FAW 2019

Stochastic Combinatorial Optimization Problems

Jian Li

Institute of Interdisciplinary Information Sciences

Tsinghua University

Combinatorial and Geometric Optimization problems

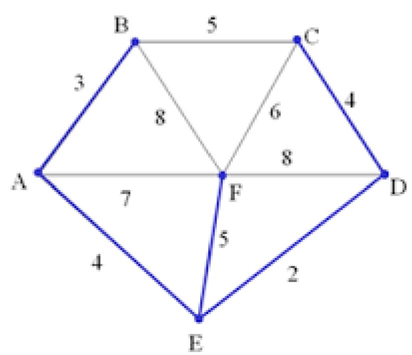

Minimum Spanning Tree

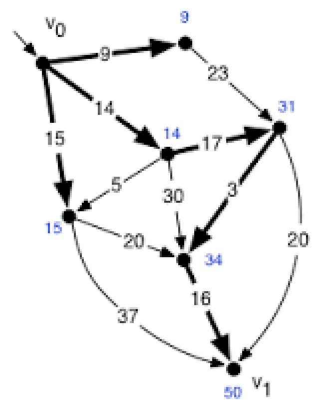

Shortest Path

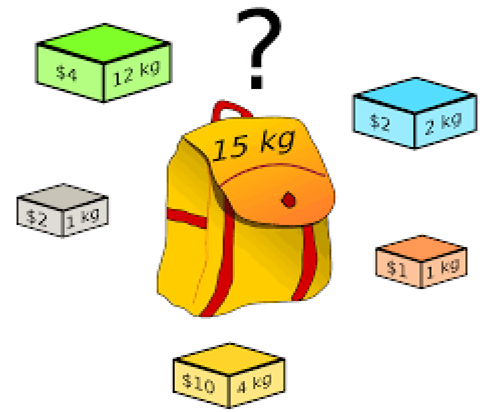

Knapsack

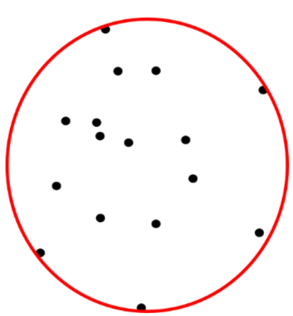

Minimum Enclosing Ball

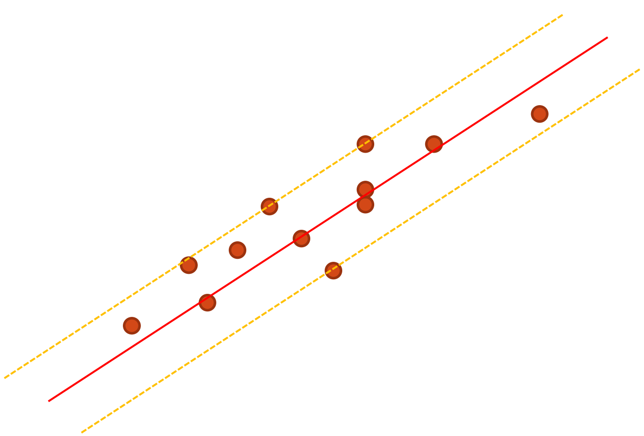

Minimum j-fat center

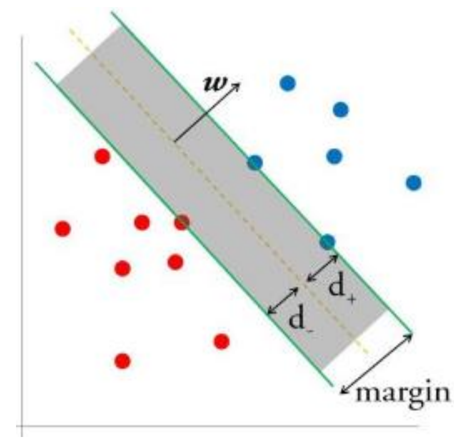

SVM

Stochastic Model Everywhere

- Data Integration and Information Extraction
- Sensor Networks; Information Networks
- Probabilistic models in machine learning
- Markov Decision Process - Reinforcement Learning

Sensor ID	Temp.
1	Gauss(40,4)
2	Gauss(50,2)
3	Gauss(20,9)
…	…

Probabilistic databases

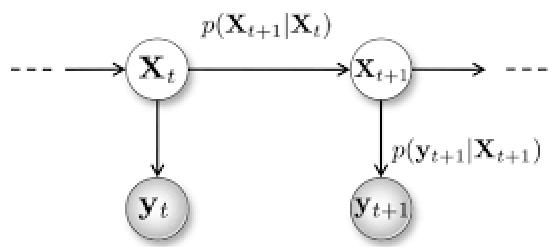

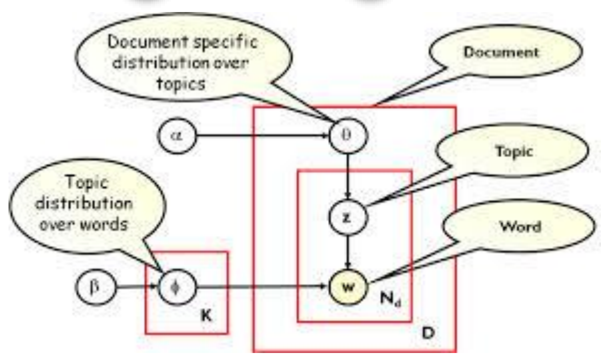

Probabilistic Models in machine learning

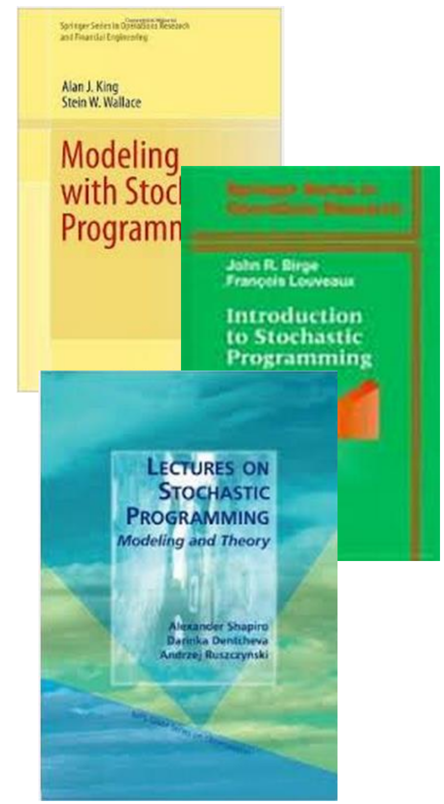

Stochastic models in operation research

Stochastic Optimization

- Danzig in 1950s (linear programming with stochastic coefficients – stochastic programming)
- Depending on how the decision process interacts with the uncertainty, we may be able to formulate different versions of stochastic optimization problems
  - Estimation (no decision)
  - Single-stage
  - 2-stage
  - Multi-stage
  - Online (adaptive/non-adaptive))

Simons Institute

- https://simons.berkeley.edu/

Algorithms and Uncertainty

Aug. 17 – Dec. 16, 2016

Workshops

Algorithms and Uncertainty Boot Camp
Aug. 22 – Aug. 26, 2016
Organizers: Avrim Blum (Carnegie Mellon University), Anupam Gupta (Carnegie Mellon University), Robert Kleinberg (Cornell University), Stefano Leonardi (Sapienza University of Rome), Eli Upfal (Brown University), Adam Wierman (California Institute of Technology)

Optimization and Decision-Making Under Uncertainty
Sep. 19 – Sep. 23, 2016
Organizers: Nikhil Bansal (Technische Universiteit Eindhoven; chair), Shipra Agrawal (Columbia University), Robert Kleinberg (Cornell University), Kamesh Munagala (Duke University), Jay Sethuraman (Columbia University), Adam Wierman (California Institute of Technology)

Learning, Algorithm Design and Beyond Worst-Case Analysis
Nov. 14 – Nov. 18, 2016
Organizers: Avrim Blum (Carnegie Mellon University; chair), Nir Ailon (Technion Israel Institute of Technology), Nina Balcan (Carnegie Mellon University), Ravi Kumar (Google), Kevin Leyton-Brown (University of British Columbia), Tim Roughgarden (Stanford University)

Organizers:
Anupam Gupta (Carnegie Mellon University; chair; co-chair), Stefano Leonardi (Sapienza University of Rome; co-chair), Avrim Blum (Carnegie Mellon University), Robert Kleinberg (Cornell University), Eli Upfal (Brown University), Adam Wierman (California Institute of Technology).

