COMBINING METAMODEL TECHNIQUES AND BAYESIAN SELECTION PROCEDURES TO DERIVE COMPUTATIONALLY EFFICIENT SIMULATION-BASED OPTIMIZATION ALGORITHMS

Carolina Osorio
Hoda Bidkhori

Massachusetts Institute of Technology
77 Massachusetts Ave
Cambridge, MA 02139, USA

ABSTRACT

This paper presents a simulation-based optimization (SO) algorithm for nonlinear problems with general constraints and computationally expensive evaluation of objective functions. It focuses on metamodel techniques. This paper proposes an SO technique that also uses metamodel information when testing the improvement of the proposed points. We use a Bayesian framework, where the parameters of the prior distributions are estimated based on probabilistic metamodel information.

In order to derive an SO algorithm that achieves a good trade-off between detail, realism and computational efficiency, the metamodel combines information from a high-resolution simulator with information from a lower-resolution yet computationally efficient analytical differentiable network model. In this paper, we use the probabilistic information from the queueing model to estimate the parameters of the prior distributions. We evaluate the performance of this SO algorithm by addressing an urban traffic management problem using a detailed microscopic traffic simulator of the Swiss city of Lausanne.

1 INTRODUCTION

1.1 Motivation

This paper considers simulation-based optimization (SO) problems that rely on stochastic and computationally expensive evaluations of the underlying objective function. This research is motivated by problems in the field of urban transportation, where the use of detailed simulators, known as microscopic simulators, is mainly limited to what-if analysis. This is due to the stochasticity and the high-computational cost of evaluating these simulators. The problems considered have continuous decision variables with nonlinear and nonconvex objective functions and constraints. These are challenging SO problems. Additionally, we focus on methods with good short-term performance, such that they can be used in a practical manner to address complex problems, that is we evaluate our algorithms under tight computational budgets (e.g., allowing only for a small and limited number of simulation runs or run-time).

The problems of interest can be formulated as:

$$\min_{x \in \Psi} f(x;q) \equiv E[F(x;q)] \tag{1}$$

The objective function f is usually the expected value of a suitable stochastic network performance measure, F. It is a function of a deterministic and continuous-valued decision or control vector x and deterministic exogenous parameters q. The feasible space Ψ consists of a set of general (typically nonlinear) constraints that link x to q and f.

For instance, a traffic signal control problem can take F as the travel time and x as the green times for the signalized lanes. Elements such as the total demand or the network topology are captured by q. Every

simulation run leads to realizations of F, and involves sampling from the numerous probability distributions that account for uncertainty in, for instance, driver behavior (e.g., route choice for individual drivers) or traffic generation (e.g., headways of vehicles entering the network).

1.2 Metamodel Methods

This paper considers metamodel techniques for SO. Metamodels are simplified models of the underlying simulation model. The most common metamodels are analytical and deterministic functions, their functional form is typically chosen based on asymptotic properties or their computational efficiency. Metamodels are typically used to approximate the stochastic objective function or constraints of the problem.

Metamodels have been classified as either functional (also called general-purpose) models and physical models ((Søndergaard 2003); (Serafini 1998)). The latter are problem-specific approximations of the objective function or constraints. They have parameters that have a physical or structural interpretation. Reviews of metamodels include Barton and Meckesheimer (2006) and Kleijnen (2008).

Functional metamodels are chosen based on their analytical and numerical properties. The most common approach is the use of quadratic polynomials. They are often a linear combination of basis functions from a parametric family. Examples of functional models include higher-order polynomials, splines, radial-basis functions. These metamodels are suitable to address a variety of problems, yet fail to provide any structural information about the problem at hand.

In past work, we have proposed a metamodel that combines a physical metamodel with a functional metamodel (Osorio and Bierlaire 2010). The functional component is a quadratic polynomial, which ensures asymptotic metamodel properties (needed to analyze asymptotic convergence properties), whereas the physical metamodel provides structural information about the underlying problem, and more importantly its functional form depends on the actual problem and objective function considered. The metamodel combines information from a low-resolution but computationally efficient analytical queueing model with high-resolution simulated data.

