



Sustainable AI-Driven Architectures for Real-Time Derivative Accounting Across Multi-Counterparty Financial Networks

Vinod Battapothu

Independent Researcher, Germany

vinod.battapothu.researcher@gmail.com

Isabella Romano

Academic Research Assistant, Mediterranean University of Rome, Italy

isabellaromanoo@outlook.com

ABSTRACT: Green AI architecture design for real-time derivative valuation across multilayer multi-counterparty networks that account for systemic and market-specific risk scenarios. A real-time data-pipeline and provenance management approach incorporates user-defined risk metrics that are used for SLA invocation and applied to distributed data streams to establish Near or Real-time apparent price and P&L: Risk vectors are also exploited for scenario generation in supporting and orchestrating stress testing. Provenance and quality parameters control the overall risk of the output data, whilst stress testing scenarios based on the same risk metrics assess the capacity of a Green AI architecture to fulfil the stipulated demands. A supporting process implements data-longevity concepts in SLA-driven regulatory reporting.

The drive towards supporting sustainable computing integrates the practices of the emerging Green AI domain with financial Information Communication Technology (ICT) infrastructures. This ushers in the consideration of Green-Computing concepts and metrics at a more detailed architectural level—directly extending the “Green” and “Minimal-Mid-Point” ICT enablers of Risk-Aware Green AI to include new components. This exploration addresses the Integrated Financial System Integration challenge, in the specific context of Real-time Apparent Pricing. The focus is the real-time accounting of retail derivatives traded over multiple counterparty pairs, where satisfactory performance characteristics need to be maintained across a specified amount of risk-based latency.

KEYWORDS: Green AI; computational sustainability; financial system integration; risk-aware AI; derivative accounting; multilateral counterparties; real-time data pipelines; systemic risk; regulatory compliance; data provenance; latency.

I. INTRODUCTION

With the growing interest in AI technologies, attention has recently turned to the carbon footprint of training and using such systems. The focus on “Green AI” is thus timely. Such attention should extend beyond the training of models to cover the provisioning of the infrastructure, its energy consumption in both the training and inferencing phases, and the value generated for the energy consumed. AI applications are gradually becoming part of critical financial infrastructures, generating reports and supporting decisions related to financial risk.

A possible branch of Green AI comprises AI systems that continuously aid humans in tasks that require constant attention, are monotonous, or are beyond human processing capacities, for example, simultaneously monitoring the credit risk of hundreds of market participants. These AI agents ingest, read, extract, reason and process information from various sources to generate proactive alerts to human operators. With such applications, it is crucial to guarantee not only that the information supplied to the operator is trustworthy but also that the use of the AI system does not increase the risk of the financial institution. Failure of such applications may thus generate news articles with “unexpected losses” following “unexpected” events. In General, Stress tests and scenario analysis compiled by prudent financial institutions address this issue even for rare events: loss distributions are built, with a certain probability level, around the worst-case losses consider all possible variations of the endogenous and exogenous variables. An extension of such tests is needed for automated setting.



A natural extension of Green AI in this context is the development of **AI-aware stress testing and scenario analysis frameworks** specifically designed for automated decision-support systems. Traditional stress tests focus on market, credit, and operational risks arising from endogenous and exogenous variables, but they rarely account for risks introduced by the AI system itself—such as data ingestion failures, reasoning errors, model drift, feedback loops, or over-reliance by human operators. For AI agents that continuously monitor complex environments and generate proactive alerts, stress testing must therefore encompass *algorithmic behavior under extreme but plausible conditions*, including corrupted data streams, delayed or missing signals, adversarial inputs, and cascading alert failures. Loss distributions should be augmented to reflect not only financial outcomes but also *model uncertainty, explainability gaps, and human–AI interaction risks*. By integrating AI-specific failure modes into scenario generation, financial institutions can better quantify tail risks attributable to automation, ensure that AI deployment reduces rather than amplifies institutional risk, and prevent “unexpected losses” that stem less from unforeseen events than from insufficiently tested automated systems.




Fig 1: Architecture for AI-Driven Financial Management

1.1. Background and Significance: Recent years have seen calls for greater responsibility and accountability in AI acquisition and use, culminating in the integration of an AI lifecycle concept into IT governance frameworks such as the Information Technology Infrastructure Library (ITIL). A shift towards Green AI, where models' footprints are and operated with awareness of environmental sustainability and carbon footprint, is one important outcome. Green AI extends prudence in machine learning transfer and resource utilization to real-time or online applications and push button deployments of trained and optimized systems. Architecture holds the center ground: designing systems that efficiently deliver (Green AI) models during deployment and minimizing ongoing carbon emissions. Domain-specific optimizations for computational processes within AI systems, focusing on operational deployment, push button modelling, and the supporting IT infrastructure, complete the Green AI picture.

Such developments are especially important for financial institutions, which consume an estimated 10% of the total energy usage of New York City. Financial system integration mechanisms, operating and regulatory reporting, auditing and compliance within multi-counterparty environments, remote data storage and processing for universal banks and their clients, and privacy and security strategies combine to define ICT Digital Deliberation. Financial AI explores increasing operational efficiency and effectiveness of such Digital Deliberation at multiple production environments for end-users at all levels of the financial services sector while reducing overall carbon footprint. Risk-aware real-time or online banking operations, with emphasis on derivative exposure management in a multi-counterparty environment, exemplify the approach.

II. BACKGROUND AND MOTIVATION

The evolution of information technology toward higher energy efficiency and lower resource usage has attracted considerable research interest, culminating in the notion of Green AI. AI models and infrastructure that support energy-efficient training and inference are classified under the Green AI umbrella. These systems are generally trained under energy constraints, deployed and operated in energy-efficient datacenters, leverage low-carbon-energy AI infrastructure, and green AI metrics such as carbon emissions are used to measure performance. High-energy consumption is often considered a burden for practically unexplored application domains beyond the realms of image processing, natural language understanding, or robotics (where massive models achieve state-of-the-art results). In finance, questions about whether it is possible to afford such massive-model energy consumption remain. However, when it comes to rich-agent ecosystems such as finance, the production and consumption of myriad small models provide opportunities to deploy



energy-efficient AI tools and closable pipelines. This is particularly the case in high-frequency trading systems that process data streams at microsecond and millisecond resolution.

Despite their heavy use and recognised importance, a lack of multi-currency and multi-countries pricing has made the proper accounting and risk valuation of derivatives instruments a difficult task. Establishing reciprocal contracts between different countries zone agents and automatically generating these agreements creates multilateral interaction models. It is believed that models supporting multilateral pricing can also be developed. Such models require data provisions concerning a derivative's risk profile configuration, conditioning it onto properly generated currents between players, public parameters, and stress testing dynamics. These innovations will enable real-time production dashboards that automatically adjust risk profiles and provide the required information.

Equation 1: Core equations implied by the paper (derived step-by-step)

Let:

- V_t = model value ("apparent price") of a derivative portfolio at valuation time t
- $V_{t-\Delta t}$ = previous valuation
- CF_t = cashflows realized in $(t - \Delta t, t]$ (coupon, settlement, fees, etc.)

Definition: incremental P&L

$$\Delta \Pi_t = (V_t - V_{t-\Delta t}) + CF_t$$

Step-by-step

1. Mark-to-market change: $V_t - V_{t-\Delta t}$
2. Add realized cashflows over the same interval: $+CF_t$
3. Result is the accounting change in value over the interval.

This is exactly the quantity the paper wants "near/real-time" via a latency-aware data pipeline.

