

Active Learning for Level Set Estimation

Alkis Gotovos
ETH Zurich

Nathalie Casati
ETH Zurich &
IBM Research – Zurich

Gregory Hitz
ETH Zurich

Andreas Krause
ETH Zurich

Abstract

Many information gathering problems require determining the set of points, for which an unknown function takes value above or below some given threshold level. We formalize this task as a classification problem with sequential measurements, where the unknown function is modeled as a sample from a Gaussian process (GP). We propose LSE, an algorithm that guides both sampling and classification based on GP-derived confidence bounds, and provide theoretical guarantees about its sample complexity. Furthermore, we extend LSE and its theory to two more natural settings: (1) where the threshold level is implicitly defined as a percentage of the (unknown) maximum of the target function and (2) where samples are selected in batches. We evaluate the effectiveness of our proposed methods on two problems of practical interest, namely autonomous monitoring of algal populations in a lake environment and geolocating network latency.

1 Introduction

Many information gathering problems require accurately determining the regions where the value of some unknown function lies above or below a given threshold level. Moreover, evaluating the function is usually a costly procedure and the measurements returned are noisy.

As a concrete example of such an application, consider the task of monitoring a lake environment for algal bloom, a phenomenon that is potentially harmful to other organisms of the ecosystem. One way to accomplish this, is by determining the regions of the lake where the levels of algae-produced chlorophyll are above some threshold value determined by field experts. These regions can be estimated by sampling at various locations of the lake using a mobile sensing device. However, each measurement is costly in terms of time and sensor battery power, therefore the sampling locations have to be picked carefully, in order to reduce the total number of measurements required.

Other example applications in the context of environmental monitoring [Rahimi *et al.*, 2004] include estimating level sets

of quantities such as solar radiation, humidity, etc., and determining the extent of hazardous phenomena, e.g. air pollution or oil spills [Galland *et al.*, 2004]. In a different category are applications that consist in determining the subset of a parameter space that represents “acceptable” hypotheses [Bryan *et al.*, 2005] or designs [Ramakrishnan *et al.*, 2005].

We consider the problem of estimating some function level set in a sequential decision setting, where, at each time step, the next sampling location is to be selected given all previous measurements. For solving this problem, we propose the Level Set Estimation (LSE) algorithm, which utilizes Gaussian processes [Rasmussen and Williams, 2006] to model the target function and exploits its inferred confidence bounds to drive the selection process. We also provide an information-theoretic bound on the number of measurements needed to achieve a certain accuracy, when the underlying function is sampled from a Gaussian process.

Furthermore, we extend the LSE algorithm to two more settings that naturally arise in practical applications. In the first setting, we do not a priori have a specific threshold level at our disposal, but would still like to perform level set estimation with respect to an *implicit* level that is expressed as a percentage of the function maximum. In the second setting, we want to select at each step a *batch* of next samples. A reason for doing so is that, in problems such as the lake sensing example outlined above, apart from the cost of actually making each measurement, we also have to take into account the cost incurred by traveling from one sampling location to the next. Traveling costs can be dramatically reduced, if we plan ahead by selecting multiple points at a time. Another reason is that some problems allow for running multiple function evaluations in parallel, in which case selecting batches of points can lead to a significant increase in sampling throughput.

Related work. Previous work on level set [Dantu and Sukhatme, 2007; Srinivasan *et al.*, 2008] and boundary [Singh *et al.*, 2006] estimation and tracking in the context of mobile sensor networks has primarily focused on controlling the movement and communication of sensor nodes, without giving much attention to individual sampling locations and the choice thereof.

In contrast, we consider the problem of level set estimation in the setting of *pool-based active learning* [Settles, 2009], where we need to make sequential decisions by choosing sampling locations from a given set. For this problem, Bryan

et al. [2005] have proposed the *straddle* heuristic, which selects where to sample by trading off uncertainty and proximity to the desired threshold level, both estimated using GPs. However, no theoretical justification has been given for the use of straddle, neither for its extension to composite functions [Bryan and Schneider, 2008]. Garnett *et al.* [2012] consider the problem of *active search*, which is also about sequential sampling from a domain of two (or more) classes (in our case the super- and sublevel sets). In contrast to our goal of detecting the class boundaries, however, their goal is to sample as many points as possible from one of the classes.

In the setting of multi-armed bandit optimization, which is similar to ours in terms of sequential sampling, but different in terms of objectives, GPs have been used both for modeling, as well as for sample selection [Brochu *et al.*, 2010]. In particular, the GP-UCB algorithm makes use of GP-inferred upper confidence bounds for selecting samples and has been shown to achieve sublinear regret [Srinivas *et al.*, 2010]. An extension of GP-UCB to the multi-objective optimization problem has been proposed by Zuluaga *et al.* [2013], who use a similar GP-based classification scheme to ours to classify points as being Pareto-optimal or not.

