

# Active Learning for Level Set Estimation

Alkis Gotovos<sup>1</sup>   Nathalie Casati<sup>1,2</sup>   Gregory Hitz<sup>1</sup>   Andreas Krause<sup>1</sup>

<sup>1</sup>Department of Computer Science  
ETH Zurich

<sup>2</sup>IBM Research – Zurich

IJCAI '13

## Swimmers of Lake Zurich, beware!



Steffen Schmidt / EPA

## Swimmers of Lake Zurich, beware!

“[...]Switzerland's Lake Zurich [...] an ideal environment for a population explosion of algae including *Planktothrix rubescens* [...]”

— *Scientific American*


## Swimmers of Lake Zurich, beware!



[www.limnobotics.ch](http://www.limnobotics.ch)

“[...]Switzerland's Lake Zurich [...] an ideal environment for a population explosion of algae including *Planktothrix rubescens* [...]”

— *Scientific American*




Flickr/Dr. Jennifer L. Graham/U.S. Geological Survey

## Swimmers of Lake Zurich, beware!



[www.limnobotics.ch](http://www.limnobotics.ch)



Flickr/Dr. Jennifer L. Graham/U.S. Geological Survey

“[...]Switzerland's Lake Zurich [...] an ideal environment for a population explosion of algae including *Planktothrix rubescens* [...]”

— *Scientific American*


“*Planktothrix rubescens* are among the most important producers of hepatotoxic microcystins in freshwaters [...]”

— *Silke Van den Wyngaert et al., ASLO, 2011*

## Swimmers of Lake Zurich, beware!



[www.limnobotics.ch](http://www.limnobotics.ch)



Flickr/Dr. Jennifer L. Graham/U.S. Geological Survey

“[...]Switzerland's Lake Zurich [...] an ideal environment for a population explosion of algae including *Planktothrix rubescens* [...]”

— *Scientific American*

“*Planktothrix rubescens* are among the most important producers of hepatotoxic microcystins in freshwaters [...]”

— *Silke Van den Wyngaert et al., ASLO, 2011*

“Microcystins [...] are cyanotoxins and can be very toxic for plants and animals including humans. Their hepatotoxicity may cause serious damage to the liver.”


— *Wikipedia*

## Autonomous surface vehicle developed by the Autonomous Systems Lab of ETH



[www.limnobotics.ch](http://www.limnobotics.ch)


Take measurements on a vertical transect of the lake




www.limrobotics.ch




Original algae concentration measurements ( $\sim 2000$ )




## Interpolated algae concentration field



Focus on accurately estimating regions of “high” concentration (e.g.  $\geq 7$ )



Classify transect into a **super-** and a **sublevel** set



Pose as a sequential decision making problem (*pool-based active learning*):

Pose as a sequential decision making problem (*pool-based active learning*):

- ▶ No measurements available in advance, just a set (pool) of possible sampling locations ( $D$ )

Pose as a sequential decision making problem (*pool-based active learning*):

- ▶ No measurements available in advance, just a set (pool) of possible sampling locations ( $D$ )

At each iteration  $t \geq 1$ :

- ▶ Decide where to measure next ( $\mathbf{x}_t \in D$ )

Pose as a sequential decision making problem (*pool-based active learning*):

- ▶ No measurements available in advance, just a set (pool) of possible sampling locations ( $D$ )

At each iteration  $t \geq 1$ :

- ▶ Decide where to measure next ( $\mathbf{x}_t \in D$ )
- ▶ Obtain noisy observation ( $y_t = f(\mathbf{x}_t) + n_t$ )



Pose as a sequential decision making problem (*pool-based active learning*):

- ▶ No measurements available in advance, just a set (pool) of possible sampling locations ( $D$ )

At each iteration  $t \geq 1$ :

- ▶ Decide where to measure next ( $\mathbf{x}_t \in D$ )
- ▶ Obtain noisy observation ( $y_t = f(\mathbf{x}_t) + n_t$ )
- ▶ Update our classification estimate

1. How do we **estimate** the underlying function from measurements?

1. How do we **estimate** the underlying function from measurements?
2. How do we **classify**?


1. How do we **estimate** the underlying function from measurements?
2. How do we **classify**?
3. Each measurement is expensive (time, battery power).  
How do we **select** “informative” measurements?

1. How do we **estimate** the underlying function from measurements?
2. How do we **classify**?
3. Each measurement is expensive (time, battery power).  
How do we **select** “informative” measurements?

