College of Natural Sciences and Mathematics
UNIVERSITY OF HOUSTON

ICLR

Distributionally Robust Optimization Via Generative Ambiguity Modeling

*Jiaqi Wen[1], Jianyi Yang**

UNIVERSITY OF HOUSTON

Motivation

- **Generalization For Generative Models**
  - **Reality Gap**
    - Generative environment ≠ Real-world
    - Unseen real-world environments introduce distribution shift
  - **Robustness Limitation**
    - Generative models fail to capture worst-case scenarios
  - **Deployment Challenge**
    - Critical barrier for safety-critical systems

Motivation

- **Limitation For Traditional DRO**

Distributionally Robust Optimization (DRO):

$$w = \min_{w \in \mathcal{W}} \max_{P \in \mathcal{B}(P_0, \epsilon)} \mathbb{E}_{x \sim P} [f(w, x)]$$

- $\boldsymbol{f(w, x)}$: the objective with the decision variable $w \in \mathcal{W}$ and input x
- $\boldsymbol{\mathcal{B}(P_0, \epsilon)}$: $\mathcal{B}(P_0, \epsilon) = \{ P \mid D(P, P_0) < \epsilon \}$ the ambiguity set with budget ϵ.
- $\boldsymbol{P_0(x)}$: the nominal distribution

- **ϕ-Divergence DRO**: Enforce absolute continuity $(P \ll P_0)$, preventing support shift and leading to weak robustness under distribution drift.

- **Wasserstein DRO:** Suffer from computational intractability and conservative approximations, limiting its effectiveness.

Motivation

- **Our Goal:**
  - Bridge generative models and DRO via Generative Ambiguity Modeling
  - Address the challenge of modeling expressive adversarial distributions
  - Maintain computational tractability

Generative Ambiguity Set (GAS)

ICLR

- **Generative Ambiguity Modeling**
  - Well-trained generative models approximate the nominal distribution, while their perturbations induce controlled deviations.
  - A generative model that can be used for generative ambiguity modeling if it satisfies:

Lemma 1 The ***inclusive*** **KL-divergence** between the nominal distribution P_0 and the sampling distribution P_θ of a likelihood-based generative model can be bounded as

$$D_{KL}(P_0||P_\theta) \leq J(\theta, P_0) + R(p', P_0) + C_1$$

- $\boldsymbol{J(\theta, P_0)}$: A reconstruction loss (e.g., Diffusion or VAE)
- $\boldsymbol{R(p', P_0)}$: A prior matching term independent of θ
- $\boldsymbol{C_1}$: A constant term unrelated to θ.

Generative Ambiguity Set (GAS)

ICLR

- **Modeling For Diffusion:**

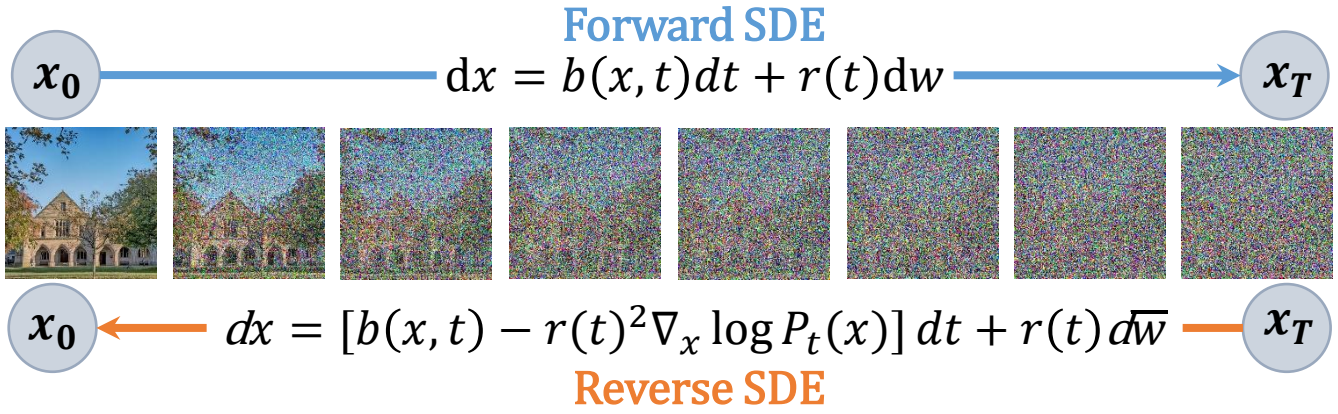

- **Score Matching Loss:** The reverse process relies on a model $s_\theta(x, t$, trained via a score-matching loss J_{SM} to approximate $\nabla_x \log P_t\,(x)$.

