Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

INTEGRATED STOCHASTIC OPTIMIZATION AND STATISTICAL EXPERIMENTAL DESIGN FOR MULTI-ROBOT TARGET TRACKING

Timothy H. Chung

Department of Systems Engineering
Naval Postgraduate School
777 Dyer Road, BU-218
Monterey, CA 93943, USA

James C. Spall

Johns Hopkins University
Applied Physics Laboratory
11100 Johns Hopkins Rd.
Laurel, MD 20723 USA

ABSTRACT

This paper presents an integrated approach for enhancing the performance of stochastic optimization processes by incorporating techniques from statistical experimental designs, such as response surface methodology. The two-stage process includes an "exploratory" phase, during which a fraction of the finite time budget is reserved for conducting informative measurements to best approximate the stochastic loss function surface, followed by execution of the optimization process for the remaining time. We formulate a representative stochastic optimization problem for the case of multiple distributed mobile sensors engaged in surveillance for one or more objects of interest. We show via simulation studies that the employment of such an exploratory phase, with the use of screening experimental designs to provide local approximations to the response surface, improves the stochastic optimization process.

1 INTRODUCTION

Many challenging problems in multi-agent systems involve optimization of an objective function in the presence of uncertainty, as evidenced by the increasing interest in the field of stochastic optimization. One such area of continued interest that intersects the optimization, controls, and operations research communities is the use of stochastic optimization techniques for resource allocation models, such as informing efficient surveillance of an area of operations with distributed mobile sensors. Such problems are not only challenging due to the spatially and temporally dynamic and stochastic environments, the difficulty in identifying universally robust and stable optimization methodologies often requires problem-specific solution approaches. Furthermore, the sensitivity to initial conditions of the optimization performance, as well as working within constraints on computational feasibility, can also hinder one-size-fits-all approaches.

The work presented in this paper investigates a framework for initialization of the stochastic optimization process that is shown through simulation studies to enhance the performance of the optimization. As an operationally relevant application, we investigate the efficient and effective employment of a team of autonomous mobile sensors performing an information-gathering mission, i.e., the coordinated state estimation of an object of interest.

1.1 Related Works

Multi-agent optimization, particular in the context of autonomous systems, has been extensively studied in numerous contexts, ranging from coverage control (Cortés, Martínez, Karatas, and Bullo 2004) to surveillance and patrolling (Bazzan and Labidi 2004) to distributed target tracking (Chung, Burdick, and Murray 2006, Yang, Freeman, and Lynch 2007) applications. The objectives in these works are nominally formulated as (stochastic) optimization problems, often addressing the minimization of a measure of uncertainty, such as entropy or estimate error covariances (Caiti, Calabro, Di Corato, Meucci, and Munafo 2013). Analogously, in recent literature, much work has been done to maximize the (expected) information gain obtained by one or more coordinated agents (Schwager, Dames, Rus, and Kumar 2011, Dames, Schwager, Kumar, and Rus 2012). However, despite recognition of computational concerns, most of these works do not address how best to efficiently allocate a computation budget, so as to limit the number of stochastic objective function calculations or measurements.

Related to this concern, various stochastic optimization techniques address the task of optimizing an objective in the presence of noise, e.g., only noisy measurements of the objective function are available, which impacts the convergence to and guarantees of optimal solutions (Spall 2003). Statistical methods to estimate the objective function (or its gradient) as a way to address the noise usually require repeated, often expensive, measurements, further impeding efficiencies. Despite these challenges, much interest in stochastic optimization, coupled with the advent of greater computational resources and broadening domain applications, have led to the development of methods such as stochastic approximation (Spall 2003) or simulation-based approaches (e.g., simulated annealing (Van Laarhoven and Aarts 1987), evolutionary computation (Back, Fogel, and Michalewicz 1997)), to include those designed for decentralized optimization problems (Bertsekas and Tsitsiklis 1989). An emerging area of research, however, remains the development of rigorous methods for the initialization of the stochastic optimization process, which often can significantly impact, sometimes quite adversely, the overall optimization performance.

Identifying the most informative collection of measurements, which could refine the stochastic optimization process and improve its performance, is well suited for the application of statistical experimental design theory. These methods for the efficient design of experiments generate the most statistically relevant observations (Montgomery 2009), and have increasingly been incorporated in numerous engineering areas, to include robotics, controls, and optimization. Though early works have explored the role of experimental design in defining, for example, optimized sensor placements in an environment (Ucinski 2009), further integration with multi-agent systems research and stochastic optimization, such as via the use of response surface methods (Myers, Montgomery, and Anderson-Cook 2009), highlight numerous potential benefits in intelligently reducing the number of measurements executed while still providing measurably good statistical models of the stochastic objective function.

The objective of the presented approach can thus be stated to be: *Given a budget on the number of measurements the system can use, identify an appropriate fraction of the total budget that should be allocated between "exploratory" sampling, that is, using an experimental design-like approach to minimize uncertainty in the loss function, and "exploitative" sampling, which is seeking to minimize the loss function itself.*

1.2 Summary of Contributions

Given the breadth of these previous works, the main contributions of this paper include the development of a framework that integrates statistical experimental design, namely response surface methodology (RSM), with the stochastic optimization of a noisy loss function designed to represent the collective sensor fusion performance of a team of distributed mobile sensors conducting a target tracking task. In this paper, we present characterization of the trade-offs between conducting "exploratory" measurements (via experimental design techniques) to better estimate the loss function and executing "exploitative" control actions to optimize the noisy loss function itself. Simulation studies provide some insights into the benefits and penalties associated with two-phase approach which seeks to address the problems in the initialization of the optimization process.

