




Data Analytics and Artificial Intelligence

Vol: 6(1), 2026

REST Publisher; ISBN: 978-81-948459-4-2

Website: <http://restpublisher.com/book-series/daai/>

DOI: <https://doi.org/10.46632/daai/6/1/3>



From Data to Decisions: Transforming Analytics with Artificial Intelligence

Omkar Dengi¹, Sheetal Shevkari², *Anita Atiwadkar³

¹MIT Arts, Commerce & Science College, Alandi, Pune, India.

^{2,3}MIT Arts, Commerce & Science College, Alandi, Pune, India.

*Corresponding Author Email: omkardengi08@gmail.com

Abstract: Artificial intelligence has moved from a peripheral tool in the analytics stack to the mechanism that now defines how organisations sense, interpret, and act upon data. This chapter traces that shift and argue that the conventional description of analytics as a four-stage progression - descriptive, diagnostic, predictive, and prescriptive - no longer captures what contemporary AI-enabled systems actually do. Rather than analysts querying data and interpreting output, increasingly the system itself proposes the interpretation, and in agentic configurations, initiates the resulting action. The chapter develops a framework for understanding this transition, examines the specific technical contributions of machine learning, deep learning, and generative AI to the analytics pipeline, and considers what is lost as well as gained when interpretive authority is delegated to a model. Specific attention is given to governance and explainability, not as norms afterthoughts but as structural requirements once decisions are automated at scale. Here are Five industry contexts - healthcare, finance, education, manufacturing, and urban infrastructure - are used to illustrate how these technologies behave differently once situated in a real operational environment, where data quality, regulatory exposure, and human trust vary considerably. The chapter closes by identifying open problems that remain underexamined in the literature: the evaluation of reasoning quality in generative analytics systems, the sustainability cost of large-scale model training for routine analytical tasks, and the question of accountability when an autonomous analytics agent acts on an inference that later proves wrong.

Keywords: artificial intelligence, data analytics, machine learning, generative AI, explainable AI, predictive analytics, prescriptive analytics, agentic AI, responsible AI, decision intelligence.

1. INTRODUCTION

Every organisation now generates more data than it possibly can read, let alone reason about, and this simple observation has quietly reshaped what the word 'analytics' is taken to mean. A decade ago, analytics referred largely to structured reporting: sheets and dashboards, aggregations, and control charts that summarised what had already happened. The analyst's task was to notice patterns a spreadsheet could not surface on its own and to translate them into a recommendation someone else would act upon. That division of labour - machine computes, human interprets - is precisely what generative and agentic AI are dismantling. Advanced AI models now handle interpretation, not just computation. What has changed is that models are now capable of producing the interpretive step itself: a large language model can read a table of quarterly figures and draft a possibly descriptive explaining the deviation, a reinforcement-learning agent can adjust a supply chain in response to a forecast without waiting for managerial approval. and a computer-vision model embedded in a diagnostic pipeline can flag which of ten thousand scans deserves a radiologist's attention first. Consequently, the line between calculating data and interpreting its meaning has blurred or entirely vanished. (Chowa et al., 2026). This chapter takes the position that the technical literature on AI in analytics has been comparatively quick to document capability - new architectures, benchmark scores, deployment case studies - but slower to interrogate what is structurally different about an analytics pipeline once inference and action are delegated to a model rather than a person. That question matters more in some domains than others. A recommendation engine suggesting a product is a low-stakes error if wrong; a clinical triage model or a credit-scoring system is not. The chapter is organised around this asymmetry. Sections 2 through 6 trace the technical evolution of the field, from early statistical analytics through machine learning, deep learning, and the recent incorporation of generative and large language models into the analytics stack. Section 7 addresses explainability and governance as a load-bearing concern rather than an addendum. Section 8 grounds the discussion in five sector-specific case contexts, and Sections 9 and 10 consider open challenges and the implications for practice and research. A note on scope: this chapter does not attempt to catalogue every algorithm or platform currently marketed as AI-powered analytics - that list changes too quickly to be useful in a printed volume, and much of it is marketing rather than a meaningful technical distinction. The intent instead is to give the reader a durable conceptual

vocabulary for evaluating whatever specific tool or vendor claim they encounter next, by focusing on the underlying shift in where interpretive and decisional authority sits within the pipeline.

