Myopic Policies for Budgeted Optimization with Constrained Experiments

Javad Azami and **Xiaoli Fern** and **Alan Fern**
School of Electrical Engineering and Computer Science, Oregon State University

Elizabeth Burrows and **Frank Chaplen** and **Yanzhen Fan** and **Hong Liu**
Dept. of Biological & Ecological Engineering, Oregon State University

Jun Jiao and **Rebecca Schaller**
Department of Physics, Portland State University

Abstract

Motivated by a real-world problem, we study a novel budgeted optimization problem where the goal is to optimize an unknown function $f(x)$ given a budget. In our setting, it is not practical to request samples of $f(x)$ at precise input values due to the formidable cost of precise experimental setup. Rather, we may request a *constrained experiment*, which is a subset r of the input space for which the experimenter returns $x \in r$ and $f(x)$. Importantly, as the constraints become looser, the experimental cost decreases, but the uncertainty about the location x of the next observation increases. Our goal is to manage this trade-off by selecting a sequence of constrained experiments to best optimize f within the budget. We introduce cost-sensitive policies for selecting constrained experiments using both model-free and model-based approaches, inspired by policies for unconstrained settings. Experiments on synthetic functions and functions derived from real-world experimental data indicate that our policies outperform random selection, that the model-based policies are superior to model-free ones, and give insights into which policies are preferable overall.

Introduction

This work is motivated by the experimental design problem of optimizing the power output of nano-enhanced microbial fuel cells. Microbial fuel cells (MFCs) (Bond and Lovley 2003) use micro-organisms to break down organic matter and generate electricity. Research shows that MFC power outputs strongly depend on the surface properties of the anode (Park and Zeikus 2003). This motivates the design of nano-enhanced anodes, where nano-structures (e.g. carbon nano-wire) are grown on the anode surface to improve the MFC's power output. The experimental design goal is to optimize the MFC power output by varying the nano-enhancements to the anodes.

Traditional experimental design approaches such as response surface methods (Myers and Montgomery 1995) commonly assume that the experimental inputs can be specified precisely and attempt to optimize a design by requesting specific experiments. For example, requesting that an anode be tested with nano-wire of specific length and density.

However, when manufacturing nano-structures it is very difficult to achieve a precise parameter setting. Instead, it is more practical to request constrained experiments, which place constraints on these parameters. For example, we may specify intervals for the length and density of the nano-structures. Given such a request, nano-materials that satisfy the given set of constraints can be produced at some cost, which will typically increase with tighter constraints.

In this paper, we study the associated budgeted optimization problem where, given a budget, we must request a sequence of constrained experiments to optimize the unknown function. Prior work on experimental design, stochastic optimization, and active learning assume the ability to request precise experiments and hence do not directly apply to constrained experiments. Further, solving this problem involves reasoning about the trade-off between using weaker constraints to lower the cost of an experiment, while increasing the uncertainty about the location of the next observation, which has not been addressed in previous work.

As the first investigation into this problem, we present non-trivial extensions to a set of classic myopic policies that have been shown to achieve good performance in traditional optimization (Moore and Schneider 1995; Jones 2001; Madani, Lizotte, and Greiner 2004), to make them applicable to constrained experiments and to take variable cost into consideration. Our experimental results on both synthetic functions and functions generated from real-world experimental data indicate that our policies have considerable advantages compared to random policies and that the model-based policies can significantly outperform the model-free policies.

Contributions. We make the following contributions. First, motivated by a real-world application, we identify and formulate the problem of budgeted optimization with constrained experiments. This problem class is not rare and for example is common in nano-structure optimization and any other design domain where it is difficult to construct precise experimental setups. Second, we extend a set of classic experimental design heuristics to the setting of constrained experiments and introduce novel policies to take cost into consideration when selecting experiments. Finally, we empirically illustrate that our novel extension to traditional model-based policies can perform better than model-free policies.

Problem Setup

Let $\mathcal{X} \subseteq R^n$ be an n-dimensional input space, where each element of $\mathcal{X}$ is an experiment. We assume an unknown real-valued function $f : \mathcal{X} \to R$, which represents the expected value of the dependent variable after running an experiment. In our motivating domain, $f(x)$ is the expected power output produced by a particular nano-structure x. Conducting an experiment x produces a noisy outcome $y = f(x) + \epsilon$, where ϵ is a noise term.

