

Distributionally Robust Optimization and Its Applications in Communication and Networking

Zhu Han, IEEE/AAAS Fellow
University of Houston

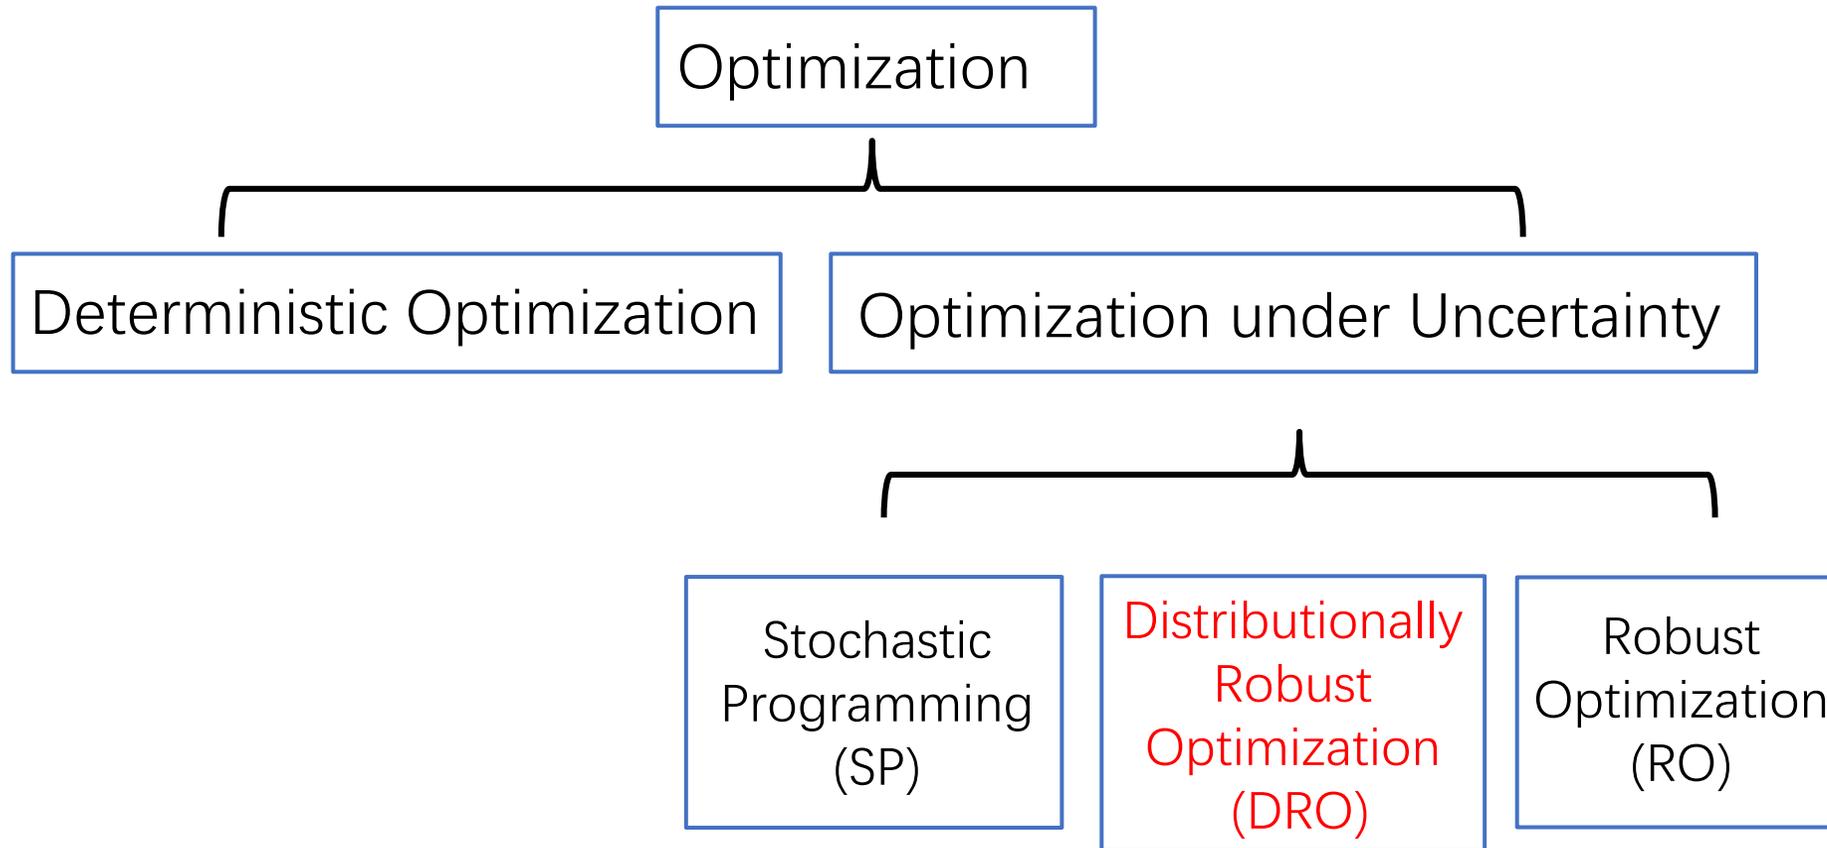


Thanks to Drs. Lei Fan, Zhuang Ling, Kai-Chu Tsai, Hongliang Zhang, Yali Chen, Yang Yang, and Yuan Zi

<http://wireless.egr.uh.edu/research.htm>

- **Introduction**
- **Case 1: Discrepancy-based DRO Optimization**
 - ✓ **Wasserstein distance**
 - ✓ **Tractability of Wasserstein DRO**
 - ✓ **Application I: Computation Offloading**
- **Case 2: DRO, Risk Aversion and CVaR**
 - ✓ **VaR and CVaR in Optimization**
 - ✓ **Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR)**
 - ✓ **Application II: Age of Information**
- **Case 3: DRO based Reinforcement Learning**
 - ✓ **Distributionally Robust Policy Iteration**
 - ✓ **Application III: Well Logging**
- **Conclusion**

✓ Optimization Classification



✓ Optimization Classification

➤ *Stochastic Programming (SP)*

assumes that the *probability distribution* of random parameters is *known* (inferred from the historical data). Objective is to find a decision x that minimizes a functional of the *expected cost*.

➤ *Robust Optimization (RO)*

assumes that the *probability distribution* of random parameter is *unknown*, but its *fluctuation range* is known. Objective is to find a decision x that minimizes the *worst-case cost* over an *uncertainty set*.

➤ *Distributionally Robust Optimization (DRO)*

The distribution is known, but the *random parameters of probability distribution is uncertain*. In most cases, random parameters can only be *estimated by the samples*. Objective is to find a decision x that minimizes the *worst-case expected cost* over an *uncertainty set*.

✓ Stochastic Programming (SP)

The *probability distribution* of random parameters is *known* (inferred from the historical data). Objective is to find a decision x that minimizes a functional of the *expected cost*.

$$\min_{x \in \chi} E_P [h(x, \xi)]$$

- x -- Decision variables
- χ -- Convex set of feasible solutions
- ξ -- Uncertain parameter follows a certain distribution P
- $h(x, \xi)$ -- Objective function in x that depends on parameters ξ

Classical assumptions in stochastic programming:

- The probability distribution of the random parameter vector is *independent of decisions*
- The *"true" probability distribution* of the random parameter vector is *known* relaxing it requires addressing *distributional uncertainty*.

✓ Robust Optimization (RO)

The probability distribution of random parameter is unknown, but its **fluctuation range** is known. Objective is to find a decision x that minimizes **the worst-case expected cost** over an **uncertainty set**.

$$\min_{x \in \chi} \max_{\xi \in U} h(x, \xi)$$

x -- Decision variables

ξ -- Uncertain Parameter

χ -- Convex set of feasible solutions

U -- Uncertain set

$h(x, \xi)$ -- Objective function in x that depends on parameters ξ

- First find a decision x that minimizes the cost.
- Then a parameter ξ which leads to the maximum cost (worst case for given decision)

✓ Distributionally Robust Optimization (DRO)

- In practice, the random parameters are *uncertain*.
- Although the exact distribution of the random variables may not be known, people usually know *partial statistic information* via *certain observed samples*.

Choose an **intermediate approach** to obtain a robust form of distributed optimization problem (*DRO*):

$$\min_{x \in \mathcal{X}} E_P [h(x, \xi)] \longrightarrow \min_{x \in \mathcal{X}} \max_{P \in \mathcal{P}} E_P [h(x, \xi)]$$

\mathcal{P} is an uncertain set of probability distributions constructed from the samples.

Objective is to find a decision x that minimizes the *worst-case expected cost* over an *ambiguity set*.

✓ Key question: How to build ambiguity sets (uncertain sets)?

The probability distribution quantifying the model parameter uncertainty is known *ambiguously*. When choosing ambiguity sets, we need to consider the following:

- Rich enough to contain the **true data-generating distribution** with high confidence.
- Small enough to exclude **pathological distributions**.
- **Tractability**
- **Practical (Statistical) Meanings**
- Performance (the potential loss comparing to the benchmark cases)

The form of **ambiguity sets** can be used to classify the distributionally robust optimization problems.

- **Moment-based ambiguity sets** : $\mathcal{P} = \{\xi: E[\xi] \leq \mu, E[\xi^T \xi] \leq \Sigma, \dots\}$
- **Discrepancy-based ambiguity sets**: $\mathcal{P} = \{P: d(\widehat{P}_N, P) \leq \rho\}$

- Introduction
- **Case 1: Discrepancy-based DRO Optimization**
 - ✓ **Wasserstein distance**
 - ✓ **Tractability of Wasserstein DRO**
 - ✓ **Application I: Computation Offloading**
- **Case 2: DRO, Risk Aversion, and CVaR**
- **Case 3: DRO Based Reinforcement Learning**
- **Conclusion**

Motivation

- In many situations, we have an **empirical** estimate of the underlying probability distributions.
- A natural way to hedge against the distributional ambiguity is to consider a neighborhood of the empirical probability distribution

Discrepancy

Ambiguity sets based on probability distance:

$$\mathcal{P} = \{P: d(\widehat{P}_N, P) \leq \varepsilon\}$$

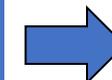
\widehat{P}_N -- Empirical probability

ε -- Radius

$d(\widehat{P}_N, P)$ -- **Metric** of the similarity of two distributions

By selecting a suitable *metric*, certain *infinite-dimensional* convex DRO problems can be transformed into *finite-dimensional* convex optimization problems

Is there a metric that is simple to calculate and suitable for discrete / continuous distributions?



**Wasserstein
distance**

Wasserstein distance

used to measure the distance between two distributions.

