

Lipschitz Constants in Experimental Optimization

Gene A. Bunin^{a,*}, Grégory François^b

^a*Ronin Institute for Independent Scholarship*

^b*Institute for Materials and Processes, School of Engineering, The University of Edinburgh,
Edinburgh EH9 3FB, UK*

Abstract

The Lipschitz constant of a response surface function upper bounds the sensitivity of a dependent variable to changes in the independent ones. Traditionally, such constants have found much implicit and abstract use in mathematically oriented applications, but their potential for explicit use in more engineering-based domains has not been explored. The latter point is the subject of this paper, where we propose several ways in which the Lipschitz constants may be used explicitly in the domain of *experimental optimization*. Specifically, we focus on how they may help ensure the satisfaction of constraints and on their potential role in reducing the negative effects of measurement and estimation uncertainty. A number of refinements to the proposed approaches are also derived, and some techniques for setting the constants are presented.

Keywords: Lipschitz constants, upper-bounding functions, experimental optimization, process constraints, real-time optimization

1. Introduction: Experimental Optimization and Lipschitz Constants

Let us consider the function $f_p : \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ to be an *experimental* relationship between the independent-variable vector, $\mathbf{u} \in \mathbb{R}^{n_u}$, and an experimental output quantity, y :

$$y = f_p(\mathbf{u}).$$

Here, the word “experimental” will imply that evaluating the corresponding y for a given \mathbf{u} requires carrying out a *physical* experiment that *cannot be performed solely with a computer*, with the evaluation intrinsic to f_p defined as a physical act and not as a series of numerical computations.

Such relationships appear virtually everywhere in the sciences (Montgomery, 2012), and are typically accompanied by systematic studies that attempt to discover the nature of the given f_p . Among the documented qualitative work

* Author to whom correspondence should be addressed (tel.: +86 13201322405).

Email addresses: gene.bunin@ronininstitute.org (Gene A. Bunin),
gregory.francois@ed.ac.uk (Grégory François)

on the subject, it is probably John Lind’s 1747 examination of scurvy patients (Dunn, 1997) that stands out as the earliest. On the quantitative front, it is likely to Ronald Fisher and his work on the design of experiments (Fisher, 1935) that modern scientific investigation is most indebted, with the foundations for obtaining a mathematical approximation of f_p laid out therein still very much in use today.

In many cases, however, simply identifying an experimental relationship is not the end goal. Instead, one is often searching to manipulate experimental conditions in such a way so as to obtain the “best” response while satisfying a number of safety or physical restrictions. In such scenarios, the *experimental optimization* problem

$$\begin{aligned} & \underset{\mathbf{u}}{\text{minimize}} && \phi_p(\mathbf{u}) \\ & \text{subject to} && g_{p,j}(\mathbf{u}) \leq 0, \quad j = 1, \dots, n_{g_p} \\ & && g_j(\mathbf{u}) \leq 0, \quad j = 1, \dots, n_g \\ & && u_i^L \leq u_i \leq u_i^U, \quad i = 1, \dots, n_u \end{aligned} \tag{1}$$

usually arises, where $\phi : \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ denote the cost function to be minimized, the functions $g : \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ denote the $n_{g_p} + n_g$ constraints, and the constants u_i^L, u_i^U denote the lower and upper limits on \mathbf{u} (now termed “decision variables”), respectively. The subscript p is used to denote explicitly those relationships that are experimental in nature, its absence indicating that the function is numerical and only requires a computer or basic algebra to evaluate. A diverse range of practical engineering problems may be cast and solved in the form of (1) (Bunin, 2016b).

Because the experimental functions ϕ_p and $g_{p,j}$ are unknown and not subject to numerical evaluation, solving (1) entails running a *series* of experiments at $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_k$, where each new experiment is chosen via some algorithmic function, Γ , of the previously applied \mathbf{u} and the corresponding experimentally evaluated function values:

$$\mathbf{u}_{k+1} = \Gamma \left(\begin{array}{c} \mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_k, \phi_p(\mathbf{u}_0), \phi_p(\mathbf{u}_1), \dots, \phi_p(\mathbf{u}_k), \\ g_{p,j}(\mathbf{u}_0), g_{p,j}(\mathbf{u}_1), \dots, g_{p,j}(\mathbf{u}_k), j = 1, \dots, n_{g_p} \end{array} \right), \tag{2}$$

the index k used to denote the latest experimental iteration. The seminal paper with regard to solving (1) is due to Box & Wilson (1951), where Γ essentially takes the form of a standard gradient-based descent method. Numerous other methods (Γ) and frameworks have been proposed in various mathematical and engineering domains since (Bunin et al., 2014b, §2).

In the present work, we explore a single facet of Problem (1) and focus on the case where the experimental functions involved are *Lipschitz continuous* over the explored domain, as we believe that making this assumption may benefit the solution procedure in multiple ways. Given the Lipschitz continuity of a function, we first state the fundamental law (Biegler, 2010, §2.2) that there exists a finite *Lipschitz constant*, κ , such that

$$f_p(\mathbf{u}_b) \leq f_p(\mathbf{u}_a) + \kappa \|\mathbf{u}_b - \mathbf{u}_a\|_2, \quad \forall \mathbf{u}_a, \mathbf{u}_b \in \mathcal{I}, \tag{3}$$

where $\mathcal{I} = \{\mathbf{u} : u_i^L \leq u_i \leq u_i^U, i = 1, \dots, n_u\}$ is the *experimental space* to which the experiments are restricted. As with any assumption, the practical validity of (3) must be questioned, with the authors’ empirical experience so far suggesting that the assumption is not unreasonable for many, if not most, practical problems. In these problems, the constant κ may be understood as a maximal *sensitivity* of the experimental function f_p to the changes in the decision variables \mathbf{u} over the experimental space \mathcal{I} .

Let us suppose now that an explicit value of κ satisfying (3) is *known* or, at the very least, that a reliable overestimate is available. As will be shown in Section 2, such information opens several new doors, thereby making it possible to (a) satisfy important experimental constraints during the optimization process, and (b) reduce the potentially harmful effects of measurement or estimation uncertainty. Of the two, (b) is seen as useful while (a) is considered crucial, as it allows the introduction of *simple* safety guarantees for problems that traditionally lack them. A number of simulated case-study examples will be presented to illustrate the potential benefits of the techniques.

Following the presentation of these core ideas, we will explore, in Section 3, the use of additional assumptions to derive an alternate version of (3) that is, in most cases, tighter and thus less conservative. Specifically, it will be shown that

- using multiple Lipschitz constants for f_p (sensitivities with respect to individual decision variables),
- using local values rather than those valid for *all* \mathbf{u}_a and \mathbf{u}_b in \mathcal{I} , and
- taking into consideration the potential convexity and concavity properties of f_p

all allow for a version of (3) that, though visually more complex, can often lead to significant gains in performance. As may be expected, the price to pay for the improvement is that of additional assumptions. However, they are not compulsory, and the user may employ as many as can reasonably be made.

Finally, Section 4 of the paper will focus on the fundamental and important issue of setting the κ values, which is not so trivial to do well – here, one must remember that since f_p is unknown, so is κ . A number of various methods that combine physical principles, model-based estimation, and data-driven refinement are discussed, with a summary of the authors’ own empirical experience with these methods’ effectiveness concluding the section.

2. The Different Uses of the Lipschitz Constants

Exploiting the Lipschitz constants will, in general, be synonymous with exploiting the corresponding lower and upper *Lipschitz bounds*,

$$f_p(\mathbf{u}_a) - \kappa \|\mathbf{u}_b - \mathbf{u}_a\|_2 \leq f_p(\mathbf{u}_b) \leq f_p(\mathbf{u}_a) + \kappa \|\mathbf{u}_b - \mathbf{u}_a\|_2, \quad \forall \mathbf{u}_a, \mathbf{u}_b \in \mathcal{I}, \quad (4)$$

of which the lower bound is a symmetrical result easily obtained by switching $\mathbf{u}_a \leftrightarrow \mathbf{u}_b$ in (3). Given the experimental nature of f_p , the practical implications of (4) are important as they allow us to analyze the worst-case behavior of the function with respect to its departure from \mathbf{u}_a . In other words, if the experiment at \mathbf{u}_a has been conducted but the one at \mathbf{u}_b has not, it becomes possible to bound the greatest possible change in f_p if \mathbf{u}_b were applied.

This section focuses on the three major uses of (4), outlining their theoretical foundations, discussing their implementation, and providing illustrations via several case-study examples.

2.1. Satisfaction of Experimental Constraints During Optimization

In general, there exists no guarantee that the iterates $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{k+1}$ generated by the algorithm Γ of (2) will satisfy the experimental constraints $g_{p,j}(\mathbf{u}_{\bar{k}}) \leq 0$ for all $\bar{k} = 1, \dots, k + 1$. The reason is simple – the functions are unknown to the user, and this makes it impossible to know in advance if the next applied set of decision variables, \mathbf{u}_{k+1} , will satisfy $g_{p,j}(\mathbf{u}_{k+1}) \leq 0$. While different methodologies for solving (1) do employ different techniques for ensuring that these constraints are met in some sense, the current state of the art is not very developed in this regard, with rigorous constraint satisfaction a notable gap in the experimental-optimization literature. When looking at the three prominent domains concerned with solving Problem (1), we notice the following:

- derivative-free optimization methods, developed mostly by the mathematical community and often dealing with numerical optimization problems, only take precautions against constraints *asymptotically* – i.e., they ensure, usually via penalty-function methods, that \mathbf{u}_{k+1} satisfies the constraints as $k \rightarrow \infty$ (Biegler et al., 2014; Liuzzi et al., 2010) (Conn et al., 2000, §14);
- response-surface optimization methods often deal with experimental systems and are used in a wide number of fields (Myers et al., 2009), but neither do they address the problem of constraint satisfaction during the experiments used to construct the response-surface model, nor do they rigorously attack the potential issue of the model optimum violating the constraints (Michaels & Pengilly, 1963; Umland & Smith, 1959);
- real-time optimization, as typically done in the chemical-engineering community (Chachuat et al., 2009; Darby et al., 2011), is arguably the most concerned with consistent constraint satisfaction, since $g_{p,j}(\mathbf{u}) > 0$ may represent a dangerous or economically disastrous operating regime; however, while both theoretical (Loeblein et al., 1999; Zhang et al., 2002) and practical (Govatsmark & Skogestad, 2005) methods have been developed for promoting the satisfaction of constraints, the former are often too complex and restrictive while the latter suffer from a lack of rigorous guarantees, with both types of methods tending to find solutions

that are significantly suboptimal because of conservatism (Li et al., 2008; Quelhas et al., 2013).

In particular, for those problems where rigorous experimental-constraint satisfaction is highly desired, the above methods are largely inadequate, and only in the real-time optimization community is the rigorous satisfaction of constraints at *every* experiment even a concern.

Some years ago, the authors proposed a Lipschitz approach to this problem (Bunin et al., 2011). Striking what we believe to be a good balance between rigor and simplicity, the method suggests to start with an initially feasible \mathbf{u}_0 , where $g_{p,j}(\mathbf{u}_0) \leq 0$, and to repeatedly couple this with the upper Lipschitz bound of (4) to generate a series of conditions that guarantee constraint satisfaction throughout:

$$\begin{aligned}
g_{p,j}(\mathbf{u}_0) + \kappa_{g_{p,j}} \|\mathbf{u}_1 - \mathbf{u}_0\|_2 \leq 0 &\Rightarrow g_{p,j}(\mathbf{u}_1) \leq 0 \\
g_{p,j}(\mathbf{u}_1) + \kappa_{g_{p,j}} \|\mathbf{u}_2 - \mathbf{u}_1\|_2 \leq 0 &\Rightarrow g_{p,j}(\mathbf{u}_2) \leq 0 \\
&\vdots \\
g_{p,j}(\mathbf{u}_k) + \kappa_{g_{p,j}} \|\mathbf{u}_{k+1} - \mathbf{u}_k\|_2 \leq 0 &\Rightarrow g_{p,j}(\mathbf{u}_{k+1}) \leq 0,
\end{aligned} \tag{5}$$

where $\kappa_{g_{p,j}}$ denotes the corresponding Lipschitz constant for $g_{p,j}$.

