UDC 621.52+511.52
DOI: 10.53297/18293336-2025.2-27

A SYSTEMATIC REVIEW OF ROBUST AND ADAPTIVE DECISION MODELS IN COMPLEX AND UNCERTAIN ENVIRONMENTS

A.H. Grigoryan

European University of Armenia

This paper reviews robust and adaptive decision models for complex environments characterized by risk, ambiguity and deep uncertainty. It contrasts three regimes of uncertainty and shows how traditional "expected value" optimization becomes fragile under distribution shifts, adversarial behavior, and model misspecification. The review covers stochastic programming with recourse and risk measures such as chance constraints and conditional value-at-risk, which explicitly manage tail events.

It then examines robust optimization and distributionally robust optimization, highlighting how uncertainty sets and ambiguity sets protect against parameter and distributional misspecification. Sequential settings are addressed through robust and risk-sensitive Markov decision processes, which embed worst-case and tail-risk considerations into dynamic control. Decision-focused learning is presented as an approach to align machine learning with downstream optimization by differentiating through solvers and, where necessary, hardening models via adversarial and distributionally robust optimization-inspired training.

The paper emphasizes that mathematical robustness alone is insufficient: effective deployment requires human–AI teaming, layered socio-technical architectures, explainability, and autonomy-preserving governance. A multi-dimensional evaluation framework integrating algorithmic, human, stakeholder, learning, and implementation metrics is outlined. The paper concludes with open challenges in explainable uncertainty, robustness–adaptivity trade-offs, scalability to large combinatorial problems, multi-agent robustness, and federated, privacy-preserving decision support.

Keywords: risk, ambiguity, deep uncertainty, robust decision-making, adaptive decision support, hybrid intelligence, stochastic programming, robust optimization.

Introduction. Decision-making nowadays faces three canonical regimes of uncertainty. Risk assumes that probabilities are known or reliably estimable. e.g., standard supply-chain lead times. Ambiguity arises when only partial information, such as bounds or moments, is available, e.g., limited demand data. Deep uncertainty occurs when there is no consensus on the appropriate probability model, requiring strategies focused on resilience rather than point optimality [1]. When environments become volatile due to model misspecification, adversarial manipulation, or severe regime shifts, policies optimized merely for expected performance can become fragile.

The first conceptual dimension is the nature of uncertainty. Under risk, decision-makers are comfortable specifying a unique distribution P, supported by rich historical data and stable conditions. Under ambiguity, decision-makers acknowledge that multiple distributions are plausible, either because data are limited, heterogeneous, or noisy, or because structural model assumptions are questionable. Under deep uncertainty, models themselves are contested, or there is an expectation of regime shifts and structural breaks, as in long-term climate adaptation or disruptive technological change [1].

The second dimension is the temporal structure of decisions. One-shot decisions involve choosing a single action with no recourse, e.g., setting a price for a unique contract. Two-stage decisions allow a recourse action after uncertainty is realized, e.g., placing an initial order now and expediting later. Multistage decisions involve a sequence of actions over time, where later actions can depend on the observed outcomes, e.g., rolling capacity planning. Infinite-horizon problems involve ongoing decisions with discounting, e.g., in control and reinforcement learning.

Robust and adaptive decision models seek to provide resilience by guaranteeing performance under adverse conditions, scenarios, or distributions, adapting quickly when new data emerge, and embedding human judgment to preserve legitimacy. These models draw from several interconnected mathematical fields: stochastic programming (SP) for risk-based recourse [2]; robust optimization (RO) for worst-case parameter protection [3]; distributionally robust optimization (DRO) for ambiguity over probability distributions [4]; dynamic programming (DP) for sequential robustness and risk sensitivity [5, 6]; and decision-focused learning (DFL) for aligning prediction with prescriptive objectives and addressing the coupling between predictive models and downstream optimization [7]; robust Markov decision processes (MDP) and risk-sensitive control for sequential decision-making with ambiguous dynamics or tail-sensitive objectives.

The goal of this paper is to synthesize the core mathematical frameworks for robust and adaptive decision-making in complex and uncertain environments, and to connect them to socio-technical design.

