Conformalized Decision Risk Assessment

Wenbin Zhou[1] **Agni Orfanoudaki**[2] **Shixiang Zhu**[1]
[1]Carnegie Mellon University [2]University of Oxford
`{wenbinz2, shixianz}@andrew.cmu.edu`
`agni.orfanoudaki@sbs.ox.ac.uk`

Abstract

We introduce `CREDO`, a framework that provides distribution-free upper bounds on the probability that any candidate decision is suboptimal. By combining inverse optimization geometry with conformal prediction and generative modeling, `CREDO` generates statistically rigorous and interpretable risk certificates. This enables decision-makers to audit and validate choices under uncertainty, bridging the gap between decision-making algorithms and human judgment.

1 Introduction

High-stakes decision-making often arises in domains such as healthcare, energy systems, emergency response, and public policy, where suboptimal decisions can lead to significant consequences [24]. In the machine learning and operations research communities, such problems are typically formalized using the *predict-then-optimize* paradigm [5, 10]: a predictive model first estimates uncertain inputs (*e.g.*, future demand or patient outcomes), and these estimates are then passed to an optimization model to determine the recommended action. This two-stage pipeline has become the foundation of many data-driven decision systems across a wide range of applications [4, 9, 19].

However, despite its widespread use, predict-then-optimize approaches have two critical limitations in high-stakes operational settings. First, these pipelines often function as black boxes that recommend decisions without providing insight into their robustness or sensitivity to the estimated parameters. This opacity makes it difficult for human decision-makers to assess confidence in the recommendations or understand when to override them with domain expertise [12, 23]. Second, and more fundamentally, pipelines can be brittle under uncertainty [16]. When the predictive model is misspecified, especially in cases involving multi-modal or non-Gaussian distributions, the optimization step may yield misleading or overconfident decisions.

In this paper, we aim to bridge the gap between human decision-makers and optimization models by introducing a tool that supports – not replaces – human judgment in high-stakes decision-making. Rather than prescribing an optimal action, we focus on evaluating the risk associated with a user-specified decision: *how likely is it to remain optimal under the true, unknown realization of uncertain parameters* (Figure 1)? This form of optimal solution auditing is valuable in practice. For instance, a clinician may wish to validate a preferred treatment plan based on historical data, or a planner may want to assess whether a proposed infrastructure investment remains robust across multiple future scenarios, maintaining control over the final decision.

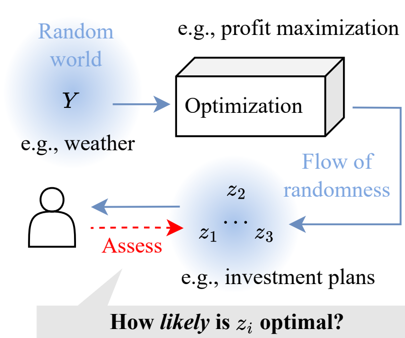

Figure 1: Decision risk assessment.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: MLxOR: Mathematical Foundations and Operational Integration of Machine Learning for Uncertainty-Aware Decision-Making.

To achieve this, this paper proposes a novel framework called `CREDO` – Conformalized Risk Estimation for Decision Optimization – that quantifies, for any candidate decision, a distribution-free upper bound on the probability that it is suboptimal, as illustrated in Figure 2. `CREDO` aligns with the spirit of scenario-based planning, a common practice in strategic investment and policy-making, while extending it with rigorous statistical guarantees. Our approach builds on two key insights: (i) in a broad class of optimization programs, the optimal decision is a deterministic function of the objective parameters [8, 13], which enables us to invert this relationship and define the precise set of outcomes for which a given decision is optimal; and (ii) using conformal prediction with generative models [15, 21], we estimate the probability mass of this region, yielding valid, data-driven upper bounds on decision risk with statistical guarantees. The resulting tool is interpretable, computationally efficient, and broadly applicable in high-stakes, uncertainty-aware decision environments.

