Machine Learning 2

More Sampling/Search

Ken Cameron

More Sampling/Search

- A/B Testing
- Multi-Arm Bandit Problem
- Thompson Sampling
- Particle Swarms
- Bayesian Quadrature

A/B Testing

- Popular approach in user interaction testing.

- Applicable to wider range of search problems.

- A special case of Multi-Arm Bandit.

A/B Testing Example

- Which version of an advert gets the most clicks?

- Solution: Randomly serve both, measure effectiveness.

Multi-Arm Bandit Problem

- Generalise this to multiple adverts...

- We’d like to maximise our overall return, but each advert has a different individual return we don’t yet know.

Multi-Arm Bandit

Trade-off

- This presents us with a dilemma.
- We need to evaluate the effectiveness of each. (Exploration.)
- But we’d like to serve the most effective immediately. (Exploitation.)

A semi-uniform strategy

- For some proportion of the time show the highest performing advert. ($1 - \varepsilon$)
- The remainder of the time, select with a uniform distribution from all the adverts. (ε)

Changing the proportion

- Epsilon-first
  - Begin with exploration only.
  - Switch to exploitation only.

- Epsilon-greedy
  - A constant proportion e.g. $\varepsilon = 0.1$

- Epsilon-decreasing
  - Start with exploration and change ε over time.

Thompson Sampling

- We can use a probability matching strategy.
- Known a Thompson sampling or Bayesian Bandits.
- Sample from the posterior of the mean value of each alternative.

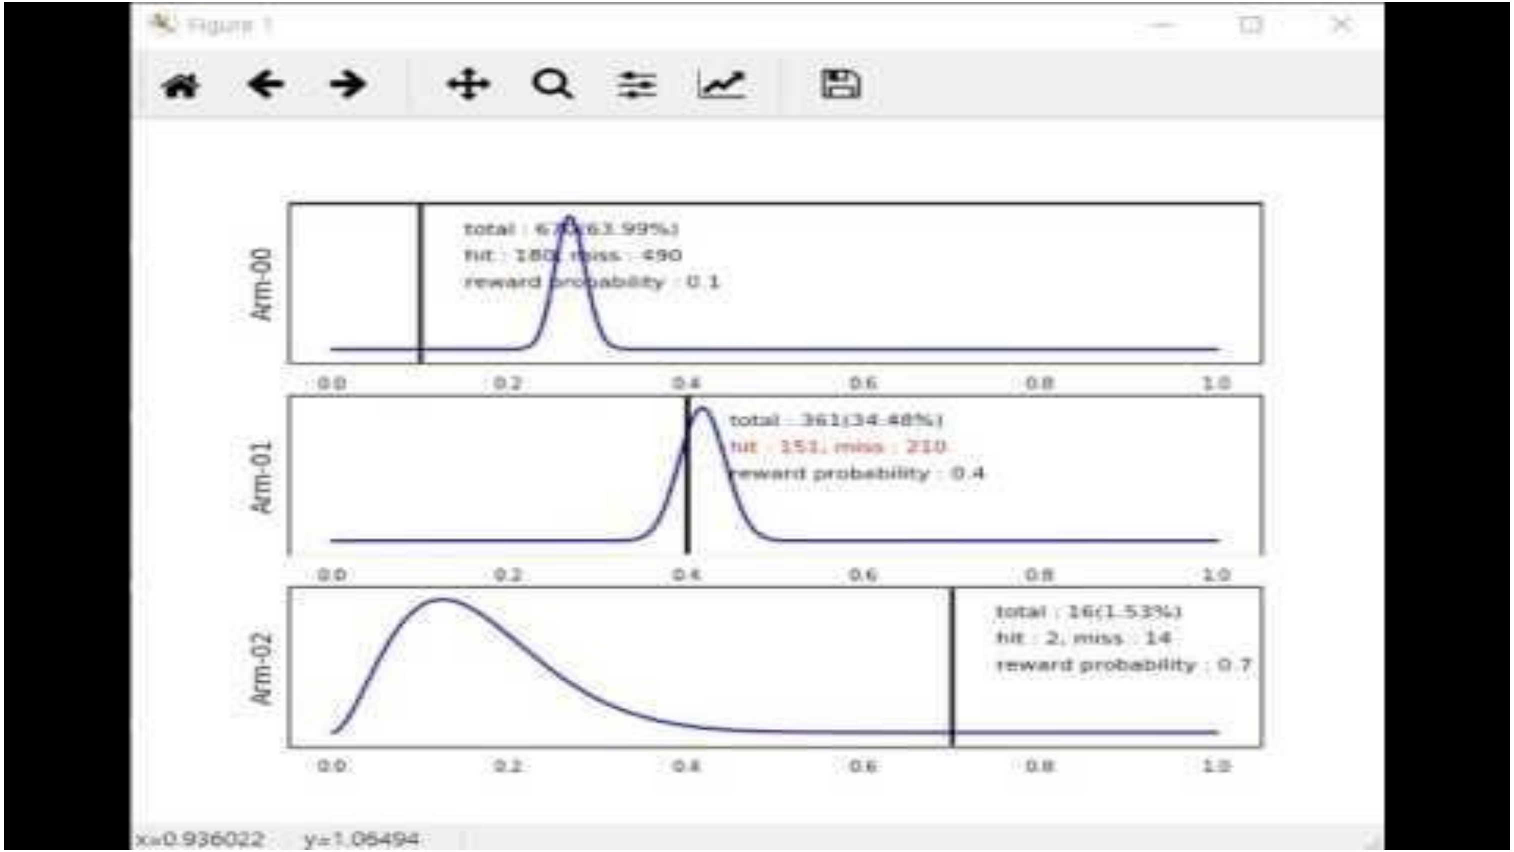
Figure 1
Arm-00
reward probability : 0.1
0.0 0.2 0.4 0.6 0.8 1.0
Arm-01
total : 361(34.48%)
hit : 151, miss : 210
reward probability : 0.4
0.0 0.2 0.4 0.6 0.8 1.0
Arm-02
total : 16(1.53%)
hit : 2, miss : 14
reward probability : 0.7
0.0 0.2 0.4 0.6 0.8 1.0
x=0.936022 y=1.06494

