Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-144-2_7

Chapter 7: Machine Intelligence for Risk: Deep Learning Models in Credit, Fraud, and Market Prediction

7.1. Introduction

Risk assessment has benefited enormously from the application of deep learning. Machine intelligence, natural-language processing, neural approximation, representation learning, and generative techniques are now routinely leveraged in the credit, fraud, and markets domains. Driven by academic curiosity, access to rich digital datasets, and some of the most powerful computing resources in history, an increasing number of researchers and practitioners are exploring deep learning for risk.

Yet many questions remain. Digital credit scores are used to evaluate the future creditworthiness of individuals, enabling faster and cheaper borrowing decisions. How effectively do such models discriminate bad and good credit risk? Is there inherent bias against protected groups? Would regulatory authorities, which ultimately evaluate the fairness of these models, prefer a complex black-box model or one built on simple, economic rules? Fraud detection systems seek to identify dubious transactions before they are completed. How do these systems balance the need to detect fraud (true positive) with the cost attached to falsely rejecting clean transactions (false positive)? In a sector where the training data is highly imbalanced, how can the risk of missing a fraudulent transaction be minimized? Rapid market prediction models are able to forecast serum glucose levels based on a limited set of prior measurements. Do these predictions hold under stress testing, when the underlying data distribution changes?

7.1.1. Overview of Deep Learning Applications in Risk Assessment

Novel forms of machine intelligence—teamwork between human and artificial agents—test and extend human intelligence for novel challenges. Recent breakthroughs in deep learning employ conceptual architectures, large reserves of digital metadata, and massive processing power to tackle ever more abstract tasks, especially in vision and language. These developments feature performance beyond previous milestones, match human ability for some specific tasks, and enable many new tasks. Applications to risk—quantification and anticipation of losses for financial or insurance contracts—are particularly significant because associated errors are non-linear and damaging, and even small relative misestimates can have catastrophic impact in critical markets.

Credit, fraud, and market prediction focus the attention. Credit scoring, selection, and pricing constitute a social and economic safeguard for releasing the natural scarcity of money as a means of economic exchange. Fraud detection serves as a vital pillar for the integrity and efficiency of digital commerce. Marked prediction attempts to anticipate future price changes, to make either direct investment decisions or decisions about associated derivative contracts that reap their profit from implicitly betting on these expectations. Each of these areas has sufficient breadth, depth, and importance to warrant separate treatment, yet they all intersect in the distribution and non-linearity of error, and in how that directs attention for successful application of contemporary deep learning.

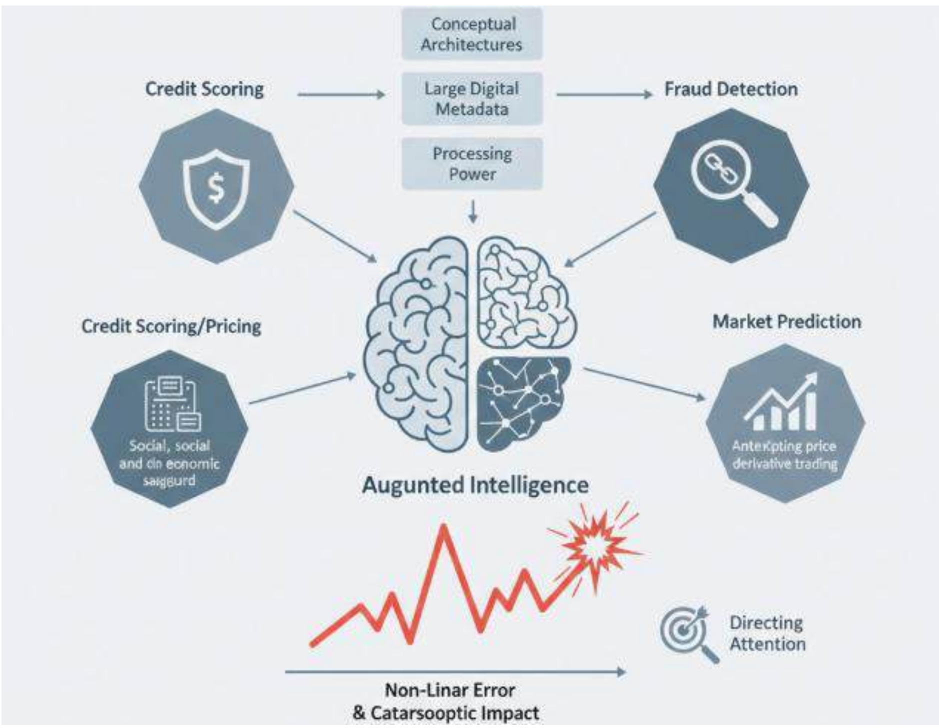

Fig 7.1: Augmented Risk Intelligence: Deep Learning Architectures for Non-Linear Error Management in Critical Financial Markets