2 Methodology

Problem Setup Let $X \in \mathcal{X}$ denote observed covariates, and let $Y \in \mathcal{Y}$ be a random outcome variable representing uncertain objective parameters. The decision-making task is formalized as solving the constrained optimization problem:

$$\pi(Y;\theta) \coloneqq \arg\min_{z \in \mathcal{Z}} \{g(z, Y; \theta) \mid z \in \mathcal{Z}(\theta)\},$$

where θ denotes known parameters that define the objective g and the feasible region $\mathcal{Z}(\theta) \subseteq \mathcal{Z}$. Our goal is to develop a distribution-free, data-driven method for estimating the probability that a prescribed decision z is suboptimal. Specifically, given a candidate decision z, we aim to develop a risk measure $\alpha(z)$ such that:

$$\mathbb{P}\{z \in \pi(Y;\theta)\} \geq 1 - \alpha(z), \quad \forall z \in \mathcal{Z}, \tag{1}$$

where we omit the dependency of $\alpha(z)$ on x for notational simplicity. Namely, given some context x, we aim to find a conservative assessment of the likelihood that the proposed decision z would be suboptimal with respect to the given decision-making task under uncertainty.

2.1 `CREDO`: Conformalized Decision Risk Assessment

We propose a framework for quantifying decision risk $\alpha(z)$ by casting this problem as a structured uncertainty quantification task over the decision space. The key idea is to reframe the decision's optimality condition as an *inverse feasible region* [8, 18] in the space of outcome Y, and then construct inner geometric approximations of this region using *conformal calibrated* regions [15, 21]. Their miscoverage rates are then used as the desired probabilistic lower bound. We detail this procedure in two steps: (i) identifying the region of outcome realizations that certify the optimality of z, and (ii) constructing a calibrated estimator for the probability that the outcome lies in this set.

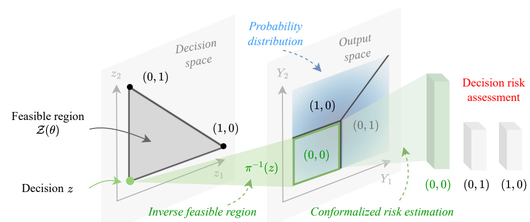

Figure 2: Illustration of the algorithm architecture.

Step 1: Reformulation with Inverse Feasible Region For any given realization y of Y, a decision z is optimal if and only if it achieves the smallest objective value among all feasible decisions:

$$z \in \pi(y;\theta) \iff g(z, y; \theta) \leq g(z', y; \theta), \quad \forall z' \in \mathcal{Z}(\theta).$$

Therefore, the inverse feasible region, which is defined as the set of outcomes y for which z is an optimal solution, can be written as

$$\pi^{-1}(z;\theta) \coloneqq \bigcap_{z' \in \mathcal{Z}(\theta)} \{y \in \mathcal{Y} \mid g(z, y; \theta) \leq g(z', y; \theta)\}. \tag{2}$$

The definition (2) allows the objective to be reformulated, as summarized in Proposition 1.

Proposition 1 (Reformulation)**.** *Let $\pi^{-1}(z;\theta)$ be defined in* (2)*, then the objective defined in* (1) *can be equivalently expressed as:*

$$\mathbb{P}\{z \in \pi(Y;\theta)\} \equiv \mathbb{P}\left\{Y \in \pi^{-1}(z;\theta)\right\}. \tag{3}$$

The important insight from Proposition 1 is that the reformulated objective decouples the random variable Y from the solution mapping π. Therefore, further analysis can be carried out within the probability space of the output as a standardized uncertainty quantification problem.

Step 2: Risk Estimation via Generative Conformal Prediction To estimate (3), we adopt a generative conformal prediction approach that constructs calibrated inner approximation sets of $\pi^{-1}(z;\theta)$. We require two core conditions to be satisfied by each constructed set: (i) It is fully contained within $\pi^{-1}(z;\theta)$, and (ii) Its coverage probability of Y can be quantified. This allows us to use this constructed set as a surrogate to bound the reformulated decision risk by:

$$\mathbb{P}\left\{Y \in \pi^{-1}(z;\theta)\right\} \overset{(i)}{\geq} \mathbb{P}\left\{Y \in \mathcal{C}(X;\alpha)\right\} \overset{(ii)}{\geq} 1-\alpha,$$

where $\mathcal{C}(X;\alpha)$ denotes the constructed set given input X, and α its coverage probability.

