



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

In This Issue

To cite this article:

(2022) In This Issue. Operations Research 70(4):iii-vii. <https://doi.org/10.1287/opre.2022.2331>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

<https://doi.org/10.1287/opre.2022.2331>

Copyright: © 2022 INFORMS

Analyzing Capacitated Two-Echelon Systems with Permutation-Dependent Separability

Capacitated multiechelon systems are common in practice due to the escalating costs of labor and advanced manufacturing technology. However, identifying the optimal replenishment policies for such systems is a largely open area of research due to the intrinsic complexity, especially when there is an upstream bottleneck. In “Technical Note—A Permutation-Dependent Separability Approach for Capacitated Two-Echelon Inventory Systems,” Shen, Yu, and Huh propose a new approach, that is, permutation-dependent separability, to tackle a capacitated two-echelon system in which the capacity of upstream stage can be the bottleneck. They show that the value function for the capacitated two-echelon system in each period is permutation-dependent separable, and that for each echelon, a permutation-dependent echelon base stock policy is optimal. The authors also develop efficient solution procedures on how to obtain the optimal policy.

Procurement Auctions with Capacity Reservation

It frequently happens that a supplier needs to invest in building capacity prior to supplying a product. Then the buyer will need to reserve capacity from suppliers and pay for this in advance of knowing the final demand. In “Capacity Games with Supply Function Competition,” Anderson, Chen, and Shao study a very general version of this problem, in which both capacity and production costs are volume dependent. They explore how an auction will operate in this setting and allow supply function bids from suppliers. The authors apply this framework to a newsvendor problem with unreliable suppliers, a portfolio procurement setting with supply options and a spot market, and a bundling problem with nonsubstitutable products. This helps to understand when we may expect the buyer to make a reservation choice that maximizes the overall supply chain profit. The key to these results is to show that the supply chain optimal profit is submodular in the set of suppliers.

On Modelling the Probability Distribution of Stochastic Sums

In “Technical Note—On Matrix Exponential Differentiation with Application to Weighted Sum Distributions,” Das, Tsai, Kyriakou, and Fusai propose an efficient methodology for approximating the unknown probability distribution of a weighted stochastic sum or time integral. Resulting from earlier contributions based on continuous-time Markov chain approximations of one-dimensional Markov processes is the Laplace transform of the unknown distribution available in exponential matrix form. The authors develop a bona fide Pearson curve-fitting approach to this distribution based on the moments, which they recover from the derivatives of the Laplace transform. Motivated by the computational hurdles toward this, they derive computationally efficient closed-form expressions for the derivatives of the matrix exponential. They then apply to pricing average-based options.

Sample Average Approximation in Data-Driven Newsvendor

In the data-driven newsvendor problem, the manager makes sequential inventory decisions while learning the unknown demand distribution based on past demand samples. How does the widely used sample average approximation approach perform in this problem? In “Technical Note—Data-Driven Newsvendor Problem: Performance of the Sample Average Approximation,” Lin, Huh, Krishnan, and Uichanco analyze the performance of the sample average approximation as the time horizon grows, which turns out to be the best possible. The authors also examine how the local flatness of the demand distribution around the optimal order quantity affects the complexity of the problem. They show that the sample average approximation has the best achievable performance in terms of not only the time horizon, but also the local flatness of the demand distribution.

Managing Queues in Online Food-Ordering Services

The shift in the restaurant industry toward digital ordering argues for major changes in how orders are managed. The main difficulty in managing queues in online food-ordering services arises from the fact that, as opposed to dine-in customers, online customers are promised a pick-up time; customers are dissatisfied if the order is not completed by that time and are also dissatisfied if the order is completed much ahead of time because the food loses freshness. In “Order Now, Pickup in 30 Minutes: Managing Queues with Static Delivery Guarantees,” Farahani, Dawande, and Janakiraman propose and analyze strategies for managing queues in online food-ordering services with the goal of keeping customer satisfaction as high as possible.