7.2. Fundamentals of Deep Learning for Risk Assessment

Fundamental concepts relevant for the application of deep learning to risk assessment and associated predictions are summarized. Deep learning methods are assessed along pertinent dimensions and their specific application in credit scoring, fraud detection, and market prediction is examined.

Considering risk-oriented prediction problems such as credit scoring and fraud detection, various well-established architectural paradigms are available. The choice of architecture reflects the nature of the predictive task, while the representation of features in the data also plays an important role. Data of different types may also benefit from specific pre-processing or embedding steps to facilitate target-oriented learning, enabling both supervised and unsupervised deep representation learning. The mathematical foundations behind representation learning also highlight a third capability beyond discriminative and importance-based modeling: the extraction of features that support down-stream learning tasks such as transfer learning.

7.2.1. Architectural Paradigms and Feature Representations

Four architectural paradigms span deep learning: convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and Transformer networks. Feature representations include manually engineered attributes, fixed-feedback embeddings, and feature extraction via embedding/presentation learning. Pre-training or data-based feedback mechanisms can also enhance representation quality. Specific domains may impose additional model requirements.

Transposed CNN architectures have achieved state-of-the-art image synthesis results, demonstrating applications beyond detection and classification tasks. Applications across computer vision, natural language processing, and audio classification often use a Transformer architecture inspired by attention models but can potentially incorporate any of the aforementioned paradigms. Model selections depend on the problem, chosen supervision type, available data, and underlying data structures. The absence of an explicit temporal dependence structure allows modern commercial photo apps to safely leverage CNNs, despite missing digital-camera metadata. Manual representation engineering is omitted in favor of feature extraction via embedding/presentation learning directly from the data. These representations can potentially enhance subsequently trained CNN architectures, either through simple model pre-training with minimal supervision or richer feedback mechanisms.

7.2.2. Evaluation Metrics and Validation Frameworks

Selecting proper evaluation metrics is critical for accurately assessing model performance and informing final decisions based on predictions. The context dictates which metrics are of particular importance; for credit risk these often include discrimination ability (the capacity to rank risky and non-risky subjects), calibration (the relationship between predicted probabilities and observed outcomes), and economic valuation (i.e. profit and/or ROC-AUC). For fraud detection, emphasis is usually placed on detection quality (considering all true events across classes), timeliness (detecting events before or close to the moment they occur), and robustness (exposure to adversarial or out-of-distribution attacks). Market predictions require particular scrutiny as temporal dynamics pose several challenges; forecasted return series are typically backtested on simulated trading strategies and profits are monitored over time.

A variety of validation strategies also exist. Cross-validation is the golden standard for performance estimation but observing the future is generally impossible. For risk-graded problems where test set sizes allow it, a stratified framework can be devised. The most common strategy for time series is splitting the data temporally into training and hold-out sets. Stress-testing — where models are evaluated out-of-sample — is also an obvious routine that enables risk quantification under extreme events. Finally, backtesting accuracy of market trades is at the core of financial predictive modelling.

7.3. Deep Learning in Credit Scoring and Underwriting

Deep Learning for Credit Scoring and Underwriting — The suitability of deep learning for risk assessment in credit scoring and underwriting has been analyzed under three lenses: performance and prediction quality, fairness along critical demographic dimensions, and considerations associated with practical deployment. Discrimination and calibration of the CCN are assessed in detail, and a transparent, interpretable follow-on model serves regulatory requirements. Detection of Model Bias and Risk Assessment in Credit Scoring and Underwriting Risk assessment in credit scoring and underwriting is the first structured, objective evaluation of deep learning discrimination and calibration properties. While convolutional neural networks have become the de-facto standard in computer vision, deep learning has found only limited application in credit, fraud, and market prediction. Credit-scoring models play an important part in financial decision making, and non-bank financial institutions require a credit score to assess creditworthiness. CCN is found to outperform benchmark methods on conventional discrimination and calibration metrics. An operating characteristic curve for profit/ROC-AUC is proposed as a model-selection criterion. A post-hoc explainability method is also applied and other aspects of predictive risk assessment considered.

