

# Approximation Algorithms for Stochastic Optimization

Lisa Hellerstein<sup>\*1</sup>, Viswanath Nagarajan<sup>\*2</sup>, and Kevin Schewior<sup>\*3</sup>

1 NYU – New York, US. [lisa.hellerstein@nyu.edu](mailto:lisa.hellerstein@nyu.edu)

2 University of Michigan – Ann Arbor, US. [viswa@umich.edu](mailto:viswa@umich.edu)

3 Universität Köln, DE. [kschewior@gmail.com](mailto:kschewior@gmail.com)

---

## Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 25132 “Approximation Algorithms for Stochastic Optimization”. In this seminar, we gathered researchers from different areas interested in combinatorial optimization problems in which there is some stochasticity in the input. The focus was on approximation algorithms for computing adaptive or non-adaptive strategies to interact with this stochastic uncertainty as well as structural measures such as the adaptivity gap.

**Seminar** March 23–28, 2025 – <https://www.dagstuhl.de/25132>

**2012 ACM Subject Classification** Mathematics of computing → Approximation algorithms; Mathematics of computing → Probability and statistics

**Keywords and phrases** adaptivity, approximation algorithms, combinatorial optimization, stochastic optimization


**Digital Object Identifier** 10.4230/DagRep.15.3.159

## 1 Executive Summary

*Kevin Schewior (Universität Köln, DE)*

*Lisa Hellerstein (NYU – New York, US)*

*Viswanath Nagarajan (University of Michigan – Ann Arbor, US)*

**License**  Creative Commons BY 4.0 International license  
© Kevin Schewior, Lisa Hellerstein, and Viswanath Nagarajan

Combinatorial optimization is a classic field, whose results are applied in numerous domains, including logistics, telecommunication, production scheduling, and health care. Many of the problems arising in this field are computationally hard (often NP-hard) to solve exactly. Therefore, approximation algorithms, i.e., efficient algorithms with provable performance guarantees, have been extensively investigated.

An aspect that is very relevant in practice, but which is not well understood, is uncertainty in the input. Stochastic models for uncertainty, where there is some probabilistic information about the uncertain parameters, are arguably the most common approach for algorithms under uncertainty. The main question that this seminar addresses is: Can we approximate (or even compute exactly) the best *strategy* to interact with stochastic uncertainty?

Such a strategy may adapt when uncertainty gets resolved. An additional, structural, question is the question for the *adaptivity gap*, i.e., to what degree adaptivity helps in the objective function. Existing and envisioned tools comprise (but are not limited to) probabilistic approaches (concentration inequalities, martingales, etc.), linear or convex optimization, rounding techniques, and dynamic programming.

---

\* Editor / Organizer



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Approximation Algorithms for Stochastic Optimization, *Dagstuhl Reports*, Vol. 15, Issue 3, pp. 159–176

Editors: Lisa Hellerstein, Viswanath Nagarajan, and Kevin Schewior



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this Dagstuhl Seminar, we gathered several researchers from different communities (approximation algorithms, algorithmic game theory, operations research, online algorithms, learning theory) that have been interested in these or related questions from different angles. The goal was to develop new techniques, to identify new models and directions, and to solve some of the many open problems related to the above question.

The specific problems considered in this seminar were stochastic function-evaluation problems, stochastic probing and selection problems, stochastic scheduling, online stochastic problems, and related problems. We also set out to investigate aspects such as sample-complexity bounds, approximate solutions with low regret, and relaxing the common independence assumption.

### Organization of the seminar

We gathered 26 participants from different communities (approximation algorithms, algorithmic game theory, operations research, online algorithms, learning theory). To develop a common background and language, we scheduled four overview talks throughout the week:

- Sahil Singla. Towards Tackling Adaptivity and Correlations in Stochastic Optimization (one hour)
- Tonguç Ünlüyurt. A Review of the Sequential Testing Problem and its Extensions (30 minutes)
- Nicole Megow. Query Minimization for Stochastic Selection Problems (one hour)
- Thomas Kesselheim. Learning for Stochastic Optimization: Samples, Bandits, Contexts (one hour)

In addition, there were 14 talks lasting 30 minutes each, covering many different topics relevant to the workshop. There were two open-problem sessions, and we kept the late afternoons free for discussions. On the last day, we held a round-up session to discuss outcomes of the workshop. On one evening, we scheduled a joint session with the parallel seminar on Weihrauch complexity, to spark cross-disciplinary discussion.

### Outcomes

The workshop was very well received. In the survey, the participants praised the research theme, the composition of the workshop, the inspiring atmosphere, and the talks. The only negative aspect mentioned several times was that the collaboration time could have been scheduled in the early (rather than late) afternoon and could have been more structured. It was mentioned once that there could have been more participants from industry.

It seems that we have identified an area that researchers from different communities are interested in and that greatly benefits from a workshop like this. Judging from the discussions started at the workshop, we expect that the workshop will have a longer-term impact in terms of results, research proposals, community building, follow-up workshops, etc. Specific directions that were discussed several times at the workshop but still seem underexplored were correlated random variables and sample-complexity bounds.

