Mathematical Optimization in Machine Learning/Decision-Making

IAS Distinguished Lecture

May 8, 2023

Yinyu Ye
Stanford University and CUHKSZ (Sabbatical Leave)

Ask ChatGPT about an Optimization Problem?

Knapsack on ChatGPT:

假设有5件物品和一个最多能承重5kg的背包。第1件物品的重量是2kg，其价值为18元；第2件物品的重量是1kg，其价值为10元；第3件物品的重量是2.5kg，其价值为20元；第4件物品的重量是3.5kg，其价值为28元；第5件物品的重量是1.5kg，其价值为15元。请问将哪些物品装入背包里，物品的价值总和最大？请使用整数规划求解该问题。

这是一个经典的背包问题，可以使用整数规划方法求解。

我们可以定义0/1变量x_i，表示是否将第i件物品放入背包中。同时，我们还需要定义一个非负整数变量z，表示物品的总价值。

然后，我们可以将问题转化为以下的线性规划问题：

最大化 $z=18x_1+10x_2+20x_3+28x_4+15x_5$

满足约束条件：

$2x_1+x_2+2.5x_3+3.5x_4+1.5x_5 \leq 5$

$x_i\in {0,1}, i=1,2,3,4,5$

由于x_i是整数变量，这是一个整数规划问题。可以使用整数规划求解器求解。

求解结果表明，将第1件、第3件、第5件物品装入背包中，可以使得物品的总价值最大，为$18+20+15=53$元。

ChatGPT 在不调用求解器的前提下，能否求解大规模优化问题？

在不调用求解器的前提下，对于大规模的优化问题，通常也比较直接求解得到最优解，因为现实中的大多数优化问题都非常复杂，存在着多个目标、多个约束、非线性、离散等多种问题，很容易遇到组合爆炸的问题。

ChatGPT:

* ChatGPT tried to formulate problem and attempted to give an answer， but could not guarantee the correctness。

* ChatGPT highly recommended using reliable optimization solvers and consulting specialists

Linear Programming and LP Giants won Nobel Prize…

$$\max \quad \sum \pi_j x_j$$

$$\text{s.t.} \quad \sum_j \boldsymbol{a}_j x_j \le \boldsymbol{b},$$

$$0 \le x_j \le 1 \quad \forall\, j = 1, \ldots, n$$

Today's Sharing

1. Online Linear Programming Algorithms and Applications

2. Accelerated Second-Order Methods for Nonlinear Optimization and Applications

3. Pre-Trained Mixed Integer Linear Programming Solvers and Applications

4. Equitable Covering & Partition – Divide and Conquer and Applications

Topic 1. Online Linear Programming
an Online Auction Example

- **There is a fixed selling period or number of buyers; and there is a fixed inventory of goods**
- **Customers come and require a bundle of goods and make a bid**
- **Decision: To sell or not to sell to each individual customer on the fly?**
- **Objective: Maximize the revenue.**

Bid #	$100	$30	….	…	…	Inventory
Decision	x1	x2				
Pants	1	0	….	…	…	100
Shoes	1	0				50
T-Shirts	0	1				500
Jackets	0	0				200
Hats	1	1	…	…	…	1000

Online Linear Programming Model and Theory

- **OLP theory and practice (Agrawal et al. 2010,14, Li&Y 2022)**

- OLP:

$$
\begin{aligned}
\max \quad & \sum \pi_j x_j \\
\text{s.t.} \quad & \sum_j \boldsymbol{a}_j x_j \le \boldsymbol{b}, \\
& 0 \le x_j \le 1 \quad \forall\, j = 1,\ldots,n
\end{aligned}
$$

  - Variables together with their data points arrive sequentially and decision makers need decide x_j on the fly, that is, before knowing the “future” data points
  - Learning-while-Doing vs Learning-First and Deciding-Second (collect and learn all relevant data, then solve for all $\boldsymbol{x}$)
  - Offline LP’s objective value is a upper bond for the online version
  - Is there an optimal online decision algorithm/mechanism

Price Mechanism for Online Auction

- Learn and compute itemized optimal prices
- Use the prices to price each bid
- Accept if it is a over bid, and reject otherwise
- There is an Optimal Online Algorithm to achieve the best you could do!

Bid #	$100	$30	….	…	…	Inventory	Price?
Decision	x1	x2					
Pants	1	0	….	…	…	100	45
Shoes	1	0				50	45
T-Shirts	0	1				500	10
Jackets	0	0				200	55
Hats	1	1	…	…	…	1000	15

App. I: Online Matching for Display Advertising

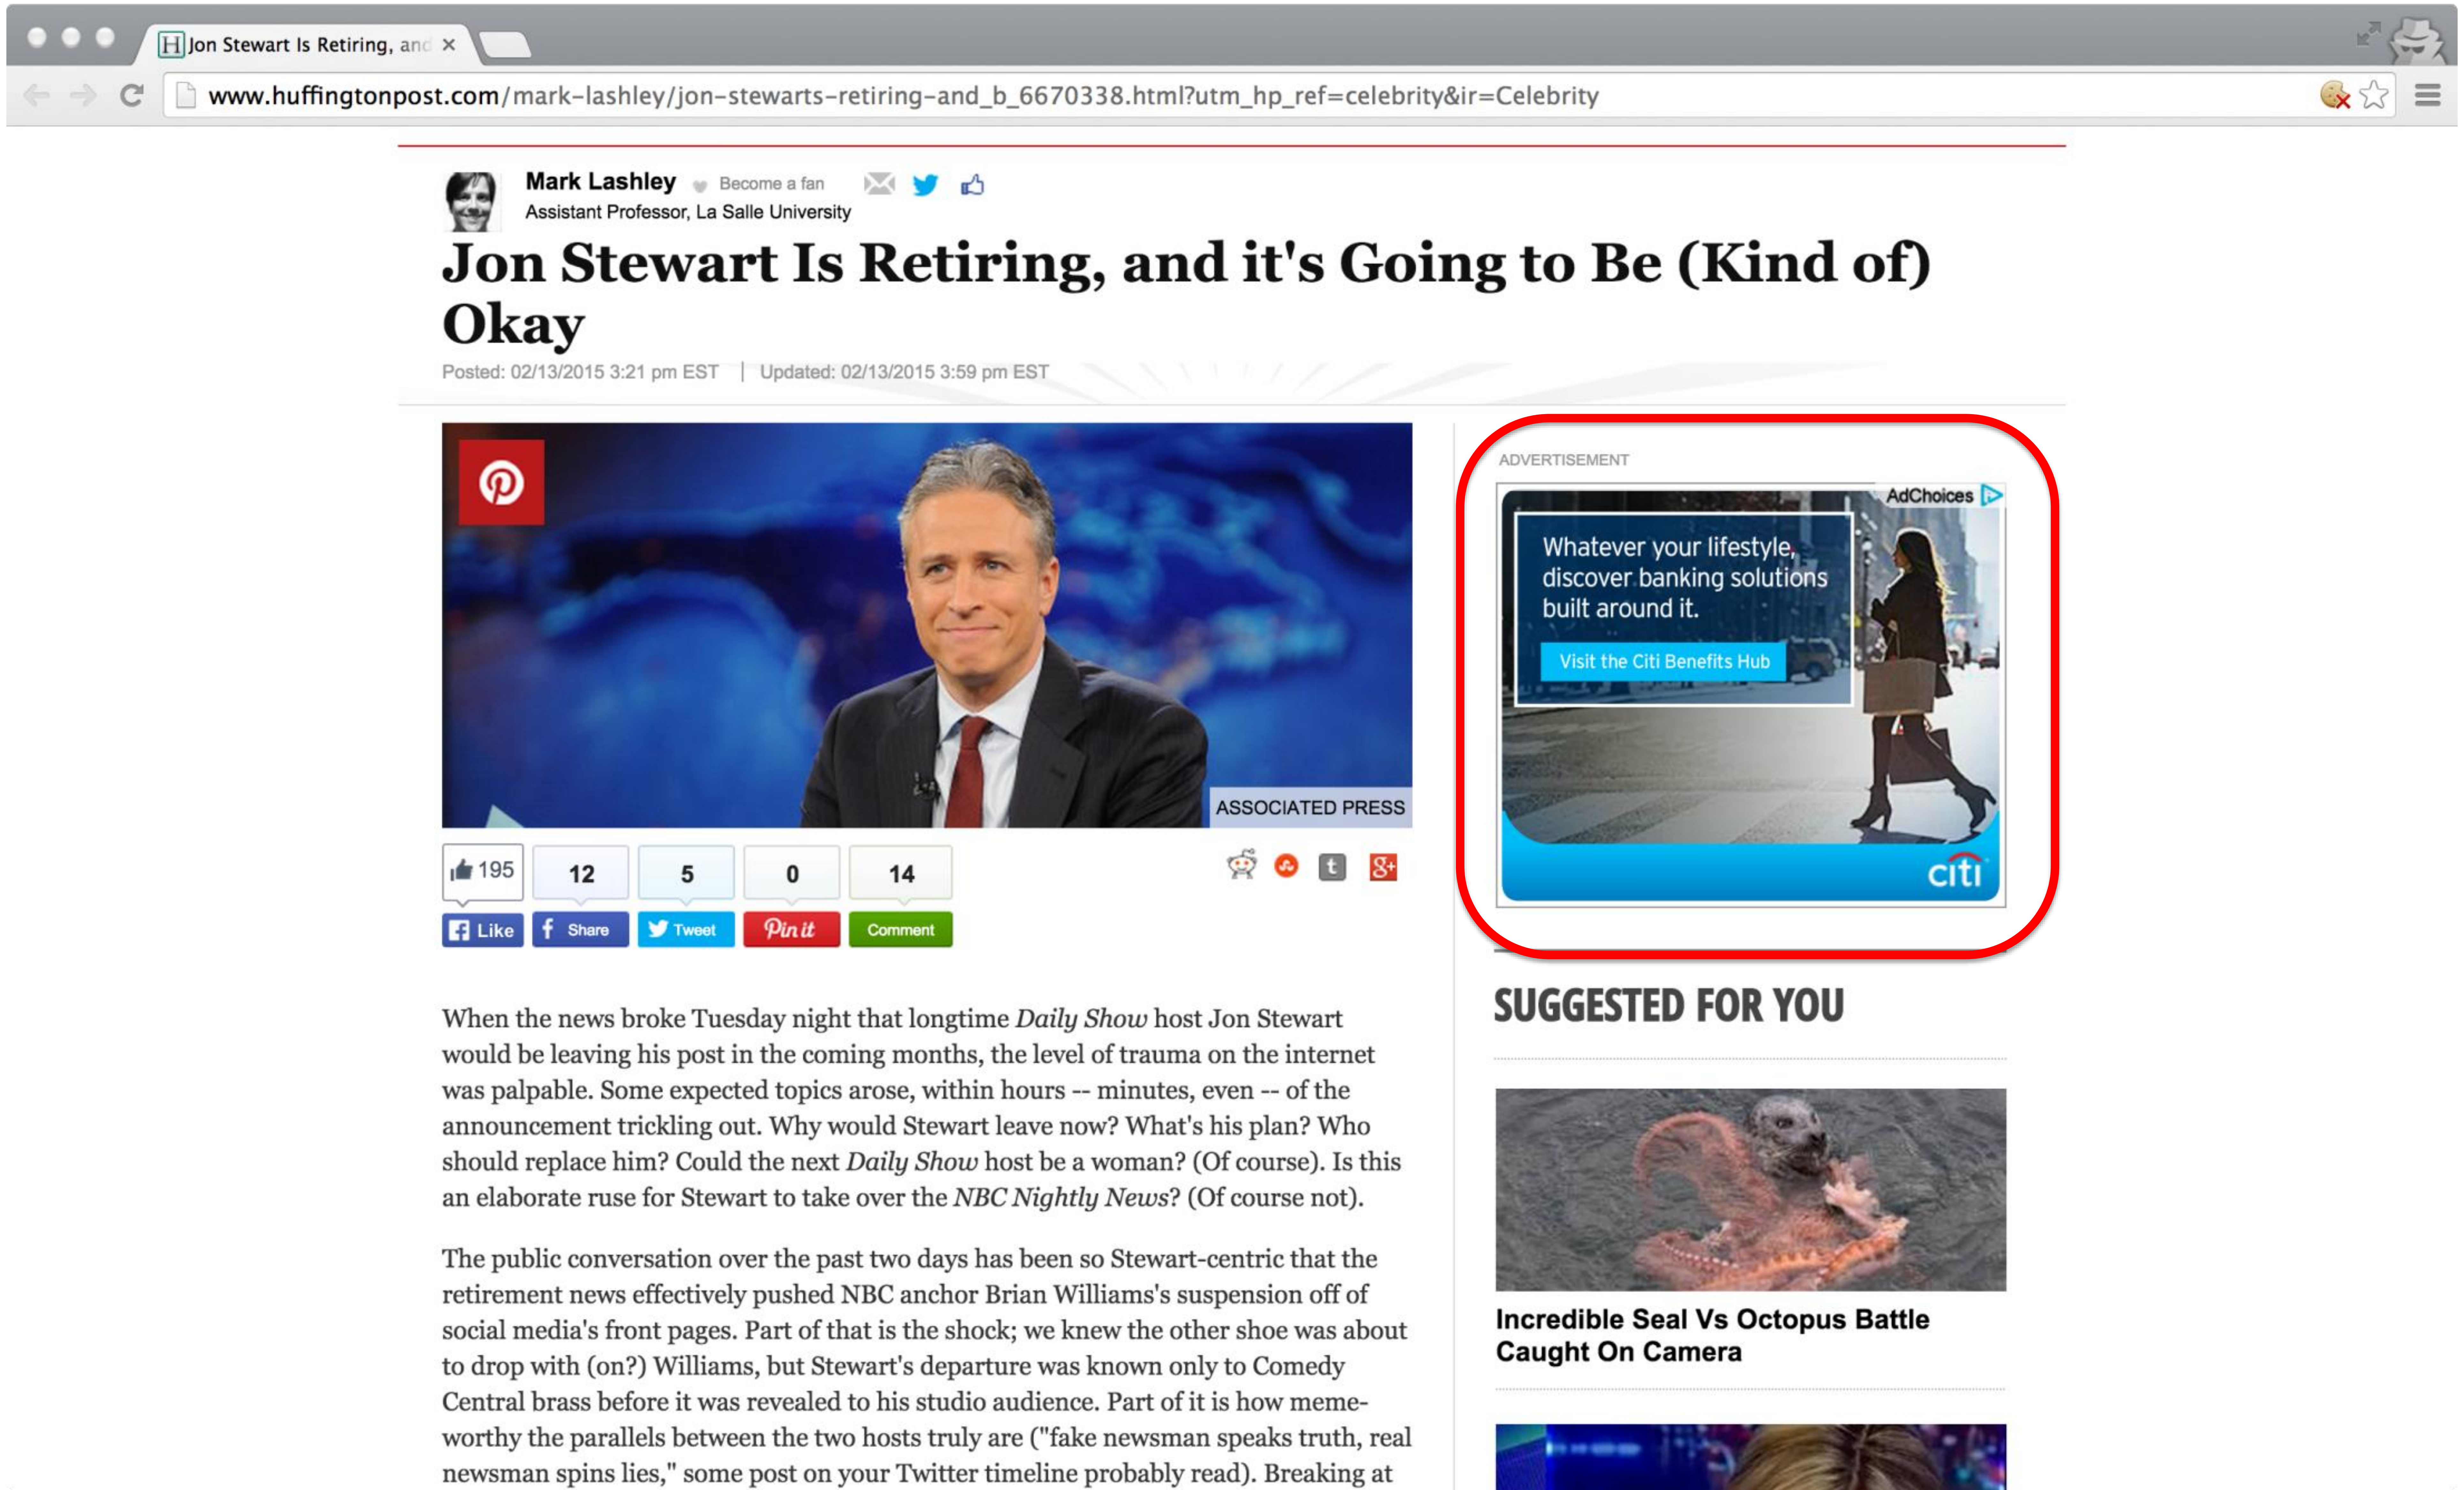

Jon Stewart Is Retiring, and

www.huffingtonpost.com/mark-lashley/jon-stewarts-retiring-and_b_6670338.html?utm_hp_ref=celebrity&ir=Celebrity

Mark Lashley Become a fan
Assistant Professor, La Salle University

Jon Stewart Is Retiring, and it's Going to Be (Kind of) Okay

Posted: 02/13/2015 3:21 pm EST | Updated: 02/13/2015 3:59 pm EST

ASSOCIATED PRESS

195 Like | 12 Share | 5 Tweet | 0 Pin it | 14 Comment

When the news broke Tuesday night that longtime *Daily Show* host Jon Stewart would be leaving his post in the coming months, the level of trauma on the internet was palpable. Some expected topics arose, within hours -- minutes, even -- of the announcement trickling out. Why would Stewart leave now? What's his plan? Who should replace him? Could the next *Daily Show* host be a woman? (Of course). Is this an elaborate ruse for Stewart to take over the *NBC Nightly News*? (Of course not).

The public conversation over the past two days has been so Stewart-centric that the retirement news effectively pushed NBC anchor Brian Williams's suspension off of social media's front pages. Part of that is the shock; we knew the other shoe was about to drop with (on?) Williams, but Stewart's departure was known only to Comedy Central brass before it was revealed to his studio audience. Part of it is how meme-worthy the parallels between the two hosts truly are ("fake newsman speaks truth, real newsman spins lies," some post on your Twitter timeline probably read). Breaking at

SUGGESTED FOR YOU

Incredible Seal Vs Octopus Battle Caught On Camera

Revenues generated by different methods

- Total Revenue for impressions in T2 by Greedy and OLP with different allocation risk functions

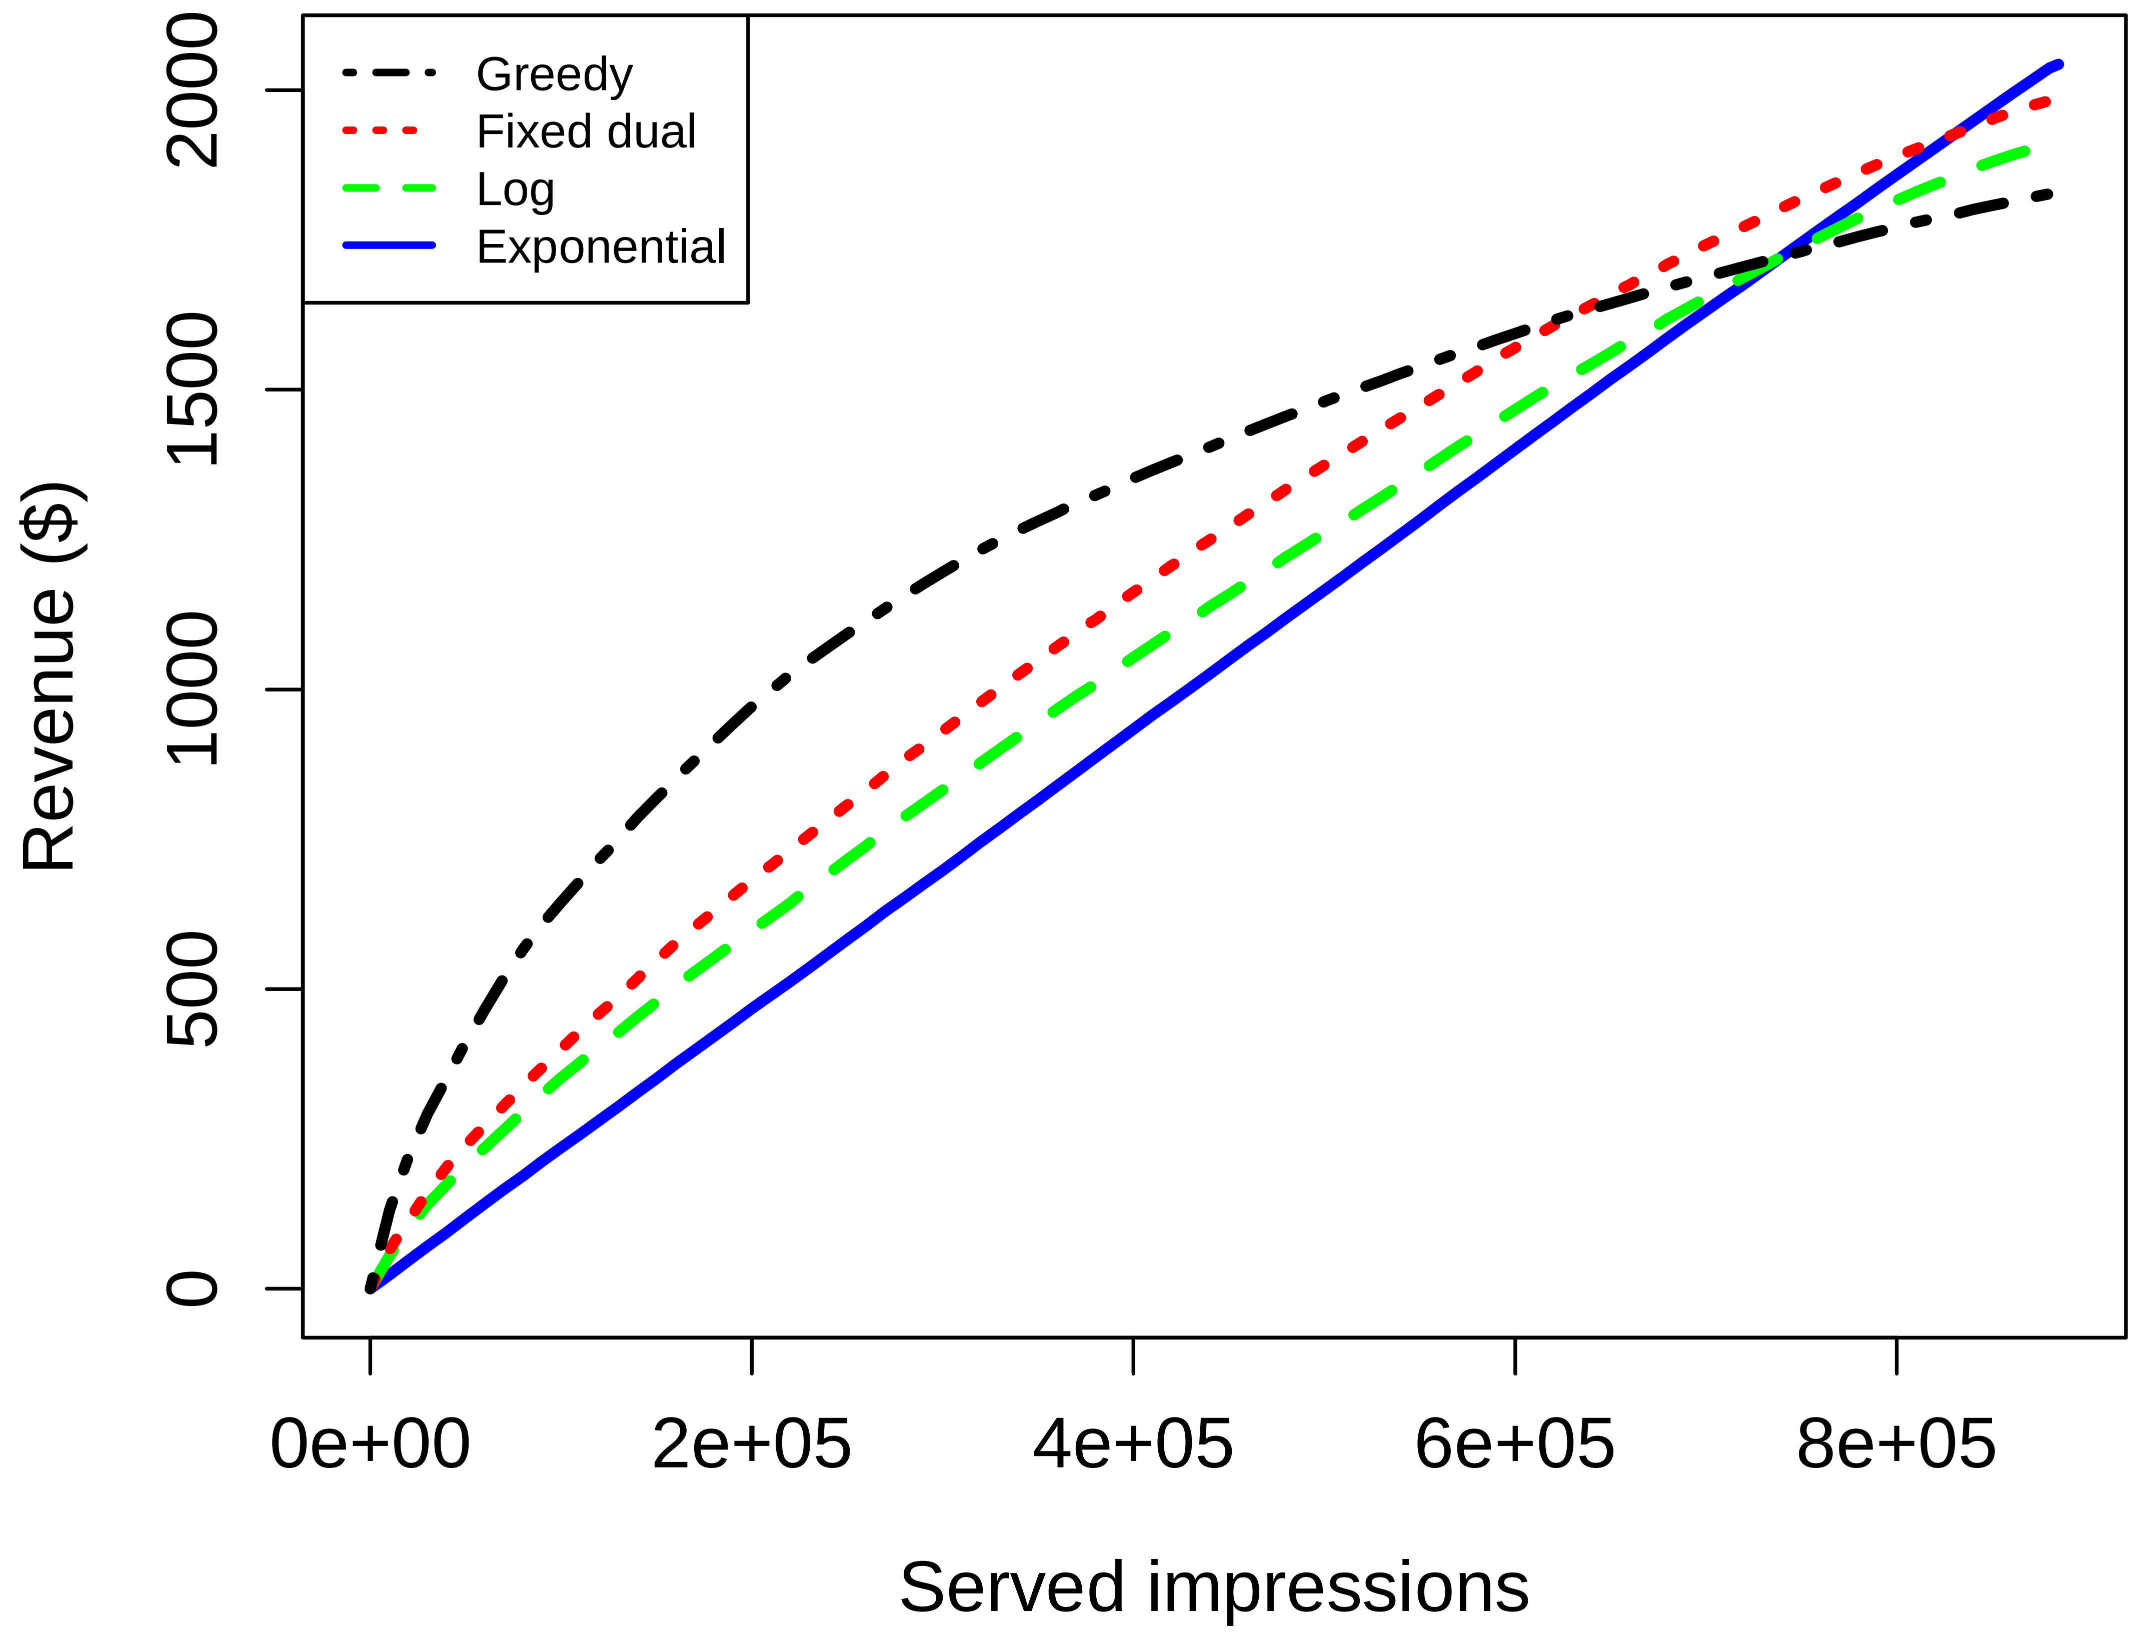

of Out-of-Budget Advertisers

- Greedy exhausts budget of many advertisers early.
- Log penalty keeps advertisers in budget but it is very conservative.
- Exponential penalty Keeps advertisers in budget until almost the end of the timeframe.

