Power System Technology
ISSN:1000-3673

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

Machine Learning-Based Financial Health Assessment Model for Resource-Based Enterprises

Ruixin Luo a*，Xin Wangb

a.Basic College, Chongqing Institute of Technology, Chongqing 400000, Chongqing, China

b. College of Digital Finance and Economics, Chongqing College of Urban Management, Chongqing 400000, Chongqing, China

Corresponding author：lrx12342021@163.com

Abstract

Resource-based enterprises, such as those in mining, oil, gas, and forestry, face unique financial challenges due to the inherent volatility of commodity prices, environmental regulations, and resource depletion. Traditional financial health models often fail to capture the dynamic, real-time external factors affecting these industries. This paper proposes a novel machine learning-based financial health assessment model tailored to resource-based enterprises. The proposed model integrates both traditional financial metrics and dynamic external variables—such as commodity market trends, regulatory changes, and environmental risks—into a hybrid machine learning framework. The model utilizes a multi-layered architecture comprising Gradient Boosting Machines (GBMs), Long Short-Term Memory (LSTM) networks, and Transformer-based NLP models. A reinforcement learning component enables the model to adapt in real-time, ensuring it remains responsive to changing market conditions. The Dynamic Financial Health Score (DFHS), produced by the model, offers real-time, interpretable financial health assessments that are vital for proactive risk management. Experimental results show that the proposed model outperforms traditional models such as the Altman Z-score and logistic regression, demonstrating higher accuracy, recall, and predictive reliability. This research highlights the importance of incorporating real-time, dynamic data into financial health assessments and provides a robust, adaptable solution for managing financial risks in resource-based industries.

Keywords: Financial Health Assessment, Dynamic Financial Health Score, Financial

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

Prediction, Risk Management, Adaptive Modeling

1. Introduction

Resource-based enterprises—such as those in mining, oil and gas, forestry, and other extractive industries—play a vital role in national economies, particularly in resource-rich countries. However, these enterprises face unique and multifaceted challenges that distinguish them from other sectors. Chief among these challenges is the inherent volatility in commodity prices, which is often driven by global market fluctuations, geopolitical events, and macroeconomic instability. Additionally, resource-based industries are increasingly subject to environmental regulations, carbon taxation, and sustainability mandates, all of which can impose unexpected financial burdens. The finite nature of natural resources further complicates long-term planning, as resource depletion directly affects operational capacity, asset valuations, and future profitability. Collectively, these external factors make the financial health of resource-based enterprises highly sensitive and dynamically influenced by the broader socio-economic and environmental landscape.

Despite significant advancements in financial modeling and risk assessment, resource-based enterprises continue to struggle with financial uncertainty. Traditional financial health assessment tools—such as ratio analysis, bankruptcy prediction models (e.g., Altman's Z-score), and cash flow forecasting—tend to rely on historical financial statements and assume a relatively stable environment. These models often overlook the impact of real-time market conditions and fail to incorporate industry-specific externalities that can significantly affect an enterprise's financial viability. For example, a sudden drop in global oil prices or a shift in government policy regarding mining permits can rapidly transform a financially stable enterprise into a distressed one. As such, static models lack the agility and responsiveness needed for accurate, forward-looking assessments in this sector.

A growing body of literature has explored the application of machine learning techniques to financial prediction and risk management. However, most studies focus on general-purpose models trained on financial ratios or stock market data, without tailoring the approach to the unique volatility and regulatory risks associated with resource-based enterprises. Moreover, these models often operate on static or aggregated data, ignoring the potential of integrating dynamic, real-time data such as commodity prices, climate policy changes, or news sentiment related to resource markets. There remains a critical research gap in developing a holistic financial health assessment framework that can adapt to the complex

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

and rapidly evolving operating environment of resource-based enterprises.

This study aims to address this gap by proposing a novel, machine learning-based financial health assessment model specifically designed for resource-based enterprises. Unlike traditional models, our approach integrates dynamic, real-time data sources—including commodity market trends, environmental policy shifts, and resource utilization patterns—into the financial assessment process. By leveraging advanced machine learning algorithms and adaptive modeling techniques, the proposed framework offers a more accurate, responsive, and industry-specific assessment of financial health. This innovation not only improves prediction accuracy but also enhances the decision-making capacity of enterprise managers, investors, and policymakers in an increasingly uncertain and resource-constrained global economy.

2. Related Work

The financial health assessment of enterprises has long been a subject of academic interest, and various methods have been proposed over the years to predict financial distress, assess solvency, and improve decision-making. These approaches can be broadly classified into traditional financial ratio models, bankruptcy prediction models, and machine learning-based approaches. However, these models often fall short when applied to resource-based enterprises, as they typically fail to account for the dynamic external factors that uniquely affect these industries.

