A Multimodal Generative AI Framework for Early Warning of Corporate Credit Rating Transitions Using Financial Analytics, Annual Report Intelligence, and News Sentiment

Chandrasekhar Subramanyam

Vijay bhoomi University

DOI: https://doi.org/10.51584/IJRIAS.2026.11070020

Received: 10 July 2026; Accepted: 15 July 2026; Published: 28 July 2026

ABSTRACT

Corporate credit ratings play a critical role in lending decisions, investment analysis, portfolio management, and regulatory compliance. Traditional credit rating methodologies rely predominantly on structured financial information and expert judgement (Altman, 1968; Hand & Henley, 1997), while the rapidly growing volume of unstructured corporate disclosures and financial news remains underutilized. Recent advances in Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs) (Devlin et al., 2019; Brown et al., 2020) provide new opportunities to automatically analyse qualitative information contained in annual reports and financial news and integrate it with conventional financial indicators.

This paper proposes a multimodal Artificial Intelligence framework for the early warning of corporate credit rating transitions by combining quantitative financial analytics with qualitative information extracted from annual reports and financial news. The proposed framework consists of three complementary modules: (i) quantitative financial analysis based on trend-weighted financial ratios using Exponential Moving Average (EMA) smoothing and sector-specific normalization, (ii) an Annual Report Intelligence Pipeline that extracts and summarizes six strategically important sections of annual reports and transforms them into sentiment-based and semantic features, and (iii) a News Intelligence Module that captures recent developments affecting corporate credit quality through dynamic news sentiment analysis. These heterogeneous information sources are integrated using a hybrid deep learning architecture to generate an explainable early warning score representing the probability of future credit rating transition.

The methodology is demonstrated through a case study of Craftsman Automation Limited, illustrating how qualitative disclosures and contemporary news complement traditional financial analysis in identifying changes in corporate credit quality. The proposed framework seeks to provide financial institutions, banks, NBFCs, and credit rating agencies with an explainable decision-support system capable of improving proactive credit surveillance, in line with recent calls for interpretable AI in credit risk management (de Lange et al., 2022; Castro Vieira et al., 2025).

Keywords: Credit Rating Transition; Early Warning System; Generative AI; Large Language Models; Financial Analytics; Annual Report Intelligence; Financial News Sentiment; Deep Learning; Explainable Artificial Intelligence; Credit Risk Assessment.

INTRODUCTION

The quality of credit risk assessment directly influences the stability of financial institutions and capital markets. Banks, non-banking financial companies (NBFCs), mutual funds, insurance companies, and pension funds depend extensively on corporate credit ratings while making lending and investment decisions. A deterioration in credit quality often precedes financial distress, making early identification of rating migration one of the most important challenges in modern financial risk management.

Traditional credit rating methodologies have evolved over several decades and have demonstrated considerable success in evaluating the long-term creditworthiness of borrowers, beginning with discriminant-analysis-based bankruptcy prediction models (Altman, 1968) and extending to statistical credit-scoring frameworks (Hand & Henley, 1997; Thomas et al., 2017). These methodologies primarily rely on financial statement analysis, business risk evaluation, industry outlook, and expert judgement. Financial ratios such as leverage, profitability, liquidity, debt servicing capacity, and cash flow adequacy form the quantitative foundation of most rating models. While these approaches have proven effective, they exhibit two important limitations.

First, conventional ratio-based models primarily analyse historical financial performance and therefore react slowly to emerging risks. Financial statements provide a backward-looking representation of corporate performance and often fail to capture rapidly changing business conditions. Credit deterioration frequently begins long before it becomes evident in published financial statements (Altman et al., 2017).

Second, significant information influencing future credit quality resides within unstructured documents such as annual reports, management discussions, auditor observations, and financial news. These qualitative sources contain valuable forward-looking information regarding business strategy, operational risks, governance practices, regulatory developments, and management expectations (Li, 2010; Loughran & McDonald, 2011). However, extracting actionable intelligence from such documents has traditionally required extensive manual analysis by experienced credit analysts.

Recent advances in Large Language Models (LLMs) and Generative Artificial Intelligence (GenAI) have significantly enhanced the capability of machines to understand, summarize, and analyse large volumes of textual information (Devlin et al., 2019; Brown et al., 2020; Touvron et al., 2023). These technologies offer the possibility of converting qualitative corporate disclosures into structured numerical representations that can complement conventional financial indicators. Rather than replacing financial analysis, GenAI enables the systematic incorporation of qualitative evidence into credit risk assessment, and recent evidence suggests that even general-purpose LLMs already show meaningful competence at financial data-analysis tasks (Chou et al., 2025).

Another important development is the increasing availability of digital financial news. News articles frequently report corporate events such as mergers, acquisitions, management changes, regulatory actions, litigation, supply chain disruptions, customer wins, and macroeconomic developments before these events are reflected in financial statements. Financial news sentiment has long been shown to carry information relevant to market behaviour (Tetlock, 2007), and more recent machine-learning evidence continues to support its predictive value, albeit with important caveats about signal strength and noise (Ding et al., 2015; Davidovic & McCleary, 2025). Consequently, financial news represents a valuable leading indicator for anticipating changes in corporate credit quality.

This research proposes that financial ratios, annual report disclosures, and financial news should not be analysed independently. Instead, they should be integrated within a unified multimodal Artificial Intelligence framework capable of producing an explainable early warning signal for corporate credit rating transitions. The proposed framework mirrors the reasoning process of experienced credit analysts who simultaneously evaluate historical financial performance, management quality, business risks, and recent corporate developments before arriving at a credit opinion.

Unlike many existing AI-based credit scoring models that operate as black-box classifiers, the proposed framework emphasizes explainability by preserving the contribution of each information source, echoing the regulatory and practical case for explainable AI (XAI) in credit assessment made by de Lange et al. (2022) and the systematic evidence on fairness and interpretability limitations surveyed by Castro Vieira et al. (2025). Financial analytics quantify historical performance, annual report intelligence captures management perspective and strategic direction, while financial news reflects the most recent external developments affecting credit quality. Their combined interpretation provides a richer understanding of future rating migration than any individual source alone.

The proposed methodology has direct applications in credit appraisal, portfolio monitoring, internal rating systems, Basel-compliant risk management (Basel Committee on Banking Supervision, 2024), Expected Credit Loss (ECL) estimation (Reserve Bank of India, 2026), and early warning systems used by banks and financial institutions. By integrating structured and unstructured information, the framework seeks to improve the timeliness and reliability of credit risk assessment while maintaining transparency in the decision-making process.

Motivation

Corporate credit rating transitions rarely occur abruptly. They are usually preceded by gradual deterioration or improvement in financial performance, management commentary, business risks, and external market developments. Conventional financial ratio analysis often identifies these changes only after they become visible in published financial statements.

Experienced credit analysts therefore supplement financial analysis with qualitative assessment of annual reports, management discussions, auditor observations, and industry developments. However, this process is labour-intensive, subjective, and difficult to scale across large corporate loan portfolios.

The emergence of Generative AI presents an opportunity to automate significant portions of this analytical process. By transforming narrative corporate disclosures into structured quantitative features and combining them with traditional financial indicators, it becomes possible to construct an intelligent early warning framework capable of continuously monitoring corporate credit quality.

Research Objectives

The primary objective of this research is to develop a multimodal Artificial Intelligence framework capable of providing early warning of corporate credit rating transitions through the integration of structured financial information and unstructured textual information. The specific objectives are:

- To develop a trend-aware quantitative financial analytics framework using multi-year financial ratios.
- To design an Annual Report Intelligence Pipeline capable of extracting credit-relevant information from annual reports.
- To incorporate financial news sentiment as a leading indicator of emerging credit risk.
- To integrate heterogeneous information sources using a hybrid deep learning architecture.
- To generate an explainable early warning score indicating the probability of future credit rating transition.

Research Questions

The study addresses the following research questions:

- RQ1: Can historical financial ratios alone provide sufficient early warning of corporate credit rating transitions?
- RQ2: Does qualitative information extracted from annual reports improve prediction of future rating changes?
- RQ3: Can financial news serve as an effective leading indicator of emerging credit risk?
- RQ4: Does a multimodal AI framework integrating financial analytics, annual reports, and financial news outperform traditional ratio-based approaches in generating explainable early warning signals?

