Report from Dagstuhl Seminar 25471

Online Algorithms beyond Competitive Analysis

Sungjin Im*[1], Nicole Megow*[2], Debmalya Panigrahi*[3], Sahil Singla*[4], and Golnoosh Shahkarami†[5]

1 **University of California – Santa Cruz, US.** sim9@ucsc.edu
2 **Universität Bremen, DE.** nicole.megow@uni-bremen.de
3 **Duke University – Durham, US.** debmalya.panigrahi@gmail.com
4 **Georgia Institute of Technology – Atlanta, US.** ssingla@gatech.edu
5 **MPI für Informatik – Saarbrücken, DE & Univerisität Bremen, DE.** golnooshshah@gmail.com

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 25471: "Online Algorithms beyond Competitive Analysis". This seminar brought 45 leading researchers in online algorithms and related areas together for a discussion of the recent progress and challenges in various aspects of online algorithms. The seminar included several talks of varying lengths and an open problem session. In addition, sufficient time was set aside for research discussions and collaborations.

Seminar November 16–21, 2025 – https://www.dagstuhl.de/25471
2012 ACM Subject Classification Theory of computation → Online algorithms
Keywords and phrases algorithms under uncertainty, beyond worst-case algorithm design, online algorithms
Digital Object Identifier 10.4230/DagRep.15.11.114

1 Executive Summary

Sungjin Im (University of California – Santa Cruz, US)
Nicole Megow (Universität Bremen, DE)
Debmalya Panigrahi (Duke University – Durham, US)
Sahil Singla (Georgia Institute of Technology – Atlanta, US)

Traditionally, algorithms are studied in a setting where the entire input is known in advance and the goal is to compute an optimal solution with respect to a given objective. In many applications, however, this paradigm is inadequate: inputs arrive over time and decisions must be made as information is revealed. This has given rise to the field of online algorithms, in which input revelation and algorithmic decisions are interleaved, creating a distinctive mix of computational and information constraints. A central theoretical tool for analyzing such algorithms is *competitive analysis*, which extends worst-case reasoning to the online setting by comparing the performance of an online algorithm to an offline optimum that has full knowledge of the complete input sequence. Over the past decades, this framework has generated a deep and influential body of results, shaped a set of core problems, and established online algorithms as a central area in algorithm design.

* Editor / Organizer
† Editorial Assistant / Collector

Online Algorithms beyond Competitive Analysis, *Dagstuhl Reports*, Vol. 15, Issue 11, pp. 114–133
Editors: Sungjin Im, Nicole Megow, Debmalya Panigrahi, Sahil Singla, and Golnoosh Shahkarami
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

At the same time, it has long been recognized that competitive analysis, despite its elegance and conceptual clarity, can be too pessimistic in practice, and can fail to separate algorithms whose empirical performance differs substantially. A canonical illustration is online caching: many deterministic policies share the same worst-case competitive ratio, yet strategies such as least-recently-used can outperform others by large margins on realistic workloads. Motivated by such gaps between worst-case guarantees and observed behavior, the community has increasingly developed and studied beyond-worst-case models for online computation. These approaches aim to preserve rigorous guarantees while incorporating additional structure or flexibility that better reflects practical environments, thereby enabling finer-grained comparisons between algorithms that standard competitive analysis treats as essentially equivalent.

The Dagstuhl Seminar "Online Algorithms beyond Competitive Analysis" (25471) was organized in the context of this broader shift. It brought together researchers working on beyond-worst-case perspectives in online algorithms, to exchange viewpoints across sub-communities, and to identify promising directions and shared challenges. Discussions were structured around three representative themes:

- **Online algorithms with predictions** (learning-augmented algorithms): models in which an online algorithm is given auxiliary information about the future (e.g., forecasts derived from past data) and is designed to benefit from accurate predictions while remaining robust when predictions are noisy or unreliable.
- **Online and dynamic algorithms with recourse**: frameworks that relax strict irrevocability by allowing limited revisions to past decisions, capturing settings where changes are possible but costly, and studying how bounded recourse improves achievable guarantees and enables dynamic variants where requirements may both arrive and depart.
- **Online stochastic input models**: models that assume stochastic structure in arrivals (such as random order or distributional assumptions) to circumvent worst-case barriers, together with emerging efforts to make such models robust to deviations such as outliers or mild correlations.

Organization of the Seminar

The seminar brought together 45 researchers working on online algorithms. The participants embraced both senior and junior researchers, including a number of postdocs and advanced PhD students. During the five days of the seminar, 28 talks of different lengths took place. Six keynote speakers gave an overview of the state-of-the art of the respective area or presented recent highlight results in 60 minutes:

- Christian Coester: A Mirror Descent Approach for Online Algorithms
- Shuchi Chawla: Pandora's Box
- Jose Correa: Prophet Inequalities
- Anupam Gupta: Online Algorithms with Samples
- Ziv Scully: Survey on Queueing Theoretic Scheduling
- David Wajc: Philosopher inequalities: online Bayesian selection beyond competitive analysis

The remaining slots were filled with shorter talks of 30 minutes and a few 8-minute spotlight talks.

Outcomes

Organizers and participants regard the seminar as a great success. The seminar achieved the goal to bring together researchers from different sub-communities addressing beyond-worst-case analyses, share the state-of-the art research and discuss the current major challenges. The talks were excellent and stimulating; participants connected and actively met in working groups in the afternoon and evenings.