2.1. Research on Green AI and Sustainable Computing AI: has been classified according to two opposed trends. The first, termed "Green AI", aims to reduce energy consumption and carbon emissions while achieving goal-related performance. The second, "Red AI", refers to achieving state-of-the-art performance in complex tasks, often at the cost of massive computational resources, with little regard for carbon emissions. To constrain energy consumption and the associated carbon footprint of ICT infrastructure, for example, research studies demonstrate the optimization of the operational utilization of physical resources and of the efficiency of photographic systems. The combined application of conservation and efficiency principles helps green the physical component of ICT both in expenditure and in operational carbon footprint.

Sustainability in Computing denotes a focus in computing research on making and using computing and affective systems products that minimize both the overall supply chain costs and existing and potential future environmental impacts and adverse social consequences throughout every stage of a system's life cycle; a growing proportion of the power consumed while a computing device is in use derives from carbon-producing sources such as coal. The increasing difficulty of finding energy-efficient solutions has motivated work on Power-Aware Design for Hardware, focusing on minimizing dynamic and static power dissipation during the utilization of physical resources through techniques such as clock gating and mainly the use of CMOS low power and Ultra Low Power technology.

2.2. Financial System Integration Challenges: The financial system's array of products and services, its associated computable and persistent risk management paradigms, their representation in informing role-based policy, and the supervision and enforcement of regulated entities is executed using numerous Information Communication Technology (ICT) subsystems. The integration of these subsystems often lacks the advanced capabilities present in more sophisticated systems, while the requisite maturity for green implementations is frequently overlooked. The proposed demonstration research emphasizes AI architectures capable of risk-aware real-time accounting for derivatives in a multi-counterparty framework. Supporting Electric-Union-architecture, Smart-Grid-and-Microgrid, and 5G-sustainable-peak-sharing arguments, these applications exemplify risk-aware regulatory compliance for active derivatives and enhance the prospects of early-energy-derivative markets.

A typical response to the spiral of risk within the financial system has centered on the formulation of additional regulations. This regulatory-crisis-recovery cycle remains a persistent feature of the edifice of the financial system, notwithstanding the explosive growth occurring in various sectors. Effective financial integration facilitates the reallocation of resources to where they are most scarce, yet integration within the financial system often remains shallow,



with the potential for fragmentation manifesting in major disruptions in some areas. An advanced set of architectures is required for the effective integration of the considerable range of financial products and services available in the system, and for optimally meeting the infrastructures envisaged in Electric Union, Smart Grid and Microgrid, Enterprise Architecture for Electric Union, Information and Communication Technology for Electric Union, Artificial Intelligence for Electric Union and Health Care, and even more focused Sustainable Peak Sharing.

2.3. Real-Time Derivative Accounting and Multilateral Counterparties: Ecosystem-based, multilateral, multi-counterparty interactions are a key part of the modern financial world, especially in relation to risk management, as derivatives help transfer and absorb risks. Most derivative transactions are settled at the end of a specified period; thus, they offer the least near-term liquidity. In case they are not netted over time, they also represent the largest exposures to a counterparty, and so they are settled daily. Daily settlement involves computing the change in the fair value of the derivative portfolio; if this change in value is positive for one party and negative for the other, cash is transferred. Such daily settlements can be performed over multilateral infrastructures, even if actual transactions take place only between two parties.

Although the values of the derivatives can be computed at any time, the actual cash flows are not available until the next business day. Thus, the exposure to each counterparty until the next business day depends on these values and can change sign. As the timing is not synchronized, it is also possible to define these daily cash flows according to the party who has computed the derivative value. Although this is a common practice, it is rarely done; indeed, no formal metric has yet been proposed to quantify the exposure resilience to these operations. A proper treatment of exposure must also consider counterparty risks induced by a deterioration of the credit profile of one or more counterparties.

Ecosystem-based, multilateral, and multi-counterparty interactions are fundamental to modern financial markets, particularly in the context of risk management, where derivatives play a central role in redistributing and absorbing risk. Because most derivative contracts settle at predefined intervals, they typically provide limited near-term liquidity and, when not netted over time, can generate substantial counterparty exposures. To mitigate this risk, derivative portfolios are commonly subject to daily settlement, in which changes in fair value are calculated and corresponding cash transfers are executed between counterparties. While these settlements may be processed through multilateral clearing infrastructures, the underlying exposures often arise from bilateral relationships. A key challenge lies in the temporal mismatch between valuation and cash settlement: although derivative values can be computed continuously, the associated cash flows are only realized on the next business day, leaving parties exposed to potentially volatile counterparty risk that can even reverse sign within this window. Moreover, the lack of synchronization in valuation timing allows settlements to be defined based on one party's computed values, a practice that is theoretically feasible but rarely implemented in a systematic way. Despite its relevance, no standardized metric currently exists to measure the resilience of counterparty exposure under such asynchronous settlement mechanisms. A comprehensive treatment of exposure therefore must go beyond valuation dynamics to incorporate the evolution of counterparty credit quality, capturing the additional risks arising from the deterioration of one or more counterparties' credit profiles.

III. THEORETICAL FRAMEWORK

The proposed Green AI architectures for financial system integration are formalized along three dimensions: risk-aware design principles for real-time derivative accounting in multi-counterparty settings, a data-driven operational paradigm for latency-sensitive product pricing, and the sustainability and compliance of risk-aware data pipelines.

Risk-Aware AI Architectures

Real-time systems for multi-party derivative accounting derive most attributions from inputs transmitted by participants' systems, enabling fast reaction to high-frequency market movements. While each institution derives liquidity, hence risk, for its own operations, the wider stress on the multilateral interaction may emerge unobserved, channeled into sudden bursts of volatility. Properly designed, such systems allow dedicated real-time derivative accounting and liquidity generation on feed-latency-sensitive products. These principles apply equally to real-time market making for listed derivatives.

Regulatory and compliance considerations become intrinsic rather than extrinsic quality for such analytics engines. IT service organization assurance reports (e.g., SOC 1, SOC 2) guarantee the suitability of an entity's controls relevant for internal control over financial reporting and the strict protection, availability and integrity of the information processed by the service organization. Well-implemented, permanent non-intrusive active detection testing of undesirable events for monitored products complements the traditional posteriori reporting.





Fig 2: Theoretical framework

3.1. Risk-Aware AI Architectures: The concept of Green AI incorporates two principal factors: a reduction in the volume of training and inference processes and a judicious use of energy-saving hardware. These two aspects contribute to lower energy consumption and greenhouse gas emissions, but they are not the only indicators of “green” systems. Another dimension of sustainability important to AI is compliance with regulations at national and international levels. In financial and other equally regulated sectors, Green AI means implementing AI systems that demonstrate compliance with relevant regulations. The proposed approach achieves compliance with the European Data Governance Act (GDPR), the Market Risk Directive (MRD), and the European Union’s Anti-Money Laundering Directive (AMLD). The resulting architecture may be labelled Risk-Aware Green AI, where compliance with regulatory and audit requirements is also a fundamental concern.

Putting this all together, an integrated financial system that supports real-time derivative accounting with risk-aware and green characteristics should have the following components: 1) sources of risk metrics for valuation and capital calculation; 2) a data pipeline for real-time delivery; 3) a mechanism to guarantee provenance and ensure data quality; 4) an architecture to optimize latency in data delivery; 5) a module for privacy preservation, security, and access control; and 6) the capacity for stress testing, scenario analysis, and regulatory reporting.

