Stochastic Optimization and Decision Recommendation with Multi-sourced Information and Unknown Trust

Siqian Shen

Professor

Department of Industrial and Operations Engineering

University of Michigan at Ann Arbor

Outline

Background

Decision Maker

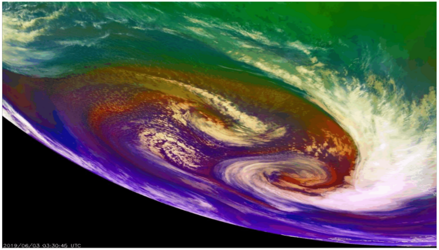

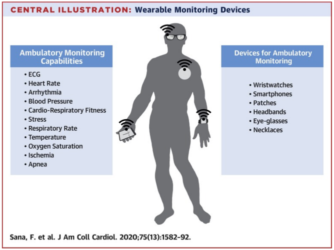

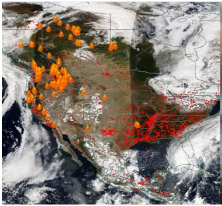

Figure: Multi-source data and information for decision-making scenarios

Figure 1-1: https://www.noaa.gov/education/resource-collections/data/real-time#data-resources; Figure 1-2: https://www.sciencedirect.com/science/article/pii/S0735109720305234?via%3Dihub; Figure 1-3 (left): https://news.wisc.edu/one-minute-data-from-uw-helps-nasa-detect-wildfires-faster/; Figure 1-3 (middle): https://www.nbcnews.com/mach/science/drones-are-fighting-wildfires-some-very-surprising-ways-ncna820966; Figure 1-3 (right): https://www.doi.gov/wildlandfire/improving-wildfire-risk-reduction-through-ecosystem-mapping

Motivation

- How should **data from heterogenous sources** be handled?
- What if data from different sources **exhibit inconsistencies**?
- If we have **varying levels of confidence in each source** based on their **past performance**, how should we approach their data?

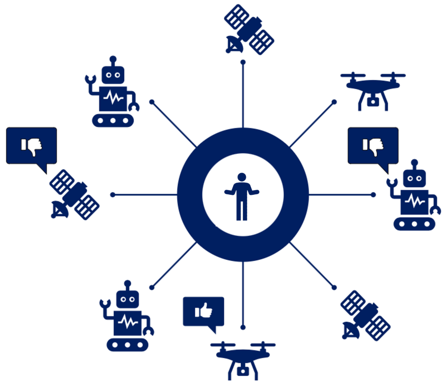

Figure: Conflicting data from multiple sources

The Notion of Trust

Lee and See (2004)[1] characterize **trust** as "**the attitude that an agent will help achieve an individual's goals in situations characterized by uncertainty and vulnerability**."

Trust as weighting factors

- **Bounded:** Self-reported trust is bounded.
- **Dynamic:** Trust can strengthen or decay due to moment-to-moment interaction with the agent.
- **Stable:** An individual's trust will stabilize over repeated interactions with the same autonomous agent.

[1] Lee JD, See KA (2004) Trust in automation: Designing for appropriate reliance. *Human Factors* 46(1):50–80.

Related Papers

This talk will cover some of our recent work on trust-informed decision-making and recommendations.

- Y. Guo, B. Zhou, R. Jiang, X. J. Yang, S. Shen, "Distributionally robust resource allocation with trust-aided parametric information fusion," in the Proceedings of the 63rd IEEE Conference on Decision and Control (CDC 2024), Milan, Italy, Dec 2024. `https://arxiv.org/pdf/2411.01428`.
- Y. Guo, R. Jiang, S. Shen,"Optimization with Multi-sourced Reference Information and Unknown Trust: A Distributionally Robust Approach," under preparation for R&R, 2025. `https://arxiv.org/abs/2501.07057`.
- Y. Guo, X. Fei, R. Jiang, X. J. Yang, S. Shen, "Learning-enhanced Route Recommendation under Unknown and Uncertain Travelers' Trust," under review, 2025.
- H. Chung, R. Jiang, S. Shen, X. J. Yang, "Crowdsourced Navigation in Mass Evacuation: A Lab Study on User Contribution," in the Proceedings of *2025 IEEE 5th International Conference on Human-Machine Systems (ICHMS)*, 2025.
- H. Chung, R. Jiang, S. Shen, X. J. Yang, "Predicting human altruistic and compliance behaviors in multiple-operator single-agent (Mosa) interaction," to appear in *International Journal of Human–Computer Interaction*, 2025.

