Efficient Learning with Knowledge Gradients

Peter Frazier
Warren Powell
Savas Dayanik

Department of Operations Research and Financial Engineering
Princeton University

January 21, 2007
EEB Lab Tea Presentation

Outline

Outline

Problem Formulation

Examples

Knowledge Gradient Definition

Numerical and Theoretical Results

Drug Discovery Application

A pharmaceutical company has a library of millions of compounds that it would like to screen for potential cancer drugs. Robots will do the initial assay by performing a fixed test one or several times on some subset of the compounds.

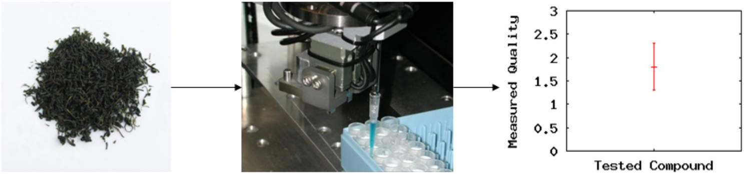

Candidate Compound — Perform Test on Compound — Noisy Measurement of Compound's Quality

Sources: http://www.paa.co.uk/img/labauto/inst_highres/ssi/mini_dispenser.jpg,

http://vitaltleaf.chainreactionweb.com/images/Angel.JPG

Drug Discovery Application: Assumptions

- We have a prior estimate of the quality of each compound.
- This prior estimate is normally distributed and is independent of estimates for other compounds.
- Each test has stationary independent unbiased normally distributed errors with known variance.
- After all tests are complete, we will choose only one compound to send to the next drug development stage.

Time 0

Suppose the number of compounds in the library is 5, and the number of tests allowed is $N = 10$. At time 0, before any tests are performed, we have a prior estimate of the quality of each compound.

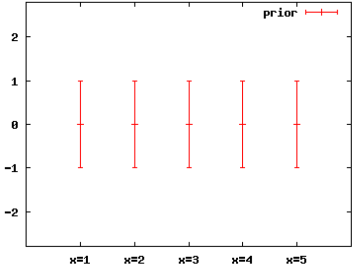

Time 1

We test compound $x = 1$ and measure its quality with noise. We update our estimate of the value of compound 1.

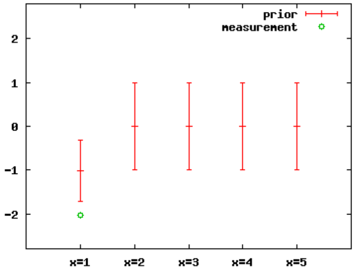

Time 2

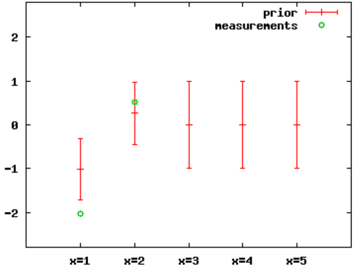

Time 5

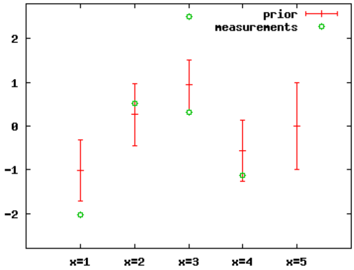

Final Time 10

After our experimental budget of $N = 10$ tests is exhausted, we choose the compound that is the best according to our current estimate, and we deliver it to the next stage of drug development.

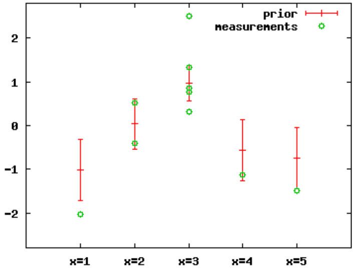

Final Time 10

After our experimental budget of $N = 10$ tests is exhausted, we choose the compound that is the best according to our current estimate, and we deliver it to the next stage of drug development.

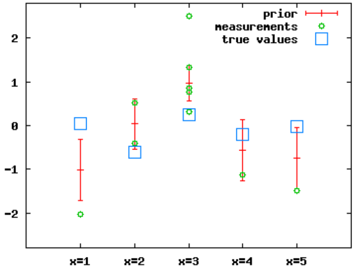

Outline

Problem Formulation

Examples

Knowledge Gradient Definition

Numerical and Theoretical Results

Example 1

One common experimental design is to spread our measurement budget equally across all alternatives.

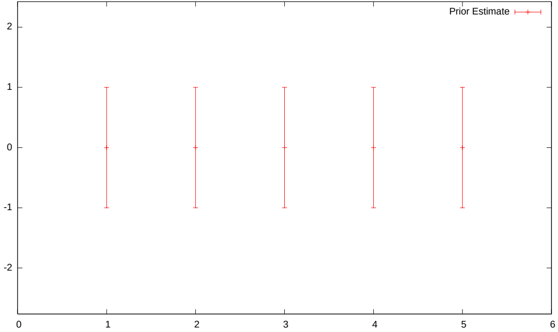

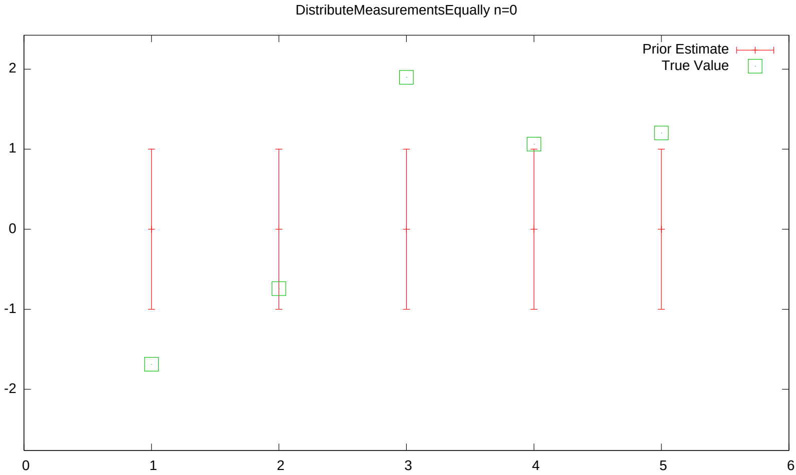
DistributeMeasurementsEqually n=0
Prior Estimate
True Value
2
1
0
-1
-2
0
1
2
3
4
5
6

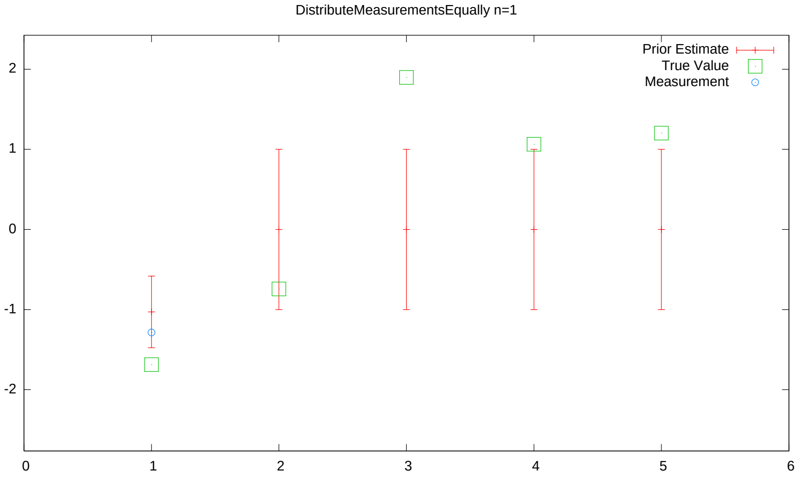
DistributeMeasurementsEqually n=1
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

