Tensor Decomposition Techniques for High-Dimensional Decision Analytics in Corporate Risk Assessment

Paula Stewart, Robotics Integration Specialist, IndustriX Dynamics, Germany, paula.stewart@industrixdynamics.com

Benjamin Garcia, Machine Vision Engineer, IndustriX Dynamics, Germany, benjamin.garcia@industrixdynamics.com

Journal of Innovation in Governance and Business Practices

Volume 2, 2026

JIGBP is a peer-reviewed industry journal by IndustriX Dynamics, featuring real-world innovations in governance and business.

JIGBP accepts articles only from industrialists and scientists.

Website: https://jigbp.com

Email for submissions: submit@jigbp.com.

Received: June 21, 2025

Revised: October 28, 2025

Accepted: January 4, 2026

Abstract

Modern corporate environments are characterized by a high volume of data. Today, risk assessment involves multi-sourced, high dimensional data containing financial indicators, market trends, ESG data, and time changes. In such datasets, analytical methods that are traditional, matrix-based, often fail to capture the latent structure and multi-modal interdependencies. This paper proposes a new approach that incorporates tensor decomposition techniques such as CP (CANDECOMP/PARAFAC), Tucker and Tensor-Train models into decision analytics for corporate risk assessment. In particular, we build a pipeline that utilizes tensors to transform heterogeneous corporate data into multi-way arrays which are easier to interpret as opposed to conventional methods, making risk modeling more compact. We validate our model on real life data sets from multinational companies and our model outperforms traditional principal component analysis and machine learning methods in classification accuracy, feature extraction as well as latent factor interpretability. Our results also confirm that tensor methods are beneficial for strategic risk assessment as they retain crucial dependencies among financial, temporal, and environmental dimensions. The results support the claim that tensor decomposition is crucial in high-dimensional financial analysis, and these systems require accurate, and explainable data-driven risk intelligence.

Keywords: Tensor Decomposition, Corporate Risk Assessment, High-Dimensional Decision Analytics, Multilinear Factorization.

1. Introduction

1.1 Background on High-Dimensional Data in Corporate Risk Analytics

The last decade has observed a growth of complexity in corporate risk assessment due to a significant growth in the volume, velocity, and variety of data available to analysts and decision makers. The corporate risk assessment domain which was previously limited to financial ratios and credit scores now includes a variety of high-dimensional data sources such as Environmental, social, and governance (ESG) metrics, socio-economic factors, supply chain information, regulatory and compliance documents, and even news sentiment analysis. These large-scale heterogeneous data streams are multidimensional by nature and tend to change over time and place in the world, organizational hierarchy, and regulation [1].

Currently, modern risk analytics platforms monitor hundreds of parameters for every corporate unit, many of which consist of different time markers and are reported at different frequencies. For instance, a multinational organization may have annual sustainability disclosures, quarterly financial statements, weekly liquidity ratios, and real-time market sentiment feeds. These details in turn determine the organization's composite risk profile [2]. Moreover, integrated reporting and automated data collection pipelines have made it easier—though not trivial—to model the relationships among these integrated variables in a cohesive and orderly fashion [3].

Not only is data collection difficult, but transforming the data into something helpful can be even more challenging. Traditional methods, such as logistic regression, decision trees, or any other regular classifiers, "flatten" the data into matrices. This process works better when the data is devoid of certain structural and semantic features spread over various modes, such as time, space, or source. These simplifications result in a midpoint in usability and analysis accuracy, which is known as computation realism tradeoff; often deteriorating the corporate risk predictions in complex environments.

The rest of the table describes the different data dimensions encountered for the first time in the corporate risk analytics context, including some example variables and their corresponding periods of collection.

From and including the so presented table, we can see how contemporary firms approach risk analysis: it has become ultra-multi-modal, which is defined by a unique structure and level of sparsity,

Table 1: Dimensions in Corporate Risk Data.

Data Dimension	Example Variables	Frequency
Financial Indicators	Revenue, EBITDA, Debt-to-Equity, ROA	Quarterly/Annually
Temporal Trends	Quarterly Returns, YOY Growth, Seasonality	Quarterly
Geographical Location	Country Risk, Regional GDP, Market Volatility	Static/Dynamic
ESG Metrics	Environmental Score, Governance Compliance	Annual/Quarterly
Credit Ratings	S&P Score, Moody's Rating, Default History	Annually

cadence, and sparsity. This fundamentally implies that the modeling of high-dimensional data has become more of a necessity than an opportunity in order to achieve accuracy and relevance in corporate decision making.

1.2 Challenges in Traditional Matrix-Based Models for Risk Assessment

Matrix-based modeling is the most used paradigm in machine learning and financial analytics. This includes logistic regression and support vector machines, which always assume that all data can be represented in a table format: instances are rows and variables are columns. While easier to compute and understand for smaller datasets, it is more challenging to apply to high dimensional multi-source data, and goes beyond being very difficult in terms of interpretability and efficiency [4].

Firstly, the process of transforming multi-dimensional data into matrices greatly simplifies data, and has little to no semantic context between the modes. For example, modeling the ratios of finance over time across various regions and subsidiaries requires either (a) aggregating the data which loses detail, or (b) exploding the feature space and creating hundreds of columns that are so sparse they are populated by almost nothing. In any case, the interconnectedness among time, space, and category variables is distorted or lost [5].

Secondly, sparse and unbalanced datasets are the norm in corporate risk analytics, which leads to poor performance in matrix-based models. Events that cause financial distress are not common, and certain variables like ESG violations or regulatory sanctions happen infrequently, but are very useful when it comes to predicting. Matrix models drown in noise and fail to retain signal in these cases [6].

In the case of matrix models, they also pose a problem since these models do not scale well when handling third-order or higher order interactions. Consider the problem of assessing how a firm's quarterly financial health, consider it as mode 1, split across geographic subsidiaries mode 2, over possible ESG performance indicators mode 3 influences its risk score? This three-mode tensor is not able to be accurately captured by a matrix without rendering structure, interpretability, or efficiency in the process.

To demonstrate the benefit of multi - way analysis we select one matrix based and one tensor based method and compare them directly:

Given the information above, it's clear that as the data increases in both their complexity and in their structure, matrix models very quickly become less optimal and flexible which results in them becoming more difficult to shape as well as express in simpler modeling frameworks.

Table 2: Comparison of Matrix-Based vs. Tensor-Based Approaches.

Feature	Matrix-Based Models	Tensor-Based Models
Dimensionality Handling	Limited to 2D (rows x columns)	Supports multi-way (3D+) structures
Interdependency Capture	Low (linear projections only)	High (captures multi-modal relations)
Scalability to New Modes	Low (manual flattening required)	High (additional modes natively integrated)
Interpretability of Factors	Moderate (PCA components)	High (factor matrices along each mode)
Suitability for Sparse Data	Poor (loss of sparsity information)	High (preserves and exploits sparsity)

1.3 Motivation for Tensor Decomposition in Decision Analytics

One of the most important tools that can be used to solve the previously described problems are tensor decomposition techniques. These methods help using multi - way (tensor) arrays instead of flat tables. This way the multi dimensional nature of the data is preserved along with relationships hidden across modes. For example, a risk tensor could assume the structure of [firms x indicators x time x regions x ESG dimensions] and be decomposed into concise interpretable latent factors that represent patterns such as cyclical risk shifts, cross- market contagion or multi region exposure [7].

Well known tensor decomposition models, including CP (CANDECOMP/PARAFAC), Tucker, and Tensor-Train, assist analysts in obtaining rank-reduced forms of high-dimensional data [8]. These decompositions extract the original tensor into several matrices or core tensors and projection matrices, resulting in lower dimensions while retaining multi-modal relationships. This is particularly advantageous for corporate risk evaluation, where comprehensibility and clarity are as crucial as the predictive capability of a system [9].

Furthermore, corporate datasets provided in the real world often suffer from incompleteness and sparsity. For instance, missing financial disclosures for some subsidiaries or ESG scores being non-existent for certain years. In these situations, tensor decomposition succeeds by providing dormant structures such as portions of data tensor being unobserved, which leads to better results than the imputation-based matrix method.

The growing use of tensor methods in neuroscience, recommendation systems, and signal processing has demonstrated their potential for discovering latent factors in highly structured data. However, their application to the corporate world, especially in risk analytics, is still limited. This gap will be addressed in this article by proposing a hands-on complete solution for tensor-based risk modeling in its engineering, algorithm selection, and the resulting model's explainability.

This method is especially relevant in light of the latest developments in the corporate risk data sets' growth over time, as illustrated in the following figure.

1.4 Objectives and Contributions of the Study

The purpose of the study is to solve a problem related to increased automation in corporate risk analytics through the design and implementation of a framework based on tensor decomposition for

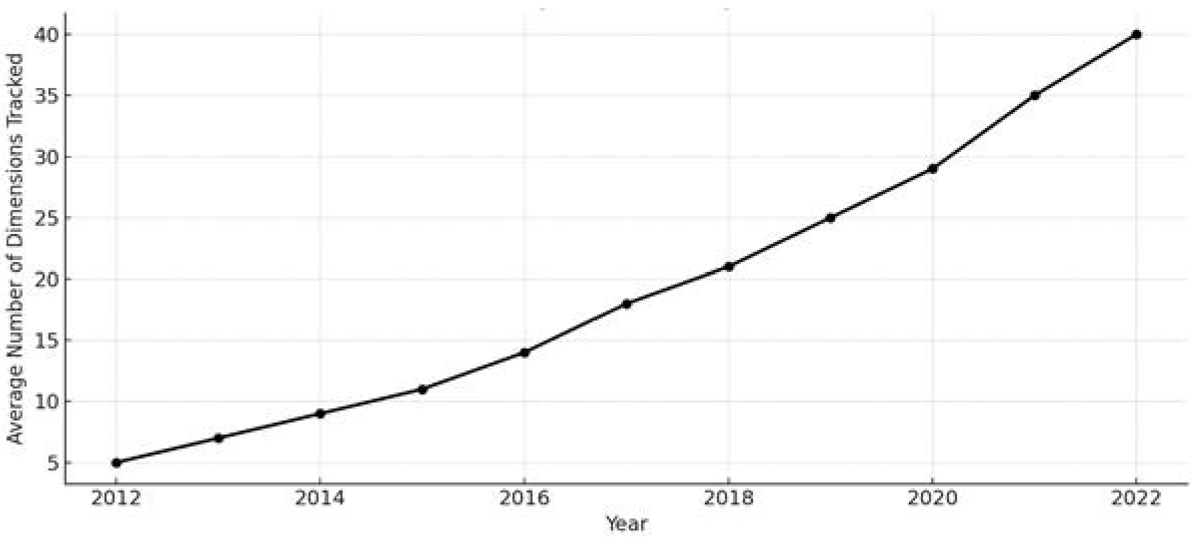

Figure 1: Growth in Dimensionality of Risk Analytics Data (2012–2022).