The algorithmic implementation of (5) is straightforward, as at the current iterate k the value of $g_{p,j}(\mathbf{u}_k)$ has been obtained, $\kappa_{g_{p,j}}$ is set, \mathbf{u}_k has been applied, and \mathbf{u}_{k+1} is controlled by the user. As such, it suffices to add the constraint

$$g_{p,j}(\mathbf{u}_k) + \kappa_{g_{p,j}} \|\mathbf{u}_{k+1} - \mathbf{u}_k\|_2 \leq 0 \tag{6}$$

to Γ when selecting the next \mathbf{u}_{k+1} to apply. Clearly, because $g_{p,j}(\mathbf{u}_k) \leq 0$, one can always bring \mathbf{u}_{k+1} sufficiently close to \mathbf{u}_k so as to ensure that (6) is satisfied. The simplicity of the implementation is particularly attractive since no complex model of the experimental functions is required, which is the case in some of the rigorous constraint-satisfaction methods previously proposed (Loeblein et al., 1999; Zhang et al., 2002).

The rigor of (5) should also be apparent. If $g_{p,j}(\mathbf{u}_0) \leq 0$ and $\kappa_{g_{p,j}}$ is such that (3) is satisfied, then (5) guarantees that the sequence of $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{k+1}$ meets the constraints at every experiment. Practically speaking, this is a very useful result.

Of course, this approach is not without its caveats, and it would be dangerous to assume otherwise. The most major is the correctness of $\kappa_{g_{p,j}}$, as setting $\kappa_{g_{p,j}}$ too small essentially invalidates the method’s rigor and makes its guarantees unreliable. The obvious solution of making $\kappa_{g_{p,j}}$ “very big” comes with the major drawback of conservatism, and so cannot be recommended as it would implicitly shrink the distances between \mathbf{u}_2 and \mathbf{u}_1 , \mathbf{u}_3 and \mathbf{u}_2 , and so on, resulting in the optimization advancing very slowly. Given the importance of this issue, we will return to it in detail in Section 4, assuming until then that a correct choice of $\kappa_{g_{p,j}}$ is available.

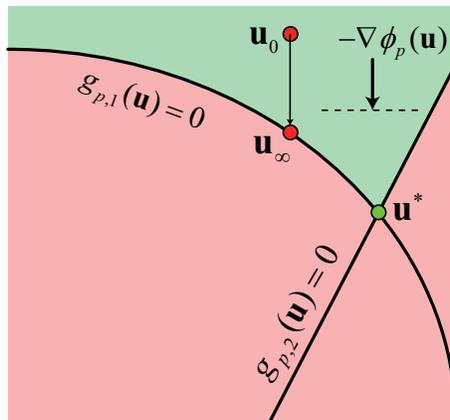


Figure 1: An example of premature convergence for a problem with a linear cost function where the applied Γ continuously takes steps in the gradient descent direction while keeping the constraint satisfied via the Lipschitz condition (6). Gradually, the algorithm converges to the point \mathbf{u}_∞ on the boundary of an experimental constraint ($g_{p,1}$), after which it can no longer make any progress towards the true optimum \mathbf{u}^* , as any step chosen would violate (6). The green and red areas represent the feasible (all constraints satisfied) and infeasible (some constraints violated) sets of the problem, respectively.

A lesser but important concern lies in the potential issue of *premature convergence* that may be experienced by any Γ that adds (6) to its definition. More concretely, this refers to the case where the experiments generated by Γ approach an experimental constraint faster than they approach an optimum, as this drives a given $g_{p,j}$ to 0 and forces $\mathbf{u}_{k+1} \approx \mathbf{u}_k$ for (6) to be met, essentially forcing the algorithm to terminate on some active experimental constraint. We illustrate the issue geometrically in Fig. 1. Although the solution to this particular problem will not be addressed in this paper, we do note that it may be attained with the addition of supplementary conditions that force Γ to not approach any constraint “too quickly”. The interested reader is referred to the unpublished document of Bunin et al. (2014b) for the details.

Finally, the method does rely on the initial experiment satisfying the constraints. However, this is not expected to be a limitation, since it is typically assumed that the experimental optimization of some given system will start at a point that is safe but suboptimal. Indeed, if \mathbf{u}_0 were not safe, then one would not expect optimization to be the user’s top priority.

2.1.1. Case-Study Example: Constraint Adaptation to Minimize the Batch Time of Polystyrene Production

To illustrate how the Lipschitz method may be successfully incorporated into the solution of an experimental optimization problem, let us take the case-study example of minimizing the batch time of a polystyrene production reactor, originally reported by Gentric et al. (1999). Here, we will use the

problem formulation of François et al. (2005), where one is interested in choosing the “switching times” of the reactor’s temperature profile so as to both (a) minimize the time required to reach the desired conversion, favored by higher temperatures, and (b) meet a terminal lower limit on the molecular weight of the product, favored by lower temperatures. The original problem is a dynamic optimization problem as it seeks to find an optimal *profile*, but the piecewise definition of the profile via prescribed “arcs” and switching times essentially turns this into an experimental optimization problem, which can then be solved in a batch-to-batch manner by varying the switching times (and thus the profile) from batch to batch until an optimal profile is hopefully found. As formulated, the problem has an experimental cost function (the time of the batch) and a single experimental constraint (the lower limit on the molecular weight). Because the molecular weight is a product specification, it would be very wasteful to run batches that violate this constraint, as all such batches may be discarded as inadequate for failing to meet specifications. It follows that an effort to satisfy this constraint for all batches should be made.

The exact problem solved is Problem P3 from the ExpOpt database (Bunin, 2016b), which the interested reader may access to find a more detailed problem description, the steps taken to place the problem into the standard form (1), and the actual code used to simulate the case study. Only the problem as written in standard form,

$$\begin{aligned}
& \underset{\mathbf{u}}{\text{minimize}} && \phi_p(\mathbf{u}) := t_b(\mathbf{u}) \\
& \text{subject to} && g_{p,1}(\mathbf{u}) := -M(\mathbf{u}) + 2 \cdot 10^6 \leq 0 \\
& && 50 \leq u_1 \leq 450 \\
& && 600 \leq u_2 \leq 1000,
\end{aligned} \tag{7}$$

is given here, where t_b denotes the final batch time to be minimized (in seconds), M denotes the number average molecular weight (in grams per mole), and u_1 and u_2 are the two switching times that define the temperature profile (in seconds). So as to avoid unnecessary numerical complications, (7) is scaled prior to being solved, with t_b divided by 8000, the constraint function divided by $4 \cdot 10^6$, and the two decision variables normalized to lie in the unit box defined by the coordinates (0, 0) and (1, 1). The initial experiment is chosen as $\mathbf{u}_0 := (242.39, 945.30)$.

To solve the problem, we apply as Γ a constraint-adaptation (Chachuat et al., 2008), or bias-update (Forbes & Marlin, 1994), algorithm. In applying this method, we suppose the availability of an approximate process model,

$$\begin{aligned}
\phi_{\hat{p}}(\mathbf{u}, \theta_1, \theta_2) &\approx \phi_p(\mathbf{u}) \\
g_{\hat{p},1}(\mathbf{u}, \theta_1, \theta_2) &\approx g_{p,1}(\mathbf{u}),
\end{aligned}$$

where θ_1 and θ_2 are the uncertain model parameters – in this case, the rate constants for propagation and transfer to monomer, respectively. Because the values of the model, $\theta_1 = 4.6 \cdot 10^6$ and $\theta_2 = 2.3 \cdot 10^{11}$ (liters per mole per second), differ from the true values of $5.7 \cdot 10^6$ and $1.5 \cdot 10^{11}$, solving the model-based optimization problem

$$\begin{aligned}
& \underset{\mathbf{u}}{\text{minimize}} && \phi_{\hat{p}}(\mathbf{u}, \theta_1, \theta_2) \\
& \text{subject to} && g_{\hat{p},1}(\mathbf{u}, \theta_1, \theta_2) \leq 0 \\
& && 50 \leq u_1 \leq 450 \\
& && 600 \leq u_2 \leq 1000
\end{aligned}$$

directly will not yield the optimal point. In constraint adaptation, one aims instead to iteratively approach a better solution by ensuring that the constraint values in the model-based optimization tend to the true measured values by applying a local bias correction term, ϵ_k , and setting each next \mathbf{u}_{k+1} as

$$\begin{aligned}
\mathbf{u}_{k+1} := \underset{\mathbf{u}}{\text{arg minimize}} && \phi_{\hat{p}}(\mathbf{u}, \theta_1, \theta_2) \\
& \text{subject to} && g_{\hat{p},1}(\mathbf{u}, \theta_1, \theta_2) + \epsilon_k \leq 0 \\
& && 50 \leq u_1 \leq 450 \\
& && 600 \leq u_2 \leq 1000,
\end{aligned} \tag{8}$$

with ϵ_k defined as a weighted sum of the current model error and the previous correction term:

$$\epsilon_k := \alpha [g_{p,1}(\mathbf{u}_k) - g_{\hat{p},1}(\mathbf{u}_k, \theta_1, \theta_2)] + (1 - \alpha)\epsilon_{k-1}. \tag{9}$$

In this example, we initialize ϵ_{-1} as 0, and apply a filter gain of $\alpha := 0.7$. Applying (8) and (9) repeatedly essentially solves the problem, and it can be seen from the top set of results in Fig. 2 that the algorithm succeeds at finding the neighborhood of the solution fairly quickly. However, the path taken results in consistent violations of the molecular weight constraint.

To incorporate the Lipschitz bound, we simply add Condition (6) as a constraint in (8):

$$\begin{aligned}
\mathbf{u}_{k+1} := \underset{\mathbf{u}}{\text{arg minimize}} && \phi_{\hat{p}}(\mathbf{u}, \theta_1, \theta_2) \\
& \text{subject to} && g_{\hat{p},1}(\mathbf{u}, \theta_1, \theta_2) + \epsilon_k \leq 0 \\
& && 50 \leq u_1 \leq 450 \\
& && 600 \leq u_2 \leq 1000 \\
& && g_{p,1}(\mathbf{u}_k) + \kappa_{g_{p,1}} \|\mathbf{u} - \mathbf{u}_k\|_2 \leq 0,
\end{aligned} \tag{10}$$

with a (correct) value¹ of 3 used for $\kappa_{g_{p,1}}$. The performance of the resulting Γ is given as the bottom set of results in Fig. 2, and one observes a trajectory of \mathbf{u} that, though approaching the optimum slower, does so without *any* violations.

As an aside, we note that this way of implementing (6) may be used for any Γ that computes \mathbf{u}_{k+1} by solving a numerical optimization problem. A more

¹The actual issue of setting correct Lipschitz constants is addressed in more detail in Section 4. For the case-study examples in this section, we will, for simplicity, use correct Lipschitz constants, as may be obtained by exhaustive numerical tests on f_p , available to us in these simulated studies.

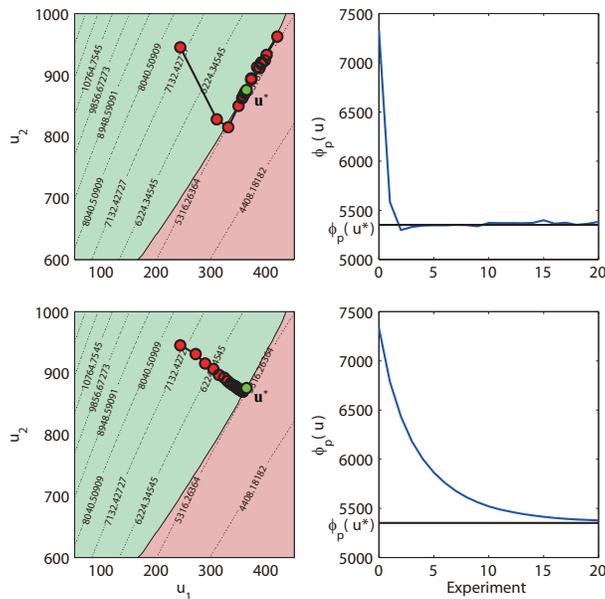


Figure 2: The solution of Problem (7) with the constraint-adaptation algorithms of (8) (top) and (10) (bottom).

general alternative of employing (6) in a line search may be used otherwise (Bunin et al., 2014b).