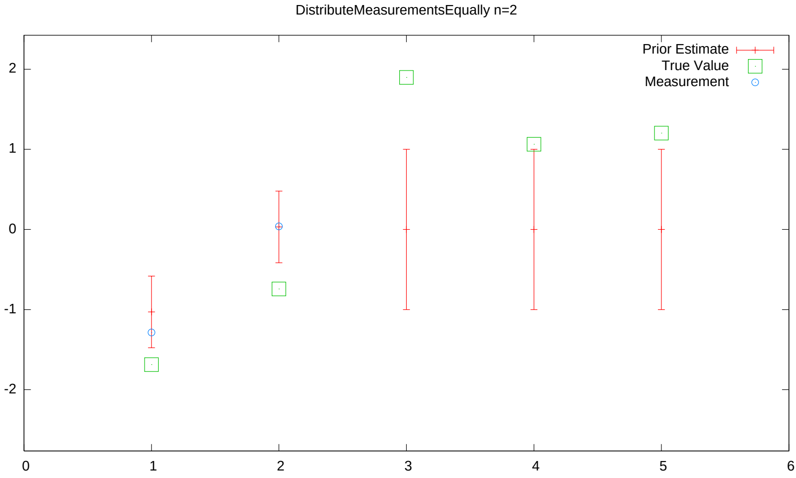
DistributeMeasurementsEqually n=2
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

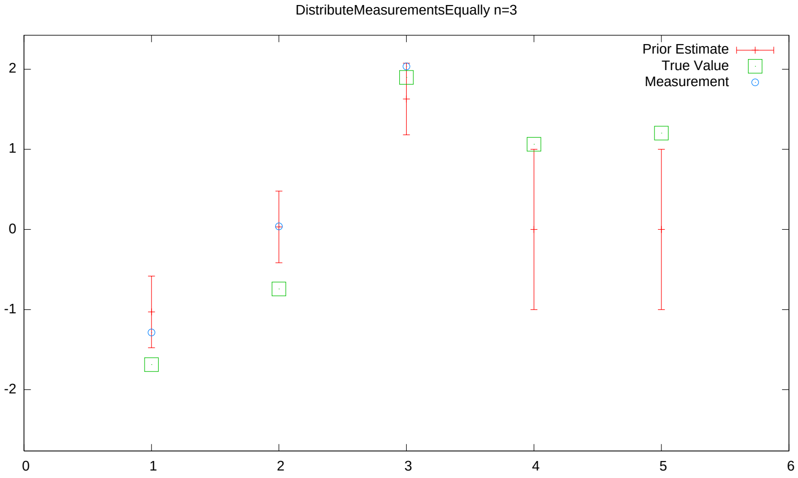
DistributeMeasurementsEqually n=3
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

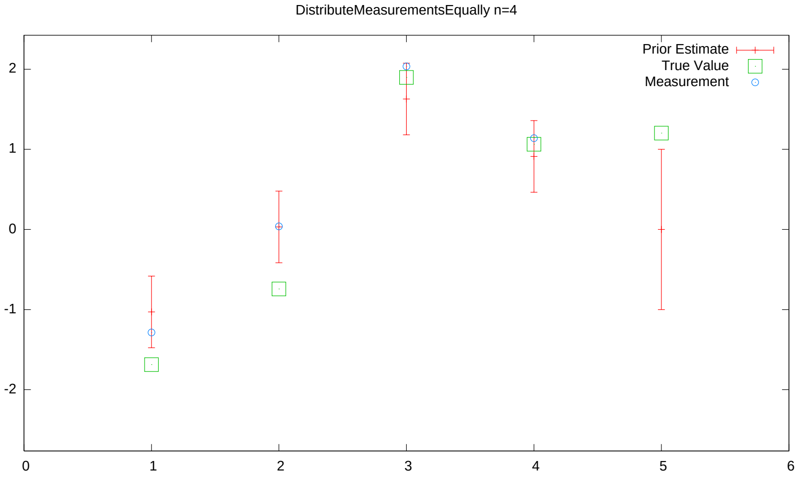
DistributeMeasurementsEqually n=4
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

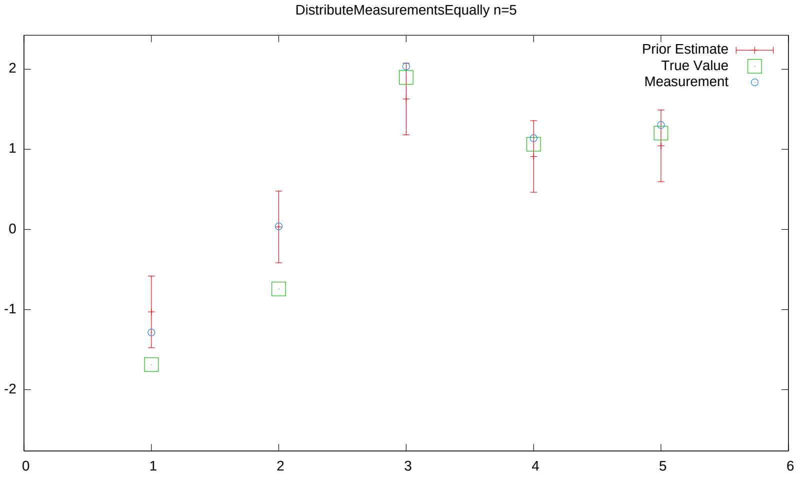
DistributeMeasurementsEqually n=5
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

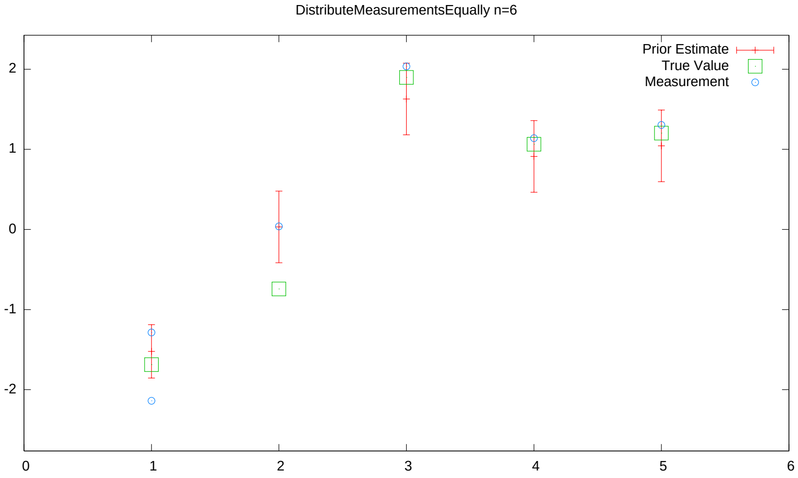
DistributeMeasurementsEqually n=6
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

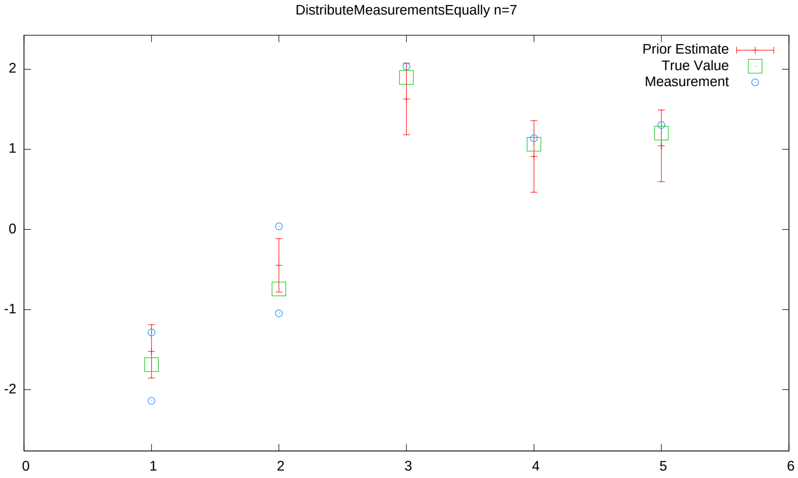
DistributeMeasurementsEqually n=7
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