Having motivated the problem in the context of previous works in this section, Section 2 describes statistical experimental design approaches for highlighting ideal initial conditions for the optimization process. Section 3 presents a formulation of a stochastic optimization model framework for the multi-agent information-gathering scenario of interest. Simulation model development and numerical studies are presented in Section 4, including relevant results and insights. Section 5 concludes the paper and highlights future avenues of research.

2 INTEGRATING STATISTICAL EXPERIMENTAL DESIGN

In many cases, due to both real-world operational considerations and computational expense, the cost of obtaining a measurement, that is, sampling the stochastic loss function, may be both high and constrained. For example, given a budget of the number of measurements allowable, one common but naïve approach for addressing the noise in such stochastic optimization problems is to take and average more measurements at fewer points in the domain, while simultaneously attempting to optimize the loss function.

However, the expense of taking many arbitrary measurements at a given point could be minimized by more intelligent selection of sampling locations to be most informative, thereby reducing the number of measurements necessary to achieve statistical confidences. This very topic is the subject of the field of statistical design of experiments.

Instead of directly optimizing the stochastic loss function initially, the approach proposed in this paper reserves a budget of k measurements (one per time step) for the application of "exploratory" observations, and using these few but well chosen observations to develop a statistical regression model to provide a model of the shape of the stochastic loss function, also known as the *response surface* in experimental design literature (Spall 2010).

2.1 Response Surface Methodology

Response surface methodology has been used in a variety of applications to experimentally optimize a response surface, i.e., a sampled loss function. The general concept, as summarized in, e.g., (Spall 2010), is to sample the loss function in the neighborhood of a given initial point in the factor space, and by locally approximating the surface at that point, one can compute an estimate of the gradient that can assist in finding optima.

More rigorously, consider the choice of a screening design that admit no more than k experiments (a.k.a. design points). Some canonical options include (fractional) factorial or central composite designs, with alternatives for more complex designs such as optimal (e.g., D-optimal) or space-filling designs (Montgomery 2009, Kleijnen, Sanchez, Lucas, and Cioppa 2005).

For a first-order model that captures only the main effects (i.e., no interactions between the factors), where main effects may collectively represent physical coordinates of the team of mobile sensing robots, one can construct a regression model based on the responses from these screening experiments that takes the form

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_n x_n = \beta_0 + \boldsymbol{\beta}^T \ \boldsymbol{x},$$

where $\boldsymbol{\beta} \in \mathbb{R}^n$ is the column vector of all coefficients corresponding to main effects. Then the gradient direction of steepest descent (in the case of trying to minimize the loss function) is given simply by $-\boldsymbol{\beta}$ (Spall 2010).

One can conduct additional experiments for factor variations in the direction of this steepest descent, continuing until no additional statistically significant reduction in the response surface is obtained. Additional tests, such as evaluation of the Hessian at these points of nearly zero gradient, can be conducted to identify the nature of the fixed point. Nominally, the response surface methodology procedure is repeated, potentially employing more complex experimental designs, such as central composite designs, at each subsequent iteration that might include higher-order terms or additional center points to capture any possible nonlinearities (Myers, Montgomery, and Anderson-Cook 2009).

However, as will be discussed in the following section, the allowable budget of experiments to be conducted may not allow for repeated application of this experimentally driven optimization process. Rather, given a fixed allotment of experiments that can be conducted in this manner, with the remaining sampling opportunities reserved for direct use in specific stochastic optimization routines, one must identify not only which class of experimental designs to select from but also how many experiments (within the budget of k measurements) to employ in the design and how many to leverage to progress in the steepest descent direction of the response surface.

2.2 Application of RSM to Multi-Agent Control

Recall that the presented approach seeks to improve the overall optimization performance by initially allotting a fraction of the total observation budget to perform exploratory samples of the loss function, leveraging statistical experimental design techniques to identify *which* samples are likely to be most informative.

In the context of the multi-agent optimization to improve target tracking performance, in keeping with response surface methodology procedures, we can apply the notion of identifying a set of locations in the neighborhood of mobile sensors' initial positions. A first-order model, which can be adequately found by use of a screening design, offers an initial approximation of the response surface. For M sensing agents in $\mathbb{R}^2$, the factor space has dimensionality $2M$, that is, deviations about the initial locations for all sensor agents in x, y directions. A number of simple experimental designs may provide adequate resolution; consider the use of a two-level, fractional factorial design, such that a subset of high and low levels of both x and y deviations are considered. For M agents, the total number of resulting design points is eight, noting that a full factorial design (for two-levels) requires $2^{2M} = 16$ experiments. Having selected a simple model to motivate the proposed approach, one can easily recognize that the number of such measurements increases exponentially with increasing numbers of sensor agents and/or state variables (e.g., velocities, headings) that might also be considered decision variables.

A number of different design approaches for screening are possible, each with potentially different number of required experiments. The trade off between reduction of the number of experiments and the ability to isolate which factors are likely to be most informative is of interest to our characterization. See Montgomery (2009) for an introductory overview.