An integrated financial system that embodies the principles of Risk-Aware Green AI brings together sustainability, regulatory compliance, and real-time performance into a unified architecture. Beyond reducing computational workload and leveraging energy-efficient hardware, such a system treats compliance with stringent regulatory frameworks as a core dimension of “greenness,” particularly in highly regulated financial environments. By aligning with European regulatory requirements—including data governance, market risk oversight, and anti-money laundering obligations—the architecture ensures that AI-driven processes are transparent, auditable, and trustworthy. To achieve this, the system must incorporate reliable sources of risk metrics to support accurate valuation and capital calculations, supported by a real-time data pipeline optimized for low latency. Robust mechanisms for data provenance and quality assurance are essential to maintain integrity across the lifecycle of financial data, while dedicated modules for privacy preservation, security, and access control ensure lawful and ethical data usage. Finally, the architecture must enable advanced stress testing, scenario analysis, and regulatory reporting, allowing institutions to respond dynamically to market conditions while meeting supervisory expectations. Together, these components define a holistic, sustainable, and regulation-aware approach to deploying AI in modern financial systems.



3.2. Real-Time Data Pipelines for Derivative Valuation Sustaining economic development in a world of limited resources requires a paradigm shift towards the ecological and social sustainability of human activities. Developing and maintaining ICT systems with the smallest possible environmental footprint is recognized as a feasible approach for achieving such a shift. For complex, resource-intensive systems—such as finance and logistics—that have a pivotal role in the economy, integrating these systems into a single system is essential for minimizing their carbon and resource footprints. The real-time risk-aware AI idea advocates integrating risk-aware AI technology with risk-aware computing in developing such systems.

Real-time data pipelines for the risk-aware real-time evaluation of financial derivatives support critical operations such as pricing, risk assessment, and portfolio evaluation. While they have a crucial role for all derivatives, they are essential for those involving multiple counterparties such as credit default swaps, over-the-counter interest rate swaps, and collateralized debt obligations, as the aggregate exposure of the counterparties to a single entity must be assessed. Maintaining these operations in a real-time context also enables the support of other AI-based management functions. These functions may depend on the assessment or accounting of risk indicators related to expected shortfall and liquidity attribution but not necessarily on the holistic valuation of derivatives.

To control the deterioration of agent performance under increasing load or demonstrate that a given architecture is in the green label space, the environment must be protocolled and continuously measured. Latency-optimized real-time data flows for the risk-aware valuation of financial derivatives are essential to accomplish this task. They can also support disclosure and assurance, which has become critical for proponents advocating against business as usual, as the failure to comply leads to minimum to no market interest.

3.3. Regulatory and Compliance Considerations: Financial enterprises are subject to a variety of legal and regulatory compliance obligations that affect their systems and operations. Some obligations apply to the enterprise, while others relate to the enterprise's trading activities with specific counterparties. The architected system design incorporates two interrelated families of regulatory requirements. The first group encompasses external reporting obligations, such as those imposed by governments and supervisory authorities, while the second group consists of internal governance obligations, including reporting requirements for senior management and governing bodies.

The obligations imposed by supervisory authorities on banks and securities firms act as a floor for similar obligations on other financial enterprises and as a ceiling for obligations on nonfinancial enterprises and individuals. Stress testing and scenario analysis provide a means for conducting "what if" examinations of risk exposure in sensitive OTC trading portfolios that afford a degree of detail beyond aggregate measures. These capabilities are useful for senior management and governing bodies.



IV. ARCHITECTURAL DESIGN

A risk-aware real-time derivative accounting system configuration for multi-counterparty environments, supporting the business functions of multiple users, is depicted in figure 4. A common data-processing cluster runs a read-optimized instance of the RDBMS in proximity of read-only replicated instances of the SDBMS, which physically provide data pipelines for joint risk metric computation. Each user, belonging to a financial institution and maintaining a client-facing service such as a trading platform, integrates the proposed solution architecture for the collective management of counterparty credit risk of over the counter (OTC) derivative instruments. Asset scope and tenor set still serve as major criteria for data subsets; User 1 being interested in FX calendar swaps while User 2 queries implied volatilities for swaptions on interest-rate swaption surfaces.

Ad-hoc analysis of implicit positions in OTC derivatives and exposure to joint counterparties is supported by the implementation of a latency-optimized, risk-metric-generating data pipeline for derivative pricing when latent state-vector Markov chains in RM functions condition joint pricing by risk-neutral SDBMS ensembles. Latencies in risk metrics for derivative-marking purposes are determined allowing for the systematic analysis of short-lived strategies with respect to value-at-risk limits. The architecture and experimental results serve as a reference and provide guidance for generic latency-optimized, risk-metric-generating pipelines that accommodate the properties of contingent-claims pricing by physical and risk-neutral measures and clarify the structure of latencies in use cases with multiple active OTC market-making counterpart in risk monitoring of positions and joint exposures.

Green computing principles indeed support the main architectural components of the risk-aware real-time derivative-accounting system. A common data-processing cluster housing a read-optimized instance of an RDBMS in proximity of read-only replicated instances of an SDBMS enables business function functions for multiple users represented by financial institutions maintaining client-facing services—e.g., trading platforms—using a latency-optimized data pipeline to support ad hoc analysis of implicit positions in OTC derivatives and exposure to joint counterparties.

Equation 2: The paper's "risk = incremental P&L effect of unit exposure shift" → sensitivities

Let a vector of risk factors be:

$$\mathbf{x} = (x_1, \dots, x_m)$$

and portfolio value be $V(\mathbf{x}, t)$. For a small factor move $d\mathbf{x}$,

Step 1: first-order Taylor expansion

$$dV \approx \sum_{i=1}^m \frac{\partial V}{\partial x_i} dx_i$$

Step 2: identify "unit exposure shift"

If we take a **unit** shift in factor x_i : $dx_i = 1$ and all other $dx_j = 0$,

$$dV \approx \frac{\partial V}{\partial x_i}$$

Step 3: map to "risk" sign convention

If the paper's risk is "negative effect on P&L", then a simple risk measure is:

$$R_i \equiv -\frac{\partial \Pi}{\partial x_i}$$

In pure mark-to-market (ignoring intra-step cashflows), $\Pi \approx V$, so:

$$R_i \approx -\frac{\partial V}{\partial x_i}$$

Common special case (equity option underlying S)

- Delta: $\Delta = \frac{\partial V}{\partial S}$
- Delta-risk (paper sign): $R_S \approx -\Delta$

4.1. System Overview: A Green AI architecture integrating risk-aware compute and data components enables real-time derivative accounting in a multi-counterparty context. Potential applications include enterprise risk management, liquidity monitoring and planning, regulatory reporting, and audit support. Financial Integration is defined as the establishment of linkages among disparate financial services with the aim of improving coverage, efficiency, or stability of the system. Risk management is one sector in which financial system integration appears feasible. If bilateral derivative transactions revolve around a straw counterparty or a substantial proportion of them are multilateral transactions, the real-time value of derivatives will remain a mirage, even if computing costs can be eliminated.



A Green AI architecture integrating risk-aware computer and data components enables real-time derivative accounting in a multi-counterparty context. Potential applications include enterprise risk management, liquidity monitoring and planning, regulatory reporting, and audit support. Financial Integration is defined as the establishment of linkages among disparate financial services with the aim of improving coverage, efficiency, or stability of the system. Risk management is one sector in which financial system integration appears feasible. As long as bilateral derivative transactions revolve around a straw counterparty or a substantial proportion of them are multilateral transactions, the real-time value of derivatives will remain a mirage, even if computing costs can be eliminated.