Existing approaches for performing multiple evaluations in parallel in the context of GP optimization, include *simulation matching* [Azimi *et al.*, 2010], which combines GP modeling with Monte-Carlo simulations, and the GP-BUCB [Desautels *et al.*, 2012] algorithm, which obtains similar regret bounds to GP-UCB, and from which we borrow the main idea for performing batch sample selection.

To our knowledge, there has been no previous work on actively estimating level sets with respect to implicitly defined threshold levels.

Contributions. The main contributions of this paper can be summarized as follows:

- We introduce the LSE algorithm for sequentially estimating level sets of unknown functions and also extend it to select samples in batches.
- We consider for the first time the problem of estimating level sets under implicitly defined threshold levels and propose an extension of LSE for this problem.
- We prove theoretical convergence bounds for LSE and its two extensions when the target function is sampled from a known GP.
- We evaluate LSE and its extensions on two real-world datasets and show that they are competitive with the state-of-the-art.

2 Background and Problem Statement

Given a function $f : D \rightarrow \mathbb{R}$, where D is a finite subset of $\mathbb{R}^d$, and a threshold level $h \in \mathbb{R}$, we define the *level set estimation problem* as the problem of classifying every point $x \in D$ into a *superlevel set* $H = \{x \in D \mid f(x) > h\}$ and a *sublevel set* $L = \{x \in D \mid f(x) \leq h\}$.

In the *strictly sequential* setting, at each step $t \geq 1$, we select a point $x_t \in D$ to be evaluated and obtain a noisy measurement $y_t = f(x_t) + n_t$. In the *batch* setting we select

B points at a time and only obtain the resulting measurements after all of the B points have been selected.

When an explicit level is not available, we can define an *implicit threshold level* with respect to the function maximum in either absolute or relative terms. We use the relative definition in our exposition with $h = \omega \max_{x \in D} f(x)$ and $\omega \in (0, 1)$.

Gaussian processes. Without any assumptions about the function f , attempting to estimate level sets from few samples is a hopeless endeavor. Modeling f as a sample from a Gaussian process (GP) provides an elegant way for specifying properties of the function in a nonparametric fashion. A GP is defined as a collection of random variables, any finite subset of which is distributed according to a multivariate Gaussian in a consistent way [Rasmussen and Williams, 2006], and is denoted as $\mathcal{GP}(\mu(x), k(x, x'))$. It is completely specified by its mean function $\mu(x)$, which can be assumed to be zero w.l.o.g., and its covariance function or kernel $k(x, x')$, which encodes smoothness properties of functions sampled from the GP.

Assuming a GP prior $\mathcal{GP}(0, k(x, x'))$ over f and given t noisy measurements $y_t = [y_1, \dots, y_t]^T$ for points in $A_t = \{x_1, \dots, x_t\}$, where $y_i = f(x_i) + n_i$ and $n_i \sim \mathcal{N}(0, \sigma^2)$ (Gaussian i.i.d. noise) for $i = 1, \dots, t$, the posterior over f is also a GP and its mean, covariance, and variance functions are given by the following analytic formulae:

$$\mu_t(x) = k_t(x)^T (K_t + \sigma^2 I)^{-1} y_t \quad (1)$$

$$\begin{aligned} k_t(x, x') &= k(x, x') - k_t(x)^T (K_t + \sigma^2 I)^{-1} k_t(x) \\ \sigma_t^2(x) &= k_t(x, x), \end{aligned} \quad (2)$$

where $k_t(x) = [k(x_1, x), \dots, k(x_t, x)]^T$ and K_t is the kernel matrix of already observed points, defined as $K_t = [k(x, x')]_{x, x' \in A_t}$.

3 The LSE Algorithm

We now present our proposed Level Set Estimation (LSE) algorithm for the strictly sequential setting with explicit thresholds. LSE is similar in spirit to the GP-UCB [Srinivas *et al.*, 2010] bandit optimization algorithm in that it uses a GP to model the underlying function and facilitates the inferred mean and variance of the GP to guide the selection of points to be evaluated.

More concretely, the inferred mean and variance of (1) and (2) can be used to construct a *confidence interval*

$$Q_t(x) = [\mu_{t-1}(x) \pm \beta_t^{1/2} \sigma_{t-1}(x)]$$

for any point $x \in D$, which captures our uncertainty about $f(x)$ after having already obtained noisy evaluations of f at points $\{x_1, \dots, x_t\}$. The parameter β_t acts as a scaling factor and its choice is discussed later. The above-defined confidence intervals serve two purposes in our algorithm: first, they allow us to judge whether a point can be classified into the super- or sublevel sets or whether the decision should be deferred until more information is available; second, they guide the sampling process towards points that are likely to be informative with respect to the desired level set.