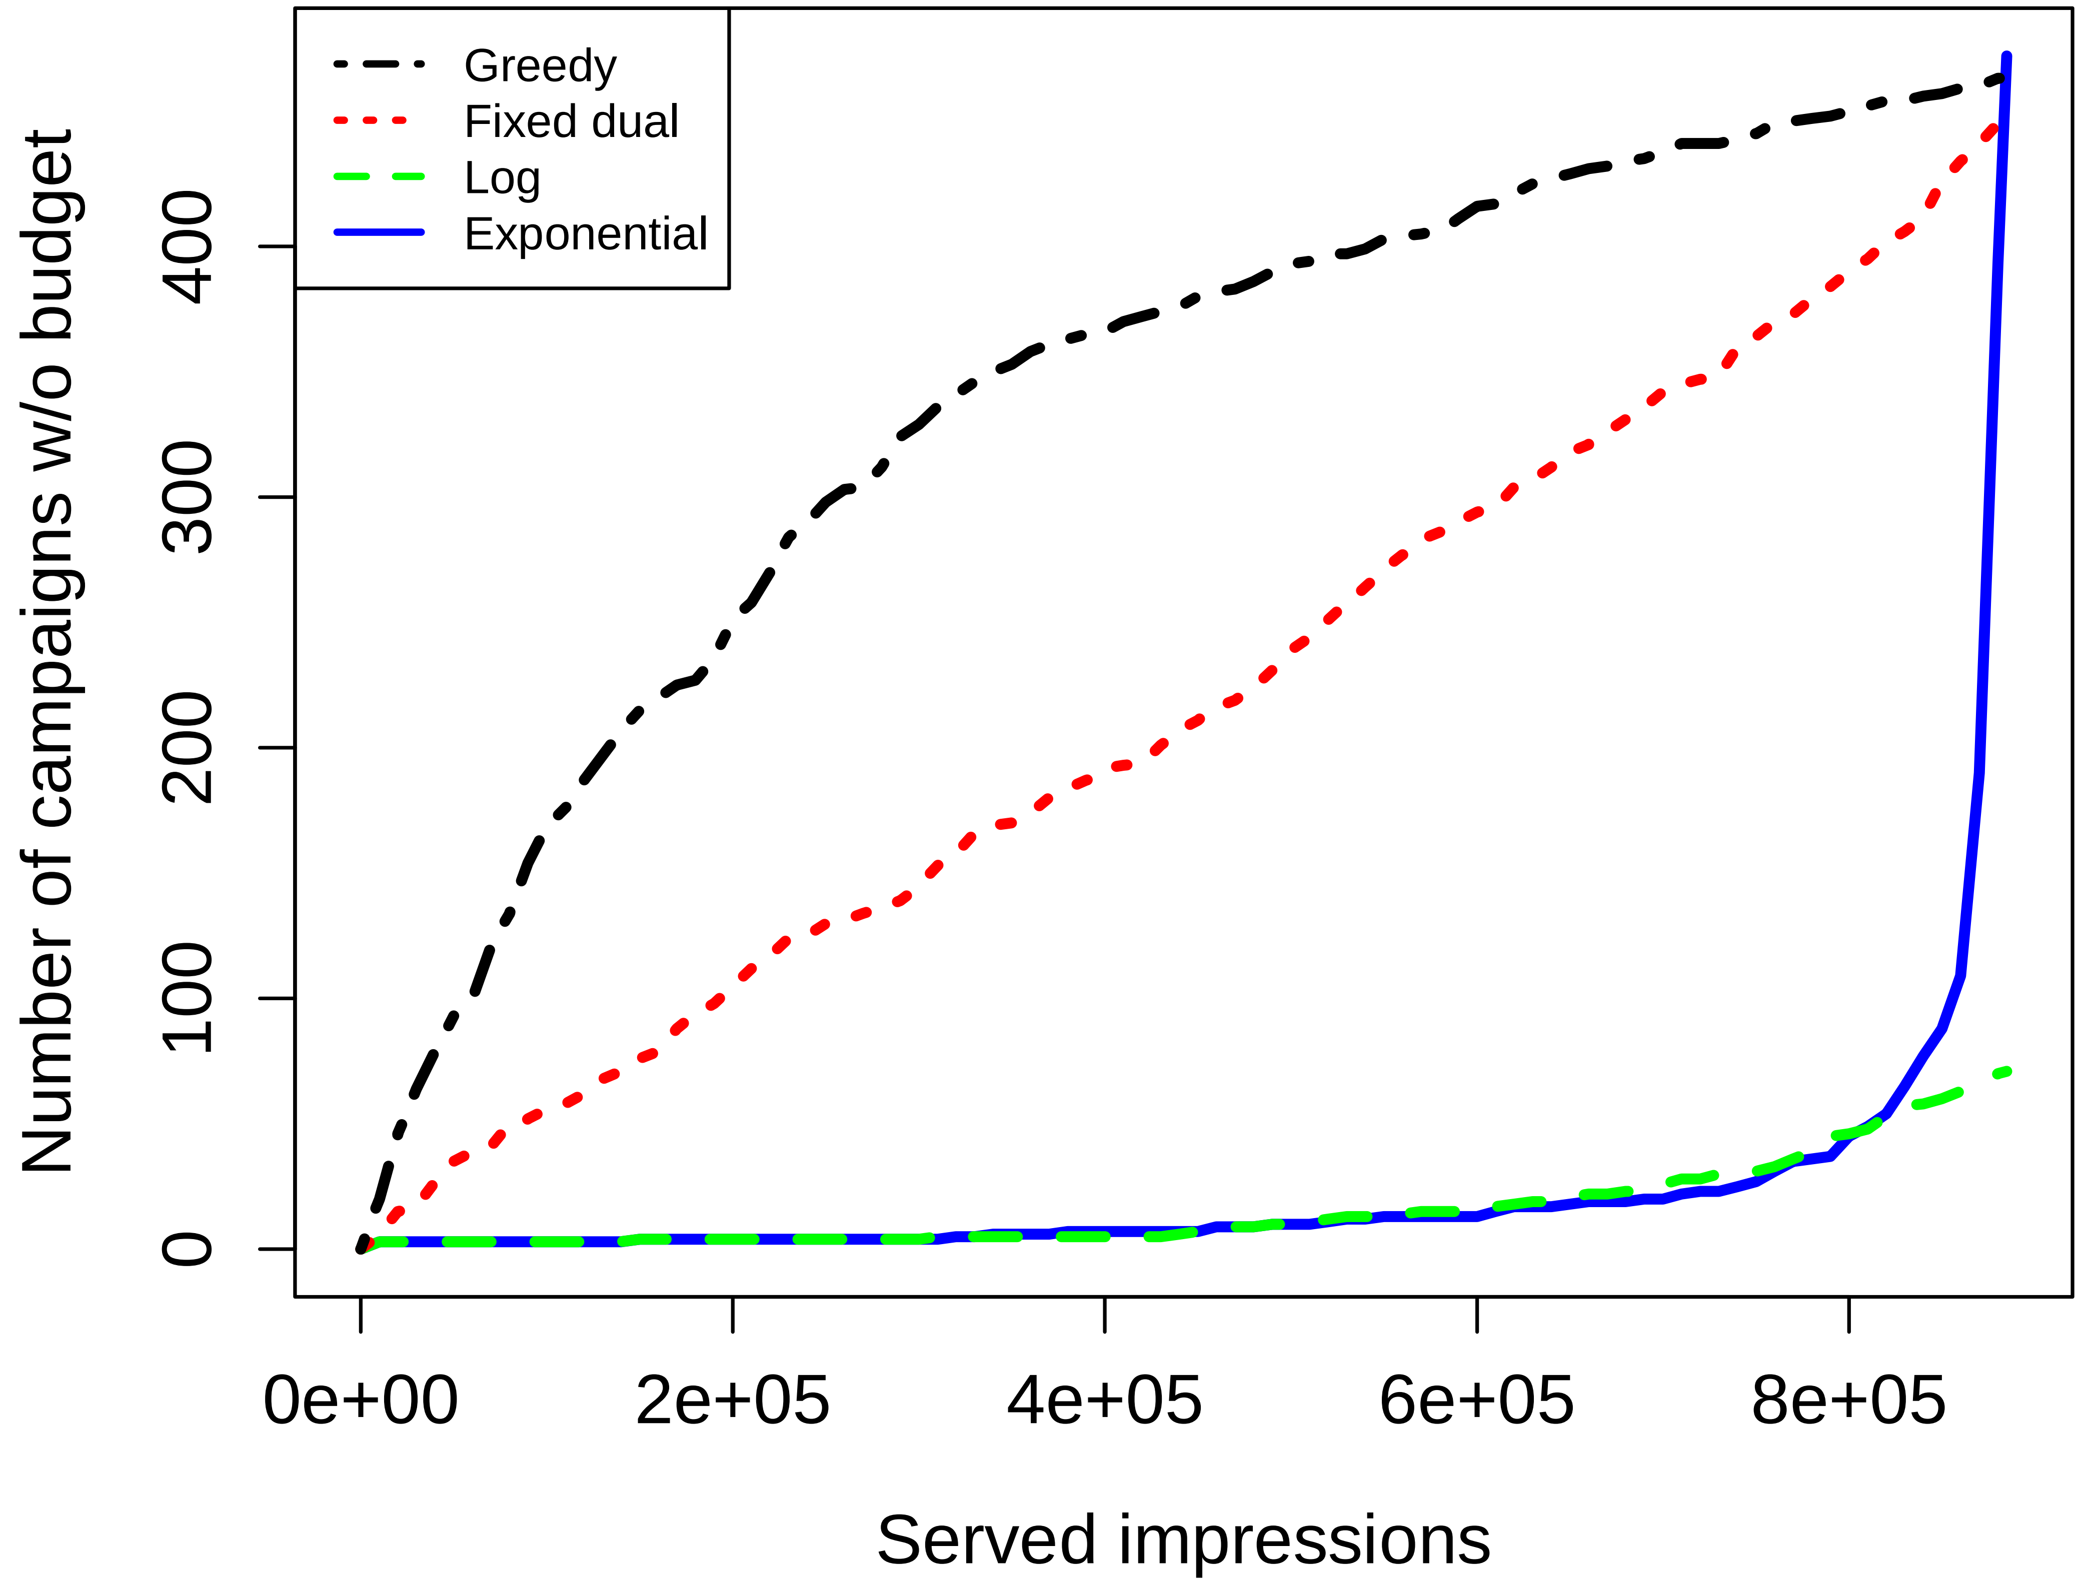

阿里巴巴在2019年云栖大会上提到在智能履行决策上使用OLP的算法

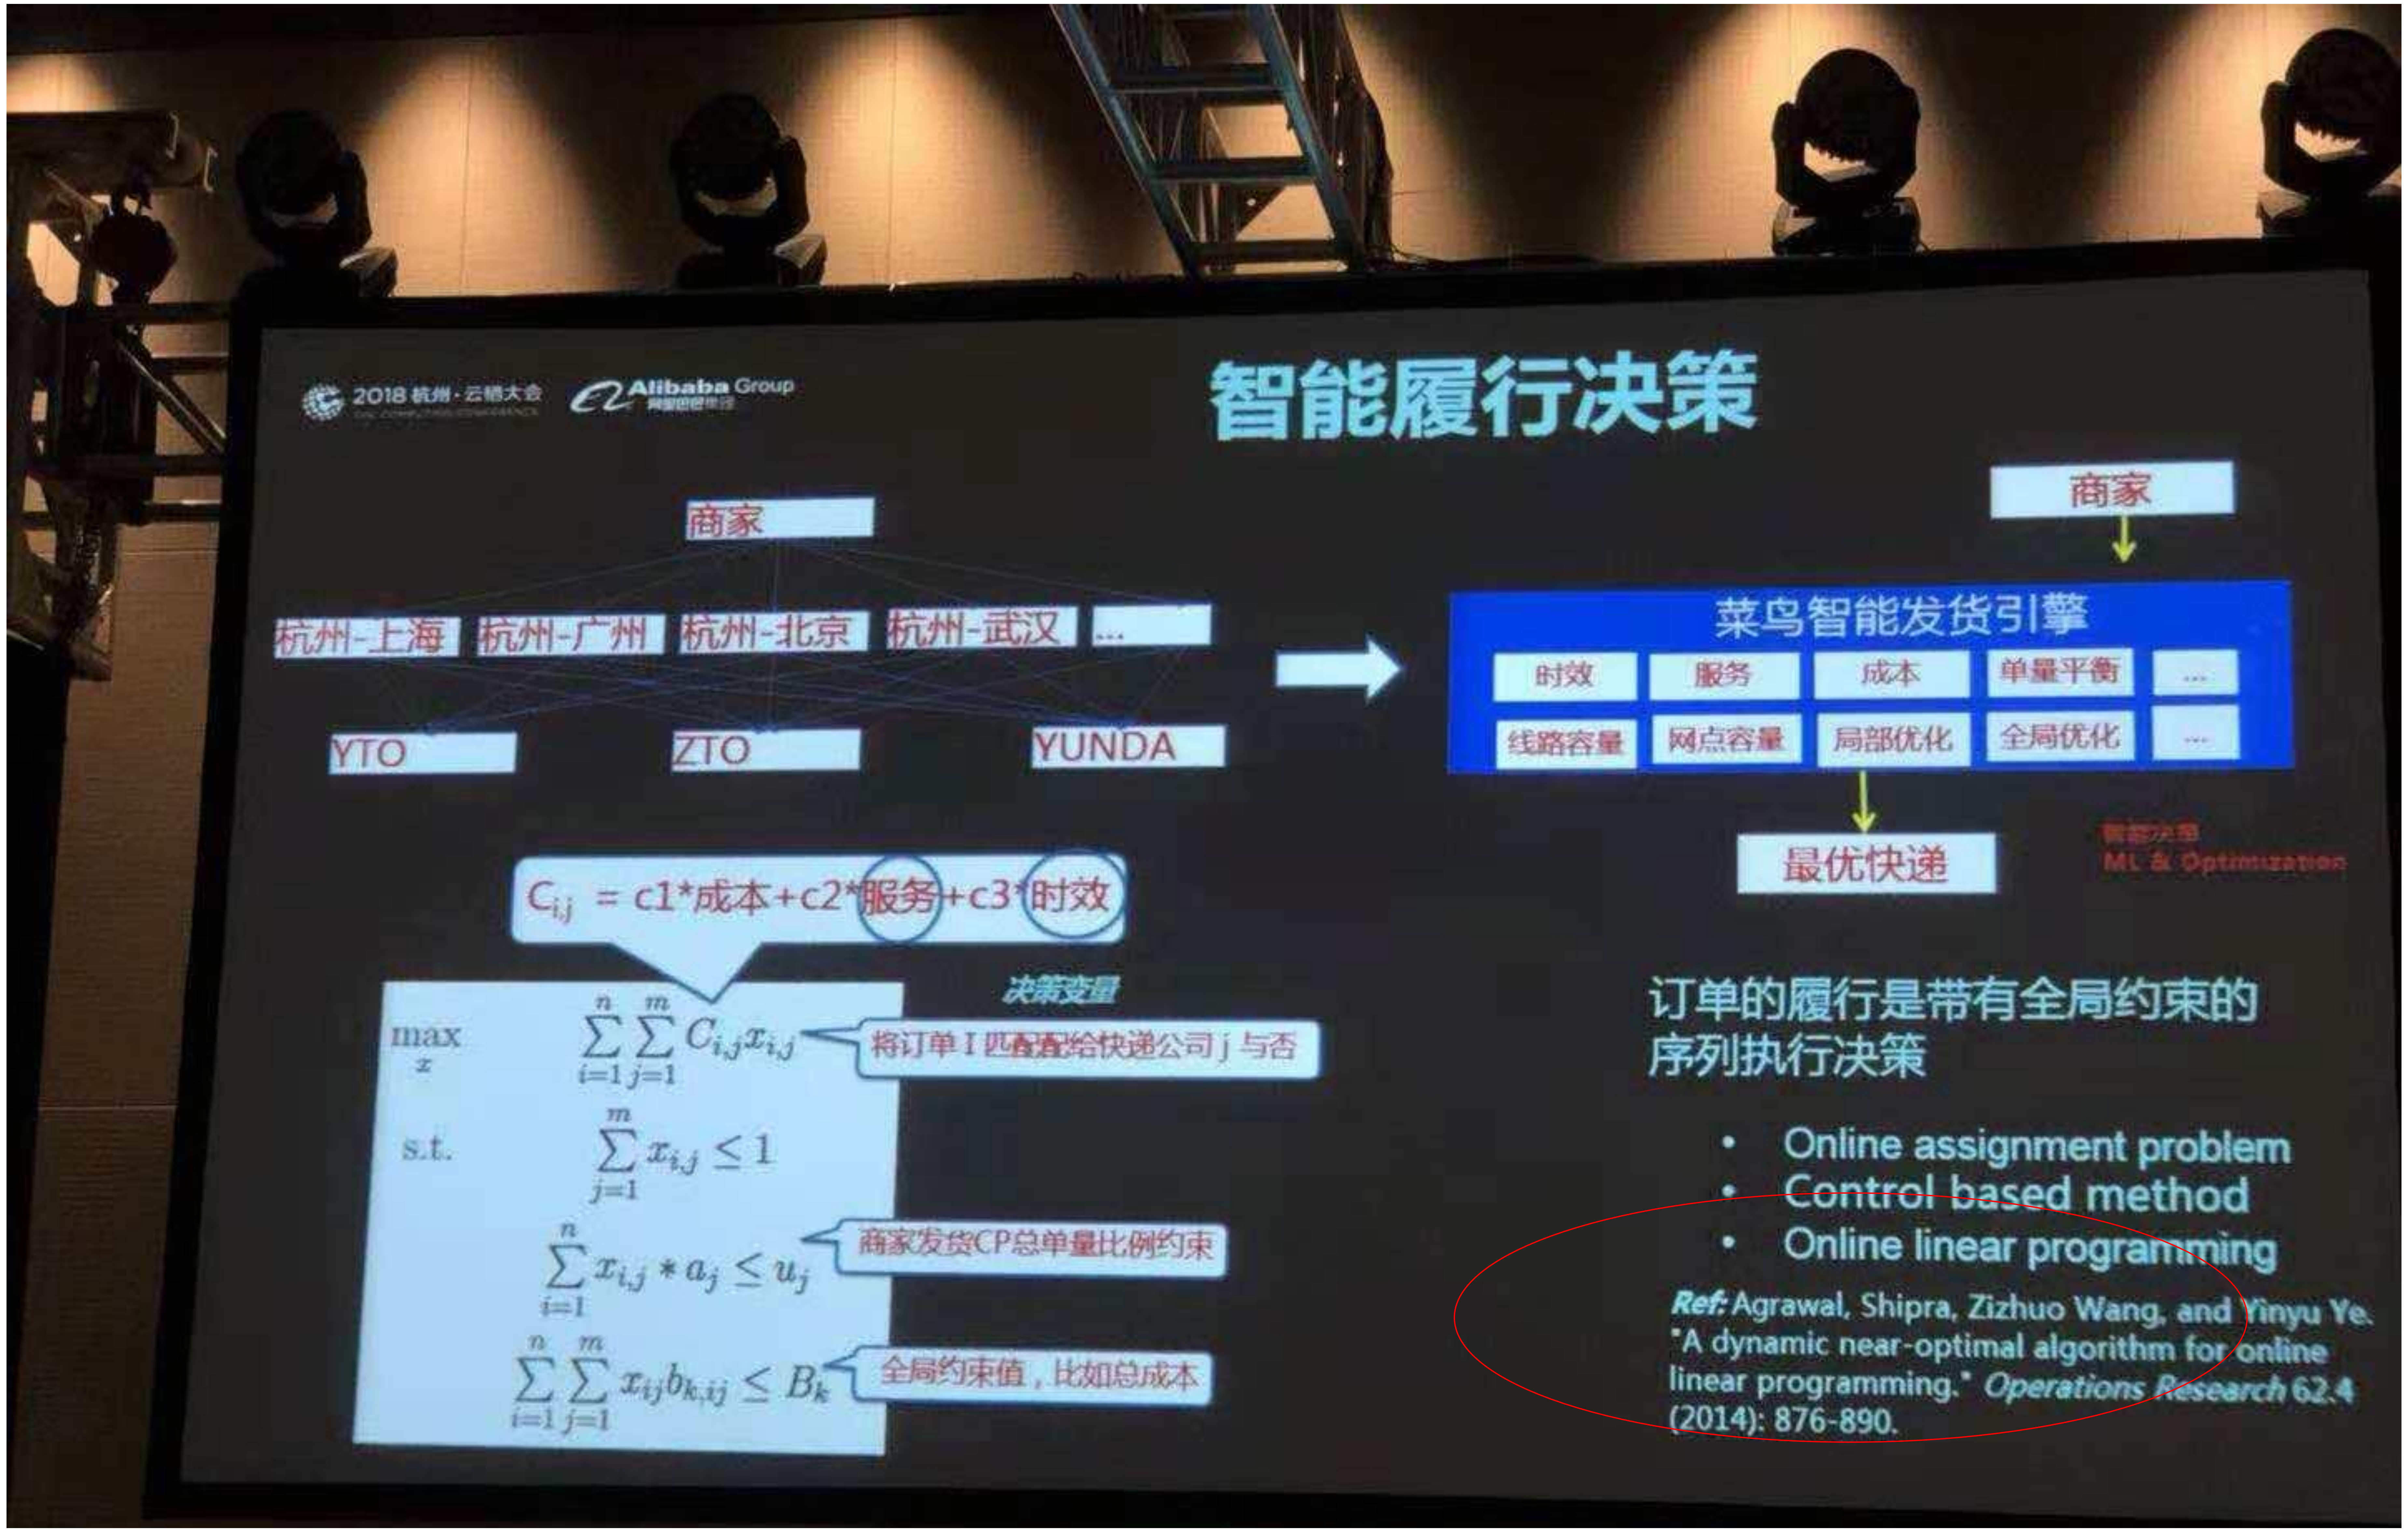

阿里巴巴团队在2020年CIKM会议论文Online Electronic Coupon Allocation based on Real-Time User Intent Detection上提到他们设计的发红包的机制也使用了OLP的方法[2]

Spending Money Wisely: Online Electronic Coupon Allocation based on Real-Time User Intent Detection

Liangwei Li*
Liucheng Sun*
leon.llw@alibaba-inc.com
liucheng.slc@alibaba-inc.com
Alibaba Group
Hangzhou, Zhejiang

Chenwei Weng
wengchenwei.pt@alibaba-inc.com
Alibaba Group
Hangzhou, Zhejiang

Chengfu Huo
chengfu.huocf@alibaba-inc.com
Alibaba Group
Hangzhou, Zhejiang

Weijun Ren
afei@alibaba-inc.com
Alibaba Group
Hangzhou, Zhejiang

$$
\begin{aligned}
\max \ & \sum_{i=1}^{M} \sum_{j=1}^{N} v_{ij} x_{ij} \\
s.t. \ & \sum_{i=1}^{M} \sum_{j=1}^{N} c_j x_{ij} \le B, \\
& \sum_{j}^{N} x_{ij} \le 1, \quad \forall i \\
& x_{ij} \ge 0, \quad \forall i, j
\end{aligned}
\tag{5}
$$

3.3 MCKP-Allocation

We adopt the primal-dual framework proposed by [2] to solve the problem defined in Equation 5. Let α and β_j be the associated dual variables respectively. After obtaining the dual variables, we can solve the problem in an online fashion. Precisely, according to the principle of the primal-dual framework, we have the following allocation rule:

$$
x_{ij} = \begin{cases} 1, & \text{where } j = \arg\max_j (v_{ij} - \alpha c_j) \\ 0, & \text{otherwise} \end{cases} \tag{9}
$$

App. II: The Online Algorithm can be Extended to Bandits with Knapsack (BwK) Applications

- For the previous problem, the decision maker first wait and observe the customer order/arm and then decide whether to accept/play it or not.
- An alternative setting is that the decision maker first decides which order/arm (s)he may accept/play, and then receive a random resource consumption vector $\mathbf{a}_j$ and yield a random reward π_j of the pulled arm.
- Known as the Bandits with Knapsacks, and it is a tradeoff exploration v.s. exploitation

- **First algorithm** to achieve the $O(\log T)$ regret bound [Li, Sun & Y 2021 ICML]

(**https://proceedings.mlr.press/v139/li21s.html**)

App. III: Online Market Place: design online posted-prices to minimize regret of the Eisenberg/Gale social welfare while achieving agent “envy-freeness” and market clearing

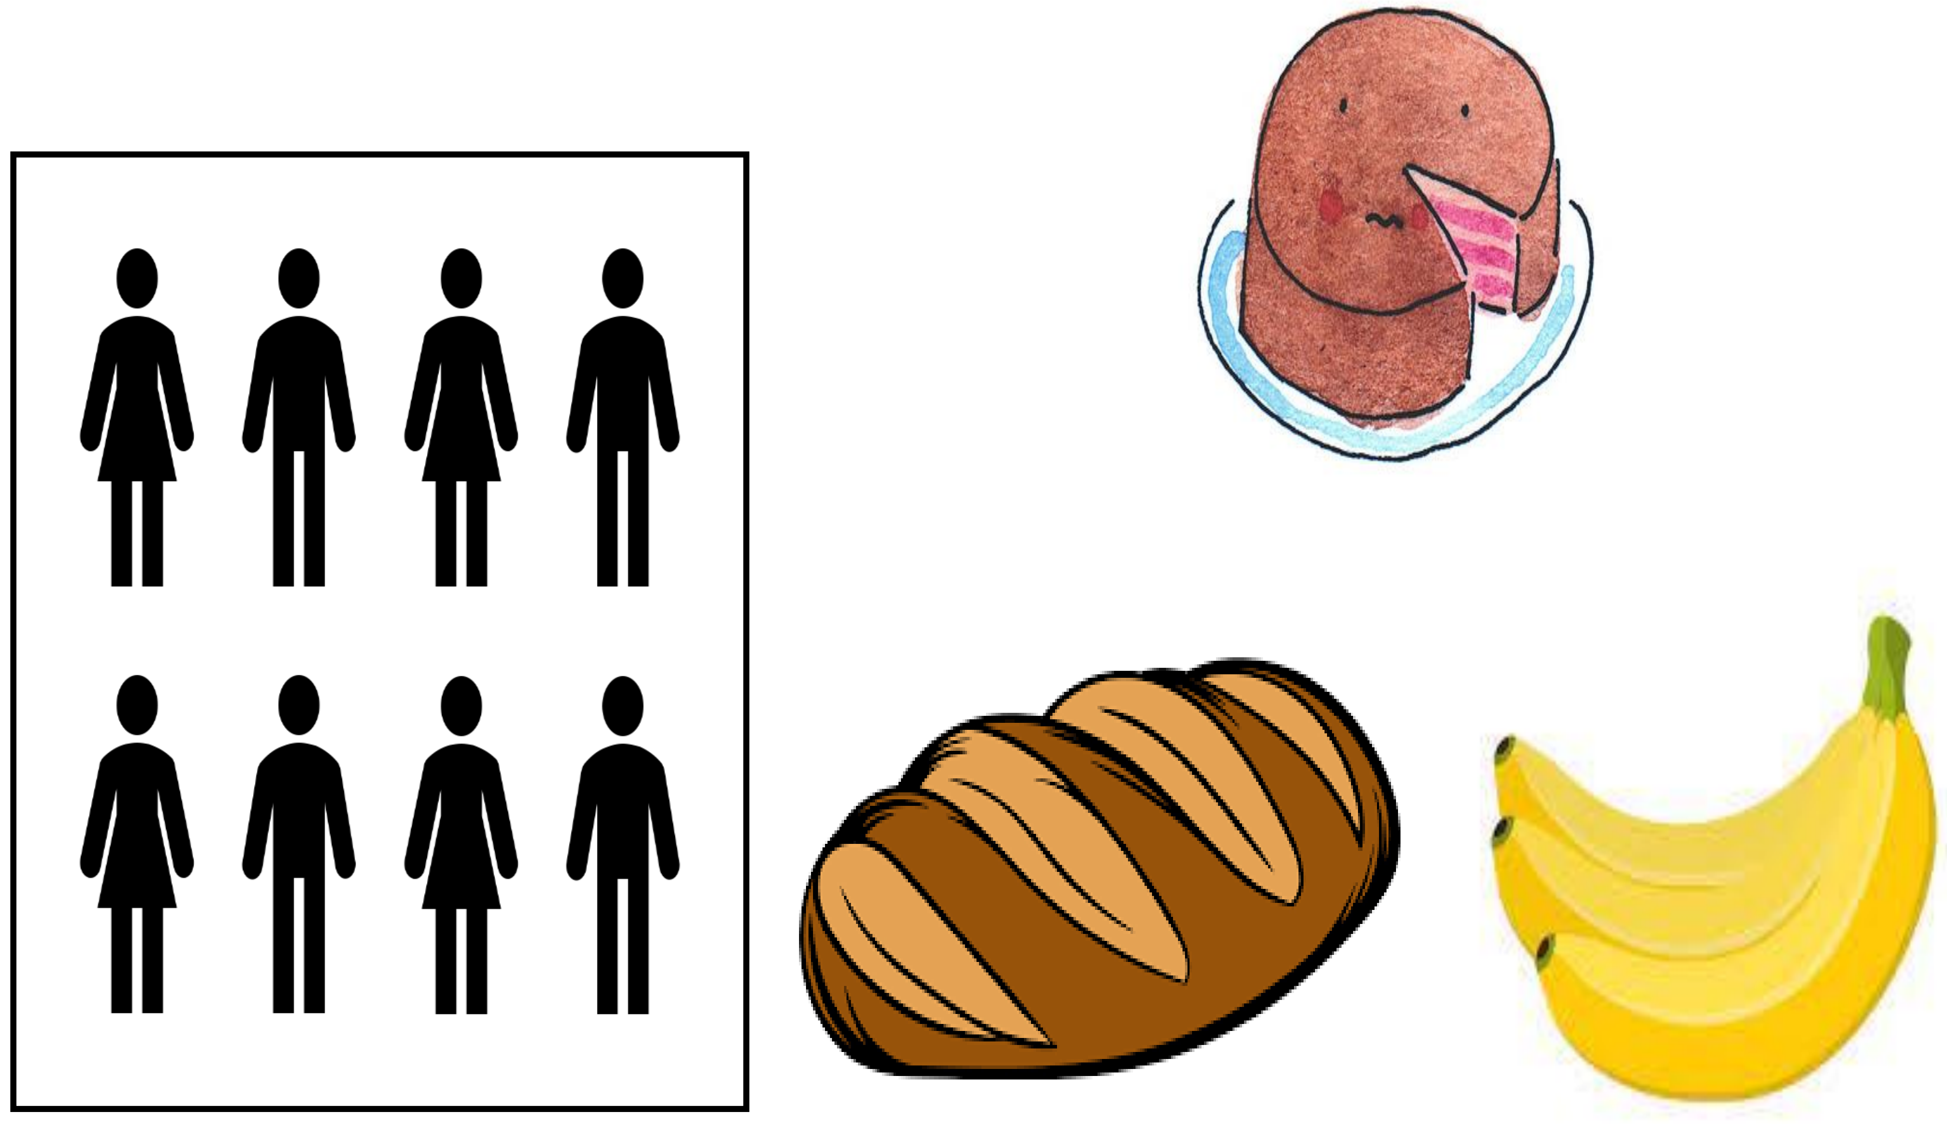

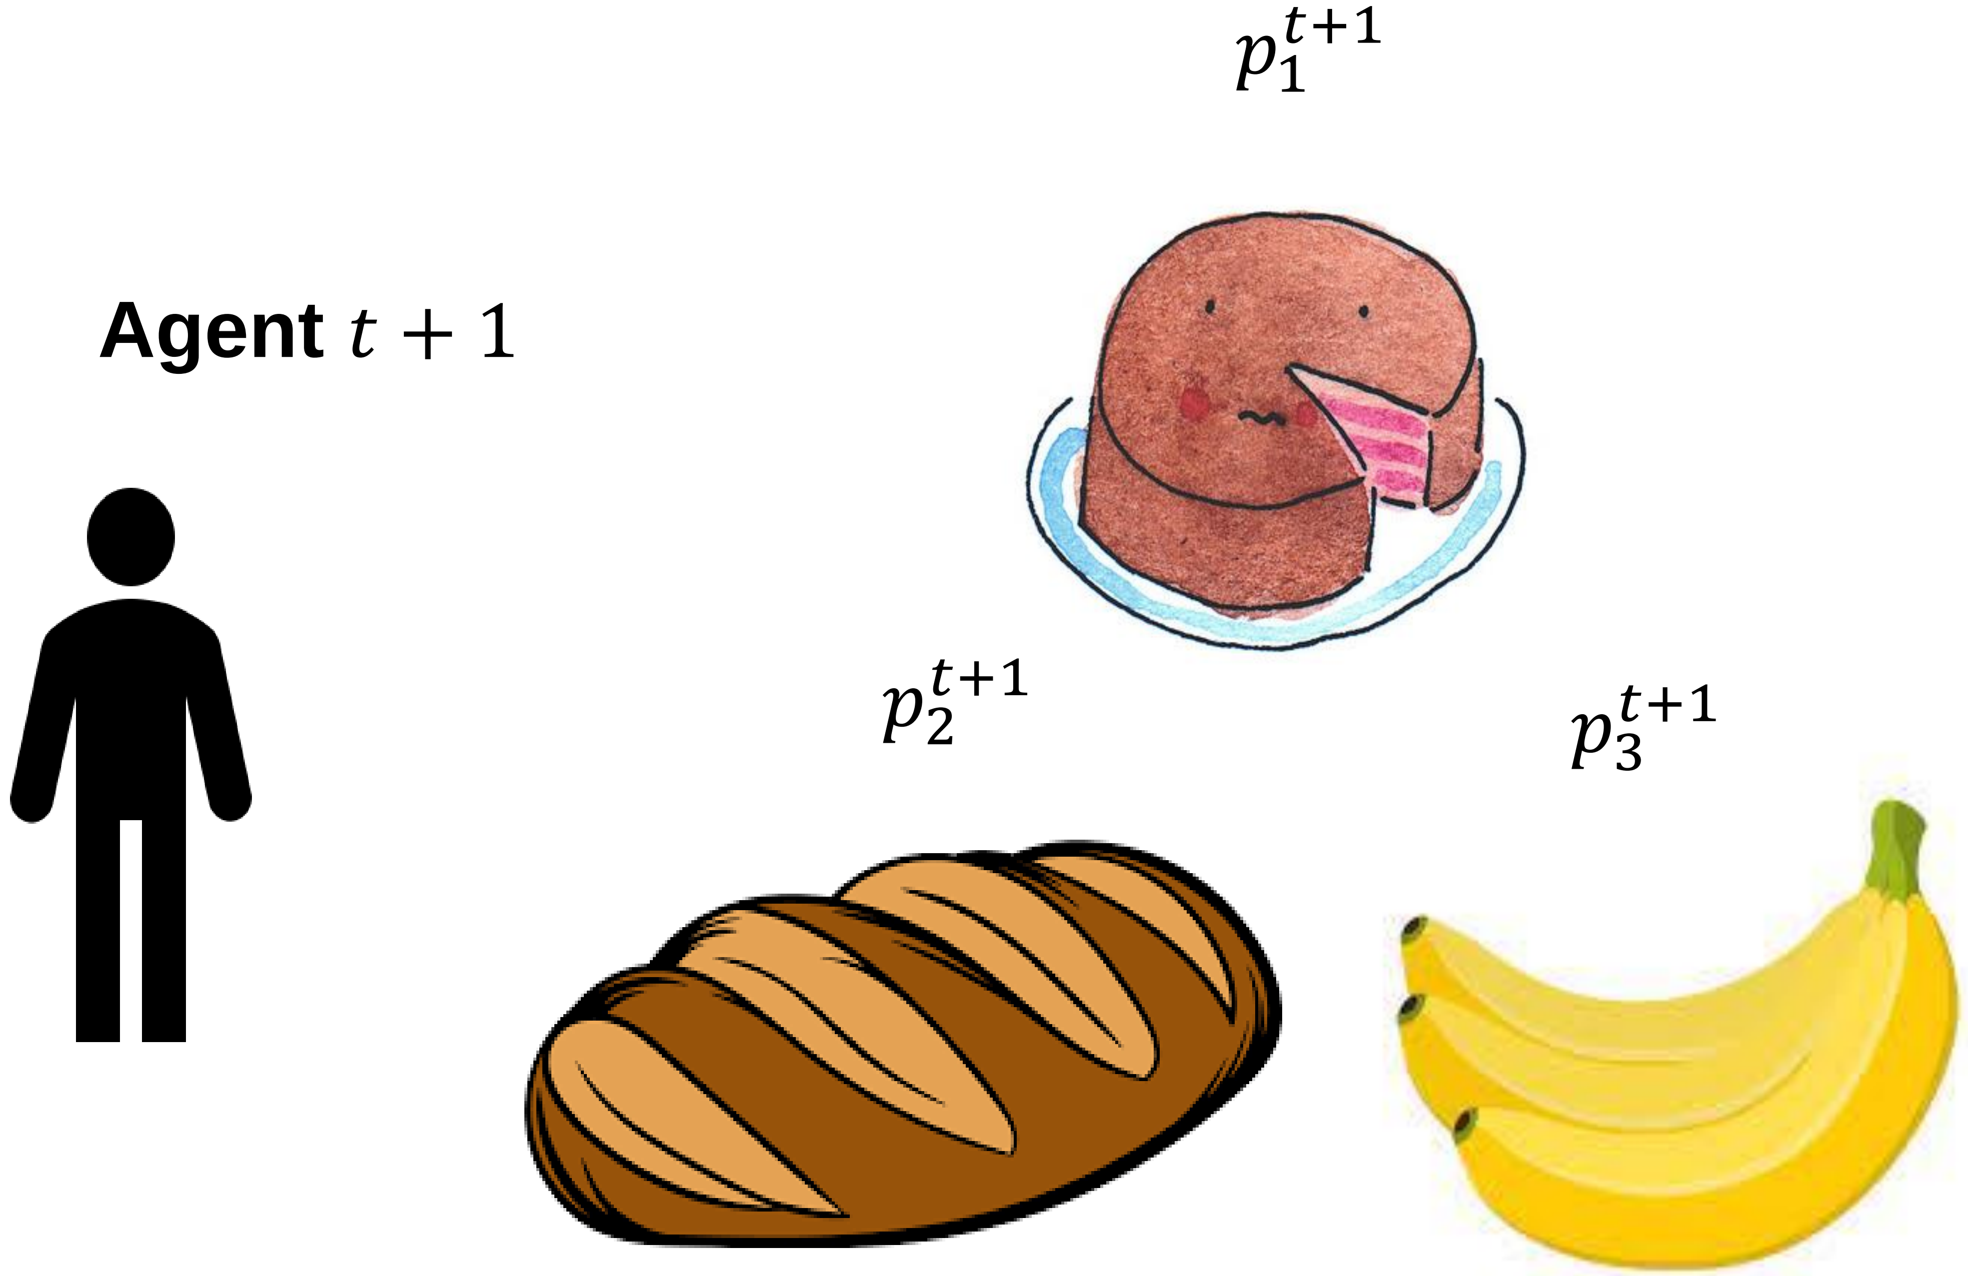

Static Fisher Market Price Equilibrium Theory and Limit

Now Agents Arrive Online

The price at time $t+1$ is reposted based on observed consumption x^t at time t

Jalota and Y

https://arxiv.org/abs/2205.00825

Topic 2. Accelerated Second-Order Methods for Unconstrained Optimization and Applications

$$\min f(x), x \in X \; in \; \mathbb{R}^n,$$

- where f is nonconvex and twice-differentiable,

  $g_k = \nabla f(x_k), H_k = \nabla^2 f(x_k)$

- Goal: find x_k such that:

  $\| g_k \| \leq \epsilon$ (primary, first-order condition)

  $\lambda_{min}(H_k) \geq -\sqrt{\epsilon}$ (secondary, second-order condition)

- First-order methods typically need **$O(n^2\epsilon^{-2})$** arithmetic operations
- Second-order methods typically need **$O(n^3\epsilon^{-1.5})$** arithmetic operations
- New? Yes, HSODM and DRSOM: a single-loop method with **$O(n^2\epsilon^{-1.75})$** operations

  (https://arxiv.org/abs/2211.08212)

App. IV: HSODM for Policy Optimization in Reinforcement Learning

- Consider policy optimization of linearized objective in reinforcement learning

$$\max_{\theta \in \mathbb{R}^d} L(\theta) := L(\pi_\theta),$$

$$\theta_{k+1} = \theta_k + \alpha_k \cdot M_k \nabla \eta(\theta_k),$$

- M_k is usually a preconditioning matrix.

- The Natural Policy Gradient (NPG) method (Kakade, 2001) uses the Fisher information matrix where M_k is the inverse of

$$F_k(\theta) = \mathbb{E}_{\rho_{\theta_k}, \pi_{\theta_k}} \left[\nabla \log \pi_{\theta_k}(s, a) \nabla \log \pi_{\theta_k}(s, a)^T \right]$$

- Based on KL divergence, TRPO (Schulman et al. 2015) uses KL divergence in the constraint:

$$\max_\theta \nabla L_{\theta_k}(\theta_k)^T (\theta - \theta_k)$$

$$\text{s.t. } \mathbb{E}_{s \sim \rho_{\theta_k}} [D_{KL}(\pi_{\theta_k}(\cdot \mid s); \pi_\theta(\cdot \mid s))] \leq \delta.$$

Homogeneous NPG:
Apply HSODM!

Preliminary Results: HSODM for Policy Optimization in RL

- A comparison of Homogeneous NPG and Trust-region Policy Optimization (Schultz, 2015)