The use of the combined metamodel has allowed us to achieve excellent short-term algorithmic performance. This SO technique been successfully used to address challenging traffic management problems ((Osorio and Chong 2012); (Osorio and Nanduri 2012); (Chen, Osorio, and Santos 2012)). With very small sample sizes the SO algorithm has identified points with improved performance. This is mainly achieved by the physical component of the metamodel, which does not need a large sample size to accurately approximate the objective function. The focus of this paper is to further enhance the short-term performance of SO algorithms. In this paper, we propose to use metamodels to go beyond the approximation of the problem formulation. We propose to use them to improve the point selection step (also known as selection procedure). At a given iteration, the algorithm must determine whether the newly identified point (called trial point) has improved performance compared to the point currently considered the best (called current iterate). This decision is called the point selection step.

This paper is organized as follows. In Section 2, we present the proposed method. We then present the optimization problem (Section 3), followed by empirical results (Section 4). We finish the paper with a brief conclusion.

2 METHOD

2.1 Algorithmic Framework

In this paper, we use the SO algorithm proposed in Osorio and Bierlaire (2010). We first describe the main steps of this algorithm, and detail its point selection step. We then present the novel point selection step.

The algorithm is a derivative-free trust region (TR) algorithm, which is based on the method proposed in Conn, Scheinberg, and Vicente (2009). For an introduction to TR methods, we refer the reader to Conn, Gould, and Toint (2000). They summarize the main steps of a TR method in the *Basic trust region algorithm*. The main idea of TR methods is to build, at each iteration, a model of the objective function

which one "trusts" in a neighborhood of the current iterate (which is the point currently considered as the best), called the *trust region*.

The method proposed by Conn, Scheinberg, and Vicente (2009) builds upon the *Basic TR algorithm* by adding two additional steps: a model improvement step and a criticality step. For a detailed description, see Conn, Scheinberg, and Vicente (2009).

A given iteration k of the algorithm considers a metamodel m_k, an iterate x_k and a TR radius Δ_k. Hereafter, the subscript k refers to the iteration. Each iteration consists of 5 steps:

- **Criticality step.** This step may modify m_k and Δ_k if the measure of stationarity is close to zero.
- **Step calculation.** Approximately solve the TR subproblem to yield a trial point, which is a point that the metamodel predicts has improved performance compared to the current iterate.
- **Point selection step: acceptance or rejection of the trial point.** The actual (i.e., simulated) reduction of the objective function is compared to the reduction predicted by the model, this determines whether the trial point is accepted or rejected. Let x_k denote the current iterate, $x_k + s_k$ the current trial point, $\hat{f}(x_k)$ (respectively, $\hat{f}(x_k + s_k)$) the performance of the current iterate (resp. trial point) estimated by the simulator, $m_k(x_k)$ (resp. $m_k(x_k + s_k)$) the performance of the current iterate (resp. trial point) approximated by the metamodel. This step of the algorithm computes

$$\hat{\rho}_k = \frac{\hat{f}(x_k) - \hat{f}(x_k + s_k)}{m_k(x_k) - m_k(x_k + s_k)}, \qquad (2)$$

  and if $\hat{\rho}_k \geq \eta_1$, then the trial point is accepted (i.e., it becomes the current iterate: $x_{k+1} = x_k + s_k$); otherwise it is rejected.
- **Model improvement.** Either certify that m_k is *fully linear* (i.e., satisfies Taylor-type bounds) in the TR or attempt to improve the accuracy of the metamodel.
- **TR radius update.**

The existing point selection step determines the trial point acceptance or rejection according to the value of $\hat{\rho}_k$ given by Equation (2). If $\hat{\rho}_k \geq \eta_1$, then the trial point is accepted (i.e., it becomes the current iterate: $x_{k+1} = x_k + s_k$); otherwise it is rejected (i.e., $x_k = x_{k+1}$). The parameter η_1 is the trial point acceptance threshold.

The current algorithm merely compares sample averages, and thus does not account for the stochasticity of $\hat{f}(x_k)$ and $\hat{f}(x_k + s_k)$ when testing for improvement. In this paper, we replace the existing point selection step with a probabilistic metric that accounts for this stochasticity.

Accounting for this stochasticity is particularly important when the sample averages are computed based on small samples. This is typically the case when SO algorithms are used under tight computational budgets, e.g., when the number of simulation runs is limited and small. This is the case of the simulation-based urban transportation problems that motivate this work.

