Nexus
Perspective

Bridging prediction and decision: Advances and challenges in data-driven optimization

Yanzhi Wang,[1] Jianxiao Wang,[2,3,*] Haoran Zhang,[4] and Jie Song[1,2,3,*]
[1]Department of Industrial Engineering and Management, College of Engineering, Peking University, Beijing 100871, China
[2]National Engineering Laboratory for Big Data Analysis and Applications, Peking University, Beijing 100871, China
[3]PKU-Changsha Institute for Computing and Digital Economy, Changsha 410000, China
[4]School of Urban Planning and Design, Peking University, Shenzhen, Guangdong 518055, China
*Correspondence: wang-jx@pku.edu.cn (J.W.); jie.song@pku.edu.cn (J.S.)
https://doi.org/10.1016/j.ynexs.2025.100057

BROADER CONTEXT

As big data technologies advance and data volumes grow, effectively leveraging these resources for complex decision-making has become a critical challenge for academia and industry. This review examines the transformative impact of big data and intelligent systems on traditional optimization paradigms, highlighting the continuum of data-driven optimization from predictive modeling to decision implementation. Key methodologies such as "sequential optimization," "end-to-end learning," and "direct learning" are analyzed, offering both theoretical insights and practical implications. Notably, we discuss breakthroughs such as implicit differentiation techniques, surrogate loss functions, and perturbation methods, which provide methodological guidance for achieving data-driven decision-making through prediction. By emphasizing the critical challenges across multiple dimensions, including data quality, model efficiency, and resilient decision-making under uncertainty, our review offers forward-looking insights to guide future research and foster the broader application of these approaches in diverse real-world scenarios.

ABSTRACT

Data-driven approaches have revolutionized traditional optimization methods by integrating prediction with decision-making. This review examines the theoretical foundations, strengths, recent advancements, and limitations of three key methods—sequential optimization, end-to-end learning, and direct learning—highlighting their practical applications in power grid scheduling, operations management, and intelligent autonomous control. A multidimensional comparison is presented, followed by a discussion of the challenges in data-centric methodology, optimization methodology, and decision-making application. This paper offers a methodological guide and outlines future directions for academia and industry to enhance decision-making in complex data environments.

INTRODUCTION

In the era of big data, the exponential growth in data volumes has opened unprecedented opportunities for refined decision-making across diverse industries and scientific fields.[1] Traditional deterministic models, constrained by their oversimplifications of complex realities, are increasingly being replaced by data-driven approaches that excel at extracting intricate patterns from large datasets.[2] These advanced models leverage the granularity of real-world data to enable precise, proactive decision-making, effectively transforming raw data into strategic value. For example, in energy systems, multi-timescale forecasting techniques are used to optimize economic dispatch. These techniques help balance supply and demand in real-time, addressing uncertainties from renewable energy sources and fluctuating loads.[3] These methods not only help operators anticipate short-term variability in solar and wind generation but also account for longer-term demand patterns, enabling precise operational strategies and ensuring system stability and cost efficiency.[4]

Prediction methodologies based on machine learning or deep learning extract features from real-world data to uncover latent patterns, supporting tasks such as diagnosis, predictive forecasting, and pattern recognition.[5] However, in practical applications such as power systems or market operations, prediction is rarely the final step. Instead, it serves as an intermediary to inform optimization models that enable smart, data-driven decision-making, such as determining optimal resource allocation under uncertainty, designing adaptive pricing strategies, or enhancing grid reliability through predictive maintenance. Given the interactions between prediction and decision-making, existing data-driven optimization methods can be broadly categorized into three frameworks, as illustrated in Figure 1, which provides a process overview, and further detailed in the comparative summary in Table 1.

(1) Sequential optimization (SO): this two-stage approach decouples and sequentially arranges prediction and decision-making. Predictive models are first trained on multi-scale data to estimate uncertain variables, with their outputs subsequently serving as inputs for optimization to derive decisions. As one of the most intuitive data-driven methods, SO offers flexibility and modularity, enabling independent advancements in predictive forecasting and optimization modeling to be seamlessly integrated. However, this two-stage coupling is subject to theoretical biases stemming from mismatches between predictive objectives—often optimized using norm-based loss functions such as mean squared error (MSE)—and decision-making goals, such as maximizing multi-dimensional social welfare under dynamic and uncertain conditions.

(2) End-to-end learning (E2E): by embedding optimization structures into the training process, E2E shifts the focus from traditional statistical learning objectives to decision-centric loss functions. During each training iteration, prediction results are processed through an optimization layer to generate decisions, with gradients subsequently backpropagated from validation in the decision-making context to refine and calibrate the predictive model. This closed-loop framework enables decision-focused learning, reducing biases inherent in traditional two-stage methods and aligning predictive tasks with optimization objectives. Techniques such as the neural network-based implicit differentiation algorithms proposed by Amos et al.[6] and the surrogate objective approach developed by Elmachtoub and Grigas[7] exemplify the practical implementation of E2E.

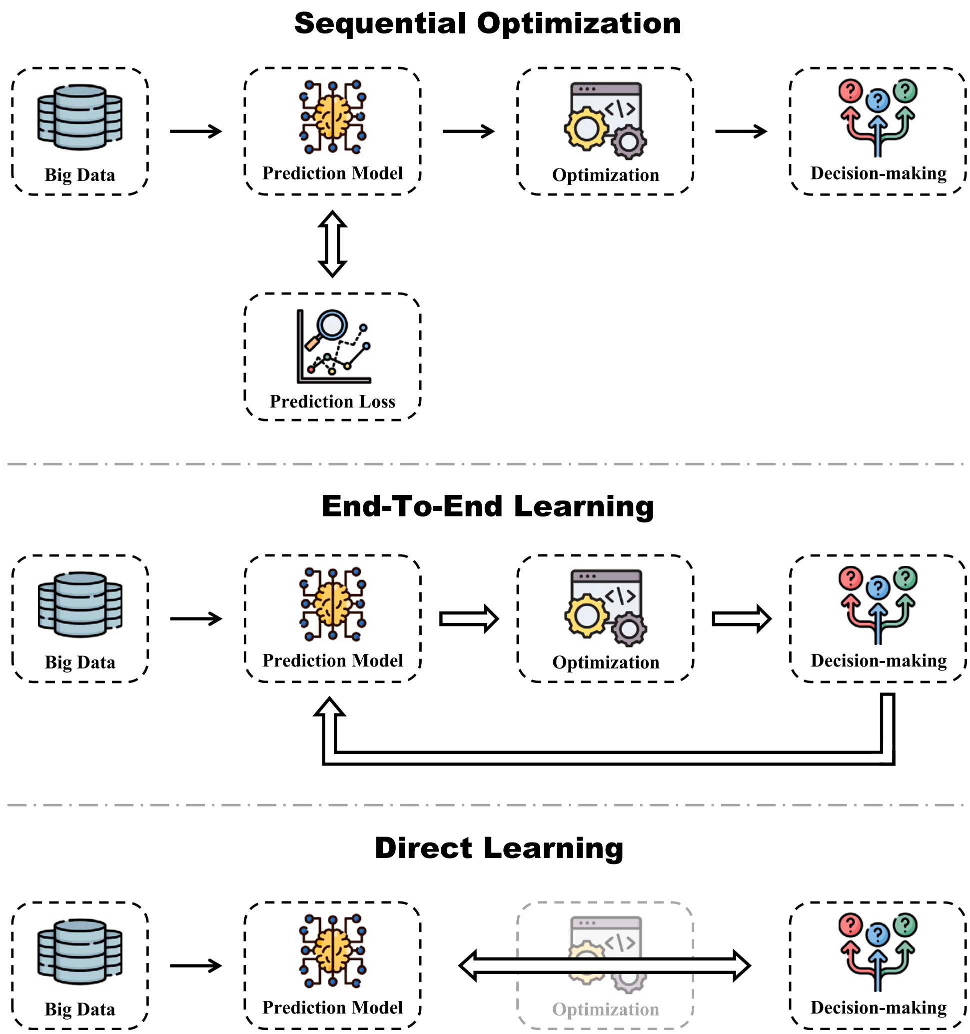

Figure 1. Frameworks for integrating prediction and decision-making

(3) Direct learning (DL): compared with SO and E2E, which relies on explicit optimization formulations, DL is tailored for scenarios with complex or implicit optimization structures, such as robotics control, where policies must be directly learned from human demonstrations to handle dynamic, unstructured environments.[8] By bypassing explicit optimization, DL focuses on aligning predictive outputs with end-goal performance metrics, enabling adaptive decision-making. Techniques such as imitation learning (IL) replicate expert behaviors, while reinforcement learning (RL) and model-free approaches dynamically adjust to evolving conditions, such as in personalized education systems, where teaching strategies adapt based on student performance.[9] These applications underscore DL's ability to integrate data directly into decisions, excelling in contexts where traditional optimization is infeasible.

Existing reviews often focus on specific aspects of optimization or data-driven methodologies. For instance, Kotary et al.[10] provide an in-depth exploration of end-to-end optimization techniques, while Qi and Shen[11] discuss the applications of integrated methods in operations management. Choi et al.[12] highlight the role of data analytics techniques—such as statistics, machine learning, and data mining—in addressing various operational challenges, and Zhang and Li[13] focus on decision-oriented learning strategies to tackle uncertainties in future power systems. While these reviews provide valuable analyses of data-driven methods within their respective domains, they fail to address the crucial interaction between predictive modeling and decision-making processes. This review bridges this gap by categorizing methodologies that integrate predictive models with decision-making frameworks, with an emphasis on their theoretical foundations and practical relevance across multiple domains. Unlike previous reviews limited to specific applications, this work adopts a broader perspective, targeting a diverse audience of researchers and practitioners in data science and optimization. By bridging optimization theory and data-driven methodologies, this review outlines strategies to advance data-driven optimization across diverse fields, offering insights to foster interdisciplinary collaboration and guide future research. It also analyzes practical challenges, including data-centric methodology, optimization methodology, decision-making application, a focus on aspects such as data quality, uncertainty modeling, and interpretability. Ultimately, this review provides perspectives that connect theoretical foundations with practical applications to support the development of data-driven optimization approaches.