$$J_{SM}(\theta) = \int_0^T \mathbb{E}_{P_t(x)}[\iota(t)\|\nabla_x \log P_t(x) - s_\theta(x,t)\|^2]\mathrm{d}t$$

- ***Inclusive KL-divergence* Bound:**

$$D_{KL}(P_0(x_0)||P_\theta(x_0))$$
$$\leq J_{SM}(\theta, P_0) + \mathbb{E}_{P_0(x)}[D_{KL}[q(x_T|x_0)||p(x_T)]] + d\log\left(\sqrt{2\pi\sigma_1}\right) - H(P_0)$$

Generative Ambiguity Set (GAS)

ICLR

- **Modeling For VAE:**
  - **Reconstruction Loss:** The reconstruction loss encourages the learned model to regenerate data that is consistent with the original data distribution.

$$J_{VAE}(\theta, P_0) = \mathbb{E}_{x\sim P_0}[\mathbb{E}_{q_\varphi(z|x)}[-\ln p_\theta(x|z)]]$$

  - ***Inclusive KL-divergence* Bound:**

$$D_{KL}(P_0(x_0)||P_\theta(x_0))$$
$$\leq \mathbb{E}_{P_0(x)}\mathbb{E}_{q_\varphi(z|x)}[-\log_{p_\theta}(x|z)] + \mathbb{E}_{P_0(x)}\left[D_{KL}[q_\varphi(z|x)\|p'(z)]\right] - H(P_0(x))$$

Generative Ambiguity Set (GAS)

- **Design Goals of GAS**
  - **Expressive Robustness:** Capture adversarial distributions beyond the nominal support
  - **Tractable Optimization:** Enable efficient inner maximization in the generative parameter space
- **GAS-DRO Objective**

$$\min_{w \in \mathcal{W}} \max_{\theta \in \Theta} \mathbb{E}_{x \sim P_\theta} [f(w, x)], \qquad \text{s.t. } J(\theta, P_0) \leq \epsilon$$

GAS-DRO

- **Generative-Based Inner Maximization**
  - **Dual Learning:**
    - We solve the inner maximization by converting the constraint into a Lagrangian relaxed objective and updating the dual variable μ to enforce $J(\theta, P_0) \leq \epsilon$.

$$\max_{\theta} \mathbb{E}_{x \sim P_\theta}[f(w,x)] - \mu \cdot J(\theta, P_0) \qquad \mu \leftarrow \max\{0, \mu + \eta \cdot (J(\theta, P_0) - \epsilon)\}$$