Particle Swarms

- Originated in simplified social system models.
- Intended to mimic flocks of birds/schools of fish.
- But it generalises to an optimiser.

Flocks of birds

- Each bird is searching for food.
- They don’t know where the food is only how far way.
- Best strategy:
  - Follow the bird who is closest.

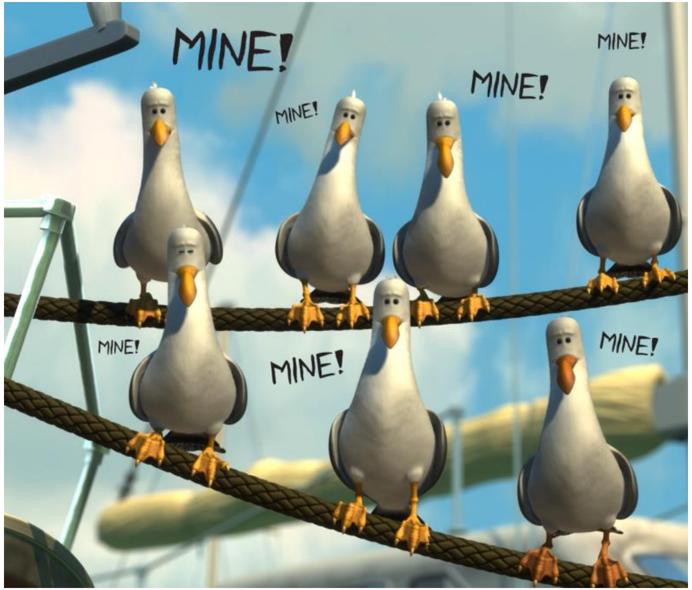

Generalise

- Birds become particles.

- On each iteration two ‘best’ values are used to update each particle.

  - Firstly, the best (optimal fitness) each particle has seen so far.
    - pbest.

  - Secondly, the best any particle has seen so far.
    - gbest.