The organizers wish to express their gratitude towards the Scientific Directorate and the administration of the Dagstuhl Center for their great support for this seminar.

2 Table of Contents

3 Overview of Talks

3.1 Randomization and Stochasticity in Learning-Augmented Algorithms

Spyros Angelopoulos (International Lab on Learning Systems – Montreal, CA)

Randomization and stochasticity introduce new challenges while also enabling meaningful extensions in the design and analysis of learning-augmented algorithms. This presentation will discuss such challenges, with a focus on three main themes: (i) the characterization of Pareto-optimal randomized strategies, illustrated through the framework of search games; (ii) the use of distributional prediction oracles, where the nature of the prediction is probabilistic; and (iii) risk-based analysis, which incorporates notions of uncertainty and risk into the design and evaluation of learning-augmented algorithms.

3.2 Competitive Bundle Trading

Yossi Azar (Tel Aviv University, IL)

Joint work of Yossi Azar, Niv Buchbinder, Roie Levin, Or Vardi
Main reference Yossi Azar, Niv Buchbinder, Roie Levin, Or Vardi: "Competitive Bundle Trading", CoRR, Vol. abs/2507.23047, 2025.
URL https://arxiv.org/abs/2507.23047

Allocating or selling a set of resources to a sequence of arriving customers has been extensively studied under a wide range of models and objectives. However, a natural extension arises when the algorithm is also allowed to purchase additional resources from suppliers. This setting arguably captures the fundamental scenario faced by any trader or retailer. A natural objective in this context is to maximize the cumulative profit over the online sequence. That is, the difference between revenue from sales and expenses from purchases. This problem of maximizing the difference between two competing quantities appears to be significantly more challenging than the sell-only case, which may explain why it has remained largely unexplored.

We initiate the study of such a trading problem and design an algorithm that achieves a logarithmic competitive ratio relative to the optimal offline solution. Our approach employs an exponential-weight–update dynamic pricing scheme, and our analysis uses a dual-fitting argument that relates the algorithm's profit to a linear programming relaxation that upper bounds the optimal offline profit. We also prove (nearly) matching lower bounds. Finally, we extend our results by designing an incentive-compatible mechanism for the setting in which customers are strategic and may misreport their true valuations.

3.3 The Power of Greedy for Online Minimum Cost Matching on the Line

Eric Balkanski (Columbia University – New York, US)

Joint work of Eric Balkanski, Yuri Faenza, Noémie Périvier
Main reference Eric Balkanski, Yuri Faenza, Noémie Périvier: "The Power of Greedy for Online Minimum Cost Matching on the Line", in Proc. of the 24th ACM Conference on Economics and Computation, EC '23, p. 185–205, Association for Computing Machinery, 2023.
URL https://doi.org/10.1145/3580507.3597794

We consider the online minimum cost matching problem on the line, in which there are n servers and, at each of n time steps, a request arrives and must be irrevocably matched to a server that has not yet been matched to, with the goal of minimizing the sum of the distances between the matched pairs. Despite achieving a worst-case competitive ratio that is exponential in n, the simple greedy algorithm, which matches each request to its nearest available free server, performs very well in practice. A major question is thus to explain greedy's strong empirical performance.

In this talk, we aim to understand the performance of greedy over instances that are at least partially random. When both the requests and the servers are drawn uniformly and independently from $[0, 1]$, we show that greedy is constant competitive and extend this constant competitive ratio to a setting with a linear excess of servers. Moreover, we show that in the semi-random model where requests are still drawn uniformly and independently but servers are chosen adversarially, greedy achieves an $O(\log n)$ competitive ratio. When the requests arrive in a random order but are chosen adversarially, it was previously known that greedy is $O(n)$-competitive. Even though this one-sided randomness allows a large improvement in greedy's competitive ratio compared to the model where requests are adversarial and arrive in a random order, we show that it is not sufficient to obtain a constant competitive ratio by giving a tight $\Omega(\log n)$ lower bound. These results invite further investigation about how much randomness is necessary and sufficient to obtain strong theoretical guarantees for the greedy algorithm for online minimum cost matching, on the line and beyond.

3.4 Machine Learning for Online Resource Allocation

Benjamin Moseley (Carnegie Mellon University – Pittsburgh, US)

Joint work of Benjamin Moseley, Andrew Li, Lin An
Main reference Lin An, Andrew A. Li, Benjamin Moseley, Gabriel Visotsky: "Best of Many in Both Worlds: Online Resource Allocation with Predictions under Unknown Arrival Model", CoRR, Vol. abs/2402.13530, 2024.
URL https://doi.org/10.48550/ARXIV.2402.13530

This talk will discuss a model for augmenting discrete optimization algorithms with machine learned predictions to improve algorithm performance. First the talk will discuss online resource allocation using predictions. Predictions will be leveraged to go beyond the worst case lower bounds in both the regret and competitive analysis frameworks.