The limited number of patterns in the dataset raises the specter of overfitting, but the conclusion on the nonbank is that the firm is neither in a position to exploit a business opportunity nor required to take measures to avert an impeding business risk. The provision of scoring models for use by credit institutions, or for use in other financial services such as the supply of insurance is a task at the heart of the EBA's work. Model risk stresses the importance of validation to confirm that a model is performing as expected. Detecting model bias increases confidence that the model is performing as expected along protected demographic dimensions.

7.3.1. Transparent and Interpretable Models in Credit

Three mutually beneficial motivations ensure transparent and interpretable credit-scoring and underwriting models. Regulatory frameworks increasingly emphasize fairness. Transparent model design can positively affect customers' acceptance of a funding decision. Effective deployment requires understanding the model that ultimately selects funding applications. Transparency and interpretability can help satisfy these requirements and risks can therefore be determined in a manner that is aligned with regulatory concerns.

Fig 7.2: Triadic Transparency: Aligning Global, Local, and Rule-Anchored Interpretability with Regulatory Compliance and Consumer Trust in Credit Underwriting

Three approaches are discussed: explicable methods that provide a global perspective on the model; methods that yield local interpretability for a specific prediction; and rule-anchored techniques that deliver a small number of comprehensive rules. Local explanations enable scrutiny and auditing. Rule-anchored models can inspire confidence in customers and offer a clear articulation of why the funding application has been approved or rejected.

7.3.2. Handling Imbalanced Data and Rare Events

Credit scoring and fraud detection often work with highly imbalanced datasets. In credit risk, the relative number of bankruptcies is historically small compared to the number of people who repay their loans. Even in datasets from credit card companies, where fraud can occur at a greater rate, the number of fraudulent transactions is low relative to the overall volume of transactions. In practical applications, the class imbalance of a dataset is only one of many considerations, but managing it well will often improve performance.

There are several strategies commonly employed when handling an imbalanced dataset, most of which alter the training dataset to (1) achieve a better balance between the number of samples in each class, or (2) inform the model of the challenges created by the class imbalance. Sampling methods can reduce the number of majority samples through under-sampling or increase the number of minority samples through over-sampling. An alternative approach is cost-sensitive learning, where greater costs are assigned to misclassifying a minority instance than a majority one. Cost-sensitive learning could also be applied through the optimization of a focal loss function during training. Another way to address imbalance is through anomaly detection algorithms, which focus solely on detecting instances in the minority class, or techniques for generating synthetic data that augment the training set with new misclassified samples.

7.4. Deep Learning for Fraud Detection

Detection of fraud in real time demands timely decision-making and detection quality is paramount. Model robustness admitting adversarial attacks is also a concern, especially for credit card transactions where the financial balance is at risk and fraud detection for target customers may be delayed until some time after the event. Behavioral profiling provides insight into normal behavior, and features that reflect the user's behavior within a session can help detect events occurring within a limited time. Active user-session features determine the pattern of normal behavior in profile models that learn the user's behavioral habits over time, take account of the mutual influence among users, and evaluate the at-risk status of a user in the current session.

Anomaly detection models can also be developed for fraud detection, under the assumption that only a tiny fraction of the data comes from fraudulent events. Since unlabeled data usually far exceeds labeled data, a combination of supervised and unsupervised methods can better take advantage of the richness of such data. Behavioral analytics models are effective but require training for each new session; a hybrid solution combining rules with exploitation of user and metadata network properties could improve the detection rate without significantly increasing the false alarm rate. Models are subject to adversarial considerations, including data integrity and information reliability, and defenses against evasion, poisoning, and data initiative leakage are crucial for bank customers' credit card transactions.