2.2. Satisfaction of Constraints During Local Perturbations

The previous section focused on modifying the optimization algorithm Γ so that the constraints remain satisfied at every future \mathbf{u}_{k+1} . However, Γ may not be the only law that dictates what experiments are run, as it is often the case that *additional experiments* are periodically required to gather information about the experimental function locally. For example, one may wish to perform experiments in the neighborhood of \mathbf{u}_k so as to identify the parameters of a locally valid model (Conn et al., 2000; Pfaff et al., 2006) or to estimate the local function derivatives (Gao & Engell, 2005; Marchetti et al., 2010), either of which may then be used to assist Γ in computing \mathbf{u}_{k+1} . Naturally, there then arises a conflict between constraint satisfaction and obtaining information when one approaches a constraint, with the information-gathering perturbations becoming a potential constraint-violation hazard.

A viable solution to this conflict has already been discussed in detail by Bunin (2016a), and what we present here is a simple variant. Namely, to guarantee that any point in the perturbation ball of $\mathcal{B}_e = \{\mathbf{u} : \|\mathbf{u} - \mathbf{u}_k\|_2 \leq \delta_e\}$ satisfy a given constraint $g_{p,j}(\mathbf{u}) \leq 0$, it is sufficient to enforce that the value $g_{p,j}(\mathbf{u}_k)$ satisfy a constraint back-off of $\delta_e \kappa_{g_{p,j}}$:

$$g_{p,j}(\mathbf{u}_k) + \delta_e \kappa_{g_{p,j}} \leq 0 \Rightarrow g_{p,j}(\mathbf{u}) \leq 0, \quad \forall \mathbf{u} \in \mathcal{B}_e \cap \mathcal{I}. \quad (11)$$

The validity of (11) is easily demonstrated. For every $\mathbf{u} \in \mathcal{B}_e$,

$$\|\mathbf{u} - \mathbf{u}_k\|_2 \leq \delta_e \Rightarrow g_{p,j}(\mathbf{u}_k) + \kappa_{g_{p,j}} \|\mathbf{u} - \mathbf{u}_k\|_2 \leq g_{p,j}(\mathbf{u}_k) + \delta_e \kappa_{g_{p,j}},$$

and thus

$$g_{p,j}(\mathbf{u}_k) + \delta_e \kappa_{g_{p,j}} \leq 0 \Rightarrow g_{p,j}(\mathbf{u}_k) + \kappa_{g_{p,j}} \|\mathbf{u} - \mathbf{u}_k\|_2 \leq 0 \Rightarrow g_{p,j}(\mathbf{u}) \leq 0.$$

Note that an analogous result may be derived for the numerical constraints g_j , with the back-off of $\delta_e \kappa_{g,j}$ resulting in the same guarantee if we assume g_j to be Lipschitz continuous with a Lipschitz constant of $\kappa_{g,j}$. However, more involved numerical techniques are also possible in this case (Bunin, 2016a).

The exact manner in which one may employ (11) to allow for perturbation while robustly satisfying the constraints would almost certainly be method-dependent. The general principles, however, are quite simple. Because one has the guarantee that satisfying the constraint with a specified back-off allows for one to perturb safely anywhere in a δ_e -ball around that point, it follows that one can choose δ_e *a priori* and to then only perturb around those experimental iterates that satisfy the corresponding back-off, since all such information-gathering experiments will satisfy the constraints. This is now illustrated for a specific case-study example.

2.2.1. Case Study: A Modifier-Adaptation Algorithm for Minimizing the Steady-State Production Cost of a Gold Cyanidation Leaching Process

One algorithm that depends heavily on perturbations is the modifier-adaptation method (Brdys & Tatjewski, 2005; Gao & Engell, 2005; Marchetti et al., 2009), which solves Problem (1) by iteratively solving a corrected model-based problem:

$$\begin{aligned} & \mathbf{u}_{k+1} := \\ & \arg \underset{\mathbf{u}}{\text{minimize}} \quad \phi_{\hat{p}}(\mathbf{u}, \boldsymbol{\theta}) + \boldsymbol{\lambda}_{\phi,k}^T \mathbf{u} \\ & \text{subject to} \quad g_{\hat{p},j}(\mathbf{u}, \boldsymbol{\theta}) + \epsilon_{j,k} + \boldsymbol{\lambda}_{j,k}^T (\mathbf{u} - \mathbf{u}_k) \leq 0, \quad j = 1, \dots, n_{g_p} \\ & \quad \quad \quad g_j(\mathbf{u}) \leq 0, \quad j = 1, \dots, n_g \\ & \quad \quad \quad u_i^L \leq u_i \leq u_i^U, \quad i = 1, \dots, n_u, \end{aligned} \quad (12)$$

with ϵ and $\boldsymbol{\lambda}$ denoting the “modifiers” used to correct the local zero- and first-order errors of the model, respectively:

$$\begin{aligned} \epsilon_{j,k} &:= \alpha [g_{p,j}(\mathbf{u}_k) - g_{\hat{p},j}(\mathbf{u}_k, \boldsymbol{\theta})] + (1 - \alpha) \epsilon_{j,k-1} \\ \boldsymbol{\lambda}_{j,k} &:= \alpha [\nabla g_{p,j}(\mathbf{u}_k) - \nabla g_{\hat{p},j}(\mathbf{u}_k, \boldsymbol{\theta})] + (1 - \alpha) \boldsymbol{\lambda}_{j,k-1} \\ \boldsymbol{\lambda}_{\phi,k} &:= \alpha [\nabla \phi_p(\mathbf{u}_k) - \nabla \phi_{\hat{p}}(\mathbf{u}_k, \boldsymbol{\theta})] + (1 - \alpha) \boldsymbol{\lambda}_{\phi,k-1}. \end{aligned} \quad (13)$$

It should be clear that this Γ is a generalization of the constraint-adaptation algorithm employed in Section 2.1.1, in that we are now correcting the model *derivatives* in addition to the constraint function values. As with constraint adaptation, the uncertain parameters of the model are not updated between iterations and are kept at some nominal, constant values, although a version of the algorithm that estimates and updates the parameters as well may certainly be used (Brdys & Tatjewski, 2005).

As in the constraint-adaptation case, the zero-order modifier values $\epsilon_{j,k}$ may be easily obtained by physically measuring the experimental function and numerically evaluating the model at \mathbf{u}_k . The first-order corrections, however, require experimental derivative estimates, which may be obtained using a number of approaches (Bunin et al., 2013a). Here, we will employ the simple method of taking the difference quotients

$$\left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{u}_k} \approx \frac{f_p(\mathbf{u}_k) - f_p(\mathbf{u}_k - \delta_e \mathbf{e}_i)}{\delta_e}, \quad (14a)$$

or, if $u_{k,i} - \delta_e < u_i^L$,

$$\left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{u}_k} \approx \frac{f_p(\mathbf{u}_k + \delta_e \mathbf{e}_i) - f_p(\mathbf{u}_k)}{\delta_e}, \quad (14b)$$

where \mathbf{e}_i is the unit vector with unity at the i^{th} spot and δ_e is some non-zero value. The choice to use δ_e as the perturbation step is, of course, not coincidental and is intended to be used in synergy with (11) so as to ensure that all perturbations performed for estimating the derivatives satisfy the constraints.

The outline of the employed algorithm, where one starts at the iteration $k := 0$ with some \mathbf{u}_0 that satisfies the constraints on $g_{p,j}$ and g_j with the additional back-offs of $\delta_e \kappa_{g_{p,j}}$ and $\delta_e \kappa_{g_j}$, is as follows:

1. Define the modifiers as in (13), with the experimental gradients obtained by carrying out an additional n_u experiments and then applying (14).
2. Compute \mathbf{u}_{k+1} by solving an augmented version of (12) that accounts for the back-offs:

$$\begin{aligned} & \mathbf{u}_{k+1} := \\ \arg \underset{\mathbf{u}}{\text{minimize}} & \quad \phi_{\hat{p}}(\mathbf{u}, \boldsymbol{\theta}) + \boldsymbol{\lambda}_{\phi,k}^T \mathbf{u} \\ \text{subject to} & \quad g_{\hat{p},j}(\mathbf{u}, \boldsymbol{\theta}) + \epsilon_{j,k} + \boldsymbol{\lambda}_{j,k}^T (\mathbf{u} - \mathbf{u}_k) \leq 0, & j = 1, \dots, n_{g_p} \\ & \quad g_j(\mathbf{u}) + \delta_e \kappa_{g_j} \leq 0, & j = 1, \dots, n_g \\ & \quad u_i^L \leq u_i \leq u_i^U, & i = 1, \dots, n_u \\ & \quad g_{p,j}(\mathbf{u}_k) + \kappa_{g_{p,j}} \|\mathbf{u} - \mathbf{u}_k\|_2 + \delta_e \kappa_{g_{p,j}} \leq 0, & j = 1, \dots, n_{g_p}, \end{aligned} \quad (15)$$

3. Apply the new decision variables to the experimental system, obtain the new measurements, augment the value of k , and return to Step 1.

Note that the modifications introduced in (15) serve to ensure that the constraint function values – experimental and numerical – meet the back-offs required to ensure that the perturbations in the ball \mathcal{B}_e around the future \mathbf{u}_{k+1} satisfy the constraints.

This algorithm is applied to Problem P11 from the ExpOpt database (Bunin, 2016b). The problem itself is adapted from the work of Jun et al. (2015) and deals with the minimization of the steady-state production cost of a gold cyanidation leaching process. Only the skeleton of the problem is provided here, with interested readers once more referred to the aforementioned references for more in-depth descriptions.

In our standard form, the problem may be written as

$$\begin{aligned} & \underset{\mathbf{u}}{\text{minimize}} && \phi_p(\mathbf{u}) := \phi_{\text{econ}}(\mathbf{u}) \\ & \text{subject to} && g_{p,1}(\mathbf{u}) := \frac{C_s(\mathbf{u}) - C_{s,0}}{C_{s,0}} + 0.75 \leq 0 \\ & && 10 \leq u_1 \leq 80 \\ & && 5 \leq u_2 \leq 20, \end{aligned} \tag{16}$$

with \mathbf{u} consisting of two variables – the flow rate of the sodium cyanide (in kilograms per hour) and the concentration of the dissolved oxygen in the liquid (in milligrams per kilogram). The function ϕ_{econ} is the economical cost function to be minimized (in Chinese RMB per hour), while C_s denotes the gold concentration in the ore (in milligrams per kilograms; $C_{s,0}$ denotes its initial value). The single experimental constraint ensures that the gold recovery is at least 75% – the solution, although not on this constraint, is nevertheless very close to it. The model used for the case study is that provided in the database and differs from the simulated “real” process by virtue of errors in the kinetic parameters. As in the previous case-study example, the decision variables are scaled down to a unit box, with the cost and constraint values divided by the scaling factors of 200 and 0.04, respectively. A δ_e value of 0.05 (for the scaled variables) is chosen, with the Lipschitz constant of the scaled problem set as $\kappa_{g_{p,1}} := 4$. The process starts operating at the suboptimal point of $\mathbf{u}_0 := (52, 18)$, which satisfies the experimental constraint with a slack greater than $\delta_e \kappa_{g_{p,1}}$. It is chosen not to filter the modifiers, with $\alpha := 1$ in (13).

The results of applying the above algorithm to this problem are given in Fig. 3. As expected, one sees that the main experimental iterates all satisfy the experimental constraint with the back-off, while the δ_e -perturbations for derivative estimation all satisfy the original constraint. In this case, we see that the modifier-adaptation algorithm converges to the neighborhood of the optimum fairly quickly.