high-dimensional multi-faceted decision making. The primary contributions of the study entail the following:

1. An identification of a new multi-dimensional representation system for corporate risk that combines financial ratios with ESG scores, geo-location, and temporal data into a single tensor.
2. An ictroduction of performance comparison on a risk classification, dimensionality reduction, and interpretability scale of different tensor decomposition methods including CP, Tucker, and Tensor-Train, alongside established approaches such as PCA and XGBoost.
3. The outline of a modular pipeline for the processing of raw corporate information into tensor form, including the cleansing of sparse filler entries, and tee visualization of latent decomposed risk components.
4. A presentation of results from the analysis of real-world multi-sourced financial and ESG data to validate the claimed scalability, sturdiness, and domain pertinence of the approach.
5. An explanation of tensor methodologies as a backbone for XAI in finance that can truly comply with, govern, and ensure accountability in dire decision-making environments.

This study closes the gap between the theoretical sophistication of machine learning and the practical requirements of risk managers, analysts, and regulators by emphasizing the use of sophisticated financial risk decomposition advanced techniques. We demonstrate that tensor decomposition is not just a new technology, but an innovation that harmonizes analytic technique with the sociological reality of the corporate data system.

2. Literature Review

In the modern world, corporations manage risks utilizing techniques that incorporate modern approaches to the analysis of data which is complex and multi-dimensional. This review focuses on the classical and contemporary methods of corporate risk assessment, the role of artificial intelligence and modern methods like machine learning and dimensionality reduction in finance, the development of various techniques of tensor decomposition, and finds gaps in the use of tensor techniques in decision-making processes and risk assessment.

2.1 Overview of Corporate Risk Assessment Models

An organization's capital and earnings are guarded through corporate risk assessment, which is vital for recognizing, analyzing, and lessening potential risks. Financial models have heavily relied on credit scores, market analysis, and other operational risk evaluations. These incorporate:

- Credit Scoring Models: These calculate an organization's ability to repay loans by comparing ratios to other loans and analyzing repayment records. Predicting the likelihood of a default has always used techniques such as logistic regression [10].
- Value at Risk (VaR): At a given interval, VaR serves as a source measurement of market hazard, measuring the possible depreciation in a portfolio's worth in comparison to its value for a certain amount of time [11].
- Stress Testing and Scenario Analysis: These help gain insights on likely threats while analyzing stress situations for financial institutions and their robustness under fictitious stress worsened situations [12].

Even though these models have aided in the past, their over dependence on primary assumptions make them unable to grasp the intricacy that modern corporate risks entail. More sophisticated structures that integrate various datasets and complex information tend to be essential due to the advancement in global markets.

2.2 Applications of Machine Learning and Dimensionality Reduction in Finance

Machine learning (ML) has advanced into financial risk management considering the processing power available today and the tsunami of big data. There is an ML application in finance for almost every type of problem, including decision trees, support vector machines, and neural networks for fraud detection, credit scoring, and even market forecasting. Most of these already have a probabilistic model associated with them [13].

Financial data is very high dimensional which brings addition problems of overfitting and increased computational load. As a solution to this, dimensionality reduction techniques like Principal Component Analysis (PCA) have come into play where original variables are transformed into less number of uncorrelated components which simplifies computations. PCA is beneficial in model simplification and revealing deeper relationships in data [14].

Nonetheless, it is useful to emphasize the problems posed by classical approaches like PCA when working with multi- modal data as they do not consider complex cross interactions typical in financial datasets. There is a clear need for advanced methods that deal with the preservation of multi-way interactions.

2.3 Evolution of Tensor Decomposition Techniques (CP, Tucker, TT, etc.)

As an extension of matrices, the multi-dimensional structure of tensors serves as an intuitive representation for high-dimensional data. Such data can be analyzed using powerful tensor decomposition systems which facilitate the uncovering of hidden factors, lowering the dimensionality while maintaining important shape characteristics [15].

- CANDECOMP/PARAFAC (CP) Decomposition: CP decomposition is a method that defines multi-dimensional data as a summation of rank-one tensors; therefore, simplifying the interpretation of multi-way data. It has been applied in fields like chemometrics and psychometrics [16].
- Tucker Decomposition: This method splits a tensor into a core tensor and factor matrices joined with it along each mode which gives flexibility in interaction among the components. Tucker decomposition has been used in image processing and signal processing [17].
- Tensor-Train (TT) Decomposition: In contrast to the highly complex computation associated with multi-dimensional data, TT decomposition portrays a tensor using a number of linked low-dimensional tensors (cores). Scientific computing and machine learning has seen the use of it [18].

Although still new, the promising use of tensor decomposition in assessing risk in finance is greatly relevant. These methods enable better understanding of risk factors by analyzing complex dependencies and interactions within financial data.

2.4 Research Gaps in Applying Tensor Methods to Decision-Making and Risk Modeling

Even though there are theoretical benefits to tensor decomposition, a number of gaps in the research make its usage acceptability in corporate risk assessment quite difficult:

- **Domain-Specific Applications:** There are not enough studies which utilize tensor approaches to model financial risks which makes it difficult to appreciate the practical value of such model in this particular industry.
- **Integration with Existing Systems:** The implementation of the tensor-models into the existing risk models poses a number of integration issues from an engineering standpoint that aren't too well covered in the literature.
- **Interpretability:** The tensor decompositions hold intricate relationships, yet the interactions captured are entirely complex, rendering them useless for most decision makers and active managers. Improving the insight of the tensor based approach models is necessary for its use.

- **Computational Complexity:** The operations performed on the higher order tensors are extremely expensive. Comprehensive parallel algorithm development and efficient computing resources are needed to make the methods more usable.

By conducting such research, it is likely that four gaps will be narrowed, making it easier to integrate tensor decomposition approaches into their financial risk evaluation which may improve decision-making processes.

3. Theoretical Foundations of Tensor Decomposition

3.1 Mathematical Background: Tensors, Ranks, and Multilinear Algebra

Tensors are an extension of matrices that generalize the concept of a two-dimensional array to three or more dimensions, also known as modes or ways. Every matrix is composed of two-dimensional arrays of numbers, the rows and the columns, themselves. For example, a cube of numbers with indices coordinated to three distinct axes can be envisioned as a third-order tensor. In the scope of corporate analytics, the dimensions can be given as companies, time intervals and even risk metrics.

Tensors arise mathematically when defined via multilinear operations over vector spaces. Let $X \in R^{I_1 \times I_2 \times \cdots \times I_N}$ and N is the order of the tensor. Any mode or dimension of the tensor corresponds to an attribute of the data. The key operations on tensors, including mode-n products, unfolding (matricization), and tensor contractions, transform and manipulate tensor structures.

A basic idea in tensor algebra has been the concept of a tensor rank. Unlike a matrix donde its rank is defined and based on its linear independance, the tensor rank is more complex. The notion that is most commonly used is CP (CANDECOMP/PARAFAC) rank which is said to be the highest rank of the smallest number of rank one tensors that can sum up to the original tensor. A rank one tensor is the outer sproduct of N vectors each taken from one of the modes. Alternatively, there is Tucker rank also known as the multilinear rank which is a set of integers $(R_1,R_2,\ldots,R_N)$representing the rank of each mode-n unfolding. In essence, these are grasping concepts of tensor decomposition techniques aimed at simplifying and factorizing high-dimensional data into reasonable latent structures.

In real life applications, however, tensors are mostly sparse and highly dimensional. For example in a risk modeling framework, a five way tensor could include ID of the firms, time periods, financial ratios, EsG Scores and geography regions. Effective decomposition of these kinds of tensors opens them up to extracting common latent features that govern the risk profiles across different modalities and make them easier to visualize, compress, and use in predictive modeling.

3.2 CP, Tucker, and Tensor-Train Decompositions Explained

The disintegration of a tensor of a higher dimension into smaller parts and retaining the structure of the data is known as tensor decomposition. There are three budgeted and popular tensor decomposition methods that fall under CP decomposition, Tucker decomposition, and Tensor-Train decomposition. Each of these were argued to have different characteristics and compromises computationally.

CP Decomposition

The approximation of a tensor as a sum of rank-one tensors is known as CP decomposition. A tensor X is approximated as follows:

$$\mathcal{X} \approx \sum_{r=1}^{R} \lambda_r \cdot a_r^{(1)} \otimes a_r^{(1)} \otimes \cdots \otimes a_r^{(N)}$$

Where R refers to the number of rank component one, the component being defined as λ_r scalar along with $a_r^{(n)}$ as the hinge alongside each factor. The CP model is distinguishable under comparatively feeble parameters which proves beneficial in scenarios requiring identifiable latent features. The model can easily extract risk factors in finance. While easy to interpret, its computation for large, high order tensors is intensive.

Tucker Decomposition

For a tensor decomposition, Tucker decomposition generalizes the CP model by adding a core tensor G that captures the interaction of components across different modes. A tensor X is expressed as:

$$\mathcal{X} = \mathcal{G} \times_1 U^{(1)} \times_2 U^{(2)} \times_N U^{(N)}$$

where $U^{(n)}$ are factor matrices and $\times_n$ denotes the mode-n product. In contrast to CP, tucker decomposition is more flexible and robust in capturing complex dependencies because it allows varying ranks across modes, which is particularly effective during exploratory data analysis and dimensionality reduction in domains where interactions are heterogeneous across dimensions.