Suppose we are able to collect samples according to the specified design. What remains is to generate a simple first-order linear regression model that can provide some insights into the response surface. In particular, consider the first-order regression model, with four factors (for $M = 2$ case)

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_1 x_2 + \cdots$$

The idea is then to use the collected response data and fit the best possible linear fit. Note that the gradient of this model as a function of a given factor can be computed analytically as well.

Consider, for example, the use of a fractional factorial design, highlighting only the direct impact (a.k.a. *main effects*) of the factors, that is, the sensor agent positions, and not any correlations between their respective positions. Then a minimalist experiment design can be generated, such as in Table 1, in the case of two mobile sensors, where -1 corresponds to the low level and 1 denotes the high level for the corresponding factor. For example, a 1 entry for x_1 indicates sampling at $x_1 + \Delta x$, or a -1 entry for y_2 suggests sampling at $y_2 - \Delta y$.

From the responses (i.e., measurements results) arising from such a design matrix, we can construct a regression model of the form:

$$\hat{L} = \beta_0 + \beta_{x_1} x_1 + \beta_{y_1} y_1 + \beta_{x_2} x_2 + \beta_{x_2} y_2,$$

where (x_1, y_1) and (x_2, y_2) are considered the factors and represent the location of the two sensor robots, and the coefficients β_{x_i} and β_{y_i} define the significance of their associated factors in impacting the response.

In the context of the measurement budget, k, allocated for locally exploring the current sensor robot locations, any remaining available measurements after completing the samples according to the experimental design can be conducted while moving in the gradient direction determined by the regression model coefficients. For example, as will be seen, if $k = 25$ and there are eight nearby locations, i.e., design points,

Table 1: Four-factor, two-level fractional factorial design, generated using MATLAB's `fracfact` command.

x_1	y_1	x_2	y_2		x_1	y_1	x_2	y_2
-1	-1	-1	-1		$x_1-\Delta x$	$y_1-\Delta y$	$x_2-\Delta x$	$y_2-\Delta y$
-1	-1	1	1		$x_1-\Delta x$	$y_1-\Delta y$	$x_2+\Delta x$	$y_2+\Delta y$
-1	1	-1	1		$x_1-\Delta x$	$y_1+\Delta y$	$x_2-\Delta x$	$y_2+\Delta y$
-1	1	1	-1	$\Longrightarrow$	$x_1-\Delta x$	$y_1+\Delta y$	$x_2+\Delta x$	$y_2-\Delta y$
1	-1	-1	1		$x_1+\Delta x$	$y_1-\Delta y$	$x_2-\Delta x$	$y_2+\Delta y$
1	-1	1	-1		$x_1+\Delta x$	$y_1-\Delta y$	$x_2+\Delta x$	$y_2-\Delta y$
1	1	-1	-1		$x_1+\Delta x$	$y_1+\Delta y$	$x_2-\Delta x$	$y_2-\Delta y$
1	1	1	1		$x_1+\Delta x$	$y_1+\Delta y$	$x_2+\Delta x$	$y_2+\Delta y$

then after completing these eight measurements and computing the response surface gradient, the remaining 17 observations are taken along the trajectories corresponding to each sensor's respective gradient direction.

3 MULTI-ROBOT TARGET TRACKING FORMULATION

Consider the multi-robot surveillance scenario in which a team of M mobile sensing agents, modeled as autonomous robots with single integrator dynamics, must cooperatively maneuver in a two-dimensional area of operations to gather, share, and fuse state information regarding one or more objects present in the environment. Let us first consider an illustrative case of a single stationary object of interest, such as an aircraft's black box or a misplaced cellphone, whose state is $\mathbf{x} \in \mathbb{R}^2$.

To illustrate a representative stochastic optimization loss function, consider the multi-agent sensor fusion and target tracking formulation presented in Chung, Burdick, and Murray (2006), for which stability analysis is performed in Yang, Freeman, and Lynch (2007), in which the loss function, L, represents the fused target state estimate error covariance, denoted $\mathbf{P}_{\text{fused}}$, with all M robots assumed to have range and bearing sensors to sense the relative location of the object:

$$L \triangleq \mathbf{P}_{\text{fused}} = \left(\sum_{i}^{M} \left(\mathbf{T}_i \bar{\mathbf{R}}_i \mathbf{T}_i^T \right)^{-1} \right)^{-1}, \tag{1}$$

$\bar{\mathbf{R}}_i \in \mathbb{R}^2$ represents the i^{th} sensor robot's uncertainty, i.e., measurement covariance matrix, in the robot's body-fixed reference frame. $\mathbf{T}_i \in \mathbb{R}^2$ is the rotation matrix to transform from the local to global reference frame, given as:

$$\bar{\mathbf{R}}_i \triangleq \begin{bmatrix} f_r(r_i) & 0 \\ 0 & f_b(r_i) \end{bmatrix}, \qquad \mathbf{T}_i \triangleq \begin{bmatrix} \cos(\phi_i) & -\sin(\phi_i) \\ \sin(\phi_i) & \cos(\phi_i) \end{bmatrix},$$

where r_i and ϕ_i are the relative distance and bearing measurements from the i^{th} robot to the object. The standard deviations of measurement errors in range and bearing, denoted $f_r(r_i)$ and $f_b(r_i)$, respectively, are assumed to depend on the relative distance, r_i (Chung, Burdick, and Murray 2006).