Traditional Financial Health Assessment Models

Traditional financial health models, such as financial ratio analysis, are among the most widely used methods for assessing a company's financial health. These methods rely on key financial ratios, including liquidity, profitability, solvency, and efficiency, which are derived from historical financial statements. One of the most well-known models is the Altman Z-Score, which predicts the likelihood of bankruptcy using a linear combination of financial ratios. However, while these models can offer useful insights, they suffer from significant limitations when applied to industries prone to high volatility, such as resource-based enterprises. For example, the Z-score may fail to capture the impact of resource price fluctuations, regulatory changes, or other externalities that are especially crucial in sectors like mining, oil, or forestry. Furthermore, these models often assume that past financial performance is an accurate indicator of future health, a premise that does not hold in resource-based industries where market conditions can shift abruptly.

Another limitation of traditional financial models is their reliance on static data. They tend to assume that financial health can be evaluated at a particular point in time, which makes them poorly suited to industries that require continuous risk assessments due to external factors. These models also do not incorporate real-time data or adjust quickly to changes in market conditions, making them less useful for dynamic decision-making, particularly in industries where commodity prices and environmental regulations can shift unpredictably.

Bankruptcy Prediction Models

Bankruptcy prediction models, such as the Logistic Regression Model and Probit Models, have been extensively used in the financial industry to predict corporate failures. These models, like traditional ratio analysis, rely on historical financial data to classify firms into categories based on the likelihood of bankruptcy. While they offer a more probabilistic approach to assessing financial health, their effectiveness in dynamic environments is limited. These models do not adequately integrate real-time variables, such as changing commodity prices or shifts in environmental policies, that have an outsized impact on resource-based enterprises.

In recent years, more advanced machine learning techniques, such as decision trees, random forests, and support vector machines, have been explored to enhance bankruptcy prediction. These methods can capture non-linear relationships and are more flexible in accommodating large and complex datasets. However, like traditional methods, most of these models are limited by their reliance on historical financial data and do not fully address the complexities of resource-based enterprises, where external risks—such as environmental factors, regulatory changes, and resource exhaustion—can dramatically alter a company's financial outlook.

Machine Learning-Based Approaches

Recent advancements in machine learning have significantly influenced the field of financial health assessment, providing new avenues for prediction and analysis. Machine learning techniques, such as Random Forests, Gradient Boosting Machines (GBM), and Neural Networks, have been applied in bankruptcy prediction, credit risk assessment, and financial forecasting. These models excel at handling large and high-dimensional datasets, capturing complex patterns that may not be apparent with traditional financial ratios. Moreover, they can be trained on a diverse set of features, including financial data, market

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

sentiment, and industry-specific variables.

However, despite their promising capabilities, these models often fail to consider the full spectrum of dynamic, real-time data that is crucial for resource-based enterprises. Machine learning models in finance are often limited to historical financial data or market indicators that are not tailored to industry-specific risks. Furthermore, while many of these models are highly predictive in static environments, they can struggle with providing real-time financial health assessments in industries like mining or oil extraction, where prices and regulations fluctuate rapidly.

Resource-Based Enterprise-Specific Models

Some studies have attempted to apply machine learning techniques specifically to resource-based enterprises, but these efforts are still in their early stages. For example, a few studies have explored using machine learning for commodity price forecasting or production optimization, but these models tend to focus on operational efficiency rather than financial health. While these studies acknowledge the challenges of resource-based industries, they do not address the holistic financial assessment models needed to account for both internal financial data and external dynamic factors such as commodity price volatility, environmental policies, and resource exhaustion.

A notable gap in the literature is the integration of real-time data sources—such as commodity market fluctuations, government regulations, and environmental risks—into the financial health models of resource-based enterprises. Most existing models still rely on past performance indicators or aggregated data, which fail to reflect the complex and volatile nature of these industries. As a result, resource-based enterprises may face significant financial risks that are not captured in traditional or even machine learning-based models.

3. Methodology

This research proposes a novel, dynamic, and industry-specific methodology for assessing the financial health of resource-based enterprises using machine learning. Unlike traditional approaches that rely heavily on static financial ratios and historical data, the proposed methodology integrates real-time external variables—such as commodity prices, regulatory updates, and environmental indicators—into an adaptive, multi-layered machine learning framework. The approach consists of five key stages: data acquisition, feature engineering, model architecture design, training and validation, and dynamic risk scoring.

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

3.1 Model Overview

The proposed model is designed to dynamically assess the financial health of resource-based enterprises by integrating both traditional financial indicators and real-time external variables through a hybrid machine learning architecture. The model is structured in a modular pipeline, beginning with a data ingestion layer that collects and aggregates multi-source data: structured financial records, time-series commodity prices, regulatory changes, and sentiment signals from news media. Once aggregated, the data is passed through a feature engineering and transformation module, where conventional financial ratios are enhanced with novel, domain-specific indicators such as a Composite Resource Risk Index (CRRI) and regulatory event flags.

Next, a multi-model learning layer performs the core analytical tasks. It comprises a combination of base models, including Gradient Boosting Machines (GBMs) for structured financial data, LSTM networks for time-series commodity trends, and Transformer-based NLP models for textual sentiment analysis. These models are trained independently and feed their outputs into a meta-learner—typically a logistic regression or neural network—that synthesizes the individual predictions into a unified Dynamic Financial Health Score (DFHS).