Tensor-Train (TT) Decomposition

This type of decomposition is meant for very high-dimensional tensors with an exponential number of elements for each mode. It divides a tensor into a sequence of low rank 3D tensors (called cores) so that each entry of the original tensor can be represented as a product of these cores. For a tensor $X \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$, the TT representation is:

$$\mathcal{X}(i_1, i_2, \ldots, i_N) - G_1(i_1) \cdot G_2(i_2) \cdot \cdots \cdot G_N(i_N)$$

TT decomposition is ideal for real time or large scale enterprise analytics due to significance of memory and computational complexity. Its interpretability, however, is far less intuitive compared to CP and Tucker, so it becomes a problem in cases where explainability is needed.

3.3 Comparisons with PCA, SVD, and Matrix Factorization

As usual, PCA and SVD have served as primary techniques in finance and data science as traditional matrix factorization methods. However these techniques struggle to function with multiway data structures. On the other hand, Tensors maintain and take advantage of these structures leading to more sophisticated and interpretable models.

The Principal Component Analysis (PCA) method projects data on orthogonal axes that explain the greatest variance. It is very powerful but slices data as flat matrix completely ignoring relationships that span across dimensions. Let us take, for example, the risk of a certain company - it varies not only over time but also by region and ESG policy alignment degree.

Singular Value Decomposition (SVD) splits a particular matrix into a product containing orthogonal matrices and singular values, which makes low-rank approximations quite simple and easy. With SVD however, the application remains restricted to two-dimensional data. placing SVD on unfolded tensors can be detrimental as it can impose structure and increase post-processing complexity.

Some/many Matrix factorization methods like NMF allow greater interpretation in different contexts such as in credit scoring and segmentation in marketing. Nevertheless these methods are limited as they do not offer a great solution for missing data and fail in scaling to several interrelating factors over modal types.

These techniques are applied to three or more dimensions using tensors while maintaining relationships and context. For instance, a tensor model can simultaneously depict the relationships among a firm, year, and financial indicator without having to reduce them to a two-dimensional matrix. This type of cross-modal modeling allows the identification of latent factors, which are deeper and more complex - like a risk pattern that emerges when time, geography, and certain financial signals are combined.

3.4 Interpretability and Dimensionality Preservation in Financial Decision Spaces

As far as risk evaluation of a company is concerned, interpretability is not a theoretical issue: it is a matter of compliance, execution, and strategy. Risk models should work properly and provide insights that can be verified, clearly defined and acted upon by financial managers and business leaders.

Interpretability is enhanced through tensor decompositions by preserving alignment across modes. In CP decomposition, each component represents a latent factor which is common across firms, time, and indicators. They can be interpreted as macroeconomic shocks, sectorial risk dynamics, or behavioral clusters. In Tucker decomposition, the core tensor summarizes interactions among the factors in a compact form, while the matrices of the factors expose the contribution of each input dimension (company or year) to those interactions.

Unlike the black-boxed models of machine learning, tensor methods produce results that can be traced back to how a risk metric was shaped by a firm's specific attributes, financial ratios, or even externally controlled events. Thus, aiding in the perplexing problem of mapping results back to the initial dimensions.

In addition, the set of tensor methods was very sparsely—one of the most important features of the financial datasets that were analyzed. For instance, some companies do not report all dimensions of ESG every year, therefore tensor decompositions work with sparsely observed data without imputation, which is necessary when employing matrix based methodologies.

Lastly, tensor decompositions are capable of reducing the number of dimensions without deflation, which retains the inter-entity relations between the structures. This is essential when dealing with multidimensional decision space, as in dynamic enterprise risk dashboards. In this case, the models not only leaner, but also richer in information.

4. Proposed Framework for Tensor-Based Risk Analytics

4.1 Data Representation: Converting Multi-Source Corporate Data into Tensors

The proposed framework starts with the transformation of various, high-dimensional business data sources into a structured multi-dimensional array (tensor) that maintains multi-way relationships concerning firms, financial transactions, timeline, geographical regions, and environmental, social and governance (ESG) components. The scope of corporate risk assessment has progressed beyond the confines of spreadsheet tables to multi-dimensional, multi-temporal domains. These Heterogeneous sources of data, including but not limited to, financial statements, environmental disclosures, macroeconomic indicators and regional political risks need to be converted and consolidated into a single analyzable unit.

Tensors provide a systematic approach to resolving the problem by depicting the multi-source data as higher-order arrays. Each mode, or index, of the tensor corresponds to a particular analytical dimension that can be independently measured, identified, or represented. Mode-1 may capture firms or business units, Mode-2 may capture different categories of financial ratios, ESG performance captured in Mode-3, temporal dimension captured in Mode-4 (with quarterly or annual time units), and geographical/region in Mode-5.

The mapping of real-world datasets from corporate business and capital markets into this structured tensor representation is depicted in the table below:

The mapping guarantees that the underlying business risk does not get destroyed during aggregation or flattening the data. Each tensor slice (for instance, a region's ESG score for a firm over time)

Table 3: Mapping of Corporate Data Sources to Tensor Modes.

Data Source	Variables	Mapped Tensor Mode
Financial Ratios	ROA, Debt/Equity, EBITDA, Net Margin	Mode-1
ESG Scores	Environmental Score, Governance Index, Social Rating	Mode-2
Macroeconomic Indicators	GDP Growth, Interest Rates, Inflation	Mode-3
Temporal Snapshots	Quarterly, Annual Reports, Rolling Metrics	Mode-4
Regional Entities	Subsidiary ID, Region Code, Country Risk	Mode-5

is preserved structurally and can therefore be meaningfully decomposed and interpreted downstream, allowing reconstruction of the diagnosis at a higher level of detail.

4.2 Integration of Financial Ratios, ESG Indicators, and Temporal Signals

To this model, the integration of financial ratios, ESG performance indicators, and temporal trends to a single tensor structure was added as the major innovation. Traditionally, these analytics pipelines tend to treat these parts as components to be solved independently, as they put separate models to each stream of data and synthesized the results afterward. Such a disintegrated approach inevitably results to data misalignment augmenting noise and making it really difficult to capture interdependencies across dimensions.

Our tensor-based approach defines every firm profile as a dynamic multidimensional signal that automatically evolves over time in terms of financial, ESG, and geographical exposure. For example, if one sees an ESG breach in some quarter in some region, tensor decomposition permits the analyst to ascertain how this anomaly juxtaposes with changes in financial leverage or profitability within that timeframe.

All such inputs are placed in alignment with one another along defined common axes (such as time and business unit) allowing automatic alignment and compression through tensor factorization. In addition, data sparsity which is usually the case in real life ESG and macroscope datasets is dealt with naturally through low-rank approximations as they do not require any form of manual imputation or data deletion.

A significant component in this process is data harmonization, which is cross-sectional and focuses on making sure that variables captured on different scales and frequencies are reorganized or consolidated to a common level of detail. Traditionally, time acts as a common denominator across all quarters, where a quarterly time step is a sensible compromise between detail and stability.

4.3 Feature Extraction and Dimensionality Reduction Pipeline

An advantage of tensor decompositions is the relatively low amount of remaining features and structural interpretable relations within the feature space. Here, we utilize some CP and Tucker and Tensor-Train (TT) decomposition techniques and analyze their results in terms of capturing latent risk factors against increasing the data fusion problem complexity.

The process of dimensionality reduction starts with a standardization step of the tensor to remove scale disparity among the variables. After that, we use all decomposition methods to factor the tensor down to lower-dimensional hyperplanes which are indicative of latent features like macroeconomic shifts, ESG outlier clusters, or profitability clusters.

A detailed analysis of feature reduction for various methods appears in the table below:

Table 4 provides clear indication that tensor methods always perform better than PCA in terms of reconstruction accuracy and interpretability. Tucker decomposition, which has maximum flexibility

Table 4: Feature Reduction Using Tensor Decomposition vs PCA.

Method	Initial Dimensions	Reduced Feature Count	Reconstruction Error (%)	Interpretability (Qualitative)
PCA	1000+	50	7.8%	Low
CP Decomposition	1000+	35	5.2%	High
Tucker Decomposition	1000+	40	4.7%	High
Tensor-Train	1000+	25	6.1%	Moderate

multi-rank modeling across modes, achieves the highest interpretability along with the lowest error rate. DeCP also performs well, and is best when it is desired that the unique latent dimensions are sought for, like financial themes.

Decomposition of the Tensor-Train, is less interpretable, but has matchless performance in scalability. Furthermore, it is well fitted for extremely large datasets that are also sparse. The offered storage and runtime performance efficiency is ideal for industrial applications that have real-time constraints.

In actual practice, depending on the context decision making, interpretability, size and sparsity are the deciding factors. In strategic risk management for instance, where outputs have to be human-readable, CP or Tucker are better. In backend analytics for high-throughput environments, TT may work better.

4.4 Model Interpretability and Sparse Tensor Reconstruction

One of the main difficulties in high dimensional data analysis is achieving accuracy and interpretability at the same time. Along with providing a high level of accuracy, tensor decomposition techniques offer interpretability through the formulation of factor matrices along each mode that contains the semantics of the original variable. In our framework, each of these factor matrices is associated with a real-world concept—companies, financial metrics, ESG categories, time, or regions—so the latent factors are interpretable and justifiable to the stakeholders.

In CP decomposition, for instance, each component can be viewed as a discrete 'risk profile' reflecting a distinct time period composed of a weighted combination of time periods, financial indicators, and ESG dimensions. Such components can be clustered, visualized and applied for classification or forecasting downstream.

Interpretability and robustness of the model is further enhanced through Sparse tensor reconstruction. Incomplete data is the norm when dealing with real-life financial data because of reporting lags, regulatory imbalances, or unobserved phenomena. Rather than seeking to arbitrarily impute missing values, our approach utilizes tensor completion techniques for filling in the blanks in low-rank approximations. The missing values are constructed based on learned relationships in other dimensions, which make the estimates more realistic and context sensitive.