Long-Term Participants (including Organizers):
Nir Ailon (Technion Israel Institute of Technology), Susanne Albers (Technische Universität München), Aris Anagnostopoulos (Sapienza University of Rome), Peter Auer (University of Leoben), Yossi Azar (Tel Aviv University), Nikhil Bansal (Technische Universiteit Eindhoven), Peter Bartlett (UC Berkeley), Eilyan Bitar (Cornell University), Avrim Blum (Carnegie Mellon University), Nicolò Cesa-Bianchi (University of Milan), Shiri Chechik (Tel Aviv University), Edith Cohen (Google Research), Artur Czumaj (University of Warwick), Amit Daniely (Google Research), Amos Fiat (Tel Aviv University), Fabrizio Grandoni (IDSIA), Anupam Gupta (Carnegie Mellon University; chair; co-chair), MohammadTaghi Hajiaghayi (University of Maryland), Longbo Huang (Tsinghua University), Sungjin Im (UC Merced), Ravi Kannan (Microsoft Research India), Sampath Kannan (University of Pennsylvania), Anna Karlin (University of Washington), Robert Kleinberg (Cornell University), Elias Koutsoupias (University of Oxford), Ravi Kumar (Google), Stefano Leonardi (Sapienza University of Rome; co-chair), Kevin Leyton-Brown (University of British Columbia), Jian Li (Tsinghua University), Na Li (Harvard University), Katrina Ligett (Hebrew University and Caltech), Aleksander Mądry (Massachusetts Institute of Technology), Yishay Mansour (Tel Aviv University), Ruta Mehta (University of Illinois, Urbana-Champaign), Jamie Morgenstern (University of Pennsylvania), Kamesh Munagala (Duke University), Viswanath Nagarajan (University of Michigan), Seffi Naor (Technion Israel Institute of Technology), Kameshwar Poolla (UC Berkeley), Kirk Pruhs (University of Pittsburgh), Ram Rajagopal (Stanford University), Satish Rao (UC Berkeley), Benjamin Recht (UC Berkeley), Rhonda Righter (UC Berkeley), Tim Roughgarden (Stanford University), Piotr Sankowski (University of Warsaw), C. Seshadhri (UC Santa Cruz), Jay Sethuraman (Columbia University), Cliff Stein (Columbia University), Chaitanya Swamy (University of Waterloo), Marc Uetz (University of Twente), Eli Upfal (Brown University), Marilena Vendittelli (Sapienza University of Rome), Maria Vlasiou (Eindhoven University of Technology), Jan Vondrák (Stanford University), Jean Walrand (UC Berkeley), Gideon Weiss (University of Haifa), Adam Wierman (California Institute of Technology), Bert Zwart (CWI Amsterdam).

Research Fellows:
Ilan Cohen (Tel Aviv University), Varun Gupta (University of Chicago), Thomas Kesselheim (Max-Planck-Institute for Informatics and Saarland University), Marco Molinaro (PUC-Rio de Janeiro; Microsoft Research Fellow), Benjamin Moseley (Washington University in St. Louis), Debmalya Panigrahi (Duke University), Xiaorui Sun (Columbia University; Google Research Fellow), Matt Weinberg (Princeton University), Qiaomin Xie (University of Illinois at Urbana-

Outline

- Estimation (no decision)
- Single-stage
- 2-stage
- Online (adaptive/non-adaptive))
- Sample Complexity

A classic problem in the stochastic graph model

- A undirected graph with n nodes
- The length of each edge: i.i.d. Uniform[0,1]

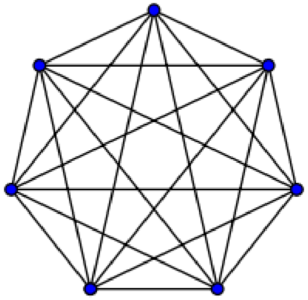

- Question: What is E[MST]? [McDiarmid, Dyer, Frieze, Karp, Steele, Bertsekas, Goemans]

A classic problem in the stochastic graph model

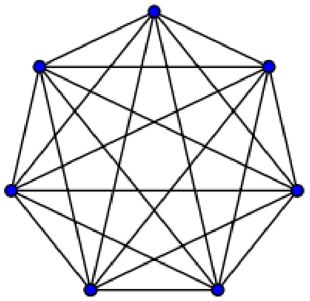

- A undirected graph with n nodes
- The length of each edge: i.i.d. Uniform[0,1]
- Question: What is E[MST]? [McDiarmid, Dyer, Frieze, Karp, Steele, Bertsekas, Goemans]
- Ignoring uncertainty ("replace by the expected values" heuristic)
  - each edge has a fixed length 0.5
  - This gives a WRONG answer 0.5(n-1)
- But the true answer is (as n goes to inf)

$$\zeta(3) = \sum_{i=1}^{\infty} 1/i^3 < 2$$

A classic problem in the stochastic geometry model

- N points: i.i.d. uniform$[0,1] \times [0,1]$

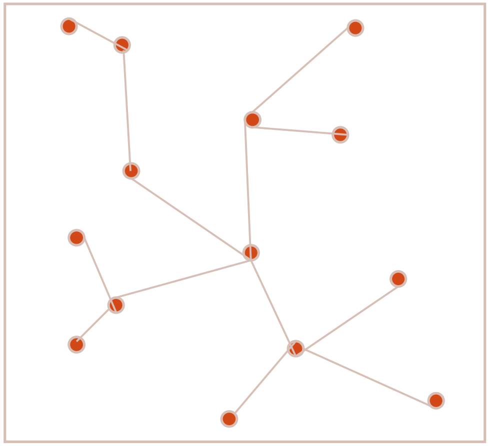

- Question: What is E[MST] ? [Frieze, Karp, Steele, …]

A classic problem in the stochastic geometry model

- N points: i.i.d. uniform$[0,1] \times [0,1]$

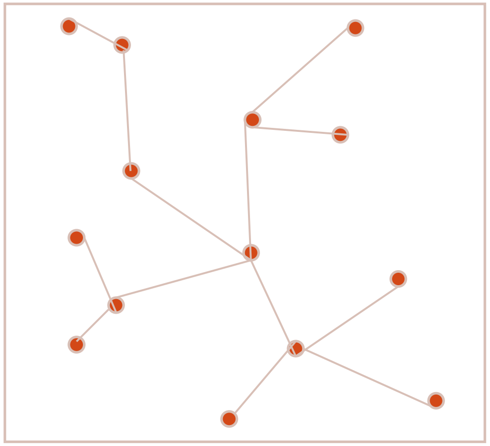

- Question: What is E[MST] ? [Frieze, Karp, Steele, …]
- Answer: $\theta(\sqrt{n})$ [Frieze, Karp, Steele, …]

The problems are similar, but the answers are not similar…………

Stochastic Graph Model

- The weight of each edge is a (discrete) random variable

  FPRAS for computing the expectation (including higher moments) for a family of problems including the diameter of G, minimum spanning tree [Emek et al. SODA'11]
- Open: shortest path, matching

- All terminal reliability problem [Moore and Shannon 56] [Valiant 79]

  Estimate Pr[the graph is (not) connected]

  FPRAS [Karger, SICOMP99]
- s-t reliability problem

  Estimate Pr[s and t are (not) connected]

  A long standing open problem

Stochastic Geometry Models

- The position of each point is random (non-i.i.d)
- All pts are independent from each other
- A popular model in wireless networks/spatial prob databases

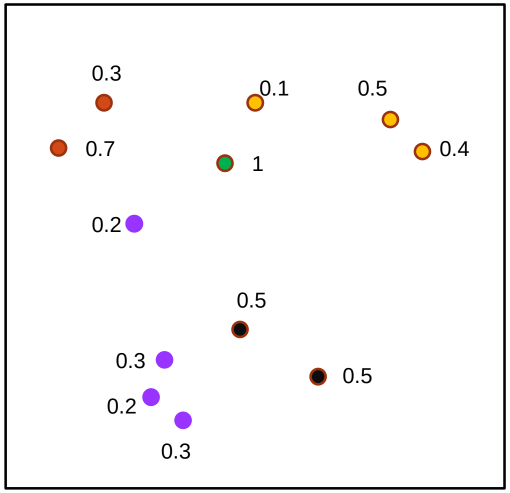

Locational uncertainty model

Stochastic Geometry Models

- A computational problem: Computing E[MST]

$$E[MST] = \sum_{r:realizations} MST(r) \times \Pr[r]$$

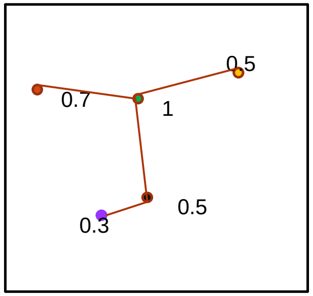

*(0.7*1*0.5*0.5*0.3) +

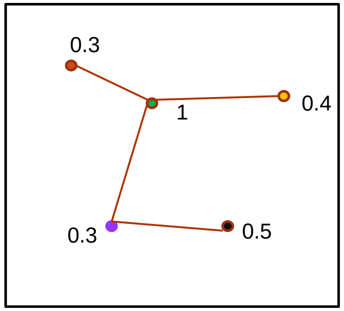

*(0.3*1*0.4*0.5*0.3) + ……

However, this does not give a polynomial time algorithm

The stochastic geometry model has been studied in several recent papers for many different problems.
[Kamousi, Chan, Suri ‘11] [Afshani, Agarwal, Arge, Larsen, Phillips. ‘11][Agarwal, Cheng, Yi. ‘12] [Abdullah, Daruki, Phillips ‘13] [Suri, Verbeek, Yıldız ‘13] [Li, Wang ‘14] [Agarwal, Har-Peled, Suri, Yıldız, Zhang 14] [Huang, Li ‘15] [Huang, Li, Phillips, Wang ‘15]

