

Explainable AI in Corporate Finance: A Bibliometric Analysis of Research Trends and Future Directions

¹Dr Deeksha Arora, ²Dr Dinesh Rawat

¹Assistant Professor, Department of Management

Jagannath International Management School, Kalkaji, New Delhi

²Assistant Professor

Department of Management

Government Professional College, Uttarakhand, India

Department of Higher Education, Government of Uttarakhand

Abstract

Explainable artificial intelligence has been on a critical methodology innovation in corporate finance, but the intellectual structure of this emerging field is poorly understood. This paper attempts a bibliometric analysis of the field by exploring 83 articles found in the Scopus database between years 2015 and March 2026. The search terms were based on key constructs of XAI, and the main themes related to corporate finance. The analysis was performed with the help of VOSviewer (version 1.6.20) and Bibliometrix package in R. The findings suggest that Computational Economics, Journal of Risk and Financial Management, and Finance Research Letters are the journals that have published the most relevant research articles. China and Italy are the countries that have made the most contributions to the field. In terms of keywords clustering, SHAP values, XGBoost, credit risk, and financial distress prediction seem to be the most common ones in the field, but there is a growing interest in such topics as ESG, peer-to-peer lending, and new regulations related to AI. Conclusively, this paper provides a list of reasons why explainable artificial intelligence has become popular among financial researchers, potential areas of applications, and research directions. To my knowledge, this bibliometric analysis was not done before, and it can help explore this field further with the help of updated literature.

Keywords: Explainable Artificial Intelligence, XAI, Bibliometric Analysis, Corporate Finance, Interpretable Machine Learning, Credit Risk, Financial Distress

1. Introduction

Machine learning has become increasingly prominent in modern finance. Credit committees, risk desks, and investment teams are increasingly relying on algorithmic outputs to make decisions that were previously the sole purview of humans. The models implementing these algorithms can often be remarkably accurate in what they do; nevertheless, many are also incredibly challenging to interpret. A model can predict a firm as a bankruptcy risk or deny a credit application without providing the analyst with anything resembling a reason. This lack of information can be problematic for the analyst who is reviewing the output, or the regulator asking questions about it.

The area of XAI (explainable artificial intelligence) was created specifically to address this issue. Various methods, such as SHAP, LIME, partial dependence plots, and counterfactuals, have been introduced to help explain the inner workings of black-box models to a human reader. In financial services, the value of such interpretability has skyrocketed: both the GDPR and Equal Credit Opportunity Act require firms to offer explanations of the decisions rendered by their algorithms. Explainability is thus not just a desirable technical property, but a regulatory requirement.

There has been a substantial increase in research analyzing the applicability of XAI techniques to credit risk, financial distress, ESG scoring, and return prediction, especially since 2020. What the literature has not produced is a consolidated view of itself. Someone new to this field would struggle to identify where the important work has appeared, whose contributions have shaped the conversation and how the thematic focus has shifted across time. Bibliometric analysis is well suited to providing exactly that kind of structural picture.

That is what this paper attempts. 83 peer-reviewed articles were retrieved from Scopus database published between 2015 and March 2026, applied a multi-stage screening process to ensure relevance, and examined the resulting corpus across several dimensions: publication volume over time, country and author productivity, journal concentration, citation influence and keyword clustering. Each of these lenses adds something different to the overall picture.

Three things follow from this work. The analysis produces the most current bibliometric account of XAI research in corporate finance, with data running through March 2026. It also synthesises what the literature collectively shows in a way that individual studies cannot as they are focused on specific applications. It also maps out research directions that emerge organically from the gaps the evidence reveals not as speculative proposals, but as extensions of what the field has demonstrably left unresolved.

The paper is organised as follows. Section 2 sets out the conceptual background, covering how XAI has evolved and applied in corporate finance. Section 3 describes the methodology. Section 4 reports the findings. Section 5 interprets the thematic patterns. Section 6 closes with a research agenda.