Major Contributions

This research makes four principal contributions.

Contribution 1

A novel multimodal framework integrating structured financial information, annual report intelligence, and financial news for early warning of corporate credit rating transitions.

Contribution 2

An Annual Report Intelligence Pipeline that transforms narrative corporate disclosures into structured quantitative credit-risk indicators through automatic section extraction, summarization, and sentiment analysis.

Contribution 3

A hybrid feature fusion methodology combining quantitative financial analytics with qualitative semantic representations generated using Large Language Models.

Contribution 4

A practical case study using Craftsman Automation Limited demonstrating how multimodal information improves the interpretability of early warning assessment compared with conventional financial ratio analysis.

LITERATURE REVIEW

Introduction

Corporate credit rating prediction has attracted significant attention from researchers, financial institutions, and rating agencies owing to its central role in credit risk management and financial stability. Traditional rating methodologies developed by major rating agencies combine quantitative financial analysis with qualitative assessment of business risk, management quality, industry outlook, and corporate governance. Although these approaches have demonstrated satisfactory predictive performance over several decades, the increasing availability of large volumes of digital financial information has motivated researchers to investigate Artificial Intelligence (AI) and Machine Learning (ML) techniques for improving the accuracy and timeliness of credit risk assessment (Lessmann et al., 2015).

Recent developments in deep learning and Generative Artificial Intelligence (GenAI) have further expanded the scope of credit analytics by enabling automated analysis of unstructured textual information such as annual reports, earnings releases, financial news, and regulatory filings. Despite these advances, most existing studies continue to examine structured financial information and unstructured textual information independently. Limited research has investigated their systematic integration within a unified and explainable framework for early warning of corporate credit rating transitions.

This section reviews the literature under six complementary themes: (i) traditional credit rating methodologies, (ii) quantitative financial analytics, (iii) machine learning approaches for credit rating prediction, (iv) natural language processing of corporate disclosures, (v) Generative AI in financial applications, and (vi) multimodal learning for financial risk assessment.

Traditional Credit Rating Methodologies

Corporate credit ratings represent independent assessments of an organization's ability to meet its long-term financial obligations. Rating agencies evaluate several dimensions of corporate performance, including financial strength, capital structure, business risk, industry position, governance quality, and management capability. Financial ratios such as leverage, profitability, liquidity, interest coverage, debt servicing capacity, and operating cash flows have historically formed the quantitative basis of these assessments, dating back to Altman's (1968) discriminant-analysis bankruptcy model and subsequently formalized in statistical credit-scoring reviews such as Hand and Henley (1997) and Thomas et al. (2017).

Traditional rating methodologies rely heavily on expert judgment. Analysts evaluate not only historical financial performance but also management quality, strategic direction, competitive positioning, regulatory environment, and macroeconomic conditions. This combination of quantitative and qualitative evaluation has enabled rating agencies to develop robust assessment frameworks that are widely accepted by lenders and investors.

However, conventional methodologies exhibit two inherent limitations. First, they depend primarily on historical financial statements that reflect past performance rather than emerging business conditions (Altman et al., 2017). Second, qualitative assessment remains largely manual, introducing subjectivity and limiting scalability when monitoring large corporate portfolios. These limitations motivate the development of AI-assisted approaches capable of automating qualitative analysis while preserving explainability.

Financial Ratio-Based Prediction Models

Financial ratio analysis remains one of the most widely used techniques for evaluating corporate creditworthiness. Ratios related to profitability, leverage, liquidity, operational efficiency, and debt servicing capacity have been extensively employed for bankruptcy prediction, default forecasting, and credit rating estimation since Altman's (1968) original Z-score model.

Early statistical approaches primarily relied on discriminant analysis and logistic regression using selected financial indicators. Subsequent research demonstrated that machine learning algorithms, including decision trees, support vector machines, random forests, gradient boosting, and artificial neural networks, often outperform traditional statistical models when complex nonlinear relationships exist among financial variables (Lessmann et al., 2015).

Although these models achieve satisfactory predictive accuracy under stable economic conditions, they exhibit several shortcomings. First, financial ratios primarily describe historical corporate performance and therefore respond slowly to rapid changes in business conditions. Second, financial statements are published periodically, creating a time lag between operational events and their representation in financial reports (Altman et al., 2017). Third, ratio-based models frequently ignore qualitative information that experienced credit analysts routinely consider during rating assessments (Thomas et al., 2017).

Consequently, financial ratios alone may not provide sufficiently early warning of impending credit rating transitions — a proposition that motivates Research Question 1 of the present study.

Artificial Intelligence in Credit Rating Prediction

The increasing availability of corporate financial data has encouraged the application of Artificial Intelligence techniques to credit risk assessment. Artificial Neural Networks were among the earliest AI models adopted for credit scoring because of their ability to model nonlinear relationships among financial variables. More recently, ensemble learning techniques and deep neural networks have demonstrated improved predictive performance across several benchmark credit-scoring datasets (Lessmann et al., 2015; Baesens et al., 2016).

Deep learning models are capable of learning complex interactions among numerous financial indicators without requiring extensive manual feature engineering. Autoencoders, recurrent neural networks, convolutional neural networks, and transformer-based architectures have all been explored for financial prediction tasks (Goodfellow et al., 2016; Vaswani et al., 2017).

Despite their predictive capability, most deep learning approaches suffer from limited interpretability. Financial institutions often require transparent explanations supporting model recommendations, particularly for regulatory compliance and internal audit. Consequently, explainable AI has emerged as an important research direction in financial risk management. De Lange et al. (2022), for instance, combine a LightGBM classifier with SHAP-based explanations for consumer loan default prediction, demonstrating that explainability and predictive accuracy need not be mutually exclusive. More broadly, Castro Vieira et al. (2025) systematically review methods for identifying and mitigating bias in AI-based credit-granting models, underscoring that transparency and fairness remain open challenges for AI adoption in regulated lending.

Another limitation is that most AI-based rating models continue to rely almost exclusively on structured numerical data. Qualitative information available in corporate reports and financial news remains insufficiently integrated into existing prediction frameworks — a gap the present study directly addresses.

Natural Language Processing of Corporate Disclosures

Corporate annual reports contain substantial qualitative information regarding business strategy, operational performance, future opportunities, governance practices, and risk management. Sections such as the Chairman's Address, Directors' Report, Management Discussion and Analysis (MD&A), Independent Auditors' Report, and Business Risk disclosures often contain forward-looking information that is not immediately reflected in financial statements (Li, 2010).

Early research in financial Natural Language Processing focused primarily on dictionary-based sentiment analysis and frequency analysis of positive and negative words (Loughran & McDonald, 2011). While these methods provided useful insights, they were unable to capture contextual meaning, semantic relationships, or nuanced managerial communication. Related work has also shown that the tone of qualitative disclosures can itself be strategically managed by firms, independent of concurrent quantitative performance (Huang et al., 2014), underscoring the need for context-aware, rather than purely lexical, sentiment extraction.

Subsequent developments in word embeddings, recurrent neural networks, and transformer architectures substantially improved contextual understanding of financial text (Devlin et al., 2019). Domain-specific language models further enhanced the ability to analyse corporate disclosures by incorporating financial vocabulary and accounting terminology.

However, most existing studies analyse only one or two sections of annual reports, typically the MD&A or Chairman's statement. Limited attention has been given to systematically integrating multiple sections of annual reports into a unified credit assessment framework.

The present study addresses this limitation by proposing an Annual Report Intelligence Pipeline that simultaneously analyses six strategically important sections of corporate annual reports and converts their qualitative content into structured numerical features suitable for multimodal learning.

Generative Artificial Intelligence in Financial Applications

Recent advances in Large Language Models (LLMs) have significantly expanded the capability of Artificial Intelligence systems to process complex financial documents. Models based on transformer architectures (Vaswani et al., 2017) demonstrate remarkable performance in text summarization, question answering, semantic retrieval, document classification, and information extraction, with large-scale few-shot models such as GPT-3 (Brown et al., 2020), open foundation models such as LLaMA (Touvron et al., 2023), and successor systems such as GPT-4 (OpenAI, 2023) substantially expanding what is practically achievable.