The new point selection step proceeds as follows. If

$$Pr(\hat{\rho}_k \geq \eta_1) \geq p_0, \qquad (3)$$

then we accept the trial point; otherwise we reject the trial point. The parameter p_0 is an exogenous threshold probability.

The probability $Pr(\hat{\rho}_k \geq \eta_1)$ is usually estimated by combining sample average and sample variance information. Nonetheless, when the SO algorithms are used under tight computational budgets the number of replications at each point is very small (e.g., only a couple of replications for each point). Thus, the sample average and variance estimates may be inaccurate, leading to SO methods with poor short-term (i.e., small sample size) performance.

To overcome this, we propose to interpret the probability in Equation (3) as that arising from a posterior distribution of a Bayesian framework. We use information from a probabilistic analytical metamodel to

estimate the parameters of the corresponding prior distributions. We expect this approach to improve the short-term performance of the algorithms.

2.2 Bayesian Framework

Consider a given iteration of an SO algorithm and a set of points that are considered candidate solutions to the optimization problem. Selection procedures determine at each step which points should remain, be included or be removed from this set of candidate solutions. Most traditional SO algorithms are based on selection procedures that are formulated for problems with a discrete and finite number of feasible points. For a recent review of such methods, see Nelson (2010). Our focus is on continuous SO problems. Nonetheless, this paper builds upon ideas initially proposed in the discrete SO literature.

In Inoue (2000) (Section 2.1.2) (alternatively see Chick and Inoue (2001)), the following Bayesian approach to compare the performance of two points (or systems) with unknown mean and variances is proposed. Suppose from a given point i, r_i independent replications are run yielding the following r_i independent observations of the underlying objective function: $y_{i1}, \cdots, y_{ir_i}$. We assume these observations to be normally distributed with unknown mean w_i and unknown variance σ_i^2. Given the Bayesian setting, we will parameterize the distribution according to its precision: $d_i = 1/\sigma_i^2$. Both w_i and d_i are unknown parameters, and are thus modeled as random variables: W_i and D_i.

The approach proposed by Inoue (2000) uses the conjugate prior distributions of a normal distribution, i.e.,

$$D_i \sim \mathscr{G}(\alpha_i, \beta_i) \tag{4}$$

$$W_i \mid d_i \sim \mathscr{N}\left(\theta_i, \frac{1}{a_i d_i}\right), \tag{5}$$

where $\mathscr{G}$ denotes the gamma distribution and $\mathscr{N}$ denotes the normal distribution. To compare the performance of two points, the posterior marginal distribution of $W_1 - W_2$ is used. This distribution is given by:

$$W_1 - W_2 | (y_1, y_2) \sim \mathbf{St}(\mathbf{z_1} - \mathbf{z_2}, \omega_{1,2}, \nu_{1,2}), \tag{6}$$

where y_i represents the vector of all replications of point i and $\mathbf{y_i}$ denote the random variables associated to the objective function of the point i. The derivation of this posterior distribution is detailed in Section 9.6 of DeGroot (1970). The parameters of this Student distribution are detailed in Inoue (2000). They are functions of the parameters of the prior distributions, as well as sample average and sample variance estimates.

The contribution of this paper, is to use prior information provided by the metamodel to fit the parameters of the prior distributions. We expect this to improve the performance of the point selection step, particularly for small sample sizes. Thus, we expect this approach to improve the short-term performance of the considered SO algorithm.

As is discussed in DeGroot (1970) and Inoue (2000), the posterior distribution of a difference $W_1 - W_2$ can be approximated as Equation (6) where the parameters can be computed as follows:

- $\nu_{1,2} = \dfrac{\nu_1^2 + \nu_2^2}{\dfrac{\nu_1^2}{2\alpha_1 + r_1} + \dfrac{\nu_2^2}{2\alpha_2 + r_2}}$

- $\nu_i = \dfrac{\beta_{\mathbf{y_i}}}{(a_i + r_i)(\alpha_i + \frac{r_i}{2})}$

- $\beta_{\mathbf{y_i}} = \beta_i + \dfrac{1}{2}\left(\sum_{j=1}^{r_i}(y_{ij} - \bar{\mathbf{y}}_\mathbf{i})^2 + \dfrac{a_i r_i}{a_i + r_i}(\theta_i - \bar{\mathbf{y}}_\mathbf{i})^2\right)$