## 2 Table of Contents

### Executive Summary

*Kevin Schewior, Lisa Hellerstein, and Viswanath Nagarajan* . . . . . 159

### Overview of Talks

Job selection and scheduling on unreliable machines <i>Alessandro Agnetis</i> . . . . .	163
Subsampling Suffices for Adaptive Data Analysis <i>Guy Blanc</i> . . . . .	163
Semi-Bandit Learning for Monotone Stochastic Optimization <i>Rohan Ghuge</i> . . . . .	164
Online and Stochastic Matching <i>Nathaniel Grammel</i> . . . . .	164
Unifying Pathwise and Expanding Search <i>Svenja M. Griesbach</i> . . . . .	165
Learning from a Sample in Online Algorithms <i>Anupam Gupta</i> . . . . .	165
Learning for Stochastic Optimization: Samples, Bandits, Contexts <i>Thomas Kesselheim</i> . . . . .	166
Search games with predictions <i>Thomas Lidbetter</i> . . . . .	166
Decomposing Probability Marginals Beyond Affine Requirements <i>Jannik Matuschke</i> . . . . .	167
Query Minimization for Stochastic Set Selection Problems <i>Nicole Megow</i> . . . . .	167
Testing whether a Partition Matroid has an Active Basis <i>Benedikt Plank</i> . . . . .	168
Towards Tackling Adaptivity and Correlations in Stochastic Optimization <i>Sahil Singla</i> . . . . .	168
Approximation Algorithms for Correlated Knapsack Orienteering <i>Chaitanya Swamy</i> . . . . .	169
Stochastic Scheduling of Bernoulli Jobs through Policy Stratification <i>Marc Uetz</i> . . . . .	169
Provably Accurate Shapley Value Estimation via Leverage Score Sampling <i>Teal Witter</i> . . . . .	170
Approximating Optimal Binary Search Trees under Uncertainty <i>Sorrachai Yingchareonthawornchai</i> . . . . .	170
Bayesian Probing on Graphs <i>Rudy Zhou</i> . . . . .	171
A review of the sequential testing problem and its extensions <i>Tonguç Ünüyurt</i> . . . . .	171

**Open problems**

Sequencing replicated jobs on parallel machines to maximize the probability of a full kit <i>Alessandro Agnetis</i> . . . . .	172
Hardness of Correlated Prophet Inequality <i>Andrés Cristi</i> . . . . .	172
Searching on a path with predictions <i>Thomas Lidbetter</i> . . . . .	173
Stochastic Combinatorial Optimization under a Budget Constraint in Expectation <i>Kevin Schewior</i> . . . . .	173
The deliberate idleness problem <i>Marc Uetz</i> . . . . .	173
Stochastic Bin Packing with Overflow <i>Rudy Zhou</i> . . . . .	174
<b>Participants</b> . . . . .	176

### 3 Overview of Talks

#### 3.1 Job selection and scheduling on unreliable machines

Alessandro Agnetis (University of Siena, IT)

**License** © Creative Commons BY 4.0 International license  
© Alessandro Agnetis

**Joint work of** Alessandro Agnetis, Leus Roel, Emmeline Perneel, Ilara Salvadori, Kevin Schewior

We address the following problem, denoted as the Unreliable Job Selection and Sequencing Problem (UJSSP). Given a set  $J$  of jobs, a subset  $S \subseteq J$  must be selected for processing on a single machine that is subject to failure. Each job  $j$  incurs a cost  $c_j$  if selected and yields a reward  $R_j$  upon successful completion. A job  $j$  is completed successfully only if the machine does not fail before or during its execution, with job-specific failure probabilities  $p_j$ . The objective is to determine an optimal subset and sequence of jobs to maximize the expected net profit. We review some known results for the case with  $c_j = 0$  (i.e., all jobs are selected), for both the single-machine and the parallel-machine cases [1, 2]. We establish the computational complexity of UJSSP, proving its NP-hardness when job costs are not identical. The relationship of UJSSP with other submodular selection problems is discussed [3, 4], showing that the special cases in which all jobs have the same cost ( $c_j = c$  for all  $j$ ) or, respectively, the same failure probability ( $p_j = p$  for all  $j$ ) can be solved in polynomial time, while the case in which all jobs have the same reward remains open.

##### References

- 1 Agnetis, A., Detti, P., Pranzo, M. and Sodhi, M.S., Sequencing unreliable jobs on parallel machines, *Journal of Scheduling*, 12(1), 45–54, 2009.
- 2 Agnetis, A., Lidbetter, T., List scheduling is 0.8531-approximate for scheduling unreliable jobs on  $m$  parallel machines, *Operations Research Letters*, 48, 405–409, 2020.
- 3 W. Stadje. Selecting jobs for scheduling on a machine subject to failure, *Discrete Applied Mathematics*, 63(3), 257–265, 1995.
- 4 Olszewski, W. and Vohra, R., Simultaneous selection, *Discrete Applied Mathematics*, 200, 161–169, 2016.

#### 3.2 Subsampling Suffices for Adaptive Data Analysis

Guy Blanc (Stanford University, US)

**License** © Creative Commons BY 4.0 International license  
© Guy Blanc

**Main reference** Guy Blanc: “Subsampling Suffices for Adaptive Data Analysis”, CoRR, Vol. abs/2302.08661, 2023.  
**URL** <https://doi.org/10.48550/ARXIV.2302.08661>

Ensuring that analyses performed on a dataset are representative of the entire population is one of the central problems in statistics. Most classical techniques assume that the dataset is independent of the analyst’s query and break down in the common setting where a dataset is reused for multiple, adaptively chosen, queries. This problem of *adaptive data analysis* was formalized in the seminal works of Dwork et al. (STOC, 2015) and Hardt and Ullman (FOCS, 2014).

We identify a remarkably simple set of assumptions under which the queries will continue to be representative even when chosen adaptively: The only requirements are that each query takes as input a random subsample and outputs few bits. This result shows that the

noise inherent in subsampling is sufficient to guarantee that query responses generalize. The simplicity of this subsampling-based framework allows it to model a variety of real-world scenarios not covered by prior work.

In addition to its simplicity, we demonstrate the utility of this framework by designing mechanisms for two foundational tasks, statistical queries and median finding. In particular, our mechanism for answering the broadly applicable class of statistical queries is both extremely simple and state of the art in many parameter regimes.