Definition:

$$d_W(P_1, P_2) = \inf_{\gamma \sim \Pi(P_1, P_2)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|]$$

$\Pi(P_1, P_2)$: the set of all possible joint distributions of P_1 and P_2 .

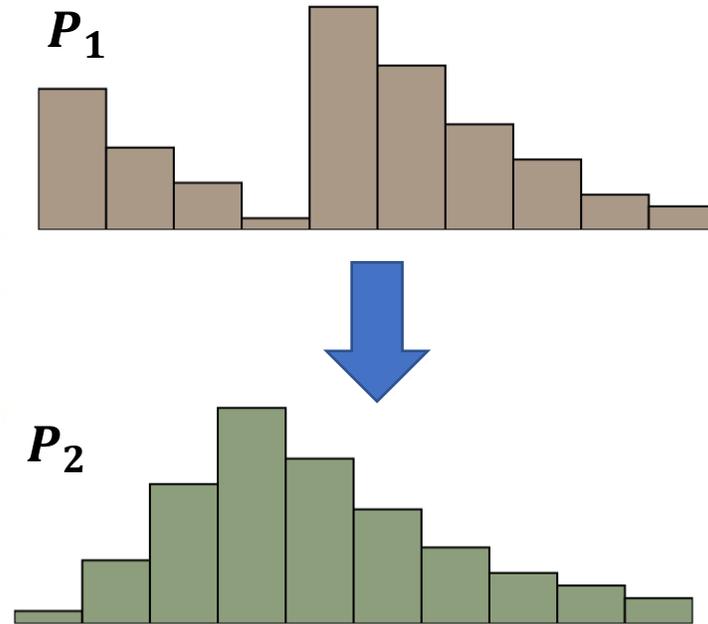
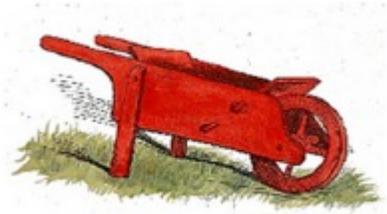
$(x, y) \sim \gamma$: samples under joint distribution γ

$\|x - y\|$: sample distance

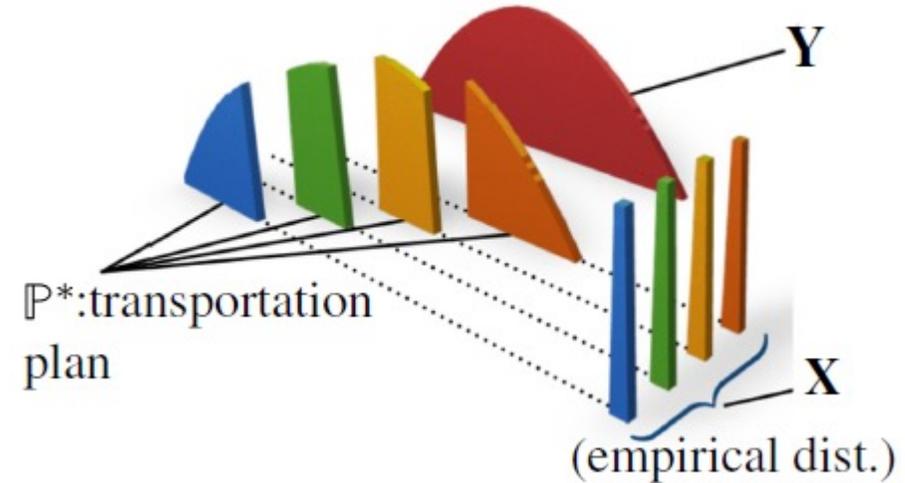
$\mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|]$: expectation of distance for sample x and y under joint distribution γ

Wasserstein distance of P_1 and P_2 : the lower bound of this expectation.

Wasserstein distance



Move mass P_1 into the shape and position of P_2 .



$\Pi(P_1, P_2)$: transportation plan

$\|x - y\|$: distance the soil moves

$Y(x, y)$: amount of moving soil from x to y

$\mathbb{E}_{(x,y) \sim \gamma}[\|x - y\|]$: bulldozing cost

Bulldozing cost : amount of moving soil multiplied by the distance the soil moves.

Wasserstein distance: the smallest bulldozing cost from P_1 to P_2 .

Wasserstein distance

Wasserstein distance-based ambiguity set:

$$\mathbb{B}_\varepsilon(\widehat{P}_N) = \{Q: d_W(\widehat{P}_N, Q) \leq \varepsilon\}$$

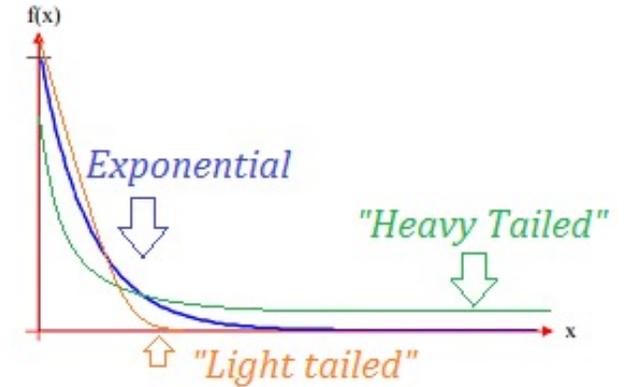
- The ambiguity set Q can be viewed as a **Wasserstein ball** which contains all probability distributions whose Wasserstein distance to the empirical distribution \widehat{P}_N is less than ε .
- Q will cover the true distribution with a higher probability with a larger value of ε .
 - There exists a trade-off between the accuracy and the complexity
 - It is important to well design the value of ε

Wasserstein distance

How to calculate ε of ambiguity set

Light-tailed distribution assumption: Distribution \mathbb{P} is call light-tailed if $a > 1$ such that

$$\mathbb{E}^{\mathbb{P}} [\exp(\|\xi\|^a)] = \int_{\Xi} \exp(\|\xi\|^a) \mathbb{P}(d\xi) < \infty.$$



This assumption requires the tail of the distribution to decay at an exponential rate. The assumption guarantees that the ambiguity set can cover **most of the possible distributions**.

Radius selection: With this assumption, suppose that $\widehat{\mathbb{P}}_N$ is the empirical distribution, m is related to the dimension and cost parameter, for $m \neq 2$ and $c_1, c_2 > 0$, under a confidence level of $1 - \beta$, we have

$$\varepsilon_N(\beta) := \begin{cases} \left(\frac{\log(c_1 \beta^{-1})}{c_2 N} \right)^{1/\max\{m, 2\}} & \text{if } N \geq \frac{\log(c_1 \beta^{-1})}{c_2}, \\ \left(\frac{\log(c_1 \beta^{-1})}{c_2 N} \right)^{1/a} & \text{if } N < \frac{\log(c_1 \beta^{-1})}{c_2}. \end{cases}$$

Number of samples

Tractability of Wasserstein DRO

With Wasserstein ball, the DRO problem can be rewritten as

$$\min_{x \in \mathcal{X}} \max_{P \in \mathbb{B}_\varepsilon(\widehat{P}_N)} E_P[h(x, \xi)]$$

Assumption:

- The uncertainty set is convex and closed
- $h(x, \xi)$ is convex with respect to ξ

Challenges to Compute Wasserstein Distances:

Computing the Wasserstein distance between two distributions \widehat{P}_N and P_N is NP-hard if one of them is continuous since the dimension will be infinite.

Solution:

Rewrite the problem into a finite-dimensional convex program by leveraging tools from robust optimizations

Tractability of Wasserstein DRO

How to transform an infinite-dimensional optimization problem into a finite-dimensional convex program:

$$\max_{P \in \mathbb{B}_\rho(\widehat{P}_N)} E_P[h(x, \xi)] = \begin{cases} \sup_{\Pi, P} \int_{\mathbb{E}} h(x, \xi) P(d\xi) \\ \text{s. t.} \quad \int_{\mathbb{E}^2} \|\xi - \xi'\| \Pi(d\xi, d\xi') \leq \varepsilon \end{cases}$$

$$\widehat{P}_N = \frac{1}{N} \sum_{i=1}^N \delta_{\widehat{\xi}_i}$$

Convex reduction

$$= \begin{cases} \sup_{P_i \in \mathcal{M}(\mathbb{E})} \frac{1}{N} \sum_{i=1}^N \int_{\mathbb{E}} h(x, \xi) P_i(d\xi) \\ \text{s. t.} \quad \frac{1}{N} \sum_{i=1}^N \int_{\mathbb{E}} \|\xi - \widehat{\xi}_i\| P_i(d\xi) \leq \varepsilon \end{cases}$$

Π is a joint distribution of ξ' and ξ with marginals distribution P and \widehat{P}_N of ξ' given $\xi' = \widehat{\xi}_i$, and conditional P_i of ξ . N is number of samples. Due to the law of total probability, we have

$$\Pi = \frac{1}{N} \sum_{i=1}^N \delta_{\widehat{\xi}_i} \otimes P_i$$

Tractability of Wasserstein DRO

Using a standard duality argument, we obtain

$$\begin{aligned} & \max_{P \in \mathbb{B}_\varepsilon(\widehat{P}_N)} E_P[h(x, \xi)] \\ &= \sup_{P_i \in \mathcal{M}(\mathbb{E})} \inf_{\lambda \geq 0} \frac{1}{N} \sum_{i=1}^N \int_{\mathbb{E}} h(x, \xi) P_i(d\xi) + \lambda \left(\varepsilon - \frac{1}{N} \sum_{i=1}^N \int_{\mathbb{E}} \|\xi - \widehat{\xi}_i\| P_i(d\xi) \right) \\ & \leq \inf_{\lambda \geq 0} \sup_{P_i \in \mathcal{M}(\mathbb{E})} \lambda \varepsilon + \frac{1}{N} \sum_{i=1}^N \int_{\mathbb{E}} (h(x, \xi) - \lambda \|\xi - \widehat{\xi}_i\|) P_i(d\xi) \\ & = \inf_{\lambda \geq 0} \lambda \varepsilon + \frac{1}{N} \sum_{i=1}^N \sup_{\xi \in \mathbb{E}} (h(x, \xi) - \lambda \|\xi - \widehat{\xi}_i\|) \end{aligned}$$

Maximum minima is
always less than
minimum maxima



\leq

The uncertainty set
contains all distributions



$=$

Tractability of Wasserstein DRO

Introducing epigraphical auxiliary variables s_i , $i \leq N$:

$$\begin{aligned} & \max_{P \in \mathbb{B}_\varepsilon(\widehat{P}_N)} E_P[h(x, \xi)] \\ & = \begin{cases} \inf_{\lambda, s_i} \lambda \varepsilon + \frac{1}{N} \sum_{i=1}^N s_i \\ \text{s. t. } \sup_{\xi \in \Xi} (h(x, \xi) - \lambda \|\xi - \widehat{\xi}_i\|) \leq s_i, \forall i \leq N \\ \lambda \geq 0 \end{cases} \end{aligned}$$

As such, the problem is transformed into a finite convex program and can be solved by existing convex optimization techniques.

