Optimal Adaptive Sampling for Boundary Estimation with Mobile Sensors

Phillip Kearns, John Lipor
Department of Electrical & Computer Engineering
Portland State University
Email: {kearns,lipor}@pdx.edu

Bruno Jedynak
Department of Mathematics & Statistics
Portland State University
Email: bruno.jedynak@pdx.edu

***Abstract*—We consider the problem of active learning in the context of spatial sampling for boundary estimation, where the goal is to estimate an unknown boundary as accurately and quickly as possible. We present a finite-horizon search procedure to optimally minimize both the final estimation error *and* the distance traveled for a fixed number of samples, where a tuning parameter is used to trade off between the estimation accuracy and distance traveled. We show that the resulting optimization problem can be solved in closed form and that the resulting policy generalizes existing approaches to this problem.**

I. Introduction

A fundamental challenge to modern science and engineering is the ability to rapidly and accurately sense the environment. Harmful algae blooms impair access to drinking water [1], traffic-related pollutants impact urban health [2], and wildfires present a persistent threat to safety and air quality in the western United States [3]. We consider the third example as a motivating problem, where our goal is to determine the spatial extent of hazardous particulate matter from wildfire smoke (see Fig. 1). In this case, all points with pollutant concentration above or below a fixed threshold can be considered as two classes in a binary classification problem, where the goal is to estimate the decision boundary as quickly as possible. Further, we consider the case where such measurements are obtained by a mobile sensor such as an unmanned aerial vehicle.

Algorithms designed to rapidly determine the decision boundary fall within the category of *active learning* or *adaptive sampling* [4, 5] and typically try to maximize information gain per sample. However, in the above example, there is a significant cost associated with both the time to take a single sample *and* the distance traveled throughout the sampling procedure. Hence, standard approaches to active learning based in search space reduction [6–8] or adaptive submodularity [9], which seek to minimize only the number of samples taken, will be accompanied by potentially dramatic drawbacks in terms of total sampling cost. Newer, bisection-style search methods like quantile search (QS) [10, 11] balance the above costs by sampling a certain fraction into the remaining interval at each step, effectively trading off between number of samples and distance traveled. Though these improve upon previous methods in terms of total sampling time, neither guarantees to find the optimal search procedure.

In this work, we propose a finite-horizon (FH) sampling procedure that optimally balances the distance traveled during the search with the final interval size after obtaining N measurements. We show that for a one-dimensional step function, fixing N allows the resulting cost to be optimized in closed form, eschewing the need for dynamic programming. The work of [10] shows we can combine a series of one-dimensional search procedures to estimate a two-dimensional boundary. We also derive the expected number of samples necessary to fall below a given final estimation error and the best policy for various search time parameters. Further, we show that QS can be viewed as an instance of the proposed FH algorithm in the case where $N = 1$ (greedy sampling). Empirical results demonstrate that our sampling strategy outperforms existing approaches and agrees with our analytical predictions in terms of the resulting distance traveled and average interval size.

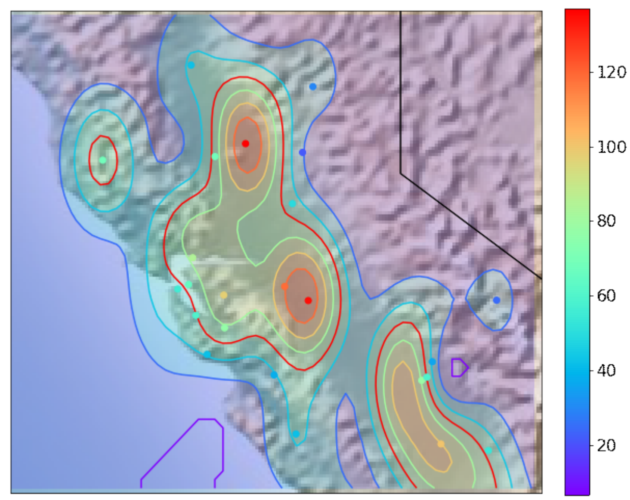

Fig. 1: Map of air quality following California's Camp Fire. Points represent air quality measurement stations and contours are generated with a Gaussian regression. The red contour represents a PM2.5 estimate of greater than 60 $\mu g/m^3$ as a decision boundary.