2.3. Reducing the Effect of Uncertainty in the Experimental Function Values

The Lipschitz constants may play an important supplementary role in the very realistic scenario where the experimental function values $f_p(\mathbf{u}_{\bar{k}})$, $\bar{k} = 0, \dots, k$ are subject to uncertainty, which may be due to (a) measurement noise

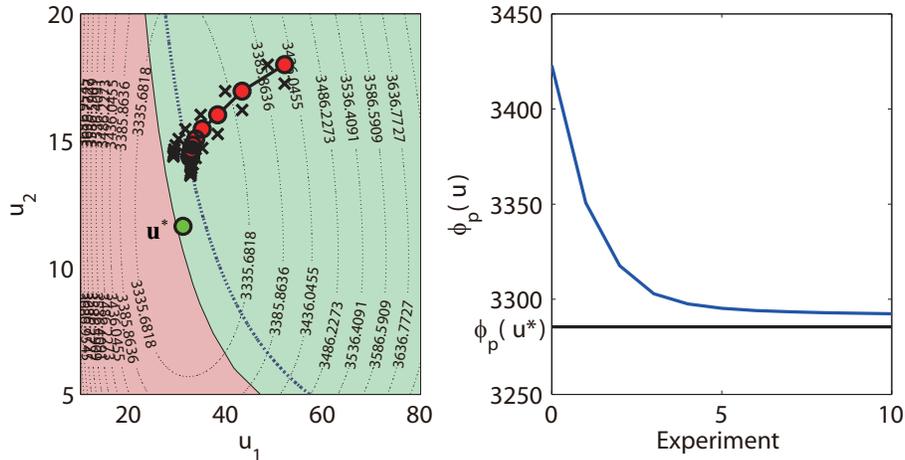


Figure 3: Results obtained from applying the modifier-adaptation algorithm to Problem (16). The dashed line in the plot on the left shows the $\delta_e \kappa_{g_p,1}$ back-off sufficient to guarantee that the additional perturbations for derivative estimation satisfy the constraint. The cross marks show the perturbations, the data for which are omitted from the plot on the right.

or (b) the experimental quantity not being directly measured but estimated. In such cases, it is typical to work with a statistical interval in which the value $f_p(\mathbf{u}_{\bar{k}})$ is expected to lie. To obtain such an interval, one needs to make certain assumptions on the statistical nature of the uncertainty.

A common uncertainty model, and the one that will be considered here, is the additive

$$\hat{f}_p^{\bar{k}} = f_p(\mathbf{u}_{\bar{k}}) + w_{\bar{k}}, \quad (17)$$

where \hat{f}_p is the observed or estimated value and w is the stochastic noise or error. By bounding the stochastic element in the probability sense²:

$$\underline{w}_{\bar{k}} \stackrel{P}{\leq} w_{\bar{k}} \stackrel{P}{\leq} \bar{w}_{\bar{k}},$$

it then becomes possible to compute a high-probability interval for $f_p(\mathbf{u}_{\bar{k}})$ by applying these bounds to a rearranged version of (17):

$$\underline{f}_p^{\bar{k}} := \hat{f}_p^{\bar{k}} - \bar{w}_{\bar{k}} \stackrel{P}{\leq} f_p(\mathbf{u}_{\bar{k}}) \stackrel{P}{\leq} \hat{f}_p^{\bar{k}} - \underline{w}_{\bar{k}} := \bar{f}_p^{\bar{k}}. \quad (18)$$

Most importantly, these bounds allow us to obtain versions of (6) and (11) that are robust and do not require the unavailable exact function values:

²As a compromise between clarity and convenience, we will use the symbol $\stackrel{P}{\leq}$ to mean “less than or equal to *with sufficiently high probability*”.

$$\bar{g}_{p,j}^k + \kappa_{g_{p,j}} \|\mathbf{u}_{k+1} - \mathbf{u}_k\|_2 \leq 0 \Rightarrow g_{p,j}(\mathbf{u}_{k+1}) \stackrel{P}{\leq} 0, \quad (19)$$

$$\bar{g}_{p,j}^k + \delta_\epsilon \kappa_{g_{p,j}} \leq 0 \Rightarrow g_{p,j}(\mathbf{u}) \stackrel{P}{\leq} 0, \quad \forall \mathbf{u} \in \mathcal{B}_\epsilon \cap \mathcal{I}, \quad (20)$$

thereby making it possible to maintain the applicability of the methods in Sections 2.1 and 2.2 for the realistic case with uncertainty.

However, while these bounds may be fairly rigorous and useful, they can also prove too conservative when the variance of w is large. A standard procedure for tightening the interval for a given $\mathbf{u}_{\bar{k}}$ is to resample the function value at $\mathbf{u}_{\bar{k}}$ multiple times and to take the mean of the values. Under reasonable assumptions on w , this then yields an estimator whose variance shrinks to 0 with an increasing number of samples (Kay, 1993, §1.3). Alternatively, one may be able to obtain tighter bounding values by assuming a locally valid structure for f_p , such as linear or quadratic, and to fit this structure to the data via maximum-likelihood estimation. Provided that the structure chosen is correct, this too can be shown to yield estimators of f_p values that go to the true value with a variance that decreases with an increasing amount of data (Kay, 1993, §7).

The alternative proposed here is to “trim off” some of the uncertainty when the resulting interval is actually larger than what would be allowed by the worst-case (Lipschitz-bound) changes in the function values from one decision-variable set to another. Consider first the robust version of (4) for some $\bar{k} \in \{0, \dots, k\}$:

$$\underline{f}_p^{\bar{k}} - \kappa \|\mathbf{u}_{\bar{k}} - \mathbf{u}_{\bar{k}}\|_2 \stackrel{P}{\leq} f_p(\mathbf{u}_{\bar{k}}) \stackrel{P}{\leq} \bar{f}_p^{\bar{k}} + \kappa \|\mathbf{u}_{\bar{k}} - \mathbf{u}_{\bar{k}}\|_2,$$

of which (18) are the particular case corresponding to $\bar{k} \rightarrow \bar{k}$. We may, however, be able to improve the bounds by considering the tightest values over all $\{0, \dots, k\}$:

$$\underline{f}_p^{\bar{k}} := \max_{\bar{k}=0, \dots, k} \left(\underline{f}_p^{\bar{k}} - \kappa \|\mathbf{u}_{\bar{k}} - \mathbf{u}_{\bar{k}}\|_2 \right) \stackrel{P}{\leq} f_p(\mathbf{u}_{\bar{k}}), \quad (21)$$

$$f_p(\mathbf{u}_{\bar{k}}) \stackrel{P}{\leq} \min_{\bar{k}=0, \dots, k} \left(\bar{f}_p^{\bar{k}} + \kappa \|\mathbf{u}_{\bar{k}} - \mathbf{u}_{\bar{k}}\|_2 \right) := \bar{f}_p^{\bar{k}}. \quad (22)$$

After using (21)-(22) to assign lower/upper bounding values to the measurements at $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_k$, one may run through the process again, since changes in some of the values may lead to additional changes in the others. The iterative scheme may be outlined as follows:

1. Set the nominal $\underline{f}_p^{\bar{k}}$ and $\bar{f}_p^{\bar{k}}$ for $\bar{k} = 0, \dots, k$ as given in (18).
2. Refine these values by applying (21)-(22) for $\bar{k} = 0, \dots, k$.

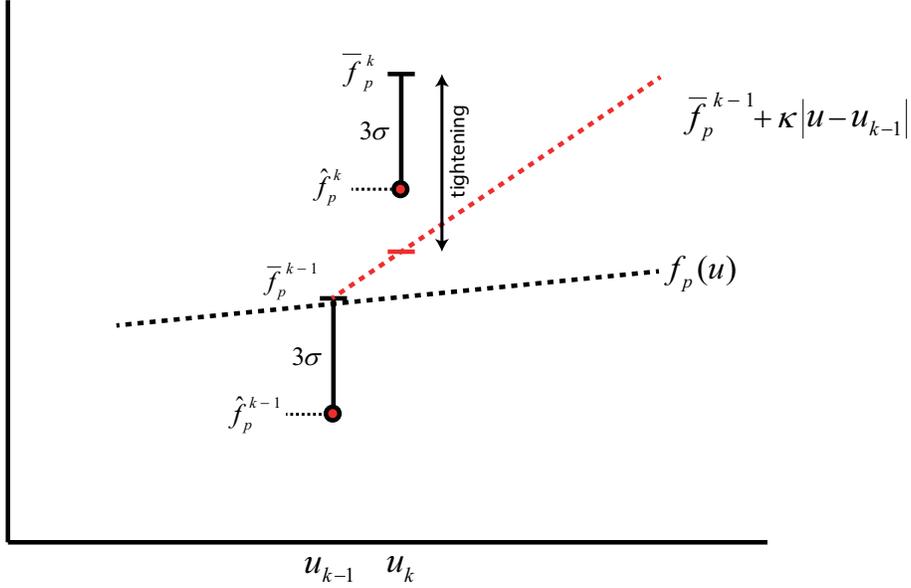


Figure 4: Illustration of how Lipschitz bounds may be used to tighten uncertainty bounds on the true function values. The red dashed line shows the Lipschitz bound.

3. If the largest refinement obtained is negligible (e.g., $< 10^{-6}$), terminate. Otherwise, return to Step 2.

A conceptual illustration of reducing the uncertainty via this procedure is given in Fig. 4, where we consider a simple one-dimensional case with white Gaussian noise of variance σ^2 , and high-probability bounds of -3σ and 3σ on the noise elements. As shown in the figure, it happens that a very poor nominal upper bounding value is obtained at u_k after a very tight upper bounding value has been obtained at the previous experiment of u_{k-1} . Because of the Lipschitz bound, the tightness at the latter may then be partially “inherited” by the former, thereby leading to significant refinement in the upper bounding value at u_k . Naturally, we may expect this method to lose its effectiveness when the decision-variable points are further apart and when the Lipschitz constants used are more conservative.

2.3.1. Case Study: Gold Cyanidation Leaching Process with Noise

Let us again consider the case-study example of Section 2.2.1, this time supposing a more realistic scenario where all of the measurements obtained, for both the cost and constraint, are corrupted by noise elements from the PDF of $\mathcal{N}(0, \sigma^2)$. A value of $\sigma := 0.07$ is used here, and the additive corruption is applied to the problem post-scaling. The nominal lower and upper bounding values on the noise elements are then chosen as $\underline{w} := -3\sigma$ and $\bar{w} := 3\sigma$ for

all measurements at all iterations, from which the nominal lower and upper bounding values on the true function values are obtained as in (18).

Two realizations of the modifier-adaptation algorithm are compared here, with both still setting the modifiers as in (13)-(14), but with the f_p function values replaced by their measured or estimated \hat{f}_p analogues. To robustify the constraint-satisfaction guarantees against the uncertainty effects, the numerical optimization problem of (15) is modified to

$$\begin{aligned}
& \mathbf{u}_{k+1} := \\
& \underset{\mathbf{u}}{\text{arg minimize}} && \phi_{\hat{p}}(\mathbf{u}, \boldsymbol{\theta}) + \boldsymbol{\lambda}_{\phi,k}^T \mathbf{u} \\
& \text{subject to} && g_{\hat{p},j}(\mathbf{u}, \boldsymbol{\theta}) + \epsilon_{j,k} + \boldsymbol{\lambda}_{j,k}^T (\mathbf{u} - \mathbf{u}_k) \leq 0, \quad j = 1, \dots, n_{g_p} \\
& && g_j(\mathbf{u}) + \delta_e \kappa_{g,j} \leq 0, \quad j = 1, \dots, n_g \\
& && u_i^L \leq u_i \leq u_i^U, \quad i = 1, \dots, n_u \\
& && \bar{g}_{p,j}^k + \kappa_{g_{p,j}} \|\mathbf{u} - \mathbf{u}_k\|_2 + \delta_e \kappa_{g_{p,j}} \leq 0, \quad j = 1, \dots, n_{g_p}.
\end{aligned} \tag{23}$$

In the first algorithm, the Lipschitz constants are *not* used to refine the bounding values, with the $\bar{g}_{p,j}^k$ value obtained from (18) used in (23). In the second algorithm, the Lipschitz-bound refinement procedure as outlined in the preceding section is run through all of the measurements – the “key” iterations $0, \dots, k$ and the additional perturbations around the key iterations used for derivative estimation. The resulting $\bar{g}_{p,j}^k$ is then used in (23). Additionally, the measured values $\hat{\phi}_p$ and $\hat{g}_{p,j}$ that are used in estimating the modifiers and derivatives are trimmed in the case that they do not satisfy the tightened bounding values (e.g., $\hat{\phi}_p^k \rightarrow \bar{\phi}_p^k$ if $\hat{\phi}_p^k > \bar{\phi}_p^k$). The Lipschitz constant for the cost is chosen as $\kappa_{\phi} := 5$.

So as to see the average effect that including the Lipschitz-based tightening has on the performance of the optimization algorithm, we run the algorithms in parallel for 1000 noise realizations and evaluate the average difference between the corresponding cost function values,

$$\Delta\phi_{\text{ave}} := \frac{1}{k_f + 1} \sum_{\bar{k}=0}^{k_f} (\phi_{p,\bar{k}}^A - \phi_{p,\bar{k}}^B),$$

for each realization, with ϕ_p^A and ϕ_p^B denoting the cost function values obtained by the original (without Lipschitz tightening) and modified (with Lipschitz tightening) algorithms, respectively. The scalar k_f is used to denote the number of main experiments run, excluding the perturbations and the initial experiment, and is set as 30 here. A positive $\Delta\phi_{\text{ave}}$ value thus suggests an improvement over the original algorithm.