Within financial services, Generative AI has been investigated for automated report generation, contract analysis, financial document summarization, customer service, regulatory compliance, and investment research. Chou et al. (2025), for example, provide an early empirical evaluation of a frontier multimodal model (ChatGPT-4o) on financial data-analysis tasks, finding performance broadly comparable to traditional statistical software but with important caveats around reproducibility and the continued need for human validation. The ability of LLMs to understand context rather than isolated keywords represents a significant improvement over traditional sentiment analysis methods.

Nevertheless, several challenges remain. General-purpose language models may lack sufficient understanding of specialized financial terminology, and unrestricted generative models may occasionally produce inaccurate or non-verifiable responses. Consequently, recent research increasingly emphasizes Retrieval-Augmented Generation (RAG) (Lewis et al., 2020), domain-specific prompting, and explainable AI techniques to improve reliability in financial applications.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

Most published studies have focused on demonstrating the capabilities of LLMs for isolated financial tasks. Comparatively little research has explored how GenAI-generated qualitative intelligence can be systematically integrated with quantitative financial analytics for predictive credit risk assessment. This integration forms one of the principal motivations of the proposed research.

Financial News as a Leading Indicator of Credit Quality

Corporate credit quality is influenced not only by internal financial performance but also by external events such as regulatory actions, mergers and acquisitions, management changes, litigation, supply chain disruptions, technological innovation, geopolitical developments, and macroeconomic shocks. These events are often reported through financial news well before they become visible in published financial statements.

Tetlock's (2007) foundational study demonstrated that media-based sentiment measures carry information relevant to asset market behaviour, and subsequent research has extended this line of enquiry using deep learning methods for event-driven stock prediction (Ding et al., 2015). More recent large-scale evidence, however, presents a more nuanced picture: Davidovic and McCleary (2025), analysing close to 1.86 million news headlines, find that explicit textual sentiment carries only modest and inconsistent predictive power for market direction, while forward-looking implied sentiment (such as the VIX) captures a substantially larger share of return variation. This suggests that naïve sentiment averaging, of the kind used in early financial NLP studies, may fail to capture emerging changes in corporate credit quality reliably.

The present research addresses this issue by introducing a News Intelligence Module that combines sentiment analysis with temporal aggregation to generate a dynamic news-based indicator representing recent changes in corporate credit risk, rather than relying on unweighted sentiment averages.

Multimodal Learning in Financial Risk Assessment

Multimodal learning refers to the integration of heterogeneous information sources within a unified learning framework. Recent advances in deep learning have demonstrated the advantages of combining structured numerical variables with unstructured textual information across several application domains.

Within finance, multimodal learning remains an emerging research area. Existing studies have primarily focused on combining financial ratios with limited textual features extracted from earnings calls or social media sentiment (Ding et al., 2015; Davidovic & McCleary, 2025). Comprehensive integration of financial statements, annual reports, financial news, and explainable AI remains relatively unexplored, and is only beginning to be addressed by studies such as de Lange et al. (2022) in the narrower context of consumer lending.

The present study extends multimodal financial learning by integrating three complementary information sources:

- Structured financial analytics, representing historical corporate performance.
- Annual report intelligence, representing management perspective, governance quality, strategic direction, and business risks.
- Financial news intelligence, representing current external developments affecting future credit quality.

This integration seeks to emulate the reasoning process of experienced credit analysts while maintaining transparency through explainable feature engineering.

Critical Research Gaps

The literature review reveals several important gaps that motivate the present research.

- Most existing credit rating models rely predominantly on structured financial ratios and underutilize qualitative corporate disclosures (Altman et al., 2017; Thomas et al., 2017).

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

- Studies employing Natural Language Processing generally analyse isolated textual sources rather than multiple complementary sections of annual reports (Li, 2010; Loughran & McDonald, 2011).
- Financial news has rarely been integrated with annual report analysis within a unified predictive framework for credit rating transition, and the predictive value of raw news sentiment is itself contested (Davidovic & McCleary, 2025).
- Existing deep learning models frequently function as black boxes, limiting their practical acceptance by regulated financial institutions (Castro Vieira et al., 2025).
- Few studies have proposed an end-to-end framework that combines quantitative financial analytics, annual report intelligence, financial news, multimodal feature fusion, and explainable AI for early warning of corporate credit rating transitions (de Lange et al., 2022).

These research gaps form the foundation of the proposed methodology presented in the next section.

Proposed Multimodal Framework for Early Warning of Corporate Credit Rating Transitions

Overview

The literature review demonstrates that existing approaches to corporate credit rating prediction generally analyse financial ratios, annual reports, or financial news independently. While each information source provides useful insights into corporate credit quality, none individually offers a comprehensive representation of the dynamic factors influencing future credit rating transitions.

This research proposes a Multimodal Generative Artificial Intelligence Framework that integrates structured financial analytics with qualitative information extracted from corporate annual reports and financial news. The underlying hypothesis is that corporate credit quality evolves gradually and is reflected simultaneously in financial performance, management disclosures, business risks, governance practices, auditor observations, and external market developments. By combining these heterogeneous information sources within a unified learning framework, the proposed methodology aims to generate an explainable early warning signal that precedes observable changes in external credit ratings.

Unlike conventional credit rating models that primarily rely on historical financial statements, the proposed framework captures both lagging indicators, represented by financial ratios, and leading indicators, represented by management disclosures and recent financial news. The resulting multimodal representation provides a richer and more timely assessment of corporate credit quality.

The overall framework consists of five sequential modules:

- Quantitative Financial Analytics
- Annual Report Intelligence Pipeline
- Financial News Intelligence
- Multimodal Feature Fusion
- Deep Learning-Based Early Warning Engine

The interaction among these modules, from raw data input through to the final explainable output, is illustrated in Figure 1.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