3.5 Competitive Recourse (Beyond Absolute Recourse)

Niv Buchbinder (Tel Aviv University, IL)

Joint work of Niv Buchbinder, Sayan Bhattacharya, Roie Levin, Thatchaphol Saranurak, Seffi Naor, David Wajc

In this talk, I will explore the concepts of absolute and competitive recourse in the analysis of low-recourse algorithms. I will argue that competitive recourse provides a more meaningful framework for evaluating such algorithms. Finally, I will survey several recent results on competitive recourse and highlight a number of intriguing open problems.

3.6 Costly Information Combinatorial Selection

Shuchi Chawla (University of Texas – Austin, US)

Joint work of Dimitris Christou, Trung Dang, Amit Harlev, Ziv Scully
Main reference Shuchi Chawla, Dimitrios Christou, Amit Harlev, Ziv Scully: "Combinatorial Selection with Costly Information", in Proc. of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2026, Vancouver, BC, Canada, January 11-14, 2026, pp. 664–717, SIAM, 2026.
URL https://doi.org/10.1137/1.9781611978971.26
Main reference Shuchi Chawla, Dimitris Christou, Trung Dang: "Commitment Gap via Correlation Gap", CoRR, Vol. abs/2508.20246, 2025.
URL https://doi.org/10.48550/ARXIV.2508.20246

Imagine you are hiring a graduate student from a large pool of candidates. How much effort should you spend reading recommendation letters? Should you interview some candidates, or none at all? Many decision-making problems involve this same dilemma: investing in information acquisition upfront can improve outcomes, but at a cost. Balancing this trade-off requires carefully weighing whether the value of additional signals justifies the resources spent to obtain them.

We capture this interplay through a model we call Costly Information Combinatorial Selection (CICS). In CICS, the algorithm is given a ground set of elements whose values are drawn independently from known distributions. The realized values can be revealed incrementally through costly probing actions, and the objective is to select a feasible subset maximizing total value minus information costs. In this talk, I will introduce the model and describe two approaches for designing approximation algorithms. Along the way, we will see how CICS connects to classical stochastic selection problems, including Pandora's Box and prophet inequalities.

3.7 Competitive Analysis via Mirror Descent

Christian Coester (University of Oxford, GB)

Mirror descent is a general technique originating from convex optimization and used extensively in online learning, which in recent years also led to the successful discovery of optimal competitive ratios for various fundamental problems in online algorithms. This talk surveys how the technique can be used to obtain competitive online algorithms.

3.8 Prophet Inequalities: Old and New

José R. Correa (University of Chile – Santiago, CL)

In this talk, I will survey some of the most widely used techniques for designing and analyzing online algorithms under stochastic input. The playground is that of the classic prophet inequality. This states that, when faced with a finite sequence of non-negative independent random variables, a gambler who knows their distribution and is allowed to stop the sequence at any time, can obtain, in expectation, at least half as much reward as a prophet who knows the values of each random variable and can choose the largest one.

We begin by reviewing the classic "balanced prices" proof of this result, due to Samuel Cahn. We then introduce online contention resolution schemes of Feldman, Svensson, and Zenklusen, and discuss how to adapt them to the prophet inequality setting. To wrap up, we dive into the data-driven approach recently introduced by Rubinstein, Wang, and Weinberg. In all cases, we provide additional references to the literature and discuss some of the more recent research and open problems. In the final part of the talk, we introduce a new technique to tackle prophet inequalities based on Ordinary Differential Equations and show how this approach provides a unified view of the problem and all known algorithms for it. This last part is based on joint work with Fichtl, Jlibene, Laraki, Livanos, Schewior, and Verdugo.

3.9 On the Informativeness of Moments in Optimal Stopping

Andrés Cristi (EPFL – Lausanne, CH)

Joint work of Andrés Cristi, José Correa, Vasilis Livanos, Víctor Verdugo, Jiechen Zhang

We study a variant of the prophet inequality with limited information, where the decision maker only has access to the first k moments of each random variable, rather than their full distributions. In this work, it is proved that even with full moment knowledge (i.e., $k = \infty$), the best possible competitive ratio is $\Theta(1/\log n)$, which we show can already be achieved with only knowledge of the first moment. While the lower bound is simple and is attained by a standard bucketing algorithm, the upper bound requires a subtle construction. This involves using Vandermonde matrices first to construct a parametrized family of distributions for which the first k moments coincide, and for which the expected maximum of n such copies varies widely across different parameter choices. Then, we use Prokhorov's Theorem to prove the existence of limit distributions, which we show have all their moments equal. Finally, we show how an adversary can select *equally looking* instances combining these distributions, making it impossible for the decision maker to obtain a factor better than $O(\log n)$ of the expected maximum. Our result, in particular, implies that to obtain improved prophet inequalities, further assumptions, beyond moment knowledge, are needed. To showcase this direction, we establish improved bounds under additional distributional assumptions, such as MHR, bounded coefficient of variation, or bounded support.