Derived in Chung, Burdick, and Murray (2006) is a gradient-based control law that provides a decentralized motion control strategy for each robot to individually maneuver to collectively minimize the loss function, given by Equation 2:

$$\mathbf{u}_i(r_i, \phi_i) = \begin{bmatrix} \left(\frac{\partial L}{\partial r_i}\right) & \left(\frac{\partial L}{\partial \phi_i}\right) \end{bmatrix}^T, \tag{2}$$

where q_i represents either the range, r_i, or bearing, ϕ_i, between the i^{th} robot and the target, such that we have

$$\frac{\partial L}{\partial q_i} = |\mathbf{P}_{\text{fused}}| \ \text{tr}\left[\mathbf{R}_i^{-1} \ \frac{\partial \mathbf{R}_i}{\partial q_i} \left(\mathbf{R}_i^{-1}\right)^T \ \mathbf{P}_{\text{fused}}\right],$$

and

$$\begin{aligned}
\frac{\partial \mathbf{R}_i}{\partial r_i} &= \mathbf{T}_i \ \frac{\partial \bar{\mathbf{R}}_i}{\partial r_i} \ \mathbf{T}_i^T = \mathbf{T}_i \ \begin{bmatrix} \frac{\partial f_r}{\partial r_i} & 0 \\ 0 & \frac{\partial f_b}{\partial r_i} \end{bmatrix} \ \mathbf{T}_i^T \\
\frac{\partial \mathbf{R}_i}{\partial \phi_i} &= \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \ \mathbf{T}_i \bar{\mathbf{R}}_i \mathbf{T}_i^T + \mathbf{T}_i \bar{\mathbf{R}}_i \mathbf{T}_i^T \ \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}
\end{aligned}$$

We present this previous analytic work so as to provide background and a means of comparison with the phased stochastic optimization solution strategy presented in this paper. For example, we want to investigate the difference between the perceived loss function (with errors due to, e.g., sensing errors) and the true loss function (which requires perfect knowledge of the true locations of targets). The closed-form expressions provides a means to perform this baseline comparison to validate the proposed methodology, with application to more complex formulations left for near-term study.

Returning to the illustrative example to be considered throughout the paper of $M = 2$ mobile sensor agents and a single stationary target, employment of the gradient-based control law leads to Figure 1, with the absence and presence of measurement noise highlighting the difficulty of the stochastic optimization problem. In this example, the range and bearing uncertainty functions, $f_r(r_i)$ and $f_b(r_i)$, respectively, are given by a quadratic dependence on range to the target, i.e.,

$$\begin{aligned}
f_r(r_i) &= a_2 \left(r_i - a_1\right)^2 + a_0, \\
f_b(r_i) &= \alpha f_r(r_i),
\end{aligned}$$

with coefficients $a_0 = 0.1528$, $a_1 = 15.625$, and $a_2 = 0.0008$, and scaling factor $\alpha = 5.0$ for the illustrative example. While the approach described in this paper does not depend on the specific form or value of these uncertainty functions, these parameter values represent a range-dependent error with a minimal error sensing range at 15.625 units from the target. Such a model would be appropriate for sensors, such as underwater sonar, for which its ideal sensing location is neither too close or too far from the target.

Naturally, the focus of much stochastic optimization literature is to apply methods for determining optimal values of the stochastic loss function, such as the model formulated above. The work presented in this paper, however, highlights the value of also optimizing the initial phases of the optimization process, and is driven by techniques to efficiently sample the stochastic loss function to improve the resulting optimization. The general idea is to mitigate the stochasticity of the loss function by allocating some of the available sensing budget, in the form of k time steps, to perform exploratory sampling prescribed by an experimental design, as was discussed previously in Section 2.

4 ANALYSIS AND RESULTS

We investigate whether there is indeed any advantage to enabling an “exploratory” period by employing experimental design techniques for computing an initial direction for all sensor agents. Several key elements of this initial study are the focus, including assessing the role of different types of screening designs, the allocation between exploratory samples, sampling along the vector indicated by the local approximation, and transitioning to standard stochastic optimization approaches, as well as ultimately the fraction of the time spent in “exploratory” versus “exploitative” phases of optimization.

We first consider the impact of varying the measurement budget (k) of the “exploratory” phase as well as that of different screening designs, to include full factorial, fractional factorial, D-optimal (Montgomery

