

FROM DIFFERENTIAL CALCULUS TO DECISION INTELLIGENCE: DERIVATIVE-BASED MODELS FOR MANAGERIAL OPTIMIZATION

Shashikumar R

Associate Professor
Department of Mathematics, Government First Grade College, Yelahanka

ABSTRACT

Differential calculus has long provided managers and economists with a compact language for reasoning about marginal change, optimality and resource allocation. Yet, contemporary managerial decision environments are characterised by high-dimensional data, nonlinear constraints, uncertainty and increasingly automated analytical pipelines. This paper develops a conceptual bridge from derivative-based managerial optimization to decision intelligence (DI). The study combines an integrative literature synthesis with a targeted bibliometric evidence map of 24 openly accessible scholarly records published from 2016 through 2024. The evidence map is deliberately transparent and is not presented as a substitute for a database-wide Scopus or Web of Science export. The corpus was coded across optimization and calculus, differentiable optimization, prescriptive analytics, artificial intelligence and decision support, decision-focused learning, and bibliometric/methodological research. The analysis indicates a clear movement from isolated optimization procedures toward integrated data-to-decision pipelines. On this basis, the paper proposes a Derivative-to-Decision Intelligence (D2DI) framework in which derivatives function as local decision signals, optimization translates those signals into feasible actions, sensitivity analysis exposes the robustness of solutions, and AI-enabled predictive or prescriptive components extend the framework to contextual and uncertain environments. The paper formalizes representative models for pricing, workforce allocation, marketing expenditure, production and investment decisions; distinguishes analytical, numerical and learned optimization; and specifies boundary conditions under which derivative-based reasoning is appropriate. The principal contribution is therefore conceptual integration rather than a claim that differential calculus is a new managerial technique. The paper argues that its enduring value lies in providing an interpretable mathematical layer within contemporary prescriptive and decision-intelligence systems.

Keywords: differential calculus; marginal analysis; managerial optimization; decision intelligence; prescriptive analytics; differentiable optimization; decision-focused learning; artificial intelligence; sensitivity analysis; mathematical decision models.

1. INTRODUCTION

Managerial decision-making has progressively moved from rule-based judgement toward quantitative, data-supported and increasingly automated forms of analysis. This transition has not displaced classical mathematics; rather, it has altered the role that mathematical reasoning plays in a decision architecture. Differential calculus is a particularly important example. A derivative expresses the local rate at which an objective changes when a decision variable is perturbed. In managerial economics, this idea underlies marginal revenue, marginal cost, marginal productivity and elasticity. In optimization, derivatives identify stationary points and local directions of improvement. In modern machine learning, gradients and Hessians are

central to parameter estimation, while automatic differentiation has made derivative computation an integral part of computational models (Baydin et al., 2018).

The contemporary significance of derivatives becomes clearer when they are viewed as components of a wider decision pipeline. Prescriptive analytics was conceptualised as a stage that moves beyond describing historical patterns and predicting future states toward recommending actions (Frazzetto et al., 2019; Lepenioti et al., 2020). At the same time, research on decision-focused learning has argued that predictive accuracy should not be treated as an adequate proxy for downstream decision quality. Wilder et al. (2019) explicitly integrated predictive models with optimization, while Elmachtoub and Grigas (2022) developed the Smart Predict-then-Optimize framework to align prediction with the structure of the subsequent optimization problem. These developments create an opportunity to reconsider differential calculus not as a historical mathematical topic but as one interpretable layer within a data-to-decision architecture.

The central research problem addressed in this paper is therefore integrative. Literature on calculus and marginal analysis explains local change and optimality; operations research formalises constraints and objective functions; prescriptive analytics connects analytical outputs with recommended actions; and AI-enabled decision support introduces learning and adaptation. These streams are highly relevant to one another but are often reviewed in disciplinary silos. The paper asks how derivative-based reasoning can be conceptually connected to decision intelligence without overstating what derivatives can accomplish in discontinuous, stochastic, nonconvex or discrete managerial environments.

1.1 Research questions

- **RQ1.** How has the open-access literature evolved from derivative-based optimization toward data-driven and intelligent decision systems?
- **RQ2.** What conceptual role can derivatives, gradients and sensitivity measures play in managerial decision intelligence?
- **RQ3.** How can derivative-based optimization be integrated with predictive analytics, prescriptive analytics and AI?
- **RQ4.** Under what conditions are derivative-based managerial models informative, and where do their assumptions become restrictive?