II. Problem Formulation & Related Work

As stated in the introduction, the full two-dimensional boundary estimation can be reduced to a series of one-dimensional search problems, where we wish to locate the change point of a one-dimensional step function, i.e., a function from the class

$$\mathcal{F} = \{f_\theta : [0,1] \to \mathbb{R} : f_\theta(x) = \mathbb{1}_{[0,\theta)}(x), \theta \in [0,1]\},$$

where $\mathbb{1}_S(x)$ denotes the set indicator function. These one-dimensional estimates may then be combined either in a piecewise-linear fashion [10] or using Gaussian process regression [12] as illustrated in Fig. 2.

Assume we obtain observations $\{Y_n\}_{n=1}^N \in \{0,1\}^N$ from the sample locations $\{X_n\}_{n=1}^N$ in the unit interval in a sequential fashion according to $Y_n = f_\theta(X_n)$, where θ is the actual, unknown, change point location. Under this model, each sample obtained reduces the interval (hypothesis space) in which the change point may lie. Our goal is then to estimate the change point location while minimizing the sampling cost for a fixed number of samples, a function of both the final expected interval size *and* expected distance traveled.

A. Related Work

Many previous approaches to finding an unknown change point are based in search space reduction [6–8] and do not permit the inclusion of general or dynamic costs such as the time between sampling points or potential need for recharging. Because these methods do not take into account these extra parameters, they tend to result in bisection-type solutions [13] that will have higher total sampling cost. Methods that seek to maximize hypothesis space reduction at each step can be classified as "greedy" search methods. Greedy methods in active learning [6, 7] lack theoretical guarantees of minimum total sampling cost, and even those that incorporate realistic costs into the algorithm formulation [14] have been shown to perform worse than the QS algorithm in [10] when applied to distance-penalized searches.

A popular greedy approach to active learning relies on the concept of adaptive submodularity (AS). AS is a diminishing returns principle that states samples are more informative early on in the search procedure, and [9] shows that a greedy procedure is optimal up to a constant factor. However, AS is a property of set functions and does not consider a sequential dependency among sampling locations. While [15] provides a theoretical analysis of greedy active learning with non-uniform costs, the authors only consider the case of query costs being fixed. In contrast, our scenario has non-uniform and dynamic costs which depend on the distance between points.

In [16], the authors introduce the idea of adaptive data collection for mobile path planning, where previous samples are used to guide the motion of the sensing vehicles for further sampling. An interesting extension is provided in [17], where the authors consider the constraint that the robot has *limited energy*. However, both works require a coarse sampling of the entire feature space to estimate a scalar field, which is not applicable to our problem. Though the energy constraint is insightful, the authors mix a network of stationary sensors with a mobile sensor to minimize integrated error, whereas we focus on minimizing the final boundary error using only a single mobile sensor.

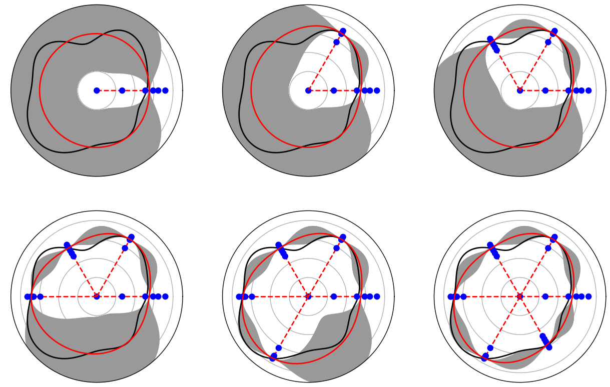

Fig. 2: Two-dimensional boundary estimate from series of one-dimensional searches. Each search is performed using FH sampling (blue dots). The black line represents the true boundary and the final estimate (red line with gray confidence bounds) is obtained using Gaussian process regression with a periodic squared exponential kernel.