2. EVOLUTION OF DATA ANALYTICS AND AI

The lineage connecting today's AI-driven analytics to earlier statistical practice is often understated in enthusiastic accounts of the field, which can give the misleading impression that generative and agentic systems represent a clean break from what preceded them. They do not. Regression analysis, control charts, and rule-based expert systems were themselves attempts to encode a form of automated judgement into data processing, and much of what distinguishes contemporary AI analytics is a difference of degree - in the volume of data handled, the dimensionality of the patterns detected, and the autonomy granted to the system - rather than a difference in kind. (LeCun et al., 2015). Three phases are useful for organising this history, though the boundaries between them are porous rather than sharp. The first, running roughly from the 1970s through the early 2000s, was dominated by descriptive analytics: relational databases, OLAP cubes, and reporting tools that answered the question of what happened. Diagnostic capability existed but was largely manual - an analyst formed a hypothesis and tested it against the data rather than the system surfacing candidate explanations itself. The second phase, spanning roughly 2005 to 2018, introduced statistical machine learning into mainstream business use. Decision trees, support vector machines, and ensemble methods such as random forests and gradient boosting made predictive analytics operationally practical outside of specialist research contexts, and platforms began shipping models trained on historical data as a native feature rather than a bespoke project. (Strategy, 2025) This is also the period in which big-data tooling - distributed storage and processing frameworks - matured enough to make training on organisation-scale datasets routine rather than exceptional. The third and current phase, roughly from 2019 onward, is characterised by the arrival of deep learning at production scale and, more recently, by generative and foundation models capable of operating across modalities - text, image, tabular, and increasingly multimodal combinations of these. What distinguishes this phase from the second is not simply larger models but a qualitative shift in what the model is asked to do: rather than predicting a single labelled outcome, contemporary systems are frequently asked to generate open-ended output - an explanation, a synthetic dataset, a recommended action sequence - that has no single correct answer against which accuracy can be trivially measured. This creates an evaluation problem that the field has not fully resolved, a point returned to in Section 9.




FIGURE 1. Three overlapping phases in the evolution of data analytics, showing the shift from descriptive reporting to predictive modelling to generative and agentic inference.

TABLE 1. Types of Analytics and Their Characteristic Questions

Analytics Type	Guiding Question	Representative Techniques	Typical AI Contribution
Descriptive	What happened?	Aggregation, dashboards, OLAP	Anomaly highlighting, automated summarisation
Diagnostic	Why did it happen?	Drill-down, correlation analysis	Causal inference, root-cause modelling
Predictive	What is likely to happen?	Regression, classification, time-series	Deep learning, ensemble forecasting
Prescriptive	What should be done?	Optimisation, simulation	Reinforcement learning, decision agents

Generative/Cognitive	What explanation or artefact can be produced?	LLM-based synthesis, scenario generation	Foundation models, retrieval-augmented generation
----------------------	---	--	---

3. AI TECHNOLOGIES TRANSFORMING DATA ANALYTICS

It is worth being precise about which technical developments are actually doing the work commonly attributed to AI in analytics, since the phrase is often used as a catch-all. Four developments deserve individual treatment.

Automated Feature Engineering and Representation Learning

Classical machine learning pipelines depended heavily on manually engineered features - an analyst deciding, for instance, that a rolling seven-day average was a more informative predictor than a raw daily figure. Deep learning architectures, particularly convolutional and transformer-based models, learn representations directly from raw or lightly processed data, which removes a substantial amount of manual effort but also removes a substantial amount of interpretive transparency, since the learned features rarely correspond to concepts, a human analyst would recognise or have chosen. (Bengio et al., 2013)

Natural Language Interfaces to Structured Data

Large language models are increasingly used as a translation layer between natural-language queries and structured query languages, allowing a non-technical stakeholder to ask a question in plain English and receive a generated SQL query, chart, or narrative summary in response. This lowers the barrier to accessing analytics considerably, but it introduces a specific and under-discussed failure mode: the model can generate a syntactically valid, semantically confident query that nonetheless answers a subtly different question than the one asked, and because the output is delivered with the same fluent tone regardless of correctness, the error is harder for a non-expert user to detect than a traditional dashboard misconfiguration would be.