Plotting the results of the thousand trials in a histogram (Fig. 5), we notice a very positive trend, with noticeably better results obtained by the second algorithm.

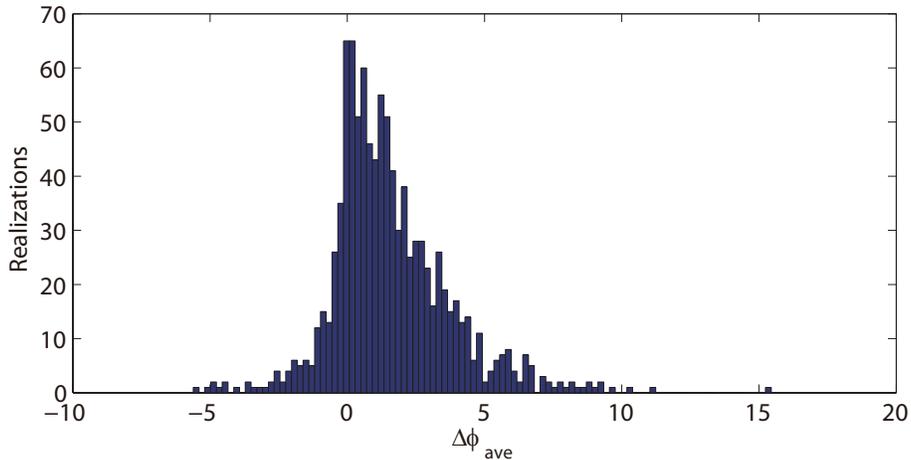


Figure 5: Distribution of the average differences between the cost function values obtained by the two modifier-adaptation algorithms.

3. Improving the Lipschitz Bounds

A significant price to pay for the use of the Lipschitz bounds of (4) is their potential conservatism. Namely, if the constant κ is chosen so that the bounds are satisfied with a very large margin, one is faced with several drawbacks:

- as already mentioned in Section 2.1, the constraint-satisfaction condition (6) may become too restrictive, forcing the algorithm to take overly small steps;
- the back-off, $\delta_e \kappa_{p,j}$, needed for safe perturbation in (11) may become too large and difficult to satisfy, and if satisfied may introduce a large degree of suboptimality;
- the Lipschitz-based tightening of the bounding values f_p and \bar{f}_p obtained in Section 2.3 becomes negligible or disappears altogether.

Fortunately, though the function f_p may be experimental and thus unknown, scattered bits of information about its behavior might be available and could be exploited to yield different Lipschitz bounds that are often much tighter.

In this section, we consider the following types of additional information:

- directional lower and upper Lipschitz constants,

$$\underline{\kappa}_i \leq \left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{u}} \leq \bar{\kappa}_i, \quad \forall \mathbf{u} \in \mathcal{I};$$

- local Lipschitz constants over the subspace

$$\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b} = \{\mathbf{u} : u_{a,i} \leq u_i \leq u_{b,i}, i = 1, \dots, n_u\} \subseteq \mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b^U} \subseteq \mathcal{I},$$

defined as

$$\underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \leq \left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{u}} \leq \bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}, \forall \mathbf{u} \in \mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}; \quad (24)$$

- the local convexity and concavity properties of f_p , where it is known that f_p is locally convex or concave in certain variables over $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$, with the indices of these variables denoted by the sets $\mathcal{I}_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}$ and $\mathcal{I}_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}$, respectively;
- lower and upper bounds on the local derivatives of f_p :

$$\underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \leq \nabla_{\underline{f}_{p,i}}^{\mathbf{u}_a} \leq \left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{u}_a} \leq \nabla_{\bar{f}_{p,i}}^{\mathbf{u}_a} \leq \bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}. \quad (25)$$

By integrating the information above, one is able to account for scenarios that would be completely ignored when using (4). For example, by using $2n_u$ Lipschitz constants instead of just one, it becomes possible to differentiate between those decision variables that have a large effect on f_p and those whose effect is known to be negligible. Additionally, by having both lower ($\underline{\kappa}_i$) and upper ($\bar{\kappa}_i$) constants, one can differentiate between the sensitivities as being either positive or negative – e.g., if it is known that increasing a certain u_i will *always* increase the value of f_p , the corresponding $\underline{\kappa}_i$ can be set as 0, which, as will be seen, may tighten the corresponding Lipschitz bound significantly.

By limiting the analysis to a subspace $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b} \subseteq \mathcal{I}$, it becomes possible to accommodate those problems where the sensitivity of f_p to some u_i is known to change drastically over \mathcal{I} . Of particular concern may be those f_p that have an inversely proportional relationship with u_i , and whose values may explode as u_i approaches 0. While it might make sense to include the 0-point in the definition of \mathcal{I} , actually operating there might be unlikely, and so it would be unnecessarily conservative to use the very large Lipschitz constants corresponding to the region around the 0-point when smaller, more reasonable values could be used instead. Considering the local subspace $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$ provides us with a convenient mathematical formulation for handling such cases.

Finally, the inclusion of convexity/concavity properties, when coupled with derivative bounds, gives greater flexibility for those problems where f_p is known with certainty to behave in a convex/concave manner when only certain variables are changed (the other variables being kept the same). The relationship between the power and current in a fuel-cell system (Marchetti et al., 2011), known to be concave, is one example that the authors have encountered. Mathematically, exploiting such relationships allows us to forgo using Lipschitz constants for certain directions and to use the less conservative derivative bounds

instead. Some results for obtaining such bounds for an experimental function f_p have been reported in Bunin et al. (2013a).

Combining the information above, let us now derive an alternative to (4).

Theorem 1 (Alternate Lipschitz Bounds). *Let the experimental function f_p be continuously differentiable (\mathcal{C}^1) over an open set containing $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$, with its derivatives bounded as in (24)-(25). The bounds*

$$f_p(\mathbf{u}_a) + \sum_{i \in I_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}} \min \left[\frac{\nabla f_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i})}{\nabla \bar{f}_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i})} \right] + \sum_{i \notin I_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}} \min \left[\frac{\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}(u_{b,i} - u_{a,i})}{\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}(u_{b,i} - u_{a,i})} \right] \leq f_p(\mathbf{u}_b), \quad (26)$$

$$f_p(\mathbf{u}_b) \leq f_p(\mathbf{u}_a) + \sum_{i \in I_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}} \max \left[\frac{\nabla f_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i})}{\nabla \bar{f}_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i})} \right] + \sum_{i \notin I_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}} \max \left[\frac{\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}(u_{b,i} - u_{a,i})}{\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}(u_{b,i} - u_{a,i})} \right] \quad (27)$$

are then satisfied for all $\mathbf{u}_a, \mathbf{u}_b \in \mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$.

PROOF. We will only prove (27) and leave (26) to the reader, as the two proofs are symmetric. First, let us write the evolution of f_p from \mathbf{u}_a to \mathbf{u}_b as occurring in “two steps”, the first step in the variables in which f_p is concave and the second step in the variables in which f_p is not. For clarity, we will partition \mathbf{u} as $\mathbf{u} = (\mathbf{v}, \mathbf{z})$, where \mathbf{v} are the variables $u_i, i \in I_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}$ and \mathbf{z} are the variables $u_i, i \notin I_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}$. The two-step evolution may be written as

$$f_p(\mathbf{v}_b, \mathbf{z}_b) = f_p(\mathbf{v}_a, \mathbf{z}_a) + [f_p(\mathbf{v}_b, \mathbf{z}_a) - f_p(\mathbf{v}_a, \mathbf{z}_a)] + [f_p(\mathbf{v}_b, \mathbf{z}_b) - f_p(\mathbf{v}_b, \mathbf{z}_a)].$$

The proof consists in upper-bounding the two differences in the brackets. For the first difference, we call upon the first-order condition of concavity (Boyd & Vandenberghe, 2008, §3):

$$f_p(\mathbf{v}_b, \mathbf{z}_a) - f_p(\mathbf{v}_a, \mathbf{z}_a) \leq \sum_{i \in I_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}} \frac{\partial f_p}{\partial u_i} \Big|_{\mathbf{u}_a} (u_{b,i} - u_{a,i}).$$

To further bound the right-hand side, we then consider two cases,

$$\begin{aligned} u_{b,i} - u_{a,i} \geq 0 &\Leftrightarrow \frac{\partial f_p}{\partial u_i} \Big|_{\mathbf{u}_a} (u_{b,i} - u_{a,i}) \leq \nabla \bar{f}_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i}) \\ u_{b,i} - u_{a,i} \leq 0 &\Leftrightarrow \frac{\partial f_p}{\partial u_i} \Big|_{\mathbf{u}_a} (u_{b,i} - u_{a,i}) \leq \nabla f_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i}), \end{aligned} \quad (28)$$

and account for the worst case via the maximum operator to obtain

$$f_p(\mathbf{v}_b, \mathbf{z}_a) - f_p(\mathbf{v}_a, \mathbf{z}_a) \leq \sum_{i \in I_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}} \max \left[\frac{\nabla f_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i})}{\nabla \bar{f}_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i})} \right]. \quad (29)$$

To place a bound on $f_p(\mathbf{v}_b, \mathbf{z}_b) - f_p(\mathbf{v}_b, \mathbf{z}_a)$, we limit our analysis to a single dimension and then apply the mean-value theorem. First, consider the function f_p over the line connecting $(\mathbf{v}_b, \mathbf{z}_a)$ and $(\mathbf{v}_b, \mathbf{z}_b)$, parameterized as

$$\tilde{f}_p(\mathbf{v}_b, \gamma) = f_p(\mathbf{v}_b, \mathbf{z}(\gamma)),$$

where

$$\mathbf{z}(\gamma) = \mathbf{z}_a + \gamma(\mathbf{z}_b - \mathbf{z}_a), \quad \gamma \in [0, 1].$$

Because f_p is \mathcal{C}^1 , \tilde{f}_p is as well, thereby allowing us to apply the Taylor series expansion between $\gamma = 0$ and $\gamma = 1$, together with the mean-value theorem (Korn & Korn, 2000, §4.7), to state that

$$\exists \tilde{\gamma} \in (0, 1) : \tilde{f}_p(\mathbf{v}_b, 1) = \tilde{f}_p(\mathbf{v}_b, 0) + \left. \frac{\partial \tilde{f}_p}{\partial \gamma} \right|_{\mathbf{v}_b, \tilde{\gamma}}.$$

The derivative term may then be defined in terms of the original function f_p by applying the chain rule:

$$\left. \frac{\partial \tilde{f}_p}{\partial \gamma} \right|_{\mathbf{v}_b, \tilde{\gamma}} = \sum_{i \notin I_{ccv}^{\mathbf{u}_a, \mathbf{u}_b}} \left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{v}_b, \mathbf{z}(\tilde{\gamma})} \left. \frac{du_i}{d\gamma} \right|_{\tilde{\gamma}} = \sum_{i \notin I_{ccv}^{\mathbf{u}_a, \mathbf{u}_b}} \left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{v}_b, \mathbf{z}(\tilde{\gamma})} (u_{b,i} - u_{a,i}).$$

Noting that $\tilde{f}_p(\mathbf{v}_b, 0) = f_p(\mathbf{v}_b, \mathbf{z}_a)$ and $\tilde{f}_p(\mathbf{v}_b, 1) = f_p(\mathbf{v}_b, \mathbf{z}_b)$, we thus obtain that

$$\exists \tilde{\gamma} \in (0, 1) : f_p(\mathbf{v}_b, \mathbf{z}_b) = f_p(\mathbf{v}_b, \mathbf{z}_a) + \sum_{i \notin I_{ccv}^{\mathbf{u}_a, \mathbf{u}_b}} \left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{v}_b, \mathbf{z}(\tilde{\gamma})} (u_{b,i} - u_{a,i}).$$

Because $(\mathbf{v}_b, \mathbf{z}(\tilde{\gamma})) \in \mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$, the constants of (24) may be used to bound the derivatives and, using the same logic as in (28), we go on to obtain

$$\left. \frac{\partial f_p}{\partial u_i} \right|_{\mathbf{v}_b, \mathbf{z}(\tilde{\gamma})} (u_{b,i} - u_{a,i}) \leq \max \left[\frac{\underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}}{\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}} (u_{b,i} - u_{a,i}), \right],$$

which then leads to

$$f_p(\mathbf{v}_b, \mathbf{z}_b) - f_p(\mathbf{v}_b, \mathbf{z}_a) \leq \sum_{i \notin I_{ccv}^{\mathbf{u}_a, \mathbf{u}_b}} \max \left[\frac{\underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}}{\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}} (u_{b,i} - u_{a,i}), \right]. \quad (30)$$

Combining (29) and (30) yields the desired result. \square

Next, we consider the analogue to (11) that follows from (27).