Stochastic Programming and Risk Measures. Stochastic programming explicitly models the ability to take corrective actions after uncertainty is realized. The prototypical two-stage stochastic programming problem minimizes the sum of the first-stage cost and expected second-stage recourse cost:

$$\min_{x \in X} \{ c^\top x + E_{\omega \sim P}[Q(x, \omega)] \}, \quad (1)$$

where x is the first-stage decision, c - its cost vector, ω - a random vector of uncertainties. In this case, (2) presents the optimal value of a second-stage problem that chooses recourse actions y in response to ω:

$$Q(x, \omega) = \min_{y \in Y(x,\omega)} d(\omega)^T y. \tag{2}$$

The expectation is generally approximated by a finite set of scenarios $\{\omega_i\}$ with probabilities p_i, and large instances are solved using decomposition methods such as Benders decomposition [2, 8].

Two common risk-control alternatives are chance constraints, which require a constraint to hold with probability at least $1 - \alpha$: $P(\alpha(\xi)^T x \leq b) \geq 1 - \alpha$, and conditional value-at-risk (CVaR), which controls the tail of the loss distribution. CVaR at a confidence level $\propto$ has a convex representation:

$$CVaR_\alpha(L) = \min_{\eta \in R}\{\eta + \frac{1}{1-\propto} E[(L - \eta)_+]\}, \tag{3}$$

where $(u)_+ = max\{u, 0\}$, making it tractable for linear losses. For many loss functions, this representation leads to convex optimization problems; for example, if L is a linear function of decision variables and random parameters, $CVaR$ minimization can often be reformulated as a linear program [9].

Robust Optimization and Distributional Ambiguity. Robust optimization replaces probabilistic assumptions with a deterministic uncertainty set U and seeks decisions that perform best in the worst case [3, 10]:

$$\min_{x \in X} \max_{u \in U} f(x, u) \quad \text{s.t.} \quad g(x, u) \leq 0; \forall u \in U. \tag{4}$$

Here, u represents uncertain parameters (such as demands, costs, or coefficients), and U is a convex uncertainty set capturing all plausible values. For linear programs with affine uncertainty, robust counterparts often preserve tractability. Consider a linear constraint with uncertain coefficients of the form $a(u)^T x \leq b$, where $a(u) = a + Au$ and $u \in U$. The robust counterpart requires:

$$\max_{u \in U} a(u)^T x \leq b. \tag{5}$$

If U is polyhedral, this translates into additional linear constraints. If U is ellipsoidal, the robust counterpart can be reformulated as a second-order cone constraint.

Adjustable robust optimization extends robust optimization to multistage settings by allowing some decisions to depend on observed realizations [3, 11].

Distributionally robust optimization acknowledges that the true probability distribution may lie in an ambiguity set $\mathcal{P}$ defined by moment information [4] or statistical distances such as the Wasserstein metric [12]. The DRO problem minimizes the worst-case expected loss:

$$\min_{x \in X} \sup_{P \in \mathcal{P}} E_{\xi \sim P} \, [\ell(x, \xi)], \tag{6}$$

where $\ell(x, \xi)$ is a loss function.

Ambiguity sets can be specified in several ways. Moment-based sets fix the mean, covariance, and perhaps higher moments of ξ, thereby including all distributions that satisfy these constraints [4]. Divergence-based sets include all distributions within a specified Kullback–Leibler, χ^2, or other $f - \text{divergence}$ distance of a nominal distribution, often the empirical distribution [13]. Wasserstein ambiguity sets define a ball around the empirical distribution in terms of optimal transport distance, yielding finite-sample performance guarantees and tractable reformulations for many convex problems [12].

An important insight is that many DRO problems are equivalent to regularized empirical risk minimization: the inner supremum over distributions corresponds, via duality, to adding a regularization term to the empirical loss that penalizes sensitivity to distributional perturbations [13]. This connection helps explain empirically observed robustness improvements: DRO effectively trains models to be robust to distribution shift within the prescribed ambiguity set.