A REVISIT TO SO

To establish a foundation in optimization, we first revisit the fundamentals of convex constrained optimization (CO), which underpins many data-driven decision-making frameworks. A general CO problem can be formulated as follows:

$$\underset{\mathbf{z}}{\text{minimize}} \quad f(\mathbf{z}, \hat{\mathbf{y}}) \qquad \text{(Equation 1)}$$

$$\text{subject to} \quad g(\mathbf{z}, \hat{\mathbf{y}}) = 0, h(\mathbf{z}, \hat{\mathbf{y}}) \leq 0 \qquad \text{(Equation 2)}$$

Table 1. Comparative summary of prediction-to-decision methodologies

Method	Training objective	Prediction-decision relationship	Theoretical applicablity
Sequential optimization	prediction loss (e.g,, MSE)	prediction outputs are treated as fixed inputs for subsequent optimization for decision-making	systems with explicit optimization formulations, but entails high computational costs and limited adaptability to decision-making feedback
End-to-end learning	decision loss (e.g., regret)	prediction models are trained through a closed-loop framework focused on optimizing decision performance	requiring decision-focused accuracy, with the added benefit of minimizing prediction bias in decision-making
Direct learning	decision loss (e.g., task-specific loss functions)	prediction models directly learn optimal decision-making strategies without relying on explicit optimization structures	with implicit decision logic or where traditional optimization formulations are infeasible

Here, f denotes the objective function, which depends on the decision variable $\mathbf{z}$ and an estimated parameter $\hat{\mathbf{y}} := p_\theta(\mathbf{x})$, where p_θ is a predictive model parameterized by θ that map input features $\mathbf{x}$ to the predicted variable $\hat{\mathbf{y}}$. Importantly, $\hat{\mathbf{y}}$ often deviates from the true value $\mathbf{y}$ due to prediction errors. The constraints g (equality constraints) and h (inequality constraints) define the feasible region for $\mathbf{z}$. To illustrate this structure with a real-world application, consider a resource allocation problem where $\mathbf{z}$ represents the allocation strategy and $\hat{\mathbf{y}}$ represents predicted resource demands, which guide the decision-making process. The objective function f quantifies the cost or efficiency of the allocation, while the constraints g and h enforce budgetary limits, resource availability, or system capacity.

This optimization formulation plays a critical role in bridging prediction to decision-making problems. In this framework, the predicted parameter $\hat{\mathbf{y}}$ acts as a signal that links the predictive model with the optimization process. It impacts both the objective function and the constraints, guiding the decision-making process based on the model's predictions. When $\hat{\mathbf{y}}$ deviates from $\mathbf{y}$, the resulting optimization may lead to suboptimal decisions, underscoring the need to align prediction accuracy with decision-making objectives. It is worth noting that this section focuses on the point prediction case. Extensions to uncertainty-aware predictions, such as stochastic optimization[14] and distributionally robust optimization,[15] will be discussed in subsequent sections. The summary of the list of the three data-driven optimization frameworks and their related methods can be found in Figure 2.

Under the SO framework, the process begins by training a predictive model using machine learning or deep learning techniques on historical resource demand data.[16] The trained model produces next-stage predicted demands, which serve as inputs to an optimization problem. Optimization solvers such as CVX or Gurobi are then used to solve the problem, incorporating specific constraints, including integer requirements or other domain-specific considerations.[17] The resulting solution represents the final decision, completing the data-driven decision-making process. This framework exemplifies the decoupled yet sequential approach to integrating prediction and optimization, a hallmark of SO methodologies. The theoretical foundation of SO hinges on the accuracy of the predictive model. In time series forecasting tasks, which are widely employed in multi-stage decision-making scenarios, such as energy scheduling, supply chain optimization, and financial portfolio management, norm-based statistical loss functions such as MSE and mean absolute percentage error (MAPE) are commonly utilized to measure the discrepancy between predicted values $\hat{\mathbf{y}}$ and true values $\mathbf{y}$.

As a result, recent research has shifted focus from traditional statistical models, such as support vector machines and k-nearest neighbors, to advanced deep learning architectures and large-scale models, representing significant progress in accurately identifying patterns in time series forecasting. In the context of energy markets, time series predictions, such as electricity price forecasting, demand estimation, and renewable energy output prediction, have become increasingly critical due to heightened market volatility driven by the growing penetration of renewable energy sources.[18] These predictions are foundational for optimizing resource allocation, managing grid stability, and developing effective bidding strategies, illustrating their role within the broader complexities of energy systems. To address these challenges, multi-horizon forecasting methodologies have been adopted to enhance predictive efficiency across various timescales.[19] A two-stage framework, which first forecasts discrete events such as price spikes and subsequently estimates continuous variables, has further improved forecasting accuracy.[20] In addition, hybrid models integrating long short-term memory (LSTM) networks with wavelet transformations have demonstrated superior forecasting precision.[21] Convolutional neural networks have proven effective in capturing local patterns and extracting high-level features,[22] while transformer models have advanced the field by managing long-range dependencies through attention mechanisms.[23] These methodological advancements collectively highlight ongoing efforts to improve predictive performance across a wide range of time series forecasting tasks, addressing the multifaceted needs of dynamic and volatile market environments.

Mismatch between prediction and decision

In practical applications, SO often encounters challenges due to the inherent limitations of achieving absolute predictive accuracy. Predictive models are typically trained using gradient descent algorithms, which terminate once the prediction loss reaches a predefined threshold to avoid overfitting, while striving to identify patterns as effectively as possible. However, decreasing prediction error does not necessarily translate into decision-making advantages, as the relationship between prediction and decision performance is often asymmetric in real-world scenarios.[24] Unlike the symmetric structure of norm-based prediction losses, decision-making criteria are rarely straightforward. For example, in load forecasting, operational costs in power systems are asymmetrical. Over-forecasting can lead to unnecessary generation, increased operational reserves, and inefficient resource allocation, while under-forecasting may result in under-supply scenarios, triggering costly balancing actions, reliance on fast-start units, and potential risks to system reliability. Norm-based statistical prediction losses treat both over- and under-forecasting errors equivalently, failing to capture these operational asymmetries.[25] This mismatch arises because decision-making relies on the practical environment to account for more implicit patterns in time series data. Valuation metrics in decision-making are often neither explicitly symmetric nor norm-based but are shaped by the nonlinear, implicit, and sometimes dynamic nature of objectives (Equation 1) and constraints (Equation 2). The inconsistent alignment between the descent direction of prediction loss and the cost objective frequently results in suboptimal decisions that deviate from the true optimal solution, potentially undermining the overall effectiveness of data-driven strategies in scenarios that place a high emphasis on decision accuracy.

E2E

To address the mismatch issue, E2E, also known as decision-focused learning, tightly couples prediction and decision-making within the training process. Unlike SO, E2E not only retains the sequential structure of prediction followed by decision-making but also integrates

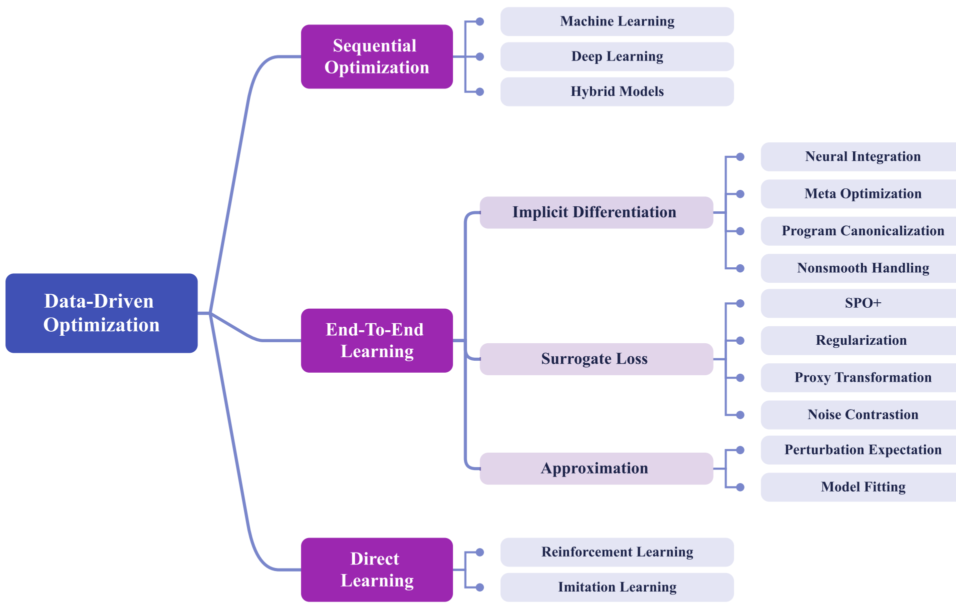

Figure 2. Taxonomy of data-driven optimization advancements

decision-making valuation directly into the loss function. This approach enables the predictive model to be trained iteratively through backpropagation of decision loss gradients, passing through the optimization structure to update the model parameters in a closed-loop manner.

For the optimal decision of the problem with objective (Equation 1) under constraints (Equation 2), denoted as $\widehat{\mathbf{z}}^* := \mathcal{M}(\widehat{\mathbf{y}})$, where $\mathcal{M}$ represent the mapping between input prediction $\widehat{\mathbf{y}}$ to optimal decision $\widehat{\mathbf{z}}^*$. An intuitive definition of decision loss quantifies the discrepancy between decisions made under the predicted scenario $\widehat{\mathbf{y}}$ and those under the true scenario $\mathbf{y}$. It is typically expressed as:

$$f(\mathcal{M}(\widehat{\mathbf{y}}), \mathbf{y}) - f(\mathcal{M}(\mathbf{y}), \mathbf{y}) \qquad \text{(Equation 3)}$$

which is commonly referred to as "regret." To accommodate diverse practical scenarios, the decision loss can also be expressed more generally as $\mathcal{L}(\mathcal{M}(\mathbf{y}), \mathcal{M}(\widehat{\mathbf{y}}))$. The general form facilitates the incorporation of theoretical constructs, such as relative regret in the newsvendor problem,[26] where customized loss functions are designed to better align with specific decision-making objectives. The primary goal of E2E is to train the predictive model with parameters θ to minimize $\mathcal{L}$, aligning predictive outputs directly with decision-making goals. Based on the chain rule, the gradient of the decision loss $\mathcal{L}$ with respect to the prediction parameter θ is expressed as follows:

$$\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial \mathcal{L}(\mathcal{M}(\mathbf{y}), \mathcal{M}(\widehat{\mathbf{y}}))}{\partial \mathcal{M}(\widehat{\mathbf{y}})} \cdot \frac{\partial \mathcal{M}(\widehat{\mathbf{y}})}{\partial \widehat{\mathbf{y}}} \cdot \frac{\partial \widehat{\mathbf{y}}}{\partial \theta} \qquad \text{(Equation 4)}$$

with a flow chart illustrated in Figure 3 for better understanding. The first term, $\frac{\partial \mathcal{L}}{\partial \mathcal{M}}$, can be computed directly since it depends solely on the predefined loss function $\mathcal{L}$. The last term, $\frac{\partial \widehat{\mathbf{y}}}{\partial \theta}$, is determined by the internal structure of the predictive model and can be obtained through standard backpropagation. The primary challenge, however, lies in computing the intermediate term, $\frac{\partial \mathcal{M}(\widehat{\mathbf{y}})}{\partial \widehat{\mathbf{y}}}$ (or $\frac{\partial \widehat{\mathbf{z}}^*}{\partial \widehat{\mathbf{y}}}$) which represents the gradient of the optimal decision $\mathcal{M}(\widehat{\mathbf{y}})$ with respect to the predicted parameter $\widehat{\mathbf{y}}$. The primary challenge lies in computing the intermediate term $\frac{\partial \mathcal{M}(\widehat{\mathbf{y}})}{\partial \widehat{\mathbf{y}}}$, which represents the gradient of the optimal decision $\mathcal{M}(\widehat{\mathbf{y}})$ with respect to the predicted parameter $\widehat{\mathbf{y}}$. For optimization tasks with analytic solutions, these gradients can be directly derived using explicit expressions.[27] In unconstrained optimization scenarios, Taylor expansion techniques can effectively simplify gradient computation.[28] However, this process becomes significantly more intricate in general CO problems, where decision outcomes are implicitly governed by the underlying optimization structure, necessitating specialized methods for accurate gradient evaluation. This section introduces three categories of approaches to E2E, with a detailed taxonomy illustrated in Figure 2, along with an analysis of computational cost and transferability challenges in E2E frameworks.