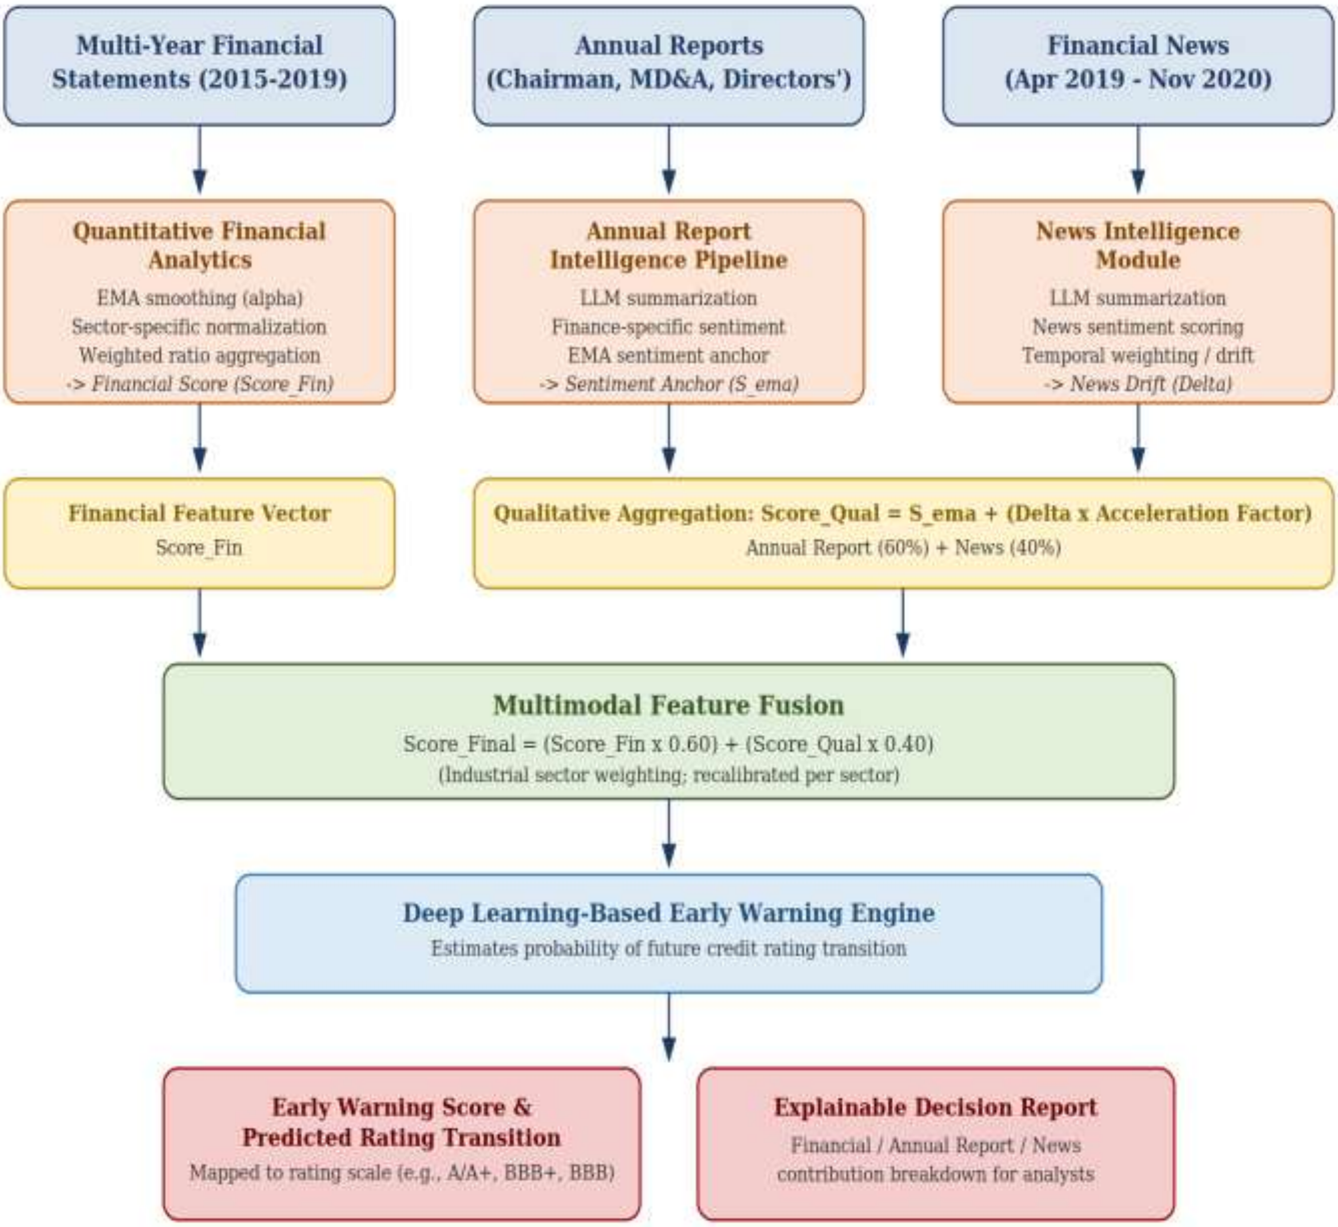

Figure 1. Overall Multimodal Framework: From Data Input (financial statements, annual reports, financial news) Through Quantitative and Qualitative Processing, Multimodal Fusion, and the Deep Learning Early Warning Engine to the Final Explainable Output.

Overall Methodology

The proposed methodology follows the same reasoning process employed by experienced credit analysts during corporate credit evaluation.

The first stage evaluates historical financial performance using trend-based financial analytics derived from multiple years of audited financial statements. The second stage analyses qualitative corporate disclosures contained in annual reports to identify managerial confidence, strategic direction, governance quality, operational risks, and industry outlook. The third stage continuously monitors financial news to detect emerging external developments that may influence future credit quality. These three complementary information sources are subsequently transformed into structured feature vectors and integrated using a multimodal deep learning architecture that estimates the probability of future credit rating transition.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

Unlike conventional AI systems that function as black-box classifiers, the proposed framework preserves the contribution of each information source, enabling the final prediction to be interpreted by credit analysts and regulatory authorities, consistent with the explainability principles advocated by de Lange et al. (2022).

Data Architecture

Structured Financial Data

Structured information consists of audited financial statements and derived financial ratios covering multiple financial years. Instead of relying on a single reporting period, the methodology analyses historical trends to reduce the impact of temporary fluctuations.

The principal financial variables employed include:

Category	Representative Variables
Profitability	ROCE, ROA, EBITDA Margin, Net Profit Margin
Leverage	Gearing Ratio, Debt–Equity Ratio
Liquidity	Current Ratio, Quick Ratio
Coverage	Interest Coverage Ratio (ICR), Debt Service Coverage Ratio (DSCR)
Cash Flow	Operating Cash Flow, NCATD, CFO/Total Debt
Efficiency	Asset Turnover, Inventory Turnover, Receivable Turnover
Growth	Revenue Growth, EBITDA Growth, Profit Growth

Table 1. Principal Financial Variables Used in Quantitative Credit Assessment.

Rather than using raw financial ratios directly, the methodology transforms these indicators into standardized scores after trend smoothing and sector-specific normalization, as described in Section 4.

Unstructured Information

Three categories of unstructured information are incorporated.

Annual Reports

Annual reports provide management's assessment of corporate performance, strategic direction, and future outlook, consistent with the guidance embedded in the IASB Management Commentary Practice Statement (IASB, 2021) and ICAI's guidance on financial statements (ICAI, 2022). Rather than analysing the document as a whole, the proposed methodology identifies six sections considered most relevant for credit assessment:

- Chairman's Address
- Directors' Report
- Management Discussion and Analysis (MD&A)
- Independent Auditors' Report
- Key Business Risks
- Demographic and Industry Outlook

Each section contributes complementary information regarding future corporate performance.

Financial News

Financial news provides near real-time information regarding corporate developments that may influence future

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

credit quality. Typical events include acquisitions, management changes, regulatory actions, customer wins, litigation, operational disruptions, technological developments, and macroeconomic events. Unlike annual reports, which are published periodically, financial news functions as a dynamic indicator of changing corporate conditions (Tetlock, 2007; Davidovic & McCleary, 2025).

Credit Rating History

Historical credit ratings, rating outlooks, and rating transition information are incorporated for model training and validation. These observations enable the framework to learn patterns associated with historical rating migration.

Annual Report Intelligence Pipeline

One of the principal contributions of this research is the development of an Annual Report Intelligence Pipeline that converts qualitative corporate disclosures into structured credit-risk features. The pipeline consists of six sequential stages.

Stage 1: Document Acquisition

Annual reports are collected in digital PDF format. Where scanned documents are encountered, Optical Character Recognition (OCR) is applied to convert images into machine-readable text.

Stage 2: Section Identification

The system automatically identifies predefined sections using document structure, headings, and semantic cues. The sections extracted include the Chairman's Address, Directors' Report, MD&A, Independent Auditors' Report, Key Business Risks, and Demographic Impact. This section-wise processing allows different aspects of corporate behaviour to be analysed independently.

Stage 3: Intelligent Summarization

Each extracted section is summarized using a finance-oriented Large Language Model (Devlin et al., 2019; Brown et al., 2020). To maintain consistency across companies, summaries are restricted to approximately five to six sentences, capturing key events, strategic initiatives, operational developments, identified risks, and future outlook. This standardized representation substantially reduces document size while preserving information relevant for credit assessment.

Stage 4: Sentiment Extraction

The summarized narrative is analysed using domain-specific sentiment models trained on financial language. Each section receives a normalized sentiment score ranging from −1 (strongly negative) to +1 (strongly positive). Unlike generic sentiment analysis, financial sentiment considers context: terms such as "debt restructuring," "capacity expansion," "working capital pressure," or "capital expenditure" are interpreted according to their financial implications rather than their everyday linguistic meaning, echoing the finance-specific dictionary approach of Loughran and McDonald (2011) and the demonstrated scope for strategic tone management identified by Huang et al. (2014).

Stage 5: Feature Engineering

Beyond overall sentiment, additional quantitative variables are generated.

Feature	Description
Sentiment Score	Overall positive/negative tone
Management Confidence	Confidence expressed regarding future performance
Governance Quality	Strength of governance disclosures
Risk Intensity	Frequency and severity of identified business risks

Growth Opportunity	Discussion of future business opportunities
Innovation Index	Technology and product innovation emphasis
Auditor Confidence	Strength of audit opinion and internal controls

Table 2. Features Generated from Annual Report Intelligence.

These variables provide structured numerical representations of qualitative disclosures.

Stage 6: Embedding Generation

Finally, semantic embeddings representing the contextual meaning of each summarized section are generated using transformer-based language models (Vaswani et al., 2017; Devlin et al., 2019). These embeddings capture relationships among concepts that cannot be represented through simple sentiment scores alone. The resulting feature vectors become inputs to the multimodal feature fusion layer.