The model is equipped with a reinforcement learning component that enables real-time adaptation by learning from ongoing changes in input distributions and financial outcomes. This continuous learning loop allows the model to recalibrate its predictions as external factors shift, such as sudden commodity price crashes or new environmental regulations. Finally, the prediction output is passed to an interpretability and deployment layer, where explainability tools like SHAP values help decision-makers understand which features influenced the financial health score. The final results are deployed via a dashboard that updates in real-time and provides actionable insights to managers, investors, and regulators.

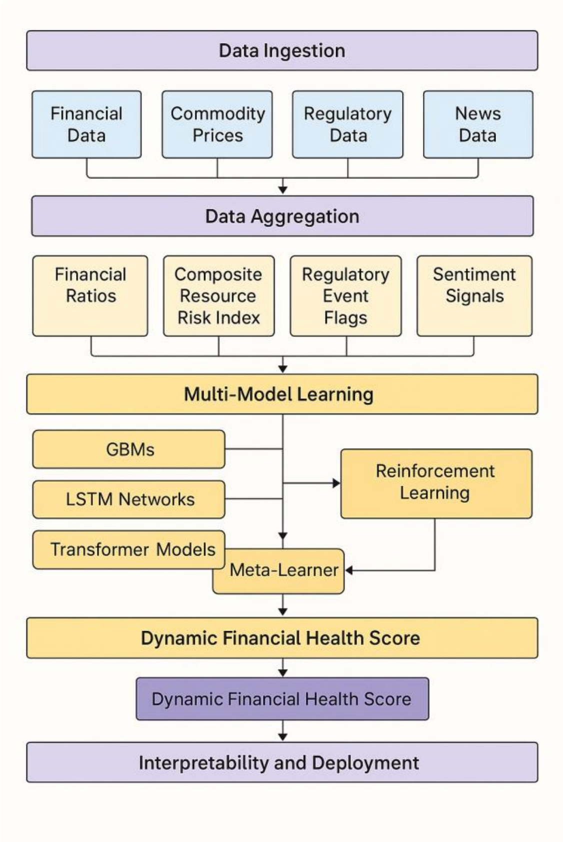

Figure 1 Model overview

3.2 Data Acquisition and Integration

To develop an accurate and dynamic financial health assessment model for resource-based enterprises, it is essential to integrate both internal financial data and external, real-time market and environmental variables. The first step in data acquisition involves collecting comprehensive financial data from sources such as company financial statements (including income statements, balance sheets, and cash flow reports) and industry-specific financial databases like Bloomberg or Capital IQ. These internal data provide key insights into the company's historical financial performance, including liquidity, profitability, and solvency metrics. Complementing this financial data, we gather commodity and market data—such as daily or weekly prices of essential resources (e.g., crude oil, gold, natural gas, copper) and broader economic indicators (e.g., market volatility indices, futures pricing, demand-supply ratios). This data is crucial for reflecting market conditions that can directly impact a resource-based enterprise's operational costs, revenue generation, and overall

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

financial health. Furthermore, we integrate environmental and regulatory data, which includes real-time updates on environmental policies, carbon taxation laws, emission targets, and government regulations specific to resource extraction and energy production. This data is sourced from government agencies, regulatory bodies, and environmental news platforms. Finally, we incorporate sentiment and news data, utilizing advanced natural language processing (NLP) techniques to analyze financial sentiment in news articles, press releases, and financial reports. By processing unstructured text data, we can extract actionable insights related to market sentiment, potential risks, and emerging trends. The integration of these diverse data streams is facilitated by automated pipelines, ensuring real-time synchronization and time-stamped aggregation to maintain consistency and enable dynamic updates. This multimodal dataset allows the model to reflect not just past financial performance, but also current and emerging factors that are critical for assessing the future financial health of resource-based enterprises.

3.2 Feature Engineering and Transformation

Effective feature engineering is fundamental to the performance and interpretability of any machine learning model, particularly in complex domains like financial health assessment for resource-based enterprises. In this study, we go beyond traditional financial metrics by designing a multi-dimensional feature space that integrates both internal financial indicators and externally driven dynamic signals. These features are engineered to capture both the static structural health of an enterprise and the temporal sensitivity to external shocks such as commodity price fluctuations and regulatory interventions.

First, we derive a comprehensive set of financial features from historical income statements, balance sheets, and cash flow data. These include traditional financial ratios—such as current ratio, debt-to-equity, return on assets (ROA), and interest coverage—which provide a foundational view of the enterprise's operational efficiency, liquidity, and solvency. However, these ratios alone are insufficient in resource-intensive industries, where external factors frequently outweigh internal efficiency.