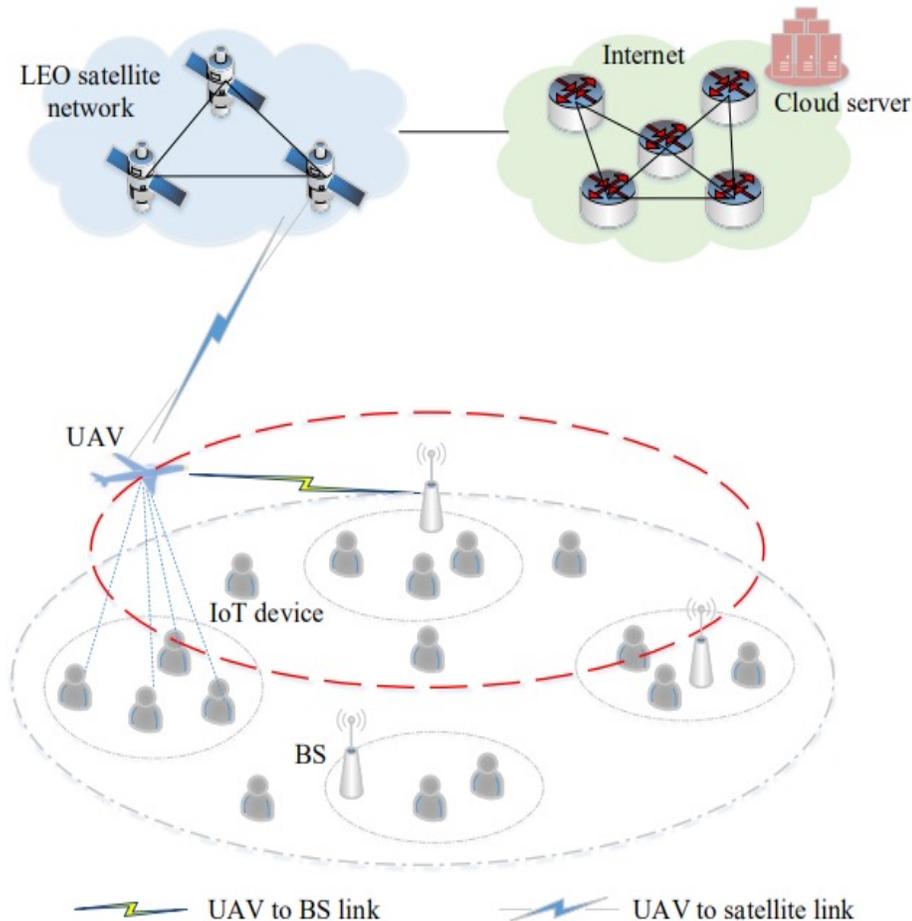
Special Case

For the case of quadratic loss (possibly non-convex and non-concave), we still can derive its exact form.

Assume that $x(\xi) = \xi^T Q \xi + 2q^T \xi$ is a (possibly indefinite) quadratic loss function. Then the worst-case risk coincides with the optimal value of a **tractable semidefinite program (SDP)**, that is,

$$\begin{cases} \max_{P \in \mathbb{B}_\varepsilon(\widehat{P}_N)} E_P [h(x, \xi)] = \\ \inf_{\gamma, \mathbf{s}_i} \gamma \varepsilon^2 + \frac{1}{N} \sum_{i=1}^N \mathbf{s}_i \\ \text{s. t.} \begin{pmatrix} \gamma - Q & q + \gamma \widehat{\xi}_i \\ q^T + \gamma \widehat{\xi}_i^T & \mathbf{s}_i + \|\widehat{\xi}_i\|_2^2 \end{pmatrix} \geq 0 \end{cases}$$

Application 1: Computation Offloading



Space-Air-Ground Integrated Networks:

- IoT devices: request services
- UAV: collect data from IoT devices
 - Determine to offload to **a nearby BS** or offload to **a certain satellite** and utilize the cloud server
 - Determine the proportion of tasks to offload, and the others will be done by the UAV
- The task request received by the UAV is **uncertain**

Yali Chen, Bo Ai, Yong Niu, Hongliang Zhang, and Zhu Han, "Energy-Constrained Computation Offloading in Space-Air-Ground Integrated Networks using Distributionally Robust Optimization," IEEE Transactions on Vehicular Technology (Volume: 70, Issue: 11, Nov. 2021)

Application 1: Problem Formulation

Objective: minimize the **expected total latency** of T time slots under the **worst-case distribution** realization in uncertainty set D

$$\min_{\mathbf{X}, \mathbf{Y}} \max_{P \in D} \mathbb{E}_P[\psi(\mathbf{X}, \mathbf{Y}, \xi_k)]$$

s.t.

$$(a) x_{bt} \in \{0, 1\}, x_{st} \in \{0, 1\}, \forall b, s, t$$

$$(b) \sum_{b=1}^N x_{bt} + \sum_{s=1}^M x_{st} = 1, \forall t$$

$$(c) y_{ut}(\xi_k) + \sum_{b=1}^N y_{bt}(\xi_k) + \sum_{s=1}^M y_{st}(\xi_k) = \xi_k, \forall t, k$$

the sum of tasks processed by UAV, all BSs, and all LEO satellites equals to the amount of tasks arrived

$$(d) y_{ut}(\xi_k), y_{bt}(\xi_k), y_{st}(\xi_k) \geq 0, \forall b, s, t, k$$

$$(e) y_{ut}(\xi_k) \leq C_u, y_{bt}(\xi_k) \leq C_b x_{bt}, y_{st}(\xi_k) \leq C_s x_{st}, \forall b, s, t, k$$

ensure the tasks assigned cannot exceed its capacity

$$(f) \psi(\mathbf{X}, \mathbf{Y}, \xi_k) = \sum_{t=1}^T \max \left(\frac{\delta y_{ut}(\xi_k)}{f_u}, \sum_{b=1}^N L_{ubt} + \sum_{s=1}^M L_{ust} \right),$$

total system latency

$$(g) E_u^{fly} + \sum_{b=1}^N \sum_{t=1}^T E_{ubt} + \sum_{s=1}^M \sum_{t=1}^T E_{ust} \leq E_{max}, \forall k$$

energy constraint

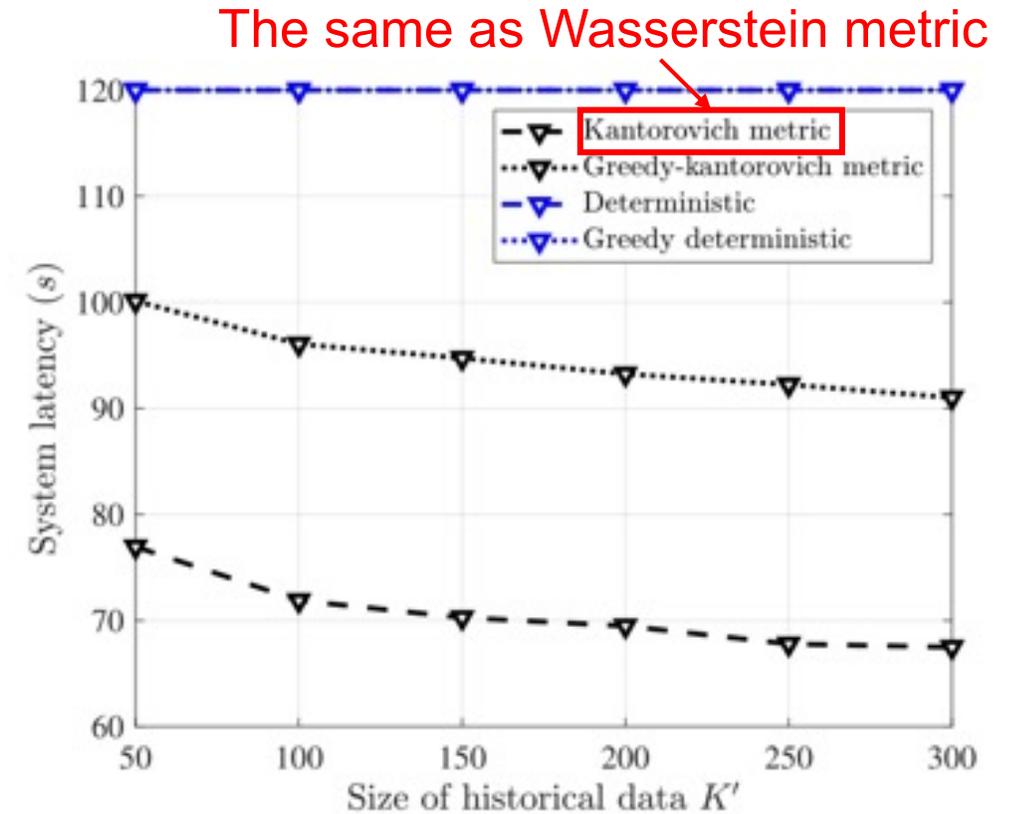
Application 1: Algorithm and Results

Algorithm 1 Distributionally Robust Latency Optimization Algorithm

Input: Sample space Ω , reference distribution P_0 based on a series of historical data

Output: \mathbf{X}, \mathbf{Y}

- 1: Solve the optimization problem (28) by the Gurobi Optimizer to obtain $P = \{p_1, p_2, \dots, p_K\}$;
 - 2: Solve the optimization problem (26) by the Gurobi Optimizer to get continuous \mathbf{X} and \mathbf{Y} ;
 - 3: **repeat**
 - 4: Select the branch variable x^* according to the rule shown in (34);
 - 5: Add $x \leq 0$ as a constraint to the problem (26) and solve it to obtain the optimal result Lat_0 ;
 - 6: Add $x \geq 1$ as a constraint to the problem (26) and solve it to obtain the optimal result Lat_1 ;
 - 7: **if** $Lat_0 < Lat_1$ **then**
 - 8: Update problem (26) by adding a constraint $x \leq 0$;
 - 9: **else**
 - 10: Update problem (26) by adding a constraint $x \geq 1$;
 - 11: **end if**
 - 12: **until** all elements in \mathbf{X} are integers.
-