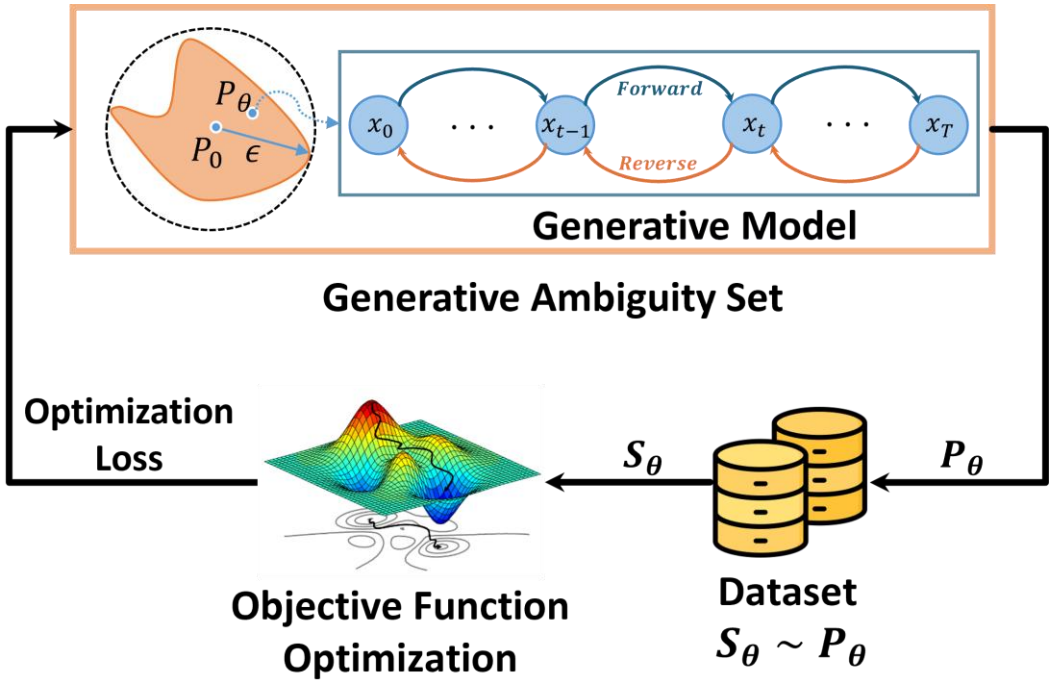

GAS-DRO

- **Generative-Based Inner Maximization**
  - **Policy Optimization:**
    - We apply Proximal Policy Optimization (PPO) to get a tractable form.

$$\max_{\theta} \widehat{\mathbb{E}}_{P_{\theta_{old}}(x_{0:T})}\left[\min\left(r_\theta(x_{0:T})f(w,x_0), \text{clip}(r_\theta(x_{0:T}), 1-k, 1+k)f(w,x_0)\right)\right] - \mu \cdot J(\theta, P_0)$$

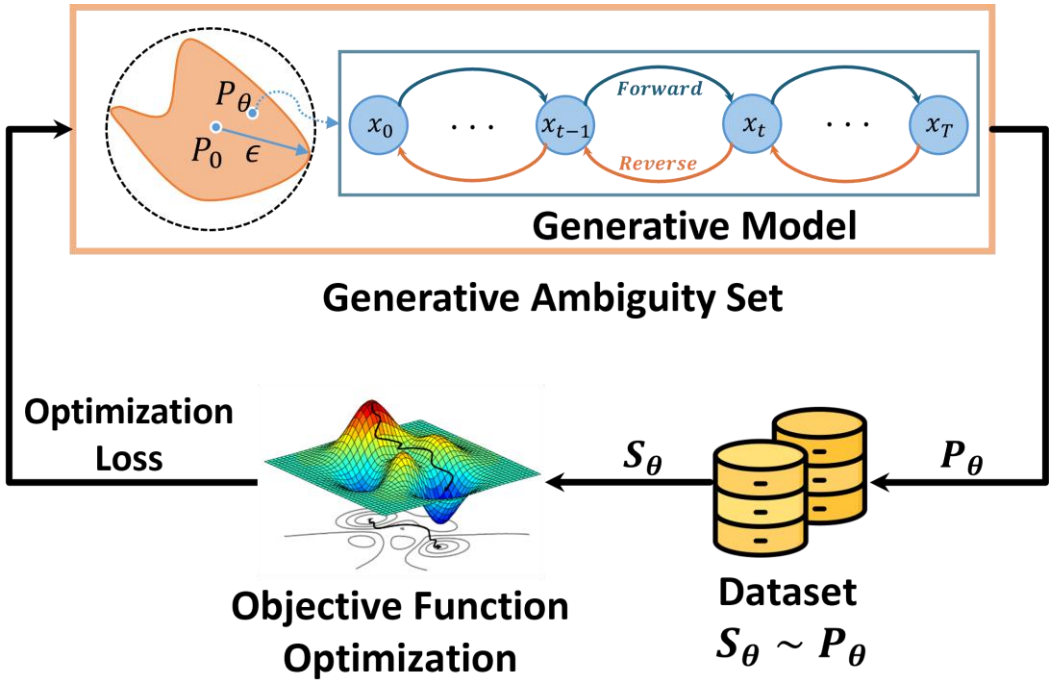

GAS-DRO

- **GAS-DRO Algorithm**
  - **Inner Maximization for Worst-Case Searching:**
    - Identify the worst-case generative model P_θ for the current w.
  - **Adversarial Update of w**
    - Generate adversarial samples from P_θ to update w, with convergence guaranteed.

