IJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

STRATEGIC MODEL SELECTION IN APPLIED ARTIFICIAL INTELLIGENCE: ALIGNING METHODS WITH PROBLEM CONTEXTS

Sunish Vengathattil [1]
https://orcid.org/0009-0003-1222-0852
Sr. Director, Software Engineering, Academia & Government, Clarivate Analytics, Philadelphia, PA
United States

Resmi Vijayan[2]
https://orcid.org/0009-0009-6120-6479
Software Engineer, Comcast, Philadelphia, PA
United States

ABSTRACT
Choosing the right machine learning model for a given problem is a critical decision that significantly impacts the success of a project. This article provides a comprehensive guide to selecting the most appropriate model based on the problem, the available data, and operational constraints. It explores key considerations such as interpretability, data quality, computational resources, and real-time requirements while emphasizing the importance of aligning the model with business goals and stakeholder needs. Through case studies in areas like fraud detection, medical image classification, and customer segmentation, the article illustrates how different models are suited to specific types of problems. It also highlights the importance of deploying, monitoring, and maintaining models to ensure their long-term effectiveness. By presenting a structured approach to model selection and ongoing improvement, this article aims to provide valuable insights for practitioners seeking to optimize the use of machine learning in real-world applications.

Keywords:
Artificial Intelligence, Machine Learning, Model Selection, AI Models

1. INTRODUCTION

Modern enterprises use artificial intelligence and machine learning in all strategic and operational aspects, such as analyzing customer journeys and interaction patterns, continuously refining their models to ensure accuracy and maintain data integrity. These technologies help organizations gain deeper insights into customer behavior and improve decision-making through adaptive, data-driven strategies (Shaffi, 2020).

One of the most frequent missteps in applying machine learning is defaulting to the most advanced or widely discussed models, often driven by trends rather than needs. Paleyes et al. (2022) clarify that, in practice, effective model selection is not about following technological buzzwords but about aligning analytical tools with the specific nature of the problem, the characteristics of the available data, and the intended outcomes. Much like selecting the right tool for a specific task, choosing a machine learning model requires careful consideration; a complex solution applied to a simple problem can be as ineffective as using a blunt instrument for delicate work

According to Hassija et al. (2023), some challenges are best addressed with interpretable, lightweight models that deliver fast and transparent results. Others may call for more sophisticated techniques, particularly those involving large, high-dimensional datasets or nuanced patterns. The central principle remains the same: clarity of purpose should guide the choice of method.

This paper seeks to provide a practical framework for navigating that decision-making process. It outlines how to assess the nature of a problem, match it with suitable model families, and evaluate real-world factors such as data quality, speed of inference, interpretability, and maintainability. In doing so, it aims to demystify model selection and support thoughtful, effective applications of machine learning.

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

2. UNDERSTANDING THE PROBLEM SPACE

Effective model selection begins with a clear understanding of the problem one aims to solve. Just as no single tool fits every task, no single machine learning model is universally optimal. A rigorous analysis of the problem space - comprising the objective, the nature of the data, and the operational context is essential to align technical solutions with real-world needs. This section outlines key considerations in framing the problem correctly before engaging in model evaluation and selection Schonle et al., 2024).

2.1 Define the business objective

The first step is to articulate the problem's core objective. Is the goal to classify data points into predefined categories, predict a continuous numerical value, identify latent structures or groupings, or generate new content? Clearly defining whether the task is one of classification, regression, clustering, recommendation, or generation allows for a more structured and efficient model selection process. Ambiguity at this stage often leads to misaligned solutions and inefficiencies downstream.

2.2 Understand data constraint

The characteristics of the available data significantly influence model choice. Wiegand et al. (2024) describe advanced methods of analyzing data constraints, including how constraints can be derived from data. Key factors include:

Volume: The size of the dataset can determine whether simpler models suffice or if complex, data-hungry approaches are viable.

Structure: Whether the data is structured (e.g., tabular), semi-structured, or unstructured (e.g., text, images).