An Integrated Approach to Managing Vessel Service in Seaports

Efficient vessel service is of utmost importance in the maritime supply chain. When serving a group of incoming vessels, berth allocation and pilotage planning are the two most important decisions made by a seaport. Although they are closely correlated, the berth allocation problem and pilotage planning problem are often solved sequentially, leading to suboptimal or even infeasible solutions for vessel services. In “Vessel Service Planning in Seaports,” Wu, Adulyasak, Cordeau, and Wang focus on a vessel service planning problem that optimizes berth allocation and pilotage planning in combination. To solve the joint problem, the authors develop an exact solution method that combines Benders decomposition and column generation within an efficient branch-and-bound framework. They also propose acceleration strategies that significantly improve the performance of the algorithm. Test instances from one of the world’s largest seaports are used to validate the effectiveness of the approach and demonstrate the value of integrated planning.

Path Dependence of Optimal Dynamic Momentum Strategies

The momentum effect has been studied by a tremendous number of papers, but surprisingly, there is little research on optimal momentum strategies. In “Optimal Dynamic Momentum Strategies,” Li and Liu explicitly solve the optimal dynamic portfolio problem when a risky asset has momentum. They show that, to optimally exploit momentum, one needs to account for path dependence, as well as momentum. In contrast, the momentum strategies discussed in most papers exploit only momentum. The optimal portfolio weight also significantly differs from that in the classic framework of Merton. Due to their path dependence, optimal portfolio weights have a wide distribution for a given level of momentum; for example, investors may short the risky asset if it has rebound price paths but leverage if it has hump-shaped price paths. This effect tends to be the most significant after large price swings. Path dependence is described through explicit formulas as well as heuristic statistics.

Capturing the Uncertainty in Long-Term Mortality Forecasts

The uncertainty in future longevity presents a substantial risk factor for insurance companies, pension funds, and retirement systems. In “Modeling the Risk in Mortality Projections,” Zhu and Bauer present novel stochastic models for analyzing this longevity risk that focus on the uncertainty associated with long-term mortality projections and capture the evolution of mortality forecasts over the past decades. They arrive at their models by analyzing time series of mortality forecasts in a forward modeling framework, which contrasts with conventional stochastic mortality models that are built on age-specific realized mortality rates. The authors showcase their models in exemplifying financial applications in both traditional life insurance markets and the emerging longevity risk transfer market. A key takeaway is that uncertainty in positions that depend on the long-term evolution of mortality is substantially greater under their models than suggested by conventional models.

Valuing Debt and Equity in Interbank Networks

Valuation adjustments to account for, for example, counterparty risk, have become an important part of derivative valuation by any bank since the 2007–2009 financial crisis. As evidenced by that crisis, considering the risk of a single firm alone can cause gross misspecification of firm health. In “Pricing of Debt and Equity in a Financial Network with Comonotonic Endowments,” Banerjee and Feinstein construct an analytical formulation for a valuation adjustment that takes the entire network of counterparty obligations into account when all institutions have comonotonic endowments. From the perspective of stress testing, such a setting corresponds to a systematic shock to all firms. This comonotonic setting is then shown to provide computationally tractable upper and lower bounds to the general network valuation problem. The authors demonstrate that the comonotonic endowment setting arises endogenously in a system of equity maximizing firms.

Easy Cases of Deadlock Detection in Train Scheduling

In a railway network, a deadlock occurs when two or more trains are preventing each other from moving forward by each occupying the tracks required by the other. Deadlocks are rare but pernicious events in railroad operations, and, in most cases, they are caused by human errors and involve only two extra-long trains missing their last potential meet location. In “Easy Cases of Deadlock Detection in Train Scheduling,” Dal Sasso,

Lamorgese, Mannino, Tancredi, and Ventura prove that the identification of two-train deadlocks can be performed in polynomial time. Moreover, they also present a pseudo-polynomial but efficient oracle that allows real-time early detection and prevention of any (potential) two-train deadlock in the Union Pacific (a U.S. class 1 rail company) railroad network. A deadlock prevention module based on the work in this paper will be put in place at Union Pacific to prevent all deadlocks of this kind.