Implicit differentiation methods

To establish an equivalence between prediction and decision-making, the most direct approach involves applying the implicit gradient theorem derived from the Karush-Kuhn-Tucker (KKT) conditions. Specifically, for convex optimization problem defined by objective (Equation 1) and constraints (Equation 2), the Lagrangian function L is expressed as:

$$L(\mathbf{z}, \lambda, \nu, \widehat{\mathbf{y}}) = f(\mathbf{z}, \widehat{\mathbf{y}}) + \lambda^T g(\mathbf{z}, \widehat{\mathbf{y}}) + \nu^T h(\mathbf{z}, \widehat{\mathbf{y}}) \qquad \text{(Equation 5)}$$

where λ and ν are the multipliers for the g and h constraints in (Equation 2), respectively. The KKT conditions, which ensure optimality, include stationarity, primal feasibility, and complementary slackness, and are represented as:

$$K(\widehat{\mathbf{z}}^*, \lambda^*, \nu^*, \widehat{\mathbf{y}}) = \begin{bmatrix} \nabla_{\mathbf{z}} L(\widehat{\mathbf{z}}^*, \lambda^*, \nu^*, \widehat{\mathbf{y}}) \\ \mathrm{diag}(\lambda^*) g(\widehat{\mathbf{z}}^*, \widehat{\mathbf{y}}) \\ h(\widehat{\mathbf{z}}^*, \widehat{\mathbf{y}}) \end{bmatrix} = \mathbf{0} \qquad \text{(Equation 6)}$$

These conditions ensure that $\widehat{\mathbf{z}}^*$ meets the optimization criteria. By applying the implicit function theorem, if the derivatives and the Jacobian matrix of K with respect to $\widehat{\mathbf{z}}^*$ ($J_{\widehat{\mathbf{z}}^*}K$) are invertible, then the derivative of the optimal decision $\widehat{\mathbf{z}}^*$ with respect to the predictive output $\widehat{\mathbf{y}}$ is expressible as:

$$J_{\widehat{\mathbf{y}}} \widehat{\mathbf{z}}^* = \frac{\partial \widehat{\mathbf{z}}^*}{\partial \widehat{\mathbf{y}}} = -[J_{\widehat{\mathbf{z}}^*} K]^{-1} [J_{\widehat{\mathbf{y}}} K] \qquad \text{(Equation 7)}$$

Thus, the key gradient calculation during each epoch of predictive model training reduces to computing the Jacobian matrix of K, followed by inverting $J_{\widehat{\mathbf{z}}^*}K$ and multiplying it with $J_{\widehat{\mathbf{y}}}K$ under the given parameters. This process efficiently captures the sensitivity of the optimal decision $\widehat{\mathbf{z}}^*$ with respect to the predictive output $\widehat{\mathbf{y}}$.

Building on this theoretical foundation, E2E has advanced substantially with methods that integrate optimization layers into neural

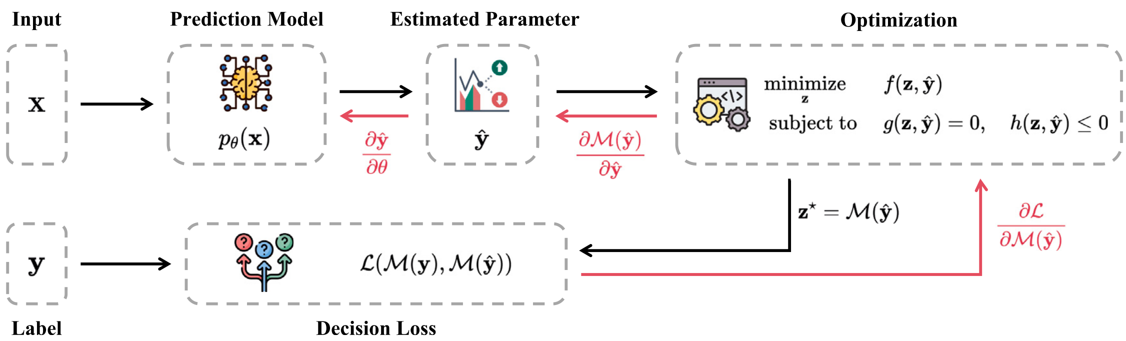

Figure 3. The flowchart of end-to-end training process

networks to enable smooth gradient propagation. A key breakthrough was made by Amos et al., who introduced OptNet, a GPU-accelerated quadratic optimization layer capable of efficiently addressing complex optimization-decision problems.[6] This layer facilitates seamless back-propagation of optimization gradients through neural networks, enhancing computational efficiency and enabling more effective decision-driven training. Further advancements have expanded the applicability of these methods. Lee et al. extended this framework to optimize embedding models for meta-learning tasks, demonstrating the utility of integrating optimization layers in improving task generalization.[29] At the same time, Agrawal et al. made a key contribution by proposing a structured approach to convex programming. Their method canonicalizes optimization problems into cone programs, enabling efficient computation of derivatives for solving complex convex optimization tasks.[30] Recent efforts have also focused on addressing challenges associated with nonsmooth objective functions. Bertrand et al. and Blondel et al. proposed efficient methods for managing modular automatic implicit subdifferentials, a crucial advancement for extending E2E to nonsmooth and large-scale optimization settings.[31,32] These innovations, combined with the rapid advancements in deep learning, notably enhance the capacity of E2E frameworks to tackle diverse and computationally demanding optimization problems.

Surrogate loss methods

Despite the advantages of implicit differentiation methods, their computational cost remains a key concern. As shown in Equation 7, each training iteration requires calculating and inverting the Jacobian matrix, which can be computationally expensive, especially for large input batches. To mitigate this, Elmachtoub and Grigas proposed simplifying the gradient calculation using surrogate loss functions.[7] For linear objective functions defined as $f(\mathbf{z},\mathbf{y}) = \mathbf{y}^T\mathbf{z}$, the regret can be expressed as $\mathcal{L} = \mathbf{y}^T(\mathcal{M}(\widehat{\mathbf{y}}) - \mathcal{M}(\mathbf{y}))$, as in Equation 3. By ensuring that the constraints in Equation 2 are independent of $\widehat{\mathbf{y}}$, a convex surrogate loss for $\widehat{\mathbf{y}}$ can be obtained through various methods, including duality, data scaling approximation, and first-order approximation, which is known as Smart "Predict, then Optimize" + (SPO+) loss and is defined as follows:

$$\mathcal{L}_{SPO+} = -(2\widehat{\mathbf{y}} - \mathbf{y})^T\mathcal{M}(2\widehat{\mathbf{y}} - \mathbf{y}) + (2\widehat{\mathbf{y}} - \mathbf{y})^T\mathcal{M}(\mathbf{y}) \qquad \text{(Equation 8)}$$

Considering the linear nature of the objective function, the subgradient of this loss is given by:

$$2\mathcal{M}(\mathbf{y}) - 2\mathcal{M}(2\widehat{\mathbf{y}} - \mathbf{y}) \in \frac{\partial \mathcal{L}_{SPO+}}{\partial \widehat{\mathbf{y}}} \qquad \text{(Equation 9)}$$

As a result, the term $\frac{\partial \mathcal{L}}{\partial \mathcal{M}(\widehat{\mathbf{y}})} \cdot \frac{\partial \mathcal{M}(\widehat{\mathbf{y}})}{\partial \widehat{\mathbf{y}}}$ in Equation 4 can be efficiently computed using two evaluations of the objective function with parameters $\mathbf{y}$ and $2\widehat{\mathbf{y}} - \mathbf{y}$, respectively. This streamlined approach eliminates the need for direct Jacobian matrix computation by leveraging the subgradient of the convex surrogate loss. Moreover, it has been theoretically proven, using Bayesian risk analysis, that minimizing the surrogate loss $\mathcal{L}_{SPO+}$ is equivalent to minimizing regret under mild conditions, offering a substantial improvement in computational efficiency for linear objectives. In addition, strategies such as problem relaxation and warm-starting have been developed to address the complexities of combinatorial optimization, further broadening the scope of E2E.[33]

Beyond SPO+, additional approaches have been developed to overcome the limitations of gradient calculation through innovative reformulations. One notable advancement is the use of smooth approximations to manage discontinuities in the objective function, combined with regularization terms to improve the differentiability of the optimization process.[34] Other methods reformulate the original optimization problems into tractable proxy tasks that preserve decision space while enhancing computational efficiency without compromising solution quality.[35] Furthermore, noise-contrastive estimation techniques treat non-optimal solutions as negative examples, leveraging learning-based surrogates to improve the discriminative power of the loss function.[36] These advancements streamline E2E gradient computation using surrogate methods, greatly improving the efficiency of decision-focused training.

Approximation methods

In addition to computational efficiency, the validity of gradient information in optimization with a polyhedral feasible domain and a linear objective is problematic, as changes in $\widehat{\mathbf{y}}$ may not always affect the optimal solution $\widehat{\mathbf{z}}^*$, and infinite solutions may occur. As shown in Figure 4A, the optimal solution lies at extreme points, where $\widehat{\mathbf{y}}$ and $\widehat{\mathbf{y}}_1$ may yield the same solution, while $\widehat{\mathbf{y}}_2$ leads to a different one. Consequently, the gradient is often undefined or zero, hindering updates to the predictive model.