4.2. Green Computing Strategies in Financial ICT: With financial information and communication technologies (ICT) contributing a significant share to global energy consumption, the industry bears a substantial responsibility for developing greener systems. Although a multitude of early approaches for reducing the power consumption of ICT systems exist, actual industrial implementations remain limited. These attempts largely focus on energy efficiency and carbon footprint minimization through improving the energy-efficiency of data centers and their network connections. Such measures rarely account for the total administrative cost in various aspects apart from energy consumption per se, let alone that of financial information processing in general.

The direct energy consumed in financial system operations is a negligible fraction when compared with the energy investments required to develop, set up and operate these systems. Administration and management thus warrant particular attention in improved green designs. Another prominent and generally neglected factor of Greenness is the hardware life cycle and utilization. Although hardware utilization metrics are improving for web and social network systems companies may for substantial amounts of hardware that are rarely used. The problem, though, is even more serious for the global financial system: The vast majority of the installed capacity remains dormant most of the time, at least as long as no crisis occurs and normal functioning normal service resumes thereafter. Enabling electricity-saving operations at the time is thus a critical requirement for any realistic future Green AI architecture in this area.

4.3. Multi-Counterparty Interaction Models: The proposed AI architecture supports two interaction models for several communicating counterparty banks where, besides bilateral positions, also the data associated to the positions of all the other banks are internally maintained. The first model serves as the simpler one, still benefiting from Green AI principles, that involve the interaction of three or more banks. For every bank the requested functionality is a real-time evaluation of its credit risk for the whole group in terms of standard metrics for counterparty risk exposure and default probability. Either the marginal risk information provided to another bank or the accumulated risk information for a bank acting as loading entity are both valid and useful, otherwise they would not be requested at that precise time. The requested functionality has all the features of a Risk-Aware AI application, as also the corresponding Risk Metrics have been previously defined. The implementation of the second model requires a further layer of interaction and coordination among the banks, but it is seldom available almost in real-time, mainly due to the fact that the three risk metrics involved are usually only provided when finally requested by a central Counterparty Clearing House (CCP). It concerns the real-time evaluation of the credit risk associated to the financial obligations towards a central CCP as a result of a netting mechanism and it is therefore naturally centred around the CCP.

In addition to the explanation of the two models supported by the proposed architecture, a discussion is provided about the possible advantages that would arise by putting in place such a systematic provision of information based on a real-time approach integrated into the public or private infrastructures of a country or multi-country area instead of an ad-hoc approach based on private communications. It also considers the case of a single one-way communication, where the information provided by one bank to another one is the only one requested or needed at that precise time. The characteristics of the communications allow the banks not to put any effort in these non-marginal exchanges other than that required for a normal operational activity, so generating indeed a pure marginal benefit that does not require to be "forced" into the architecture by the introduction of some incentive scheme.

Symbol	Meaning	Typical unit
V_t	Portfolio value ("apparent price") at time t	currency
$\Delta \Pi_t$	Incremental P&L over Δt	currency
$\mathbf{x}$	Risk-factor vector	varies
Σ	Covariance of risk-factor changes	varies
$E_k(t)$	Positive exposure to counterparty k	currency
$EPE_k(t)$	Expected positive exposure	currency
CVA_k	Credit valuation adjustment	currency



Symbol	Meaning	Typical unit
L_{e2e}	End-to-end latency	seconds
PUE	P_{total}/P_{IT}	ratio
CUE	carbon per IT energy	kgCO ₂ /kWh

Table: Symbol table (pricing + risk + exposure + green)

V. DATA MANAGEMENT AND PROCESSING

Data provenance and quality assurance strategies enable users to fulfil their own requirements while maintaining supporting traceability and reliability accounts at minimal effort. Proximity, bandwidth, and latency influences drive the construction of a favorably oriented monomial, with reduced-latency pipelines generated from the infrastructure scaling and resource requirements. Coupled with privacy, security, and access-control dimensions in the accordance-preferences space, the constructed data-flows are fused with quality contrast and direction-matching for privacy- and security-perferring outcome-hoistings.

Data provenance covering the source, ownership, life-cycle status, and freshness is gained from the data service taxonomy, allowing users to examine provenance records during selection. The complete origination chains are further preserved in the quality assurance process. Users choose threshold levels of quality attributes; provenance, for example, can drive their selection of origin locations if mismatched contrast settings exist between provenance and the current privacy-sensitive identity. Quality checks, directed by available resources and bandwidth, enable selective control of data sources with the least degradation cost, with the remaining requirements satisfyingly filled. In accurately controlled cases, however, origin selection occurs only after quality check failure.




Fig 3: Big data in financial risk management

5.1. Data Provenance and Quality Assurance: Risk-aware real-time financial ICT application architectures require comprehensive data provenance alongside traditional functional validation. Provenance records data flows for accountability, enabling high-frequency audits that increase user trust and reduce monitoring burdens, thus optimising resource consumption. Real-time assurance ensures compliance with user-defined constraints governing latency-sensitive metrics such as deposited funds and balance sheet totals. Strategies for quality assurance and privacy protection bolster provenance management.

Economic ICT applications necessarily handle extremely high data volumes, with low latencies required by the local equilibrium theorem. Provenance reduces implementation and management burdens across functionally disparate applications regardless of fluorometric surveying requirements. Real-time financial processes entail bidirectional data flows with content-varying volumes. Dynamic verification, implemented locally yet orchestrated externally, permits high-frequency flowspecific monitoring, verifying compliance with user-specified data constraints. Provenance sensors, examining multiple information streams, test high-frequency quantitative features, enabling detection of latent faults in multidimensional data patterns.

5.2. Latency-Optimized Data Flows: Managing the latency of messages flowing through any data pipeline is crucial to its acceptance by users. Pipelines that inject high latencies into data at source, before data propagate further downstream, are particularly undesirable, as they delay the availability of data not only at its immediate consumers, but also at all later stages of processing. Consequently, such messages are typically generated and processed in bulk. In this light, a realistic




latency-target for the pipeline between the company's derivative-systems and their sources of market-reference data is of the order of normal publishing intervals for such reference data, within the range of a few seconds.

The commodity-, cash-, and swap-market data-sources of interest typically publish snapshots of such normal fluctuations (i.e., noise) at publishing-interval of a few Hz (even up to several Hz in the case of high-volume currency-pairs and rates), but such high internal rates of noise-fluctuation are typically of little interest for realistic derivative valuations. Monitoring the real-time noise-fluctuations of all those reference-data streams is therefore very wasteful. The alternative is to replace this nonessential monitoring overhead by noise-fluctuation lines, each monitored at the real-time rates of change supplied by the respective market-data-source and whenever available at these market-data-sources. These noise-fluctuation lines can then be combined into the best-of-all-available-derivative-valuations. Such combination, in turn, needs to happen at such low latencies that they are negligible compared to the publishing latencies of the market-data-sources.

5.3. Privacy, Security, and Access Control: Real-time data pipelines are bridges between input data and the consumer applications that utilize them. Latency-optimized data flow patterns help ensure timely data delivery, but such patterns typically disregard security and privacy. Latency-optimized process chains transform the data but do not consider aspects such as Data Provenance or Data Quality Assurance. Privacy concerns and the goal of mitigating Cybersecurity risks in general add additional complexity. The data is subject to privacy legislation such as GDPR and, in some cases, jurisdictional constraints.