Labels: The presence or absence of labeled outcomes will inform whether supervised, unsupervised, or semi-supervised learning is appropriate.

Quality: The completeness, consistency, and cleanliness of the dataset play a pivotal role in determining the feasibility and performance of the chosen approach.

Understanding these dimensions helps in setting realistic expectations and constraints for model performance.

2.3 Consider the Operational and Business Context

As the operational and business landscapes surrounding artificial intelligence grow more complex, researchers are placing greater emphasis on understanding how to effectively integrate AI into organizational processes. The focus is on identifying strategies that not only streamline operations but also drive innovation and create measurable business value. By doing so, they aim to help organizations gain a sustainable competitive edge, adapt to technological advancements, and stay ahead in an increasingly data-driven and AI-influenced market. Model selection should be guided not only by technical fit but also by the broader context in which the solution will be deployed (Rana et al., 2021). Considerations include:

Latency Requirements: Does the solution require real-time inference, or can it operate in batch mode?

Interpretability Needs: Are stakeholders expecting transparent, explainable outputs?

Update Frequency: Will the model require regular retraining to adapt to evolving data (e.g., due to concept drift)?

End-User Environment: Is the model embedded in a cloud platform, edge device, or legacy system?

Aligning the model with these practical constraints ensures greater long-term viability and stakeholder adoption.

3. MODEL FAMILIES

Once the problem has been clearly defined and the data context is well understood, the next step involves selecting a suitable family of models. Tufail et al. (2023) dive deeper into different model families and the associated problems they solve. Each class of machine learning algorithms is built upon distinct assumptions and mathematical foundations, making them more effective for specific tasks than others. This section presents a high-level overview of commonly used model families, along with their typical strengths and ideal application contexts.

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

3.1 Rule-Based Systems

Rule-based systems operate on explicitly defined logic, typically using "if–then" statements. Though not traditionally considered machine learning, they can be effective in domains with well-understood processes or clearly defined thresholds (Sarker et al., 2024).

Strengths:

- High interpretability and transparency
- Fast inference and low computational cost
- Easy to debug and audit

Limitations:

- Poor adaptability to changing patterns
- Not suited for complex or high-dimensional data

Use Cases: Eligibility screening, threshold-based alerts, deterministic workflows.

3.2 Classical Machine Learning Models

Classical machine learning models include algorithms such as logistic regression, decision trees, support vector machines (SVMs), k-nearest neighbors (k-NN), and naive Bayes. These models are well-suited for structured, tabular data with limited feature sets and are often chosen for their effectiveness in handling tasks like regression, classification, clustering, and dimensionality reduction. Their adaptability and efficiency make them valuable tools in various data-driven analytical and predictive modeling applications (Tufail et al., 2023, Mukhamediev et al., 2021).

Strengths:

- Require less data and computational power than deep learning
- Often easier to interpret, especially linear models and decision trees
- Proven performance in many business-critical use cases

Limitations:

- Limited performance with unstructured data (e.g., images or raw text)
- May plateau in accuracy on more complex tasks

Use Cases: Credit scoring, churn prediction, basic classification and regression tasks.

3.3 Ensemble Methods

Ensemble learning combines multiple models to improve overall performance. Examples include random forests, gradient boosting machines (e.g., XGBoost, LightGBM), and bagging methods (Mukhamediev et al., 2021).

Strengths:

- High predictive accuracy
- Robust to overfitting in many cases
- Often effective out-of-the-box with minimal feature engineering

Limitations:

- Lower interpretability compared to simpler models
- Longer training times, especially with large datasets

Use Cases: Competitive modeling tasks, structured data competitions (e.g., Kaggle), high-stakes decision systems.

3.4 Deep Learning Models

Deep learning encompasses neural networks and their architectures, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformers. These are particularly well-suited for complex, high-dimensional, and unstructured data types (Tufail et al., 2023, Mukhamediev et al., 2021).