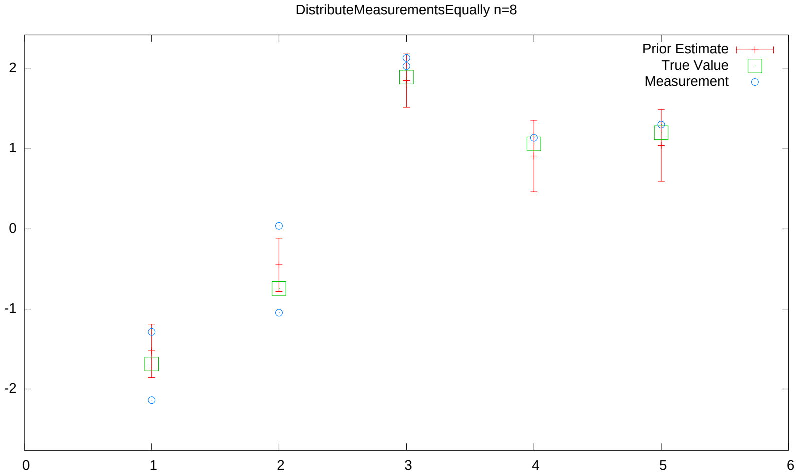
DistributeMeasurementsEqually n=8
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

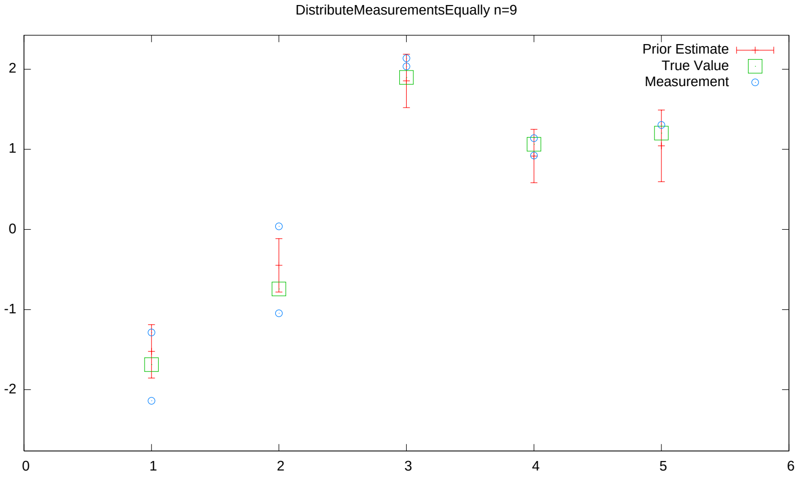
DistributeMeasurementsEqually n=9
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

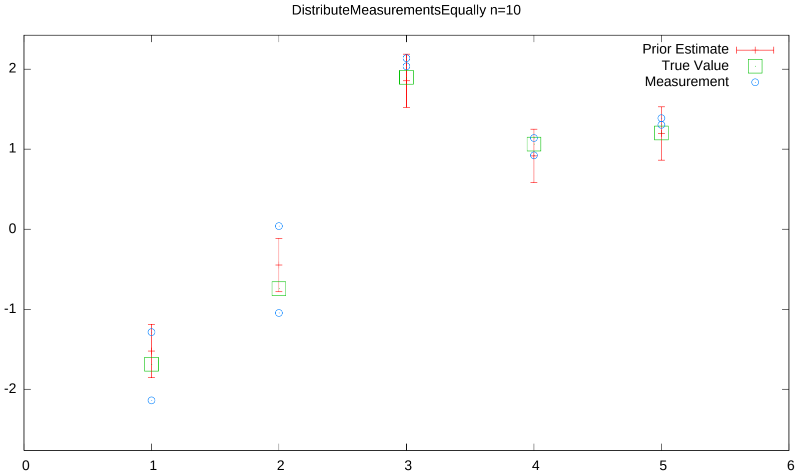
DistributeMeasurementsEqually n=10
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

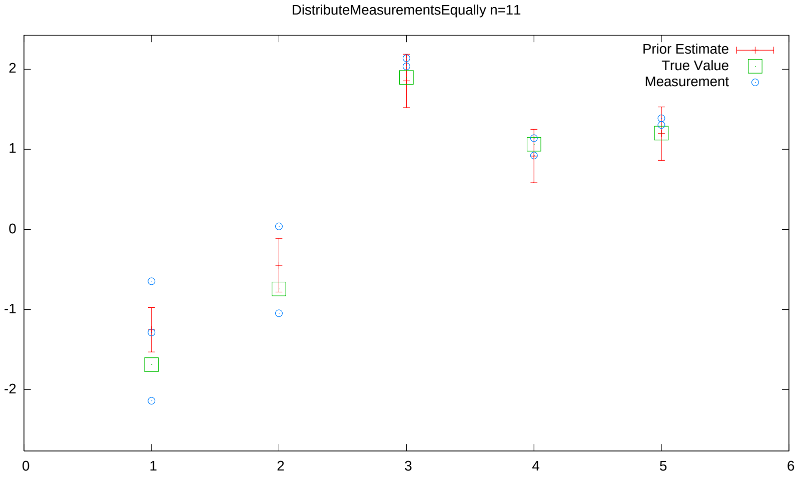
DistributeMeasurementsEqually n=11
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

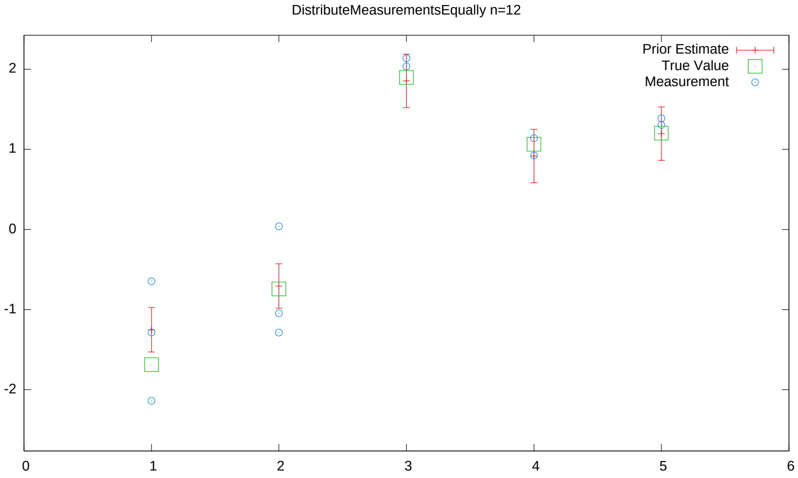
DistributeMeasurementsEqually n=12
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

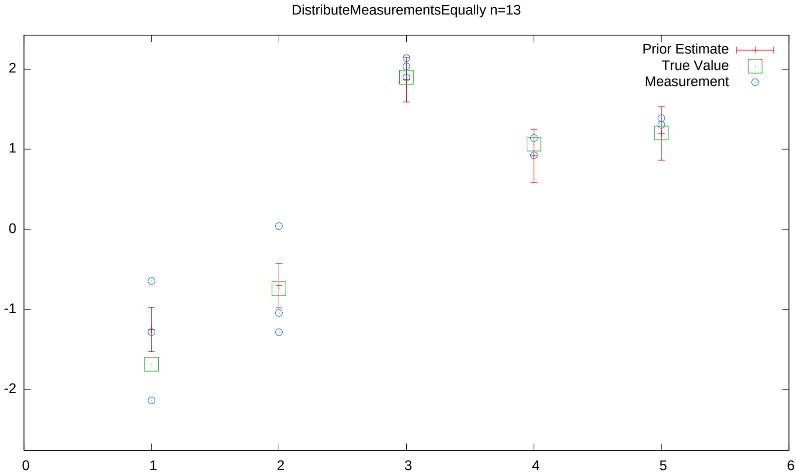
DistributeMeasurementsEqually n=13
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