3.10 The Cost of Consistency: Submodular Maximization with Constant Recourse

Paul Dütting (Google – Zürich, CH)

Joint work of Paul Dütting, Federico Fusco, Silvio Lattanzi, Ashkan Norouzi-Fard, Ola Svensson, Morteza Zadimoghaddam
Main reference Paul Dütting, Federico Fusco, Silvio Lattanzi, Ashkan Norouzi-Fard, Ola Svensson, Morteza Zadimoghaddam: "The Cost of Consistency: Submodular Maximization with Constant Recourse", in Proc. of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025, pp. 1406–1417, ACM, 2025.
URL https://doi.org/10.1145/3717823.3718131

In this work, we study online submodular maximization, and how the requirement of maintaining a stable solution impacts the approximation. In particular, we seek bounds on the best-possible approximation ratio that is attainable when the algorithm is allowed to make at most a constant number of updates per step. We show a tight information-theoretic bound of $2/3$ for general monotone submodular functions, and an improved (also tight) bound of $3/4$ for coverage functions. Since both these bounds are attained by non poly-time algorithms, we also give a poly-time randomized algorithm that achieves a 0.51-approximation. Combined with an information-theoretic hardness of $1/2$ for deterministic algorithms from prior work, our work thus shows a separation between deterministic and randomized algorithms, both information theoretically and for poly-time algorithms.

3.11 $O(1/\epsilon)$ is the answer in online weighted throughput maximization

Franziska Eberle (TU Berlin, DE)

Main reference Franziska Eberle: "$O(1/\epsilon)$ Is the Answer in Online Weighted Throughput Maximization", in Proc. of the 41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024), Leibniz International Proceedings in Informatics (LIPIcs), Vol. 289, pp. 32:1–32:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.
URL https://doi.org/10.4230/LIPIcs.STACS.2024.32

We study a fundamental online scheduling problem where jobs with processing times, weights, and deadlines arrive online over time at their release dates. The task is to preemptively schedule these jobs on a single or multiple (possibly unrelated) machines with the objective to maximize the weighted throughput, the total weight of jobs that complete before their deadline. To overcome known lower bounds for the competitive analysis, we assume that each job arrives with some slack $\epsilon > 0$; that is, the time window for processing job j on any machine i on which it can be executed has length at least $(1 + \epsilon)$ times j's processing time on machine i. Our contribution is a best possible online algorithm for weighted throughput maximization on unrelated machines: Our algorithm is $O(1/\epsilon)$-competitive, which matches the lower bound for unweighted throughput maximization on a single machine. Even for a single machine, it was not known whether the problem with weighted jobs is "harder" than the problem with unweighted jobs. Thus, we answer this question and close weighted throughput maximization on a single machine with a best possible competitive ratio $\Theta(1/\epsilon)$.

3.12 Learning-augmented algorithms for MTS with bandit access to multiple predictors

Marek Eliáš (Bocconi University – Milan, IT)

Joint work of Marek Eliáš, Matei Coşa
Main reference Matei Gabriel Coşa, Marek Eliáš: "Learning-Augmented Algorithms for MTS with Bandit Access to Multiple Predictors", CoRR, Vol. abs/2506.05479, 2025.
URL https://doi.org/10.48550/ARXIV.2506.05479

We consider the following problem: We are given ℓ heuristics for Metrical Task Systems (MTS), where each might be tailored to a different type of input instances. While processing an input instance received online, we are allowed to query the action of only one of the heuristics at each time step. Our goal is to achieve performance comparable to the best of the given heuristics. The main difficulty of our setting comes from the fact that the cost paid by a heuristic at time t cannot be estimated unless the same heuristic was also queried at time $t-1$. This is related to Bandit Learning against memory bounded adversaries (Arora et al., 2012 [1]). We show how to achieve regret of $O(OPT^{2/3}$) and prove a tight lower bound based on the construction of Dekel et al. (2013) [2].

References

1 Arora, R., Dekel, O., and Tewari, A. Online bandit learning against an adaptive adversary: from regret to policy regret. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, ICML'12, pp. 1747–1754, Madison, WI, USA, 2012. Omnipress. ISBN 9781450312851.
2 Dekel, O., Ding, J., Koren, T., and Peres, Y. Bandits with switching costs: T2/3 regret. *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, 2013. URL `https://api.semanticscholar.org/CorpusID:9425790`.

3.13 Algorithmic Contract Design

Michal Feldman (Tel Aviv University, IL)

Joint work of Michal Feldman, Paul Dütting, Tomer Ezra, Yoav Gal-Tzur, Thomas Kesselheim, Maya Schlesinger
Main reference Michal Feldman: "Algorithmic Contract Design (Keynote)", in Proc. of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, p. 2, Association for Computing Machinery, 2024.
URL https://doi.org/10.1145/3618260.3664273

Algorithmic contract design is a new frontier at the interface of economics and computation, studying scenarios where a principal delegates the execution of a costly project to an agent or a team of agents, and incentivizes them through a contract that specifies payments contingent on the project's success. This domain has gained increasing interest from the theoretical computer science community, particularly in the realm of combinatorial contracts. In this talk, I will survey two distinct models of combinatorial contracts, each illustrating unique sources of complexity encountered in contract design. The first model allows an agent to select from a set of possible actions, examining the intricate dependencies among these actions. The second model involves motivating a team of agents, focusing on the interdependencies within various agent combinations. I will present both (approximation) algorithms and hardness results concerning the optimal contract problem in these settings.

3.14 Online Algorithms with a Sample

Anupam Gupta (New York University, US)

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? A recent line of work has explored this research direction, and we give a survey of some results and techniques in the area.