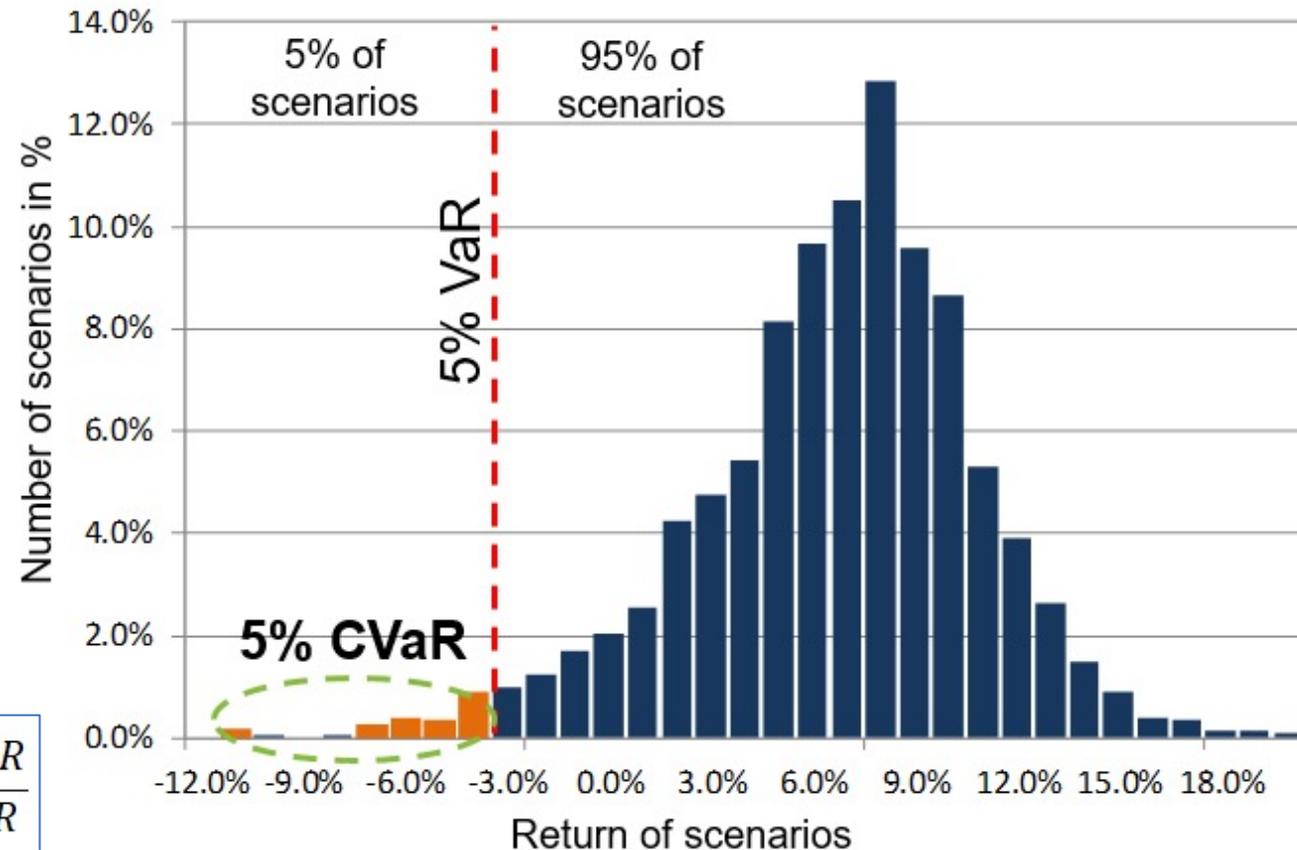
For deterministic scheme, many tasks cannot be processed and **retransmission** is required, leading to a higher latency

- Introduction
- Case 1: Discrepancy-based DRO Optimization
- **Case 2: DRO, Risk Aversion and CVaR**
 - ✓ Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR)
 - ✓ VaR and CVaR in Optimization
 - ✓ Application II: Age of Information
- Case 3: DRO based Reinforcement Learning
- Conclusion

VaR & CVaR Representation

- Risk Management is a procedure for evaluating loss distribution in Risk Aversion.
- Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) are popular function for measuring risk.

Risk functions: graphical representation of VaR, VaR Deviation, CVaR, CVaR Deviation, Max Loss, and Max Loss Deviation.



$$5\% CVaR = \frac{\sum \text{return of scenarios} < 5\% VaR}{\text{number of scenarios} < 5\% VaR}$$

VaR & CVaR

Let X be a random variable, with the X may have meaning of loss or gain.

$$F_X(z) = P\{X \leq z\}$$

VaR Definition A lower α -percentile of the random variable X .

$$\text{VaR}_\alpha(X) = \min\{z \mid F_X(z) \geq \alpha\}$$

$\text{VaR}_\alpha(X)$ is nonconvex and discontinuous function of the confidence level α for discrete distributions.

Difficult to control/optimize for non-normal distributions: VaR has many extremums for discrete distributions.

VaR& CVaR

CVaR Definition

The CVaR of X with confidence level $\alpha \in [0, 1]$ is the mean of the generalized α -tail distribution:

$$CVaR_{\alpha}(X) = \int_{-\infty}^{+\infty} z dF_X^{\alpha}(z)$$

where the distribution in question is the one with distribution function defined by

$$F_X^{\alpha}(z) = \begin{cases} 0 & \text{when } z < VaR_{\alpha}(X) \\ \frac{F_X(z) - \alpha}{1 - \alpha} & \text{when } z \geq VaR_{\alpha}(X) \end{cases}$$

For random variables with continuous distribution functions, CVaR $_{\alpha}(X)$ equals the conditional expectation of X subject to $Z \geq VaR_{\alpha}(X)$.

CVaR is continuous with respect to α . CVaR is convex in X

VaR & CVaR Example

Main focus: Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) [Rockafellar and Uryasev, 2000].

if $\text{VAR}(95) = 3\%$

$$5 = 100 - 95$$

5% chance to lose 3%
or more on a given day

if $\text{CVAR}(95) = 4.5\%$

$$5 = 100 - 95$$

In the worst
5% of returns

Your average
loss will be 4.5%

CVaR gives us an average expected loss
VaR gives us a range of potential losses

- If more than lost 3% happens with 5%, but we still do not know average lost.
- CVAR describes the average lost conditioned on VaR happens.

Equivalence of Chance Constraints and VaR Constraints

In **portfolio management**, often it is required that portfolio loss at a certain future time is, with high reliability, at most equal to a certain value.

Let $f(x, y)$ be the loss associated with **the decision vector** x and **the random vector** y .

$$\text{Prob}\{f_i(x, y) \leq 0\} \geq p_i, i = 1, \dots, m.$$

Let $VaR_\alpha(x)$ be the VaR_α of a loss function $f(x, y)$

$$VaR_\alpha(x) = \min \left\{ \zeta : \text{Prob}\{f(x, y) \leq \zeta\} \geq \alpha \right\}$$

Then the following holds:

$$\text{Prob}\{f(x, y) \leq \zeta\} \geq \alpha \iff \text{Prob}\{f(x, y) > \zeta\} \leq 1 - \alpha \iff VaR_\alpha(x) \leq \zeta$$

In general, $VaR_\alpha(x)$ is **nonconvex** w.r.t. x , (e.g., discrete distributions)

Minimization and CVaR Constraints

The underlying probability distribution of y will be assumed for convenience to have density $p(y)$. The α -CVaR of the loss associated with a decision x is the value

$$CVaR(x) = \frac{1}{1-\alpha} \int_{f(x,y) \geq VaR(x)} f(x,y) p(y) dy$$

The main idea is to define a function that can be used in connection with VaR and CVaR

$$F_\alpha(x, \zeta) = \zeta + \frac{1}{1-\alpha} \int_{y \in \mathcal{Y}} [f(x,y) - \zeta]^+ p(y) dy$$



$$F_\alpha(x, \zeta) = \zeta + \frac{1}{1-\alpha} E\{[f(x,y) - \zeta]^+\}, \quad \text{where } [t]^+ = \max\{0, t\}.$$

1. $F_\alpha(x, \zeta)$ is convex with respect to α ;
2. $VaR_\alpha(x)$ is a minimum point of function $F_\alpha(x, \zeta)$ w.r.t. ζ ;
3. As a function of $\zeta \in \mathbb{R}$, $F_\alpha(x, \zeta)$ is finite and convex (hence continuous), :

Chance constrained programs

It is known that **distributionally robust chance constraints** can be conservatively approximated by **Worst-Case Conditional Value-at-Risk (CVaR)** constraints.

$$\begin{aligned} & \underset{x \in \mathbb{R}^n}{\text{minimize}} && c^T x \\ & \text{subject to} && \Pr \left(a_i(\tilde{\xi})^T x \leq b_i(\tilde{\xi}) \quad \forall i = 1, \dots, m \right) \geq 1 - \epsilon \\ & && x \in \mathcal{X}, \end{aligned}$$

\Pr can be unknown arbitrary distribution

x is the **decision vector** and c is a **cost vector**

The **chance constraint** requires a set of m **uncertainty-affected inequalities** to be jointly satisfied with a probability of at least $1 - \epsilon$.

$$\begin{aligned} a_i(\tilde{\xi}) &= a_i^0 + \sum_{j=1}^k a_i^j \tilde{\xi}_j \quad \text{and} \quad b_i(\tilde{\xi}) = b_i^0 + \sum_{j=1}^k b_i^j \tilde{\xi}_j. \\ y_i^j(x) &= (a_i^j)^T x - b_i^j, \quad i = 1, \dots, n, \quad j = 0, \dots, k. \end{aligned}$$

$$\Pr \left(y_i^0(x) + y_i(x)^T \tilde{\xi} \leq 0 \quad \forall i = 1, \dots, m \right) \geq 1 - \epsilon$$

Distributionally robust approach

A natural way to *immunize* the chance constraint against uncertainty in the probability distribution is to adopt the **following ambiguous or distributionally robust chance constraint**.

$$\begin{aligned} & \underset{x \in \mathbb{R}^n}{\text{minimize}} && c^T x \\ & \text{subject to} && \inf_{\substack{\mathbb{P} \in \mathcal{P} \\ x \in \mathcal{X}}} \mathbb{P} \left(y_i^0(x) + y_i(x)^T \tilde{\xi} \leq 0 \quad \forall i = 1, \dots, m \right) \geq 1 - \epsilon \end{aligned}$$

$\mathbb{P} \in \mathcal{P}$ denotes the distribution \mathbb{P} belongs to an uncertainty set \mathcal{P} with certain known structural properties.

Distributionally robust chance constraint means **the worst case** satisfies the probability at least $1 - \epsilon$ in the presence of **channel uncertainties**.