Traditional optimization aims to find an x that approximately optimizes f by requesting a sequence of experiments and observing their outcomes. In this paper, we request *constrained experiments*, which specify the constraints that define a subset of experiments in $\mathcal{X}$. Specifically, a constrained experiment is defined by a hyper-rectangles $r = \langle r_1, \ldots, r_n \rangle$, where the r_i are non-empty intervals along input dimension i. Given a constrained experiment request r, the experimenter will return a tuple (x, y, c), where x is an experiment in r, y is the noisy outcome of x, and c is the cost incurred. The units of cost will vary across applications. In our motivating application, costs are dominated by the time required to construct an experiment satisfying r.

The inputs to our problem include a budget B, a conditional density $p_c(c|r)$ over the cost c of fulfilling a constrained experiment r, and a conditional density $p_x(x|r)$ over the specific experiment x generated for a constrained experiment r. Given the inputs, we then request a sequence of constrained experiments whose total cost is within the budget. This results in a sequence of m experiment tuples $(x_1, y_1, c_1), \ldots, (x_m, y_m, c_m)$ corresponding to those requests that were filled before going over the budget. After the sequence of requests we must then output an $\hat{x} \in \{x_1, \ldots, x_m\}$, which is our prediction of which *observed* experiment has the best value of f. This formulation matches the objective of our motivating application to produce a good nano-structure $\hat{x}$ using the given budget, rather than to make a prediction of what nano-structure might be good. In this paper we assume that p_c and p_x are part of the problem inputs and use simple forms for these. In general, these functions can be more complicated and learned from experience and/or elicited from scientists.

Related Work

This work is related to active learning, which focuses on requesting training examples to improve learning efficiency (Cohn et al. 1996). To our knowledge all such algorithms request specific training examples and are thus not applicable to our setting. Also, our objective of function optimization is quite different. Our work is more closely related to stochastic optimization and experimental design (Myers and Montgomery 1995), where many myopic policies have been developed (Moore and Schneider 1995; Lizotte, Madani, and Greiner 2003; Madani, Lizotte, and Greiner 2004). Given the current set of experiments D, these policies leverage statistical models of the data to select the next *unconstrained* experiment. There are two reasons that these existing policies cannot be directly applied here. First, they operate on individual experiments rather than constrained experiments.

Second, such policies do not account for variable experimental cost, and assume that each experiment has uniform cost. Indeed all of the experimental evaluations in (Moore and Schneider 1995; Lizotte, Madani, and Greiner 2003; Madani, Lizotte, and Greiner 2004), to name just a few, involve uniform cost experiments.

Myopic Policies

Drawing on existing work in traditional optimization (Moore and Schneider 1995; Jones 2001; Madani, Lizotte, and Greiner 2004), we extend existing myopic policies to our setting of constrained experiments. We consider both model-free policies that do not use predictive models of data, and model-based policies that use statistical data models.

Model-Free Policies.

Our model-free policies, *round robin (RR)* and *biased round robin (BRR)*, are motivated by previous work on budgeted multi-armed bandit problems (Madani, Lizotte, and Greiner 2004). In the multi-armed bandit setting the RR policy seeks to keep the number of pulls for each arm as uniform as possible. Here we apply the same principle to keep the experiments as uniformly distributed as possible. Given an experimental state (D, B), where D is the current set of observed experiments and B is the remaining budget, RR returns the largest hyper-rectangle r (i.e., with the least cost) that does not contain any previous experiment in D. If the expected cost of r exceeds budget B, we return the constrained experiment with expected cost B that contains the fewest experiments in D. Ties are broken randomly.

The *BRR* policy is identical to *RR* except that it repeats the previously selected r as long as it results in an outcome that improves over the current best outcome in D and has expected cost less than B. Otherwise, *RR* is followed. This policy is analogous to *BRR* for multi-armed bandit problems where an arm is pulled as long as it has a positive outcome.