Of primary relevance to the work presented in this paper is the work of [10], which introduces the Quantile Search (QS) algorithm for determining the change point of a one-dimensional step function while balancing sampling and travel costs. QS is a generalization of binary bisection [13, 18, 19], where the idea is that by successively sampling a fixed fraction $1/m$ into the remaining hypothesis space (defined by an interval), the desired tradeoff between number of samples and distance traveled can be achieved. This work was extended in [11], where the key observation is that QS can be improved by allowing the fraction $1/m$ to grow as the hypothesis space shrinks. Yet, neither algorithm provides guarantees of optimality in terms of the total sampling cost. We believe that this work is the first to provide a theoretical guarantee of optimal search procedure for an environment with non-uniform, dynamic sampling costs.

III. Proposed Algorithm & Analysis

It is convenient, while not restrictive, to define search strategies in terms of the *fraction of the remaining interval* to move at each step, whether forward or backward, in an analogous fashion to [10, 11]. The resulting class of policies is adaptive to the unknown location of θ and non-restrictive in the sense that any optimal policy will not sample in locations with probability zero (locations outside the remaining interval).

Begin with a uniform prior on the change point θ, and let the N fractions be $\{x_n\}_{n=1}^N$. A straightforward Bayesian update yields the posterior distribution after each sample. Let H_N be the entropy of the posterior distribution after N observations, D_N be the total distance traveled, and $\lambda > 0$ be a tuning parameter that governs the tradeoff between these costs. We define the total sampling cost, J, after N observations

$$J(x_1, \ldots, x_N) = \mathbb{E}_\theta \left[e^{H_N} + \lambda D_N \right]. \quad (1)$$

Note that for a uniform distribution on an interval of length a, $e^{H_N} = e^{\log(a)} = a$; thus, (1) is equivalent to minimizing a weighted combination of the (expected) final interval length and expected distance traveled.

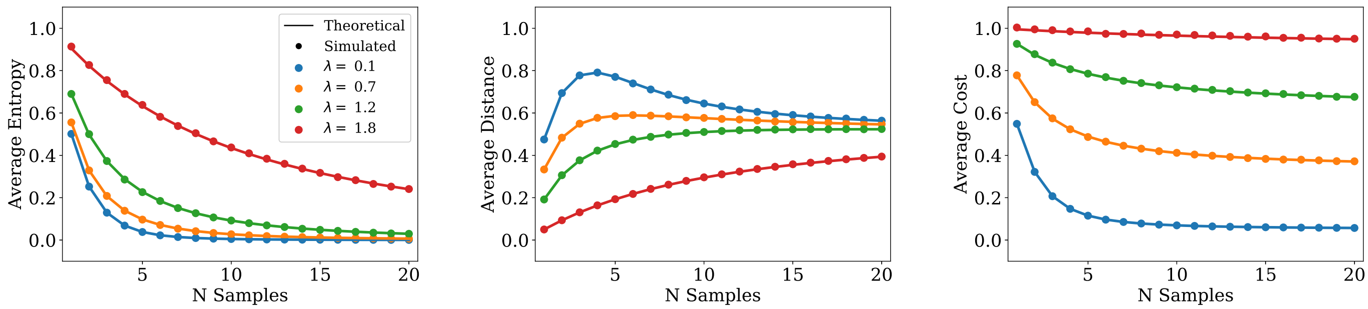

Fig. 3: Performance of proposed FH algorithm for spatial sampling for fixed N samples. Each data point represents the final value obtained after following the optimal N-step policy. Left-to-right: average entropy of hypothesis space, average distance traveled, and average total cost after last sample.