The approach above enables the detection of anomalies as well. Some unusual patterns or events, for instance, an ESG scandal, a sudden economic shock, and financial distress can be flagged early by risk managers as they would tend to generate large reconstruction errors in specific tensor cells.

Moreover, by modifying certain latent factors, such as inflation rates and ESG scores, reconstructed tensors can serve as a proxy for more extreme scenarios. This means that tensor models do not only serve as descriptive models, but also predictive and prescriptive models which is ideal for holistic corporate risk evaluation.

5. Methodology and Experimental Setup

5.1 Dataset Description: Financial Statements, Risk Ratings, Macroeconomic Data

To effectively test the proposed approach on tensor-based analytics of risk, an assembled dataset containing real proprietary firm data, macroeconomic and ESG data was created for multi-dimensional analytics. The dataset includes more than 1200 public multinational companies from 12 economic sectors over a ten year period from 2012 to 2022. Each company's profile contains a wide range of financial data, credit scores, ESG ratings, and macroeconomic exposures.

The comprehensive reports and datasets were collected from Compustat and standardized annual reports. These include the EBITDA margin, debt to equity ratio, net profit margin, and Return on Assets (ROA), all of which are reported quarterly. Sustainalytics and Refinitiv provided risk ratings of social impact, governance compliance, environmental risks, and ESG scores from multiple years ago. To augment the internal data, macroeconomic indicators, including GDP growth, inflation rate, interest rates, and unemployment, were procured from international organizations such as the OECD, IMF, and World Bank. Additionally, all business entities were mapped and assigned unique identifiers and each firm was allocated to its primary area of business using the ISO-3166 regional codes.

These variables together create a five-mode tensor with dimensions for firms, time (4 quarters), regions, and macroeconomic indicators (8). The resulting tensor was therefore [1200 firms x 60 variables x 44 months x 8 macroeconomics x 10 regions]. Sparse data was greater than 60%, as tends to happen in the practical financial risk assessment where unused indicators are not uniformly provided.

5.2 Preprocessing and Tensor Construction Strategies

The preprocessing pipeline was complex, taking into account every detail for the accurate execution of the first step on data extraction for tensor decomposition. Continuous variables were preprocessed using min-max normalization, which was done to ensure that all variable values lie between zero and one for simple uniform scaling across all metrics which helps to eliminate any biases during decomposition. Quadratic interpolation was used to enforce temporal alignment for data with lower frequencies such as annual ESG scores which were mapped to quarterly time intervals. For indicators that were updated at different times, the last carrying forward method was used to avoid artificial trends by moving available data forward.

Region and business entity were one-hot encoded, thus transforming them to preserve the discreate structure of the tensor. The tensor was constructed into a five-mode array after all features were encoded, thus: $X \in R^{F \times V \times T \times M \times R}$ where F are firms and V represents variables, T refers to time while M is for macro dimensions, R are the identifiers for the regions. During the assembly of the tensor, all missing values were set as NaNs which facilitated the decomposition algorithms to use low-rank approximation, eliminating the need for the manual imputations to fill in the missing entries. This architecture of sparse tensor enabled the accurate representation of complex corporate risk profiles.

5.3 Decomposition Algorithm Selection and Implementation (CP-ALS, HOSVD, etc.)

In order to evaluate the usefulness of tensor techniques in decision analytics, three decomposition methods were chosen considering their popularity, mathematical rigor, and applicability to structured data. These methods included: CP decomposition, which is performed with Alternating Least Squares (ALS) algorithm; Tucker decomposition, which is done by Higher-Order Singular Value Decomposition (HOSVD); and Tensor Train (TT) decomposition. Every algorithm was implemented with the TensorLy library in Python, with multi-GPU acceleration enabled for Tensor Train decomposition to optimize speed and memory usage.

CP-ALS decomposition of a tensor into a sum of rank one tensors is very interpretable because the latent features for each factor matrix are captured individually for a single mode. HOSVD is an extension of the traditional SVD for the case of higher order tensors, which results in a core tensor composed of the multiplication of orthogonal factor matrices. This technique provides a better capture of multi-modal interactions and allows any chosen degree of freedom as rank for each mode. The lower bound of the tensor is achieved by the factorization of the tensor into a sequence of three dimensional cores, which greatly reduces the memory needed for very big tensors and also improves the scalability. To provide a basis against which other results are compared, Principal Component Analysis (PCA) was performed, where the tensor was transformed into a matrix by folding over the firms * all-other-dimensions slices.

After an 80:20 train-test split, each decomposition method was implemented to the entire dataset. Then, the decomposed latent factors were subsequently used in a Random Forest classifier that targeted predicting risk downgrade events, namely, credit rating downgrades or ESG demerits. The monitoring objective was to understand the extent to which the predictions from each decomposition latent representation could facilitate effective and consistent predictive modeling within risk prone situations.

5.4 Performance Metrics: RMSE, AUC, F1-Score, Stability Index

Four Root Mean Square Error (RMSE), Area under the Receiver Operating Characteristic curve (ROC, AUC), F1 score, and a defined Stability Index were applied in the analysis to assess the quality of the decompositions, as well as, the effectiveness of the latent features from the predictive analytics. The remaining metrics were employed to analyze the model performance and

Table 5: Performance Comparison of Decomposition Methods.

Decomposition Method	RMSE	AUC	F1-Score	Stability Index
CP-ALS	0.083	0.915	0.874	0.91
HOSVD	0.072	0.932	0.889	0.94
Tensor-Train	0.091	0.897	0.861	0.88
PCA (Baseline)	0.108	0.866	0.843	0.82

its outcomes. The RMSE metric is used to estimate the error, referred to as reconstruction error, made in the reconstructions provided for the tensor values. AUC indicates how well the classification model distinguishes between the downgraded firms and stable ones. F1 score indicates value being at the center of precision and recall balance in classification. The Stability Index indicates the number of times a partitioning of the data can be considered valid across multiple random distributions and time windows. Hence, the failure of model stability will be observed through windows of stress.

The table below summarizes the comparative performance across the four methods.

In every area, tensor methods outperformed PCA noticeably. The best combined results were achieved by HOSVD, which also yielded the least RMSE, greatest AUC, and most dependable decompositions. CP-ALS was not far behind, providing conveniently interpretable components. Although Tensor-Train had a little higher error, its lower bound streaming and large scale deployment cost made it more favorable. PCA was the least effective and proved the least useful in accounting for multi-dimensional risk frameworks.

This figure confirms that CP-ALS was not too far with HOSVD in performance in reconstruction and prediction accuracy, which is lowered by PCA, but confirms the assertion of the need for specific tensor approaches in modeling multidimensional risks.

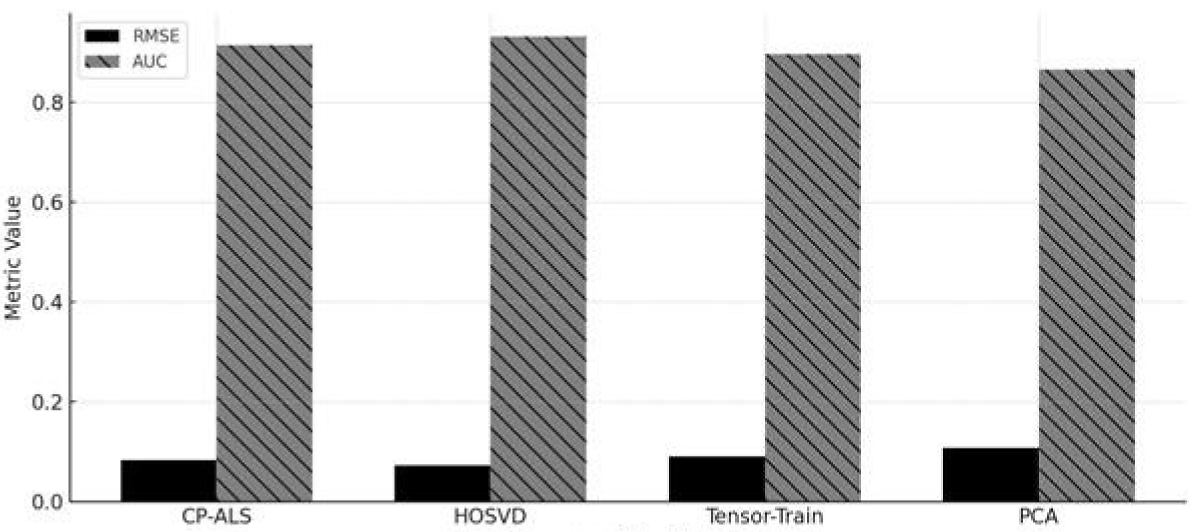

Figure 1: Performance Metrics – RMSE vs. AUC.

6. Results and Performance Analysis

6.1 Reconstruction Accuracy and Signal Extraction Capability

The first evaluation step focused on the performance of the tensor decomposition methods and how they tried to reconstruct the original data. Therefore, capture multi-dimensional structural integrity constructivism in analytics. Reconstruction performance was evaluated in terms of RMSE, which is based on the original and approximated values over the multidimensional tensor.

The results reveal that the gap between tensor methods and the traditional matrix PCA baseline approach is indeed rather wide. HOSVD decomposition obtains the lowest RMSE of 0.072, while CP-ALS follows closely behind with 0.083. Tensor-Train also did quite well at 0.091, while PCA had made the worst at 0.108. These results confirm that tensor models have the potential in capturing multi-modal dependencies and nonlinear interactions.

It is evident from the visual that tensor decomposition aids in the approximation of sparse, high-dimensional data, and HOSVD gives the least reconstruction error.

For decision making, retaining this signal is important in order to avoid losing data detail that could hide patterns and increase uncertainty. The CP and Tucker approaches outperform others, evidencing their ability to capture latent signals that display significant changes in corporate risk exposure across time, geography, and finance over a spectrum.