Financial News Intelligence

While annual reports summarize management's perspective over an entire financial year, financial news captures rapidly changing external developments. The proposed News Intelligence Module continuously collects financial news from trusted business news sources and performs three analytical tasks: automatic summarization, finance-specific sentiment analysis, and temporal aggregation using weighted averaging.

Instead of treating every news article equally, recent news receives greater importance than older observations. This temporal weighting enables the framework to detect emerging changes in corporate credit quality before they become visible in annual financial statements. This design choice is informed by evidence that unweighted, naïvely averaged sentiment carries only modest predictive power on its own (Davidovic & McCleary, 2025), which motivates the temporal weighting and drift-based aggregation adopted here rather than simple averaging.

A concise AI-generated news summary is created for each observation period, together with an aggregated sentiment score and a News Drift Indicator, representing the direction and magnitude of recent changes in market perception.

Hybrid Information Representation

The proposed framework transforms heterogeneous information sources into a common numerical representation suitable for deep learning. Three complementary feature blocks are constructed: F_Financial, representing structured financial ratios; F_AnnualReport, representing qualitative annual report intelligence; and F_News, representing recent external developments. These feature vectors are subsequently integrated through a multimodal fusion layer described in the following section.

The hybrid representation enables the framework to combine historical financial performance with forward-looking qualitative information, thereby improving the timeliness and interpretability of credit rating transition prediction.

Novelty of the Proposed Framework

The novelty of the proposed methodology lies not in the use of any single AI technique but in the systematic integration of complementary information sources within a unified credit risk assessment framework. Compared with existing approaches, the proposed framework introduces the following innovations:

- Trend-aware quantitative financial analytics using multi-year smoothing and sector-specific normalization.
- An Annual Report Intelligence Pipeline that analyses six credit-relevant sections rather than treating the annual report as a single document.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

- A dynamic News Intelligence Module that captures recent external developments through temporally weighted sentiment aggregation, addressing the known limitations of naïve sentiment averaging (Davidovic & McCleary, 2025).
- Multimodal feature fusion integrating structured financial variables, qualitative disclosures, and financial news.
- Explainable early warning scores supporting interpretation by credit analysts, banks, and regulatory authorities, in the spirit of de Lange et al. (2022) and Castro Vieira et al. (2025).

Together, these innovations provide a comprehensive framework for early identification of potential corporate credit rating transitions and establish the methodological foundation for the quantitative analytics presented in the following section.

Quantitative Financial Analytics

Introduction

Financial statement analysis remains the foundation of corporate credit assessment. Rating agencies evaluate financial strength using a combination of profitability, leverage, liquidity, debt servicing capacity, operational efficiency, and cash flow adequacy. These financial indicators collectively reflect the borrower's ability to generate sustainable earnings and meet long-term financial obligations (Altman, 1968; Thomas et al., 2017).

Traditional credit rating models generally analyse financial ratios for a single reporting period or compute simple averages across multiple years. Such approaches often fail to distinguish temporary fluctuations from sustained changes in financial performance. Corporate financial performance evolves continuously, and recent observations usually provide more relevant information regarding future credit quality than older data.

To address this limitation, the proposed framework introduces a Trend-Aware Quantitative Financial Analytics Module that combines multi-year financial information using Exponential Moving Average (EMA) smoothing, sector-specific normalization, and weighted financial aggregation. Rather than relying on isolated financial ratios, the methodology evaluates the long-term trajectory of corporate performance, thereby providing a more stable and informative representation of financial strength.

Financial Ratio Selection

The financial variables employed in this study were selected based on commonly accepted credit assessment practices adopted by banks and credit rating agencies (Lessmann et al., 2015; Baesens et al., 2016). These variables collectively represent profitability, leverage, liquidity, debt servicing capacity, operational efficiency, cash generation, and business growth.

Category	Financial Variables
Profitability	Return on Capital Employed (ROCE), Return on Assets (ROA), EBITDA Margin, Net Profit Margin
Leverage	Gearing Ratio, Debt–Equity Ratio
Liquidity	Current Ratio, Quick Ratio
Debt Servicing	Interest Coverage Ratio (ICR), Debt Service Coverage Ratio (DSCR)
Cash Flow	Cash Flow from Operations (CFO), Net Cash Accrual to Total Debt (NCATD), CFO / Total Debt
Operational Efficiency	Asset Turnover Ratio, Inventory Turnover Ratio, Receivable Turnover Ratio
Growth Indicators	Revenue Growth, EBITDA Growth, Profit After Tax Growth

Table 3. Financial Variables Used in the Proposed Framework.

These indicators are computed for multiple financial years, enabling trend analysis rather than single-period assessment.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

Multi-Year Financial Trend Analysis

One of the key innovations of the proposed framework is the use of multi-year financial smoothing. Financial ratios frequently exhibit short-term fluctuations arising from economic cycles, exceptional events, accounting adjustments, or one-time expenditures. Direct use of such ratios may therefore produce unstable credit assessments.

To reduce the influence of temporary variations while preserving recent trends, each financial ratio is transformed using an Exponential Moving Average (EMA). For a financial ratio R_t observed during year t, the smoothed value is computed as:

$$EMA_t = \alpha \cdot R_t + (1 - \alpha) \cdot EMA_{(t-1)}$$

where R_t represents the financial ratio for year t, EMA_t denotes the smoothed financial indicator, and α is the smoothing coefficient ($0 < \alpha < 1$).

The EMA assigns greater importance to recent financial performance while retaining information from previous years. Consequently, persistent improvement or deterioration in corporate performance becomes more evident than under conventional averaging techniques.

Sector-Specific Financial Normalization

Financial ratios vary considerably across industries. A leverage ratio considered acceptable within capital-intensive manufacturing industries may indicate elevated financial risk in service-oriented businesses. Similarly, profitability benchmarks differ substantially across sectors.

To ensure meaningful comparison, the proposed framework maps each smoothed financial ratio to a standardized score between 0 and 1 using sector-specific thresholds. For example, higher ROCE corresponds to higher normalized scores; lower Gearing Ratio corresponds to higher financial strength; higher Interest Coverage Ratio indicates stronger debt servicing capability; and higher Debt Service Coverage Ratio reflects improved repayment capacity.

Each normalized metric therefore represents the relative financial strength of a company within its own industrial sector rather than against a universal benchmark. The mapping thresholds may be calibrated using historical industry data or rating agency guidelines and can be updated periodically to reflect structural changes within specific sectors.

Weighted Financial Scoring

Not all financial variables contribute equally to corporate credit quality. Debt servicing capability generally has greater significance than inventory turnover, while sustained profitability is more informative than short-term revenue growth.

Accordingly, the proposed methodology aggregates normalized financial variables using a weighted scoring framework. The overall Financial Score is computed as:

$$FS = \Sigma (w_i \cdot S_i), \text{ for } i = 1 \text{ to } n$$

where S_i denotes the normalized score for the i-th financial variable, w_i represents its corresponding weight, and $\Sigma w_i = 1$.

For industrial companies, the initial weighting scheme adopted in this study is shown in Table 4.

Financial Variable	Weight (%)
Return on Capital Employed (ROCE)	35
Gearing Ratio	30

Interest Coverage Ratio	20
Debt Service Coverage Ratio	15

Table 4. Example Weight Distribution for Financial Variables.

These weights are based on their relative importance in long-term credit assessment and can be recalibrated for different industrial sectors. The framework is therefore sufficiently flexible to accommodate sector-specific credit evaluation policies.

Financial Stability Assessment

The proposed framework extends beyond static financial scoring by examining the consistency of corporate financial performance over time. Companies exhibiting steadily improving profitability, declining leverage, strengthening cash flows, and increasing debt servicing capacity generally demonstrate lower credit risk than firms experiencing large fluctuations in financial indicators.

Financial stability is therefore evaluated through three complementary dimensions: Trend Direction (improving, stable, deteriorating), Trend Magnitude (rate of improvement or decline), and Trend Consistency (volatility across successive reporting periods). The combination of these measures enables the framework to distinguish temporary financial disturbances from sustained structural changes in corporate performance.

Financial Feature Engineering

Following trend analysis and normalization, each financial variable is transformed into a structured feature suitable for multimodal learning. The final financial feature vector includes smoothed profitability indicators, smoothed leverage indicators, liquidity measures, debt servicing indicators, cash flow adequacy measures, operational efficiency indicators, growth indicators, a Financial Stability Index, and an overall Financial Score. Collectively, these features provide a compact numerical representation of historical corporate financial performance.