Strengths:

- Exceptional performance in vision, language, and audio tasks
- Capable of automatic feature extraction from raw data
- Scalable for large datasets and distributed computing

Limitations:

- Requires substantial computational resources and large labeled datasets

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

- Often considered “black boxes” due to limited explainability
- More complex to maintain and tune

Use Cases: Image classification, natural language processing, speech recognition, recommendation engines.

3.5 Probabilistic and Bayesian Models

These models explicitly account for uncertainty in predictions, including Bayesian networks and hidden Markov models. They are often used in domains where probabilistic reasoning or temporal dynamics are critical (Ni et al., 2021).

Strengths:

- Handle uncertainty in a principled way
- Offer interpretable probabilistic outputs
- Useful in low-data environments when prior knowledge is available

Limitations:

- Computationally intensive in high-dimensional settings
- May be less performant in terms of raw predictive accuracy

Use Cases: Risk modeling, time-series analysis, causal inference, decision support under uncertainty.

By understanding the strengths and trade-offs of each model family, practitioners can better align model capabilities with the specific demands of their application. The next section will explore how to make these decisions in context, balancing technical feasibility with practical considerations.

4. KEY CONSIDERATIONS WHEN CHOOSING A MODEL

Selecting an appropriate machine learning model involves more than matching algorithms to data types. It also requires balancing a range of practical considerations that influence the model’s effectiveness, usability, and sustainability in production settings. This section outlines the most critical factors to weigh during the model selection process, particularly when working under real-world constraints (Schonle et al., 2024).

4.1 Interpretability

In many domains, particularly healthcare, finance, and scientific publishing, stakeholders demand that model decisions be transparent and explainable. Linear models, decision trees, and rule-based systems offer high levels of interpretability, allowing users to trace predictions back to specific features or rules. In contrast, more complex models, such as deep neural networks, may achieve higher accuracy but at the cost of reduced transparency (Hassija et al., 2023).

When it matters: Regulatory compliance, risk-sensitive decisions, end-user trust, and model validation processes.

4.2 Data Availability and Quality

Model performance is fundamentally constrained by the volume, variety, and cleanliness of available data. Classical models generally require less data and are more tolerant of noisy or incomplete datasets. In contrast, deep learning models demand large, well-labeled datasets and are highly sensitive to inconsistencies (Lones, 2021).

When it matters: Early-stage projects, domains with limited data access, or datasets with missing or unstructured elements.

4.3 Computational Resources

AI-driven automation and generative models enhance efficiency by streamlining operations, reducing errors, and optimizing resource use. This leads to lower operational costs and improved productivity, allowing organizations to allocate resources more effectively and maintain smoother workflows with minimal manual intervention (Balamurugan & Ramamoorthy, 2025). However, these models will come at their own costs. Resource requirements vary widely across model families. Simpler models can run efficiently on standard hardware and are suitable for edge or embedded deployments. More sophisticated approaches, particularly ensemble and deep learning models, require powerful GPUs, high memory capacity, and potentially distributed computing infrastructure (Khan et al., 2022).

When it matters: Budget limitations, deployment in low-power environments, or strict latency constraints.

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

4.4 Speed and Latency

Some applications demand real-time or near-real-time responses (e.g., fraud detection, industrial automation, personalized recommendations). In such cases, model inference speed is a critical factor. Simpler models typically offer faster predictions and can be more easily optimized for latency-sensitive use cases (Fantinuoli & Montecchio, 2022).
When it matters: Time-critical systems, interactive applications, and embedded AI solutions.

4.5 Maintainability and Lifespan

Models do not operate in static environments. Data distributions may drift, business objectives may evolve, and integration needs may shift. Models that are easier to retrain, monitor, and debug tend to offer greater long-term value. Overly complex systems may be difficult to update or troubleshoot, increasing technical debt over time (Schonle et al., 2024).

When it matters: Long-term deployments, rapidly changing domains, and teams with limited machine learning support.

4.6 Business Alignment and Stakeholder Buy-In

The best technical solution is not always the right business solution. A model must serve the organization's larger strategic goals and be understandable and acceptable to its stakeholders. Involving domain experts early in the modeling process and communicating results in accessible language can be crucial for adoption and trust (Rana et al., 2021).
When it matters: Cross-functional teams, executive-level reporting, or projects involving non-technical decision-makers.