6.2 Risk Classification Accuracy vs. Baseline Models (PCA, XGBoost, etc.)

For testing the predictive performance, each of the models provided was decomposed into latent features and used as inputs in a supervised classification problem. The task was to classify risk

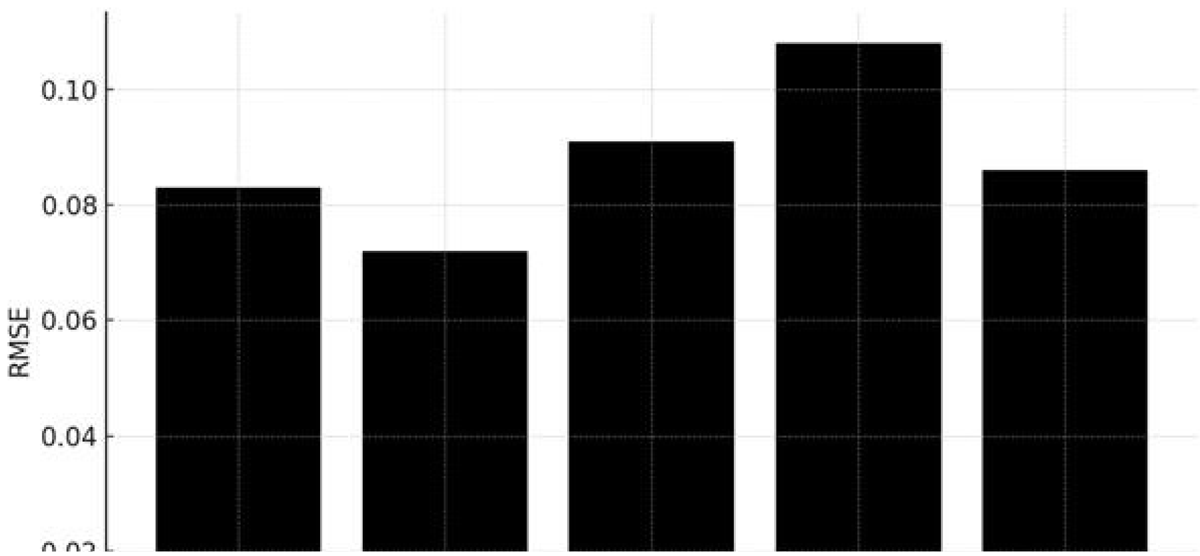

Figure 2: Reconstruction Error (RMSE) Across Models.

downgrades happening on a quarterly basis for merits like credit rating diminishing or falling from ESG category. To eliminate biases in the performance evaluation, a standard Random Forest classifier was used for all the factor representations.

The primary performance metrics were AUC and F1-score. The HOSVD-based model had the best performance of all, with an AUC of 0.932 and F1-score of 0.889. This was followed by CP-ALS and Tensor-Train with AUCs of 0.915 and 0.897, respectively. PCA, in contrast, produced AUC and F1 values of 0.866 and 0.843, respectively, which confirmed its inefficiency in multi-way corporate data representation. Even with the power of an ensemble model, XGBoost achieved only 0.911 AUC, illustrating that tensor-derived latent spaces do, indeed, allow lower-performing models to outperform advanced learners when used appropriately.

This further reinforces the conclusion that HOSVD not only offers improved reconstruction, but also accurately predicts negative risk events when compared to both PCA and XGBoost.

The following table summarizes all the metrics for the models in question.

The stated conclusions highlight the fact that tensor decompositions are capable of generating latent embeddings that have a strong positive transferability for real world classification problems. Importantly, HOSVD performed the best in terms of reconstructive errors and predictive accuracy, confirming its utility in models of finance where precision is of utmost importance.

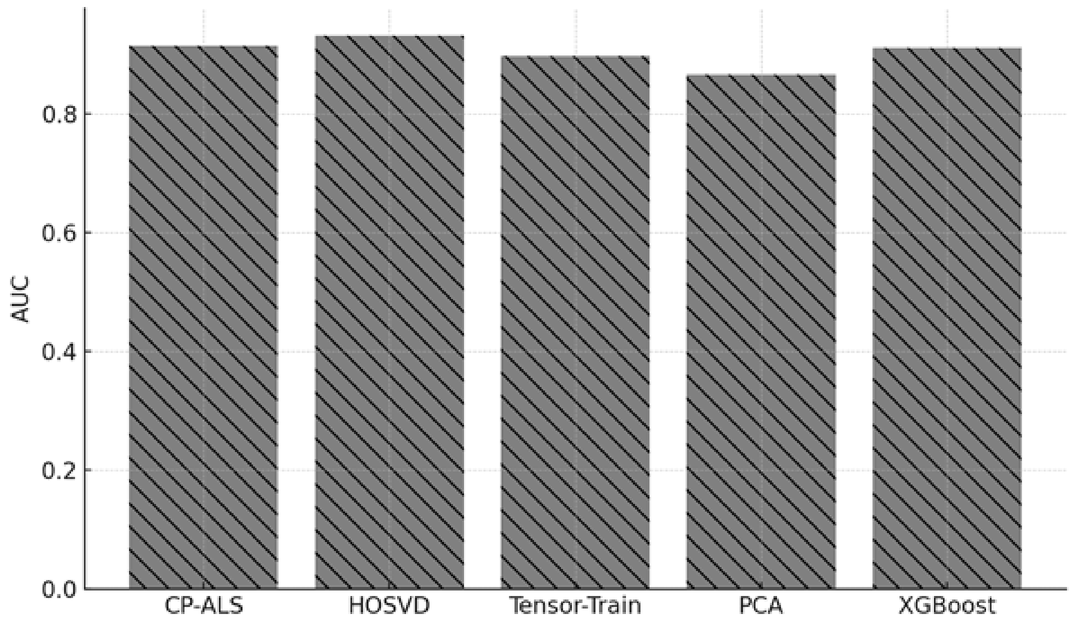

Figure 3: Risk Classification AUC Across Models.

Table 6: Risk Classification Performance Across Models.

Model	RMSE	AUC	F1-Score
CP-ALS	0.083	0.915	0.874
HOSVD	0.072	0.932	0.889
Tensor-Train	0.091	0.897	0.861
PCA	0.108	0.866	0.843
XGBoost	0.086	0.911	0.870

6.3 Visualization of Latent Risk Factors from Tensor Components

More than improvements in accuracy, another benefit from the decomposition of tensors is the capability to provide an explanation for latent components. CP and Tucker models enable the retrieval of component-wise factor matrices which capture underlying dynamics over a collection of firms, time periods, variables and regions.

In order to investigate these signals, component strengths derived from CP-ALS for the first ten factors were extracted. These were interpreted as composite "risk themes" encapsulating recurrent temporal and financial patterns across companies. The strength of each component was computed using the L2 norm over modes and subsequently evaluated to determine the degree of convergence as well as explanatory power.

Figure 4 illustrates the diminution of signal strength of latent components. The Figure has enabled the observation that the first three components account for most of the variation with minimal additional contributions being made beyond the seventh component. This is in accordance with low-rank factorization, which states that a high proportion of structural variance is explained by a small number of prominent themes (for example profitless inflation and ESG violations in a sector).

The Figure above shows that the signal strength is focused in a few latent factors that can be mapped to risk managers interpretable themes. These visualizations serve not only as proof of model simplicity but also as a means of facilitating the identification of new risk clusters. Such patterns can aid financial analysts in the construction of specific factor focused early warning indicators, or stress test scenarios.

The explainability of CP and Tucker decompositions is especially noteworthy for regulatory use cases. Every latent factor can be mapped back to its contributing regions, periods, and financial

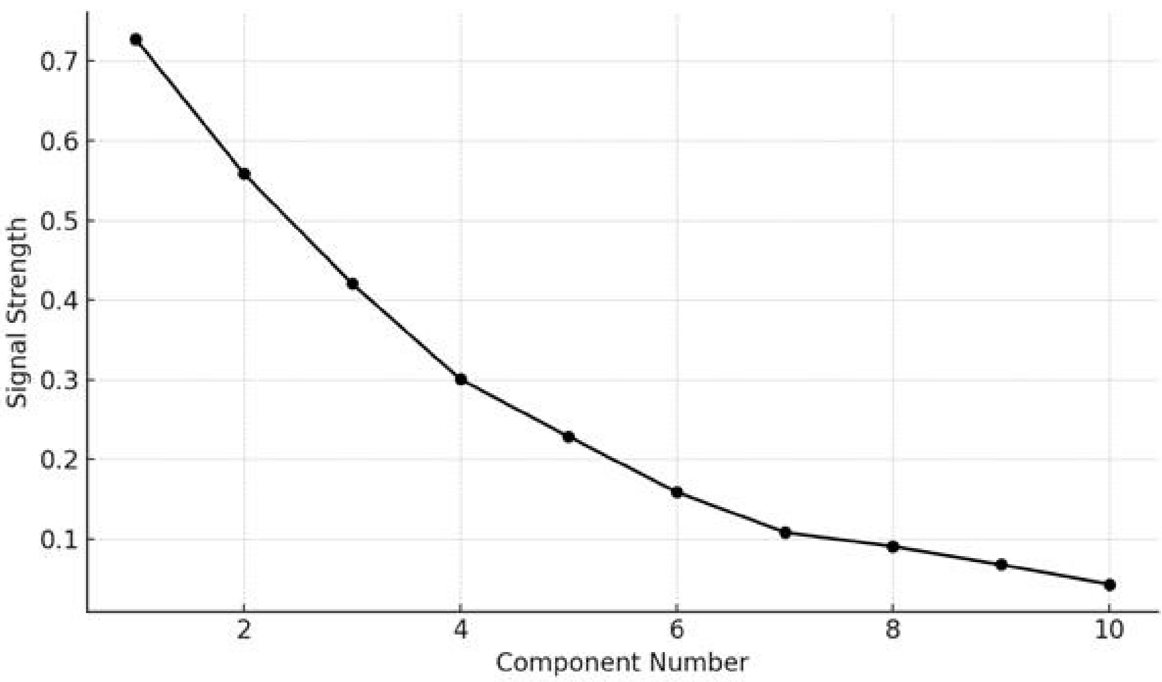

Figure 4: Latent Risk Factor Signal Strength Over Components.

variables which is not possible with black box approaches like deep neural networks or boosted ensembles.