These considerations imply that the data has to be sufficiently protected such that access cannot be granted indiscriminately to external parties. Consequently, a Data Provenance approach with access control capabilities is recommended, combined with the Data Privacy concept for the reduced leakage of sensitive or private content from data sets. Nevertheless, it is important to treat Data Provenance and Data Privacy as two tightly coupled components. Privacy rules enforced by a Data Provenance scheme will inevitably introduce additional access control overhead when granting fine-grained access to data. In addition, access control policies on other elements of Data Provenance may often lead to restrictions on Data Privacy settings as well.



IV. RISK MANAGEMENT AND COMPLIANCE

Risk management, compliance, and regulatory requirements represent the context within which actual values of financial derivatives should be calculated. Derivative processes with excessively large positive or negative values at the valuation instant should be flagged. Unexpected parameter values typically provide early warning signs of underlying market events and are therefore crucial for the construction of early warning systems to support management and traders. Stress-testing and scenario-testing tools should also be implemented. If these metrics are ignored, one or more of the counterparties may be obliged to record large losses, or risk ratings may be downgraded, resulting in reduced access to financing, increased costs of funds, and reduced profitability. Stress, scenario, and sensitivity testing are mandated by the majority of regulators for all banks under normal operating conditions. All major regulatory reports are submitted to



regulatory bodies on a tight timetable, typically within two months. Banks need an independent IT capability to generate these regulatory reports, presenting results that are as reliable as possible and need to be adequately audited.

The cumulative use of cross-derivatives by all market participants results in a significant decrease in market liquidity. Major banks, hedge funds, and pension funds frequently hedge their portfolios with bonds in one or two currencies. Nevertheless, the market information provided by all these financial movements is scant. Opcredit and Opdeb Credit represent the most obvious measure of regional risk for lending operations with major banks, as these counterparties are in contact with market movements in every possible direction. They are obliged to consummate the operations because they are always ready for sale or purchase. Nonetheless, changes in national income are less readily reflected in indicator numbers, even though they remain the main determinants of exchange rates in the interbank market. The market, however, is powerless to stem the appreciation or depreciation of a currency that does not correspond to the flow of transactions. When this situation occurs, the market is sometimes favored by the banks in the direction opposed to the tendency indicated by changing differences in national incomes.

Equation 3: Risk aggregation with covariance (the “must monitor covariance/correlation” part)

Let gradient (sensitivities):

$$\nabla V = \left(\frac{\partial V}{\partial x_1}, \dots, \frac{\partial V}{\partial x_m} \right)$$

Step 1: P&L approximation

$$dV \approx \nabla V \cdot d\mathbf{x}$$

Step 2: variance of dV

$$\text{Var}(dV) \approx \text{Var}(\nabla V d\mathbf{x}) = \nabla V \Sigma (\nabla V)^\top$$

Step 3: volatility

$$\sigma_v \approx \sqrt{\nabla V \Sigma (\nabla V)^\top}$$

This is the standard bridge from “monitor correlation/covariance” to an actionable real-time risk scalar.

6.1. Risk Metrics for Real-Time Derivative Accounting: The risk of a derivative, or indeed any financial position, may be defined as the negative incremental effect of unit exposure shift on the profit & loss of the position; for a long Call option the loss from an increase in the underlying matches the profit from a decrease, so in this case risk is defined as the profit from a decrease but the principle is general and can be expressed mathematically. In a multi-counterparty scenario with an individual dealer or other transaction counterparty, such real-time risk can be formulated for each counterparty. However, when the derivative is part of a portfolio involving another multilateral settlement layer, supported with proper signalling and information, the risk formulation must reflect a portfolio view consistent with the bilateral valuations. Such complexity demands dynamic expression reconciliation as a separate concern. It is at this point that parallel variant evaluation becomes difficult to organise optimally. In any event the cross-sections of risk-important variables underlying the derivatives must be monitored for consistency with market-based covariance and correlation information over the full forecast horizon; any emergence of new market-based covariance information adding public risk sources, or other changes in perceived risk, must naturally trigger a matching reaction in the risk-aware evaluation of all relevant derivatives.

Risk evaluation for a market-facing position is, moreover, very much a multi-source risk evaluation requirement. The very notion of a market view – of which the central category for derivative management is credit spread – is one blended across many diverse sources. Convergence prediction for a mean-reverting process provides the simplest model; in practice the empirical evidence that spreads trade within bands offers greater robustness both in calibration and application. Other drivers are also naturally blended live – for example the underlying asset drift; drivers of cross-asset correlation – particularly the direction of correlation for foreign exchange – but many other market-based correlation parameters are worth monitoring and simple band tests provide an easy consistency check.

6.2. Stress Testing and Scenario Analysis: Stress testing evaluates the financial system's resilience and risk management processes, aiming to ensure stability even under extreme adverse conditions. Stress tests typically undergo scenario analysis based on forecasts or historical estimates, utilizing proprietary taxonomies and realistic criteria. The scenario generator addresses econometric data availability and identifies suitable algorithms for predicting time series likelihoods. Risk-aware systems may exploit such synthetic probability distributions to anticipate reaction-based scenarios and conduct low-complexity parallel simulation analyses.

The data pipeline facilitates the generation of risk metrics for market positions in a state of inertia, aligning with the valence requirements of derivatives revaluation. Scenarios derived from the scenario generator's specifications either



modify specific latent variables, inducing direct effects, or emulate macroeconomic perspectives that exert indirect influences across all latent factors and triggers. To ensure risk-aware management systems remain in supine stance latency, a key requirement for the real-time calculations employed by the majority of AI architecture types, the entire signalling capability, encapsulated as a data flow, supports both stress-testing needs and the basic services of signalling risk-management processes. Given the architectural support for parallel simulation analysis, none of the responsible entities need apply creative or DJ-Motley operational roles.

Moreover, creative stress-testing scenario analyses can take place during low-activity periods of DJ-Motley-type architectures that need not maintain their signalling elements in a state of inertia. The scenario generator may either hint at potential innovations in existing reaction-based-scene models or take user-defined directions. The provisions made within data management guarantee a continuous feed of the most recent historical data used for the creation of all newly required scenario analyses. These corner-stone features of risk-aware AI architectures guarantee the possibility of burden-free compliance with current stress-testing regulations and ad-hoc requests from entities responsible for the complete stability of the financial system.

6.3. Regulatory Reporting and Auditability: Government and industry agencies impose various disclosure and reporting requirements to address the impact of derivative portfolios on financial stability. This trend, which initially focused on large financial institutions deemed systemically important, is also affecting smaller market players due to the detection of hidden risks arising from a lack of transparency. Consequently, regulators are extending the scope of the Basel III framework by defining additional capital requirements for banks with large over-the-counter derivative positions, stressing the importance of mitigating procyclical effects.

Real-time derivative accounting, by enabling timely provisioning of required metrics, supports regulatory compliance by addressing risk assessment but may still fall short for complex portfolios. Non-trivial positions, for instance, require more than just delta and vega capital charges for calculation of risk metrics. Moreover, whether conditions can signal market movements beyond the range predicted by the models needs to be investigated to achieve comprehensive risk awareness. Sufficient interrogation and explanation of risk metrics allow sounder and earlier decision support for traders and risk managers. Stress tests and scenario analyses, though typically run intermittently, can also be incorporated into a borderline-vulnerability module to enable additional appropriate provisioning.

