INFORMS 2022 Conference Notes

Enrique Areyan Viqueira and Fiete Krutein

October 2022

Contents

1 Introduction

The following notes document the talks we attended at INFORMS 2022. Any mistake or misrepresentation of the work is our own and not the presenters' fault. Please contact us if we misunderstood your work and would like us to correct it!

2 Sunday, October 16, 2022

2.1 Session: Delegated Sequential Search

2.1.1 Optimal Presentation of Alternatives

1. Authors: Zeya Wang, Morvarid Rahmani, Karthik Ramachandran, Georgia Institute of Technology.
2. Working paper.

Abstract

In many contexts such as technology and management consulting, clients seek the expertise of providers to find solutions for their problems. When there are multiple alternatives that could potentially solve the client's problem, providers can lead the client's exploration by choosing which alternative(s) to present and in what sequence. In this paper, we develop a dynamic game-theoretic model where the provider chooses how to present alternative solutions, and the client chooses which solution to try. Our analysis reveals that it is generally optimal for the provider to offer alternatives sequentially. Following a failed trial, the provider should readily offer a new alternative if the client's capability is either very high or very low. Otherwise, the provider should allow the client to try the same solution multiple times, especially when the project duration is long.

3. Gamma Therapeutics Example
   (a) Pharmaceutical company to treat cancer
   (b) Founder is a brilliant scientist but limited expertise in getting the therapeutic approved by the FDA (brief time window of time/runways)
   (c) The entrepreneur is under time pressure since the company needs to report advances to investors
   (d) Imperfect operations, compared to more established companies
4. (a) Some of these observations apply to other startups
   (b) Most of them fail.
   (c) How to increase the chances of startup success?
5. Mentoring
   (a) External advisors/experts
   (b) Incubators (increases the success ratio of startups significantly).
   (c) Accelerators
6. Mentoring entrepreneurs:
   (a) The context is:
      i. has a problem that needs to be solved but no direction
      ii. Advisor is aware of solutions
7. Research question

   (a) How should an advisor recommend alternative solutions to an entrepreneur?
   (b) Taking into account the constraints of the particular startup

8. Solution: a sequential decision model
   (a) At each round advisor selects a set of options
   (b) Entrepreneur selects a solution from the set

9. Outcomes:
   (a) Solution implemented
   (b) Entrepreneurs choose not to follow the advice and terminate their relationship with the advisor

10. Alternatives
    (a) Only two alternatives $\theta_a, \theta_b \in [0, 1]$, viable or not
    (b) entrepreneur operational capability, $\gamma \in [0, 1]$
    (c) Success probability $\gamma\theta$

11. Evolution of beliefs
    (a) Posterior belief updates using Bayes Rule $Pr(\text{option } a \text{ being viable}|a \text{ failed}) = \frac{(1-\gamma)\theta_1}{1-\gamma\theta_a}$
    if $\gamma = 0$, there in no learning, that is $\theta_a^1 = \theta_a$
    if $\gamma = 1$, there is complete learning, that is $\theta_a^1 = 0$
    (b) In each round, both update their beliefs

12. Analytic results
    (a) Author shows equilibrium strategies (μ_1^*, μ_2^*)
    (b) The equilibrium depends on the effectiveness of the alternatives presented by the advisor and the entrepreneur's operational capability.

13. Other results
    (a) Effect of trial cost
    (b) Effect number of alternative solutions

14. Takeaways
    (a) Advisor should strategically sequence their recommendation.
    (b) First recommends inferior before superior alternative.
    (c) Entrepreneurs' operational capability matters.

2.2 Plenary: Cynthia Rudin. Do Simpler Machine Learning Models Exist and How Can We Find Them?

Abstract

While the trend in machine learning has tended towards building more complicated (black box) models, such models are not as useful for high stakes decisions - black box models have led to mistakes in bail and parole decisions in criminal justice, flawed models in healthcare, and inexplicable loan decisions in finance. Simpler, interpretable models would be better. Thus, we consider questions that diametrically oppose the trend in the field: for which types of datasets would we expect to get simpler models at the same level of accuracy as black box models? If such simpler-yet-accurate models exist, how can we use optimization to find these simpler models? In this talk, I present an easy calculation to check for the possibility of a simpler (yet accurate) model before computing one. This calculation indicates that simpler-but-accurate models do exist in practice more often than you might think. Also, some types of these simple models are (surprisingly) small enough that they can be memorized or printed on an index card.

1. COMPAS, a proprietary model for recidivism.
   (a) Wondering how accurate COMPAS was.
   (b) Compared COMPAS vs. CORELS.
   (c) CORELS produces a tiny model, a simple tree. Basically, as a function of priors.
   (d) The main point is that this model was as accurate as COMPAS, which is a more complicated model.
   (e) given the high-stakes nature of the application, is it natural to ask why do we need proprietary models in this case?
2. 2HELPS2B: another example is preventing brain damage in Critically ill patients.
   (a) Tiny model (fits in a slide)
   (b) Just as accurate as black box models
   (c) doctors can decide whether to trust it
3. Problem spectrum: we think about the following two categories of data/problems very differently!
   (a) Tabular
      i. categorical data; counts mostly
      ii. more robust to small changes
   (b) Raw
      i. pixels, words, parts of sound waves
      ii. change one pixel, and you no longer have a realistic example
4. Rudin's theory: the Rashomon Set theory
   (a) Assume: there are a lot of good models
   (b) If the set of simple models is a good cover for good models, there should be at least one simple, good model.
   (c) Claim: Rashomon set large in many problems of interest.

   (a) On the existence of simpler machine learning model (insert link), they found that Large Rashomon sets are correlated with
      i. the existence of simple models and

ii. many different ML methods having the same performance
iii. that is, if you run many ML models and they all have similar performance, then think of a large Rashomon set.

(b) Implication: if the theory holds Optimizing for simplicity won't sacrifice performance. This is really important in real-world, high-stakes applications.

5. Explainable Machine Learning Challenge (HELOC dataset)
   (a) Data about 10k application, many factors/features
   (b) Best black-box accuracy (boosted decision trees 73%)
   (c) Best black box AUC (2-layer NN) .8
6. Their method is a generalized additive model.
   (a) Fast Sparse Classification for Generalized Linear and Additive Models (insert link)
   (b) Tested on the FICO dataset. Similar performance but again...
   (c) resulting model is tiny (compared to black-box) and fits in one slide.
      i. consists of 21 total step features
      ii. moreover, is created in under 3.85 seconds.
   (d) Their methods in basically a Sparse Logistic Regression
      i. trying to set coefficients to 0 as much as possible
      ii. using exponential loss so that there is an analytical formula for each coefficient
7. Summary so far
   (a) Rashomon set:
      i. simpler model exists when there are many good models.
      ii. If you have a Rashomon set, some algorithms can find sparse (tiny) accurate model
      iii. Even on competition (complicated) datasets, they find an interpretable model.
8. Observation:
   (a) ML community has mostly been concerned with building complex models and worrying about overfitting
   (b) Rudin wants to go the opposite way and build simpler models. Under learning theory, we wouldn't have to worry about overfitting with simple models.
9. Optimal Sparse Decision Tress (GOSDT algorithm, insert link)
   (a) Optimized for misclassification error plus sparsity of the tree
   (b) Example usage on the Broward Country Florida re-arrest data
   (c) Another example is on the FICO dataset.
10. The universal paradigm of machine learning
   (a) Training Set $->$ Algorithm $->$ Predictive Model
   (b) Claim: this is wrong! (at least for high-stakes applications)
11. Proposal:
   (a) If there are large Rashomon sets, let experts choose among the simple ones from them
   (b) Show the entire Rashomon set to experts!
   (c) Paper: Exploring the whole Rashomon set of sparse Decision Trees. TreeFARMS returns all almost-optimal trees. (insert link).
   (d) TimberTrek: Visualizing all Tress
   (e) Exploring the Whole Rashomon Set of Sparse Decision Trees

2.3 Session: Machine Learning in Finance

2.3.1 Data Driven Security Selection for Wealth Management

1. Authors: Sikun Xu, Ali Hirsa, Miao Wang, Federico Klinkert

2. Paper?

Abstract

In this research we established a data driven system for wealth management. The wealth management system consists of four major modules: security selection, asset allocation, portfolio optimization and risk management. It utilizes a large variety of data including fund performance, alternative data, macroeconomic indicators, etc., to assist dynamic and robust portfolio decisions. We designed a dynamic, explainable and automated pipeline of machine learning and deep learning models to process the large amount of hetero-structured data. We tested our system in the U.S. mutual fund market and in comparison to traditional wealth management methodologies, we are able to achieve superior performance.

3. Wealth Management How do you invest $100k? To answer this question, we need to answer other questions.

   what is your risk-return expectation what are the available assets, etc.

   Focus for the talk on mutual funds

4. There are five components

   security clustering security selection (*) asset allocation portfolio construction risk management

5. Focus on security selection

   Goal: select the top-performing funds within an asset category.

   Input/data: macroeconomic indicators mutual fund performance mutual fund alternative data (e.g., who is the fund's manager, location, etc.?)

6. Framework

   embedding dimension reduction

   forecasts times series

   Explanation: we need to understand black-box predictions

   ranking

7. training process is different:

   Expanding-window training process instead of traditional cross-validation

8. embedding categories time-series dimension reduction apply PCA rolling to over 600 macroeconomic indicators The top 15 pricing components explain over 90% of the variance time embedding dimension reduction for cross-sectional features

9. forecast

   overlapping data samples

   takeaways reduce overlapping data use simpler models generate synthetic data to enrich the training set

10. Results not so good out of sample

11. Explainability Local Interpretable Model-Agnostic Explanation (LIME)

12. Dynamic Model adjustment Given a portfolio of predictors, we can dynamically take a weighted average of them based on their out-of-sample performance. Every n days (n is a hyper-parameter), update weights. Not clear how to choose n.

    Improve overall performance and stabilize forecasts

13. Regime detection

    Assume that a discrete number of regimes exist in the market, and the performance of mutual funds is similar within a regime.

    Use HMM with Gaussian distribution rolling window.

14. Regime evolution

    Gaussian distribution allows us to visualize the evolution of regimes under different macroeconomic situations.

15. Conclusion, end-to-end data-driven framework for selecting the top-performing funds

    automating the decision process for wealth management

2.4 Session: Black-box Optimization: Algorithms and Applications

2.4.1 Linewalker: Line Search for Black Box Derivative-Free Optimization and Surrogate Model Construction

1. Authors: Dimitri Papageorgiou et. al. ExxonMobil Research & Engineering

2. Paper: ?

Abstract

We present a simple, but effective sampling method for learning the extrema of a discrete approximation of a multi-dimensional function along a one-dimensional line segment of interest. The method does not rely on derivative information and the function to be learned can be a "black box" function that must be queried via simulation or other means. We assume that the underlying function being approximated is noise-free and smooth. However, the Algorithm can still be effective when the underlying function is non-differentiable and possibly discontinuous. Numerous examples are shown to illustrate the Algorithm's competitiveness and potential superiority relative to state-of-the-art methods like NOMAD and Bayesian optimization.

3. Given a smooth function
   (a) Goal, find global optimal and
   (b) create a good surrogate
   (c) applications
       1 D line search, a.k.a. learning rate

4. Addressing the elephant in the room who cares about 1D functions? Isn't this too limiting?

5. author first shows results, then details of the algorithm linewalker.

   author shows a galley of test functions nasty, non-convex functions also, non-smooth functions sanity check on the performance of algorithms!

6. Fixing a test bead: NOMAD, one of the best derivative-free optimization solvers author's Algorithm: linewalker Bayes Optimization work very well

7. Surrogate approximations

8. total absolute scaled error metric how much of the (total absolute) error of the Mediocre method is explained by the Alternative Method? Insert Link.

9. The Linewalker Algorithm

   (a) Visual explanation.

   (b) Build and approximate function.

   (c) Look at the extrema (minima) of the approximate function build a new approximation then choose the next extrema (very similar to Bayesian optimization)

   (d) after a number of samples, break out of the regime and picks the maximum

   (e) More details on the Algorithm

      i. again, start with an approximation

      ii. pick a point

         A. TABU structure:
         do not revisit recently-sampled regions too frequently

         B. explore (before it was always exploiting)

         C. sample slightly away from an extremum
         This exploits the structure of the 1D case, where it is easy to come up with sampling a bit to the left or to the right.

10. Extensions

    Handling Noisy functions. Still a work in progress.

11. Questions: do we actually want to find the global optima of some of the nasty functions on the test bead? It depends on the application!

2.4.2 Branch-and-Model: A Derivative-free Global Optimization Algorithm

1. Authors: Kaiwen Ma et al. Carnegie Mellon University

2. Paper: ?

Abstract

Derivative-free optimization (DFO) is an important class of optimization algorithms that solve problems based on objective and function evaluations. DFO methods have enormous practical potential to address problems where derivatives are unavailable, unreliable, or only available at a significant cost. In this work, we present a novel derivative-free global algorithm Branch-and-Model (BAM). The BAM algorithm utilizes a flexible partition scheme and model-based search techniques, which exploit the local trend and speed up the convergence in solution refinement. The BAM algorithm is guaranteed to converge to the globally optimal function value under mild assumptions. Extensive computational experiments over 500 publicly open-source test problems show that BAM outperforms state-of-the-art DFO algorithms, especially for higher-dimension problems.