- $\omega_{1,2} = \left(\dfrac{\beta_{\mathbf{y_1}}}{(a_1 + r_1)(\alpha_1 + \frac{r_1}{2})} + \dfrac{\beta_{\mathbf{y_2}}}{(a_2 + r_2)(\alpha_2 + \frac{r_2}{2})}\right)^{-1}$

- $\mathbf{z_1} - \mathbf{z_2} = \dfrac{a_1 \theta_1 + r_1 \bar{\mathbf{y}}_1}{a_1 + r_1} + \dfrac{a_2 \theta_2 + r_2 \bar{\mathbf{y}}_2}{a_2 + r_2}$

2.3 Metamodel

The metamodel combines information from a simulation model, which is in our context a microscopic urban traffic simulator (Dumont and Bert 2006), with information from an analytical probabilistic traffic model (see Osorio and Bierlaire (2009)). The latter is a network model based on finite capacity queueing theory. In this paper, this metamodel is used both to approximate the objective function of the problem, and to provide prior information by estimating the parameters of the prior Bayesian distributions described in the previous section. For a detailed description of the metamodel, we refer the reader to Osorio and Bierlaire (2010).

We use the queueing model to approximate the expectations of each of the two prior distributions. According to the notation of Equation (1), the objective function is given by $E[F(x_i; q_i)]$. Let e_i denote the approximation of this objective function provided by the queueing model. We then use e_i to estimate the expectation of the prior distributions.

We use uniform non-informative prior to estimate the the prior values of the gamma distributions α_i and β_i, we set $\alpha_i = 0.01$ and $\beta_i = 0.001$. Even though we use non-informative prior, still the Bayesian framework allows us to take the variance of simulator observations into account which helps us to make a more accurate decision.

The remaining two parameters of the prior distributions are estimated based on the weight that we want to give prior information from the metamodel compared to simulated information. These two parameters are therefore set according to the problem formulation at hand.

3   OPTIMIZATION PROBLEM

To evaluate the short-term performance of this approach, we consider a fixed-time traffic signal control problem as formulated in detail in Osorio and Bierlaire (2010) and in Osorio and Bierlaire (2009). In order to formulate the problem, we introduce the following notation.

b_i	available cycle ratio of intersection i;
$x(j)$	green split of phase j;
x_L	vector of minimal green splits;
$\mathscr{I}$	set of intersection indices;
$\mathscr{P}_I(i)$	set of phase indices of intersection i.

The problem is formulated as follows.

$$\min_{x} E[F(x; q)] \tag{7}$$

subject to

$$\sum_{j \in \mathscr{P}_I(i)} x(j) = b_i, \ \forall i \in \mathscr{I} \tag{8}$$

$$x \geq x_L. \tag{9}$$

The decision vector x consists of the green splits which are normalized green times for each phase, the objective is to minimize the expected travel time. Equation (9) ensures minimal green times for each phase. Equation (8) ensures that for each intersection the available green time is fully allocated across all phases. At each iteration of the algorithm, the trust region subproblem is solved. The detailed formulation of this subproblem is given in Osorio and Bierlaire (2010).
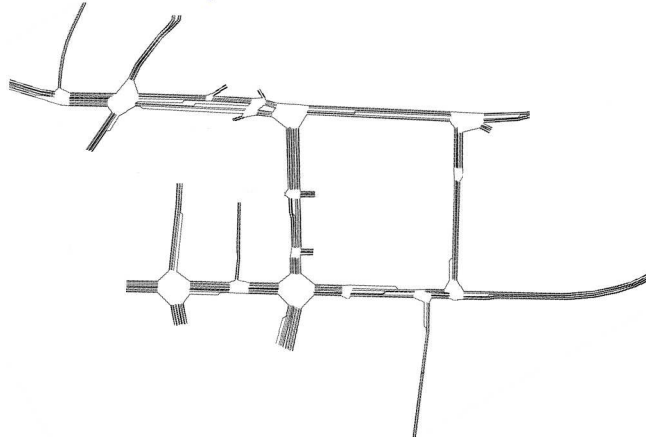
Figure 1: Subnetwork of the Lausanne city center.