### 3.3 Semi-Bandit Learning for Monotone Stochastic Optimization

*Rohan Ghuge (Georgia Institute of Technology – Atlanta, US)*

**License** © Creative Commons BY 4.0 International license  
© Rohan Ghuge

**Joint work of** Arpit Agarwal, Rohan Ghuge, Viswanath Nagarajan  
**Main reference** Arpit Agarwal, Rohan Ghuge, Viswanath Nagarajan: “Semi-Bandit Learning for Monotone Stochastic Optimization”, in Proc. of the 65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024, pp. 1260–1274, IEEE, 2024.  
**URL** <https://doi.org/10.1109/FOCS61266.2024.00083>

Stochastic optimization is a widely used approach for optimization under uncertainty, where uncertain input parameters are modeled by random variables. Exact or approximation algorithms have been obtained for several fundamental problems in this area. However, a significant limitation of this approach is that it requires full knowledge of the underlying probability distributions. Can we still get good algorithms if these distributions are unknown, and the algorithm needs to learn them through repeated interactions? In this talk, I will discuss a generic online learning algorithm that obtains optimal regret bounds relative to the best algorithms (under known distributions) for a large class of “monotone” stochastic problems. This class includes fundamental problems like single-resource revenue management, Pandora’s box, and stochastic knapsack. Notably, our online algorithm works in a semi-bandit setting, where in each period, the algorithm only observes samples from the random variables that were actually probed.

### 3.4 Online and Stochastic Matching

*Nathaniel Grammel (University of Maryland – College Park, US)*

**License** © Creative Commons BY 4.0 International license  
© Nathaniel Grammel

**Joint work of** Nathaniel Grammel, Brian Brubach, Will Ma, Aravind Srinivasan  
**Main reference** Nathaniel Grammel, Brian Brubach, Will Ma, Aravind Srinivasan: “Follow Your Star: New Frameworks for Online Stochastic Matching with Known and Unknown Patience”, in Proc. of the The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event, Proceedings of Machine Learning Research, Vol. 130, pp. 2872–2880, PMLR, 2021.  
**URL** <http://proceedings.mlr.press/v130/grammel21a.html>

We consider variants of the classical graph matching problem that exhibit various forms of uncertainty. Online Bipartite Matching considers the problem of finding a maximum weight matching in a bipartite graph when the vertices of one side arrive one by one and, at each arrival, the algorithm must make an immediate and irrevocable decision about whether to match the vertex to one of its neighbors. Variants of the problem consider stochastic arrival models (e.g., arriving vertices are sampled IID from a known distribution).

Stochastic Matching (or Matching with Stochastic Edges) considers the standard graph matching problem in settings where the existence of each edge is (initially) unknown to the algorithm and only discovered after querying the edge. Variants of the problem consider the case where each vertex has a “patience”, indicating the maximum number of its incident edges that may be queried. Typically, this patience value is deterministic and known to the algorithm. Online Stochastic Matching combines both of these problems, so that vertices on one side of a bipartite graph arrive one by one, and in each step the algorithm may make multiple attempts to match the vertex until one succeeds, or until the patience is exhausted. We discuss results for some of these variants. A final variant considers the case of unknown (stochastic) patience, i.e. where the patience is drawn from some known distribution and is only learned by the algorithm once it becomes exhausted.

### 3.5 Unifying Pathwise and Expanding Search

*Svenja M. Griesbach (University of Chile – Santiago de Chile, CL)*

**License** © Creative Commons BY 4.0 International license

© Svenja M. Griesbach

**Joint work of** Svenja M. Griesbach, Felix Hommelsheim, Max Klimm, Kevin Schewior

We establish a framework unifying the pathwise search problem and the expanding search problem. We consider a graph  $G = (V, E)$  with non-negative vertex weights and a designated start vertex  $s$ . Furthermore, each edge  $e$  is equipped with a non-negative cost and a discount factor  $\alpha_e \in [0, 1]$  such that for the second and further traversals of this edge, its cost is multiplied by  $\alpha_e$ . For a path that starts in  $s$ , the latency of a vertex is the total cost of that path until the vertex is visited for the first time. The goal is to find a path that starts in  $s$  and visits all vertices with positive weight such that the weighted sum of the latencies of all vertices is minimized. If  $\alpha_e = 0$  for all  $e \in E$ , the problem corresponds to the expanding search problem, and if  $\alpha_e = 1$  for all  $e \in E$ , it corresponds to the pathwise search problem. We give a polynomial time algorithm that yields a constant approximation factor for all choices of  $\alpha \in [0, 1]^{|E|}$ . For  $\alpha = 0$  and  $\alpha = 1$ , our factor attains the same ratio as the so far best factors for expanding and pathwise search, respectively.

### 3.6 Learning from a Sample in Online Algorithms

*Anupam Gupta (New York University, US)*

**License** © Creative Commons BY 4.0 International license

© Anupam Gupta

**Joint work of** C. J. Argue, Alan M. Frieze, Anupam Gupta, Christopher Seiler


**Main reference** C. J. Argue, Alan M. Frieze, Anupam Gupta, Christopher Seiler: “Learning from a Sample in Online Algorithms”, in Proc. of the Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 – December 9, 2022, 2022.

**URL** [http://papers.nips.cc/paper\\_files/paper/2022/hash/5a093120ff4776b4f0dc452e3e3b6652-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/5a093120ff4776b4f0dc452e3e3b6652-Abstract-Conference.html)

While analyzing algorithms in the worst-case has long served us well, recent years have seen an exciting surge in analyzing algorithms in models that go beyond the worst case. We consider the classical problem of load-balancing, where jobs arrive online and must be assigned to collections of machines to minimize the maximum load. Can we get results better than the adversarial setting if we have a small sample of the upcoming data? The results in this talk are based on work with C.J. Argue, Alan Frieze, and Chris Seiler.