Figure 1: (a) Example of the three possible configurations of confidence regions. (b) Ambiguities of two example points. (c) Chlorophyll concentration in relative fluorescence units (RFU) inferred from 2024 measurements on a vertical transect plane of Lake Zurich (the level set at $h = 1.3$ RFU is shown dashed). (d) LSE after $t = 50$ iterations on a grid of 100×100 points: regions of already classified points in orange (H_t) and blue (L_t), of yet unclassified points (U_t) in black, and observed points ($\{x_i\}_{1 \leq i \leq t}$) as white marks. Note how the sampling process focuses on the ambiguous regions around the desired level set.

The pseudocode of Algorithm 1 depicts in detail the operation of LSE. Our algorithm maintains a set of yet unclassified points U_t , as well as a superlevel set H_t and a sublevel set L_t , which are updated at each iteration. Furthermore, the algorithm maintains for each x a monotonically decreasing *confidence region* $C_t(x)$, which results from intersecting successive confidence intervals, i.e.

$$C_t(x) = \bigcap_{i=1}^t Q_i(x) = C_{t-1}(x) \cap Q_t(x).$$

Initially, all points $x \in D$ are unclassified and the confidence regions have infinite range (line 1). At each iteration, the confidence regions of all unclassified points are updated (line 6) and each of these points is either classified into one of H_t or L_t , or is left unclassified (lines 7–10). Then, the next point is selected and evaluated (lines 11–12) and the new GP posterior is computed (line 13). The algorithm terminates when all points in D have been classified, in which case the estimated super- and sublevel sets $\hat{H}$ and $\hat{L}$ are returned (line 16).

Classification. The classification of a point x into H_t or L_t depends on the position of its confidence region with respect to the threshold level h . Intuitively, if all of $C_t(x)$ lies above h , then with high probability $f(x) > h$ and x should

be moved into H_t . Similarly, if $C_t(x)$ lies below h , then x should be moved into L_t . Otherwise, we are still uncertain about the class of x , therefore it should, for the moment, remain unclassified. As can be seen in the classification rules of lines 7 and 9, we relax these conditions by introducing an accuracy parameter ϵ , which trades off classification accuracy for sampling cost. The resulting classification scheme is illustrated by the example of Figure 1a, in which point x would be classified into H_t and point x'' into L_t , while point x' would remain in U_t as unclassified. Note that LSE uses a monotonic classification scheme, meaning that once a point has been classified, it stays so until the algorithm terminates.

Sample selection. For selecting the next point to be evaluated at each iteration, we define the following quantity

$$a_t(x) = \min\{\max(C_t(x)) - h, h - \min(C_t(x))\},$$

which we call *ambiguity* and, as it names suggests, quantifies our uncertainty about whether x belongs to H_t or L_t (see Figure 1b). The intuition of sampling at areas of the sample space with large classification uncertainty, expecting to gain more information about the problem at hand when sampling at those areas, manifests itself in LSE by choosing to evaluate at each iteration the point with the largest ambiguity amongst the yet unclassified (see Figures 1c and 1d).

We can make an interesting observation at this point. If we use the confidence intervals $Q_t(x)$ instead of the confidence regions $C_t(x)$ in the definition of ambiguity, we get

$$\begin{aligned} a'_t(x) &= \min\{\max(Q_t(x)) - h, h - \min(Q_t(x))\} \\ &= \beta_t^{1/2} \sigma_{t-1}(x) - |\mu_{t-1}(x) - h|. \end{aligned}$$

For $\beta_t^{1/2} = 1.96$, this is identical to the *straddle* [Bryan *et al.*, 2005] selection rule, which can thus be intuitively explained in terms of classification ambiguity.

