Technometrics, ISSN 0040-1706
Volume 58, number 1 (february 2016)

Modeling an Augmented Lagrangian for Blackbox Constrained Optimization P. 1-11

Robert B. Gramacy, Genetha A. Gray, Sébastien Le Digabel, Herbert K. H. Lee, Pritam Ranjan, Garth Wells & Stefan M. Wild

Abstract

Constrained blackbox optimization is a difficult problem, with most approaches coming from the mathematical programming literature. The statistical literature is sparse, especially in addressing problems with nontrivial constraints. This situation is unfortunate because statistical methods have many attractive properties: global scope, handling noisy objectives, sensitivity analysis, and so forth. To narrow that gap, we propose a combination of response surface modeling, expected improvement, and the augmented Lagrangian numerical optimization framework. This hybrid approach allows the statistical model to think globally and the augmented Lagrangian to act locally. We focus on problems where the constraints are the primary bottleneck, requiring expensive simulation to evaluate and substantial modeling effort to map out. In that context, our hybridization presents a simple yet effective solution that allows existing objective-oriented statistical approaches, like those based on Gaussian process surrogates and expected improvement heuristics, to be applied to the constrained setting with minor modification. This work is motivated by a challenging, real-data benchmark problem from hydrology where, even with a simple linear objective function, learning a nontrivial valid region complicates the search for a global minimum.

Rejoinder P. 26-29

Robert B. Gramacy, Genetha A. Gray, Sébastien Le Digabel, Herbert K. H. Lee, Pritam Ranjan, Garth Wells & Stefan M. Wild

Abstract

We are grateful for the many insightful comments provided by the discussants. One team politely pointed out oversights in our literature review and the subsequent omission of a formidable comparator. Another made an important clarification about when a more aggressive variation (the so-called NoMax) would perform poorly. A third team offered enhancements to the framework, including a derivation of closed-form expressions and a more aggressive updating scheme; these enhancements were supported by an empirical study comparing new alternatives with old. The last team suggested hybridizing the statistical augmented Lagrangian (AL) method with modern stochastic search. Here we present our responses to these contributions and detail some improvements made to our own implementations in light of them. We conclude with some thoughts on statistical optimization using surrogate modeling and open-source software.

Sequential Design of Computer Experiments for the Assessment of Fetus Exposure to Electromagnetic Fields P. 30-42

Marjorie Jala, Céline Levy-Leduc, Éric Moulines, Emmanuelle Conil & Joe Wiart

Abstract

In this article, we describe four sequential sampling strategies for estimating the quantile of $Y = f(X)$, where X has a known distribution in and f is a deterministic unknown, expensive-to-evaluate real-valued function. These approaches all consist in modeling *f* as a sample of a well-chosen Gaussian process and aim at estimating the quantile by using as few evaluations of *f* as possible. The different methodologies are first compared through various numerical experiments. Then, in the framework of the ANR-JST FETUS project, we apply our strategies to a real example corresponding to the exposure of a Japanese pregnant-woman model and her 26-week-old fetus to a plane wave. Finally, we compare our methodologies on a simplified geometric model designed for modeling the fetus exposure to plane waves.

Abstract

Supersaturated designs (SSDs) are often used to reduce the number of experimental runs in screening experiments with a large number of factors. As more factors are used in the study, the search for an optimal SSD becomes increasingly challenging because of the large number of feasible selection of factor level settings. This article tackles this discrete optimization problem via an algorithm based on swarm intelligence. Using the commonly used $E(s^2)$ criterion as an illustrative example, we propose an algorithm to find $E(s^2)$-optimal SSDs by showing that they attain the theoretical lower bounds found in previous literature. We show that our algorithm consistently produces SSDs that are at least as efficient as those from the traditional CP exchange method in terms of computational effort, frequency of finding the $E(s^2)$-optimal SSD, and also has good potential for finding D_3-, D_4-, and D_5-optimal SSDs.