3. Black-box (Data-Drive) Optimization. Why?

   (a) Simulators

   (b) Experimental facilities

   (c) Hyperparameter tuning

4. Setting: we have a box that we can query without much more information.

5. Timeline on work around black-box optimization

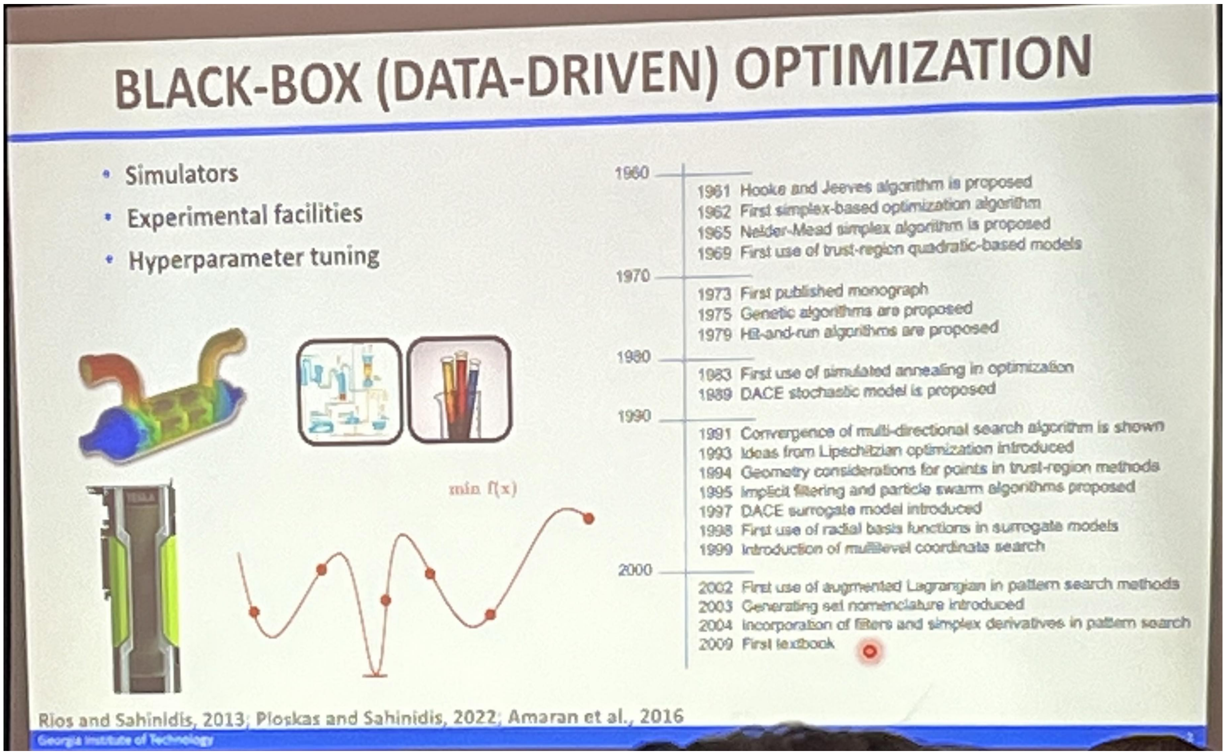

6. Experimental setup
   (a) Limit of 2500 function evaluations
   (b) solved if within 1% of global optimum
   (c) 10 random starting points for each problem
   (d) run through algebraic solver to get actual optimum

   Look at the fraction of problems solved as the probability that the Algorithm finds the optimal solution.

7. Global BBO algorithms invoke a search element to escape from local optima

   DIRECT (Jones et al. 1993) How to subdivide? Where to sample? Partition search space into boxes. Sample "middle of the box" Decide which box to sample next. We need to ensure the boxes shrink in every dimension.

   This framework works in simulations but not in experiments.

8. Branch and Model (as opposed to Branch and model)

   Start with a Latin Hypercube Sampling (LHS) We don't want a priori search boxes. We want boxes to be dynamic and use representative given the data. Project each sample point into its axis. You can show this will yield one point per box.

   Assume that there is a Lipschitz kind of condition within each box. Now, pick a box whose lower bound is below all other boxes.

9. the final Algorithm, BAM, performs well in different benchmarks with respect to the different experimental criteria. Scales to start solving problems around 30 dimensions.

10. Other resources, look at ALAMO algorithm to exploit local trends BARON)best subset selection).

11. Questions: constraints? No constraints other than penalizing the objective function.

2.5 Session: Bayesian Learning under Strategic Interactions

2.5.1 Learning to Lend Under Adverse Selection

1. Authors: Yifan Feng, Jussi Keppo, National University of Singapore, Singapore, Singapore.

2. Paper: ?

Abstract

We consider a dynamic pricing problem for a lender who repeatedly interacts with a borrower who has private information about his own default probability and strategically decides whether to accept the loan offers. We show that if the lender can commit to a simple markup policy, then asymptotically the information rent can be entirely offset by the benefit of learning. In addition, if the lender is sufficiently patient, she always lends the first loan at a low rate, which contrasts the implications of a static Akerlof-type model.

3. Personal Loans via Big data The missing data challenge Not all data is available We can link missing data to the strategic behavior of participants

4. Example Akerlof (1970) Information asymmetry leading to information rent and market inefficiency Sellers and buyers Only sellers know the quality At equilibrium, high-quality-product sellers are driven out of the market. This is known as adverse selection without information assymetry the market could be more efficient. This work is about overcoming adverse selection

5. Model in the context of lending Borrow has private information $\theta \in \theta_L \; \theta_H$, low is good, high is bad (high chance of default).

   Borrower Borrow money for a project Probability fail p_θ If it fail, the borrow defaults Otherwise, the borrower pas in full at maturity

   $v_\theta(r)$: expected utility of the loaned project with interest rate.

   Assume existence of reservation interest rate (makes borrow indifferent between borrowing or not r_θ

   Lender

   Only knows the distribution of borrowers.

   $u(r, x)$: utility of the loan for project with interest rate r and outcome $x \in$ (success, fail)

   assume reservation rate $\tilde{r}_\theta$

6. Key Assumptions

   Trading is efficient without information asymmetry. For every type θ, $r_\theta > \tilde{r}_\theta$

   Adverse election, $p_\theta > p_{\theta'}$ if and only if borrower's reservation price $r_\theta > r_{\theta'}$

7. Adverse Select and Market inefficient

   Suppose the fraction of high-type borrower is $b \in (0, 1)$.

   When b is sufficiently high, only high-type borrowers will rate.

   Hence, the market is inefficient; it drives low-type borrowers out of it.

8. Research question

   Can historical data be used to Mitigate market inefficiency Improve the lender's profit performance

9. Dynamic Model

   Lender interacts with Borrower repeatedly. Each time getting some noisy signals.

   Data Project: t Interest rate: r_t Accept? a_t no, $a_t = 0$ nothing happens yes, $a_t = 1$, then observe success or fails Repeat over time

10. Formulation The lender now has to think about a policy π: history $\mapsto r_t$.

    Utility: total discounted average utility (conditional on borrowers' type.

11. Lender wants to solve an optimization problem to find the optimal policy. Also, measure trading intensity as a metric for market efficiency.

12. they study a family of "low-then-high" markup policies.

13. At optimality, a markup policy quotes interest rates based on a threshold b^* such that

    Keep offering r_L first, the threshold optimally balances the type I/II errors.

    There are also asymptotic optimality properties of the optimal markup policy compared to the full-information policy, that is, the policy that knows exactly what price to quote (knows borrowers' types).

14. summary the markup policy can use to break the curse of adverse selection (asymptotically).

    Lenders value: achieve almost the full-information value up to a constant market efficiency: almost always trade up to a constant difference

2.6 Session: Negotiation Models

2.6.1 E-negotiation Model to Assess Offers and Select a Supplier in Agribusiness

1. Author: Over Manuel Causil, Danielle C. Costa Morais, Universidade Federal de Pernambuco

2. Paper ?

Abstract

The negotiation process covers the exchanging of offers and counteroffers between the negotiators. In order to aid this process of making offers and counteroffers, and evaluating tradeoffs among conflicting negotiation objects an electronic negotiation support system (e-NSS) can be used. We propose a negotiation model using an e-NSS applied in agribusiness to select a packaging supplier. Price, delivery time, and green manufacturing practices were used as negotiation objects. The negotiation process was carried out with five potential suppliers where a final compromise was achieved with a single provider.

3. Problem

    (a) Significant amount of water, energy, and material used to package food.
    (b) There are Lots of negative impacts to using plastic but necessary to preserve the quality of food.
    (c) Agri-food wants to find ways to minimize the negative impact caused by this industry.

4. Sustainable supplier selection.

5. Goal: An e-negotiation model to asses offers from different suppliers and select the provider that fits the sustainable supplier selection program.

6. Solution: Negotiation

    (a) Model
        i. Pre-negotiation: Negotiation issues
            A. Price
            B. Delivery Time
            C. Packages P of issues. Valuation functions as a linear combination of issues.
        ii. Negotiation. Two sub-stages.
            A. Pre-agreements

B. Comparisons of pre-agreements

iii. Post agreement

(b) FITRADEOFF (flexible and interactive tradeoff)

i. Can visualize the position of packages against other packages.

ii. Dynamic set of packages

A. Iteratively eliminate packages

B. Package P is proposed by a negotiator if proposed but rejected

2.7 Bayesian Optimization

2.7.1 Multi-step Budgeted Bayesian Optimization with Unknown Evaluation Costs

1. Author: Daniel Jiang et.al, Meta.

2. Paper

Abstract

Most Bayesian optimization algorithms ignore how evaluation costs, which are often unknown, may change over the optimization domain. An unknown cost function with a budget constraint introduces a new dimension to the exploration-exploitation trade-off, where learning about the cost incurs the cost itself. We propose a new dynamic programming-based acquisition function for this problem setting.

3. MDP formulation of the problem

4. Find policy that maximizes value of information between starting data set and final dataset

5. Challenges of MDP formulation
   (a) highly intractable
   (b) continuous
   (c) grows with the number of observed points so far
   (d) high-dimensional..

6. Approach: decision trees
   (a) at any t, simulate future (expand trees)
   (b) have a history representation of the future

7. Optimizing via differentiable tree

8. based on the notion of fantasizing from the GP
   (a) suppose we want to know the effect on our knowledge of measuring at t
   (b) sample "fantasy observation"
   (c) add fantasy observation to GP training data
   (d) each time you measure, you get a new "fantasy model"

9. The contribution here is a new acquisition function that plugs into regular BO.
   (a) Additionally, some tricks using linear algebra to speed up computation with fantasy GPs.

10. Drawbacks
    (a) decision trees do not scale with time horizon
    (b) requires a expensive re-optimization at each step

2.7.2 Bayesian Optimization for Heterogeneous Functions

1. Authors: Mohit Malu et al., Arizona State University

2. Paper: ?

Abstract

Bayesian Optimization with Gaussian Process prior typically assumes stationarity of the underlying function over the search space, but many real-world applications require optimizing non-stationary function. Non-stationary function can be considered as a set of stationary functions over the input space divided into multiple partitions with one class of stationary function in each partition. Often in control system setting we have access to class information along with the function evaluation. In this work, we propose a novel optimization technique Class-BO (Class Bayesian Optimization) for the non-stationary functions. We compare the empirical performance of Class-BO and show that it outperforms other non-stationary methods.

3. Usual goal, maximize black-box functions

4. Usual example, hyper-parameter optimization

5. BO sequential optimization strategy

   (a) statistical modeling - typically GP
   (b) Why? analytical tractability
   (c) intelligent sampling - using acquisition functions

   (a) default
      i. stationary kernel
      ii. performs poorly if function is non-stationary of heterogeneous
      iii. variation across input space

6. Motivation

   (a) Real-world problem requires to think about non-stationarity, that is, functions that are:
      i. Locally stationary
      ii. global non-stationary

7. The author proposes

   (a) Class-GP: to model the non-stationarity
   (b) Class-ICV new acquisition function

8. Model

   (a) X is a compact input space with p partitions
   (b) Each partition has a class label
   (c) Heterogeneous function $F : X \mapsto \mathbb{R}$, partition into a set of locally stationary functions
   (d) Each stationary function have the same class label and same level of variability

9. Observation Model

   (a) Query input point x and get back
      i. Function evaluation (can be noisy)
      ii. class label z
      iii. distance from closest boundary

(b) Initial data set consists of N samples, each sample as above

10. Approach
    (a) First, learning partition, then learn function in each partition.
    (b) Learning Partition.
    Using a tree algorithm that uses closest boundary information
    (c) Learning function in each partition
    new log marginal likelihood function
11. Classification Tree Algorithm
    (a) First, the best feature and split threshold are selected from closest boundary information
    (b) Second, use conventional CART algorithm to further grow the tree
12. GP Modeling
    (a) Assume GP with stationary kernel within each partition.
13. Experiments
    (a) Synthetic data generated by a sinusoidal function
    (b) MSE on 5000 test data points
    (c) MSE averaged over 50 runs for each set of parameters
    (d) the new methods perform better with more classes

2.7.3 Achieving Metric Diversity for Sample-efficient Search of Multiobjective Optimization Problems

1. Author: Eric Hans Lee, SigOpt
2. Paper

Abstract

Performing multi-objective optimization of important scientific applications such as materials design is becoming an increasingly important research topic. This is due largely to the high costs of said applications, and the resulting need for sample-efficient, multimetric optimization methods that efficiently explore the Pareto frontier to expose a promising set of design solutions. We propose moving away from using explicit optimization to identify the Pareto frontier and instead suggest searching for a diverse set of outcomes that satisfy user-specified performance criteria. This presents decision makers with a robust pool of promising designs and helps them better understand the space of good solutions. To achieve this outcome, we present the Likelihood of Metric Satisfaction acquisition function and demonstrate its viability on various problems.