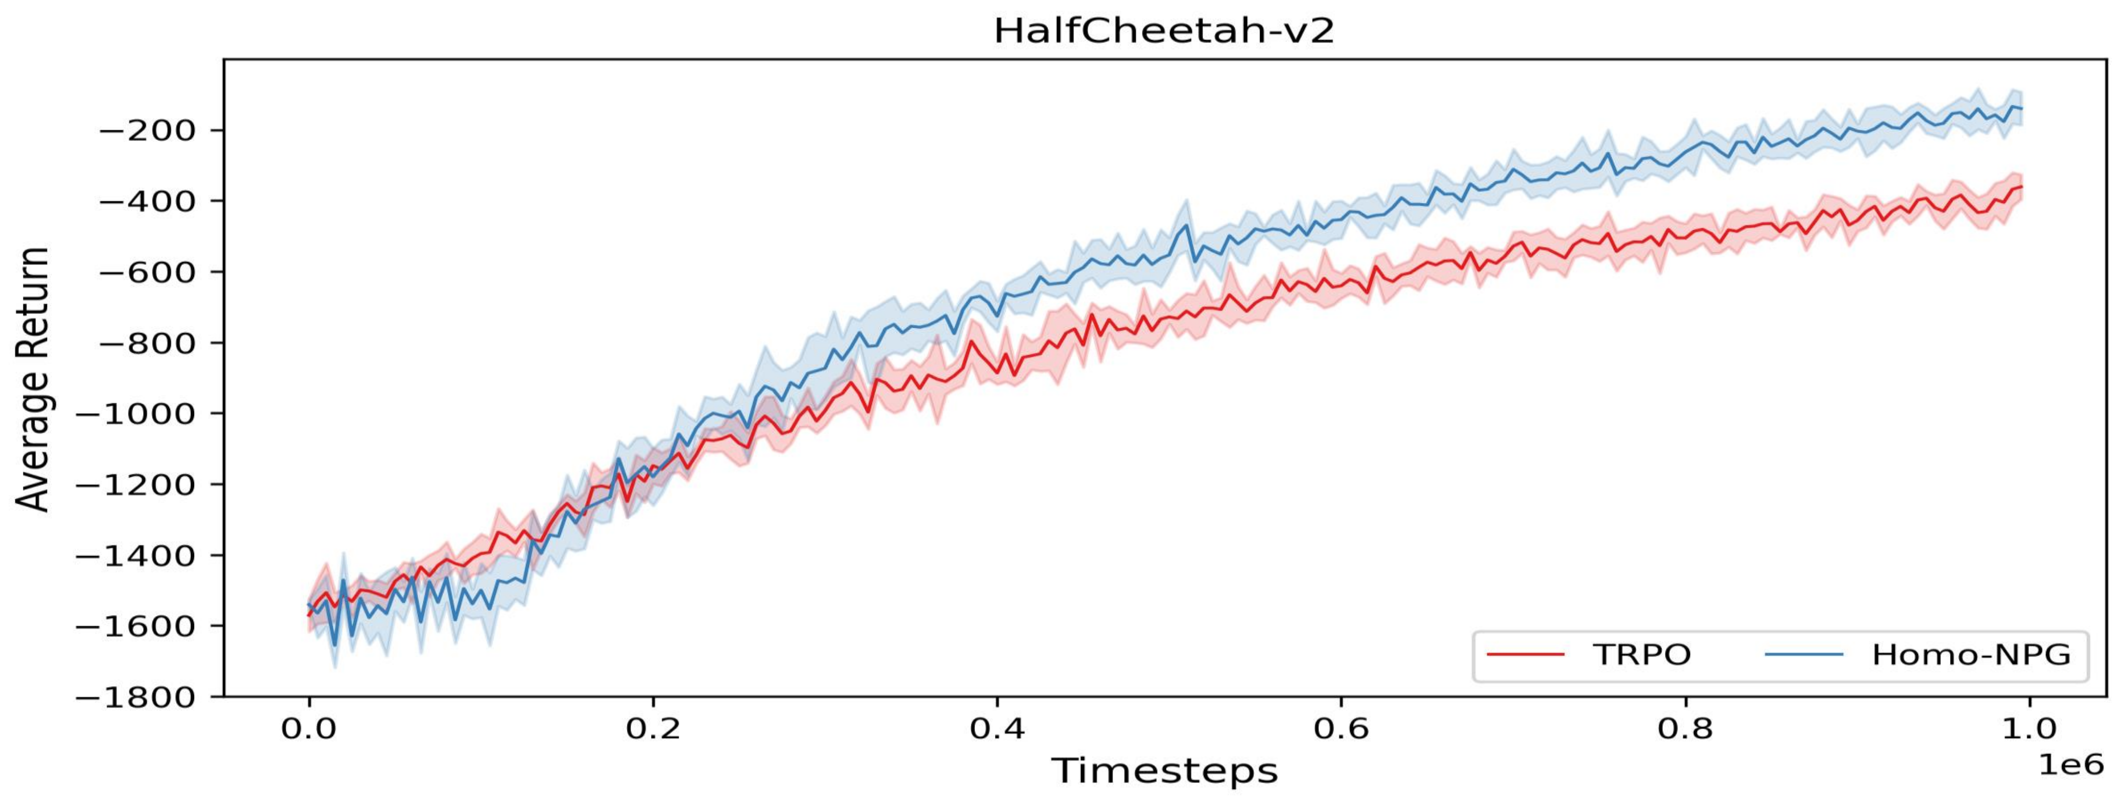

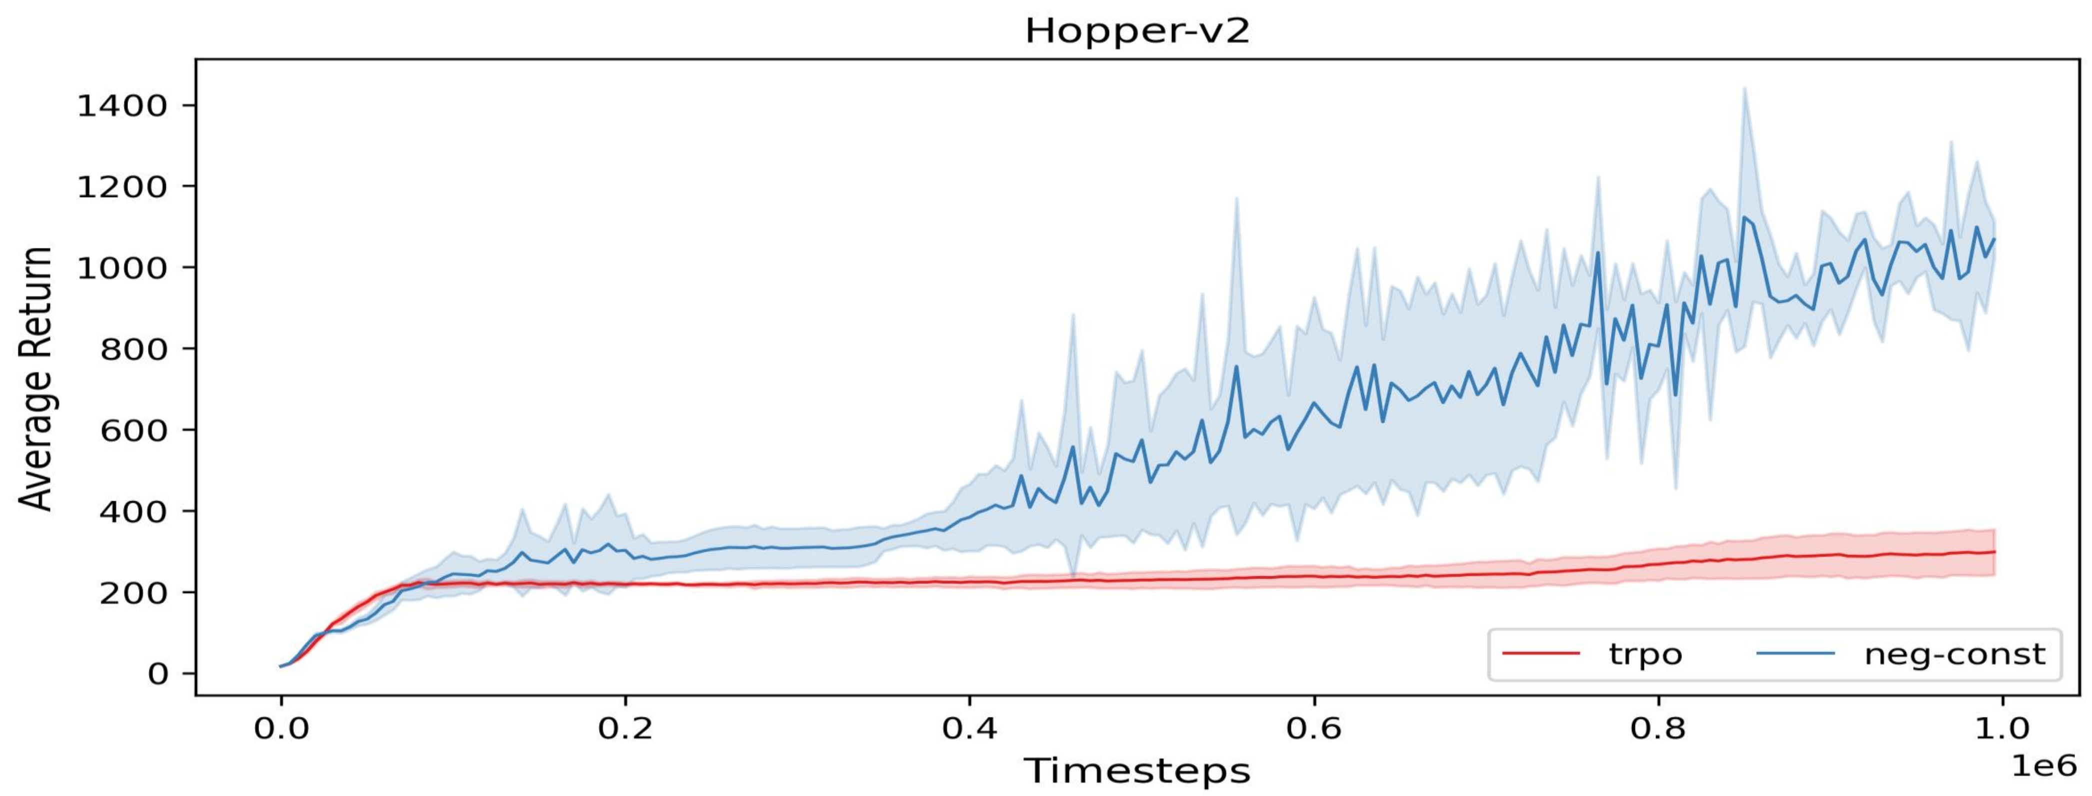

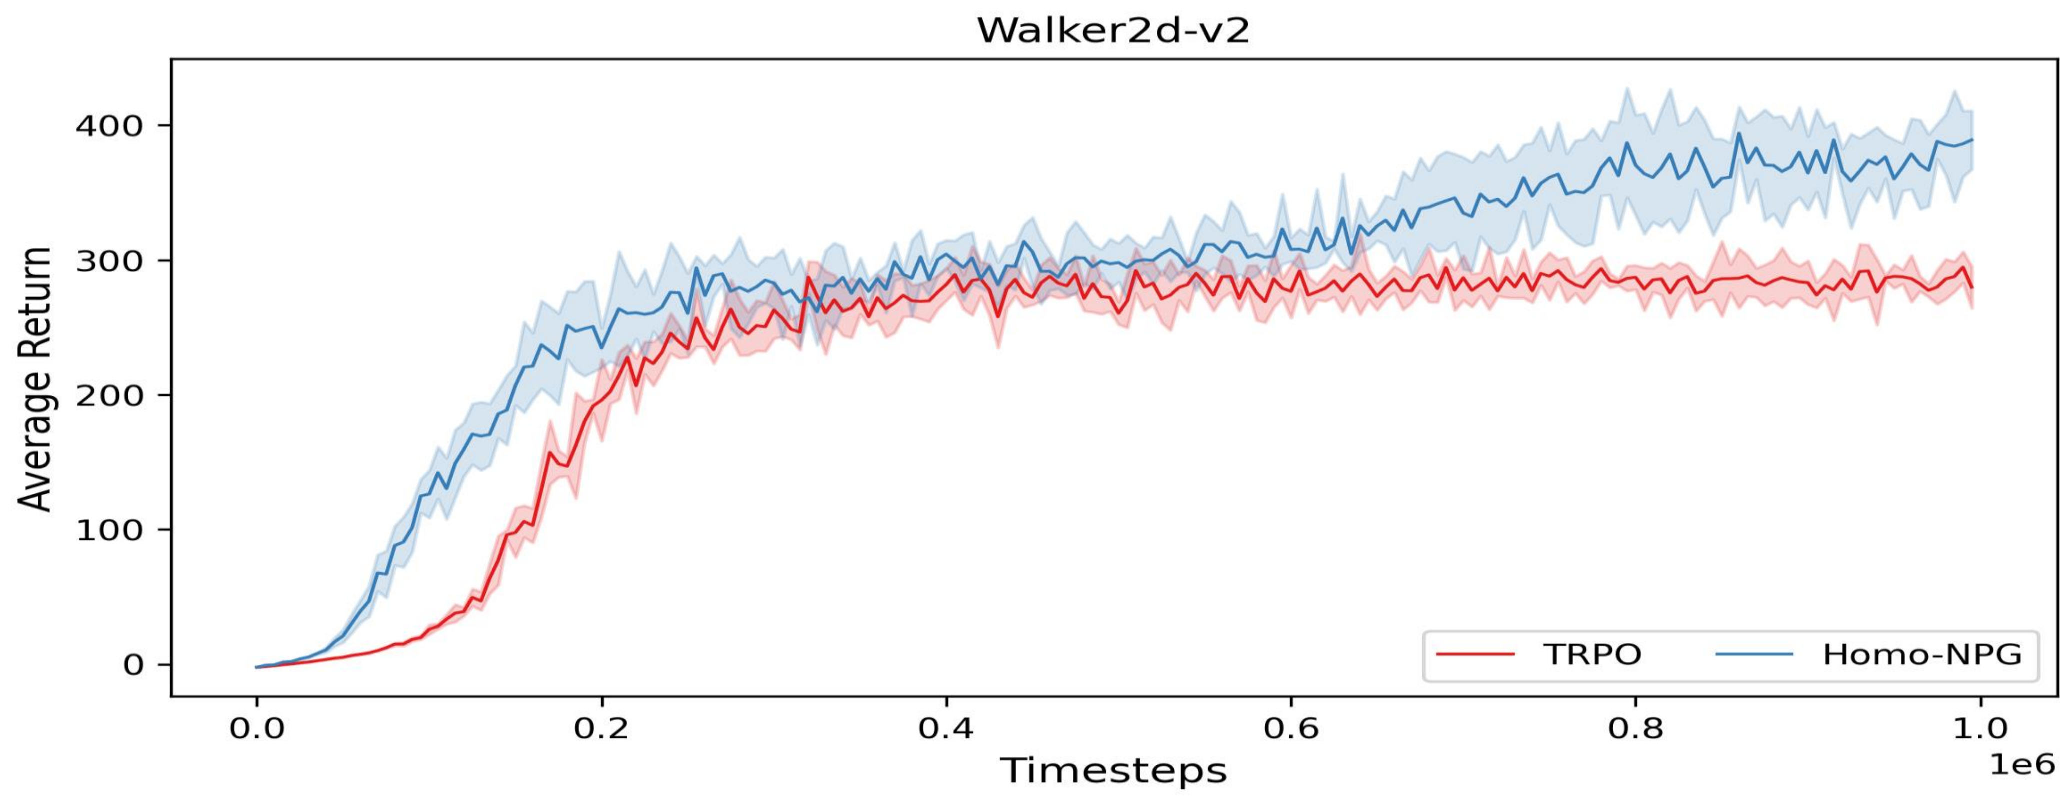

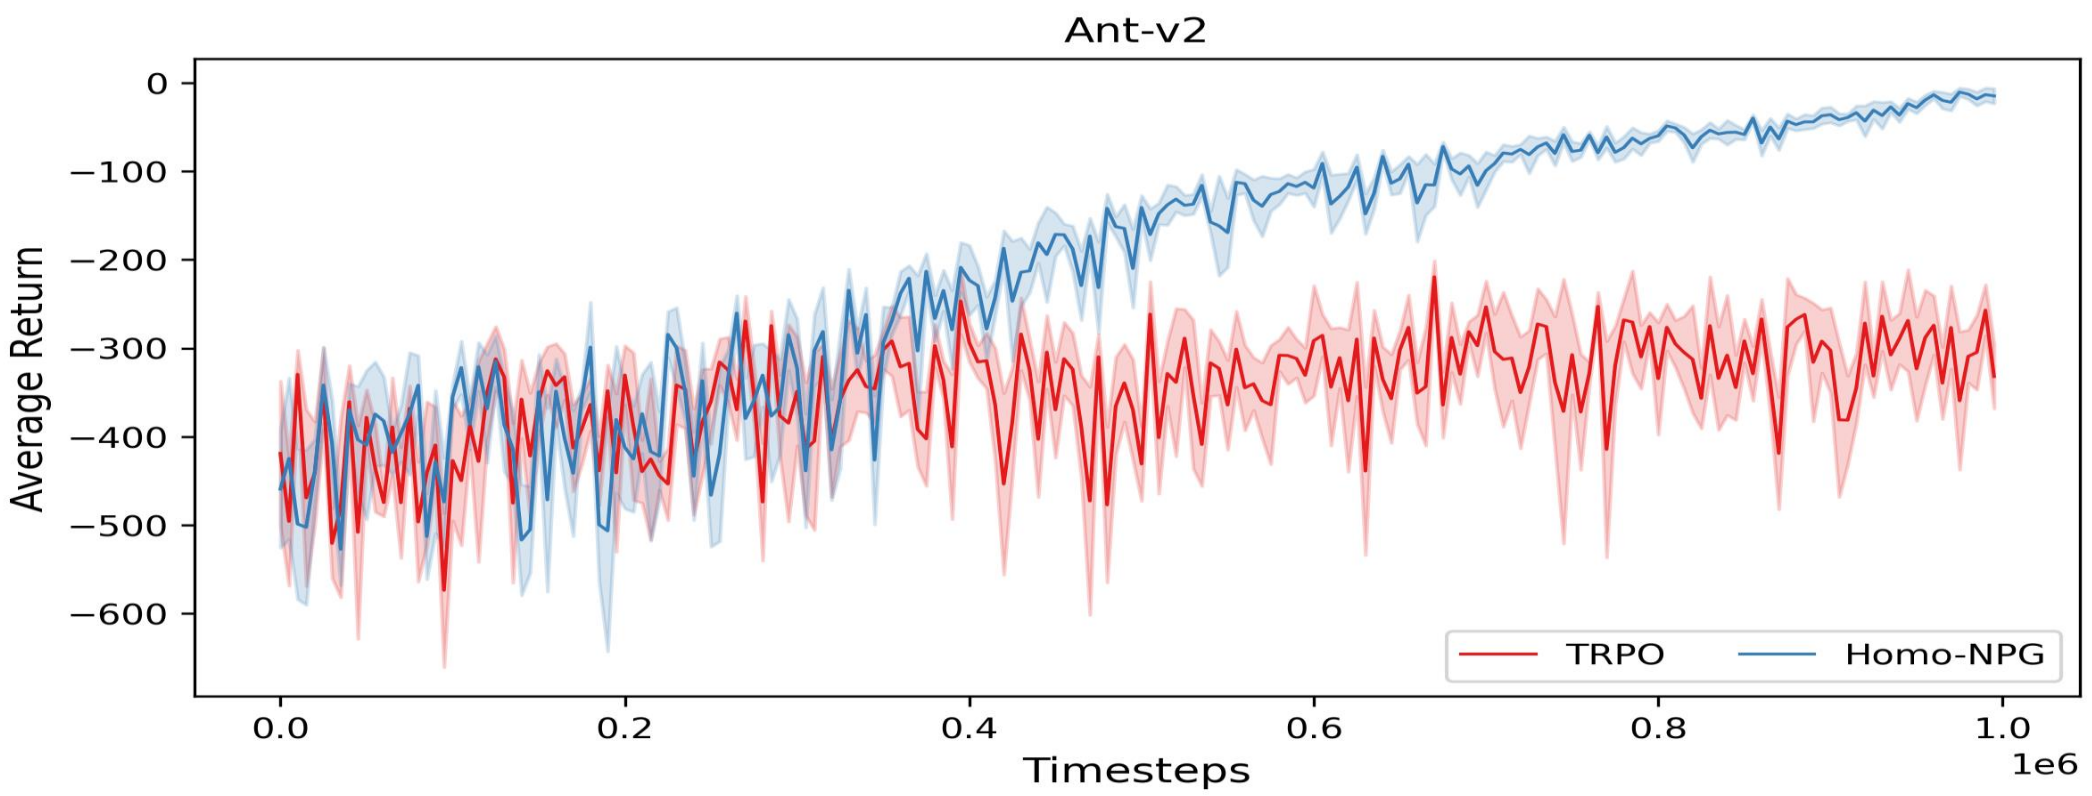

- HSODM provides significant improvements over TRPO
- Ongoing: second-order information of L?
- **Further reduce the computation cost per step**

Dimension Reduced Second-Order Method (DRSOM)

- Motivation from Multi-Directional FOM and Subspace Method, such as CG and ADAM, DRSOM applies the trust-region method in low dimensional subspace.
- This results in a low-dimensional quadratic sub-minimization problem:
- Typically, DRSOM adopts two directions $d = -\alpha^1 \nabla f(x_k) + \alpha^2 d_k$

  where $g_k = \nabla f(x_k), H_k = \nabla^2 f\left(x^k\right), d_k = x_k - x_{k-1}$
- Then we solve a 2-d quadratic minimization problem:

$$\min\ m_k^{\alpha}(\alpha) \coloneqq f(x_k) + (c_k)^T \alpha + \frac{1}{2} \alpha^T Q_k \alpha$$

$$||\alpha||_{G_k} \leq \Delta_k$$

$$G_k = \begin{bmatrix} g_k^T g_k & -g_k^T d_k \\ -g_k^T d_k & d_k^T d_k \end{bmatrix}, Q_k = \begin{bmatrix} g_k^T H_k g_k & -g_k^T H_k d_k \\ -g_k^T H_k d_k & d_k^T H_k d_k \end{bmatrix}, c_k = \begin{bmatrix} -||g_k||^2 \\ g_k^T d_k \end{bmatrix}$$

App. V: Sensor Network Location (SNL)

- Localization
  - –Given partial pair-wise measured distance values
  - –Given some anchors' positions
  - –Find locations of all other sensors that fit the measured distance values

This is also called graph realization on a fixed dimension Euclidean space

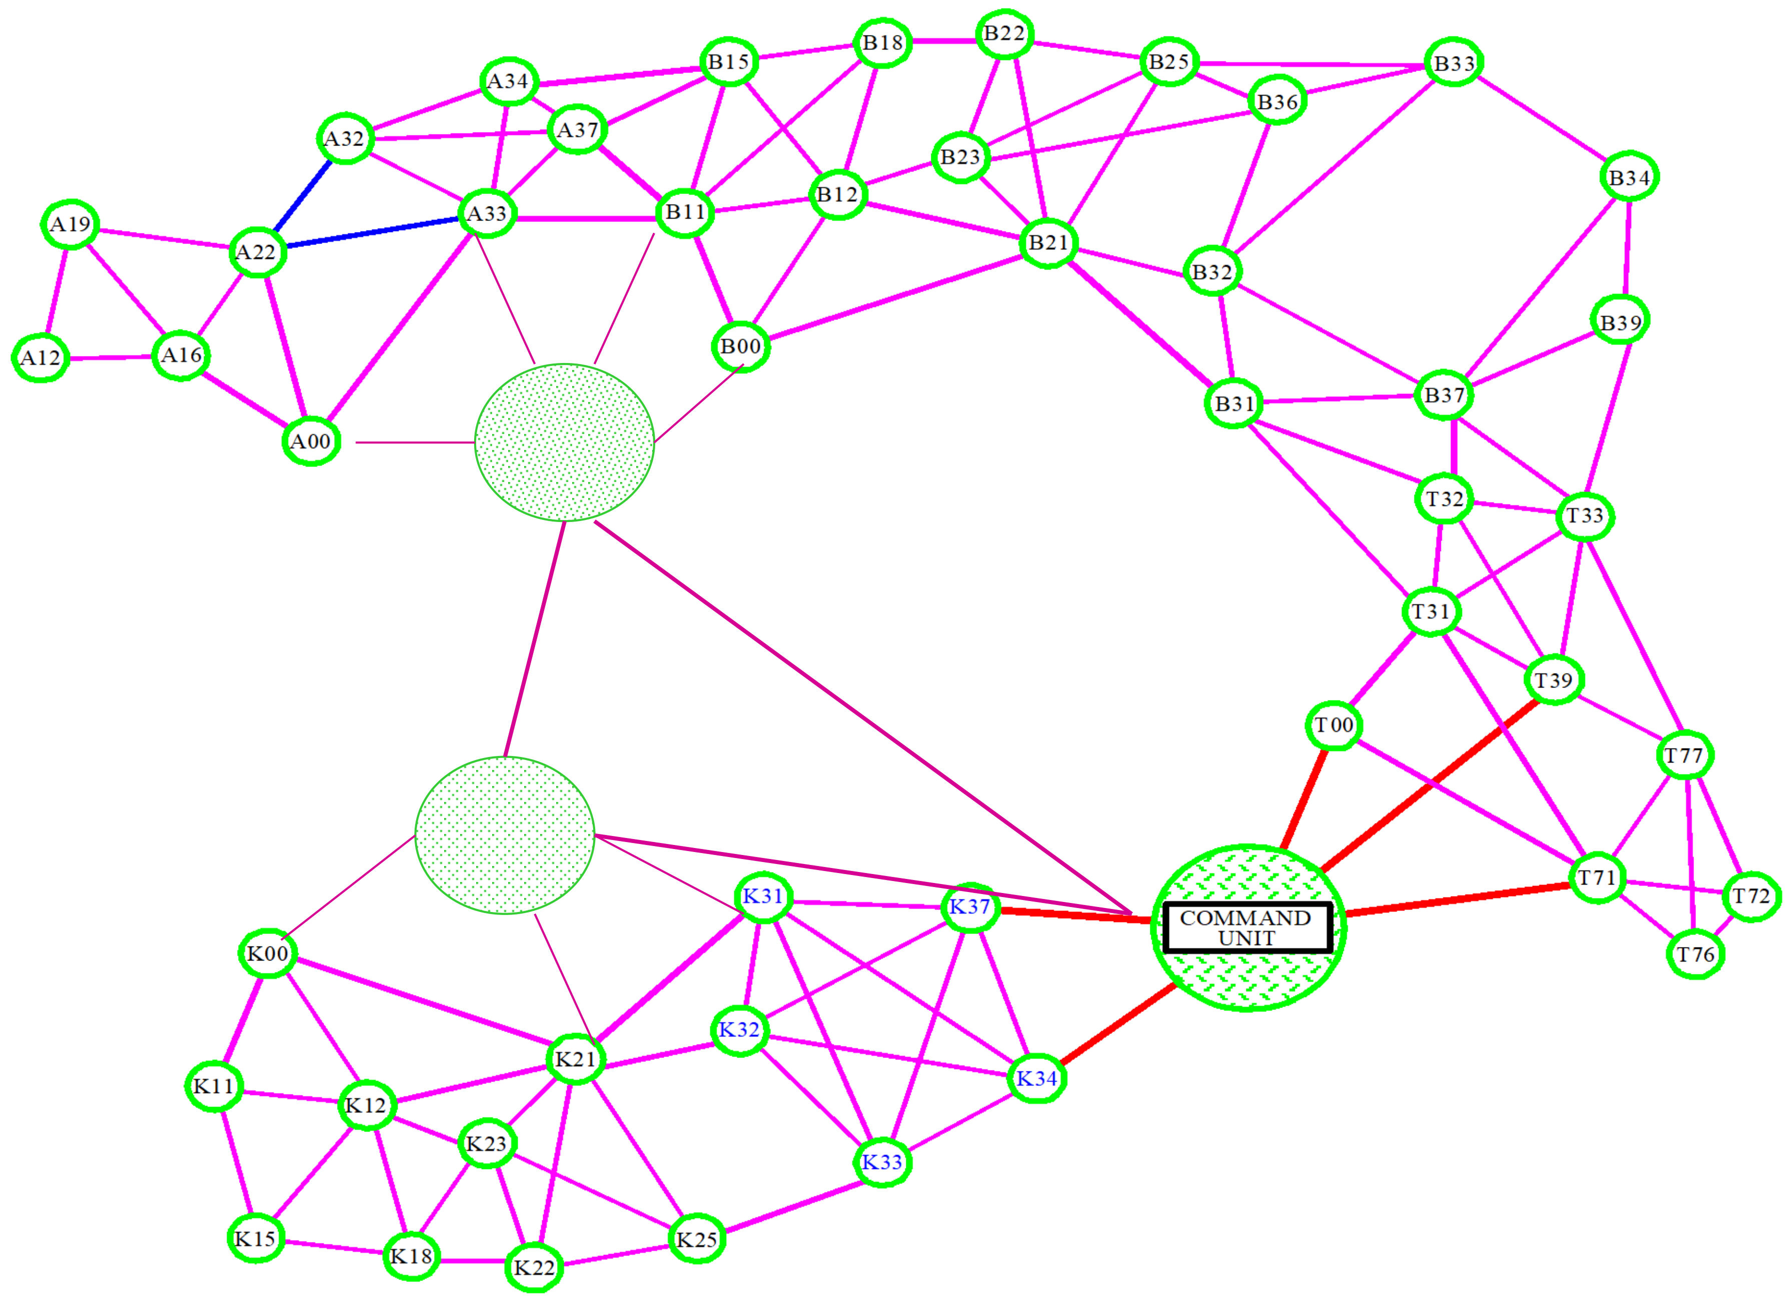

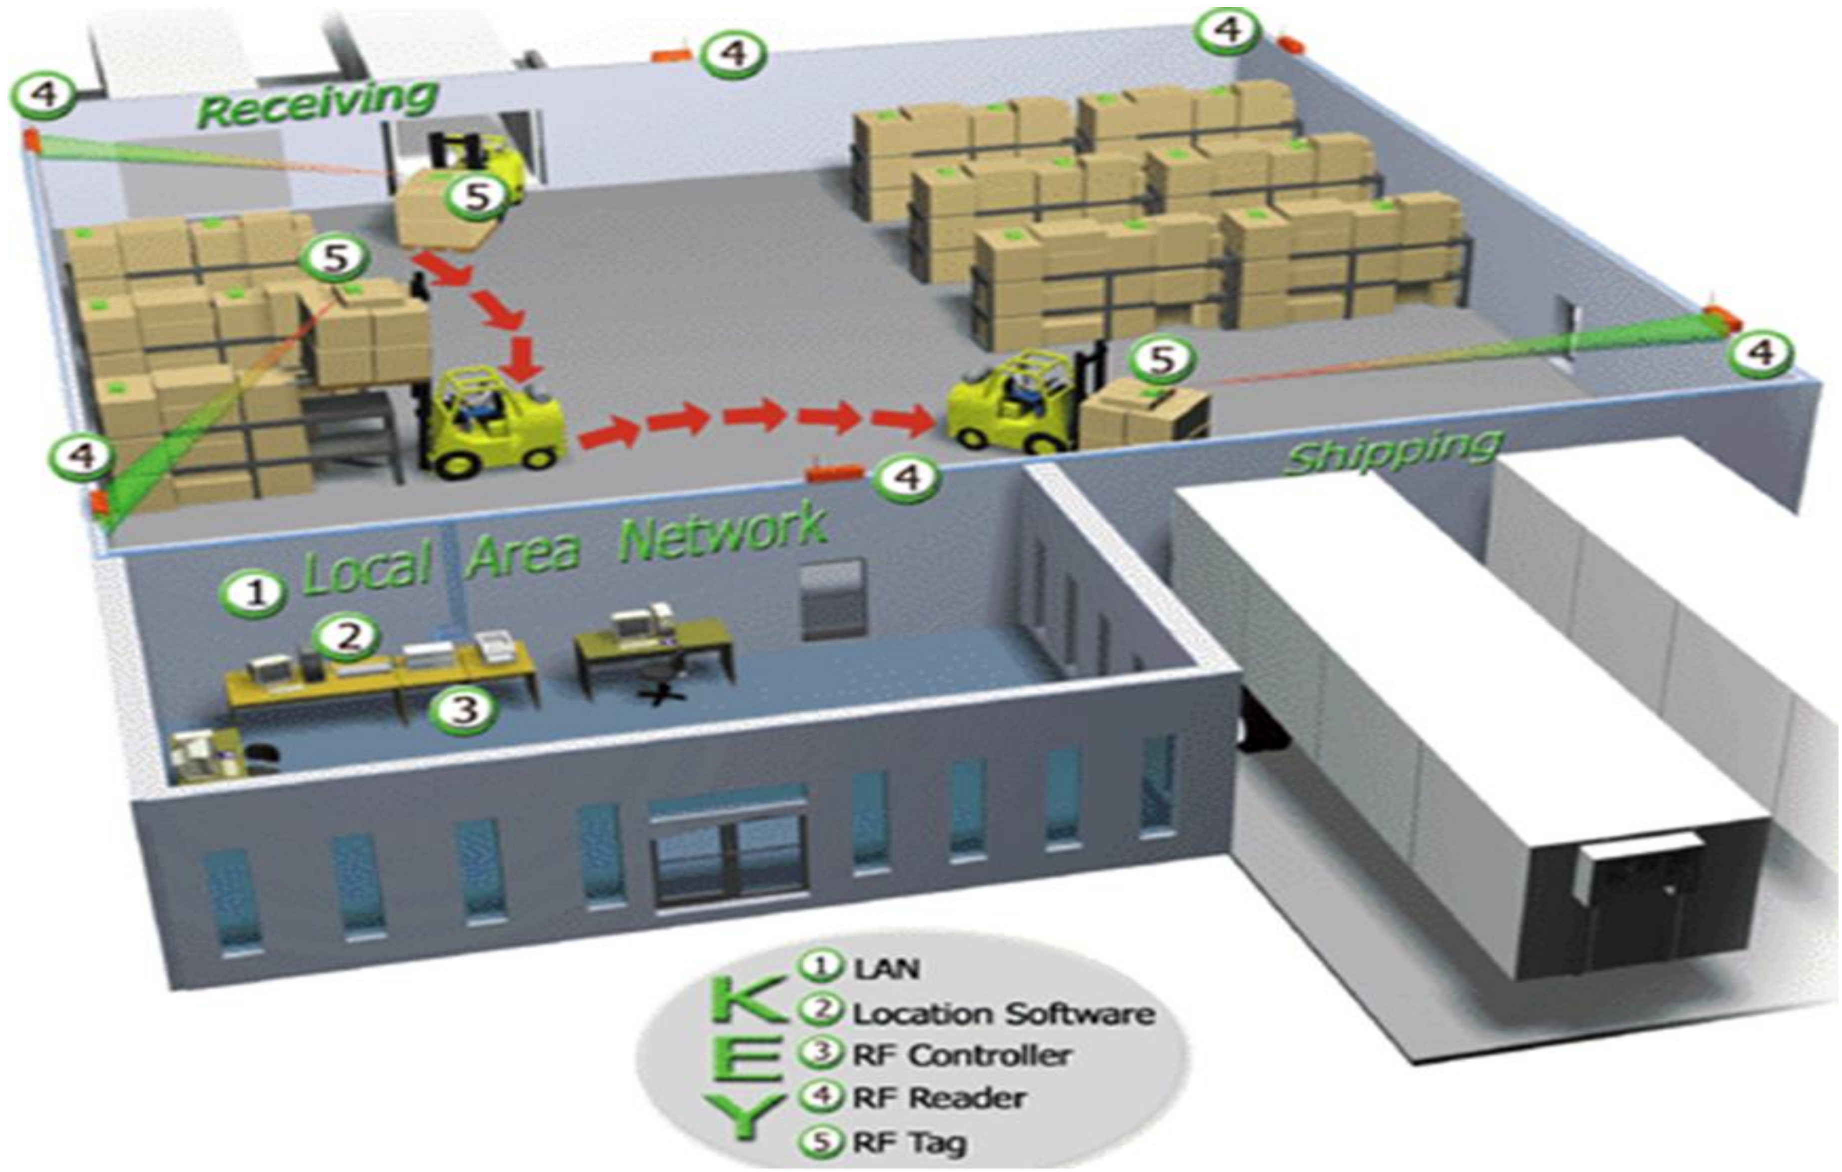

Mathematical Formulation of Sensor Network Location (SNL)

- Consider Sensor Network Location (SNL)

$$N_x = \{(i,j) : \|x_i - x_j\| = d_{ij} \le r_d\},\ N_a = \{(i,k) : \|x_i - a_k\| = d_{ik} \le r_d\}$$

where r_d is a fixed parameter known as the radio range. The SNL problem considers the following QCQP feasibility problem,

$$\|x_i - x_j\|^2 = d_{ij}^2, \forall (i,j) \in N_x$$
$$\|x_i - a_k\|^2 = \bar{d}_{ik}^2, \forall (i,k) \in N_a$$

- Alternatively, one can solve SNL by the nonconvex nonlinear least square (NLS) problem

$$\min_X \sum_{(i<j,j)\in N_x} (\|x_i - x_j\|^2 - d_{ij}^2)^2 + \sum_{(k,j)\in N_a} (\|a_k - x_j\|^2 - \bar{d}_{kj}^2)^2.$$

Sensor Network Location (SNL)