1.2 Contributions

- It synthesizes mathematical optimization, prescriptive analytics, differentiable optimization and decision-focused learning into one managerial framework.
- It provides a transparent targeted bibliometric evidence map based on 24 open-access scholarly records published through 2024.
- It proposes the Derivative-to-Decision Intelligence (D2DI) framework, linking local marginal information to optimization, sensitivity analysis and managerial action.
- It develops domain-level model templates for pricing, marketing, human resources, operations, supply chain and finance.
- It identifies boundary conditions and a research agenda for integrating interpretable mathematical optimization with AI-enabled decision systems.

2. THEORETICAL FOUNDATIONS

2.1 Derivatives as local decision signals

Let an objective function be represented by $y = f(x)$, where x is a controllable managerial variable. The derivative $f'(x)$ represents the instantaneous change in the objective associated with an infinitesimal change in x . In a managerial setting, this can be interpreted as a marginal decision signal. For example, if $\pi(q)$ denotes profit from quantity q , then $\pi'(q)$ is the marginal change in profit associated with an incremental change in output. An interior optimum requires a stationary condition $\pi'(q^*) = 0$, while a second-order condition such as $\pi''(q^*) < 0$ supports a local maximum. The mathematical structure is elementary, but its managerial interpretation remains powerful because it makes the mechanism of an optimum explicit.

The distinction between a decision level and a decision signal is important. The derivative does not by itself select an action. It identifies the direction and local magnitude of change. An optimization procedure combines that information with the objective, constraints and feasible set. This distinction becomes increasingly important in data-rich environments, where an AI system may estimate demand or risk but still require an optimization layer to translate estimates into an actionable allocation.

2.2 First- and second-order conditions

For an unconstrained differentiable maximization problem, an interior candidate x^* satisfies $f'(x^*) = 0$. The second derivative provides local curvature information. If $f''(x^*) < 0$, the point is a strict local maximum; if $f''(x^*) > 0$, it is a strict local minimum. For a multivariable objective $f(x)$, the gradient $\nabla f(x)$ generalises the first derivative and the Hessian $H(x)$ describes local curvature. These concepts form the mathematical core of many modern optimization algorithms.

The practical managerial value lies in interpretation. A gradient tells a manager which direction would most rapidly increase an objective under a local approximation. A Hessian describes whether the local landscape is flat, steep, saddle-shaped or strongly curved. In machine learning, gradient descent operationalises the same logic at scale; gradient and stochastic-gradient methods remain fundamental to large-scale learning, although nonconvexity creates theoretical and practical complications (Jin et al., 2021).

2.3 Constrained optimization

Real managerial decisions are rarely unconstrained. A general formulation is:

$$\max f(x) \text{ subject to } g_i(x) \leq 0, \quad i = 1, \dots, m$$

When regularity conditions hold, the Lagrangian approach combines the objective and constraints. The resulting Karush–Kuhn–Tucker conditions provide a more general optimality framework. This is conceptually important for management because resources, budgets, labour availability, capacity, regulatory requirements and service-level constraints frequently define the feasible region. Derivative-based reasoning therefore becomes most useful when it is embedded within a complete optimization model rather than treated as a stand-alone calculation.

2.4 From calculus to differentiable optimization

A major conceptual transition occurred when optimization problems themselves became differentiable components of learning systems. OptNet demonstrated that quadratic programs can be embedded as layers in neural networks, with sensitivity analysis and implicit

differentiation enabling gradients to pass through the optimization layer (Amos & Kolter, 2017). Agrawal et al. (2019) extended this idea to disciplined convex programs, making it possible to differentiate analytically through a broad class of optimization programs. Automatic differentiation provides the computational infrastructure for efficiently obtaining derivatives of complex computational graphs (Baydin et al., 2018).

This development is directly relevant to the proposed D2DI framework. The derivative is no longer only the final calculation performed by an analyst; it can become a computational interface between data-driven models and explicit decision constraints. Such architectures allow the objective of a decision problem to influence the learning process rather than being imposed only after a predictive model has been trained.

3. METHODOLOGY

3.1 Research design

The study adopts an integrative conceptual review supported by a targeted open-access bibliometric evidence map. The purpose is not to estimate a causal effect or statistically pool empirical results. Instead, the methodology combines structured literature identification, bibliometric descriptive mapping, thematic coding and conceptual synthesis. The reporting logic follows the transparency principles associated with PRISMA 2020, while the bibliometric component follows established science-mapping guidance (Donthu et al., 2021; Page et al., 2021).