3. Beyond the Pareto efficient frontier: constraint active search for multi-objective experiment design
4. Review of multi-objective optimization instead on one objective function, there are m functions
    (a) now we have to deal with the Pareto frontier
    (b) non improving property: cannot improve one function without degrading another
5. Objective enumerate the Pareto front

6. Motivating Questions is the mathematically optimal answer always the "best" for the user? what does the user think "best" means anyways? even if "best" == optimal, can we reliably identify it in practical solutions? probably no, unless there are strong assumptions

7. Constraint Active search. Overview

   (a) Application Goal: through simulations, find glass designs with low reflectance and high transparency

   (b) First try BO. Many Challenges

      i. Limited physical precision
      ii. points you thought were Pareto optimal in theory, are not in practice
      iii. Auxiliary objectives
      iv. more constraints are needed to address issues such as:heat resistance anti-fogging, etc.
      v. these are not feasible to evaluate on simulations

   (c) To address these... constraint active search! think of it as a search problem instead of pure optimization

      i. deals with limited physical precision by finding designs as distinct from each other as possible.
      ii. deals with auxiliary by distribution points evenly in a satisfactory region maximize chance that at the end you have something that works!

   (d) BO produces points too cluster together, so chances of having a feasible design are low The proposed method founds more disperse points so more chances of finding a good solution

8. Technical Details Same procedure as regular BO but for multiple objective Multiple GPs acquisition function evaluate, update model

9. New Acquisition function expected coverage improvement likelihood of metric satisfaction (LMS)

10. LMS quantifies how likely a parameter x will improve diversity in metric space Bottom line sample feasible points that are faraway from other samples

11. Results Measuring success Fill distance (down is good) Number of satisfactory observations (up is good) Number of neighbors closer than radius r (down in good) Hyper-volume of observed Pareto Frontier (up is good)

2.8 Keynote: Modeling Systemic Risk in Supply-Demand Networks

1. David D. Yao

Abstract

Recent events (the pandemic, geo-political conflicts, climate change, etc) call for studies on systemic risk in supply-demand networks (SDNs). An SDN is a network with nodes (or "agents") representing resources with processing and/or storage capabilities and arcs representing their supply-demand relations. Systemic risks in the SDN arise from its interconnectedness, such that disruption (or "shock") at one node may quickly propagate to other nodes and possibly lead to a system-wide disaster. There are similarities to systemic risk in the financial system, but also fundamental differences. We will discuss how stochastic networks can play an essential role in modeling and analyzing systemic risk in the SDN, along with certain risk-hedging tools and other technologies such as digital twins and reinforcement learning.

2. What is systemic risk?, Why systemic risk?, Why supply and demand?

3. Recent events: pandemic, geo-political conflicts, etc., causing worldwide supply chain disruptions

4. What is NOT systemic risk?

(a) Yossi Sheffi and Barry C Lynn, “Systemic Supply Chain Risk”, NAE/Bridge Sep 2014 Examples: Japanese triple disaster, earthquake, tsunami, and radioactive release in 2011 Thaalind flood in 2011

(b) Systematic risk is not systemic risk Osadchiy, Gaur, and Seshadri: “Systematic Risk in S...” Systematic risk is used in fianance and refer more to correlation as opposed to non-systematic (idiosyncratic) risk: variance, with noise/forecast error

5. Examples of Systemic risk (pre-pandemic, 2014)

   (a) Yossi Sheffi and Barry C Lynn, “Systemic Supply Chain Risk”, NAE/Bridge Sep 2014
   (b) Capital and Credit Risks, 2008-09 financial crisis
   (c) geographic concentration
   (d) emergence of “super” suppliers (over concentration)
   (e) multi-tiered Supply Chains (supply chain gets complex, with many layers)

6. Outline Part 1 Yao’s research studying financial systemic risk. Part 2 stochastic processing network Part 3 mitigation strategies: risk hedging we cannot completely avoid all failures, can we mitigate the consequences?

7. Part 1

   (a) financial institutions knit a complex network if one party defaults, this can propagate through the network. literature: Eisnberg and Noe (2001), Cifuenter, Ferrucci, and ?
   (b) in addition to the network effect, there is a market effect financial institutions are directly interconnected via the market (even if not directly connected) the institutions might owned common assets
   (c) To study a crisis like 2008 crisis, study both effects above.

8. Stress test Darrell Duffie’s 10-by-10-by-10 proposal: 10 stressful scenarios, each applied to 10 systematically important financial intuitions along with 10..?

9. Part 2

   (a) stochastic processing network to develop a stress test tool.
   (b) take publicly available data from BIS or ECB
   (c) develop a stochastic network
   (d) try to bring the two effects: network effect and market effect

10. SDN supply demand network

    (a) SND as Stochastic Processing Network
    (b) A set of nodes: resources that provides a processing/storage/distribution function (e.g., factory, warehouse, transportation)
    (c) each node has a processing capacity: service rate
    (d) some nodes that interact with the outside world will have an exogenous demand: external input rate or arrival rate
    (e) a routing transition matrix $P = [p_{ij}]$ where p_{ij} rate of output from node i to j.

11. Key Idea

    (a) think of the interconnected resources or interdependent critical infrastructures as processing networks
    (b) focus on

i. performance under normal operating conditions
ii. performance degradation under some kind of extreme event (for example, natural disaster) estimate required time and resources for recovery

12. Mathematical formulation
    (a) Dynamics: Skorohod Problem
    (b) low maintenance model in terms of required data
    (c) it is dynamic
13. Performance Measures:
    (a) You can identify bottleneck and non-bottleneck nodes
    (b) Throughout
    (c) Congestion/inventory
    (d) Delay/response time
    (e) Recovery times (once you can measure, you can design)
14. Under normal conditions:
    (a) focus on steady-state, equilibrium (design for long-term operation)
    (b) but we are also interest in extreme conditions
15. Under extreme events, for example,
    (a) contagion dynamics
    (b) recovery times

2.9 Session: Learning and Optimization in Pricing

2.9.1 Policy Optimization Using Semi-parametric Models for Dynamic Pricing

1. Author: Mengxin Yu, Yongyi Guo, Jianqing Fan, Princeton University
2. Paper

Abstract

In this paper, we study the contextual dynamic pricing problem where the market value of a product is linear in some observed features plus some market noise (with unknown distribution). Products are sold one at a time, and only a binary response indicating the success or failure of a sale is observed. We propose a dynamic statistical learning and decision-making policy that combines semi-parametric estimation and online decision-making to minimize regret (maximize revenue). Under mild conditions, we show that for a market noise c.d.f. F with m-th order derivative, our policy achieves a sublinear regret. The upper bound is further reduced to sqrt root of T if F is super smooth whose Fourier transform decays exponentially. These upper bounds are close to the lower bound where F belongs to a parametric class.

3. Online shopping Use specific recommendation then revenue management. Statistical learning + Decision Making
4. Example: Uber Customer arrive according to the current conditions personal preference platform evaluates the price according to private preference finally, the platform received a signal: buy or no buy
5. Idea: $(context_t, decision_t, reward_t)$

6. Main goal: maximize cumulative reward (as opposed to instant reward) Hence, can't be myopic decision affects future as well take into account interaction with environment

7. Feature-based dynamic pricing context = features decision: price assume (unknown) market value feedback: $1_{p_t<v_t}$

   assume v_t is a linear function $v_t = \theta^T x_t + z_t$, where z_t is unknown market noise based on this assumptions: $y_T = 1$ w.p $1 - F(p_t\theta_0^T x_t)$ or 0 w.p $F(p_t\theta_0^T x_t)$

   $rev_t(p) = p(1-p)...$

8. So, what is the best decision we can make? and what is the benchmark?

$$\text{Oracle: } p_t^* \in \arg\max_p rev_t(p) = g(\theta_0^T x_t)$$

9. So, leverage past information plus current context to determine posted prices.

10. Main goal

    (a) decide $p_t = p_t(x_t, (x_i, p_i, y_i)_{i<t})$ for $t = 1, \ldots, T$, to minimize regret

    (b) Regret = expectation(oracle revenue - algorithm revenue)

    (c) Need to balance learning θ_0, F, with decision, adaptively optimize p_t

11. Statistical Learning

    Assume total time horizon is now (generalize this later)

    Two phases

    (a) Exploration

       i. Strategy 1: a steps
          i.i.d p_t uniformly over range $(0, B)$
          At the end of the exploration period, estimate θ_0, F, given data $(x_i, p_i, y_i)_{i<t}$

       ii. good guarantees with theoretical assumptions (enough exploration time)

       iii. but, we need non-asymptotic bounds in practice

       iv. get regret of O(a)

    (b) Exploitation: $T - a$ Strategy 2: given θ_0, F, offer price according to demand get regret $O(Ta^{-1})$

    Total regret $\leq a + Ta^{-1}$

12. What if we have unknown time horizon?

    Divide by episodes, each episode explore then exploit as above Double the exploit time in subsequent episodes

13. Main Result Non-asymptotic bound on regret over T.

2.9.2 Linear Contextual Dynamic Pricing

1. Author: Jianyu Xu, University of California Santa Barbara, Santa Barbara, CA,

2. Paper

Abstract

Feature-based dynamic pricing are formally studied as an online learning problem where a seller sets appropriate prices for a sequence of products (described by their features) on the fly and learns from the binary feedbacks ("Sold" if valuation ¿= price and "Not Sold" otherwise). We study this problem by making a linear use of these features. In specific, we consider the following two models: (a) a "linear valuation" problem where customers' valuations are a linear mapping of features adding iid noises, and (b) a "linear policy" problem where we are agnostic of the valuation mechanism and only aim at competing with the best linear pricing policy. For both of these two problems, we design algorithms with provable regret upper bounds and propose information-theoretic lower bounds under a variety of assumptions.

3. Problem setting
   (a) Online-fashion contextual pricing
   (b) At time $t = 1, 2, \ldots, T$
      i. Feature $x_t \in \mathbb{R}^d$ is revealed
      ii. A customer generates a private valuation y_t
      iii. The seller (we) propose a price v_t
      iv. customer makes a decision $1_t = 1[v_t \le y_t]$
      v. We receive a reward $r_t = v_t \cdot 1_t$
   (c) Main feature, v_t is private, not available to seller only decisions are available
   (d) Compare with Contextual Bandit.
   (e) Similarities online-learning process partial-information feedback interactive decision
   (f) Differences: infinite Actions and non-continuous reward half-space feedback if price x is accepted then any prices below x is accepted if price x is rejected then any price above x is rejected
   (g) Performance Metric: Regret Compared to oracle

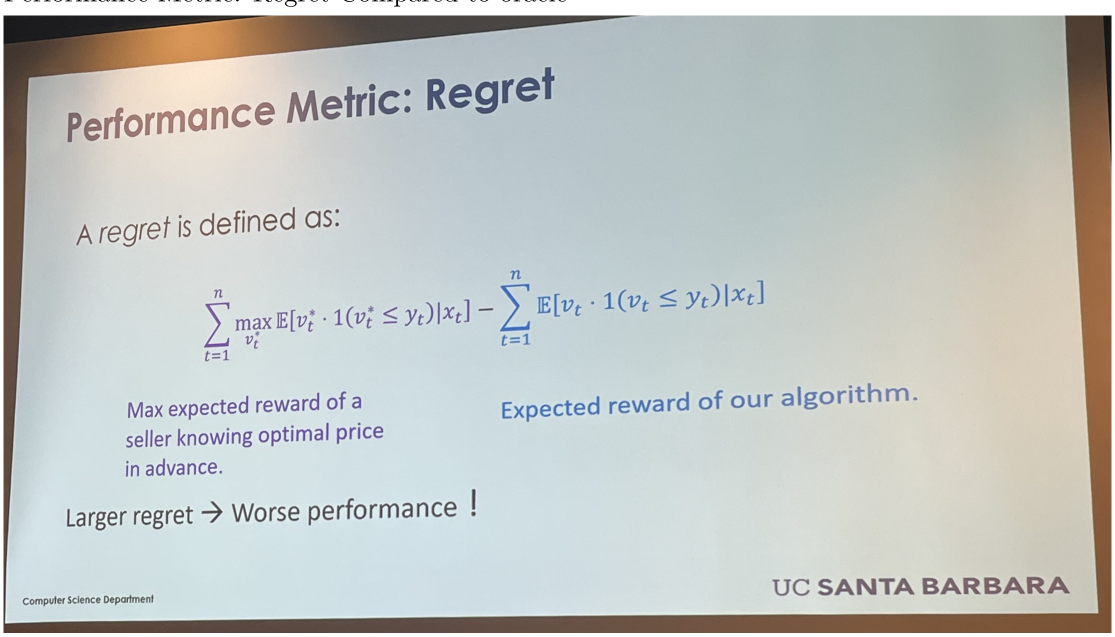

4. A linear Valuation Problem: assume customer's valuation y_t is
   (a) linear (on feature x_t)
   (b) noisy (added as N_t)