- DRSOM vs the First-Order-Gradient-Descent Method

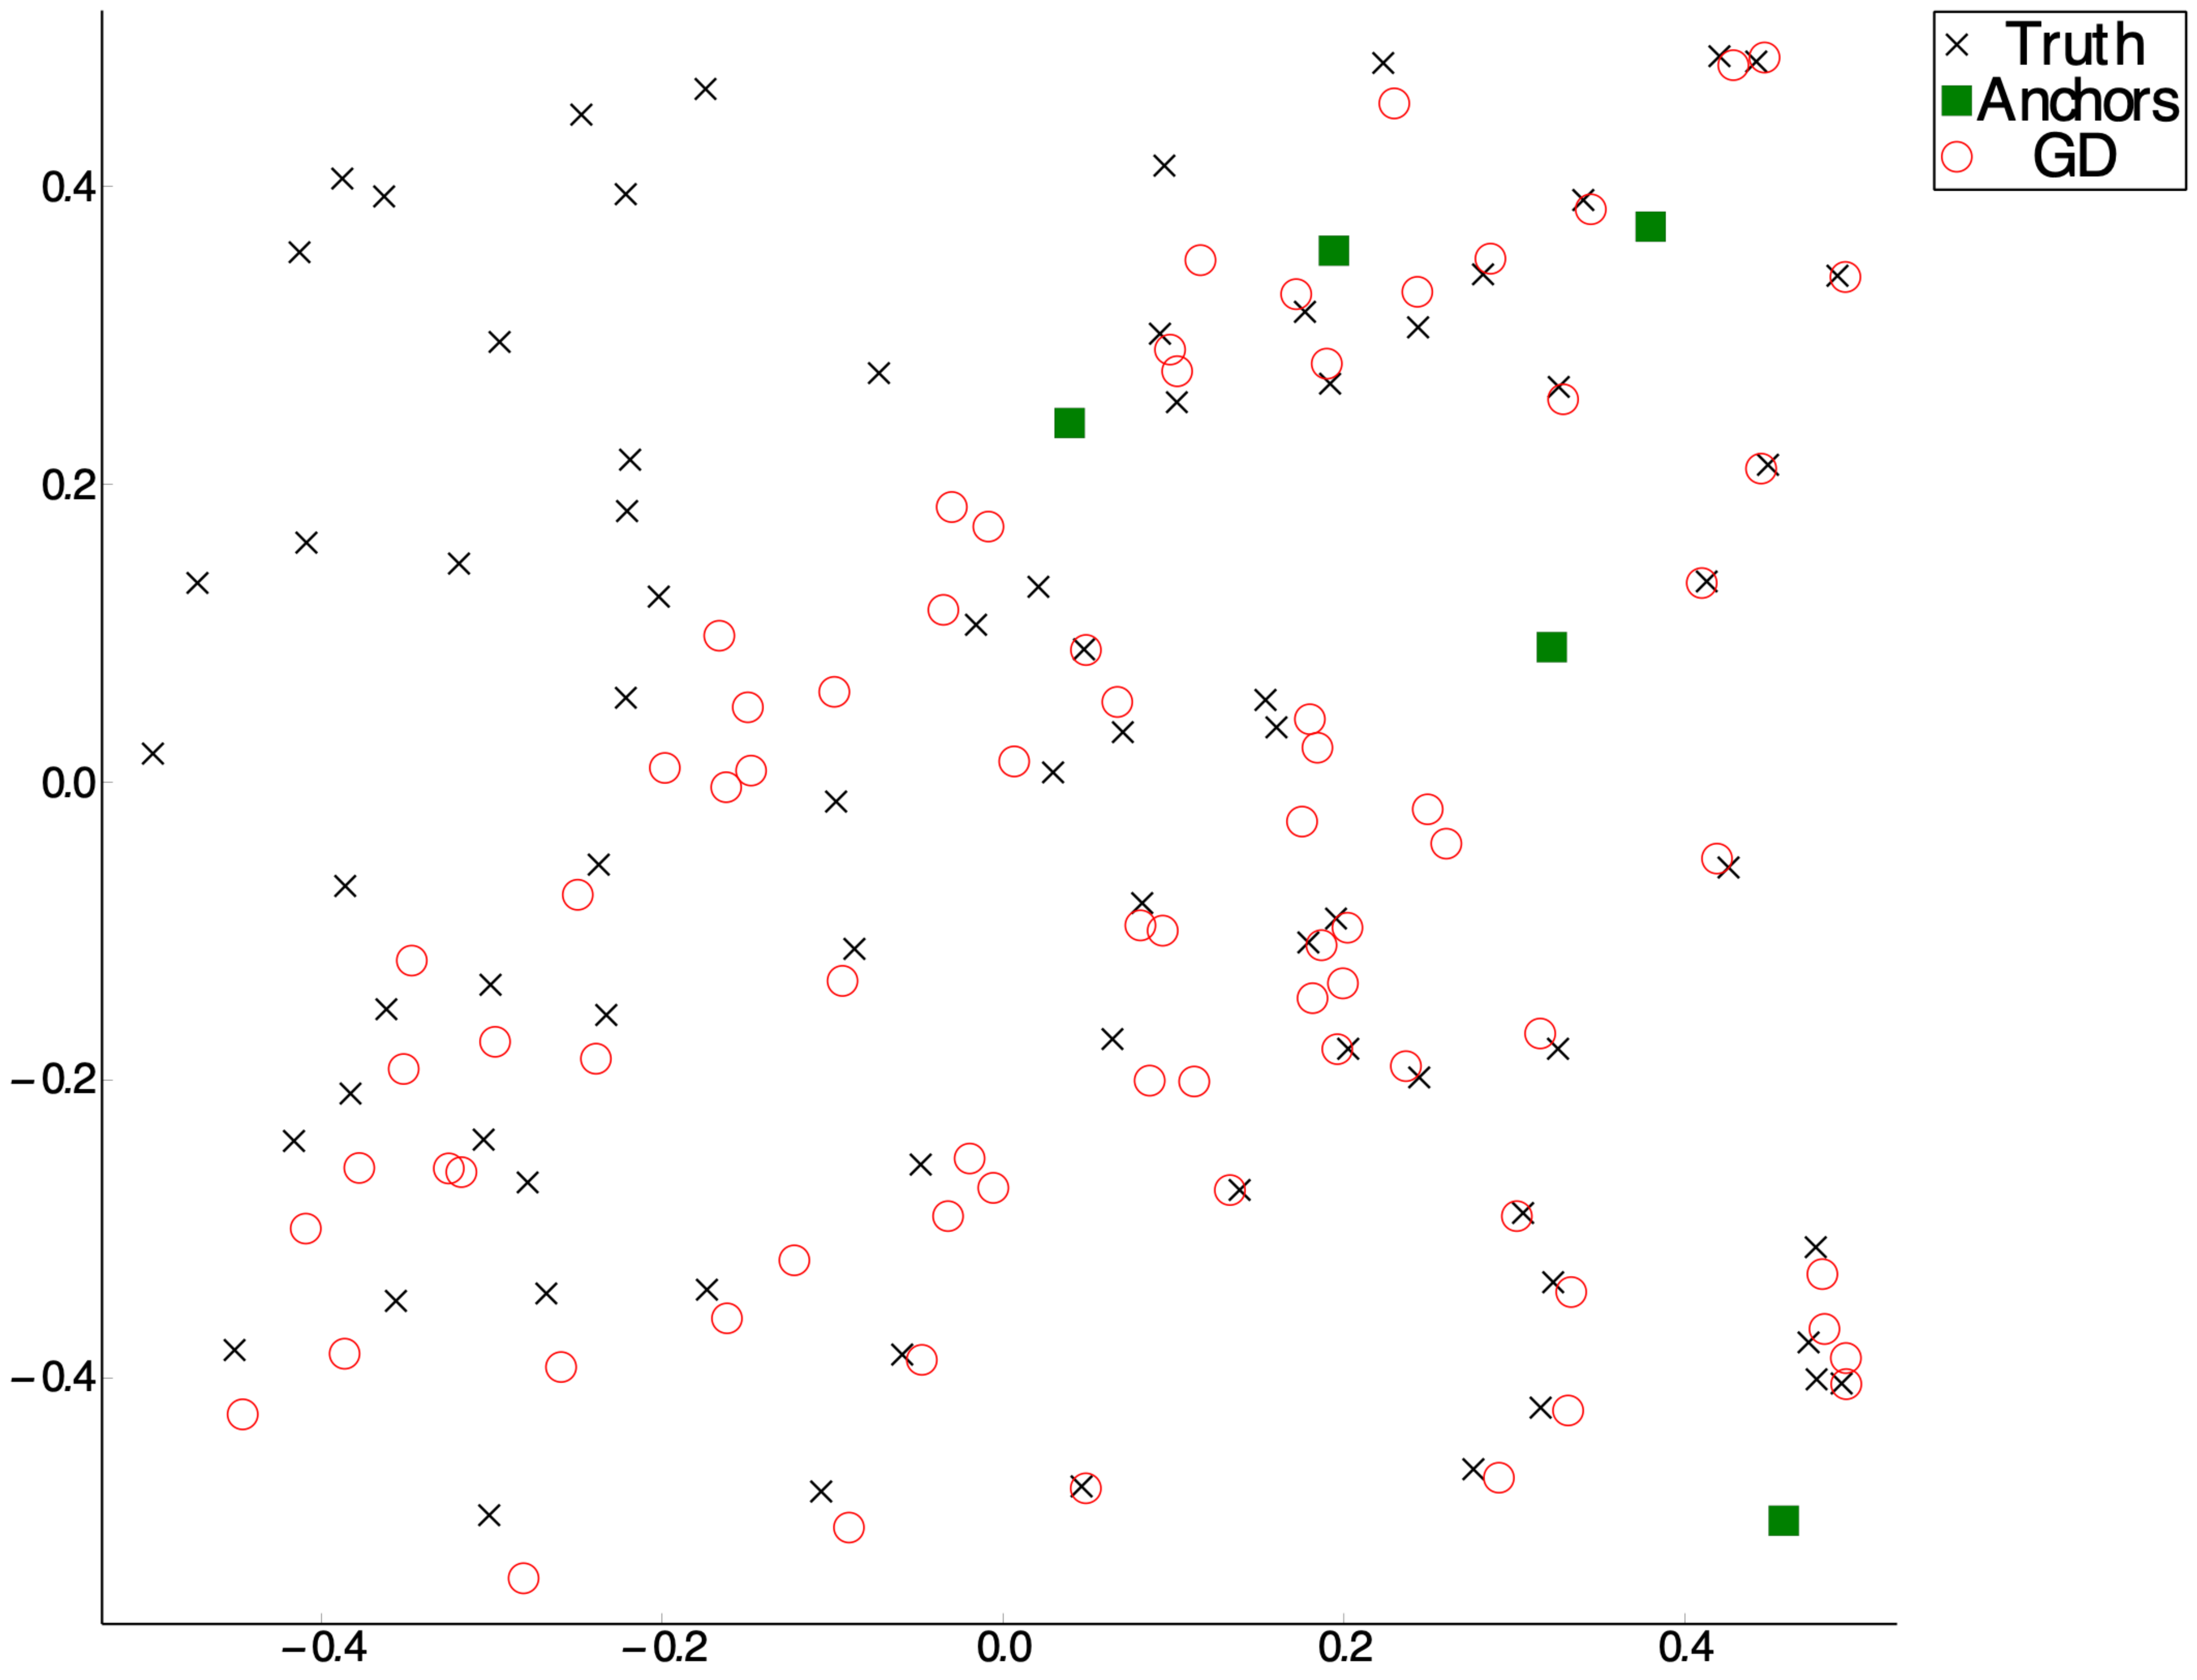

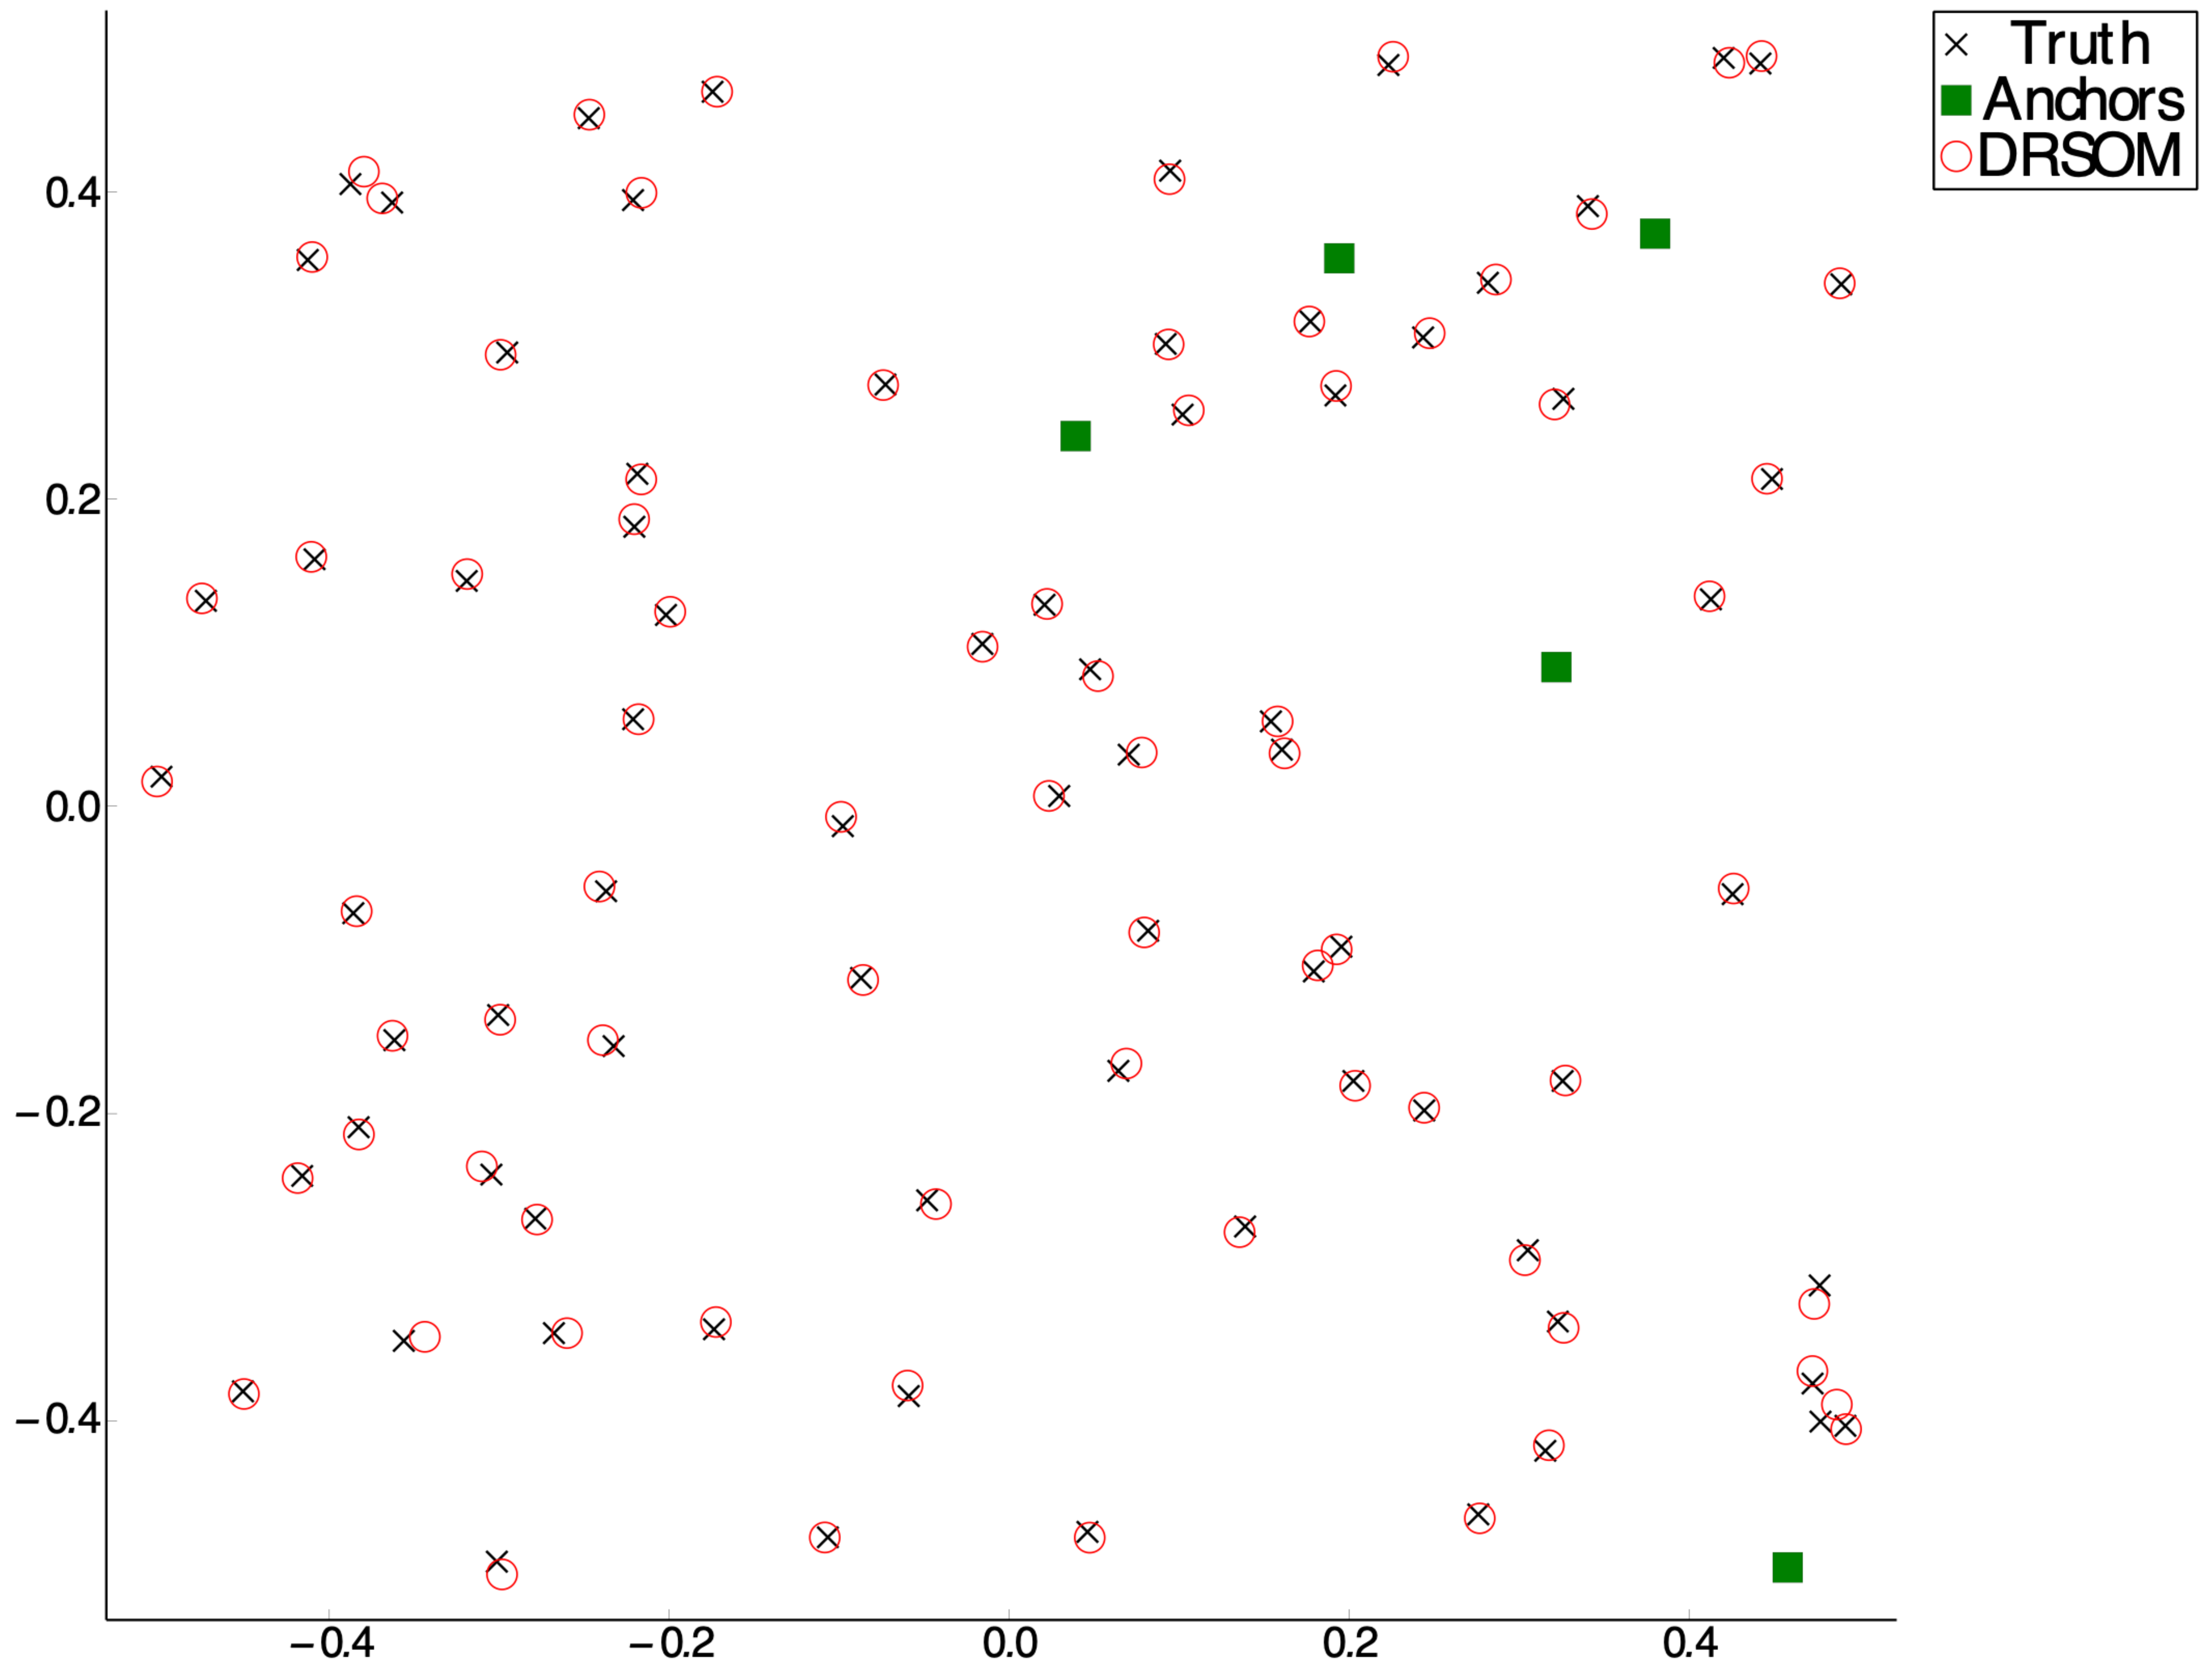

Sensor Network Location, Large-Scale Instances

- Graphical results with 10,000 nodes and 1000 anchors (no noise) **within 3,000 seconds**

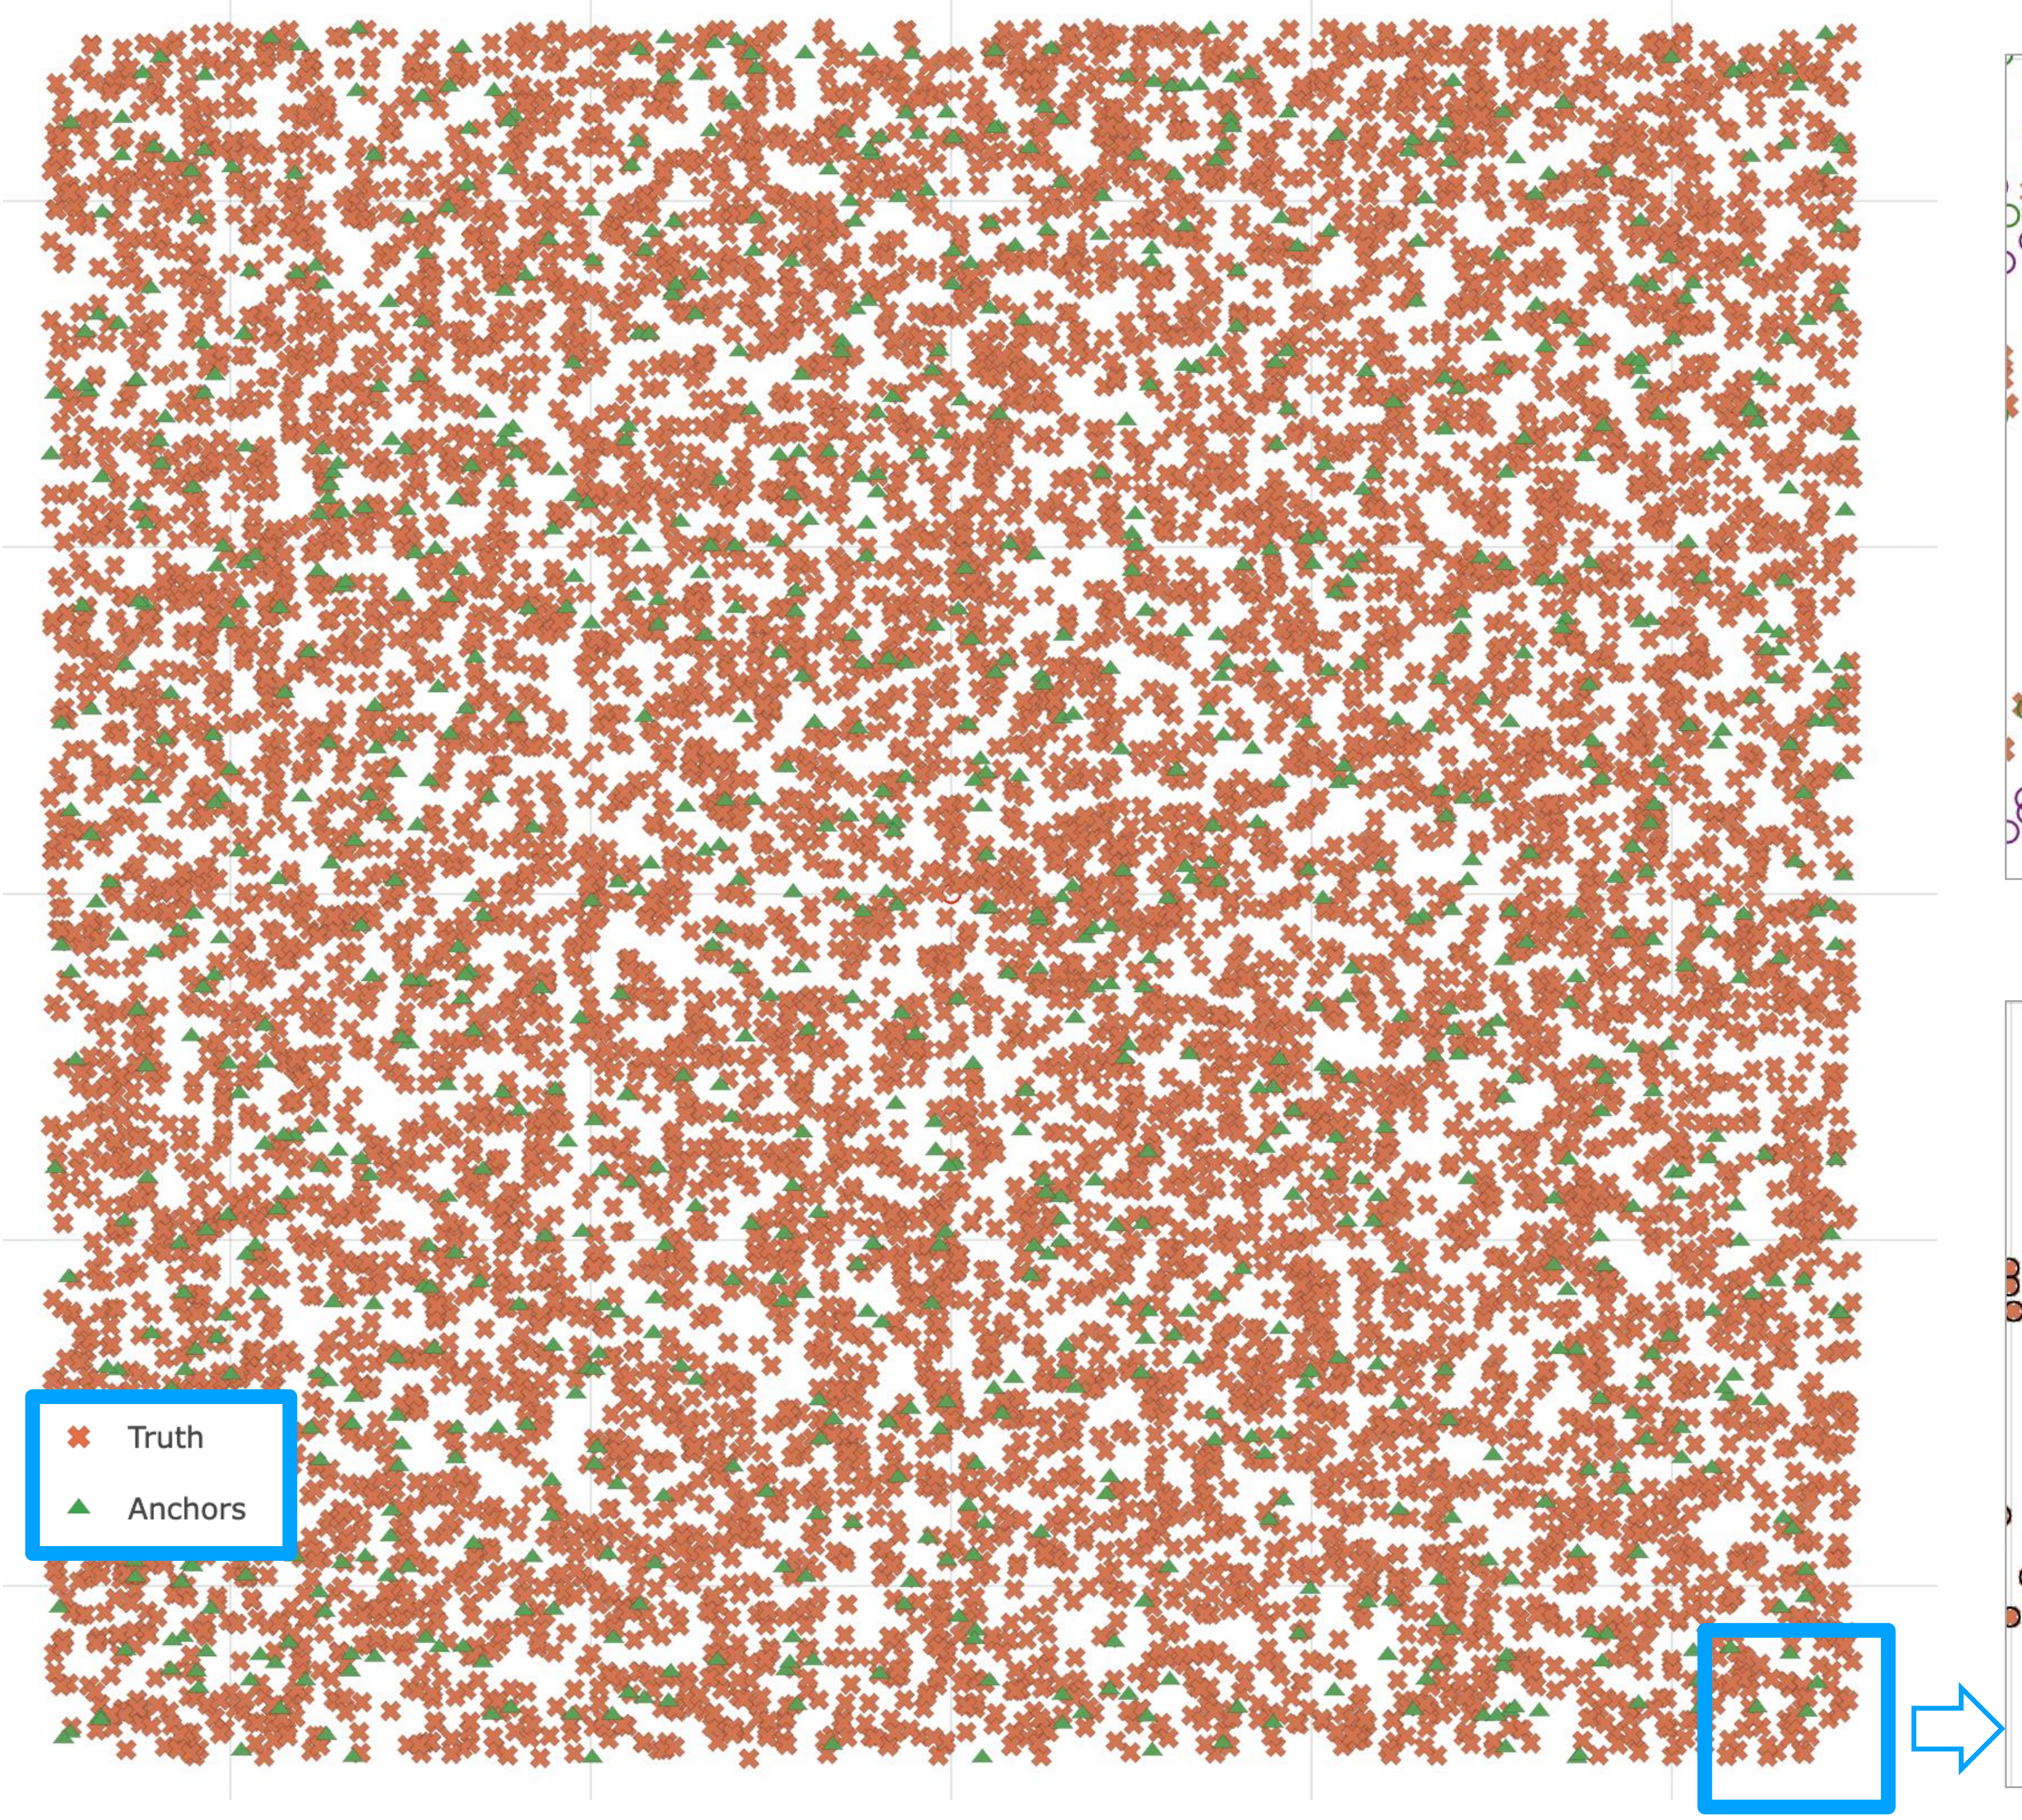

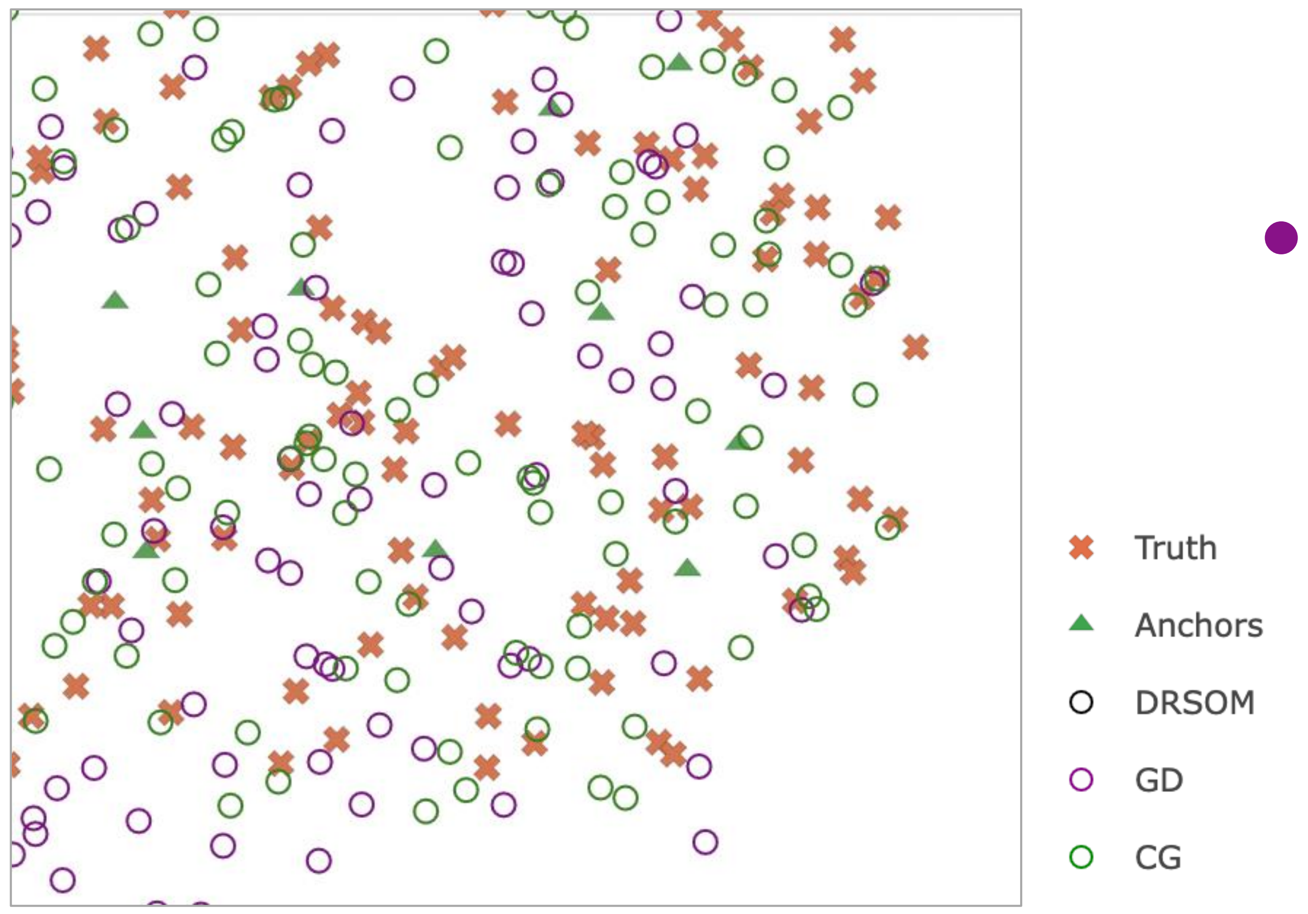

- **GD with Line-search and Hager-Zhang CG both timeout**

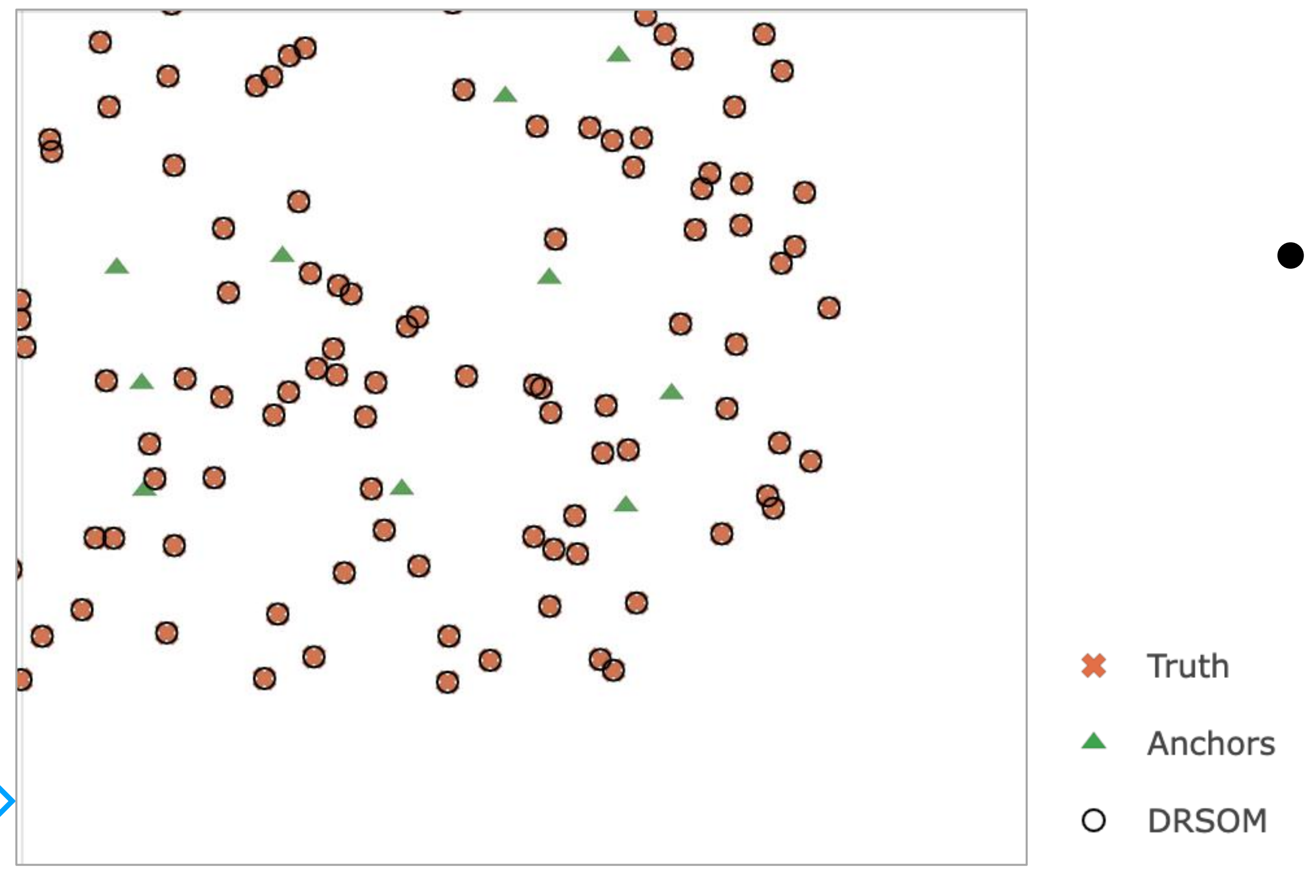

- **DRSOM can converge to $|g_k| \le 1e^{-5}$ in 2,200s**

Sensor Network Online Tracking, 2D and 3D

Theory: Dimension Expansion via Semidefinite Programming

$$\min_X \sum_{(i<j,j)\in N_x} (\|x_i - x_j\|^2 - d_{ij}^2)^2 + \sum_{(k,j)\in N_a} (\|a_k - x_j\|^2 - \bar{d}_{kj}^2)^2.$$

Local minima always exist regardless how much edge information

<u>Step 1</u>: Linearization and dimension expansion

$$\|x_i - x_j\|^2 = x_i^T x_i - 2x_i^T x_j + x_j^T x_j$$

$$Y_{ii} \qquad Y_{ij} \qquad Y_{jj}$$

$$\|a_k - x_j\|^2 = a_k^T a_k - 2a_k^T x_j + x_j^T x_j$$

$$Y_{jj}$$

Tighten: $Y = X^T X$, $X=[x_1,\ldots,x_n]$

<u>Step 2:</u> Relax and Rank Constraint

$$Y \geq X^T X \Leftrightarrow Z = \begin{bmatrix} I & X \\ X^T & Y \end{bmatrix} \geq PSD$$ **with rank d**

This is a conic linear program with rank constraint.

Theorem: its Convex Relaxation has a unique solution when edge information is sufficient: "Bless of Dimensionality" (Biswas and Y 2004, So and Y 2005)