Retrieval-Augmented and Grounded Analytics

A partial mitigation for the above is retrieval-augmented generation, in which a model's output is constrained or supplemented by passages retrieved from a verified internal knowledge base rather than generated purely from the model's parametric memory. In an analytics context this typically means grounding a generated explanation in the organisation's own historical reports, metric definitions, or documented business rules, which reduces, though does not eliminate, the risk of a fluent but ungrounded explanation being mistaken for a verified one. (Elastic, 2024; IBM, 2024)

Agentic Orchestration

The most recent development is the use of AI agents that do not simply answer a query but plan and execute a sequence of analytical steps autonomously - retrieving data, selecting a model, running it, evaluating the output, and in some deployments, initiating a downstream action such as adjusting inventory levels or flagging a transaction for review. This is a meaningful departure from earlier automation because the sequence of steps is not fixed in advance by a human designer; the agent determines it dynamically based on intermediate results. The governance implications of this shift are substantial and are addressed in Section 7.

4. AI-DRIVEN ANALYTICS FRAMEWORK

It is useful to lay out, in schematic form, how these technologies compose into an operational pipeline, since much of the confusion in vendor literature comes from treating AI analytics as a single undifferentiated capability rather than a layered system. A reasonably general framework can be described in five stages. The first stage, data ingestion and integration, is unglamorous but decisive: a model trained on incomplete or biased source data will propagate that bias forward regardless of how sophisticated the downstream architecture is, a point that is easy to state and consistently underweighted in practice. The second stage, preprocessing and representation, prepares data for modelling, increasingly through learned rather than manually specified transformations, as discussed in Section 3.1. The third stage, inference, is where the traditional analytics literature has concentrated most of its attention - the model itself, whether a gradient-boosted tree or a transformer. The fourth stage, explanation and contextualisation, generates a human-readable account of what the inference means, a task now frequently performed by a generative layer sitting atop the predictive model rather than by a human analyst. The fifth stage, action, is where prescriptive and agentic systems close the loop by executing or recommending a specific intervention. The framework's value lies less in the individual stages, which are not novel, than in making explicit where human oversight has historically been concentrated and where it is now being displaced. In the earlier era of analytics, human judgement sat almost entirely at stages four and five - a human read the output and decided what to do. Contemporary systems increasingly automate stage four and, in agentic deployments, stage five as well, which means the point at which a human can meaningfully intervene has moved earlier in the pipeline, typically to stage one or two: deciding what data the system is allowed to see and what constraints bound its downstream actions. Organisations that continue to place their governance effort at the old checkpoint - reviewing final output - are, in effect, checking a decision that has often already been executed.




FIGURE 2. A five-stage AI-driven analytics pipeline, showing the historical location of human oversight (stages 4-5) and its displacement toward earlier stages (1-2) in agentic deployments.