7.4.1. Behavioral Analytics and Anomaly Detection

Detection quality, timeliness, and robustness are key criteria for fraud detection. User behavior changes over time. A sudden deviation from previously learned behavior can be indicative of fraud. Fraudsters, then, commonly exhibit behavior inconsistent with that of normal users. Focus on methods designed to analyze customer behavior or model user-session features. Detecting fraudsters before they act can reduce a company's losses; hence, prevention is important. Incorporate the temporal dimension naturally by putting a joint focus on user behavior and the sequence of actions taken by a user session. Behavioral analytics, user-session features, network effects, and anomaly detection techniques can be considered separately or jointly. Unsupervised and supervised approaches are often complementary.

Fraud detections based on user behavior or user-session features are potentially capable of detecting fraudsters before they "strike"—i.e., before actual fraud has occurred. Detect fraud in a fraud-centric way, using supervision based on labeled fraudsters. Behavioral models of user actions or physiological models of user-session features from normal customers can enable anomaly detection.

7.4.2. Adversarial Robustness and Model Security

The unprecedented success of deep learning has led to its heightened adoption in mission-critical applications, including those involving security, safety, and finance, where a smart decision can save lives and enormous resources. Deep learning model behavior is often unintuitive, posing vulnerability and security concerns, especially for classification and generative tasks, wherein malicious users could attempt to evade an authentication or detection system. Trust reallocation of intelligent systems, which was formerly based on expert sifting of false positives, is more challenging. Successful scripts for performance undermining have been demonstrated on image, speech, and text classification models. Misleading samples might be stored for later poisoning of a model, or even crafted for decision triggering when one aims for increasing the likelihood of bank loan approval by evading the underlying decision-making model. Furthermore, the growing complexity of deep neural networks may engender inadvertent leakage of sensitive training data.

Consider the task of fraud detection, where fast reactions are needed and models are vulnerable to evasion risk. Instead of asking how to protect models against evasion, the risk associated with evasion attacks and the effectiveness of various protective options should be studied. Robust frameworks for detecting such evasion attacks involve data labeling, security mechanisms, and architecture adapted to anomaly detection. In addition to defending against evasion attacks, it is likewise necessary to protect against poisoning attacks (reducing a model's resistance against evasion) while ensuring model security. Data integrity must also be validated in terms of quality, provenance, and trust by leveraging data integrity techniques or developing detection models to uncover data integrity breaches.

7.5. Market Prediction with Deep Learning

Predicting the future direction of market indices is a complex task because financial markets aggregate information about a wide range of factors, including economic trends, monetary policies, business cycles, supply and demand, weather patterns, elections, and geopolitical events. The time-series aspect of financial data further complicates prediction. However, some researchers have sought to improve predictive performance by incorporating a variety of alternative data sources—often related to sentiment—with varying degrees of success. Nonetheless, no data source is without its flaws, and even state-of-the-art models, built on years of research, have proven brittle in a changing environment.

Although recurrent neural networks (RNNs)—which can be thought of as an extension of hidden Markov models—have been the dominant approach to sequential data, some studies have found that simple fully connected architectures or 1D convolutional networks can outperform LSTMs when trained on a sufficiently large dataset. More recently, attention-based architectures have gained prominence. To allow RNNs to benefit from global information rather than just the preceding time step, a hybrid architecture combining RNNs and attention mechanisms has been proposed. It has also been suggested to apply sequence classification instead of horizon-aware prediction when predicting a time-series label from temporally correlated features. Achieving a clear distinction between the seen and unseen periods of a time series can be difficult when using a sliding-window approach. Different sliding-window strategies have been proposed to facilitate time-series classification with a shallow architecture.

7.5.1. Time Series Modeling and Sequence Learning

Market-influencing events (such as data privacy scandals, elections, natural disasters, and pandemics) and periodic events (such as seasons and holidays) disrupt and change underlying trends, which can be predicted by market data (news sentiment, supply and demand); these events are reflected in social media. A temporal prediction algorithm can be extended to support horizon-aware forecasting, which improves prediction performance. In algorithmic trading, shorter forecast horizons are favored because prediction error elevation is a common phenomenon. Models whose underlying learning mechanisms are based on prediction are suitable. Sequence-to-sequence models are often based on recurrent-neural-network (RNN) architecture; thus, multi-step predictions can be approached directly. Sequence-to-sequence models based on attention mechanisms can be used for longer- and horizon-aware forecasting tasks; these models do not decay with increases in the prediction horizon, and prediction error elevation is alleviated.