To improve gradient smoothness and differentiability, scholars have introduced conceptual approximations using noise. By adding Gaussian noise ϵZ to $\widehat{\mathbf{y}}$, the perturbed solution is defined as:

$$\mathbf{z}^\star_{perturb} = \mathbb{E}_Z[\mathcal{M}(\widehat{\mathbf{y}} + \epsilon Z)] \qquad \text{(Equation 10)}$$

This perturbed expectation replaces the exact solution mapping $\mathcal{M}(\widehat{\mathbf{y}})$. As shown in Figure 4B, $\mathbf{z}^\star_{perturb}$ represents the expected value of extreme points, with probabilities influenced by $\widehat{\mathbf{y}}$ and the noise ϵZ. Changes in the probabilities of extreme points induce smooth variations in the perturbed solution, making the gradient continuous.[37] For specific scenarios, such as ensuring positivity in shortest path problems, researchers have adjusted the noise distribution (e.g., exponential noise) to maintain $\widehat{\mathbf{y}}$ positivity under perturbation, ensuring stability.[38] In addition to statistical approximations, continuous functional fitting offers another approach. Interpolating functions can ensure continuity in optimization.[39]

Computational cost and transferability in E2E

The methods discussed above focus on improving the feasibility and efficiency of E2E gradient computation. However, despite various simplifications, iterative optimization remains computationally intensive, particularly for integer-constrained problems. For example, mixed-integer linear programming (MILP), widely used in applications such as determining the switching status of power generation units, poses critical challenges. To address this, Bertsimas and Stellato proposed

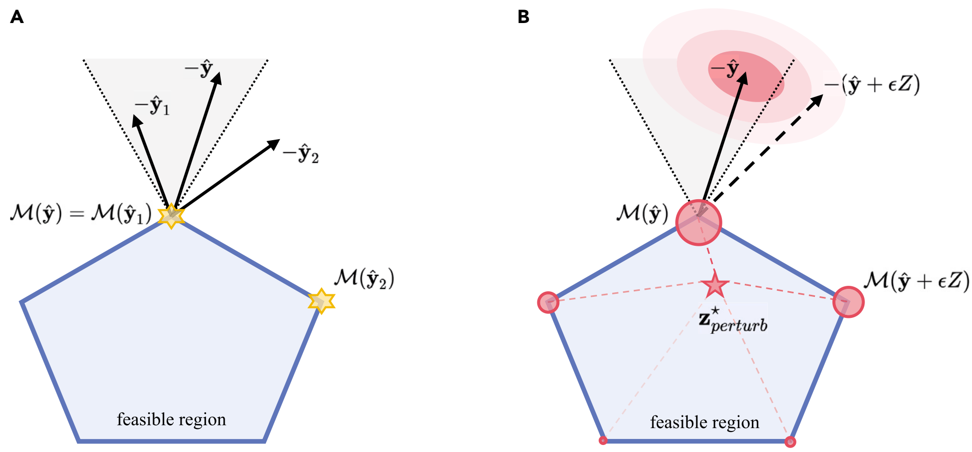

Figure 4. Geometric illustration of polyhedral feasible region with linear objective: (A) origin method and (B) perturbation approximation

an online learning algorithm with a feedforward neural network to accelerate convergence,[40] while Tang et al. used deep RL to adaptively select cutting planes, improving solver performance and efficiency.[41] These advancements are pivotal for enabling E2E methods to tackle increasingly complex optimization problems effectively. On the other hand, improving training efficiency offers a complementary approach to reducing computational costs. Techniques such as grid search, stochastic search, and Bayesian optimization are commonly used to identify optimal hyperparameters.[42] Multi-fidelity optimization further enhances efficiency by periodically pruning underperforming hyperparameter combinations based on performance metrics,[43] as demonstrated in wind power dispatch scenarios.[44] These strategies streamline the training process, reducing trial-and-error and accelerating convergence for E2E methods.

Transferability evaluates whether decisions made in one context can be effectively applied to others. In SO, the predictive model is relatively adaptable to related decision-making tasks when conceptual mismatches can be tolerated. However, E2E frameworks often sacrifice statistical accuracy to prioritize decision performance, resulting in inconsistent or nonsensical predictions.[7] Figure 5 illustrates this challenge. The decision problem involves selecting the edge (1 or 2) with the lowest cost, based on input feature x. In Figure 5A, the black line represents the optimal decision boundary, where x is less than (selects edge 2) or greater than (selects edge 1) than the threshold. Using SO (Figure 5B) leads to a mismatch between the orange line (suboptimal decision) and the black line. By contrast, E2E cases (Figures 5C and 5D) align decisions with the optimal boundary but produce predictions that deviate considerably due to overfitting to the decision task, reducing interpretability and hindering transferability to similar tasks. Despite attempts to improve transferability through hybrid losses combining decision and prediction objectives,[45] extending E2E methods to broader domains remains a critical challenge.

DL

The core process of E2E involves transferring the gradient of the optimization problem to the prediction model p_θ for unbiased training. However, methods such as implicit differentiation, surrogate losses, and approximations rely heavily on preserving the optimization structure, which may not always be feasible in complex, real-world scenarios. DL further integrates prediction and decision-making by bypassing the requirement for optimization structure preservation. Instead, it directly updates predictive model parameters based on the calculated decision loss $\mathcal{L}$, eliminating the need for optimization gradients.

RL

In DL, decision-making tasks are reformulated within a RL framework to eliminate reliance on explicit optimization structures. The model p_θ maps the input $\mathbf{x}$ to a decision $\mathbf{z}$. The RL framework optimizes the expected cumulative reward $J(\theta)$, defined as:

$$J(\theta) = \mathbb{E}_{\pi_\theta}\left[\sum_{t=0}^{\infty} \gamma^t r_t(\mathbf{x}_t, \mathbf{z}_t)\right] \qquad \text{(Equation 11)}$$

where t represents the time index in a Markov decision process $(\mathcal{S}, \mathcal{A}, \mathcal{T}, r, \gamma)$, which can also be formulated in implicit contexts, such as solving combinatorial optimization problems.[46] With $\mathbf{x} \in \mathcal{S}$ as the state, $\mathbf{z} \in \mathcal{A}$ as the action, and $\mathcal{T}(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{z}_t)$ defining the system's dynamics. The immediate reward r_t reflects decision quality, such as cost savings or reduced operational risks, while $\gamma \in [0, 1]$ is the discount factor balancing immediate and future rewards. The goal of RL-based DL is to find a policy $\pi_\theta(\mathbf{z}|\mathbf{x})$ parameterized by θ that maximizes $J(\theta)$.

RL provides a flexible framework for integrating prediction and optimization in data-driven decision-making, treating them as interdependent processes. By directly leveraging observed data, RL enables dynamic and adaptive strategies tailored to complex systems.[9] For instance, RL has been applied to error calibration in sequential day-ahead forecasts of demand load and wind power, improving downstream unit commitment decisions.[47] It also facilitates smart energy management by optimizing battery and thermal storage systems based on household-level indicators.[48] To enhance scalability and robustness, advanced RL techniques have been introduced. Proximal policy optimization enhances stability by constraining policy updates through a clipped surrogate objective, ensuring reliable convergence in high-dimensional optimization.[49] Deep deterministic policy gradient extends policy gradients to continuous action spaces using an actor-critic framework, enabling precise and efficient decision-making.[50] These methods highlight RL's ability to tackle complex, large-scale optimization problems requiring stability and precision.

While RL-based DL provides a robust framework for integrating prediction and decision-making, it faces key challenges. Balancing exploration and exploitation remain difficult, as discovering new strategies often comes at the expense of leveraging known policies. Sparse or delayed rewards, such as in long-term energy management, further complicate policy learning and convergence. In addition, scalability is a remarkable issue, as training RL policies for large-scale systems demands substantial computational resources, limiting their practical application. Promising solutions include hybrid approaches such as pretraining policies with supervised or adopting distributed RL algorithms to enhance sample efficiency and scalability.

IL

IL, a key method in DL, focuses on directly fitting the relationship between inputs $\mathbf{x}$ and decisions $\mathbf{z}$ using expert data collected during early stages.[51] Unlike RL, which balances exploration and exploitation

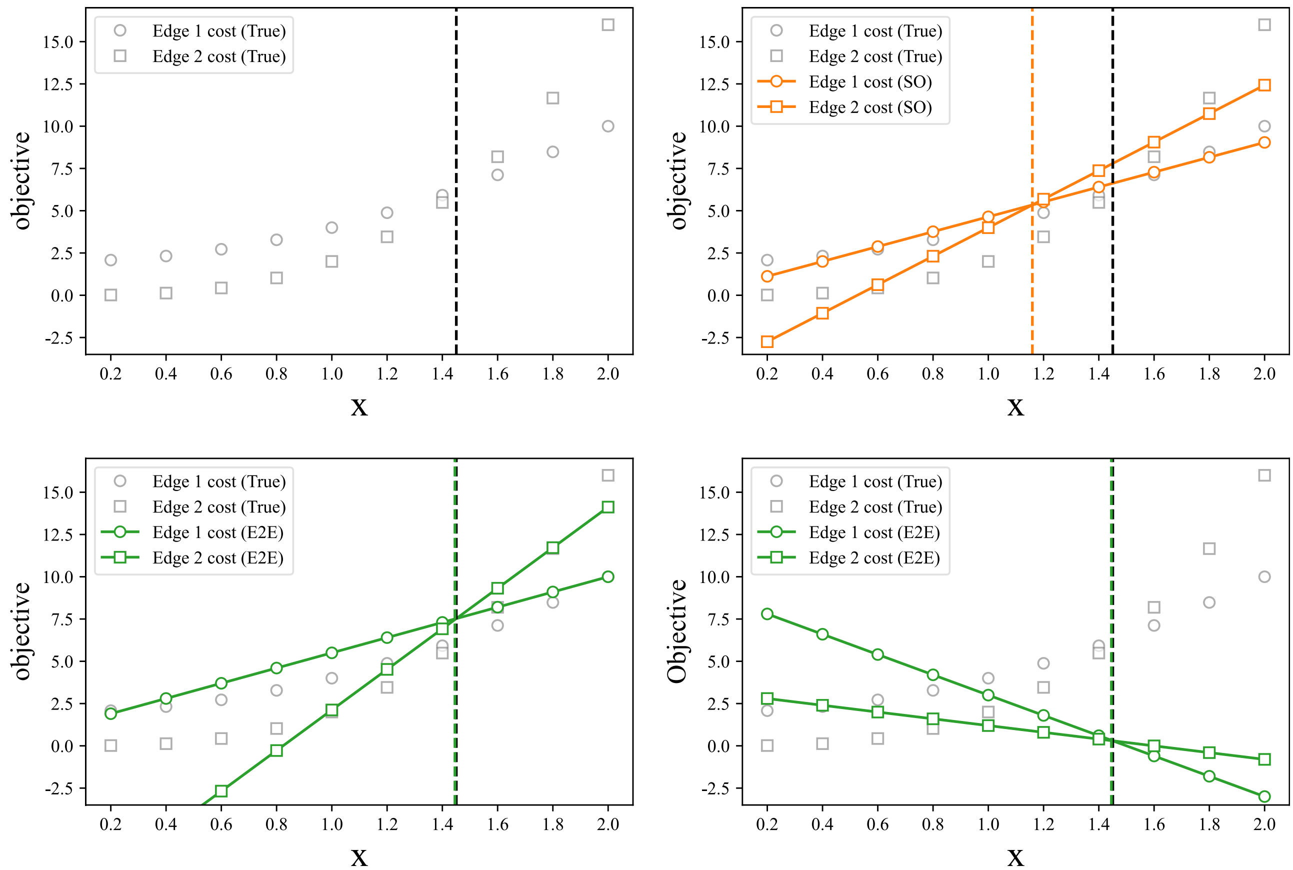

Figure 5. Untransferability from prediction randomness in end-to-end learning (adapted from Elmachtoub and Grigas[7])

through implicit optimization, IL bypasses optimization entirely, focusing solely on accurately mapping prediction to decision-making using expert demonstrations. IL is widely applied in real-world scenarios such as urban autonomous driving, where crafting generalized decision-making rules is challenging.[52]

In addition, localized IL methods have been proposed as a compromise to simulate gradient propagation in optimization processes. These methods retain the predictive model while directly fitting the gradient mapping, bypassing explicit optimization structures. Deep learning models h_ϕ are employed to approximate the decision loss using inputs $\hat{\mathbf{y}}$ and $\mathbf{y}$,[53] optimizing the following objective:

$$\underset{\phi}{\text{minimize}} \| h_\phi(\hat{\mathbf{y}}, \mathbf{y}) - f(\mathcal{M}(\hat{\mathbf{y}}), \mathbf{y}) \|^2 \qquad \text{(Equation 12)}$$

This approach treats the mapping from input features to decision outcomes as a "black box," allowing h_ϕ to approximate gradients. Consequently, the optimization task for θ is simplified to $h_\phi(p_\theta(\mathbf{x}), \mathbf{y})$, replacing the explicit gradient computation in Equation 4 with the internal gradient propagation of h_ϕ. While this method improves computational efficiency, its reliance on black box approximations introduces risks of local optima during auxiliary model training. Scenario-specific adjustments are often necessary to enhance robustness and ensure accurate decision-making.