4   EMPIRICAL ANALYSIS

We evaluate and illustrate the use of this framework with a case study based on the road network of the Swiss city of Lausanne. We use a calibrated microscopic traffic simulation model of the Lausanne city center (Dumont and Bert 2006). This model is implemented with the AIMSUN simulator (TSS 2008). Details regarding the Lausanne network are given in Osorio and Bierlaire (2009). This network considers demand for the evening peak period.

We consider a set of 48 roads and 15 intersections, these are displayed in Figure 1. The signalized intersections have a cycle time of either 90 or 100 seconds. Nine intersections are signalized and control the flow of 30 roads. There are 102 lanes, 60 of which are signalized. There are a total of 51 phases that are considered variable (i.e., the dimension of the decision vector is 51). The queueing model of this subnetwork consists of 102 queues (one queue for each lane). The TR subproblem consists of 621 endogenous variables with their corresponding lower bound constraints, 408 nonlinear equality constraints, 171 linear equality constraints and 1 nonlinear inequality constraint. Note that this problem is considered a large-scale problem for existing unconstrained derivative-free methods, not to mention the added complexity of the nonlinear constraints.

We consider two different initial points, which are randomly drawn signal plans. For each initial point, we proceed as follows. We consider a tight computational budget, which is defined as a maximum number of simulation runs that can be carried out. The computational budget is set to 150 runs. We run each algorithm 3 times, and each time allow for these 150 simulation runs. This yields three different "optimal" solutions (i.e., solutions which are expected to have improved performance). We then use the simulator to evaluate in detail the performance of these solutions.

To evaluate the performance of a given signal plan, we run 50 replications of the simulation model, and plot the empirical cumulative distribution function (cdf) of the average travel times over these 50 runs. The empirical cdf's of the different signal plans are then compared.

Figures 2 and 3 display results for two different initial points. Figure 2 displays the cdf's of the average travel time in the controlled subnetwork. This figure contains two plots, one for each initial signal plan. Each plot displays 4 empirical cdf's. The solid curves are the cdf's of the signal plans proposed by our method. The dashed curve is the cdf of the initial signal plan. For both initial points, the proposed methodology systematically identifies signal plans with improved performance even under tight computational budgets.

Figure 3 represents the results for the same 2 initial points, but displays the average travel time of the full network (i.e., accounting also for the roads outside the subnetwork). This figure indicates that the proposed plans systematically provide improvement of travel times at the full city-scale.
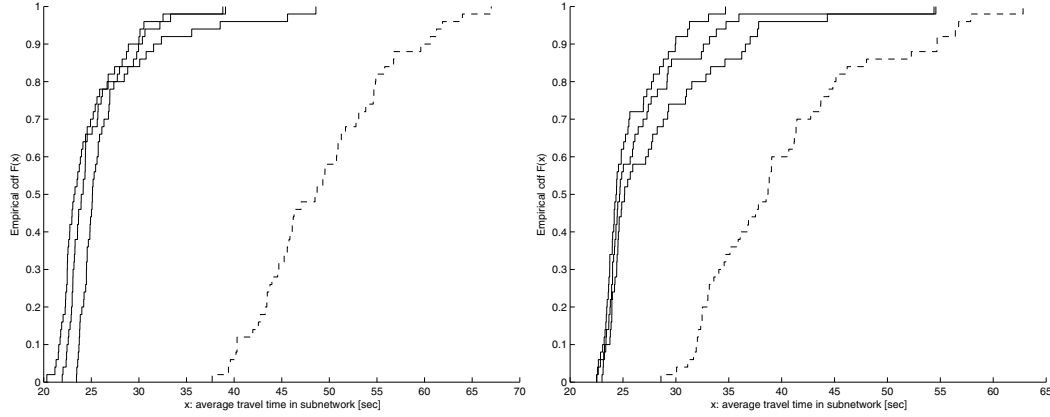
Figure 2: Empirical cdf's of the average travel times in the subnetwork which obtained by both approaches.
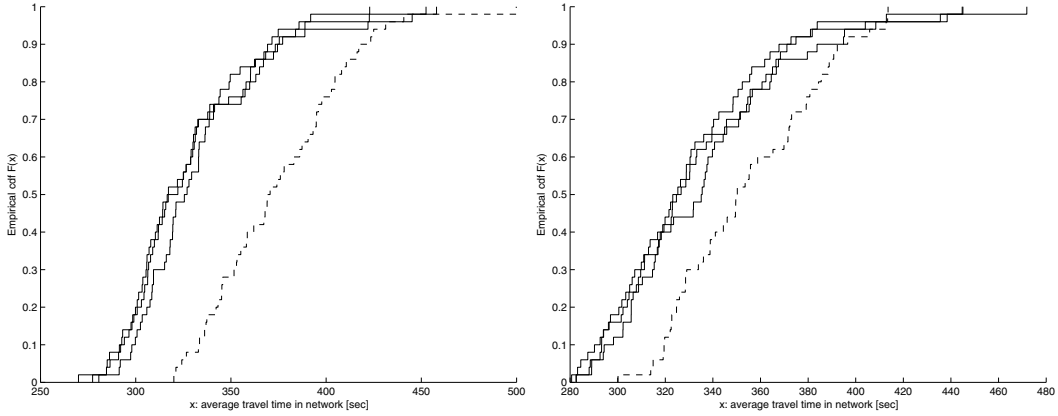