Topic 3: Pre-Trained Mixed Integer Linear Programming Solvers

Application VI: Unit Commitment and Power Grid Optimization

COPT, Cardinal Operations 2022

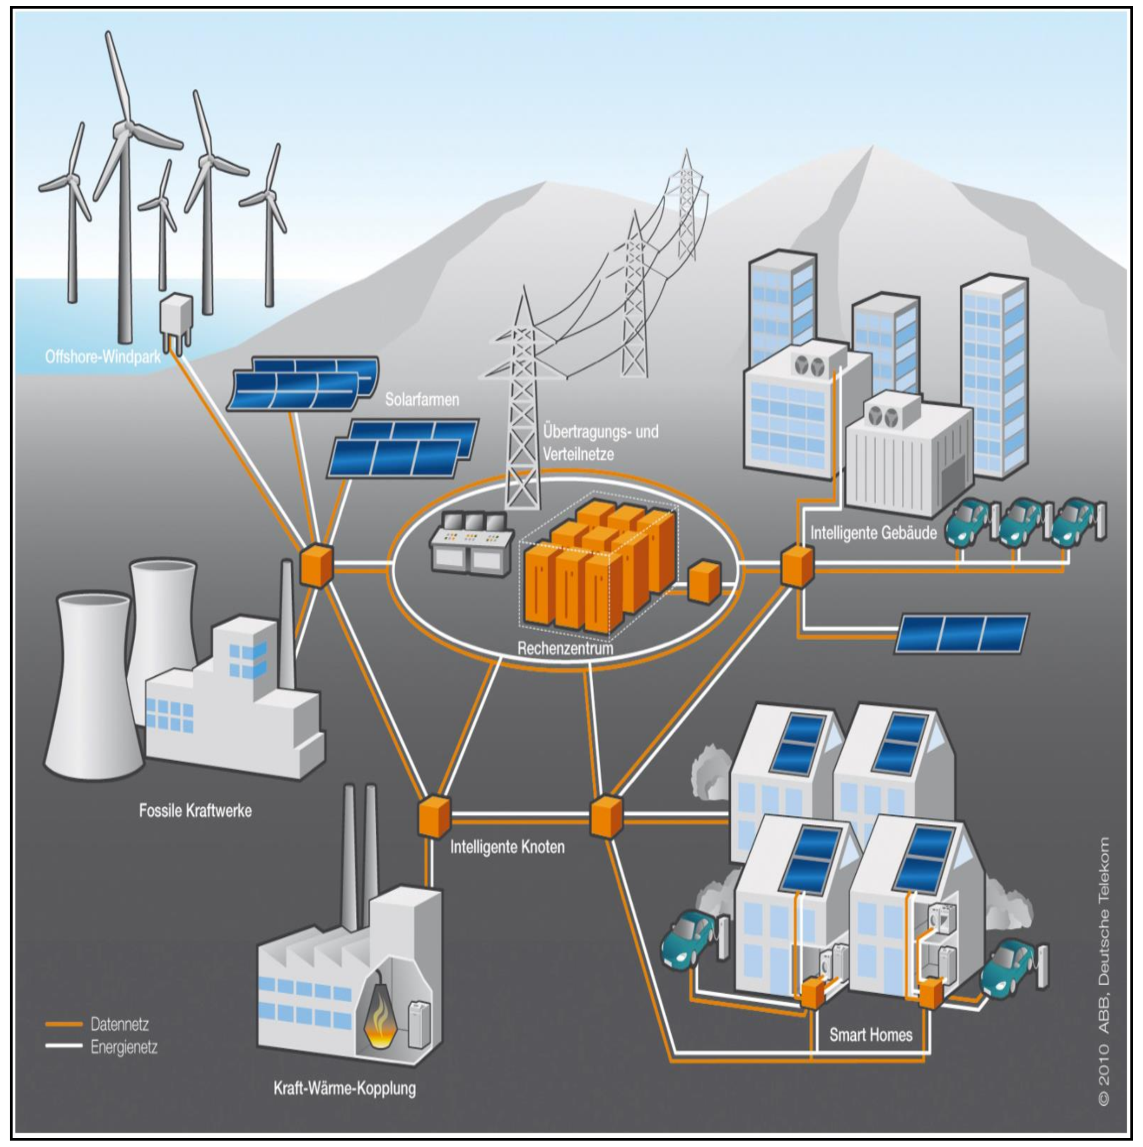

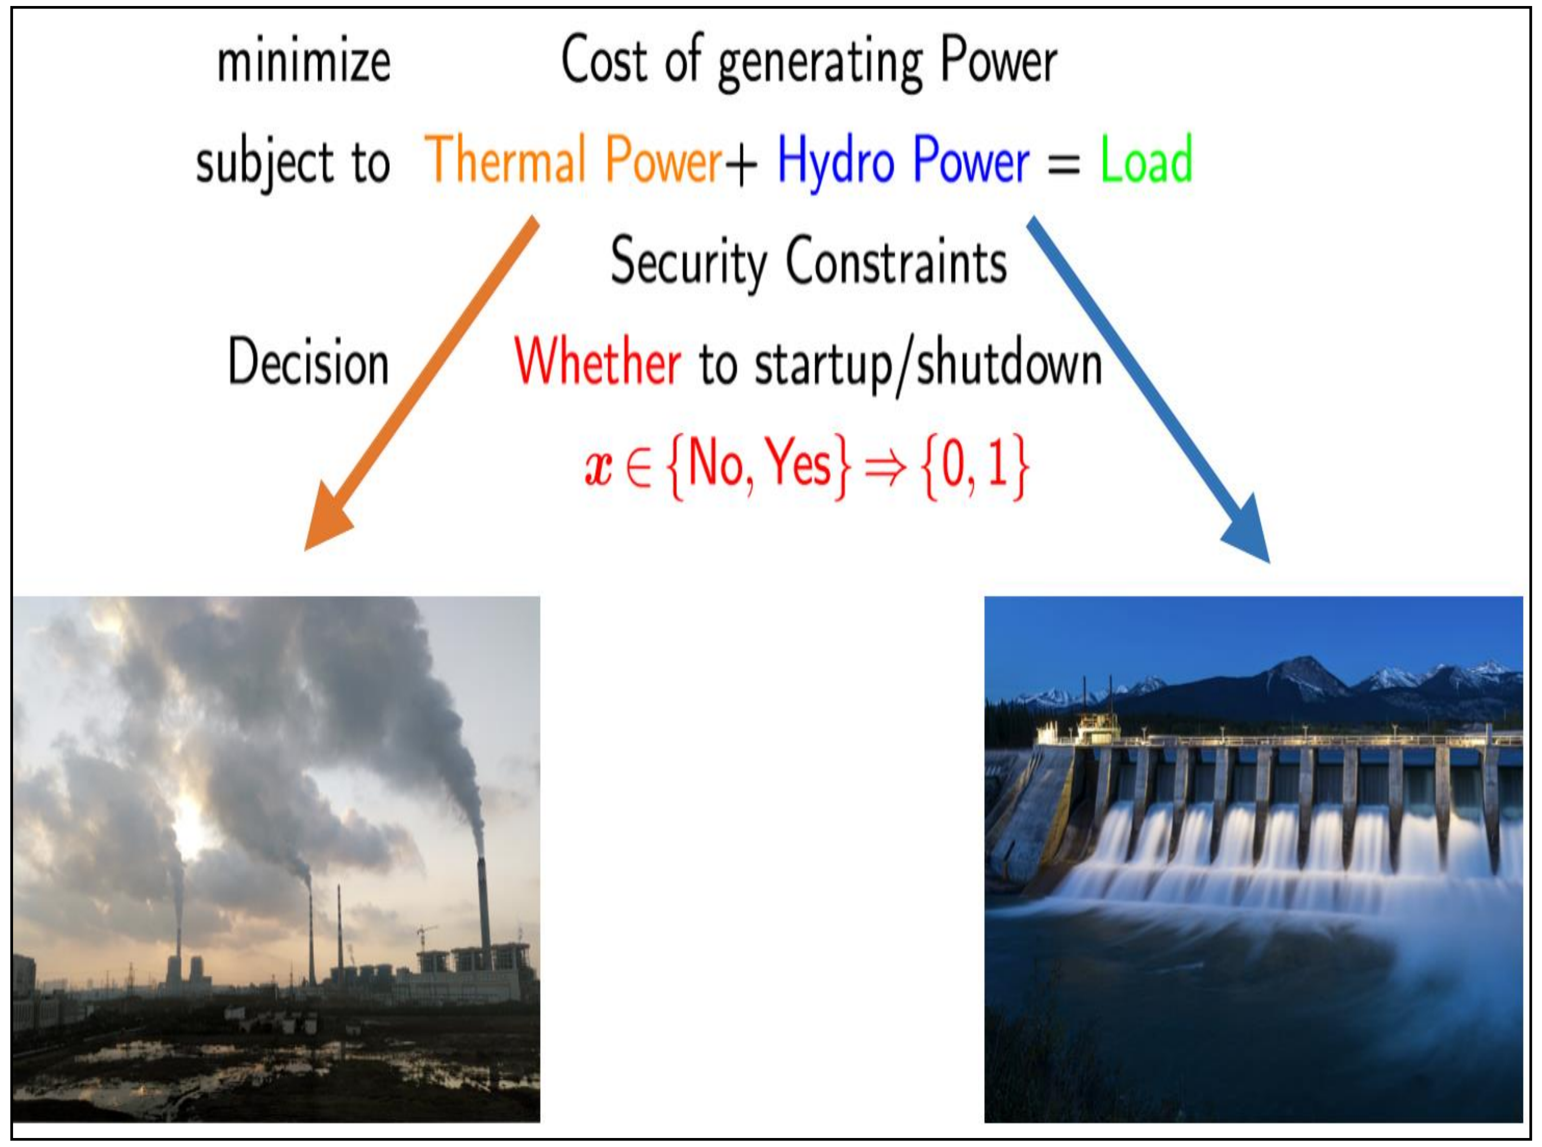

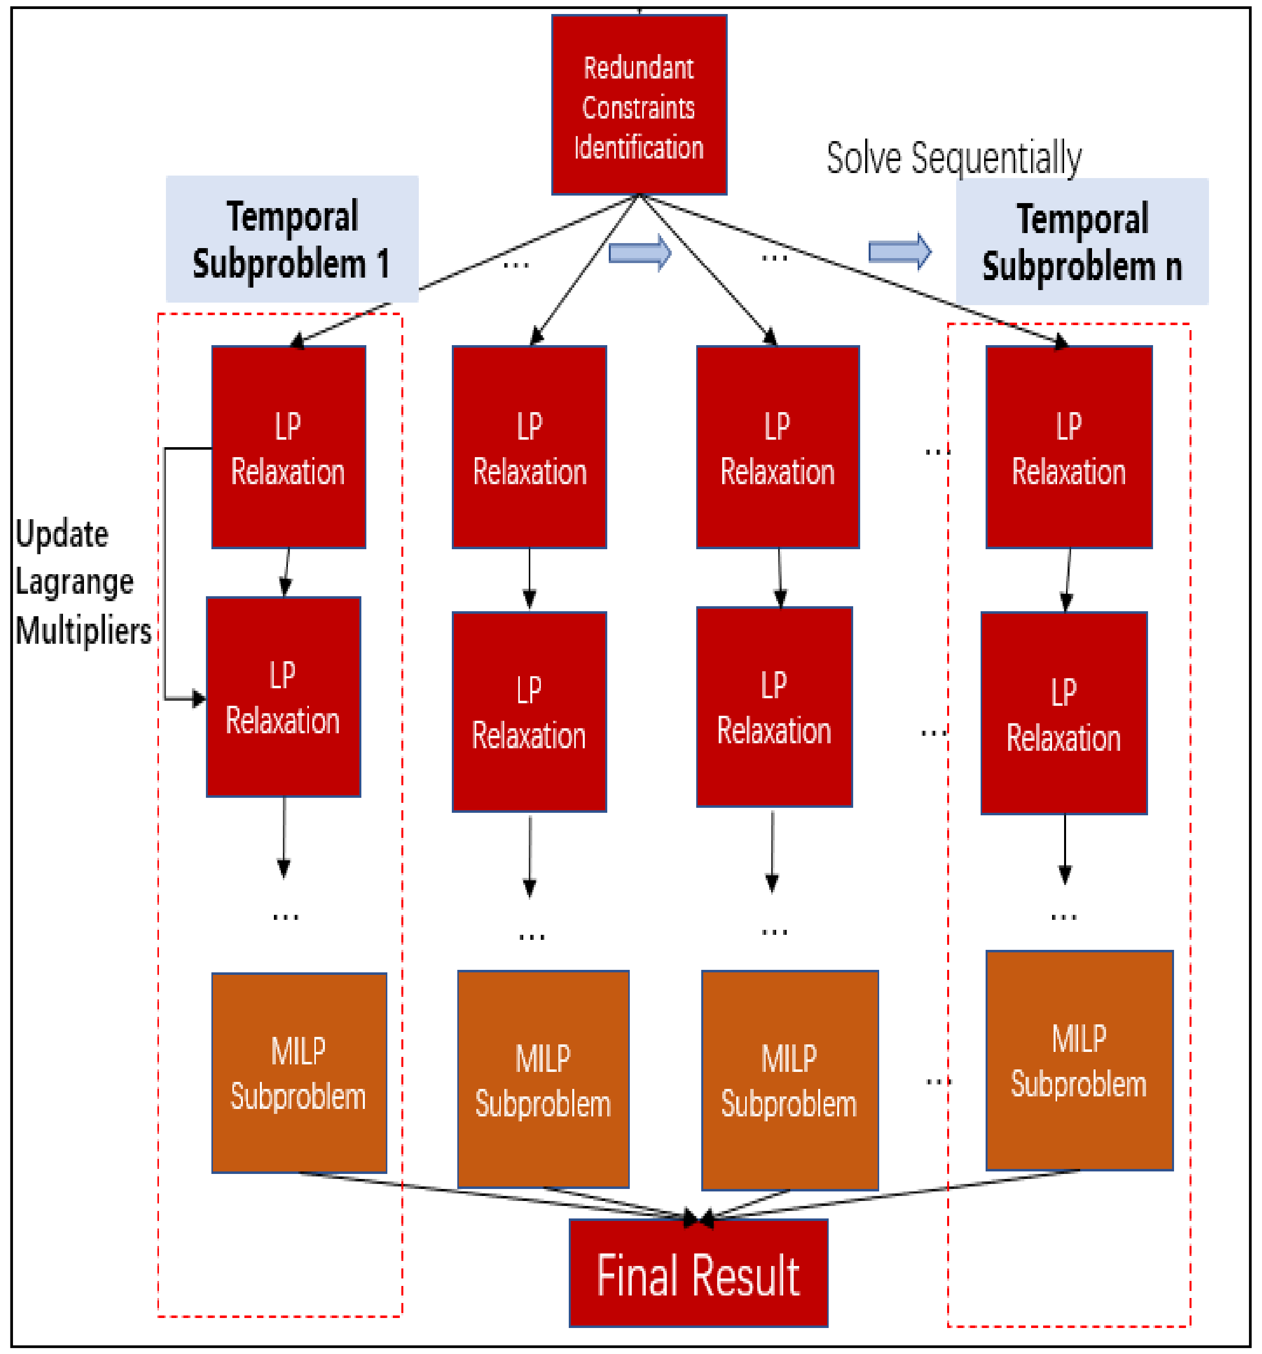

Unit Commitment Problem

- Electricity is generated from units (various generators)
- Transmitted *safely* and *stably* through power grids
- Consumed at minimum (reasonable) price

Optimization has its role to play

minimize	Cost of electricity
subject to	Safety and Stability
	Adaptivity to various units

Unit commitment problem dispatches the units
safely and stably at minimum cost

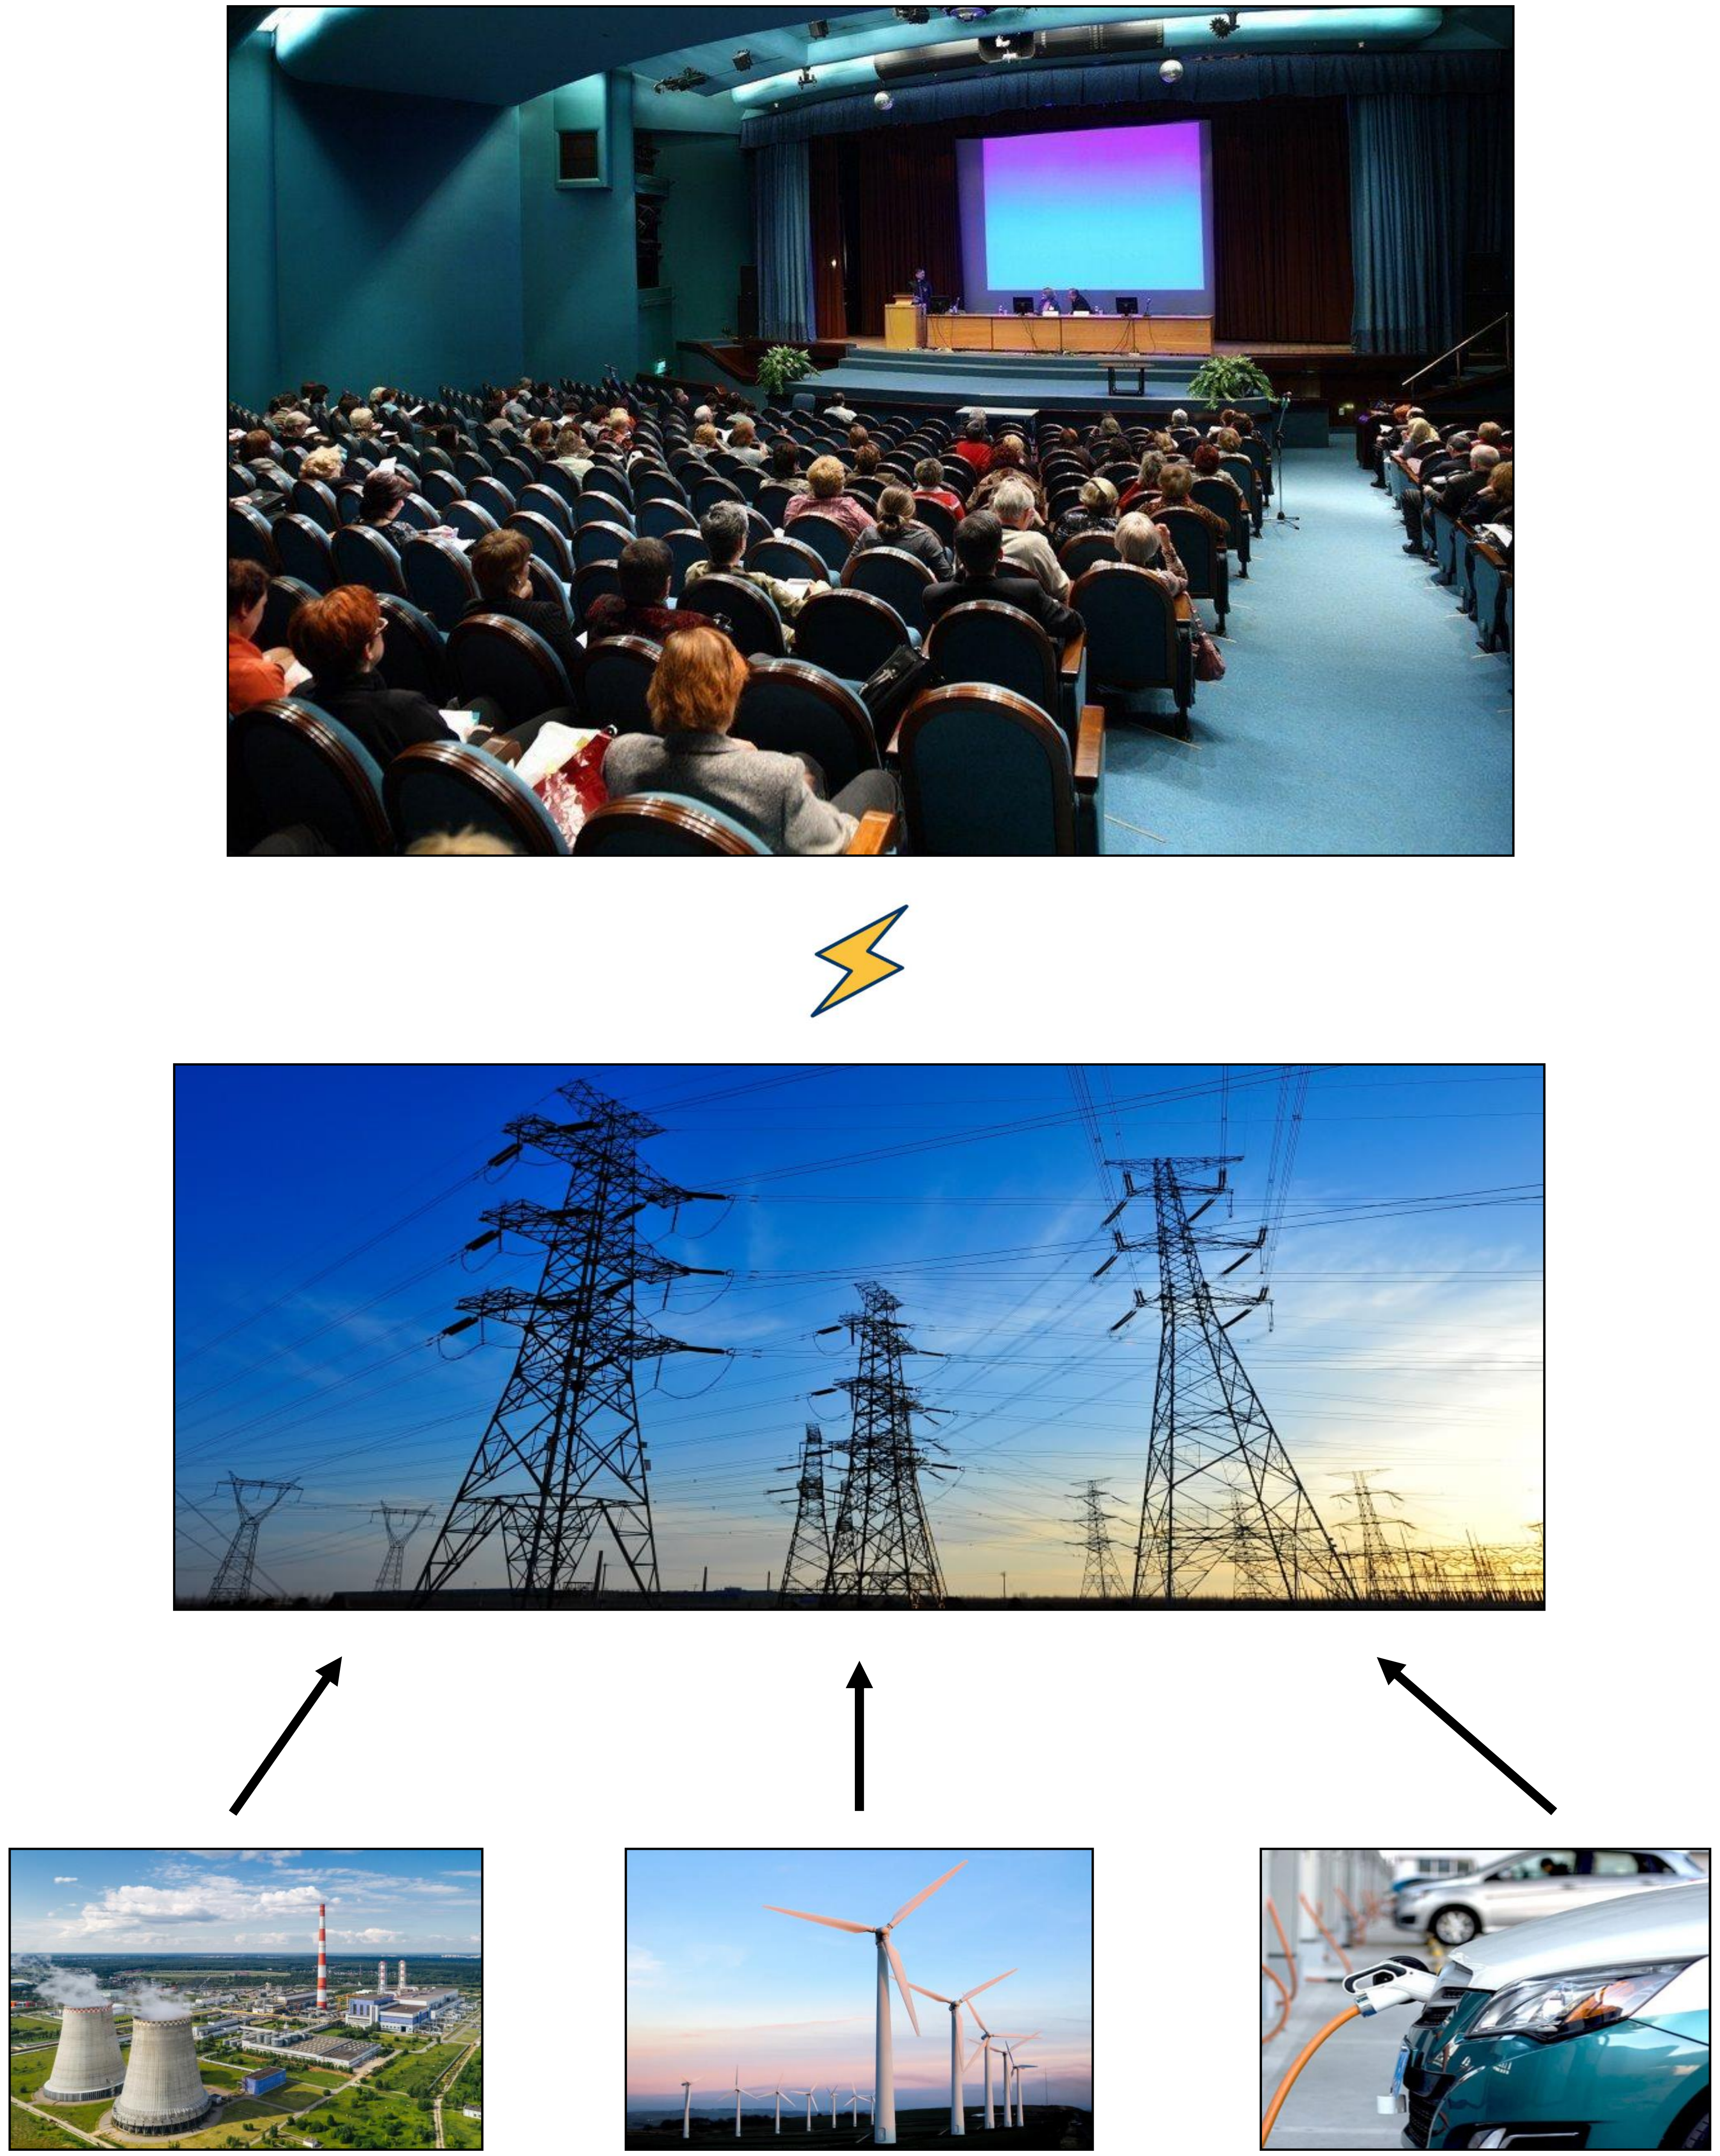

Successively Implemented in a Larger Region

- A very larger UC problem with security constraint
- With many (millions of) constraints and variables
- More than 1000 units of Thermal, Hydro and New energy
- Consider interaction between regions and time periods

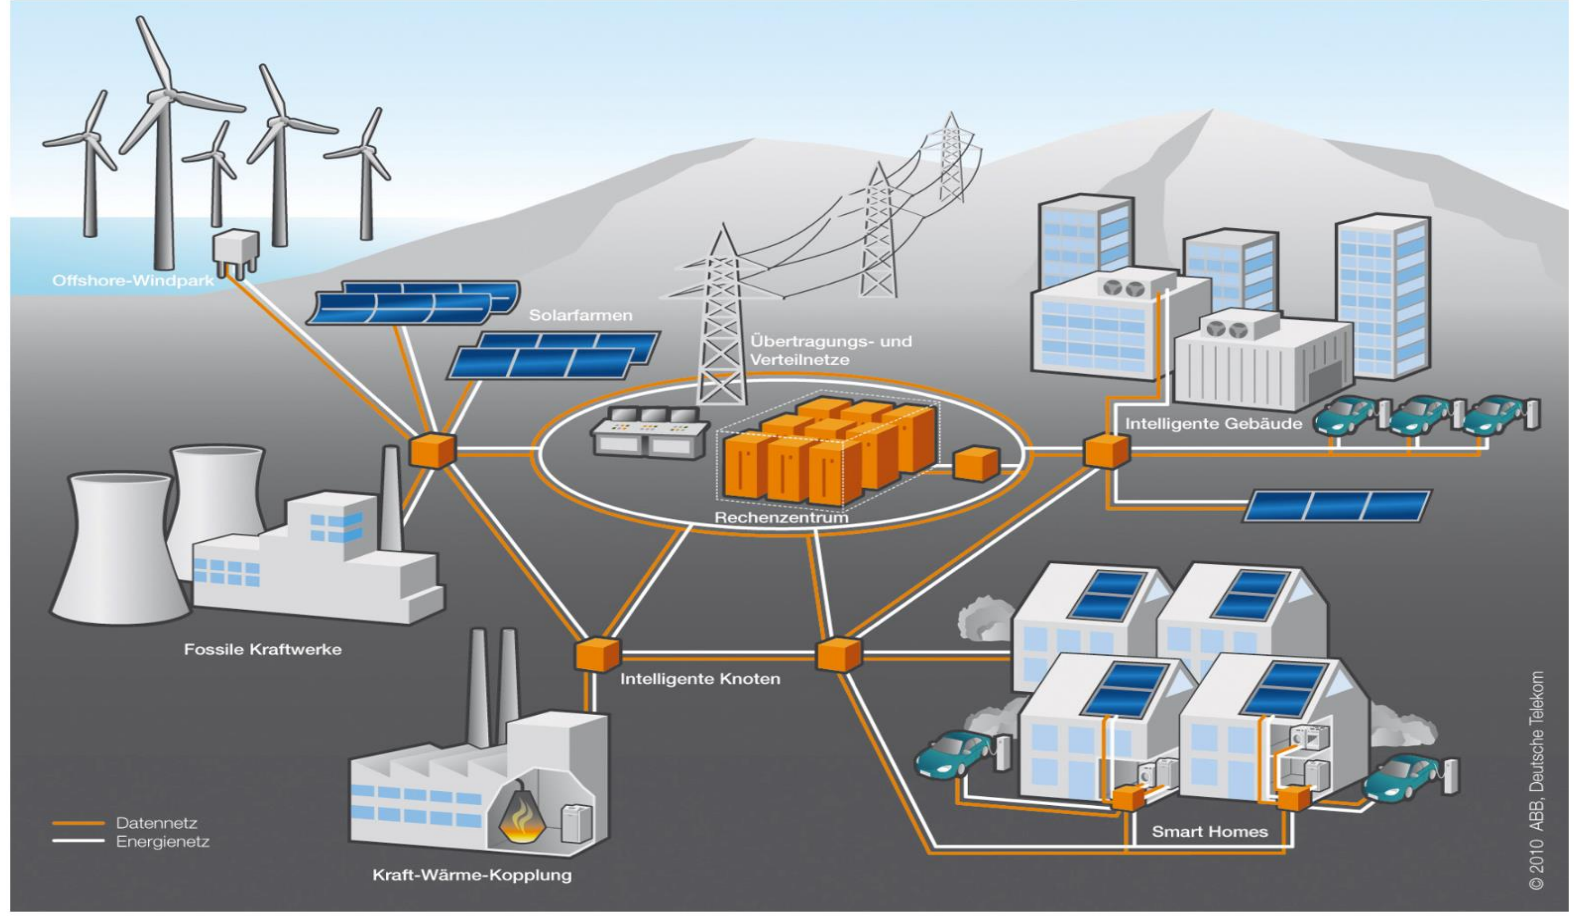

Huge size + Various business logic + Complicated coupling constraints

- Intractable without exploring structure
- Accurate and succinct model helps
- Domain specific algorithms matter a lot
- ML/AI has a big role to play

Model, Algorithm and ML/AI together make it tractable

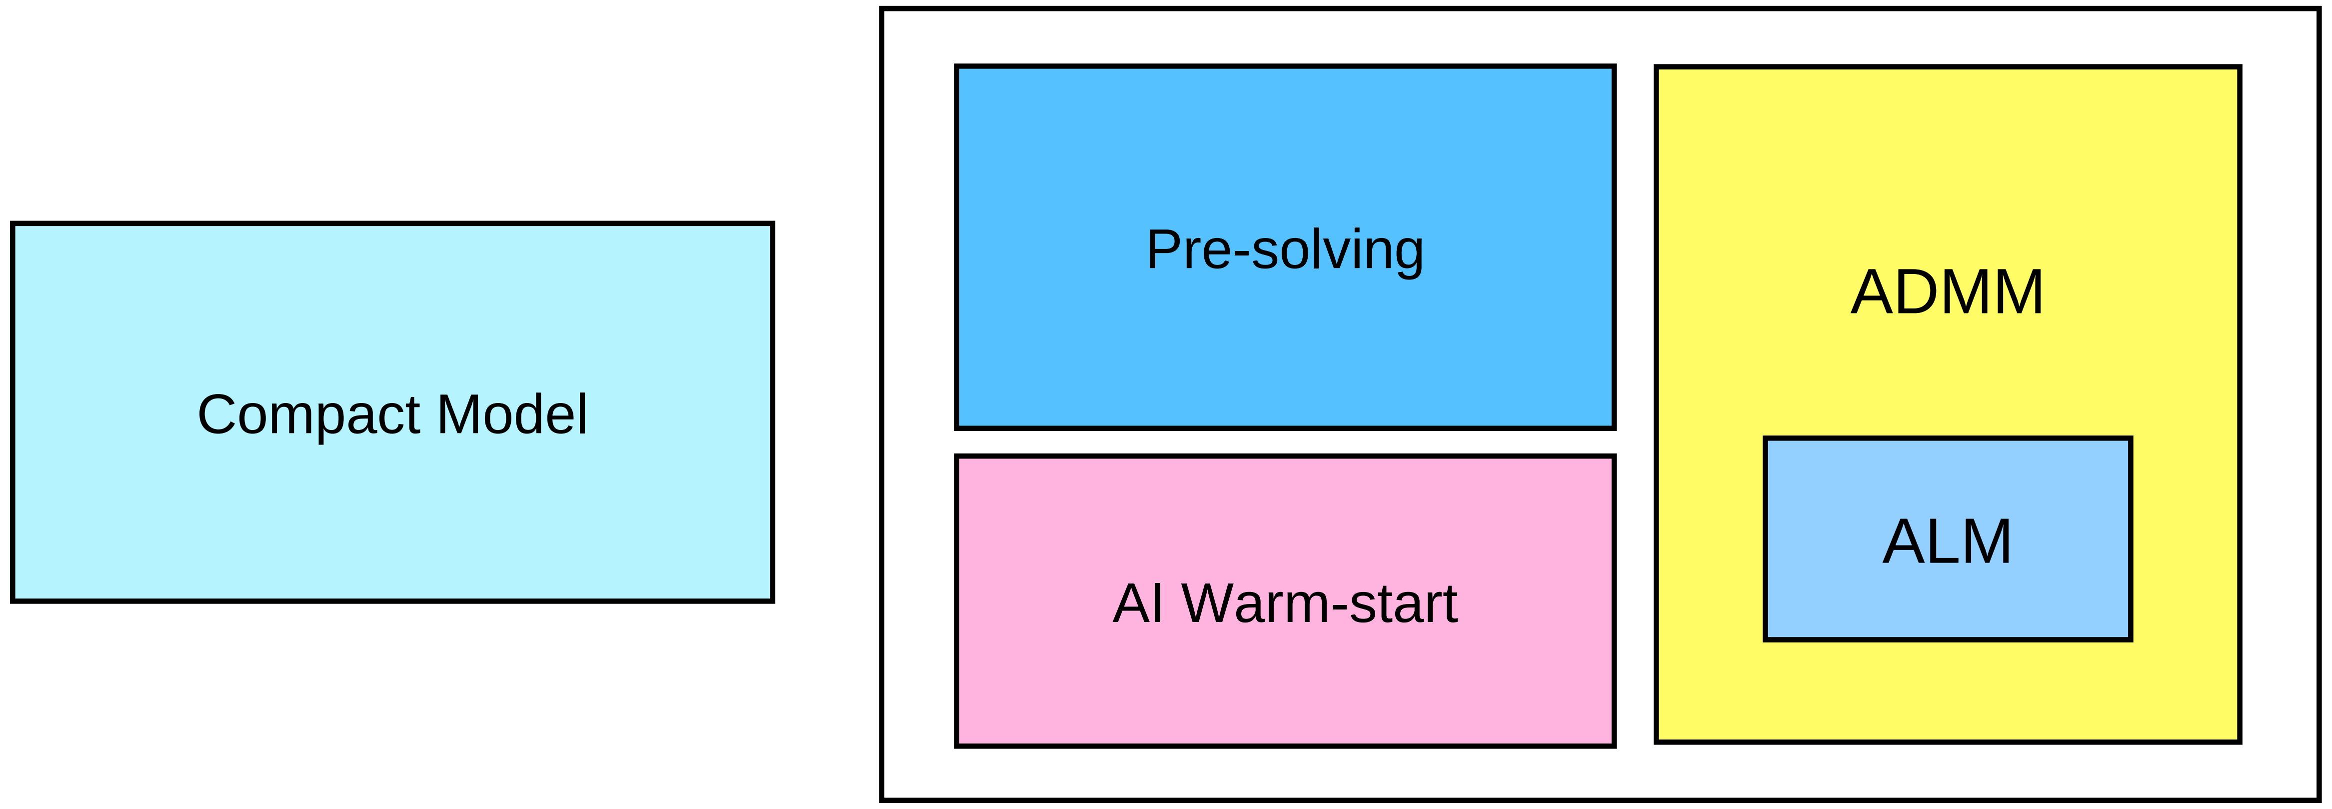

App. VII: Beijing Public Transport Intelligent Urban Bus Operations Management with Mixed Fleet Types and Charging Schedule

北京公交
Beijing Public Transport

杉数科技
Cardinal Operations

Kickoff 2022.8

Peak Reduction due to Smart Charging and Discharging

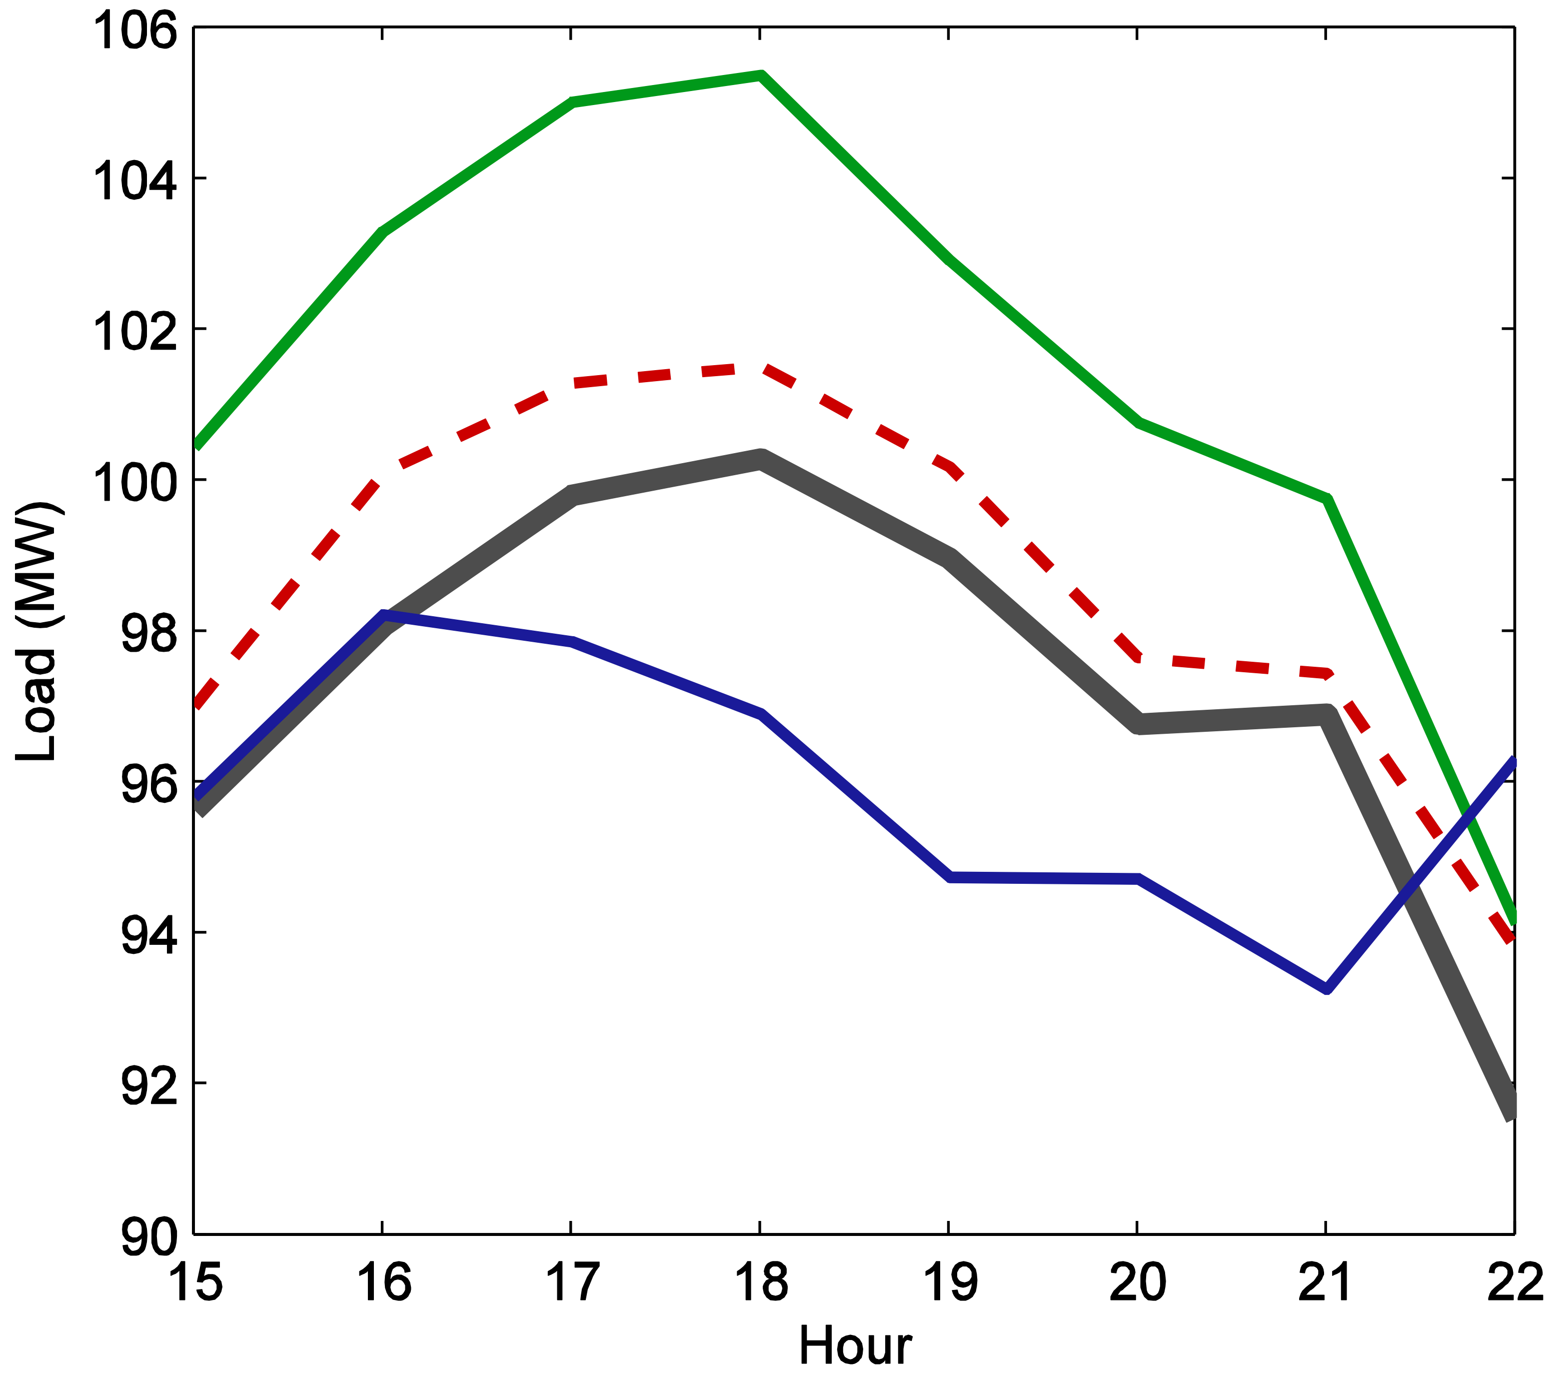

	Standard	Low PGE	Linear Progr.
Total Fleet ($)	97,678	83,695	65,349
Mean Cost / Mile	0.068	0.044	0.0054
Increase in Peak	5.1%	1.4%	-0.25%

Background: Decision Intelligence in the case of Beijing Public Transport

北京公交 Beijing Public Transport

杉数科技 Cardinal Operations

最大化工作效率
最小化总体运营成本

新能源车购车选型、车线匹配、能源布局、保养计划

运筹优化、求解器、机器学习等智能决策技术

北京市“十四五规划”目标

加快构建“综合、绿色、智能、安全”的立体化现代化城市交通系统

- 到2025年，中心城区绿色出行比例提高至 76.5%
- 全面推进智慧城市建设，重点发展智慧交通
- 围绕轨道交通优化地面公交线网，减少长距离、长时间运行线路，提高车辆利用率