Importantly, the present bibliometric component should be interpreted as a targeted evidence map rather than as a database-census claim. The corpus was assembled from open scholarly sources and openly accessible repositories and publisher/repository records including PMLR, JMLR, AAAI, PMC, Springer, Oxford Academic, INFORMS, arXiv and related scholarly repositories. A full Scopus/WoS export would be required for a journal article that reports database-wide publication counts, citation rankings or formal co-citation statistics.

3.2 Eligibility and coding

Table 1. Eligibility criteria applied to the targeted evidence corpus.

Criterion	Inclusion rule	Exclusion rule
Accessibility	Open-access full text or openly accessible scholarly version	Records with no accessible scholarly version
Topical relevance	Optimization, derivatives/gradients, differentiable optimization, prescriptive analytics, AI/DSS, decision-focused learning	Unrelated mathematical or AI applications
Publication type	Peer-reviewed journal articles, conference papers, scholarly reviews and open research papers	Blogs, commercial reports and non-scholarly webpages
Language	English	Other languages
Analytical relevance	Contains a method, review, framework or conceptual result relevant to the data-to-decision chain	Purely descriptive technology commentary

3.3 Bibliometric procedures

The selected records were coded by publication year, methodological orientation and dominant conceptual theme. Because the corpus was intentionally targeted rather than generated from a single database API, the analysis focuses on descriptive temporal and thematic patterns rather than citation-count league tables. The thematic categories are non-exclusive: a paper may simultaneously contribute to optimization and prescriptive analytics, for example. This approach is consistent with the broader use of bibliometric mapping to reveal thematic evolution and research structure (Donthu et al., 2021; van Eck & Waltman, 2010).

4. BIBLIOMETRIC EVIDENCE MAP

The targeted corpus contains 24 anchor records distributed across the 2016–2024 period. The temporal pattern is noteworthy: the early part of the corpus is dominated by gradient-based optimization and automatic differentiation, whereas the later period increasingly includes prescriptive analytics, decision-focused learning, data-driven optimization and AI-enabled decision support.




Figure 1. Temporal distribution of the 24-record targeted open-access/openly accessible evidence corpus.

The pattern should not be read as an estimate of the entire literature. It is an evidence-map signal showing the growing convergence of optimization and learning-oriented decision research. The 2020–2024 concentration reflects the expansion of data-driven optimization, prescriptive analytics and learning-to-optimize research.




Figure 2. Thematic distribution of the targeted evidence corpus. Coding categories are non-exclusive.

Thematic coding indicates that optimization and calculus remain the foundational layer, while differentiable optimization and prescriptive analytics provide the principal bridges to contemporary AI and decision systems. The emergence of decision-focused learning is particularly relevant because it changes the objective of predictive modelling: a model can be judged by the quality of the decision it enables rather than only by predictive accuracy (Wilder et al., 2019; Elmachoub & Grigas, 2022).

4.1 Evolution of the intellectual structure

Table 2. Temporal evolution of the dominant evidence stream, 2016–2024.

Period	Dominant evidence stream	Interpretation for managerial optimization
2016–2018	Gradient methods, online optimization, automatic differentiation	Calculus becomes computational infrastructure for large-scale learning and adaptive optimization.
2019	Prescriptive analytics, decision-focused learning, differentiable optimization	Prediction and optimization begin to be treated as a connected pipeline.
2020–2021	Bilevel/data-driven optimization, implicit differentiation, prescriptive analytics	Optimization increasingly incorporates data, uncertainty and learned parameters.
2022–2023	Learning to optimize, predictive-to-prescriptive pipelines, gradient computation	The boundary between algorithm design and learning becomes increasingly permeable.
2024	AI/business bibliometrics, learned optimization, context-aware decision generation	Research shifts toward adaptive decision generation and broader organizational applications.

4.2 Conceptual positioning




Figure 3. Conceptual positioning of the derivative-to-decision research stream.

Figure 3 is a conceptual positioning map rather than a statistical scale. It illustrates the central argument of the paper: classical calculus sits at the foundation, while contemporary decision intelligence increases both the mathematical complexity of the decision representation and the degree of integration between analysis and action.