A Computational Problem

- The position of each point is random (non-i.i.d)

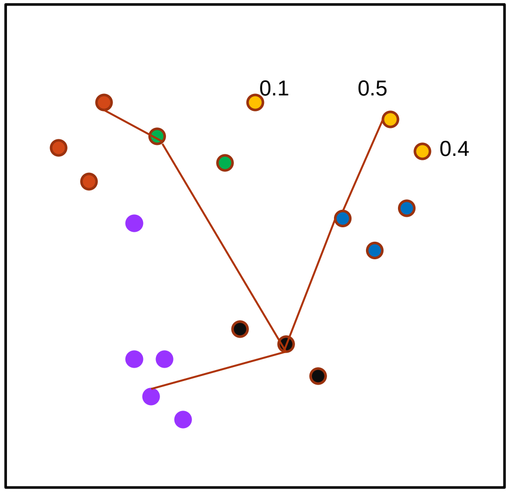

- Question: What is E[MST] ?
- Of Course, there is no close-form formula
- We need efficient algorithms to compute E[MST]

MST over Stochastic Points

- The problem is #P-hard [Kamousi, Chan, Suri. SoCG'11]
- So, let us focus on approximating the value
- FPTAS for existential model, constant approx for locational model [Kamousi, Chan, Suri. SoCG'11]
  - FPRAS: Solution $\in (1 \pm \epsilon) \times$ true value in time Poly$(n, \frac{1}{\epsilon})$
- Other problems:

[Huang, L. ICALP'15]

Problems		Existential	Locational
Closest Pair (§2)	$\mathbb{E}[\mathrm{C}]$	FPRAS	FPRAS
	$\Pr[\mathrm{C} \leq 1]$	FPRAS	FPRAS
	$\Pr[\mathrm{C} \geq 1]$	Inapprox	Inapprox
Diameter (§2)	$\mathbb{E}[\mathrm{D}]$	FPRAS	FPRAS
	$\Pr[\mathrm{D} \leq 1]$	Inapprox	Inapprox
	$\Pr[\mathrm{D} \geq 1]$	FPRAS	FPRAS
Minimum Spanning Tree (§4)	$\mathbb{E}[\mathrm{MST}]$	FPRAS[25]	FPRAS
k-Clustering (§3)	$\mathbb{E}[\mathrm{kCL}]$	FPRAS	Open
Perfect Matching (§5)	$\mathbb{E}[\mathrm{PM}]$	N.A.	FPRAS
kth Closest Pair (§B.1)	$\mathbb{E}[\mathrm{kC}]$	FPRAS	Open
Cycle Cover (§6)	$\mathbb{E}[\mathrm{CC}]$	FPRAS	FPRAS
kth Longest m-Nearest Neighbor (§7)	$\mathbb{E}[\mathrm{kmNN}]$	FPRAS	Open

Outline

- Estimation (no decision)
- Single-stage
- 2-stage/Multi-stage
- Online (adaptive/non-adaptive))
- Sample Complexity

(Single stage) Stochastic shortest path

- Find an s-t path P such that $Pr[w(P)<1]$ is maximized
- Route planning: maximize the prob that one can reach the destination in 1 hour

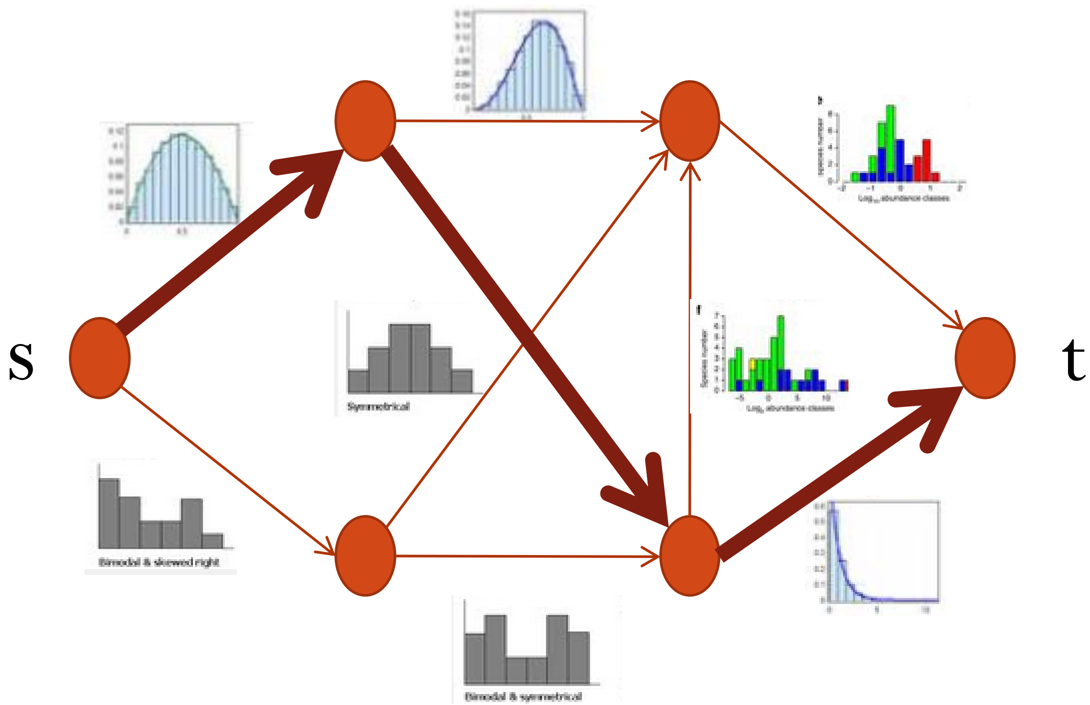

Threshold Probability Maximization

- Deterministic version:
  - A set of element $\{e_i\}$, each associated with a weight w_i
  - A solution S is a subset of elements (that satisfies some property)
  - **Goal:** Find a solution S such that the total weight of the solution $w(S)=\Sigma_{i\epsilon S}w_i$ is minimized
  - E.g. shortest path, minimal spanning tree, top-k query, matroid base
- Stochastic version:
  - w_is are independent positive random variables
  - **Goal:** Find a solution S such that the *threshold probability*

$$\Pr[w(S) \leq 1] \quad \text{is maximized.}$$

Even computing the threshold prob is #P-hard in general! (generalizes #knapsack)
FPTAS exists [**L**, Shi, ORL'14]

Our Result

If the deterministic problem is “easy”, then for any $\epsilon > 0$, we can find a solution S such that

$$\Pr[w(S) \leq 1 + \epsilon] > OPT - \epsilon$$

“Easy”: there is a PTAS for the corresponding O(1)-dim packing problem:

- Shortest path, MST, matroid base, matroid intersection, min-cut

- Related work: Special distributions [Nikolova, Kelner, Brand, Mitzenmacher. ESA’06] [Nikolova. APPROX’10] [Goel, Indyk. FOCS’99] [Goyal, Ravi. ORL’09] [Bhalgat, Goel, Khanna. SODA’11] [**L**, Deshpande. MOR’18]

Algorithm

- Step 1: Discretizing the prob distr

(Similar to [Bhalgat, Goel, Khanna. SODA'11], but simpler)