Outline

Trust-aided Multi-reference DRO [2], [3]

Wasserstein Distributionally Robust Optimization (DRO) problem:

$$\inf_{\boldsymbol{x}\in\mathbb{X}} \left\{ \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}}_I)} \mathbb{E}_{\mathbb{Q}}[\ell(\boldsymbol{x},\boldsymbol{\xi})] \right\}$$

with **ambiguity set** $\mathbb{B}_\epsilon(\widehat{\mathbb{P}}_I) := \{\mathbb{Q}\in\mathcal{M}(\Xi): d_W(\widehat{\mathbb{P}}_I,\mathbb{Q})\le\epsilon\}$. d_W is the **Wasserstein metric** measuring **distance between two distributions.**

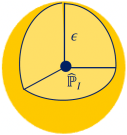

Parametric Method

$\mathbb{P}_1: \mathcal{N}(\mu_1, \sigma_1^2)$

......

$\mathbb{P}_H: \mathcal{N}(\mu_H, \sigma_H^2)$

↓

Trust $t_1, \ldots, t_H$ as weights for fusing $\mu_1, \ldots, \mu_H$ and $\sigma_1^2, \ldots, \sigma_H^2$ to obtain $\mathbb{P}_e$

→ Sample from $\mathbb{P}_e$ to obtain $\widehat{\mathbb{P}}_e$ and construct ambiguity set $\mathbb{B}_\epsilon(\widehat{\mathbb{P}}_e)$

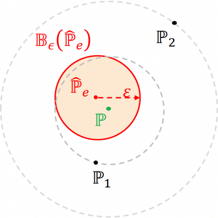

Non-parametric Method

$\widehat{\mathbb{P}}_1 \to \mathbb{P}_1$

......

$\widehat{\mathbb{P}}_H \to \mathbb{P}_H$

→ Trust $t_1, \ldots, t_H$ as weights to obtain $\widehat{\mathbb{P}}_{HI}$

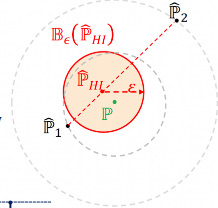

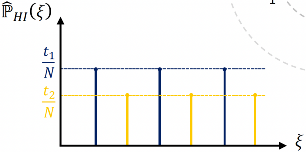

[2] Y. Guo, B. Zhou, R. Jiang, X. J. Yang, S. Shen, "Distributionally robust resource allocation with trust-aided parametric information fusion," in the Proceedings of the 63rd IEEE Conference on Decision and Control (CDC 2024), Milan, Italy, Dec 2024. https://arxiv.org/pdf/2411.01428.

[3] Y. Guo, R. Jiang, S. Shen, "Optimization with Multi-sourced Reference Information and Unknown Trust: A Distributionally Robust Approach," under R&R, 2025. https://arxiv.org/abs/2501.07057.

Multi-reference Distributionally Robust Model (MR-DRO)

A single-stage stochastic program for our problem is stated as:

$$\inf_{\mathbf{x}\in\mathbb{X}} \mathbb{E}_{\mathbb{P}}[\ell(\mathbf{x},\boldsymbol{\xi})]. \quad (1)$$

