Advances in Reliability, Safety and Security

ESREL 2024
Collection of Extended Abstracts

Distributionally Robust Optimization For Decision Making Under Uncertainty For RAMS

Pascal Quach[a], Yiping Fang[a,b], Anne Barros[a,b]

[a] *Laboratoire Génie Industriel, CentraleSupélec, Université Paris-Saclay, France*
[b]*Chair Risk and Resilience of Complex Systems, Laboratoire Génie Industriel, CentraleSupélec, Université Paris-Saclay, France*

Keywords: distributional ambiguity, distributionally robust optimization, maintenance optimization, complex system resilience

Uncertainty quantification is a key component of modern reliability, and safety engineering. Most real-world decision-making problems in maintenance planning include uncertain parameters. The uncertainty model might be fitted to existing historical data, or expert knowledge, but remains ambiguous (Ellsberg, 1961). When data are scarce, or the system evolves in real conditions, the model may end up inaccurate. The impact of model misspecification is well-studied in the literature through perturbation and stability analysis (Rahimian and Mehrotra, 2022). To address model misspecification, we propose to introduce in RAMS analysis a framework from decision-making under uncertainty: the distributionally robust optimization framework.

A general theory of robust and distributionally robust optimization is presented in (Zhen et al., 2023). We quickly recall the setting of distributionally robust optimization. Let us consider a minimization problem with a decision-dependent objective function $g: \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, with decisions variables $\mathbf{x} \in \mathbb{R}$ and uncertain variables $\mathbf{z} \in \mathbb{R}^m$.

Table 1. Optimization under Uncertainty Paradigms.

Name	Formulation	Knowledge
Nominal Problem (NP)	$\min_x g(\mathbf{x}, \mathbf{z_0})$	$\mathbf{z_0}$ is known
Robust Optimization (RO)	$\min_x \max_{\mathbf{z} \in \mathcal{U}} g(\mathbf{x}, \mathbf{z})$	$\mathbf{z}$ is unknown, $\mathcal{U}$ is known
Stochastic Programming (SP)	$\min_x \mathbb{E}_{Z\sim\mathbb{P}_0}\, g(\mathbf{x}, \mathbf{Z})$	$\mathbb{P}_0$ is known, $\mathbf{Z}$ is a random vector
Distributionally Robust Optimization (DRO)	$\min_x \max_{\mathbb{P}\in\mathcal{P}} \mathbb{E}_{Z\sim\mathbb{P}}\, g(\mathbf{x}, \mathbf{Z})$	$\mathbb{P}$ is unknwon, $\mathcal{P}$ is known, $\mathbf{Z}$ a random vector

In robust optimization, we consider a ***worst-case approach***, i.e. we seek to find the least worst decision $\mathbf{x}_{RO}$ given a known uncertainty set $\mathcal{U}$. In stochastic optimization, we consider a known probabilistic model $\mathbb{P}_0$, and we seek to find the best decision $\mathbf{x}_{RO}$ under $\mathbb{P}_0$. Distributionally robust optimization can be seen as a combination of both robust and stochastic optimization, where we consider a family of probabilistic models $\mathcal{P}$, and we seek to find the worst-case optimal decision $\mathbf{x}_{DRO}$. Mathematically, the different approaches are summarized in Table 1.

We focus on the inner problem of DRO to better understand the role of the ambiguity set $\mathcal{P}$. Let equation denote the uncertainty quantification problem: $max_{\mathbb{P}\in\mathcal{P}}\, \mathbb{E}_{\mathbb{P}}\, [g(\mathbf{Z})]\#(\mathbf{1})$, with: (i) ambiguity set $\mathcal{P}$ containing all distributions that satisfy $J \in \mathbb{N}$ moment conditions $\mathcal{P} = \{\mathbb{P} | \mathbb{E}_{\mathbf{Z}\sim\mathbb{P}}[h(\mathbf{Z})] \leq \boldsymbol{\mu}\}$; (ii) decision-independent objective function $g: \mathbb{R}^n \to \mathbb{R}$; (iii) known vector of moment bounds $\boldsymbol{\mu} \in \mathbb{R}$; (iv) J known moment functions $h_j: \mathbb{R}^m \to \mathbb{R}, j \in \{1, \dots, J\}$.