### 3.7 Learning for Stochastic Optimization: Samples, Bandits, Contexts

*Thomas Kesselheim (Universität Bonn, DE)*

**License**  Creative Commons BY 4.0 International license  
© Thomas Kesselheim

Commonly, in stochastic optimization, one assumes to know the probability distribution(s) the input is coming from. This is often motivated by the availability of historic data. In this talk, we survey different recent results formalizing this aspect and attempting to understand how one can learn the distribution well enough. For example, we might get a number of samples from the distribution. Can we still show that the optimal policy on the empirical distribution has a good performance? And how should we behave in repeated settings, where we get one sample from the distribution per day?

### 3.8 Search games with predictions

*Thomas Lidbetter (Rutgers University – Newark, US)*

**License**  Creative Commons BY 4.0 International license  
© Thomas Lidbetter

**Joint work of** Spyros Angelopoulos, Thomas Lidbetter, Konstantinos Panagiotou  
**Main reference** Spyros Angelopoulos, Thomas Lidbetter, Konstantinos Panagiotou: “Search Games with Predictions”, CoRR, Vol. abs/2401.01149, 2024.  
**URL** <https://doi.org/10.48550/ARXIV.2401.01149>

We introduce the study of search games between a mobile Searcher and an immobile Hider in a new setting in which the Searcher has some potentially erroneous information, i.e., a prediction on the Hider’s position. The objective is to establish tight tradeoffs between the consistency of a search strategy (i.e., its worst case expected payoff assuming the prediction is correct) and its robustness (i.e., the worst case expected payoff with no assumptions on the quality of the prediction). Our study is the first to address the full power of mixed (randomized) strategies; previous work focused only on deterministic strategies, or relied on stochastic assumptions that do not guarantee worst-case robustness in adversarial situations. We give Pareto-optimal strategies for three fundamental problems, namely searching in discrete locations, searching with stochastic overlook, and searching in the infinite line. As part of our contribution, we provide a novel framework for proving optimal tradeoffs in search games which is applicable, more broadly, to any two-person zero-sum games in learning-augmented settings.



### 3.9 Decomposing Probability Marginals Beyond Affine Requirements

Jannik Matuschke (KU Leuven, BE)

**License** © Creative Commons BY 4.0 International license  
© Jannik Matuschke

**Main reference** Jannik Matuschke: “Decomposing Probability Marginals Beyond Affine Requirements”, in Proc. of the Integer Programming and Combinatorial Optimization – 25th International Conference, IPCO 2024, Wrocław, Poland, July 3-5, 2024, Proceedings, Lecture Notes in Computer Science, Vol. 14679, pp. 309–322, Springer, 2024.

**URL** [https://doi.org/10.1007/978-3-031-59835-7\\_23](https://doi.org/10.1007/978-3-031-59835-7_23)

Consider the triplet  $(E, \mathcal{P}, \pi)$ , where  $E$  is a finite ground set,  $\mathcal{P} \subseteq 2^E$  is a collection of subsets of  $E$  and  $\pi : \mathcal{P} \rightarrow [0, 1]$  is a requirement function. Given a vector of marginals  $\rho \in [0, 1]^E$ , our goal is to find a distribution for a random subset  $S \subseteq E$  such that  $\Pr[e \in S] = \rho_e$  for all  $e \in E$  and  $\Pr[P \cap S \neq \emptyset] \geq \pi_P$  for all  $P \in \mathcal{P}$ , or to determine that no such distribution exists.

Generalizing results of Dahan, Amin, and Jaillet, we devise a generic decomposition algorithm that solves the above problem when provided with a suitable sequence of admissible support candidates (ASCs). We show how to construct such ASCs for numerous settings, including supermodular requirements, Hoffman-Schwartz-type lattice polyhedra, and abstract networks where  $\pi$  fulfils a conservation law. The resulting algorithm can be carried out efficiently when  $\mathcal{P}$  and  $\pi$  can be accessed via appropriate oracles. For any system allowing the construction of ASCs, our results imply a simple polyhedral description of the set of marginal vectors for which the decomposition problem is feasible. Finally, we characterize balanced hypergraphs as the systems  $(E, \mathcal{P})$  that allow the perfect decomposition of any marginal vector  $\rho \in [0, 1]^E$ , i.e., where we can always find a distribution reaching the highest attainable probability  $\Pr[P \cap S \neq \emptyset] = \min \sum_{e \in P} \rho_e, 1$  for all  $P \in \mathcal{P}$ .

### 3.10 Query Minimization for Stochastic Set Selection Problems

Nicole Megow (Universität Bremen, DE)

**License** © Creative Commons BY 4.0 International license  
© Nicole Megow

**Main reference** Evripidis Bampis, Christoph Dürr, Thomas Erlebach, Murilo Santos de Lima, Nicole Megow, Jens Schlöter: “Orienting (Hyper)graphs Under Explorable Stochastic Uncertainty”, in Proc. of the 29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference), LIPIcs, Vol. 204, pp. 10:1–10:18, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

**URL** <https://doi.org/10.4230/LIPICS.ESA.2021.10>

**Main reference** Nicole Megow, Jens Schlöter: “Set Selection Under Explorable Stochastic Uncertainty via Covering Techniques”, in Proc. of the Integer Programming and Combinatorial Optimization – 24th International Conference, IPCO 2023, Madison, WI, USA, June 21-23, 2023, Proceedings, Lecture Notes in Computer Science, Vol. 13904, pp. 319–333, Springer, 2023.