- velocity[] = velocity[]
  + c1 * (pbest[] – present[])
  + c2 * rand() * (gbest – present[])

- present[] = present[] + v[]

- Typically c1= c2 = 2

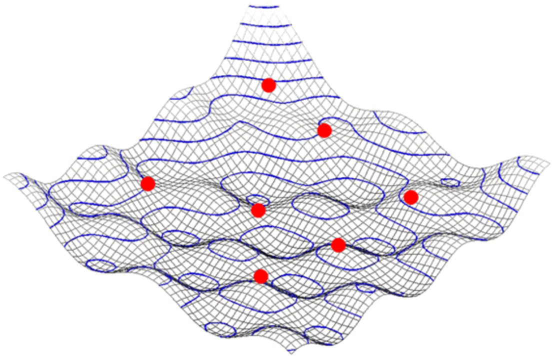

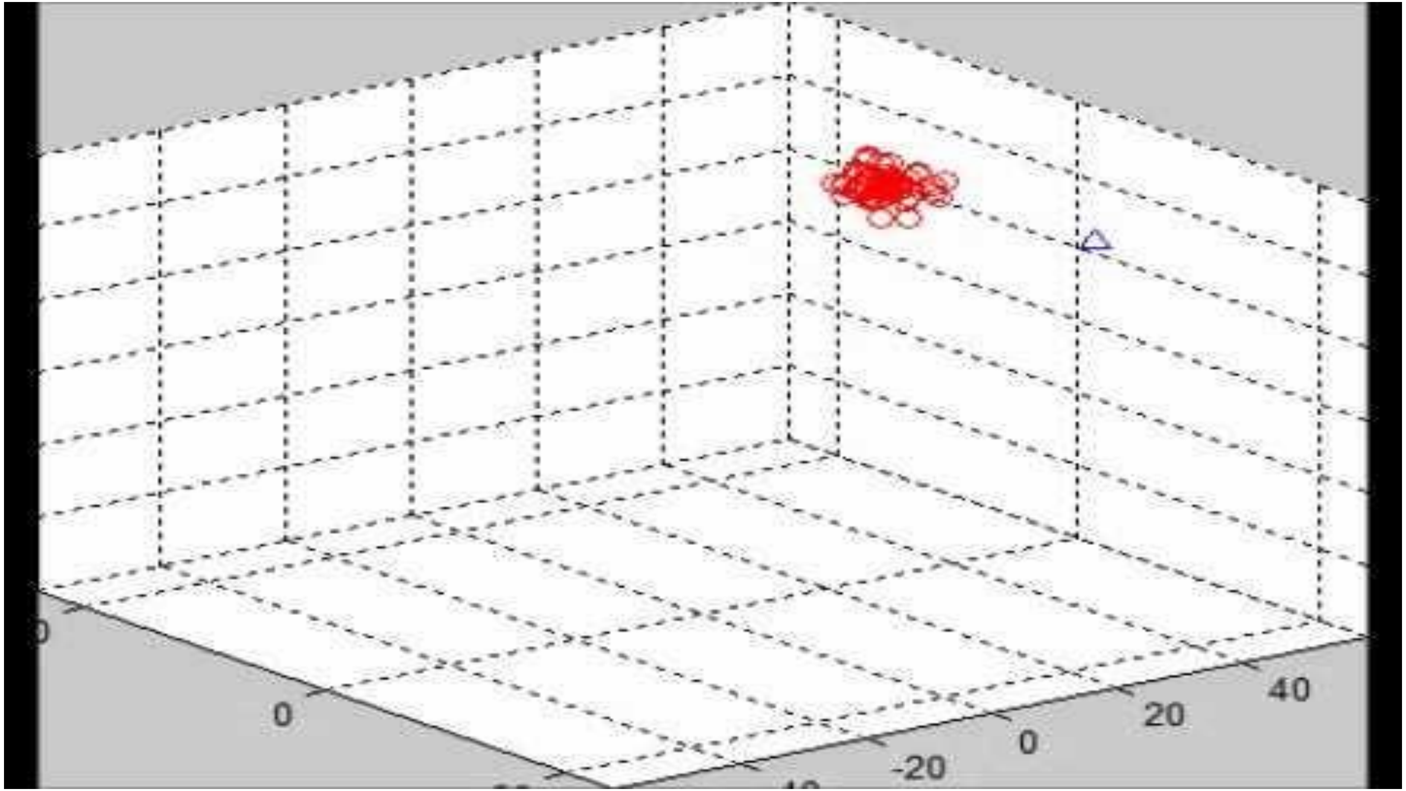
0
40
20
0
-20

Bayesian quadrature: Introduction

Imagine trying to find the value of the following definite integral

$$\int_0^1 \exp\left(\frac{(x-0.35)^2}{2(0.1)^2}\right)+\frac{\sin(10x)}{3}\,dx$$

. . . and you forgot most of calculus!