The crux of DRO lies in modelling the ambiguity using $\mathcal{P}$, capturing it from partial knowledge, e.g., mean, variance, historical data, expert knowledge, etc., to hedge an adversarial model, i.e., the worst-case distribution. For example, existing data may suggest a component has an MTTF of 10^5 hours. In which case, we can quantify this uncertainty with the set of all time-to-failure distributions with mean 10^5 hours, i.e. $\mathcal{P} = \{\mathbb{P} | \mathbb{E}_{\mathbb{P}}[Z] = 10^5\}$, for Z the random variable representing the component lifetime. Additional information may be added to $\mathcal{P}$, e.g., standard deviation, skewness, etc. The complete decision problem is to find the worst-case optimal decision $\mathbf{x}_{\mathbf{DRO}}$, i.e. $\min_x \max_{\mathbb{P}\in\mathcal{P}} \mathbb{E}_{Z\sim\mathbb{P}}\, g(\mathbf{x}, \mathbf{Z})$, where $\mathbf{x}$ may denote a maintenance policy, operating conditions, system investment, etc. The DRO paradigm benefits from many advantages over traditional approaches (Kuhn et al., 2019). Most importantly, the ambiguity set is a powerful tool to model distributional ambiguity from historical data, leading to less risk without sacrificing computational tractability or optimality (Van Parys et al., 2021).

DRO guarantees better decisions in case of model misspecification in exchange for additional cost.

Basic maintenance problems often assume the maintained system lifetime distribution is: (i) known *a priori* or estimated from historical data; (ii) stationary; (iii) independent from environmental factors. In practice, these assumptions are restrictive, and may lead to ineffective or costly maintenance policies (Gan et al., 2023; Sidibé et al., 2016). These assumptions may be relaxed by assuming partial knowledge of the lifetime distribution. DRO is particularly well-suited to model this ambiguity and is proven to be effective in practice (Kuhn et al., 2019; Zhen et al., 2023), particularly in finance.

Let us a consider a portfolio of N independent assets. These N assets are subject to a stochastic degradation process. Unfortunately, the only reliable information we have from historical data are the mean and variance $\boldsymbol{\mu} \in \mathbb{R}_+^{\mathrm{N}}$ and $\boldsymbol{\Sigma} \in \mathbb{R}_+^{\mathrm{N}}$ of the lifetime $\mathbf{Z}$, where $\mathbf{Z}(\boldsymbol{\omega}) \in \mathbb{R}_+^{\mathrm{N}}$. This information may come from expert knowledge, or other estimation methods, e.g., physics models, data-driven prediction models, etc. We can model the ambiguity set $\mathcal{P}$ as the set of all distributions $\mathbb{P}$ with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\Sigma}$, i.e., $\mathcal{P} = \{\mathbb{P} | \mathbb{E}_{\mathbb{P}}[\mathbf{Z}] = \boldsymbol{\mu}, \mathbb{E}_{\mathbb{P}}[(\mathbf{Z} - \boldsymbol{\mu})^{\mathbf{2}}] = \boldsymbol{\Sigma}\}$. Let us assume, we want to minimise the downtime cost $\mathbf{c}$ of operating this system, where $\mathbf{x} \in \mathcal{A}^N$ is the maintenance policy, i.e., the actions $a \in \mathcal{A}$ to be taken for each asset. Then, we may formulate this decision problem as such $min_x\, max_{\mathbb{P} \in \mathcal{P}}\, \mathbb{E}_{\mathbb{P}}\, \mathbf{c}(\mathbf{x}, \mathbf{Z})\ s.t. operating\ constraints$. Notice that no assumption has been made on the distribution of $\mathbf{Z}$, but only on the first and second moments. The ambiguity set $\mathcal{P}$ can be further refined by considering a confidence set around the moments (Delage and Ye, 2010).

Where dependability management is concerned with performance under required conditions, resilience aims to defend against unexpected changes. As such, little data is available to model the adversary. The inherent multi-stage nature of resilience problems, i.e. planning, response, and recovery, makes them particularly well-suited for combinatorial scenario-based approaches. However, these approaches are limited by the curse of dimensionality and computational complexity of the solution methods. Existing literature in resilience enhancement is extensive, and a wide variety of optimization methods have been used. Applications of DRO to resilience problems are few, and often limited to power systems. However, the general framework of DRO is well-suited to generic infrastructures under random contingency (Bellè et al., 2023).