**URL** [https://doi.org/10.1007/978-3-031-32726-1\\_23](https://doi.org/10.1007/978-3-031-32726-1_23)

I will give an overview of basic set selection problems in a model where querying uncertain data incurs a cost. Given subsets of uncertain values, we study two tasks: identifying the minimum element in each set and selecting the subset of minimum total value while querying as few values as possible. These problems fall under the umbrella of explorable uncertainty. In the adversarial setting, strong lower bounds on query complexity extend to a wide range of classical problems such as knapsack, matchings, and linear programming. We then introduce a stochastic variant, where each weight is drawn independently from a known distribution, and present algorithms that, in expectation, beat these adversarial bounds. Our approach builds on a careful analysis of the underlying offline problems, exploiting connections to vertex covers and LP formulations. Finally, I will outline further research directions involving parallelization, robustification, and other extensions.

### 3.11 Testing whether a Partition Matroid has an Active Basis

*Benedikt Plank (Berlin, DE)*


License  Creative Commons BY 4.0 International license  
© Benedikt Plank

Joint work of Lisa Hellerstein, Benedikt Plank, Kevin Schewior

We consider the following Stochastic Boolean Function Evaluation problem, which is closely related to several problems from the literature. A matroid  $\mathcal{M}$  (in compact representation) on ground set  $E$  is given, and each element  $i \in E$  is active independently with known probability  $p_i \in (0, 1)$ . The elements can be queried, upon which it is revealed whether the respective element is active or not. The goal is to find an adaptive querying strategy for determining whether there is a basis of  $\mathcal{M}$  in which all elements are active, with the objective of minimizing the expected number of queries. When  $\mathcal{M}$  is a uniform matroid, this is the problem of evaluating a  $k$ -of- $n$  function, first studied in the 1970s. This problem is well-understood, and has an optimal adaptive strategy that can be computed in polynomial time. Interestingly, already when  $\mathcal{M}$  is a partition matroid, we show that the standard approaches fail to give even a constant-factor approximation. Our main result is a randomized polynomial-time constant-factor approximation algorithm for this problem. Our algorithm adaptively interleaves solutions to several instances of a novel type of stochastic querying problem, with a constraint on the expected cost. We believe that this problem is of independent interest and that several of our techniques have the potential for more general applications.

### 3.12 Towards Tackling Adaptivity and Correlations in Stochastic Optimization

*Sahil Singla (Georgia Institute of Technology – Atlanta, US)*

License  Creative Commons BY 4.0 International license  
© Sahil Singla

In this survey talk, we will consider discrete optimization problems where the inputs include probability distributions, and the goal is to maximize the expected reward. Two key benchmarks for these problems are the hindsight (offline) optimum and the optimal (online) policy. A central challenge, regardless of the chosen benchmark, is that optimal algorithms often adapt to the realizations of random variables. This adaptivity can lead to decision trees of exponential size, making the problem computationally intractable. We will explore techniques that focus on non-adaptive algorithms, which offer simpler and more efficient solutions, with only a small loss in performance. We will also discuss advances in stochastic discrete optimization models that incorporate correlations, while maintaining tractability in algorithm design.

### 3.13 Approximation Algorithms for Correlated Knapsack Orienteering

Chaitanya Swamy (University of Waterloo, CA)

**License** © Creative Commons BY 4.0 International license  
© Chaitanya Swamy

**Joint work of** Chaitanya Swamy, Davis Alemán Espinosa

**Main reference** David Alemán Espinosa, Chaitanya Swamy: “Approximation Algorithms for Correlated Knapsack Orienteering”, in Proc. of the Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2024, August 28-30, 2024, London School of Economics, London, UK, LIPIcs, Vol. 317, pp. 29:1–29:24, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

**URL** <https://doi.org/10.4230/LIPICS.APPROX/RANDOM.2024.29>

We consider the *correlated knapsack orienteering* (CSKO) problem: we are given a travel budget  $B$ , processing-time budget  $W$ , finite metric space  $(V, d)$  with root  $\rho \in V$ , where each vertex is associated with a job with possibly correlated random size and random reward that become known only when the job completes. Random variables are independent across different vertices. The goal is to compute a  $\rho$ -rooted path of length at most  $B$ , in a possibly adaptive fashion, that maximizes the reward collected from jobs that are processed by time  $W$ . To our knowledge, CSKO has not been considered before, though prior work has considered the uncorrelated problem, *stochastic knapsack orienteering*, and *correlated orienteering*, which features only one budget constraint on the *sum* of travel-time and processing-times.

We show that the *adaptivity gap* of CSKO is not a constant, and is at least  $\Omega(\max\{\sqrt{\log B}, \sqrt{\log \log W}\})$ . Complementing this, we devise *non-adaptive* algorithms that obtain: (a)  $O(\log \log W)$ -approximation in quasi-polytime; and (b)  $O(\log W)$ -approximation in polytime. We obtain similar guarantees for CSKO with cancellations, wherein a job can be cancelled before its completion time, foregoing its reward. We also consider the special case of CSKO, wherein job sizes are weighted Bernoulli distributions, and more generally where the distributions are supported on at most two points (2-CSKO). Although weighted Bernoulli distributions suffice to yield an  $\Omega(\sqrt{\log \log B})$  adaptivity-gap lower bound for (uncorrelated) *stochastic orienteering*, we show that they are easy instances for CSKO. We develop non-adaptive algorithms that achieve  $O(1)$ -approximation in polytime for weighted Bernoulli distributions, and in  $(n + \log B)^{O(\log W)}$ -time for the more general case of 2-CSKO.

### 3.14 Stochastic Scheduling of Bernoulli Jobs through Policy Stratification

Marc Uetz (University of Twente – Enschede, NL)

**License** © Creative Commons BY 4.0 International license  
© Marc Uetz

**Joint work of** Antonios Antoniadis, Ruben Hoeksma, Kevin Schewior, Marc Uetz

**Main reference** Antonios Antoniadis, Ruben Hoeksma, Kevin Schewior, Marc Uetz: “Stochastic scheduling with Bernoulli-type jobs through policy stratification”, CoRR, Vol. abs/2505.03349, 2025.