5. MACHINE LEARNING AND DEEP LEARNING IN ANALYTICS

Machine learning and deep learning remain the computational core beneath most of what is marketed as AI analytics, even where a generative or agentic layer sits on top, and it is worth distinguishing their respective contributions rather than treating them as interchangeable terms. Classical machine learning - logistic regression, decision trees, random forests, gradient boosting machines such as XGBoost and LightGBM - continues to dominate structured, tabular analytics tasks, and for good reason: these methods typically require less data to train reliably, produce more directly interpretable outputs, such as a feature-importance ranking, and are considerably cheaper to deploy and retrain than deep architectures. (Fayaz et al., 2022; Schwartz-Ziv & Armon, 2022) The persistent finding in much of the applied literature that gradient-boosted trees outperform deep neural networks on tabular business data is one of the more consistent, if underappreciated, results in the field, and it should temper the assumption that deeper always means better. Deep learning earns its advantage where the data itself has spatial, sequential, or otherwise high-dimensional structure that hand-engineered features struggle to capture: image and video data via convolutional or vision-transformer architectures, sequential and time-series data via recurrent networks and transformers, and multimodal data combining several of these. In analytics specifically, deep learning has been most transformative in forecasting problems with long-range temporal dependencies, in anomaly detection within high-volume sensor or transaction streams, and in extracting structured signal from unstructured inputs such as free-text customer feedback or medical imaging. A genuine methodological tension runs through this section that is worth naming directly rather than glossing over: the field's evaluation culture rewards accuracy metrics computed on held-out test sets, but a model that performs well on a static benchmark can still degrade in production as the underlying data distribution shifts - a phenomenon usually called concept drift. Analytics deployments are unusually exposed to this problem because the environments they model, such as consumer behaviour, financial markets, and epidemiological trends, are themselves non-stationary, and a chapter on this subject that did not flag drift monitoring as a first-class operational requirement, rather than a secondary concern, would understate the practical difficulty of sustaining model performance over time. (Jourdan, 2024; "Understanding Data Drift and Concept Drift," 2025)

6. GENERATIVE AI AND LARGE LANGUAGE MODELS FOR ANALYTICS

Generative AI's incorporation into analytics workflows is recent enough that the literature describing it is still largely case-study and preprint material rather than settled consensus, and this chapter treats it accordingly - with more explicit hedging than the preceding sections. Three distinct roles for large language models within an analytics pipeline are worth separating, since conflating them leads to overstated claims about capability. The first is interface: the LLM as a natural-language front end translating a question into a query, as discussed in Section 3.2. The second is synthesis: the LLM generating a narrative summary, report draft, or explanation from structured output the model itself did not produce - essentially acting as a communication layer between the numerical result and the human reader. The third, considerably more ambitious role, is the LLM as an analytical reasoner in its own right, asked to identify patterns, propose hypotheses, or evaluate the plausibility of a causal claim directly from data presented in its context window. The third role is where the greatest caution is warranted. Large language models are not statistical inference engines in the formal sense; they generate text that is locally coherent and often plausible, but the underlying mechanism is pattern completion over training data rather than a calculation grounded in the specific dataset provided. A model asked to identify the driver of a sales decline from a table of figures may produce a confident, well-written, and entirely wrong explanation, and because the explanation is fluent, it is more likely to be trusted than an equally wrong output from a traditional statistical test would be - a risk sometimes described in the literature as automation bias compounded by generative fluency. (Kalai et al., 2026) This is not an argument against using generative

models in analytics; it is an argument for using them at the roles they are suited to, interface and synthesis, with explicit verification steps wherever they are used for the third, reasoning, a distinction the marketing literature around AI analysts and autonomous data scientists frequently elides.

TABLE 2. AI Algorithms Commonly Applied Across the Analytics Pipeline

Category	Representative Algorithms	Primary Analytics Use
Tree-based ensembles	Random Forest, XGBoost, LightGBM	Tabular prediction, feature importance, credit scoring
Sequential / temporal	LSTM, GRU, Temporal Fusion Transformer	Demand forecasting, time-series anomaly detection
Convolutional	CNN, ResNet variants	Image-based quality inspection, medical imaging
Transformer-based LLMs	GPT-family, open-weight LLMs (e.g. Llama)	Natural-language querying, report synthesis
Reinforcement learning	Deep Q-Networks, policy-gradient methods	Dynamic pricing, resource allocation agents
Unsupervised / clustering	k-means, DBSCAN, autoencoders	Customer segmentation, outlier and fraud detection