We begin the construction of these sets by training a generative model $\hat{f}:\mathcal{X}\to\mathcal{Y}$ to approximate the conditional distribution of $Y|X$. For any $\alpha\in[0,1]$, we define the conformalized set as:

$$\mathcal{C}(x_{n+1};\alpha) = \left\{y \in \mathcal{Y} \mid \|y-\widehat{y}_{n+1}\|_2 \leq \frac{\sum_{i=1}^{n}\|\hat{y}_i - y_i\|_2}{\alpha(n+1)-1}\right\} \text{ if } \alpha \geq \frac{1}{n+1} \text{ else } \mathcal{Y}, \tag{4}$$

where $\hat{y}_i \sim \hat{f}(x_i)$ are the in-sample prediction residuals on the calibration dataset $\{(x_i,y_i)\}_{i=1}^n$. Here, the miscoverage rate α is identified as the smallest value that would allow the set to lie entirely within the inverse feasible region [17], which can be formulated as solving:

$$\hat{\alpha}(z) = \min_{\alpha\in[0,1]} \left\{\alpha \mid \mathcal{C}(x_{n+1};\alpha) \subseteq \pi^{-1}(z;\theta)\right\}. \tag{5}$$

In the case that the prediction $\widehat{y}_{n+1}$ does not fall inside the inverse feasible region $\pi^{-1}(z;\theta)$, we conservatively map $\hat{\alpha}(z)$ to one, as the optimization in (5) yields no solution. The above procedure is then repeated for K times to obtain a list of risk estimates $\hat{\alpha}^{(k)}(z)$, which are then averaged to produce the final risk estimator:

$$\hat{\alpha}(z) = \frac{1}{K}\sum_{k=1}^{K}\hat{\alpha}^{(k)}(z). \tag{6}$$

3 Theoretical Results

This section establishes the theoretical properties of our method, focusing on two key criteria: *conservativeness*, which ensures a valid upper bound of the decision risk, and *computational efficiency*, which ensures that the algorithm can be efficiently implemented in practice under certain conditions.

First, we can prove that our method yields a valid upper bound on the decision risk in expectation, under the standard exchangeability assumption [2, 3, 14] in the conformal prediction literature, summarized in Theorem 1. Its proof relies on a key argument that the designed conformalized set is constructed by an e-value statistic [20, 11], which ensures proper coverage of the set even when α is data-dependent. This result underscores the trustworthiness of our approach, particularly for applications where conservativeness is critical.

Theorem 1 (Conservativeness)**.** *Suppose the data $\{(x_i,y_i)\}_{i=1}^{n+1}$ is exchangeable, then the proposed risk estimator* (6) *satisfies the following guarantee:*

$$\mathbb{P}\left\{z \in \pi(Y;\theta)\right\} \geq 1 - \mathbb{E}\left[\hat{\alpha}(z)\right], \quad \forall z \in \mathcal{Z}.$$

Next, we show that our algorithm admits a closed-form estimator for a broad class of optimization problems, enabling efficient practical implementation. We first state the following assumption.

Assumption 1 (Separable objective)**.** *There exist $\phi:\mathcal{Y}\to\mathbb{R}^d$ and $\psi:\mathcal{Z}\to\mathbb{R}^d$ such that for any $z\in\mathcal{Z}$ and $y\in\mathcal{Y}$, the objective function can be written as $g(z,y;\theta) = \langle\phi(y),\psi(z)\rangle$.*

Assumption 1 covers a wide range of problems such as linear programming, quadratic programming, and discrete optimization [7, 22]. Under Assumption 1, we prove that a closed-form estimator exists, enabling efficient algorithm execution, without relying on iterative or approximate procedures.

Proposition 2 (Closed-form estimator)**.** *Under Assumption 1, the algorithm can be modified to yield the following closed-form estimator:*

$$\hat{\alpha}(z) = 1 - \frac{1}{K}\sum_{k=1}^{K}\left(w^{(k)}(z,x_{n+1}) \cdot \prod_{v\in\mathcal{V}(\theta)} \mathbb{1}\left\{\langle\phi(\hat{y}_{n+1}^{(k)}), \psi(z)-v\rangle \leq 0\right\}\right), \tag{7}$$

where $\mathcal{V}(\theta)$ is the set of extreme points of the convex hull of $\{\psi(z) \mid z \in \mathcal{Z}(\theta)\}$, and $w(z, x_{n+1})$ is the conformalized weight, defined as

$$w^{(k)}(z, x_{n+1}) = \min \left\{ 0, \max \left\{ \frac{n\hat{D}^{(k)} - \sum_{i=1}^{n} \|\phi(y_i) - \phi(\hat{y}_i)\|_2}{(n+1)\hat{D}^{(k)}}, \frac{n}{n+1} \right\} \right\},$$

where $\hat{D}^{(k)} = \min_{v \in \mathcal{V}(\theta)\setminus\{\psi(z)\}} |\langle \phi(\hat{y}_{n+1}^{(k)}), \psi(z) - v\rangle| / \|\psi(z) - v\|_2$.