Temporal dynamics can be modeled by representing time series as sequences that follow a natural ordering. Sliding windows are one of the classical approaches used for time-series representation—by partitioning a sequence into overlapping windows. Each sample is a sequence of sample size W with time-step size S, and samples projected onto the same time axis can be arranged as a 3D tensor. A suitable approach for time-series modeling, with emphasis on probability distribution, is sequence generation supported by a generative model, which generates a sequence by determining one element in the sequence at one time.

7.5.2. Integrating Alternative Data Sources

Although financial prediction models are generally based on information contained in traditional price, volume, and technical indicator data, it is reasonable to conjecture that the use of external data covering additional aspects of investor behavior can aid forecast accuracy. Three different data sets are constructed and used in experiments: emotion data derived from social media text analysis, forecast data published by well-known market experts, and news sentiment data. The forecast accuracy is improved through combination of these alternative data sources with the use of the orthogonalized partial least squares method. Financial forecasting is then treated as a multi-source and multi-output learning problem by training a unified, joint, cross-source and cross-output model, which further benefits accuracy.

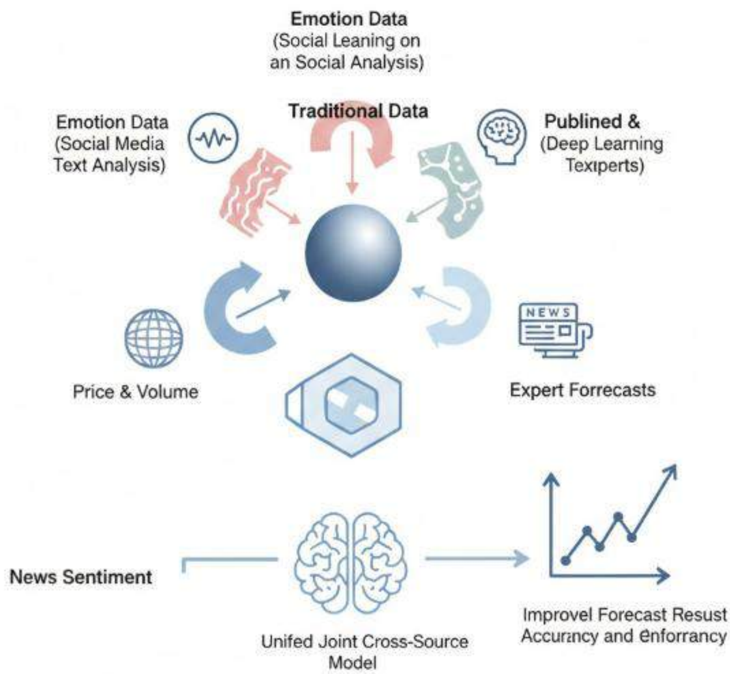

Fig 7.3: Multi-Source Latent Feature Fusion: Orthogonalized PLS and Joint Cross-Output Learning in Alternative Data-Driven Financial Forecasting

Combining data from multiple sources increases market prediction performance; however, the results depend on both the data source used and how the prediction task is formulated. First, market sentiment extracted from social media text using deep learning-based techniques offers an improvement in accuracy when integrated with traditional feature sets. Second, the use of expert forecasts provides a valuable additional signal that can drive accuracy higher. Third, fusing independent news sentiment assessment data may not produce the obvious accuracy improvement expected when incorporated into the prediction challenge. Finally, a joint cross-source model that considers various alternative data streams together serves as a powerful multi-source learning formulation but may suffer when sources are not diverse or complementary.

7.6. Data Governance, Privacy, and Compliance

Fundamental to risk is the quality of data and adherence to civil and data regulatory compliance. Use of deep learning models is beyond risk-positive business cases; detrimental effects can produce wholly uncorrelated external business risk and operations risk. Data governance concerns and privacy requirements—e.g., GDPR, CCPA—are paramount for success, from initial data acquisition to model deployment. Resulting data practices should incorporate the principles of data quality, provenance, lineage, and lifecycle management. While deep learning models generally occupy the external business risk domain, the associated data practices influence operational risk wishes.

A strong connection to civil rights legislation is key to risk-positive deep learning for business case and model risk. Data-minimization practices, such as federated learning, aid compliance. Achieving true AI capability, beyond mere automation, demands not only accurate and responsible models but also Data Minimized, Respected AI. Data-positive AI requires the adoption of data-quality principles such as governing trained models, inspection readiness, and auditability. Privacy-preserving methods may reduce adverse impact in domains with protected classes, especially when akin to prediction equity—however, the cost arising from misleading labels usually cannot be wholly eliminated. Privacy-preserving techniques encompass data minimization, federated learning, and differentially private machine learning.