Theoretical analysis. The convergence analysis of LSE rests on quantifying the complexity of the GP prior for f in information-theoretic terms. The information gain [Cover and Thomas, 2006] about f from observing t noisy measurements $y_t = (y_i)_{1 \leq i \leq t}$ is

$$I(y_t; f) = H(y_t) - H(y_t | f).$$

Srinivas *et al.* [2010] used the *maximum information gain* over all possible sets of t observations

$$\gamma_t = \max_{y_t} I(y_t; f)$$

Algorithm 1 The LSE algorithm

Input: sample set D , GP prior ($\mu_0 = 0, k, \sigma_0$), threshold value h , accuracy parameter ϵ

Output: predicted sets $\hat{H}, \hat{L}$

```

1:  $H_0 \leftarrow \emptyset, L_0 \leftarrow \emptyset, U_0 \leftarrow D, C_0(x) \leftarrow \mathbb{R}$ , for all  $x \in D$ 
2:  $t \leftarrow 1$ 
3: while  $U_{t-1} \neq \emptyset$  do
4:    $H_t \leftarrow H_{t-1}, L_t \leftarrow L_{t-1}, U_t \leftarrow U_{t-1}$ 
5:   for all  $x \in U_{t-1}$  do
6:      $C_t(x) \leftarrow C_{t-1}(x) \cap Q_t(x)$ 
7:     if  $\min(C_t(x)) + \epsilon > h$  then
8:        $U_t \leftarrow U_t \setminus \{x\}, H_t \leftarrow H_t \cup \{x\}$ 
9:     else if  $\max(C_t(x)) - \epsilon \leq h$  then
10:       $U_t \leftarrow U_t \setminus \{x\}, L_t \leftarrow L_t \cup \{x\}$ 
11:    $x_t \leftarrow \operatorname{argmax}_{x \in U_t} (a_t(x))$ 
12:    $y_t \leftarrow f(x_t) + n_t$ 
13:   Compute  $\mu_t(x)$  and  $\sigma_t(x)$  for all  $x \in U_t$ 
14:    $t \leftarrow t + 1$ 
15:  $\hat{H} \leftarrow H_{t-1}, \hat{L} \leftarrow L_{t-1}$ 

```

for bounding the regret of the GP-UCB algorithm. Even though the problem we consider is different, we can use the same quantity to bound the number of LSE iterations required to achieve a certain classification quality.

To quantify the quality of a solution $(\hat{H}, \hat{L})$ with respect to a single point $\mathbf{x} \in D$ we use the misclassification loss

$$\ell_h(\mathbf{x}) = \begin{cases} \max\{0, f(\mathbf{x}) - h\}, & \text{if } \mathbf{x} \in \hat{L} \\ \max\{0, h - f(\mathbf{x})\}, & \text{if } \mathbf{x} \in \hat{H} \end{cases}.$$

The overall quality of a solution can then be judged by the largest misclassification loss among all points in the sample space, i.e. $\max_{\mathbf{x} \in D} \ell_h(\mathbf{x})$. Intuitively, having a solution with $\max_{\mathbf{x} \in D} \ell_h(\mathbf{x}) \leq \epsilon$ means that every point $\mathbf{x}$ is correctly classified with respect to a threshold level that deviates by at most ϵ from the true level h ; we call such a solution ϵ -accurate. The following theorem establishes a convergence bound for LSE in terms of γ_t for any given accuracy ϵ .

Theorem 1. *For any $h \in \mathbb{R}$, $\delta \in (0, 1)$, and $\epsilon > 0$, if $\beta_t = 2 \log(|D|\pi^2 t^2 / (6\delta))$, LSE terminates after at most T iterations, where T is the smallest positive integer satisfying*

$$\frac{T}{\beta_T \gamma_T} \geq \frac{C_1}{4\epsilon^2},$$

where $C_1 = 8 / \log(1 + \sigma^{-2})$.

Furthermore, with probability at least $1 - \delta$, the algorithm returns an ϵ -accurate solution, that is

$$\Pr \left\{ \max_{\mathbf{x} \in D} \ell_h(\mathbf{x}) \leq \epsilon \right\} \geq 1 - \delta.$$

The proof of Theorem 1 can be outlined as follows. The choice of β_t guarantees the inclusion of $f(\mathbf{x})$ in the confidence regions $C_t(\mathbf{x})$ for all $\mathbf{x} \in D$. From the monotonic classification scheme and the maximum ambiguity selection rule, it follows that the ambiguities of the selected points, $a_t(\mathbf{x}_t)$, are decreasing with t and, using γ_t , they can be shown to decrease as $\mathcal{O}((\frac{\beta_t \gamma_t}{t})^{\frac{1}{2}})$. Finally, the classification rules of LSE connect the number of iterations to the accuracy parameter ϵ and guarantee an ϵ -accurate solution with high probability.

Note that bounds on γ_T have been established for commonly used kernels [Srinivas *et al.*, 2010] and can be plugged into Theorem 1 to obtain concrete bounds on T . For example, for a d -dimensional sample space and a squared exponential GP kernel, $\gamma_T = \mathcal{O}((\log T)^{d+1})$, and the expression in the bound of Theorem 1 becomes $T / (\log T)^{d+2} \geq C / \epsilon^2$, where, for any given sample space and kernel hyperparameters, C depends only on the choice of δ .