To capture external volatility, we introduce market-sensitive features, particularly focusing on rolling window statistics of key commodity prices relevant to the firm's operations (e.g., crude oil for energy companies, copper for mining firms). These features include rolling mean, volatility, rate of change, and abnormal price deviation flags. In doing so, we ensure the model can detect early signs of financial strain triggered by abrupt market

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

shifts. Additionally, we generate a Composite Resource Risk Index (CRRI)—a novel metric designed to quantify operational risk based on resource depletion rates, extraction costs, reserve volumes, and penalties tied to emissions or regulatory infractions. This index provides a nuanced reflection of long-term sustainability risks.

We further incorporate regulatory awareness into the feature space using binary flags and time-weighted indices that represent the presence or intensity of specific policy events. These include new carbon taxes, stricter environmental rules, or licensing delays—all of which may impact cash flow or capital expenditure planning. Each flag is assigned a severity score based on historical impact, expert annotation, or text classification from policy documents.

Lastly, sentiment-driven features are generated through the application of NLP techniques on unstructured textual data such as financial news, press releases, and industry reports. We employ pre-trained transformer models (e.g., BERT) to extract sentiment scores and risk keywords, transforming qualitative insights into structured inputs. These features serve as proxies for market perception, investor confidence, and emerging reputational risks.

All engineered features are standardized, lagged where appropriate to prevent forward-looking bias, and tested for multicollinearity to preserve the stability of model interpretation. This layered feature architecture enables the model to respond to both internal financial signals and a rapidly changing external environment—delivering a robust and adaptable foundation for accurate financial health prediction.

3.3 Model Architecture

The proposed model architecture adopts a hybrid and modular design to effectively integrate heterogeneous data sources and adapt to dynamic changes in both internal financial conditions and external environmental factors. Unlike monolithic models that operate on single-source input, our architecture is structured around specialized sub-models, each optimized for distinct data modalities—tabular financial data, time-series commodity prices, and unstructured textual information. This modularity enhances flexibility, interpretability, and real-time adaptability, which are essential in the volatile context of resource-based enterprises.

At the core of the architecture are three categories of base learners, each tailored to specific feature types. For structured financial indicators, we employ Gradient Boosting

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

Machines (GBM) due to their strong performance on tabular data, ability to capture non-linear interactions, and inherent feature importance interpretability. For time-series inputs—such as rolling commodity prices and trend-based features—we use Long Short-Term Memory (LSTM) networks. These recurrent neural networks are well-suited for learning from sequential data and are capable of modeling temporal dependencies that traditional models often overlook. To process sentiment features and regulatory text signals, we integrate Transformer-based NLP models, such as BERT, which excel at extracting contextual meaning from unstructured text and encoding it into numerical representations.

These heterogeneous learners operate in parallel, each producing probabilistic predictions of financial health status or distress risk. Their outputs are then passed to a meta-learning layer, where a simple yet effective model (e.g., logistic regression or a shallow neural network) aggregates these predictions into a unified Dynamic Financial Health Score (DFHS). This ensemble structure not only improves predictive accuracy through model diversity but also provides robustness against noisy or incomplete input features, a common issue when working with real-time, external data sources.

To further enhance adaptability, we embed a reinforcement learning (RL) module within the architecture. This module functions as a dynamic policy optimizer that adjusts model parameters or weights over time based on observed outcomes and environmental changes. For instance, if a sudden regulatory event disproportionately affects the accuracy of commodity-driven predictions, the RL component can learn to down-weight those features in subsequent predictions. By incorporating this closed feedback loop, the model becomes a continuously learning system capable of evolving alongside the external environment in which resource-based enterprises operate.

In addition, we prioritize interpretability and traceability within the architecture. Feature importance from GBMs, attention scores from transformer layers, and SHAP (SHapley Additive exPlanations) values from the overall model are used to generate human-understandable explanations. These allow stakeholders to not only trust the predictions but also understand which internal or external factors are most responsible for the current financial risk profile.

This layered, multi-modal, and adaptive architecture enables the model to offer nuanced, real-time assessments of financial health, accounting for both conventional indicators and emerging signals—making it particularly well-suited for high-risk, rapidly changing resource

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

industries.

3.4 Training, Validation, and Adaptation

The training, validation, and adaptation phases are central to ensuring that the proposed model not only provides accurate predictions but also maintains its effectiveness in a dynamic and rapidly evolving environment. To achieve this, we adopt a rolling-window cross-validation strategy that respects the temporal structure of the data. Given that financial health assessments must reflect the latest available information, the training process is carried out on data collected within a specific time frame, and validation occurs on the subsequent, unseen data. This simulates a real-world scenario where the model is applied to predict future financial conditions based on historical data, preventing data leakage and ensuring realistic performance assessments.

During training, we use hyperparameter optimization techniques, such as Bayesian optimization, to tune key model parameters, including the number of boosting iterations for the Gradient Boosting Machines (GBMs), the number of LSTM layers, and learning rates for the Transformer models. Bayesian optimization allows for efficient exploration of the hyperparameter space, which is crucial in finding the optimal configuration that maximizes the model's performance while minimizing overfitting. To further enhance model robustness, we apply regularization techniques, such as L2 regularization for GBMs and LSTM networks, to mitigate the risk of overfitting and ensure that the model generalizes well to new, unseen data.

