

RESEARCH STATEMENT

Yilun Chen*

1 Motivation

My research focuses on understanding and addressing the fundamental challenges of *sequential decision-making* in the big data era. Such decision-making tasks arise frequently in important practical settings ranging from pricing, inventory control and revenue management to public healthcare and pandemic management. They have long been studied in the operations literature, often boiling down to solving stochastic dynamic programs (DP). Recently, amidst the ongoing data revolution, new, profound challenges stem from the increasing possibilities and expectations for companies/governments to fully exploit the information extracted from various relevant data streams (i.e. data-driven) in their decision-making processes. Such data-driven problems generally demand large-scale, more realistic and flexible models in different application domains, which renders the associated DPs computationally intractable due to *curse of dimensionality* issues. *Coming up with simple, tractable and provably good decision-making rules* is a pressing task faced by many practitioners, which keeps driving the development of solutions in academia, and motivates my own research. My research vision is to

- i. develop new methodologies that lead to simple decision-making solutions with provable performance guarantees; and*
- ii. devise (near) optimal and implementable algorithms for practitioners to improve their daily operations.*

2 Research Accomplishments

2.1 High-dimensional optimal stopping

The fundamental problem of optimal stopping has wide applications across operations management, economics, computer science and control, and is of central importance to the fields of applied probability and mathematical finance, with a long and rich history. To list a few examples, modern retailers consistently gather sales data/customer information to ensure launching a promotion at the right time; and governments collect data (at various levels) on a daily basis in support of the vital, real-time decisions of whether to shutdown during a pandemic. Over two decades ago, the then-hot topics of pricing complicated financial derivatives motivated the early studies of *high-dimensional* optimal stopping, which has since aroused considerable research interest, and remains highly relevant in modeling practical decision-making tasks in this big data era. Such

*School of Operations Research and Information Engineering, Cornell University, Email: yc2436@cornell.edu

high-dimensional optimal stopping problems, unlike their low dimensional, i.i.d. secretary-type counterparts, are notoriously challenging and long considered intractable.

In my work [1], we prove for the first time that, contrary to popular belief, high-dimensional optimal stopping is *fundamentally tractable*. More precisely, we develop a new methodology for approximately computing the optimal value/stopping policy, that can *provably* achieve any fixed error tolerance *efficiently* (with runtime scaling polynomially in the time horizon and effectively independent of the dimension). This surprising result is built on our novel discovery (in the same paper) of an equivalent formulation (in the form of an infinite expansion) of the optimal stopping DP, contrasting existing heuristic approaches (ADP, duality, deep learning etc.) that typically do not come with theoretically-strong joint-performance-and-efficiency guarantees. Our discovery also draws a novel connection between optimal stopping and the theory of network flows, which opens up the possibility of leveraging the rich tools in combinatorial optimization to solve challenging sequential decision-making problems. The work was awarded first place in the *2019 INFORMS Nicholson Student Paper Competition*.

2.2 High-dimensional sequential decision-making

In practice, underlying high dimensionality broadly exists in complex decision-making settings that are beyond the modelling capability of optimal stopping, imposing severe computational issues. Motivated by the fact that in many such settings, frequent switches from one action to another can be highly costly/undesirable/impossible, I study a fairly general class of high-dimensional sequential decision-making problems subject to a “limited action change” constraint in [4], and prove they are in fact tractable. Our model captures a large number of important real-world decision-making examples, including the dynamic pricing of a product that requires many raw materials/commodities as input (typically with highly volatile prices) under the *limited price change* business rule, and the evaluation of high-dimensional financial derivatives (swing options) with a restricted number of exercising opportunities. Viewing each time point of action change as an optimal stopping time, we leverage the results of [1] to build efficient and provably good approximation algorithms with runtime polynomial in the time horizon and independent of the problem dimension for any fixed error tolerance and fixed action-change budget.

2.3 Online maximum weighted independent set

Sequential decision-making problems with underlying network/graph structure can model various operational tasks faced by online platforms, and have been widely considered in the revenue management literature. In these applications, the underlying graph nodes often represent arriving customers, who, in reality, are characterized by high-dimensional data streams with possibly non i.i.d., history-dependent demands. The combination of such high-dimensionality/path dependency and the network structure renders the decision-making problems intractable, yet is underexplored in the literature. In [2] we overcome this crucial computational challenge for the first time in the context of high-dimensional *online maximum weighted independent set* in bounded-degree bipartite graphs. Combining several sophisticated techniques, including a novel blocking flow algorithm inspired by [1] and classical network flow algorithms such as Dinic’s algorithm, we propose a novel polynomial time approximation scheme (PTAS) that exhibits an interesting correlation decay prop-

erty. Our algorithm’s runtime scales only *linearly* in the time horizon for any fixed error tolerance and max degree, and is effectively *independent* of the underlying dimension.