Shaping Opinions in Social Networks

How can we place persuasion agents in a social network to influence a population? In “Optimizing Opinions with Stubborn Agents,” Hunter and Zaman present an algorithm based on opinion dynamics models that shows where to place these agents in a network for maximum persuasive effect. Using this algorithm, one can shape opinions in a variety of interesting ways. For instance, one can use the agents to maximize the opinion mean in order to increase support for an issue. More interestingly, one can use the algorithm to shape the opinion variance in order to decrease, or even increase, the polarization in a network. Simulations on a variety of real Twitter networks showed that, with their algorithm, a small number of strategically placed agents can create significant opinion shifts.

Joint Online Learning and Resource Allocation

Joint online learning and resource allocation is a fundamental problem inherent in many applications. In a general setting, heterogeneous customers arrive sequentially, each of which can be allocated to a resource in an online fashion. Customers stochastically consume the resources, allocations yield stochastic rewards, and the system receives feedback outcomes with delay. In “Online Resource Allocation with Personalized Learning,” Zhalechian, Keyvanshokoh, Shi, and Van Oyen introduce a generic framework to solve this problem. It judiciously synergizes online learning with a broad class of online resource allocation mechanisms, where the sequence of customer contexts is adversarial, and the customer reward and resource consumption are stochastic and unknown. The authors propose online algorithms that strike a three-way balance between exploration, exploitation, and hedging against adversarial arrival sequence. A performance guarantee is provided for each online algorithm, and the efficacy of their algorithms is demonstrated using clinical data from a health system.

The Sequential MNL Model: Algorithmic Frameworks for Product Recommendation Displays and Data-Driven Case Studies

In “Technical Note—The Multinomial Logit Model with Sequential Offerings: Algorithmic Frameworks for Product Recommendation Displays,” Feldman and Segev consider a sequential assortment problem that has applications ranging from appointment scheduling in hospitals, restaurants, and fitness centers to product recommendation in e-commerce settings. Their main contribution comes in the form of a strongly polynomial-time approximation scheme for the most general form of the problem. The authors also conduct an extensive case study in which they fit their sequential model to historical search data from Expedia’s hotel booking platform. They observe substantial gains in fitting accuracy when their model is benchmarked against other well-known choice models designed for the setting at hand.

Optimal Dynamic Pricing for a Product Portfolio

Dynamic pricing is a well-known revenue management technique used by firms to maximize their revenues. In industries

such as fashion retail and airlines, it is observed in practice that firms vary the prices of their products, through time, only among certain pre-fixed discrete price points; for example, in fashion retail, products are first sold at a base price and then later at, say, 10%, 20%, 30% discounted prices. In such discrete-price practices, when the number of products offered by a firm is large, it is mathematically challenging to determine the optimal pricing decisions. In “Multiproduct Pricing with Discrete Price Sets,” Manchiraju, Dawande, and Janakiraman design pricing algorithms that are fast and result in near-optimal revenues.

Scheduling High-Capacity Shared Rides via Hypergraph Matchings

Sharing economies, such as Uber Pool and Lyft Shared Saver, are becoming ubiquitous in modern transportation services. In “Technical Note—Online Hypergraph Matching with Delays,” Pavone, Saberi, Schiffer, and Tsao study an online hypergraph matching problem inspired by on-demand scheduling of shared rides in ridehailing systems. They formulate a ridesharing problem through a weighted hypergraph in which riders are represented by vertices and hyperedges and the associated weights represent the utility of serving the associated vertices groups of riders with a single vehicle. A ridesharing algorithm is, thus, a matching algorithm for the hypergraph. To model the on-demand nature of ride-sharing, riders arrive sequentially and must either be quickly and irrevocably matched or discarded. In this work, the authors present a polynomial-time algorithm based on randomized batching and prove that it has the optimal worst-case performance.