7.6.1. Data Quality, Provenance, and Lifecycle

Risk mitigation for the use of data in assessing credit, detecting fraud, and predicting the market requires thorough governance, management, and monitoring, particularly focused on quality throughout the lifecycle, including management of sensitive content, accurate activities and changes, and external provenance.

Data quality is the fulcrum for responsible machine-intelligent credit scoring, underwriting, fraud detection, and market prediction. Accurate, comprehensive, current, and credible data is essential, but hard to attain. Bias, noise, corruption, incompleteness, and missing values can together nullify the most sophisticated learning frameworks. Applied risk practices thus need preventive, detective, and corrective activities for quality assurance and improvement throughout the model lifecycle. Operations should span the data spectrum and include provenance, lineage, activity, and lifecycle management. Governance by ML risk, audit, and privacy officers must also be reinforced by collaboration with engineering. Transparent, reproducible, and explainable risk are improbable without validation of internal and external data.

Provenance captures the origin and complete chain of creation or modification during a data object's existence. It thus outlines whether, how, and by whom the delivered data had been created or changed. Timely indication of problems that corrupt the integrity of the data is crucial in all stages: preparation, training, testing, production, support, and obsolescence. Risk teams deal with sensitive data (such as those protected under the General Data Protection Regulation) by minimizing data, federated learning with local aggregators, and differential privacy through random noise augmentation to obscure the exact value of records. Ample documentation, management, and storage of the responsible use of data during the assessment of credit, detection of fraud, and prediction of the market are vital for effective audit.

7.6.2. Privacy-Preserving Techniques in Risk Models

Risk-data governance mandates compliance with data privacy laws, which support data subjects' rights over their information to mitigate the adverse impact of data collection. For financial service providers, the implications are substantial; credit and fraud risk are intimately intertwined with sensitive personal data, and market risk may involve information about trading activity or product purchases that could identify customers. While non-public information is often enriched with additional variables before model training, comply with data minimization strategies, by mitigating the quantity of sensitive information stored. Policies limiting the amount of processing are meshed with the industry-specific principle of purpose limitation; risks tied to data use may be lessened by not keeping data longer than necessary, deleting information as soon as it is no longer required to fulfill a specific, legitimate purpose.

Privacy-preserving techniques strive to address the restrictions imposed by legislation through controlled procedures and crafting the risk assessment without directly leveraging the sensitive personal information of customers. These methods may also serve to alleviate the data governance implications, because the information used to predict the risk considered is not stored and made available for non-tertiary usage. Federated learning enables computers to collaboratively model a risk task without joint sharing of the data, by moving the processing to the sites where the data are located without exposing the information being processed. When credit models are necessarily trained using sensitive information, the risk may still be lessened by introducing differential privacy, such as integrating noise in model training to guarantee sufficient privacy even if users voluntarily disclose the classification scores.

7.7. Conclusion

Framing risk with machine learning—particularly with advanced deep learning models—revolves around three core questions. How successfully and reliably can machine-learning models predict creditworthiness, unusual or fraudulent behavior, or future market price movement? Like any artificial-intelligence system, a machine-learning model's success depends largely on the quality and suitability of the input data. How good, timely, and relevant is the information used by the models? Yet a third question emerges from the demand for greater societal equity, fairness, and privacy; it looms increasingly large in discussions of automated systems and their deployment: Are the decisions made by these systems transparent, explainable, and auditable, and do they function in ways that do not discriminate among certain demographic or social groups?

The answers streamline the many applications of deep learning to credit, fraud, and market prediction. For credit risk, three aspects are paramount: predictive performance, fairness or freedom from discrimination, and the ability for external regulation and audit to build trust. For fraud detection, the focus is on detection quality (correctness and speed), and the need to go beyond classification to encompass other aspects related to risk. For market prediction, the supply and integration of specialized data—such as news, tweets, social media, or a company's credit rating—are critical, since one or more variables might alter market behavior long before they flow into the purely market-critical part of the data, with important implications for prediction quality.