A. Analysis of Sampling Cost

We now demonstrate that the cost (1) can be minimized in closed form, resulting in an optimal policy that subsumes the QS algorithm. The theorem below provides a closed-form analysis of the cost (1). This simplified form allows us to compute the optimal sampling fractions in linear time.

Theorem 1. *Let $\lambda \in [0,2]$ and assume the change point has distribution $\theta \sim$ Unif$([0,1])$. Further, assume the N measurements are defined via N fractions $x_1,\dots,x_N$ denoting the proportion of the current hypothesis space to sample. Define*

$$\xi_i = x_i^2 + (1-x_i)^2, \quad i = 1,\dots,N.$$

Then,

$$J(x_1,\dots,x_N) = \prod_{i=1}^{N}\xi_i + \lambda\sum_{i=1}^{N} x_i \prod_{j=0}^{i-1}\xi_j. \tag{2}$$

Proof. By Lemma 1 below, we have that

$$\mathbb{E}\left[e^{H_N}\right] = \prod_{i=1}^{N}\xi_i.$$

Let D_N be the distance traveled after N samples. Note that

$$D_n = \sum_{i=1}^{n} x_i e^{H_{i-1}}.$$

Therefore

$$\mathbb{E}\left[D_N\right] = \mathbb{E}\left[\sum_{i=1}^{N} x_i e^{H_{i-1}}\right] = \sum_{i=1}^{N} x_i \mathbb{E}\left[e^{H_{i-1}}\right].$$

Applying Lemma 1 then yields

$$\mathbb{E}\left[D_N\right] = \sum_{i=1}^{N} x_i \prod_{j=0}^{i-1}\xi_j.$$

The proof is completed by applying linearity of expectation. □

Lemma 1. *Let H_N be the length of the hypothesis space after N measurements. Under the conditions of Theorem 1, we have*

$$\mathbb{E}\left[e^{H_N}\right] = \prod_{i=1}^{N}\left(x_i^2 + (1-x_i)^2\right). \tag{3}$$

Proof. First note that under the uniform distribution on the unit interval, the exponentiated differential entropy is the length of the hypothesis space after N samples. The proof will proceed by induction on N. Consider the base case, $N = 1$, for which it is trivial to show that

$$\mathbb{E}\left[e^{H_1}\right] = x_1^2 + (1-x_1)^2 = \xi_1.$$

Now assume that (3) holds for some $N \in \mathbb{N}$. Sampling some fraction x_{N+1} into the remaining hypothesis space e^{H_N} results in two potential entropies

$$e^{H_{N+1}} = \begin{cases} x_{N+1}e^{H_N} & \text{w/ probability} \quad x_{N+1} \\ (1-x_{N+1})e^{H_n} & \text{w/ probability} \quad 1-x_{N+1}. \end{cases}.$$

Therefore

$$\begin{aligned} \mathbb{E}[e^{H_{N+1}}] &= x_{N+1}^2\mathbb{E}[e^{H_N}] + (1-x_{N+1})^2\mathbb{E}[e^{H_N}] \\ &= \left(x_{N+1}^2 + (1-x_{N+1})^2\right)\mathbb{E}[e^{H_N}] \\ &= \prod_{i=1}^{N+1}\left(x_i^2 + (1-x_i)^2\right). \end{aligned}$$

□

Thm. 1 shows that the entropy and distance components of the sampling cost can both be written in terms of the expected reduction in interval size ξ_i. This form allows us to compute the minimum of (1) in closed form.

Theorem 2. *Under the same conditions as Thm. 1, the optimal sampling fractions are of the form*

$$x_k = \frac{1}{2} - \lambda\frac{1}{4\rho_k}, \quad k = 1,\dots,N. \tag{4}$$

where $\rho_N = 1$ and

$$\rho_k = \prod_{i=k+1}^{N}\xi_i + \lambda\sum_{i=k+1}^{N} x_i \prod_{j=k+1}^{i-1}\xi_j, \quad k = 1,\dots,N-1,$$

depends only on the fractions $x_{k+1}, \ldots, x_N$.