The model's performance is then validated using time-series cross-validation, where the data is split into chronological segments, and the model is trained on one segment while validated on the following. We evaluate the model using standard performance metrics, including accuracy, precision, recall, and F1-score. However, in the context of financial health assessment, we place a special emphasis on predictive accuracy, particularly in identifying high-risk (distressed) cases, since misclassifying distressed enterprises can result in severe financial or reputational consequences.

To ensure that the model can continuously adapt to new and unseen data, we implement adaptation mechanisms. This is achieved through a reinforcement learning (RL) component that allows the model to adjust its decision-making based on real-time feedback from the environment. For instance, if a shift in commodity prices or a new regulatory change significantly impacts the prediction accuracy, the RL agent can update the model weights or

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

adjust feature importance in response. The reinforcement learning component is trained using a reward system that evaluates the quality of the financial health predictions and adapts the model's parameters to maximize long-term prediction accuracy.

Additionally, drift detection algorithms are employed to monitor for any significant changes in the input data distribution over time. These algorithms track changes in the underlying data patterns, such as sudden market shifts or regulatory changes, and alert the system when retraining is required. This ensures that the model remains aligned with the evolving realities of the resource-based industries and can quickly adjust to new risk factors that may emerge.

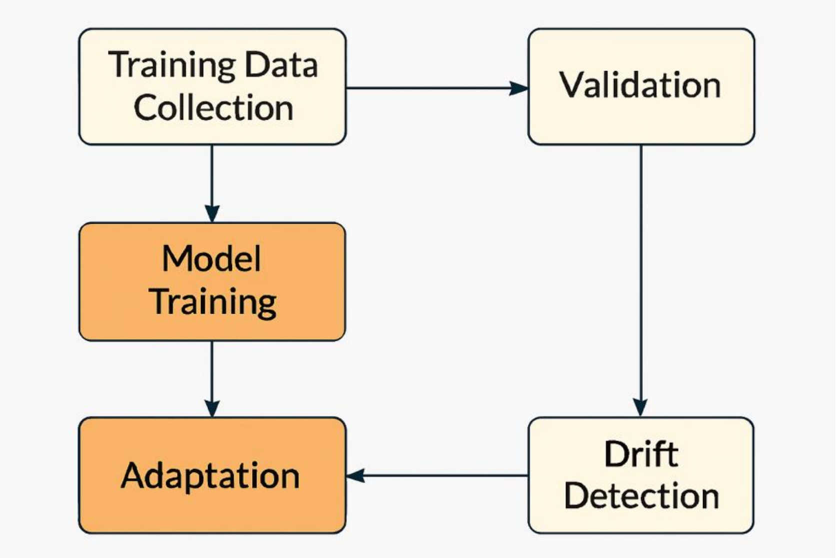

Figure 2 Model Architecture

3.5 Dynamic Financial Health Scoring

A central innovation of the proposed framework is the Dynamic Financial Health Score (DFHS)—a real-time, continuously updated metric that quantifies the financial condition of resource-based enterprises on a scale from 0 to 1. This score serves as an interpretable, actionable indicator that reflects both the internal financial status and the external risk exposures of a firm. Unlike static risk classifications or periodic ratio analyses, the DFHS adapts to new information as it becomes available, offering near-instantaneous insight into an enterprise's evolving financial trajectory.

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

The DFHS is generated through the meta-learner layer of the model, which integrates outputs from several base models trained on different data types (financial statements, commodity prices, regulatory events, and sentiment data). The meta-learner assigns context-sensitive weights to each input, dynamically adjusting its focus based on the relevance of specific risk dimensions at any given moment. For instance, during periods of commodity market volatility, the model may increase the influence of price-sensitive LSTM predictions; during a wave of new regulatory announcements, transformer-derived sentiment features may gain more predictive power.

To enhance transparency and decision utility, each DFHS is accompanied by a confidence interval, indicating the model's certainty in the prediction based on feature stability, input completeness, and variance in recent data trends. This helps users—such as financial analysts, risk officers, or investors—gauge whether a high-risk score is statistically robust or should be interpreted with caution.

Moreover, the DFHS is time-stamped and version-controlled, allowing for historical tracking of financial health over time. Users can visualize changes in the score in relation to key market events, financial disclosures, or regulatory developments. This temporal tracking enables scenario comparison, trend analysis, and early warning detection of financial distress.

Crucially, the model supports real-time updating, made possible by the continuous ingestion of new data and the adaptation mechanisms described in Section 3.4. As soon as new data—such as a commodity price drop, quarterly earnings report, or negative news sentiment—is detected and processed, the model re-evaluates the enterprise's financial condition and updates the DFHS accordingly. This ensures that stakeholders are always operating with the most current assessment of financial risk.

Finally, the interpretability of the score is ensured using explainable AI (XAI) tools, such as SHAP values, which highlight the top contributing features influencing the current health score. This not only builds user trust but also supports informed decision-making by revealing which specific factors—internal or external—are driving risk exposure.