**URL** <https://doi.org/10.48550/ARXIV.2505.03349>

This talk addresses the problem of computing a scheduling policy that minimizes the total expected completion time of a set of  $N$  jobs with stochastic processing times on  $m$  parallel identical machines. When all processing times follow Bernoulli-type distributions, Gupta et al. (SODA '23) exhibited approximation algorithms with an approximation guarantee  $\tilde{O}(\sqrt{m})$ , where  $m$  is the number of machines and  $\tilde{O}(\cdot)$  suppresses polylogarithmic factors in  $N$ , improving upon an earlier  $O(m)$  approximation by Eberle et al. (OR Letters '19) for a special case. The present paper shows that, quite unexpectedly, the problem with

Bernoulli-type jobs admits a PTAS whenever the number of different job-size parameters is bounded by a constant. The result is based on a series of transformations of an optimal scheduling policy to a “stratified” policy that makes scheduling decisions at specific points in time only, while losing only a negligible factor in expected cost. An optimal stratified policy is computed using dynamic programming. Two technical issues are solved, namely (i) to ensure that, with at most a slight delay, the stratified policy has an information advantage over the optimal policy, allowing it to simulate its decisions, and (ii) to ensure that the delays do not accumulate, thus solving the trade-off between the complexity of the scheduling policy and its expected cost. Our results also imply a quasi-polynomial  $\tilde{O}(\log N)$ -approximation for the case with an arbitrary number of job sizes.

### 3.15 Provably Accurate Shapley Value Estimation via Leverage Score Sampling

*Teal Witter (New York University, US)*

**License** © Creative Commons BY 4.0 International license  
© Teal Witter

**Joint work of** Christopher Musco, R. Teal Witter  
**Main reference** Christopher Musco, R. Teal Witter: “Provably Accurate Shapley Value Estimation via Leverage Score Sampling”, in Proc. of the The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025, OpenReview.net, 2025.  
**URL** <https://openreview.net/forum?id=wg3rBImn3O>

Originally introduced in game theory, Shapley values have emerged as a central tool in explainable machine learning, where they are used to attribute model predictions to specific input features. However, computing Shapley values exactly is expensive: for a general model with  $n$  features,  $O(2^n)$  model evaluations are necessary. To address this issue, approximation algorithms are widely used. One of the most popular is the Kernel SHAP algorithm, which is model agnostic and remarkably effective in practice. However, to the best of our knowledge, Kernel SHAP has no strong non-asymptotic complexity guarantees. We address this issue by introducing Leverage SHAP, a light-weight modification of Kernel SHAP that provides provably accurate Shapley value estimates with just  $O(n \log n)$  model evaluations. Our approach takes advantage of a connection between Shapley value estimation and agnostic active learning by employing leverage score sampling, a powerful regression tool. Beyond theoretical guarantees, we show that Leverage SHAP consistently outperforms even the highly optimized implementation of Kernel SHAP available in the ubiquitous SHAP library [Lundberg & Lee, 2017].

### 3.16 Approximating Optimal Binary Search Trees under Uncertainty

*Sorrachai Yingchareonthawornchai (ETH Zürich, CH)*

**License** © Creative Commons BY 4.0 International license  
© Sorrachai Yingchareonthawornchai

**Joint work of** Parinya Chalermsook, Zahra Hadizadeh, Wanchote Jiamjitrak, Sorrachai Yingchareonthawornchai

Constructing an optimal binary search tree (BST) has long been a foundational problem in data structures. Given a fixed, known probability distribution over keys, Knuth’s seminal result (1971) provides an efficient method for computing the optimal static BST. Later, Mehlhorn (1975) showed that a near-optimal BST can be approximated in linear time.

We initiate the study of the robust optimization variant of the classical BST problem by considering settings where the underlying distribution is uncertain. Instead of a single known distribution, we are given  $k$  different distributions  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_k$ . The goal is to construct a single BST  $T$  that performs well across all of them – minimizing the worst-case expected access cost.

In this talk, I will give a simple 1.5-approximation algorithm for  $k = 2$ , which can be generalized to a  $0.804k$ -approximation for arbitrary  $k$ . This improvement is achieved by carefully combining two distinct 2-approximation algorithms, leveraging their strengths to refine the approximation ratio. I will also show a hardness result under the Unique Games Conjecture: when  $k$  is large, computing an optimal BST is NP-hard, even when each distribution has support of size at most two.

### 3.17 Bayesian Probing on Graphs

*Rudy Zhou (Carnegie Mellon University – Pittsburgh, US)*

License © Creative Commons BY 4.0 International license  
© Rudy Zhou

Joint work of Anupam Gupta, Benjamin Moseley, Rudy Zhou

We introduce a stochastic probing problem with correlated items, which we call Bayesian probing, where the correlations are modeled by an underlying graph. In our model, each vertex has a known probability of being active or inactive. Each item is an edge in the graph, and its distribution is the product of its two endpoints. The goal is to adaptively probe items/edges subject to a knapsack constraint to maximize the expected total reward obtained from all probed edges. This problem is a special case of stochastic knapsack with correlations across items and Bayesian active search problems considered in machine learning.

The distinguishing feature of Bayesian probing is that the probing an edge reveals the outcome of both of its endpoints, which induces a Bayesian update on the expected reward of all other incident edges. We design approximation algorithms computing policies that are either fully non-adaptive, or they make a single Bayesian update after using half of the knapsack budget. Our approximation ratios depend on natural graph parameters of the underlying correlation graph.