A Near-Optimal Patient Admission Policy in Post-acute Care

Motivated by applications in the postacute healthcare industry, in “Technical Note—A Near-Optimal Algorithm for Real-Time Order Acceptance: An Application in Postacute Healthcare Services,” Qu, Dawande, and Janakiraman study an infinite-horizon, stochastic optimization problem with a set of long-term capacity investment decisions and a sequence of real-time order acceptance/rejection decisions. To maximize the long-run average expected profit per period, the firm accepts/rejects stochastically arriving referrals in real time. Referrals differ in the revenue they offer to the firm, the resource requirements and the frequency of usage of the resources, and the stochastic duration of the episode. The authors develop a simple policy, derive a worst case guarantee on its optimality gap, and show that the policy is asymptotically optimal. They also illustrate an impressive numerical performance of the policy using public healthcare data.

A General Framework for the Dynamic Allocation of Reusable Resources

In “Technical Note—Bifurcating Constraints to Improve Approximation Ratios for Network Revenue Management with Reusable Resources,” Baek and Ma provide a unified framework for dynamically allocating reusable resources over time in the general network revenue management setting. That is, each incoming query is allowed to request an arbitrary combination of resources for different unknown durations, capturing applications like cloud computing and the dispatching of human resources. Through this framework, they derive simple algorithms that achieve the same approximation ratios as existing algorithms, which only allow queries to request one resource. Their framework also suggests a meta-algorithm for “bifurcating” the resources depending on the network structure and then using two different algorithms for controlling the capacities of two different groups of resources. This contrasts existing “one-size-fits-all”

approaches to network revenue management and is demonstrated to improve both theoretical guarantees and empirical performance over a wide variety of instances.

Modelling Choices of Price Sensitive Customers

Customers are sensitive to price when they shop from a group of substitutable products. Specifically, they are often attracted to the cheapest product in the group, even when the price difference to the next cheapest product is small. Such a phenomenon has been long noticed by the airline industry when customers purchase airline tickets. In “Network Revenue Management Under a Spiked Multinomial Logit Choice Model,” Cao, Kleywegt, and Wang consider a new choice model to incorporate customers’ preference toward the cheapest product, which extends the popular multinomial logit choice model. They find that using this new choice model may improve airlines’ revenue. Moreover, (near) optimal revenue management policies under this choice model can be computed in an efficient manner.

Revenue Prediction for Network Goods with Uncertain Network Effects

What is the revenue prediction of a seller in selling digital goods when the seller does not have full information about the extent of externality of users? In “Technical Note—Revenue Volatility Under Uncertain Network Effects,” Baron, Hu, and Malekian address this question by characterizing how the seller’s revenue prediction depends on the underlying network structure and the available information to the seller.

Persuasion in Networks: Public Signals and Cores

In many social networks, individuals’ actions depend on the actions of their peers and their information. How can a designer design informative signals to induce a desired outcome in a networked system? In “Persuasion in Networks: Public Signals and Cores,” Candogan answers this question in a setting where the designer is restricted to using public signaling mechanisms. He provides a convex programming framework for obtaining optimal public signaling mechanisms in networks where agents’ actions exhibit local strategic complementarities and characterizes the (double-interval) structure of these mechanisms. The framework he develops is useful in other settings (where the designer’s payoff is an increasing step function of the posterior mean). The author also provides an approach for designing asymptotically optimal public signaling mechanisms in large random networks that relies only on using the (limiting) degree distribution information. Finally, he sheds light on another fundamental question: Which networks are more amenable to persuasion?