Figure 3: Empirical cdf's of the average travel times in the fullnetwork which obtained by both approaches.

We compare the performance of the proposed approach to that of the initial approach (where the point selection step consists of a comparison of sample averages). We consider an initial point which is a randomly drawn signal plan. Again, we consider a tight computational budget, the computational budget is set to 150 runs. We run each algorithm 4 times and each time allow for 150 simulation runs. Each of these algorithms yields four different "optimal" solutions. We then use the simulator to evaluate in detail the performance of these solutions. To evaluate the performance of a given signal plan, we run 50 replications of the simulation model, and plot the empirical cumulative distribution function (cdf) of the average travel times over these 50 runs. The empirical cdf's of the different signal plans are then compared.

Figure 4 displays the cdf's of the average travel time in the controlled subnetwork. This figure displays 4 empirical cdf's for Bayesian framework optimal plans and also 4 empirical cdf's for the initial SO algorithm. The dashed curves are the cdf's of the signal plans proposed by our method. The solid curves is the cdf's of the original trust region framework. As we see in Figure 4 for a given initial point, our proposed methodology systematically identifies signal plans with improved performance even under tight computational budgets.
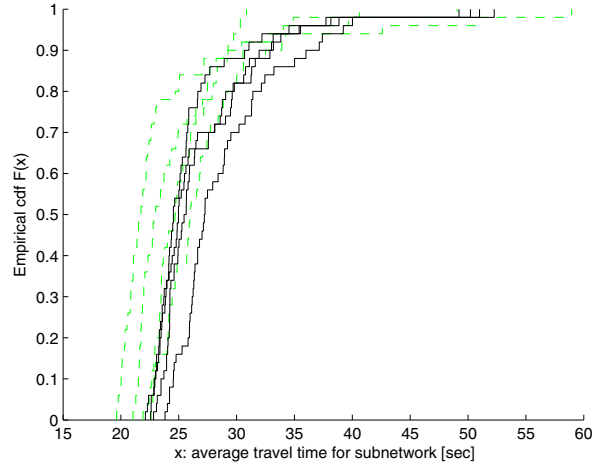
Figure 4: Compare the performance of the proposed and the initial approach.

5 CONCLUSION

This paper proposes a simulation-based optimization algorithm suitable to address complex transportation problems under tight computational budgets. The approach is a metamodel approach. This paper goes beyond the traditional use of metamodels, which is to approximate the problem formulation. It uses information from the metamodel to fit prior distributions that are used when comparing the performance of different points. We address a signal control problem. The proposed method identifies signal plans that improve local and as well as city-wide travel times, and does so within tight computational budgets. As part of this ongoing research, we are evaluating the performance of this method to that of other metamodel methods, which do not use metamodel information in the selection procedure.