(c) formulation: $y_t = x_t\theta^* + N_t$

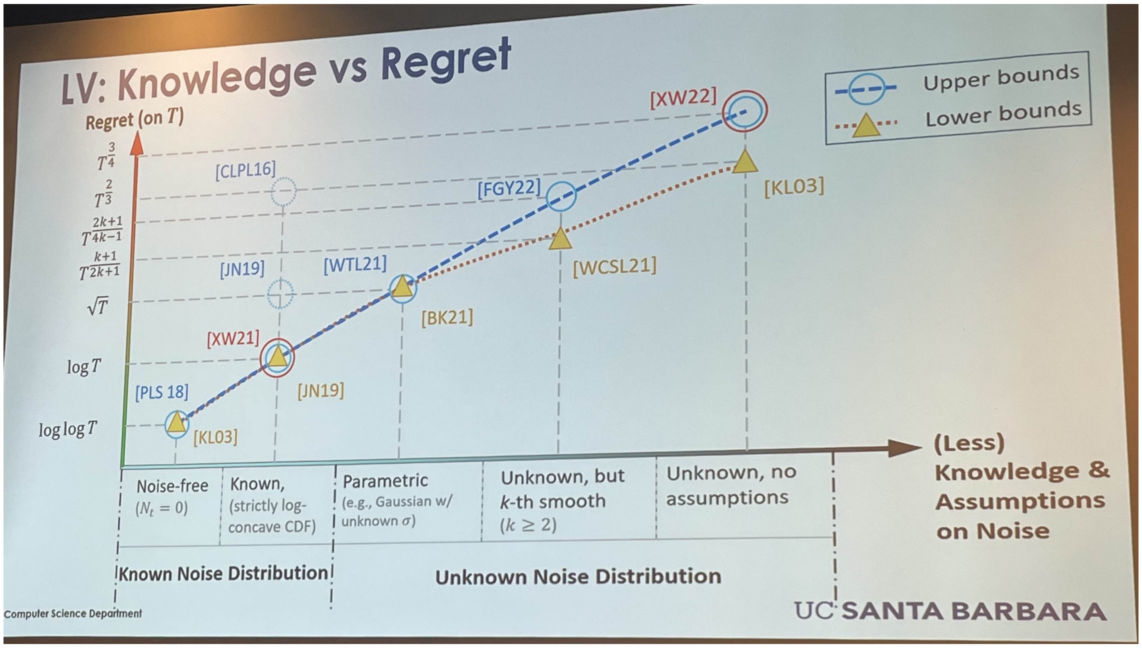

(a) Regret for known noise distributions

i. $O(d \log T)$

ii. Key is that we know the noise distribution

iii. Prove a $\Omega(\sqrt{T})$ lower bound for $Normal(0, \sigma^2)$ noise.

(b) Regret for Unknown Noise Distributions

i. $\tilde{O}(T^{3/4})$

ii. D2-EXP4 algorithm

iii. Half-space => lipschitz continuity

5. A linear policy Problem

(a) Agnostic valuation model

(b) unknown noise distribution

(c) no assumptions on valuation y_t at all!
not possible to learn anything

(d) change objective and instead of comparing to Oracle, compare with best linear pricing policies? (for some definition of linear policy)

(e) with these assumptions, achieve $\tilde{O}(T^{2/3}d^{1/3})$ regret

6. Conclusion

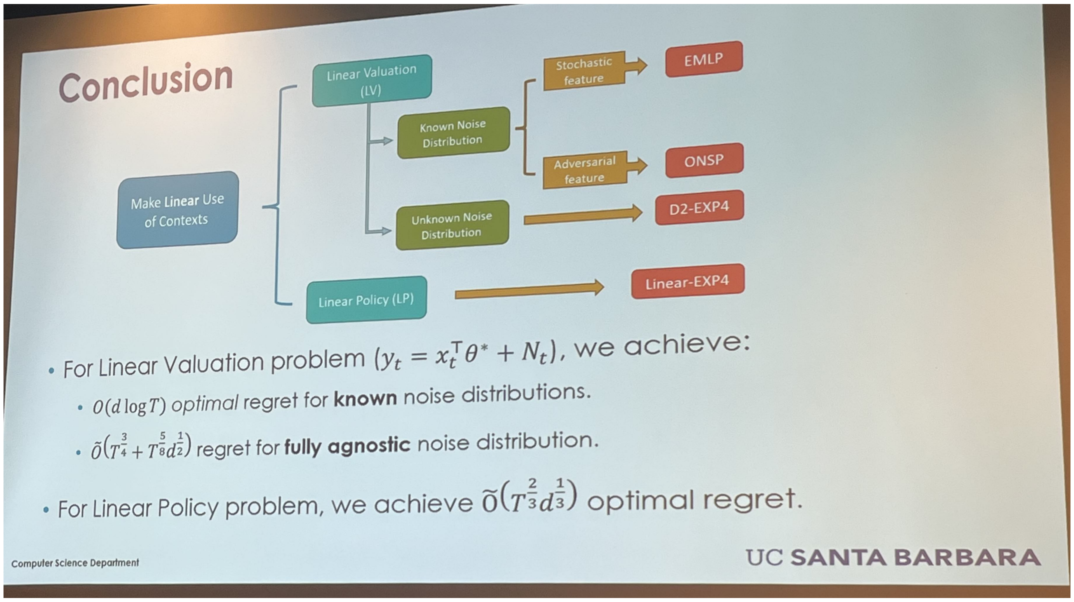
Conclusion
Make Linear Use of Contexts
Linear Valuation (LV)
Known Noise Distribution
Stochastic feature
EMLP
Adversarial feature
ONSP
Unknown Noise Distribution
D2-EXP4
Linear Policy (LP)
Linear-EXP4
For Linear Valuation problem ($y_t = x_t^\top \theta^* + N_t$), we achieve:
$O(d \log T)$ optimal regret for known noise distributions.
$\tilde{O}(T^{\frac{3}{4}} + T^{\frac{5}{8}} d^{\frac{1}{2}})$ regret for fully agnostic noise distribution.
For Linear Policy problem, we achieve $\tilde{O}(T^{\frac{2}{3}} d^{\frac{1}{3}})$ optimal regret.
Computer Science Department
UC SANTA BARBARA

3 Monday, October 17, 2022

3.1 Plenary: Maxine Bèdat. The life and death of your Jeans

Abstract

Maxine Bédat, founder and Director of think and do tank, New Standard Institute and author of Unraveled: The Life and Death of a Garment, an FT Business Book of the Year, will share the hidden world behind our clothing, highlighting key problems within our long supply chains, research that needs to be undertaken, and systems solutions that management science can champion.

1. Journey of our jeans
   (a) how they get to be
   (b) what happens once we get rid of them
2. In 6 units
   (a) Fiber
   (b) Yarn to Textile
   (c) Cut and Sew
   (d) Retail
   (e) Disposal
   (f) Future Research and Systems Solutions
3. Fashion industry perceived as fun and frivolous
   (a) this has delayed paying attention (from the OR point of view) to a $2.5 trillion industry
   (b) very energy intensive process
   (c) lots of inefficiencies - major contributor of gas emissions
4. Unit 1: fiber
   (a) Farm, Ranch, or more likely today in an oil well.
   chemical intensive agriculture
   impressive increase in output
   not without consequence
   health of farmers
   unhealthy soil
   (b) why not using all organic?
   three year process to certified field as organic
   cost goes up during this period
   no financial upside during this period
   no guarantee that the final outcome is positive
   incentive then to not switch
   also, organic doesn't exactly mean sustainable
   (c) Blue River Technology
   alternative for sustainable agriculture
   Using AI to reduce herbicide use by 90%.

(d) Since the 2000s, Polyester now dominates Cotton and Wool.
Rapid changes in fashion industry
Single piece of clothing might become outdated after a short period of time.

5. Unit 2: Yarn to Textile
   (a) To turn raw fiber into yarn and then into a colored textile that we will wear is a laborious process
   many steps in this process
   output is an anonymous amount of textile
   (b) This is by far the more intensive energy process
   76% of the carbon footprint
   eventually, textile industry could take 1/4 of all energy consumption!

6. Unit 3: Cut and Sew

   Inside the garment factory, we find horrible conditions

7. Unit 4: Retail

8. Unit 5: Disposal
   (a) How do get rid of our clothing?
   (b) what happens to the clothing after we get rid of? land-toss, plastic seal at the bottom until fill (20, 30 years) toady, clothing is not biodegradable, regardless of what it is made of
   (c) in the US throw away 80% of the clothing 80 pound per person per year

9. Unit 6: Future Research and Systems Solutions

   How to still benefit from the benefits of advances in textile technology while minimizing negative externalities?

3.2 Session: Public Sector OR: Community Resilience

3.2.1 Equitable Access to Public Transportation in Times of Covid-19

1. Author: Gabriela Gorgova-Svartzman

Abstract

Over the last couple of years, transportation services have experienced major impacts due to changes in demand, scarcity of drivers, differences in people's mobility patterns, and riders' hesitance to different mandates and policies related to the Covid-19 pandemic. The need for transportation has become different depending on the population around and the services available. This work proposes a framework for displacement and gentrification with a layer on transportation that explores the impacts of Covid-19 on mobility. Future work suggests how to re-think community resilience as we rebuild mobility in this "new normal".

2. Motivation: There was widespread fear in the public transit community that the pandemic will kill trust and reliability on public transport, and especially reduce access for underserved communities. Once stuff is shut down it will be hard to set it back up reliably.

3. Approach: Reserach was conducted in NYC since the survey structures in the city allow to get granularity on sub-district data.

4. Results from regression modeling demonstrated that it can be shown that in most regions gentrification has increased significantly over the last years. Despite efforts to monitor such a change, this is not visible by monitoring common factors like mean salaries in a district, ethnicity distributions, etc.

3.2.2 The Isolated Community Evacuation Problem for Response Purposes

1. Author: Klaas Fiete Krutein

Abstract

During responses to evacuation notices, decisions on resource allocation need to be made quickly. Frequently, information about the location and exact numbers of evacuees is incomplete. This is especially relevant for evacuations that require the coordination of evacuation resources. While the recently introduced ICEP has provided an evacuation planning tool for isolated areas, it relies on accurate data to provide a good solution. This research solves this problem through a robust optimization (R-ICEP) and a rolling-horizon optimization (RH-ICEP) variant of the ICEP. Computational results demonstrate the value of using evolving information which can help emergency coordinators to respond more efficiently to isolated community evacuations.

2. Motivation: More and more areas in the US and around the world are threatened by natural hazards. Particularly vulnerable are isolated areas like islands and mountain valleys that have limited evacuation routes.

3. Key question: how to evacuate the population of an isolated community as quick as possible if the number and distribution of evacuees is not exactly known?

4. Approaches are modifications of previous papers on deterministic versions of this problem

5. Approach one: cardinality-constrained robust optimization

6. Approach two: rolling-horizon optimization

7. Both approaches demonstrated superiority over deterministic variants and the rolling-horizon implementation is generally performing best.

8. Research demonstrated the value of incorporating evolving information in optimization problems, particularly in relation to disasters.

3.3 Session: Risk Behaviors

3.3.1 Dynamic Moral Hazard with Adverse Selection - A pontryagin Approach

1. Author: Fifan Zhang (Duke)

2. Paper: ?

3. Motivation agent (employee) supposed to exert effort and obtain results over time for a principal (a company) agent cares about its own welfare how to guarantee employees are working on the project

4. Examples R&D department funds research to generate 'breakthrough; results' Company hire sales agencies to attract new customers Firm hires lobbying agencies to influence politicians

5. In all these applications, two dimensions of uncertainty Dynamic Moral Hazard with Adverse Selection 1) unobservables efforts (form the point of view of employee) - Moral Hazard 2) unknown capabilities (operating cost) - Adverse Selection Based on observables (published papers, attracted customers, etc) need to design a contract to address these issues

6. Model

   (a) Customers arrive and brings R to the principal, Poisson process with arrival rate $v_t \in \{0, \mu\}$
   (b) Agent's effort process: $v, v_t \in \{0, \mu\}$. (No effort, no arrival of customers)
   (c) Effort cost: c per unit of time $c \in C = \{good, bad\}$, high-quality, low-quality agents, respectively
   (d) another model is continuous on agent's quality, $[c_{Low}, c_{High}]$

(e) principal reimburses the agent's operating cost (Limited Liability)

(f) Time discount r

(g) What can the principal do to incentives agents?

i. Contracts (contingent upon past arrival times)

ii. Payments, based on past performance (deliver n papers in x years and get y payment)

iii. Potentially termination time (get fired if not performing)

7. Goal Design a menu of contracts such that the agent truthfully reports the cost (Revelation principal) (under adverse selection)

8. Literature

(a) continuous time moral hazard

i. Brownian Motion [Sannikov (2008)]

ii. Poisson Process

A. Bad Arrivals [Myerson (2015)]

B. Good arrivals [Mason and Valimaki (2015)]

(b) dynamic moral hazard and adverse selection

i. Brownian motion [Cvitanic et al (2013)]

ii. Poisson model [Mayer (2020)]

