Learning Scenario Reduction for Two-Stage Robust Optimization with Discrete Uncertainty

Tianjue Lin[1], Jianan Zhou[1], Jieyi Bi[1], Yaoxin Wu[2], Wen Song[3], Zhiguang Cao[4], Jie Zhang[1]

[1] Nanyang Technological University [2] Eindhoven University of Technology
[3] Shandong University [4] Singapore Management University

Two-Stage Robust Optimization

- **Uncertain parameters:** Parameters such as travel times and costs vary with weather and traffic, so their values are uncertain when decisions are made.
- **Decision risk:** A decision that works for the expected parameter values may become costly or even infeasible when the actual values differ.
- **Robust protection:** Robust optimization guards against the worst case within an uncertainty set, but fixing all decisions upfront can be conservative.
- **Two-stage flexibility:** Some decisions must be made in advance, while others can adapt once the uncertain parameters are observed. This flexibility can reduce conservatism.

Motivation

- **Computational challenge:** With more scenarios, the solver must consider more outcomes and recourse decisions, which increases computation time.
- **Speeding up via scenario reduction:** Retain fewer scenarios to solve a smaller problem.
- **Prior work:** Problem-agnostic reduction methods use scenario data without accounting for the optimization objective, constraints or recourse.
- **Question:** How can we select scenarios so the smaller problem preserves solution quality?
- **Key idea:** Select scenarios based on their impact on the optimization objective.
- **Machine learning:** Train a model to predict high-impact scenarios at lower selection cost.

Contributions

- **PRISE:** Achieves a smaller optimality gap than non-oracle baselines in most settings.
- **NeurPRISE:** Achieves 17–200× speedup over PRISE by learning to select scenarios.
- **Performance & Generalization:** Experiments on Selection (SEL), Vertex Cover (VC) and Capacitated Facility Location (CFLP) show competitive quality at a fraction of PRISE's cost, generalizing to larger instances, more scenarios and new distributions.
- **Lower-cost supervision:** Higher MIP gap tolerances make labels cheaper to generate, and NeurPRISE remains effective with noisy gain scores.

A Few Scenarios Can Be Enough

- **Stopping steps:** PRISE usually selects just a few scenarios, one per step.
- **Scenario compression:** Each bar shows the average percentage of scenarios used to target a 1% gap to the full objective.

Figure 1. PRISE stopping steps.

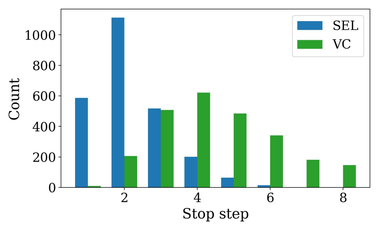
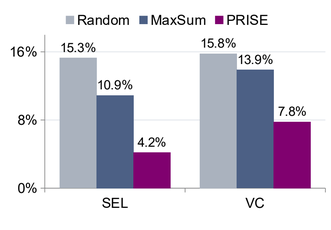


Figure 2. Scenario compression.



Methodology

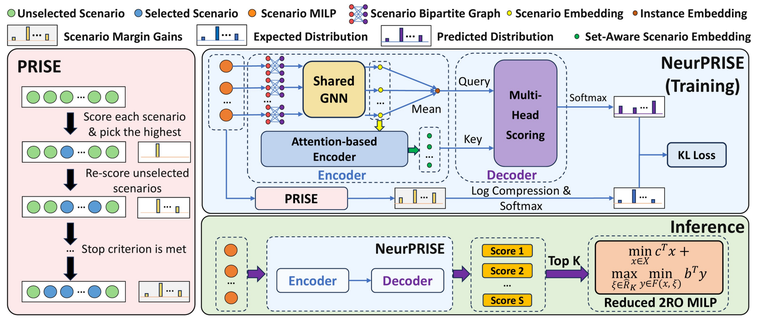


Figure 3. Overview of PRISE and NeurPRISE.

PRISE

Input: scenario set $\Xi=\{\xi_1,\ldots,\xi_S\}$ with $S=|\Xi|$; maximum budget K; tolerance $\epsilon\geq 0$
Output: compression budget $\hat{k}$; reduced set R; supervision data $\mathcal{D}$
Value function: $V(R)$ is the optimal objective value of the two-stage robust problem restricted to scenarios in R.
Function PRISE(Ξ, K, ϵ):
  Initialize $R \leftarrow \emptyset$, $\mathcal{D} \leftarrow \emptyset$, and $v_{\text{prev}} \leftarrow 0$
  while $|R| < K$ **do**
    for each unselected scenario $\xi_j \in \Xi \setminus R$ **do**
      $\text{score}(\xi_j) \leftarrow V(R \cup \{\xi_j\})$
    end for
    $\xi^\star \leftarrow$ scenario with highest score
    $\Delta \leftarrow \text{score}(\xi^\star) - v_{\text{prev}}$
    if $\Delta \leq \epsilon$ **then**
      break
    end if
    Add $(R, \xi^\star, \Delta)$ to supervision data $\mathcal{D}$
    Add $\xi^\star$ to the reduced set: $R \leftarrow R \cup \{\xi^\star\}$
    $v_{\text{prev}} \leftarrow \text{score}(\xi^\star)$
  end while
  $\hat{k} \leftarrow |R|$
  return $\hat{k}$, R, $\mathcal{D}$

- **Intuition:** PRISE selects the scenario that most increases the reduced problem's optimal objective. A larger increase suggests a greater impact on worst-case cost.