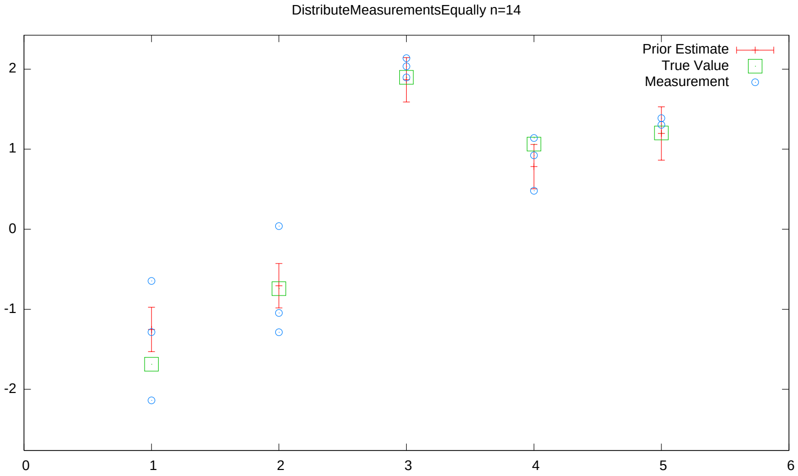
DistributeMeasurementsEqually n=14
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

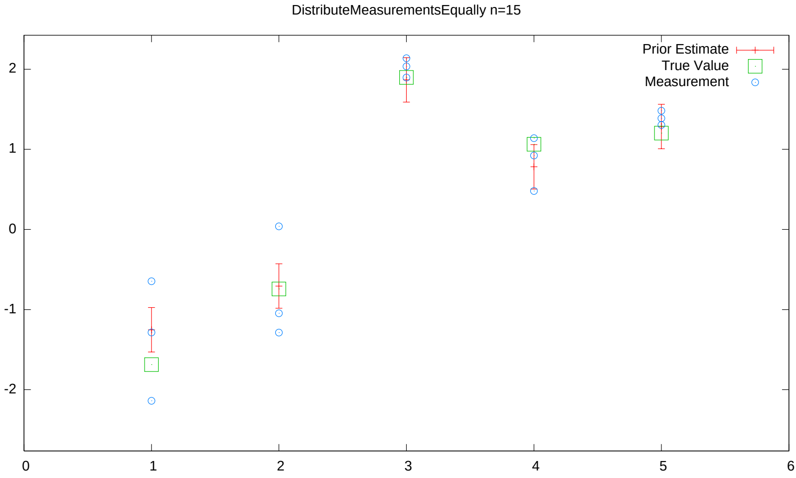
DistributeMeasurementsEqually n=15
Prior Estimate
True Value
Measurement
2
1
0
-1
-2
0
1
2
3
4
5
6

Example 2

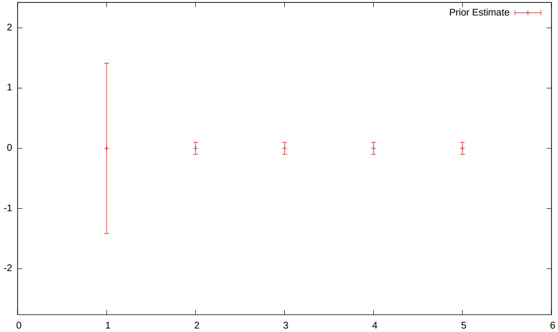

Example 2

When some compounds have larger variance, we want to measure these more often.

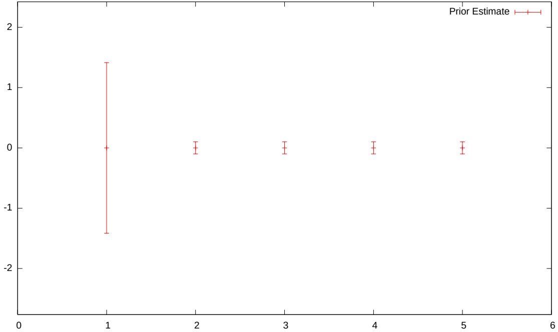

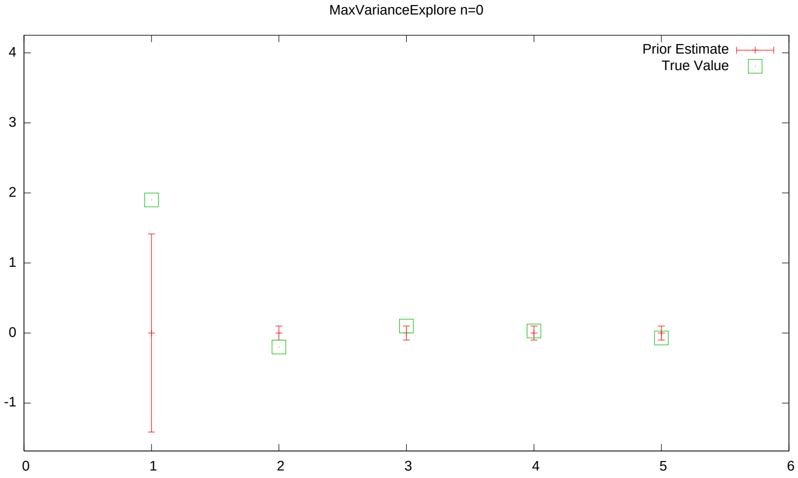
MaxVarianceExplore n=0
Prior Estimate
True Value
4
3
2
1
0
-1
0
1
2
3
4
5
6

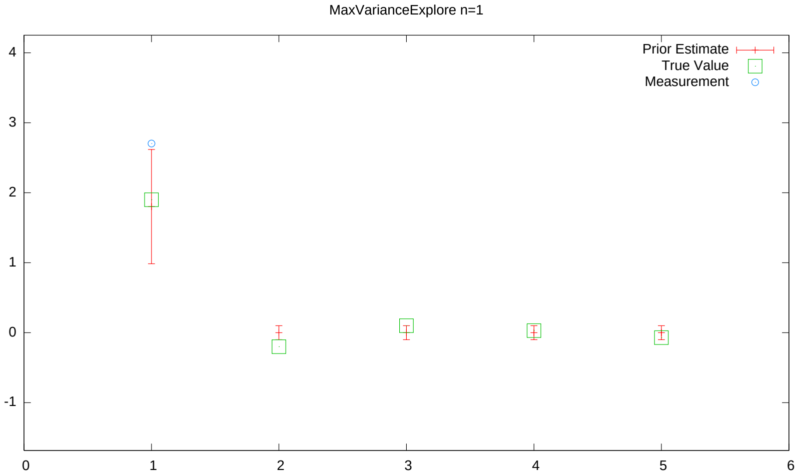
MaxVarianceExplore n=1
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