2. Literature Review

2.1 The Rise of Machine Learning in Corporate Finance

The story begins with bankruptcy prediction. Altman's (1968) discriminant analysis model set the standard for decades, but linear boundaries struggle with financial data that does not behave linearly. Neural networks, decision trees, and support vector machines each improved on this, though none displaced traditional approaches decisively. Gradient boosting algorithms, namely, XGBoost, LightGBM, CatBoost changed that, building models iteratively with each step correcting the last. Machine learning now underpins credit risk, financial distress prediction, ESG scoring, and stock return forecasting across the industry, delivering accuracy gains consistently at the cost of comprehensibility.

2.2 The Explainability Imperative

SHAP values, rooted in cooperative game theory, have become the dominant explainability technique in finance, decomposing model outputs into feature-level contributions at both the individual prediction and dataset level. LIME offers a computationally lighter alternative but produces locally unstable explanations which is a material concern in credit scoring where consistency is a regulatory requirement (García-Céspedes et al., 2025).

2.3 XAI Applications Across the Corporate Finance Landscape

Credit Risk has attracted the most research. Bussmann et al. (2021) combined SHAP with network analysis to cluster borrowers by risk explanation similarity, and De Lange et al. (2022) demonstrated that LightGBM with SHAP outperformed a Norwegian bank's existing scoring system. Work since then spans consumer credit (Davis et al., 2023; Nwafor and Nwafor, 2023), mortgage defaults (Ozturkkal and Wahlstrøm, 2025), auto lending (Xie and Shingadia, 2025), digital finance (Ying et al., 2025), SME early warning (Akhmetova et al., 2026), sovereign risk (Giraldo et al., 2023; Anand et al., 2023), agricultural cooperatives (Kanta et al., 2026), and graph-based credit modelling (Pourkhoshgoftar et al., 2025; Dwivedi et al., 2026). Class imbalance, heterogeneous borrower populations, and MSME default modelling have been addressed by Ali and Patel (2025), Arias-Serna et al. (2025), Chen et al. (2025), Wang et al. (2025), Sunagatullin and Bahrami (2025), and Bitetto et al. (2025).

Bankruptcy and distress prediction have been similarly active. Liu et al. (2023) used TreeSHAP on Chinese firms, uncovering nonlinear relationships that linear models had been obscuring. Nguyen et al. (2025) tested gradient boosting models on French firms across multi-year horizons. Further contributions cover Nordic contexts

(Abrahamsen et al., 2024), climate risk (Stef et al., 2025), network augmentation (Kadkhoda and Amiri, 2025), multi-horizon forecasting (Dong and Li, 2026), voting ensembles (Phothong et al., 2026), stress testing (Zhao et al., 2026), COVID-19 distress features (Bozkurt and Kaya, 2023), multiperiod prediction (Beade et al., 2024), interpretability conflicts in boosted trees (Zhang et al., 2026), and decision support systems (Sun et al., 2023; Deng et al., 2025; Peng et al., 2025; Wang et al., 2024).

ESG applications are growing quickly. Teplova et al. (2023) found that investor weighting of environmental factors shifted during COVID-19. Helhel and Mengüç (2025) showed that mandatory ESG disclosure predicts financial performance more strongly than voluntary disclosure. Seow (2025), Cil and Yildiz (2025), Dossa et al. (2025), Ben Jabeur et al. (2025), Baklanova et al. (2024), and Mu et al. (2025) extend this across systematic reviews, ESG-performance links, geopolitical risk, NFT sentiment, and household financial participation.

In asset pricing and stock returns, Kang et al. (2025) found that tree-based models dominated Korean equity prediction with 36-month momentum as the key SHAP-identified driver. Ferrara and Ciano (2024) bridged algorithmic return predictions with factor explanations. Jang (2025), Chan et al. (2026), Li et al. (2025), Kim et al. (2024), Jacobs (2024), Fasano et al. (2024), Franco and Laurini (2025), Batebi and Elnahas (2026), Jiang et al. (2025), and Conti and Morelli (2025) extend XAI across analyst recommendations, activist targeting, default probability, cryptocurrency portfolios, insider trading, and Bayesian default modelling.