By systematically considering these dimensions, interpretability, data constraints, computational demands, speed, maintainability, and stakeholder alignment, practitioners can make informed trade-offs that go beyond technical performance metrics alone. These trade-offs often define whether a machine learning solution succeeds in practice, not just in experimentation.

5. CASE STUDIES: MATCHING PROBLEMS TO MODELS

Understanding the theoretical aspects of model selection is crucial, but practical application often requires navigating the nuances of real-world problems (Anand et al., 2025). This section presents a series of hypothetical case studies that illustrate how various machine learning models can be matched to different types of problems. Each case study emphasizes the importance of aligning problems, data, and operational constraints with the most suitable modeling approach. Please note that these are hypothetical examples, and practical cases may be more complex and may take an alternate decision process.

5.1 Fraud Detection in Financial Transactions

Problem: A bank wants to detect fraudulent transactions in real-time based on user activity and transaction history.
Model Selection: Fraud detection is a classification problem, where the goal is to identify fraudulent transactions (positive class) from legitimate ones (negative class). Due to the real-time nature of the task, speed and accuracy are critical. Classical machine learning models such as logistic regression, decision trees, and ensemble methods like random forests or gradient boosting machines (e.g., XGBoost) are well-suited for this problem. These models offer a good balance of interpretability and performance, essential in high-stakes applications.
Considerations:

- Data Quality: The bank may have access to a rich dataset, including user behavior, transaction details, and historical fraud cases. However, the dataset may also be imbalanced, with significantly more legitimate transactions than fraudulent ones.

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

- Performance Metrics: Precision and recall are of particular importance in this case, as false negatives (failing to flag fraud) may be more costly than false positives (incorrectly flagging legitimate transactions).
- Interpretability: Being able to explain the reasons behind a fraud detection decision can build trust with users and regulatory authorities.

Outcome: An ensemble method like XGBoost provides the necessary predictive power while maintaining a level of interpretability for auditing purposes.

5.2 Image Classification for Medical Diagnosis

Problem: A healthcare provider wants to automatically classify medical images (e.g., X-rays or MRIs) to diagnose diseases such as pneumonia.
Model Selection: This is an image classification problem, where deep learning models, particularly convolutional neural networks (CNNs), excel. CNNs are designed to automatically extract relevant features from image data, making them highly effective in recognizing patterns in pixel-based data.
Considerations:

- Data Availability: High-quality labeled image datasets are essential. The model needs to be trained on a substantial number of labeled examples to achieve high accuracy. Data augmentation techniques may be employed to artificially increase the dataset size.
- Performance Metrics: Accuracy is critical, but so is interpretability. In medical applications, understanding why a model classifies an image in a particular way can help practitioners trust and validate the results.
- Computational Requirements: Training CNNs requires substantial computational resources, including GPUs for both training and inference.

Outcome: A well-tuned CNN model achieves superior accuracy in classifying medical images, but interpretability techniques like Grad-CAM are employed to highlight relevant image regions, ensuring transparency for medical professionals.

5.3 Customer Segmentation for Marketing

Problem: A retail company wants to segment its customers based on purchasing behavior to target marketing efforts more effectively.
Model Selection: Customer segmentation is typically approached as a clustering problem. Unsupervised learning techniques, such as k-means clustering or hierarchical clustering, are well-suited to this task. These models group customers based on similarities in their purchasing patterns, without requiring labeled data.
Considerations:

- Data Structure: The dataset may include transactional history, demographic data, and customer interactions with the company's website. Features such as frequency, recency, and monetary value (RFM) could be used to generate meaningful clusters.
- Performance Metrics: Since clustering is unsupervised, the evaluation metric is often qualitative. Business experts must assess whether the resulting clusters align with actionable marketing strategies.
- Interpretability: Understanding the characteristics of each cluster is important for marketing teams to craft targeted strategies. K-means clustering provides an intuitive approach, but other algorithms like DBSCAN or Gaussian Mixture Models (GMMs) could offer more flexibility.