Corollary 1 (Alternate Back-off for Local Perturbations). Let $g_{p,j}$ be \mathcal{C}^1 over an open set containing $\mathcal{I}_{\mathbf{u}_k - \delta_e}^{\mathbf{u}_k + \delta_e}$, where $\delta_e \in \mathbb{R}^{n_u}$ is the vector of δ_e values, and let the derivatives be bounded as in (24)-(25) with $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b} \rightarrow \mathcal{I}_{\mathbf{u}_k - \delta_e}^{\mathbf{u}_k + \delta_e}$ and $f_p \rightarrow g_{p,j}$. Defining the mixed vector of greatest absolute Lipschitz constants and derivative bounds as

$$\tilde{\kappa}_{p,ji} = \begin{cases} \max \left(\left| \nabla \underline{g}_{p,ji}^{\mathbf{u}_k} \right|, \left| \nabla \bar{g}_{p,ji}^{\mathbf{u}_k} \right| \right), & i \in I_{\text{ccv},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}, \\ \max \left(\left| \underline{\kappa}_{p,ji}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e} \right|, \left| \bar{\kappa}_{p,ji}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e} \right| \right), & i \notin I_{\text{ccv},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}, \end{cases}$$

it follows that

$$g_{p,j}(\mathbf{u}_k) + \delta_e \|\tilde{\kappa}_{p,j}\|_2 \leq 0 \Rightarrow g_{p,j}(\mathbf{u}) \leq 0, \quad \forall \mathbf{u} \in \mathcal{B}_e. \quad (31)$$

PROOF. Applying (27), it follows that

$$\begin{aligned} g_{p,j}(\mathbf{u}) \leq g_{p,j}(\mathbf{u}_k) + & \sum_{i \in I_{\text{ccv},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \max \left[\begin{array}{c} \nabla \underline{g}_{p,ji}^{\mathbf{u}_k}(u_i - u_{k,i}), \\ \nabla \bar{g}_{p,ji}^{\mathbf{u}_k}(u_i - u_{k,i}) \end{array} \right] \\ & + \sum_{i \notin I_{\text{ccv},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \max \left[\begin{array}{c} \underline{\kappa}_{p,ji}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i}), \\ \bar{\kappa}_{p,ji}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i}) \end{array} \right] \end{aligned} \quad (32)$$

holds for all $\mathbf{u} \in \mathcal{B}_e \subseteq \mathcal{I}_{\mathbf{u}_k - \delta_e}^{\mathbf{u}_k + \delta_e}$. Let us use this bound to define a *Lipschitz polytope* centered at \mathbf{u}_k :

$$\mathcal{L}_j = \left\{ \mathbf{u} : \begin{array}{l} g_{p,j}(\mathbf{u}_k) + \sum_{i \in I_{\text{ccv},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \max \left[\begin{array}{c} \nabla \underline{g}_{p,ji}^{\mathbf{u}_k}(u_i - u_{k,i}), \\ \nabla \bar{g}_{p,ji}^{\mathbf{u}_k}(u_i - u_{k,i}) \end{array} \right] \\ + \sum_{i \notin I_{\text{ccv},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \max \left[\begin{array}{c} \underline{\kappa}_{p,ji}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i}), \\ \bar{\kappa}_{p,ji}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i}) \end{array} \right] \leq 0 \end{array} \right\}.$$

The remainder of the proof is essentially a “busier” version of the proof of Theorem 1 in Bunin (2016a), with the same concepts at work. From (32), it should be clear that any \mathbf{u} belonging to both \mathcal{B}_e and \mathcal{L}_j will satisfy $g_{p,j}(\mathbf{u}) \leq 0$, and so we proceed to identify the back-off on $g_{p,j}(\mathbf{u}_k)$ that would enforce $\mathcal{B}_e \subseteq \mathcal{L}_j$, and thus $\mathbf{u} \in \mathcal{B}_e \Rightarrow \mathbf{u} \in \mathcal{L}_j$ (i.e., by inscribing \mathcal{B}_e in \mathcal{L}_j). This is done analytically by building progressively smaller subsets of \mathcal{L}_j , all of which include \mathcal{B}_e .

From the algebraic relation $\max(\underline{xy}, \bar{xy}) \leq \max(|\underline{x}|, |\bar{x}|)|y|$, we have

$$\max \left[\begin{array}{c} \nabla \underline{g}_{p,ji}^{\mathbf{u}_k}(u_i - u_{k,i}), \\ \nabla \bar{g}_{p,ji}^{\mathbf{u}_k}(u_i - u_{k,i}) \end{array} \right] \leq \max \left(\left| \nabla \underline{g}_{p,ji}^{\mathbf{u}_k} \right|, \left| \nabla \bar{g}_{p,ji}^{\mathbf{u}_k} \right| \right) |u_i - u_{k,i}|,$$

$$\max \left[\frac{\underline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i})}{\overline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i})} \right] \leq \max \left(\left| \frac{\underline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}}{\overline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \right|, \left| \frac{\overline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}}{\underline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \right| \right) |u_i - u_{k,i}|,$$

whence follows

$$\begin{aligned} \sum_{i \in I_{\text{cev},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \max \left[\begin{array}{c} \nabla \underline{g}_{p,j}^{\mathbf{u}_k}(u_i - u_{k,i}), \\ \nabla \overline{g}_{p,j}^{\mathbf{u}_k}(u_i - u_{k,i}) \end{array} \right] &+ \sum_{i \notin I_{\text{cev},j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}} \max \left[\begin{array}{c} \underline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i}), \\ \overline{\kappa}_{p,j}^{\mathbf{u}_k - \delta_e, \mathbf{u}_k + \delta_e}(u_i - u_{k,i}) \end{array} \right] \\ &\leq \sum_{i=1}^{n_u} \tilde{\kappa}_{p,j} |u_i - u_{k,i}| = \tilde{\kappa}_{p,j}^T \Delta \mathbf{u}, \end{aligned} \quad (33)$$

with $\Delta \mathbf{u} = (|u_1 - u_{k,1}|, |u_2 - u_{k,2}|, \dots, |u_{n_u} - u_{k,n_u}|)$. Defining the polytope

$$\tilde{\mathcal{L}}_j = \{ \mathbf{u} : g_{p,j}(\mathbf{u}_k) + \tilde{\kappa}_{p,j}^T \Delta \mathbf{u} \leq 0 \},$$

we see that $\tilde{\mathcal{L}}_j \subseteq \mathcal{L}_j$ by virtue of (33).

Consider now the Cauchy-Schwarz inequality, which states that

$$\tilde{\kappa}_{p,j}^T \Delta \mathbf{u} \leq \|\tilde{\kappa}_{p,j}\|_2 \|\Delta \mathbf{u}\|_2 = \|\tilde{\kappa}_{p,j}\|_2 \|\mathbf{u} - \mathbf{u}_k\|_2, \quad (34)$$

and define the next (ball) subset as

$$\hat{\mathcal{L}}_j = \{ \mathbf{u} : g_{p,j}(\mathbf{u}_k) + \|\tilde{\kappa}_{p,j}\|_2 \|\mathbf{u} - \mathbf{u}_k\|_2 \leq 0 \}.$$

That $\hat{\mathcal{L}}_j \subseteq \tilde{\mathcal{L}}_j$ follows from (34). It now remains to show that $\mathcal{B}_e \subseteq \hat{\mathcal{L}}_j$, as this implies $\mathcal{B}_e \subseteq \mathcal{L}_j$ and completes the proof. For every $\mathbf{u} \in \mathcal{B}_e$, we have, by definition, $\|\mathbf{u} - \mathbf{u}_k\|_2 \leq \delta_e$, and so it follows that any $\mathbf{u} \in \mathcal{B}_e$ will also be in $\hat{\mathcal{L}}_j$ if

$$g_{p,j}(\mathbf{u}_k) + \delta_e \|\tilde{\kappa}_{p,j}\|_2 \leq 0$$

holds, which is precisely what is ensured by (31). \square

Employing the above theorem and corollary, we may now state the alternate results of Sections 2.1-2.3:

Constraint Satisfaction During Optimization (Section 2.1)

$$\begin{aligned} g_{p,j}(\mathbf{u}_k) &+ \sum_{i \in I_{\text{cev},j}^{\mathbf{u}_k, \mathbf{u}_{k+1}}} \max \left[\begin{array}{c} \nabla \underline{g}_{p,j}^{\mathbf{u}_k}(u_{k+1,i} - u_{k,i}), \\ \nabla \overline{g}_{p,j}^{\mathbf{u}_k}(u_{k+1,i} - u_{k,i}) \end{array} \right] \\ &+ \sum_{i \notin I_{\text{cev},j}^{\mathbf{u}_k, \mathbf{u}_{k+1}}} \max \left[\begin{array}{c} \underline{\kappa}_{p,j}^{\mathbf{u}_k, \mathbf{u}_{k+1}}(u_{k+1,i} - u_{k,i}), \\ \overline{\kappa}_{p,j}^{\mathbf{u}_k, \mathbf{u}_{k+1}}(u_{k+1,i} - u_{k,i}) \end{array} \right] \leq 0 \Rightarrow g_{p,j}(\mathbf{u}_{k+1}) \leq 0. \end{aligned} \quad (35)$$

Constraint Satisfaction During Perturbation (Section 2.2)

$$g_{p,j}(\mathbf{u}_k) + \delta_e \|\tilde{\boldsymbol{\kappa}}_{p,j}\|_2 \leq 0 \Rightarrow g_{p,j}(\mathbf{u}) \leq 0, \quad \forall \mathbf{u} \in \mathcal{B}_e. \quad (36)$$

Bounding Values in the Presence of Uncertainty (Section 2.3)

$$\underline{f}_{-p}^{\bar{k}} := \max_{\bar{k}=0,\dots,k} \left(\begin{aligned} &\underline{f}_{-p}^{\bar{k}} + \sum_{i \in I_{\text{cvx}}^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}} \min \begin{bmatrix} \nabla \underline{f}_{-p,i}^{\bar{k}}(u_{\bar{k},i} - u_{\bar{k},i}), \\ \nabla \bar{f}_{-p,i}^{\bar{k}}(u_{\bar{k},i} - u_{\bar{k},i}) \end{bmatrix} \\ &+ \sum_{i \notin I_{\text{cvx}}^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}} \min \begin{bmatrix} \underline{\kappa}_i^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}(u_{\bar{k},i} - u_{\bar{k},i}), \\ \bar{\kappa}_i^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}(u_{\bar{k},i} - u_{\bar{k},i}) \end{bmatrix} \end{aligned} \right) \stackrel{P}{\leq} f_p(\mathbf{u}_{\bar{k}}), \quad (37)$$

$$f_p(\mathbf{u}_{\bar{k}}) \stackrel{P}{\leq} \min_{\bar{k}=0,\dots,k} \left(\begin{aligned} &\bar{f}_p^{\bar{k}} + \sum_{i \in I_{\text{ccv}}^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}} \max \begin{bmatrix} \nabla \bar{f}_{-p,i}^{\bar{k}}(u_{\bar{k},i} - u_{\bar{k},i}), \\ \nabla \underline{f}_{-p,i}^{\bar{k}}(u_{\bar{k},i} - u_{\bar{k},i}) \end{bmatrix} \\ &+ \sum_{i \notin I_{\text{ccv}}^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}} \max \begin{bmatrix} \underline{\kappa}_i^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}(u_{\bar{k},i} - u_{\bar{k},i}), \\ \bar{\kappa}_i^{\mathbf{u}_{\bar{k}}, \mathbf{u}_{\bar{k}}}(u_{\bar{k},i} - u_{\bar{k},i}) \end{bmatrix} \end{aligned} \right) := \bar{f}_p^{\bar{k}}. \quad (38)$$

Finally, it is worthwhile to show how the simpler bound of (4) may be obtained from the alternate (26)-(27). In doing so, we also present a sufficient condition for which (26)-(27) provide tighter bounds over $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$.