3.15 Optimal Stopping with a Predicted Prior

Zhiyi Huang (University of Hong Kong, HK)

Joint work of Tian Bai, Zhiyi Huang, Chui Shan Lee, Dongchen Li
Main reference Tian Bai, Zhiyi Huang, Chui Shan Lee, Dongchen Li: "Optimal Stopping with a Predicted Prior", CoRR, Vol. abs/2511.03289, 2025.
URL https://doi.org/10.48550/ARXIV.2511.03289

There are two major models of value uncertainty in the optimal stopping literature: the secretary model, which assumes no prior knowledge, and the prophet inequality model, which assumes full information about value distributions. In practice, decision makers often rely on machine-learned priors that may be erroneous. Motivated by this gap, we formulate the model of optimal stopping with a predicted prior to design algorithms that are both consistent, exploiting the prediction when accurate, and robust, retaining worst-case guarantees when it is not.

Existing secretary and prophet inequality algorithms are either pessimistic in consistency or not robust to misprediction. A randomized combination only interpolates their guarantees linearly. We show that a family of bi-criteria algorithms achieves improved consistency-robustness trade-offs, both for maximizing the expected accepted value and for maximizing the probability of accepting the maximum value. We further prove that for the latter objective, no algorithm can simultaneously match the best prophet inequality algorithm in consistency, and the best secretary algorithm in robustness.

3.16 Prophet and Secretary at the Same Time

Thomas Kesselheim (Universität Bonn, DE)

Joint work of Thomas Kesselheim, Gregory Kehne
Main reference Gregory Kehne, Thomas Kesselheim: "Prophet and Secretary at the Same Time", CoRR, Vol. abs/2511.09531, 2025.
URL https://doi.org/10.48550/ARXIV.2511.09531

Many online problems are studied in stochastic settings for which inputs are samples from a known distribution, given in advance, or from an unknown distribution. Such distributions model both beyond-worst-case inputs and, when given, partial foreknowledge for the online algorithm. But how robust can such algorithms be to misspecification of the given distribution? When is this detectable, and when does it matter? When can algorithms give good competitive ratios both when the input distribution is as specified, and when it is not?

We consider these questions in the setting of optimal stopping, where the cases of known and unknown distributions correspond to the well-known prophet inequality and to the secretary problem, respectively. Here we ask: Can a stopping rule be competitive for the i.i.d. prophet inequality problem and the secretary problem at the same time? We constrain the Pareto frontier of simultaneous approximation ratios (α, β) that a stopping rule can attain. We introduce a family of algorithms that give nontrivial joint guarantees and are optimal for the extremal i.i.d. prophet and secretary problems. We also prove impossibilities, identifying (α, β) unattainable by any adaptive stopping rule. Our results hold for both n fixed arrivals and for arrivals from a Poisson process with rate n. We work primarily in the Poisson setting, and provide reductions between the Poisson and n-arrival settings that may be of broader interest.

3.17 Online Bin Packing under Random-Order and IID model

Arindam Khan (Indian Institute of Science – Bangalore, IN)

We consider the online bin packing problem under two stochastic settings. First, we study bin packing under the i.i.d. model, where item sizes are sampled independently and identically from a distribution in $[0, 1]$. We provide a linear time $(1 + \epsilon)$-competitive algorithm for online bin packing under i.i.d. Model. We then study the random-order model, where an adversary chooses the instance, but the order of arrival of items in the instance is drawn uniformly at random from the set of all permutations of the items. Kenyon [SODA'96] conjectured the random-order ratio is 1.15. We make some progress towards the conjecture.

3.18 Optimal Online Algorithm for Robust Flow Time Scheduling

Amit Kumar (Indian Institute of Technology – New Delhi, IN)

Joint work of Amit Kumar, Anupam Gupta, Debmalya Panigrahi, Zhaozi Wang
Main reference Anupam Gupta, Amit Kumar, Debmalya Panigrahi, Zhaozi Wang: "An Optimal Online Algorithm for Robust Flow Time Scheduling", in Proc. of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2026, Vancouver, BC, Canada, January 11-14, 2026, pp. 1214–1238, SIAM, 2026.
URL https://doi.org/10.1137/1.9781611978971.47

The problem of minimizing the total flow time on a single machine is one of the few problems for which we can give an optimal online algorithm: just schedule the job with the shortest remaining processing time (SRPT). However, this requires knowledge of the true running time p_j of each job j. Azar, Leonardi, and Touitou recently asked: what if we are given estimates p'_j for each job, such that the multiplicative error between p_j and p'_j (called the distortion) is at most μ? It is easy to construct examples where no algorithm can be $o(\mu)$ competitive; can we get $O(\mu)$ competitiveness?

We show how to achieve this asymptotically optimal result, improving on the previous best result of $O(\mu \log \mu)$. Moreover, we give a very simple algorithm to get this tight bound of $O(\mu)$; the previous ZigZag algorithm was relatively more involved. Our proof is via a dual-fitting argument based on the idea of a reduced instance: we consider an LP relaxation based on knapsack-cover inequalities, and show a solution with a large dual value. Our ideas can also be extended to give alternative dual-fitting arguments for previously analyzed algorithms (like SRPT, ZigZag, and others for minimizing the total flow time, and the Bansal-Dhamdhere algorithm for weighted flow-time minimization).