IL's direct approach to mapping inputs to decisions makes it a computationally efficient alternative, while its ability to sidestep explicit optimization highlights its potential for streamlined decision-making in complex systems. However, the use of black box structures, such as neural networks, to approximate this mapping can lead to mismatches with optimization constraints, resulting in infeasible or suboptimal solutions and raising concerns about robustness and security. To address these issues, adjustments can be made during training to align predictions with decision constraints.[54] For instance, regularization terms or additional constraint-aware training objectives can be introduced to penalize infeasible predictions. In addition, strict architectural constraints can be embedded within the data-driven model to ensure consistent production of valid solutions.[55]

METHODS COMPARISON

To complement Table 1, we provide a structured comparison of SO, E2E, and DL based on their alignment between prediction and decision-making, training efficiency, and transparency, as shown in Table 2.

SO prioritizes predictive accuracy, typically minimizing statistical losses such as MSE or MAE, which measure the norm-based discrepancy between predicted values $\hat{\mathbf{y}}$ and ground truth $\mathbf{y}$. E2E, by contrast, aligns prediction objectives with decision-making goals, using decision loss as the optimization target to ensure that the predictive model directly serves decision-making needs. In DL, the training loss is abstracted, with reinforcement learning relying on predefined decision losses and imitation learning focusing on replicating expert decision rules without explicitly solving optimization problems.

Training efficiency largely hinges on the computational cost of gradient calculations in the optimization process. SO is computationally efficient as it optimizes only the predictive model, making it suitable for large-scale problems. E2E, however, is resource-intensive, requiring the optimization problem to be solved and gradients to be computed at each epoch. DL's efficiency depends on the availability of expert training data: with sufficient samples, training costs are primarily limited to predictive model updates, while insufficient data necessitate costly iterative optimization.

SO retains high transparency, as the prediction and optimization processes are distinct and interpretable. Although E2E sacrifices some interpretability by embedding decision objectives into the prediction

Table 2. Multidimensional comparison of different methodologies

Method	Objective	Consistency	Efficiency	Transparency
SO	$\lVert\hat{\mathbf{y}} - \mathbf{y}\rVert$	–	+	+
E2E	$\mathcal{L}(\hat{\mathbf{z}}^*, \mathbf{z}^*)$	+	–	+
DL	$\mathcal{L}(\hat{\mathbf{z}}^*, \mathbf{z}^*)$ or $\lVert\hat{\mathbf{z}}^* - \mathbf{z}^*\rVert$	+	○	–

Method	implementation contexts
SO	best suited for scenarios where retraining predictive models is resource-intensive, such as in large-scale models or diverse application contexts, where the system can tolerate some degree of decision suboptimality
E2E	ideal for tasks that require high decision accuracy, especially where predictive models can be easily retrained and closely aligned with decision-making objectives, ensuring optimal decision outcomes
DL	most effective in scenarios with complex decision-making problems that lack explicit optimization structures, where the primary goal is to accurately capture decision patterns directly from training data, as seen in autonomous systems such as self-driving cars

training, which complicates the tracing of decision patterns, it still retains the complete formulation of the optimization. DL, by bypassing explicit optimization structures, often operates as a black box, raising concerns about the reasoning and robustness of its decisions.

Each method has distinct advantages depending on the application context. SO is well suited for scenarios requiring scalable predictive models with low computational overhead, such as meteorological forecasting or large-scale natural language models, although it may compromise decision optimality. E2E is ideal for tasks demanding high decision accuracy, where predictive models are explicitly trained to optimize decision performance. DL, which abstracts optimization structure in favor of pattern recognition from training data, is particularly advantageous for complex, dynamic decision-making environments, such as autonomous driving.

APPLICATIONS IN DATA-DRIVEN OPTIMIZATION

Compared with traditional decision-making methods based on prior knowledge and models, integrating data and predictive techniques greatly enhances practical applications. For instance, by leveraging the fusion of large-scale meteorological data and geographic information, more accurate and efficient prediction of renewable energy can be achieved. This enables smarter and more efficient distributed power generation in microgrids through the use of multi-scale artificial intelligence models. In addition, in asset-intensive industries, to address common management challenges faced by power generation companies and grid operators, industry giants such as IBM and Oracle have developed big data-driven smart energy management products and services.[56] Overall, the advent of technologies such as data input and predictive analytics has injected new vitality into decision-making applications, offering substantial improvements in industry practices and standards.

This section highlights three mini case studies showcasing the versatility of data-driven optimization methods across distinct domains, including power grid scheduling, market operations, and broader operation management. These examples emphasize the cross-domain applicability and effectiveness of various techniques in addressing complex decision-making challenges.

Power grid scheduling

As one of the most complex systems in modern society, power grids integrate cross-domain data, including meteorological, energy, equipment, and network information. E2E has demonstrated significant advantages over SO in optimizing scheduling within such systems.

For instance, Wahdany et al.[44] developed an E2E framework for the optimal power flow problem using implicit gradient techniques derived from KKT conditions. Their method effectively reduced congestion and overloads in power systems under high wind capacity by up to 8.5% while decreasing error variance by 70%, remarkably outperforming SO. Similarly, Zhou et al.[57] introduced an LSTM-based load forecasting model coupled with a two-stage E2E approach using MILP optimization. Their multi-energy system optimization, which integrates electricity, heat, and cooling loads, achieved a 0.40% reduction in operational costs and annual savings of 124.66 kCNY.

Surrogate loss methods further enhance E2E computational efficiency in large-scale systems. Sang et al.[45] applied SPO+ to electricity price arbitrage for energy storage systems. Despite higher RMSE and MAPE (up to 10.24 times), their E2E approach reduced regret by 40.3% and improved daily economic benefits by \$1.72 per MWh (approximately 6.11%) compared with SO using a multilayer perceptron. Chen et al.[58] extended the applicability of SPO+ to MILP within a network-constrained unit commitment problem. By employing Lagrangian relaxation to simplify constraints and enhance feature utilization efficiency, their E2E framework achieved a remarkable 0.35% reduction in daily operation costs on an ISO-scale 5655-bus system.

Operations management

Operations management focuses on optimizing processes and decision-making to enhance efficiency and performance in manufacturing and service sectors. Advances in information technology have revolutionized operations management by enabling data-driven methods for precise resource allocation, demand management, and cost optimization, significantly improving operational effectiveness.[59]

Chu et al.[60] proposed an E2E framework for last-mile delivery services that combines SPO+ with efficient mini-batching gradient techniques and heuristic algorithms. Their approach integrates order allocation and route optimization, achieving approximately a 5% reduction in travel costs compared with SO. Tian et al.[61] introduced an E2E approach using SPO tree ensembles to optimize port state control (PSC) inspections. By considering inspection, repair, and risk costs, their method reduced a ship's total operating expenses by approximately 1% compared with SO and improved port logistics efficiency by minimizing resources needed for PSC inspections and alleviating port congestion.

As further advancements, Donti et al.[62] developed an E2E method based on stochastic optimization for inventory stock problems with uncertain demand, enabling more informed decision-making with observable feature data. Moreover, Qi et al.[63] proposed a DL framework for multi-period inventory replenishment under uncertain demand and vendor lead times, where deep learning models directly output suggested replenishment quantities from input features without intermediate predictions. Compared with SO, this approach reduced holding costs by 26.1% and average stockout costs by 51.7%, demonstrating substantial improvements in operational efficiency.

Intelligent autonomous control

Intelligent autonomous control refers to systems capable of making real-time decisions and performing tasks, such as autonomous driving, drone control, and robotics, in dynamic and uncertain environments. Due to the complexity of these tasks and the variability in operating conditions, formulating explicit optimization objectives is often impractical.

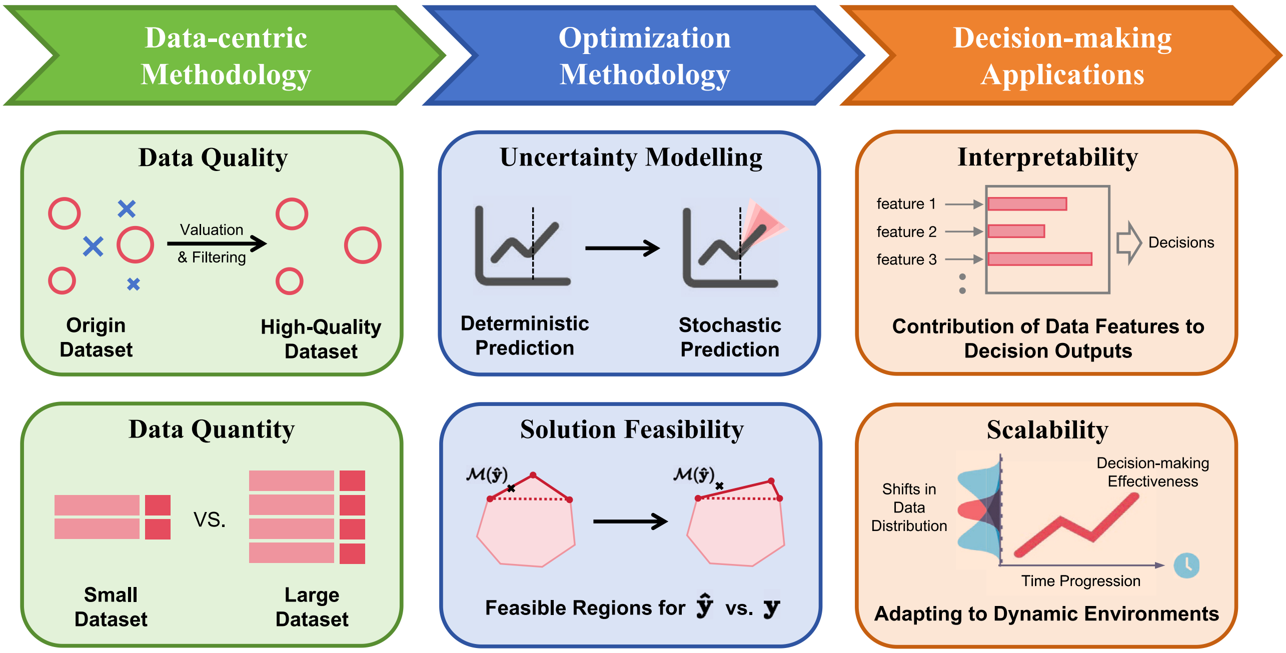

Figure 6. Overview of challenges in data-driven optimization

As a result, DL offers an effective approach to adapt to changing conditions and enhance data-driven decision-making.