Proof. The gradient of the total cost (2) is

$$\frac{\partial J}{\partial x_l} = \left(\prod_{i=1}^{l-1} \xi_i\right) \left(\left(4x_l - 2\right)\rho_l + \lambda\right),$$

and setting the gradient to zero yields:

$$x_l = \frac{1}{2} - \lambda \frac{1}{4\rho_l}.$$

□

Thm. 2 shows further that the optimal policy may be computed in linear time, beginning with x_N and proceeding backwards. Further, taking $N = 1$ results in a policy that samples a constant fraction of the remaining interval at each step, which is exactly the strategy of the QS algorithm. Hence, QS may be considered an instance of our proposed method with $N = 1$. When $\lambda \geq 2$, the cost of performing even a single non trivial sample, λx_1, is larger than the expected entropy reduction, $1 - \xi_1$; thus the trivial sample which requires no displacement is preferred.

B. Samples Needed for Fixed Estimation Error

By Lemma 1, we have the expected length of the hypothesis space after N samples. In certain instances, it is desirable to use a threshold on the final interval size to terminate the search procedure. When this is the case, we apply the results of Theorem 2 and eq. (4) to calculate the optimal fraction for the final step and then proceed backwards, calculating the fractional policy at each preceding step until an interval smaller than the desired final error is expected. Pseudocode for finding the optimal policy starting with a given interval of length L and subject to a desired final estimation error is given in Algorithm 1.

Algorithm 1 Calculating Policy for Expected Convergence

```
Input: interval length L, penalty λ, stopping error ε
Initialize: x_N ← 1/2 − λ/4, l ← 1
while L ∏_i ξ_i > ε do
    x_{N−l} ← 1/2 − λ/(4ρ_{N−l})
    l ← l + 1
end while
N ← l
```

C. Search Procedure Description

In the case where a search terminates only after a certain estimation error has been obtained, we follow a two-phase procedure. Before the search begins, we use the method presented in Algorithm 1 to calculate the N steps such that the expected interval size is less than ε. In the first stage of the search, samples are taken according to this optimal N-step policy. If the desired error has been met before all samples have been taken, the search terminates. However, if the desired error has not been met after these N samples, the algorithm then performs a greedy search (optimal 1-step policy, line 7) until the interval is sufficiently small. Pseudocode is provided in Algorithm 2.

Algorithm 2 Finite Horizon Search

```
1:  Input: policy x, stopping error ε
2:  Initialize: X_0 ← 0, Y_0 ← 1, a ← 0, b ← 1, n ← 1
3:  while b − a > ε do
4:      if n ≤ N then
5:          x ← x_n
6:      else
7:          x ← 1/2 − λ/4
8:      end if
9:      if Y_{n−1} = 1 then
10:         X_n ← X_{n−1} + x(b − a)
11:     else
12:         X_n ← X_{n−1} − x(b − a)
13:     end if
14:     Y_n ← f(X_n)
15:     a = max {X_i : Y_i = 1, i ≤ n}
16:     b = min {X_i : Y_i = 0, i ≤ n}
17:     θ̂_n ← (a+b)/2
18:     n ← n + 1
19: end while
```

IV. Simulations

In this section, we verify the performance of the proposed finite-horizon sampling policy. To obtain a profile of performance as a function of λ, we range θ over 100 uniformly-spaced values in the interval $(0, 1)$ for 5 values of $\lambda \in [0.1, 1.8]$. Fig. 3 shows the resulting average entropy, distance traveled, and final cost for each corresponding N-step policy. The figures demonstrate that our proposed method indeed obtains a tradeoff between average final entropy and distance traveled via the tuning parameter λ. Further, comparing with the expected entropy and distance calculated in Section III-A, we see that our theoretical results closely match the empirical values obtained. Additionally, we see that as the number of samples increases the total cost decreases. While our cost function trades off final entropy against distance traveled (and thus prefers a longer procedure in which more, less aggressive samples will result in both less error *and* less potential overshoot), another potential cost function to minimize is the total sampling *time*.