[Roman Garnett]

Trapezoid rule

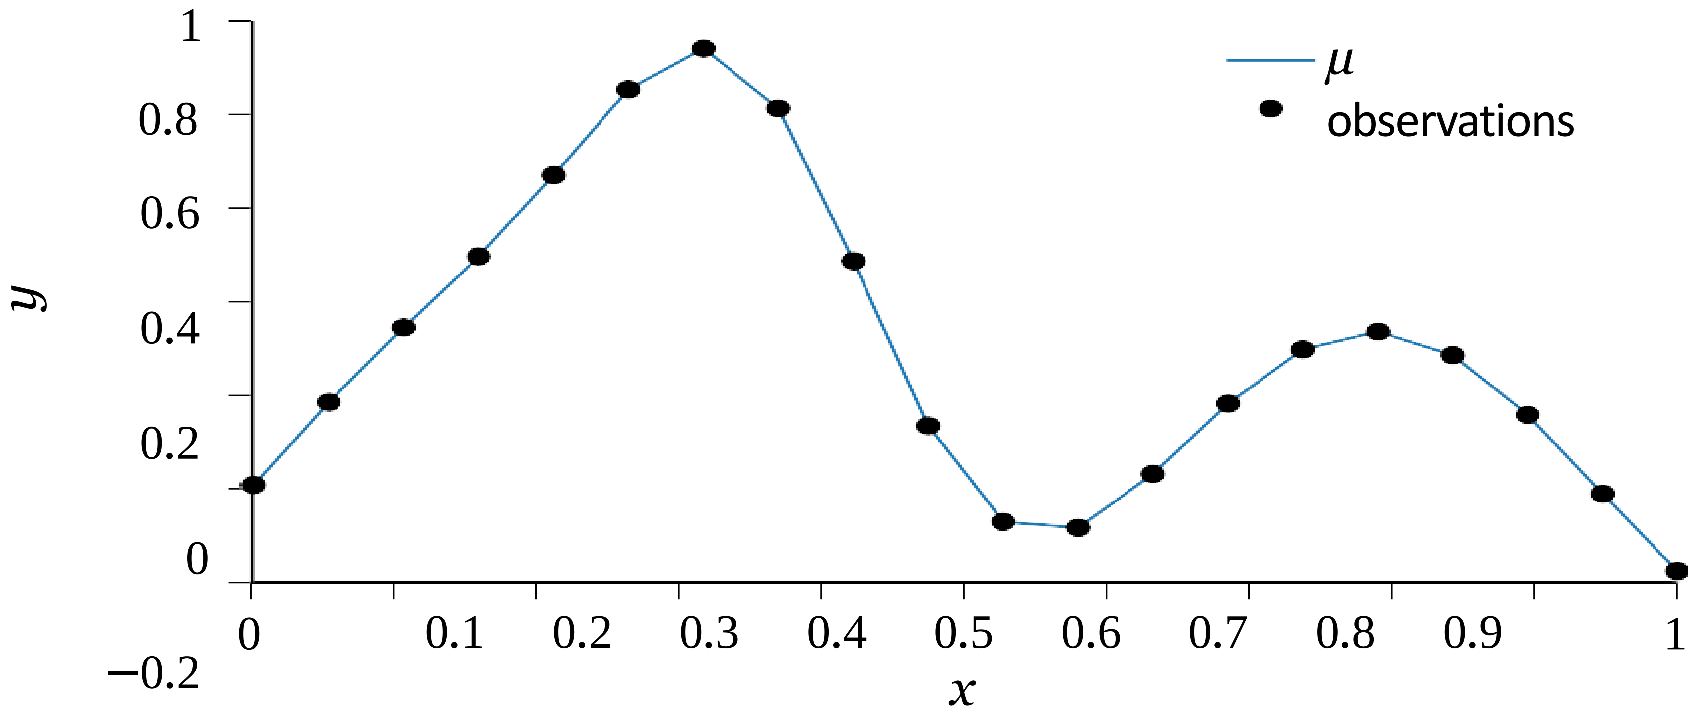

Questions

Here the trapezoid rule gives

$$\int_0^1 f(x)dx \approx 0.3104$$

(the true answer is ≈ 0.3119).

Questions:

- When should I stop?
- Where should I measure the function?

A Bayesian approach

Let's try a Bayesian approach.

Here we will treat the value of the integral

$$Z = \int_0^1 f(x)dx$$

as a random variable.

We will choose a prior for Z and use Bayes' rule to find the posterior distribution given our observations.

A Bayesian approach: Prior

It turns out that placing a prior on f rather than on Z directly is sometimes easier.