Cao et al.[64] introduced a dynamic confidence-aware RL approach to address uncertainty in the performance of self-driving vehicles under extreme conditions. By leveraging the worst confidence value, this method minimizes value function estimation errors, enabling stable improvement to autonomous driving strategies. Kim et al. applied RL to train a drone control system, integrating wing strain sensors, drone posture, and wind data to optimize flight altitude. In indoor fight experiment, the system's robustness improved as training progressed.[65]

IL further enhances autonomous decision-making by training on high-quality expert samples.[66] Behavior cloning (BC), inverse reinforcement learning (IRL), and generative adversarial imitation learning (GAIL) align an agent's strategy with expert demonstrations, with BC matching state-action trajectories, IRL improving generalization through reward functions, and GAIL generating policies that closely follow the expert's, ensuring stable and robust robotic control.[67]

CHALLENGES

This section explores the challenges and potential solutions in applying the data-driven optimization approaches discussed earlier. The insights are organized into three key areas: data-centric methodology, optimization methodology, and decision-making applications, as shown in Figure 6. By addressing these aspects, we aim to provide researchers and practitioners with a deeper understanding of the critical considerations and advancements needed to enhance the integration of prediction and decision-making across domains.

Data-centric methodology

As artificial intelligence and related technologies continue to evolve, data have become a critical resource driving innovation in AI systems.[68] This shift has given rise to data-centric AI, which focuses on optimizing data to maximize its value in intelligent applications. Likewise, the quality and quantity of data play a pivotal role in determining the efficiency and effectiveness of decision-making.

Data quality

With the continuous advancement of digital technologies, industries are increasingly driving intelligent data collection and storage, leading to a surge in data volumes. For example, in energy systems, petabytes of data are generated annually.[69] However, this increase in data does not necessarily lead to improved decision-making efficiency. In fact, the role of data in tasks is not uniform; without proper preprocessing, low-quality data can hinder training, obstruct the extraction of key patterns, and ultimately reduce the effectiveness of applications.[70] These negative impacts not only reduce prediction accuracy but also diminish decision-making efficiency, underscoring the importance of accurately identifying data value—a key challenge for data-driven optimization. While most data-driven techniques mentioned above rely on model-based advancement, unlocking the potential of deterministic data usage at the front end can further enhance decision-making performance.[71]

Current methods for evaluating or identifying high-quality data can be broadly categorized into unsupervised and supervised approaches. Unsupervised techniques include general data cleaning, consistency analysis,[72] Shannon entropy,[73] and outlier detection,[74] which have been effectively applied to multi-scale time series datasets, such as wind power prediction in day-ahead scheduling.[75] These methods can efficiently distinguish data computationally, but they also face a trade-off with valuation accuracy. This is because static, one-size-fits-all evaluation methods are insufficient for decision-focused needs, as the importance of the same data may vary under different conditions, such as typical versus extreme climate scenarios. This underscores the need for supervised methods. While most current supervised data valuation approaches target prediction tasks,[76] integrating decision-making insights into data valuation remains an underexplored and highly promising avenue.

As actionable insights, to further enhance the efficiency of data-driven optimization, we need more decision-focused data valuation methods. Wang et al.[77] proposed a general learning-based data valuation method that calibrates the valuation distribution through dynamic sampling and pre-defined data analytics validation. This approach achieves an unbiased estimate that maximizes target value, improving prediction accuracy by 7.69% in weather-coupled power forecasting tasks, and provides a solid foundation for decision-focused data valuation.[78] Other methods, such as dynamic distribution-based approach,[79] should also be extended to the cross-domain decision-making field through the theoretical expansion of E2E frameworks.

Data quantity

In iterative training, insufficient data can lead to model overfitting, while excessive data can substantially increase training costs and computational time. This challenge is particularly pronounced in data-driven decision-making tasks involving E2E methods such as implicit differentiation, where overly large datasets result in prohibitive computational overheads, yet overly small datasets fail to capture the full spectrum of critical decision scenarios. Achieving a balance by selecting an appropriately scaled training dataset remains a critical and complex challenge. Research has shown that a data volume two orders of magnitude smaller than the theoretical upper bound is often sufficient to achieve high-accuracy decision-making in practice.[26] This highlights the potential for optimizing data scale to enhance computational efficiency, suggesting that leveraging smaller data samples can still deliver substantial decision-making benefits.

Techniques such as data augmentation and generative adversarial networks have been shown to extend training datasets and improve data utility.[80,81] However, these methods primarily focus on enriching text and image data, and their applicability to data-driven decision-making remains an area for further exploration. Other strategies, such as few-shot learning,[82] help generalize key information from limited datasets and enhance predictive performance, offering potential solutions for reducing redundant data usage in data-driven optimization.

Optimization methodology

In previous sections, we have focused on point forecasting among data-driven optimization methods. This section further explores more complex prediction to decision-making couplings in real-world applications, along with the associated risks in parameter transmission.

Uncertainty modeling

As optimization theory and decision-making requirements evolve, the modeling of decision-making in complex systems has become increasingly sophisticated. Predictive analytics derived from industry data often transcends deterministic frameworks due to inherent uncertainties in forecasting unknown variables. This has led to the adoption of probabilistic and statistical optimization techniques, such as expectation-based and min-max approaches, to address these complexities. For instance, in power systems, the characterization of renewable energy sources such as wind and solar has shifted from deterministic outputs to stochastic distributions or non-parametric, feature-driven probabilistic models, enabling a more comprehensive representation of random variability. This paradigm shift necessitates advancements in stochastic and robust optimization frameworks, driving the expansion of data-driven optimization methodologies to better accommodate the uncertainties inherent in real-world systems. Donti et al.[62] pioneered the use of expectation-based implicit differentiation in E2E frameworks, effectively addressing uncertain problems such as price forecasting and battery storage optimization.

Moreover, uncertainty modeling has evolved to include more flexible approaches, such as chance constraints, which transform deterministic constraints into probabilistic ones, ensuring they are met with a specified probability under a given distribution.[83] These methods have been applied to safety-critical issues, such as renewable energy grid dispatching,[84] but challenges remain in the theoretical formulation of E2E gradient propagation. In addition, robust optimization techniques use uncertainty to guide decision-making and mitigate worst-case scenarios,[85] with successful applications in power system problems such as unit commitment.[86] Moreover, approaches such as distributionally robust optimization[15] have matured and been applied in multi-source operations.[87] In summary, data-driven optimization methods must further develop their theoretical foundations for uncertainty modeling and gradient computation to meet the growing demands of high-dimensional optimization in both academia and industry.

Solution feasibility

Beyond theoretically minimizing decision-making suboptimality through advanced E2E methods, it is crucial to ensure that decisions based on predicted parameters remain feasible when applied to actual parameters. This requires maintaining continuity in the training process to avoid the applied loss (Equation 3) leading to infeasible or meaningless solutions. However, the feasible region defined by predicted parameters often diverges from the true feasible region. For instance, renewable energy output predictions inherently form balance constraints, and misalignment between the feasible regions in the predictive and actual operational environments can potentially threaten the efficiency of unbiased power management training.

To mitigate the risks to system security, researchers have explored strategies such as imposing high penalty costs for violating safety constraints[88] and incorporating these constraints into the objective function as risk measures.[89] In addition, post-hoc corrections, where constraint terms are adjusted through parameter iterations, have been used to ensure feasibility before minimizing decision loss.[62] Other techniques, such as projection-based mapping to a feasible region, have also been employed. However, more comprehensive theoretical frameworks to ensure feasibility within a closed-loop training process are still under development.

Decision-making application

Compared with forecasting problems, decision-making is more constrained by the specific nature of the optimization content, which introduces potential risks in practical applications. This section discusses two critical challenges: interpretability—how well the decision can be understood and explained, and scalability—how to adapt to changing environments and non-stationary data.

Interpretability

Interpretability is essential for enabling decision-makers to trust and act confidently on predictions by understanding the underlying factors and rationale. In the context of black box prediction models (e.g., machine learning, deep learning), interpreting the fitting logic can provide valuable insights into better predictions, prediction-to-decision processes, and direct decision-making for frameworks such as SO, E2E, and DL. For boosting tree models, feature importance metrics, based on training process splits and the Gini index, quantify the contribution of different input features to the prediction model.[90] Other interpretability techniques, such as local interpretable model-agnostic explanations,[91] which focuses on learning a locally interpretable classification model, and integrated gradients,[92] which ensures feature sensitivity and implementation invariance, offer deeper insights into individual feature contributions. These methods enable decision-makers to better understand how each variable influences predictions training in complex models.

In addition, Shapley value, a method for allocating contributions in cooperative game theory, has been applied to interpret input feature contributions.[93] Shapley additive explanations[94] decomposes model outputs based on feature contributions and analyzes feature interactions at multiple scales. This has been particularly useful in domains such as healthcare, providing patient-physician decision analysis.[95] Furthermore, physics-informed neural networks enhance interpretability by integrating neural networks with physical laws, ensuring that predictions align with domain-specific principles, with applications in energy systems.[96] These approaches not only increase interpretability but also improve model reliability by aligning learned patterns with established domain knowledge.

Despite these advances, applying interpretability to decision-making remains a critical challenge. The causal logic linking data to decisions requires further clarification to help decision-makers understand the factors influencing outcomes and potential decision losses within the predictive and optimization framework. This necessitates extending interpretability frameworks traditionally applied in ML or DL to subsequent decision-making processes. Moreover, for theoretically unbiased methods such as E2E, suboptimal predictions that sacrifice local statistical optimality must provide deeper insights and explanations aligned with decision-level optimality. Such understanding must be tailored to the specific optimization context, where interpretability facilitates the intelligent utilization of data and model training.

Scalability

Decision scalability in data-driven decision-making refers to a model's ability to adapt to dynamic environments and evolving data distributions, a challenge particularly pronounced in fields such as finance and environmental monitoring. In such contexts, the underlying data distributions shift over time, rendering traditional models that assume stationarity less effective. To address this, prediction-based decision frameworks must adjust dynamically during the training process to accommodate these changes. This adaptability represents a key challenge for data-driven optimization in dynamic environments and is essential for unlocking its sustained applicability across diverse industries. One effective approach to address this challenge is transfer learning, which allows models to transfer knowledge from related tasks or domains, enabling adaptation to new conditions without retraining from scratch.[97] This can include parameter inheritance or structural fine-tuning to better capture the nuances of new environments. In addition, model-agnostic meta-learning algorithms optimize representations that facilitate rapid adaptation to novel tasks, offering a flexible framework for evolving environments.[98] These methods enable models to generalize across changing contexts, ensuring scalable decision-making.