NeurPRISE

- **Representation:** A shared GNN embeds each scenario's problem structure.
- **Scoring:** Attention refines keys; a mean-pooled GNN query yields multi-head scores.
- **Training:** Score-Weighted KL loss learns from PRISE's log-compressed gains.

Learning from Cheaper, Noisy Labels

PRISE labeling	Time saved (%) ↑	Gap (%) ↓
Original tolerance	0.0	3.49
2% MIP gap	19.7	3.64
5% MIP gap	**63.7**	3.73

Table 1. VC labeling cost (6 selected scenarios).

Training supervision	Gap (%) ↓
Original supervision	2.68
Noisy gain scores	2.64
Random scenario replacement	7.73

Table 2. Noisy labels (VC small, 6 selected scenarios).

- **Labeling:** A 5% MIP gap saves 63.7% of label time with a minor loss in solution quality.
- **Robustness:** Gain noise has little effect, but replacing scenarios increases the gap.

Experiments

Performance

k = number of selected scenarios

Method	Small — 20 items/nodes, 50 scenarios								Medium — 50 items/nodes, 50 scenarios							
	SEL · Gap (%) ↓				VC · Gap (%) ↓				SEL · Gap (%) ↓				VC · Gap (%) ↓			
	k=1	k=2	k=4	k=6	k=1	k=2	k=4	k=6	k=1	k=2	k=4	k=6	k=1	k=2	k=4	k=6
Exact	0.00	—	—	—	0.00	—	—	—	0.00	—	—	—	0.00	—	—	—
PRISE	5.93	2.14	0.92	0.82	13.85	7.91	3.00	1.04	6.64	2.11	0.90	0.59	16.35	9.17	3.83	2.28
Random	22.11	16.13	12.29	9.72	16.11	12.49	10.58	9.65	13.56	9.74	8.98	7.77	17.10	12.51	10.15	9.39
SOR	5.45	4.57	3.44	2.83	11.15	8.52	6.71	5.63	8.01	6.00	5.37	4.91	13.77	11.39	9.06	8.40
MaxSum	6.93	4.50	2.47	1.96	11.75	**8.25**	5.76	4.40	6.66	6.27	4.24	3.33	13.65	**9.56**	6.90	5.20
GK-RF	**4.00**	2.73	1.74	1.16	14.43	9.45	5.00	3.29	4.62	**3.91**	2.36	2.09	16.57	9.79	**6.16**	4.50
Neur2RO	5.88	—	—	—	**11.13**	—	—	—	5.56	—	—	—	**12.87**	—	—	—
NeurPRISE	5.59	**2.18**	**0.79**	**0.42**	14.75	9.36	**4.20**	**2.34**	**4.38**	4.12	**2.19**	**1.76**	15.26	9.69	6.50	**3.65**

Table 3. Optimality gap comparison on small and medium instances.

Generalization

Gap (%) ↓		5× size			4× scenarios	
Method	k	SEL	VC	k	SEL	VC
Neur2RO	—	6.85	14.79	—	5.09	10.52
	—			—		
MaxSum	4	3.47	7.06	4	2.84	5.59
	6	2.96	5.76	8	1.99	3.56
NeurPRISE	4	**2.12**	**5.89**	4	**1.24**	**4.78**
	6	**0.81**	**4.20**	8	**0.55**	**2.48**

Table 4. Scale generalization.

Gap (%) ↓		Normal		Multimodal	
Model	k	SEL	VC	SEL	VC
Neur2RO	—	2.96	4.69	1.55	7.08
	—				
MaxSum	4	2.13	**2.35**	1.08	5.19
	6	1.62	**0.70**	0.77	4.18
NeurPRISE: Uniform-trained	4	**1.89**	2.39	**0.46**	**4.66**
	6	**1.56**	0.89	**0.28**	**3.84**

Table 5. Distribution shift.

Faster Solutions, Slightly Lower Quality

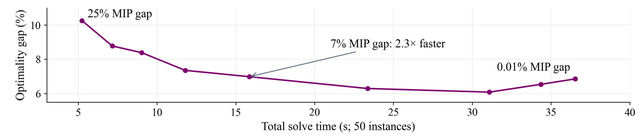


Figure 4. Solve time and optimality gap under different MIP gap tolerances.

- After scenario selection, relaxing the solver's MIP gap tolerance can shorten solve time with a small loss in solution quality.

Summary

Takeaway: NeurPRISE maintains solution quality with just a few scenarios, even on larger problems and new distributions. It also learns well from less precise labels.

Limitations: Generating training labels still takes time. Keeping more scenarios also makes the final optimization problem slower to solve.

Contacts & Links

Email: tianjue002@e.ntu.edu.sg github.com/lintianjue/NeurPRISE

WeChat WhatsApp LinkedIn