5. Derivative-to-Decision Intelligence (D2DI) Framework

The paper proposes the Derivative-to-Decision Intelligence framework as a synthesis of six streams: differential calculus, mathematical optimization, sensitivity analysis, predictive analytics, prescriptive analytics and AI-enabled decision support. The framework treats derivatives as interpretable local signals rather than as complete decision rules.




Figure 4. Derivative-to-Decision Intelligence (D2DI) framework.

The framework has eight stages. First, the managerial problem is translated into controllable variables. Second, an objective function formalises the outcome to be maximised or minimised. Third, derivatives or gradients identify local changes. Fourth, constraints and optimality conditions determine feasible solutions. Fifth, sensitivity analysis examines how the solution changes when parameters move. Sixth, predictive and optimization models are integrated when future states are uncertain. Seventh, the output is interpreted through a decision-intelligence layer. Finally, the recommended action is implemented and observed, generating feedback for model updating.

5.1 MATHEMATICAL CORE

For a generic managerial problem, let $x \in \mathbb{R}^n$ denote the vector of controllable decisions and θ the vector of contextual parameters. The proposed representation is:

$$x^*(\theta) = \arg \max_x F(x, \theta) \quad \text{subject to} \quad g(x, \theta) \leq 0$$

The derivative layer computes $\nabla_x F$ and, where appropriate, the Hessian $H_{xx}F$. The optimization layer solves for x^* . The sensitivity layer studies $\partial x^* / \partial \theta$ or related comparative-static quantities. The decision-intelligence layer combines the resulting solution with business rules, risk preferences, forecasts and human judgement.

This separation is important because a derivative is local whereas a managerial decision may be global. A zero gradient does not guarantee a global optimum in a nonconvex landscape. Likewise, a mathematically optimal solution may be infeasible from a governance, ethical, contractual or implementation perspective. The D2DI framework therefore places feasibility, sensitivity and managerial interpretation after derivative calculation rather than treating the stationary point as the final answer.

5.2 Sensitivity as an interpretability mechanism

Sensitivity analysis is the bridge between mathematical optimality and managerial interpretation. Suppose the optimal decision depends on demand, wage rates, interest rates or capacity. A small change in these parameters can alter the optimal decision. The manager therefore needs more than x^* : the manager needs an understanding of how x^* responds to the environment. In this sense, derivatives provide not only an optimization mechanism but also a language for explaining the stability and responsiveness of recommendations.

6. MANAGERIAL APPLICATIONS

6.1 Pricing and revenue management

Let demand be $q(p)$, where p is price. Revenue is $R(p) = p \cdot q(p)$, while profit may be represented as $\pi(p) = p \cdot q(p) - C(q(p))$. The derivative $\pi'(p)$ identifies the local effect of a price change on profit. A manager can therefore use the derivative to identify a candidate optimum and the second derivative to evaluate local curvature. In practice, however, $q(p)$ must be estimated from data and may be nonlinear, segmented or affected by competitors. The modern D2DI interpretation is therefore: estimate demand $\rightarrow$ construct objective $\rightarrow$ differentiate or optimize $\rightarrow$ test sensitivity $\rightarrow$ recommend price.

6.2 Workforce and HR optimization

Workforce planning can be formulated as an optimization problem in which labour quantity, skill mix, training investment and overtime are decision variables. Let L denote workforce size and T training expenditure. A simplified productivity objective can be written as $P(L, T)$, subject to budget and service constraints. Partial derivatives P_L and P_T represent marginal productivity effects, while cross-partial derivatives can reveal complementarity or substitution between staffing and training. Such a formulation creates a natural link between HR analytics and prescriptive workforce decisions.

The managerial caution is that labour decisions are not fully reducible to a smooth objective function. Employment law, fairness, employee preferences, collective arrangements and discrete staffing decisions may create constraints that require mixed-integer, multi-objective or human-in-the-loop methods. Derivative analysis should therefore support, rather than replace, institutional judgement.

6.3 Marketing expenditure

Let A denote advertising expenditure and $S(A)$ the expected sales response. Profit can be represented as $\pi(A) = m \cdot S(A) - A - C$, where m is contribution margin and C captures other costs. The marginal return to advertising is $m \cdot S'(A) - 1$. The interior optimum occurs where the marginal contribution equals the marginal expenditure, subject to budget and strategic constraints. With data-driven demand models, the function $S(A)$ can be estimated using machine learning, after which the optimization layer translates the estimated response into an expenditure decision.