Financial Analytics Output

The output of the Quantitative Financial Analytics Module consists of two complementary components: (1) the Financial Feature Vector, which serves as input to the multimodal deep learning architecture; and (2) the Financial Score, the weighted financial score representing the overall financial strength of the company after incorporating multi-year trends and sector-specific normalization.

Unlike conventional financial ratio analysis, the proposed methodology captures both the absolute level and temporal evolution of corporate financial performance, thereby providing a stronger basis for predicting future credit rating transitions.

DISCUSSION

The proposed quantitative framework differs from traditional ratio-based models in several important respects. First, financial performance is analysed over multiple years rather than a single reporting period, reducing sensitivity to temporary fluctuations. Second, Exponential Moving Average smoothing emphasizes recent financial performance while retaining valuable historical information. Third, sector-specific normalization enables meaningful comparison among companies operating within similar industrial environments. Finally, weighted aggregation produces an interpretable financial score that can be combined seamlessly with qualitative information extracted from annual reports and financial news.

The resulting Financial Score does not directly predict future credit rating transitions; rather, it forms one of the three principal information streams used by the proposed multimodal framework. Subsequent sections demonstrate how qualitative intelligence extracted from corporate disclosures and financial news complements quantitative financial analysis to generate a comprehensive early warning indicator.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

Hybrid Early Warning Framework: Annual Report Intelligence, News Analytics, Multimodal Fusion, and Case Study

Introduction

Corporate credit deterioration is rarely caused by changes in financial ratios alone. Experienced credit analysts routinely complement financial statement analysis with qualitative assessment of management disclosures, auditor observations, industry developments, and current business news before arriving at a credit opinion. The objective of the proposed framework is to emulate this analytical process using Generative Artificial Intelligence.

The Quantitative Financial Analytics module presented in the previous section provides a trend-aware representation of historical corporate performance. However, historical financial information alone cannot capture management expectations, emerging operational risks, or significant business events that may influence future credit quality. Consequently, the proposed methodology incorporates qualitative intelligence extracted from annual reports and financial news to construct a comprehensive early warning framework.

Annual Report Intelligence

Annual reports constitute one of the richest publicly available sources of qualitative information describing a company's business environment, strategic priorities, and future outlook (IASB, 2021; ICAI, 2022). Rather than analysing the entire document as a single text corpus, the proposed framework extracts six sections that are routinely examined by credit analysts.

Chairman's Address

The Chairman's Address provides management's strategic perspective on business performance, major achievements, and future opportunities.

Illustrative Summary (case study company): Revenue growth remained strong despite a challenging macroeconomic environment. Management highlighted expansion in precision engineering and aluminium manufacturing while emphasizing continued investment in technology and production capacity. The report expressed confidence regarding future opportunities arising from infrastructure development and increasing demand for lightweight automotive components. Business risks were acknowledged but considered manageable through operational excellence and engineering capability. Overall management sentiment remained strongly optimistic regarding sustainable long-term growth.

Sentiment Score: +0.91

Directors' Report

The Directors' Report primarily reflects governance quality, statutory compliance, and financial performance.

Illustrative Summary: The Board reported healthy financial performance supported by capacity expansion and improved operational efficiency. Corporate governance practices, internal controls, and regulatory compliance were strengthened during the reporting period. Continued investment in manufacturing capabilities and technology remained central to long-term strategy. Dividend policy and shareholder value creation reflected management's confidence in future performance. Overall, the report portrays disciplined governance with a positive outlook.

Sentiment Score: +0.74

Management Discussion and Analysis (MD&A)

Illustrative Summary: Management identified continued opportunities across automotive, engineering, and industrial businesses while recognizing short-term economic uncertainties. Operational efficiency, product innovation, and customer diversification were highlighted as important competitive advantages. Expansion into

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

higher value-added engineering products was expected to support future profitability. The discussion reflects balanced optimism supported by measurable strategic initiatives. Overall management confidence remains high.

Sentiment Score: +0.84

Independent Auditors' Report

Illustrative Summary: The auditors issued an unqualified opinion indicating that the financial statements present a true and fair view of the company's financial position. Internal financial controls were considered adequate and operating effectively. No material weaknesses affecting financial reporting were identified. Statutory compliance requirements were satisfied without qualification. Consequently, the audit report contributes positively to overall credit assessment while remaining professionally neutral.

Sentiment Score: +0.36

Key Business Risks

Illustrative Summary: The report identifies cyclical demand, raw material price volatility, customer concentration, and regulatory changes as principal business risks. Expansion projects introduce execution risk but also create opportunities for future growth. Management describes mitigation strategies involving operational flexibility, technology investments, and customer diversification. No material going-concern concerns are reported. Overall risk assessment remains moderate.

Sentiment Score: +0.22

Demographic and Industry Outlook

Illustrative Summary: Urbanization, infrastructure development, and increasing demand for lightweight engineering products are expected to create favourable long-term market opportunities. Government investment in transportation, logistics, and manufacturing is likely to stimulate industrial demand. Expansion of the automotive sector and increasing aluminium substitution provide additional growth potential. Overall industry outlook remains positive.

Sentiment Score: +0.79

Annual Report Feature Generation

The extracted summaries are transformed into structured quantitative variables suitable for multimodal learning.

Feature	**Description**
Section Sentiment	Overall sentiment of each section
Management Confidence	Forward-looking optimism
Governance Quality	Board effectiveness and compliance
Auditor Confidence	Audit quality and internal controls
Risk Intensity	Severity of identified business risks
Industry Opportunity	Future growth prospects
Innovation Index	Technology and R&D orientation

Table 5. Features Generated from Annual Report Intelligence.

These variables complement traditional financial ratios by representing qualitative dimensions of corporate credit quality (Li, 2010; Huang et al., 2014).

News Intelligence

Financial news frequently captures emerging developments well before they appear in audited financial statements. Consequently, news sentiment functions as a leading indicator within the proposed framework, subject to the caveats on raw sentiment's predictive strength discussed in Section 2.7 (Davidovic & McCleary, 2025).

For each company, news articles published during the observation period are collected from reliable financial news sources. Duplicate articles and non-material events are removed before processing. The Large Language Model performs three tasks: news summarization, financial sentiment analysis, and event classification.

Example News Summary (case study company): During the observation period the company continued strengthening manufacturing capabilities and expanding customer relationships. Business operations experienced temporary disruption due to the COVID-19 pandemic and associated supply chain constraints. Despite these challenges, management maintained confidence regarding long-term business prospects supported by strong engineering capability and diversified product offerings. Market perception remained generally positive, although short-term uncertainty moderated investor expectations. Overall news sentiment indicates temporary operational disruption rather than structural deterioration in credit quality.

News Sentiment: +0.63

Hybrid Early Warning Algorithm

The proposed methodology integrates quantitative financial analytics, annual report intelligence, and financial news within a unified prediction framework.

Algorithm 1. Hybrid Early Warning Framework for Corporate Credit Rating Transition

Input: Multi-year financial ratios; Annual reports; Financial news.

- Step 1 — Compute trend-aware financial indicators using Exponential Moving Average and sector-specific normalization.
- Step 2 — Extract six predefined annual report sections.
- Step 3 — Generate concise AI summaries and finance-specific sentiment scores for each section.
- Step 4 — Collect recent financial news and compute temporally weighted sentiment.
- Step 5 — Construct structured feature vectors representing financial performance, annual report intelligence, and news intelligence.
- Step 6 — Integrate heterogeneous features through multimodal feature fusion.
- Step 7 — Estimate the probability of future credit rating transition using the deep learning prediction model.

Output: Financial Score; Qualitative Score; Early Warning Score; Predicted Rating Transition; Explainable Decision Report.

This algorithm synthesizes the trend-aware financial aggregation, annual report sentiment anchoring, news drift analysis, and weighted financial/qualitative integration developed in the expanded methodology described above.

Craftsman Automation Case Study

To demonstrate the proposed methodology, the framework was applied to Craftsman Automation Limited, an Indian precision-engineering and automotive-components manufacturer, using publicly available financial statements, annual report disclosures, and financial news as a proof-of-concept rather than a live deployment.

INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN APPLIED SCIENCE (IJRIAS)
ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

Historical financial ratios were transformed using EMA smoothing and sector-specific normalization to generate a quantitative financial score. Annual reports were processed using the proposed Annual Report Intelligence Pipeline to produce section-wise summaries and sentiment scores. Contemporary financial news was analysed to generate a News Sentiment Indicator representing recent market developments.

Quantitative Financial Analytics Results

Table 6 reports the computed results of the Quantitative Financial Analytics Module (Section 4) for the case study company, using the latest available ratio for FY2020, its EMA-smoothed value (α = 0.40), the resulting normalized score, the assigned sector weight, and the weighted contribution to the overall Financial Score.

Financial Parameter	Latest Ratio (FY2020)	EMA (α = 0.40)	Normalized Score (0–1)	Weight (%)	Weighted Contribution	Financial Trend
Return on Capital Employed (ROCE)	0.1205	0.1399	0.0494	35	0.0173	Declining
Gearing Ratio*	1.7915	1.5284	0.4165	30	0.1250	Moderate Risk
Interest Coverage Ratio (ICR)	2.6901	2.8804	0.0267	20	0.0053	Declining
Debt Service Coverage Ratio (DSCR)	1.2291	1.2238	0.7851	15	0.1178	Stable
Current Ratio	0.8147	0.7587	0.9398	—	—	Improving
Net Cash Accrual to Total Debt (NCATD)	0.2227	0.2496	0.7343	—	—	Stable
Profit After Tax (PAT) Margin	0.0247	0.0390	0.0469	—	—	Declining
Overall Financial Score	—	—	—	100	0.2654	Moderate Credit Strength

Table 6. Quantitative Financial Analytics Results for Craftsman Automation Limited.

*Note: For the Gearing Ratio, lower values indicate stronger financial health; reverse normalization was therefore applied before weighted aggregation.

Converting the Overall Financial Score to a 100-point scale for ease of interpretation, Financial Score = 0.2654 × 100 = 26.54. Table 7 presents the interpretation scheme used to classify this score.

Financial Score (0–100)	Interpretation
80–100	Excellent Financial Strength
60–80	Strong Financial Position
40–60	Satisfactory
20–40	Moderate Credit Risk
Below 20	High Credit Risk

Table 7. Interpretation Scheme for the Overall Financial Score.

On this basis, Craftsman Automation Limited receives a Financial Score of 26.54/100, placing it in the Moderate Credit Risk band based solely on the quantitative financial module. The Debt Service Coverage Ratio and Current Ratio show relatively strong normalized performance, reflecting satisfactory debt servicing capability and short-term liquidity, while profitability indicators — ROCE and PAT Margin — show declining trends over the observation period, pulling down the aggregate score.

This financial-only classification is deliberately conservative: it is subsequently combined with the qualitative evidence from annual reports and financial news (Sections 5.2–5.4) to arrive at the final multimodal early warning assessment, which may revise this classification upward or downward.

Integration with Qualitative Evidence

The multimodal framework combines this Financial Score with the Qualitative Score derived from Annual Report Intelligence and News Intelligence (Section 5.5, Algorithm 1) using the sector-specific overall weighting (60% Financial / 40% Qualitative for the industrial sector, per Appendix A). While the quantitative financial indicators point to moderate credit risk on a standalone basis, the annual report and news sentiment evidence (Sections 5.2–5.4) reflect materially higher management confidence and industry outlook. This divergence between the two information streams is itself informative: it illustrates why an early warning system should report the financial and qualitative components separately for analyst review rather than collapsing them into a single opaque figure, consistent with the explainability principle underlying the proposed framework (de Lange et al., 2022).

The multimodal framework successfully combined these complementary information sources into a unified Early Warning Score. News sentiment moderated the qualitative assessment by capturing temporary uncertainty associated with pandemic-related disruptions, illustrating the advantage of integrating historical financial information with current qualitative intelligence when assessing future credit quality — consistent with the broader finding that news-based signals are most useful when combined with, rather than substituted for, structured financial and disclosure-based indicators (Davidovic & McCleary, 2025).

Practical Applications

The proposed framework has immediate applications in banking and financial services. Potential applications include early warning systems for corporate loan portfolios; internal credit rating systems; Basel III/IV and Expected Credit Loss (ECL) monitoring (Basel Committee on Banking Supervision, 2024; Reserve Bank of India, 2025, 2026); continuous surveillance of large borrower portfolios; decision support for credit committees; corporate bond portfolio monitoring; NBFC credit risk management; and portfolio-level rating migration analysis.

Unlike traditional scoring systems, the framework continuously incorporates newly available information from financial statements, annual reports, and financial news, enabling proactive credit surveillance.

Limitations and Future Research

The present study has several limitations. The quality of qualitative analysis depends on the completeness and transparency of corporate disclosures. Financial news sentiment may also be influenced by reporting bias or temporary market reactions, and recent large-sample evidence suggests that raw sentiment signals are considerably noisier than early studies implied (Davidovic & McCleary, 2025). In addition, sector-specific financial thresholds require periodic recalibration to reflect structural changes in industry conditions.

Future research may extend the framework by incorporating earnings call transcripts, ESG disclosures, alternative data sources, and multilingual corporate reports. Integration of graph neural networks and agentic AI architectures may further enhance explainability and predictive performance, and empirical back-testing across a larger, cross-sector panel of rated companies — rather than a single illustrative case study — would materially strengthen the evidentiary basis of the framework.

CONCLUSIONS

This paper presented a multimodal Generative Artificial Intelligence framework for early warning of corporate credit rating transitions by integrating quantitative financial analytics, annual report intelligence, and financial news sentiment. Unlike conventional credit rating models that rely primarily on historical financial ratios (Altman, 1968; Altman et al., 2017), the proposed methodology combines structured financial information with qualitative corporate disclosures and dynamic market intelligence to generate an explainable early warning signal, in line with recent calls for explainable and fair AI in credit decisioning (de Lange et al., 2022; Castro Vieira et al., 2025).

The proposed Annual Report Intelligence Pipeline transforms six strategically important sections of annual reports into structured credit-risk features, while the News Intelligence Module captures emerging developments that may influence future credit quality, while remaining mindful of the limited standalone predictive power of unweighted news sentiment (Davidovic & McCleary, 2025). These complementary information sources are integrated through a hybrid multimodal framework that mirrors the reasoning process of experienced credit analysts.

The Craftsman Automation case study demonstrates the practical applicability of the proposed methodology and illustrates how qualitative information enhances traditional financial analysis. The framework offers a scalable and transparent decision-support system for banks, NBFCs, financial institutions, and credit rating agencies engaged in proactive credit surveillance. As a proof-of-concept built on publicly available data, the framework's predictive performance and generalizability now warrant systematic empirical validation across a larger, multi-sector sample of rated corporates.

Conflicts of Interest

The author declares no conflict of interest.

Funding

This research received no external funding.

Data Availability Statement

The datasets used in this study consist of publicly available financial statements, annual reports, regulatory disclosures, and publicly accessible corporate filings of Indian listed companies. Additional processed datasets are available from the author upon reasonable request.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

REFERENCES

1. Altman, E. I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. Journal of Finance, 23(4), 589–609.
2. Altman, E. I., Iwanicz-Drozdowska, M., Laitinen, E. K., & Suvas, A. (2017). Financial distress prediction in an international context: A review and empirical analysis of Altman's Z-Score model. Journal of International Financial Management & Accounting, 28(2), 131–171.
3. Baesens, B. (2014). Analytics in a Big Data World. Wiley.
4. Baesens, B., Rösch, D., & Scheule, H. (2016). Credit Risk Analytics: Measurement Techniques, Applications, and Examples in SAS. Wiley.
5. Basel Committee on Banking Supervision. (2024). Basel Framework. Bank for International Settlements.
6. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901.
7. Castro Vieira, J. R. de, Barboza, F., Cajueiro, D., & Kimura, H. (2025). Towards Fair AI: Mitigating Bias in Credit Decisions—A Systematic Literature Review. Journal of Risk and Financial Management, 18(5), 228. https://doi.org/10.3390/jrfm18050228
8. CFA Institute. (2023). Financial Statement Analysis. CFA Program Curriculum, CFA Institute.
9. CFA Institute Research Foundation. (2024). Artificial Intelligence in Investment Management. CFA Institute Research Foundation.