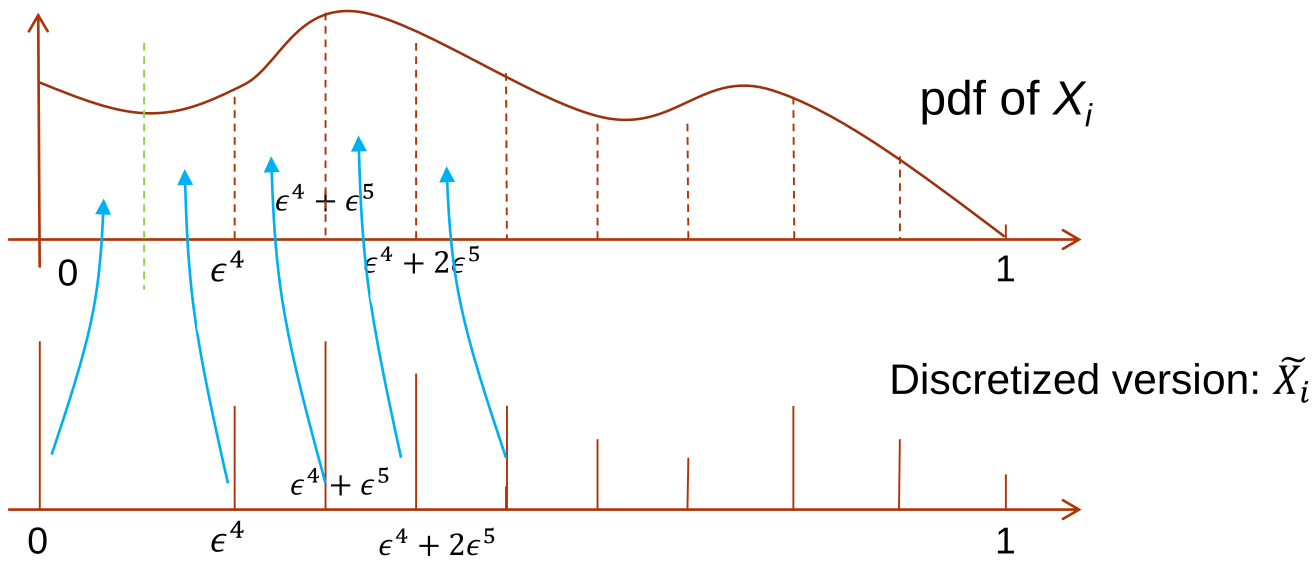

The behaviors of $\widetilde{X}_i$ and X_i are close:

1. $\Pr[X(S) \le \beta] \le \Pr[\widetilde{X}(S) \le \beta + \epsilon] + O(\epsilon)$;
2. $\Pr[\widetilde{X}(S) \le \beta] \le \Pr[X(S) \le \beta + \epsilon] + O(\epsilon)$.

Algorithm

- Step 2: Reducing the problem to the multi-dim problem
  - Heavy items: $E[X_i]>poly(\epsilon)$
    - At most $O(1/poly(\epsilon))$ many heavy items, so we can afford enumerating them

Algorithm

- Step 2: Reducing the problem to the multi-dim problem
  - Heavy items: $E[X_i] > \mathrm{poly}(\epsilon)$
    - At most $O(1/\mathrm{poly}(\epsilon))$ heavy items, so we can afford enumerating them
  - Light items:
    - Fix the set H of heavy items
    - Each X_i can be represented as a O(1)-dim vector **Sg(i)** (signature)

$$\mathbf{Sg}(i) = (\Pr[\tilde{X}_i = \epsilon^4], \Pr[\tilde{X}_i = \epsilon^4 + \epsilon^5], \ldots \ldots)$$

    - Enumerating all O(1)-dim (budget) vectors B
      - Find a set S such that $S \cup H$ is feasible and

$$\mathbf{Sg}(S) = \sum_{i \in S} \mathbf{Sg}(i) \leq (1+\epsilon)B \quad \text{(using the multi-dim PTAS)}$$

      (or declare there is none S s.t. $\mathbf{Sg}(S) \leq B$)
    - Return $S \cup H$ for which $\Pr[w(S \cup H) \leq 1 + \epsilon]$ is largest

Poisson Approximation

Well known: Law of small numbers

n Bernoulli r.v. X_i (1-p, p)

$np = const$

As $n \to \infty, \sum X_i \sim Poisson(np)$

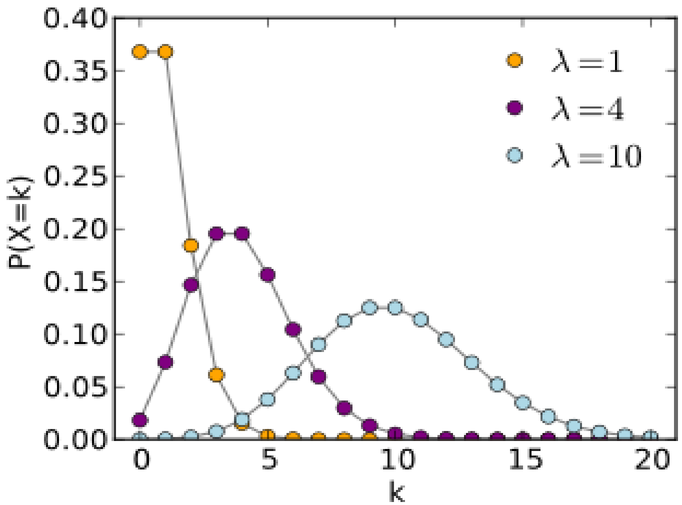

Poisson Approximation

Le Cam's theorem (rephrased):

n r.v. X_i (with common support (0,1,2,3,4,…)) with signature

$$\mathbf{sg}_i = (\Pr[X_i = 1], \Pr[X_i = 2], \dots)$$

Let $\mathbf{sg} = \sum_i \mathbf{sg}$

Y_i are i.i.d. r.v. with distr $\mathbf{sg}/|\mathbf{sg}|_1$

Y *follows the* ***compound Poisson distr (CPD)*** *corresponding to* $\mathbf{sg}$

$$Y = \sum_{i=1}^{N} Y_i \text{ where } N \sim \text{Poisson}(|\mathbf{sg}|_1)$$

Then, $\Delta(\sum X_i, Y) \le \sum p_i^2$ where $p_i = \Pr[X_i \neq 0]$

Variational distance:
$\Delta(X, Y) = \sum_i |\Pr[X = i] - \Pr[Y = i]|$

Poisson Approximation

- **Le Cam's theorem:** $\Delta(\sum X_i, Y) \leq \sum p_i^2$
- Ob: If S_1 and S_2 have the same signature, then they correspond to the same CPD
- So if $\sum_{i \in S_1} p_i^2$ and $\sum_{i \in S_2} p_i^2$ are sufficiently small, the distributions of $X(S_1)$ and $X(S_2)$ are close
- Therefore, enumerating the signature of light items suffices (instead of enumerating subsets)

Summary

- The #dimension needs to be $L = poly(1/\epsilon)$
- We solve an $poly\left(\frac{1}{\epsilon}\right)$-dim optimization problem
- The overall running time is polynomial (PTAS) [L, Yuan STOC'13]
- Can be easily extended to the multi-dimensional case, other combinatorial constraints etc.

- Open problem:

Remove the first ϵ in $\Pr[w(S) \leq 1 \boxed{+\epsilon}] > OPT - \epsilon$

Remove this?

Related results:

Bernoulli random variables, FPTAS [De SODA '18] (Boolean function analysis)

fault tolerant storage problem [Daskalakis et al. SODA'14] (using results from linear threshold function)

Outline

- Estimation (no decision)
- Single-stage
- 2-stage/Multi-stage
- Online (adaptive/non-adaptive))
- Sample Complexity

2-stage facility location

- First stage:
  - We know a distribution of demand
  - We can build a set of facilities
- Second stage
  - The set of demand realizes
  - We can build some extra facilities (but more expansive, inflation factor $\gamma > 1$)
- GOAL: minimize the expected total cost

A large body of literature. Constant approximation for many problems. Extensive studied.

Some general technique: boosted sampling [Gupta et al. STOC'04]

Outline

- Estimation (no decision)
- Single-stage
- 2-stage/Multi-stage
- Online (adaptive/non-adaptive))
- Sample Complexity

Weitzman's Pandora problem

- Proposed by Weitzman in Econometrica 79
- A typical stochastic optimization problem

Econometrica, Vol. 47, No. 3 (May, 1979)

OPTIMAL SEARCH FOR THE BEST ALTERNATIVE

BY MARTIN L. WEITZMAN[1]

This paper completely characterizes the solution to the problem of searching for the best outcome from alternative sources with different properties. The optimal strategy is an elementary reservation price rule, where the reservation prices are easy to calculate and have an intuitive economic interpretation.

Pandora's Boxes

- Pandora has n boxes.
- Box *i* contains an unknown value x_i, distributed with known c.d.f. F_i.
- At known fixed cost c_i, she can open box i and discover x_i.
- Pandora may choose the order to open the boxes, and stop at will
- She (adaptively) opens a subset of boxes $S \subseteq [n]$, and wish to maximize the expected value of

$$R = \max_{i \in S} x_i - \sum_{i \in S} c_i.$$




Adaptive Policies

$x_1 = 1$ w.p. 0.5
$x_1 = 5$ w.p. 0.5
$c_1 = 0$

$x_2 = 3$ w.p. 1.0

$c_2 = 1$

- Policy 1: first open box 1. If $x_1 = 1$, then open box 2.
  - E[reward]=0.5*(3-1)+0.5*5=3.5
- Policy 2: first open box 2. Then open box 1.
  - E[reward]=0.5*(3-1)+0.5*(5-1)=3.0

Pandora rule

A surprisingly simple index policy, a so-called **Pandora rule**.

- SELECTION RULE: If a box is to be opened, it should be that closed box with highest reservation price.
- STOPPING RULE: Terminate search whenever the maximum sampled reward exceeds the reservation price of every closed box.