(c) dynamic contracting in OR/OM

i. Zorc et al (2019)

9. Agent

(a) Agent's promised utility

i. if you are worker, you are thinking forward, that is, if I keep working on this project, what is my payoff in the future

ii. at each time working on the project, incur a cost

iii. utility: promised payment - cost

(b) contract design problem boils down to the design of the optimal future reward for agent

i. future reward for each arrival

ii. need to introduce an incentive compatibility constraint

10. Principal

(a) A simple contract

i. Naive principal thought: tell agent that every time there is a breakthrough, you are going to get a payment

A. this is not the optimal contract - the principal is not using all its available information

B. e.g.,not using termination

Principal's utility under a contract is the revenue under arrival minus payment to agents Principal looks for contracts that maximizes its utility subject to truth telling constraints

Ultimate game: design the menu of contracts so that agents self-select, that is, select the contract that matches their ability

11. Implementable Contracts

(a) sign-on-bonus contract

i. Payment B to the agent at time 0

ii. after time 0, the agent's utility starts from w and (IC) constraint is always binding
iii. the agent's utility diminishes over time, until it finds a new client or achieves a breakthrough
A. consider someone that is unlucky, works on a project for a long time, but fails to deliver
outcome: termination
B. if the agent keeps delivering, the model caps the utility at some upper-bound
outcome: no longer fear of termination

(b) probation contract (looks like an internship)

i. no payment at time 0
ii. before the first arrival, a shorter probation period ((IC) is not binding)
iii. after the first arrival, binding contract
sorting out high from low types

12. Optimal Contracts (Bad agent is inefficient - not ok to fire)

when bad type is really bad (not even covering their cost) bad type we want them out of the system, but will they leave? pay-to-leave contract information rent paid to the bad agent

good agent probation contract a shorter probation period weeds out the bad agents

13. Optimal Contracts (Bad agent is efficient - ok to hire)

(a) bad agent should select sign-on-bonus
(b) good agent should select probation contract
(c) why sign-on-bonus?
i. we want bad agent to work
ii. we want to make the bad agent's contract less appealing

14. Optimal Contracts (Continuum-type) (not enough time in the presentation to cover this part)

15. Conclusion

(a) business need long term contracts to manage incentive over time
(b) there are easy to implement and compute contracts
i. two-types
ii. continuous-type

3.4 Session: Matchmaking in Two-sided Marketplaces

3.4.1 Capacity Planning in Stable Matching: An Application to School Choice

1. Authors: Ignacio Rios et. al. The University of Texas at Dallas

2. Paper: ??

3. Centralized assignment mechanism. Use widely, for example

(a) School choice
(b) refugee resettlement
(c) job markets

4. key properties

(a) stability
(b) crucial for long-term sustainability
(c) eliminate justifiable envy

5. basic assumption: fixed capacities but this is not always the case.

6. How to jointly decide capacities and the assignment?

   (a) with fixed capacities: use deferred acceptance algorithm

   (b) Now, suppose there is an additional seat (say school assignment problem)

      i. to which school give the extra assignment? example shows that this is not trivial, introduces a tradeoff

         A. improvement (original allocation of a student improved)

         B. vs access (originally unmatched, now matched)

7. Contributions Problem Formulation Solution Approach exact and heuristic methods empirical application, data from Chile

8. Model Input Schools $c \in C$, priorities $>_c$, initial seats q_c budget B (overall extra seats) Students $s \in S$, preferences $>_s$, valuations $r_{s,c}$, penalty $r_{s,\emptyset}$ (cost for unassigned students)

9. Decisions $x_{s,c} = 1$ if student s assigned to c. $t_c =$ seats added to school $c \in C$

10. Constraints Matching definition Stability constraints Budget

11. Objective obtain a budget allocation and a stable matching to minimize the sum of preference of assignments. (taking penalty into account)

12. Observation, the ensuing optimization problem is computationally hard.

13. What is the effect $r_{s,\emptyset}$? low $r_{s,\emptyset}$, then improve the assignment of as many students as possible. high $r_{s,\emptyset}$, then we want to find schools that guarantees more access.

    Theorem 1, if $r_{s,\emptyset}$ is sufficiently low, then the number of students who improve...

    Theorem 2, if $r_{s,\emptyset}$ is sufficiently large, then we obtain maximum cardinality student optimal stable matching

14. Incentives Is in the best interest of students to report their preferences truthfully? depends on the information student have on extra seats the paper consider the ex-ante case, students report before knowing the budget for extra seats

    Theorem 3: strategy-proof in the ex-ante and ex-post Theorem 4: the mechanism is not strategy-proof for students if they have interim knowledge (know the budget but not where it is allocated)

    The resulting paper results in a non-linear formulation, idea use linearization McCormick envelop

15. Heuristics linearization might take too long. So, use heuristics. Greedy Update one seat at the time maximizing marginal improvement LPH obtain optimal capacity decisions relaxing stability, and then apply DA

16. Empirical Application School Choice in Chile this is a good setting for this problem, Ministry of education trying to find a good assignment and can expand seats in school data from the Magallanes region, south most region perform simulations varying budget and penalty measures of improvement students that "enter" the match access students that "improve" their match fairness

3.4.2 Learning Equilibria in Matching Markets from Bandit Feedback

1. Authors: Alexander Wei et al. UC Berkeley (Michael Jordan group)

2. Paper: Appeared at NeurIPS 2021. (insert link)

3. Think about Data-driven matching platforms like Uber, tinder, taskrabbit, doordash, airbnb, upwork, etc. Learn to match from data without knowing fully agent preferences.

4. Example ride sharing platform two-sided, riders and drivers Have heterogeneous preferences riders may prefer closer drives drivers may prefer certain destinations

   platform proposes market outcome: matching + $ transfer

5. Challenges preferences must be learning from data agents must be matched online at large scale suggested market outcomes should be incentive-aware to retain users, platform should offer desirable matches at fair prices

6. Model we need to learn agents' preferences from repeated feedback/interaction with the platform so, propose match, receive feedback, repeat

7. Two-sided matching with transfers (Shapley, Shubik)

   Customers, Provider Market outcome = (matching, Transfers)

   assume utilities of customer C for matching with providers P

8. Contribution of paper incentive-aware learning framework for matching market algorithms design simple+optimal no-regret algorithms for learning stable outcomes

9. Existing approaches

   classical economic theory focus on known agent preferences deferred acceptance (etc)

   classical bandit theory combinatorial bandits maximizes total welfare of the matching however, unaware of incentives (e.g., price-setting)

   The idea is to merge these two things into one framework

10. Framework Set of I customers Set of J of providers

    Market outcome (μ, τ) Matching μ transfers τ (positive if receive payment, negative if they have to pay)

11. Matching with transfers — Stability No blocking pairs no pair of agents can renegotiate and match outside the proposed matching Individual rationality

    Transfers $->$ stable outcome maximize social welfare (Shapley-Shubik 1971) but this doesn't take transfers into account

12. Learning and feedback Learning takes place over T rounds. In the t-th round

    (a) A subset of Agents arrive
    (b) Platform selects outcome
    (c) platform sees noisy payoff of each agent (bandit feedback)
    (d) platform incurs loss given by a measure of instability

    Regret is the sum over instabilities over all time periods

13. Why measure instability?

    (a) stability is a binary notion: unsuitable for optimization!
        with uncertainty, impossible to guarantee exact stability
        can only hope to optimize for "approximate" stability
    (b) Goal: design a measure of distance from stability that is:
        tractable
        economically meaningful
        quantifying lack of stability via subset instability

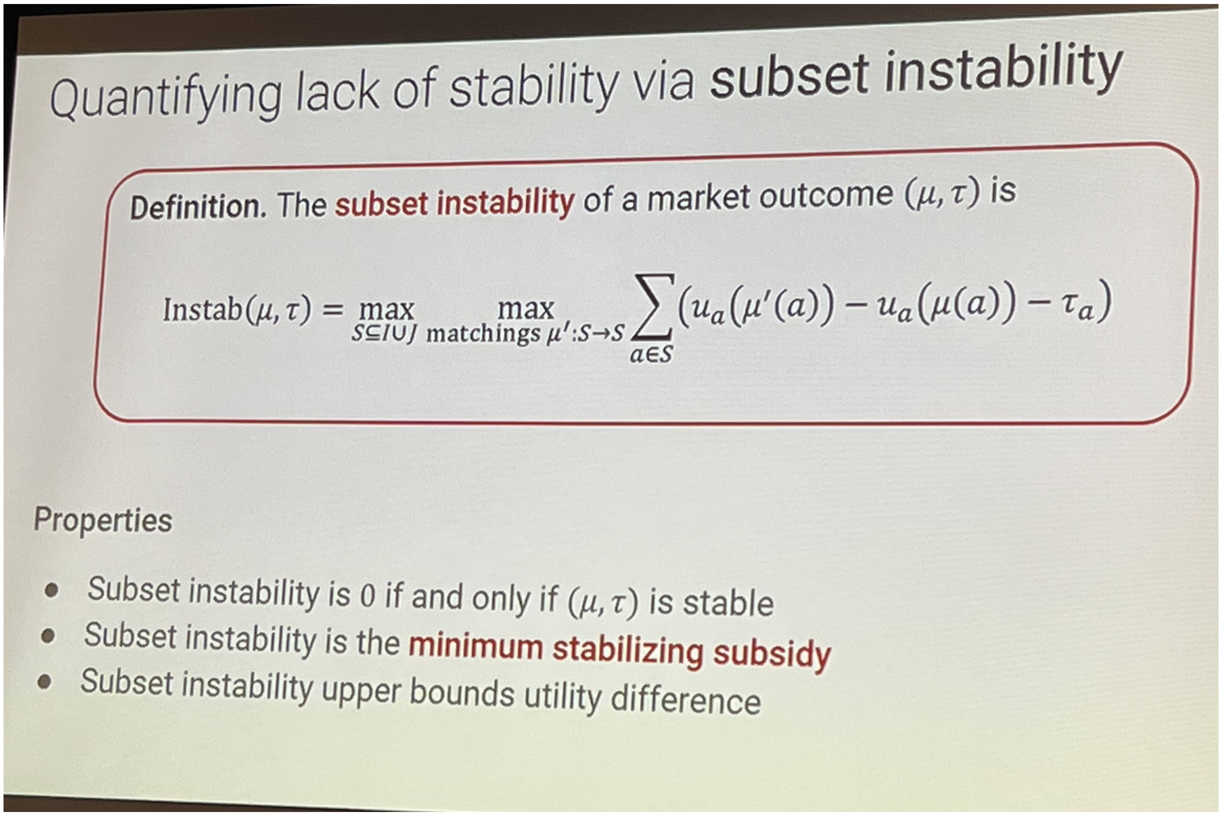

14. Results: Theorem: there exists an algorithm with $N^{3/2}T^{1/2}$ instance-independent regret with N agents over T round

15. Algorithm MatchUCB

    (a) Idea: maintain confidence intervals over all utilities that contain the true utilities w.h.p.
    Then, at each round

    (b) compute UCB estimates for all utilities

    (c) choose market outcome w.r.t upper-bound utilities
    Lemma: instability is bounded by length of confidence intervals.

16. Preference Structure for arbitrary preferences, regret grows super-linearly in the size N of the market when can we do better?

    typed preferences each agent belongs to one of finitely many types linear preferences utilities are linear.

17. Next Steps stable matching in more complicated environments)e.g., changing utilities, adversarial settings? more general market equilibria? competitive equilibria?

3.5 Session: Optimizing Matchmaking in Platforms

3.5.1 The Cost of Impatience in Dynamic Matching

1. Author: Angela Kohlenberg, Northwestern University

2. Trade-off in dynamic matching

    (a) match immediately, as agents become available

    (b) wait, maybe more optimal matches, but at the cost of waiting and potential abandonment, that is, agents leave the market

(c) how to balance these two?

3. How does impatience impact the quality of the market? Identify optimal strategy when agents are impatient.

4. Research focus How do we optimally match heterogeneous agents with finite patience?

5. Model

   (a) Three agent types, Two matches
   (b) Agents arrives as Poisson process $\lambda_a, \lambda_b, \lambda_s$
   (c) Serve can match
   (d) agents and server will abandon queue with exponential rates
   (e) Objective: maximize long-run average reward
   (f) Matching policy π: specifies matches to perform, based on the current number in each queue.