Sequential Decision-Making and Control. Markov decision processes model sequential decision-making under stochastic dynamics. Robust MDP treat transition probabilities as belonging to ambiguity sets $\mathcal{P}(s, a)$ and solve the robust Bellman equation [5]:

$$V(s) = \max_{a} \left\{ r(s, a) + \gamma \inf_{p \in \mathcal{P}(s,a)} \sum_{s'} p(s'|s, a) V(s') \right\}. \tag{7}$$

When ambiguity sets are rectangular, robust value iteration converges to an optimal stationary policy [14]. Risk-sensitive MDP minimize the CVaR of cumulative discounted cost, leading to policies that trade expected performance for tail-risk protection [6]. Distributional reinforcement learning techniques are employed to estimate the full return distribution needed for CVaR optimization [15].

Decision-Focused Learning (Learning to Decide). Traditionally, predictive models ("predict-then-optimize" pipelines) are trained to minimize a statistical loss such as mean-squared error, independent of the downstream optimization. However, improving statistical accuracy does not necessarily improve decision quality; prediction errors that are benign from a statistical perspective can be costly in terms of decisions [16].

DFL addresses this misalignment by training the predictive model to minimize a decision loss that reflects downstream performance [7, 16, 17]. Let $\widehat{\theta_{\phi}}(Z)$ be a predictor with parameters ϕ and inputs Z, and let θ denote the true parameters.

Let $\pi(\hat{\theta})$ denote the solution of an optimization problem parameterized by $\hat{\theta}$. The DFL objective can be written as:

$$\min_{\phi} E_{Z,\theta}\left[\mathcal{L}_{dec}\left(\pi\left(\widehat{\theta_{\phi}}(Z)\right),\theta\right)\right], \tag{8}$$

where $\mathcal{L}_{\text{dec}}$ measures the decision quality (for example, cost or regret), and $\pi(\,.\,)$ denotes the solution of an optimization layer. To optimize this objective using gradient-based methods, one must differentiate through the optimization layer.

For convex optimization problems with differentiable objectives, implicit differentiation of the Karush–Kuhn–Tucker (KKT) conditions provide gradients of the optimal solution with respect to parameters [7, 16]. Specialized differentiable layers have been developed for quadratic programs, linear programs, and certain conic programs. In these architectures, the solver becomes a differentiable component in the computation graph of a deep learning model.

DFL can be made robust by integrating DRO or adversarial training ideas. One approach is to minimize worst-case decision loss over a set of plausible data perturbations:

$$\min_{\phi} E_{Z,\theta}\left[\max_{\|\delta\|\le\epsilon}\mathcal{L}_{dec}\left(\pi\left(\widehat{\theta_{\phi}}(Z+\delta)\right),\theta\right)\right], \tag{9}$$

where δ represents adversarial or worst-case noise in the inputs [18, 19]. If the inner maximization is solved approximately and its effects backpropagated, the resulting model tends to be more robust to perturbations at deployment time. Such robust decision-focused learning is particularly relevant when decisions are safety-critical.

Human–AI Teaming and Socio-Technical Integration. Robust and adaptive decision models are not deployed in a vacuum; they are embedded in socio-technical systems, in which humans retain ultimate responsibility. Mathematical guarantees only translate into real-world resilience if human operators understand, appropriately rely on, and can override the system when necessary.

A commonly recommended architecture for human–AI decision-support systems consisting of three interacting layers [20]. The situational awareness layer ingests data, performs preprocessing and uncertainty quantification, and produces a compact representation of the environment, such as scenario sets, belief states, or risk maps. Uncertainty quantification may involve ensembles, Bayesian models, or residual monitoring. The collaborative decision layer runs robust optimization, DRO, stochastic programs, or robust MDP, generating recommended actions and associated risk metrics, such as worst-case costs, CVaR values, or robustness gaps. The reflective practice layer logs decisions, overrides, rationales, and outcomes; this information feeds back into model retraining, system redesign, and governance audits.

Dynamic control transfer is an important design element in such systems. Depending on risk, uncertainty, and workload, control may shift between humanled, machine-led, and joint decision modes. In low-risk, high-confidence situations, automation may predominate; in high-risk, low-confidence situations, the system may recommend deferring to human judgment, perhaps providing scenario analyses rather than prescriptive recommendations.