Abstract

We propose an approach for constructing a new type of design, called a sliced orthogonal array-based Latin hypercube design. This approach exploits a slicing structure of orthogonal arrays with strength two and makes use of sliced random permutations. Such a design achieves one- and two-dimensional uniformity and can be divided into smaller Latin hypercube designs with one-dimensional uniformity. Sampling properties of the proposed designs are derived. Examples are given for illustrating the construction method and corroborating the derived theoretical results. Potential applications of the constructed designs include uncertainty quantification of computer models, computer models with qualitative and quantitative factors, cross-validation and efficient allocation of computing resources.

Abstract

Unreplicated factorial designs have been widely used in scientific and industrial settings, when it is important to distinguish “active” or real factorial effects from “inactive” or noise factorial effects used to estimate residual or “error” terms. We propose a new approach to screen for active factorial effects from such experiments that uses the potential outcomes framework and is based on sequential posterior predictive model checks. One advantage of the proposed method is its ability to broaden the standard definition of active effects and to link their definition to the population of interest. Another important aspect of this approach is its conceptual connection to Fisherian randomization tests. Extensive simulation studies are conducted, which demonstrate the superiority of the proposed approach over existing ones in the situations considered.

Abstract
In earlier work, Jones and Nachtsheim proposed a new class of screening designs called definitive screening designs. As originally presented, these designs are three-level designs for quantitative factors that provide estimates of main effects that are unbiased by any second-order effect and require only one more than twice as many runs as there are factors. Definitive screening designs avoid direct confounding of any pair of second-order effects, and, for designs that have more than five factors, project to efficient response surface designs for any two or three factors. Recently, Jones and Nachtsheim expanded the applicability of these designs by showing how to include any number of two-level categorical factors. However, methods for blocking definitive screening designs have not been addressed. In this article we develop orthogonal blocking schemes for definitive screening designs. We separately consider the cases where all of the factors are quantitative and where there is a mix of quantitative and two-level qualitative factors. The schemes are quite flexible in that the numbers of blocks may vary from two to the number of factors, and block sizes need not be equal. We provide blocking schemes for both fixed and random blocks.

A Decomposition Strategy for the Variational Inference of Complex Systems P. 84-94
José M. Laínez-Aguirre, Linas Mockus, Seza Orçun, Gary Blau† & Gintaras V. Reklaitis

Abstract
Markov chain Monte Carlo approaches have been widely used for Bayesian inference. The drawback of these methods is that they can be computationally prohibitive especially when complex models are analyzed. In such cases, variational methods may provide an efficient and attractive alternative. However, the variational methods reported to date are applicable to relatively simple models and most are based on a factorized approximation to the posterior distribution. Here, we propose a variational approach that is capable of handling models that consist of a system of differential-algebraic equations and whose posterior approximation can be represented by a multivariate distribution. Under the proposed approach, the solution of the variational inference problem is decomposed into three steps: a maximum a posteriori optimization, which is facilitated by using an orthogonal collocation approach, a preprocessing step, which is based on the estimation of the eigenvectors of the posterior covariance matrix, and an expected propagation optimization problem. To tackle multivariate integration, we employ quadratures derived from the Smolyak rule (sparse grids). Examples are reported to elucidate the advantages and limitations of the proposed methodology. The results are compared to the solutions obtained from a Markov chain Monte Carlo approach. It is demonstrated that significant computational savings can be gained using the proposed approach.

Tail Estimation for Window-Censored Processes P. 95-103
Holger Rootzén & Dmitrii Zholud

Abstract
This article develops methods to estimate the tail and full distribution of the lengths of the 0-intervals in a continuous time stationary ergodic stochastic process that takes the values 0 and 1 in alternating intervals. The setting is that each of many such 0–1 processes has been observed during a short time window. Thus, the observed 0-intervals could be noncensored, right-censored, left-censored, or doubly-censored, and the lengths of 0-intervals that are ongoing at the beginning of the observation window have a length-biased distribution. We exhibit parametric conditional maximum likelihood estimators for the full distribution, develop maximum likelihood tail estimation methods based on a semiparametric generalized Pareto model, and propose goodness-of-fit plots. Finite sample properties are studied by simulation, and asymptotic normality is established for the most important case. The methods are applied to estimation of the length of off-road glances in the 100-car study, a big naturalistic driving experiment.