For $m = 1$, the feasible set is denoted by

$$\mathcal{X}^{\text{ICC}} = \left\{ x \in \mathbb{R}^n : \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P} \left(y^0(x) + y(x)^T \tilde{\xi} \leq 0 \right) \geq 1 - \epsilon \right\}$$

CVaR chance constraints

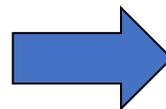
The CVaR at ϵ level with respect to \mathbb{P} is defined as

$$\mathbb{P}\text{-CVaR}_\epsilon(L(\tilde{\xi})) = \inf_{\beta \in \mathbb{R}} \left\{ \beta + \frac{1}{\epsilon} \mathbb{E}_{\mathbb{P}} \left((L(\tilde{\xi}) - \beta)^+ \right) \right\}$$

CVaR essentially evaluates the **conditional expectation** of loss above the $(1 - \epsilon)$ -quantile of the loss distribution. It can be shown that CVaR represents a convex functional of the random variable $L(\tilde{\xi})$

$$\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon \left(y^0(x) + y(x)^T \tilde{\xi} \right) \leq 0 \implies \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P} \left(y^0(x) + y(x)^T \tilde{\xi} \leq 0 \right) \geq 1 - \epsilon.$$

$$\begin{aligned} & \text{minimize } c^T x \\ & \quad x \in \mathbb{R}^n \\ & \text{subject to } \Pr \left(a_i(\tilde{\xi})^T x \leq b_i(\tilde{\xi}) \quad \forall i = 1, \dots, m \right) \geq 1 - \epsilon \\ & \quad x \in \mathcal{X}, \end{aligned}$$

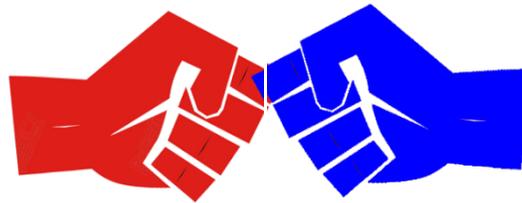


$$\begin{aligned} & \text{minimize } c^T x \\ & \quad x \in \mathbb{R}^n \\ & \text{subject to } \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-CVaR}_\epsilon \left(y^0(x) + y(x)^T \tilde{\xi} \right) \leq 0 \end{aligned}$$

Age of Information Minimization in Healthcare IoT Using Distributionally Robust Optimization [2]

How can I deal with channel state information (CSI) errors in Healthcare IoT system ?

- ❑ The perturbations in CSI are modeled to be statistically unbounded according to **Gaussian distribution** and **Rayleigh distribution**.
- ❑ These statistical channel assumptions **may not match** the healthcare IoT applications perfectly.



- ❑ The **retransmission scheme** is investigated against CSI errors to reduce the Aol.
- ❑ The retransmission will lead to the **energy consumption** of the IoT device **increase dramatically**, especially when the IoT device is used to retransmit the same updates for infinite times

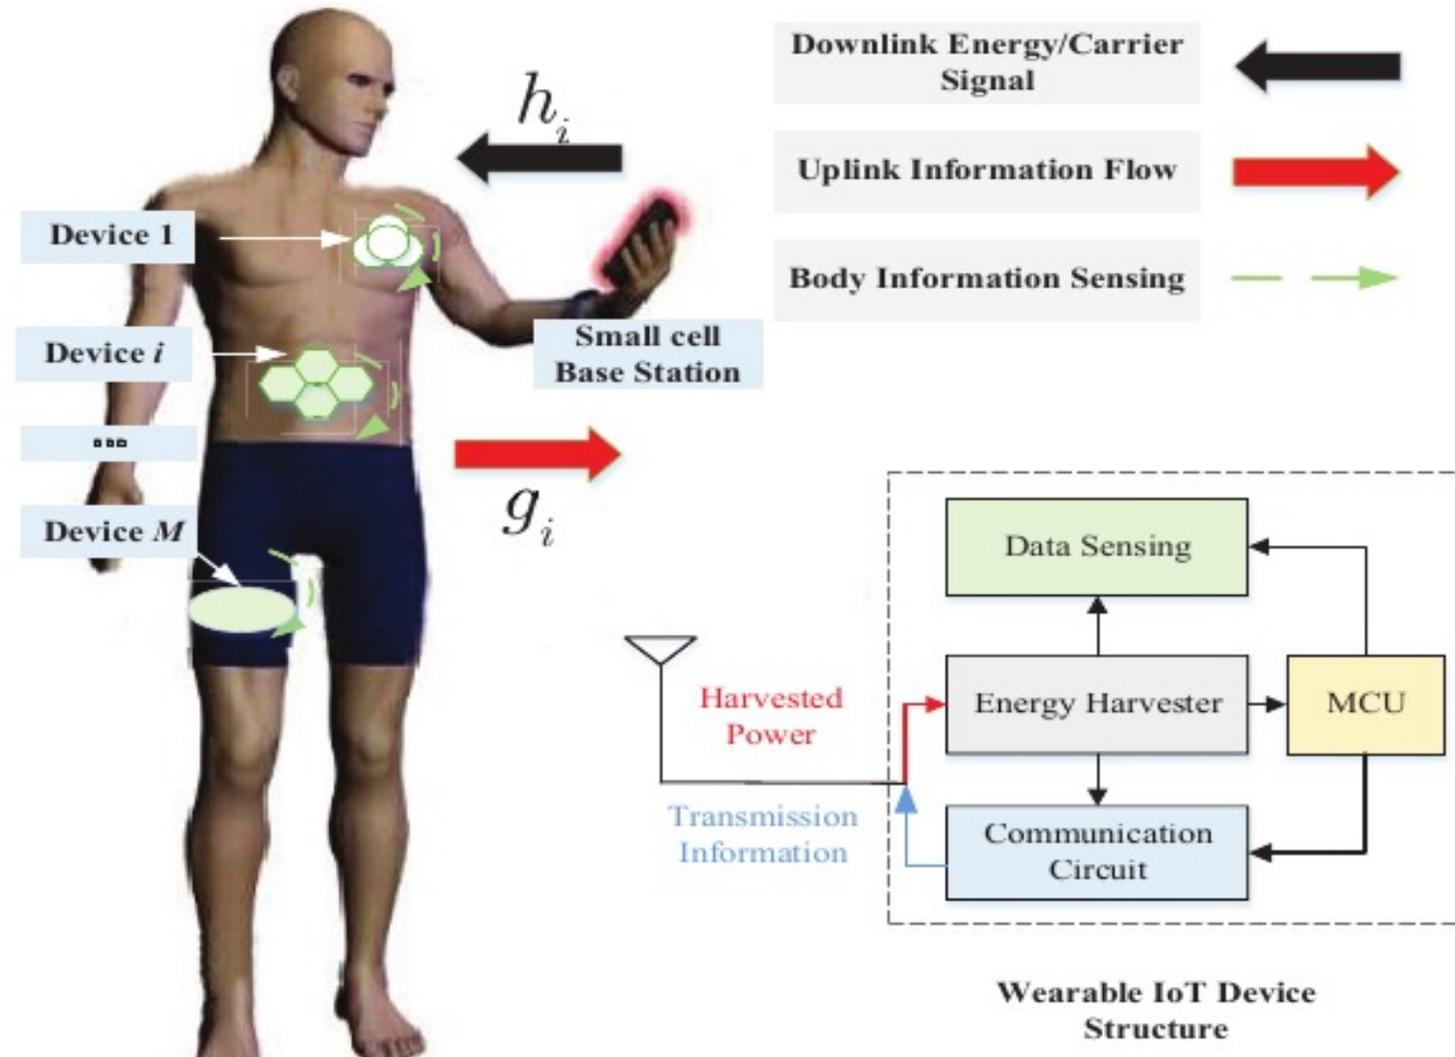
A key open problem is to consider **the case of imperfect CSI** and investigate **how CSI error effects may be mitigated through quantitative designs.**

Application 2: Body Area Networks

System model

As shown in this picture, we consider **a single-hop cellular IoT** to support healthcare IoT applications:

1. the wearable IoT devices **harvest wireless RF energy** from a dedicated power cell BS
2. the wearable IoT devices send **their sensing physiological information signals** to a separate information receiving cell BS.



Age of Information = Time – Timestamp



Sensor



Destination

Measure the “age of information” that destination know about a source node.

Potential application: UAV, Uplink, Down link system in wireless communication IoT network, etc.

Application 2: Problem Formulation



$$\begin{aligned}
 & \min_{\substack{\alpha_i(t), p_i(t), \gamma_i(t), \tau_i(t) \\ \mathbf{M}_1, \mathbf{M}_2, \beta_1, \beta_2}} \sup_{\mathbb{P} \in \mathcal{P}} \bar{\Delta}_i(t) \\
 s.t. & \begin{cases} \beta_1 + \frac{1}{\gamma_i(t)_{\max}} \text{Tr}(\mathbf{\Omega}_1 \mathbf{M}_1) \leq 0, \\ \mathbf{M}_1 \in \mathbb{S}^2, \beta_1 \in \mathbb{R}, \\ \mathbf{M}_1 \succcurlyeq \mathbf{0}, \mathbf{M}_1 - \begin{bmatrix} q_1 & 0 \\ 0 & E_{C,i} - \beta_1 \end{bmatrix} \succcurlyeq \mathbf{0}, \end{cases} \\
 & \begin{cases} \beta_2 + \frac{1}{\tau_i(t)_{\max}} \text{Tr}(\mathbf{\Omega}_2 \mathbf{M}_2) \leq 0, \\ \mathbf{M}_2 \in \mathbb{S}^3, \beta_2 \in \mathbb{R}, \\ \mathbf{M}_2 \succcurlyeq \mathbf{0}, \mathbf{M}_2 - \begin{bmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & r_{U,\min} \sigma_A^2 - \beta_2 \end{bmatrix} \succcurlyeq \mathbf{0}, \end{cases} \\
 & 0 < \alpha_i(t) < \alpha_{\max} \\
 & 0 < p_i(t) \leq P_{\max}, \\
 & 0 < \gamma_i(t) < 1, \\
 & 0 < \tau_i(t) < 1, \\
 & \forall (\Delta_{h_i}, \Delta_{g_i}) \sim \mathbb{P} \in \mathcal{P}, \\
 & \forall i \in \{1, 2, \dots, M\}, \forall t \in \{1, 2, \dots, K\},
 \end{aligned}$$

CVaR constraint.

CVaR constraint.

Power constraint

Outage probability constraint for energy harvesting.

Outage probability constraint for transmission

To get the optimal average Aol, we propose a **low-complexity upper bound of Aol minimization (UBAM)** algorithm in an iterative manner to address the **distributionally robust optimization problem**.