Gaussian processes to the rescue!


# Gaussian processes

- Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$




# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates



# Gaussian processes


- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$





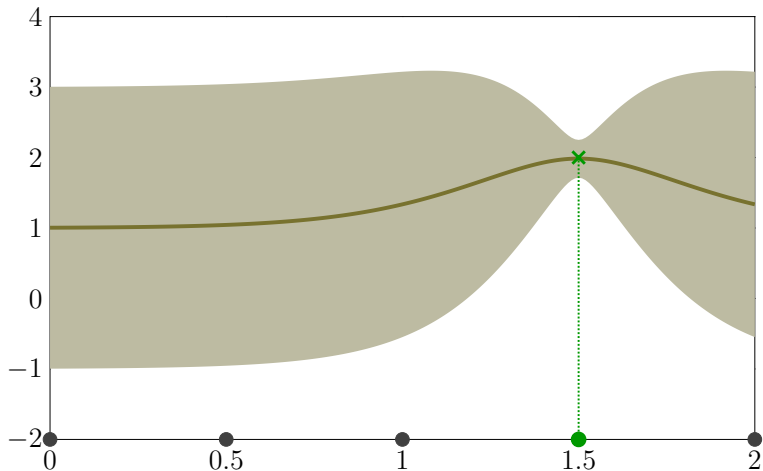
# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$



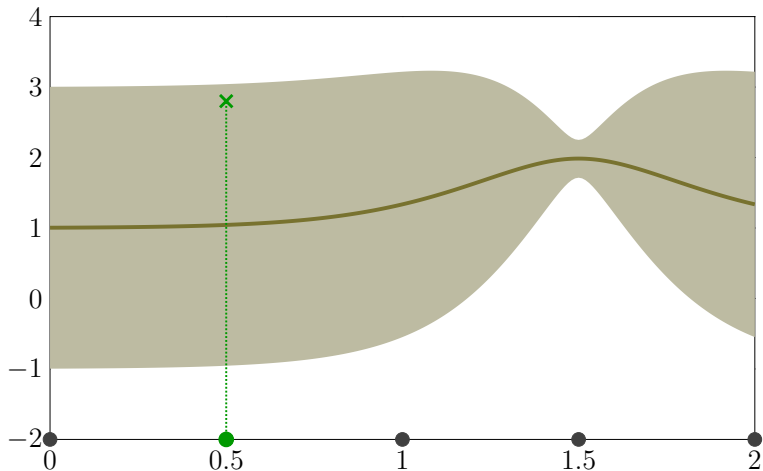
# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$




# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$




# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$




# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$




# Gaussian processes

- ▶ Mean **and variance** estimates: construct confidence intervals  $C(\mathbf{x})$
- ▶ **Bayesian**, yet **efficient**: suitable for step-by-step updates
- ▶ Impose prior “smoothness” assumptions via kernel  $k(\mathbf{x}, \mathbf{x}')$




# 1. How do we **estimate** the function?

# 1. How do we **estimate** the function?







1. How do we **estimate** the function? ✓




1. How do we **estimate** the function? ✓
2. How do we **classify**?



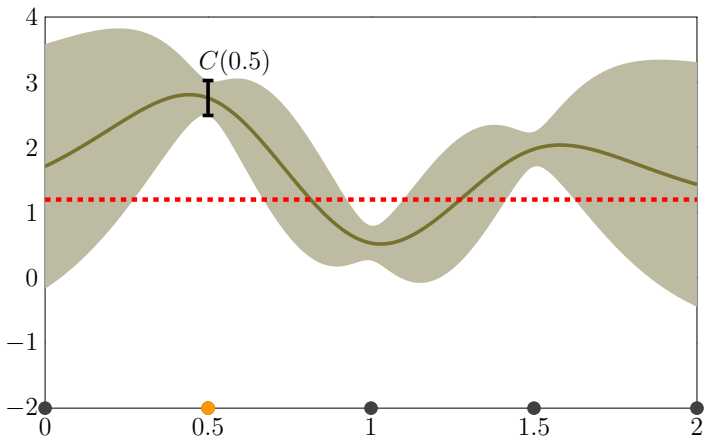
1. How do we **estimate** the function? ✓
2. How do we **classify**?