加快建立科学、高效的“城市智能运行决策管理体系”

- 高水平推动城市交通的数字转型和智慧升级，形成城市交通整体解决方案
- 加快建设公共交通网络化智能调度体系，让公交出行越来越可靠，时间有保证

More efficient and intelligent decision-making to achieve 14th Five-Year Plan goals

Beijing Public Transport Line 7 is selected as the Key Pilot Unit of the intelligent transformation of Beijing Public Transport

App. VIII: Beijing-Shanghai High-speed Railway Scheduling Optimization

COPT, **Cardinal Operations** 2022

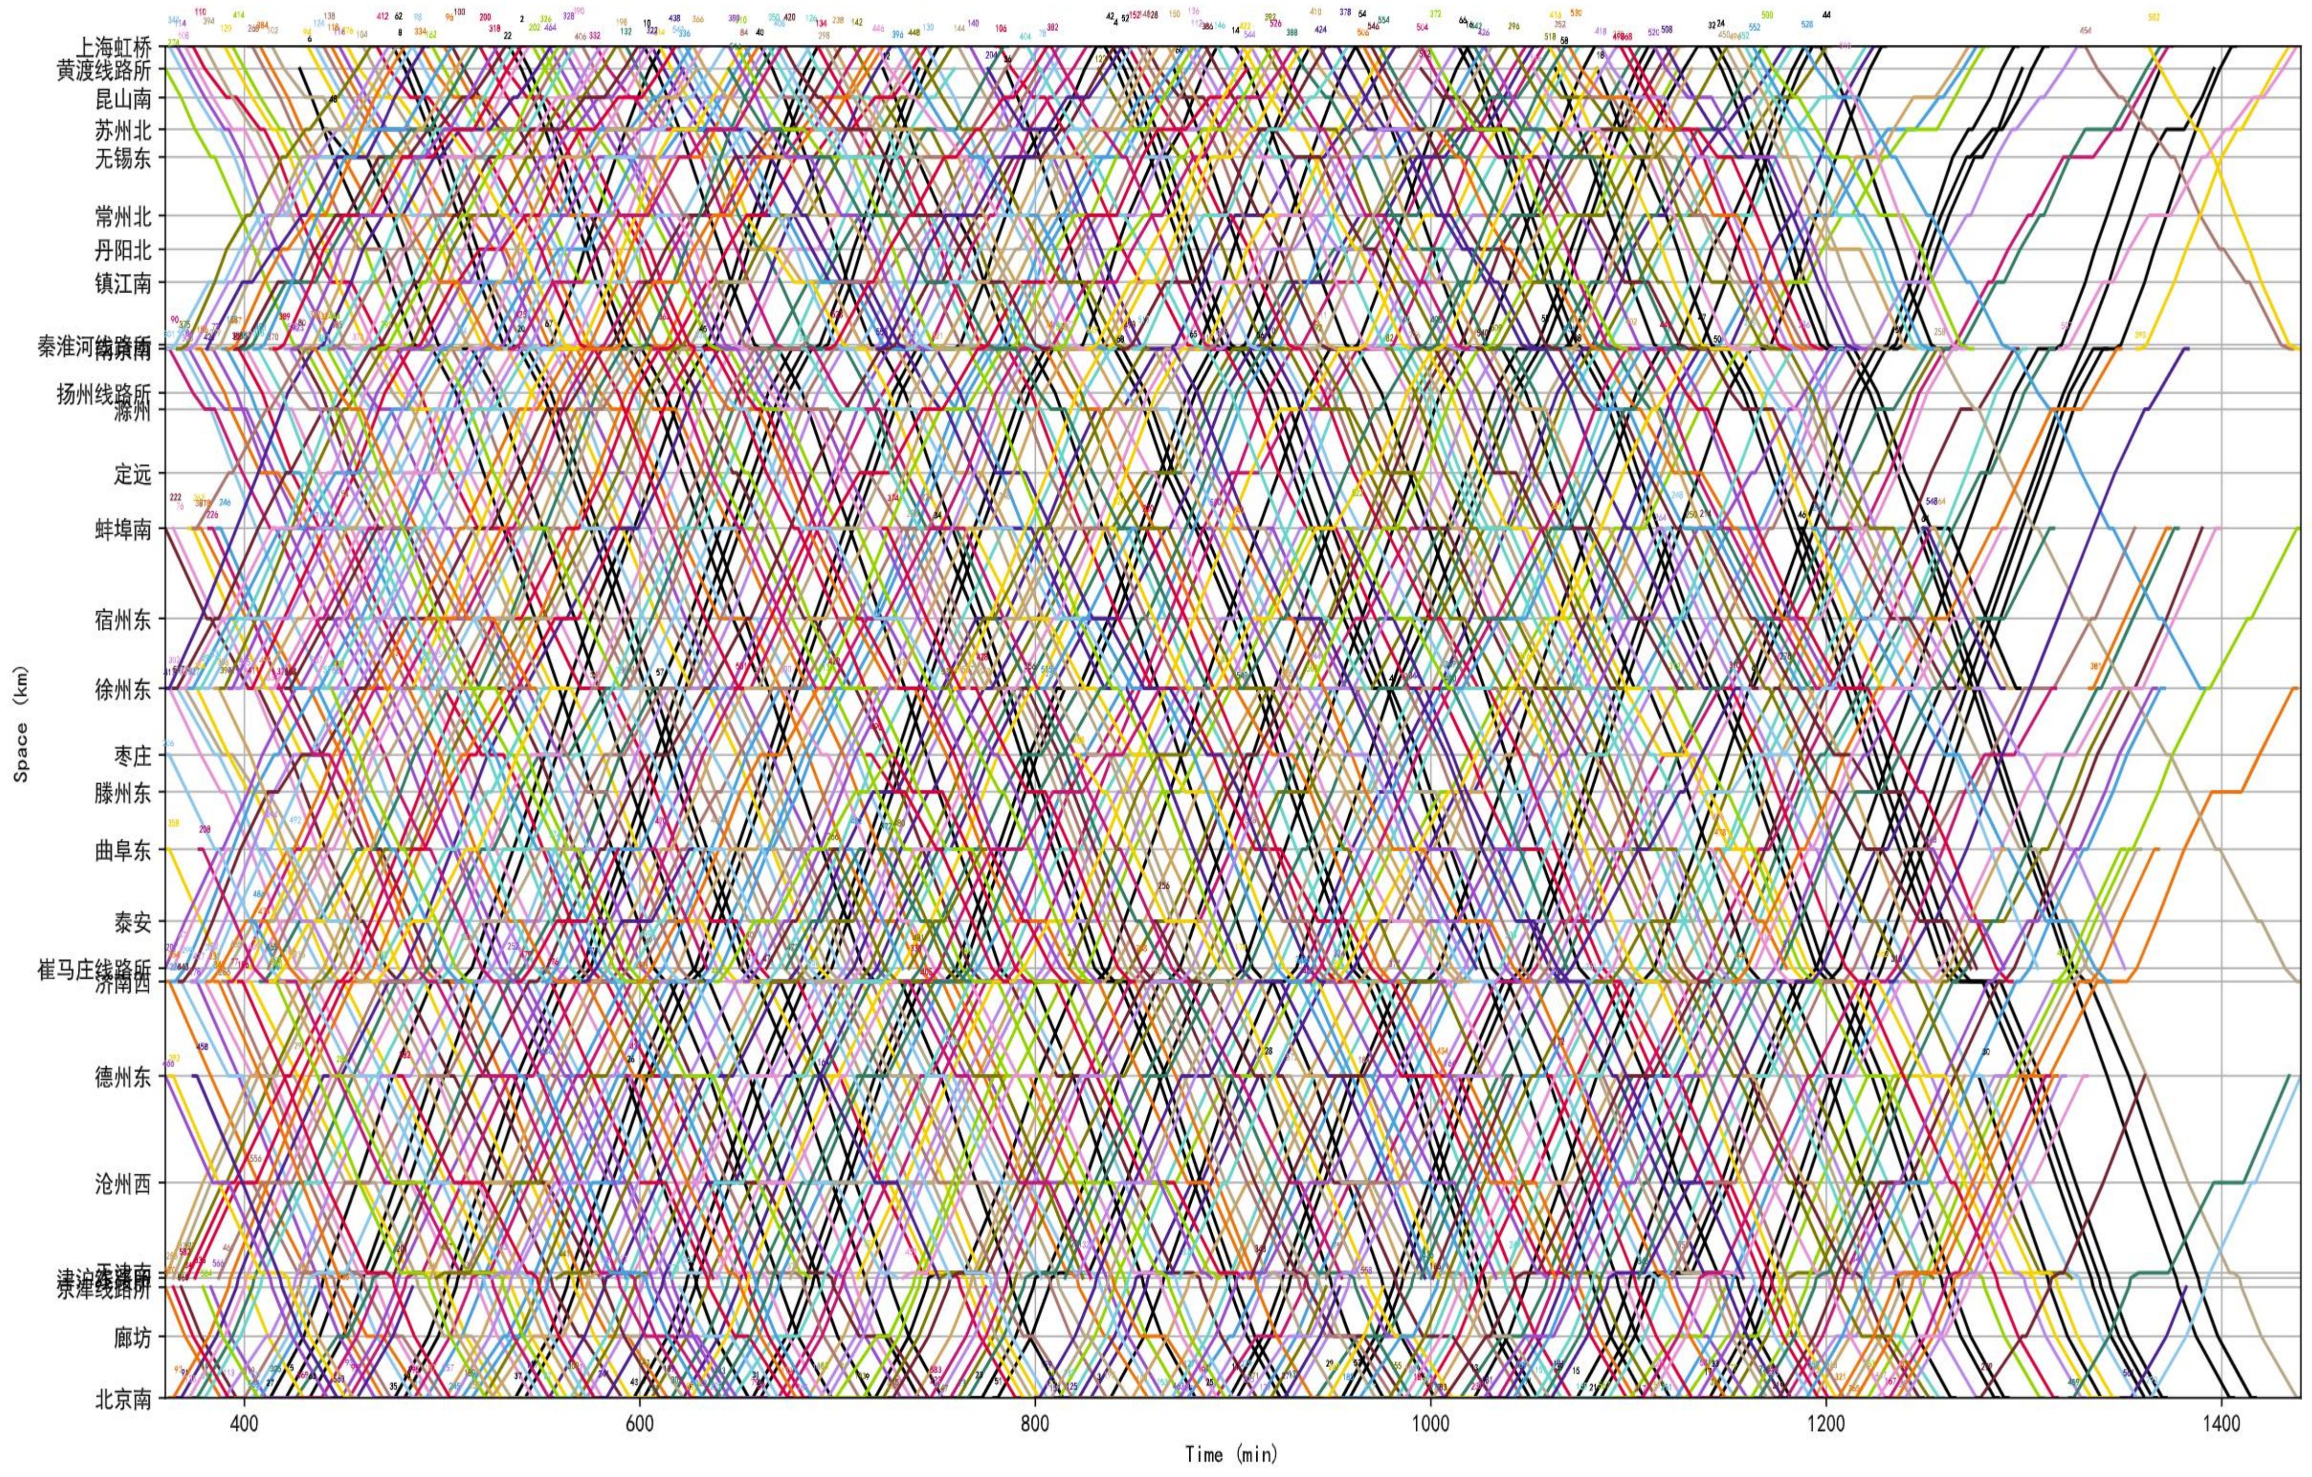

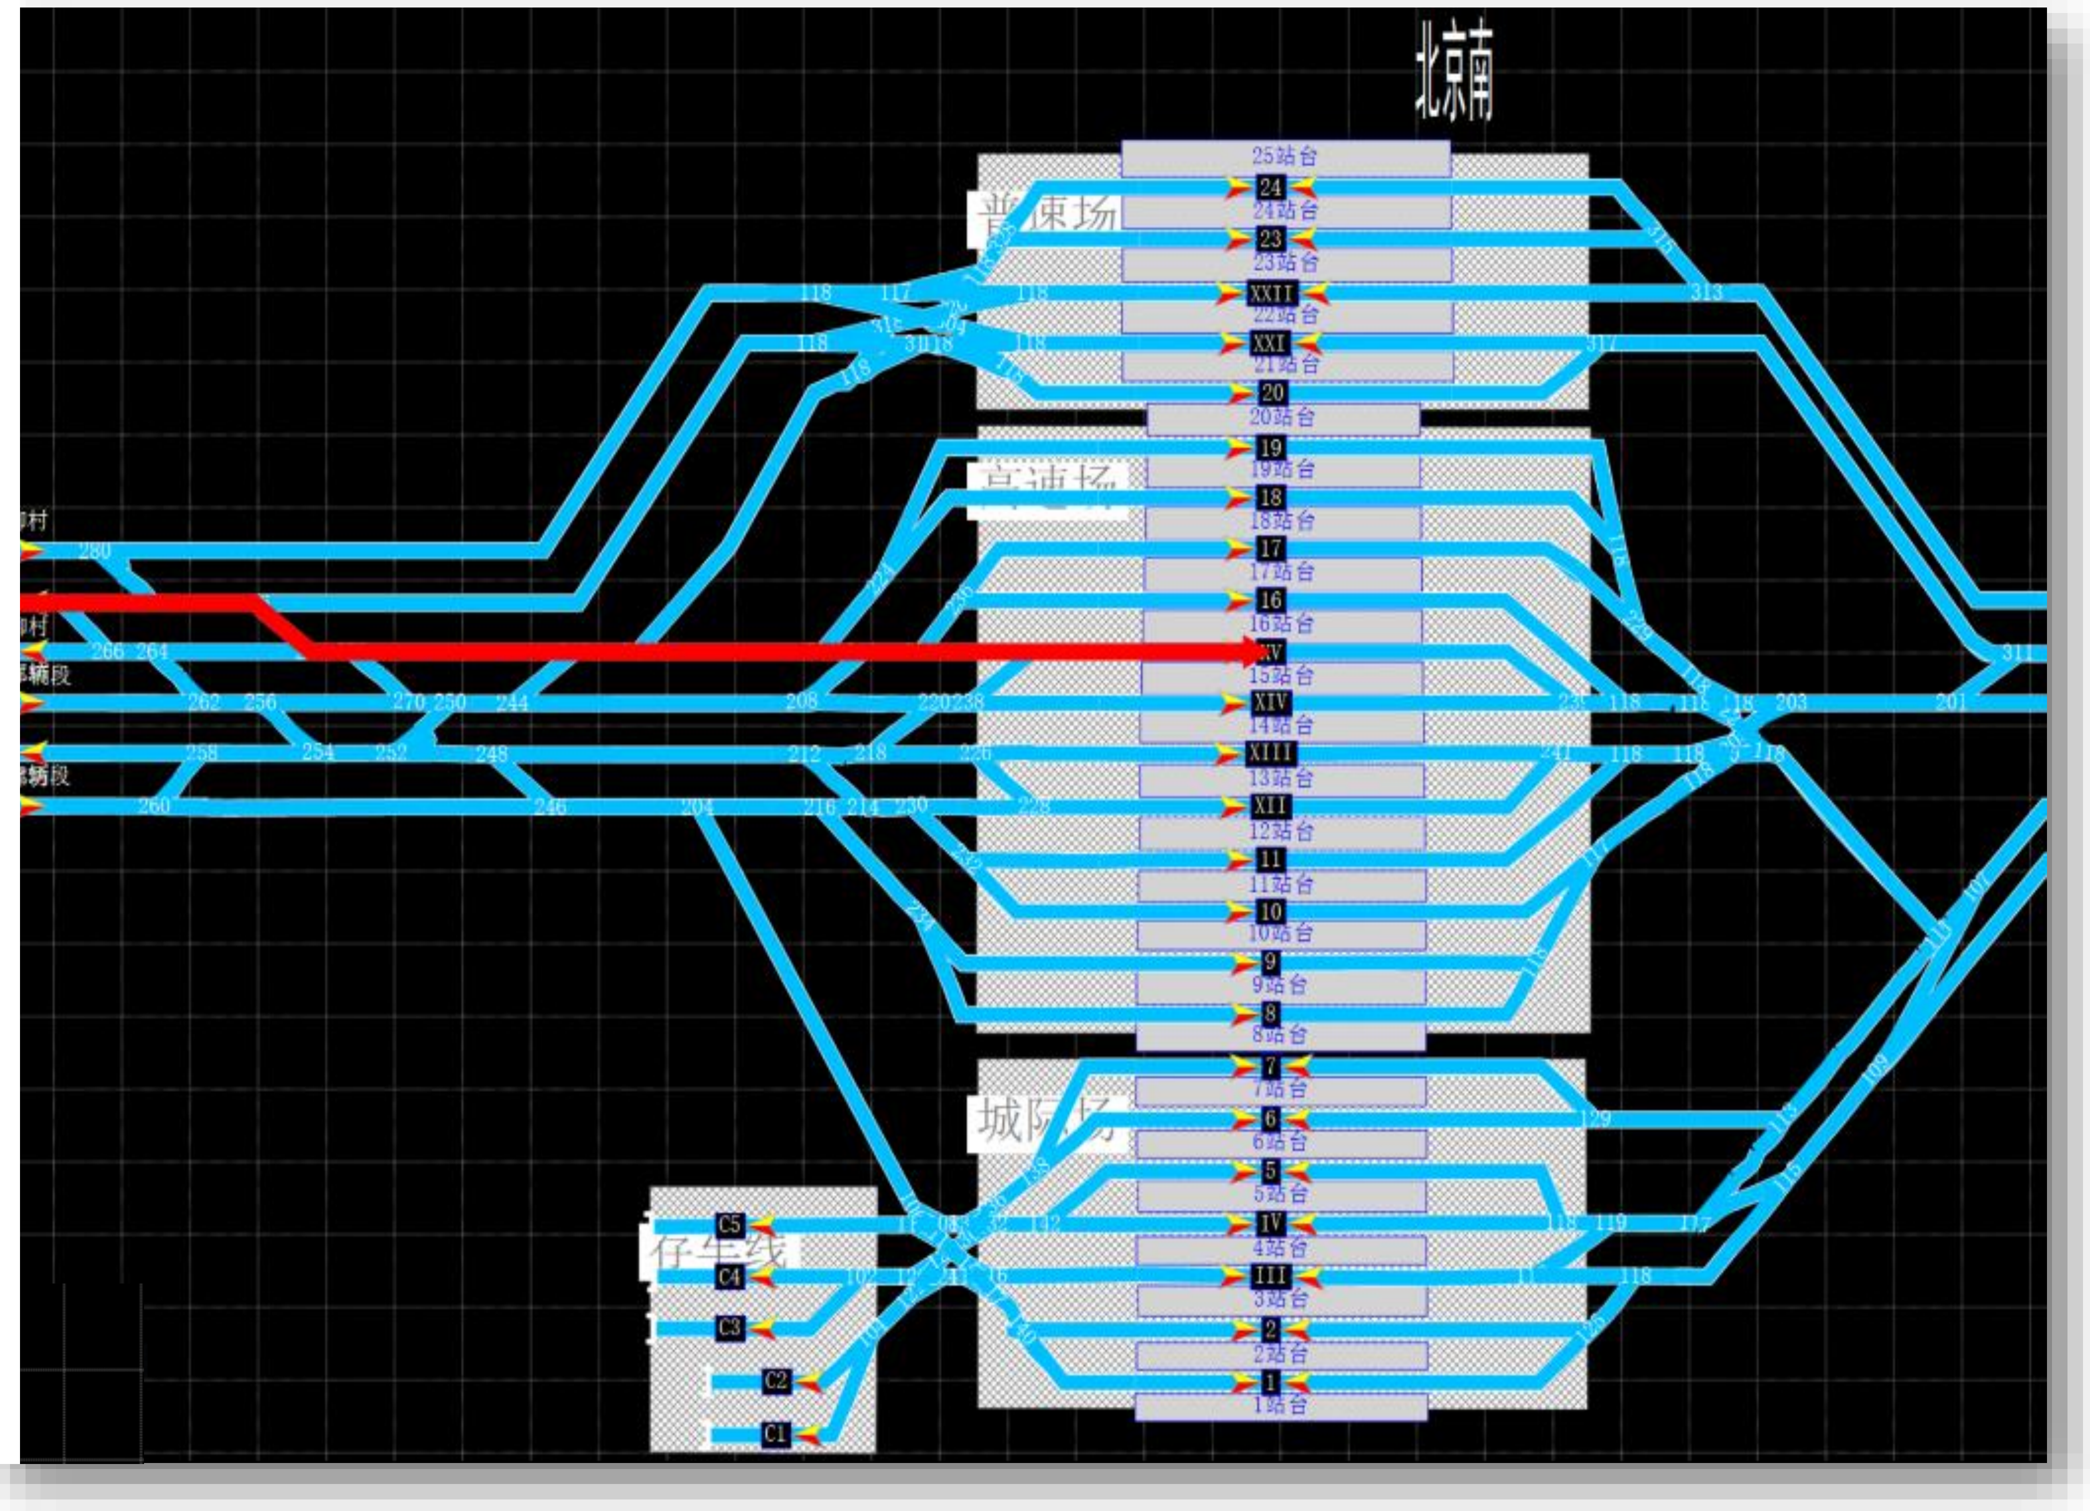

Background

- China High-speed Railway has been committed to providing high-quality transportation services to passengers, and the formulation of train scheduling is a key link in the operation.
- **At present, train scheduling is based on human experience**, which becomes increasingly difficult to handle the growing network. Therefore, both industry and academia are seeking ways to **automate train scheduling**.
- The train scheduling problem can be divided into **Train Timetabling Problem (TTP)** and **Train Platforming Problem (TPP)**.

Optimization Model:

- **Objective:** maximize the number of trains placed in the train scheduling, thereby maximizing operating revenue；
- **Constraints:** describe the running behavior of trains and prevent train collisions；

- The project mainly solves **TTP for Beijing-Shanghai High-speed Railway** and **TPP at Beijingnan Railway Station**.
  - **Beijing-Shanghai High-speed Railway** is the busiest high-speed railway with the largest number of passengers in China. It is 1,318 km in total and passes 29 stations.
  - **Beijingnan Railway Station** is the largest railway station in Beijing, with the largest area and the largest number of trains.
  - Both problems are challenging scheduling tasks, which can be formulated as Mixed Integer Programming (MIP).

Numerical Results: TTP for Beijing-Shanghai

- We solve the TTP for Beijing-Shanghai high-speed railway using Cardinal Optimizer (COPT).
- COPT is the first fully independently developed mathematical programming solver in China with strong solving ability of MIP problem. It also has excellent performance in solving this problem.
- The result is presented in the following figure. We only need about **1000 seconds** to schedule 584 train in two directions.

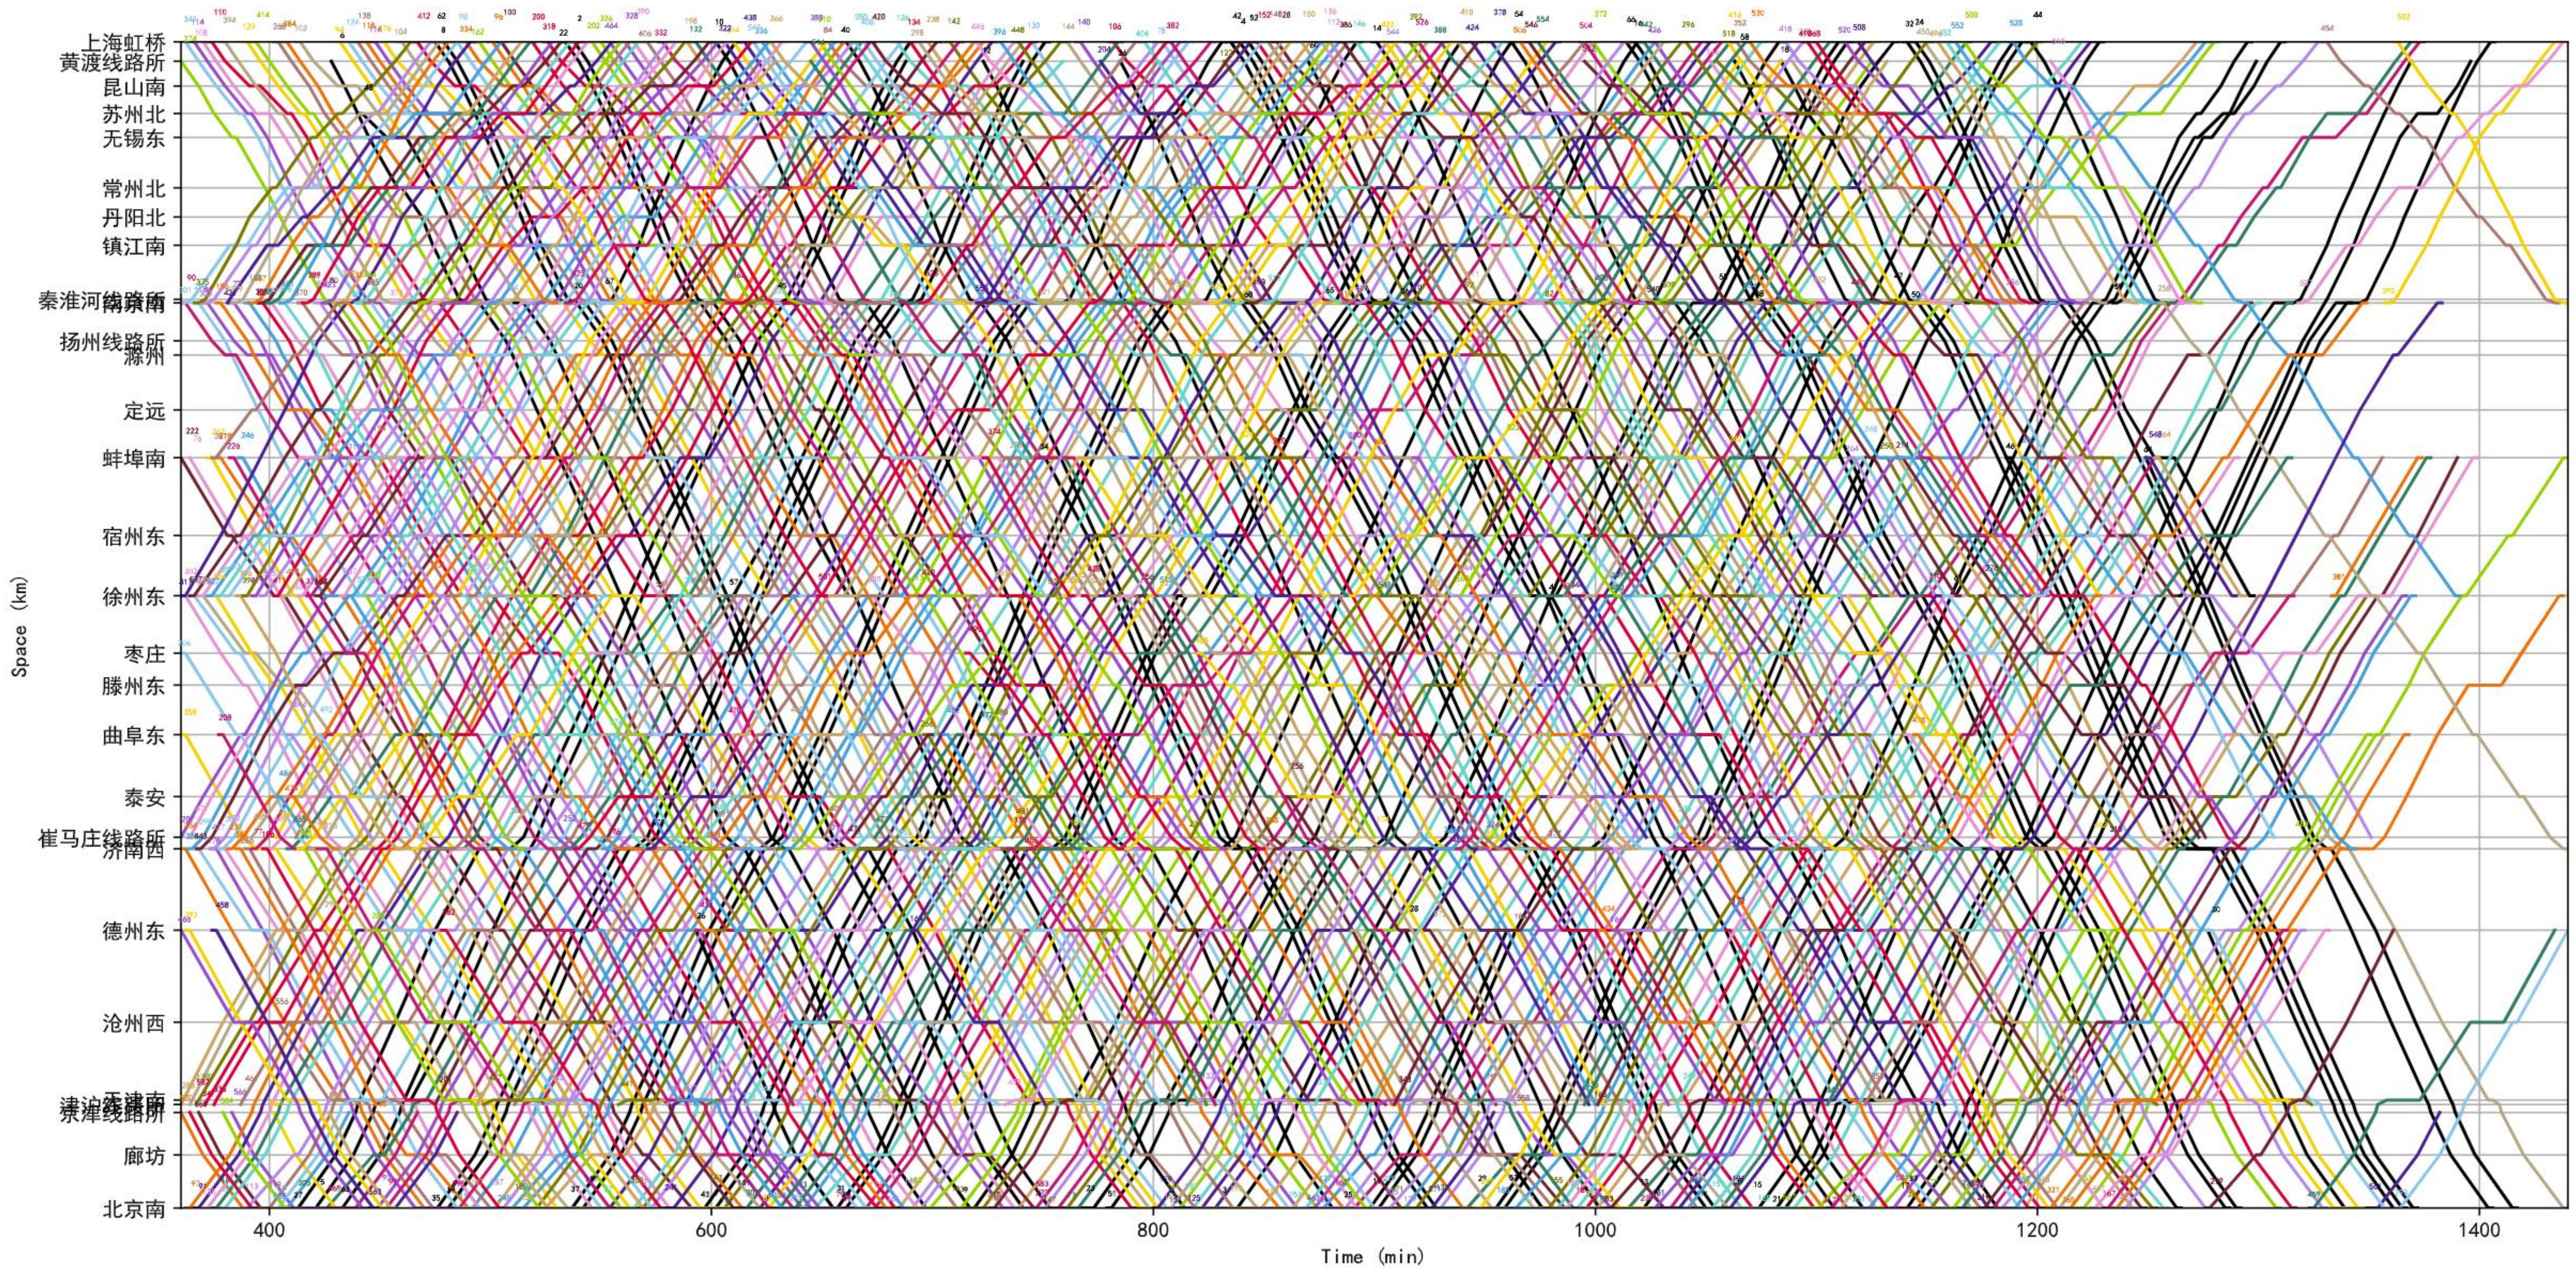

Old Technology: Pre-Trained Heuristics for Mixed Integer Optimization Solvers for solving a specific problem instance

Most pre-solving heuristics for mixed integer optimization works as follows, where *y* are binary variables,

- **Given history instances from the same distribution** $\{(\mathcal{P}(\xi_1), y^*(\xi_1)), \ldots \mathcal{P}(\xi_m, y^*(\xi_m))\}$
- **Use some learning models to capture the relation** $\hat{y}(\xi) \approx y^*(\xi)$

  regression models, trees, deep networks …
- **Use the information** $\hat{y}(\xi)$ **to speedup MIO solving by**

  fixing variables, faster primal solution finding, faster branching

We try to address it

Good News: the learning models have been quite mature

Bad News: how to reduce RISK of the learned information remains challengy

New Technology: Pre-Trained Optimization via Statistical Cut Generation

While most pre-trained optimizers

- **Either resort to complicated implementation:**

  many require a nontrivial modification of the branch and bound procedure

- **And/or lacks explainability and is not provably optimal:**

  there is a risk of missing the true optimal solution

We propose a statistical cut generation procedure that is *simple, efficient* and *provable*

New Technology: Pooling the Risk via Variance Reduction

- Given a new problem, assume we have in hand a set of binary predictions

$$\begin{pmatrix} \hat{y}_1(\xi) \\ \hat{y}_2(\xi) \\ \vdots \\ \hat{y}_n(\xi) \end{pmatrix} = \begin{pmatrix} 0.99 \\ 0.12 \\ \vdots \\ 0.38 \end{pmatrix}$$

- Intuitively we can interpret each $\hat{y}$ as the likelihood a variable takes 1.0 in the optimal solution

- Each variable introduces some risk/variance of such rounding

  so that dealing them separately results in extremely risk outcomes

Q: What should we do to deal with a bunch of risky guesses? A: Put them in a pool!