### 3.18 A review of the sequential testing problem and its extensions

*Tonguç Ünlüyurt (Sabanci University – Istanbul, TR)*


License © Creative Commons BY 4.0 International license  
© Tonguç Ünlüyurt

In the sequential testing problem, the goal is to evaluate a Boolean (or discrete) function with the minimum expected cost where the values of the variables can be learned by paying a cost. The variables take values independent of each other with known probabilities. The problem has been studied in different domains for various applications. In this talk, we concentrate on works that have been published in the last 20 years and provide a general review of the results that have been obtained for different special cases and extensions of the problem. We also provide insights to explore potential research areas for future research.

## 4 Open problems

### 4.1 Sequencing replicated jobs on parallel machines to maximize the probability of a full kit

*Alessandro Agnetis (University of Siena, IT)*

License  Creative Commons BY 4.0 International license  
© Alessandro Agnetis

Consider  $n$  job types,  $m$  machines and  $m$  copies of each job type. When processed by a machine, a job of type  $j$  is successfully carried out with probability  $\pi_j$ , while with probability  $(1 - \pi_j)$  it fails. In the latter case, the machine halts and cannot process all subsequently scheduled jobs. The problem is to decide how to sequence the  $n$  job copies on each of the  $m$  machines, in order to maximize the probability of having at least one “full kit” of jobs, i.e., at least one copy of each job successfully carried out.

For  $m = 2$ , the optimal solution is simply found by arbitrarily sequencing the  $n$  job copies on the first machine, and by sequencing them in the reverse order on the second machine. Another easily solvable case is when  $n = 2$ . In this case the optimal solution can be found in  $O(\log m)$ . These results can be found in [1].


As far as I know, the complexity of the problem is open for  $m \geq 3$ . Even the case  $m = n$ , or the case with identical probabilities do not seem obvious...

#### References

- 1 Agnetis, A., Benini, M., Detti, P., Hermans, B., Pranzo, M., Replication and sequencing of unreliable jobs on parallel machines, *Computers and Operations Research*, 139, 105634, 2022, doi:10.1016/j.cor.2021.105634.

### 4.2 Hardness of Correlated Prophet Inequality

*Andrés Cristi (EPFL – Lausanne, CH)*

License  Creative Commons BY 4.0 International license  
© Andrés Cristi

In the prophet inequality we observe a sequence of  $n$  random variables one by one and we want to decide when to stop. The reward of a online stopping policy is the variable it stops with. When the random variables are independent and the distributions are known, the optimal stopping policy can be computed in polynomial time via backward induction. A large body of literature is dedicated to bound the ratio between the expected reward of the optimal online policy and the expected offline optimum, under many variants of the problem. When the variables are correlated, no constant approximation with respect to the offline optimum is possible. Moreover, when the variables are correlated, the computation of the optimal online policy becomes challenging, because a naive backward induction algorithm would need to “remember” at each time all previous realizations. Here I propose two models to study the computational complexity of correlated optimal stopping, one easy and one apparently hard:

- (1) We are given a set of  $m$  scenarios. Each scenario is a deterministic sequence of  $n$  values. A realization of the  $n$  random variables is drawn as a uniformly chosen scenario. We can compute the optimal online algorithm via backward induction, by noticing that at any given time, the observed values must agree with at least one of the scenarios. Therefore, we can encode past observations by just choosing one scenario that is compatible with them. Thus we can compute the optimal online policy in time  $\text{poly}(n, m)$ .

(2) We are given a set of  $m$  scenarios, but now a scenario is a sequence of  $n$  distributions, all supported in a set of  $k$  different values. A realization of the  $n$  random variables is drawn by first choosing uniformly one scenario, and then drawing each variable independently according to the  $n$  distributions of the chosen scenario. Can we compute the optimal online policy in time  $\text{poly}(m, n, k)$ ? If not, can we approximate it?

### 4.3 Searching on a path with predictions


Thomas Lidbetter (*Rutgers University – Newark, US*)

License  Creative Commons BY 4.0 International license  
© Thomas Lidbetter

A target is located at one of the two ends of a discrete path according to a known probability distribution. A Searcher begins at a root node  $O$  at time 0. In each time step, she moves to an adjacent node. Upon reaching a node, she receives a signal that points either left or right. (We assume there is no signal at  $O$  at time 0.) With probability  $p \geq 1/2$ , the signal points towards the target, otherwise it points in the opposite direction. We wish to find a policy to minimize the expected time to reach the target. Attached is a solution for a path of length 3. Is there a closed form policy in general?

### 4.4 Stochastic Combinatorial Optimization under a Budget Constraint in Expectation


Kevin Schewior (*Universität Köln, DE*)

License  Creative Commons BY 4.0 International license  
© Kevin Schewior  
Joint work of Lisa Hellerstein, Benedikt Plank, Kevin Schewior

Motivated by Benedikt Plank's talk, I propose to relax stochastic combinatorial optimization problems in which there is a hard budget constraint by considering them under a budget constraint *in expectation*. One can then ask (i) how to efficiently compute or approximate the best strategy and (ii) by what factor the objective-function value changes in the worst case. Such problems may be interesting by itself, but Benedikt Plank's talk has shown another motivation. For the  $k$ -of- $n$  SFBE problem considered in that talk the answer to (ii) was non-constant. What else can we say, e.g., for the ProbeMax problem?

### 4.5 The deliberate idleness problem

Marc Uetz (*University of Twente – Enschede, NL*)

License  Creative Commons BY 4.0 International license  
© Marc Uetz

The *deliberate idleness problem* is a problem in stochastic machine scheduling. In stochastic machine scheduling, we are concerned with the question how to optimally schedule  $n$  jobs with stochastic processing requirements on  $m$  machines. More specifically, the processing times follows distributions  $p_j \sim X_j$ ,  $j = 1 \dots, n$ . The jobs are nonpreemptive, all available

at time 0, and have to be scheduled on  $m$  parallel, identical machines. Each machine can only do one job at a time, and a job can go on any of the  $m$  machines. Moreover, each job has a weight  $w_j$ , and we want to find a scheduling policy that minimizes the expected value of the weighted sum of completion times,  $E[\sum_j w_j C_j]$ . An instance consists of the input of jobs  $(w_j, X_j)$ ,  $j = 1, \dots, n$ , and the encoding of the number of machines  $m$ .