Explainability is significant for appropriate reliance. Users need to understand not only what action is recommended, but why, and how confident the system is. In robust and adaptive systems, this includes explaining which scenarios or uncertainty sets drive the recommendation, which constraints are binding, and what trade-offs are being made between robustness and performance [21, 22]. Counterfactual explanations, which show what changes in inputs would lead to different recommendations, can help users explore the decision space.

Sectoral Applications and Evaluation. Robust and adaptive decision models have been applied across a range of sectors with differing levels of maturity. In supply chain and operations management, robust optimization and DRO have been used for facility location, inventory control, and production planning under uncertain demand and lead times [23, 24].

In finance and energy, mean–CVaR portfolio optimization seeks to control tail risk in asset returns [19]. DRO has been employed to construct portfolios that are robust to statistical estimation error and regime shifts. In power systems, robust unit commitment models hedge against forecast errors in renewable generation and against equipment failures, often using scenario-based stochastic programming augmented with robust constraints [12].

In healthcare and emergency management, stochastic programming and robust MDPs have been applied to capacity planning, triage, and epidemic control. For example, models have been used to allocate intensive care unit beds, ventilators, or vaccines under uncertain demand and disease progression, often combining risk-based models for short-term dynamics with scenario-based approaches for deep uncertainty in long-term evolution [1]. Robust policies that perform well across a range of epidemiological scenarios have been investigated.

In infrastructure and cyber-physical systems, robust network design and restoration models aim to improve resilience against natural disasters and attacks. For instance, robust optimization and stochastic programming have been used to plan investments in flood defenses, transportation networks, and electricity grids under uncertain hazards. Cyber defense has been modeled as a robust MDP or partially observable game, in which defenders face strategic attackers with uncertain capabilities [10].

Across these domains, successful deployments typically combine mathematical robustness with adaptive learning and human-centric interfaces. For example, a robust scheduling algorithm may be integrated into an operator dashboard with scenario exploration tools, explanations of robustness margins, and override options.

Evaluating robust and adaptive decision systems requires going beyond traditional algorithmic metrics such as expected cost, regret, or convergence speed. A more comprehensive evaluation considers multiple dimensions. At the algorithmic level, metrics include regret under distribution shift, robustness gaps between nominal and worst-case performance, tail risk measures like CVaR, fairness metrics, and computational efficiency. At the human factors level, measures of calibrated trust, cognitive load, decision time, and appropriate reliance are important. Stakeholder impact can be accessed via perceived fairness, participation, and alignment with values. Learning and adaptation can be evaluated by the speed of drift detection, the frequency and quality of policy updates, and the extent to which lessons are institutionalized. Implementation sustainability encompasses adoption rates, cost–benefit ratios, resilience under stress tests, and compliance with legal and regulatory standards.

These dimensions can be integrated into a composite evaluation model. Let A_t denote algorithmic performance, H_t human factors, S_t stakeholder impact, L_t learning and adaptation, and I_t implementation sustainability at the time t. A composite outcome index might be written as:

$$O_t = f(A_t, H_t, S_t, L_t, I_t), \tag{10}$$

where f reflects organizational priorities and trade-offs. For example, an organization might require minimum thresholds on H_t and S_t to ensure acceptable human and stakeholder outcomes, while optimizing O_t subject to these constraints.

Conclusion. Despite substantial progress, several open challenges remain in the design and deployment of robust and adaptive decision systems.

Explainable uncertainty is a central challenge. While ambiguity sets, robust policies, and risk measures are precise mathematical constructs, translating them into narratives and visualizations that are meaningful to diverse users is non-trivial. Research is needed on interfaces and explanation techniques that make robustness and risk constraints transparent without overwhelming users.

Managing the trade-off between robustness and adaptivity is another challenge. Highly robust models can be overly conservative, sacrificing performance in typical scenarios to guard against unlikely extremes. Highly adaptive models can overfit to transient patterns or recent data, increasing vulnerability to rare events. Formalizing and optimizing trade-offs between worst-case protection and adaptation speed, possibly through hybrid objectives that penalize both worst-case loss and adaptation lag, is an active area of research.

