Executive Summary, Dagstuhl Seminar 09181 on Sampling-Based Optimization in the Presence of Uncertainty

Organizers:
Jürgen Branke[1], Barry Nelson[2], Warren Powell[3], and Tom Santner[4]

[1] University of Warwick, UK, `branke@aifb.uni-karlsruhe.de`
[2] Northwestern University, USA, `nelson@northwestern.edu`
[3] Princeton University, USA, `powell@princeton.edu`
[4] Ohio State University, USA, `santner.1@osu.edu`

1 Motivation

There are numerous industrial optimization problems in manufacturing, transportation and logistics, security, energy modeling, finance and insurance, and the sciences where decisions have to be evaluated by a process that generates a noisy result. The process might be a discrete-event simulation, a Monte Carlo evaluation of a complex function, or a physical experiment (e.g., how many cancer cells were killed by a particular compound?). There might be a small number of discrete decisions (the location of an emergency response facility, the design of a compound, or a set of labor work rules), or a large vector of decision variables (the allocation of a fleet of vehicles, choosing a set of research projects or allocating assets among investments). There are applications in virtually any area of business, government, science and engineering. Algorithms to support decisions in these diverse environments are urgently needed. This Dagstuhl seminar focused primarily on problems where this measurement is expensive (for example, some computer models can take a day or more for a single data point), in which case the number of samples that could possibly be generated is rather limited. When the goal is to efficiently identify an optimal (or at least a very good) solution, the search for good solutions, and the collection of information to guide the search, are tightly coupled. It is necessary to strike a balance between collecting information (exploration or global search) and making decisions that appear to be the best given what we know (exploitation or local search). This is particularly true when measurements are expensive (long simulations, field experiments). Because of its wide-ranging applications, sampling-based optimization has been addressed by different communities with different methods, and from slightly different perspectives. Currently, communities are largely tied to problem categories (e.g., finite vs. infinite number of alternatives; discrete vs. continuous decision variables; desired statement at termination). This Dagstuhl seminar brought together researchers from statistical ranking and selection; experimental design and response-surface modeling; stochastic programming; approximate dynamic programming; optimal learning; and the design and analysis of computer experiments with the goal of attaining a much better mutual understanding of the commonalities and differences of the various approaches to sampling-based optimization, and to take first steps toward an overarching theory, encompassing many of the topics above.

2 Seminar week

The seminar brought together 31 internationally renowned researchers from 11 countries. After an introductory session the seminar started with four **tutorials** on the various involved communities:

– Jack Kleijnen: Design and Analysis of Experiments: An Overview
– Steven Chick: Ranking and Selection Tutorial
– Barry L. Nelson: A Brief Introduction to Optimization via Simulation
– Warren Powell: Tutorial on optimal learning

Other planned events included two **feature talks** (Pierre LEcuyer and David Morton), 15 **regular talks**, and a **panel discussion** on 'Barriers to Application' (panelists Steve Chick, Genetha Gray, Tom Santner and Warren Powell).

A significant amount of time of the seminar was spent in **working groups**. Based on suggestions made by the participants, four working groups were formed to discuss some important and cutting-edge research questions in more detail:

1. Multiobjective optimization under uncertainty
2. Optimization with expensive function evaluations
3. Approximate dynamic programming/optimal learning
4. Cross-fertilization of experimental design, ranking & selection and optimization.

These working groups met each day for 1-2 hours, and presented their results to the general audience on the last day.

Besides the official programme, there were plenty of opportunities for informal discussions, e.g., during lunches, a short hike on Wednesday afternoon and a wine& cheese party on Wednesday evening.

Overall, the seminar was a great success and offered many possibilities for cooperation. It was generally agreed that such a workshop should be repeated in two years time.

3 Seminar programme

(There was no programme on Friday, as 1 May was a holiday).

Monday
 9:00 Opening and getting to know each other
10:45 Design and Analysis of Experiments: An Overview (Jack Kleijnen)
11:45 Optimization of MRI Sampling Trajectories by Bayesian Experimental Design
      (Matthias Seeger)
14:00 Tutorial on Ranking and Selection (Steve Chick)
15:00 Vote on working groups
15:45 A Brief Introduction to Optimization via Simulation (Barry L. Nelson)
16:45 Working group formation
17:00 First working group session

Tuesday
 9:00 Feature Talk: Simulation-Based Optimization for Staffing and Scheduling in Call Centers (Pierre l'Ecuyer)
 9:45 Some problems in need of Sampling-based Optimization Under Uncertainty Methods (Genetha Gray)
10:45 Tutorial on optimal learning (Warren Powell)
11:45 Approximate Dynamic Programming in the presence of rare events (Shie Mannor)
14:30 Sequential Parameter Optimization (Thomas Bartz-Beielstein)
15:00 Modelling and Optimization under Uncertainty in Vienna (Arnold Neumaier)
16:00 A Sequential Design for Approximating the Pareto Front Using an Expected Fitness Improvement Function (Tom Santner)
16:30 Working groups
20:00 Panel discussion: Barriers to application

Wednesday
 9:00 Feature Talk: Assessing Solution Quality in Stochastic Programs (David Morton)
 9:45 Robust Planning Using Approximate Linear Programming (Marek Petrik)
10:45 Stochastic Gradient Estimation: An Overview (Michael Fu)
11:15 Ranking and Selection of Many Alternatives Using Knowledge Gradients and Correlated Beliefs (Peter Frazier)
11:45 Probabilistic Interference for Fast Learning in Control (Marc Deisenroth)
13:30 Hiking trip
16:30 Working groups
20:00 Wine and Cheese party

Thursday
 9:00 Aggregation and the Sensitivity-Based Optimization (Xi-Ren Cao)
 9:30 Particle Learning and Optimization (Matt Teddy)
10:00 Sequential Convex Approximations to Joint Chance Constraint Programms (Liu Hong)
10:45 Working groups
14:30 Learning and Anticipation in Online Dynamic Optimization with Evolutionary Algorithms (Peter Bosman)
15:45 Presentation of Working Group Results
17:45 Closing of Workshop