Another insight from Proposition 2 is that the proposed estimator can be viewed as a weighted version of a Monte Carlo estimator based on the conditional generative models. The weights take values between 0 and 1, and are determined by the proposed conformalized procedure to guarantee conservatism (Theorem 1). Without the weights, Theorem 1 typically does not hold unless further assumptions are placed on the approximation quality of $f(Y|X)$ to the ground-truth conditional PDF of $Y|X$, which is not assumed in our setting. This further highlights the robustness of our algorithm.

4 Numerical Results

In this section, we present one experiment across multiple datasets where we assess decision quality based on risk estimates produced by CREDO across both synthetic and real-world settings. We show that CREDO can be used to select decisions with consistently higher confidence, demonstrating the effectiveness of CREDO's risk estimates in guiding practical decision-making.

We adopt *empirical confidence ranking* as our primary evaluation metric: Given a decision policy π and a test dataset $\{(x_i, y_i)\}_{i=1}^m$, we apply π to each input x_i to generate predicted decisions $\{z_i\}_{i=1}^m$. We then compute the score $\sum_{i=1}^m h(z_i)$, where h maps each prediction to a discrete rank based on its frequency among the ground-truth optimal decisions $\{z_i^*\}_{i=1}^m$ in the test set. This metric is designed to capture a method's tendency to select decisions that are most likely to be optimal—an aspect valued in practice but not fully reflected by the standard notion of regret.

We define the decision policy for CREDO as selecting the action with the lowest estimated risk. It is compared with four decision-making baselines: predict-then-optimize (PTO) [5], robust optimization (RO) [6], smart predict-then-optimize (SPO+) [10], and decision-focused learning (DFL) [1].

Table 1: Evaluated empirical confidence ranking (↓) for different methods across three datasets.

Method	Setting I			Setting II			Real Data
	$\sigma = 0.1$	$\sigma = 1$	$\sigma = 10$	$\sigma = 0.1$	$\sigma = 1$	$\sigma = 10$	
PTO	$\mathbf{1.00 \pm 0.00}$	2.76 ± 0.59	2.24 ± 0.79	3.55 ± 0.50	3.36 ± 0.48	2.04 ± 1.65	1.75 ± 1.69
RO	$\mathbf{1.00 \pm 0.00}$	2.98 ± 0.14	3.00 ± 0.00	4.99 ± 0.10	6.00 ± 0.00	3.98 ± 0.80	3.00 ± 1.29
SPO+	$\mathbf{1.00 \pm 0.00}$	2.68 ± 0.65	2.02 ± 0.82	3.95 ± 1.20	4.67 ± 1.56	3.56 ± 1.50	2.67 ± 1.43
DFL	2.44 ± 0.64	1.83 ± 0.81	2.06 ± 0.79	3.60 ± 1.52	3.96 ± 2.07	3.66 ± 2.48	1.92 ± 1.04
CREDO	1.75 ± 0.77	$\mathbf{1.61 \pm 0.56}$	$\mathbf{1.48 \pm 0.52}$	$\mathbf{1.05 \pm 0.22}$	$\mathbf{1.00 \pm 0.00}$	$\mathbf{2.03 \pm 0.96}$	$\mathbf{1.75 \pm 0.92}$

Table 1 presents the comparison results. It can be seen that CREDO achieves the smallest ranking metric value across most datasets, on average selecting the top two most likely decisions across all datasets. Though it might seem concerning that the in Setting I ($\sigma = 0.1$), PTO, RO, and SPO+ all achieve better performance than CREDO, this is because when σ is small, the data becomes highly concentrated around the mean, rendering the problem nearly deterministic and can be best dealt with point-prediction baselines. These results highlight the capability of our method for high-quality decision-making, especially under highly uncertain environments.

5 Conclusion

We proposed CREDO, a distribution-free framework for decision risk assessment that combines inverse optimization with conformal prediction to estimate the probability that a candidate decision is suboptimal. We theoretically derived that CREDO provides statistically valid risk certificates, and also admits a computationally efficient closed-form solution under a broad class of separable optimization problems. We present empirical comparisons that demonstrate CREDO consistently delivers high-quality decisions in uncertain settings where existing methods often fall short.