7. EXPLAINABLE AI AND RESPONSIBLE ANALYTICS

Explainability is frequently framed in the literature as a compliance obligation - something a system needs in order to satisfy a regulator or an audit. That framing is not wrong, but it understates the operational case for explainability, which holds even in the complete absence of regulatory pressure. An analytics system whose reasoning cannot be inspected cannot be effectively debugged when it fails, cannot have its biases identified before they cause harm, and cannot earn the sustained trust of the human operators who are meant to supervise it. Explainability, in other words, is a precondition for reliable deployment, not an ethical addendum bolted onto an otherwise complete system. Post-hoc explanation techniques - SHAP and LIME being the most widely adopted - approximate a complex model's local behaviour with a simpler, interpretable surrogate, and these have become close to standard practice in regulated analytics contexts such as credit underwriting. It is worth noting a limitation that is sometimes underemphasized: a post-hoc explanation describes what the surrogate model says is important, which is not guaranteed to be identical to what the original complex model actually relied upon, and there is documented evidence that these explanation methods can themselves be manipulated to present a misleadingly benign account of a model's behavior. (Slack et al., 2020) This does not make the techniques useless, but it does mean explainability output should be treated as informative rather than authoritative. Responsible analytics extends beyond explainability to encompass fairness auditing, data provenance tracking, and, increasingly relevant as agentic systems take autonomous action, clear accountability structures for when an automated decision causes harm. The question of who is responsible when an autonomous analytics agent, acting within its designed parameters, produces a harmful outcome is not adequately settled by existing frameworks, most of which were designed for a human-in-the-loop model of decision-making that agentic AI is specifically designed to move away from.

TABLE 3. Comparison of Machine Learning, Deep Learning, and Generative AI in Analytics

Dimension	Machine Learning	Deep Learning	Generative AI
Typical data need	Moderate (thousands of rows)	Large (often 10^5+ samples)	Pretrained; fine-tuning needs modest data
Interpretability	Comparatively high	Low without post-hoc methods	Low; output is fluent but not verifiable by construction
Compute cost	Low	Moderate to high	High (training); moderate (inference/API use)
Best suited to	Structured, tabular prediction	Unstructured, high-dimensional data	Language interface, synthesis, drafting
Principal risk	Underfitting on complex patterns	Overfitting; opacity	Confident but ungrounded output

8. INDUSTRY CASE STUDIES

The behaviour of these technologies changes meaningfully once situated within a specific operational and regulatory environment. The following five sketches are intended as illustrative context for the preceding argument rather than as exhaustive sector reviews, and each note where a full treatment would require verified, sector-specific citation.

Healthcare

Clinical analytics has been an early and cautious adopter of deep learning, particularly in diagnostic imaging, where convolutional and transformer-based models have demonstrated performance competitive with specialist clinicians on narrowly defined tasks such as detecting particular lesion types in radiographic scans. (Aggarwal et al., 2021) The caution is warranted: healthcare data is unusually heterogeneous across institutions, meaning a model validated at one hospital can degrade substantially when deployed at another with different imaging equipment or patient demographics, and the consequence of a missed or false positive is materially different from a comparable error in, say, retail demand forecasting.

Finance

Financial services apply AI analytics across credit scoring, algorithmic trading, and fraud detection, and is also among the most heavily regulated environments for model deployment, which has made it something of a proving ground for explainability requirements that other sectors are only now adopting. Real-time fraud detection systems in particular illustrate the concept-drift problem discussed in Section 5 acutely: fraud patterns adapt deliberately and quickly in response to detection methods, which means a static model trained on historical fraud data has a comparatively short effective shelf life. (Tax et al., 2021)

Education

Learning analytics - the use of AI to model student engagement, predict at-risk learners, and personalize instructional content - has grown considerably, and is a domain the author has direct research experience with through agentic early-warning systems built on institutional academic records. A recurring practical difficulty in this domain is that the signal available to a model, such as attendance, assignment submission timing, and engagement with a learning management system, is a proxy for the outcome actually of interest, namely learning or risk of academic failure, and the gap between proxy and outcome is not always made explicit in system design or evaluation.