Reservation price:

$$x_i^* = \inf\Big\{y : y \geq -c_i + E[\max\{x_i, y\}]\Big\}$$
$$= \inf\Big\{y : c_i \geq E[\max\{x_i - y, 0\}]\Big\}.$$

Pandora rule – we are lucky

- "That such an elementary decision strategy as Pandora's Rule is optimal depends more crucially than might be supposed on the simplifying assumptions of the model. There does not seem to be available a sharp characterization of an optimal solution when certain features of the present model are changed. Pandora's Rule does not readily generalize." (Weitzman, 1979)

- Like the famous Gittin's index for Markovian Bandit

Probe the MAX value

- Almost the same setting as Pandora's problem, except that
- Boxes have no cost, but she can open at most a set S of k boxes
- Goal: maximize the expect value $E[\max_{i \in S} x_i]$

- Seems easy: pick the boxes with highest reservation prices??
- The reservation price technique doesn't work here!
- Unfortunately, no simple optimal policy is known
- Probably the optimal policy is an exponentially large decision tree
- Hardness? (could be PSPACE-hard. Open.)

A 1-1/e approximation

- Consider function defined over subsets of [n]

$$f(S) = E[\max_{i \in S} X_i]$$

- It can be shown that $f(S)$ is a submodular function

  ($f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$ for any S, T)

- By known result for online submodular optimization [Asadpour *et al.*], the greedy algorithm is a 1-1/e approximation
- IS there a PTAS??

A PTAS for ProbeMAX

- Decision Tree

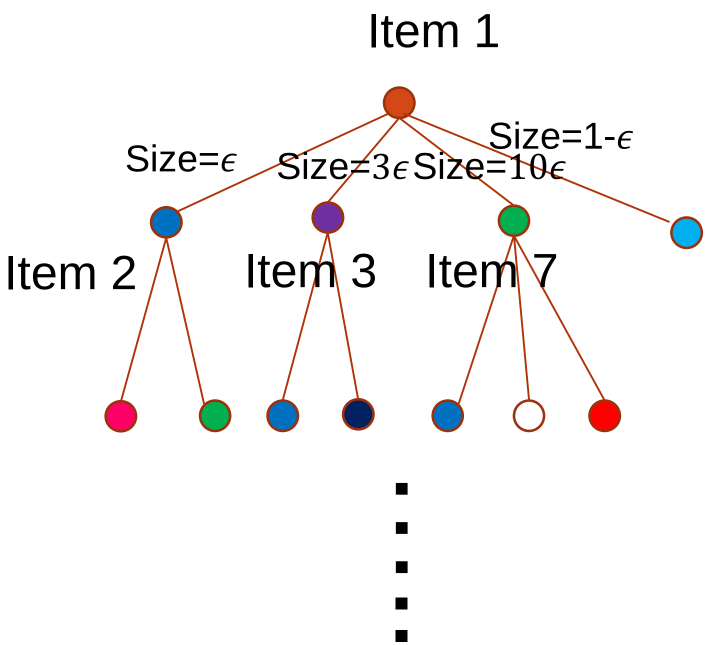

Exponential size!!!! (depth=n)

How to represent such a tree? Compact solution?

A PTAS for ProbeMAX

- By discretization, we make some simplifying assumptions:
  - Support of the size distribution: $(0, \epsilon, 2\epsilon, 3\epsilon, \ldots\ldots, 1)$.

Still way too many possibilities, how to narrow the search space?

Block Adaptive Policies

- Block Adaptive Policies: Process items block by block

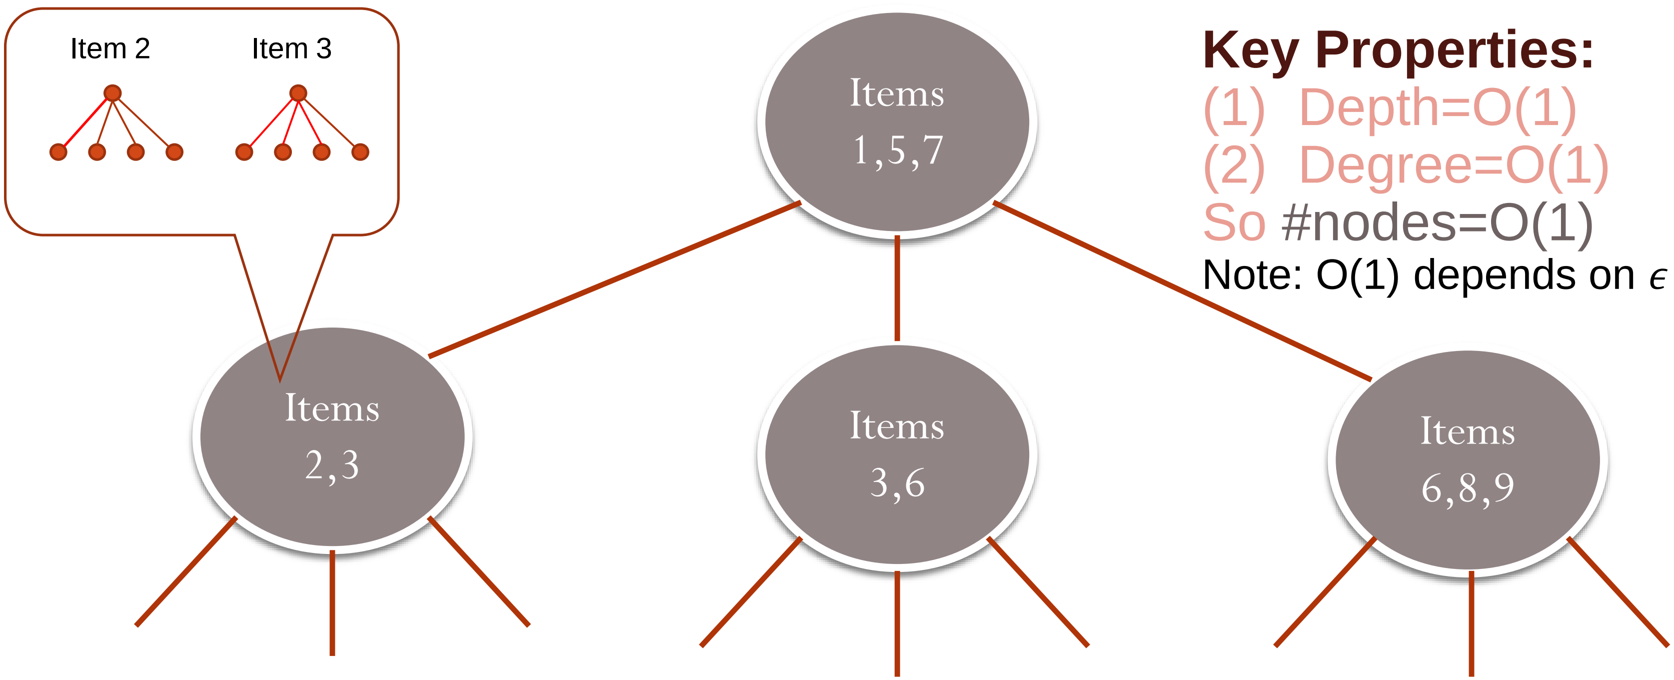

Key Properties:
(1) Depth=O(1)
(2) Degree=O(1)
So #nodes=O(1)
Note: O(1) depends on ϵ

LEMMA: There is a block adaptive policy that is nearly optimal

Block Adaptive Policies

- Block Adaptive Policies: Process items block by block

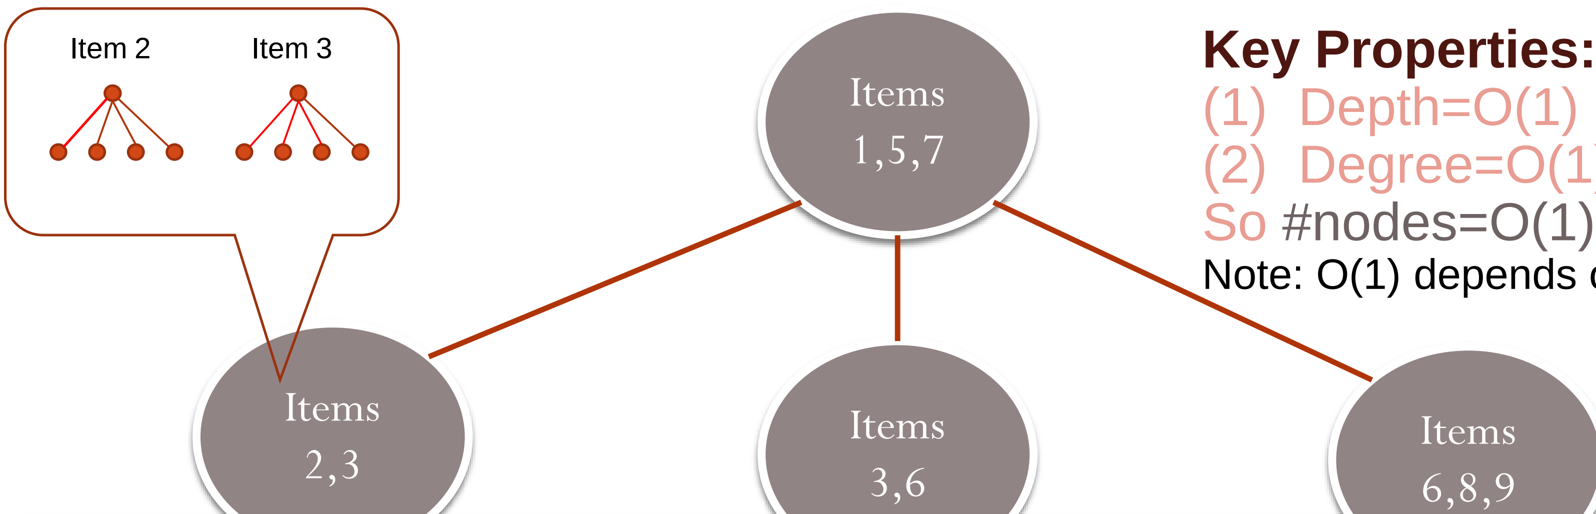

Key Properties:
(1) Depth=O(1)
(2) Degree=O(1)
So #nodes=O(1)
Note: O(1) depends on ϵ