Efficient and Profit-Maximizing Dynamic Double Auctions for Two-Sided Markets

Two-sided markets that enable sellers and buyers to trade have received considerable attention in the past decade. Prominent examples include online advertising, freelancing, and ride-hailing. In these markets, trade is coordinated by an intermediating platform that determines which parties should trade, collects payments from buyers, and transfers payments to sellers. How should trading mechanisms be designed when agents repeatedly trade over time? In “Dynamic Double Auctions: Toward First Best,” Balseiro, Mirrokni, Paes Leme, and Zuo design dynamic double auctions that satisfy the following practical requirements: no positive transfers, that is, the platform never asks sellers to make payments nor are buyers ever paid; and

periodic individual rationality, that is, agents should derive a nonnegative utility from every trade opportunity. This work provides mechanisms satisfying these requirements that, as the number of trading opportunities grows, are asymptotically efficient, budget balanced, and allow to extract the welfare generated by the market as profit.

Targeting Interventions in Networks Without Measuring the Whole Network

In the presence of contagion, decision makers strategize about where in a network to intervene (e.g., seeding a new product). A large literature has developed methods for approximately optimizing the choice of k seeds to cause the largest cascade of, for example, product adoption. However, it is often impractical to measure an entire social network. In “Seeding with Costly Network Information,” Eckles, Esfandiari, Mossel, and Rahimian develop and analyze algorithms for making a bounded number of queries of a social network and then selecting k seeds. They prove hardness results for this problem and provide almost tight approximation guarantees for their proposed algorithms under widely used models of contagion. One proposed algorithm is practical for both querying online social networks and structuring in-person surveys. This framework further allows reasoning about tradeoffs between spending budget on collecting more network data versus increasing the number of seeds.

Dynamic Stochastic Matching Under Limited Time

In “Dynamic Stochastic Matching Under Limited Time,” Aouad and Saritaç analyze the design of matching policies in dynamic markets such as carpooling platforms and kidney exchange schemes. A crucial distinction with previous literature is that the agents’ arrivals and departures are fully dynamic. The demand and supply side are constantly replenished; each market participant remains available for potential matches during a limited period of time. Specifically, the authors formulate a general dynamic matching model over edge-weighted graphs, where the agents’ arrivals and abandonments are stochastic and heterogeneous. The platform controls how long each agent waits and whom s/he is matched with. These decisions are subject to a fundamental tradeoff between increasing market thickness and mitigating the risk of abandonments from certain participants. The authors’ main contribution is to devise simple matching algorithms with strong performance guarantees for a broad class of networks. In contrast, they show that widely used batching algorithms have an arbitrary bad performance on certain graph-theoretic structures. Their analysis involves novel techniques including linear programming benchmarks, value function approximations, and proxies for continuous-time Markov chains, which may be of broader interest. Extensive simulations on real-world taxi demand data demonstrate that the newly developed algorithms can significantly improve cost efficiency against batching algorithms.

Taming High-dimensional Markov models

In “Learning Markov Models Via Low-Rank Optimization,” Zhu, Li, Wang, and Zhang focus on learning a high-dimensional Markov model with low-dimensional latent structure from a single trajectory of states. To overcome the curse of high dimensions, the authors propose to equip the standard MLE (maximum-likelihood estimation) with either nuclear norm regularization or rank constraint. They show that both approaches can estimate the full transition matrix accurately using a trajectory of length that is merely proportional to the number of states. To solve the rank-constrained MLE, which is a nonconvex problem, the authors develop a new DC (difference) programming algorithm. Finally,

they apply the proposed methods to analyze taxi trips on Manhattan Island and partition the island based on the destination preference of customers; this partition can help balance supply and demand of taxi service and optimize the allocation of traffic resources.