Model-Based Policies.

For model-based policies, we learn a posterior distribution $p(y|x, D)$ over the outcome y of each individual experiment x from the current set of observed experiments D. Existing model-based policies for traditional experimental design typically select the experiment x that maximizes certain heuristics computed from statistics of the posterior (e.g. the mean or the upper confidence interval of y, see (Jones 2001)). The heuristics provide different mechanisms for trading-off exploration (probing unexplored regions) and exploitation (probing areas that appear promising given D). However, since they are defined for precise experiments they are not directly applicable to constrained experiments.

To define analogous heuristics for constrained experiments, we note that it is possible to use $p(y|x, D)$ to define a posterior distribution $\bar{p}$ over the outcome of any constrained experiment. In particular, for a constrained experiment r, we have $\bar{p}(y|r, D) \triangleq E_{X|r}[p(y|X, D)]$ where $X \sim p_x(\cdot|r)$. This definition corresponds to the process of drawing an experiment x from r and then drawing an outcome for x from $p(y|x, D)$. $\bar{p}$ effectively allows us to treat constrained experiments as if they were individual experiments in a traditional optimization problem. Thus, we can define heuristics

for constrained experiments by computing the same statistics of the posterior $\bar{p}$, as used in traditional optimization. In this work we consider four such heuristics.

Definition of Heuristics. The *maximum mean (MM)* heuristic computes the expected outcome according to the current posterior $\bar{p}$. That is, $\mathrm{MM}(r|D) = E[y|r, D]$ for $y \sim \bar{p}(\cdot|r, D)$. MM is purely exploitive, which has been shown to be overly greedy and prone to getting stuck in poor local optima (Moore and Schneider 1995).

The *maximum upper interval (MUI)* heuristic attempts to overcome the greedy nature of MM by exploring areas with non-trivial probability of achieving a good result as measured by the upper bound of the 95% confidence interval of y. The MUI heuristic for constrained experiments $\mathrm{MUI}(r|D)$ is equal to the upper bound of the 95% confidence interval of the current posterior $\bar{p}(y|r, D)$. Initially, MUI will aggressively explore untouched regions of the experimental space since y in such regions will have high posterior variance. As experimentation continues and uncertainty decreases, MUI will focus on the most promising areas and behave like MM.

The *maximum probability of improvement (MPI)* heuristic computes the probability of generating an outcome y that improves on the current best outcome y^* in D. The basic MPI only considers the probability of improvement and does not consider the *amount* of improvement. This deficiency is addressed by using a margin m with MPI, which we will refer to as $\mathrm{MPI}(m)$. Let y^* be the current best outcome in D, then $\mathrm{MPI}_m(r|D)$ is equal to the probability that the outcome of constrained experiment r will be greater than $(1 + m)y^*$. Increasing the margin m causes the heuristic to change its behavior from exploitive to explorative.

The *maximum expected improvement (MEI)* heuristic aims to improve on basic MPI without requiring a margin parameter. Specifically, it considers the expected value of improvement according to the current posterior. In particularly let $I(y, y^*) = max\{y - y^*, 0\}$, the MEI heuristic is defined as $\mathrm{MEI}(r|D) = E[I(y, y^*)|r]$ for $y \sim \bar{p}(y|r, D)$.

Note that these heuristics do not take the variable cost into account. If we assume that the cost of a constrained experiment r monotonically decreases with the size of its region of support, which is the most natural case to consider, it is easy to show that for all of our heuristics, the value of a constrained experiment r is monotonically non-decreasing with respect to the cost of r. If a heuristic does not take the variable cost into account, it will select the constraint experiments that are maximally constraint and centered around the individual experiment maximizing the heuristics. Therefore, it might consume more budget than is warranted. This shows that there is a fundamental trade-off between the heuristic value and cost of constrained experiments that we must address. Below we introduce two novel cost-sensitive policies to address this tradeoff.