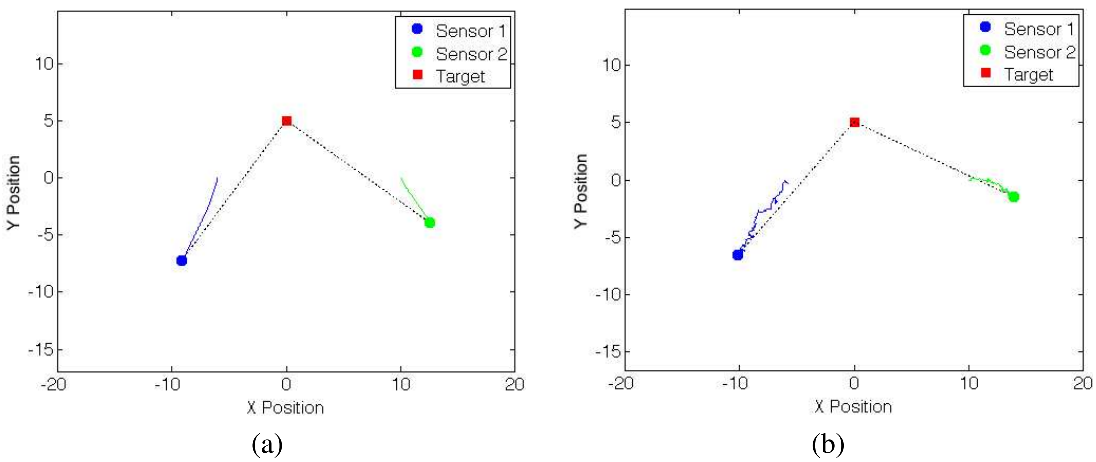

Figure 1: Example output of the gradient-based control law. Trajectories shown are two sensor agents in the (a) absence and (b) presence of measurement noise in range and bearing.

2009), and other space-filling designs such as the Nearly-Orthogonal Latin Hypercube (Kleijnen, Sanchez, Lucas, and Cioppa 2005, Sanchez 2011). These alternatives are constructed, for example, using MATLAB's `fracfact` and `rowexch` functions for fractional factorial and *D*-optimal designs, respectively, which can all be compared against the cases of perfect (i.e., in the absence of noise) gradient calculations using Equation 2 as well as the case where the gradient-based control law is executed immediately, that is, no budget for exploration is provided.

For the illustrative two-sensor, single-target example, we investigate the value of the loss function achieved after 1000 time steps in Monte Carlo simulation studies, with summary statistics from 500 runs shown in Table 2. With the given sensing parameters in this example, the true minimum loss function value can be computed to be 0.0162. These simulations highlight that the median stochastic loss function value for nearly all variations that involve an exploratory phase prior to commencing direct stochastic optimization techniques (in this example, via the gradient-based control law) is *lower* than the studies without exploration of the response surface. In particular, the simple fractional factorial with only eight design points and a modest budget of 50 measurements total for the exploratory phase results in the minimized value, as highlighted in Table 2.

Table 2: Summary of loss function value statistics at the end of 1000 time steps in Monte Carlo simulation (500 runs) for different screening designs and configurations.

Design name	k	# DPs	med	max
truth (no noise)	n/a	n/a	0.0162	0.0162
no exploration	0	n/a	0.0208	0.0293
full factorial	50	16	0.0200	0.0294
frac factorial	25	8	0.0201	0.0294
frac factorial	**50**	**8**	**0.0196**	**0.0297**
frac factorial	100	8	0.0207	0.0293
frac factorial	200	8	0.0210	0.0295
D-optimal	50	25	0.0204	0.0293
D-optimal	100	50	0.0210	0.0296
D-optimal	150	100	0.0206	0.0294
NOLH	25	17	0.0204	0.0296
NOLH	50	17	0.0201	0.0294

From the visualization of the evolution of the two-sensor trajectories illustrated in Figure 2, the effect of conducting exploratory measurements in the direction of the gradient as estimated by the RSM-based procedure can be seen. For the same number of design points for a given screening design (e.g., fractional factorial with eight design points), an increasing measurement budget, k, leads to a longer exploratory trajectory.