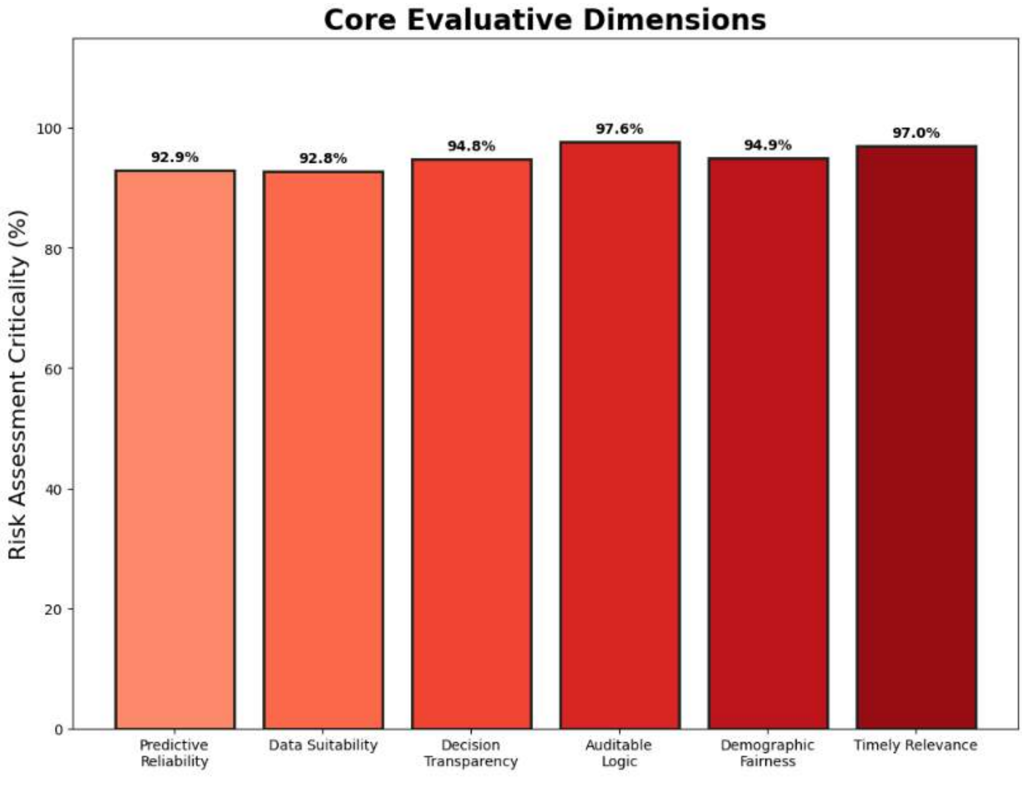

Fig 7.4: Core Evaluative Dimensions

7.7.1. Final Thoughts and Future Directions in Deep Learning for Risk Assessment

Data-driven machine intelligence approaches have demonstrated remarkable yet uneven success in risk assessment for credit, fraud, and markets. Recent years have witnessed rapid growth in deployment possibilities and operational constraints, especially for neobanks and non-financial institutions engaged in payment processing or credit issuance. As such, the maturity and scale of these applications create a unique opportunity for machine-intelligent deep learning models for risk assessment to inform development directions, guide future research, and generate practical deployment principles in an application area that is often viewed through the lens of much more specialized model architectures.

A positive result in any individual task area—credit, fraud, or market risk—is certainly worthwhile. But it is the collective assemblage of these results that is especially significant. Credit scoring and underwriting affect the allocation of pool liquidity, determining who should receive a loan, under what conditions, and from which institution. Fraud detection within those products and services helps protect both institutions and users from harms arising from fraud. Market prediction helps institutions adapt to external factors, acts as a signal to decision-makers, and constitutes an investment in its own right. The ability for an institution to invest in one of these areas—and potentially all three—represents an opportunity cost and competitive advantage.

References

Dal Pozzolo, A., Bontempi, G., Snoeck, M., & Snoeck, M. (2020). Adversarial drift detection in credit card transactions. Pattern Recognition Letters, 136, 252–258.

Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. Information Sciences, 479, 448–455.

Nagubandi, A. R. (2023). Advanced Multi-Agent AI Systems for Autonomous Reconciliation Across Enterprise Multi-Counterparty Derivatives, Collateral, and Accounting Platforms. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 653-674.