1. How do we **estimate** the function? ✓
2. How do we **classify**?



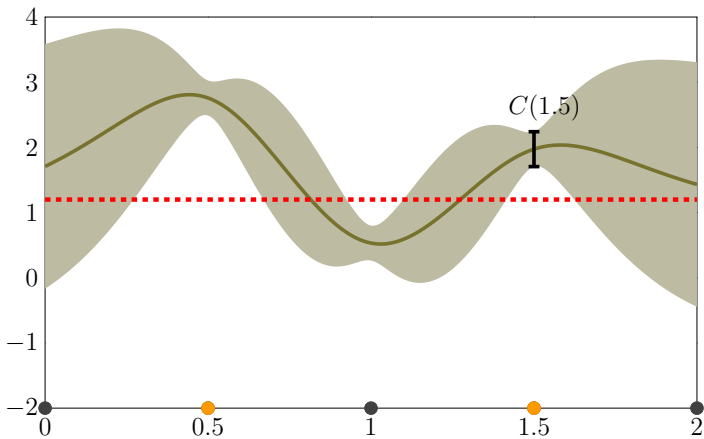
1. How do we **estimate** the function? ✓
2. How do we **classify**?




1. How do we **estimate** the function? ✓
2. How do we **classify**?



1. How do we **estimate** the function? ✓
2. How do we **classify**?




1. How do we **estimate** the function? ✓
2. How do we **classify**?







1. How do we **estimate** the function? ✓
2. How do we **classify**?




1. How do we **estimate** the function? ✓
2. How do we **classify**?




1. How do we **estimate** the function? ✓
2. How do we **classify**?




1. How do we **estimate** the function? ✓
2. How do we **classify**? ✓




1. How do we **estimate** the function? ✓
2. How do we **classify**? ✓
3. How do we **select** “informative” measurements?




1. How do we **estimate** the function? ✓
2. How do we **classify**? ✓
3. How do we **select** “informative” measurements?
  - ▶ Pick among the yet unclassified...



1. How do we **estimate** the function? ✓
2. How do we **classify**? ✓
3. How do we **select** “informative” measurements?
  - ▶ Pick among the yet unclassified...
  - ▶ ...the most “ambiguous” point




1. How do we **estimate** the function? ✓
2. How do we **classify**? ✓
3. How do we **select** “informative” measurements?
  - ▶ Pick among the yet unclassified...
  - ▶ ...the most “ambiguous” point





1. How do we **estimate** the function? ✓
2. How do we **classify**? ✓
3. How do we **select** “informative” measurements? ✓
  - ▶ Pick among the yet unclassified...
  - ▶ ...the most “ambiguous” point



## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

**while**  $\exists$  unclassified points in  $D$  **do**

**end while**

## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

**while**  $\exists$  unclassified points in  $D$  **do**

**for all** unclassified points  $x \in D$  **do**

**if**  $C_t(x)$  lies above  $h$  **then**

      classify  $x$  into superlevel set

**else if**  $C_t(x)$  lies below  $h$  **then**

      classify  $x$  into sublevel set

**else**

      leave  $x$  unclassified

**end if**

**end for**

← classify

**end while**

## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

**while**  $\exists$  unclassified points in  $D$  **do**

**for all** unclassified points  $x \in D$  **do**

**if**  $C_t(x)$  lies above  $h$  **then**

      classify  $x$  into superlevel set

**else if**  $C_t(x)$  lies below  $h$  **then**

      classify  $x$  into sublevel set

**else**

      leave  $x$  unclassified

**end if**

← classify

**end for**

$x_t \leftarrow \operatorname{argmax}\{a_t(x) \mid x \in U_t\}$

← sample

$y_t \leftarrow f(x_t) + n_t$

**end while**

## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

**while**  $\exists$  unclassified points in  $D$  **do**

**for all** unclassified points  $x \in D$  **do**

**if**  $C_t(x)$  lies above  $h$  **then**

      classify  $x$  into superlevel set

**else if**  $C_t(x)$  lies below  $h$  **then**

      classify  $x$  into sublevel set

**else**

      leave  $x$  unclassified

**end if**

**end for**

$\mathbf{x}_t \leftarrow \operatorname{argmax}\{a_t(\mathbf{x}) \mid \mathbf{x} \in U_t\}$