3.19 A Model for Data Filtering

Ravi Kumar (Google – Mountain View, US)

Joint work of Ravi Kumar, Aditya Bhaskara, Ashok Cutkosky, Manish Purohit

Selecting high quality data is a crucial prerequisite for training and fine-tuning large language models. We propose a rigorous theoretical framework for data filtering that lets us reason about the alignment of the filtered data to the task of interest while being fully oblivious to the specifics of the downstream learning model. Formally, given a target distribution d and access to a distribution D that is a mixture of d and some other unknown distribution D', we model the data filtering problem as designing a filter such that the distribution D conditioned on passing the filter is close in total variation distance to the target distribution d. Our main result is a simple filter that achieves a total variation distance of roughly $\tilde{O}\big(\frac{1}{N_d} + \frac{1}{pN_D}\big)$, where N_D is the number of samples from D, N_d is the number of samples from d, and p is the mixing probability. We also show nearly matching information-theoretic lower bounds.

3.20 Random Order Online Learning

Stefano Leonardi (Sapienza University of Rome, IT)

Joint work of Stefano Leonardi, Martino Bernasconi, Andrea Celli, Riccardo Colini-Baldeschi, Federico Fusco, Matteo Russo
Main reference Martino Bernasconi, Andrea Celli, Riccardo Colini-Baldeschi, Federico Fusco, Stefano Leonardi, Matteo Russo: "Online Learning in the Random Order Model", CoRR, Vol. abs/2510.02820, 2025.
URL https://doi.org/10.48550/ARXIV.2510.02820

In the random-order model for online learning, the sequence of losses is chosen upfront by an adversary and presented to the learner after a random permutation. While algorithms for adversarial inputs naturally maintain their regret guarantees in random order, simple no-regret algorithms exist for the stochastic model that fail against random-order instances. In this paper, we propose a general template to adapt stochastic learning algorithms to the random-order model without substantially affecting their regret guarantees.

3.21 Open problems in Knapsack Prophet Inequality and Contention Resolution

Will Ma (Columbia University – New York, US)

Main reference Will Ma, Calum MacRury, Cliff Stein: "Forward-backward Contention Resolution Schemes for Fair Rationing", in Proc. of the 26th ACM Conference on Economics and Computation, EC 2025, Stanford University, Stanford, CA, USA, July 7-10, 2025, p. 945, ACM, 2025.
URL https://doi.org/10.1145/3736252.3742654
Main reference Jiashuo Jiang, Will Ma, Jiawei Zhang: "Tight Guarantees for Multi-unit Prophet Inequalities and Online Stochastic Knapsack", in Proc. of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022, pp. 1221–1246, SIAM, 2022.
URL https://doi.org/10.1137/1.9781611977073.51

A sequence of items $i = 1, ..., n$ arrive with random sizes S_i and weights W_i drawn from known distributions, and must be irrevocably accepted or rejected, subject to the constraint that the total size accepted is at most 1. What is the maximum total weight that an online algorithm can accept in expectation, compared to a prophet who sees the realizations in advance? Surprisingly, the best-known algorithmic guarantee for this problem relies on the (loose) knapsack polytope followed by online contention resolution, and better guarantees are difficult to achieve even under random-order or IID arrivals. This talk surveys some of these challenges and puzzling open problems.

3.22 Single-Sample Prophet Inequalities

Rebecca Reiffenhäuser (University of Amsterdam, NL)

Joint work of Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, Rebecca Reiffenhäuser, Thomas Kesselheim, Brendan Lucier, Sahil Singla
Main reference Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, Rebecca Reiffenhäuser: "Single-Sample Prophet Inequalities via Greedy-Ordered Selection", in Proc. of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022, pp. 1298–1325, SIAM, 2022.
URL https://doi.org/10.1137/1.9781611977073.54
Main reference Paul Dütting, Thomas Kesselheim, Brendan Lucier, Rebecca Reiffenhäuser, Sahil Singla: "Online Combinatorial Allocations and Auctions with Few Samples", in Proc. of the 65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024, pp. 1231–1250, IEEE, 2024.
URL https://doi.org/10.1109/FOCS61266.2024.00081

The classic prophet inequality uses knowledge of the distributions that produce all online values to circumvent the impossibility of obtaining constant competitive ratios. However, in many problem versions, complete distributional knowledge is unnecessary to achieve this: instead, it often suffices to obtain a single sample (value) from each distribution before start of the algorithm. The talk introduces a general technique to obtain such Single-Sample Prophet inequalities, and demonstrates its generality by applying to different combinatorial versions of the problem.

3.23 Approximating Matroid Basis Testing for Partition Matroids using Budget-In-Expectation

Kevin Schewior (Universität Köln, DE)

Joint work of Lisa Hellerstein, Benedikt M. Plank, Kevin Schewior
Main reference Lisa Hellerstein, Benedikt M. Plank, Kevin Schewior: "Approximating Matroid Basis Testing for Partition Matroids using Budget-In-Expectation", in Proc. of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2026, Vancouver, BC, Canada, January 11-14, 2026, pp. 2588–2616, SIAM, 2026.
URL https://doi.org/10.1137/1.9781611978971.94

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.

Taking $\mathcal{M}$ to instead be a partition matroid, we show that previous approaches fail to give a constant-factor approximation. Our main result is a polynomial-time constant-factor approximation algorithm producing a randomized strategy for this partition matroid problem. We obtain this result by combining a new technique with several well-established techniques. 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 type of problem is of independent interest, will spark follow-up work, and has the potential for additional applications.