The solution to such a problem is not a schedule, but a scheduling *policy* which tells us, at any point in time  $t$  (typically when a machine falls idle, but possibly also at other points in time), which job(s) to schedule next. This decision may depend on the input of the problem, and the state of the system at time  $t$ . The latter is given by time  $t$ , the set of jobs already completed, the set of jobs currently running together with their conditional distribution of remaining processing time, and the set of jobs not yet started.

**Question:** Assume  $m \geq 3$  machines, and assume that all jobs follow an exponential distribution,  $p_j \sim \exp(\lambda_j)$ , that is, the processing times are memoryless. Does there always exist an optimal policy that avoids deliberate idleness? That is, as long as there are unprocessed jobs, it would never leave a machine deliberately idle. Some background information follows.

- For arbitrary  $p_j \sim X_j$  there are examples showing that deliberate idleness can be necessary, even on  $m = 2$  machines. See Uetz: When Greediness Fails: Examples from Stochastic Scheduling, OR Lett. 31, 2003, 413-419. (Franziska Eberle might have an example even when all  $w_j = 1$ ).
- For  $m = 2$  machines and  $p_j \sim \exp(\lambda_j)$  an optimal policy always exists that avoids deliberate idleness. This is not totally trivial, but not too difficult either, using an exchange argument (and induction).
- WSEPT (greedily schedule jobs in order of ratios  $w_j/Ep_j$ ) has a performance guarantee of  $2 - 1/m$ , whenever the distribution have coefficient of variation  $\leq 1$ , including exponential distributions. See Möhring, Schulz, Uetz: Approximation in stochastic scheduling: the power of LP-based priority policies, J. ACM 46 (1999), 924-942. For a more recent improvement of this upper bound to  $4/3$ , see Jäger, Skutella: Generalizing the Kawaguchi-Kyan bound to stochastic parallel machine scheduling, 35th STACS, LIPIcs no. 43, vol. 96, 2018
- When all  $w_j = 1$ , the problem is solved optimally by the SEPT rule, greedily schedule jobs with shortest expected processing time first. See Bruno, Downey, Frederickson: Sequencing Tasks with Exponential Service Times to Minimize the Expected Flow Time or Makespan, J. ACM 28 (1981), 100-113.

## 4.6 Stochastic Bin Packing with Overflow

Rudy Zhou (Carnegie Mellon University – Pittsburgh, US)

**License**  Creative Commons BY 4.0 International license  
© Rudy Zhou

**Joint work of** Sebastian Perez-Salazar, Mohit Singh, Alejandro Toriello, Rudy Zhou

**Main reference** Sebastian Perez-Salazar, Mohit Singh, Alejandro Toriello: “Adaptive Bin Packing with Overflow”, Math. Oper. Res., Vol. 47(4), pp. 3317–3356, 2022.

**URL** <https://doi.org/10.1287/MOOR.2021.1239>

In the stochastic bin packing problem, we must adaptively pack items with random sizes into unit-sized bins. Each item size is independent with known distribution but unknown realized value. Because of this, we may place an item into a bin such that its realized size



overflows the bin capacity. Overflowing a bin incurs an additive penalty. Thus, our objective is to pack all items to minimize the expected number of bins opened plus the penalties for any overflowed bins.

This problem was introduced by Sebastian Perez-Salazar, Mohit Singh, and Alejandro Toriello (<https://arxiv.org/abs/2007.11532>). Among other things, they prove that in the online setting (items arrive one-by-one and must be packed irrevocably upon arrival), one can achieve a  $O(1)$ -approximation if all items are drawn i.i.d. from a known distribution and a  $O(\log C)$  approximation if all items are exponentially distributed, where  $C$  is a function of the parameters of the exponentials.

It would be interesting to give approximation algorithms for more general distributions or consider the case where we do not know the item size distributions.

## Participants

- Alessandro Agnetis  
University of Siena, IT
- Guy Blanc  
Stanford University, US
- Andrés Cristi  
EPFL – Lausanne, CH
- Franziska Eberle  
TU Berlin, DE
- Rohan Ghuge  
Georgia Institute of Technology – Atlanta, US
- Nathaniel Grammel  
University of Maryland – College Park, US
- Svenja M. Griesbach  
University of Chile – Santiago de Chile, CL
- Anupam Gupta  
New York University, US
- Lisa Hellerstein  
NYU – New York, US
- Thomas Kesselheim  
Universität Bonn, DE
- Rebecca Lehming  
Universität Bonn, DE
- Thomas Lidbetter  
Rutgers University – Newark, US
- Naifeng Liu  
Universität Mannheim, DE
- Jannik Matuschke  
KU Leuven, BE
- Nicole Megow  
Universität Bremen, DE
- Viswanath Nagarajan  
University of Michigan – Ann Arbor, US
- Benedikt Plank  
Berlin, DE
- Kevin Schewior  
Universität Köln, DE
- Daniel Schmand  
Universität Bremen, DE
- Sahil Singla  
Georgia Institute of Technology – Atlanta, US
- Chaitanya Swamy  
University of Waterloo, CA
- Tonguç Ünliüyt  
Sabanci University – Istanbul, TR
- Marc Uetz  
University of Twente – Enschede, NL
- Teal Witter  
New York University, US
- Sorrachai Yingchareonthawornchai  
ETH Zürich, CH
- Rudy Zhou  
Carnegie Mellon University – Pittsburgh, US