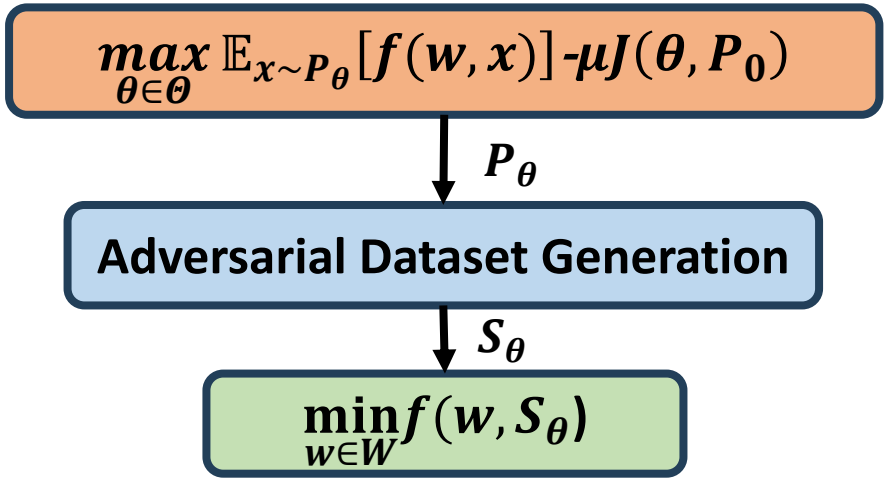

Convergence Analysis

ICLR

- **Convergence of Inner Maximization**

Theorem 1 [informal] Let P_{θ^*} be the optimal generative model for the inner maximization problem and P_{θ_k} be the generative model in the $k-th$ loop, the expected reward gap between P_{θ^*} and P_{θ_k} diminishes with K. The inner maximization error holds:

$$\Delta' := \mathbb{E}_{P_{\theta^*}}[f] - \mathbb{E}_{k\in[K]}\mathbb{E}_{P_{\theta_k}}[f] \sim O(\frac{1}{\sqrt{K}})$$

The reversed KL divergence between P_0 and P_{θ_k} is bounded:

$$\mathbb{E}_{k\in[K]}\, D_{KL}(P_0 \,\|\, P_{\theta_k}) \leq \epsilon + D_{KL}(P_T \| \pi) + C_1 + O(\frac{1}{\sqrt{K}})$$

- **Convergence**: With sufficient iterations K, the inner maximization error becomes negligible.
- **Consistent Worst-Case:**
  - $D_{\mathrm{KL}}(P_0 \parallel P_\theta)$is bounded by: ϵ, prior matching error $R(\pi, P_0)$,and constant C_1
  - GAS-DRO finds worst-case distributions that remain consistent with the nominal distribution via ϵ

Convergence Analysis

ICLR

- **Convergence of GAS-DRO**

Theorem 2 [informal] If $f(w,x)$ is $\beta-smooth$, $L-Lipschitz$ and is upper bounded by $\bar{f}$, GAS-DRO converges to a near-stationary point of the smoothed robust objective $\phi(w) = \max_{\theta} \mathbb{E}_{P_\theta}[f(w,x)]$. The average norm of Moreau envelope of $\phi(w)$ satisfies:

$$\mathbb{E}_{j,k}\left[\left\|\nabla\phi_{\frac{1}{2\beta}}(w)\right\|^2\right] \leq 4\beta\Delta' + \frac{V_1}{\sqrt{n}} + \frac{V_2}{\sqrt{H}}$$

$$V_1 = 8\beta\bar{f}\sqrt{\log(2/\delta)} \qquad V_2 = 4L\sqrt{(\phi_{\frac{1}{2\beta}}(w_1) - \min_w \phi(w))\beta}$$

- **Bounded Gradient**: With large outer minimization iteration H and sample size n, the gradient norm $\| \nabla\phi_{1/(2\beta)}(w) \|$is controlled by the inner error Δ'.
- **Vanishing Error:** As inner maximization iteration $K \uparrow$, $\Delta' \downarrow$ ⇒gradient norm becomes sufficiently small.
- **Approximate Stationarity:** Small gradient ⇒ existence of nearby $\hat{w}$ that is an approximately stationary solution

Evaluation

ICLR

- **Task 1 – Carbon Intensity Time Series Prediction**
  - **Datasets:** Original & Corrupted carbon-intensity data from Electricity Maps.
  - **Results:**
    - All methods outperform ML, with GAS-DRO achieving the largest gain (63.7%), followed by DRAGEN (48.9%), P-DRO (42.4%), and DML (39.7%).
    - KL-DRO and W-DRO are limited by support constraints or overly conservative adversaries.
    - The ablation study on GAS-DRO, DML, and ML shows that the diffusion contribute a 39.7% performance gain, while DRO provides an additional 24% improvement.