Still exponential many possibilities, even in a single block

LEMMA: There is a block adaptive policy that is nearly optimal

Algorithm

- Outline: Enumerate all block structures with a signature (similar to that in Poisson approximation) associated with each node

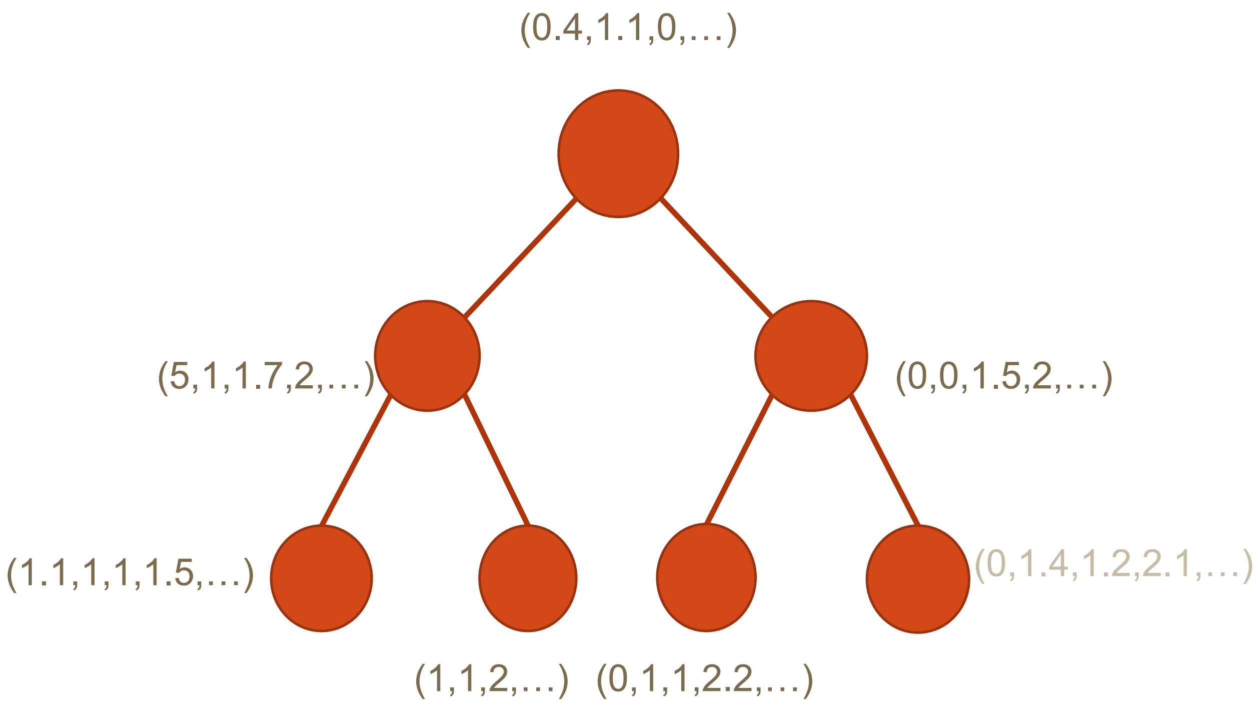

- O(1) nodes
- Poly(n) possible signatures for each node
- So total #configuration =poly(n)

Algorithm

2. Find an assignment of items to blocks that matches all signatures
   – (this can be done by standard dynamic programming)

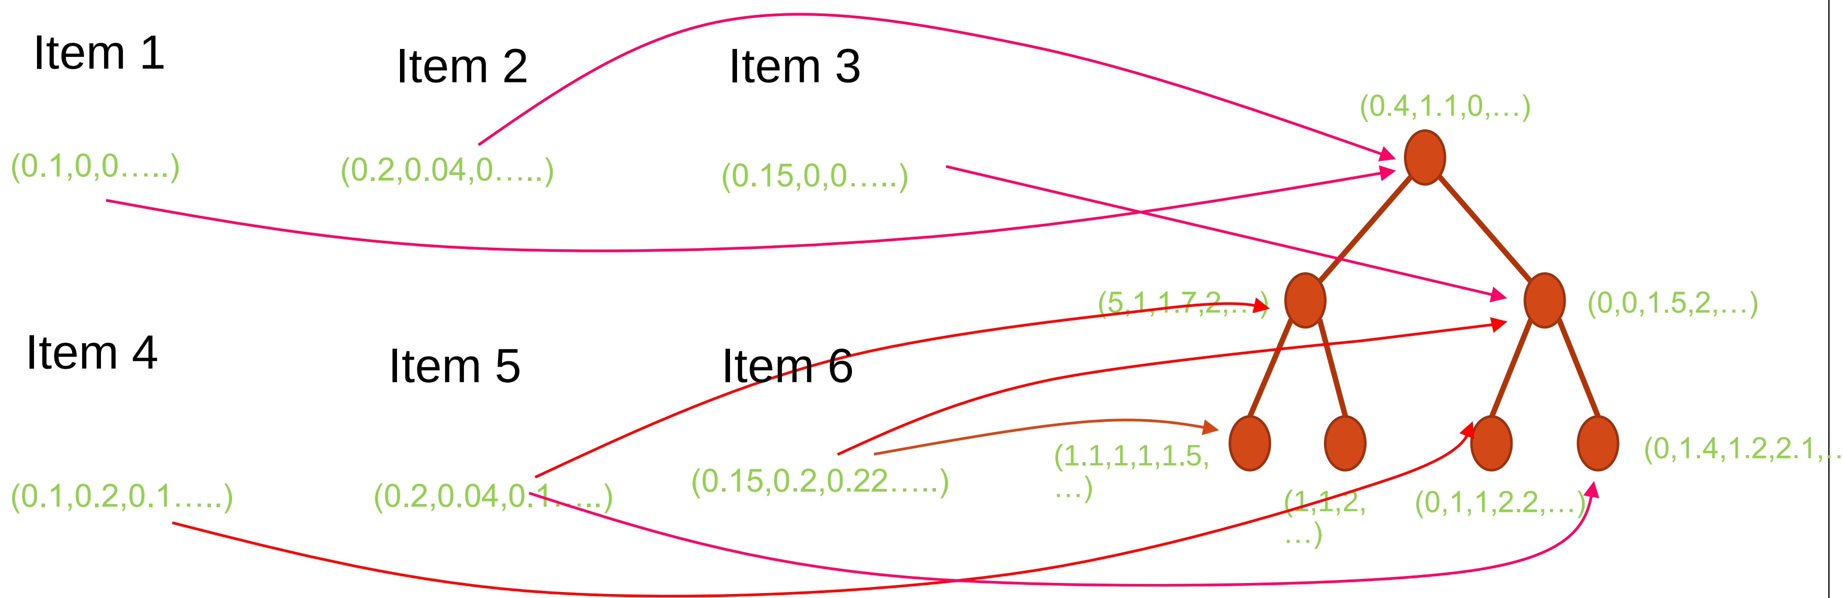

On any root-leaf path, each item appears at most once

Stochastic Knapsack

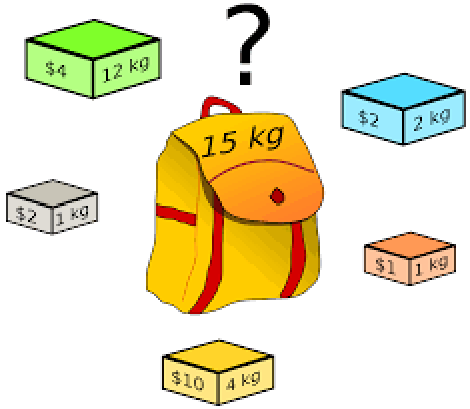

- A knapsack of capacity C
- A set of items, each having a fixed profit
- Known: Prior distr of size of each item.
- Each time we choose an item and place it in the knapsack irrevocably
- The actual size of the item becomes known after the decision
- Knapsack constraint: The total size of accepted items <= C
- Goal: maximize E[Profit]

[Dean et al. FOCS05] [Dean et al. MOR08] [Bhalgat et al.SODA 11] [Gupta et al.FOCS 13] [LY STOC13] [Will SODA14] …..