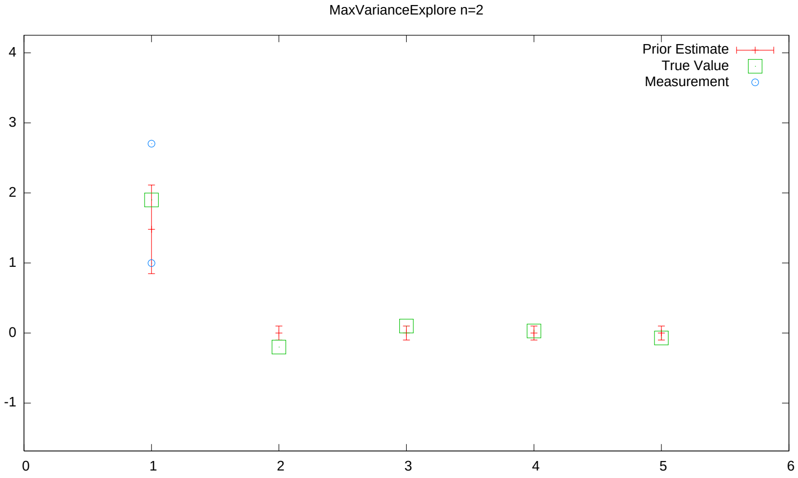
MaxVarianceExplore n=2
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

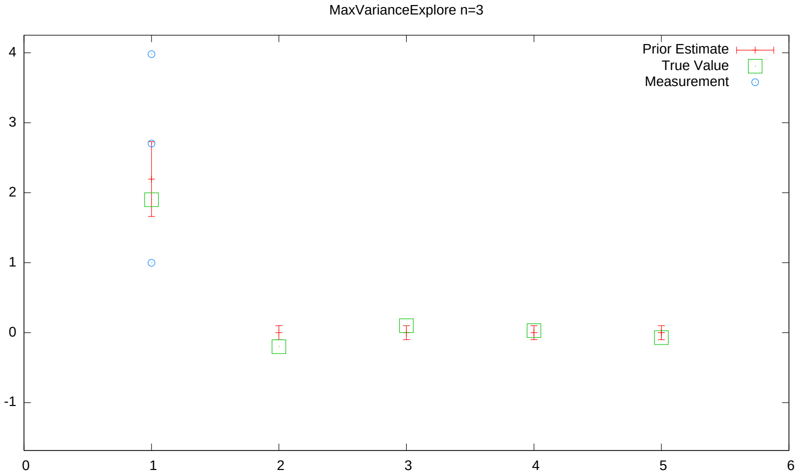
MaxVarianceExplore n=3
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

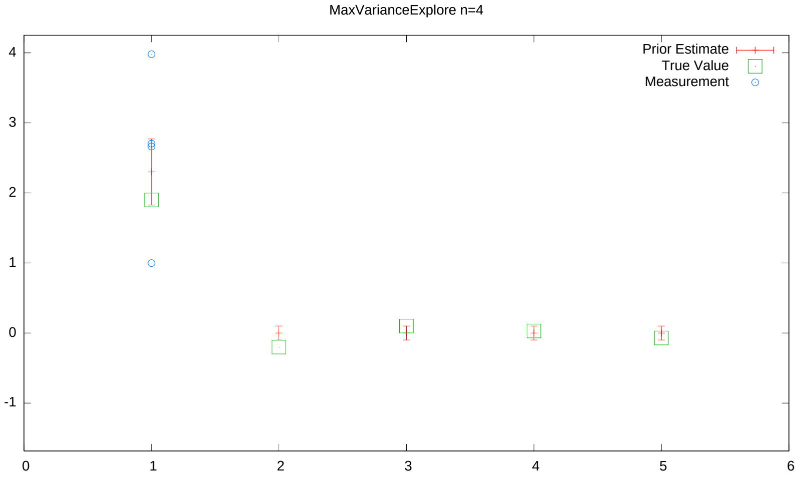
MaxVarianceExplore n=4
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

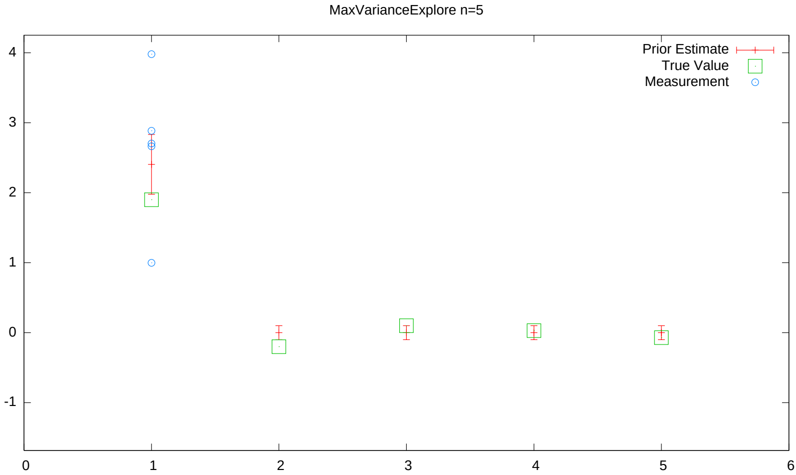

MaxVarianceExplore n=5
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

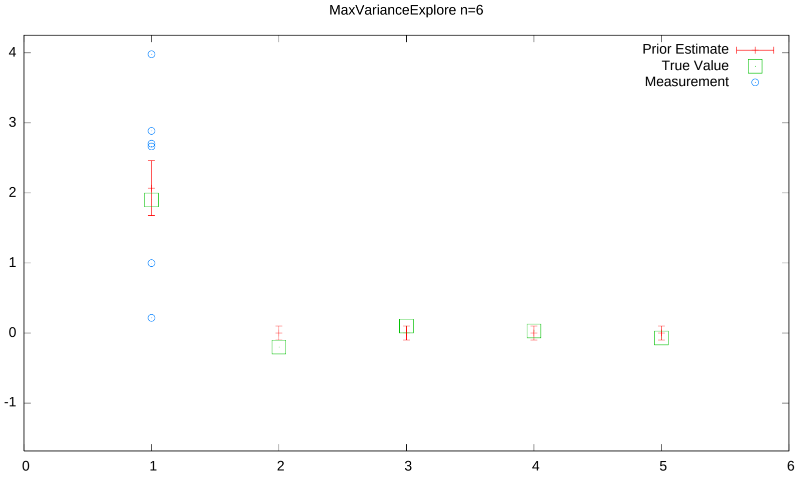
MaxVarianceExplore n=6
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

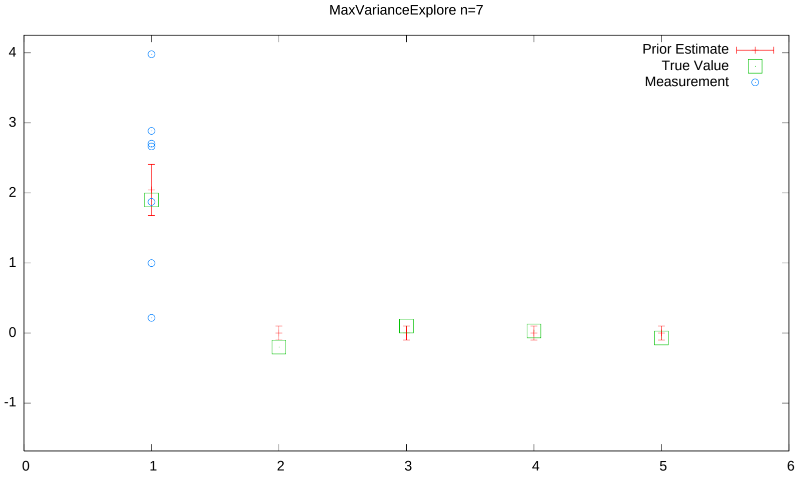
MaxVarianceExplore n=7
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