TABLE 4. Illustrative Industrial Applications of AI-Driven Analytics

Sector	Representative Application	Dominant Technique	Principal Constraint
Healthcare	Diagnostic image triage	CNN / vision transformer	Cross-institutional generalisation
Finance	Real-time fraud detection	Ensemble models, anomaly detection	Adversarial concept drift
Education	At-risk student identification	Classification, agentic monitoring	Proxy-outcome validity
Manufacturing	Predictive maintenance	Time-series deep learning	Rare-event labelling
Smart cities	Traffic and grid optimization	Reinforcement learning, sensor fusion	Public accountability, data equity

Manufacturing

Predictive maintenance, using sensor and time-series data to forecast equipment failure before it occurs, is among the more mature and economically well-justified applications of AI analytics in industrial settings, and one where the return on investment is comparatively easy to quantify against the baseline cost of unplanned downtime. (Ismail et al., 2025) The technical challenge here is less about model architecture than about sensor data quality and the cost of labelling rare failure events for supervised training.

Smart Cities and Infrastructure

Urban analytics applications - traffic flow optimisation, energy grid balancing, public safety resource allocation - increasingly rely on real-time, multi-source data fusion, combining IoT sensor streams with historical patterns. This domain surfaces governance questions with particular sharpness, since the populations affected by an urban analytics system's decisions are not

customers who have opted in but residents subject to a public authority's deployment choices, which raises accountability considerations distinct from those in a commercial context.

9. CHALLENGES AND FUTURE RESEARCH DIRECTIONS

Several problems raised across this chapter remain genuinely open rather than merely underdiscussed, and are worth stating plainly as a research agenda rather than folded quietly into the preceding sections. The evaluation of generative reasoning in analytics contexts is the most immediate. Standard machine learning evaluation assumes a well-defined ground truth against which accuracy can be measured; a generated explanation or narrative summary frequently does not have one, and the field currently lacks consensus metrics for judging whether a generated analytical explanation is not merely fluent but actually correct and non-misleading. (Kalai et al., 2026) Second, the computational and environmental cost of applying large foundation models to routine analytical tasks that a far smaller, task-specific model could handle adequately is a sustainability concern that receives comparatively little attention relative to its practical significance, particularly as organisations default to large general-purpose models for tasks that do not require that scale. Third, accountability frameworks for agentic analytics systems - where responsibility lies when an autonomous agent's action, taken within its designed authority, produces harm - remain underdeveloped relative to the pace at which such systems are being deployed, and existing regulatory instruments were largely drafted with a human-in-the-loop assumption that agentic architectures are specifically designed to relax.

TABLE 5. Key Challenges in AI-Driven Analytics and Associated Mitigation Directions

Challenge	Underlying Cause	Potential Mitigation Direction
Evaluating generative reasoning	No fixed ground truth for open-ended output	Consensus benchmarks combining factual grounding and coherence checks
Concept drift	Non-stationary real-world data distributions	Continuous monitoring, scheduled retraining triggers
Explanation reliability	Post-hoc methods approximate, do not verify, model logic	Combine multiple explanation methods; treat as diagnostic, not proof
Agentic accountability	Regulatory frameworks assume human-in-the-loop	Action-scoped permissions, mandatory audit trails for autonomous actions
Data and compute sustainability	Default use of large foundation models for routine tasks	Task-appropriate model sizing; smaller fine-tuned models where sufficient
Data equity	Training data skewed toward well-resourced contexts	Distributional fragility testing beyond aggregate accuracy metrics

A fourth and less frequently discussed issue concerns data equity: AI-driven analytics systems trained predominantly on data from well-resourced institutions or regions risk performing systematically worse for underrepresented populations, a concern that compounds rather than resolves as models are deployed more widely across heterogeneous global contexts. Future research would benefit from evaluation benchmarks that explicitly test for this kind of distributional fragility rather than reporting only aggregate accuracy.