6. Cost of Impatience: $\bar{v} - v^{\pi}$

   (a) Long-run average:
      i. $\bar{v}$: with no abandonment
      ii. v^{π} from a policy π
      iii. v^{*} optimal policy π^{*}

7. Results

   (a) cost of impatience scaling
      i. lower bound on $\bar{v} - v^{\pi}$ as a function of the network parameters
   (b) different policies in different regimes
      i. impatient server regime: greedy policy
      ii. low load regime: threshold policy
      iii. high load regime: dedicated policy

3.5.2 Optimizing Free-to-play Multiplayer Games with Premium Subscription

1. Author: Yunke Mai, University of Kentucky

2. Paper: ?

3. Focus on video games industry. Bigger than music and films.

4. Evolving business models

   (a) Retail model (1970 -)
   (b) Game as a service (2010 -)
   (c) Free to play (2018 -)
      i. free-to-play
      ii. sales of virtual items

5. Focus on free-to-play multiplayer games for example, apex, fortnite, league legends, etc.

6. Sales of virtual items direct selling fee for item/bundle
   loot box
   premium subscription flat free to unlock all virtual items offered during a period of time

7. Premium subscription ex fornit's battle pass (2018) \$1.8 billion revenue cost only \$10
8. Motivations for paying virtual items direct satisfaction improve performance/appearance

   social comparison video games as virtual social environments downward comparisons: owner feels superior to non-owners upward comparison: the non-owners feel inferior to owners
9. Paper consider F2P online multiplayer game with premium subscriptions

   model social comparisons

   characterize the optimal dynamic advertising and pricing policies
10. Model
    (a) active player base $n \in [0, 1]$
    (b) players derive a base utility v (per unit of time)
    (c) premium subscription $s \in [0, 1]$ proportion of subscribed players
    (d) subscription price p
    (e) when a premium player sees a free player: additional utility $\alpha_d > 0$ (downward)
11. Analysis get proportion of player that will subscribe at equilibrium
12. Player base growth model a hybrid model of the Bass diffusion model and the replicator equation
13. Optimal advertising and pricing strategies
    (a) First, game is totally free, invest heavily on advertising
    (b) Once a sizable player base is established
        i. decrease advertising
        ii. introduce the premium plan
        iii. start with a high price, then decrease (more premium player is going to hurt everybody)

3.6 Session: Learning and Inference of Preferences

3.6.1 Learning Stochastically Revealed Preference

1. Author: Chunlin Sun et al., Stanford University
2. Paper: ?
3. Goal: infer the unknown preference of a customer through the observations of this customer's choices
4. This goal has a long history in economics (Varian, 2006)
   (a) most past work assumes that there is a single vector consistent with all observations
   (b) however, in practice, utility changes
   (c) many methods for the deterministic case fail for the stochastic case (not surprisingly)
5. Applications
   (a) investor belief analysis
   (b) market analysis
   (c) pricing
   (d) etc

6. Contributions

   explore the revealed preference problem in a stochastic setting with a linear utility function new standard to measure the performance of algorithms in a stochastic setting Bayesian methods that provably achieve a sub-linear prediction error we can extend methods to inverse optimization problems

7. Model

   (a) Customer Assume the customer is rational, The utility is captured by the following LP:

   $$LP(u, a, b) = max_x \sum_{i=1}^{n} u_i x_i, s.t., \sum_{i=1}^{n} a_i x_i \leq b, 0 \leq x_1 \leq 1, i = 1, \ldots, n$$

   where (u, a, b) is sampled from some distribution.
   The action will be x^*, which is the optimal solution to the LP.
   Assume the true distribution of the utility is in the set $\{\mathcal{P}_{u,\theta}\}_{\theta\in\Theta}$ parameterized by θ, and denote θ^* as the true hyperparameter.

   (b) Learner
      i. Known sample size T
      ii. i.i.d observations $\mathcal{D}_T = \{(x_t, a_t, b_t)\}$, u_t's are hidden!
      iii. known the distribution of the price and budget pair

   (c) Goal: find hyperparameters to fit utilities

8. Bayesian Method uses samples to update our belief of the real hyperparameter

9. Theorem

   Under some assumptions on the probability distributions, we have a w.p. bound on the posterior distribution Posterior distribution concentrates over the true hyperparameter.

10. Extensions to the case when the distribution is not perfectly known. For this case, there is an approximate Bayesian Method

11. Take-ways

    (a) The stochastic setting is important for the revealed preference problem preference can change
    (b) the Bayesian method applicable in this setting
    (c) We can still use the approximate Bayes method if a small data amount is polluted.
    (d) Bayesian method still suffers from the curse of dimensionality.

3.7 Tutorial: Yao Zhao. Supply Chain Analytics from problem solving to problem discovery

1. "If you define the problem correctly you almost have the solution" Steve Jobs

2. Descriptive, Diagnostic, Predictive, Prescriptive (all Analytics)

3. Usually, most of the time is taken to teach Prescriptive, But in reality, the problem is never well-defined.

4. (Insert Albert Einstein quote 55 minutes to ask question, 5 to answer it)

5. "A problem well stated is a problem half solved" (Insert author)

6. "If you define the problem correctly you almost have the solution" Steve Jobs

7. Point of this tutorial

   A sub-optimal solution for the right problem is better than an optimal solution for the wrong problem.

8. Why is defining data so hard?

   Data Interpretation - What does the data mean? Transforming data into insights

9. Example 1: Compaq vs. Dell (2001)

   Compaq was the market leader healthy cash flow, many patents but Compaq sold itself, abandon its strong brand, why?

   The exercise here is to look at data to answer this question.

10. Example 2: Taiwan Semi conductor company to open 5nm Factory in Arizona

11. Example 3: Airlines "The worst sort of business is one that grows rapidly, requires capital to engender growth, and earns little to no money" warren buffet

12. Example 4: The health care sector

4 Tuesday, October 18, 2022

4.1 Session: Interpretable Machine Learning via Mixed-integer and Robust Optimization

4.1.1 Adaptive Robust Ensemble Modeling for Time Series Forecasting

1. Author: Leonard Boussioux, MIT

2. Paper:

Abstract

Time series forecasting plays a crucial role in a wide range of problems with a temporal component. Since time series data is prone to distribution shifts, a single predictive model's performance can vary significantly across time. Therefore, ensemble modeling proposes leveraging several available models to improve accuracy further. We contribute a new methodology for robust ensemble modeling of time series forecasting models. We develop an adaptive robust optimization (ARO) approach to formulate a linear regression ensemble where the models' weights change over time. We show the impact of our ensemble method on a range of real-world challenges, including tropical cyclone intensity forecasting, pollution management, and energy consumption forecasting, and hope to open the door for further use of ARO in machine learning.

3. Satellite data holds a lot of potential (Multi-modality)

4. Intensity of hurricanes is really hard to predict

   (a) Even the best models of the NHC regularly fail to predict rapid intensification of hurricanes

   (b) Adaptive Robust Optimization as a way to make methods more robust
       that is, using ensemble models
       take average of many available models

5. The principle of ensemble models

   (a) Given some data

   (b) Train many models, Trees, NN, etc.

   (c) Obtain many predictions for the same data

   (d) average all output's model

6. Time-Series Ensemble Forecasting Setting

(a) first idea: static weighted average of models' predictions
(b) second idea: dynamic weighted models.
consider this problem under robust optimization

7. Each model comes with an uncertainty

   Weight models as a function of their errors and uncertainty

8. Robust Ensemble Formulation

$$\min_{\beta(\Delta)} \max_{\Delta} ||y - (X + \Delta)\beta(\Delta)||_2$$

   under a certain constraint set

   this formulation is too hard to solve directly, so use an equivalent formulation using adaptive ridge regression

9. the formulation performs much better than extant models

10. Leo Breiman (creator of trees) quote on how trees are unstable the methods in these paper can help with this issue

4.1.2 Slowly Varying Machine Learning

1. Author: Vassilis Digalakis, MIT

2. Paper

Abstract

We introduce the framework of slowly varying machine learning, which aims at building machine learning that vary slowly over time, or space, or any other dimension. In the first half of the talk, we consider the problem of parameter estimation in slowly varying regression models with sparsity constraints, and make both theoretical and algorithmic contributions. The second half of the talk is focused on slowly varying classification trees.

Work based on two papers by Vassilis, insert links

Slowly varying classification trees Slowly varying regression under sparsity

What is slowly varying regression?

Example: smart building we want to predict energy consumption collect features, e.g. whether conditions temperature, humidity inside the building bedroom office estimate sparse linear regression model say the most important feature is the temperature in the office

however, you might spend more time in the office during the day makes sense this is the most important features but not at night this is the point: want to predict coefficients that can vary over time

Impact Estimate spares LR models where regression coefficients vary slowly and sparsely under graph-based temporal or spatial structure factor models in stock market over consecutive periods meteorological models over adjacent spatial ares voting patterns in the US electorate over adjacent spatial ares. code available

How do we formulate SSVR and what are the challenges? Main problem: non-convexity of objective functions

Setting Assume a feature matrix and outcome vector for each time step, $X^1, \ldots, X^T$ and $y^1 \ldots y^T$ Assume a similarity graph where nodes are connected if they are similar

Mixed-Integer Optimization (MIO) Formulation

$$\min_{\beta^1,\ldots\beta^T} \sum_t ||y^t - X^t\beta^t||_2 + \text{robustness term} + \text{slow variation term}$$

s.t.,

(a) local sparsity
(b) global sparsity
(c) sparsely varying support

this is a combinatorial problem, solved by introducing binary variable $z : \beta = 0$ if $z = 0$

replace β with $z \times \beta$

with this, all constraints can be linearized

3. Literature
   (a) Sparsity: Hazimeh et al (2021)
   (b) Slow Variation: Hazle and Tibh (1993)
   (c) Sparse variation: Blakley and vert (2011)
   (d) This work combines all three.

4. how do we overcome non-convexity?

   After some transformation, the inner minimization can be solved in closed-form

5. how do we solve the problem?
   (a) Solve to optimality using the outer approximation method
   (b) Iteratively tighten piecewise linear lower approximation of objective function

6. separable heuristic algorithm
   (a) to obtain good starting points
   (b) fit separate univariate regressions per vertex per feature
   (c) approximate slow variation penalty with new regularization term that depends on degree of each vertex

4.2 Plenary: Jaillet Patrick (MIT) Online Optimization and Learning for Sequential Decision-Making

Abstract

resources allocated with incomplete knowledge of the future. It is not clear in this setting how to measure the quality of a proposed decision strategy. Online optimization compares the performance of a strategy that operates with no knowledge of the future (on-line) with the performance of an optimal strategy that has complete knowledge of the future (off-line). In some cases, probabilistic information about the future may be available or learned. In this talk, we provide an overview of some results obtained from that perspective on various classical sequential decision-making problems.

1. Typically, from the perspective of the supplier, you see customer coming one at a time and you have to make a decision upon arrival.

2. Dynamic market can be modeled as a graph with supply on one side and demand on the other but again, demand has online arrivals

3. Making sequential decision with uncertainty commit resources now on the face of an uncertain future

4. fundamental question how to evaluate strategies operating under uncertainty? probability models =¿ stochastic optimization no probabilities assumptions =¿ online optimization a bit of both =¿ online optimization and learning the last one is the focus on the plenary

5. online concepts
   (a) online optimization problem = instance incrementally revealed over time
   (b) online algorithm: the quality measured by its competitive ratio measure against clairvoyant algorithm with full knowledge of the instance

$$\text{ratio, } c = \inf_{\text{instances } I} ALG(I)/OPT(I)$$

   (c) in the learning case, one uses the notion of regret instead of competitive ratio
   (d) combining the two can lead to better algorithms
6. Outline
   (a) First Part: Online Resource Allocation Problems with Partial Learnability
   (b) Second Part: Online Resource Allocation Problems with Samples
   (c) Third Part: Online Linear Programming without Knowing the Horizon
   (d) Fourth Part: Universal Online Learning and Regression
7. First Part:
   (a) Context
      i. The simplest basic online resource allocation
      ii. A firm with b identical units of a product to sell over n periods
      iii. At each period at most one customer arrives demanding one unit
      iv. Two types of customers
         A. type 1: each willing to pay $v = 1$
         B. type 2: each willing to pay $v = a$, where $0 < a < 1$
      v. at each period, decided to allocate or not.
   (b) What can the firm achieve
      i. Offline solution (OPT): select b customers with the highest value.
      ii. Online solution: depends on what the firm knows about arrivals what ware we willing to assume? adversarial approach: no information stochastic approach can learn patterns of future arrivals
      iii. Comparing the two approaches Adversarial $\frac{1}{2-a}OPT$ Stochastic $[1 - O(1/\sqrt{b})]OPT$
      iv. Example suppose $a = 0.75$. Adversarial only guarantees 80% of hindsight revenue.
      v. Limitations: Adversarial: results in too conservative policies Stochastic: demand has unpredictable components that cannot be learned Shamsi et al.(2014l Esfandiar et al (2015))
      vi. Author introduces new model that has both adversarial and stochastic components can be partially learned
      vii. design adaptive and non-adaptive dynamic thresholds.
8. A partially learnable model
   (a) Adversary determines the initial sequence of demand arrivals $I' = (v'_1, \ldots, v'_n), (v'_i = 0)$ means no customer
   (b) stochastic group : a subset of customers does not follow this sequence
      i. customer belong to that group independently with probability p
      ii. customers in that group arrive at "random" times

   Other groups => adversarial

9. The algorithms the non-adaptive algorithm use two thresholds: one for stochastic group, one for adversarial group the adaptive algorithm. use info of what we have seen so far to estimate number of types of customers arriving in the future

10. Second Part:

    (a) Context Almost same problem as before, but when a customer comes, we don't know its willingness to pay Two types of customers type 1 with unknown expected reward r_1 type 2 with unknown expected reward r_2 We still want to maximize reward

    (b) Model discussion if r_1 and r_2 are both know, go to part 1 Otherwise, adversarial setting: no hope to get meaningful results with online algorithm, need more info what is the minimum info we would need to get meaningful results?