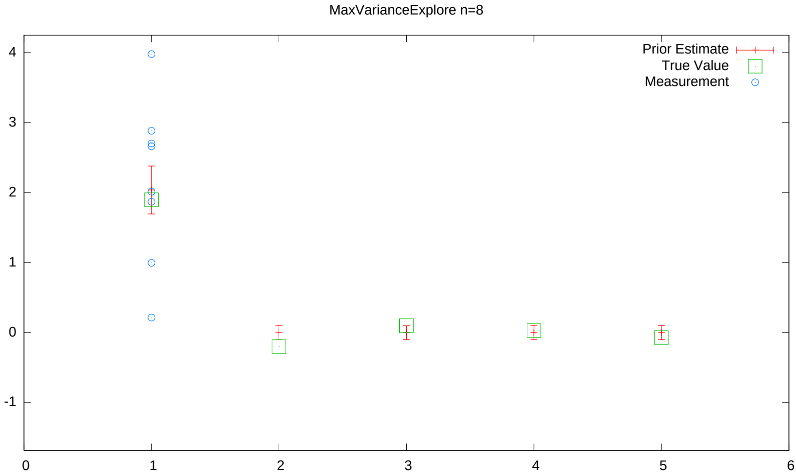
MaxVarianceExplore n=8
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

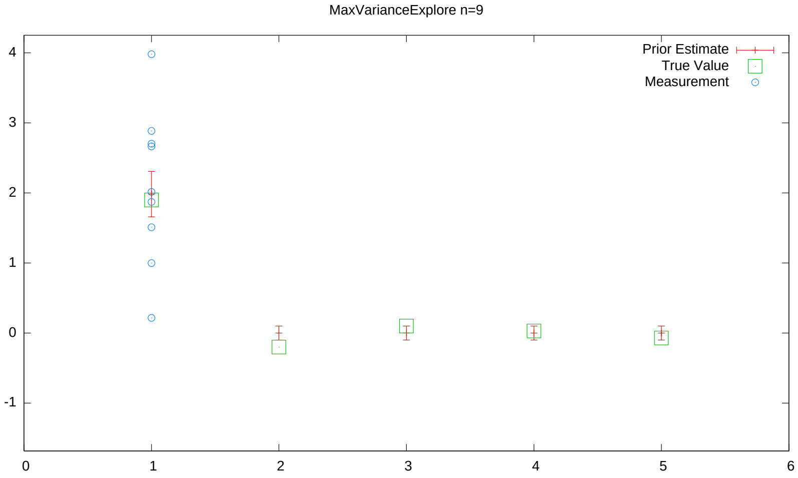
MaxVarianceExplore n=9
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

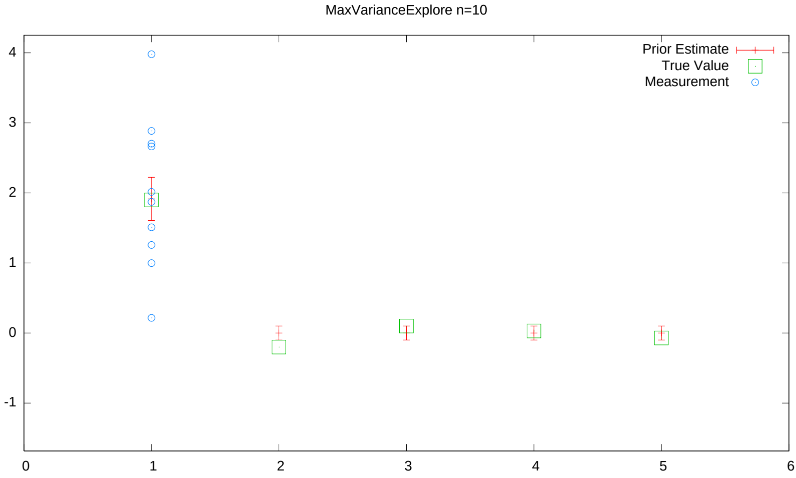
MaxVarianceExplore n=10
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

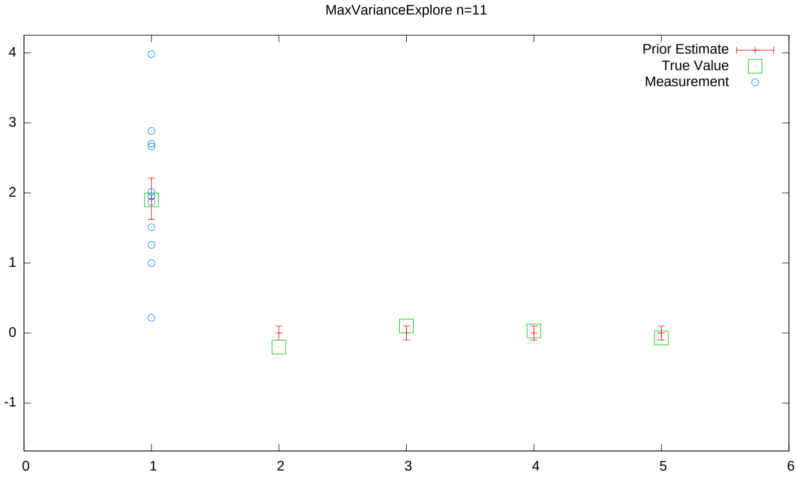
MaxVarianceExplore n=11
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

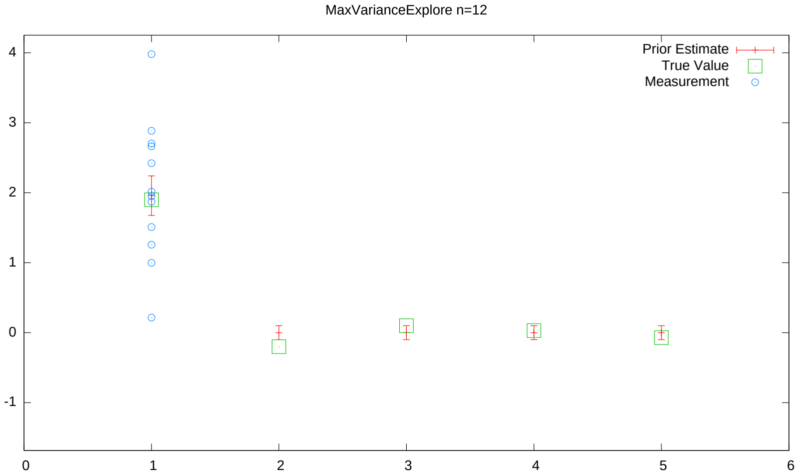
MaxVarianceExplore n=12
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

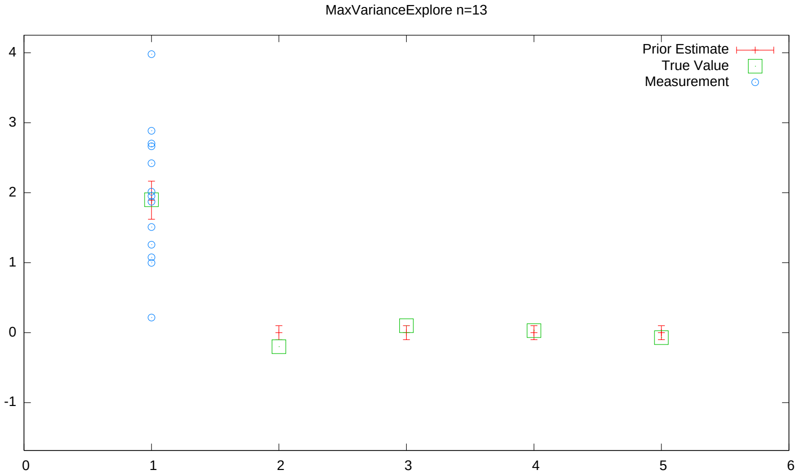
MaxVarianceExplore n=13
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