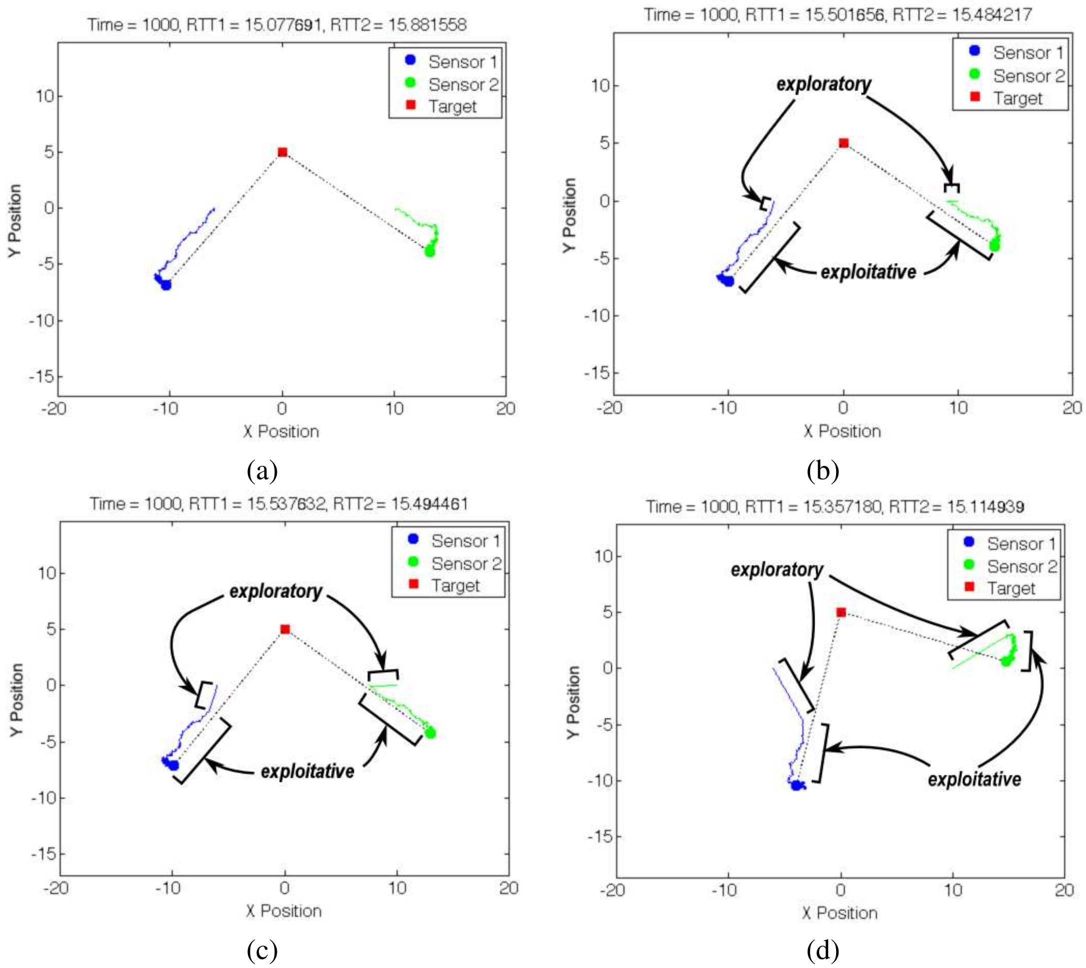

Figure 2: Evolution of the multi-sensor target tracking trajectories for (a) $k = 0$, (b) $k = 25$, (c) $k = 50$, and (d) $k = 100$, using a fractional factorial screening design (using a common pseudo-random number seed value for consistency). As annotated, the initial portions of the trajectories are in the direction of the estimated gradient from the screening measurements, followed by computation of the gradient-based motion control laws directly for the remainder of the tracking mission.

Further, by inspecting the initial evolution of the loss function values as computed by the two-sensor team tracking the target, the effect of the exploratory measurements can further be investigated, as shown in Figure 3. In particular, plotted against the ideal case of using the gradient-based control law given by Equation 2 without any measurement noise, measurements during the exploratory phase provide little improvement to the loss function value. However, as evidenced in Table 2, this initial "delay" in cost reduction still can lead to achieving a more favorable steady-state loss function value, as appears to be true for $k = 50$ in the fractional factorial case.

In addition to showing asymptotic approach to the optimal loss function value, an additional measure of performance of interest is the number of time steps (i.e., measurements) necessary to converge to within