Constrained Minimum Cost (CMC) Policies. For any heuristic on constrained experiments $H(r|D)$, which assigns a score to constrained experiment r given a set of observed experiments D, we can define an associated CMC policy. Informally, the principle behind CMC is to select the least cost constrained experiment that both 1) "approximately optimizes" the heuristic value, and 2) its expected improvement is no worse than spending the same amount of budget on a set of random experiments. The first condition encourages constrained experiments that look good according to the heuristic. The second condition helps to set a limit on the cost we are willing to pay to achieve a good heuristic value. In particular, we will only be willing to pay a cost of c for a single constrained experiment r if the expected improvement achieved by r is at least as good as the expected improvement achieved by a set of random experiments with total cost c. Note that the second condition can always be satisfied since a random experiment is equivalent to a constrained experiment of the least possible cost.

To formalize this policy we first define "approximately optimize" more precisely by introducing a parameterized version of the CMC policy and then showing how the parameter will be automatically selected via condition 2. For a given heuristic H, let h^* be the value of the highest scoring constrained experiment. For a given parameter $\alpha \in [0, 1]$, the $\mathrm{CMC}_{H,\alpha}$ policy selects the least-cost constrained experiment that achieves a heuristic value of at least $\alpha \cdot h^*$. Formally, this is defined as:

$$\mathrm{CMC}_{H,\alpha}(D, B) = \arg \min_{r \in R}\{c(r) : H(r|D) \geq \alpha h^*, c(r) \leq B\}$$

where $c(r)$ is the expected cost of r.

The value of α controls the trade off between the cost of r and its heuristic value. When α is close to one then $\mathrm{CMC}_{H,\alpha}$ must choose constrained experiments that are close to the maximum heuristic value, which will be costly. Rather, when α is small, it will choose lower cost constrained experiments with lower heuristic values, which increases the uncertainty about the individual experiment that will be observed, providing some constrained exploration.

In preliminary work, we experimented with the $\mathrm{CMC}_{H,\alpha}$ policy and found that it was difficult to select a value of α that worked well across a wide range of optimization problems, cost structures, and budgets. In particular, when α is too small the effect of the heuristic guidance is minimized and when it is too large, the policy would often choose extremely costly constrained experiments, leaving little budget for the future. This motivated the introduction of the condition 2 above to help adaptively select an appropriate α at each call of the policy.

In particular, we select the largest value of α such that the experiment suggested by $\mathrm{CMC}_{H,\alpha}$ satisfies condition 2. To formalize this, let r_α be the constrained experiment returned by $\mathrm{CMC}_{H,\alpha}$. Let $\mathrm{EIR}(D, C)$ be the expected improvement (as defined for the MEI policy above) of selecting random experiments that total a cost of C, and $c(r_\alpha)$ be the expected cost of r_α. Our parameter-free CMC policy is defined as:

$$\mathrm{CMC}_H(D, B) = \mathrm{CMC}_{H,\alpha^*}(D, B)$$
$$\alpha^* = \max\{\alpha \in [0, 1] : \mathrm{MEI}(r_\alpha|D) \geq \mathrm{EIR}(D, c(r_\alpha))\}$$

In practice we compute $\mathrm{EIR}(D, C)$ via Monte Carlo simulation and α^* via a discretized line search. We provide further details on implementations in the next section. In our experiments, we have observed that this approach for setting α is robust across different problems and heuristics.

Cost Normalized MEI (CN-MEI). Here we define an alternative parameter-free policy, which selects the constrained experiment achieving the highest expected improvement *per unit cost*. In particular, the policy is defined as:

$$\text{CN-MEI}(D, B) = \arg \max_{r \in R} \left\{ \frac{\text{MEI}(r|D)}{c(r)} : c(r) \leq B \right\}.$$

In concept we could define CN policies for heuristics other than MEI. To limit our scope we choose to consider only MEI due to it nice interpretation as a rate of improvement.

This principle of normalizing an improvement measure by cost is a natural baseline and has been suggested in the context of other optimization problems. However, in most prior work the actual empirical evaluations involved uniform cost models and thus there is little empirical data regarding the performance of normalization. Thus, our empirical evaluation necessarily provides a substantial evaluation of this normalization principle.