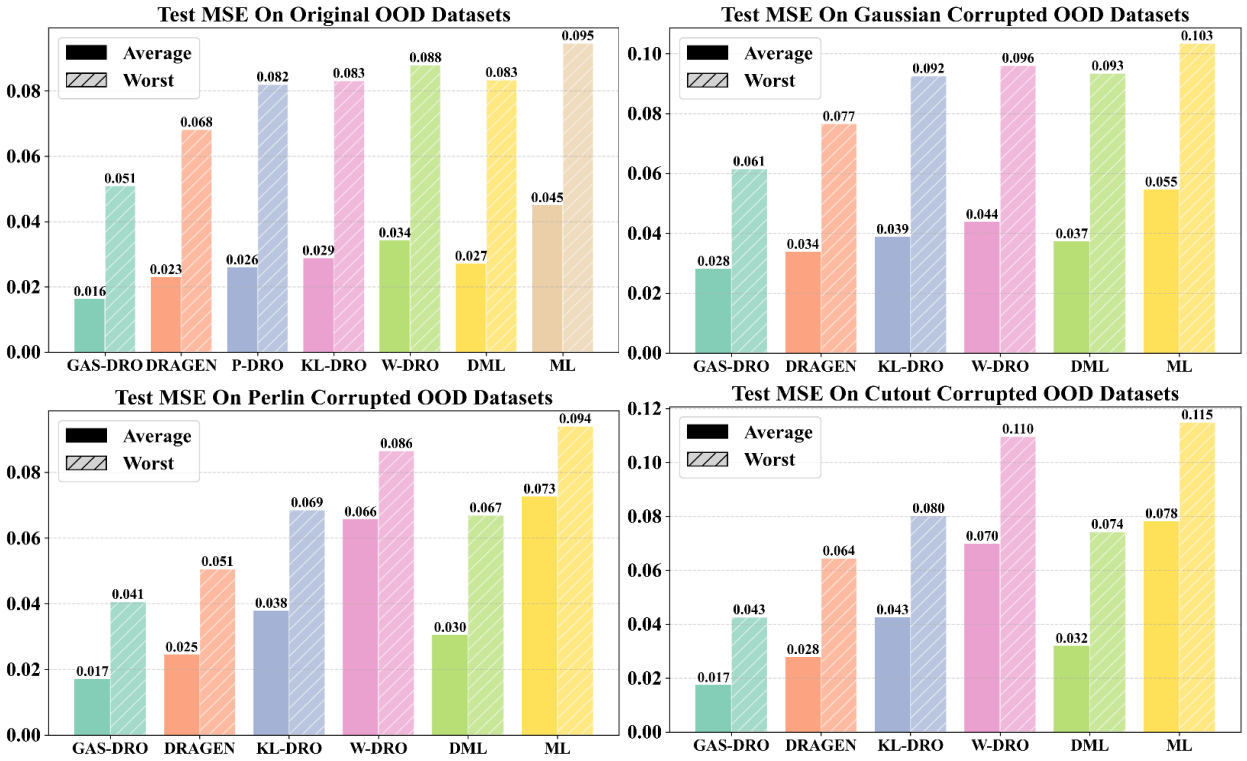

Evaluation

- **Task 1 – Carbon Intensity Time Series Prediction**
  - **Datasets:** Original & Corrupted carbon-intensity data from Electricity Maps.
  - **Results:**
    - GAS-DRO consistently outperforms all baseline across varying levels of Gaussian noise corrupted datasets
    - GAS-DRO performs best at our chosen budget of $\epsilon = 0.015$, which strikes the optimal balance—smaller values hurt generalization while larger ones lead to overly conservative optimization.