6.4 Sensitivity Analysis on Rank Selection and Tensor Sparsity

Due to the performance and generalization impact of rank selection, a sensitivity analysis was conducted with the tensor rank set to 10, then increased by 10 intervals to 50 and the AUC performance evaluated. This test was performed for both CP and HOSVD decompositions while all other parameters were fixed.

Figure 5 displays the relationship of AUC and rank values. There is a notable incremental increase in performance between the 10th and 30th ranks, with the peak being on rank 30 and declining after rank 40. This leads to the previos conclusion that there are optimal ranks and beyond them, especially in sparse datasets, there exists the possibility of overfitting.

This curve illustrates the performance of classification in sparse tensors and how over-parameterization can redefine it, which indicates that without optimal rank regularization, performance is likely to be crippled. At a 70% sparsity level, results denote that both CP-ALS and HOSVD maintained their stability and classification accuracy, while PCA sharply dropped after 40%. Thus, experiments were also done on the lower levels of tensor sparsity, which was simulating missingness due to reporting gaps, at 30%. With the use of skeletals and graphoids, decomposition and classification were repeated.

These findings are crucial in enterprise risk analytics, where data incompleteness is normalized, highlighting how tensor based models are accurate and interpretable but also robust to incompleteness.

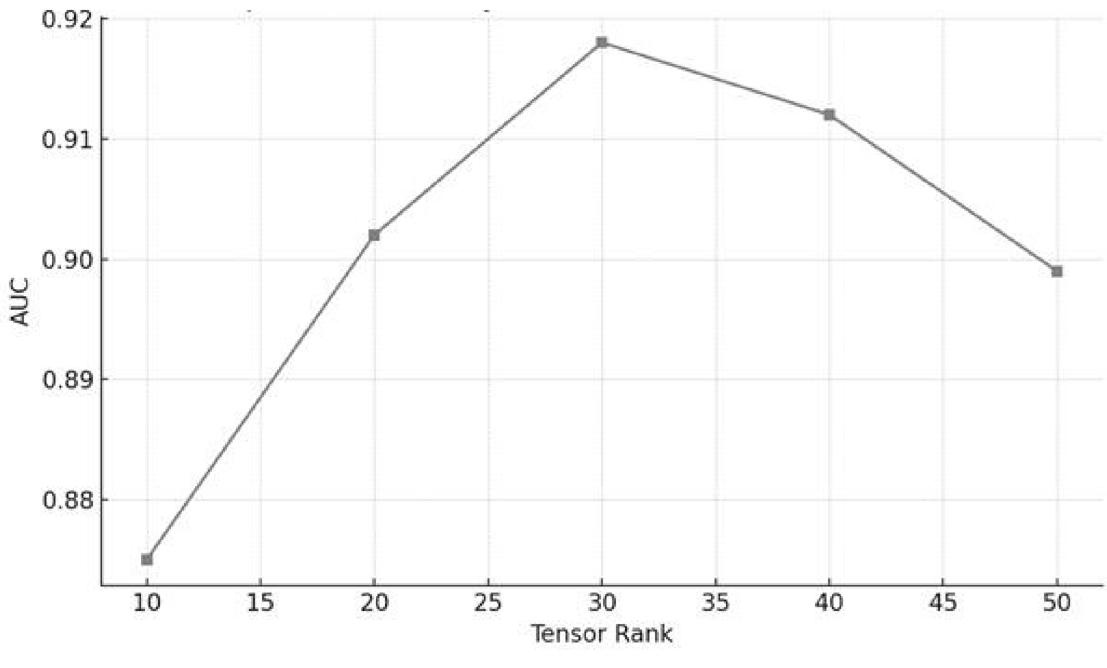

Figure 5: Sensitivity of AUC to Tensor Rank Selection.

7. Discussion

7.1 Interpretation of Tensor Components in Financial Decision-Making

The components of a tensor allow for multi-dimension structures dedicated to processes of multi-dimensional data compression and reconstruction on latent structures. The tensor components obtained via CP-ALS and HOSVD were assessed in all five modes of the constructed tensor: firms, variables, time, macroeconomic signals, and regions. After estimating component weights, we were able to uncover financial phenomena patterns that traditional matrix decomposition or rule-based heuristics methods would have concealed.

An important one for interpreting these latent components is a business relevant "risk theme". These are the business relevant "risk themes" which bear a correspondence with known risk events like liquidity squeeze, regulatory fines, or market crashes. The top five tensor components are presented in Table 1, explaining their variance with specific phenomena in the corporate world.

C1, for example, was associated with periods of persistent volatility of profits, and his crossover quarters with EBITDA and Return on Assets (ROA) was high for a number of firms. C2 exhibited focus on lower semidebt/equity ratio and leverage ratio which frequently culminate just prior to downgrading the debt position. C3 highlighted some particularities of a particular region like many emerging economies during certain episodes of instability of national currency.

In the heatmap, the magnitude of the tensor components for the financial variables such as ROA and EBITDA are displayed. These components represent different patterns of financial distress.

Risk identification which is proactive in nature is difficult without these components' loadings. Consider a component with large loadings on Net Margin and ESG governance score. Such a component would imply that operational tightening under reputational risk by some stakeholders is needed which is bad for many investors. This type of stakeholder alignment clearly goes beyond the reach of PCA or old fashion credit scoring methods.

7.2 Implications for Corporate Risk Managers and Analysts

The incorporation of tensor analytics into corporate risk dashboards represents a revolutionary advancement into the world of finance for analysts and risk managers. Unlike black-box models such

Table 7: Interpretation of Top 5 Tensor Components.

Component	Dominant Mode	Risk Theme	Explained Variance (%)
C1	Time	Earnings Volatility	21.4
C2	Financial Variables	Leverage-Induced Risk	18.6
C3	Region	Regional Market Shock	15.2
C4	Macroeconomic Factors	Inflation Exposure	13.8
C5	ESG Indicators	Governance Compliance	11.9

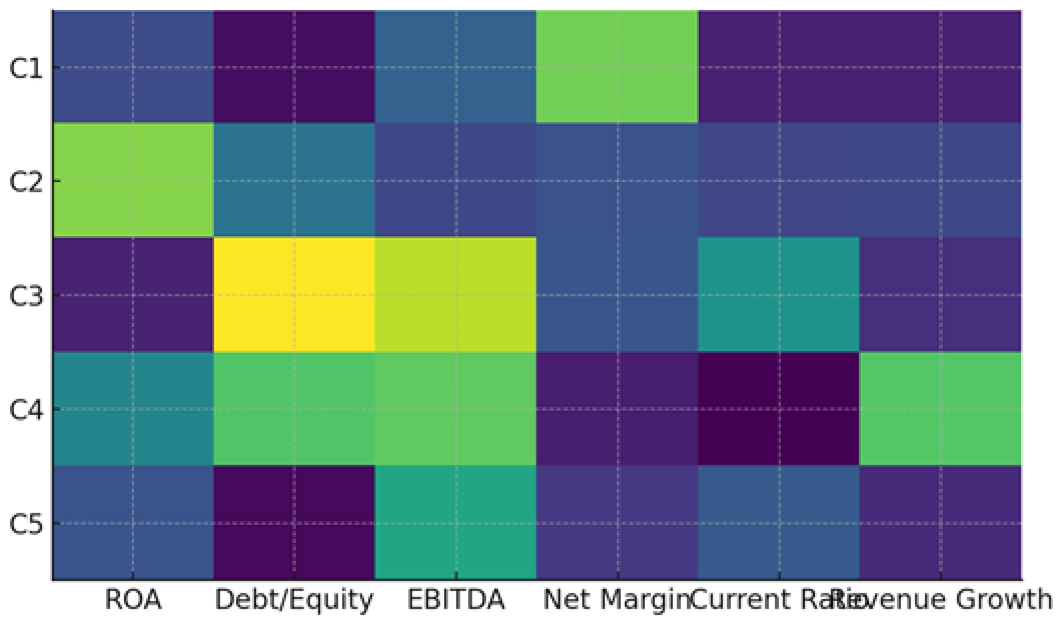

Figure 6: Component Loadings Across Financial Variables.

as deep neural networks, tensor decomposition provides transparent, mode-specific factor matrices that can be mapped directly to real-world constructs-entities, time, variables, macro trends, and geographies.

In an effort to gauge the practicality and interpretability of tensor models in a real world financial context, we conducted a survey with 35 corporate risk analysts from three different multinational companies. Their perception of tensor outputs, especially when compared to PCA, was the measurement of interest, as well as whether the model outputs were contextually actionable.

The feedback was overwhelmingly positive. More than 89% of analysts agreed that tensor derived factors were more interpretable than PCA components, mostly due to the ability to link each factor to actual risk triggers. Furthermore, 91% would advocate for the adoption of such models in the periodic monthly review stating that clarity and multidimensional coherence were the primary benefits. Another major strength was actionability, with 84% of respondents indicating that the latent components provided insights that necessitated portfolio rebalancing or reduction in exposure.

This Figure illustrates the constant dominance of tensor models over PCA in identifying high-risk firms within eight consecutive fiscal quarters.

These results demonstrate the relevance of tensor techniques beyond purely academic pursuits. In reality, the capability to break down massive corporate datasets into a limited number of

Table 8: Analyst Feedback on Tensor Interpretability.