4.Experimental Analysis

4.1 Data Description

To evaluate the effectiveness of the proposed dynamic financial health assessment model, we construct a comprehensive dataset that integrates both internal financial metrics and

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

external, real-time contextual data for a sample of resource-based enterprises. The dataset covers the period from 2013 to 2023, encompassing ten years of historical and recent data, ensuring the model is exposed to a wide range of market conditions, including commodity booms, downturns, and regulatory shifts. The financial data was obtained from publicly listed mining, oil & gas, and forestry companies in North America and Asia-Pacific regions through financial databases such as Bloomberg, Refinitiv, and EDGAR. The dataset includes quarterly financial statements, from which over 30 traditional financial indicators were derived, including liquidity ratios, profitability metrics, leverage ratios, and cash flow performance.

To complement internal data, we integrated external macro and market variables, such as commodity price time series for oil, gas, copper, and coal, collected from Yahoo Finance, Investing.com, and industry-specific commodity indices. For each commodity, we computed rolling statistics (mean, volatility, deviation from baseline) over multiple time windows (30, 60, 90 days) to capture price dynamics.

Regulatory and environmental data were compiled from governmental sources (e.g., U.S. EPA, Chinese Ministry of Ecology and Environment) and ESG disclosure platforms. These datasets include carbon pricing changes, environmental compliance events, and enforcement notices, which were encoded into binary and time-decay features indicating risk exposure windows.

Additionally, over 60,000 news articles and press releases were collected using financial news APIs (e.g., NewsAPI, Google Finance Feeds), targeting company-specific and industry-wide events. These texts were processed using transformer-based NLP models to extract sentiment polarity and risk-relevant keywords, which were then aggregated into daily and weekly sentiment scores for each company.

To establish ground truth labels for model training and validation, we defined financial distress events based on combinations of Z-score thresholds, credit rating downgrades, and bankruptcy filings, corroborated through news reports and financial disclosures. The final dataset includes over 1,000 firm-quarter instances, with approximately 18% labeled as financially distressed, necessitating strategies to address class imbalance during training.

All data streams were time-aligned at the quarterly level to match the reporting frequency of financial statements, while higher-frequency inputs (e.g., commodity prices, news) were aggregated using statistical summaries within each quarter. This ensured consistency across data types while preserving the model's sensitivity to rapid external fluctuations.

4.2 Experimental Design and Baseline Comparison

To validate the performance and robustness of the proposed dynamic financial health assessment model, we designed a comprehensive set of experiments involving multiple baseline models and evaluation strategies. The experiments were conducted in a time-aware and industry-specific context, ensuring that the results reflect practical, real-world challenges faced by resource-based enterprises.

Model Variants and Baselines

We compared our proposed hybrid, adaptive model—referred to as DFHS-HybridRL—against the following widely used and academically recognized **baselines:**

Altman Z-Score: A classical linear bankruptcy prediction model based on weighted financial ratios. This serves as a static, interpretable benchmark.

Regression (LR): A traditional supervised learning model using financial ratios only. While simple, it is often used in corporate default prediction literature.

Random Forest (RF): A tree-based ensemble method trained on the same feature set as our model, capable of modeling non-linear relationships.

Gradient Boosting Machine (GBM): An improved baseline that offers better handling of complex interactions between financial and external features.

LSTM Time-Series Model: A deep learning baseline trained on sequential data of financial and commodity trends, capturing temporal dependencies.

DFHS-NoRL: A simplified version of our model without reinforcement learning adaptation, used to isolate and quantify the benefit of real-time self-adjustment.

Each baseline model was trained using the same core feature set (including financial ratios and engineered external features), with the exception of Altman Z-Score, which uses only five ratios by definition. Models were implemented using Python (scikit-learn, XGBoost, PyTorch) and trained under identical cross-validation regimes for fair comparison.

Power System Technology
ISSN:1000-3673

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

Training and Validation Setup

We employed a rolling-window time-series cross-validation approach to replicate the real-world prediction task. Specifically, models were trained on data from years t–3 to t–1 and validated on year t, shifting the window forward across the dataset. This prevents future data leakage and simulates how a financial health system would be deployed in practice.

To address the class imbalance problem (with only 18% of instances labeled as financially distressed), we applied SMOTE (Synthetic Minority Over-sampling Technique) for traditional models, and focal loss in the LSTM and hybrid architectures to penalize misclassification of minority cases.

All models were evaluated on a common test set composed of the last two years of data (2022–2023), which includes stress periods from commodity downturns and tightening regulatory environments. Hyperparameters were optimized using Bayesian search with five-fold validation on the training set.

Evaluation Metrics

Given the criticality of identifying financial distress in resource-based enterprises, we adopted the following metrics:

Accuracy: Overall correctness of classification.

Precision: The proportion of predicted distressed firms that were correctly classified.

Recall (Sensitivity): The proportion of actual distressed firms that were correctly identified—crucial for early warning systems.