$y_t \leftarrow f(\mathbf{x}_t) + n_t$

  perform GP inference

**end while**

← classify

← sample

## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

```
while  $\exists$  unclassified points in  $D$  do  
  for all unclassified points  $x \in D$  do  
    if  $C_t(x)$  lies above  $h$  then  
      classify  $x$  into superlevel set  
    else if  $C_t(x)$  lies below  $h$  then  
      classify  $x$  into sublevel set  
    else  
      leave  $x$  unclassified  
    end if  
  end for  
   $\mathbf{x}_t \leftarrow \operatorname{argmax}\{a_t(\mathbf{x}) \mid \mathbf{x} \in U_t\}$   
   $y_t \leftarrow f(\mathbf{x}_t) + n_t$   
  perform GP inference  
end while
```

### ► Monotonicity of

1. confidence intervals
2. classification

## The Level Set Estimation (LSE) algorithm

**Input:** sample space  $D$ , threshold level  $h$

**Output:** predicted super- and sublevel sets

```
while  $\exists$  unclassified points in  $D$  do
  for all unclassified points  $x \in D$  do
    if  $C_t(x)$  lies above  $h$  then
      classify  $x$  into superlevel set
    else if  $C_t(x)$  lies below  $h$  then
      classify  $x$  into sublevel set
    else
      leave  $x$  unclassified
    end if
```

```
end for
```

```
 $x_t \leftarrow \operatorname{argmax}\{a_t(x) \mid x \in U_t\}$ 
```


```
 $y_t \leftarrow f(x_t) + n_t$ 
```

```
perform GP inference
```


```
end while
```


- ▶ Monotonicity of
  1. confidence intervals
  2. classification

- ▶ Relax classification rules by an accuracy parameter  $\epsilon$













$t = 0$







$t = 20$







$t = 40$







$t = 60$







$t = 80$






$t = 100$







$t = 120$










$t = 140$







$t = 160$







$t = 180$







$t = 200$






$t = 260$





$t = 354$



## Theorem (Convergence of LSE)

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  $\epsilon$ -accurate solution, that is

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

## Theorem (Simplified)

*If we choose  $\beta$  appropriately (large enough), then:*



## Theorem (Simplified)

*If we choose  $\beta$  appropriately (large enough), then:*

- ▶ *LSE terminates after a number of iterations  $T$*

## Theorem (Simplified)

If we choose  $\beta$  appropriately (large enough), then:

- ▶ LSE terminates after a number of iterations  $T$ 
  1. smoother kernel  $\Rightarrow T \downarrow$

## Theorem (Simplified)

If we choose  $\beta$  appropriately (large enough), then:

- ▶ LSE terminates after a number of iterations  $T$ 
  1. smoother kernel  $\Rightarrow T \downarrow$
  2.  $\sigma \uparrow \Rightarrow T \uparrow$

## Theorem (Simplified)

If we choose  $\beta$  appropriately (large enough), then:

- ▶ LSE terminates after a number of iterations  $T$ 
  1. smoother kernel  $\Rightarrow T \downarrow$
  2.  $\sigma \uparrow \Rightarrow T \uparrow$
  3.  $\epsilon \uparrow \Rightarrow T \downarrow$

## Theorem (Simplified)

If we choose  $\beta$  appropriately (large enough), then:

- ▶ LSE terminates after a number of iterations  $T$ 
  1. smoother kernel  $\Rightarrow T \downarrow$
  2.  $\sigma \uparrow \Rightarrow T \uparrow$
  3.  $\epsilon \uparrow \Rightarrow T \downarrow$
- ▶ The solution returned is  $\epsilon$ -accurate with high probability

## Theorem (Simplified)

If we choose  $\beta$  appropriately (large enough), then:

- ▶ LSE terminates after a number of iterations  $T$ 
  1. smoother kernel  $\Rightarrow T \downarrow$
  2.  $\sigma \uparrow \Rightarrow T \uparrow$
  3.  $\epsilon \uparrow \Rightarrow T \downarrow$
- ▶ The solution returned is  $\epsilon$ -accurate with high probability

# Experiments

## 1. LSE

# Experiments

1. LSE
2. Maximum variance sampling:

$$\mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in D} \sigma_{t-1}(\mathbf{x})$$




## Experiments



1. LSE
2. Maximum variance sampling:




$$\mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in D} \sigma_{t-1}(\mathbf{x})$$


3. State of the art “straddle” heuristic (Bryan *et al.*, 2005):



$$\mathbf{x}_t \approx \operatorname{argmax}_{\mathbf{x} \in D} a_{t-1}(\mathbf{x}) \quad (\text{for } \beta_t^{1/2} = 1.96)$$















## Implicit threshold level

- ▶ What if we don't know the threshold level  $h$ ?