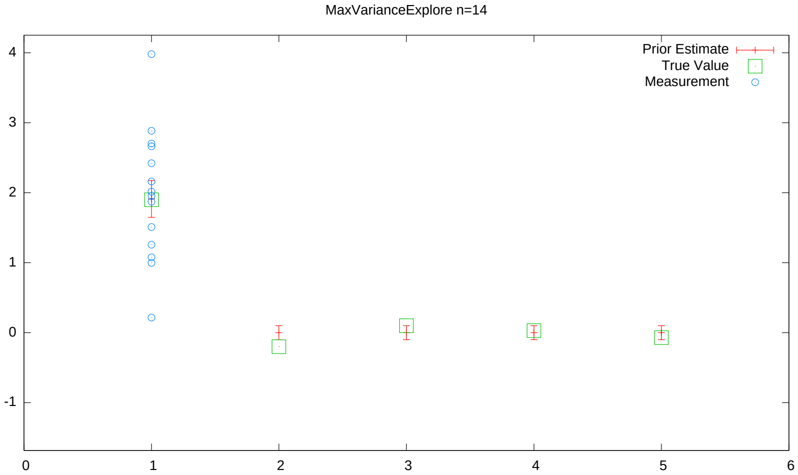
MaxVarianceExplore n=14
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

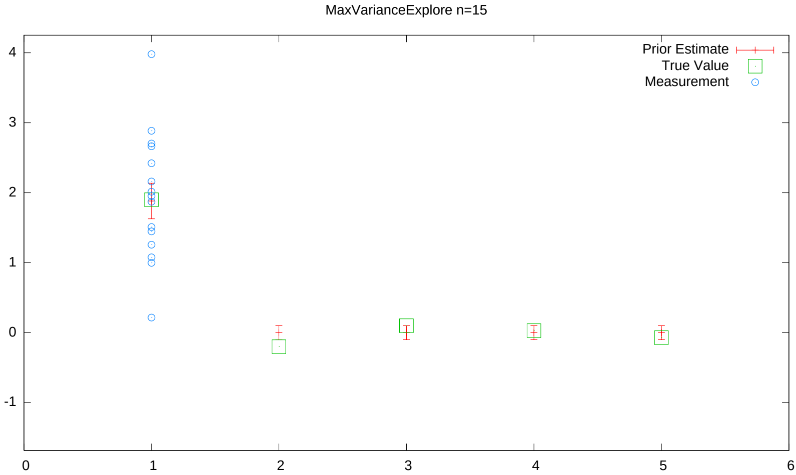
MaxVarianceExplore n=15
Prior Estimate
True Value
Measurement
4
3
2
1
0
-1
0
1
2
3
4
5
6

Example 3

Allocating measurements based soley on variance can also be inefficient. There is less value in measuring compounds whose estimated quality is very poor.

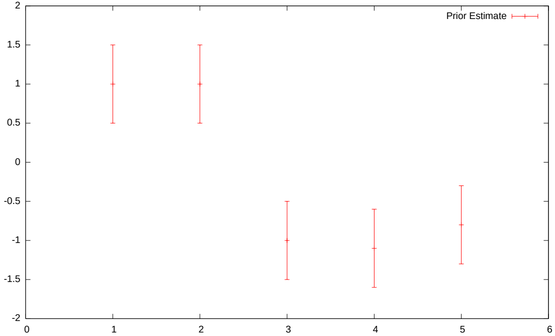

Example 4

In this example, compound 5 has the largest variance but compound 4 has the largest mean. Which compound should we test?

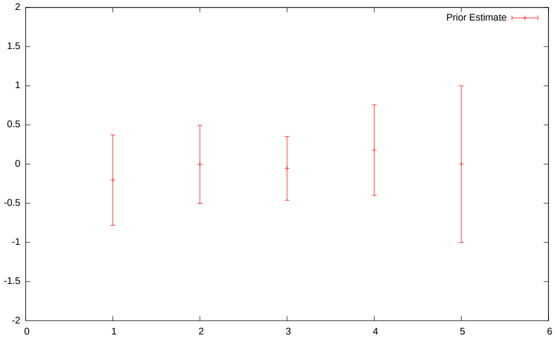

Goal

- Can we find the best measurement strategy?
  - Not in general, though we can in some special cases.
  - Dynamic programming tells us what the best measurement strategy is in general, but the "curse of dimensionality" prevents us from computing it.
- If we cannot find the best, can we find a measurement strategy that is "good enough"?

Outline

Problem Formulation

Examples

Knowledge Gradient Definition

Numerical and Theoretical Results

The Information State

- Our estimate of the quality of compound x after n tests have been performed takes the form of a normal distribution.
- Define μ_x^n to be the estimate's mean and Σ_{xx}^n its variance.
- Then, our estimate of the quality of all compounds is completely represented by the vector S^n

$$S^n := (\mu_1^n, \ldots, \mu_M^n, \Sigma_{11}^n, \ldots, \Sigma_{MM}^n)$$

We call S^n the "information state."

Measurement's Effect on the Information State

If we test compound x^n at time n, our estimate of the value of compound x changes randomly but our other estimates remain the same.

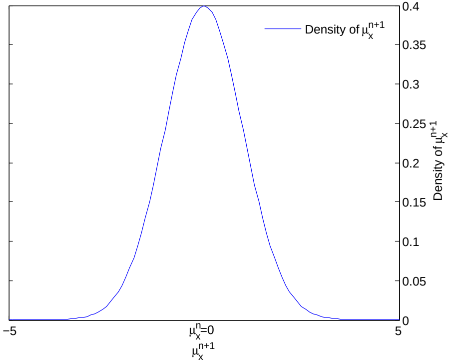

Knowledge Definition

We define the amount of "knowledge" contained in our information state S^n as

$$K(S^n) := \max_x \mu_x^n.$$

When the experimental budget is exhausted, $K(S^N)$ is the expected value of our drug assay. This is because we choose for further development the compound with maximum mean according to our estimate.

Knowledge Gradient Definition

Consider the change in knowledge resulting from our measurement x^n at time n,

$$K(S^{n+1}) - K(S^n) = \max_x \mu_x^{n+1} - \max_x \mu_x^n$$

The knowledge gradient policy is the one that chooses its measurements to **maximize the expected gain in knowledge**.

$$X^{KG}(S^n) := \operatorname*{arg\,max}_{x^n \in \{1 \ldots M\}} \mathbb{E}_n \left[K(S^{n+1}) - K(S^n) \right]$$

Computation of the Knowledge Gradient

The Knowledge Gradient policy may be computed using the formula

$$X^{KG}(S^n) = \arg\max_{x \in \{1 \ldots M\}} \tilde{\sigma}(\Sigma^n_{xx}) f(\zeta_x(S^n)).$$

where φ is the normal pdf, Φ is the normal cdf, and

$$\tilde{\sigma}(\Sigma_{xx}) := \Sigma_{xx} / \sqrt{\Sigma_{xx} + \left(\sigma^{(\varepsilon)}\right)^2}$$

$$\zeta_x(S^n) := -\left| \frac{\mu^n_x - \max_{x' \neq x} \mu^n_{x'}}{\tilde{\sigma}(\Sigma^n_{xx})} \right|$$

$$f(z) := z\Phi(z) + \varphi(z)$$

Numerical Example 1

$$X^{KG}(S^n) = \arg\max \tilde{\sigma}_x^n f(\zeta_x^n) = 1$$

x	μ_x	Σ_{xx}	$\tilde{\sigma}_x = \frac{\Sigma_{xx}}{\sqrt{\Sigma_{xx}+\sigma^2}}$	$\Delta_x = \lvert \mu_x - \max_{x' \neq x} \mu_{x'} \rvert$	$\zeta_x = \frac{-\Delta_x}{\tilde{\sigma}(\Sigma_{xx})}$	$\tilde{\sigma}_x f(\zeta_x)$
1	0	4	1.789	0	0	**0.714**
2	0	2	1.155	0	0	0.461
3	0	1	0.707	0	0	0.282