Scaling robust and decision-focused methods to large, combinatorial problems remains difficult. While differentiable optimization layers and convex surrogates like SPO+ have broadened the scope of DFL, many real-world decision problems involve mixed-integer programs or nonconvex constraints. Efficient approximate methods, structure-exploiting algorithms, and problem decompositions will be essential to scale robust DFL to industrial applications.

Multi-agent and game-theoretic robustness is another frontier. Many systems involve strategic interactions among multiple actors, each with their own objectives and information. Extending robust optimization and DRO ideas to stochastic games, where both dynamics and payoffs may be ambiguous, is complex but important, particularly in cybersecurity, financial markets, and competitive logistics.

Finally, federated and privacy-preserving robustness poses both technical and governance challenges. In domains like healthcare and finance, data are distributed across institutions and subject to strict privacy constraints. Developing federated DRO, robust reinforcement learning, and decision-focused learning methods that respect data locality and legal constraints, while providing robustness guarantees, is a promising direction.

In conclusion, robust and adaptive decision models offer a rich set of tools for reasoning under uncertainty in complex environments. Stochastic programming provides risk-based recourse under known distributions; robust optimization and DRO guard against parameter and distributional ambiguity; risk-sensitive and robust MDPs extend robustness to sequential control; and decision-focused learning aligns prediction with downstream optimization. When embedded in human-centered architectures with explainability, autonomy-preserving governance, and multi-dimensional evaluation, these methods can significantly enhance resilience and performance. Continued research at the intersection of optimization, machine learning, control, and human–computer interaction will be critical for realizing the full potential of robust and adaptive decision-making in practice.

References

1. **Walker W.E., Haasnoot M., Kwakkel J.H.** Adapt or Perish: A Review of Planning Approaches for Adaptation Under Deep Uncertainty // Sustainability. – 2013. – Vol. 5, No. 3. – P. 955-979.
2. **Birge J. R., Louveaux F.** Introduction to Stochastic Programming.- Springer, 2011. (2nd ed.).
3. **Ben-Tal A., Nemirovski A.** Robust Optimization.- Princeton University Press, 2009.
4. **Delage E., Ye Y.** Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems // Operations Research. – 2010. – Vol. 58, No. 3. – P. 595-612.
5. **Iyengar G. N.** Robust Dynamic Programming // Mathematics of Operations Research. – 2005. – Vol. 30, No. 2. – P. 257-280.
6. **Chow Y., Tamar A., Mannor S., Pavone M.** Risk-Sensitive and Robust Decision-Making: A CVaR Optimization Approach // Advances in Neural Information Processing Systems. – 2015. – P. 1522-1530.

7. **Amos B., Kolter J. Z.** OptNet: Differentiable Optimization as a Layer in Neural Networks // Proceedings of the 34th International Conference on Machine Learning. – 2017. – P. 136-145.
8. **Shapiro A., Dentcheva D., Ruszczyński A.** Lectures on Stochastic Programming: Modeling and Theory (2nd ed.) // SIAM. – 2014.
9. **Rockafellar R. T., Uryasev S.** Optimization of Conditional Value-at-Risk// Journal of Risk. – 2000. – Vol. 2, No. 3. – P. 21-41.
10. **Bertsimas D., Brown D. B., Caramanis C.** Theory and Applications of Robust Optimization // SIAM Review. – 2011. – Vol. 53, No. 3. – P. 464-501.
11. **Ben-Tal A., Goryashko A., Guslitzer E., Nemirovski A.** Adjustable Robust Solutions of Uncertain Linear Programs // Mathematical Programming. – 2004. – Vol. 99, No. 2. – P. 351-376.
12. **Esfahani P.M., Kuhn D.** Data-Driven Distributionally Robust Optimization Using the Wasserstein Metric: Performance Guarantees and Tractable Reformulations // Mathematical Programming. – 2018. – Vol. 171, No. 1-2. – P. 115-166.
13. **Namkoong H., Duchi J. C.** Variance-Based Regularization with Convex Objectives // Advances in Neural Information Processing Systems. – 2017. – P. 2971-2980.
14. **Nilim A., El Ghaoui L.** Robust Markov Decision Processes with Uncertain Transition Matrices // Operations Research. – 2005. – Vol. 53, No. 5. – P. 780-798.
15. **Tamar A., Chow Y., Ghavamzadeh M., Mannor O.** Policy Gradient for Coherent Risk Measures // Advances in Neural Information Processing Systems. – 2015. – P. 1468-1476.
16. **Elmachtoub A. N., Grigas P.** Smart Predict, then Optimize // Management Science. – 2021. – Vol. 67, No. 9. – P. 5752-5780.
17. Differentiable Convex Optimization Layers / **A. Agrawal, B. Amos, S. Barratt, S. Boyd, et al** // Advances in Neural Information Processing Systems. – 2019. – P. 9562-9574.
18. Towards Deep Learning Models Resistant to Adversarial Attacks / **A. Madry, A. Makelov, L. Schmidt, D. Tsipras, et al** // International Conference on Learning Representations. – 2018.
19. **Sinha A., Namkoong H., Duchi J.C.** Certifiable Distributional Robustness with Principled Adversarial Training // International Conference on Learning Representations. – 2018.
20. Guidelines for Human-AI Interaction / **S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, et al** // Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. – 2019. – P. 1-13.
21. **Doshi-Velez F., Kim B.** Towards a Rigorous Science of Interpretable Machine Learning // arXiv:1702.08608. – 2017.
22. **Miller T.** Explanation in Artificial Intelligence: Insights from the Social Sciences // Artificial Intelligence. – 2019. – Vol. 267. – P. 1-38.
23. **Kouvelis P., Yu G.** Robust Discrete Optimization and its Applications.- Springer, 1997.