A Gaussian process (GP) is a convenient choice:

$$p(f) = GP(f; \mu, K).$$

Note : GPs are closed under affine transformations

A Bayesian approach: Example

Let's revisit our example.

We choose Brownian motion as our GP prior for f,

and estimate our integral by integrating the posterior mean. . .

Trapezoid rule

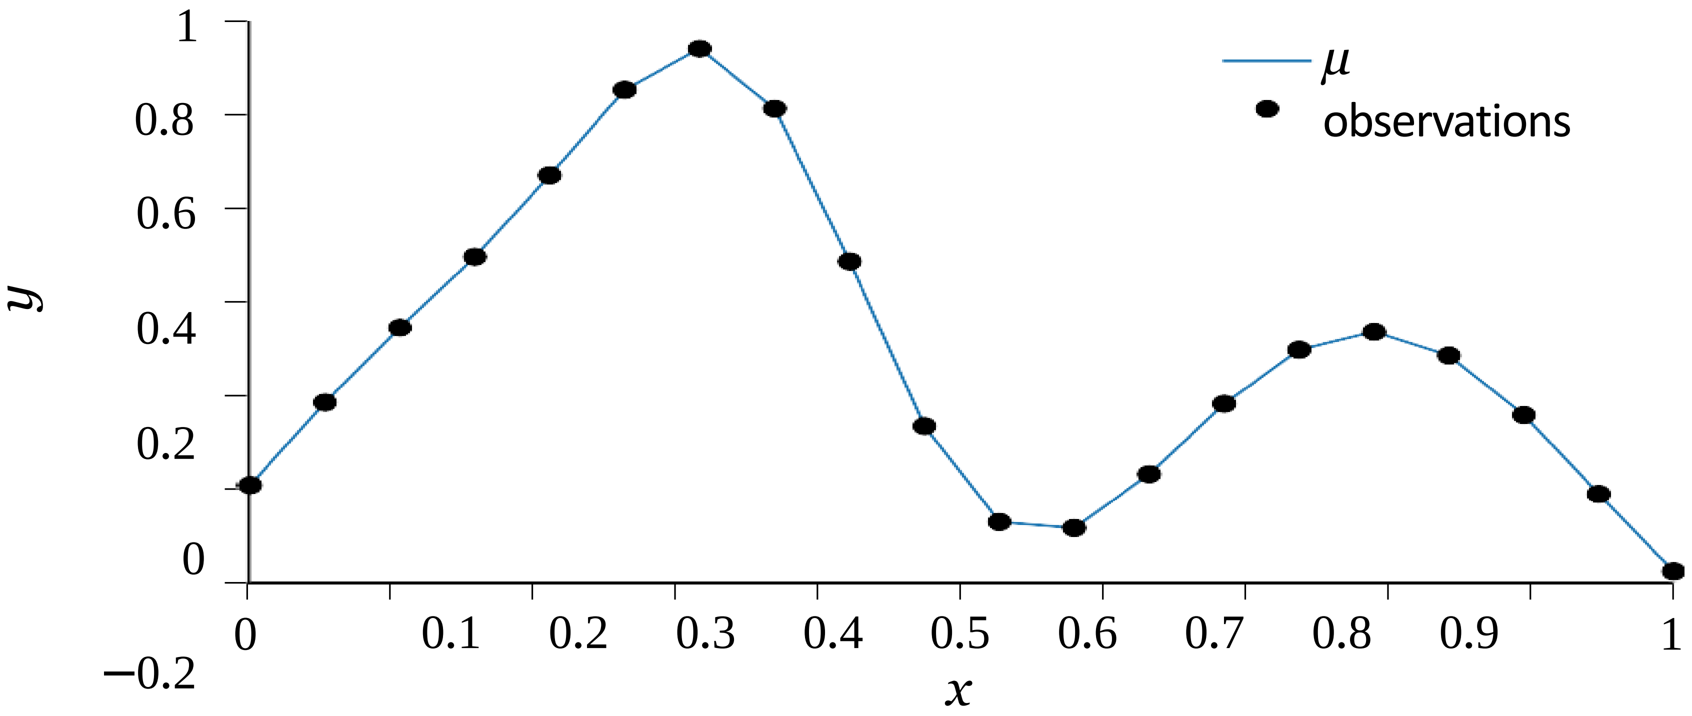

Why?