4 Extensions

We now extend LSE to deal with the two problem variants introduced in Section 2. We highlight the key differences in the extended versions of the algorithm and the resulting implications about the convergence bound of Theorem 1.

4.1 Implicit Threshold Level

The substitution of the explicit threshold level by an implicit level $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ requires modifying the classification rules as well as the selection rule of LSE, which results in what we call the LSE_{imp} algorithm.

Algorithm 2 The LSE_{imp} extension

Input: sample set D , GP prior $(\mu_0 = 0, k, \sigma_0)$, threshold ratio ω , accuracy parameter ϵ

Output: predicted sets $\hat{H}, \hat{L}$

- 1: // initialization as in LSE; in addition: $Z_0 \leftarrow D$
- 2: **while** $U_{t-1} \neq \emptyset$ **do**
- 3: // new set definitions as in LSE
- 4: **for all** $\mathbf{x} \in U_{t-1}$ **do**
- 5: $C_t(\mathbf{x}) \leftarrow C_{t-1}(\mathbf{x}) \cap Q_t(\mathbf{x})$
- 6: $h_t^{opt} \leftarrow \omega \max_{\mathbf{x} \in Z_{t-1}} \max(C_t(\mathbf{x}))$
- 7: $f_t^{pes} \leftarrow \max_{\mathbf{x} \in Z_{t-1}} \min(C_t(\mathbf{x}))$, $h_t^{pes} \leftarrow \omega f_t^{pes}$
- 8: **if** $\min(C_t(\mathbf{x})) + \epsilon \geq h_t^{opt}$ **then**
- 9: $U_t \leftarrow U_t \setminus \{\mathbf{x}\}$
- 10: **if** $\max(C_t(\mathbf{x})) < f_t^{pes}$ **then** $H_t \leftarrow H_t \cup \{\mathbf{x}\}$
- 11: **else** $M_t^H \leftarrow M_t^H \cup \{\mathbf{x}\}$
- 12: **else if** $\max(C_t(\mathbf{x})) - \epsilon \leq h_t^{pes}$ **then**
- 13: $U_t \leftarrow U_t \setminus \{\mathbf{x}\}$
- 14: **if** $\max(C_t(\mathbf{x})) < f_t^{pes}$ **then** $L_t \leftarrow L_t \cup \{\mathbf{x}\}$
- 15: **else** $M_t^L \leftarrow M_t^L \cup \{\mathbf{x}\}$
- 16: $Z_t \leftarrow U_t \cup M_t^H \cup M_t^L$
- 17: $\mathbf{x}_t \leftarrow \operatorname{argmax}_{\mathbf{x} \in Z_t} (w_t(\mathbf{x}))$
- 18: // GP inference as in LSE
- 19: $\hat{H} \leftarrow H_{t-1} \cup M_{t-1}^H$, $\hat{L} \leftarrow L_{t-1} \cup M_{t-1}^L$

Since h is now an estimated quantity that depends on the function maximum, we have to take into account the uncertainty associated with it when making classification decisions. Concretely, we can obtain an *optimistic estimate* of the function maximum as $f_t^{opt} = \max_{\mathbf{x} \in U_t} \max(C_t(\mathbf{x}))$ and, analogously, a *pessimistic estimate* as $f_t^{pes} = \max_{\mathbf{x} \in U_t} \min(C_t(\mathbf{x}))$. The corresponding estimates of the implicit level are defined as $h_t^{opt} = \omega f_t^{opt}$ and $h_t^{pes} = \omega f_t^{pes}$, and can be used in a similar classification scheme to that of LSE. However, for the above estimates to be correct, we have to ensure that U_t always contains all points that could be maximizers of f , i.e. all points that satisfy $\max(C_t(\mathbf{x})) \geq f_t^{pes}$. For that purpose, points that should be classified, but are still possible function maximizers according to the above inequality, are kept in two sets M_t^H and M_t^L respectively, while a new set $Z_t = U_t \cup M_t^H \cup M_t^L$ is used in place of U_t to obtain the optimistic and pessimistic estimates h_t^{opt} and h_t^{pes} . The resulting classification rules are shown in Algorithm 2, where the conditions are again relaxed by an accuracy parameter ϵ .

In contrast to LSE, which solely focuses on sampling the most ambiguous points, in LSE_{imp} it is also of importance to have a more exploratory sampling policy in order to obtain more accurate estimates f_t^{opt} and f_t^{pes} . To this end, we select at each iteration the point with the largest confidence region *width*, defined as

$$w_t(\mathbf{x}) = \max(C_t(\mathbf{x})) - \min(C_t(\mathbf{x})).$$

If confidence intervals were not intersected, this would be equivalent to maximum variance sampling (within Z_t).