Exploiting Bilevel Optimization Techniques to Disconnect Graphs Into Small Components

In order to limit the spread of possible viral attacks in a communication or social network, it is necessary to identify critical nodes, the protection of which disconnects the remaining unprotected graph into a bounded number of shores (subsets of vertices) of limited cardinality. In “Casting Light on the Hidden Bilevel Combinatorial Structure of the Capacitated Vertex Separator Problem,” Furini, Ljubic, Malaguti, and Paronuzzi provide a new bilevel interpretation of the associated capacitated vertex separator problem and model it as a two-player Stackelberg game in which the leader interdicts (protects) the vertices, and the follower solves a combinatorial optimization problem on the resulting graph. Thanks to this bilevel interpretation, the authors derive different families of strengthening inequalities and show that they can be separated in polynomial time. The ideas exploited in their framework can also be extended to other vertex/edge deletion/insertion problems or graph partitioning problems by modeling them as two-player Stackelberg games to be solved through bilevel optimization.

Supply Risk Mitigation Through Informative Sourcing and Expediting

How to use real-time supply information to effectively deploy and manage a multisource, stochastic, and capacitated supply system is challenging. In “Smart Policies for Multisource Inventory Systems and General Tandem Queues with Order Tracking and Expediting,” Song, Xiao, Zhang, and Zipkin construct a class of ordering and expediting policies with easy-to-understand structures but using more supply information than usual. They identify two special policies that possess simple closed-form performance measures, which greatly enhance policy evaluation and optimization. The authors show that combining sourcing flexibility with some, but limited, dynamic expediting is valuable. In contrast, the upstream congestion information is of limited value.

Risk Averse Stochastic Programming: Time Consistency and Optimal Stopping

Decision making under uncertainty includes reassessing and reevaluating risk after initial decisions. To this end, it is essential to consider a governing value process and to track its evolution over time. In “Risk-Averse Stochastic Programming: Time Consistency and Optimal Stopping,” Pichler, Liu, and Shapiro develop a consistent framework by consolidating risk and optimal stopping. The authors develop the notion of time consistency for the stochastic multistage optimization problem. Supermartingales and envelopes characterizing optimal decisions are given explicitly. With that, dynamic equations are derived, which gradually reveal the optimal policy. Taking risk into account requires updating optimal policies, as an explicit example on American options demonstrates.

Stability of Parallel Server Systems

A parallel server system is a queueing system in which jobs are routed upon their arrivals to one of several buffers, each handled by a different server. The main operational and theoretical problem in such systems is to find an efficient routing

policy that maximizes their throughput (or minimizes waiting times). In “Stability of Parallel Server Systems,” Moyal and Perry consider a large class of routing policies, which includes the most prevalent policies studies in the literature, under the assumption that routing errors may occur because of incomplete information about the state of the system at decision epochs. The stability region for this class of policies is studied as a function of the error probability, and it is shown that the standard stability condition, namely, that the traffic intensity is smaller than one, does not guarantee that the system is stable.

Routing and Scheduling Platooning Vehicles

In “A Repeated Route-then-Schedule Approach to Coordinated Vehicle Platooning: Algorithms, Valid Inequalities and Computation,” Luo and Larson propose a novel repeated route-then-schedule algorithmic framework to efficiently solve a complex vehicle routing and scheduling problem arising in the intelligent transportation system. The goal is to maximize the collective savings of a set of vehicles (especially heavy-duty vehicles) by utilizing the fact that platooning vehicles save energy due to reduced aerodynamic drag. In the algorithm, the original simultaneous route-and-schedule approach is decomposed into the routing stage and scheduling stage with a sophisticated learning-like feedback mechanism to update the presumed fuel cost for each vehicle traversing through each road segment. This leads to an iterative change of objective function in the routing problem and thereby changes the routes that are fed to the scheduling problem. This approach helps identify high-quality solution. The algorithmic framework leads to a very tight formulation of subproblems that can be solved in a timely manner.