Algorithm 1 UBAM Algorithm

Input: Initialize $\alpha^{\{0\}}$, $\mathbf{p}^{\{0\}}$, $\tau^{\{0\}}$, $\gamma^{\{0\}}$, the terminated threshold $\Theta = 10^{-3}$, and $j = 0$.

1: **repeat**

2: $j = j + 1$.

3: Solve the power allocation subproblem for given $\gamma^{\{j\}}$ and $\tau^{\{j\}}$, to obtain the feasible solution as $\alpha^{\{j+1\}}$ and $\mathbf{p}^{\{j+1\}}$.

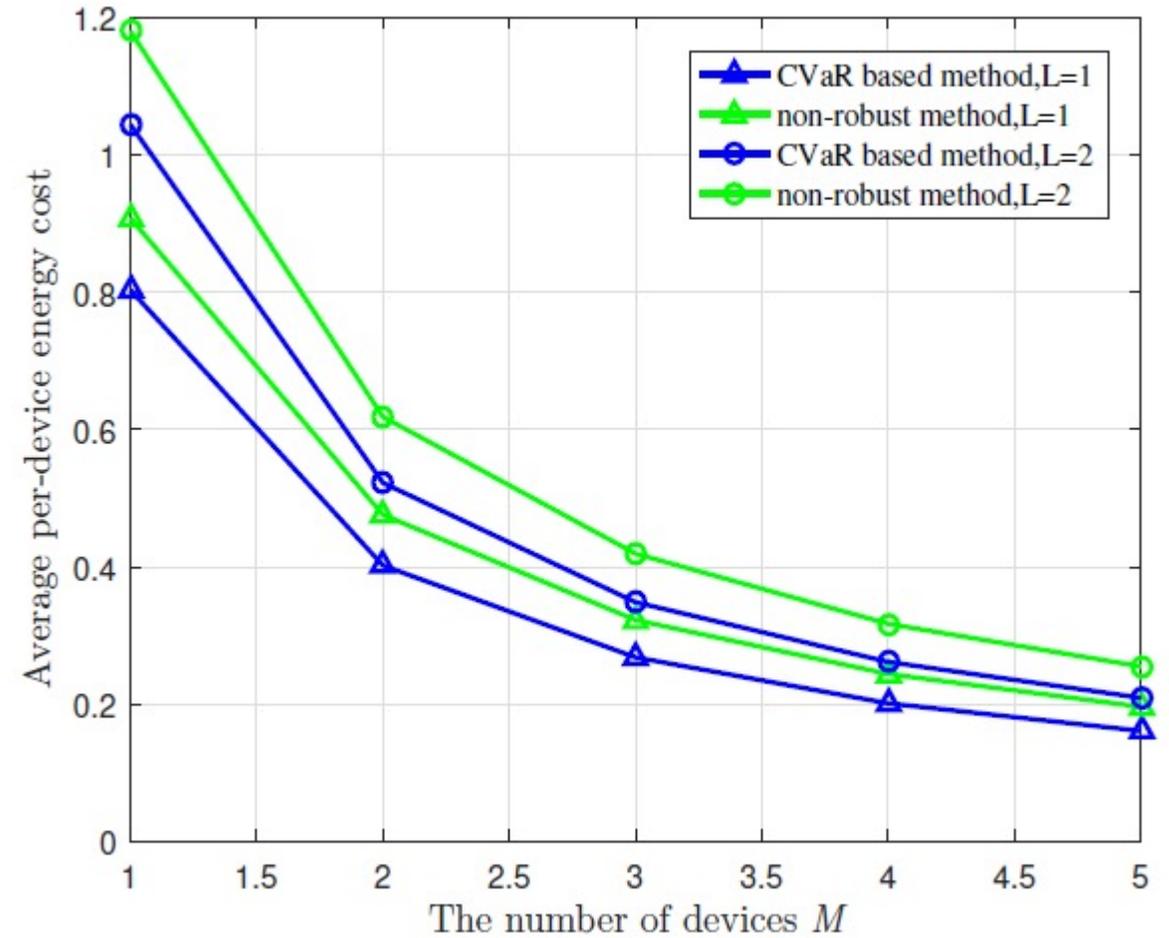
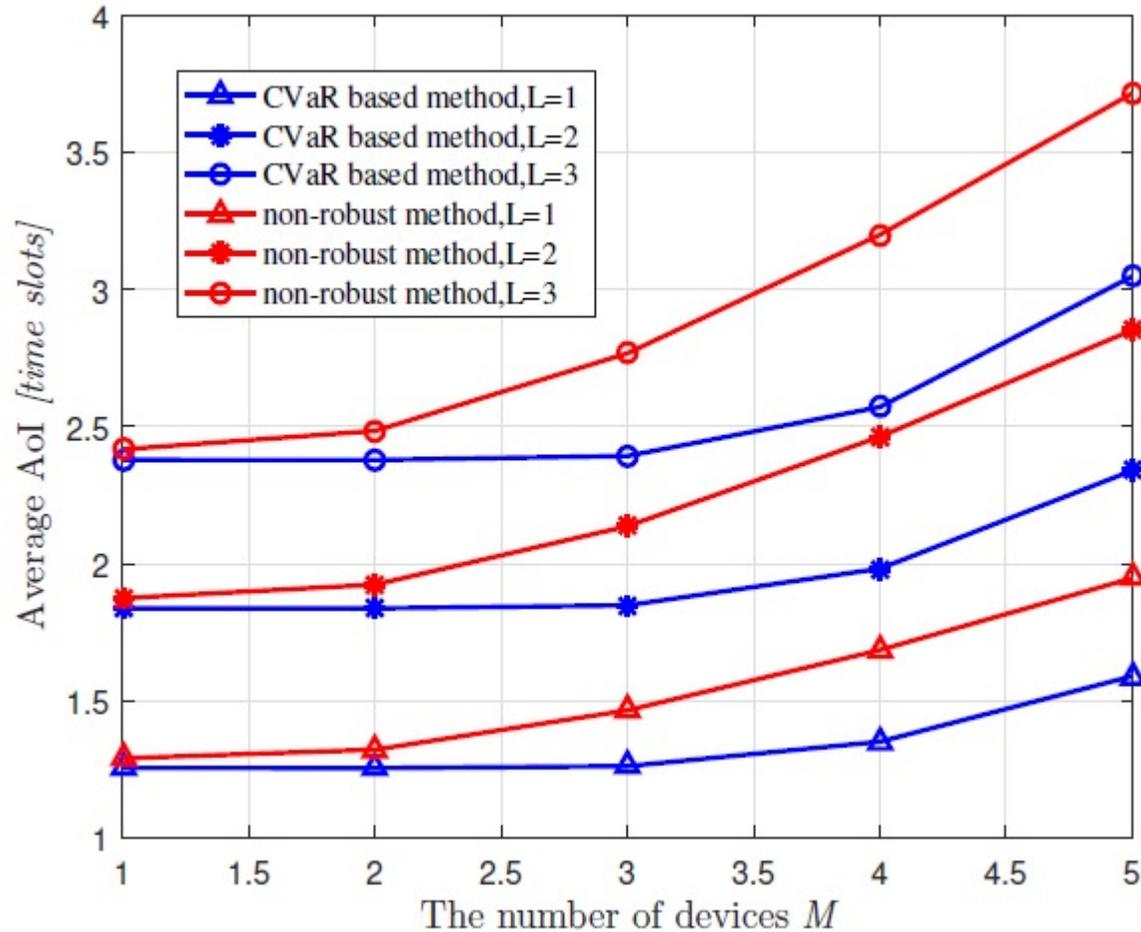
4: Solve the outage probability subproblem for given $\alpha^{\{j+1\}}$ and $\mathbf{p}^{\{j+1\}}$, to obtain the optimal solution as $\gamma^{\{j+1\}}$ and $\tau^{\{j+1\}}$.

5: **until** $|\sup_{\mathbb{P} \in \mathcal{P}} \bar{\Delta}_i(t)^{\{j+1\}} - \sup_{\mathbb{P} \in \mathcal{P}} \bar{\Delta}_i(t)^{\{j\}}|$ satisfies termination criterion Θ .

6: **return** $\mathbf{p}^{\{j-1\}}$, $\alpha^{\{j-1\}}$, $\gamma^{\{j-1\}}$, and $\tau^{\{j-1\}}$.

Output: The optimal solution α^* , \mathbf{p}^* , γ^* , and τ^* .

Application 2: Results



The energy harvesting and information transmission successful opportunities for each link become lower with more wearable IoT devices. In this sense, we find an **AoI-energy tradeoff** in the healthcare IoT system.

- Introduction
- Case 1: Discrepancy-based DRO Optimization
- Case 2: DRO, Risk Aversion, and CVaR
- **Case 3: DRO Based Reinforcement Learning**
 - ✓ **Distributionally Robust Policy Iteration**
 - ✓ **Application III: Well Logging**
- Conclusion

- Previous works considering robustness in Reinforcement Learning
 - Mainly focus on **uncertain environment** for searching strategies to achieve best performance
 - ✓ Robust reinforcement learning
 - Guard against common systematic perturbation situations
 - ✓ Robust Adversarial Reinforcement Learning
 - Conquer rare, catastrophic events

➤ DRO RL

- Not only consider uncertainty in environment
- Limited collected samples when training
- Affect how to transit to new state based on current state and action
 - ✓ **Estimation error in Policy Iteration**

$$\pi_{t+1} \in \mathcal{G}(\tilde{V}_t)$$
$$\tilde{V}_{t+1} = \mathcal{T}^{\pi_{t+1}} \tilde{V}_t + \delta_t \text{ Estimation error}$$

- Concentrate on learning process for agent itself
 - ✓ Conservative policy in unknown environment
 - ✓ Optimistic policy in familiar environment

Distributionally robust policy evaluation

➤ Uncertainty set:

$$\mathcal{U}_\epsilon(\pi) := \{ \tilde{\pi} \in \Delta_{\mathcal{A}}^S \mid D_{\text{KL}}(\tilde{\pi}(\cdot|s) \parallel \pi(\cdot|s)) \leq \epsilon(s), \forall s \in \mathcal{S} \}$$

Considered possible policies
Best policy

➤ Adversarial Bellman operator:

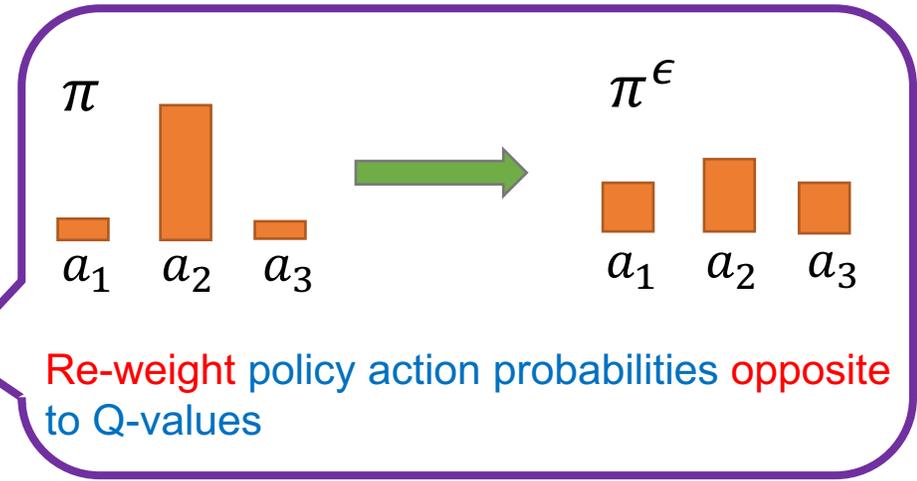
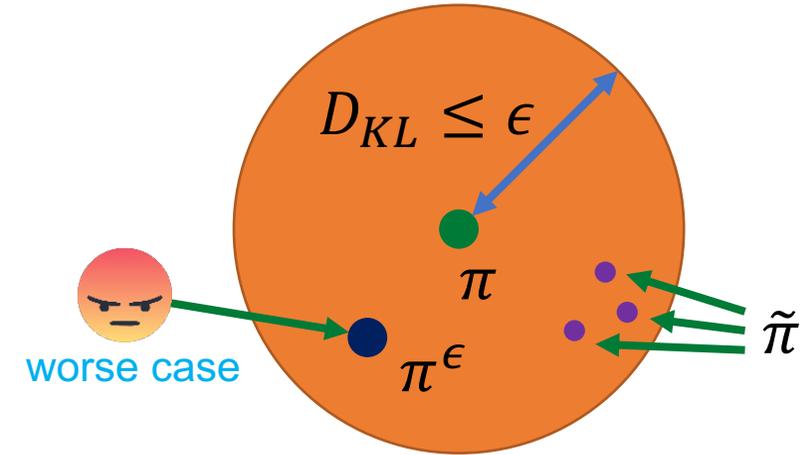
- Evaluate with an adversarial policy $\mathcal{T}^{\pi^\epsilon} V := \min_{\tilde{\pi} \in \mathcal{U}_\epsilon(\pi)} \mathcal{T}^{\tilde{\pi}} V$.

➤ Rewrite the maximum problem by

$$\begin{aligned} [\mathcal{T}^{\pi^\epsilon} V](s) &= \max_{\lambda(s) > 0} \min_{\tilde{\pi} \in \Delta(\mathcal{A})} [\mathcal{T}^{\tilde{\pi}} V](s) \\ &\quad + \lambda(s) D_{\text{KL}}(\tilde{\pi}(\cdot|s) \parallel \pi(\cdot|s)) - \lambda(s) \epsilon(s) \\ &= \min_{\lambda(s) > 0} \max_{\pi \in \Delta(\mathcal{A})} [-\mathcal{T}^\pi V](s) \\ &\quad - \lambda(s) D_{\text{KL}}(\tilde{\pi}(\cdot|s) \parallel \pi(\cdot|s)) + \lambda(s) \epsilon(s) \end{aligned}$$



$$\begin{aligned} \pi^\epsilon(a|s; \lambda) &\propto \exp(-Q_V(s, a) / \lambda(s)) \pi(a|s) \quad \text{adversarial policy} \\ \lambda^*(s) &:= \arg \min_{\lambda(s) > 0} \Omega^*(-Q_V / \lambda(s)) + \lambda(s) \epsilon(s) \quad \text{Adversarial temperature} \end{aligned}$$



➤ Distributionally robust modified policy iteration (MPI)

- Policy improvement: $\pi_{t+1} \leftarrow \mathcal{G}(\tilde{V}_t)$
- Policy evaluation: $\tilde{V}_{t+1} \leftarrow (\mathcal{T}^{\pi_{t+1}^\epsilon})^m \tilde{V}_t$

➤ *Rewritten of Bellman operator for value function*

- $[T^\pi V](s) := E_{a \sim \pi(\cdot|s)} [\sum_{t \geq 0} \gamma^t r_t | s_0 = s, \pi]$
 $= E_{a \sim \pi(\cdot|s)} [r(s, a) + \gamma E_{s' \sim p(s'|s, a)} [V(s')]]$ $Q_V(s, a)$
 $= \langle \pi(\cdot |s), Q_V(s, a) \rangle \dots$ inner product of the policy and Q-function

➤ *Regularized Bellman optimality operator*

- $[T^{\pi, \Omega} V](s) := T^\pi V - \Omega(\pi)$ is a **Fenchel dual function** and can be seen as the normal value function minus the baseline.
- *Set of optimal policies*

$$\blacksquare \mathcal{G}^\Omega(V) := \left\{ \pi : \pi \in \arg \max_{\pi \in \Pi} T^{\pi, \Omega} V = \nabla \Omega^*(Q_V) \right\}$$

➤ *Lengendre-Fenchel duality (Conjugate function)*

- Fenchel dual : $[\Omega^\star V](Q_V) = \max_{\pi \in \Delta A} (\langle \pi, Q_V \rangle - \Omega(\pi))$
- Gradient form: $[\nabla \Omega^\star V](Q_V) = \arg \max_{\pi \in \Delta A} (\langle \pi, Q_V \rangle - \Omega(\pi))$

➤ Soft distributionally robust modified policy

$$\Omega(\pi(\cdot|s)) = \alpha \mathcal{H}(\pi(\cdot|s)) \quad \forall s \in \mathcal{S}$$

➤ Distributionally robust soft actor-critic

$$\Omega_\lambda(\pi(\cdot|s)) = \lambda \times (-D_{\text{KL}}(\pi(\cdot|s) || \mu(\cdot|s)))$$

➤ Soft distributionally robust modified policy

- Policy improvement: $\pi_{t+1} \leftarrow \mathcal{G}^\Omega(\tilde{V}_t)$
- Policy evaluation: $\tilde{V}_{t+1} \leftarrow (\mathcal{T}^{\pi_{t+1}^\epsilon, \Omega})^m \tilde{V}_t$

➤ Soft adversarial Bellman operator

$$\mathcal{T}^{\pi^\epsilon, \Omega} V := \min_{\tilde{\pi} \in \mathcal{U}_\epsilon(\pi)} \mathcal{T}^{\tilde{\pi}, \Omega} V, \quad \Omega(\pi(\cdot|s)) = \underbrace{\alpha \mathcal{H}(\pi(\cdot|s))}_{\text{entropy form}} \quad \forall s \in \mathcal{S} \quad \mathbb{E}_{a \sim \pi(\cdot|s)} \log[\pi(a|s)]$$

- Here, we take **entropy form** into consideration to obtain **robustness guarantees on exploration process**.

➤ Adversarial entropy-regularized policy

$$\pi^\epsilon(a|s; \alpha, \lambda) \propto \exp((1/\alpha - 1/\lambda(s))Q_V(s, a))$$

- Recall the adversarial policy $\pi^\epsilon(a|s; \lambda) \propto \exp(-Q_V(s, a)/\lambda(s)) \underbrace{\pi(a|s)}_{\substack{\uparrow \\ \text{Regularized policy}}} \propto \exp(Q_V(s, a)/\alpha)$

Extended to Continuous Control

- Apply this distributional robust MPI to the **continuous control problem**.
- The formulation is intractable and we get the **approximate** formula as follows:

- Regularized Bellman operator can be written in terms of Fenchel conjugate

$$[\mathcal{T}^\Omega V](s) = \Omega^*(Q_V(s, \cdot))$$

- Regularization parameter

$$\Omega_\lambda(\pi(\cdot|s)) := \lambda \Omega(\pi(\cdot|s))$$

Prior policy

$$= \lambda \times (-D_{\text{KL}}(\pi(\cdot|s) || \underline{\mu}(\cdot|s)))$$

- Fenchel conjugate

$$\Omega_\lambda^*(Q_V(s, \cdot)) = \lambda \log \mathbb{E}_{a \sim \mu(\cdot|s)} \exp(Q_V(s, a)/\lambda) \quad \text{Logarithm of moment-generating function}$$

Taylor expansion ↩

$$= \mathbb{E}_{a \sim \mu} [Q_V(s, a)] + \frac{1}{2\lambda} \text{Var}_{a \sim \mu}(Q_V(s, a)) + \mathcal{O}\left(\frac{1}{\lambda^2}\right)$$

reward shaping

$$\simeq Q_V(s, a) + \frac{1}{2\lambda} \text{Var}_{a \sim \mu}(Q_V(s, a))$$

Potential-based
reward shaping ↩

$$= \mathbb{E}_{a \sim \mu, s' \sim p(\cdot|s, a)} [\underline{r}^\Omega(s, a, s') + \gamma V(s')] \quad \uparrow$$

$$\underline{r}^\Omega(s, a, s') := r(s, a) + \gamma \underline{\Phi}(s') - \underline{\Phi}(s)$$

$$\underline{\Phi}(s) := \frac{1}{2\lambda} \text{Var}_{a \sim \mu}(Q_V(s, a))$$

- When it come to the implementation, it's quite simple. ➡ Just change the reward
 - encourage the agent to visit states **with smaller variance**

➤ In this **continuous control case**: the distribution of the policy is **Normal distribution**.