A. Minimizing Sampling Time

If we seek to minimize the total time that a vehicle takes to complete a search, we need to consider a cost function of the form

$$J_T(x_1, \ldots, x_N) = T_s N + T_t D.$$

where T_s and T_t represent the time per sample and time per unit distance travelled. In order to minimize this cost in expectation, we follow a three-step procedure. First, calculate the number of samples, N_λ, and total distance, D_λ, expected

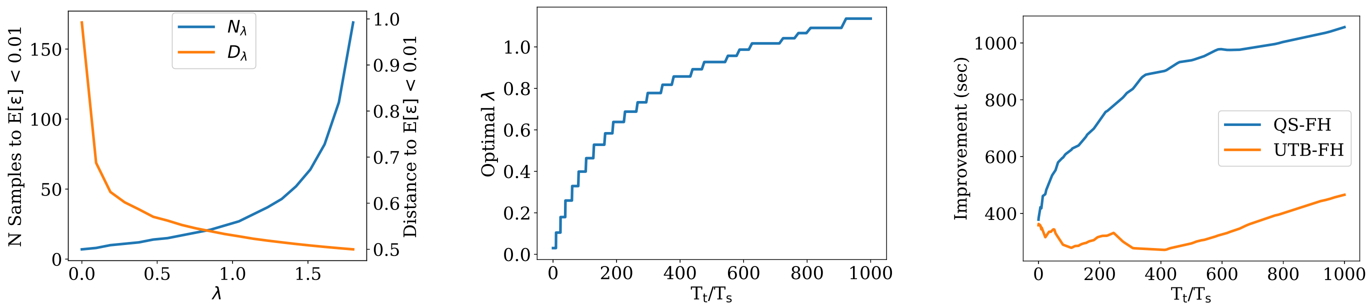

Fig. 4: Components of the time-penalized FH search procedure. Left: N_λ and D_λ to reach an interval of 0.01 for each value of λ. Middle: values of λ^* for each ratio of T_t/T_s. Right: average improvement of FH search over QS and UTB algorithms.

for each value of λ to converge below a desired error ε. Then, for every ratio of T_t/T_s, select λ^* according to

$$\lambda^* = \arg\min_{\lambda} T_s N_\lambda + T_t D_\lambda.$$

Finally, perform the N-step FH search for the selected value of λ according to section III-C and Algorithm 2.

We compare the performance of this method with the existing QS [10] and Uniform-to-Binary search (UTB) [11] algorithms. We consider the same grid over θ for 1000 different ratios of T_t/T_s in the range of 1×10^{-4} to 1×10^{3} with $T_s = 100$ as the base sampling cost. Fig. 4 shows the resulting improvement in sampling time obtained via the proposed FH policy. When the T_t/T_s ratio is small and the total search time is short, we see an improvement of approximately 34% over both UTB and QS. At the highest ratios of T_t/T_s where the search is much longer, the relative improvement decreases to roughly 1-2%.

V. Conclusion

We have presented a novel active learning algorithm for spatial sampling that optimally balances the final estimation error and the distance traveled for a fixed number of samples. We have derived the closed-form solution and to the best of our knowledge, believe that this work is the first to provide a theoretical guarantee of optimal search procedure for an environment with non-uniform, dynamic sampling costs. We have also shown how our solution generalizes existing approaches to this problem, and empirical results indicate the performance benefits of finite-horizon search over existing methods in the literature.