Outcome: K-means clustering creates distinct customer segments based on purchasing behavior, which are used to design targeted promotions and product recommendations. The segmentation process is continuously refined based on evolving customer data.

5.4 Predicting Equipment Failure in Manufacturing

Problem: A manufacturing plant wants to predict when critical machinery is likely to fail in order to schedule preventative maintenance.
Model Selection: This is a time-series forecasting problem. Recurrent neural networks (RNNs), particularly long short-term memory (LSTM) networks, are well-suited for tasks that involve sequential data with temporal dependencies, such as sensor data from machines.
Considerations:

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

- Data Quality: Sensor data may be noisy or incomplete, requiring robust preprocessing and possibly filling in missing values. Additionally, labeled failure events may be sparse, making it important to balance predictive performance and class distribution.
- Performance Metrics: The model's ability to predict failure with sufficient lead time is critical. Metrics such as mean absolute error (MAE) or root mean square error (RMSE) are commonly used in regression tasks.
- Real-time Operation: The model must be integrated into a production environment where it can predict failures in real time and trigger maintenance alerts.

Outcome: An LSTM-based model provides reliable predictions of equipment failure, enabling the plant to reduce unplanned downtime and improve maintenance scheduling.

These case studies demonstrate how understanding the problem, data characteristics, and practical constraints informs the choice of model. By mapping the problem space to the right model family, organizations can maximize the effectiveness of machine learning solutions and avoid common pitfalls. The next section will delve into the final steps in the model selection process, including implementation, monitoring, and continuous improvement.

6. IMPLEMENTING, MONITORING, AND IMPROVING THE MODEL

Once a model has been selected and trained, the next step involves integrating it into production environments and ensuring it continues to perform effectively over time. This section discusses the key steps in implementing a model, monitoring its performance in real-world conditions, and ensuring its continuous improvement. These practices are essential for maintaining the long-term viability of machine learning solutions and maximizing their value (Feng et al., 2022).

6.1 Model Deployment and Integration

After selecting and fine-tuning the model, it is deployed into production. This phase involves integrating the model with the organization's existing infrastructure and ensuring it can interact seamlessly with other systems, databases, and applications (Sjödin et al., 2021). There are several deployment strategies to consider, including:
Batch vs. Real-time: Deciding whether the model will make predictions in real-time or in batches (e.g., nightly updates).
Edge Deployment: In some cases, models may need to be deployed directly on edge devices (e.g., sensors or mobile apps), where computational resources are limited.
Scalability: The deployed system should be capable of handling increasing loads as data volumes grow over time, often necessitating cloud-based or distributed architectures.

Effective deployment also includes establishing an infrastructure for version control and model retraining, ensuring that updates to the model can be rolled out smoothly without disrupting operations.

6.2 Monitoring Model Performance

Monitoring a model after deployment is crucial to ensure that it continues to deliver accurate and reliable results (Sjödin et al., 2021). Model performance can degrade over time due to various factors, such as:
Data Drift: Variations in data distribution across time can lead to reduced accuracy in the model's predictions. Regular monitoring helps detect such changes early.
Model Drift: Even if data distribution remains stable, the model's performance can degrade due to issues such as outdated features or evolving business objectives.
External Factors: Environmental changes, such as new regulations or market conditions, can also impact model effectiveness.
Key performance indicators (KPIs) to track include accuracy, precision, recall, and more specific domain-related metrics (e.g., false positives in fraud detection or AUC in medical diagnoses). Regular validation against a holdout dataset or through A/B testing can help ensure the model is performing as expected.