6.4 Operations and capacity

For production quantity q , total cost $TC(q)$ and revenue $R(q)$ define profit $\pi(q) = R(q) - TC(q)$. The first-order condition $R'(q) = TC'(q)$ recovers the familiar marginal-revenue/marginal-cost rule. In modern production environments, the problem may include capacity, machine availability, demand uncertainty and setup costs. These conditions motivate extensions to constrained nonlinear programming, mixed-integer optimization, simulation optimization and learned policies. Data-driven and bilevel methods illustrate how

optimization can be adapted when the classical analytical form is insufficient (Beykal et al., 2020).

6.5 Finance and investment

Investment decisions can be represented through objective functions involving return, risk and liquidity. For a continuous portfolio vector w , an illustrative formulation is $\max \mu'w - \lambda w' \Sigma w$ subject to $1'w = 1$ and additional constraints. The gradient identifies marginal effects of portfolio weights, while the Hessian is linked to curvature and risk structure. Contemporary learning-to-optimize research demonstrates how contextual information can be integrated into portfolio decisions, although such approaches must be assessed for robustness and domain shift (Babier et al., 2024).

6.6 Supply chain and maintenance

The predict-then-optimize paradigm is particularly relevant to supply chains because decisions depend on uncertain demand, machine failures and inventory states. van Staden et al. (2022), for example, combined predictive failure estimation with optimization of preventive maintenance timing. Such work illustrates the transition from a static calculus problem to a sequential decision problem in which estimates of future states become inputs to an optimization policy.

7. FROM PREDICTIVE ANALYTICS TO DECISION INTELLIGENCE

A major conceptual distinction in the literature is the difference between predicting an outcome and selecting an action. Predictive analytics estimates what may happen; prescriptive analytics seeks to identify what should be done given objectives and constraints. Frazzetto et al. (2019) describe prescriptive analytics as the phase that closes the gap between data and decisions, while Lepenioti et al. (2020) identify optimization as a major mechanism for prescriptive decision-making.

Decision-focused learning takes this logic further. Wilder et al. (2019) argue that the predictive model should be trained in relation to the downstream optimization task. Elmachtoub and Grigas (2022) formalise this through the SPO framework, which incorporates the structure of the optimization problem into prediction. These contributions suggest that decision intelligence should be evaluated at the level of decision quality, not merely predictive accuracy.

The D2DI framework adds a mathematical interpretability layer to this pipeline. Instead of treating optimization as an opaque final-stage algorithm, it exposes gradients, marginal effects, constraints and sensitivity. The resulting architecture can be summarised as:

***Data* → *Prediction* → *Derivative/Gradient* → *Optimization* → *Sensitivity* →
Recommendation → *Action* → *Feedback***

8. THE ROLE OF AI AND LEARNING-TO-OPTIMIZE

Learning-to-optimize research demonstrates a second direction of integration: instead of using optimization to train a predictive model, machine learning can be used to improve the optimization algorithm itself. Chen et al. (2022) surveyed this paradigm and highlighted its potential efficiency on repeated problem distributions while noting generalisation limitations outside the training distribution. Tang and Yao (2024) similarly describe learn-to-optimize as a data-driven approach for accelerating algorithm configuration.

This creates a two-way relationship. First, optimization can become a layer inside an AI model, as in OptNet and differentiable convex optimization layers (Amos & Kolter, 2017;

Agrawal et al., 2019). Second, AI can learn how to configure or accelerate optimization algorithms (Chen et al., 2022; Tang & Yao, 2024). The D2DI framework places both directions inside a managerial decision architecture.

9. COMPARATIVE METHODOLOGICAL POSITIONING

Table 3. Comparative positioning of optimization and analytics approaches within the D2DI framework.

Approach	Primary strength	Typical limitation	D2DI interpretation
Classical calculus	Interpretability and analytical transparency	Requires smooth functional relationships	Core marginal-signal layer
Nonlinear optimization	Handles complex objectives and constraints	May require numerical solvers and careful initialization	General optimization layer
Simulation optimization	Useful for stochastic and complex systems	Computational cost and noisy objective evaluation	Scenario/robustness layer
Machine learning	Flexible prediction from high-dimensional data	Prediction accuracy may not equal decision quality	Context and prediction layer
Prescriptive analytics	Connects analytics with recommended actions	Can be methodologically broad and heterogeneous	Decision-output layer
Differentiable optimization	Allows optimization to interact with learning	Requires differentiability/regularity and computational care	Learning–optimization interface
Decision-focused learning	Aligns learning with downstream decisions	Can be computationally demanding and task-specific	Decision-aware learning layer
Human-in-the-loop DI	Adds contextual, ethical and governance judgement	May reduce automation and require structured interaction	Final governance/action layer