However, we cannot solve (1) directly because of the unknown true distribution $\mathbb{P}$. As a remedy, we formulate the problem as an MR-DRO problem:

$$\inf_{\mathbf{x}\in\mathbb{X}} \left\{ \sup_{\mathbb{Q}\in\mathcal{B}_\epsilon(\widehat{\mathbb{P}}_{HI})} \mathbb{E}_{\mathbb{Q}}[\ell(\mathbf{x},\boldsymbol{\xi})] \right\}, \quad (2)$$

where the definition of the ambiguity set $\mathcal{B}_\epsilon(\widehat{\mathbb{P}}_{HI})$ is given by:

$$\mathcal{B}_\epsilon(\widehat{\mathbb{P}}_{HI}) := \left\{ \mathbb{Q} \in \mathcal{M}(\Xi) : d_W(\widehat{\mathbb{P}}_{HI},\mathbb{Q}) \le \epsilon \right\}.$$

Here, d_W is the Wasserstein metric defined via:

$$d_W(\widehat{\mathbb{P}}_{HI},\mathbb{Q}) = \inf \left\{ \int_{\Xi^2} \|\boldsymbol{\xi}-\boldsymbol{\xi}'\| \Pi(d\boldsymbol{\xi}, d\boldsymbol{\xi}') \right\},$$

in which Π is a joint distribution of $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ with marginals $\widehat{\mathbb{P}}_{HI}$ and $\mathbb{Q}$, respectively. $\|\cdot\|$ represents an arbitrary norm on $\mathbb{R}^M$.

Convex Reformulation

Under a non-parametric ambiguity set, we can reformulate the MR-DRO model (2) as follows: [4]:

$$
\begin{aligned}
\inf_{\mathbf{x},\lambda,s_{hi},\gamma_{hij}} \quad & \lambda\epsilon + \sum_{h=1}^{H}\sum_{i=1}^{I} \frac{t_h^{(I)}}{I} s_{hi} && \text{(3a)}\\
\text{s.t.} \quad & \mathbf{x} \in \mathbb{X}, && \text{(3b)}\\
& b_j + \langle \mathbf{a}_j, \hat{\boldsymbol{\xi}}_h^{(i)} \rangle + \langle \gamma_{hij}, \mathbf{g} - \mathbf{C}\hat{\boldsymbol{\xi}}_h^{(i)} \rangle \le s_{hi}, \quad \forall h \in [H], i \in [I], j \in [J], && \text{(3c)}\\
& \|\mathbf{C}^T \gamma_{hij} - \mathbf{a}_j\|_* \le \lambda, \quad \forall h \in [H], i \in [I], j \in [J], && \text{(3d)}\\
& \gamma_{hij} \ge 0, \quad \forall h \in [H], i \in [I], j \in [J]. && \text{(3e)}
\end{aligned}
$$

where

- λ, s_{hi}, γ_{hij}: Auxiliary variables introduced in the reformulation process.
- $\mathbf{C}$, $\mathbf{g}$: $\mathbf{C}$ is a matrix and $\mathbf{g}$ is a vector of appropriate dimension that appear in the assumption of the uncertainty set as a polytope: $\Xi = \{\boldsymbol{\xi} \in \mathbb{R}^M : \mathbf{C}\boldsymbol{\xi} \le \mathbf{g}\}$.
- $\mathbf{a}_j$, b_j: For every fixed $\mathbf{x} \in \mathbb{X}$, each elementary function $\ell_j(\mathbf{x}, \boldsymbol{\xi}) = \langle \mathbf{a}_j, \boldsymbol{\xi} \rangle + b_j$ in the objective function is an affine function with regard to $\boldsymbol{\xi}$.

[4] Proof follows Mohajerin Esfahani P, Kuhn D (2018) Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* 171(1):115–166.