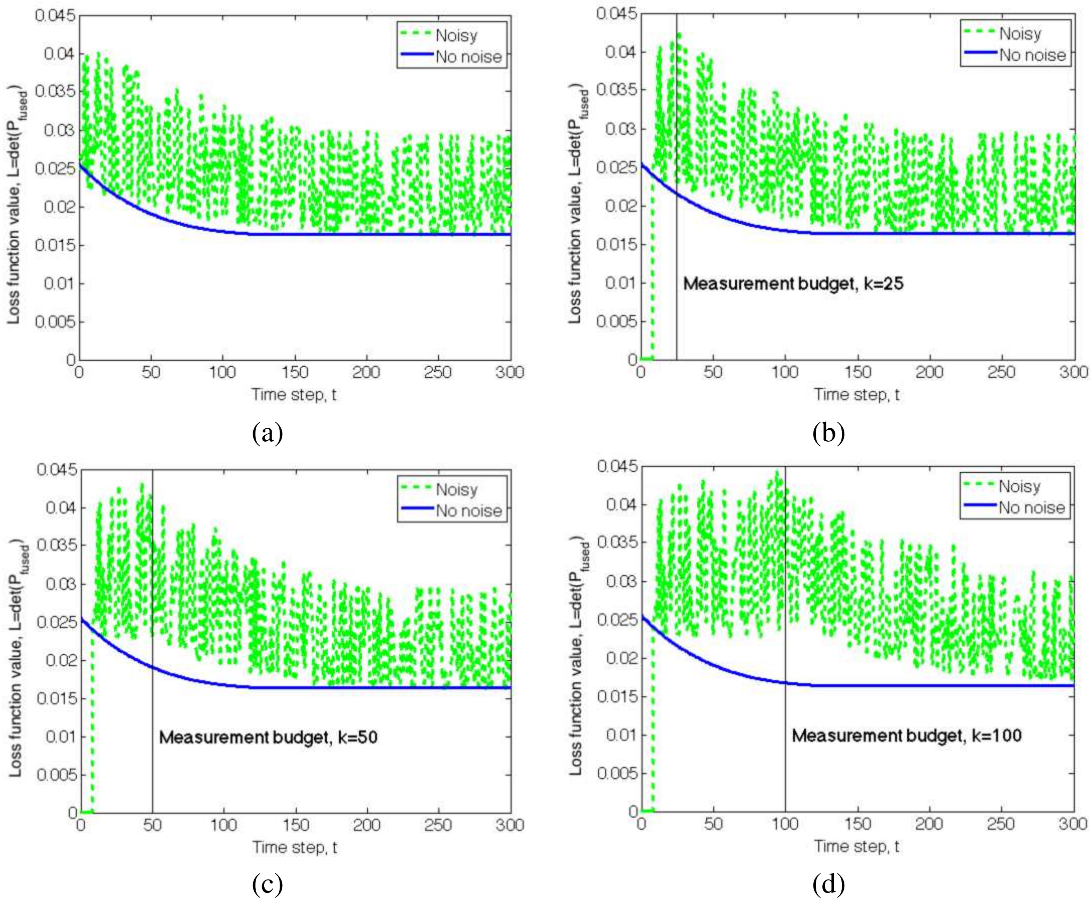

Figure 3: Cost evolution traces for varying exploratory measurement budgets, (a) $k = 0$, (b) $k = 25$, (c) $k = 50$, and (d) $k = 100$, using a fractional factorial screening design. For cases where the screening design is executed (b-d), the eight design points correspond to the first eight time steps, after which exploratory measurements in the direction of the estimated gradient extend for the remaining measurement budget.

a tolerance of the optimal solution. The objective is to identify any dependence on experimental design of the speed at which the sensors first arrive at their minimal cost configurations. We see that over 1000 Monte Carlo simulation runs, the summary statistics of each alternative's convergence time distributions are provided by Table 3, with the case of using a fractional factorial design with a budget of $k = 50$ measurements illustrated in Figure 4.

From these results, interestingly, the *average* time until the two-sensor team arrives as the minimal loss function value is generally *higher* when incorporating an exploratory phase to the integrated stochastic optimization problem. However, the *minimum* time to reach this minimum does depend on the choice of design and allotment of exploratory measurements; in some cases, it is possible to arrive at the minimum faster, even accounting for time spent performing the exploratory measurements, than if no exploration was conducted. This interesting relationship highlights the trade-off between exploration and exploitation precisely of interest in this study.

5 CONCLUSIONS AND FUTURE WORK

This paper proposes an approach that integrates multi-disciplinary methods for enhancing the performance and outcome of stochastic optimization procedures, showcasing an application of a state estimation task of an object of interest by a team of mobile sensor agents. The presented approach leverages fundamental

Table 3: Summary of loss function value statistics for convergence times to the minimal loss function value (0.0162) over 1000 Monte Carlo simulation runs for different screening designs and configurations.

Design name	***k***	**# DPs**	**min**	**mean**	**std**
truth (no noise)	n/a	n/a	127	127	n/a
no exploration	0	n/a	184	280	49
frac factorial	25	8	189	296	55
frac factorial	50	8	166	306	75
frac factorial	75	8	155	325	96
frac factorial	100	8	141	344	133
frac factorial	150	8	151	404	154
frac factorial	200	8	210	474	146
D-optimal	100	25	168	295	88
D-optimal	100	50	201	292	52

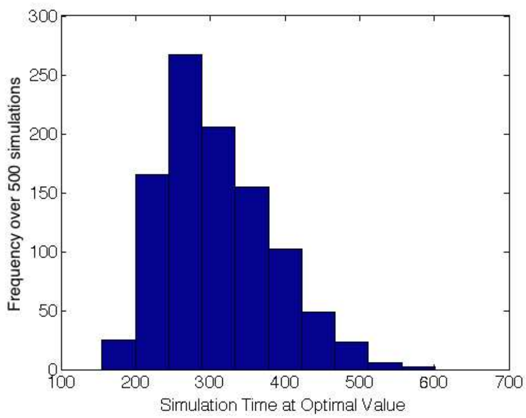

Figure 4: Frequency distribution of convergence times to within a given tolerance of the optimal loss function value, shown for the case of using a fractional factorial design possessing eight design points with a budget of $k = 50$ measurements (c.f. Table 3)

principles of statistical experimental design, namely, response surface methods and screening experimental designs to inform the trade-offs between expending measurements in support of "exploring" the stochastic loss function versus directly proceeding with the stochastic optimization in the presence of performance-degrading uncertainty. We show, through simulation studies and modeling of a stochastic optimization problem representing a multi-robot target tracking application, that employing an exploratory phase as the initial phase of the optimization process can lead to improved convergence to the optimal solution. Further, the modeled loss function representing target tracking uncertainty has a gradient that can be analytically derived, and by using it for rigorous baseline comparisons, a canonical fractional factorial design is found to be sufficient to approach the optimal minimal value.

Given the breadth of mathematical tools employed from different disciplines, there are many avenues of future research to expand upon this initial study. In particular, different classes of stochastic loss functions may benefit from different screening design methodologies. We investigate a number of the canonical designs; however, further investigation into the possible benefit from customized designs will highlight whether automated procedures are feasible. Also related is the issue of the step size used to define the neighborhood of the initial starting points of the sensors. Initial investigation of $\Delta x = \Delta y$ equal to 0.5 or 2.0 (vice a step size of 1.0 studied in this paper) indicates that there is little dependence on step size, but more extensive study is warranted.

Further, this work assumed that the initial locations of the sensor agents were dictated due to operational or other considerations; one can likely incorporate one of the many sensor deployment algorithms in the literature that may address specific requirements, potentially accounting for prior information regarding the object location within the search area. Sensor deployment models, such as in Martinez and Bullo (2006), may provide additional advantage in these integrated stochastic optimization scenarios.