Implementation Details

To compute the conditional posterior $p(y|x, D)$ required by our model-based heuristics, we use Gaussian process regression with a zero mean prior and a covariance specified by a Gaussian Kernel $K(x, x') = \sigma \exp -\frac{1}{2l}|x - x'|^2$, with kernel width $l = 0.02$ and signal variance $\sigma = y_{\max}^2$, where $y_{\max}$ is an upper bound on the outcome value. We did not optimize these parameters, but did empirically verify that the GPs behaved reasonably for our test functions.

Given $p(y|x, D)$ we need to compute the various heuristics defined in the previous section, which were defined over $\bar{p}(y|r, D)$ and need to be computed via integration over the region of a constrained experiment r. Because these integrals do not have closed forms, we approximate the integrals by discretizing the space of experiments and using summation in place of integration. In our experiments we used a discretization level of 0.01 (1/100 of the input range).

This implementation is most appropriate for low dimensional optimization. For higher dimensional applications, one can use Monte Carlo sampling to estimate all of the expectations, the complexity of which can be independent of the dimension, and use empirical gradient descent with random restarts to find an optimizing constrained experiment.

Experimental Results

Our goal is to evaluate the performance of the proposed policies in scenarios that resemble typical real-world scientific applications. In particular, we focus on low-dimensional optimization problems for two reasons. First, with typical budgets the number of total experiments is often limited, which makes optimizing over many dimensions impractical. Thus, the scientists typically select only a few key dimensions to consider. Second, in real-world applications, such as our motivating problem, it is prohibitively difficult to satisfy constraints on more than a couple of experimental variables. Thus, the most relevant scenarios for us and many other problems are moderate numbers of experiments and small dimensionality.

Table 1: Benchmark Functions	
Cosines	$1 - (u^2 + v^2 - 0.3cos(3\pi u) - 0.3cos(3\pi v))$ $u = 1.6x - 0.5, v = 1.6y - 0.5$
Rosenbrock	$10 - 100(y - x^2)^2 - (1 - x)^2$
Discont	$1 - 2((x - 0.5)^2 + (y - 0.5)^2)$ if $x < 0.5$ 0 otherwise

Test Functions

We evaluate our policies using five functions over $[0, 1]^2$. The first three functions: *Cosines*, *Rosenbrock*, and *Discontinuous* (listed in Table 1) are benchmarks widely used in previous work on stochastic optimization (Brunato, Battiti, and Pasupuleti 2006). Here we use a zero-mean Gaussian noise model with variance equal to 1% of the function range. Exploratory experiments with larger noise levels indicate similar qualitative performance, though absolute optimization performance of the methods degrades as expected.

The other two functions are derived from real-world experimental data involving hydrogen production and our MFC application. For the former data was collected in a study on biosolar hydrogen production (Burrows et al. 2009), whose goal was to maximize the hydrogen production of the cyanobacteria *Synechocystis sp. PCC 6803* by optimizing the growth media's pH and Nitrogen levels. The data set contains 66 samples uniformly distributed over the 2D input space. This data is used to simulate the true function by fitting a GP to it, picking kernel parameters via standard validation techniques. We then simulated the experimental design process by sampling from the posterior of this GP to obtain noisy outcomes for requested experiments.

For the MFC application, we utilize data from some initial MFC experiments using anodes coated with gold nanoparticles under different fabrication conditions. The construction of each anode is costly and requires approximately two days. Each anode was installed in an MFC whose current density was recorded at regular intervals. The temporally averaged current density is taken to be the dependent variable to be optimized by modifying two features of the nano-structure on each anode: average area of individual particles, and average circularity of individual particles. We select these features because they can be roughly controlled during the fabrication process and appear to influence the current density. Due to the high cost of these experiments, our data set currently contains just 16 data points, which are relatively uniformly distributed over the experimental space. Considering the sparsity we choose polynomial Bayesian regression to simulate the true function, using standard validation techniques to select model complexity. The simulated function is then visually confirmed by our collaborator.