Theoretical analysis. The main challenge in extending the results of Theorem 1 to the implicit threshold level setting, stems from the fact that achieving a certain level of classification accuracy, now also depends on having accurate estimates

of h . This translates into appropriately bounding $h_t^{opt} - h_t^{pes}$, while, at the same time, guaranteeing that classification is still being performed correctly (by showing that $h_t^{opt} > h$ and $h_t^{pes} < h$). The following theorem expresses the resulting convergence bound for LSE_{imp} .

Theorem 2. *For any $\omega \in (0, 1)$, $\delta \in (0, 1)$, and $\epsilon > 0$, if $\beta_t = 2 \log(|D| \pi^2 t^2 / (6\delta))$, LSE_{imp} terminates after at most T iterations, where T is the smallest positive integer satisfying*

$$\frac{T}{\beta_T \gamma_T} \geq \frac{C_1(1 + \omega)^2}{4\epsilon^2},$$

where $C_1 = 8 / \log(1 + \sigma^{-2})$.

Furthermore, with probability at least $1 - \delta$, the algorithm returns an ϵ -accurate solution with respect to the implicit level $h = \omega \max_{x \in D} f(x)$, that is

$$\Pr \left\{ \max_{x \in D} \ell_h(x) \leq \epsilon \right\} \geq 1 - \delta.$$

Note that the sample complexity bound of Theorem 2 is a factor $(1 + \omega)^2 \leq 4$ larger than that of Theorem 1, and that $\omega = 0$ actually reduces to an explicit threshold of 0.

4.2 Batch Sample Selection

In the batch setting, the algorithms are only allowed to use the observed values of previous batches when selecting samples for the current batch. A naive way of extending LSE (resp. LSE_{imp}) to this setting would be to modify the selection rule so that, instead of picking the point with the largest ambiguity (resp. width), it chooses the B highest ranked points. However, this approach tends to select “clusters” of closely located samples with high ambiguities (resp. widths), ignoring the decrease in the estimated variance of a point resulting from sampling another point nearby.

Fortunately, we can handle the above issue by exploiting a key property of GPs, namely that the predictive variances (2) depend only on the selected points x_t and not on the observed values y_t at those points. Therefore, even if we do not have available feedback for each selected point up to iteration t , we can still obtain the following useful confidence intervals

$$Q_t^b(x) = \left[\mu_{fb[t]}(x) \pm \eta_t^{1/2} \sigma_{t-1}(x) \right],$$

which combine the most recent available mean estimate ($fb[t]$ being the index of the last available observation) with the always up-to-date variance estimate. Confidence regions $C_t^b(x)$ are defined as before by intersecting successive confidence intervals and are used without any further changes in the algorithms. However, to guarantee convergence we must compensate for using outdated mean estimates, by employing a more conservative (i.e., larger) scaling parameter η_t compared to β_t , in order to ensure that the resulting confidence regions $C_t^b(x)$ still contain $f(x)$ with high probability.

5 Experiments

In this section, we evaluate our proposed algorithms on two real-world datasets and compare them to the state-of-the-art. In more detail, the algorithms and their setup are as follows.

LSE/ LSE_{imp} : Since the bound of Theorem 1 is fairly conservative, in our experiments we used a constant value of $\beta_t^{1/2} = 3$, which is somewhat smaller than the values suggested by the theorem.

$LSE_{batch}/LSE_{imp-batch}$: We used $\eta_t^{1/2} = 4$ and $B = 30$.

STR: The state-of-the-art straddle heuristic, as proposed by Bryan *et al.* [2005], with the selection rule $x_t = \operatorname{argmax}_{x \in D} (1.96\sigma_{t-1}(x) - |\mu_{t-1}(x) - h|)$.

STR_{imp} : For the implicit threshold setting, we have defined a variant of the straddle heuristic that uses at each step the implicitly defined threshold level with respect to the maximum of the inferred mean, i.e. $h_t = \omega \max_{x \in D} \mu_{t-1}(x)$.

STR_{rank}/STR_{batch} : We have defined two batch versions of the straddle heuristic: STR_{rank} selects the $B = 30$ points with the largest straddle score, while STR_{batch} follows a similar approach to LSE_{batch} by using the selection rule $x_t = \operatorname{argmax}_{x \in D} (1.96\sigma_{t-1}(x) - |\mu_{fb[t]}(x) - h|)$.

VAR: The max. variance rule $x_t = \operatorname{argmax}_{x \in D} \sigma_{t-1}(x)$.