For non-stationary data, where data patterns continuously shift over time, online learning provides a key strategy. By enabling models to update continuously as new data become available, online learning ensures models remain relevant and effective even in the face of unpredictable changes. This approach maximizes decision-making accuracy by refining the model based on previous learning tasks and new observations, maintaining relevance in fast-changing scenarios.[99] Furthermore, few-shot learning can accelerate adaptation to new tasks with minimal labeled data, a particularly useful tool in environments where data are sparse or expensive to acquire.[100] When combined with robust optimization techniques, these strategies enhance the scalability of data-driven decision-making systems, ensuring they remain accurate and efficient, even under uncertainty and shifting data distributions.

CONCLUSION AND DISCUSSION

As digital technologies continue to advance and decision-making demands grow increasingly specialized and diverse, the need for more efficient solution algorithms has driven the rapid evolution of data-driven optimization. This review provides a comprehensive analysis of the key methodologies in this field, including theoretical frameworks, comparative assessments, and practical applications. It highlights critical challenges and outlines key areas for future research, offering valuable insights for both academic researchers and industry practitioners. By identifying emerging trends and unresolved issues, this review serves as a crucial resource for guiding future innovations in data-driven optimization.

This review provides a systematic examination of three key approaches to data-driven optimization, with a focus on the integration of prediction and decision-making. SO follows a two-stage framework that trains predictive models to optimize for norm-based accuracy, achieving rapid convergence with output parameters for subsequent decision-making. However, this approach is subject to the theoretical risk of suboptimality, influenced by the mismatch of the optimization objective. E2E addresses this limitation by incorporating decision loss into the training of predictive models, ensuring consistency between predictions and decisions. To tackle challenges in gradient transmission through the optimization structure, implicit differentiation employs Lagrangian functions to compute the Jacobian matrix, surrogate loss methods simplify computation by optimizing well-behaved objectives, and approximation methods estimate gradients using expectation-based techniques. However, these methods need to consider computational cost and transferability during training. DL focuses on directly generating decision outputs from predictive models, making it particularly suited for scenarios with complex or implicit optimization structures, as exemplified in reinforcement and imitation learning. This review offers an in-depth comparative analysis of these approaches, exploring their theoretical foundations, structural characteristics, and practical applications. Case studies in power grid scheduling, operations management, and intelligent autonomous control illustrate the real-world relevance of these methods.

Finally, we identify key challenges in data-driven optimization and propose directions for future research. For data-centric approaches, future work should focus on enhancing both data quality and quantity, exploring unsupervised valuation methods such as entropy-based and outlier detection techniques, as well as learning-based methods with adaptive strategy. In addition, data quantity considerations should encompass usage necessity, including strategies such as data augmentation and few-shot learning. From an optimization modeling perspective, incorporating more realistic uncertainty representations, such as stochastic and robust optimization, could better address the needs of complex systems, especially when driven by uncertainty distributions and statistical characteristics. Moreover, potential interference between predicted and true values could lead to decision infeasibility in ground truth validation, necessitating further work on ensuring training coherence. Finally, from an application perspective, dissecting predictive models trained for varying levels of decision-making penetration can enhance decision interpretability. Given the challenges posed by changing environments and non-stationary data streams, future research should also explore methods such as transfer learning and meta-learning to maintain the practical efficiency of data-driven optimization.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (T2121002, 724B2002), the Peking University - BHP Carbon and Climate Wei-Ming PhD Scholars Program (WM202407). Some of the icons used in the images are sourced from Iconfont.cn and Flaticon.com.

AUTHOR CONTRIBUTIONS

Y.W., J.W., and J.S. conceived the study and designed the research framework. Y.W. conducted the literature review, wrote the manuscript, and edited the content. J.W., H.Z., and J.S. supervised the work and provided critical feedback. All authors contributed to the final manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

1. Provost, F., and Fawcett, T. (2013). Data science and its relationship to big data and data-driven decision making. Big Data 1, 51-59.
2. Sarker, I.H. (2021). Data science and analytics: an overview from data-driven smart com-puting, decision-making and applications perspective. SN Comput. Sci. 2, 377.
3. Ahmad, T., Zhang, D., Huang, C., et al. (2021). Artifi-cial intelligence in sustainable energy industry: Status quo, challenges and opportunities. J. Clean. Prod. 289, 125834.
4. Antonanzas, J., Osorio, N., Escobar, R., et al. (2016). Review of photovoltaic power forecasting. Sol. Energy 136, 78-111.
5. Shrestha, A., and Mahmood, A. (2019). Review of deep learning algorithms and architec-tures. IEEE Access 7, 53040-53065.
6. Amos, B., and Kolter, J.Z.O. (2017). Optnet: Differentiable optimization as a layer in neural networks. In International Conference on Machine Learning (PMLR), pp. 136-145.
7. Elmachtoub, A.N., and Grigas, P. (2022). Smart "predict, then optimize". Manag. Sci. 68, 9-26.
8. Hussein, A., Gaber, M.M., Elyan, E., et al. (2017). Imitation learning: A survey of learning methods. ACM Comput. Surv. 50, 1-35.
9. Sutton, R.S., and Barto, A.G. (2018). Reinforcement Learning: An Introduction (MIT press).
10. Kotary, J., Fioretto, F., Van Hentenryck, P., et al. (2021). End-to-end constrained optimization learning: A survey. Preprint at arXiv. https://doi.org/10.48550/arXiv.2103.16378.
11. Qi, M., and Shen, Z.-J. (2022). Integrating prediction/estimation and optimization with applications in operations management. In Tutorials in Operations Research: Emerging and Impactful Topics in Operations (INFORMS), pp. 36-58.

12. Choi, T.-M., Wallace, S.W., and Wang, Y. (2018). Big data analytics in operations man-agement. Prod. Oper. Manag. 27, 1868-1883.

13. Zhang, H., and Li, R. (2024). Decision-oriented learning for future power system decision-making under uncertainty. Preprint at arXiv. https://doi.org/10.48550/arXiv.2401.03680.

14. Zakaria, A., Ismail, F.B., Lipu, M.H., et al. (2020). Uncertainty models for stochastic optimization in renewable energy applications. Renew. Energy 145, 1543-1571.

15. Rahimian, H., and Mehrotra, S. (2019). Distributionally robust optimization: A review. Preprint at arXiv. https://doi.org/10.48550/arXiv.1908.05659.

16. Ferreira, K.J., Lee, B.H.A., and Simchi-Levi, D. (2016). Analytics for an online retailer: Demand forecasting and price optimization. Manuf. Serv. Oper. Manag. 18, 69-88.

17. Burer, S., and Letchford, A.N. (2012). Non-convex mixed-integer nonlinear programming: A survey. Surveys Oper. Res. Manag. Sci. 17, 97-106.

18. Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. Int. J. Forecast. 30, 1030-1081.

19. Chitsaz, H., Zamani-Dehkordi, P., Zareipour, H., et al. (2018). Electricity price forecasting for operational scheduling of behind-the-meter storage systems. IEEE Trans. Smart Grid 9, 6612-6622.

20. Shi, W., Wang, Y., Chen, Y., et al. (2021). An effective two-stage electricity price forecasting scheme. Elec. Power Syst. Res. 199, 107416.

21. Chang, Z., Zhang, Y., and Chen, W. (2019). Electricity price prediction based on hybrid model of adam optimized lstm neural network and wavelet transform. Energy 187, 115804.

22. Koprinska, I., Wu, D., and Wang, Z. (2018). Convolutional neural networks for energy time series forecasting. In 2018 International Joint Conference on Neural Networks (IJCNN) (IEEE), pp. 1-8.

23. Li, S., Jin, X., Xuan, Y., et al. (2019). Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. Adv. Neural Inf. Process. Syst. 32.

24. Zareipour, H., Canizares, C.A., and Bhattacharya, K. (2010). Economic impact of electric-ity market price forecasting errors: A demand-side analysis. IEEE Trans. Power Syst. 25, 254-262.

25. Zhang, J., Wang, Y., and Hug, G. (2022). Cost-oriented load forecasting. Elec. Power Syst. Res. 205, 107723.

26. Besbes, O., and Mouchtaki, O. (2023). How big should your data really be? data-driven newsvendor: Learning one sample at a time. Manag. Sci. 69, 5848-5865.

27. Tappen, M.F., Liu, C., Adelson, E.H., et al. (2007). Learning gaussian conditional random fields for low-level vision. In 2007 IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 1-8.

28. Domke, J. (2012). Generic methods for optimization-based modeling. In Artificial Intelligence and Statistics (PMLR), pp. 318-326.

29. Lee, K., Maji, S., Ravichandran, A., et al. (2019). Meta-learning with differentiable convex optimization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10657-10665.

30. Agrawal, A., Amos, B., Barratt, S., et al. (2019). Differ-entiable convex optimization layers. Adv. Neural Inf. Process. Syst. 32.

31. Bertrand, Q., Klopfenstein, Q., Massias, M., et al. (2022). Implicit differentiation for fast hyperparameter selection in non-smooth convex learning. J. Mach. Learn. Res. 23, 1-43.

32. Blondel, M., Berthet, Q., Cuturi, M., et al. (2022). Efficient and modular implicit differentiation. Adv. Neural Inf. Process. Syst. 35, 5230-5242.

33. Mandi, J., Stuckey, P.J., Guns, T., et al. (2020). Smart predict-and-optimize for hard combinatorial optimization problems. In Proceedings of the AAAI Conference on Artificial Intelligence, *34*, pp. 1603-1610.

34. Xie, Y., Dai, H., Chen, M., et al. (2020). Differentiable top-k with optimal transport. Adv. Neural Inf. Process. Syst. 33, 20520-20531.

35. Wilder, B., Ewing, E., Dilkina, B., et al. (2019). End to end learning and optimization on graphs. Adv. Neural Inf. Process. Syst. 32.

36. Mulamba, M., Mandi, J., Diligenti, M., et al. (2020). Contrastive losses and solution caching for predict-and-optimize. Preprint at arXiv. https://doi.org/10.48550/arXiv.2011.05354.

37. Berthet, Q., Blondel, M., Teboul, O., et al. (2020). Learning with differentiable pertubed optimizers. Adv. Neural Inf. Process. Syst. 33, 9508-9519.

38. Dalle, G., Baty, L., Bouvier, L., et al. (2022). Learning with combinatorial optimization layers: a probabilistic approach. Preprint at arXiv. https://doi.org/10.48550/arXiv.2207.13513.

39. Pogančić, M.V., Paulus, A., Musil, V., et al. (2019). Differentiation of blackbox combinatorial solvers. In International Conference on Learning Representations.