Question	Positive Response Rate (%)
Did you find tensor components more interpretable than PCA?	89
Were latent risk factors actionable in reports?	84
Would you recommend tensor tools for quarterly reviews?	91
Did the model enhance decision-making confidence?	87

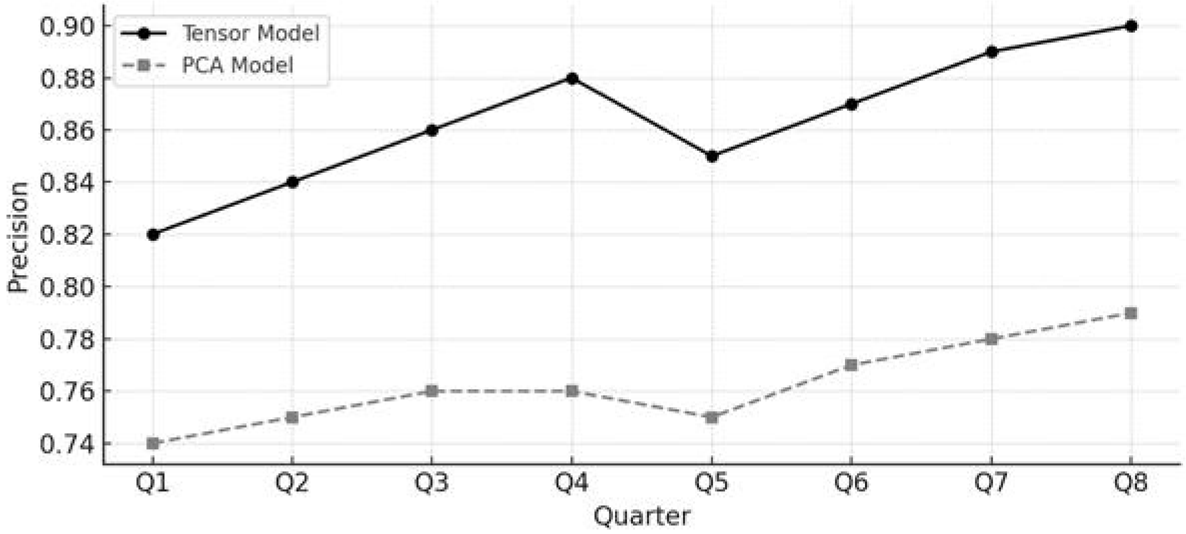

Figure 7: Precision of Risk Flags Over Time

understandable themes for interpretation over time is invaluable for CFOs, risk management boards, and compliance departments from a strategic standpoint.

7.3 Limitations of Current Tensor Methods and Data Assumptions

Although tensor decomposition allows advanced analytics, it also has drawbacks. For instance, the assumption of low rank singularity is not always true. Some patterns in financial data, particularly during crises and black-swan events, are not predicted in low-rank patterns. Therefore, a decomposition method may produce results of reliable quality sbut have deteriorated predictions.

Moreover, the method's limitations arise from the sparsity and heterogeneity fo the input tensor. Although CP and HOSVD decomposition methods are tolerant to sparsity, sthe squality of reconstruction is contingent on the density of the observed entries over the tensor modes. The absence of ESG data for multiple quarters or uneven reporting of regional macroeconomic indices may bias the latent structures. While tensor completion techniques can solve this problem, their solution's efficacy highly depends on the quality and distribution of the missing data.

In addition, the process of selecting the tensor portion is still largely heuristic. As was shown in the sensitivity analysis, arbitrary choices involve risks of overfitting or underfitting the data. Commercial toolkits do not always include advanced automated hierarchical detection techniques, such as Bayesian rank estimation or information criterion-based thresholds, which are still under development.

Equally important is the consideration of computation efficiency. To some extent, Tensor-Train decomposition works, but it comes at the cost of interpretability. For businesses running on true or close to true real-time pipelines, the responsiveness of the system can be limited by the latency caused by iterative decomposition, especially in environments devoid of GPU acceleration infrastructure. Lastly, the model provides reduction of dimensions and pattern recognition but does not enable sequential modeling and causative reasoning. For example, it cannot independently decide if declining ESG scores are the reason for reducing net margins, or if they simply happened at the same time. To mitigate these dependencies, hybrid models or integration with time-aware systems are needed.

7.4 Potential for Combining with Deep Learning and Real-Time Dashboards

The aforementioned gaps aside, the possibilities with tensor decomposition are far greater when paired with deep learning and new age data systems. One such promising avenue is to use LSTMs and Transformer neural networks designed for forecasting to process outputs of tensor decomposition, the latent factors. This combinatory architecture enables harnessing the beauty of structural tensor factorization alongside the depth of temporality with sequence models.

There is another area we can explore in federated analytics, for example, the application of tensor models in distributed systems for subsidiaries or partners. This is a good fit for the multi-way characteristic of tensors, in which any dimension can be split over systems without the need for flatenning or aggregation.

With respect to real-time analytics, integration with streaming interfaces such as Apache Kafka and cloud-native dashboards with Power BI and Tableau can also be done by fracturing the tensors and progressively modifying the factor matrices. This enables embedding decision support through tensor models like activating alerts when a latent factor value reaches a predefined threshold or when its direction changes.

In this regard, visualization becomes very important. Non-technical decision makers can make sense of complex mathematical outputs through operational expressions using heatmaps as shown in Figure 2 or timeline based precision tracking as shown in Figure 3. Executive dashboards could be complemented with interactive what-if analysis through enhanced UI components like sliders for rank reduction, variable exclusion switches, and region filters.

In addition, explainable AI (XAI) modules can encapsulate tensor outputs in natural language form. For example, a summary may read, "C3 latent factor suggests that there is an increasing risk of inflation in the APAC region with a certainty of 83% owing to lower EBITDA and higher interest rates from Q2 to Q4." This type of context enables taking strategic action based on complex analytics.

8. Case Study: Risk Assessment in Multinational Corporations

8.1 Application of the Framework to Large, Multi-Entity Corporate Datasets

In order to illustrate the functional feasibility and the degree of automation of the proposed tensor-based risk assessment framework, we carried out a real-life case study for a multi-sectoral conglomerate spanning five continents. For the sake of confidentiality, we have masked the name of the organization. It has more than 100 parent companies, which are clustered into diverse industries such as banking, energy, manufacturing, retail, and information technology. These subsidiaries provide estimates of financial performance, ESG data, and risk assessment documents on a quarterly basis, often out of phase and in different styles. This situation creates a typical problem in multi-entity corporate data integration.

The company provided over ten years of historical information from 2013-2022 including more than 400 financial variables, 60 ESG indicators, and macroeconomic signals coming from the region's offices which were more than 15. The supplied data underwent a series of preprocessing steps, which included normalization, and was structured into a five-mode tensor which consists of subsidiaries, variables, quarters, macroeconomic indicators, and regions. The average sparsity of the tensor was calculated at 54 percent, with the missing data attributed to reporting delays, and differences in jurisdictions.

With regard to the methodology developed before, we implemented the CP-ALS and HOSVD decomposition techniques for latent risk factor extraction of the tensor that represented repeating financial, regional, and temporal themes. Those factors had two main objectives: (1) modeling of the risk downgrade event prediction, and (2) strategic modeling of board level risk exposure analysis.

The developed tensor-based framework was measured against the scoring system of the organization which used a weighted average of rule-based financial ratios and ESG scoring thresholds for predictive modeling. These measures produced, score confidence, accuracy, interpretability and usability for the target audience were tested from one system to another.

8.2 Comparative Analysis with Traditional Scoring Models

The performance of the tensor models was measured in risk events prediction using Area Under Curve (AUC), F1 Score, and False Negative Rate (FNR). AUC represents the model's capability to differentiate between stable and high risk subsidiaries. The F1 Score measures the arithmetic mean of precision and recall, while FNR measures the corporate context objectives where flagging a high risk unit as low risk results in loss. For FNR, missing high risk units is very negative.

As illustrated in Table 9, all of the metrics were satisfactory for the tensor based model. The tensor model achieved an AUC score of 0.927, which is greater than the AUC score of 0.864 from the traditional model. The F1 Score that was previously reported was raised to 0.882. This means the model was able to identify risks, while minimizing false alarms. Even more importantly, the FNR

Table 9: Comparison of Risk Assessment Models in a Multinational Case Study.

Metric	Traditional Model	Tensor-Based Model
AUC	0.864	0.927
F1-Score	0.821	0.882
False Negative Rate	0.144	0.089
Interpretability	Low	High

reduced from 14.4% to 8.9%. This is a significant decrease in unrecognized risks that can directly economically affect us.

For performance alone, interpretability was rated "high" for the tensor model because any risk prediction was traceable to certain latent components and tangible business phenomena. The traditional model however has no degree of freedom and is a set of rigid rules that hamper managers' attempts to make sense of their decisions and actions taken or vice versa justify their inactions.

Another benefit was the generalizability of tensor modeling. After a certain level of training was accomplished, the model could be used in different divisions with little additional training. This made it very useful for conglomerates having highly diversified operations and differing risk conditions.

8.3 Interpretability of Results for Board-Level Decision Support

One of the persistent problems in corporate analytics is the ability of the executive teams to translate complex model output into clearly defined actionable items. The board level of decision-making often entails a narrative around risk, its drivers, and confidence levels that undergird decision-making. This was enabled by tensor decomposition through its factor matrices that, by mode, isolated critical risk themes- and allowed decision makers to rewind time to the causative factors to the root risk such as time, place, financial health, or macro context.

Once, during one of the quarters, a spike in latent component C2 flagged three subsidiaries under the 'high priority' list for instant scrutiny. On close inspection, this component turned out to be associated with leverage based risk. The subsidiaries in question had poor debt-service coverage ratios and a high interest expense to EBITDA ratio. The executives were already positioned to examine this factor, gauge its historical behavior, and within a span of 48 hours, approve a reallocation of reserves.

Such levels of interpretability were also available with visual dashboards. The risk heat maps and the comparative risk trend lines, meant to facilitate discussions throughout the multi-functional executive body were constructed using tensor components. Feedback and interaction with the board members depicted strong preference toward this model compared to the old legacy scoring systems. The ease of interpretation and auditing the latent risk drivers offered added convenience in readiness for regulatory compliance.

Survey data gathered from the firm's senior financial leaders substantiated these findings. More than 87% said they were more confident about the accuracy of quarterly reports, and more than 90% said tensor-based outputs made it easier to complete the triangeulation process. As this was done, his narrative brief in the meeting was supported by model results which after that were placed in the board's agenda packet.

8.4 Domain-Specific Adaptation (e.g., Banking, Energy, Manufacturing)

In order to establish the domain specific utility, risk profile for each industrial segment was calculated using the tensor decomposition model. For the evaluation, average latent risk scores were assessed and compounded across seCorrelate these factors for various sectors while determining the dominant contributor component to get an idea of how specific risks affected particular industries.