Nagabhyru, K. C. (2023). From Data Silos to Knowledge Graphs: Architecting CrossEnterprise AI Solutions for Scalability and Trust. Available at SSRN 5697663.

Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. Expert Systems with Applications, 240, 122156.

Kaggle, M., Dal Pozzolo, A., Bontempi, G., & Snoeck, M. (2021). Calibrating probability with undersampling for unbalanced classification. Pattern Recognition Letters, 140, 224–231.

Guntupalli, R. (2025, June). AI-Powered Data Analytics in Cloud Computing. In International Conference on Data Analytics & Management (pp. 280-289). Cham: Springer Nature Switzerland.

Zhang, Z., & Zhou, Z.-H. (2021). Adversarial learning for generative models: A survey. IEEE Transactions on Knowledge and Data Engineering, 33(3), 1026–1040.

Varri, D. B. S. (2022). AI-Driven Risk Assessment And Compliance Automation In Multi-Cloud Environments. Available at SSRN 5774924.

Carcillo, F., Dal Pozzolo, A., Bontempi, G., & Snoeck, M. (2021). Scarff: A scalable framework for streaming credit card fraud detection with concept drift adaptation. Information Fusion, 71, 182–197.

Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. International Journal of Scientific Research and Modern Technology, 1(12), 227–237. https://doi.org/10.38124/ijsrmt.v1i12.1111

Juszczak, P., Adams, N. M., Hand, D. J., Whitrow, C., & Weston, D. J. (2008). Off-the-peg and bespoke classifiers for fraud detection. Computational Statistics & Data Analysis, 52(9), 4521–4532.

Rao, A. N., Garapati, R. S., Suganya, R. T., Kaliappan, A., & Kamaleshwar, T. (2025, August). Smart Solar Harvesting and Power Management in IoT Nodes Through Deep Learning Models. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

Bahnsen, A. C., Aouada, D., & Ottersten, B. (2015). Cost-sensitive decision trees for fraud detection. Expert Systems with Applications, 42(3), 977–986.

Vadisetty, R., Polamarasetti, A., Goyal, M. K., Rongali, S. K., kumar Prajapati, S., & Butani, J. B. (2025, May). Cloud-Based Immersive Learning: The Role of Virtual Reality, Big Data, and Generative AI in Transformative Education Experiences. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-6). IEEE.

Wang, S., Cao, J., Yu, P. S., et al. (2022). Deep learning for anomaly detection: A survey. ACM Computing Surveys, 54(2), 1–38.

Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.

Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. ACM Computing Surveys, 52(1), 1–38.

Amistapuram, K. (2022). Fraud Detection and Risk Modeling in Insurance: Early Adoption of Machine Learning in Claims Processing. Available at SSRN 5741982.

Raghavan, S., & Manchanda, P. (2021). Behavioral analytics for fraud detection. Journal of Marketing Analytics, 9(2), 73–89.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

Siva Hemanth Kolla. (2022). Knowledge Retrieval Systems for Enterprise Service Environments. International Journal of Intelligent Systems and Applications in Engineering, 10(3s), 495–506. Retrieved from https://ijisae.org/index.php/IJISAE/article/view/8037

Goodfellow, I., Bengio, Y., & Courville, A. (2020). Deep learning. MIT Press.

Zhou, C., Paffenroth, R. C., et al. (2017). Anomaly detection with robust deep autoencoders. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 665–674.

Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows.

Chen, Y., Kou, G., Peng, Y., & Alsaadi, F. E. (2021). A multi-layer ensemble framework for credit risk prediction. Expert Systems with Applications, 174, 114762.

Eling, M., Nuessle, D., & Staubli, J. (2022). The impact of artificial intelligence along the insurance value chain and on the workforce. Journal of Risk and Insurance, 89(2), 1–38.

Xu, Y., Sun, J., & Liu, J. (2023). Fairness-aware machine learning for insurance risk prediction. IEEE Access, 11, 24501–24513.

Inala, R. Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective.
Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 17(3), e0265480.
Richman, R., & Wüthrich, M. V. (2021). A neural network extension of the classical chain-ladder model. Insurance: Mathematics and Economics, 99, 331–347.
Serrano-Guerrero, J., Herrera-Viedma, E., & Olivas, J. A. (2021). Fuzzy and machine learning models in insurance risk classification. Applied Soft Computing, 102, 107070.