A Multivariate EWMA Controller for Linear Dynamic Processes P. 104-115
Sheng-Tsaing Tseng, Hsin-Chao Mi & I.- Chen Lee

Abstract
Most research of run-to-run process control has been based on single-input and single-output processes with static input–output relationships. In practice, many complicated semiconductor manufacturing processes have multiple-input and multiple-output (MIMO) variables. In addition, the effects of previous process input recipes and output responses on the current outputs might be carried over for several process periods. Under these circumstances, using conventional controllers usually results in unsatisfactory performance. To overcome this, a complicated process could be viewed as dynamic MIMO systems with added general process disturbance and this article proposes a dynamic-process multivariate exponentially weighted moving average (MEWMA) controller to adjust those processes. The long-term stability conditions of the proposed controller are derived analytically. Furthermore, by minimizing the total mean square error (TMSE) of the process outputs, the optimal discount matrix of the proposed controller under vector IMA(1,□1) disturbance is derived. Finally, to highlight the contribution of the proposed controller, we also conduct a comprehensive simulation study to compare the control performance of the proposed controller with that of the single MEWMA and self-tuning controllers. On average, the results demonstrate that the proposed controller outperforms the other two controllers with a TMSE reduction about 32% and 43%, respectively.

Abstract
To model the tail of a distribution, one has to define the threshold above or below which an extreme value model produces a suitable fit. Parameter stability plots, whereby one plots maximum likelihood estimates of supposedly threshold-independent parameters against threshold, form one of the main tools for threshold selection by practitioners, principally due to their simplicity. However, one repeated criticism of these plots is their lack of interpretability, with pointwise confidence intervals being strongly dependent across the range of thresholds. In this article, we exploit the independent-increments structure of maximum likelihood estimators to produce complementary plots with greater interpretability, and suggest a simple likelihood-based procedure that allows for automated threshold selection.

Abstract
Quantile functions are important in characterizing the entire probability distribution of a random variable, especially when the tail of a skewed distribution is of interest. This article introduces new quantile function estimators for spatial and temporal data with a fused adaptive Lasso penalty to accommodate the dependence in space and time. This method penalizes the difference among neighboring quantiles, hence it is desirable for applications with features ordered in time or space without replicated observations. The theoretical properties are investigated and the performances of the proposed methods are evaluated by simulations. The proposed method is applied to particulate matter (PM) data from the Community Multiscale Air Quality (CMAQ) model to characterize the upper quantiles, which are crucial for studying spatial association between PM concentrations and adverse human health effects.

Abstract

Turbine operations in a wind farm benefit from an understanding of the near-ground behavior of wind speeds. This article describes a probabilistic spatial-temporal model for analyzing local wind fields. Our model is constructed based on measurements taken from a large number of turbines in a wind farm, as opposed to aggregating the data into a single time-series. The model incorporates both temporal and spatial characteristics of wind speed data: in addition to using a time epoch mechanism to model temporal nonstationarity, our model identifies an informative neighborhood of turbines that are spatially related, and consequently, constructs an ensemble-like predictor using the data associated with the neighboring turbines. Using actual wind data measured at 200 wind turbines in a wind farm, we found that the two modeling elements benefit short-term wind speed forecasts. We also investigate the use of regime switching to account for the effect of wind direction and the use of geostrophic wind to account for the effects of meteorologic factors other than wind. These at best provide a small performance boost to speed forecast.

Technometrics, ISSN 0040-1706
Volume 58, number 1 (february 2016)