A number of open questions remain. In this work, we have considered only the case of noiseless measurements drawn from a uniform distribution. Extending to noisy measurements as done in [10] and non-uniform priors are important next steps. While the search parameter λ allows for various search costs (sampling time, travel time, recharging, etc.) to be approximated in a compact state space notation, we have only provided a formula for converting travel and sample time into our notation. Finally, by following [10, 18] we have generalized a one-dimensional search to a two-dimensional problem, but optimal methods for two-dimensional search continues to be an interesting problem.

Acknowledgment

J. Lipor was supported by NSF CRII award CCF-1850404.

References

[1] Oregon Health Authority. (2018, Jun.) Harmful algae blooms, environmental public health. [Online]. Available: https://www.oregon.gov/OHA/PH/HealthyEnvironments/Recreation/HarmfulAlgaeBlooms/pages/index.aspx

[2] E. T. Gall, L. A. George, R. B. Cal, and A. Laguerre, "Indoor and outdoor air quality at harriet tubman middle school and the design of mitigation measures: Phase i report," Portland State University, Tech. Rep., 2018.

[3] State of Oregon. (2018, Aug.) Fire information and statistics. [Online]. Available: https://www.oregon.gov/ODF/Fire/pages/FireStats.aspx

[4] B. Settles, *Active Learning*. Morgan & Claypool, 2012.

[5] R. Castro and R. Nowak, "Active learning and sampling," in *Foundations and Applications of Sensor Management*, 1st ed. New York, NY: Springer, 2008, ch. 8.

[6] R. Nowak, "Generalized binary search," in *Proc. Allerton Conference on Communication, Control, and Computing*, 2008.

[7] S. Dasgupta, "Analysis of a greedy active learning strategy," in *Proc. Advances in Neural Information Processing Systems*, 2005.

[8] R. Willett, R. Nowak, and R. M. Castro, "Faster rates in regression via active learning," in *Advances in Neural Information Processing Systems*, 2006, pp. 179–186.

[9] D. Golovin and A. Krause, "Adaptive submodularity: Theory and applications in active learning and stochastic optimization," *Journal of Artificial Intelligence Research*, vol. 42, pp. 427–486, 2011.

[10] J. Lipor, B. P. Wong, D. Scavia, B. Kerkez, and L. Balzano, "Distance-penalized active learning using quantile search," *IEEE Transactions on Signal Processing*, vol. 65, no. 20, pp. 5453–5465, 2017.

[11] J. Lipor and G. Dasarathy, "Quantile search with time-varying search parameter," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2018, pp. 1016–1018.

[12] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT Press Cambridge, MA, 2006, vol. 2, no. 3.

[13] M. V. Burnashev and K. S. Zigangirov, "An interval estimation problem for controlled observations," *Problems in Information Transmission*, vol. 10:223-231, 1974, translated from Problemy Peredachi Informatsii, 10(3):51-61, July-September, 1974.

[14] P. Donmez and J. G. Carbonell, "Proactive learning: cost-sensitive active learning with multiple imperfect oracles," in *Proceedings of the 17th ACM conference on Information and knowledge management*. ACM, 2008, pp. 619–628.

[15] A. Guillory and J. Blimes, "Average-case active learning with costs," in *Proc. Algorithmic Learning Theory*, 2009.

[16] A. Singh, R. Nowak, and P. Ramanathan, "Active learning for adaptive mobile sensing networks," in *Proc. Information Processing in Sensor Networks*, 2006.

[17] B. Zhang and G. S. Sukhatme, "Adaptive sampling for estimating a scalar field using a robotic boat and a sensor network," in *Proc. IEEE International Conference on Robotics and Automation*, 2007.

[18] R. Castro and R. Nowak, "Minimax bounds for active learning," *IEEE Trans. Inf. Theory*, vol. 54, pp. 2339–2353, May 2008.

[19] M. Horstein, "Sequential decoding using noiseless feedback," *IEEE Trans. Inf. Theory*, vol. 9, 1963.