Corollary 2 (Relation Between (4) and (26)-(27)). *Let $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b} \subseteq \mathcal{I}$ and let f_p be \mathcal{C}^1 over an open set containing $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$, with its derivatives bounded as in (24)-(25). Defining the mixed vectors of greatest absolute Lipschitz constants and derivative bounds as*

$$\tilde{\boldsymbol{\kappa}}_{\bar{f},i} = \begin{cases} \max \left(\left| \nabla \underline{f}_{-p,i}^{\mathbf{u}_a} \right|, \left| \nabla \bar{f}_{-p,i}^{\mathbf{u}_a} \right| \right), & i \in I_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}, \\ \max \left(\left| \underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \right|, \left| \bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \right| \right), & i \notin I_{\text{ccv}}^{\mathbf{u}_a, \mathbf{u}_b}, \end{cases}$$

$$\tilde{\boldsymbol{\kappa}}_{\underline{f},i} = \begin{cases} \max \left(\left| \nabla \underline{f}_{-p,i}^{\mathbf{u}_a} \right|, \left| \nabla \bar{f}_{-p,i}^{\mathbf{u}_a} \right| \right), & i \in I_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}, \\ \max \left(\left| \underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \right|, \left| \bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \right| \right), & i \notin I_{\text{cvx}}^{\mathbf{u}_a, \mathbf{u}_b}, \end{cases}$$

it follows that $\|\tilde{\boldsymbol{\kappa}}_{\bar{f}}\|_2 \leq \kappa$ implies that (27) is tighter than the upper bound of (4). Likewise, $\|\tilde{\boldsymbol{\kappa}}_{\underline{f}}\|_2 \leq \kappa$ implies that (26) is tighter than the lower bound of (4).

PROOF. For the upper bound, we borrow the analysis already carried out in the proof of Corollary 1 to state that the bound

$$f_p(\mathbf{u}_b) \leq f_p(\mathbf{u}_a) + \|\tilde{\boldsymbol{\kappa}}_{\bar{f}}\|_2 \|\mathbf{u}_b - \mathbf{u}_a\|_2$$

is a looser, more conservative bound than (27). Clearly, $\|\tilde{\kappa}_{\bar{f}}\|_2 \leq \kappa$ then implies that the upper bound of (4) is even more conservative, whence the proof of the first half of the result.

For the lower bound, we note that

$$\min \left[\begin{array}{c} \nabla \underline{f}_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i}), \\ \nabla \bar{f}_{p,i}^{\mathbf{u}_a}(u_{b,i} - u_{a,i}) \end{array} \right] \geq -\max \left(\left| \nabla \underline{f}_{p,i}^{\mathbf{u}_a} \right|, \left| \nabla \bar{f}_{p,i}^{\mathbf{u}_a} \right| \right) |u_{b,i} - u_{a,i}|,$$

$$\min \left[\begin{array}{c} \underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}(u_{b,i} - u_{a,i}), \\ \bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b}(u_{b,i} - u_{a,i}) \end{array} \right] \geq -\max \left(\left| \underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \right|, \left| \bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} \right| \right) |u_{b,i} - u_{a,i}|,$$

which allows us to apply the same analysis as before to obtain

$$f_p(\mathbf{u}_b) \geq f_p(\mathbf{u}_a) - \|\tilde{\kappa}_f\|_2 \|\mathbf{u}_b - \mathbf{u}_a\|_2$$

as a looser, more conservative alternative to (26). By the same reasoning, $\|\tilde{\kappa}_f\|_2 \leq \kappa$ then implies that the lower bound of (4) is even more conservative, thus implying that (26) is tighter. \square

Clearly, we can obtain a valid version of (4) from (26)-(27) by setting $\kappa := \max(\|\tilde{\kappa}_f\|_2, \|\tilde{\kappa}_{\bar{f}}\|_2)$.

4. The Setting and Estimation of Lipschitz Constants

Setting the Lipschitz constants in implementation so that they satisfy (3) or (24) is usually not trivial and requires some care. As already mentioned, choosing constants that do not satisfy these inequalities potentially makes the aforementioned techniques invalid because of the possibility that the corresponding Lipschitz bounds no longer hold. At the same time, as discussed at the start of Section 3, satisfying the inequalities with too much conservatism may affect performance undesirably.

It is thus the goal of this section to offer some insight into how one may set and refine the estimates of the Lipschitz constants intelligently. We begin with the special and desirable case where these constants are easily known from the physical laws governing a given experimental system (Section 4.1). Unfortunately, as there exists no means to rigorously guarantee that the Lipschitz constants chosen for a *general experimental function* be correct, one must inevitably turn to heuristic methods in the general case, and some of these are discussed in Sections 4.2-4.3. Despite lacking the desired rigor, the methods proposed are not doomed to failure for this reason, and can certainly be applied with success, as has been demonstrated with some empirical evidence from both simulated and experimental case studies. This last point is discussed in Section 4.4.

4.1. Setting Lipschitz Constants by Exploiting Physical Laws

Virtually all experimental relationships are subject to physical laws, and for some relationships such laws are so well known and documented that they may be assumed to hold. A simple example previously employed by the authors is the exchange of heat between a heating and heated element – e.g., a jacket (heating element) surrounding a reactor (heated element). If the temperature of the heated element is an experimental function, and the temperature of the heating element is one of the decision variables under the user’s control, then much can be said about the sensitivity relating the changes in the temperatures of the two. For very many systems, it would be expected that raising the temperature of the heating element will raise the temperature of the heated element, and that, symmetrically, lowering the temperature of one will lead to a lower temperature in the other. As such, the sensitivity would always be positive, and the lower Lipschitz constant, $\underline{\kappa}_i$, could safely be set to 0. Additionally, heat losses are likely to ensure that a change in the temperature of the heating element will lead to a *smaller* change in the temperature in the heated one, thus allowing us to set $\bar{\kappa}_i$ as 1.

While this particular example is somewhat trivial, there are others that are less so. For example, in the solid oxide fuel-cell system where one manipulates the current while needing to respect a lower-limit constraint on the cell voltage (Marchetti et al., 2011), it is known that the relationship between voltage and current is always inverse proportional, thus allowing one to set the corresponding $\bar{\kappa}_i$ as 0. And although the physical laws could not tell us the value of $\underline{\kappa}_i$ in this case, one could nevertheless rely on the multitude of available experimental data to come up with a good estimate. For the polymerization example of Section 2.1.1, one may know that extended heating (increased u_1), together with shortened cooling (decreased u_2) and thus increased temperature, will generally lead to a favorization of the termination reactions and shorter polymer chains, thus resulting in a final product with a lower average molecular weight. This knowledge may then be used to set the Lipschitz constants for the molecular-weight constraint function accordingly.

Identifying such relationships and their relevance to the setting of Lipschitz constants may be an important step to include in the experimental and theoretical work that normally goes into understanding a given system prior to optimization. Note, however, that the physical laws are harder to link to the lumped κ in (3), which is more abstract than the individual $\underline{\kappa}_i, \bar{\kappa}_i$ of (24).

4.2. Model-Based Estimation

Many of the relationships that appear in experimental optimization problems will have undergone a fair amount of theoretical investigation, resulting in first-principles parametric models,

$$f_p(\mathbf{u}) \approx f_{\hat{p}}(\mathbf{u}, \boldsymbol{\theta}),$$

being available for them, where the uncertain parameters $\boldsymbol{\theta}$ are assumed to belong to a bounded set, Θ .

If the assumption of parametric uncertainty is not too egregious, and if the uncertainty set Θ is not too erroneously defined, then a reasonable choice of Lipschitz constants in (24) may be obtained by minimizing and maximizing the derivatives of the model over $\mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b}$ and Θ :

$$\underline{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} := \arg \underset{\substack{\mathbf{u} \in \mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b} \\ \boldsymbol{\theta} \in \Theta}}{\text{minimize}} \left. \frac{\partial f_{\hat{p}}}{\partial u_i} \right|_{\mathbf{u}, \boldsymbol{\theta}}, \quad (39)$$

$$\bar{\kappa}_i^{\mathbf{u}_a, \mathbf{u}_b} := \arg \underset{\substack{\mathbf{u} \in \mathcal{I}_{\mathbf{u}_a}^{\mathbf{u}_b} \\ \boldsymbol{\theta} \in \Theta}}{\text{maximize}} \left. \frac{\partial f_{\hat{p}}}{\partial u_i} \right|_{\mathbf{u}, \boldsymbol{\theta}}, \quad (40)$$

with the analogous formulation

$$\kappa := \arg \underset{\substack{\mathbf{u}_a, \mathbf{u}_b \in \mathcal{I} \\ \mathbf{u}_a \neq \mathbf{u}_b \\ \boldsymbol{\theta} \in \Theta}}{\text{maximize}} \frac{f_{\hat{p}}(\mathbf{u}_b, \boldsymbol{\theta}) - f_{\hat{p}}(\mathbf{u}_a, \boldsymbol{\theta})}{\|\mathbf{u}_b - \mathbf{u}_a\|_2} \quad (41)$$

for the lumped κ . While solving these optimization problems may not be computationally easy, it is conceivable that they would be solved off-line prior to any Γ being applied, thus making the associated computational burdens of lesser concern. If an additional layer of safety is needed, one can heuristically lower and increase the estimates as required.

For the very general case where no model is available, one could still apply the above approach but in a data-driven fashion. One could, for example, construct a linear or quadratic model from several obtained measurements, choose $\boldsymbol{\theta}$ as the coefficients of the model, let their confidence intervals define Θ (Montgomery, 2012, §10.5), and then apply either (39)-(40) or (41).

4.3. Refinement via Data-Driven Consistency Checks

The previous two sections have addressed the problem of setting Lipschitz constants when none are initially provided. Now, let us consider a technique for refining these initial estimates so that they are *consistent with the obtained data*.

Using (3) as an example, suppose now that we have obtained an estimate of a Lipschitz constant, denoted by $\hat{\kappa}$, and as such have the hypothetically valid bound

$$f_p(\mathbf{u}_b) \leq f_p(\mathbf{u}_a) + \hat{\kappa} \|\mathbf{u}_b - \mathbf{u}_a\|_2, \quad \forall \mathbf{u}_a, \mathbf{u}_b \in \mathcal{I}. \quad (42)$$

Although there exists no way to prove the validity of (42) for an experimental function f_p in the general case, we can and should confirm that this bound is at least satisfied for those experimental measurements that have been collected.

Letting k_1 and k_2 denote two (different) indices, a basic consistency-check algorithm would consist in the following two steps being carried out for every combination (k_1, k_2) of $k_1, k_2 \in [0, \dots, k]$:

(i) Check if the inequality

$$f_p(\mathbf{u}_{k_2}) \leq f_p(\mathbf{u}_{k_1}) + \hat{\kappa} \|\mathbf{u}_{k_2} - \mathbf{u}_{k_1}\|_2 \quad (43)$$

is satisfied.

(ii) If (43) is satisfied, proceed to the next (k_1, k_2) combination unless no combinations are left, in which case terminate. Otherwise, increase the value of $\hat{\kappa}$ by a preset, strictly positive quantity and return to (i).

It is easy to show that such an algorithm will terminate with a $\hat{\kappa}$ value that is consistent with the obtained data. For the more complex case of (24), one could apply the same concept by lowering $\underline{\kappa}_i$ and raising $\bar{\kappa}_i$ until (26)-(27) hold. While there is no apparent “best way” to increase the $\underline{\kappa}_i, \bar{\kappa}_i$ values when inconsistency occurs, one could employ heuristics, such as doubling the values. More elegant schemes may also be proposed (Bunin et al., 2014a, §3.5.3).

4.4. Empirical Evidence for Lipschitz-Constant Estimation

Clearly, since Lipschitz constants are often unknown, it is difficult to test, in any sort of closed-loop manner, whether different estimation methods are actually successful at identifying proper values for these constants in practice. Instead, one may apply the methods during the solution procedure of (1) and see if the problem is solved in a satisfactory manner. When the problem comes with experimental constraints, the strongest criterion is to verify that the constraints were met during optimization – if they were not, then this is strong evidence for poor estimates, as it suggests that the constraint-satisfaction condition (6) was not valid for the constraints violated. The speed of decrease in the cost function value is another telling criterion – if progress is very slow, then this may be a sign of the estimates being too conservative. In other cases, the quality of the estimates may be difficult to evaluate, although different “strange” behaviors during optimization may be an indicator that the constants are not being estimated properly. To date, much of the available empirical evidence for the effectiveness of Lipschitz-constant estimation methods comes from problems that have been tackled by the SCFO solver (Bunin, 2015), as this solver incorporates almost all of the techniques discussed in this paper and relies heavily on the Lipschitz constants to operate.

