28] R. Marchant and F. Ramos, "Bayesian optimisation for intelligent environmental monitoring," in *proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 2242–2249.

[29] L.-F. Cheng, B. Dumitrascu, G. Darnell, C. Chivers, M. Draugelis, K. Li, and B. E. Engelhardt, "Sparse multi-output gaussian processes for online medical time series prediction," *BMC medical informatics and decision making*, vol. 20, no. 1, pp. 1–23, 2020.

[30] H. J. Nussbaumer, "The fast fourier transform," in *Fast Fourier Transform and Convolution Algorithms*. Springer, 1981, pp. 80–111.

[31] A. Wilson and R. Adams, "Gaussian process kernels for pattern discovery and extrapolation," in *proc. International Conference on Machine Learning*. PMLR, 2013, pp. 1067–1075.

[32] A. Tompkins and F. Ramos, "Fourier feature approximations for periodic kernels in time-series modelling," in *proc. AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

[33] S. Russell and P. Norvig, "Artificial intelligence: a modern approach," 2002.

[34] E. Staaterman *et al.*, "Bioacoustic measurements complement visual biodiversity surveys: preliminary evidence from four shallow marine habitats," *Marine Ecology Progress Series*, vol. 575, pp. 207–215, 2017.

[35] N. R. Lomb, "Least-squares frequency analysis of unequally spaced data," *Astrophysics and space science*, vol. 39, no. 2, pp. 447–462, 1976.

[36] J. D. Scargle, "Studies in astronomical time series analysis," *The astrophysicak Journal*, vol. 263, pp. 835–853, 1982.

[37] J. Griffié, L. Boelen, G. Burn, A. P. Cope, and D. M. Owen, "Topographic prominence as a method for cluster identification in single-molecule localisation data," *Journal of biophotonics*, vol. 8, no. 11-12, pp. 925–934, 2015.

[38] C. B. Browne *et al.*, "A survey of monte carlo tree search methods," *IEEE Transactions on Computational Intelligence and AI in games*, vol. 4, no. 1, pp. 1–43, 2012.