10. BOUNDARY CONDITIONS AND RISKS

The proposed framework should not be interpreted as a universal replacement for other optimization paradigms. Derivative-based models are most informative when the decision variables are sufficiently continuous, the objective and constraints are differentiable or can be appropriately approximated, and local information is meaningful for the decision. Several boundary conditions require alternative or hybrid approaches.

- **Discrete decisions:** staffing counts, facility opening decisions and binary investments may require integer programming.
- **Non-smooth objectives:** penalties, thresholds and piecewise policies can make classical derivatives inadequate.

- **Nonconvexity:** stationary points may be local minima, local maxima or saddle points. Gradient methods can therefore require globalisation strategies.
- **Uncertainty:** stochastic objectives and uncertain constraints may require stochastic, robust or distributionally robust optimization.
- **Model risk:** an analytically elegant optimum is only as credible as the objective function and parameter estimates.
- **Ethical and governance constraints:** fairness, privacy, labour rights and accountability may not be adequately represented by a numerical objective.
- **Distribution shift:** learned optimization policies may perform poorly when future instances differ materially from the training distribution.

These limitations reinforce rather than weaken the proposed framework. The purpose of D2DI is not to make every managerial decision differentiable; it is to provide a disciplined route for determining when derivative-based reasoning can be embedded in a larger decision architecture.

11. MANAGERIAL IMPLICATIONS

For managers, the practical value of the framework lies in changing the question from 'What does the model predict?' to 'What action should change, by how much, under which constraints, and how robust is that recommendation?' This shift encourages a more explicit connection between analytics and resource allocation.

Table 4. Managerial questions reinterpreted through the D2DI framework.

Managerial question	Derivative-based interpretation	Decision-intelligence extension
Should price increase?	Marginal profit from price change	Demand prediction + constrained price optimization
Should staffing increase?	Marginal productivity/cost	Skill mix, service level and fairness constraints
Should advertising increase?	Marginal contribution of expenditure	Customer-response modelling + budget optimization
Should production expand?	Marginal revenue versus marginal cost	Demand uncertainty + capacity optimization
Should maintenance be advanced?	Marginal cost/risk trade-off	Failure prediction + dynamic policy optimization
Should portfolio weights change?	Marginal return/risk effect	Contextual prediction + constrained portfolio optimization

12. RESEARCH AGENDA

- Develop empirical D2DI models using real managerial datasets rather than illustrative functions alone.
- Compare derivative-based, derivative-free and hybrid optimization under identical managerial decision tasks.
- Develop explainable sensitivity dashboards that translate gradients and shadow prices into managerial language.
- Investigate human-AI collaboration in optimization decisions, particularly when objectives contain ethical or qualitative constraints.
- Study robustness of learned optimization policies under distribution shift and changing market conditions.
- Integrate causal inference with derivative-based optimization so that recommendations distinguish correlation from intervention effects.
- Explore agentic AI architectures in which specialized agents construct, validate, stress-test and explain optimization models under human oversight.
- Extend the framework to multi-objective managerial decisions where profit, risk, sustainability and stakeholder outcomes must be jointly considered.

13. DISCUSSION

The literature examined in this paper suggests that the contemporary significance of differential calculus is less about performing symbolic differentiation manually and more about preserving a mathematically interpretable representation of change within complex computational systems. Automatic differentiation makes gradients available at scale; differentiable optimization makes optimization itself part of a learning architecture; prescriptive analytics converts analytical outputs into recommended actions; and decision-focused learning aligns model training with downstream decision quality.

The resulting research trajectory can be described as a progression from local mathematics to integrated decision systems. In the classical setting, the analyst solves a function and interprets a stationary point. In the modern setting, the analyst may estimate an objective from data, differentiate through a computational graph, impose constraints, solve an optimization problem, quantify sensitivity and then present a recommendation to a manager. The mathematical core remains recognisable, but its position in the workflow has changed.