40. Bertsimas, D., and Stellato, B. (2022). Online mixed-integer optimization in milliseconds. Inf. J. Comput. 34, 2229-2248.

41. Tang, Y., Agrawal, S., and Faenza, Y. (2020). Reinforcement learning for integer programming: Learning to cut. In International Conference on Machine Learning (PMLR), pp. 9367-9376.

42. Feurer, M., and Hutter, F. (2019). Hyperparameter optimization. In Automated Machine Learn-Ing: Methods, Systems, Challenges (Springer International Publishing), pp. 3-33.

43. Jamieson, K., and Talwalkar, A. (2016). Non-stochastic best arm identification and hyperparame-ter optimization. In Artificial Intelligence and Statistics (PMLR), pp. 240-248.

44. Wahdany, D., Schmitt, C., and Cremer, J.L. (2023). More than accuracy: end-to-end wind power forecasting that optimises the energy system. Elec. Power Syst. Res. 221, 109384.

45. Sang, L., Xu, Y., Long, H., et al. (2022). Electricity price prediction for energy storage system arbitrage: A decision-focused approach. IEEE Trans. Smart Grid 13, 2822-2832.

46. Mazyavkina, N., Sviridov, S., Ivanov, S., et al. (2021). Reinforcement learning for combinatorial optimization: A survey. Comput. Oper. Res. 134, 105400.

47. Ajagekar, A., and You, F. (2023). Deep reinforcement learning based unit commitment scheduling under load and wind power uncertainty. IEEE Trans. Sustain. Energy 14, 803-812.

48. Pinthurat, W., Surinkaew, T., and Hredzak, B. (2024). An overview of reinforcement learning-based approaches for smart home energy management systems with energy storages. Renew. Sustain. Energy Rev. 202, 114648.

49. Schulman, J., Wolski, F., Dhariwal, P., et al. (2017). Proximal policy optimization algorithms. Preprint at arXiv. https://doi.org/10.48550/arXiv.1707.06347.

50. Lillicrap, T. (2015). Continuous control with deep reinforcement learning. Preprint at arXiv. https://doi.org/10.48550/arXiv.1509.02971.

51. Le Mero, L., Yi, D., Dianati, M., et al. (2022). A survey on imitation learning techniques for end-to-end autonomous vehicles. IEEE Trans. Intell. Transport. Syst. 23, 14128-14147.

52. Hawke, J., Shen, R., Gurau, C., et al. (2020). Urban driving with conditional imitation learning. In 2020 IEEE International Conference on Robotics and Automation (ICRA) (IEEE), pp. 251-257.

53. Shah, S., Wang, K., Wilder, B., et al. (2022). Decision-focused learning without decision-making: Learning locally optimized decision losses. Adv. Neural Inf. Process. Syst. 35, 1320-1332.

54. Pan, X., Zhao, T., Chen, M., et al. (2021). Deepopf: A deep neural network approach for security-constrained dc optimal power flow. IEEE Trans. Power Syst. 36, 1725-1735.

55. Kervadec, H., Dolz, J., Yuan, J., et al. (2022). Constrained deep networks: Lagrangian optimization via log-barrier extensions. In 2022 30th Euro-Pean Signal Processing Conference (EUSIPCO) (IEEE), pp. 962-966.

56. Zhou, K., Fu, C., and Yang, S. (2016). Big data driven smart energy management: From big data to big insights. Renew. Sustain. Energy Rev. 56, 215-225.

57. Zhou, Y., Wen, Q., Song, J., et al. (2024). Load data valuation in multi-energy systems: An end-to-end approach. IEEE Trans. Smart Grid.

58. Chen, X., Yang, Y., Liu, Y., et al. (2022). Feature-driven economic improvement for network-constrained unit commitment: A closed-loop predict-and-optimize framework. IEEE Trans. Power Syst. 37, 3104-3118.

59. Gunasekaran, A., and Ngai, E.W. (2012). The future of operations management: an out-look and analysis. Int. J. Prod. Econ. 135, 687-701.

60. Chu, H., Zhang, W., Bai, P., et al. (2023). Data-driven optimization for last-mile delivery. Complex Intell. Systems 9, 2271-2284.

61. Tian, X., Yan, R., Liu, Y., et al. (2023). A smart predict-then-optimize method for targeted and cost-effective maritime transportation. Transp. Res. Part B Methodol. 172, 32-52.

62. Donti, P., Amos, B., and Kolter, J.Z. (2017). Task-based end-to-end model learning in stochastic optimization. Adv. Neural Inf. Process. Syst. 30.

63. Qi, M., Shi, Y., Qi, Y., et al. (2023). A practical end-to-end inventory management model with deep learning. Manag. Sci. 69, 759-773.

64. Cao, Z., Jiang, K., Zhou, W., et al. (2023). Continuous improve-ment of self-driving cars using dynamic confidence-aware reinforcement learning. Nat. Mach. Intell. 5, 145-158.

65. Kim, T., Hong, I., Im, S., et al. (2024). Wing-strain-based flight control of flap-ping-wing drones through reinforce-ment learning. Nat. Mach. Intell. 6, 992-1005.

66. Bojarski, M. (2016). End to end learning for self-driving cars. Preprint at arXiv. https://doi.org/10.48550/arXiv.1604.07316.

67. Hua, J., Zeng, L., Li, G., et al. (2021). Learning for a robot: Deep reinforcement learning, imitation learning, transfer learning. Sensors 21, 1278.

68. Ferruz, N., Zitnik, M., Oudeyer, P.-Y., et al. (2024). Anniversary ai reflections. Nat. Mach. Intell. 6, 6-12.

69. Wang, J., Gao, F., Zhou, Y., et al. (2023). Data sharing in energy systems. Adv. Appl. Energy 10, 100132.

70. Jain, A., Patel, H., Nagalapatti, L., et al. (2020). Overview and importance of data quality for ma-chine learning tasks. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 3561-3562.

71. Zha, D., Bhat, Z.P., Lai, K.-H., et al. (2023). Data-centric ai: Perspectives and challenges. In Proceedings of the 2023 SIAM International Conference on Data Mining (SDM) (SIAM), pp. 945-948.

72. Cichy, C., and Rass, S. (2019). An overview of data quality frameworks. IEEE Access 7, 24634-24648.

73. Lin, J. (1991). Divergence measures based on the shannon entropy. IEEE Trans. Inf. Theor. 37, 145-151.

74. Blázquez-García, A., Conde, A., Mori, U., et al. (2021). A review on out-lier/anom-aly detection in time series data. ACM Comput. Surv. 54, 1-33.

75. Yu, M., Wang, J., Yan, J., et al. (2022). Pricing information in smart grids: A qual-ity-based data valuation paradigm. IEEE Trans. Smart Grid 13, 3735-3747.

76. Liang, W., Tadesse, G.A., Ho, D., et al. (2022). Advances, challenges and oppor-tunities in creating data for trustworthy ai. Nat. Mach. Intell. 4, 669-677.

77. Wang, Y., Wang, J., Gao, F., et al. (2024). Unveiling value patterns via deep rein-forcement learning in heterogeneous data analytics. Patterns 5, 100965.

78. Wang, Y., and Song, J. (2025). Dissecting renewable uncertainty via decon-structive analysis-based data valuation. IEEE Trans. Ind. Appl. 61, 1626-1635.

79. Ghorbani, A., Kim, M., and Zou, J. (2020). A distributional framework for data valuation. In International Conference on Machine Learning (PMLR), pp. 3535-3544.

80. Lata, K., Dave, M., and KN, N. (2019). Data augmentation using generative ad-versarial network. In Proceedings of 2nd International Conference on Advanced Computing and Software Engineering (ICACSE).

81. Mumuni, A., and Mumuni, F. (2022). Data augmentation: A comprehensive sur-vey of modern approaches. Array 16, 100258.

82. Wang, Y., Yao, Q., Kwok, J.T., et al. (2020). Generalizing from a few examples: A survey on few-shot learning. ACM Comput. Surv. 53, 1-34.

83. Pr'ekopa, A. (2013). Stochastic Programming, 324 (Springer Science & Business Media).

84. Wang, Q., Guan, Y., and Wang, J. (2012). A chance-constrained two-stage sto-chastic program for unit commitment with uncertain wind power output. IEEE Trans. Power Syst. 27, 206-215.

85. Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). Robust Optimization, 28 (Princeton university press).

86. Xiong, P., Jirutitijaroen, P., and Singh, C. (2017). A distributionally robust opti-mization model for unit commitment considering uncertain wind power gener-ation. IEEE Trans. Power Syst. 32, 39-49.

87. Mieth, R., Morales, J.M., and Poor, H.V. (2024). Data valuation from data-driven optimiza-tion. IEEE Trans. Control Netw. Syst.

88. Garcia, J.D., Street, A., Homem-de Mello, T., et al. (2021). Application-driven learning: A closed-loop prediction and optimization approach applied to dy-namic reserves and demand forecasting. Preprint at arXiv. https://doi.org/10.48550/arXiv.2102.13273.

89. Zhang, H., Li, R., Chen, Y., et al. (2023). Risk-aware objective-based forecasting in inertia management. IEEE Trans. Power Syst.

90. Chen, T., and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining, pp. 785-794.

91. Ribeiro, M.T., Singh, S., and Guestrin, C. (2016). "Why should i trust you?" ex-plaining the predic-tions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1135-1144.

92. Sundararajan, M., Taly, A., and Yan, Q. (2017). Axiomatic attribution for deep networks. In Inter-national Conference on Machine Learning (PMLR), pp. 3319-3328.

93. Ghorbani, A., and Zou, J. (2019). Data shapley: Equitable valuation of data for machine learning. In International Conference on Machine Learning (PMLR), pp. 2242-2251.

94. Lundberg, S.M., and Lee, S.-I. (2017). A unified approach to interpreting model predic-tions. Adv. Neural Inf. Process. Syst. 30.

95. Wang, Y., Zhao, Y., Song, J., et al. (2022). What drives patients to choose a physician online? a study based on tree models and shap values. In 2022 IEEE 18th International Conference on Automation Science and Engineering (CASE) (IEEE), pp. 1676-1683.

96. Misyris, G.S., Venzke, A., and Chatzivasileiadis, S. (2020). Physics-informed neural networks for power systems. In 2020 IEEE Power & Energy Society General Meeting (PESGM) (IEEE), pp. 1-5.

97. Weiss, K., Khoshgoftaar, T.M., and Wang, D. (2016). A survey of transfer learning. J. Big Data 3, 9.

98. Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In International Conference on Machine Learning (PMLR), pp. 1126-1135.

99. Hoi, S.C., Sahoo, D., Lu, J., et al. (2021). Online learning: A comprehensive sur-vey. Neurocomputing 459, 249-289.

100. Parnami, A., and Lee, M. (2022). Learning from few examples: A summary of approaches to few-shot learning. Preprint at arXiv. https://doi.org/10.48550/arXiv.2203.04291.