- What's the point?
- We can also quantify our uncertainty in the integral!

$$\mathrm{var}[Z \mid D] = \iint K_{f \mid D}\, dx\, dx'$$

- This can help us answer the previous questions:
  - When should I stop?
  - Where should I measure the function?

Uncertainty!

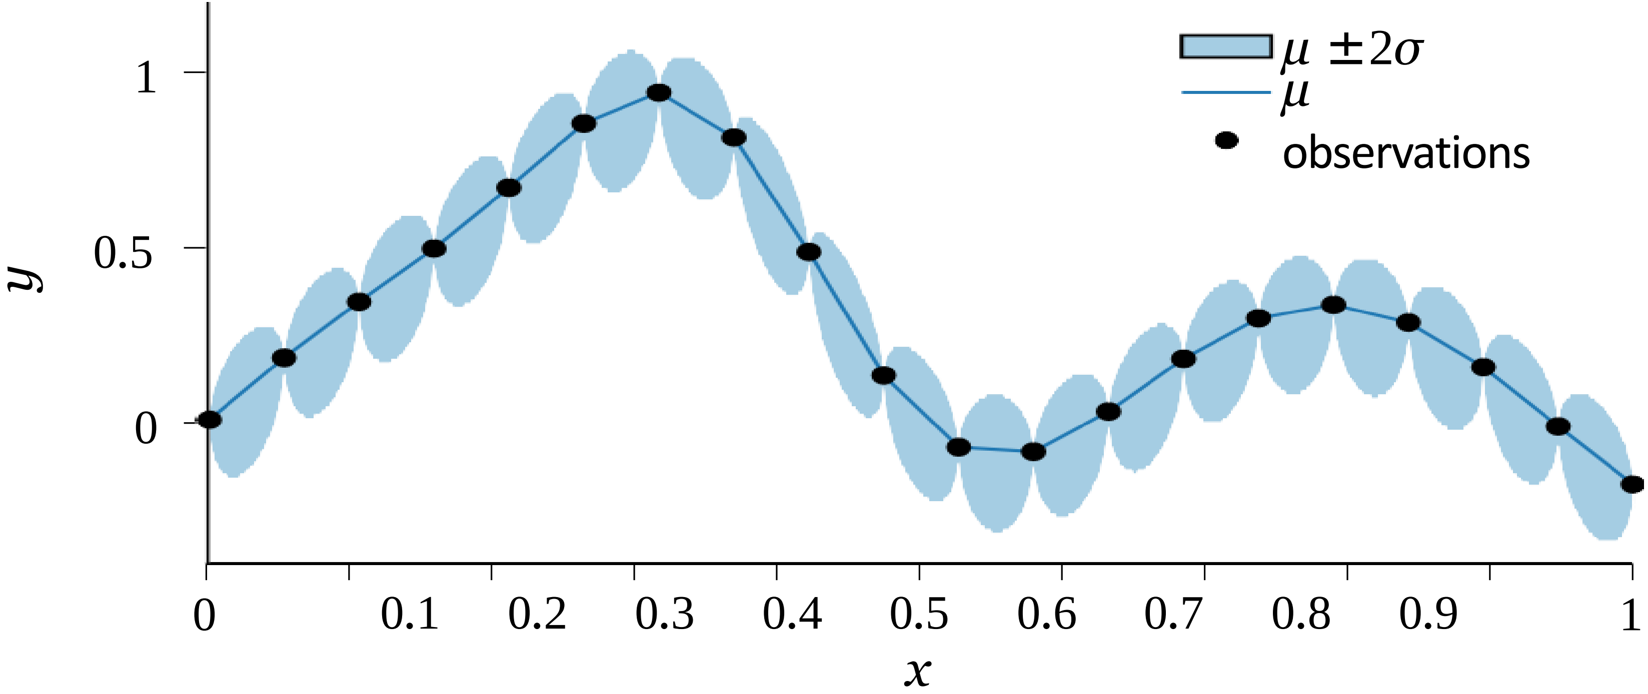

When should I stop?

- The magnitude of the uncertainty in Z can help us decide when to stop.

- For this example, we have

$$\text{sqrt}(\text{var}[Z \mid D]) \approx 0.015.$$

Where should I measure?

- The potential reduction in uncertainty in Z can also help us decide where to measure.

Summary

- A/B Testing
- Multi-Arm Bandit Problem
- Thompson Sampling
- Particle Swarms
- Bayesian Quadrature