One Iteration of Trust Update Process

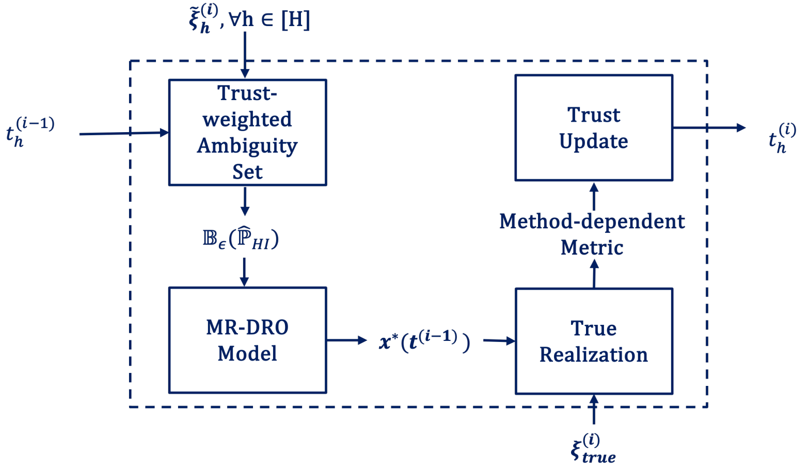

Trust Update Based on Loss

$$L^{(i)} = \ell\left(\boldsymbol{x}^*\left(\boldsymbol{t}^{(i-1)}\right), \boldsymbol{\xi}_{\text{true}}^{(i)}\right)$$

We then update each $t_h^{(i-1)}$ for all $h \in [H]$ based on the partial derivative:

$$t_h^{(i)} = t_h^{(i-1)} - w \times \frac{\partial L^{(i)}}{\partial t_h^{(i-1)}},$$

where w is a small step size.

$\frac{\partial L^{(i)}}{\partial t_h^{(i-1)}} < 0 \longrightarrow$ an increase in $t_h^{(i)}$

$\frac{\partial L^{(i)}}{\partial t_h^{(i-1)}} < 0 \longrightarrow$ keep $t_h^{(i)}$ unchanged

$\frac{\partial L^{(i)}}{\partial t_h^{(i-1)}} < 0 \longrightarrow$ a decrease in $t_h^{(i)}$

Trust Update Based on Prediction Error

$$t_h^{(i)} = \begin{cases} t_h^{(i-1)} + \Delta t, & \text{if } h = \underset{h\in[H]}{\operatorname{argmin}} \parallel \Delta \boldsymbol{\xi}_h^{(i)} \parallel, \\ t_h^{(i-1)} - \Delta t, & \text{if } h = \underset{h\in[H]}{\operatorname{argmax}} \parallel \Delta \boldsymbol{\xi}_h^{(i)} \parallel, \\ t_h^{(i-1)}, & \text{otherwise.} \end{cases}$$

Example error-based update methods:

- (i) Min-max error trust update;
- (ii) Exponential error trust update;
- (iii) Variable-share error trust update.

$$\Delta \boldsymbol{\xi}_h^{(i)} = \tilde{\boldsymbol{\xi}}_h^{(i)} - \boldsymbol{\xi}_{\text{true}}^{(i)}, \qquad i \in [I], h \in [H]$$

$$\hat{\boldsymbol{\xi}}_h^{(i)} = \tilde{\boldsymbol{\xi}}_h^{(I+1)} - \Delta \boldsymbol{\xi}_h^{(i)}, \qquad i \in [I], h \in [H]$$

$$\hat{\boldsymbol{\xi}}_h = [\hat{\boldsymbol{\xi}}_h^{(1)}, \dots, \hat{\boldsymbol{\xi}}_h^{(I)}]^{\mathrm{T}}, \qquad h \in [H]$$

$$\widehat{\mathbb{P}}_{HI}(\boldsymbol{\xi}) := \sum_{h=1}^{H} \frac{t_h^{(I)}}{I} \sum_{i=1}^{I} \delta(\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_h^{(i)})$$

Resource Allocation under Unknown Demand

Consider the stochastic resource allocation problem in, e.g., disaster control. The problem can be modeled via a distributionally robust newsvendor problem.