    (c) First a tested period to learn types, then use this information

    (d) Algorithm design Main idea: use the customers in the sample to learn the expected rewards of the two types and information about the initial sequence

        Sampling Calculate the empirical expected reward Make a decision based on which type have higher empirical expected reward

    (e) Results As you get more resources to learn at the beginning of the algorithm, the algorithm perform better compared to not using the information

    (f) Takeaway from first two parts even limited stochastic info significantly improves performance for large b

11. Third Part: Motivation: online packing LP classical primal-dual LP assume that the columns of the constraint matrix arrive one at a time once a decision is made, we cannot go back and change it

    adversarial matrix decides the order of the column of the matrix this problem has been study in many communities [Kesselheim et al (2014)] [Agrawal et al (2014)] [Li et al (2002)]

    Question: what if the number of columns in unknown?

    Simplest case matrix coefficient is 1, $m = 1, a_{1t} = b_1 = 1$ reduces to the classical secretary problem Studied by many communities well know technique: reject first $\lceil n/e \rceil$ candidates, then accept the first candidate that is the best seen so far $p \approx 1/e \approx 0.37$ if we assume that n is random with known distribution [Presman and Sonin (1973)] much more complicated strategy even when you know exactly the distribution

    Fundamental questions given some constraints on the strategy space (say, only thresholds) how bad would the optimal strategy in this family compared against the optimal strategy?

    $$\sup_{\text{strategies}} \inf_{\text{distributions}} \text{Performance(strategy, distribution)}$$

    Results (1) Simple strategies are not too bad. Limiting to threshold looses between $1/e$ and 1. (2) assuming upper bound on n, there is an upperbound on the performance of the strategies if you have moment information, the result refines to take into the account the moment (3) Beyond worst case: most distributions are good the bad case is actually extremely rare, that is, distributions over which we get the worst performance are rare

12. Fourth Part:

    Universal Regression with Adversarial Responses, Blanchard et.al (2022)

    Question: what is learnable?

    Motivation and setup: supervised learning: prediction task with data samples online learning: data arrives in streams, predictions on the fly using data (kind of SL one at a time) can we characterize online learnability with provably minimal assumptions?

    Setup At each step t

(a) observe new instance X_t
(b) predict response $\hat{Y}_t$ using all historical data and X_t
(c) observe true response Y_t, incur loss $l(\hat{Y}_t, Y_t)$

Goal learn the best function $f^* : \mathcal{X} \mapsto \mathcal{Y}$ (so that $Y_t \approx f^*(X_t)$)

assume no noise: classical online learning Challenges: Noisy data, corrupted, or even adversarial "Best" f may be complex; which function class? how to construct robust algorithms?

Classical assumptions restrict input space X to be i.i.d, ergodic, stationary, law of large number restrict output space Y to be a restricted function class with regularity, low complexity, etc.

What do we mean by universal learning? no assumptions on responses

Main questions: what are the minimum assumption on input space sequence so that universal learnability is achievable? can we construct algorithms whenever possible?

4.3 Session: Pricing

4.3.1 Selling Hope with Uncertain Pricing

1. Author: Yongqin Lei, Western University
2. Paper
3. What is uncertain pricing? Traditional price discount promotion: the discount is fixed Uncertain price promotion: the final discount is ex ante uncertain to consumers, then draw a lottery, then revealed discount. example: discount ranging from 15% to 100%, need to go to the store, scratch ticket, see actual discount usually short period of times
4. Observations: Some optimistic consumers, they think they will get the best discount Some pessimist consumers, they think will get lowest discount

   After finding discount, customers can walk away without buying
5. Question: why would this scheme be a good idea for seller? 100% discount, 0 revenue low discount, customer walks away, 0 revenue
6. How do consumers beliefs affect firms' prices?
7. Literature review
   (a) Price obfuscation
   (b) Uncertain price promotion
   (c) all prior work empirical
   (d) Main contribution of this work: provides an analytical model where consumers are allowed to have subjective beliefs that may or may not be correct in equilibrium.
8. Model Simple Hotelling model consumer are uniformly distributed on a Hotelling line: two firms $u_A(x) = v - p_A - tx^2$ and $u_B(x) = v - p_B - t(1-x)^2$ firms As price p_A is ex ante uncertain but customer knows in bounded on an interval.

   two types of consumers optimistic, θ_O pessimistic, θ_p consumers perceived price of firm A is $\theta_O \bar{P}$ if optimistic, similarly for pessimistic.

   consumer decision-making process first decide which firm to visit, as a function of prices after visiting, customers can accept the firm's price of switch to other firm (assume no transportation cost to move from one firm to another)
9. Lemma (in the last period) in equilibrium, firm A sets price be equal to the upper bound upon consumers' arrival.

   Then solve for first period

10. Key takeaways firms can increase profit by switching to uncertain prices, subject to constraints on customers' beliefs

4.3.2 Feature Centralized Multiproduct Newsvendor with Substitution

1. Author: Alba V. Olivares Nadal, University of Zürich
2. Incorporating exogenous features to improve newsvendor model.

   In practice, replace empirical distribution as the demand distribution into the classical newsvendor model.

   Literature Ban and Rudin (2019)

   Goal Extend the previous setting into multiple products When one product is not available, we will have substitutes.

   even when demand is not deterministic, this problem is not convex.

   Conditional expected loss framework assume demand is dependent on some external feature

   $$D \mid (X = x) \sim P_x$$

   Empirical Risk Minimization framework minimize average on training set

   CMNS formulation with random demand D^* Netessine and Rudi (2003) centralized and competitive inventory models with demand substitution

   MINL reformulation Zhang, Xie, Sarin, (2021), multi-product newsvendor problem with customer-driven demand substitution: A stochastic integer program perspective.

4.4 Event Ticket Pricing with Capacity Constraints and Price Restrictions

1. Author: Yunke Li, University of Miami
2. Motivation A real ticket pricing problem box office manager for live concerts facing problem the perceived value of seats: different perceived quality of seat locations and idiosyncratic characteristics of the customers (in a concert hall, some people might prefer different places as a function of how music is perceived)
3. Practical restrictions: capacity limitation customer loyalty
4. Short selling period they prefer to make static pricing decisions dynamic pricing is expensive in this case
5. Objective multi-product static pricing problem to maximize revenue
6. Contribution find close form solution under different demand models
7. Demand Models

   two models, both special case of (McFadden and Train, 2000)

   $$u_j = \theta q_j - p_j + \mu \xi_j$$

   $q = (q_1, \ldots, q_N)$ quality levels of seat category $\xi = (\xi_1, \xi_2, \ldots \xi_N)$ consumers idiosyncratic utility, μ strength of idiosyncrasy θ sensitivity to product quality $p = (p_1, \ldots, p_n)$ price vector, $\alpha = (\alpha_1, \ldots, \alpha_n)$ demand vector
8. Problem formulation

   Maximize revenue s.t. capacity constraints average price constraint ceiling constraints (the price of the lowest quality product has a ceiling)
9. solve this problem is challenging.
10. Strategy solve unconstrained problem then add constraints, one at a time, and understand the impact of the constraints

4.5 Session: Data-driven Sequential Decision Making: Bandits and Reinforcement Learning

4.5.1 Learning Temporally-extended Actions with Risk-sensitive Q-learning

1. Author: Joongkyu Lee, Seoul National University
2. Temporal Abstraction Consider a robot in a hallway wasn't to go into a room Traditional RL would take action at every time step so robot might do go right, go left, go right, etc. But people don't make these decision, they make higher level decision walk to door walk forward ten times etc.
3. Why Action Repetition?
   (a) action repetition induces a deeper exploration
   (b) two main streams
      1) repeating an action only for exploration
      2) learning a repeating action
4. This is not enough, is agent picks a sub-optimal action, and repeats it many times, this is a problem.
5. Why does Action Repetition fail?

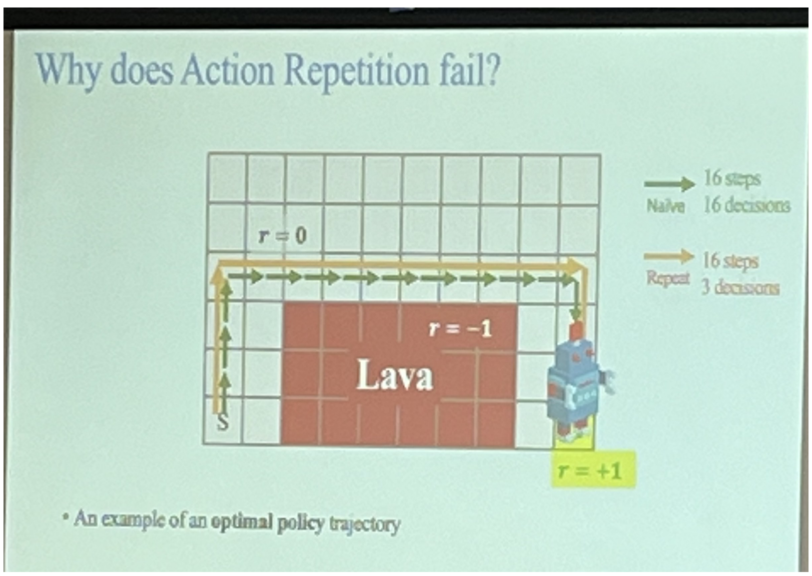

6. Action Repeating Option

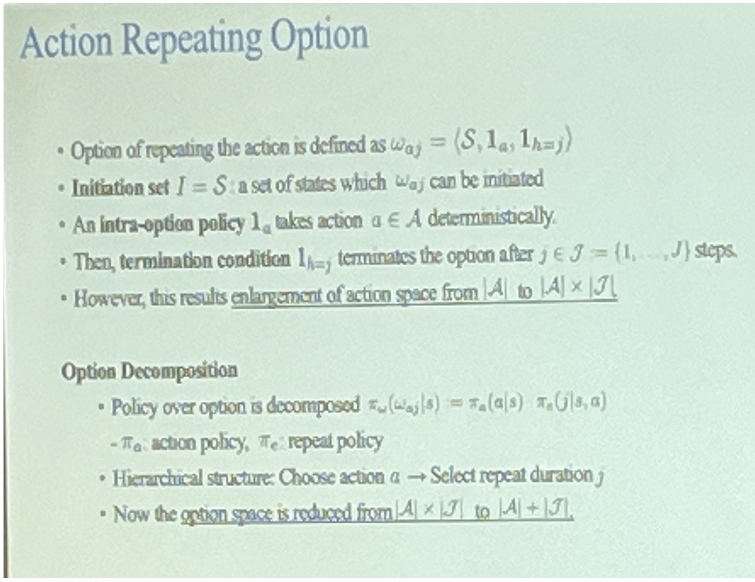

7. Algorithm RiSAR: Risk-Sensitive Action Repetition

   Usual algorithm but when choice repeat duration j,

   a risk parameter λ is used

   the parameter control the relative preference between explore and exploit

8. Experimental results on Chain MDP, GridWorlds, and Atari 2600

4.6 Model-based Reinforcement Learning with Multinomial Logistic Function Approximation

1. Author: Taehyun Hwang, Seoul National University
2. Paper: ?
3. RL with function approximation has made significant advances in empirical study (go, StarCraft)
4. However, theoretical understanding of this method is not well understood most work is still linear function approximation trying to gap bridge between practice and theory
5. Usual model, MDP, value functions.
6. Performance measure: $Regret_\pi(K) = (V_1^* - V_1^\pi)$
7. RL Landscape (Insert picture)
8. Challenges going beyond tabular methods (Insert picture)
9. Limitation of Linear Transition Model

   for an arbitrary set of features, a linear transition model cannot induce a proper probability distribution over next states

10. Multinomial Logistic Transition Model

    Motivation

    MNL model is a natural choice after linear model generalization of the linear model

    no prior work on RL with MNL model!

11. Upper Confidence Model-based RL for MNL

(a) this is the algorithm proposed by the authors
(b) model parameters updated using ridge penalization MLE.
(c) author shows a w.p. upper bound on regret.

4.7 Near-optimal Algorithm for Linear Contextual Bandits with Hybridization by Randomization

1. Author: Wonyoung Kim, Columbia University
2. Paper: ?
3. Contextual Bandit Problem
   (a) Used for sequential and adaptive decision making.
   (b) Examples:
      i. User using a laptop, the algorithm wants to show best ad for the user
      ii. Algorithm picks an Ad to show, get a reward
      iii. learn users preferences.
   (c) The algorithm observed the context which is related in an unknown way with rewards
   (d) Need to balance between exploration vs exploitation
4. Linear Contextual Bandit problem $X_{i,t} \in \mathbb{R}^d$, where d-dimensional context
   reward $Y = X\beta + \eta$
5. Goal Minimize cumulative regret difference between optimal arm and chosen arm
6. there are many regret upper bounds from many algorithms.
7. A long-standing question
   (a) is there an algorithm that achieves $\sqrt{dT}$ regret bound that does not discard the data?
   (b) this paper gives answer to this question
8. Data Structure
   (a) observe all context, but rewards only for selected arms
   (b) how about we view this problem from the missing data point of view?
   impute missing rewards
   (c) there are methods that do this an improve regret bound
9. Hybridization by Randomization (HyRan) Algorithm uses estimated rewards to update its model
10. Novel regret decomposition
    improves state of the art without discarding data and achieves $O(\sqrt{dT \log T}$ regret bound

4.8 Bayesian Design Principles for Frequentist Bandit and Reinforcement Learning

1. Author: Yunbei Xu
2. Part 1: Motivation and Contribution
3. Online learning with partial observation sequential decision to accumulate reward usual problem: only observe reward of chosen action and not the rewards of other actions

   MAB RL dynamic pricing, only observe single point price instead of whole demand curve
4. Frequentist vs. Bayesian
   (a) Frequentist: UCB bound, sample average or linear regression to estimate mean of reward
   (b) Bayesian: Thompson Sampling with fixed, pre-specific prior, updates posterior at each round
5. Neither are applicable to non-stationary environments.
6. Frequentist vs. Bayesian: Characteristics
   (a) Frequentist

       pros: does not require knowledge of environment

       cons: heavily depends on case-by-case design and special structure existing estimating methods only applicable to simple model classes
   (b) Bayesian

       pros: updating posteriors is general, intuitive and often optimal

       cons: knowledge of the environment (prior) not accessible in

       complex settings maintain posteriors is computationally expensive
7. Main Research Question can we design general-purpose structured algorithms, which know of the environment that are computationally efficient?
8. Contributions Algorithmic beliefs $->$ synergizing Frequentist and Bayesian approaches
9. Part 2: Framework and Algorithms

   Problem instance (Π, M), where Π decisions space, M model space

   Model M maps each π into a distribution over the observation space. Non-stochastic environment

   Regret, as usual, but may change at each round in the stochastic environment.