2.4 Ongoing works

2.4.1 Fast and implementable optimal stopping

One of my immediate goals is to convert the theoretical progress of [1] into practical methods that can help practitioners deal with real-world tasks. Further building on the connection between optimal stopping and network flow theory, we devise a family of fast and implementable algorithms. The algorithms combine the core ideas of [1] with several classical network flow algorithms such as Push-Relabel. We perform numerical tests on multiple benchmark examples (default test examples in the optimal stopping/option pricing literature) to justify the practicality of the proposed algorithms. For instance, our algorithm outputs almost matching upper and lower bounds of the true prices of the benchmark *high-dimensional Bermudan knock-out options*, improving upon the state-of-the-art methods by an order of magnitude. The results are presented in working paper [3].

2.4.2 Bayesian regret of multi-armed bandit

Multi-armed bandit (MAB) is an important sequential decision-making problem with numerous applications in operations management (including dynamic price experimentation, online advertising etc). Over the years, various algorithms (e.g. methods based on upper confidence bounds (ucb), Thompson Sampling (TS)) have been proposed for MAB problems in the literature. In certain asymptotic regimes, many of these algorithms are proved to be “optimal.” However, how to effectively compare the performance of these “optimal” algorithms is a fundamental question that remains unanswered. We propose a solution in the simplified yet canonical case of Bayesian one-and-a-half arm bandits with Bernoulli rewards. Based on a higher order asymptotic analysis of the algorithm’s regret, we quantify the suboptimality of ucb and TS (compared to the exact-optimal, Gittins index policy), providing a framework for properly comparing MAB algorithms. This is a work in progress.

2.4.3 High-dimensional inventory control

Inventory control with lost sales and large lead times has long been recognized as intractable due to the high dimensionality introduced by the pipeline vector. Building on the recent progress of [5], we investigate the connection between several simple policies that overcome the curse of dimensionality and enjoy optimal performance guarantees in the asymptotic regime where lead time approaches infinity. This is a work in progress.

2.4.4 Generalization of machine learning models

Today, machine learning models of high complexity/expressiveness (e.g. deep neural networks) have achieved a huge empirical success in identifying the patterns in all kinds of data. These models often lead to “interpolating” solutions that fit the training data, yet do not exhibit “overfitting” issues in practice, contradicting the classical statistical theory. During my summer internship at IBM Research in 2020, my collaborators and I investigated this disagreement between theory and

practice by studying the impact of interpolation on a model’s generalization ability. This is an ongoing project.

3 Future plans

Moving forward, I would like to continue pursuing deep theoretical insights and simple practical solutions to address new decision-making challenges emerging from real-world applications. Specifically, I have the following plans to accomplish in the future.

Implementable algorithms and real-world applications. Given the broad range of high-impact real-world problems in which optimal stopping and high-dimensional control play a fundamental role, I am excited to find opportunities to collaborate with practitioners, leveraging our methods to make a significant impact. My concrete vision is to devise implementable algorithms/commercialized software based on the theory we developed, to help our partners make better decisions in complex environments, and ultimately improve revenue/social wellbeing.

Advancing the theory of sequential decision-making. My Ph.D works shed light on a number of important and interesting cutting-edge questions in the theory of sequential decision-making that remain to be explored. These questions often lie at the interface of control, optimization, applied probability and computer science, whose resolution requires diverse tools. My vision is to harness my broad set of skills developed during my Ph.D studies to pursue the answers to these questions. For example, I’m particularly interested in understanding the information-theoretical limit on the sample complexity of optimal stopping and sequential decision-making under various computational models (classical, parallel, quantum etc.). Such results can greatly advance our understanding of the fundamental hardness of certain sequential decision-making problems, and can potentially lead to more efficient algorithms.

References

- [1] Y. Chen and D. A. Goldberg. Beating the curse of dimensionality in optimal stopping. *Under Revision for Mathematics of Operations Research*, 2018.
First Place in the 2019 INFORMS Nicholson Student Paper Competition
Finalist in the 2018 INFORMS Section on Finance Best Student Paper Competition.
- [2] Y. Chen and D. A. Goldberg. Polynomial time algorithms for the weighted online bipartite independent set problem. *Preprint*, 2020.
- [3] Y. Chen and D. A. Goldberg. Fast and practical algorithms for high-dimensional optimal stopping. *Working paper*, 2020.
- [4] Y. Chen and D. A. Goldberg. A new approach to high-dimensional online decision-making. *Preprint to be submitted to Management Science*, 2020.
- [5] D. A. Goldberg, D. A. Katz-Rogozhnikov, Y. Lu, M. Sharma, and M. S. Squillante. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Mathematics of Operations Research*, 41(3):898–913, 2016.