- parametrized Gaussian policies

$$\pi_{\theta}(a|s) = \mathcal{N}(\mu_{\theta}(s), \Sigma_{\theta}^2(s))$$

- Approximation the variance of Q-values by using the 1st order **Taylor approximation of Q-values** around the mean action

$$\text{Var}_{a \sim \pi_{\theta}}(Q(s, a)) \simeq \underline{g_0(s)^T \Sigma_{\theta}(s) g_0(s)}$$

1st order Taylor approximation of Q-values around the mean action

Independent actions

$$\Sigma_{\theta}(s) = \text{diag}(\sigma_{1,\theta}(s), \dots, \sigma_{K,\theta}(s))$$

$$g_0(s) = \nabla_a Q(s, a)|_{a=\mu_{\theta}(s)}$$

$$\simeq \|g_0(s)\|_{\text{diag}(\sigma_{1,\theta}(s), \dots, \sigma_{K,\theta}(s))}^2$$

- Combined with the soft-actor critic algorithm

$$\pi^* = \underset{\pi}{\operatorname{argmax}} \sum_{t=0}^T E_{(s_t, a_t) \sim \tau_{\pi}} [\gamma^t (r(s_t, a_t) + \alpha \mathcal{H}(\pi(.|s_t)))]$$

$$\mathcal{H}(\pi(.|s_t)) = -\log \pi(.|s_t).$$

Importance: entropy V.S. reward

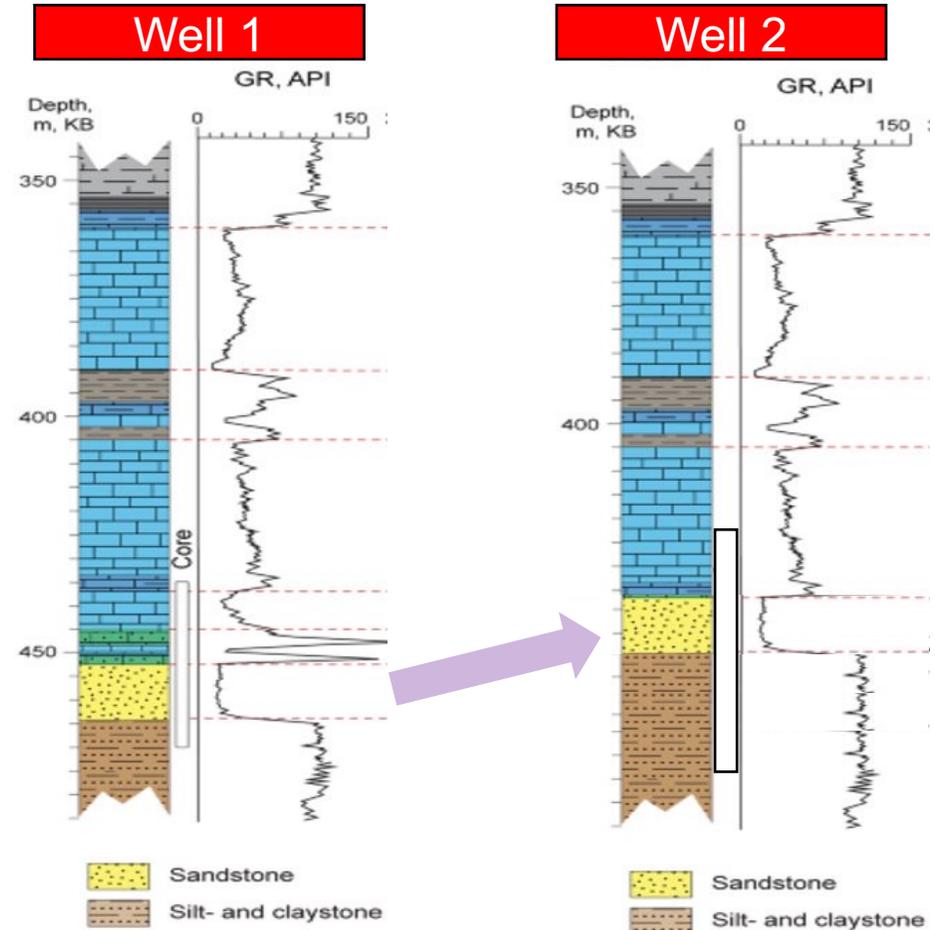
Application 3: Signal Localization

Gamma-Ray Signal Localization Task

Many 1D signals containing the same signal pattern are collected by the well-logging process and sent to human-expert to do hand-picking for oil&gas localization.

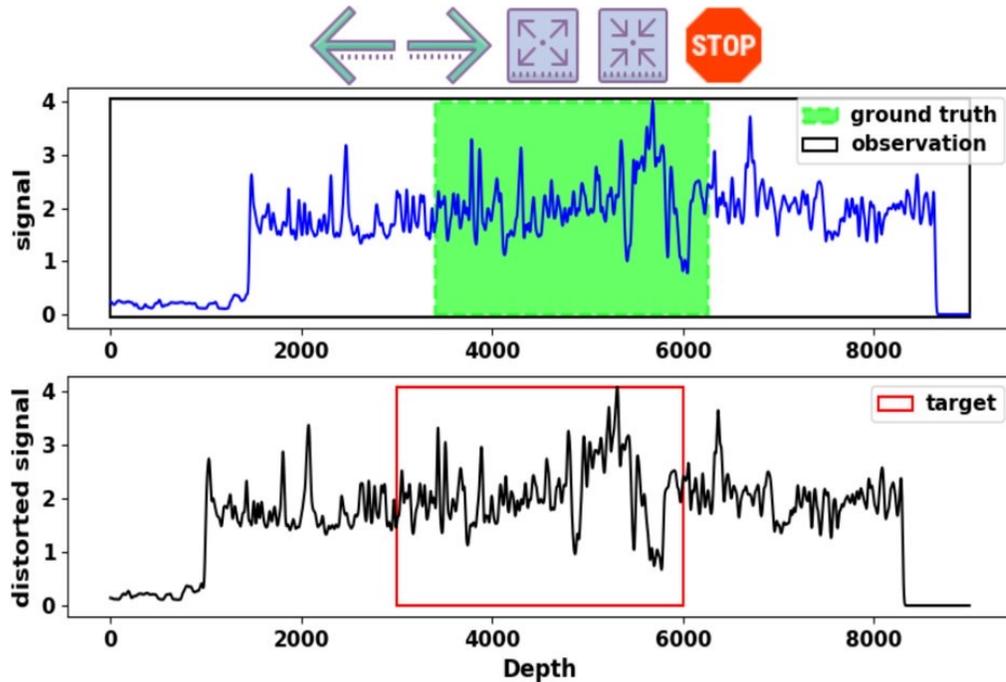
Traditional methods

- The human experts pick the patterns associated with interested rock based on one reference signal:
 - Recognize the pattern with domain prior knowledges.
 - Roughly matching with correlation matching methods.



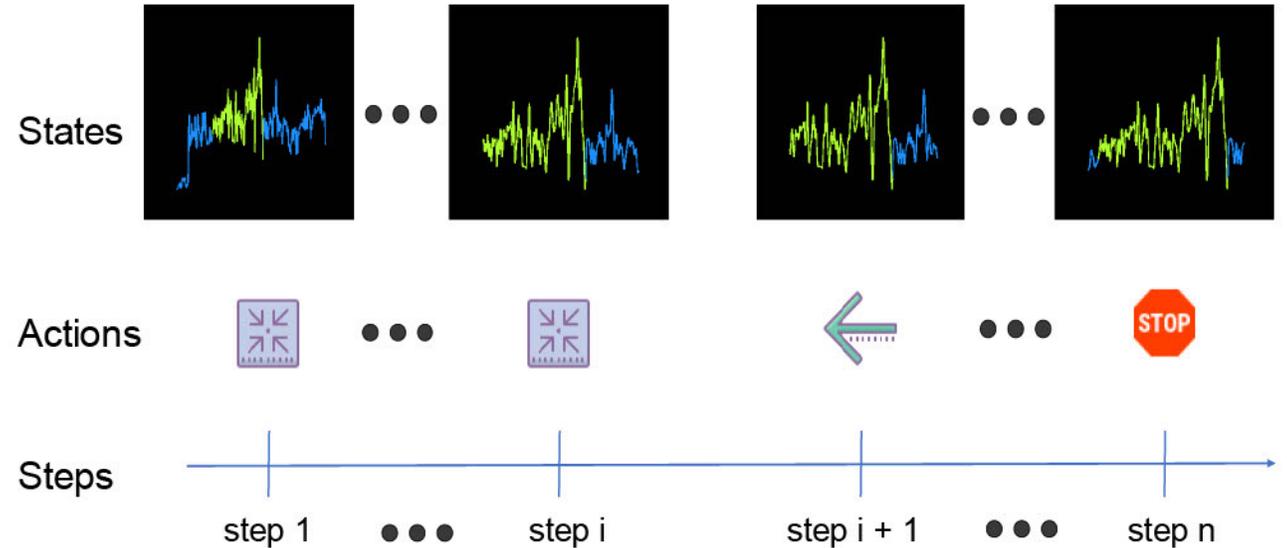
Yuan Zi, Lei Fan, Jiefu Chen, and Zhu Han, "Active Gamma-ray Log Pattern Localization with Distributional Robust Reinforcement Learning," submitted to IEEE Transactions on Neural Networks and Learning Systems.

Reinforcement Learning Signal Localization Scheme



- Thoroughly studied log with interested signal fragment as reference/target.
- In the new log, there is a signal fragment that has the same pattern as the reference

Sequence of attended regions to localize the object States



- Initial the whole new log trace as the agent's observation.
- Let the agent move (left, right, expand, shrink) to search the reference pattern.

- Motivation
 - DRO accounts for the fact that one is never able to exactly specify a **probability distribution** in practice.
 - Due to the **finite experience sample**, the distribution of the real experience is hard to get.
- DRO Optimization
 - The **Wasserstein distributionally robust problems** can often be reformulated as (or tightly approximated by) finite convex programs within a certain **Wasserstein distance** from the empirical distribution constructed from training samples.
- CVAR
 - **Distributionally robust chance constraints** can be conservatively approximated by **worst-case CVaR constraints**.
- DRO RL
 - The risk-averse exploration in approximate RL setting is required and we can use the distributionally robust modified policy iteration scheme that implements **safety in policy evaluation step w.r.t. estimation errors** to avoid.
 - Implement DRO to the Reinforcement Learning is just add one safe regularizer to **encourage agent visit the state with smaller variance**.

Main References



Rahimian H, and Mehrotra S.

Distributionally robust optimization: A review. 2019



Zymler S, Kuhn D, and Rustem B.

Distributionally robust joint chance constraints with second-order moment information. 2013



Esfahani P M, and Kuhn D.

Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. 2018



T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine.

Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. 2018



B. Scherrer, M. Ghavamzadeh, V. Gabillon, B. Lesner, and M. Geist.

Approximate modified policy iteration and its application to the game of tetris. Journal of Machine Learning Research, 16:1629–1676, 2015



E. Smirnova, E. Dohmatob, J. Mary.

Distributionally robust reinforcement learning. 2019

**Thank You Very Much
for
Your Attention !**

Conditional value-at-risk in optimization

In problems of optimization under **uncertainty**, CVaR can be used in the objective or the constraints, or both.

Optimization shortcut

$$\min_x CVaR_\alpha(x) = \min_{(x,\zeta)} F_\alpha(x, \zeta)$$

$$(x^*, \zeta^*) \in \underset{(x,\zeta) \in X \times \mathbb{R}}{\operatorname{argmin}} F_\alpha(x, \zeta) \iff x^* \in \underset{x \in X}{\operatorname{argmin}} CVaR(x) \quad \zeta^* \in \underset{\zeta \in \mathbb{R}}{\operatorname{argmin}} F_\alpha(x^*, \zeta)$$

CVaR accounts for losses exceeding VaR. **So when CVaR is considering, VaR can be ignored**