- **Set**
  - $[K_a]$: K_a is the number of subregions.
- **Parameter**
  - $\tilde{\boldsymbol{d}} \in \mathbb{R}^{K_a}$: uncertain demand.
  - $\boldsymbol{c^u}$ or $\boldsymbol{c^o} \in \mathbb{R}^{K_a}$: unit penalty cost of unmet demand or overserved demand.
  - $B \in \mathbb{R}$: resource allocation budget, $B > 0$.
- **Decision Variable**
  - $\boldsymbol{x} \in \mathbb{X}$: $\mathbb{X} = \left\{\boldsymbol{x} \in \mathbb{R}_+^{K_a}: \sum_{k=1}^{K_a} x_k \leq B\right\}$

$$J^* = \inf_{\boldsymbol{x}\in \mathbb{X}} \left\{ \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}}_{HI})} \mathbb{E}_{\mathbb{Q}}\left[(\boldsymbol{c}^u)^T(\tilde{\boldsymbol{d}} - \boldsymbol{x})^+ + (\boldsymbol{c}^o)^T(\boldsymbol{x} - \tilde{\boldsymbol{d}})^+\right]\right\}$$

$$= \inf_{\boldsymbol{x}\in \mathbb{X}} \left\{ \sup_{\mathbb{Q}\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}}_{HI})} \mathbb{E}_{\mathbb{Q}}\left[\sum_{k=1}^{K_a} \max\left\{c_k^u\left(\tilde{d}_k - x_k\right)^+, c_k^o\left(x_k - \tilde{d}_k\right)^+\right\}\right]\right\}$$

Numerical Setup

- $K = 4$: Number of subregions, $H = 3$: Number of data sources, $I = 200$: Number of iterations.
- $\mathbf{c}_u = \begin{bmatrix} 5000 \\ 10000 \\ 15000 \end{bmatrix}$, $\mathbf{c}_o = \begin{bmatrix} 1000 \\ 1000 \\ 1000 \end{bmatrix}$, $B = 200$.
- Initial trust $t_{h,k}^{(0)} = 1/H, \quad \forall k \in [K],\ h \in [H]$.
- $\tilde{d}_{\text{true},k}^{(i)} \sim U(10, 20),\ \forall i,\ k$.
- $\tilde{d}_{h,k}^{(i)}$ is sampled from a truncated normal distribution with mean: $\tilde{d}_{\text{true},k}^{(i)} + \mu_{h,k}$, and variance: $\sigma_{h,k}^2$.
- $\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \boldsymbol{\mu}_3)^\mathsf{T} = ((0,0,0,0)^\mathsf{T}, (0,5,0,5)^\mathsf{T}, (0,-5,5,2)^\mathsf{T})^\mathsf{T}$ and $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2, \boldsymbol{\sigma}_3)^\mathsf{T} = ((1,1,1,5)^\mathsf{T}, (2,1,5,5)^\mathsf{T}, (5,1,1,2)^\mathsf{T})^\mathsf{T}$.
- 30 replications

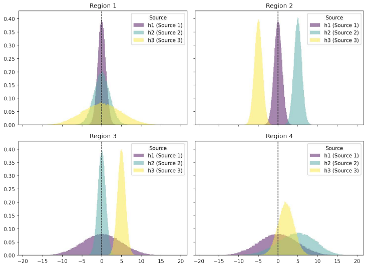

Nonparametric vs. Parametric Ambiguity Set for MR-DRO

We show the trust update process based on non-parametric and parametric data-fusion.

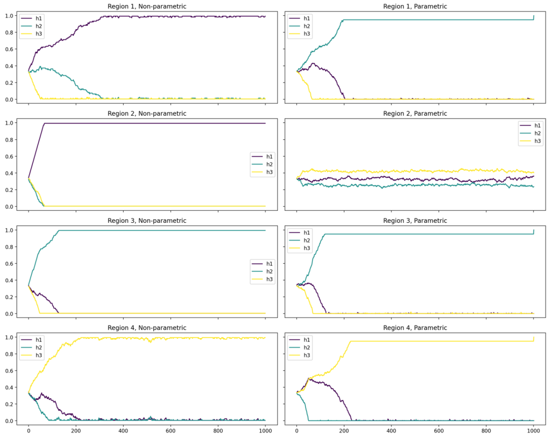

Overall, the nonparametric way of building the Wasserstein ambiguity set can result in similar in-sample objective loss using 1/3 of CPU time, and can lead to 50% out-of-sample loss reduction.