This perspective also clarifies the relationship between classical managerial economics and AI. AI does not make derivative-based reasoning obsolete; instead, AI can estimate the functions to be optimized, learn optimization strategies, or provide contextual information that changes the objective or feasible set. Conversely, mathematical optimization can provide structure and constraints that reduce the risk of producing analytically impressive but operationally infeasible recommendations. The strongest future systems are therefore likely to be hybrid rather than purely mathematical or purely data-driven.

14. CONCLUSION

This paper has developed an integrative account of how differential calculus can contribute to contemporary decision intelligence. The central argument is that derivatives should be understood as local decision signals embedded within a broader optimization and governance architecture. The targeted open-access bibliometric evidence map shows an increasing

convergence between optimization, automatic differentiation, prescriptive analytics, decision-focused learning and AI-enabled decision support during 2016–2024.

The proposed D2DI framework formalises this convergence through a sequence of managerial problem formulation, objective construction, derivative or gradient analysis, constrained optimization, sensitivity assessment, predictive and prescriptive integration, decision interpretation and feedback. Its principal theoretical contribution is the bridging of a classical mathematical language with a contemporary data-to-decision architecture. Its managerial contribution is to make the logic connecting a change in a managerial variable to a recommended action more transparent.

The paper nevertheless emphasises that derivatives are not universally sufficient. Discrete choices, non-smooth objectives, uncertainty, nonconvexity, distribution shift and governance requirements necessitate broader optimization and human-in-the-loop approaches. Future empirical research should therefore test D2DI against alternative methods on real managerial datasets and evaluate not only predictive accuracy or computational performance, but also decision quality, interpretability, robustness and implementation outcomes.

REFERENCES

1. Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., & Kolter, J. Z. (2019). Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems* (Vol. 32, pp. 9558–9570). https://web.stanford.edu/~boyd/papers/diff_cvxpy.html
2. Amos, B., & Kolter, J. Z. (2017). OptNet: Differentiable optimization as a layer in neural networks. *Proceedings of the 34th International Conference on Machine Learning*, 70, 136–145. <https://proceedings.mlr.press/v70/amos17a.html>
3. Babier, A., Chan, T. C. Y., Diamant, A., & Mahmood, R. (2024). Learning to optimize contextually constrained problems for real-time decision generation. *Management Science*. Advance online publication. <https://doi.org/10.1287/mnsc.2020.03565>
4. Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2018). Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, 18(153), 1–43. <https://www.jmlr.org/papers/v18/17-468.html>
5. Beykal, B., Avraamidou, S., Pistikopoulos, I. P. E., Onel, M., & Pistikopoulos, E. N. (2020). DOMINO: Data-driven optimization of bi-level mixed-integer nonlinear problems. *Journal of Global Optimization*, 78, 1–36. <https://doi.org/10.1007/s10898-020-00890-3>
6. Bertsimas, D., & Koduri, N. (2022). Data-driven optimization: A reproducing kernel Hilbert space approach. *Operations Research*, 70(1), 454–471. <https://doi.org/10.1287/opre.2020.2069>
7. Bertrand, Q., Klopfenstein, Q., Blondel, M., Vaiter, S., Gramfort, A., & Salmon, J. (2020). Implicit differentiation of Lasso-type models for hyperparameter optimization. *Proceedings of the 37th International Conference on Machine Learning*, 119, 810–821. <https://proceedings.mlr.press/v119/bertrand20a.html>
8. Chen, T., Chen, X., Chen, W., Heaton, H., Liu, J., Wang, Z., & Yin, W. (2022). Learning to optimize: A primer and a benchmark. *Journal of Machine Learning Research*, 23(189), 1–59. <https://www.jmlr.org/papers/v23/21-0308.html>

9. Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
10. Elmachetoub, A. N., & Grigas, P. (2022). Smart 'Predict, then Optimize'. *Management Science*, 68(1), 9–26. <https://doi.org/10.1287/mnsc.2020.3922>
11. Frazzetto, D., Nielsen, T. D., Pedersen, T. B., & Siksnys, L. (2019). Prescriptive analytics: A survey of emerging trends and technologies. *The VLDB Journal*, 28, 575–595. <https://doi.org/10.1007/s00778-019-00539-y>
12. Ganie, A. G., & Dadvandipour, S. (2023). From big data to smart data: A sample gradient descent approach for machine learning. *Journal of Big Data*, 10, 162. <https://doi.org/10.1186/s40537-023-00839-9>
13. Gupta, S., Modgil, S., Bhattacharyya, S., & Bose, I. (2022). Artificial intelligence for decision support systems in the field of operations research: Review and future scope of research. *Annals of Operations Research*, 308, 215–274. <https://doi.org/10.1007/s10479-020-03856-6>
14. Jin, C., Netrapalli, P., Ge, R., Kakade, S. M., & Jordan, M. I. (2021). On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *Journal of the ACM*, 68(2), Article 11, 1–29. <https://doi.org/10.1145/3418526>
15. Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. *International Journal of Information Management*, 50, 57–70. <https://doi.org/10.1016/j.ijinfomgt.2019.04.003>
16. Lorraine, J., Vicol, P., & Duvenaud, D. (2020). Optimizing millions of hyperparameters by implicit differentiation. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 108, 1540–1552. <https://proceedings.mlr.press/v108/lorraine20a.html>
17. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
18. Patra, A. K., Praharaj, A., Sudarshan, D., & Chhatoi, B. P. (2024). AI and business management: Tracking future research agenda through bibliometric network analysis. *Heliyon*, 10(1), e23902. <https://doi.org/10.1016/j.heliyon.2023.e23902>
19. Poornima, S., & Pushpalatha, M. (2020). A survey on various applications of prescriptive analytics. *International Journal of Intelligent Networks*, 1, 76–84. <https://doi.org/10.1016/j.ijin.2020.07.001>
20. Russo, D., & Van Roy, B. (2018). Learning to optimize via information-directed sampling. *Operations Research*, 66(1), 230–252. <https://doi.org/10.1287/opre.2017.1663>
21. Shah, S., Wilder, B., Perrault, A., & Tambe, M. (2024). Leaving the nest: Going beyond local loss functions for predict-then-optimize. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(13), 14902–14909. <https://doi.org/10.1609/aaai.v38i13.29410>
22. Tang, K., & Yao, X. (2024). Learn to optimize — a brief overview. *National Science Review*, 11(8), nwae132. <https://doi.org/10.1093/nsr/nwae132>

23. van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, 84, 523–538.
<https://doi.org/10.1007/s11192-009-0146-3>
24. van Staden, H. E., Deprez, L., & Boute, R. N. (2022). A dynamic 'predict, then optimize' preventive maintenance approach using operational intervention data. *European Journal of Operational Research*, 302(3), 1079–1096.
<https://doi.org/10.1016/j.ejor.2022.01.037>
25. Wilder, B., Dilkina, B., & Tambe, M. (2019). Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(1), 1658–1665.
<https://doi.org/10.1609/aaai.v33i01.33011658>

APPENDIX

Anchor Corpus Used for the Targeted Evidence Map

The following 24 records constitute the anchor corpus used for the descriptive evidence map. They were selected because they provide openly accessible scholarly versions and directly illuminate one or more links among calculus, optimization, differentiable computation, prescriptive analytics, decision-focused learning and AI-enabled decision support.

Table A1. Anchor corpus (n = 24) coded by year, representative source, and primary theme.

Year	Representative source(s)	Primary coding
2016	Ruder	Gradient optimization
2017	Amos & Kolter	Differentiable optimization
2017	Russo & Van Roy	Online learning/optimization
2018	Baydin et al.	Automatic differentiation
2019	Agrawal et al.	Differentiable convex optimization
2019	Frazzetto et al.	Prescriptive analytics
2019	Wilder et al.	Decision-focused learning
2020	Bertrand et al.	Implicit differentiation
2020	Beykal et al.	Data-driven bilevel optimization
2020	Lepenioti et al.	Prescriptive analytics
2020	Lorraine et al.	Implicit differentiation
2020	Poornima & Pushpalatha	Prescriptive analytics applications
2021	Bertsimas & Koduri	Data-driven optimization
2021	Donthu et al.	Bibliometric methodology
2021	Jin et al.	Nonconvex optimization

Year	Representative source(s)	Primary coding
2021	Page et al.	Systematic-review reporting
2022	Chen et al.	Learning to optimize
2022	Elmachtoub & Grigas	Predict-then-optimize
2022	Gupta et al.	AI + DSS + OR
2022	van Staden et al.	Predict-then-optimize maintenance
2023	Ganie & Dadvandipour	Gradient descent
2024	Babier et al.	Contextual optimization
2024	Patra et al.	AI/business bibliometrics
2024	Shah et al.; Tang & Yao; Salehin et al.	Decision-aware / learn-to-optimize / AutoML