Motivation

- Stochastic Scheduling
  - Jobs, each having an uncertain length, and a fixed profit
  - You have C hours
  - How to (adaptively) schedule them (maximize E[profit])

Jobs:

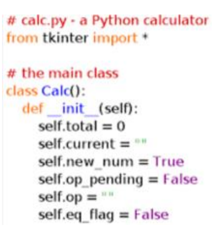

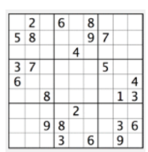

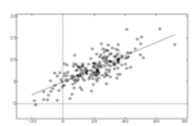

Profits: 20$ 5$ 10$ 50$

Running time:

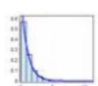

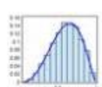

C=5 hours

A unified approach

- We realize that the stochastic knapsack problem can be solved by similar technique
- Later, we found other variants of Pandora's box problem can be solved by similar technique (have to change several places)

- Very tedious…..

- A unified approach
- Dynamic programming recursion:

$$\mathrm{DP}_t(\mathcal{A}_t, I_t) = \max_{a_t \in \mathcal{A}_t} \mathbb{E}\Big[\mathrm{DP}_{t+1}(\mathcal{A}_t \setminus a_t, f(\cdot \,|\, I_t, a_t)) + g(\cdot \,|\, I_t, a_t)\Big], \quad t = 1, \ldots, T.$$

Our result

THM(informal) [Fu, L, Xu, ICALP'18] Under some assumption of the number of states, and the properties of the transitions, we obtain a PTAS for solving such stochastic dynamic program.

Example:

- ProbeMax (best known 1-1/e [Asadpour *et al., Management science 15*])
- ProbeMax-(m,k) (constant approx. [Munagala 16])
- Committed Pandora's box (constant approx. by known technique)
- Stochastic Knapsack (recover results in [Balghat et al., SODA11] [L,Yuan, STOC13])
- Threshold Probability Stochastic Knapsack
  (previously only heuristic [İlhan *et al. Operation Research 11*])
- Bayesian online selection with knapsack constraint

Outline

- Estimation (no decision)
- Single-stage
- 2-stage/Multi-stage
- Online (adaptive/non-adaptive))
- Sample Complexity

The Stochastic Multi-armed Bandit

- We don't know the distribution. We can only take samples.
- Stochastic Multi-armed Bandit
  - Set of n arms
  - Each arm is associated with an unknown reward distribution supported on $[0,1]$ with mean θ_i
  - Each time, sample an arm and receive the reward independently drawn from the reward distribution

classic problems in stochastic control, stochastic optimization and online learning

Best Arm Identification

- Best-arm Identification: Find the best arm out of n arms, with means $\mu_{[1]}, \mu_{[n]},.., \mu_{[n]}$ (for simplicity, assume they follows Gaussian distr with unit variance)
- Goal: use as few samples as possible
- Formulated by Bechhofer in 1954

- Applications: medical trails, A/B test, crowdsourcing, team formation, many extensions….
- Close connections to regret minimization

Stochastic Multi-armed Bandit

- Statistics， medical trials (Bechhofer, 54) ,Optimal control，Industrial engineering (Koenig & Law, 85), evolutionary computing (Schmidt, 06), Simulation optimization (Chen, Fu, Shi 08),Online learning (Bubeck Cesa-Bianchi,12)

- [Bechhofer, 58] [Farrell, 64] [Paulson, 64] [Bechhofer, Kiefer, and Sobel, 68],…., [Even-Dar, Mannor, Mansour, 02] [Mannor, Tsitsiklis, 04] [Even-Dar, Mannor, Mansour, 06] [Kalyanakrishnan, Stone 10] [Gabillon, Ghavamzadeh, Lazaric, Bubeck, 11] [Kalyanakrishnan, Tewari, Auer, Stone, 12] [Bubeck, Wang, Viswanatha, 12]….[Karnin, Koren, and Somekh, 13] [Chen, Lin, King, Lyu, Chen, 14]

- Books:
  - Multi-armed Bandit Allocation Indices, John Gittins, Kevin Glazebrook, Richard Weber, 2011
  - Regret analysis of stochastic and nonstochastic multi-armed bandit problems S. Bubeck and N. Cesa-Bianchi., 2012
  - ……

Applications

- **Clinical Trails**
  - **One arm – One treatment**
  - **One pull – One experiment**

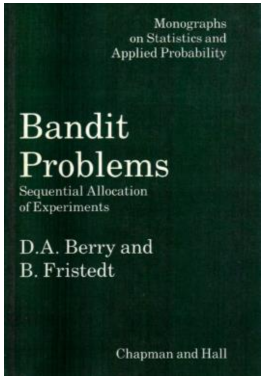

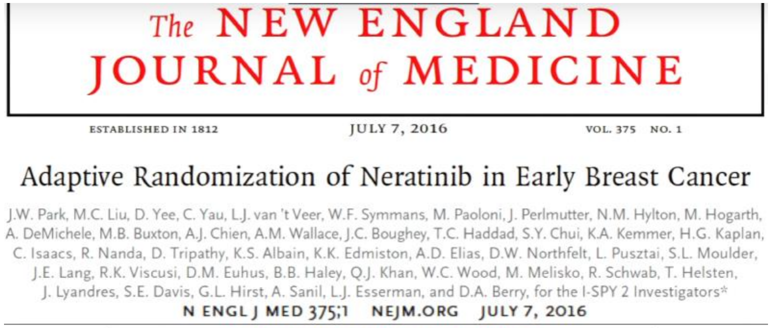

The NEW ENGLAND JOURNAL of MEDICINE

ESTABLISHED IN 1812 JULY 7, 2016 VOL. 375 NO. 1

Adaptive Randomization of Neratinib in Early Breast Cancer

J.W. Park, M.C. Liu, D. Yee, C. Yau, L.J. van 't Veer, W.F. Symmans, M. Paoloni, J. Perlmutter, N.M. Hylton, M. Hogarth, A. DeMichele, M.B. Buxton, A.J. Chien, A.M. Wallace, J.C. Boughey, T.C. Haddad, S.Y. Chui, K.A. Kemmer, H.G. Kaplan, C. Isaacs, R. Nanda, D. Tripathy, K.S. Albain, K.K. Edmiston, A.D. Elias, D.W. Northfelt, L. Pusztai, S.L. Moulder, J.E. Lang, R.K. Viscusi, D.M. Euhus, B.B. Haley, Q.J. Khan, W.C. Wood, M. Melisko, R. Schwab, T. Helsten, J. Lyandres, S.E. Davis, G.L. Hirst, A. Sanil, L.J. Esserman, and D.A. Berry, for the I-SPY 2 Investigators*

N ENGL J MED 375;1 NEJM.ORG JULY 7, 2016

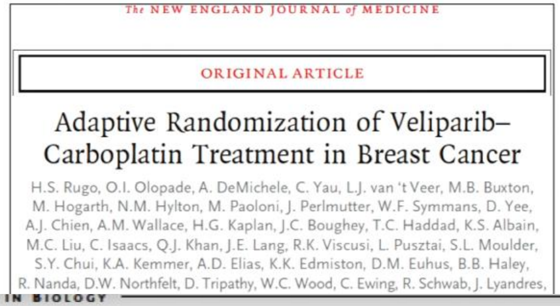

The NEW ENGLAND JOURNAL of MEDICINE

ORIGINAL ARTICLE

Adaptive Randomization of Veliparib–Carboplatin Treatment in Breast Cancer

H.S. Rugo, O.I. Olopade, A. DeMichele, C. Yau, L.J. van 't Veer, M.B. Buxton, M. Hogarth, N.M. Hylton, M. Paoloni, J. Perlmutter, W.F. Symmans, D. Yee, A.J. Chien, A.M. Wallace, H.G. Kaplan, J.C. Boughey, T.C. Haddad, K.S. Albain, M.C. Liu, C. Isaacs, Q.J. Khan, J.E. Lang, R.K. Viscusi, L. Pusztai, S.L. Moulder, S.Y. Chui, K.A. Kemmer, A.D. Elias, K.K. Edmiston, D.M. Euhus, B.B. Haley, R. Nanda, D.W. Northfelt, D. Tripathy, W.C. Wood, C. Ewing, R. Schwab, J. Lyandres,