4.9 Key Note: Brian Macdonald. Sports Analytics

Abstract

Data has been meticulously collected in sports for decades. As technology has improved, so has the data, allowing for more appropriate statistical modeling and machine learning techniques to be used to answer interesting questions. Today, many leagues use computer vision, remote sensing, and machine learning techniques to collect player and ball location data multiple times per second throughout the duration of every game, and to automatically detect events of interest during a game, both of which allow for numerous new types of analyses. We will give a brief history of analytics in sports, the questions that teams, leagues, media organizations and fans try to answer using data, the current state of analytics in the sports industry, and how sports analytics is being used in education to inspire student interest in data science.

1. What is sports analytics?

2. what is analytics?

   INFORMS definition: the application of scientific & mathematical methods to the study and analysis of problems involving complex systems

   That, applied to sports

3. Industry: multi-billion dollar business. Data helps inform decisions.

4. Academia: great data and problems for research and teaching.

5. Lots of public data, freely available to anyone. wide variety of data, problems, and methods

6. Why sports analytics?

   Problems solve here transfer to other settings Sports are a controlled environment Sports are popular Real-life validation

7. What sports organizations use data? Teams Team operations (focus of the key note) Business operations Leagues Sports and business analytics Media ESPN FiveThirtyEight Betting Companies

8. Types of data Game summary data basic information about a game what teams? location of game. scores. box score data play-by-play more detailed data, summary of events that hapend during the game, with timestamp and players involved in the play also substitution information, can infer what players where on the field at any moment of the game player tracking data location of each player at any moment in time pose data

   Data acquisition sources of data leagues internal data 3rd party companies internet, e.g., crowdsource data collection projects

9. Team operations Front office pro player evaluation, projections, trade analysis scouting amateur player projections, analysis of scouting reports coaching analysis of in-game decisions, strategy, pre- and post-game reports conditioning/medical nutrition, workload, sleep

10. Team Ratings/Projections How good is Team A? Will Team A make the playoffs? Win division? How can we estimate how good team A is? consider results of games, opponents, travel, rest, etc.

11. Regression-based team rating Outcome: score differential in a game Predictors: team, opponent, home advantage, travel diff, rest diff, altitude diff Interpret regression coefficient

12. Team ratings used for game predictions season simulations expected win totals, Prob(Make the playoffs)

13. Game Predictions W/L, Prob(Win) Score differential Predicted score distributions of possible outcomes

    what influences the game outcomes? teams, venue, etc

4.10 Session: Large Markets and Mechanism Design

4.10.1 Equilibrium Learning and Bilateral Bargaining

1. Author: Martin Bichler et al. Technical University of Munich

2. Paper: ?

Abstract

Bilateral bargaining of a single good among one buyer and one seller describes the simplest form of trade, yet Bayes-Nash equilibrium strategies are largely unknown. Recent advances in equilibrium learning provide a numerical approach to auction games, which can push the boundaries of existing results. We analyze Neural Pseudogradient Ascent (NPGA) and Simultaneous Online Dual Averaging (SODA), two new equilibrium learning algorithms for Bayesian games with continuous type and action spaces. We show that both algorithms consistently learn equilibrium even in this challenging environment. In our analysis, we derive equilibrium bid functions for non-uniform priors, risk-averse traders, and markets with multiple traders on each side, which has been impossible so far. Besides, we provide a new convergence result for NPGA.

Monotonic Algorithms for PDEs

Motivation is to understand out come of interactive decision making for Bayesian auction games that can be modeled as a system of non-linear PDEs. Computational complexity is high.

New equilibrium learning algorithms for auction games, particularly bilateral trades

Convergence for a wide variety of auctioning models

Convergence in bilateral trade odels wintou analytical solution Robust and pur equilibria finding learning algorithms

NPGA conbverges locally for uniformly distributed priors.

Key question: Maybe Equilibrium computation is tractable.

PDE solvers are instable and analytical solutions only exist for specific rules

Many step function equilibria exist.

Monotonicity can be achieved for bilateral trade games with linear strategies

Interestingly equilibria are formed as step function if the priors are not known

Conclusions: If a game is monotone algorithms converge Infinte equilibria for bilateral bargaining

NPGA and SODA find equilibrium in bilateral trade with gaussian priors

4.11 Ascending-price Mechanism for General Multi-sided Markets

1. Author: Rica Gonen et. al. Ariel University, Ariel, Israel
2. Paper:

Abstract

We present an ascending-price mechanism for a multi-sided market with a variety of participants, such as manufacturers, logistics agents, insurance providers, and assemblers. Each deal in the market may consist of a combination of agents from separate categories, and different such combinations are simultaneously allowed. This flexibility lets multiple intersecting markets be resolved as a single global market. Our mechanism is obviously- truthful, strongly budget-balanced, individually rational, and attains almost the optimal gain-from- trade when the market is sufficiently large. We evaluate the performance of the suggested mechanism with experiments on real stock market data and synthetically produced data.

Ascending - price mechanisms for general multi-sided markets

Motivation: Many buyers and sellers and match up lead to pretty complex markets so many assumptions in games that are well known do not hold Value in efficient markets is generated by the perfect matches created between buyers and sellers Ideal propertioes: Efficient Rational actors Truthfulness about prices These assumptions are not particularly realistic Another property to consider: Weak budget balance: auctioner does not add money to market Strong budget balance: auctioneer does not receive money What do we do if we have N sided markets? New concept: external competition The best about this is that you can calculate the market competition in one go for an n sided market instead of as a chain of two sided markets. General multi sided markets: Where customers can purchase from multiple sources Ascending-price mechanism for general muti-sided markets algorithm to compute the problem and achieve the equilibrium Scheme runs until market is in balance (all prices summed together are zero) Price initialization is created dependent on the depth of the market

4.12 Minimum Price Equilibrium in the Assignment Market: The Serial Vickrey Mechanism

1. Author: Shigehiro Serizawa.
2. Paper: ?

Abstract

We study an assignment market where multiple heterogeneous objects are sold to unit-demand agents who have general preferences that accommodate income effects and market frictions. The minimum price equilibrium (MPE) is one of the most important equilibrium notions in such settings. Nevertheless, none of the well-known mechanisms that find the MPEs in quasi-linear environment can identify or even approximate the MPEs for general preferences. We establish novel structural characterizations of MPEs and design the "Serial Vickrey (SV) mechanism" based on the characterizations. The SV mechanism finds an MPE for general preferences in a finite number of steps. Moreover, the SV mechanism only requires agents to report finite-dimensional prices finitely many times, and also has nice dynamic incentive properties.

3. Minimum price equilibrium in the assignment market: serial vickrey mechanisms
4. Objective: Design of an adjustment process for non-quasi-linear preferences

4.13 Deeds for Speed: Rewarding Innovation with Transferable Regulatory Speed

Last Talk (neglected disease incentive development policies)

1. Author: David Ridley, Duke University.
2. Paper: ?

Abstract

Title: Deeds for Speed: A Mechanism for Rewarding Innovation with Transferable Regulatory SpeedAuthors: Sandro Brusco, Giuseppe Lopomo, David Ridley Abstract: Under the priority review voucher program, the developer of a drug for a neglected disease receives a transferable voucher for faster regulatory review for a different drug. Vouchers have sold for more than $300 million, but as voucher supply increased, voucher prices fell to $100 million. We show how mechanisms like priority review vouchers can replace monetary rewards. We also show how to refine the program. Rather than awarding a voucher conditional on regulatory approval, we show that the regulator should award a voucher conditional on a high drug quantity sold. This mechanism would drive down drug prices, drive up voucher prices, encourage innovation, and increase consumer surplus.

Basic idea comes from a previous paper that became the law:

Neglected disease drugs allow pharmaceutical companies to receive a voucher to shorten review times for another drug The idea is to reward good deed with speed

The problems with this law are: Some of the neglected disease problems are expensive Second problem is the low price for the voucher, which weakens the incentive

If the price potential is too high, allowing a voucher is not worth it, it also creates opportunity cost since other drugs will now be neglected

Idea: use a price cap

Downside: this may hurt people as there are some drugs that otherwise would never be developed

5 Wednesday, October 19, 2022

5.1 Plenary: Morse Lectureship - Alvin E. Roth. Market Design: The Dialog Between Simple Abstract Models and Practical Implementation

Abstract

I'll review some of the elegantly simple models that underlie the initial designs for matching processes like the medical residency Match, school choice and kidney exchange, and the modifications, complications and computations that were needed to get new designs adopted, implemented and maintained over the years.

1. Chat before talk

   (a) Best Advise:

      i. when things don't go well, keep pushing
      ii. find problems worth studying
      iii. using tools that you enjoy using
      most times there is no progress, so if you don't enjoy the tools, you won't have enough motivation to keep going

2. Connection between simple analytic models and practical implementations

3. Talk about One hand, Models and Algorithms On the Other Applications and Complications

4. Gale and Shapley (1962), college stable matching, two-sided matching, deference acceptance algortihm

5. applied to the medical market

   practical complication: not only individuals looking for job, but couples looking for jobs

6. another application is schools admissions

   practical complication:explaining and persuading

7. Gale and Shapley (1962)

   Firms multiple positions preferences over workers Workers preferences over firms

   Outcome: matching firms and workings that respect capacities of firms matching is stable if there isn't any blocking pair

   We see in practice that theoretical stability actually matters for the life of the match

8. The Algorithm

   Input: preferences in rank order (both firms and workers)

   Step 1: Each worker "applies" to its first choice. Each firm assigns its positions to its applicants one at a time in their priority order until all positions are tentatively filled. Any remaining applicants are rejected.

   ...

   Step k: each worker who was rejected in the previous step applies to her next choice if one remains. Each firm considers the workers it has been holding together with its new applicants and tentatively assigns its seats to these workers on a time in priority order. Any remaining applicants are rejected

   The algorithm terminates when no worker is rejects, and each worker is assigned its final tentative assignment.

   Properties: Dominant strategy to reveal your true preferences

9. Applying algorithm to doctors A complicated market. For example, applying as couple: needs to position not just one.

   Work on couple matching started in the 1970s. Couples would report their preferences as individual, but then defect.

   Details of the algorithm Get certified as a couple Select a "leading member" Submit a separate rank order list for each member The leading member went through the match as if single the other member then had his/her rank order list edited to remove positions not in the "same community" as the one the leading member This did not work well. Why?

   The iron law of marriage: You can't be happier than your spouse.

   Couple consumes pairs of jobs! The algorithm didn't have the right language to deal with these preferences.

   Why is the couples a hard problem?

   A modification of DA that accounts for couples was successful (1990)

10. School choice presented a complication Schools don't have fine grained preferences over students, in fact, lots of indifferences. Tie breaking is important in this application and turns out it matters how you break ties. Tie introduce a "fictitious" preference.

    DA with tie breaking (for School Assignment)

    Step 0 Priorities of students at schools are augmented with random lottery numbers, for tie-breaking when needed. How lotteries are implemented matters! Apply DA

11. Review of top trading cycles

12. Fairness rationale for strategy-proof mechanisms The system will be more fain since those who cannot strategize will not be penalized That was a rationale for strategy-proof from the school point of view

13. new work on better ways to explain to family and students why is good to reveal your true preferences when running top trading cycles (Strategy-proofness Exposing Mechanism Descriptions Yannai et al, 2022)

14. Kidney Exchange

    Kidney disease is a fatal disease but you can get a transplant Someone that loves you can give you a kidney, but only if they are a match

    One can think of using top trading cycle, but there are practical complications physical limitation on how exchanges can be implemented in practice, if donors are not in the same place. exchanges need to happen at the same time, but a simple two-way cycle needs four operation rooms and four teams! number of resources grows exponentially

    if you can only do pairs, you will lose on a lot of matches

15. Kidney Exchanges is powered via Mixed-integer programming, even with hard instances. This is a good example of dialogue between simple models and applications.

    When you design a mechanism, you play the role of creator, you design the game and assume participants will play. In reality, when you deploy your mechanism, it is part of a larger market environment that introduces complications and constraints that then feeds into new theoretical problems. Rinse and Repeat.