## Implicit threshold level

- ▶ What if we don't know the threshold level  $h$ ?
- ▶ Implicitly define  $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ ,  $0 < \omega < 1$

## Implicit threshold level

- ▶ What if we don't know the threshold level  $h$ ?
- ▶ Implicitly define  $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ ,  $0 < \omega < 1$
- ▶ No existing algorithms for this problem (to our knowledge)

## Implicit threshold level

- ▶ What if we don't know the threshold level  $h$ ?
- ▶ Implicitly define  $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ ,  $0 < \omega < 1$
- ▶ No existing algorithms for this problem (to our knowledge)
- ▶ We extend LSE (and its theory!) to this setting

## Implicit threshold level


- ▶ What if we don't know the threshold level  $h$ ?
- ▶ Implicitly define  $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ ,  $0 < \omega < 1$
- ▶ No existing algorithms for this problem (to our knowledge)
- ▶ We extend LSE (and its theory!) to this setting
- ▶ **LSE<sub>imp</sub>** algorithm:


## Implicit threshold level

- ▶ What if we don't know the threshold level  $h$ ?
- ▶ Implicitly define  $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ ,  $0 < \omega < 1$
- ▶ No existing algorithms for this problem (to our knowledge)
- ▶ We extend LSE (and its theory!) to this setting
- ▶ **LSE<sub>imp</sub>** algorithm:
  - ▶  $h$  is now an estimated quantity  $\rightarrow$  modified classification rules


## Implicit threshold level

- ▶ What if we don't know the threshold level  $h$ ?
- ▶ Implicitly define  $h = \omega \max_{\mathbf{x} \in D} f(\mathbf{x})$ ,  $0 < \omega < 1$
- ▶ No existing algorithms for this problem (to our knowledge)
- ▶ We extend LSE (and its theory!) to this setting
- ▶ **LSE<sub>imp</sub>** algorithm:
  - ▶  $h$  is now an estimated quantity  $\rightarrow$  modified classification rules
  - ▶ Need to accurately estimate the maximum  $\rightarrow$  modified selection rule







$t = 0$










$t = 20$







$t = 40$







$t = 60$







$t = 80$







$t = 100$







$t = 120$






$t = 140$







$t = 160$










$t = 180$







$t = 200$








$t = 340$





$t = 486$





## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample.  
What about the traveling distance between measurements?

## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample.  
What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step


## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements




## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements




## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements



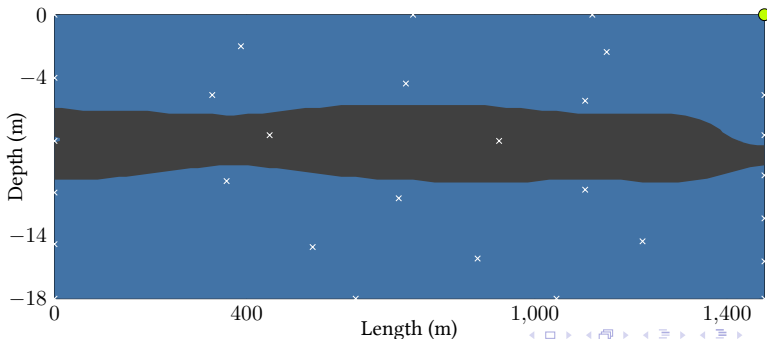
# Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements




## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements




# Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements




# Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements




## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\text{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements



## Batch sampling


- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\text{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements






## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements



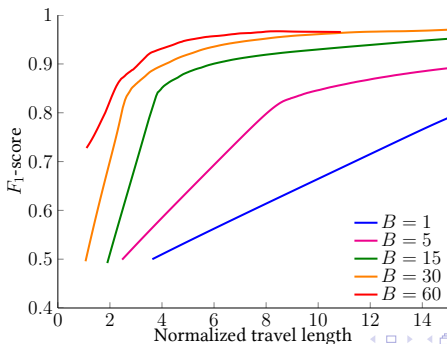
## Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\text{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements



# Batch sampling

- ▶ Up to this point we have assumed a fixed cost per sample. What about the traveling distance between measurements?
- ▶ We extend LSE to select a *batch* of sampling locations at each step
- ▶ Plan ahead:
  - ▶ Use  $\mathbf{LSE}_{\text{batch}}$  to select a batch of sampling locations
  - ▶ Connect them using a Euclidean TSP path
  - ▶ Traverse path and collect measurements