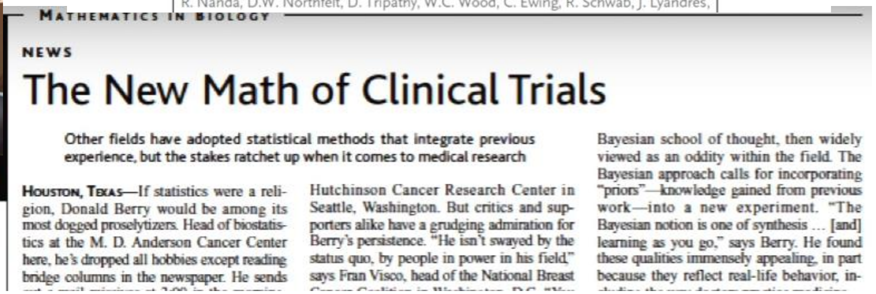

MATHEMATICS IN BIOLOGY

NEWS

The New Math of Clinical Trials

Other fields have adopted statistical methods that integrate previous experience, but the stakes ratchet up when it comes to medical research

HOUSTON, TEXAS—If statistics were a religion, Donald Berry would be among its most dogged proselytizers. Head of biostatistics at the M. D. Anderson Cancer Center here, he's dropped all hobbies except reading bridge columns in the newspaper. He sends

Hutchinson Cancer Research Center in Seattle, Washington. But critics and supporters alike have a grudging admiration for Berry's persistence. "He isn't swayed by the status quo, by people in power in his field," says Fran Visco, head of the National Breast

Bayesian school of thought, then widely viewed as an oddity within the field. The Bayesian approach calls for incorporating "priors"—knowledge gained from previous work—into a new experiment. "The Bayesian notion is one of synthesis ... [and] learning as you go," says Berry. He found these qualities immensely appealing, in part because they reflect real-life behavior, in-

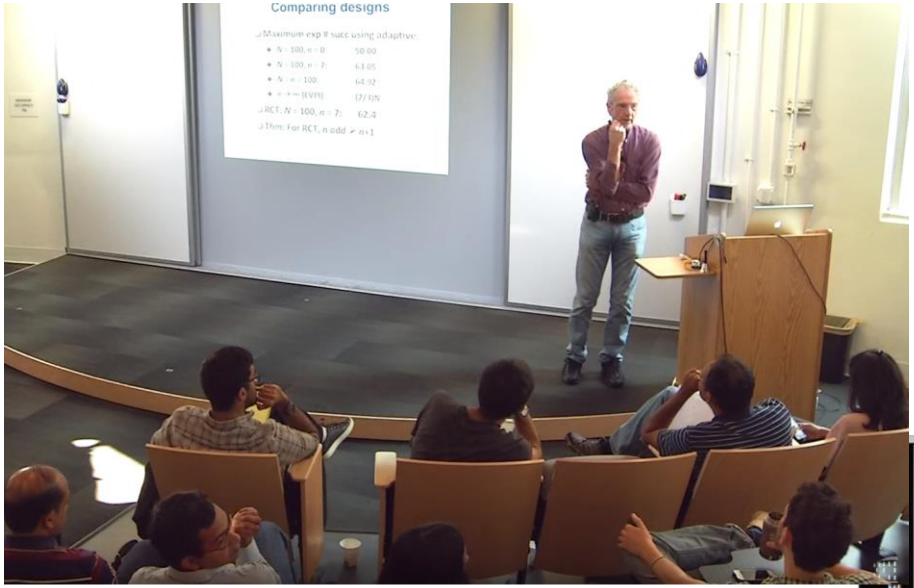

Don Berry, University of Texas
MD Anderson Cancer Center

Best-1-Arm: – a misclaim

Exact version:

Source	Sample Complexity
Even-Dar et al. [12]	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \left(\ln \delta^{-1} + \ln n + \ln \Delta_{[i]}^{-1} \right)$
Gabillon et al. [16]	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \left(\ln \delta^{-1} + \ln \sum_{i=2}^{n} \Delta_{[i]}^{-2} \right)$
Jamieson et al. [19]	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \left(\ln \delta^{-1} + \ln \ln \left(\sum_{j=2}^{n} \Delta_{[j]}^{-2} \right) \right)$
kalyanakrishnan et al. [23]	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \left(\ln \delta^{-1} + \ln \sum_{i=2}^{n} \Delta_{[i]}^{-2} \right)$
Jamieson et al. [19]	$\ln \delta^{-1} \cdot \left(\ln \ln \delta^{-1} \cdot \sum_{i=2}^{n} \Delta_{[i]}^{-2} + \sum_{i=2}^{n} \Delta_{[i]}^{-2} \ln \Delta_{[i]}^{-1} \right)$
Karnin et al.[24], Jamieson et al.[20]	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \left(\ln \delta^{-1} + \ln \ln \Delta_{[i]}^{-1} \right)$
This paper (Thm 2.5)	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \left(\ln \delta^{-1} + \ln \ln \min(n, \Delta_{[i]}^{-1}) \right) + \Delta_{[2]}^{-2} \ln \ln \Delta_{[2]}^{-1}$
This paper (clustered instances) Thm B.22	$\sum_{i=2}^{n} \Delta_{[i]}^{-2} \ln \delta^{-1} + \Delta_{[2]}^{-2} \ln \ln \Delta_{[2]}^{-1}$

Mannor-Tsitsiklis lower bound: $\Omega\left(\sum_{i=2}^{n} \Delta_{[i]}^{-2} \ln \delta^{-1} \right)$

Farrell's lower bound (2 arms): $\Delta_{[2]}^{-2} \ln \ln \Delta_{[2]}^{-1}$

Attempting to believe : Karnin's upper bound is optimal

misclaim

Instance Optimal Sample Complexity

- (almost) Instance optimal algorithm for best arm

$$H_i = \sum_{u \in G_i} \Delta_{[u]}^{-2}$$

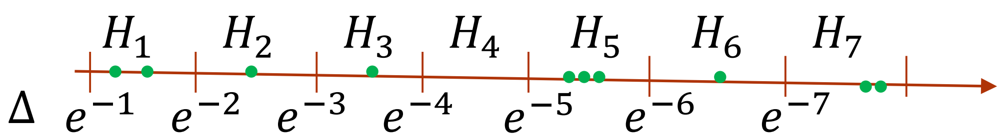

- Gap Entropy: $\mathsf{Ent}(I) = \sum_{G_i \neq \emptyset} p_i \log p_i^{-1}.$ $\qquad p_i = H_i / \sum_j H_j.$

- An instance-wise lower bound $\mathcal{L}(I,\delta) = \Theta\left(H(I)(\ln \delta^{-1} + \mathsf{Ent}(I))\right).$
- An algorithm with sample complexity: $H(I) = \sum_{i=2}^{n} \Delta_{[i]}^{-2}.$

$$O\left(\mathcal{L}(I,\delta) + \Delta_{[2]}^{-2} \ln \ln \Delta_{[2]}^{-1}\right).$$

We almost achieve the above bounds, modulo some small additive term (getting rid of it is an OPEN question) (for the upper bound), and some mild assumption (for the lower bound).

[Chen, Li. ArXiv 15]

[Chen, Li., Qiao COLT 17]

Combinatorial Pure Exploration

Combinatorial Pure Exploration in multi-armed bandit

- A general combinatorial constraint on the feasible set of arms
  - Best-k-arm: the uniform matroid constraint
  - First studied by [Chen et al. NIPS14]
  - E.g., we want to build a MST. But each time get a noisy estimate of the true cost of each edge

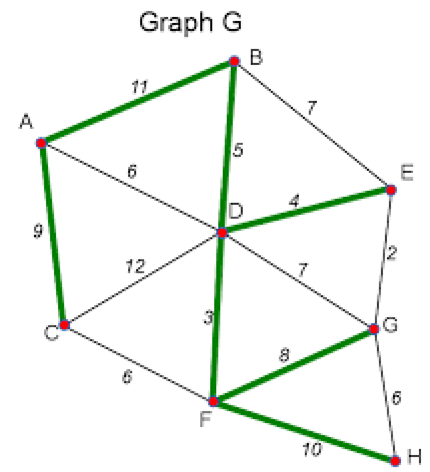

- More general combinatorial constraints
  - [Chen et al. NIPS 14][CG**L**. COLT'16] [CG**L**QW. COLT'17] [Cao et al. COLT'18]
  - Optimal upper lower bounds for general constraints: Still OPEN.

Conclusion

- Bayesian mechanism design (essentially stochastic optimization problems)
- Learning+Optimization
  - We don't have to first learn the distributions first, and then solve the stochastic optimization problem. We can do it together and use less samples!
- Interesting connections to many subareas in TCS (counting, coresets, geometry, VC theory, Boolean functions, bandits, online algorithms, mechanism design,….) and probability theory/statistics
- A lot more interesting problems to be studied
- Many open problems
- A Survey: Jian Li, Yu Liu. Approximation Algorithms for Stochastic Combinatorial Optimization Problems.

Thanks

lapordge@gmail.com

Wechat: lapordge