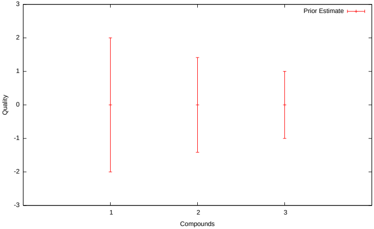

Numerical Example 2

$$X^{KG}(S^n) = \arg\max \tilde{\sigma}_x^n f(\zeta_x^n) = 2$$

x	μ_x	Σ_{xx}	$\tilde{\sigma}_x = \frac{\Sigma_{xx}}{\sqrt{\Sigma_{xx}+\sigma^2}}$	$\Delta_x = \lvert\mu_x - \max_{x' \neq x} \mu_{x'}\rvert$	$\zeta_x = \frac{-\Delta_x}{\tilde{\sigma}(\Sigma_{xx})}$	$\tilde{\sigma}_x f(\zeta_x)$
1	1	0	0	1	$-\infty$	0
2	0	1	0.707	1	-1.414	**0.025**
3	-1	1	0.707	2	-2.828	0.001

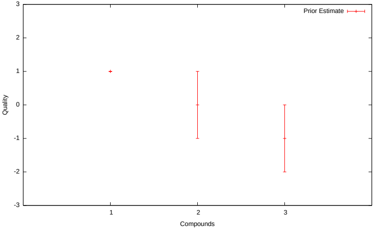

Numerical Example 3

$$X^{KG}(S^n) = \arg\max \tilde{\sigma}_x^n f(\zeta_x^n) = 2$$

x	μ_x	Σ_{xx}	$\tilde{\sigma}_x = \frac{\Sigma_{xx}}{\sqrt{\Sigma_{xx}+\sigma^2}}$	$\Delta_x = \lvert\mu_x - \max_{x' \neq x} \mu_{x'}\rvert$	$\zeta_x = \frac{-\Delta_x}{\tilde{\sigma}(\Sigma_{xx})}$	$\tilde{\sigma}_x f(\zeta_x)$
1	1	1	0.707	1	-1.41	0.0251
2	0	4	1.789	1	-0.56	**0.3223**
3	-1	1	0.707	2	-2.83	0.0005

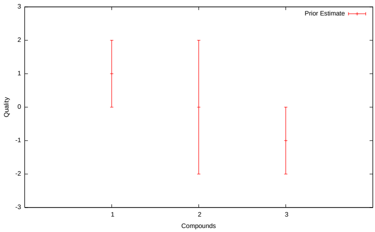

Outline

Vary N

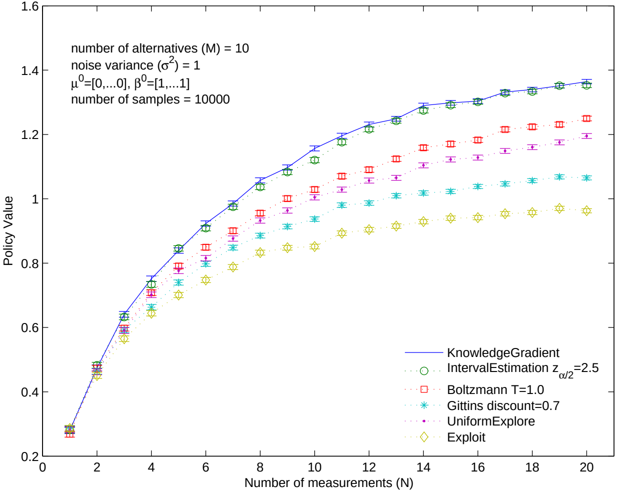
number of alternatives (M) = 10
noise variance (σ^2) = 1
μ^0=[0,...0], β^0=[1,...1]
number of samples = 10000
Policy Value
Number of measurements (N)
KnowledgeGradient
IntervalEstimation z_α/2=2.5
Boltzmann T=1.0
Gittins discount=0.7
UniformExplore
Exploit
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
2
4
6
8
10
12
14
16
18
20

Vary M

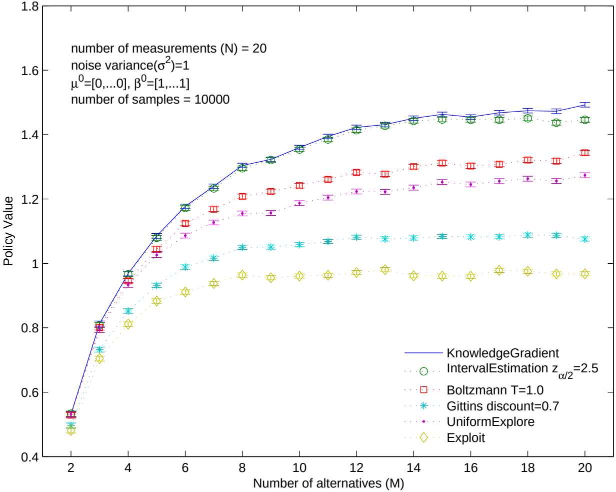

number of measurements (N) = 20
noise variance(σ^2)=1
μ^0=[0,...0], β^0=[1,...1]
number of samples = 10000
Policy Value
Number of alternatives (M)
KnowledgeGradient
IntervalEstimation z_α/2=2.5
Boltzmann T=1.0
Gittins discount=0.7
UniformExplore
Exploit
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
2
4
6
8
10
12
14
16
18
20

Vary N and M

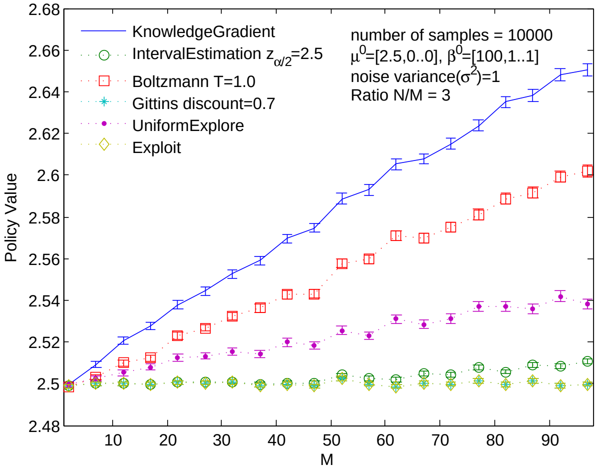
KnowledgeGradient
IntervalEstimation $z_{\alpha/2}$=2.5
Boltzmann T=1.0
Gittins discount=0.7
UniformExplore
Exploit
number of samples = 10000
μ^0=[2.5,0..0], β^0=[100,1..1]
noise variance(σ^2)=1
Ratio N/M = 3
Policy Value
M
2.68
2.66
2.64
2.62
2.6
2.58
2.56
2.54
2.52
2.5
2.48
10
20
30
40
50
60
70
80
90

Cases in which KG is Optimal

KG is optimal if **any** of the following conditions are met:

- $N = 1$
- In the limit as $N \to \infty$
- $M = 2$
- There is no measurement noise ($\sigma^2 = 0$) and the prior is ordered in both mean and variance ($\mu_1^0 \geq \mu_2^0 \geq \ldots \geq \mu_M^0$ and $\Sigma_{11}^0 \geq \Sigma_{22}^0 \geq \ldots \geq \Sigma_{MM}^0$).

Further Information Collection Applications

- A polling agency may wish to determine which candidate is most preferred by the electorate.
- An agronomist studying a number of varieties of wheat may seek to select the variety that will produce the largest number of bushels per acre.
- The Department of Homeland Security may wish to allocate their surveillance budget across some subset of targets to minimize the probability of a successful terrorist attack.
- During a disease outbreak, a doctor may wish to quickly determine which drug is most effective in treating the disease.
- A scientist running a Monte Carlo simulation may wish to find the set of input parameters that best fit an observed natural phenomenon.