Probability Dominance and Dominate Trust

For showing how probability dominance relates to trust dominance, consider only two data sources in each subregion, and report the corresponding trust update results for the 30 trials.

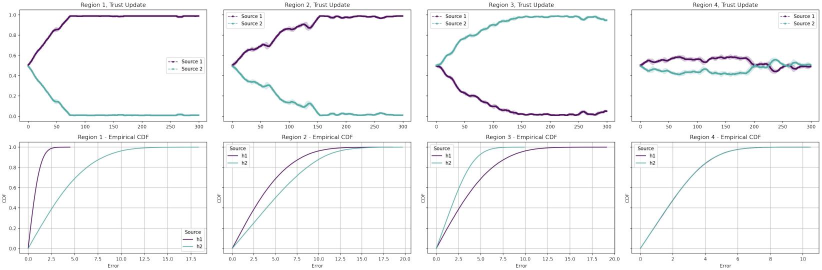

If one data source dominates the other ⇒ converges to dominating trust of the former. Otherwise, we do not fully trust either source[5].

[6] See Theorems 4–7 in Y. Guo, R. Jiang, S. Shen, "Optimization with Multi-sourced Reference Information and Unknown Trust: A Distributionally Robust Approach," https://arxiv.org/abs/2501.07057.

Outline

Problem Description

- Example: route recommendation (e.g., based on shortest path) to mass travelers.
- Applications: autonomous navigation, evacuation planning, shared mobility, delivery and logistics, etc.

Figure: Evacuation route sign

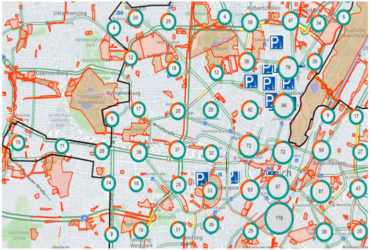

Figure: Shared mobility tracking

Travelers' Trust on Recommended Routes

Characteristics:

- Not 100% acceptance of recommended (optimal) solutions
- High level of trust $\Rightarrow$ better route optimization
- Unknown trust and travel environment $\Rightarrow$ Mixed endogenous and exogenous uncertainties

Challenges:

- To incorporate trust in optimization models for decision recommendation
- To dynamically learn unknown trust of travelers (dynamic trust update)

Motivating Example

Each arc has capacity 1, if exceeding the capacity, the travel time will double

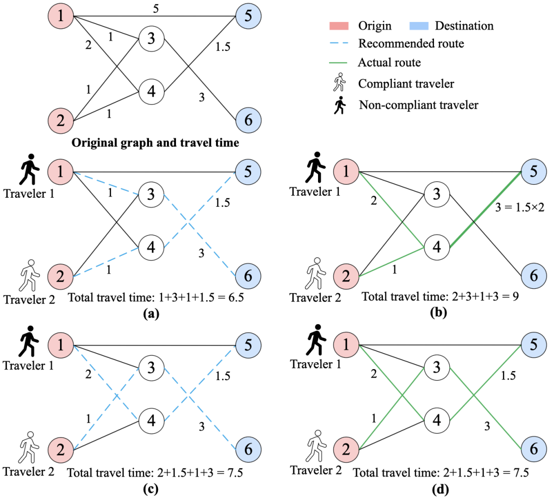

Parameters & Variables

- Parameters
  - N, O, D: set of all the nodes, origin nodes, and destination nodes.
  - A: set of arcs.
  - P_o: shortest path from origin node $o \in O$ to the closest shelter.
  - d_o: demand of each origin node $o \in O$.
  - α_o: trust rate on recommendation of travelers from origin node $o \in O$.
  - c_{ij}: capacity of road segment $(i, j) \in A$.
  - t_{ij}^0: travel time of road segment $(i, j) \in A$ without congestion.
  - a, b: pre-defined parameters in the travel time function in terms of congestion.
- Variables
  - x_{ij}: number of travelers who follow recommended routes on arc $(i, j) \in A$.
  - z_{ij}: total number of travelers on arc $(i, j) \in A$.