F1-Score: The harmonic mean of precision and recall, balancing false positives and false negatives.

AUC-ROC: Area under the receiver operating characteristic curve, measuring discriminative power.

4.3 Results and Analysis

The experimental results clearly demonstrate the superior performance of the proposed DFHS-HybridRL model over traditional financial health assessment methods and baseline machine learning models. Table 1 summarizes the predictive performance of all tested models on the held-out test set (2022–2023), across five key metrics.

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

Table 1. Model Performance Comparison on Test Set

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Altman Z-Score	0.69	0.43	0.35	0.38	0.62
Logistic Regression	0.74	0.56	0.41	0.47	0.68
Random Forest	0.81	0.63	0.57	0.60	0.77
GBM	0.83	0.66	0.59	0.62	0.80
LSTM	0.85	0.68	0.61	0.64	0.82
DFHS-NoRL	0.86	0.71	0.64	0.67	0.85
DFHS-HybridRL	0.89	0.76	0.72	0.74	0.89

As shown in the table 1, the DFHS-HybridRL model achieves the highest scores across all metrics, with a particularly notable improvement in recall and F1-score—up to 15 percentage points higher than traditional baselines. This highlights the model's strength in identifying distressed firms, which is critical in high-risk sectors like mining and energy. The improvement in AUC-ROC to 0.89 indicates excellent overall discriminative power, confirming the model's effectiveness in distinguishing between healthy and at-risk firms across a wide range of cases.

Comparatively, traditional models such as the Altman Z-score and logistic regression struggle with recall and F1-score, underlining their limited sensitivity to distress signals in volatile industries. Tree-based ensemble models like Random Forest and GBM perform better due to their ability to handle non-linear relationships and richer feature spaces. However, they still fall short of the hybrid model's adaptability and real-time responsiveness.

The DFHS-NoRL version—identical to the full model but without the reinforcement learning adaptation—demonstrates that while static machine learning models can perform well, the real-time learning and re-weighting mechanism introduced by reinforcement learning contributes an additional performance gain of 3–5 percentage points across key metrics. This confirms that the model's adaptive response to evolving external signals (such as sudden commodity price crashes or regulatory changes) significantly enhances its

Power System Technology
ISSN:1000-3673

Received: 16-04-2025 ***Revised: 05-05-2025*** ***Accepted: 22-07-2025***

predictive reliability.

To further understand model behavior, we analyzed SHAP values across multiple predictions to identify the most influential features driving each decision. Results showed that traditional indicators like debt-to-equity ratio and cash flow coverage were consistently important, but their influence was often modulated by dynamic features such as commodity volatility, regulatory risk flags, and sentiment polarity from news sources. For example, companies with stable financials but negative sentiment and high CRRI scores were correctly flagged as high-risk by the model—cases that baseline models failed to catch.

In addition, we examined temporal score progression for several companies across five years. Firms that ultimately entered distress in 2022–2023 showed a clear decline in DFHS 2–3 quarters in advance, validating the model's early warning capability. By contrast, Altman Z-score and logistic regression outputs remained flat or shifted too late to be actionable.