Modeling Assumptions

Our experiments assume that the density $P_x(x|r)$ over experiments given a request r is uniform over r. We further assume that the density $P_c(c|r)$ over the cost given a request is a Gaussian distribution, with a standard deviation of 0.01 and a mean equal to $1 + \frac{\text{slope}}{l(r)}$, where $l(r) = \sum_{i=1}^{n} |r_i|$ measures the size of the constrained experiment r. In this setup, each experiment requires at least one unit cost, with the ex-

pected additional cost being inversely proportional to the size of the constrained experiment r. The value of *slope* dictates how fast the cost increases as the size of a constrained experiment decreases. This cost model in inspired by the fact that all experimental setups require a minimal amount of work, here unit cost, but that the amount of additional work grows with increased constraints given to the experimenter. All of our policies easily extend to other cost models.

Results

Given a budget and function f to be optimized, each policy run results in a set of observed experiments D. Let $\hat{x}$ be the experiment in D that is predicted to have the maximum expected outcome. The *regret* of the policy for a particular run is defined to be $y^* - f(\hat{x})$, where y^* is the maximum value of f. For each test function and choice of budget and cost structure (i.e. choice of slope), we evaluate each policy by averaging the regret over 200 runs. Each run starts with five randomly selected initial points and then policies are used to select constrained experiments until the budget runs out, at which point the regret is measured. To compare the regret values across different functions we report normalized regret values, which are computed by dividing the regret of each policy by the mean regret achieved by a random policy (i.e., a policy always selecting the entire input space). A normalized regret less than one indicates that an approach outperforms random, while a value greater than one suggests it is worse than random.

In our first set of experiments we fix the budget to be 15 units and consider the effect of varying the cost-model slope parameter over values 0.1, 0.15, and 0.3.

For the MPI heuristic we set the margin $m = 0.2$, which based on a small number of trials performed well compared to other values. Table 2 shows the results for individual functions, giving both the average normalized regrets over 200 trials and the standard deviations. To assess the overall performance, it also gives the normalized regrets for each policy averaged across our five functions at the bottom of the table. Below we review our main empirical findings.

Model-Free Policies. The overall average performance of *RR* and *BRR*, as shown in Table 2, improves over random by approximately 10% across different slopes. This shows that the heuristic of evenly covering the space pays off compared to *random*. *BRR* is observed to perform slightly better than *RR*, suggesting that the additional exploitive behavior of BRR pays off overall. Looking at the individual results, however, we see that for hydrogen and fuel-cell both BRR and RR perform worse than random. Further investigation reveals that this poor performance is because RR/BRR have a bias toward experiments near the center of the input space, and the hydrogen and fuel-cell functions have their optimal points near corners of the space. This bias is because constrained experiments (rectangles) are required to fall completely within the experimental space of $[0, 1]^2$ and there are fewer such rectangles that contain points near the edges and corners.

Model-Based Policies. We first observe from the averaged overall results that the model-based policies perform better

Table 2: Normalized regrets.