New Technology: Risk Pooling through Theory of Concentration

- **Pooling the risk means putting binary variables together such as**

$$\sum_{i \in \mathcal{U} = \{j : \hat{y}_j(\xi) \geq 0.9\}} y_i^*(\xi) \geq \alpha \cdot |\mathcal{U}| \qquad \sum_{i \in \{j : \hat{y}_j(\xi) \leq 0.1\}} y_i^*(\xi) \leq \beta \cdot |\mathcal{L}|$$

- **Intuitively we know that the above two inequalities are expectedly to hold for $\alpha \to 0.9$ and $\beta \to 0.1$**
- **These two inequalities are exactly *cutting planes* for MIO**
- **The last issue is how to choose α, β to increase the confidence level:**

interpret y^* as Bernoulli random variables with expectation $\hat{y}$, then apply the concentration inequality theory

New Technology: Pre-Trained Statistical (Confidence) Cut Generation

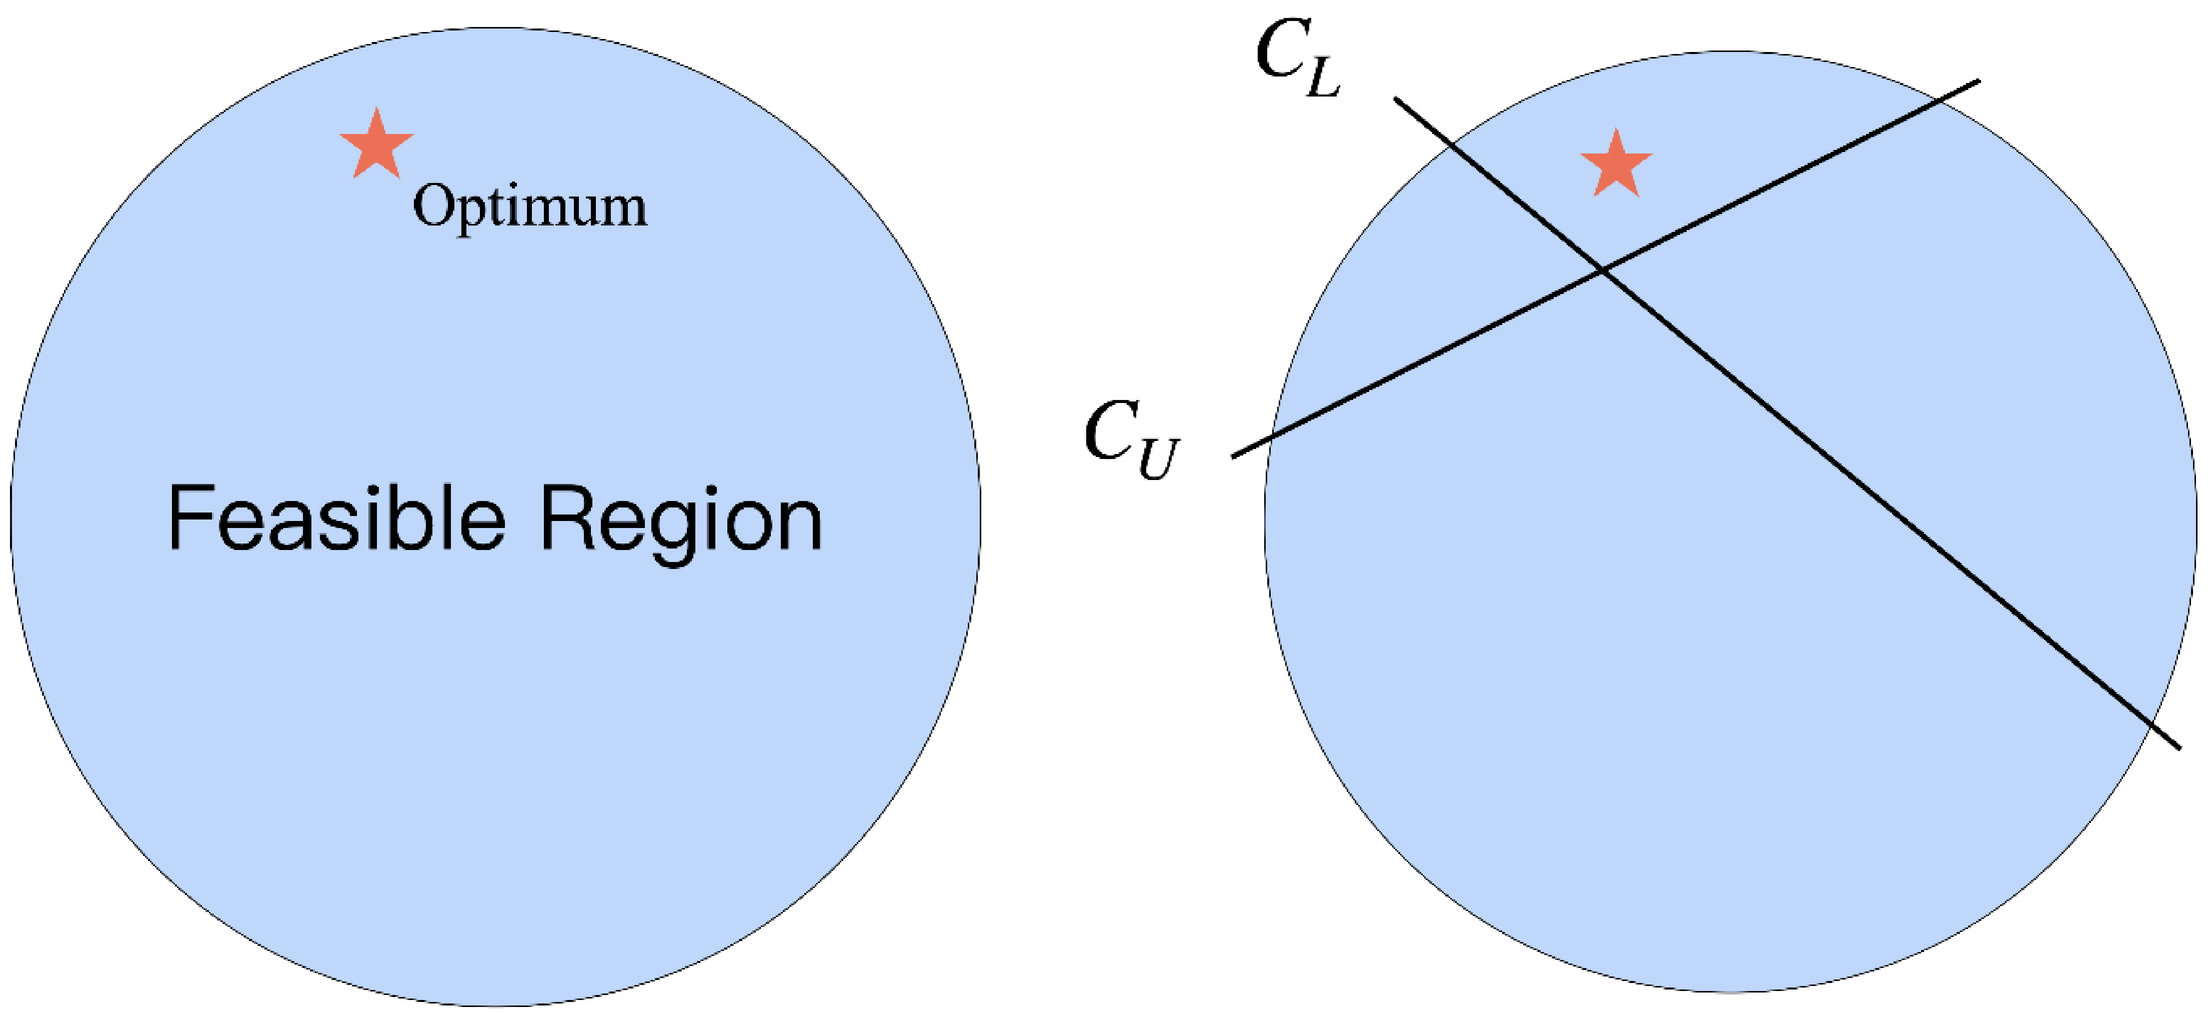

- **Overall, the two cuts (and their complement) split the whole feasible region into four regions**
- **Solving the most likelihood region of two cuts often gives a satisfying solution with confidence**
- **Enumerating over the four regions independently will not miss the optimal solution**

New Technology: Numerical Experiments

- The method is tested on multi-knapsack, set-covering and unit-commitment problems

  Train on 500 instances and test on 20 instances

- Measure the speedup of finding a good solution on in the region formed by two cuts

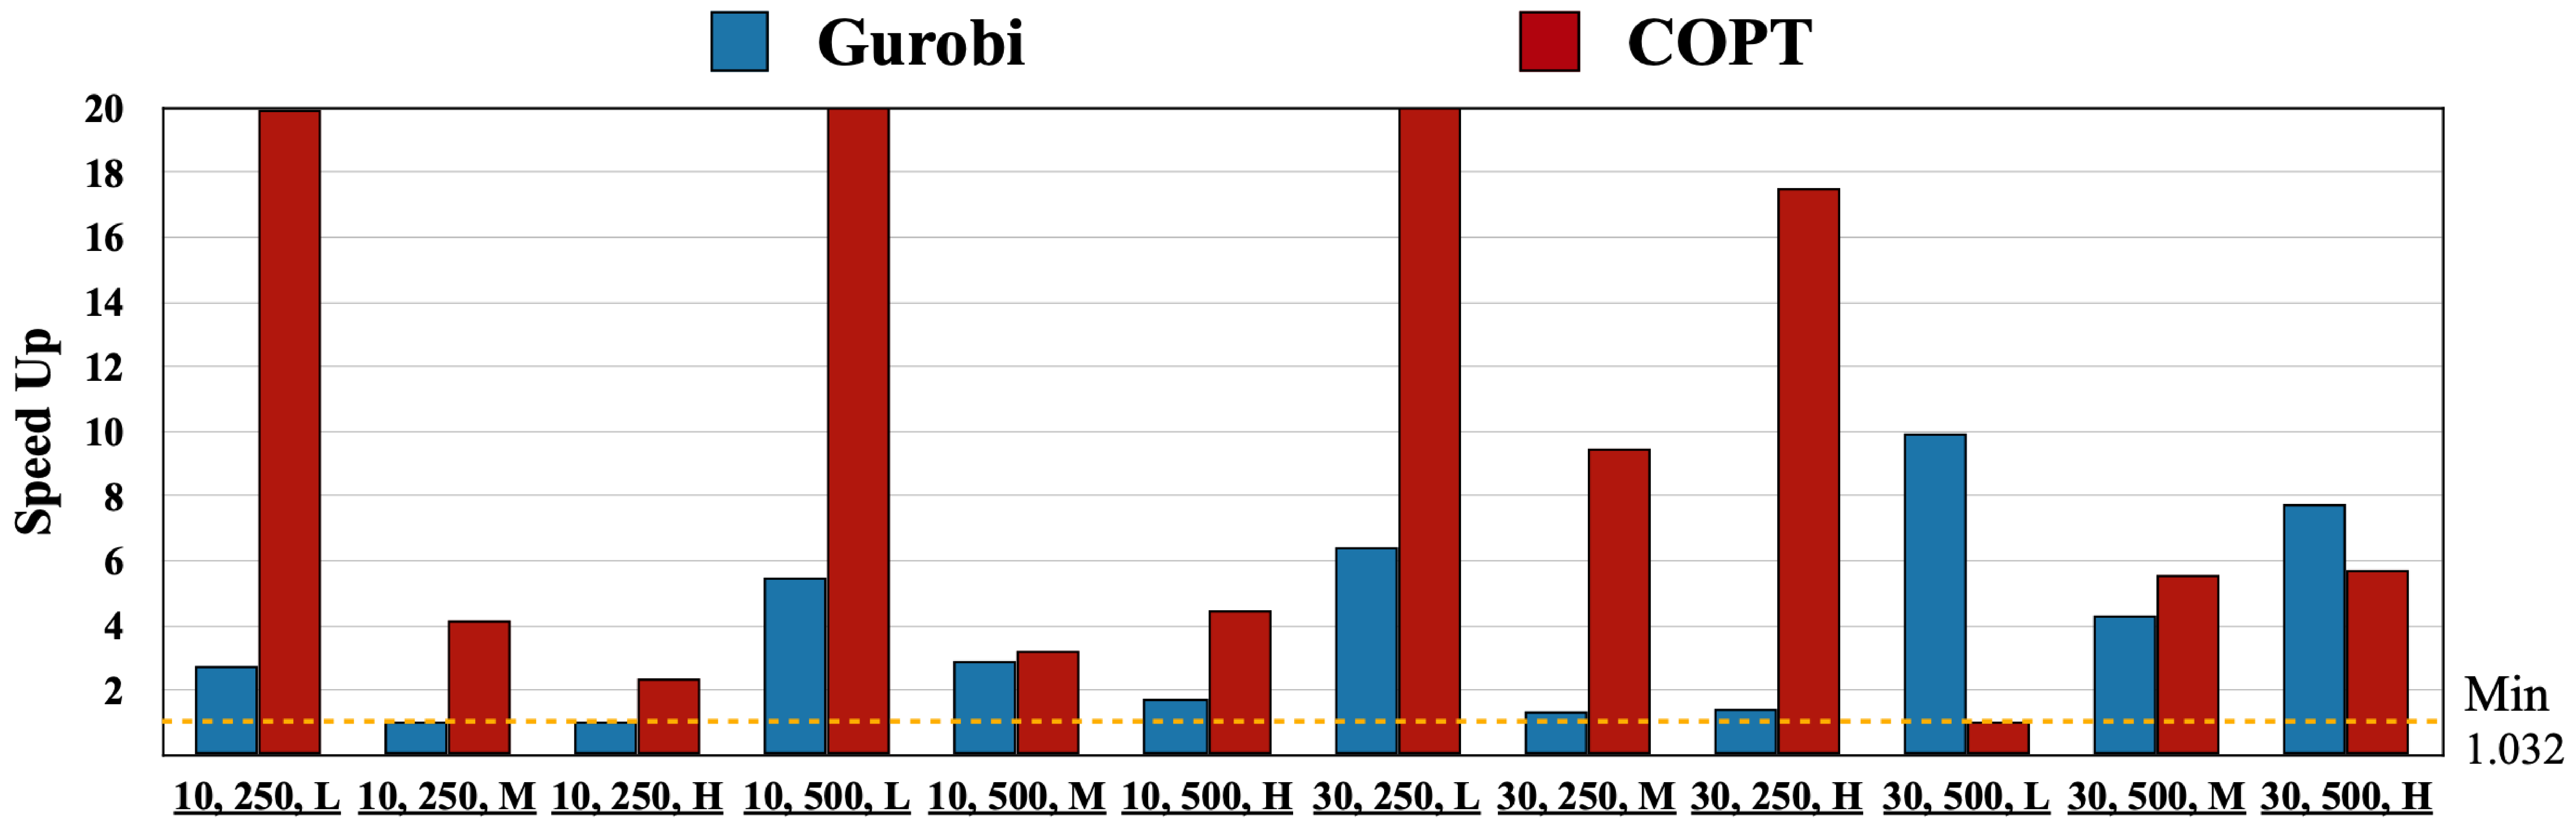

Average speedup on knapsack instances

New Technology: Numerical Experiments

- **Acceleration by two lines of code**
- **Remarkable speedup on primal solution finding**

  Both on Gurobi and COPT
- **No loss of optimality**

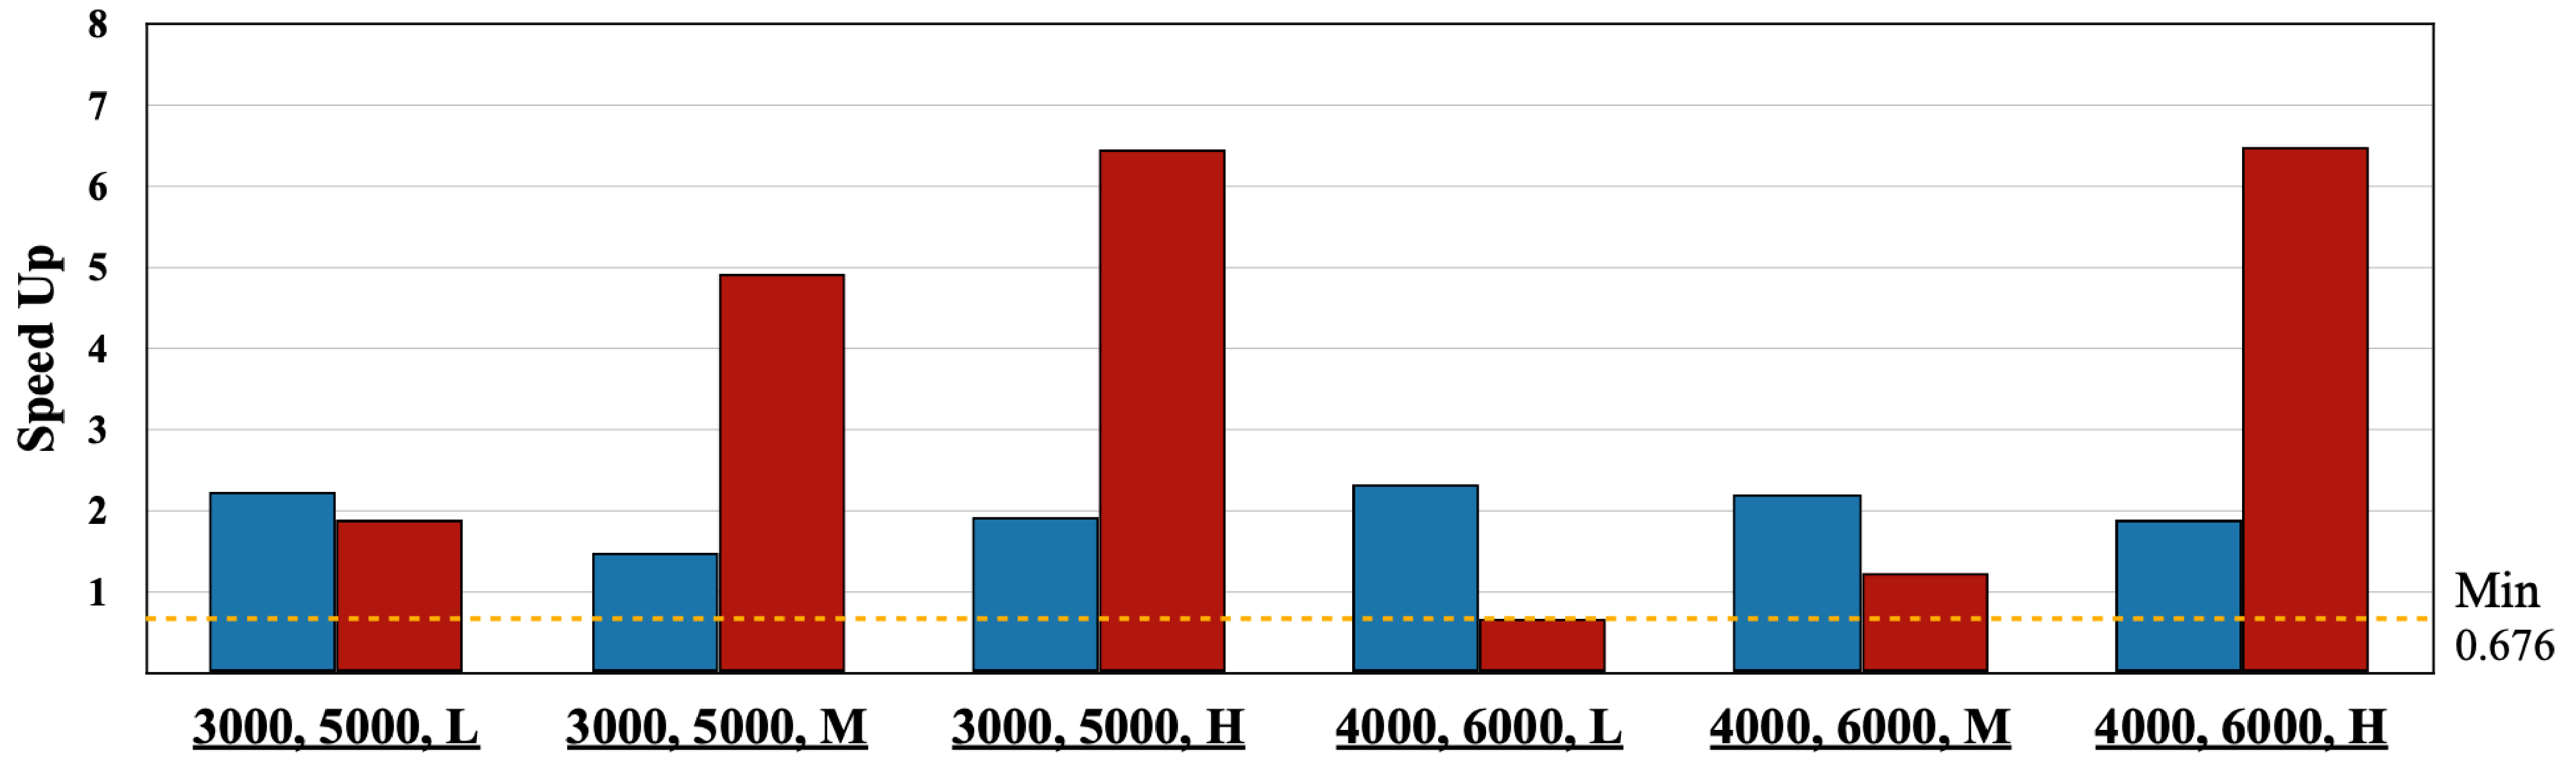

Unit Commitment

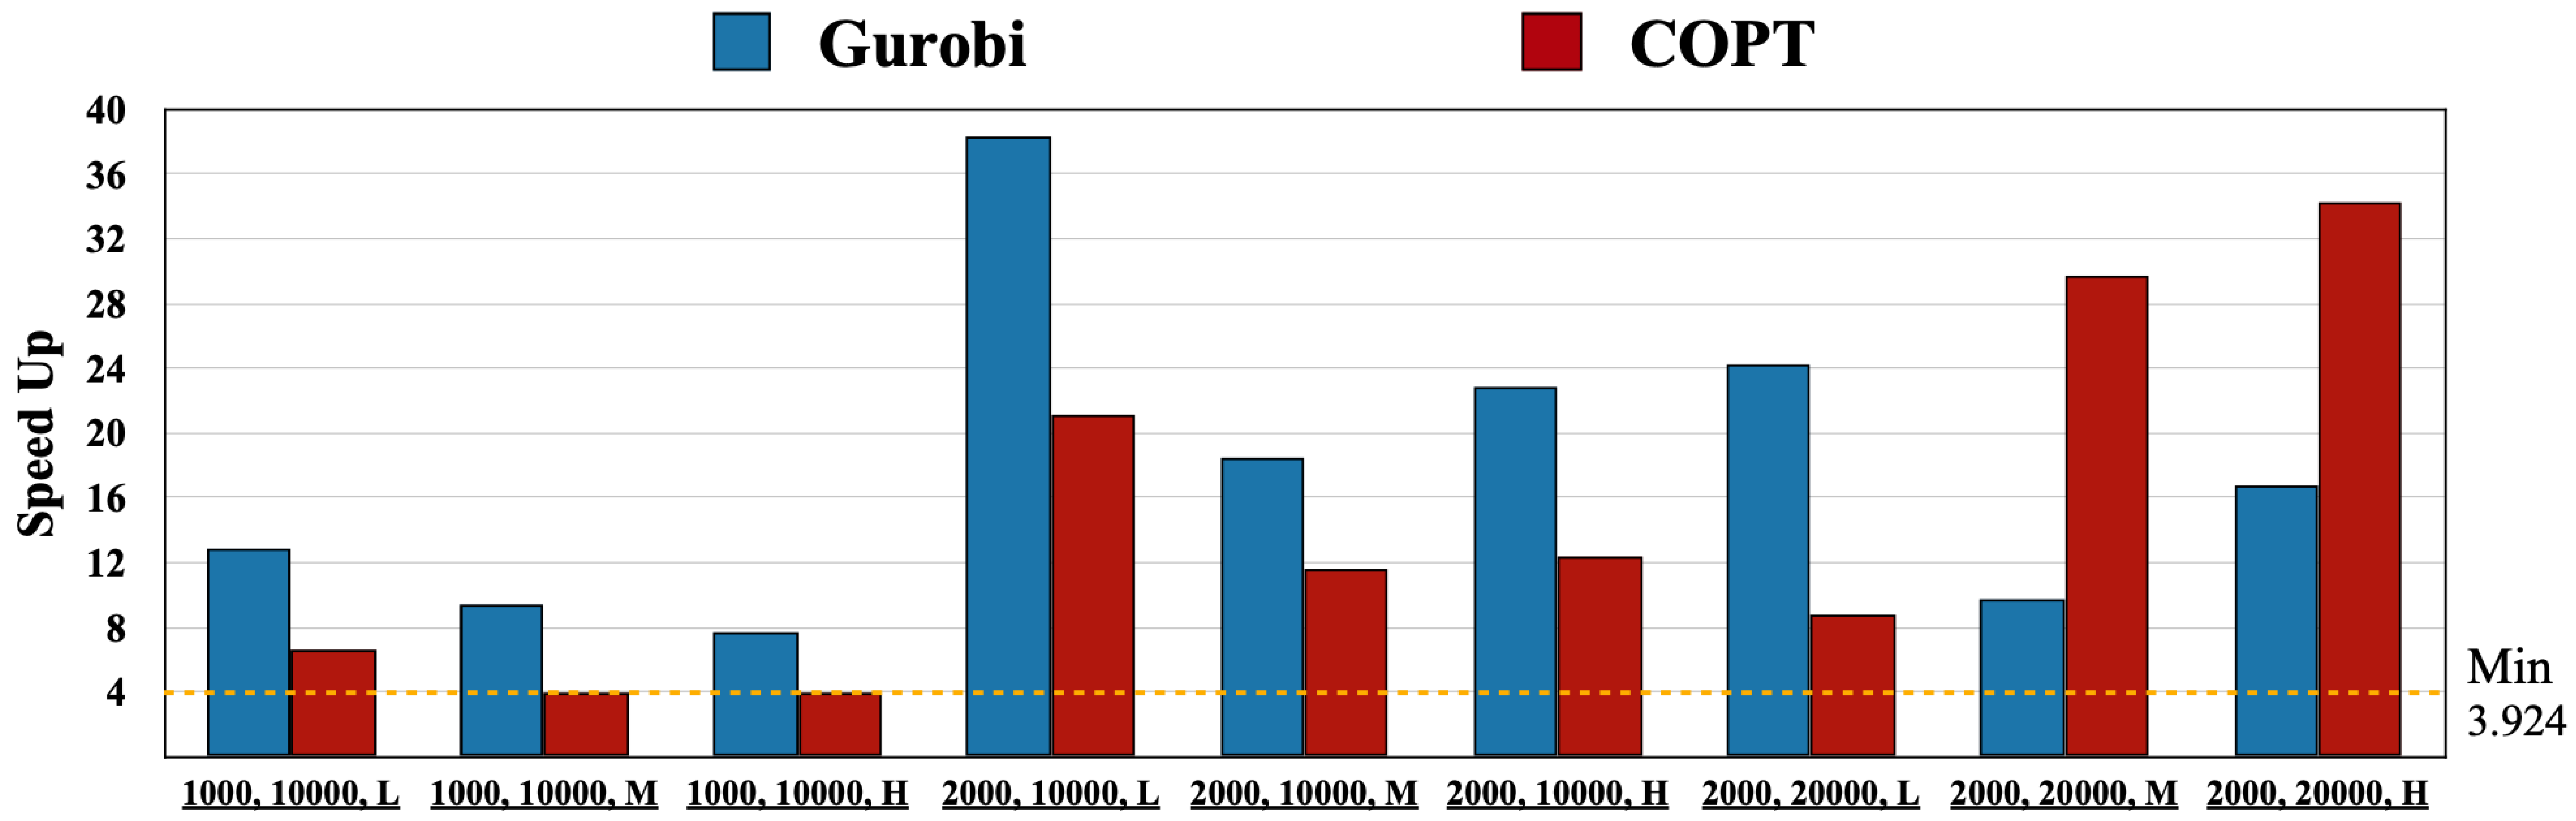

Set-Covering

Topic 4: Equitable Covering & Partition – Divide and Conquer (Carlsson et al. 2009)

Problem Statement: Divide-Conquer

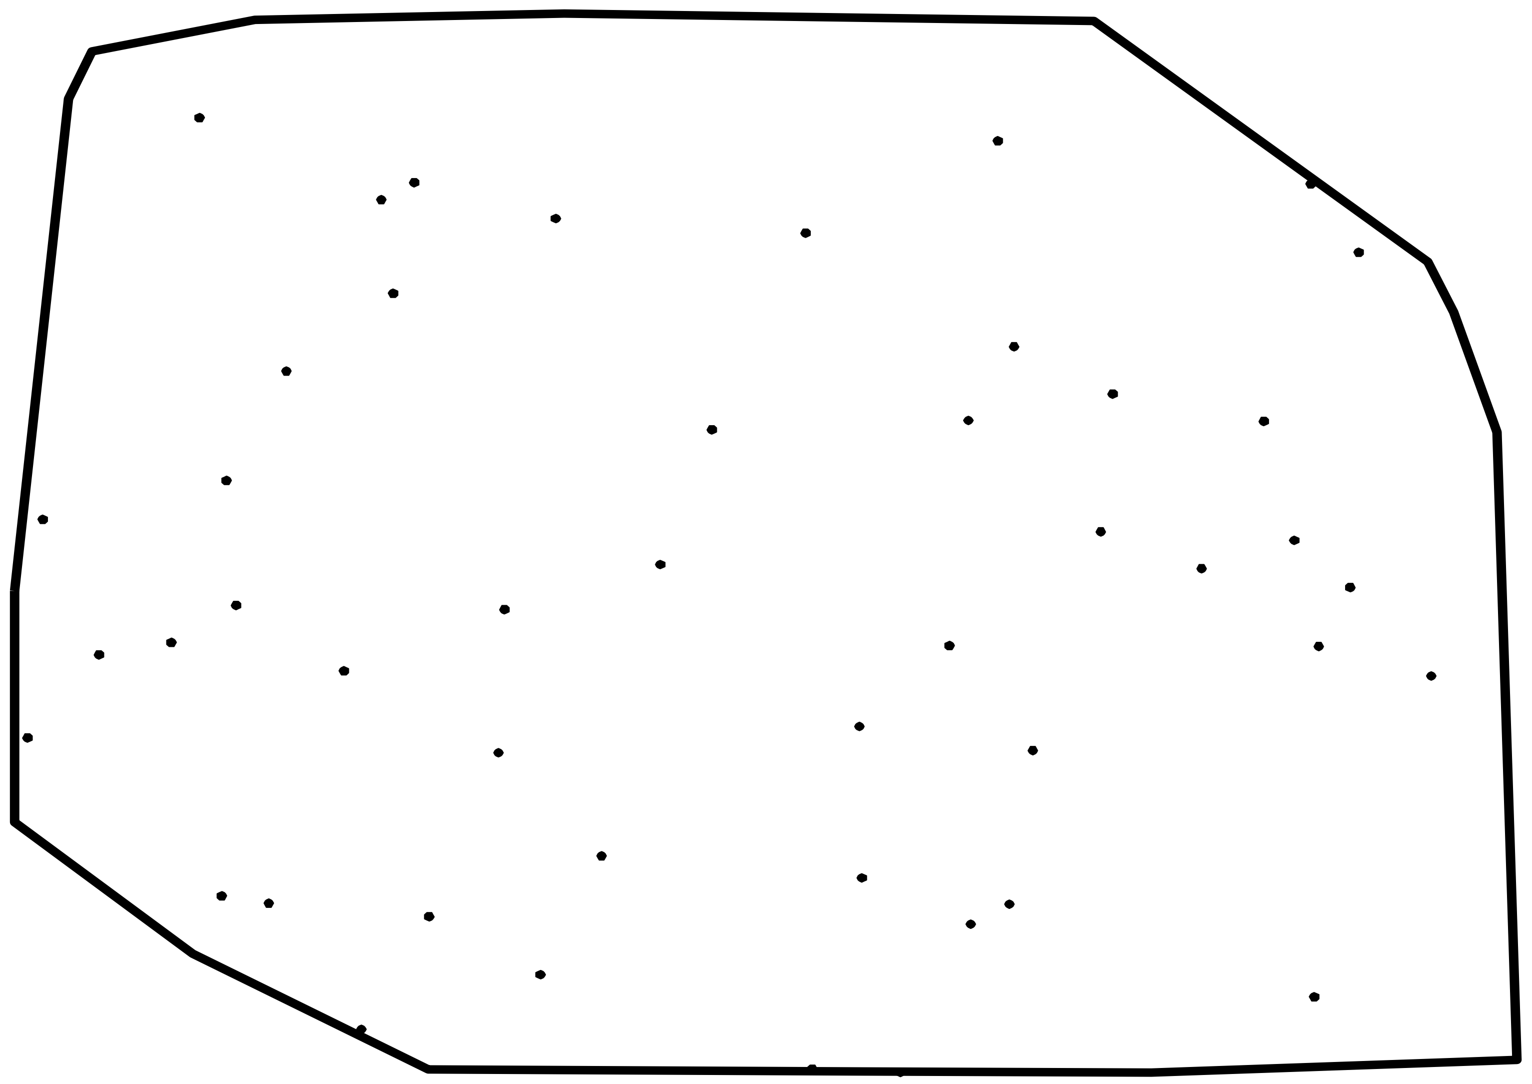

n points are scattered inside a convex polygon *P* (in 2D) with *m* vertices. Does there exist a partition of *P* into *n* sub-regions satisfying the following:

- Each sub-region is a convex polygon
- Each sub-region contains one point
- All sub-regions have equal area

Related ML Problem: Voronoi Diagram

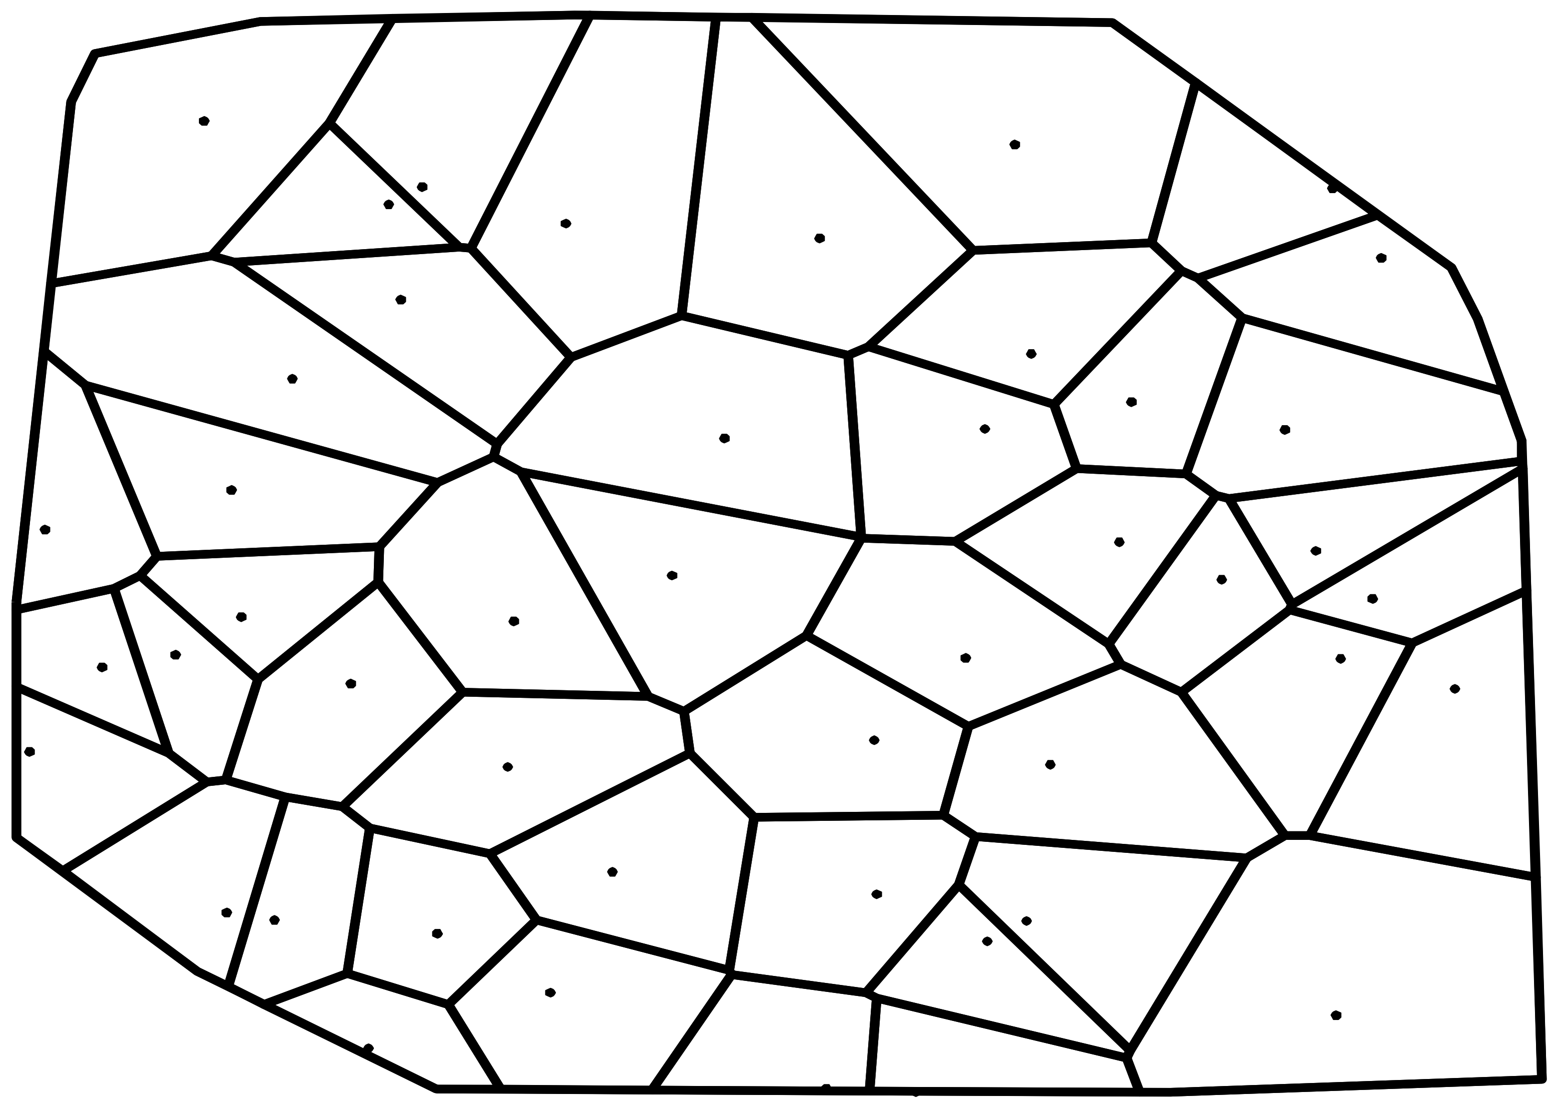

In the *Voronoi Diagram*, we satisfy the first two properties (each sub-region is convex and contains one point), but the sub-regions have different areas.

Our Theoretical Result (Carlsson et al. 2007

Not only such an equitable partition always exists, but also we can find it exactly in running time O(Nn log N), where $N = m + n$.