Of the experimental results available, perhaps the most telling are those obtained in the experimental optimization of a laboratory solid-oxide fuel cell stack (François & Bonvin, 2014), where no theoretical model was used and the Lipschitz constants were initialized using a data-driven model and the methods of Section 4.2, before being refined by the consistency checks of Section 4.3 as more data became available. From the limited results obtained, one does see that the estimation is sufficiently good from the safety perspective – the constants are conservative enough so as to keep the system from violating its two experimental constraints, and a close-to-optimal system efficiency is achieved

without requiring too many iterations, which suggests that the constants are not *too* conservative.

Other experimental results have been obtained by applying an earlier version of the SCFO to the problem of run-to-run controller tuning, where the tuning parameters of the controllers were treated as the decision variables, with different trajectory tracking metrics defined to judge controller performance for a given setpoint profile (Bunin et al., 2013b). The solver managed to successfully autotune a model-predictive controller for a water-tank system and a fixed-order controller for a mechanical torsional plant – however, these results are less validating since no experimental constraints were present and so it was impossible to evaluate the quality of the Lipschitz-constant estimates with regard to their ability to enforce constraints. At the same time, the constants for the experimental cost function *were* used to help reduce the effects of variance in the cost-function values (Section 2.3). While this is likely to have aided in the optimization, there is no way to say how much, as testing the quality of the Lipschitz-constant estimators was not the goal of this largely proof-of-concept study. Once again, no *a priori* models were used – the solver constructed a data-driven model from $n_u + 1$ initial experiments, and then refined the initial estimates via the consistency checks.

In simulated tests, a good quantity of results is now available for a number of case-study test problems, including six with experimental constraints, courtesy of the ExpOpt database (Bunin, 2016b). Examining the results of the SCFO solver for these problems shows that the current estimation methods employed by the solver are certainly not perfect, and one can indeed have Lipschitz-constant estimates that are not valid and that allow for constraint violations, although this varies significantly from problem to problem. For certain problems, the estimates are consistently good and violations are minimal, with the violations being of very small magnitude even when they are present. For other problems, violations may be larger and more frequent, but it is nevertheless rare to see consistently large constraint violations, suggesting that consistency checks are eventually able to improve on poor initial estimates.

With regard to results obtained for methods other than the SCFO, an application of a data-driven approach of Section 4.2 has recently been reported by Bunin (2016a). Here, a modified version of G. E. P. Box’s evolutionary-operation method (Box & Draper, 1969) was proposed, with the Lipschitz bound used to ensure that the experiments carried out by the algorithm would never violate the problem constraints. After each set of $2n_u$ axially distributed experiments, the Lipschitz constants were estimated locally as the derivatives of a linear model for the $2n_u + 1$ axial-plus-center data points, with additional conservatism added to the estimates to account for noise and nonlinearity effects. The end result was very satisfactory – in testing the algorithm for three problems from the ExpOpt database, the estimated Lipschitz constants were always sufficiently conservative and the experimental constraints were not violated once.

5. Concluding Remarks

It is the authors' hope that the present document has helped convince the reader that Lipschitz constants – despite being an implicit, age-old mathematical concept – may be made explicit and thereby bring multiple benefits to the experimental optimization domain. Most importantly, the use of these constants appears to offer an extremely simple way to guarantee constraint satisfaction in problems with experimental constraints, which is something that may be exploited during both the optimization and information-gathering (perturbation) phases of the solution process. Because constraint satisfaction may be immensely important in experimental settings, this alone presents a very strong argument for the implementation of Lipschitz constants and bounds. Additionally, it has been shown that these constants may play a useful role in reducing the effect of measurement or estimation uncertainty, which, depending on the optimization algorithm, may lead to significant performance improvements. There are no major implementation difficulties involved with the techniques described, and so it is the authors' recommendation that these techniques be given consideration when approaching experimental optimization problems.

This paper has also addressed the two major concerns regarding the use of Lipschitz constants in practice – the concern that the set constants may be too conservative and the concern that they may be improperly set or estimated. Both are valid and represent potential issues. However, as has been shown in Section 3, there do exist several ways to reduce conservatism – namely, one may tighten the bounds by exploiting direction, locality, and the potential convexity/concavity properties of the function. With regard to estimation, Section 4 has outlined a number of methods to set and refine the Lipschitz-constant estimates, and empirical evidence has shown that such methods can, in fact, be sufficient for many problems.

Acknowledgements

The authors would like to thank Professor Dominique Bonvin of the Laboratoire d'Automatique (École Polytechnique Fédérale de Lausanne) for his insights and input.

References

- Biegler, L. T. (2010). *Nonlinear Programming: Concepts, Algorithms, and Applications to Chemical Processes*. Society for Industrial and Applied Mathematics and the Mathematical Optimization Society.
- Biegler, L. T., Lang, Y., & Lin, W. (2014). Multi-scale optimization for process systems engineering. *Comput. Chem. Eng.*, *60*, 17–30.
- Box, G., & Draper, N. (1969). *Evolutionary Operation: A Statistical Method for Process Improvement*. John Wiley & Sons.

- Box, G., & Wilson, K. (1951). On the experimental attainment of optimum conditions. *J. Royal Stat. Society. Series B (Methodological)*, *13*, 1–45.
- Boyd, S., & Vandenberghe, L. (2008). *Convex Optimization*. Cambridge University Press.
- Brdys, M., & Tatjewski, P. (2005). *Iterative Algorithms for Multilayer Optimizing Control*. Imperial College Press.
- Bunin, G. A. (2015). *The SCFO Experimental Optimization Solver: Users' Guide (version 0.91.2)*. [Http://www.ccapprox.info/#soft](http://www.ccapprox.info/#soft).
- Bunin, G. A. (2016a). Constraint back-offs for safe, sufficient excitation: A general theory with application to experimental optimization. *Comput. Chem. Eng.*, *93*, 353–360.
- Bunin, G. A. (2016b). Test problem database for experimental optimization (ver. 1.3). [Http://www.ccapprox.info/expopt/](http://www.ccapprox.info/expopt/). Accessed February 2016.
- Bunin, G. A., François, G., & Bonvin, D. (2013a). From discrete measurements to bounded gradient estimates: A look at some regularizing structures. *Ind. Eng. Chem. Res.*, *52*, 12500–12513.
- Bunin, G. A., François, G., & Bonvin, D. (2013b). A real-time optimization framework for the iterative controller tuning problem. *Processes*, *1*, 203–237.
- Bunin, G. A., François, G., & Bonvin, D. (2014a). Implementation techniques for the SCFO experimental optimization framework. *arXiv:1406.3997 [math.OC]*, (pp. 1–66).
- Bunin, G. A., François, G., & Bonvin, D. (2014b). Feasible-side global convergence in experimental optimization. *arXiv:1406.4063 [math.OC]*, (pp. 1–35).
- Bunin, G. A., François, G., Srinivasan, B., & Bonvin, D. (2011). Input filter design for feasibility in constraint-adaptation schemes. In *18th World Congress of the International Federation of Automatic Control (IFAC) (Milan)* (pp. 5585–5590).
- Chachuat, B., Marchetti, A., & Bonvin, D. (2008). Process optimization via constraints adaptation. *J. Process Control*, *18*, 244–257.
- Chachuat, B., Srinivasan, B., & Bonvin, D. (2009). Adaptation strategies for real-time optimization. *Comput. Chem. Eng.*, *33*, 1557–1567.
- Conn, A. R., Gould, N. I. M., & Toint, P. L. (2000). *Trust-Region Methods*. SIAM.
- Darby, M. L., Nikolaou, M., Jones, J., & Nicholson, D. (2011). RTO: An overview and assessment of current practice. *J. Process Control*, *21*, 874–884.

- Dunn, P. M. (1997). James Lind (1716-94) of Edinburgh and the treatment of scurvy. *Arch. Dis. Child. Fetal Neonatal Ed.*, *76*, F64–F65.
- Fisher, R. A. (1935). *The Design of Experiments*. Oliver and Boyd.
- Forbes, J. F., & Marlin, T. E. (1994). Model accuracy for economic optimizing controllers: The bias update case. *Ind. Eng. Chem. Res.*, *33*, 1919–1929.
- François, G., & Bonvin, D. (2014). Real-time optimization: optimizing the operation of energy systems in the presence of uncertainty and disturbances. In *13th International Conference on Sustainable Energy Technologies (Geneva)* (p. E40137).
- François, G., Srinivasan, B., & Bonvin, D. (2005). Use of measurements for enforcing the necessary conditions of optimality in the presence of constraints and uncertainty. *J. Process Control*, *15*, 701–712.
- Gao, W., & Engell, S. (2005). Iterative set-point optimization of batch chromatography. *Comput. Chem. Eng.*, *29*, 1401–1409.
- Gentric, C., Pla, F., Latifi, M. A., & Corriou, J. P. (1999). Optimization and non-linear control of a batch emulsion polymerization reactor. *Chem. Eng. J.*, *75*, 31–46.
- Govatsmark, M., & Skogestad, S. (2005). Selection of controlled variables and robust setpoints. *Ind. Eng. Chem. Res.*, *44*, 2207–2217.
- Jun, Z., Zhi-zhong, M., & Run-da, J. (2015). Real-time optimization based on SCFO for gold cyanidation leaching process. *Chem. Engin. Sci.*, *134*, 467–476.
- Kay, S. M. (1993). *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall.
- Korn, G. A., & Korn, T. M. (2000). *Mathematical Handbook for Scientists and Engineers*. Dover Publications.
- Li, P., Arellano-Garcia, H., & Wozny, G. (2008). Chance constrained programming approach to process optimization under uncertainty. *Comput. Chem. Eng.*, *32*, 25–45.
- Liuzzi, G., Lucidi, S., & Sciandrone, M. (2010). Sequential penalty derivative-free methods for nonlinear constrained optimization. *SIAM J. Optim.*, *20*, 2614–2635.
- Loeblein, C., Perkins, J. D., Srinivasan, B., & Bonvin, D. (1999). Economic performance analysis in the design of on-line batch optimization systems. *J. Process Control*, *9*, 61–78.

- Marchetti, A., Chachuat, B., & Bonvin, D. (2009). Modifier-adaptation methodology for real-time optimization. *Ind. Eng. Chem. Res.*, *48*, 6022–6033.
- Marchetti, A., Chachuat, B., & Bonvin, D. (2010). A dual modifier-adaptation approach for real-time optimization. *J. Process Control*, *20*, 1027–1037.
- Marchetti, A., Gopalakrishnan, A., Chachuat, B., Bonvin, D., Tsikonis, L., Nakajo, A., Wuillemin, Z., & van Herle, J. (2011). Robust real-time optimization of a solid oxide fuel cell stack. *J. Fuel Cell Sci. Tech.*, *8*, 051001/1–051001/11.
- Michaels, S. E., & Pengilly, P. J. (1963). Maximum yield for specified cost. *J. Royal Stat. Soc. Series C (App. Stat.)*, *12*, 189–193.
- Montgomery, D. C. (2012). *Design and Analysis of Experiments*. Wiley.
- Myers, R., Montgomery, D., & Anderson-Cook, C. (2009). *Response Surface Methodology*. John Wiley & Sons.
- Pfaff, G. F., Forbes, J. F., & McLellan, P. J. (2006). Generating information for real-time optimization. *Asia-Pacific Journal of Chemical Engineering*, *1*, 32–43.
- Quelhas, A., Castro, N., & Pinto, J. (2013). Common vulnerabilities of RTO implementations in real chemical processes. *Can. J. Chem. Eng.*, *91*, 652–668.
- Umland, A. W., & Smith, W. N. (1959). The use of Lagrange multipliers with response surfaces. *Technometrics*, *1*, 289–292.
- Zhang, Y., Monder, D., & Forbes, J. F. (2002). Real-time optimization under parametric uncertainty: a probability constrained approach. *J. Process Control*, *12*, 373–389.