24. **Snyder L.V.** Facility Location Under Uncertainty: A review // IIE Transactions. – 2006. – Vol. 38, No. 7. – P. 547-564.

Received on 15.10.2025.
Accepted for publication on 29.01.2026.

ԿՈՄՊԼԵՔՍ ԵՎ ԱՆՈՐՈՇ ՄԻՋԱՎԱՅՐԵՐՈՒՄ ՌՈԲԱՍՏ ԵՎ ԱԴԱՊՏԻՎ ՈՐՈՇՈՒՄՆԵՐԻ ՄՈԴԵԼՆԵՐԻ ՀԱՄԱԿԱՐԳԱՅԻՆ ԱԿՆԱՐԿ

Ա.Հ. Գրիգորյան

Ներկայացվում է ռոբաստ և ադապտիվ որոշումների մոդելների ակնարկ կոմպլեքս միջավայրերի վերաբերյալ, որոնք բնութագրվում են ռիսկով, «ոչ միանշանակությամբ» (ambiguity) և խոր անորոշությամբ (deep uncertainty): Առանձնացվում են անորոշության երեք ռեժիմներ, և ցույց է տրվում, թե ինչպես է դասական «սպասվող արժեքի» օպտիմալացումը դառնում փխրուն՝ բաշխումների փոփոխության, թշնամական (հակընդդեմ) վարքի և մոդելի սխալ որոշման պայմաններում: Դիտարկվում է ստոխաստիկ ծրագրավորումը՝ ռեկուրսիայի և ռիսկի վրա հիմնված չափանիշներով, ինչպիսիք են հավանականային սահմանափակումները և պայմանական ռիսկային արժեքը (CVaR)՝ «պոչային ռիսկերը» (tail risk) կառավարելու համար:

Հաջորդիվ ներկայացվում են ռոբաստ օպտիմալացումն ու բաշխումով ռոբաստ օպտիմալացումը (distributionally robust optimization), որոնք օգտագործում են անորոշ և ոչ միանշանակ բազմություններ՝ պարամետրերի և բաշխման սխալ գնահատումներից խուսափելու համար: Հաջորդական որոշումների համար ուսումնասիրվում են ռոբաստ և ռիսկի նկատմամբ զգայուն որոշումների մարկովյան գործընթացները՝ դինամիկ կառավարման մեջ ներառելով վատագույն դեպքը և պոչային ռիսկը: Որոշումներին ուղղված ուսուցումը նկարագրվում է որպես մոդելի ուսուցումը հետագա օպտիմալացման հետ համապատասխանեցնելու միջոց՝ լուծողների միջև տարբերակման և, անհրաժեշտության դեպքում, հակընդդեմ (adversarial) ուսուցման և բաշխումից ոգեշնչված ուսուցման օպտիմալացման միջոցով՝ ռոբաստությունը բարելավելու համար: Որոշումահեն ուսուցումը (decision-focused learning) նկարագրվում է որպես մոդելի ուսուցումը հետագա օպտիմալացման հետ համապատասխանեցնելու միջոց՝ որոշիչների միջև տարբերակման և, անհրաժեշտության դեպքում, հակընդդեմ ուսուցման և բաշխումից ոգեշնչված ուսուցման օպտիմալացման միջոցով՝ ռոբաստությունը օպտիմալացնելու համար:

Ընդգծվում է, որ միայն մաթեմատիկական ռոբաստությունը բավարար չէ. անհրաժեշտ են նաև մարդ-արհեստական բանականություն փոխգործակցություն, բազմաշերտ սոցիալ-տեխնիկական ճարտարապետություններ, բացատրելիություն և մարդկային ինքնավարությունը պահպանելու մեխանիզմներ: Առաջարկվում է բազմաչափ գնահատման շրջանակ, որը ինտեգրում է ալգորիթմական չափանիշները, մարդկային գործոնները, շահագրգիռ կողմերի ազդեցությունը, կազմակերպչական ուսուցումը և ներդրման կայունությունը: Քննարկվում են հետազոտության ենթակա խնդիրներ, ինչպիսիք են՝ բացատրելի անորոշությունը, ռոբաստությունը և հարմարվողականության միջև հավասարակշռությունը, մասշտաբավորումը և բազմագենտ կայունությունը:

Առանցքային բառեր. ռիսկ, ոչ միանշանակություն, խոր անորոշություն, ռոբաստ որոշումների կայացում, ադապտիվ որոշումների կայացում, հիբրիդային բանականություն, ստոխաստիկ ծրագրավորում, ռոբաստ օպտիմալացում:

СИСТЕМАТИЧЕСКИЙ ОБЗОР РОБАСТНЫХ И АДАПТИВНЫХ МОДЕЛЕЙ РЕШЕНИЙ В КОМПЛЕКСНЫХ И НЕОПРЕДЕЛЁННЫХ СРЕДАХ

А.Г. Григорян

Представлен обзор робастных и адаптивных моделей принятия решений для комплексных сред, характеризующихся риском, неоднозначностью и глубокой неопределённостью. Выделяются три режима неопределённости и показывается, что классическая оптимизация по ожидаемому значению становится хрупкой при смещении распределений, враждебных действиях и ошибочной спецификации моделей. Рассматривается стохастическое программирование с рекурсией и риск-ориентированными мерами, такими как вероятностные ограничения и условная стоимость риска (CVaR), позволяющими управлять хвостовыми рисками.

Далее анализируются робастная оптимизация и дистрибутивно-робастная оптимизация, где множества неопределённости и неоднозначности защищают от ошибок в параметрах и распределениях. Для последовательных решений изучаются робастные и риск-чувствительные марковские процессы принятия решений, включающие наихудший случай и хвостовой риск в динамическое управление. Подход обучения, направленный на принятие решений (DFL), описывается как способ согласовать обучение моделей с последующей оптимизацией путём дифференцирования через решатели и, при необходимости, использования противодействующей и оптимизации тренировки, инспирированной с точки зрения распределения для повышения робастности.

Подчёркивается, что одной математической робастности недостаточно: необходимы человек–искусственный интеллект-команды, многоуровневые социотехнические архитектуры, объяснимость и механизмы сохранения человеческой автономии. Предлагается многомерная схема оценки, объединяющая алгоритмические показатели, человеческие факторы, влияние на заинтересованные стороны, обучение организации и устойчивость внедрения. Обсуждаются открытые задачи: объяснимая неопределённость, баланс между робастностью и адаптивностью, масштабирование и многоагентная робастность.

Ключевые слова: риск, неоднозначность, глубокая неопределённость, робастное принятие решений, адаптивное принятие решений, гибридный интеллект, стохастическое программирование, робастная оптимизация.