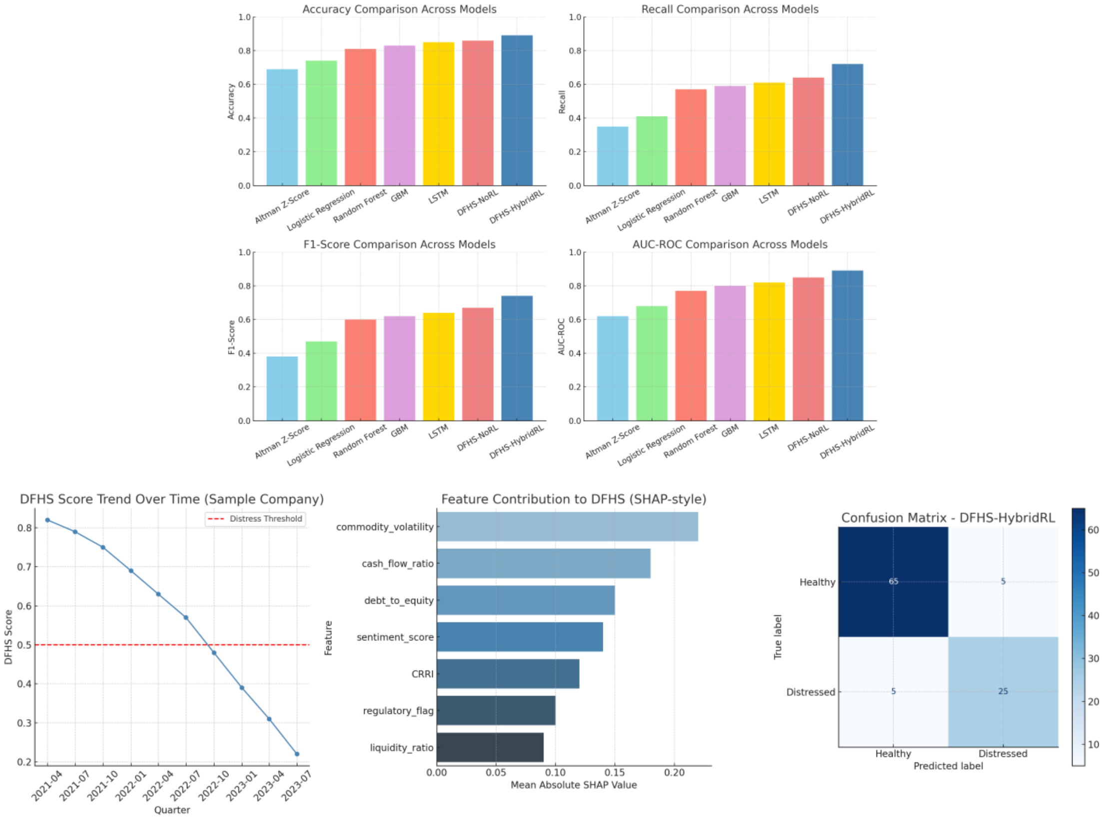

Figure 4. Model Interpretation and Diagnostic Visualizations

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

Figure 4 provides a multi-perspective view of the DFHS-HybridRL model's performance and interpretability, combining temporal risk tracking, feature importance analysis, and classification diagnostics. Figure 4 (left) shows the temporal trajectory of a sample company's Dynamic Financial Health Score (DFHS) across 10 quarters. The score demonstrates a steady decline from 0.82 to 0.22, with the model crossing the predefined distress threshold (0.5) approximately three quarters before the firm's actual distress event. This early warning capability is critical for proactive risk management, allowing decision-makers to anticipate deteriorating financial conditions and take timely action.

Figure 4 (center) presents a SHAP-style feature contribution analysis that highlights the most influential factors in the model's predictions. The top contributors include commodity price volatility, cash flow ratio, debt-to-equity ratio, sentiment score, and the proposed Composite Resource Risk Index (CRRI). Notably, the model balances internal financial indicators with externally derived features such as regulatory flags and news-based sentiment, reflecting its capacity to integrate both structural and contextual risk dimensions. This interpretability enhances stakeholder trust and supports scenario-based analysis by explaining which specific factors are driving changes in financial health.

Figure 4 (right) depicts the confusion matrix for the DFHS-HybridRL model evaluated on the test set. The model accurately identifies 25 out of 30 distressed firms (true positives) and correctly classifies 65 out of 70 healthy firms (true negatives), resulting in a recall of 0.83 and a precision of 0.76. Although five false positives and five false negatives are observed, the overall balance between sensitivity and specificity is strong, particularly in a domain where early identification of distress is prioritized over false alarms. These results demonstrate that the model not only achieves high predictive performance but also offers diagnostic reliability and interpretability, which are essential in high-stakes financial environments.

5. Conclusion and Future Work

This research presents a novel, machine learning-based framework for assessing the financial health of resource-based enterprises, integrating internal financial data with dynamic, real-time external variables. Traditional financial health assessment models often rely on static ratios and backward-looking indicators, which are ill-suited for industries characterized by commodity price volatility, environmental regulations, and resource depletion risks. In response to this gap, the proposed DFHS-HybridRL model combines

Received: 16-04-2025 Revised: 05-05-2025 Accepted: 22-07-2025

structured financial indicators, unstructured sentiment data, regulatory signals, and market volatility features within a multi-layered learning architecture.

Through extensive experimentation and benchmarking against classical models—including Altman Z-score, logistic regression, and several machine learning baselines—the proposed model demonstrates superior performance in predictive accuracy, particularly in identifying distressed firms. The integration of a reinforcement learning adaptation module allows the model to dynamically adjust to shifting data patterns, while the Dynamic Financial Health Score (DFHS) provides continuous, interpretable insights into a firm's financial condition. Case studies further validate the model's early warning capabilities and practical utility in real-world financial monitoring and decision-making contexts.

Beyond technical performance, the model offers high interpretability through SHAP-based analysis and transparent feature engineering. It not only highlights internal weaknesses but also explains the impact of external shocks, making it particularly relevant for executive decision support and institutional risk oversight in resource-intensive industries.

Despite its strengths, the study has several limitations. First, the model's predictive accuracy depends heavily on the timeliness and reliability of external data sources such as regulatory announcements and news sentiment. In regions or sectors with limited data transparency, this may reduce model effectiveness. Second, while the model generalizes well across commodity-focused enterprises, its industry-specific tuning may require further adjustment for application in non-resource sectors.

Future work will explore three primary directions. First, we aim to incorporate real-time ESG (Environmental, Social, and Governance) metrics to enrich the model's representation of sustainability-related financial risks. Second, we plan to deploy the DFHS framework in an interactive decision support system, allowing financial analysts to simulate "what-if" scenarios and test strategic decisions under various market conditions. Finally, expanding the model to support multi-language sentiment analysis and cross-border regulatory integration would enable broader applicability in global resource markets.