10. CONCLUSION

The central argument of this chapter has been that AI's contribution to data analytics is not accurately described as simply faster or more accurate analytics, though it is frequently both. The more consequential change is structural: interpretive and, increasingly, decisional authority that once sat with a human analyst is migrating into the system itself, and the point at which meaningful human oversight can be exercised is moving earlier in the pipeline - from reviewing a final recommendation to constraining what data and actions a system is permitted to access in the first place. Organisations and researchers who continue to locate their governance effort at the old checkpoint are likely to find it arrives too late to matter. None of this argues against the adoption of AI-driven analytics, which offers genuine and substantial capability gains across every sector examined in this chapter. It argues instead for a more precise vocabulary - distinguishing where a model is serving as interface, as synthesizer, or as autonomous reasoner - and for evaluation and governance practices that keep pace with a field whose technical capability has, in recent years, moved considerably faster than its accompanying frameworks for accountability. That gap, more than any specific algorithmic limitation, is the most pressing unresolved problem this chapter has identified.

REFERENCES

- [1]. Aggarwal, R., Sounderajah, V., Martin, G., Ting, D. S. W., Karthikesalingam, A., King, D., Ashrafian, H., & Darzi, A. (2021). Diagnostic accuracy of deep learning in medical imaging: A systematic review and meta-analysis. *npj Digital Medicine*, 4, 65. <https://doi.org/10.1038/s41746-021-00438-z>
- [2]. Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828. <https://doi.org/10.1109/TPAMI.2013.50>
- [3]. Chowa, S. S., Alvi, R., Rahman, S. S., Rahman, M. A., Raiaan, M. A. K., Islam, M. R., Hussain, M., & Azam, S. (2026). From language to action: A review of large language models as autonomous agents and tool users. *Artificial Intelligence Review*, 59, 71. <https://doi.org/10.1007/s10462-025-11471-9>
- [4]. Elastic. (2024, June 17). RAG platform solutions for enterprise search. <https://www.elastic.co/elasticsearch/rag>
- [5]. Fayaz, S. A., Zaman, M., Kaul, S., & Butt, M. A. (2022). Is deep learning on tabular data enough? An assessment. *International Journal of Advanced Computer Science and Applications*, 13(4). <https://doi.org/10.14569/IJACSA.2022.0130454>
- [6]. IBM. (2024, October 20). What is RAG (retrieval augmented generation)? <https://www.ibm.com/think/topics/retrieval-augmented-generation>
- [7]. Ismail, L., Abdelmoti, A., Basu, A., Berini, A. D. E., & Naouss, M. (2025). A systematic review of digital twin-driven predictive maintenance in industrial engineering: Taxonomy, architectural elements, and future research directions. *arXiv preprint*. <https://arxiv.org/abs/2509.24443>
- [8]. Jourdan, N. (2024). Addressing concept drift in machine learning-based monitoring of manufacturing processes (Doctoral thesis). Technical University of Darmstadt. <https://tuprints.ulb.tu-darmstadt.de/28043/>
- [9]. Kalai, A. T., Nachum, O., Vempala, S. S., & Zhang, E. (2026). Evaluating large language models for accuracy incentivizes hallucinations. *Nature*, 653, 1047–1051. <https://doi.org/10.1038/s41586-026-10549-w>
- [10]. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- [11]. Shwartz-Ziv, R., & Armon, A. (2022). Tabular data: Deep learning is not all you need. <https://www.ravid-shwartz-ziv.com/publication/tabular/>
- [12]. Slack, D., Hilgard, S., Jia, E., Singh, S., & Lakkaraju, H. (2020). Fooling LIME and SHAP: Adversarial attacks on post hoc explanation methods. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 180–186. <https://doi.org/10.1145/3375627.3375830>
- [13]. Strategy. (2025, July 1). The state of AI+BI analytics global 2025 report. <https://www.strategy.com/it/software/survey>
- [14]. Tax, N., de Vries, K. J., de Jong, M., Dosoula, N., van den Akker, B., Smith, J., Thuong, O., & Bernardi, L. (2021). Machine learning for fraud detection in e-commerce: A research agenda. *arXiv preprint*. <https://arxiv.org/abs/2107.01979>
- [15]. Understanding data drift and concept drift in machine learning systems. (2025). *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*. <https://ijsrseit.com/index.php/home/article/view/CSEIT25111239>