6.3 Model Maintenance and Retraining

Machine learning models should not be viewed as static; they require continuous maintenance and periodic retraining to adapt to new data and evolving requirements (Sjödin et al., 2021). Key activities in maintaining a model include:

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

Regular Retraining: The model should be retrained periodically with new data to account for changes in patterns or underlying distributions. This can be done on a scheduled basis (e.g., monthly or quarterly) or triggered by performance degradation.
Incremental Learning: In some applications, it is possible to implement incremental learning techniques, where the model learns continuously as new data arrives without needing complete retraining from scratch.
Feedback Loops: Feedback from users and stakeholders can help inform model updates. For example, user corrections to predictions in a recommender system can be incorporated into the next iteration of the model.
Effective retraining also involves careful testing and validation before rolling out updated models to ensure that new versions outperform previous iterations.

6.4 Ethical Considerations and Bias Mitigation

As machine learning models become integrated into decision-making processes, it is essential to consider their ethical implications. Models can unintentionally propagate biases present in training data, leading to unfair or discriminatory outcomes. Ethical considerations should guide model selection, particularly in sensitive domains like hiring, lending, and criminal justice (Sreerama & Krishnamoorthy, 2022).

To mitigate bias:
Bias Audits: Regular audits should be conducted to ensure the model's predictions are not disproportionately harmful to specific groups.
Fairness Constraints: Incorporating fairness constraints during model development or optimization can help reduce unintended bias.
Transparency: Providing clear explanations of model decisions (e.g., through explainable AI techniques) can help ensure that decisions are made in a just and transparent manner.

6.5 Continuous Improvement and Iteration

Machine learning is an iterative process, and even the best models can always be improved. Continuous feedback, coupled with advancements in model architecture, algorithms, and data sources, presents ongoing opportunities for improvement. By fostering a culture of constant refinement and adaptation, organizations can ensure their machine learning systems stay relevant and performant in the face of new challenges and opportunities (Feng et al., 2022).

Key practices for continuous improvement include:

- Regularly reviewing model performance with stakeholders to ensure it aligns with evolving business objectives.
- Staying up to date with new developments in machine learning research and incorporating cutting-edge methods when appropriate.
- Engaging in cross-disciplinary collaboration to incorporate diverse insights into model updates.

7. CONCLUSION

Selecting, deploying, and maintaining a machine learning (ML) model is a nuanced process that requires alignment between technical capabilities and practical demands. It begins with a thorough understanding of the problem space—clarifying business objectives, data constraints, and operational requirements. Based on these factors, practitioners must identify an appropriate class of models, whether rule-based, classical ML, ensemble methods, or deep learning architectures, each with unique advantages and limitations.

Effective model selection goes beyond choosing the most sophisticated algorithm. Instead, it emphasizes suitability—considering interpretability, computational efficiency, scalability, and ease of integration within existing systems. Real-world constraints such as latency, retraining needs, infrastructure compatibility, and end-user interaction must inform this decision.

Deployment marks only the midpoint of the journey. Ensuring long-term success requires continuous monitoring, evaluation, and retraining to adapt to evolving data and maintain relevance. Tools for performance tracking, version control, and feedback loops become essential components of the model's lifecycle.

Ultimately, the success of applied machine learning depends not on novelty, but on contextual fit. By selecting models that align with both technical and business objectives, organizations can ensure sustainable value creation and actionable insights from their AI initiatives.

iJETRM
International Journal of Engineering Technology Research & Management
Published By:
https://www.ijetrm.com/