References

[1] Brandon Amos and J Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International conference on machine learning*, pages 136–145. PMLR, 2017.

[2] Anastasios N Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*, 2021.

[3] Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023.

[4] Dimitris Bertsimas, Leonard Boussioux, Ryan Cory-Wright, Arthur Delarue, Vassilis Digalakis, Alexandre Jacquillat, Driss Lahlou Kitane, Galit Lukin, Michael Li, Luca Mingardi, et al. From predictions to prescriptions: A data-driven response to covid-19. *Health care management science*, 24:253–272, 2021.

[5] Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

[6] Dimitris Bertsimas and Aurélie Thiele. Robust and data-driven optimization: modern decision making under uncertainty. In *Models, methods, and applications for innovative decision making*, pages 95–122. INFORMS, 2006.

[7] Timothy Chan, Erick Delage, and Bo Lin. Conformal inverse optimization for adherence-aware prescriptive analytics. *Available at SSRN*, 2024.

[8] Timothy CY Chan, Rafid Mahmood, and Ian Yihang Zhu. Inverse optimization: Theory and applications. *Operations Research*, 73(2):1046–1074, 2025.

[9] Shuyi Chen, Ferdinando Fioretto, Feng Qiu, and Shixiang Zhu. Global-decision-focused neural odes for proactive grid resilience management. *arXiv preprint arXiv:2502.18321*, 2025.

[10] Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022.

[11] Etienne Gauthier, Francis Bach, and Michael I Jordan. E-values expand the scope of conformal prediction. *arXiv preprint arXiv:2503.13050*, 2025.

[12] Michael Lingzhi Li and Shixiang Zhu. Balancing optimality and diversity: Human-centered decision making through generative curation. *arXiv preprint arXiv:2409.11535*, 2024.

[13] Bo Lin, Erick Delage, and Timothy Chan. Conformal inverse optimization. *Advances in Neural Information Processing Systems*, 37:63534–63564, 2024.

[14] Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *Machine learning: ECML 2002: 13th European conference on machine learning Helsinki, Finland, August 19–23, 2002 proceedings 13*, pages 345–356. Springer, 2002.

[15] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.

[16] Melvyn Sim, Qinshen Tang, Minglong Zhou, and Taozeng Zhu. The analytics of robust satisficing: predict, optimize, satisfice, then fortify. *Operations Research*, 2024.

[17] Sukrita Singh, Neeraj Sarna, Yuanyuan Li, Yang Lin, Agni Orfanoudaki, and Michael Berger. Distribution-free risk assessment of regression-based machine learning algorithms. In *The 13th Symposium on Conformal and Probabilistic Prediction with Applications*, pages 44–64. PMLR, 2024.

[18] Onur Tavaslıoğlu, Taewoo Lee, Silviya Valeva, and Andrew J Schaefer. On the structure of the inverse-feasible region of a linear program. *Operations Research Letters*, 46(1):147–152, 2018.

[19] Xuecheng Tian, Ran Yan, Yannick Liu, and Shuaian Wang. A smart predict-then-optimize method for targeted and cost-effective maritime transportation. *Transportation Research Part B: Methodological*, 172:32–52, 2023.

[20] Vladimir Vovk. Conformal e-prediction. *Pattern Recognition*, page 111674, 2025.

[21] Zhendong Wang, Ruijiang Gao, Mingzhang Yin, Mingyuan Zhou, and David Blei. Probabilistic conformal prediction using conditional random samples. In *International Conference on Artificial Intelligence and Statistics*, pages 8814–8836. PMLR, 2023.

[22] Pedro Zattoni Scroccaro, Bilge Atasoy, and Peyman Mohajerin Esfahani. Learning in inverse optimization: Incenter cost, augmented suboptimality loss, and algorithms. *Operations Research*, 2024.

[23] Zhiqiang Zhang, Pengyi Shi, and Amy Ward. Admission decisions under imperfect classification: An application in criminal justice. *Available at SSRN 5214197*, 2025.

[24] Shixiang Zhu, He Wang, and Yao Xie. Data-driven optimization for atlanta police-zone design. *INFORMS Journal on Applied Analytics*, 52(5):412–432, 2022.