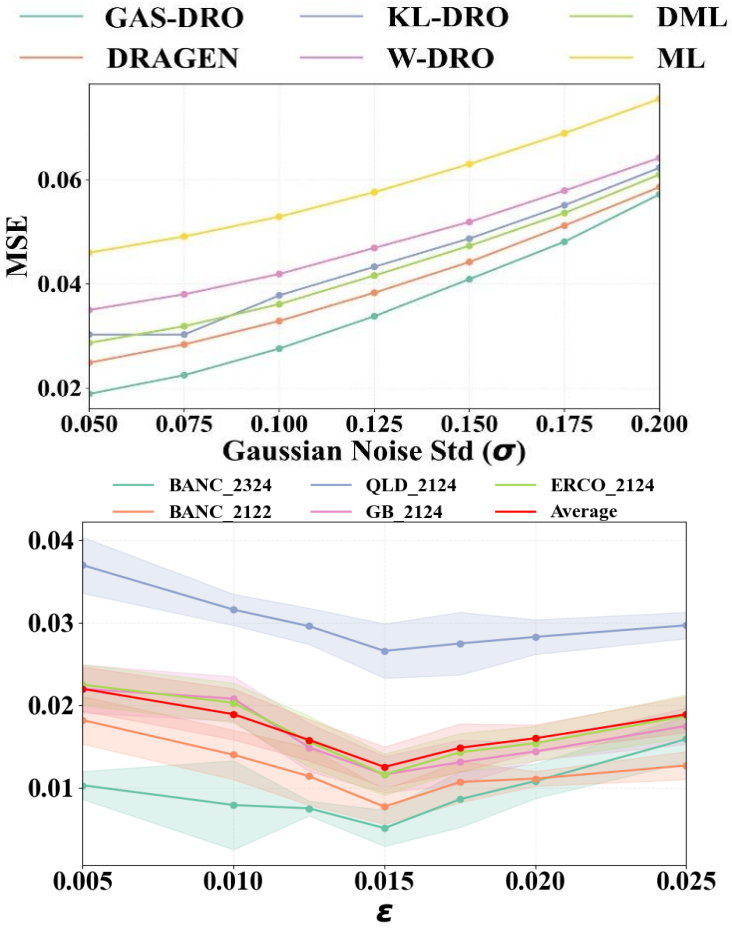

Evaluation

ICLR

- **Task 2 – Handwritten Digit Classification**
  - **Datasets:** Original & Corrupted MINST & USPS datasets.

    Note: Here, GAS-DRO employs a VAE as the generative model, the same as DRAGEN.
  - **Results:**
    - GAS-DRO significantly outperforms all baselines, demonstrating strong generalization and robustness.
    - GAS-DRO outperforms prior methods even with the same VAE backbone as DRAGEN, demonstrating that its gains stem from ambiguity set design rather than the generative model itself.

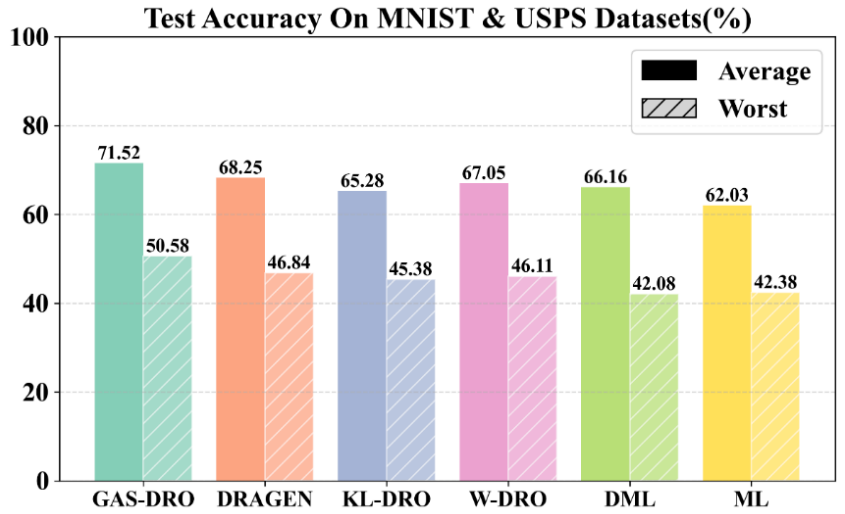

Conclusion

We propose GAS-DRO, which models ambiguity sets in the parameter space of generative models:

- It enable consistent, expressive, and tractable optimization.
- It explores diverse worst-case distributions with theoretical guarantees and strong performance.

➤ Advantages

- **Capture the data distribution, ensuring consistency with the nominal distribution.**
- **Generate diverse samples beyond the nominal support to discovery worst-case distribution.**
- **Enables tractable optimization in a finite, parameterized space via policy optimization.**

College of Natural Sciences and Mathematics
UNIVERSITY OF HOUSTON

- Thank You -

UNIVERSITY OF HOUSTON

Paper Link