10. Chou, W.-H. (J.), Feng, Z., Li, B., & Liu, F. (2025). A First Look at Financial Data Analysis Using ChatGPT-4o. Journal of Risk and Financial Management, 18(2), 99. https://doi.org/10.3390/jrfm18020099
11. Davidovic, M., & McCleary, J. (2025). News Sentiment and Stock Market Dynamics: A Machine Learning Investigation. Journal of Risk and Financial Management, 18(8), 412. https://doi.org/10.3390/jrfm18080412
12. de Lange, P. E., Melsom, B., Vennerød, C. B., & Westgaard, S. (2022). Explainable AI for Credit Assessment in Banks. Journal of Risk and Financial Management, 15(12), 556. https://doi.org/10.3390/jrfm15120556
13. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT 2019, 4171–4186.
14. Ding, X., Zhang, Y., Liu, T., & Duan, J. (2015). Deep learning for event-driven stock prediction. Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), 2327–2333.
15. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
16. Hand, D. J., & Henley, W. E. (1997). Statistical classification methods in consumer credit scoring: A review. Journal of the Royal Statistical Society: Series A, 160(3), 523–541.
17. Huang, X., Teoh, S. H., & Zhang, Y. (2014). Tone management. The Accounting Review, 89(3), 1083–1113.
18. International Accounting Standards Board (IASB). (2021). IFRS Practice Statement 1: Management Commentary. IFRS Foundation.
19. Institute of Chartered Accountants of India (ICAI). (2022). Guidance Note on Audit of Financial Statements. ICAI.
20. Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. European Journal of Operational Research, 247(1), 124–136.
21. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474.
22. Li, F. (2010). The information content of forward-looking statements in corporate filings—A naïve Bayesian machine learning approach. Journal of Accounting Research, 48(5), 1049–1102.
23. Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. Journal of Finance, 66(1), 35–65.
24. OpenAI. (2023). GPT-4 Technical Report. arXiv:2303.08774.
25. Reserve Bank of India. (2025). Master Circular on Prudential Norms on Income Recognition, Asset Classification and Provisioning Pertaining to Advances. RBI.
26. Reserve Bank of India. (2026). Draft Directions on Expected Credit Loss (ECL) Framework for Provisioning by Banks. RBI.
27. Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. Journal of Finance, 62(3), 1139–1168.
28. Thomas, L. C., Crook, J., & Edelman, D. (2017). Credit Scoring and Its Applications (2nd ed.). SIAM.
29. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., et al. (2023). LLaMA: Open and efficient foundation language models. arXiv:2302.13971.
30. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.
31. World Economic Forum. (2024). The Future of Artificial Intelligence in Financial Services. World Economic Forum.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VII July 2026

Appendix A. Detailed Computational Algorithm for the Hybrid Early Warning Score

This appendix specifies, at implementation level, the three-phase computational procedure summarized as Algorithm 1 in Section 5.5: (1) quantitative financial computation, (2) qualitative sentiment computation, and (3) final rating aggregation and interpretation. The weights and thresholds below correspond to the industrial sector calibration used in the Craftsman Automation case study (Section 5.6) and should be recalibrated for other sectors.

Appendix A.1 Phase 1: Quantitative Financial Computation (Sector Weight: 60%)

Step A.1.1 — Multi-Year Smoothing (EMA)

For each financial ratio, an Exponential Moving Average is computed over the historical observation window rather than a simple average, so that recent years carry greater weight than older years:

$$EMA\_current = (Value\_current \times \alpha) + (EMA\_previous \times (1 - \alpha))$$

The smoothing coefficient α is calibrated per application; higher values of α (e.g., 0.5) place more emphasis on the most recent reporting period, while lower values (e.g., 0.4) retain more historical information. The Craftsman Automation case study in Section 5.6 uses $\alpha = 0.40$.

Step A.1.2 — Sector-Specific Mapping

Each EMA-smoothed ratio is mapped to a standardized Metric Score (S_m) between 0 and 1 using sector-specific thresholds. For the industrial sector used in this study, illustrative thresholds are:

- Gearing Ratio: ≤ 0.50 → Score 1.00; ≤ 1.20 → Score 0.85; ≤ 1.80 → Score 0.65 (reverse-scored, since lower gearing indicates lower financial risk).
- ROCE: ≥ 0.22 → Score 1.00; ≥ 0.14 → Score 0.85; ≥ 0.10 → Score 0.65.

Analogous threshold bands are defined for the Interest Coverage Ratio and Debt Service Coverage Ratio, and are recalibrated for non-industrial sectors using sector-specific benchmark data.

Step A.1.3 — Weighted Financial Block

The normalized metric scores are aggregated using the internal financial weights defined for the sector:

$$Score\_Fin = \Sigma\ (S\_m \times W\_m)$$

For the industrial sector: ROCE (35%), Gearing Ratio (30%), Interest Coverage Ratio (20%), Debt Service Coverage Ratio (15%), consistent with Table 4 and Table 6.

Appendix A.2 Phase 2: Qualitative Sentiment Computation (Sector Weight: 40%)

This phase uses a Large Language Model to process annual report narratives and financial news headlines.

Step A.2.1 — Annual Report "Anchor" Sentiment

Narrative text is extracted from the Chairman's Address, MD&A, and Directors' Report for each available reporting year. The LLM assigns a sentiment score (S_ar) between −1 and +1 for each year (Section 5.2). The same EMA logic used in Phase 1 is then applied to these yearly scores to obtain a Long-Term Sentiment Anchor (S_ema), so that more recent disclosures are weighted more heavily.

Step A.2.2 — News "Drift" (the Leading Indicator)

Key news items are extracted for a defined recent window (e.g., from the end of the last audited financial year through the current assessment date). The average sentiment score for this window (Score_News) is compared against the Long-Term Sentiment Anchor to obtain the News Drift:

$$\Delta = Score\_News - S\_ema$$

This Drift term is the mechanism by which the framework captures emerging developments that have not yet been reflected in the annual report narrative (Section 5.4).

Step A.2.3 — Final Qualitative Aggregation

The Long-Term Sentiment Anchor is adjusted using the Drift term, scaled by a sector-specific Acceleration Factor (A_f) that controls how strongly recent news is allowed to move the qualitative assessment away from the annual-report-based anchor:

$$Score\_Qual = S\_ema + (\Delta \times A\_f)$$

For the industrial sector used in this study, an illustrative Acceleration Factor of $A\_f = 0.45$ is adopted; this parameter should be reduced for sectors where news flow is noisier relative to fundamentals, consistent with the caution urged by Davidovic and McCleary (2025) regarding the standalone reliability of news sentiment. The Annual Report score and News score are combined using sector-specific sub-weights (e.g., 60% Annual Report / 40% News for the industrial sector) prior to this adjustment.

Appendix A.3 Phase 3: Final Rating Aggregation and Interpretation

Step A.3.1 — Final Score

The Financial Block and Qualitative Block are combined using the sector-specific overall weighting. For the industrial sector:

$$Score\_Final = (Score\_Fin \times 0.60) + (Score\_Qual \times 0.40)$$

Step A.3.2 — Rating Scale Mapping

The Final Score is mapped to an indicative credit rating band, calibrated against historical rating transitions for the relevant sector, for example:

- $Score\_Final \geq 0.75 \rightarrow$ indicative rating band A / A+
- $0.65 \leq Score\_Final < 0.75 \rightarrow$ indicative rating band BBB+
- $0.60 \leq Score\_Final < 0.65 \rightarrow$ indicative rating band BBB

These bands are illustrative of the industrial sector calibration used in this study and require recalibration against a larger historical rating-transition dataset before operational deployment, consistent with the future-research direction noted in Section 5.8.

Note on interpretation: as discussed in Section 5.6, the Financial Score reported in Table 6 for the case study company (26.54/100, i.e., Score_Fin = 0.2654) reflects the standalone quantitative module using the FY2020 EMA parameters described in Appendix A.1. The final Score_Final reported for the case study is obtained by combining this value with the Qualitative Score via Step A.3.1, and does not, by construction, need to independently exceed the illustrative Phase 3 rating thresholds above for the framework's diagnostic purpose to hold — the framework's value lies in exposing the components (Score_Fin, Score_Qual, Δ) separately for analyst review rather than in the single blended figure alone.