	slope = 0.1	slope = 0.15	slope = 0.30
	Cosines		
cn-MEI	0.569 (.34)	0.714 (.41)	0.826 (.42)
cmc-MEI	0.417 (.28)	0.514 (.40)	0.794 (.46)
cmc-MPI(0.2)	0.535 (.39)	0.584 (.42)	0.616 (.43)
cmc-MUI	0.797 (.40)	0.804 (.41)	0.817 (.40)
cmc-MM	0.708 (.48)	0.767 (.48)	0.736 (.45)
RR/BRR	0.842 / 0.833 (.40/.40)	0.866 / 0.862 (.40/.41)	0.897 / 0.889 (.41/.41)
	Discontinuous		
cn-MEI	0.527(.44)	0.497(.44)	0.626(.60)
cmc-MEI	0.564(.42)	0.677(.57)	0.779(.65)
cmc-MPI(0.2)	0.954(.82)	0.940(.74)	0.951(.78)
cmc-MUI	0.710(.69)	0.709(.79)	0.693(.64)
cmc-MM	1.289(1.1)	1.225(1.1)	1.116(1.1)
RR/BRR	0.602/0.587 (.51/.50)	0.617/0.602 (.53/.53)	0.634/0.634 (.56/.57)
	Rosenbrock		
cn-MEI	0.602(.40)	0.665(.49)	0.736(.60)
cmc-MEI	0.547(.35)	0.556(.39)	0.630(.47)
cmc-MPI(0.2)	0.503(.36)	0.594(.44)	0.608(.48)
cmc-MUI	0.805(.79)	0.974(.13)	0.913(1.0)
cmc-MM	0.721(.68)	0.740(.73)	0.662(.61)
RR/BRR	0.897/0.881 (.85/.83)	0.930/ 0.927 (.87/.90)	0.967/0.955 (1.0/.98)
	Hydrogen		
cn-MEI	0.176(.26)	0.354(.46)	0.852(.66)
cmc-MEI	0.129(.26)	0.233(.43)	0.420(.53)
cmc-MPI(0.2)	0.408(.63)	0.449(.71)	0.613(.73)
cmc-MUI	0.716(.60)	0.695(.61)	0.868(.65)
cmc-MM	0.728(.80)	0.605(.73)	0.691(.77)
RR/BRR	1.107/ 1.065 (.68/.68)	1.161/1.123 (.69/.70)	1.173/1.148 (.67/.67)
	Fuel Cell		
cn-MEI	0.929(.15)	0.950(.17)	0.986(.22)
cmc-MEI	0.931(.11)	0.908(.14)	0.940(.16)
cmc-MPI(0.2)	0.932(.16)	0.930(.19)	0.943(.20)
cmc-MUI	0.971(.19)	0.973(.23)	0.995(.21)
cmc-MM	0.945(.22)	0.963(.32)	0.963(.33)
RR/BRR	1.038/1.029 (.21/.20)	1.049/1.044 (.22/.21)	1.046/1.041 (.22/.22)
	Overall Average		
	slope = 0.1	slope = 0.15	slope = 0.30
cn-MEI	0.560	0.636	0.805
cmc-MEI	0.517	0.578	0.712
cmc-MPI(0.2)	0.666	0.698	0.746
cmc-MUI	0.800	0.831	0.857
cmc-MM	0.874	0.889	0.834
RR/BRR	0.897 / 0.879	0.925 / 0.911	0.944 / 0.934

than random. Looking at the results on individual functions, we see that all model-based policies outperform *random*, with the exception of CMC-MM on the Discontinuous function. This shows that our two approaches (CMC and CN) for considering cost in our policies are able to make sensible choices and avoid expending the budget more quickly than is warranted. In addition, the averaged overall results suggest that all model-based approaches outperform the model-free approaches. This indicates that the heuristics we are considering and our GP regression model are providing useful in-
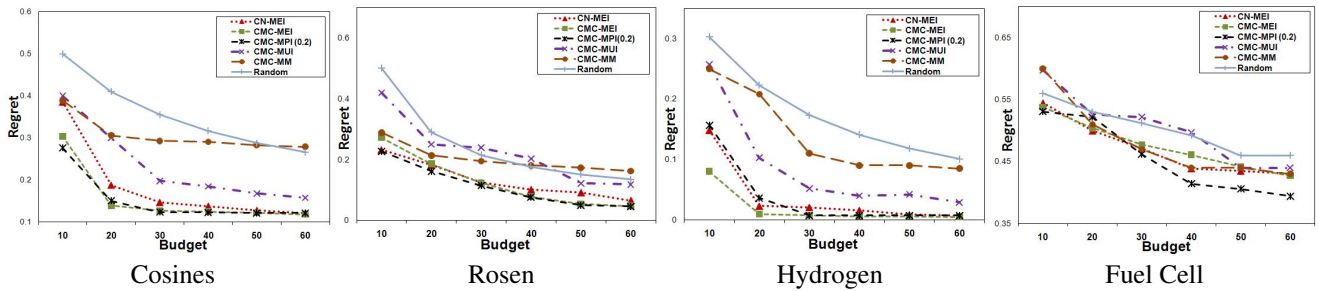
Figure 1: Regret *(unnormalized)* vs. Cost curves for 4 proposed functions at slope=0.1.

formation for guiding the constrained experiment selection.