A Baseline Model

- Travel time of each traveler: the Bureau of Public Roads (BPR) function
- Objective function: total travel time

$$(P_\alpha) \quad \min \sum_{(i,j)\in A} t_{ij}^0 \left(1 + a\left(\frac{z_{ij}}{c_{ij}}\right)^b\right) z_{ij}$$

s.t. $\sum_{(i,j)\in A} x_{ij} - \sum_{(j,i)\in A} x_{ji} = \alpha_o d_o,\ \forall o \in O$	Flow amount following recommendation = trust rate × total demand
$\sum_{(i,j)\in A} x_{ij} - \sum_{(j,i)\in A} x_{ji} = 0,\ \forall i \in N\backslash O\backslash D$	Flow balance constraints for ordinary nodes
$z_{ij} = x_{ij} + \sum_{o\in O:(i,j)\in P_o} (1-\alpha_o)d_o,\ \forall (i,j) \in A$	Relationship between total flow and flow following recommendation
$z_{ij}, x_{ij} \geq 0,\ \forall (i,j) \in A.$	Non-negative constraints

This can be equivalent to solving a second-order conic program:

$$(SOCP_\alpha) \quad \min \sum_{(i,j)\in A} t_{ij}^0 \left(z_{ij} + \frac{a}{c_{ij}^b}\gamma_{1,ij}\right)$$

$$\text{s.t.} \quad \text{Original constraints}$$

$$\text{Constraints set } \Gamma.$$

Learn Travelers' Trust

- Goal: learn the true travelers' trust $\hat{\alpha}$ via dynamic simulation.
- Assumption: true trust rate $\hat{\alpha}$ follows a Normal distribution with mean α.
- Method: Bandit-based learning process
  - Bandit arm l: route planning model $SOCP_{\alpha_l}$ with estimated trust rate α_l.
  - Pull an arm l: run a simulation under a sampled true trust rate $\hat{\alpha}$ and recommendation from $SOCP_{\alpha_l}$.
  - Learn: which arm has the trust rate α_l closest to α? (represented by l^*)

Algorithmic Procedure

- Obtain route planning of a given trust rate α_l: solve $SOCP_{\alpha_l}$.
- Choose arm:

$$\hat{l} = \arg\max_{l} \; -\frac{T_{kl}}{n_l} + \sqrt{\frac{2\log(n)}{n_l}}$$

where $T_{kl} = \sum_{k=1}^{n} T_k \mathbf{1}_{\{arm_k = l\}}$: cumulative travel time of arm l by iteration k; n_l: total pulls number of arm l.

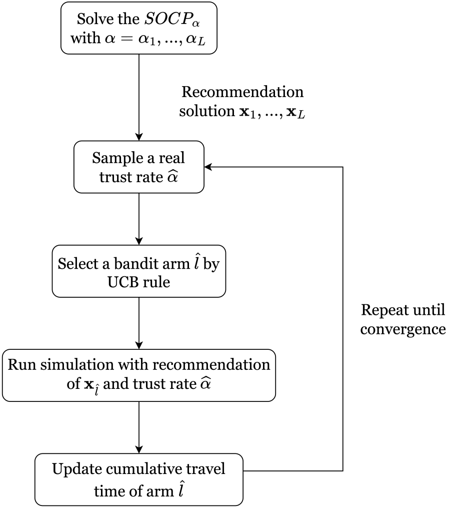

Convergence Results

- Recall: l^* is the arm with trust rate closest to the true trust rate mean α.

Theorem 1

After iteration n, the expectation of the cumulative evacuation time and evacuation time under plan with true trust rate is bounded by $O(\log n)$:

$$\mathbb{E}\left[\frac{1}{n}\sum_{k=1}^{n}\left(T(\alpha^{(k)}) - T_{l^*}(\alpha^{(k)})\right)\right] \leq O(\frac{\log n}{n}).$$

Theorem 2

For any arm such that $\alpha_l \neq \alpha_{l^*}$, the expectation of simulation times is bounded by $O(\log n)$, therefore after enough simulations, we will always provide the recommendation with the most accurate estimation.