We assess the classification accuracy for all algorithms using the F_1 -score, i.e. the harmonic mean of precision and recall, by considering points in the super- and sublevel sets as positives and negatives respectively. Finally, in all evaluations of LSE and its extensions, the accuracy parameter ϵ is chosen to increase exponentially from 2% up to 20% of the maximum value of each dataset.

Dataset 1: Network latency. Our first dataset consists of round-trip time (RTT) measurements obtained by “pinging” 1768 servers spread around the world. The sample space consists of the longitude and latitude coordinates of each server, as determined by a commercial geolocation database¹. Example applications for geographic RTT level set estimation, include monitoring global networks and improving quality of service for applications such as internet telephony or online games. Furthermore, selecting samples in batches results in significant time savings, since sending out and receiving multiple ICMP packets can be virtually performed in parallel.

We used 200 randomly selected measurements to fit suitable hyperparameters for an anisotropic Matérn-5 [Rasmussen and Williams, 2006] kernel by maximum likelihood and the remaining 1568 for evaluation. The threshold level we chose for the experiments was $h = 200$ ms.

Dataset 2: Environmental monitoring. Our second dataset comes from the domain of environmental monitoring of inland waters and consists of 2024 *in situ* measurements of chlorophyll concentration within a vertical transect plane, collected by an autonomous surface vessel in Lake Zurich [Hitz *et al.*, 2012]. Since chlorophyll levels can vary throughout the year, in addition to having a fixed threshold concentration, it can also be useful to be able to detect relative “hotspots” of chlorophyll, i.e. regions of high concentration with respect to the current maximum. Furthermore, selecting batches of points can be used to plan sampling paths and reduce the required traveling distances.

In our evaluation, we used 10,000 points sampled in a 100×100 grid from the GP posterior that was derived using the 2024 original measurements (see Figure 1c). Again, an

¹<http://www.maxmind.com>




Figure 2: Results for the network latency [N] and environmental monitoring [E] datasets. (a), (b) LSE is competitive with STR, while both clearly outperform VAR. (c), (d) LSE_{batch} and STR_{batch} are only slightly worse than STR, while the naive STR_{rank} is far inferior. (e) LSE_{imp} and $LSE_{imp-batch}$ achieve high accuracy at a slower rate than the explicit threshold algorithms, while STR_{imp} fails to do so. (f) Using larger batch sizes for planning dramatically reduces the traveled path lengths.

anisotropic Matérn-5 kernel was used and suitable hyperparameters were fitted by maximizing the likelihood of a second available chlorophyll dataset from Lake Zurich. We used an explicit threshold level of $h = 1.3$ RFU and chose $\omega = 0.74$ for the implicit threshold case, so that the resulting implicit level is identical to the explicit one, which enables us to compare the two settings on equal ground. We evaluated the effect of batch sampling on the required traveling distance using an approximate Euclidean TSP solver to create paths connecting each batch of samples selected by LSE_{batch} .

Results and discussion. Figures 2a and 2b compare the performance of the strictly sequential algorithms on the two datasets. In both cases, LSE and STR are comparable in performance, which is expected given the similarity of their selection rules (see Section 3). Although VAR is commonly used for estimating functions over their entire domain, it is clearly outperformed by both algorithms and, thus, deemed unsuitable for the task of level set estimation.

In Figures 2c and 2d we show the performance of the batch algorithms on the two datasets. The LSE_{batch} and STR_{batch} algorithms, which use the always up-to-date variance estimates for selecting batches, achieve similar performance. Furthermore, there is only a slight performance penalty when compared to the strictly sequential STR, which can easily be outweighed by the benefits of batch point selection (e.g. in the network latency example, the batch algorithms would have about $B = 30$ times higher throughput). As expected, the STR_{rank} algorithm, performs significantly worse than the other two batch algorithms, since it selects a lot of redundant samples in areas of high straddle score (cf. Section 4.2).

Figure 2e shows the results of the implicit threshold experiments on the environmental monitoring dataset. The difficulty of estimating the function maximum at the same time as performing classification with respect to the implicit threshold level is manifested in the notably larger sampling cost of LSE_{imp} required to achieve high accuracy compared to the

explicit threshold experiments. As before, the batch version of LSE_{imp} is only slightly worse than its sequential counterpart. More importantly, the naive STR_{imp} algorithm completely fails to achieve high accuracy, as it gets stuck with a wrong estimate of the maximum and never recovers, since the straddle rule is not sufficiently exploratory.

Finally, Figure 2f displays the dramatically lower required travel length by using batches of samples for path planning: to achieve an F_1 -score of 0.95 using sequential sampling requires more than six times larger traveling distance than planning ahead with $B = 30$ samples per batch.