REFERENCES

[1] Anand, A. S., Sawant, S., Reinhardt, D., & Gros, S. (2025). All AI Models are Wrong, but Some are Optimal. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2501.06086
[2] Balamurugan, M., & Ramamoorthy, L. (2025). Transformative Intelligence: AI and generative models as catalysts for creative Problem-Solving in complex environments. International Journal for Multidisciplinary Research, 7(2). https://doi.org/10.36948/ijfmr.2025.v07i02.38404
[3] Fantinuoli, C., & Montecchio, M. (2022). Defining maximum acceptable latency of AI-enhanced CAI tools. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2201.02792
[4] Feng, J., Phillips, R. V., Malenica, I., Bishara, A., Hubbard, A. E., Celi, L. A., & Pirracchio, R. (2022). Clinical artificial intelligence quality improvement: towards continual monitoring and updating of AI algorithms in healthwsadcare. Npj Digital Medicine, 5(1). https://doi.org/10.1038/s41746-022-00611-y
[5] Hassija, V., Chamola, V., Mahapatra, A., Singal, A., Goel, D., Huang, K., Scardapane, S., Spinelli, I., Mahmud, M., & Hussain, A. (2023). Interpreting Black-Box Models: A review on Explainable Artificial intelligence. Cognitive Computation, 16(1), 45–74. https://doi.org/10.1007/s12559-023-10179-8
[6] Khan, T., Tian, W., Zhou, G., Ilager, S., Gong, M., & Buyya, R. (2022). Machine learning (ML)-centric resource management in cloud computing: A review and future directions. Journal of Network and Computer Applications, 204, 103405. https://doi.org/10.1016/j.jnca.2022.103405
[7] Lones, M. A. (2021). How to avoid machine learning pitfalls: a guide for academic researchers. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2108.02497
[8] Mukhamediev, R. I., Symagulov, A., Kuchin, Y., Yakunin, K., & Yelis, M. (2021). From classical machine learning to deep neural networks: A simplified scientometric review. Applied Sciences, 11(12), 5541. https://doi.org/10.3390/app11125541
[9] Ni, P., Li, J., Hao, H., Han, Q., & Du, X. (2021). Probabilistic model updating via variational Bayesian inference and adaptive Gaussian process modeling. Computer Methods in Applied Mechanics and Engineering, 383, 113915. https://doi.org/10.1016/j.cma.2021.113915
[10] Paleyes, A., Urma, R., & Lawrence, N. D. (2022). Challenges in Deploying Machine Learning: A survey of case studies. ACM Computing Surveys, 55(6), 1–29. https://doi.org/10.1145/3533378
[11] Rana, N. P., Chatterjee, S., Dwivedi, Y. K., & Akter, S. (2021). Understanding dark side of artificial intelligence (AI) integrated business analytics: assessing firm's operational inefficiency and competitiveness. European Journal of Information Systems, 31(3), 364–387. https://doi.org/10.1080/0960085x.2021.1955628
[12] Sarker, I. H., Janicke, H., Ferrag, M. A., & Abuadbba, A. (2024). Multi-aspect rule-based AI: Methods, taxonomy, challenges and directions towards automation, intelligence and transparent cybersecurity modeling for critical infrastructures. Internet of Things, 25, 101110. https://doi.org/10.1016/j.iot.2024.101110
[13] Schonle, D., Reich, C., & Ould-Abdeslam, D. (2024). Streamlining AI: Techniques for Efficient Machine Learning Model Selection. The International Journal on Advances in Intelligent Systems, 17(12), 9.
[14] Shaffi, S. M. (2020). Comprehensive Digital Forensics and Risk Mitigation Strategy for modern enterprises. International Journal of Science and Research (IJSR), 9(12), 1885–1892. https://doi.org/10.21275/sr201211165829
[15] Sjödin, D., Parida, V., Palmié, M., & Wincent, J. (2021). How AI capabilities enable business model innovation: Scaling AI through co-evolutionary processes and feedback loops. Journal of Business Research, 134, 574–587. https://doi.org/10.1016/j.jbusres.2021.05.009
[16] Sreerama, J., & Krishnamoorthy, G. (2022). Ethical Considerations in AI Addressing bias and fairness in machine learning models. Journal of Knowledge Learning and Science Technology ISSN 2959-6386 (Online), 1(1), 130–138. https://doi.org/10.60087/jklst.vol1.n1.p138
[17] Tufail, S., Riggs, H., Tariq, M., & Sarwat, A. I. (2023). Advancements and Challenges in Machine Learning: A comprehensive review of models, libraries, applications, and algorithms. Electronics, 12(8), 1789. https://doi.org/10.3390/electronics12081789
[18] Wiegand, B., Klakow, D., & Vreeken, J. (2024). What Are the Rules? Discovering Constraints from Data. Proceedings of the AAAI Conference on Artificial Intelligence, 38(8), 8182–8190. https://doi.org/10.1609/aaai.v38i8.28658