Category	Metric	Equation / definition
Sensitivities	Delta, Vega, etc.	$\partial V / \partial x_i$
Covariance risk	Factor-variance P&L	$\nabla V \Sigma \nabla V^T$
Tail risk	VaR_α	$\inf\{l: \mathbb{P}(\text{Loss} \leq l) \geq \alpha\}$
Tail risk	ES_α	$\mathbb{E}[\text{Loss} \mid \text{Loss} \geq \text{VaR}_\alpha]$
Counterparty risk	EPE/PFE	See above
Counterparty risk	CVA	integral form above
Stress testing	scenario s	revalue $V^{(s)}$ under shocked $\mathbf{x}^{(s)}$

Table: “Risk metrics” menu consistent with the paper’s section 6 framing

VII. IMPLEMENTATION AND EVALUATION

The multi-faceted systems integration of risk-aware real-time derivative accounting in a multi-counterparty environment based on Green AI has been designed, primarily through theoretical considerations and architectural design, and provides a foundation for experimental implementation. Empirical evaluation requires a Multi-Platform and a Multi-Brokered setup, based on diverse technologies and languages, to assess real-world latency conditions, together with a separate risk-aware version of Broker1, which incorporates regulation-mandated, but inter-coupled, risk metrics normally computed at the end of the day. These risk-aware architectures are examined within the Multi-Platform but without actual latency consideration, since this is only relevant as a test of speed in the Green Computing context.

Different systems, technologies, programming languages, tools, capabilities, and resource consumption are position benchmarks to operationalize experimental support for the use case and reinforce the analysis by deploying supporting testing of the Green AI integration frame. Latency is of special interest in considering the Green AI Quality Dimension of Energy Utilization because any increase in elapsed time or resource usage reduces Energy Efficiency, but using a multi-brokered setup with involved latency introduces the additional challenge of maintaining Risk Utility. Consequently,



the stress-testing of Broker1 required for controlling the risk-aware aspect of Broker1 is conducted separately within the Multi-Brokered setup, where more precise definitions are possible, while the Data-Driven Architecture plays its part within the Multi-Platform.

7.1. Experimental Setup and Benchmarks: The architectural design is implemented using a cloud computing provider that supports serverless data pipelines. Multiple services run in the Amazon Web Services (AWS) environment, including an Amazon S3 storage account and the AWS Lambda service that executes pipeline components. For different pieces of functionality, an internal support service written in Python with the bottle framework serves as an API backend, while the Lambda execution environment provides the task-specific functions. The underlying database system used for persistent storage of asset prices is Amazon Aurora with MySQL compatibility. Pipelines for standard classic computing tasks, such as data provenance and quality control, are set up to run regularly.

The experimental benchmark incorporates two applications: one that focuses on classic Green ICT for economic computing and one that emphasizes economic computing for risk-aware operation in the Green AI paradigm. The first application generates a standard benchmark that qualifies the deployed architecture space for classic economic metrics, such as economic efficiency or Green metrics, versus conventional alternative implementations. Following this qualification, the second application tests the support service as an economic computing segment for enhancing intelligent risk-aware capabilities. In line with the operational objectives, these segment capabilities are triggered and benchmarked externally by an arbitrary external requester.




Fig 4: Green AI: Creating a Sustainable Artificial Intelligence

7.2. Green Metrics and Resource Utilization: Experimental analysis of the influence of different resource provisioning strategies on the overall resource consumption of AI-based architectures is still ongoing. Measurements are being recorded from an experimental setup operating in a public cloud datacenter using an architecture performing risk-aware real-time derivative accounting in a multi-counterparty environment. Currently, the Greenness of the architecture is quantified according to the Power Utilization Efficiency (PUE) and Carbon Usage Effectiveness (CUE) metrics of the underlying datacenter infrastructure. These metrics are necessary, yet not sufficient, for completeness. The Green AI definition identifies a rich set of relevant objectives. For example, the design-inherent support for Turing-complete model execution allows control over distinct aspects of resource provisioning separately, such as sufficient memory bandwidth versus sufficient network bandwidth for data movement-intensive workloads, and adaptation of the workload resource requirements at runtime based on workload characteristics. Such decoupling can provide significant energy savings.

7.3. Performance, Accuracy, and Robustness: The integration of computing and telecommunications technologies has established cloud computing as one of the leading computing paradigms for delivering and accessing IT resources. For Data Centre Networks (DCNs) cloud resources can be allocated dynamically in a flexible and convenient manner according to the actual demand, becoming a better energy consumption solution. However, at a global scale, the energy consumption and greenhouse gas emissions of DCNs are increasing year by year, leading to concerns about the sustainability of cloud computing. Therefore, Green Clouds have been proposed for the efficient provisioning of resources in terms of energy, cost, and performance. While at the service level, the concepts of sustainability of the cloud services have received much attention in recent years, some work in literature discusses Green Clouds focusing on parameters of latency instead of energy efficiency. Towards this goal in this section the emphasis is on measuring the energy sustainability of the overall cloud services, based on their latency level, often neglected in terms of market development.

Accuracy and robustness are important aspects of any machine learning system care about model performance. Accuracy is defined as a measure of correctness of predictions made by the model while robustness is a measure of model's sensitivity with respect to the noise present in the training data. In the case of ML-based systems, the noise can be defined



as any perturbation in the training data such for example in domain adaptation when the data distributions of the training and test domains are not exactly the same but close.

VIII. CONCLUSION

The proposals for risk-aware real-time accounting of derivative portfolios with respect to regulatory and compliance requirements in the presence of multiple counterparties fit well into a Green AI architecture. Research on the three facets of sustainable computing illustrates the connection—these facets extend naturally to eco-friendly computing. Risk-aware real-time accounting of derivative portfolios provides a rich set of metrics that permits environment-related testing and evaluative mission compliance reporting.

A wide domain of problems and business methods may derive unique requirements for research and management that likely expand dynamically decades into the future beyond the currently planned business missions. The capability for appropriate stress testing and scenario analysis throughout the planning decade is just one element of the total mission support environment. In addition to concept verification, such investigations indicate the fundamental architecture and any requirements for specific improvement.

8.1. Future Trends: Ongoing research is seeking to shape risk-aware AI architectures for other financial operations. Approaches to stock-and-flow accounting, both for physical deliveries and for liquidities, depend on classically established systems, but concentrates on their real-time deployment within the context of multi-IA systems. Similar reasoning is used in prototype implementation that shows how the impact of paraphed products on a clearing authority's banks can be properly reflected without direct involvement of the authority itself.

As risk-awareness becomes greater for financial instruments with lesser risk-understanding than commonly used derivatives (stock futures and option spreads), we can expect further symbology-based models to be built into the modelling of real-time derivative valuations. The sympathy of an operation's model with that of the position tracked by the bank performing the operation should limit any potential communication or data-adjustment latency involved in incorporating that model into the risk-aware processing. In further enabling support for liquidity and cash-flow modelling, attention is being accorded to address-space definitions for agent banks' and banks' operational models in risk-aware architectures for multi-scale transaction processing.