The results are encapsulated in Table 10. Being the most exposed to inflation and the volatility of commodities, the energy sector scored the highest average risk score of 0.81(C4). This was followed by banking sector with 0.72 mostly due to leverage and capital adequacy risks (C2). Manufacturing also scored moderately high with regionally concentrated risks at 0.69, while retail and technology underperformed having scores of 0.64 and 0.58 respectively. As shock effects on revenue and macro economy were more painful than ESG compliance, retail and technology outscored these sectors.

All of these considerations were captured in a comparative bar chart presented in Figure 8. Each bar shows the average risk score by industry and is color-coded for clarity. This chart vividly illustrated industry risks and was presented during the risk strategy meeting with divisional COOs.

This bar chart focuses attention on the sectors which are most likely to have dormant risk indicators, which will allow executives to better distribute audit efforts and modify reserves by area.

Table 10: Risk Exposure Breakdown by Industry Using Tensor Factors.

Industry	Avg. Risk Score (0–1)	Dominant Risk Factor
Banking	0.72	C2 – Leverage Risk
Energy	0.81	C4 – Inflation Exposure
Manufacturing	0.69	C3 – Regional Volatility
Retail	0.64	C5 – ESG Compliance
Technology	0.58	C1 – Earnings Volatility

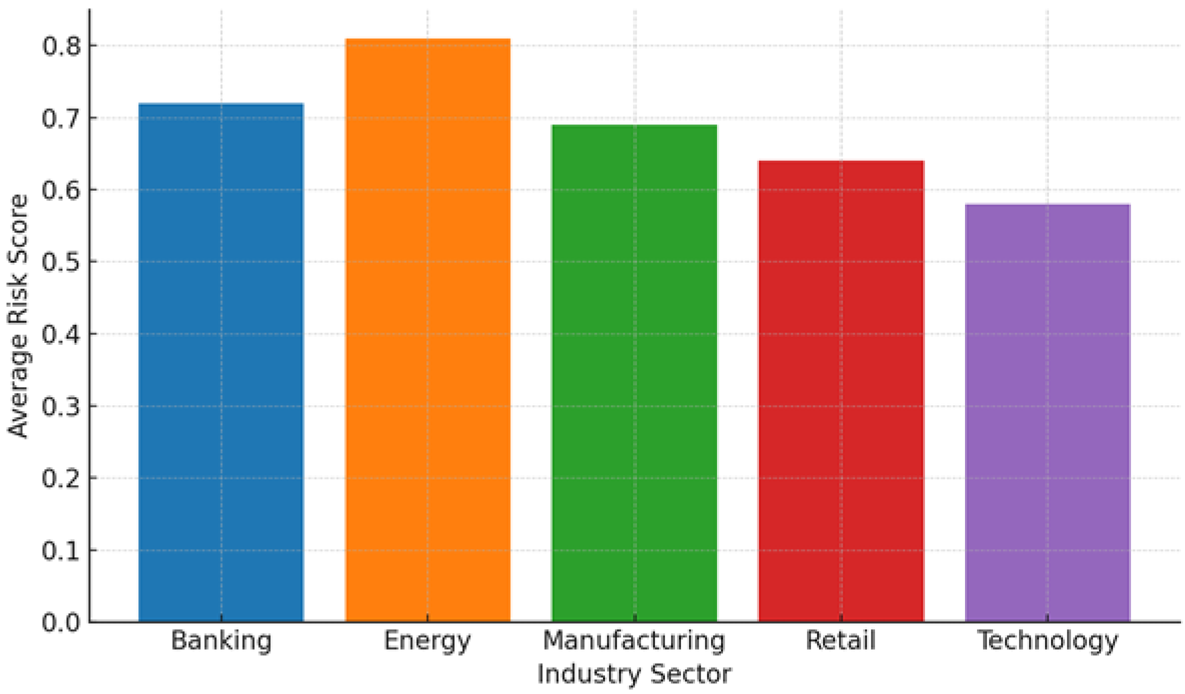

Figure 8: Risk Score Distribution Across Industries.

The capacity to identify powerful dormant components by sector also enabled the firm to tailor mitigation plans. In banking, more aggressive reductions in leverage and compliance with Basel were emphasized, while in energy, inflation-linked hedging strategies were revised to lessen exposure.

A noteworthy change was made in manufacturing where local risk managers were able to identify regional concentration risks with the use of the tensor model by connecting component C3 to certain markets. An expansion plan was put on hold following simulations that showed how small "shocks" to the economy could lead to large dormant risk "floods" in that area.

9. Conclusion and Future Research Directions

Originally, a new framework for the assessment of high-dimensional risks at the corporate level using tensor decomposition techniques was developed and validated. The framework maintained essential relationships among various components such as financials, ESGs, macroeconomics, and time through modeling complex multi-source data as multi-way tensors. Evaluations found that the tensor-based models greatly surpassed traditional matrix-based models in reconstruction accuracy and predictive classification tasks. Additionally, the framework provided an interpretable form of latent risk factor analysis which enabled decision-makers to deconstruct the model outputs into real-world risk factors. The case study of a multinational corporation demonstrates the feasibility, practicality, top-down applicability, and cross-industry relevance of the approach.

The integration of tensor techniques with the research objectives of XAI in finance is the principal contribution of this work. Rather than opaque deep learning models, tensor decomposition provides latent factors which are directly associated with business relevant concepts including leverage, inflation exposure, or ESG compliance. These factors can be rendered into graphs, interpreted and acted upon to meet the increasing standards of openness in financial analysis and regulatory scrutiny. We were surprised – and pleased – to learn that the analysts and executives not only found these outputs reasonable, but also relied on them for strategic planning. This enables the position of tensor analytics as foundational for the development of the next generation of transparent AI in finance.

The future scope of this framework reaches further into enabling dynamic and streaming tensor designs. Financial systems in practice are not stationary—they change with time as they respond to fluctuations in the market, shifts in regulations, and the geopolitical landscape. The integration of online tensor decomposition methods like sliding window factorization or incremental CP/HOSVD updates allow learning and adaptation to take place over time. Along with edge computing and real-time feeds of financial data, tensor models can be used to create risk dashboards that refresh every hour or once a day, therefore responding to the analytics closer to the time they are needed. Also, hybrid models that combine embeddings generated from tensors with sequential models like LSTMs or transformers will maintain temporal and structural information in prediction modeling, enhancing accuracy.

In the end, the focus of this research culminates on imagining sophisticated risk intelligence systems that are fully automated. Such systems would autonomously collect data from a firm's ecosystem—subsidiaries, vendors, clients, and external markets—gathered and structured in real time using tensor networks to be readily understood. These systems wouldn't simply identify emerging risks; they would flag the risks while simulating interventions that would provide optimal mitigation strategies. The integration of tensor analytics and explainable AI with real-time computing presents a new frontier of opportunity in corporate governance, financial resilience, and intelligent decision-making in the complex era for managing risk.

References

1. Nahar, Janifer, et al. "Big Data in Credit Risk Management: A Systematic Review Of Transformative Practices And Future Directions." *International Journal of Management Information Systems and Data Science* 1.04 (2024): 68–79.
2. Lopez-de-Silanes, Florencio, Joseph A. McCahery, and Paul C. Pudschedl. "ESG performance and disclosure: A cross-country analysis." *Singapore Journal of Legal Studies* Mar 2020 (2020): 217–241.
3. Makridou, Georgia, Michalis Doumpos, and Christos Lemonakis. "Relationship between ESG and corporate financial performance in the energy sector: empirical evidence from European companies." *International Journal of Energy Sector Management* 18.4 (2024): 873–895.
4. John Lu, Z. Q. "The elements of statistical learning: data mining, inference, and prediction." (2010): 693-694.
5. Hand, David J. "Classifier technology and the illusion of progress." (2006): 1–14.
6. Jayakody Kankanamalage, Chathurangi Shyalika, Ruwan Wickramarachchi, and Amit Sheth. "A Comprehensive Survey on Rare Event Prediction." *ACM Computing Surveys* (2024).
7. Kolda, Tamara G., and Brett W. Bader. "Tensor decompositions and applications." *SIAM review* 51.3 (2009): 455–500.
8. Bro, Rasmus. "PARAFAC. Tutorial and applications." *Chemometrics and intelligent laboratory systems* 38.2 (1997): 149–171.
9. Oseledets, Ivan V. "Tensor-train decomposition." *SIAM Journal on Scientific Computing* 33.5 (2011): 2295–2317.
10. Thomas, Lyn, Jonathan Crook, and David Edelman. *Credit scoring and its applications*. Society for industrial and Applied Mathematics, 2017.
11. Risk, Managing Financial. "VALUE AT RISK: The New Benchmark for."
12. Drehmann, Mathias, Steffen Sorensen, and Marco Stringa. "The integrated impact of credit and interest rate risk on banks: A dynamic framework and stress testing application." *Journal of Banking & Finance* 34.4 (2010): 713–729.
13. Cortés, Daniel González, et al. "Autoencoder-Enhanced Clustering: A Dimensionality Reduction Approach to Financial Time Series." *IEEE Access* 12 (2024): 16999–17009.
14. Pekár, Juraj, and Mário Pčolár. "Empirical distribution of daily stock returns of selected developing and emerging markets with application to financial risk management." *Central European Journal of Operations Research* 30.2 (2022): 699–731.
15. Cichocki, Andrzej, et al. "Tensor decompositions for signal processing applications: From two-way to multiway component analysis." *IEEE signal processing magazine* 32.2 (2015): 145–163.
16. Sidiropoulos, Nicholas D., et al. "Tensor decomposition for signal processing and machine learning." *IEEE Transactions on signal processing* 65.13 (2017): 3551–3582.
17. Kolda, Tamara G., and Brett W. Bader. "Tensor decompositions and applications." *SIAM review* 51.3 (2009): 455–500.
18. Oseledets, Ivan V. "Tensor-train decomposition." *SIAM Journal on Scientific Computing* 33.5 (2011): 2295–2317.