Numerical Setup

- Instance: Sioux Falls network[6]
- 24 nodes & 74 links
- Origins: 6, 10, 23
- Destinations: 1, 13, 18, 20
- BPR parameters: $a = 0.15, \ b = 4$

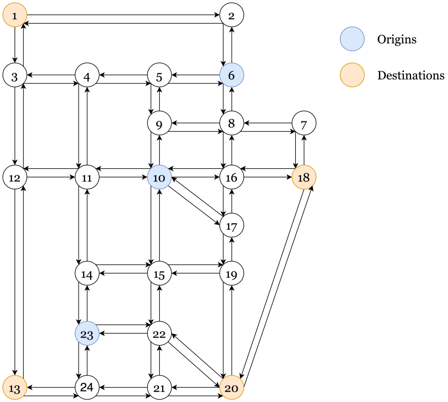

[6] Source: https://github.com/bstabler/TransportationNetworks

Average Travel Time Under Varying Known Trust Rate

- Average travel time results of model P_α with given trust rate α.
- A higher trust rate leads to a lower travel time.
- With trust rate $\geq 80\%$, the travel time is similar.

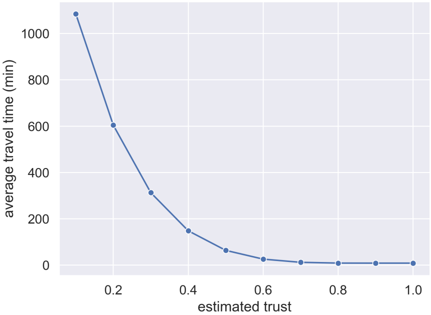

Average Travel Time Under Inaccurate Trust Estimation

- (x, y): solve the *SOCP* with y and run simulation under true rate x.
- $x = y$: best results under a given x.
- Accurate realization of trust rate improves travel time.
- A higher true trust rate requires more accurate estimation.

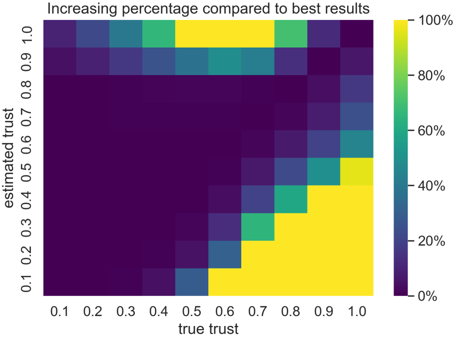

Dynamic Trust Learning

- Regret = cumulative average travel time - optimal travel time
- Nearly linear convergence
- If the underlying true trust is higher, the learning result is more stable and performs better.

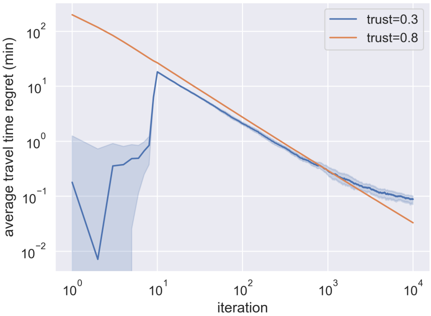

Outline

Conclusion

- We discuss two ways of introducing the notion of “trust” into decision models under uncertainty.
- In MR-DRO, trust is used as a parameter to weigh on multiple data sources, and we can dynamically update the trust parameter based on estimation errors or objective losses (if real data cannot be obtained).
- In trust-informed recommendation, trust is part of the uncertainty vector in the stochastic optimization model, and we introduce a bandit setting to learn true user trust by simulating our solutions iteratively.

Future research:

- Study decision-making problems with endogenous trust.
- Study multistage stochastic optimization problems with unknown trust in different data sources or decisions.
- Explore different learning methods for updating trust.

Q & A

Thank you!