REFERENCES

1. Polineni, T. N. S., Teja Ganti, V. K. A., Maguluri, K. K., & Rani, P. R. (2024). AI-Driven Analysis of Lifestyle Patterns for Early Detection of Metabolic Disorders. *Journal of Computational Analysis & Applications*, 33(8).
2. Polineni, T. N. S., Kumar, A. S., Maguluri, K. K., Koli, V., Valiki, D., & Ravikanth, S. (2024, November). A Scalable and Robust Framework for Advanced Semi Supervised Learning Supporting Universal Applications. In *Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024)*.
3. Pamisetty, V. (2024). AI-Driven Decision Support for Taxation and Unclaimed Property Management: Enhancing Efficiency through Big Data and Cloud Integration. Available at SSRN 5250776.
4. Kalisetty, S., Pandugula, C., Sondinti, L. R. K., Mallesham, G., & Rani, P. S. (2024). AI-driven fraud detection systems: enhancing security in card-based transactions using real-time analytics. *Journal of Electrical Systems*, 20(11s).
5. Shukla, A., Dubey, S., Nithya, P., Shankar, B., Vankayalapati, R. K., & Khatana, K. (2024, November). Edge-optimized and explainable deep learning framework for real-time intrusion detection in industrial IoT. In *Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024)*.
6. Ramanakar Reddy Danda, Z. Y., Mandala, G., & Maguluri, K. K. (2024). Smart Medicine: The Role of Artificial Intelligence and Machine Learning in Next-Generation Healthcare Innovation.
7. Pandugula, C., Ganti, V. K. A. T., & Mallesham, G. (2024). Predictive Modeling in Assessing the Efficacy of Precision Medicine Protocols. EDUCATIONAL ADMINISTRATION: THEORY AND PRACTICE Учредители: Green Publication БИБЛИОМЕТРИЧЕСКИЕ ПОКАЗАТЕЛИ: Входит в РИНЦ; на рассмотрении Цитирований в РИНЦ: 0 Входит в ядро РИНЦ; нет Цитирований из ядра РИНЦ; 0 Рецензии: нет данных Процентиль журнала в рейтинге SI: ТЕМАТИЧЕСКИЕ НАПРАВЛЕНИЯ:.
8. Kalisetty, S. (2024). Agentic AI and predictive analytics: Revolutionizing retail supply chain management for next-gen resilience and efficiency. Available at SSRN 5231377.
9. Danda, R. R., Nampalli, R. C. R., Sondinti, L. R. K., Vankayalapati, R. K., Syed, S., Maguluri, K. K., & Yasmeen, Z. (2024). Harnessing Big Data and AI in Cloud-Powered Financial Decision-Making for Automotive and Healthcare Industries: A Comparative Analysis of Risk Management and Profit Optimization.



10. Kalisetty, S. (2024). Revolutionizing Retail Operations: Srinivas's AI-Driven Approach to Demand Forecasting and Seamless Supply Chain Management. Available at SSRN 5203708.
11. Venkatasubramanian, K., Yasmeen, Z., Reddy Kothapalli Sondinti, L., Valiki, S., Tejpal, S., & Paulraj, K. (2024, November). Unified deep learning framework integrating CNNs and vision transformers for efficient and scalable solutions. In Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024).
12. Nampalli, R. C. R., & Syed, S. (2024, November). Planet 2050 and Big Vehicle Manufacturing: Leveraging Data Analytics for Zero Emissions and Sustainable Manufacturing Practices. In Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024).
13. Syed, S. (2024). Transforming Manufacturing Plants for Heavy Vehicles: How Data Analytics Supports Planet 2050's Sustainable Vision. *Nanotechnology Perceptions*, 20(6), 10-62441.
14. Vijayarajeswari, R., Pushpa, N., Chandra Rao Nampalli, R., Koli, V., & Murali, E. (2024, November). Optimized Deep Learning Techniques for Scalable Transparent and Real-Time Secure IoT Networks. In Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024).
15. Ganti, V. K. A. T. (2024). Designing Bias-Aware AI Models to Address Health Disparities in Multi-Ethnic Healthcare Systems. *American Journal of Analytics and Artificial Intelligence (ajaii)* with ISSN 3067-283X, 2(1).
16. Vankayalapati, R. K., Yasmeen, Z., Bansal, A., Dileep, V., & Abhireddy, N. (2024, December). Advanced fault detection in semiconductor manufacturing processes using improved AdaBoost RT model. In 2024 9th International Conference on Communication and Electronics Systems (ICCES) (pp. 467-472). IEEE.
17. Maguluri, K. K., Pandugula, C., & Yasmeen, Z. (2024). Neural Network Approaches for Real-Time Detection of Cardiovascular Abnormalities.
18. Nampalli, R. C. R., Syed, S., Bansal, A., Vankayalapati, R. K., & Danda, R. R. (2024, December). Optimizing Automotive Manufacturing Supply Chains with Linear Support Vector Machines. In 2024 9th International Conference on Communication and Electronics Systems (ICCES) (pp. 574-579). IEEE.
19. Pandugula, C., Ganti, V. K. A. T., & Malleshm, G. (2024). Predictive Modeling in Assessing the Efficacy of Precision Medicine Protocols. *EDUCATIONAL ADMINISTRATION: THEORY AND PRACTICE* Учредители: Green Publication БИБЛИОМЕТРИЧЕСКИЕ ПОКАЗАТЕЛИ: Входит в РИНЦ: на рассмотрении Цитирований в РИНЦ: 0 Входит в ядро РИНЦ: нет Цитирований из ядра РИНЦ: 0 Рецензии: нет данных Процентиль журнала в рейтинге SI: ТЕМАТИЧЕСКИЕ НАПРАВЛЕНИЯ:.
20. Syed, S., Jayalakshmi, S., Kumar Vankayalapati, R., Mandala, G., Yadav, O. P., & Yadav, A. K. (2024, December). A Robust and Scalable Deep Learning Framework for Real-Time Iot Intrusion Detection with Adaptive Energy Efficiency and Adversarial Resilience. In Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024).
21. Krishna AzithTejaGanti, V., Senthilkumar, K. P., Robinson L, T., Karunakaran, S., Pandugula, C., & Khatana, K. (2024, November). Energy-Efficient Real-Time Hybrid Deep Learning Framework for Adaptive Iot Intrusion Detection with Scalable and Dynamic Threat Mitigation. In Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024).
22. Mandala, G., Reddy, R., Nishanth, A., Yasmeen, Z., & Maguluri, K. K. (2023). Ai and ml in healthcare: redefining diagnostics, treatment, and personalized medicine. *International Journal of Applied Engineering & Technology*, 5(S6).
23. Azith Teja Ganti, V. K., Maguluri, K. K., & Sudha Rani, P. R. (2024). Neural Network Applications in Understanding Neurodegenerative Disease Progression. *Frontiers in Health Informatics*, 13(8), 471.
24. Pamisetty, A. (2024). Integration Of Artificial Intelligence And Machine Learning In National Food Service Distribution Networks.
25. Sriram, H. K., Challa, S. R., Challa, K., & ADUSUPALLI, B. (2024). Strategic Financial Growth: Strengthening Investment Management. Secure Transactions, and Risk Protection in the Digital Era. *Secure Transactions, and Risk Protection in the Digital Era* (November 10, 2024).
26. Sriram, H. K. (2024). Leveraging AI and machine learning for enhancing secure payment processing: A study on generative AI applications in real-time fraud detection and prevention. Available at SSRN 5203586.
27. Kannan, S. (2023). The Convergence of AI, Machine Learning, and Neural Networks in Precision Agriculture: Generative AI as a Catalyst for Future Food Systems. *JOURNAL FOR REATTACH THERAPY AND DEVELOPMENTAL DIVERSITIES*.
28. Fonton Moffo, A. M. (2024). A machine learning approach in stress testing US bank holding companies. *International Review of Financial Analysis*, 95, 103476.
29. Fu, J. (2024). Real-time data analytics and predictive maintenance in commercial bank credit risk management. *Frontiers in Artificial Intelligence and Applications*, 388, 115–124.
30. Moreno T., J. F., Zapata, C., & Aragon, D. (2024). Neural networks and the nonlinear Feynman–Kac theorem applied to financial options pricing. *SN Computer Science*, 5, 697.



31. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon emissions and large neural network training. *Communications of the ACM*, 64(11), 39–47.
32. Kannan, S., & Seenu, A. (2024). Advancing Sustainability Goals with AI Neural Networks: A Study on Machine Learning Integration for Resource Optimization and Environmental Impact. *management*, 32(2).
33. Kannan, S. (2024). Revolutionizing Agricultural Efficiency: Leveraging AI Neural Networks and Generative AI for Precision Farming and Sustainable Resource Management. *Utilitas Mathematica*.
34. Recharla, M., Chava, K., Chakilam, C., & Suura, S. R. (2024). Postpartum Depression: Molecular Insights and AI-Augmented Screening Techniques for Early Intervention. *International Journal of Medical Toxicology and Legal Medicine*, 27(5), 935-957.
35. Suura, S. R. (2026). Integrating Generative AI into Non-Invasive Genetic Testing: Enhancing Early Detection and Risk Assessment. *Transcriptome*, 1(1), 22-33.
36. Annapareddy, V. N. (2024). Leveraging Artificial Intelligence, Machine Learning, and Cloud-Based IT Integrations to Optimize Solar Power Systems and Renewable Energy Management. *Machine Learning, and Cloud-Based IT Integrations to Optimize Solar Power Systems and Renewable Energy Management* (December 06, 2024).
37. Srinivas Kalisetty, D. A. S. (2024). Leveraging Artificial Intelligence and Machine Learning for Predictive Bid Analysis in Supply Chain Management: A Data-Driven Approach to Optimize Procurement Strategies.
38. Shyamala Anto Mary, P., Kalisetty, S., & Mandala, V. M. (2024, November). Advancing IoT Data Forecasting with Deep Learning Framework for Resilience Scalability and Real-World Applications. In *Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024)*.
39. Kalisetty, S., Pandugula, C., Sondinti, L. R. K., Malleshham, G., & Rani, P. S. (2024). AI-driven fraud detection systems: enhancing security in card-based transactions using real-time analytics. *Journal of Electrical Systems*, 20(11s).
40. Mason, T., & Jarvis, S. (2023). Omnichannel retail: How to build winning stores in a digital world. Kogan Page Publishers.
41. Mason, T., & Jarvis, S. (2023). Omnichannel retail: How to build winning stores in a digital world. Kogan Page Publishers.
42. Vankayalapati, R. K. (2025). Zero-trust security models for cloud data analytics: Enhancing privacy in distributed systems. Available at SSRN 5121185.
43. Zhang, C., Amici, G., & Morandotti, M. (2024). Calibrating the Heston model with deep differential networks. *Quantitative Finance*, 24(12), 1789–1806.
44. Polineni, T. N. S., Teja Ganti, V. K. A., Maguluri, K. K., & Rani, P. R. (2024). AI-Driven Analysis of Lifestyle Patterns for Early Detection of Metabolic Disorders. *Journal of Computational Analysis & Applications*, 33(8).
45. Maguluri, K. K., Ganti, V. K. A. T., & Subhash, T. N. (2024). Advancing patient privacy in the era of artificial intelligence: a deep learning approach to ensuring compliance with HIPAA and addressing ethical challenges in healthcare data security. *International Journal of Medical Toxicology & Legal Medicine*, 27(5), 10-47059.
46. Danda, R. R. (2024). AI and deep learning techniques for health plan satisfaction analysis and utilization patterns in group policies. *International Journal of Medical Toxicology & Legal Medicine*.
47. Huge, B., & Savine, A. (2020). Differential machine learning: The shape of things to come. *Risk*, 33(9), 42–47.
48. Danda, R. R., Nampalli, R. C. R., Sondinti, L. R. K., Vankayalapati, R. K., Syed, S., Maguluri, K. K., & Yasmeen, Z. (2024). Harnessing Big Data and AI in Cloud-Powered Financial Decision-Making for Automotive and Healthcare Industries: A Comparative Analysis of Risk Management and Profit Optimization.
49. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. *Int. J. Innov. Sci. Res.*, 9, 593-595.
50. Andersson, K., & Oosterlee, C. W. (2021). Deep learning for CVA computations of large portfolios of financial derivatives. *Applied Mathematics and Computation*, 409, 126399.
51. Agram, N., Øksendal, B., & Rems, J. (2024). Deep learning for quadratic hedging in incomplete jump market. *Digital Finance*, 6, 463–499.
52. Ahmad, W., Tiwari, S. R., Wadhwani, A., Khan, M. A., & Bekiros, S. (2023). Financial networks and systemic risk vulnerabilities: A tale of Indian banks. *Research in International Business and Finance*, 65, 101962.
53. Bahoo, S., Cucculelli, M., Goga, X., & Mondolo, J. (2024). Artificial intelligence in finance: A comprehensive review through bibliometric and content analysis. *SN Business & Economics*, 4, 23.
54. Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance*, 19(8), 1271–1291.
55. Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. *Artificial Intelligence Review*, 57, 216.
56. Cheney, J., Chiticariu, L., & Tan, W. C. (2009). Provenance in databases: Why, how, and where. *Foundations and Trends in Databases*, 1(4), 379–474.
57. Gnoatto, A., Picarelli, A., & Reisinger, C. (2023). Deep xVA solver: A neural network-based counterparty credit risk management framework. *SIAM Journal on Financial Mathematics*, 14(1), 314–352.



57. Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, 32, 100577.
58. Hutchinson, J. M., Lo, A. W., & Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance*, 49(3), 851–889.
59. Khosravi, A., Sandoval, O. R., Taslimi, M. S., Sahrakorpi, T., Amorim, G., & Garcia Pabon, J. J. (2024). Review of energy efficiency and technological advancements in data center power systems. *Energy and Buildings*, 323, 114834.
60. Lannelongue, L., Grealey, J., & Inouye, M. (2021). Green Algorithms: Quantifying the carbon footprint of computation. *Advanced Science*, 8(12), 2100707.
61. Li, B., & Zhang, X. (2024). Systemic risk and financial networks. *The Quarterly Review of Economics and Finance*, 94, 25–36.
62. Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273.
63. Mahadevan, S. (2024). Clouding the Future: Innovating Towards Net-Zero Emissions. *International Journal of Computing and Engineering*, 6(2), 17-23.
64. Sirignano, J., & Cont, R. (2019). Universal features of price formation in financial
65. Wang, Y., Tong, L., & Zhao, Y. (2024). Revolutionizing hedge fund risk management: The power of deep learning and LSTM in hedging illiquid assets. *Journal of Risk and Financial Management*, 17(6), 224.
66. Weber, P., Carl, K. V., & Hinz, O. (2024). Applications of explainable artificial intelligence in finance—A systematic review of finance, information systems, and computer science literature. *Management Review Quarterly*, 74, 867–907.
67. Yao, J., Li, Y., & Tan, C. L. (2000). Option price forecasting using neural networks. *Omega*, 28(4), 455–466.
68. Ruf, J., & Wang, W. (2020). Neural networks for option pricing and hedging: A literature review. *Journal of Computational Finance*, 24(1), 1–46.
69. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63.
70. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.
71. Tabbakh, A., Al Amin, L., Islam, M., Mahmud, G. M. I., Chowdhury, I. K., & Mukta, M. S. H. (2024). Towards sustainable AI: A comprehensive framework for Green AI. *Discover Sustainability*, 5, 408.