Additionally, demonstration of the presented methods on more complex stochastic optimization functions, such as that presented in Peterson, Newman, and Spall (2014), would better highlight the strengths and limitations of the proposed approach.

REFERENCES

Back, T., D. B. Fogel, and Z. Michalewicz. 1997. *Handbook of Evolutionary Computation*. IOP Publishing Ltd.

Bazzan, A. L. C., and S. Labidi. (Eds.) 2004. *Advances in Artificial Intelligence – SBIA 2004*, Volume 3171 of *Lecture Notes in Computer Science*. Berlin, Heidelberg: Springer Berlin Heidelberg.

Bertsekas, D. P., and J. N. Tsitsiklis. 1989. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc.

Caiti, A., V. Calabro, F. Di Corato, D. Meucci, and A. Munafo. 2013. "Cooperative Distributed Algorithm for AUV Teams: A Minimum Entropy Approach". In *Proceedings of the IEEE OCEANS Conference*, 1–6: IEEE.

Chung, T. H., J. W. Burdick, and R. M. Murray. 2006. "A Decentralized Motion Coordination Strategy for Dynamic Target Tracking". In *Proc. of 2006 IEEE Int'l. Conf. on Robotics and Automation*, 2416–2422. Orlando, Florida.

Cortés, J., S. Martínez, T. Karatas, and F. Bullo. 2004. "Coverage Control for Mobile Sensing Networks". *IEEE Transactions on Robotics and Automation* 20 (2): 243–255.

Dames, P., M. Schwager, V. Kumar, and D. Rus. 2012. "A Decentralized Control Policy for Adaptive Information Gathering in Hazardous Environments". In *Proceedings of the IEEE Conference on Decision and Control*, 2807–2813.

Kleijnen, J., S. Sanchez, T. Lucas, and T. Cioppa. 2005. "State-of-the-Art Review: A User's Guide to the Brave New World of Designing Simulation Experiments". *INFORMS Journal on Computing* 17 (3): 263–289.

Martinez, S., and F. Bullo. 2006. "Optimal Sensor Placement and Motion Coordination for Target Tracking". *Automatica* 42 (4): 661–668.

Montgomery, D. 2009. *Design and Analysis of Experiments*. 7th ed. John Wiley and Sons Inc.

Myers, R. H., D. C. Montgomery, and C. M. Anderson-Cook. 2009. *Response Surface Methodology: Process and Product Optimization using Designed Experiments*, Volume 705. John Wiley & Sons.

Peterson, C. K., A. J. Newman, and J. C. Spall. 2014. "Simulation-based Examination of the Limits of Performance for Decentralized Multi-agent Surveillance and Tracking of Undersea Targets". In *Proceedings of SPIE 9091, Signal Processing, Sensor/Information Fusion, and Target Recognition XXIII*, 90910F. International Society for Optics and Photonics.

Sanchez, S.M. 2011. "NOLHdesigns spreadsheet". Available online via http://harvest.nps.edu. [Accessed 04/08/2015].

Schwager, M., P. Dames, D. Rus, and V. Kumar. 2011. "A Multi-robot Control Policy for Information Gathering in the Presence of Unknown Hazards". In *Proceedings of the International Symposium on Robotics Research*.

Spall, J. 2010. "Factorial Design for Efficient Experimentation". *IEEE Control Systems* 30 (5): 38–53.

Spall, J. C. 2003. *Introduction to Stochastic Search and Optimization*. John Wiley & Sons, Inc.

Ucinski, D. 2009. *Optimal Measurement Methods for Distributed Parameter System Identification*. First ed. CRC Press.

Van Laarhoven, P. J. M., and E. H. L. Aarts. 1987. *Simulated Annealing*. Springer.

Yang, P., R. A. Freeman, and K. M. Lynch. 2007. "A General Stability Condition for Multi-Agent Coordination by Coupled Estimation and Control". In *2007 American Control Conference*, 723–728: IEEE.

AUTHOR BIOGRAPHIES

TIMOTHY H. CHUNG is an Assistant Professor of Systems Engineering at the Naval Postgraduate School in Monterey, California. Dr. Chung also serves as the Deputy Director of the Secretary of the Navy initiative, the Consortium for Robotics and Unmanned Systems Education and Research (CRUSER). His research interests include modeling and analysis of operational settings involving unmanned systems, notably information gathering for probabilistic search and detection missions. Combining algorithm development with field experimentation, active research pursuits include systems design of large teams of cooperating and adversarial robots. He received his doctorate (2007) and M.S. (2002) at the California Institute of Technology in mechanical engineering and his B.S. (2001) in mechanical and aerospace engineering at Cornell University. His email address is thchung@nps.edu.

JAMES C. SPALL is a member of the Principal Professional Staff at the Johns Hopkins University, Applied Physics Laboratory (JHU/APL) and is a Research Professor in the Johns Hopkins Department of Applied Mathematics and Statistics. He is involved further in educational and leadership activities through serving as a Program Chair, instructor, and adviser in the Applied and Computational Mathematics Program within the Johns Hopkins University Engineering and Applied Science Programs for Professionals. Dr. Spall's technical and educational activities are focused on stochastic systems, statistical analysis, and computational algorithms in optimization and related areas. His email address is james.spall@jhuapl.edu.