3.24 Incremental Graph Algorithms with Predictions

Shikha Singh (Williams College – Williamstown, US)

Joint work of Shikha Singh, Samuel McCauley, Benjamin Moseley, Aidin Niaparast, Helia Niaparast
Main reference Samuel McCauley, Benjamin Moseley, Aidin Niaparast, Helia Niaparast, Shikha Singh: "Incremental Approximate Single-Source Shortest Paths with Predictions", in Proc. of the ICALP, LIPIcs, Vol. 334, pp. 117:1–117:20, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
URL https://doi.org/10.4230/LIPICS.ICALP.2025.117

In an incremental graph problem, the vertices of the graph G are known ahead of time and the edges are revealed over time. The goal is to efficiently maintain some property (such as connectivity or shortest path) of the current graph G_t. Predictions about future inserts in such problems help interpolate smoothly between the fully online and fully offline settings. This talk describes recent results that use a recursive decomposition to leverage predictions and design faster incremental graph algorithms for shortest path and strong connectivity.

3.25 Improved Approximation Algorithms Using "Edge Predictions"

Sandeep Silwal (University of Wisconsin-Madison, US)

Joint work of Sandeep Silwal, Anders Aamand, Justin Y. Chen, Siddharth Gollapudi, Hao Wu, Yuhao Liu

Augmenting constraint satisfaction problems (CSPs) with predictions has been a recent popular theme in learning-augmented algorithms. In this talk I will describe improved approximation algorithms for CSPs using "a prediction per constraint." In the context of NP-hard graph problems, each edge provides two bits of information, corresponding to predictions about whether each of its endpoints belongs to an optimal solution. In other contexts, such as Max3SAT, every clause gives information about the variables participating in the clause. Even with weak predictions where each bit is only ϵ-correlated with the true solution, this information allows us to break approximation barriers in the standard setting. We develop algorithms with improved approximation ratios for MaxCut, Vertex Cover, Set Cover, and Maximum Independent Set problems (among others). Across these problems, our algorithms share a unifying theme, where we separately satisfy constraints related to high-degree vertices (using predictions) and low-degree vertices (without using predictions) and carefully combine the answers. We will also discuss the many open problems related to the "strengths" of different prediction models for CSPs and potential future directions.

3.26 Learning-Augmented Online Algorithms for Nonclairvoyant Joint Replenishment Problem with Deadlines

Seeun William Umboh (The University of Melbourne, AU)

Joint work of Seeun William Umboh, Jeremy Fineman, Michael Dinitz
Main reference Michael Dinitz, Jeremy T. Fineman, Seeun William Umboh: "Learning-Augmented Online Algorithms for Nonclairvoyant Joint Replenishment Problem with Deadlines", CoRR, Vol. abs/2511.16094, 2025.
URL https://doi.org/10.48550/ARXIV.2511.16094

This talk considers using predictions in the context of the online Joint Replenishment Problem with Deadlines. Prior work includes asymptotically optimal competitive ratios of $O(1)$ for the clairvoyant setting and $O(\sqrt{n})$ for the nonclairvoyant setting, where n is the number of items. Our goal is to significantly reduce the competitive ratio in the nonclairvoyant case by leveraging predictions.

3.27 Philosopher Inequalities: Online Bayesian Selection Beyond Competitive Analysis

David Wajc (Technion – Haifa, IL)

Joint work of David Wajc, Mark Braverman, Mahsa Derakhshan, Seffi Naor, Christos Papadimitriou, Tristan Pollner, Amin Saberi, Aravind Srinivasan

In online Bayesian selection problems, a seller is faced with a stream of buyers arriving sequentially. Each arriving buyer makes, for each item on sale, a take-it-or-leave-it offer drawn from some known distribution, which the seller must immediately either accept or decline. It is common to analyze the seller's gain via the lens of competitive analysis, comparing the seller's gain to that of the optimal offline solution, obtained by a "prophet" who knows the future. However, the seller might not even be able to compete with the optimal online algorithm, which may be computationally intractable to run. In this case the optimal online algorithm is obtained by a "philosopher" with sufficient time to think (compute). In this talk I will discuss some developments on online algorithms efficiently approximating this philosopher.

4 Open problems

4.1 Revisiting the free-order prophet for one-dimensional parameter classes of distributions

Vianney Perchet (ENSAE – Palaiseau, FR)

In their 1985 paper, Hill and Hordijk [1] defined, in the now classical prophet inequality problem, as "descending optimal" any class of distribution, parameterized by a one-dimensional parameter, such that the value of the online optimal algorithm (i.e., dynamic programming) is the highest when the distributions are observed in decreasing order with respect to their parameter.

We wonder if we can give a simple characterization of classes that are descending optimal, maybe when those random variables have the same distribution up to some (known) multiplicative factor, the one-dimensional parameter. It is known that Bernoulli, uniform and exponential distributions are indeed descending optimal.

An alternative consists in the same question when the dynamic programming is replaced with the best "single-threshold" algorithm.

References

1 Hill, T. P., Hordijk, A. (1985). *Selection of order of observation in optimal stopping problems.* Journal of applied probability, 22(1), 177-184.