Taking Advantage of the Lead Time Randomness in Supply Chains

Randomness in lead times is a major—and increasingly important—issue of inventory management, as a variety of risk factors motivate companies to diversify their supply sources and rely on distributed networks of suppliers. In “Exploiting Random Lead Times for Significant Inventory Cost Savings,” Stolyar and Wang show that, surprisingly, instead of being a damaging factor to supply chain performance, randomness may be harnessed for potentially very substantial reductions of inventory costs. Specifically, the theoretical analysis and simulation results in the paper demonstrate that, under certain conditions, appropriately designed novel policies can significantly outperform the conventional base stock policies.

Dynamic Resource Allocation in the Cloud with Near-Optimal Efficiency

Motivated by maximizing spot instances in cloud shared systems, in this work, the authors consider the problem of taking advantage of unused resources in highly dynamic cloud environments while preserving users’ performance. In “Dynamic Resource Allocation in the Cloud with Near-Optimal Efficiency,” Perez-Salazar, Menache, Singh, and Toriello introduce an online model for sharing resources that captures basic properties of cloud systems, such as unpredictable users’ demand patterns, very limited feedback from the system, and service level agreement (SLA) between the users and the cloud provider. They provide a simple and efficient algorithm for the single-resource case. For any demand patterns, their algorithm guarantees near-optimal resource utilization as well as high users’ performance

compared with their SLA baseline. In addition to this, the authors validate empirically the performance of their algorithm using synthetic data and data obtained from Microsoft’s systems.

Inferring Objective Functions of a Linear Program from Noisy Decision Data: A Stability Perspective

Although inverse linear programming (LP) has received increasing attention as a technique to identify an LP that can reproduce observed decisions that are originally from a complex system, the performance of the linear objective function inferred by existing inverse LP methods is often highly sensitive to noise, errors, and uncertainty in the underlying decision data. Inspired by robust regression techniques that mitigate the impact of noisy data on the model fitting, in “Quantile Inverse Optimization: Improving Stability in Inverse Linear Programming,” Shahmoradi and Lee propose a notion of stability in inverse LP and develop an inverse optimization model that identifies objective functions that are stable against data imperfection. Although such a stability consideration renders the inverse model a large-scale mixed-integer program, the authors analyze the connection between the model and well-known bilinear problems and propose an efficient exact algorithm as well as heuristics.

Preconditioning and Regularization Enable Faster Reinforcement Learning

Natural policy gradient (NPG) methods, in conjunction with entropy regularization to encourage exploration, are among the most popular policy optimization algorithms in contemporary reinforcement learning. Despite the empirical success, the theoretical underpinnings for NPG methods remain severely limited. In “Fast Global Convergence of Natural Policy Gradient Methods with Entropy Regularization,” Cen, Cheng, Chen, Wei, and Chi develop nonasymptotic convergence guarantees for entropy-regularized NPG methods under softmax parameterization, focusing on tabular discounted Markov decision processes. Assuming access to exact policy evaluation, the authors demonstrate that the algorithm converges linearly at an astonishing rate that is independent of the dimension of the state-action space. Moreover, the algorithm is provably stable vis-à-vis inexactness of policy evaluation. Accommodating a wide range of learning rates, this convergence result highlights the role of preconditioning and regularization in enabling fast convergence.

A Fluid-Diffusion-Hybrid Limiting Approximation for Priority Systems with Fast and Slow Customers

In “A Fluid-Diffusion-Hybrid (FDH) Approximation for Priority Systems with Fast and Slow Customers,” Yu, Iravani, and Perry propose a new heavy-traffic asymptotic regime for a two-class priority system in which the high-priority customers require substantially larger service times than the low-priority customers. In the FDH limit, the high-priority queue is a diffusion, whereas the low-priority queue operates as a (random) fluid limit, whose dynamics are driven by the former diffusion. A characterizing property of the authors limit process is that, unlike other asymptotic regimes, a non-negligible proportion of the customers from both classes must wait for service. This property allows them to study the costs and benefits of de-pooling, and prove that a two-pool system is often the asymptotically optimal design of the system.