Consider a networked infrastructure of size N. The goal is to invest in additional components to ensure that the network remains operational under any scenario of K simultaneous random failures. Each failure scenario can be modelled by a binary vector $\mathbf{u} \in \{0,1\}^N$, where $u_i = 0$ indicates that component i has failed. The set of all feasible scenarios is denoted by $\mathcal{A} = \left\{\mathbf{u} \in \{\mathbf{0}, \mathbf{1}\}^{\mathbf{N}} \middle| ||\mathbf{u}||_1 \geq N - K\right\}$. Distributional ambiguity around $\mathbf{u}$ can be modelled by considering partial knowledge of marginal failure probabilities in the form of upper bounds $\boldsymbol{\pi}^{max}$ $\in \mathbb{R}_+^{\mathrm{N}}$. The ambiguity set $\mathcal{P}$ is defined as the set of all distributions $\mathbb{P}$ with mean component failure probabilities $\mathbb{E}[\mathbf{1} - \mathbf{U}]$ lower than $\boldsymbol{\pi}^{max}$ with $\mathcal{M}(\mathcal{A})$ the set of all distributions defined on the σ-algebra of $\mathcal{A}$: $\mathcal{P} = \mathbb{P} \in \mathcal{M}(\mathcal{A}) | 0 \leq \mathbb{E}_{\mathbb{P}}\, [\mathbf{1} - \mathbf{U}] \leq \boldsymbol{\pi}^{max}$. The decision problem is then written as the minimization of the investment and failure cost $\mathbf{c}$ with investment decisions $\mathbf{x}$: $min_x\, max_{\mathbb{P} \in \mathcal{P}}\, \mathbb{E}_{\mathbb{P}}\, \mathbf{c}(\mathbf{x}, \mathbf{U})\ s.t. operating\ constraints$.

We have briefly introduced the distributionally robust optimization framework and discussed its potential applications in RAMS analysis through two basic examples. We intend to demonstrate its attractive computational and interpretability properties by solving the previous toy examples.

Acknowledgements

This work was supported by funding from the *Agence Nationale de la Recherche (ANR)* under the project "Robust and Scalable Prescriptive Analytics for the Resilience of Critical Infrastructure Networks – RoScaResilience" (ANR-22-CE39-0003) and the CentraleSupélec industrial chair "Risk and Resilience of Complex Systems (RRCS)", 2020-2024.

References

Bellè, A., Abdin, A.F., Fang, Y.-P., Zeng, Z., Barros, A., 2023. A data-driven distributionally robust approach for the optimal coupling of interdependent critical infrastructures under random failures. European Journal of Operational Research. doi: /10.1016/j.ejor.2023.01.060

Delage, E., Ye, Y., 2010. Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems. Operations Research 58, 595–612. https://doi.org/10.1287/opre.1090.0741

Ellsberg, D., 1961. Risk, Ambiguity, and the Savage Axioms*. The Quarterly Journal of Economics 75, 643–669. doi: /10.2307/1884324

Gan, S., Hu, H., Coit, D.W., 2023. Maintenance optimization considering the mutual dependence of the environment and system with decreasing effects of imperfect maintenance. Reliability Engineering & System Safety 235, 109202. doi: 10.1016/j.ress.2023.109202

Kuhn, D., Esfahani, P.M., Nguyen, V.A., Shafieezadeh-Abadeh, S., 2019. Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning, in: Operations Research & Management Science in the Age of Analytics, INFORMS TutORials in Operations Research. INFORMS, pp. 130–166. https://doi.org/10.1287/educ.2019.0198

Rahimian, H., Mehrotra, S., 2022. Distributionally Robust Optimization: A Review. Open Journal of Mathematical Optimization 3, 1–85. https://doi.org/10.5802/ojmo.15

Sidibé, I.B., Khatab, A., Diallo, C., Adjallah, K.H., 2016. Kernel estimator of maintenance optimization model for a stochastically degrading system under different operating environments. Reliability Engineering & System Safety 147, 109–116. doi: 10.1016/j.ress.2015.11.001

Van Parys, B.P.G., Esfahani, P.M., Kuhn, D., 2021. From Data to Decisions: Distributionally Robust Optimization Is Optimal. Management Science 67, 3387–3402. https://doi.org/10.1287/mnsc.2020.3678

Zhen, J., Kuhn, D., Wiesemann, W., 2023. A Unified Theory of Robust and Distributionally Robust Optimization via the Primal-Worst-Equals-Dual-Best Principle. Operations Research. https://doi.org/10.1287/opre.2021.0268

Monitoring and early warning systems