4.2 Applications of Correlated Sampling from the Hypersimplex

David Wajc (Technion – Haifa, IL)

Joint work of Joseph (Seffi) Naor, Nitya Raju, Abhishek Shetty, Aravind Srinivasan, Renata Valieva, David Wajc
Main reference Joseph (Seffi) Naor, Nitya Raju, Abhishek Shetty, Aravind Srinivasan, Renata Valieva, David Wajc: "Dimension-Free Correlated Sampling for the Hypersimplex", in Proc. of the 17th Innovations in Theoretical Computer Science Conference, ITCS 2026, Bocconi University, Milan, Italy, January 27-30, 2026, LIPIcs, pp. 104:1–104:20, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026.
URL https://doi.org/10.4230/LIPICS.ITCS.2026.104

▶ **Definition 1.** Let $\Delta n, k \triangleq x \in [0,1]n :\| x \| 1 \leq k$ be the convex hull of origin and the hypersimplex (the same polytope with the inequality replaced with an equality).

▶ **Definition 2.** A distribution D over deterministic algorithms $A : \Delta n, k \to 2[n]$ mapping points in $\Delta n, k$ to subsets of $[n]$ of cardinality at most k is an α-stretch correlated sampling algorithm for $\Delta n, k$ if it satisfies the following properties for all $x, y \in \Delta n, k :$

(P1) (Cardinality) $|A(x)| \in \lfloor \| x \| 1 \rfloor, \lceil \| x \| 1 \rceil$ (hence $|A(x)| \leq k$) for all A in the support of D.

(P2) (Marginals) PrA$\sim$ $D[i \in A(x)] = x_i$ for all $i \in [n]$.

(P3) (Stretch) EA$\sim$ $D[|A(x) \oplus A(y)|] \leq \alpha. \| x - y \| 1$, where $\oplus$ denotes the symmetric difference.

For $k = 1$, it is known that $\alpha = 2$ is best possible. Algorithms for this problem, referred to as correlated sampling, have found applications to approximation algorithms, parallel repetition, replicability and privacy, parallel sampling and parallelizing next-token generation, cryptography, and locality sensitive hashing.

For $k > 1$, a work of Chen et al. ICALP17 shows that $\alpha = O(\log n)$ is attainable, improved to a dimension-free bound of $O(\log k)$ by Naor et al. ITCS26. The latter algorithm is also implementable efficiently in parallel and dynamic settings, and furthermore preserves submodular objectives. The recent work shows some simple applications, but it seems plausible that this machinery has further applications. What applications does this generalization of correlated sampling enable?

Participants

- Anders Aamand
University of Copenhagen, DK
- Spyros Angelopoulos
International Lab on Learning Systems – Montreal, CA
- Antonios Antoniadis
University of Twente – Enschede, NL
- Yossi Azar
Tel Aviv University, IL
- Eric Balkanski
Columbia University – New York, US
- Niv Buchbinder
Tel Aviv University, IL
- Shuchi Chawla
University of Texas – Austin, US
- Christian Coester
University of Oxford, GB
- José R. Correa
University of Chile – Santiago, CL
- Andrés Cristi
EPFL – Lausanne, CH
- Sami Davies
University of California – Berkeley, US
- Paul Dütting
Google – Zürich, CH
- Franziska Eberle
TU Berlin, DE
- Marek Eliáš
Bocconi University – Milan, IT
- Michal Feldman
Tel Aviv University, IL
- Rohan Ghuge
University of Texas – Austin, US
- Anupam Gupta
New York University, US
- Zhiyi Huang
University of Hong Kong, HK
- Sungjin Im
University of California – Santa Cruz, US
- Thomas Kesselheim
Universität Bonn, DE
- Arindam Khan
Indian Institute of Science – Bangalore, IN
- Amit Kumar
Indian Institute of Technology – New Delhi, IN
- Ravi Kumar
Google – Mountain View, US
- Stefano Leonardi
Sapienza University of Rome, IT
- Roie Levin
Rutgers University – New Brunswick, US
- Vasilis Livanos
University of Chile – Santiago, CL
- Will Ma
Columbia University – New York, US
- Nicole Megow
Universität Bremen, DE
- Marco S. Molinaro
Microsoft Research – Redmond, US & PUC – Rio de Janeiro, BR
- Benjamin J. Moseley
Carnegie Mellon University – Pittsburgh, US
- Seffi Naor
Technion – Haifa, IL
- Debmalya Panigrahi
Duke University – Durham, US
- Vianney Perchet
ENSAE – Palaiseau, FR
- Rebecca Reiffenhäuser
University of Amsterdam, NL
- Adi Rosén
CNRS – Paris, FR & Université Paris Cité, FR
- Barna Saha
University of California – San Diego, US
- Kevin Schewior
Universität Köln, DE
- Ziv Scully
Cornell University – Ithaca, US
- Golnoosh Shahkarami
MPI für Informatik – Saarbrücken, DE
- Sandeep Silwal
University of Wisconsin-Madison, US
- Shikha Singh
Williams College – Williamstown, US
- Sahil Singla
Georgia Institute of Technology – Atlanta, US
- Ola Svensson
EPFL – Lausanne, CH
- Seeun William Umboh
The University of Melbourne, AU
- David Wajc
Technion – Haifa, IL