Comparing different heuristics, we see that the MEI-based methods are the top contenders among all methods. Looking at the individual results, we can see that this holds for all functions except for Rosenbrock, where the CMC-MPI is slightly better than MEI-based methods. A closer look into the MPI and MEI heuristics reveals that MPI tends to select slightly fewer experiments than MEI, which we conjecture is because the MEI heuristic tends to be smoother than MPI over the experimental space. Based on these observations and the fact that it does not require selecting a margin parameter, we consider MEI as the most preferable heuristic to use.

Next, we compare the two different strategies for dealing with cost, i.e., CMC and CN. Specifically, we compare CMC-MEI with CN-MEI and observe that CMC-MEI outperforms CN-MEI for all but the Discontinuous function. When examining the behaviors of these two policies, we observed that CMC-MEI tends to be more conservative regarding the cost it selects early on in the budget, choosing constrained experiments that were closer to random. This behavior is desirable since exploration is the primary issue when few experiments have been performed and random is a reasonably good and very cheap exploration policy.

Varying Budgets. In the last set of experiments, we fixed the cost model slope to 0.1 and varied the budget from 10 to 60 units in increments of 10. The goal is to observe the relative advantage of the model-based policies compared to random as the budget increases. In addition, we want to see how the model-based policies work and what is the difference among them given larger budgets.

Figure 1 plots the *absolute* regret (unnormalized) versus cost for our model based policies on four test functions. As the main observation, we see that policies based on MEI and MPI maintain a significant advantage over random across a wide range of budgets. The CMC-MEI and CMC-MPI policies are roughly comparable for all functions except for the fuel cell function. In that case CMC-MPI slightly outperforms CMC-MEI for large budgets.

Overall, given the results from the previous experiment, CMC-MEI can still be considered as a recommended method, due to its combination of good performance, robustness and the fact that it is parameter-free.

Conclusion

This paper makes the following key contributions: 1) Motivated by a real-world application we formulated the novel budgeted optimization problem with constrained experiments; 2) we extended a number of classic experimental design heuristics to the case of constrained experiments; 3) we introduced novel strategies (CMC and CN) to take cost into consideration in selecting constrained experiments; 4) we empirically evaluated the proposed heuristics and cost-sensitive policies. Our results demonstrate that the proposed cost-sensitive policies significantly outperform both random selection as well as the model-free policies. Overall we found that the CMC-MEI policy demonstrated robust performance and is parameter-free, making it a recommended method. For future work, we plan to investigate non-myopic policies (e.g., monte-carlo planning techniques) and richer cost and action models that more closely match the needs of real budgeted optimization applications.

References

Bond, D., and Lovley, D. 2003. Electricity production by geobacter sulfurreducens attached to electrodes. *Applications of Environmental Microbiology* 69:1548–1555.

Brunato, M.; Battiti, R.; and Pasupuleti, S. 2006. A memory-based rash optimizer. In *AAAI-06 Workshop on Heuristic Search, Memory Based Heuristics and Their applications*.

Burrows, E. H.; Wong, W.-K.; Fern, X.; Chaplen, F. W.; and Ely, R. L. 2009. Optimization of ph and nitrogen for enhanced hydrogen production by synechocystis sp. pcc 6803 via statistical and machine learning methods. *Biotechnology Progress* 25(4):1009–1017.

Cohn, D.; ; Ghahramani, Z.; and Jordan, M. 1996. Active learning with statistical models. *Journal of Artificial Intelligence Research* 4:129–145.

Jones, D. 2001. A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization* 21:345–383.

Lizotte, D.; Madani, O.; and Greiner, R. 2003. Budgeted learning of naive-bayes classifiers. In *UAI*.

Madani, O.; Lizotte, D.; and Greiner, R. 2004. Active model selection. In *UAI*.

Moore, A., and Schneider, J. 1995. Memory-based stochastic optimization. In *NIPS*.

Myers, R., and Montgomery, D. 1995. *Response surface methodology: process and product optimization using designed experiments*. Wiley.

Park, D., and Zeikus, J. 2003. Improved fuel cell and electrode designs for producing electricity from microbial degradation. *Biotechnol.Bioeng.* 81(3):348–355.

Rasmussen, C. E., and Williams, C. K. I. 2006. *Gaussian Processes for Machine Learning*. MIT.