6 Conclusion

We presented LSE, an algorithm for estimating level sets, which operates based on confidence bounds derived by modeling the target function as a GP. We considered for the first time the challenge of implicitly defined threshold levels and extended LSE to this more complex setting. We also showed how both algorithms can be extended to select samples in batches. In addition, we provided theoretical bounds on the number of iterations required to obtain an ϵ -accurate solution when the target function is sampled from a known GP. The experiments on two real-world applications showed that LSE is competitive with the state-of-the-art, while its extensions are successful in handling the corresponding problem variants and perform significantly better than naive baselines. We believe our results provide an important step towards addressing complex real-world information gathering problems.

Acknowledgments. We would like to thank the Limnological Station of University of Zurich and the Autonomous Systems Lab of ETH Zurich for providing the environmental monitoring datasets. This research was supported in part by SNSF grants 200021 137971 and CR2212-130023, ERC StG 307036, and a Microsoft Research Faculty Fellowship.

References

- [Azimi *et al.*, 2010] Javad Azimi, Alan Fern, and Xiaoli Fern. Batch bayesian optimization via simulation matching. In *Neural Information Processing Systems (NIPS)*, 2010.
- [Brochu *et al.*, 2010] Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv:1012.2599*, 2010.
- [Bryan and Schneider, 2008] Brent Bryan and Jeff Schneider. Actively learning level-sets of composite functions. In *International Conference on Machine Learning (ICML)*, 2008.
- [Bryan *et al.*, 2005] Brent Bryan, Jeff Schneider, Robert Nichol, Christopher Miller, Christopher Genovese, and Larry Wasserman. Active learning for identifying function threshold boundaries. In *Neural Information Processing Systems (NIPS)*, 2005.
- [Cover and Thomas, 2006] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2006.
- [Dantu and Sukhatme, 2007] Karthik Dantu and Gaurav S. Sukhatme. Detecting and tracking level sets of scalar fields using a robotic sensor network. In *International Conference on Robotics and Automation (ICRA)*, 2007.
- [Desautels *et al.*, 2012] Thomas Desautels, Andreas Krause, and Joel Burdick. Parallelizing exploration-exploitation tradeoffs with gaussian process bandit optimization. In *International Conference on Machine Learning (ICML)*, 2012.
- [Galland *et al.*, 2004] Frédéric Galland, Philippe Rófrógier, and Olivier Germain. Synthetic aperture radar oil spill segmentation by stochastic complexity minimization. *IEEE Geoscience and Remote Sensing Letters (GRSL)*, 2004.
- [Garnett *et al.*, 2012] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying. In *International Conference on Machine Learning (ICML)*, 2012.
- [Hitz *et al.*, 2012] Gregory Hitz, François Pomerleau, Marie-Eve Garneau, Cédric Pradalier, Thomas Posch, Jakob Pernthaler, and Roland Y. Siegwart. Autonomous inland water monitoring: Design and application of a surface vessel. *Robotics & Automation Magazine (RAM)*, 2012.
- [Rahimi *et al.*, 2004] Mohammad Rahimi, Richard Pon, William J. Kaiser, Gaurav S. Sukhatme, Deborah Estrin, and Mani Srivastava. Adaptive sampling for environmental robotics. In *International Conference on Robotics and Automation (ICRA)*, 2004.
- [Ramakrishnan *et al.*, 2005] Naren Ramakrishnan, Chris Bailey-Kellogg, Satish Tadepalli, and Varun N. Pandey. Gaussian processes for active data mining of spatial aggregates. In *SIAM International Conference on Data Mining (SDM)*, 2005.
- [Rasmussen and Williams, 2006] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [Settles, 2009] Burr Settles. Active Learning Literature Survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.
- [Singh *et al.*, 2006] Aarti Singh, Robert Nowak, and Parmesh Ramanathan. Active learning for adaptive mobile sensing networks. In *International Conference on Information Processing in Sensor Networks (IPSN)*, 2006.
- [Srinivas *et al.*, 2010] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning (ICML)*, 2010.
- [Srinivasan *et al.*, 2008] Sumana Srinivasan, Krithi Ramamritham, and Purushottam Kulkarni. Ace in the hole: Adaptive contour estimation using collaborating mobile sensors. In *International Conference on Information Processing in Sensor Networks (IPSN)*, 2008.
- [Zuluaga *et al.*, 2013] Marcela Zuluaga, Andreas Krause, Guillaume Sergent, and Markus Püschel. Active learning for multi-criterion optimization. In *International Conference on Machine Learning (ICML)*, 2013.