This research is motivated by complex problems that need to be addressed in practice by running few runs of the underlying simulator. Going beyond the traditional use of metamodels, and more generally using prior information about the problem, can allow us to improve the short-term performance of simulation-based optimization algorithms. Such information can be used, for instance, to inform various sampling decisions, and more generally to improve the way we allocate our computational budgets.

REFERENCES

Barton, R. R., and M. Meckesheimer. 2006. "Metamodel-based simulation optimization". In *Handbooks in operations research and management science: Simulation*, edited by S. G. Henderson and B. L. Nelson, Volume 13, Chapter 18, 535–574. Amsterdam: Elsevier.

Chen, C. Osorio, and B. F. Santos. 2012. "A simulation-based approach to reliable signal control". *International Symposium on Transportation Network Reliability (INSTR)*.

Chick, S. E., and K. Inoue. 2001. "New two-stage and sequential procedures for selecting the best system". *Operations Research* 49 (5): 732–743.

Conn, A. R., N. I. M. Gould, and P. L. Toint. 2000. *Trust-region methods*. MPS/SIAM Series on Optimization. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics and Mathematical Programming Society.

Conn, A. R., K. Scheinberg, and L. N. Vicente. 2009. "Global Convergence of General Derivative-Free Trust-Region Algorithms to First- and Second-Order Critical Points". *SIAM Journal on Optimization* 20 (1): 387–415.

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.

Dumont, A. G., and E. Bert. 2006, May. "Simulation de l'agglomération Lausannoise SIMLO". Technical report, Laboratoire des voies de circulation, ENAC, Ecole Polytechnique Fédérale de Lausanne.

Inoue, K. 2000. *Decision-theoretic comparison of alternate system configurations using stochastic simulation*. Ph. D. thesis, University of Michigan, Ann Arbor MI, USA.

Kleijnen, J. P. C. 2008. *Design and Analysis of Simulation Experiments*. New York, USA: Springer.

Nelson, B. L. 2010. "Optimization via Simulation Over Discrete Decision Variables". *Tutorials in Operations Research* 7:193–207.

Osorio, C., and M. Bierlaire. 2009. "A surrogate model for traffic optimization of congested networks: an analytic queueing network approach". Technical Report 090825, Transport and Mobility Laboratory, ENAC, Ecole Polytechnique Fédérale de Lausanne.

Osorio, C., and M. Bierlaire. 2010. "A simulation-based optimization approach to perform urban traffic control". In *Proceedings of the Triennial Symposium on Transportation Analysis (TRISTAN)*. Tromsø, Norway.

Osorio, C., and L. Chong. 2012. "An efficient simulation-based optimization algorithm for large-scale transportation problems". *Winter Simulation Conference (WSC), Berlin, Germany.*.

Osorio, C., and K. Nanduri. 2012. "Energy-efficient traffic management: a microscopic simulation-based approach". *International Symposium on Dynamic Traffic Assignment (DTA), Marthas Vineyard, USA*.

Serafini, D. B. 1998. *A framework for managing models in nonlinear optimization of computationally expensive functions*. Ph. D. thesis, Rice University, Houston, Texas, USA.

Søndergaard, J. 2003. *Optimization using surrogate models - by the Space Mapping technique*. Ph. D. thesis, Technical University of Denmark, Lyngby, Denmark.

TSS 2008, May. *AIMSUN NG and AIMSUN Micro Version 5.1*. Transport Simulation Systems.

AUTHOR BIOGRAPHIES

CAROLINA OSORIO is an Assistant Professor of Civil and Environmental Engineering (CEE) at the Massachusetts Institute of Technology (MIT). She received her Ph.D from EPFL in 2010. She develops operations research techniques to inform the design and operations of transportation systems. She is interested in the techniques that combine ideas from the fields of probability theory, simulation, simulation-based optimization, derivative-free optimization, nonlinear optimization, statistics, traffic control and traffic flow theory. Her email address is osorioc@mit.edu.

HODA BIDKHORI received her Ph.D. degree in Applied Mathematics from the Massachusetts Institute of Technology, Cambridge, in 2010. Since Dec 2011, she has been a postdoctoral associate in the Department of Civil and Environment Engineering at the Massachusetts Institute of Technology working on operations Research. Her research interests include applications of optimization and mathematical methods to operations research and transportation problems. Her email address is bidkhori@mit.edu.