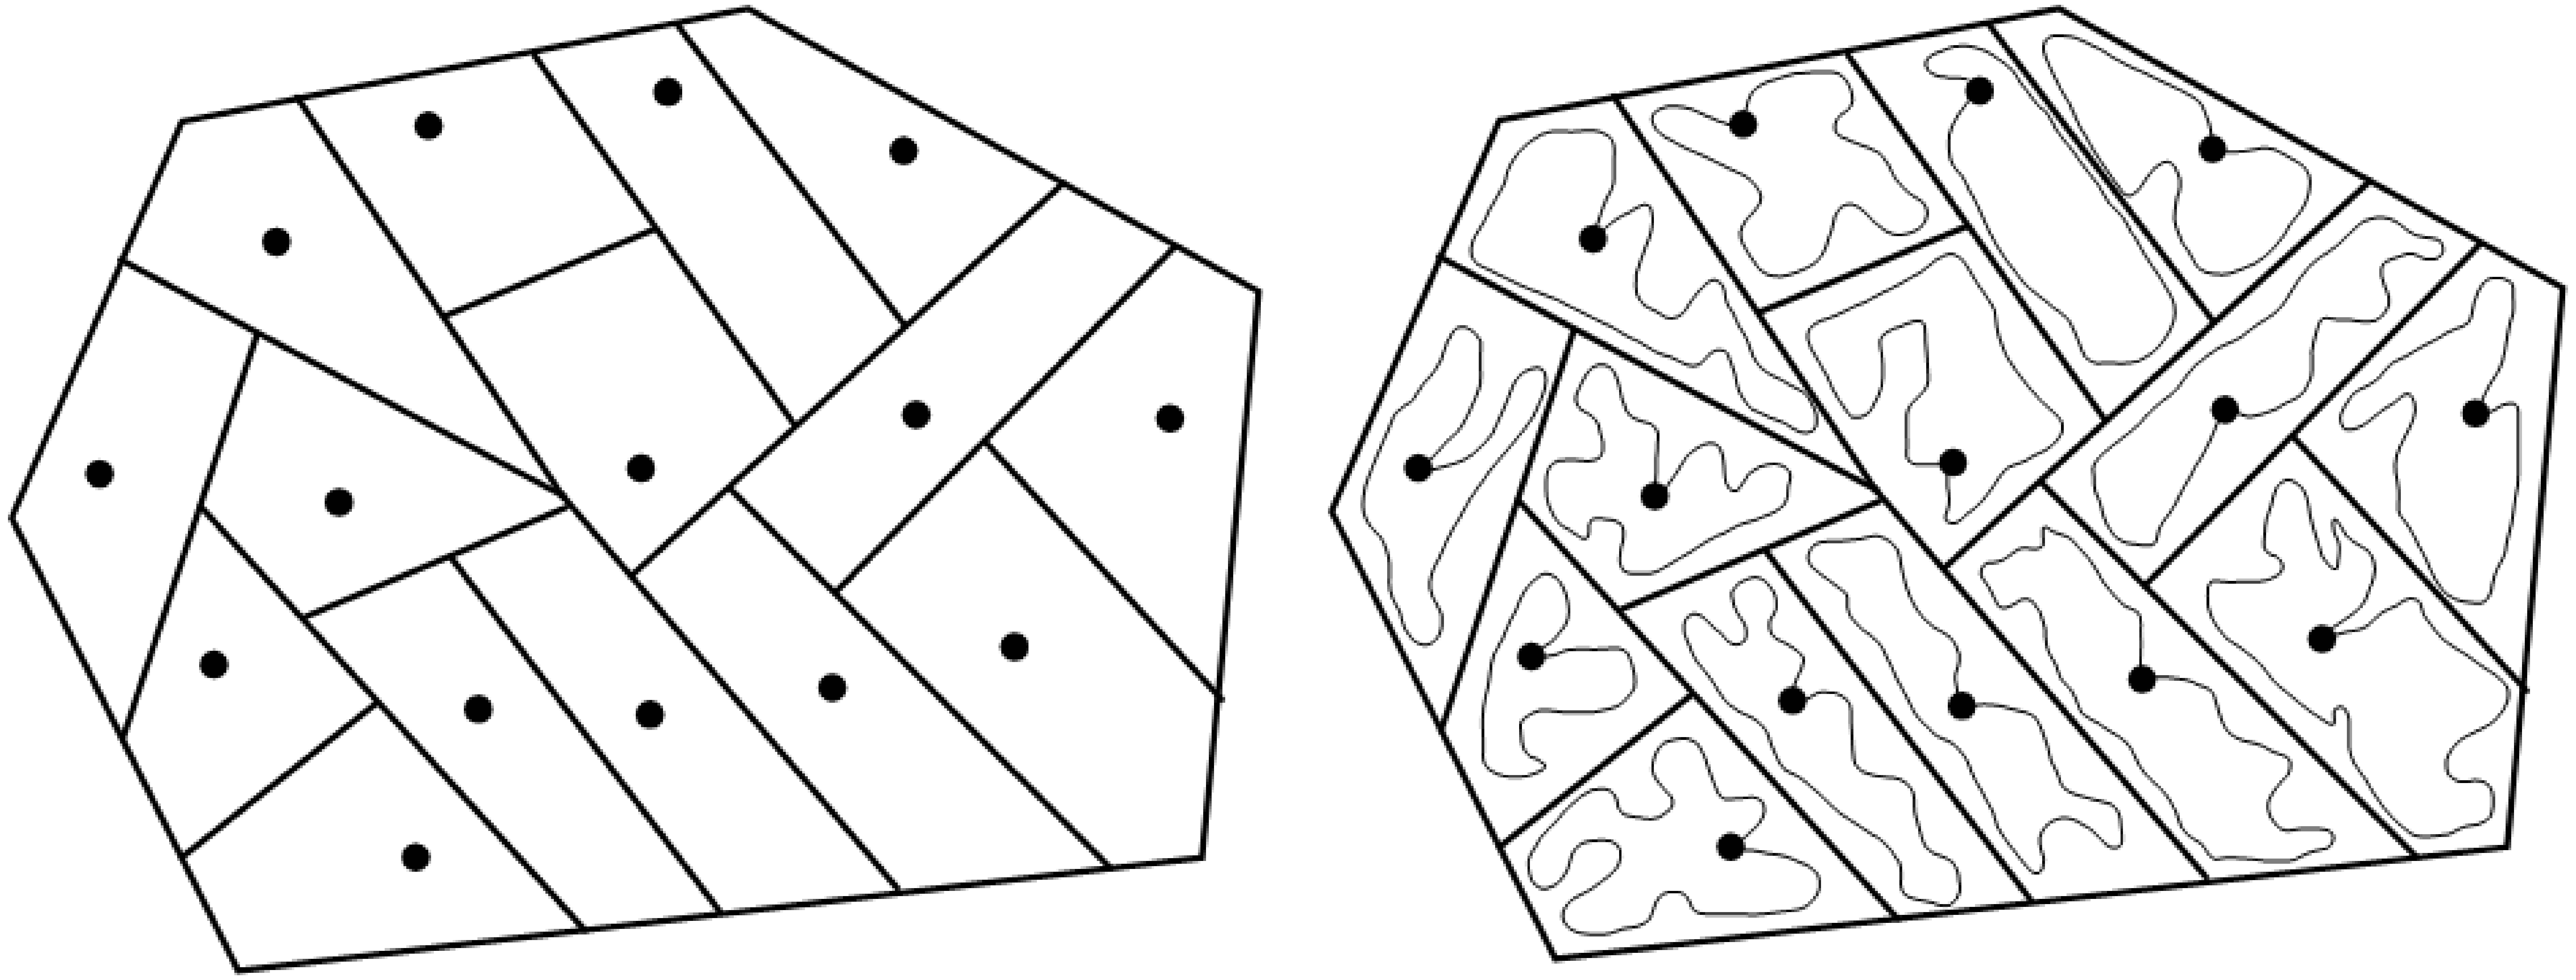

App. IX: Wireless Tower – Dynamic Resource Allocation

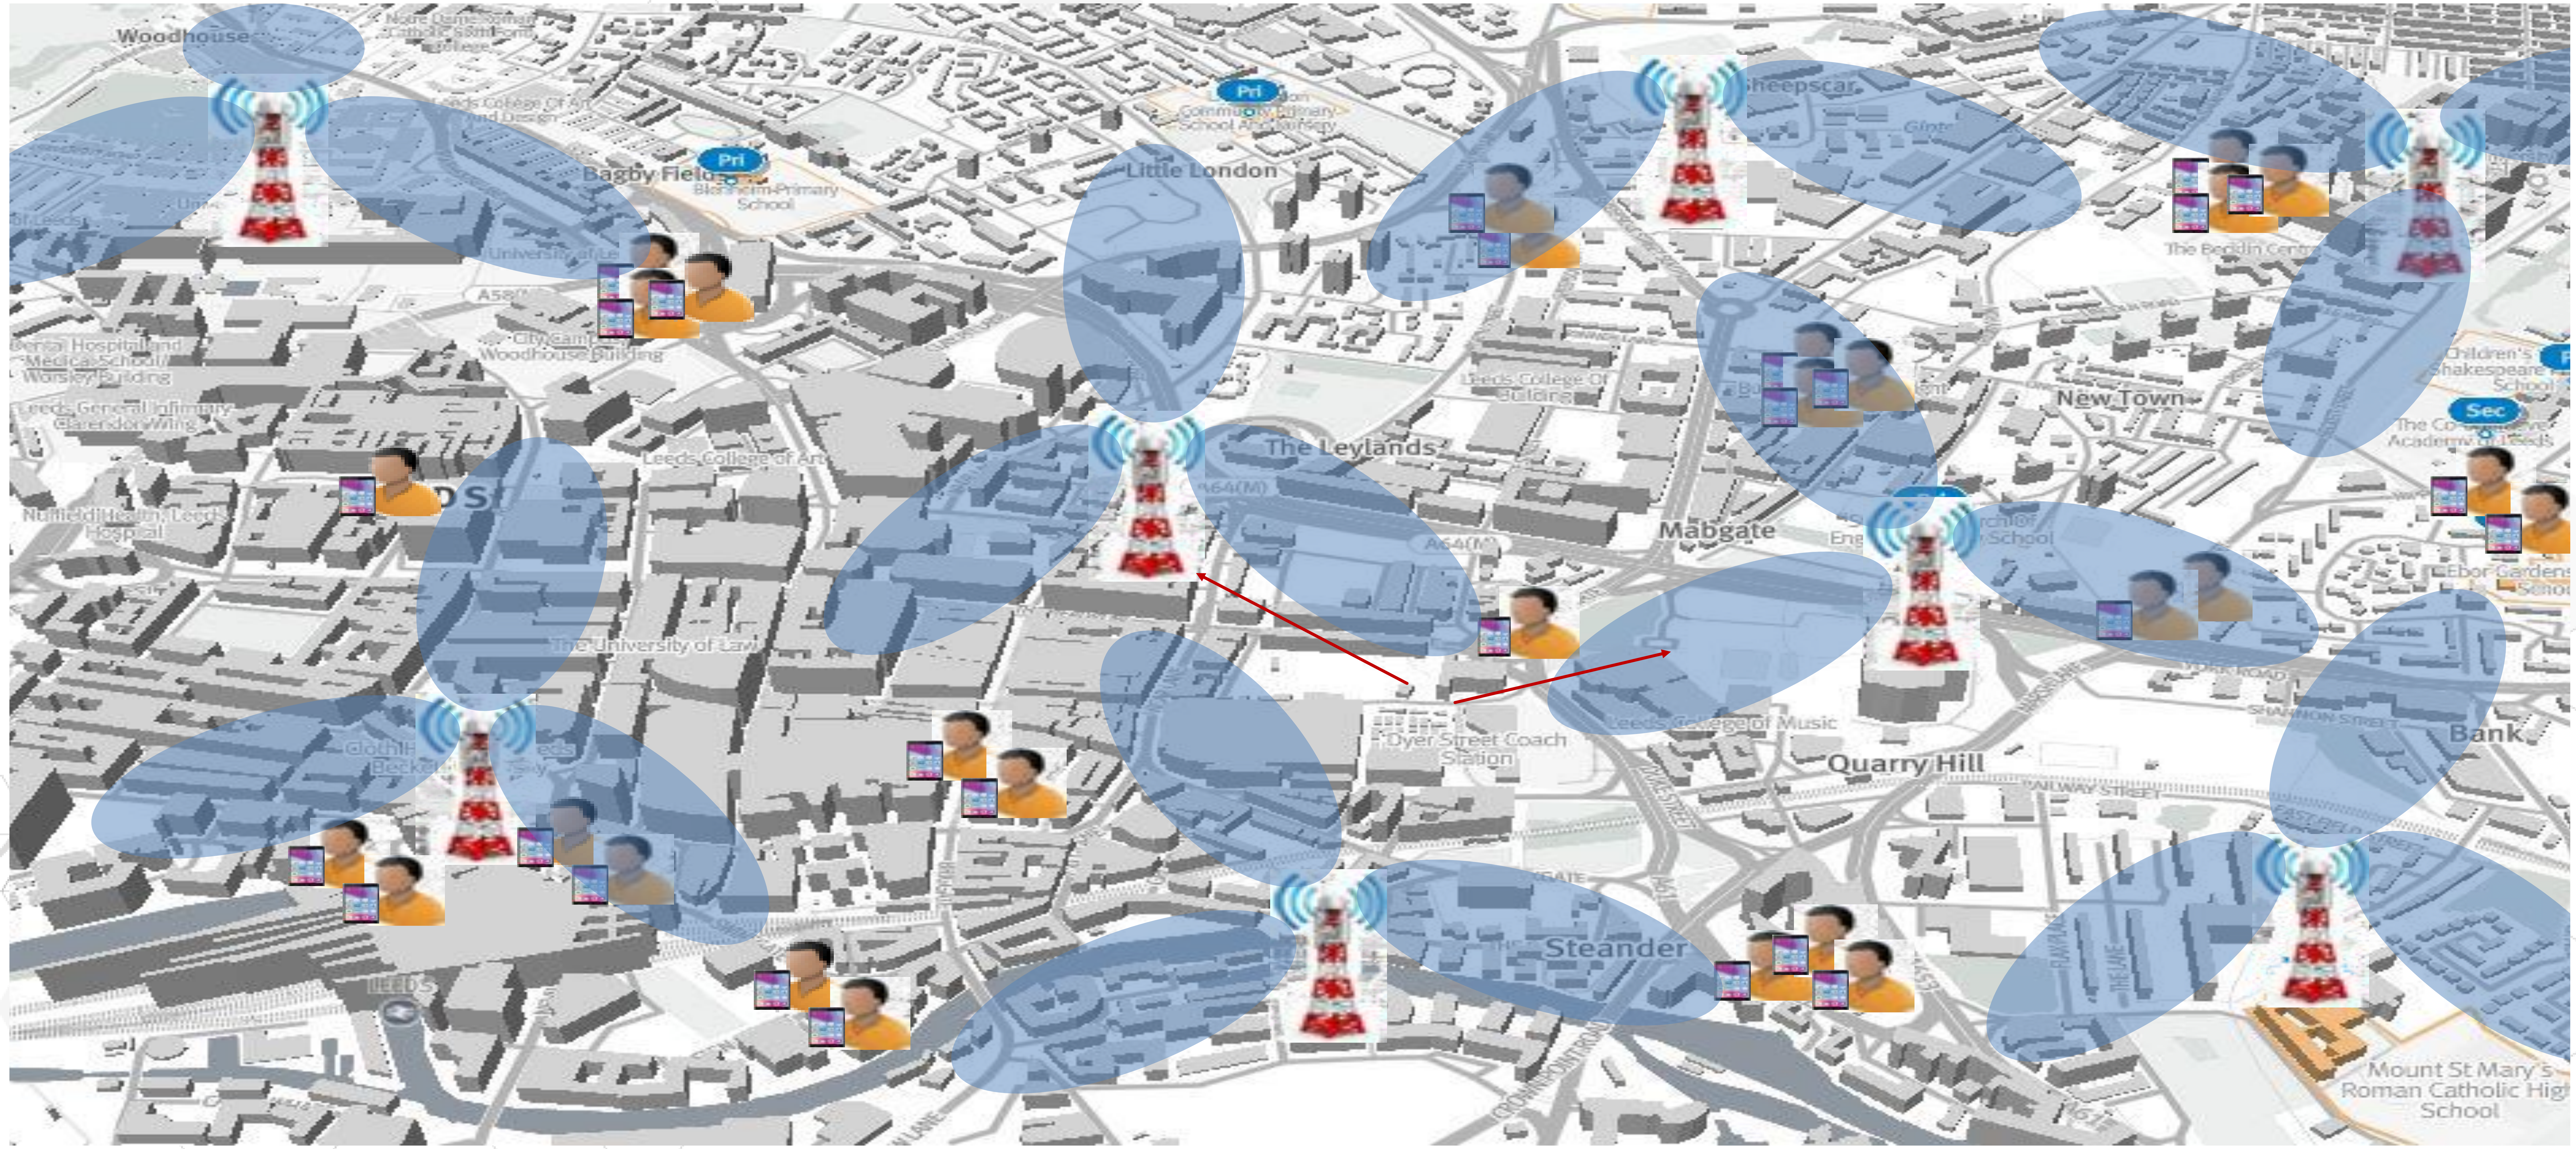

Preliminary Test Result—Effectiveness

基于真实商用网络进行模型优化效果的测试验证验证统计结果：

	小区数	时段	区域平均负载	区域平均吞吐率(Mb/S)	高负载小区负载	高负载小区吞吐率(Mb/S)
优化前	27	中午及晚共6小时	31%	5.3	68%	2.3
优化后			30%	6.12(提升15%)	66%	2.8(提升22%)
优化前		晚7时话务高峰	37%	3.9	77%	1.6
优化后			33%	5.2(提升33%)	68%	2.1(提升32%)

优化前吞吐率

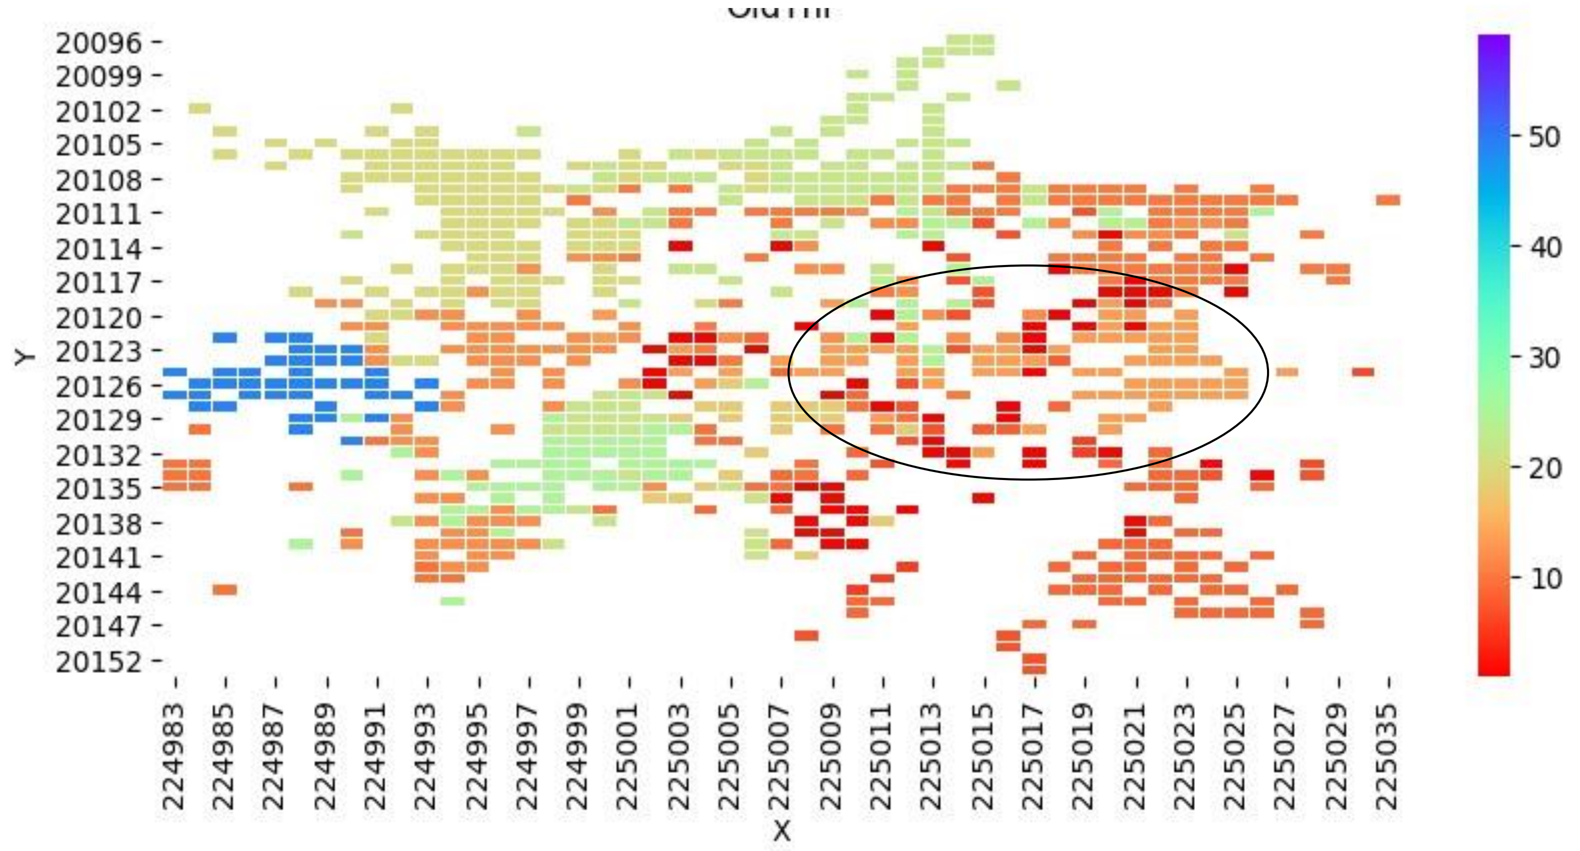

优化后吞吐率

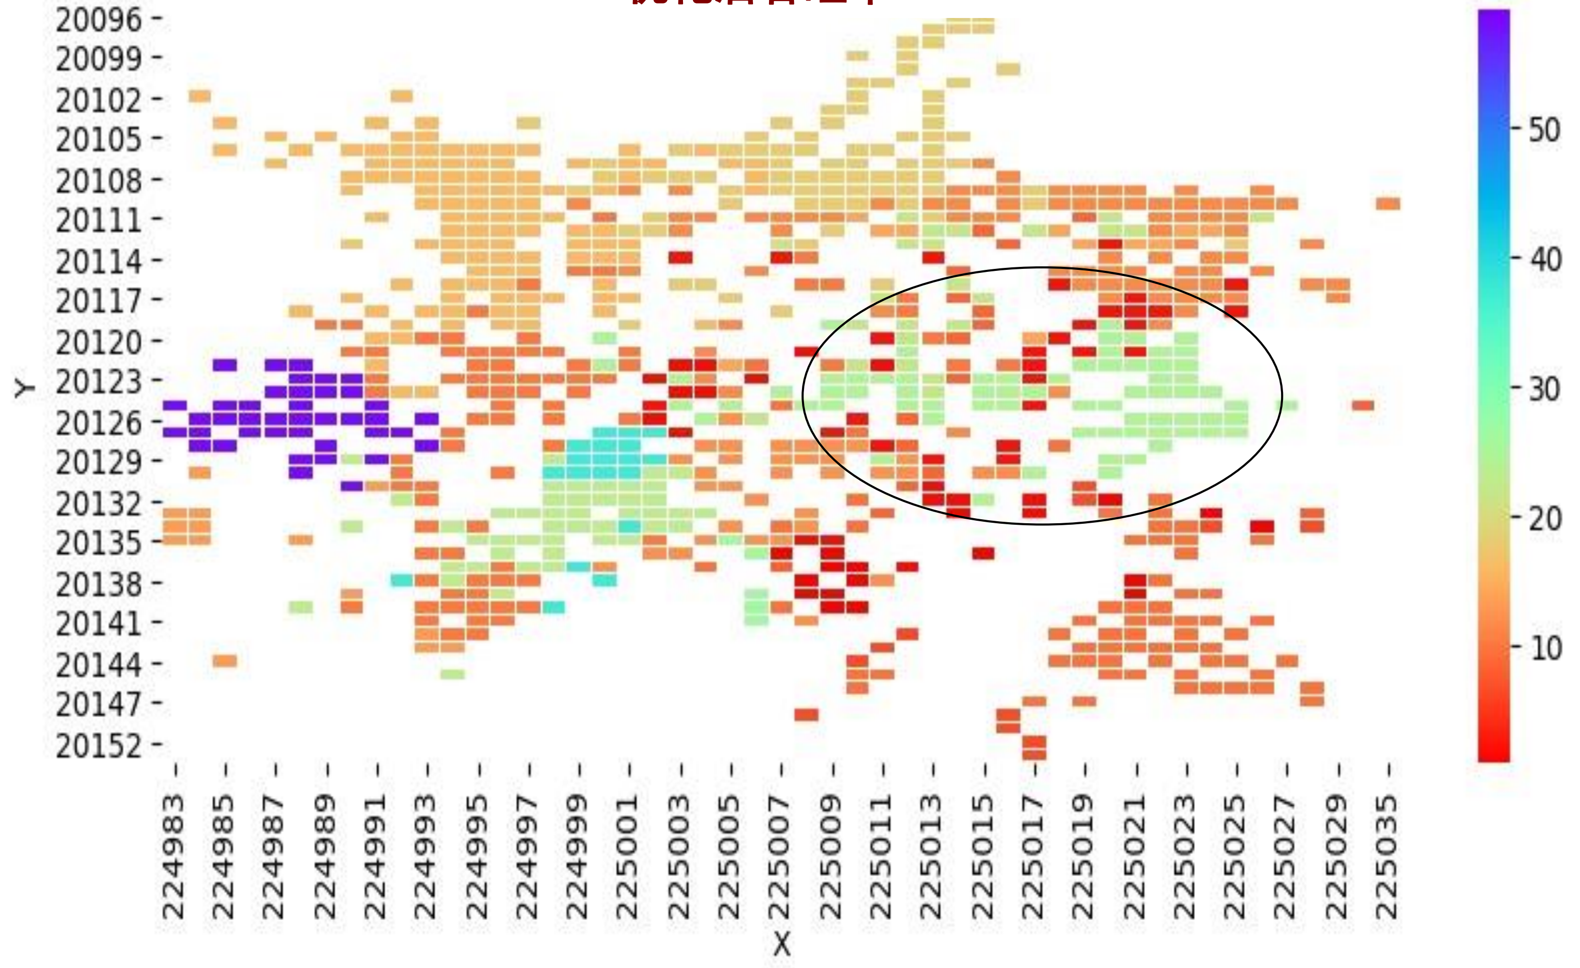

优化前负载

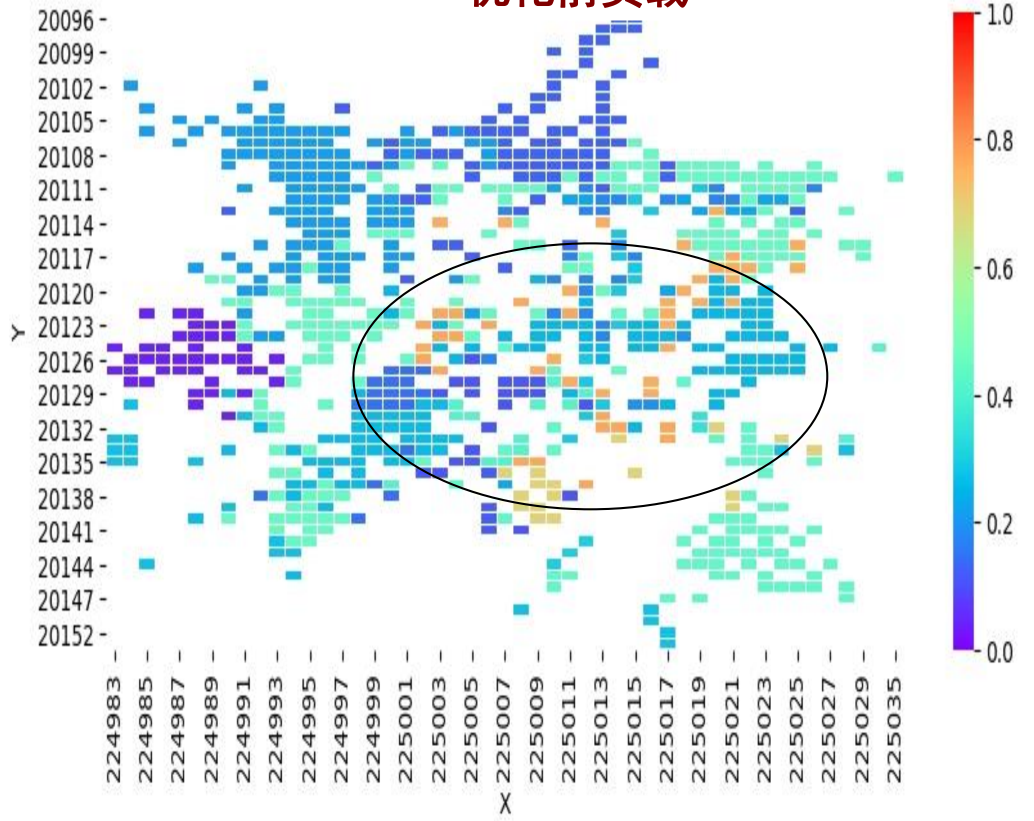

优化后负载

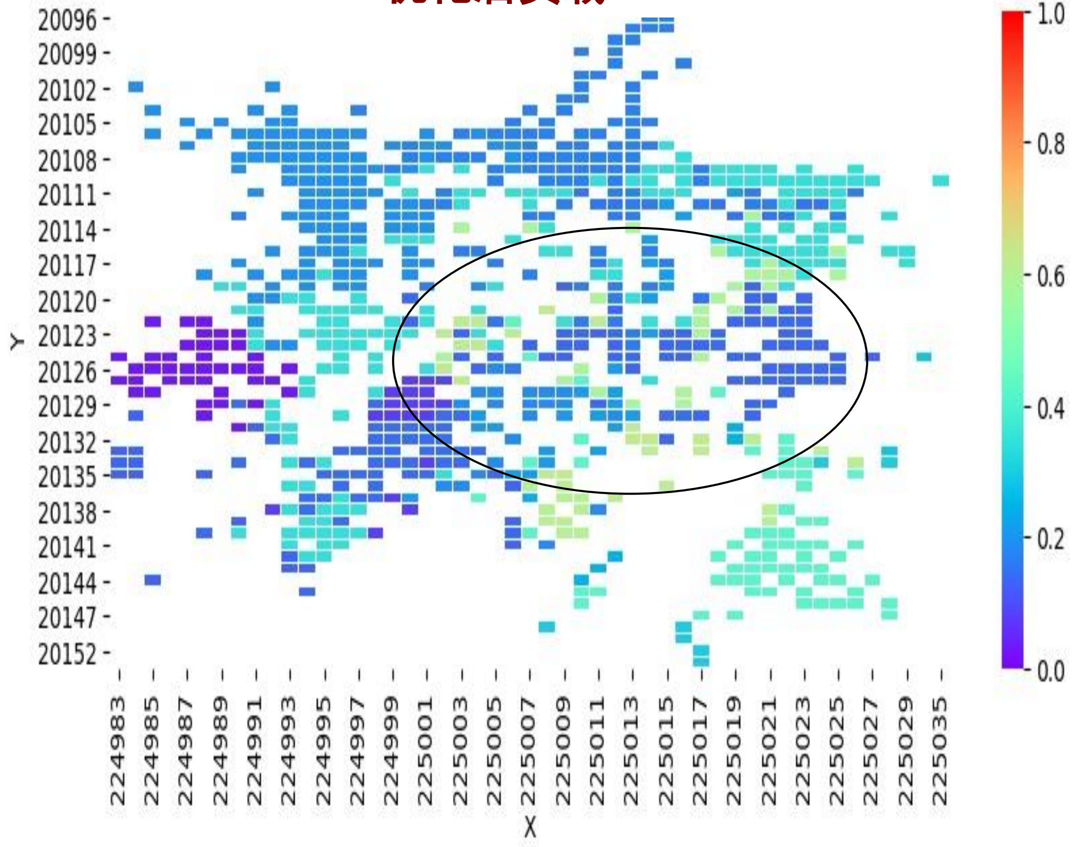

App. X: Street View Application Map-Making

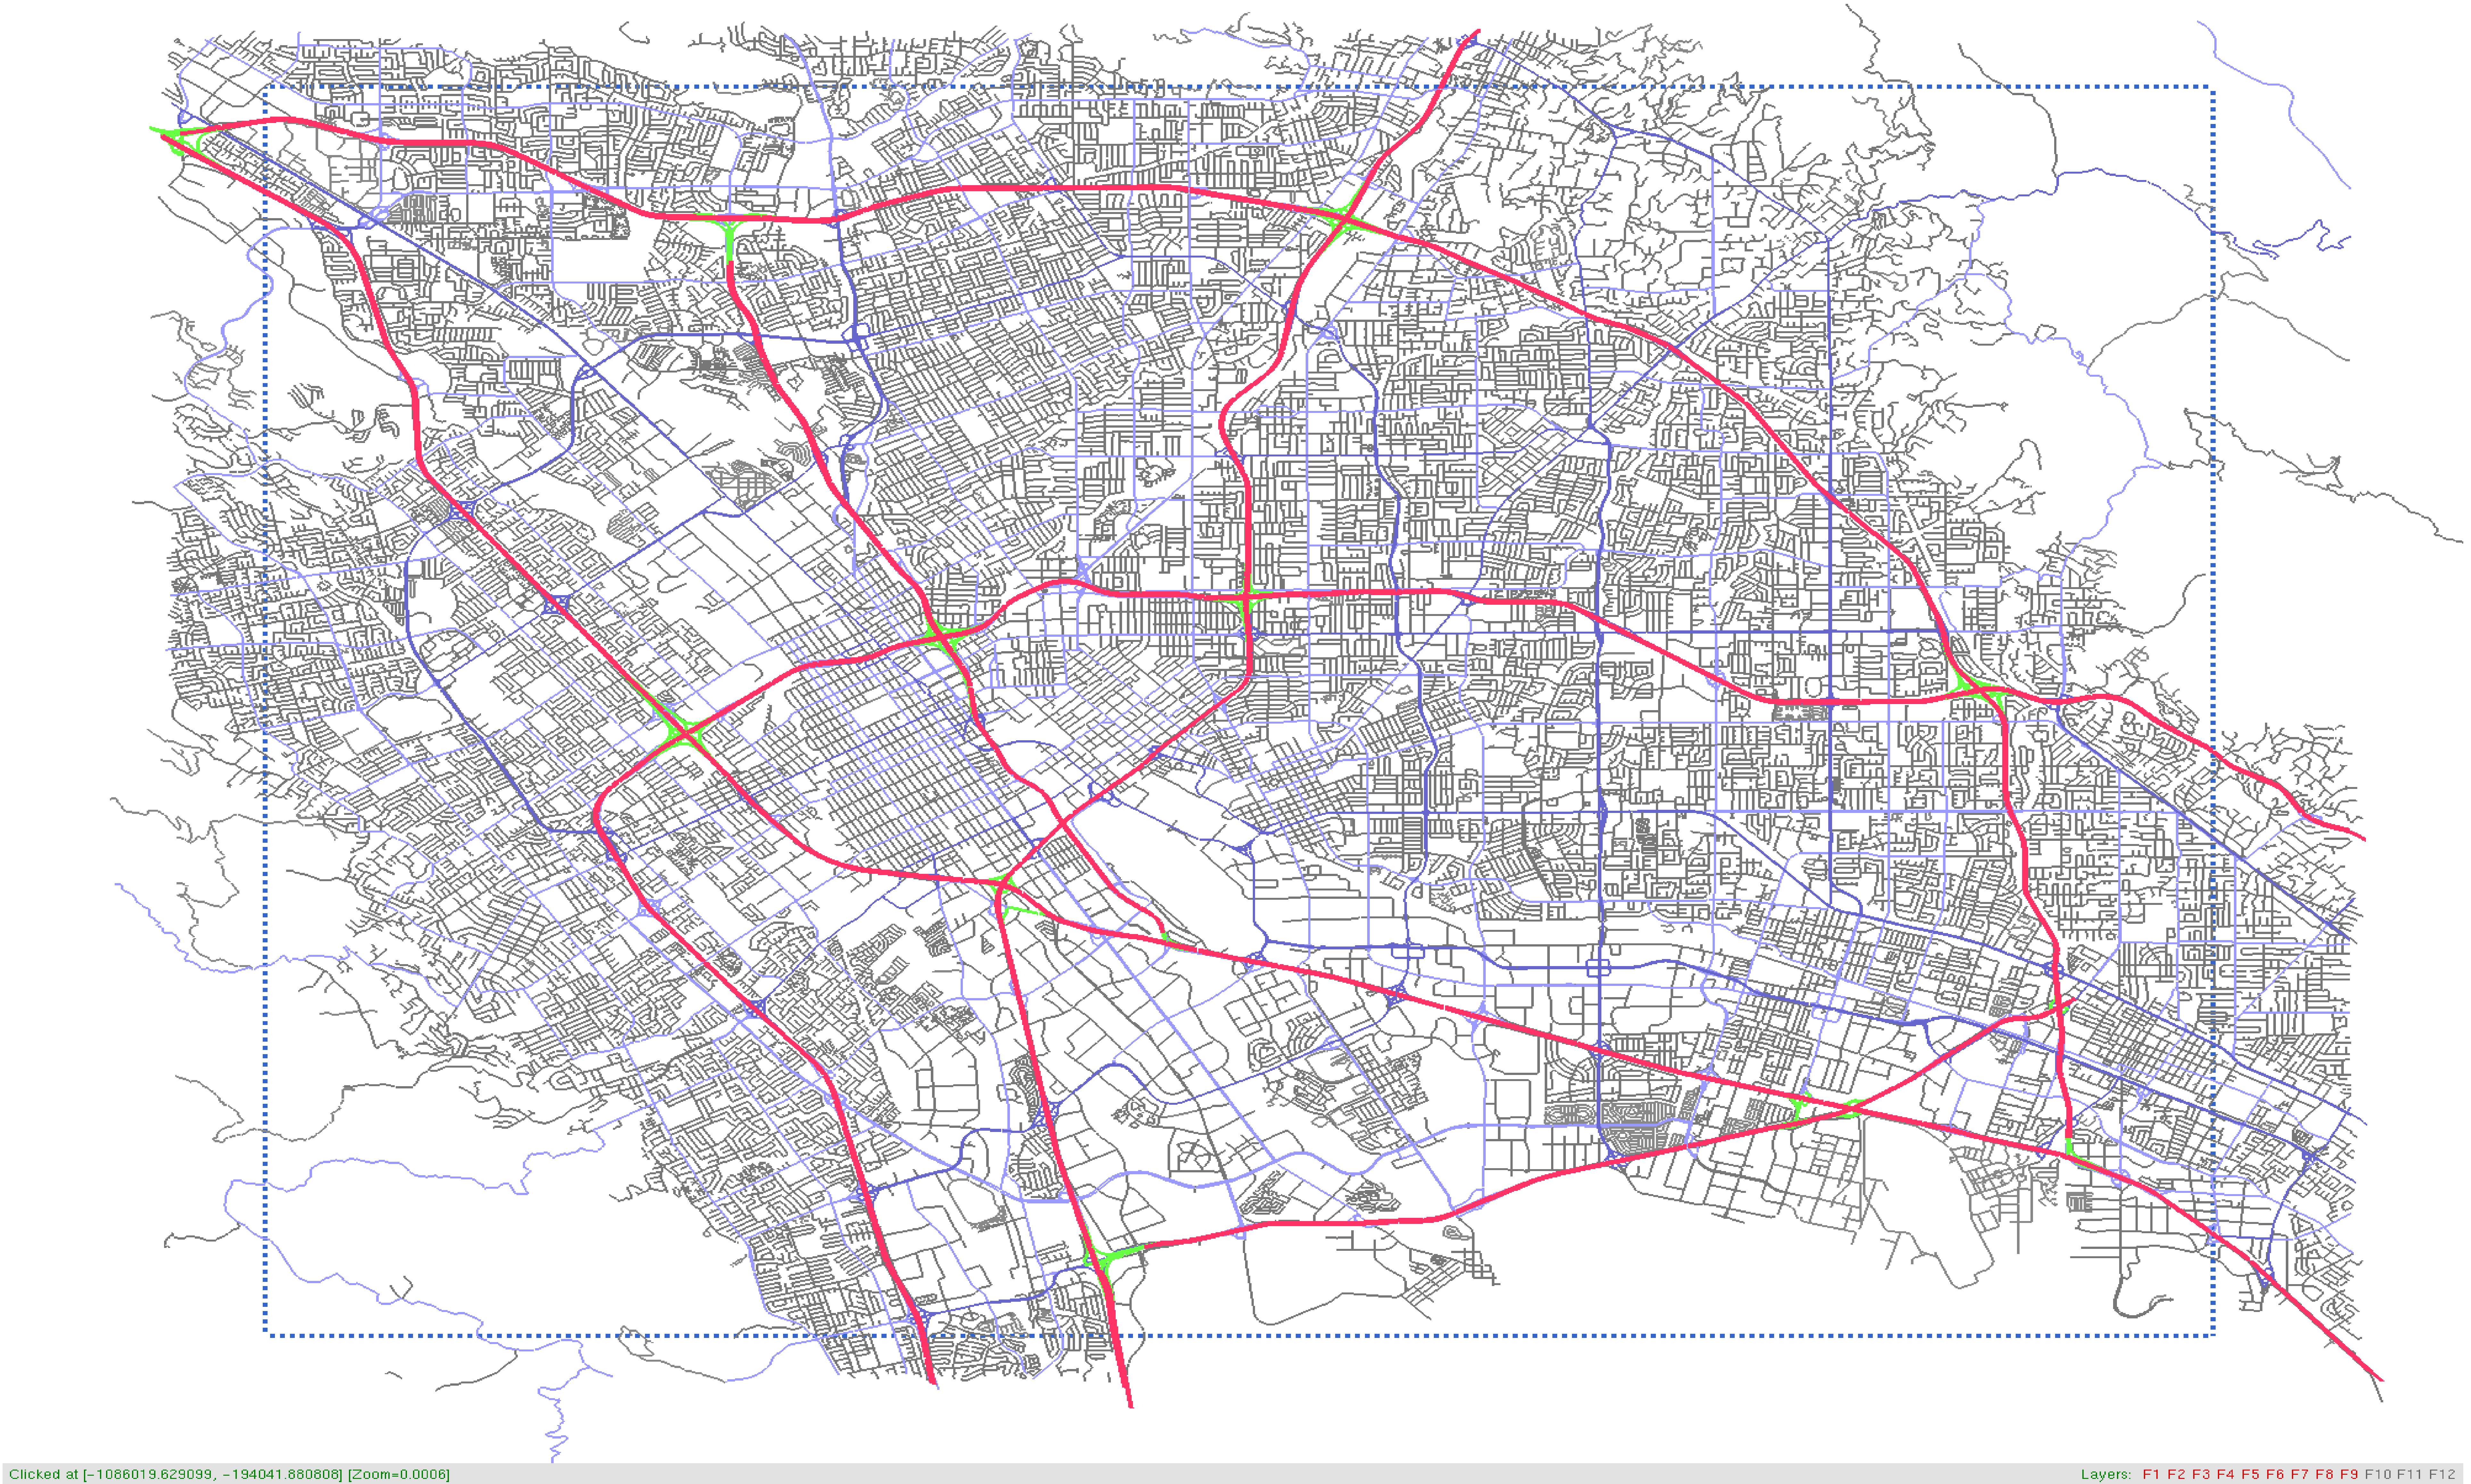

Overall Takeaways

It is possible to maker online decisions for quantitative decision models with performance guarantees close to that of the offline decision-making with perfect/complete information

Second-Order Derivative information matters and better to integrate FOM and SOM for nonlinear optimization!

Pre-training greatly improves Mixed Integer LP solvers that benefit real economy

Decomposition (Divide and Conquer) helps solving very large-scale scientific decision/computation problems

- THANK YOU