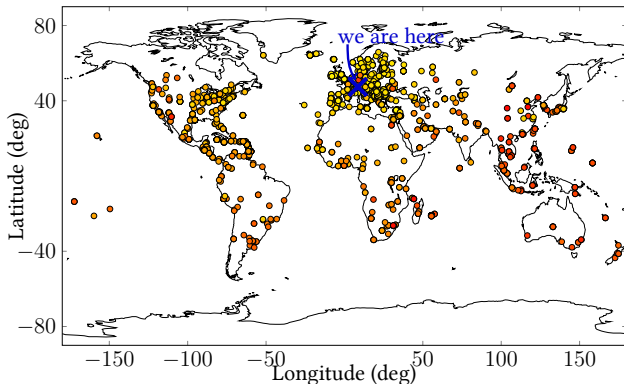
In our paper, more...

## In our paper, more...

- ▶ ...theory: sample complexity bounds and their proofs in more detail

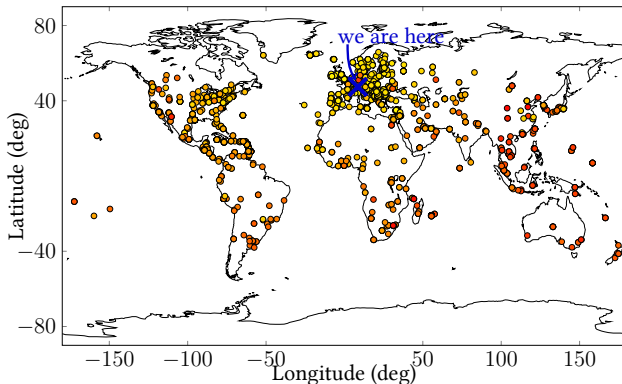
## In our paper, more...

- ▶ ...theory: sample complexity bounds and their proofs in more detail
- ▶ ...applications: estimate world regions of low internet latency



## In our paper, more...

- ▶ ...theory: sample complexity bounds and their proofs in more detail
- ▶ ...applications: estimate world regions of low internet latency



- ▶ ...experimental results

# Summary



# Summary

## ► LSE algorithm:

### Theoretical guarantees

#### Theorem (Convergence of LSE)

For any  $h \in \mathbb{R}$ ,  $\delta \in (0, 1)$ , and  $\epsilon > 0$ , if  $\beta_1 = 2 \log(|D| \pi^2 \ell^2 / (6\delta))$ , LSE terminates after at most  $T$  iterations, where  $T$  is the smallest positive integer satisfying


$$\frac{T}{\beta_1 \gamma_1} \geq \frac{C_1}{4\epsilon^2},$$

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

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

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

### Competitive with the state of the art



# Summary

## ► LSE algorithm:

### Theoretical guarantees

#### Theorem (Convergence of LSE)

For any  $h \in \mathbb{R}$ ,  $\delta \in (0, 1)$ , and  $\epsilon > 0$ , if  $\beta_1 = 2 \log(|D| \pi^2 \ell^2 / (6\delta))$ , LSE terminates after at most  $T$  iterations, where  $T$  is the smallest positive integer satisfying


$$\frac{T}{\beta_1 \gamma_1} \geq \frac{C_1}{4\epsilon^2},$$

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

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


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

### Competitive with the state of the art




## ► Two useful extensions:

### Implicit threshold level ( $LSE_{imp}$ )



### Batch sampling ( $LSE_{batch}$ )



# Summary

## ► LSE algorithm:

### Theoretical guarantees

#### Theorem (Convergence of LSE)

For any  $h \in \mathbb{R}$ ,  $\delta \in (0, 1)$ , and  $\epsilon > 0$ , if  $\beta_1 = 2 \log(|D| \pi^2 \ell^2 / (6\delta))$ , LSE terminates after at most  $T$  iterations, where  $T$  is the smallest positive integer satisfying


$$\frac{T}{\beta_1 \gamma_1} \geq \frac{C_1}{4\epsilon^2},$$

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

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


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

### Competitive with the state of the art




## ► Two useful extensions:

### Implicit threshold level ( $LSE_{imp}$ )



### Batch sampling ( $LSE_{batch}$ )



## ► Look out for algae when swimming in Lake Zurich! 😊