Methodologically, Philips et al. (2021) argued for auditable machine learning in factor investing, Ramteke et al. (2022) compared machine learning and statistical interpretability, and Lin and Wang (2025) raised concerns about SHAP stability across training runs with direct regulatory implications. Shoko et al. (2025), Chen et al. (2026), Lim and Jeong (2025), Pollestad and Oust (2025), and Radebe et al. (2026) extend interpretable methods to Bayesian optimisation, post-M&A default risk, cash holdings, real estate valuation, and municipal distress collectively signalling the diffusion of XAI well beyond its original banking context.

3. Research Methodology

3.1 Bibliometric Analysis as a Research Method

Bibliometric analysis works by treating scientific literature as data. Publications, citations, keywords, and authorship records become the raw material, and statistical methods are applied to these to reveal how a field is organised. In this study, XAI has been applied to corporate finance which is recent enough that no established review work exists for it. Researchers working here are doing so without much in the way of synthesized prior literature to orient themselves against. Also, output in this space has grown fast particularly from 2020 onward. Four methods have been used in the study each having its crucial role in the development of this research work. Co-occurrence mapping of author keywords picks up which concepts tend to appear together and gives a rough map of the field's thematic terrain. Citation counts point toward work that has had genuine downstream influence, as opposed to work that was simply published. Productivity profiling records where output has concentrated: which countries, which journals, which individual authors account for the most. None of these alone would be sufficient. Run together, they cover the field from several angles simultaneously.

3.2 Data Collection and Search Strategy

The data was gathered from the Elsevier Scopus database. The search was performed on 27 March 2026, which serves as the cut-off date for this study. The search string required to perform the query in the database included combinations of terms that do not naturally combine. For the XAI side: "explainable AI," "explainable artificial intelligence," "XAI," "interpretable machine learning," "model explainability," "model interpretability," "SHAP values," "LIME explainability." For the corporate finance side: "corporate finance," "corporate governance," "capital structure," "firm performance," "financial reporting," "earnings management," "investment decisions," "corporate risk," "credit risk," "financial distress," "bankruptcy prediction," "audit quality," "stock return." The Boolean AND operator connected the two groups, so only articles where both sides of the intersection are genuinely present would pass through. Titles, abstracts, and author keywords were the fields searched. The dataset starts from 2015 because before that point, XAI barely existed.

Table 1: Search Strategy and Screening Results

Search / Screening Stage	Articles (Scopus)
Initial keyword search (27 March 2026) -XAI terms AND corporate finance terms (TITLE-ABS-KEY)	370
Refined by language: English only Chinese (3), Korean (2), Spanish (1)	364
Restricted to source type: Journals only Conference proceeding (88) , Book series (50), Book (11), Trade journal (1)	214
Restricted to document type: Articles and Reviews Article (206), Review (6), Erratum (1), Retracted (1)	212
Restricted to subject areas: Economics, Econometrics & Finance AND Business, Management & Accounting Computer Science (34), Decision Sciences (22), Mathematics (13), Social Sciences (7), Psychology (4), Arts and Humanities (3), Environmental Science (2), Engineering (2), Energy (2)	101
Final sample after manual screening for relevance	83

3.3 Screening and Inclusion Criteria

The initial search yielded 370 results as presented in Table 1. The sample was next cut down step by step, each time eliminating the records that failed particular inclusion criteria. The first procedure removed all the non-English publications, resulting in 364 articles. The next exclusion step limited the sample to journal publications only; thus, conference publications, book items, and trade publications were discarded, leaving 214 items. Another screening process focused on document type, removing such items as errata and retracted work, retaining only research articles and review papers; thus, 212 items were left. One more step applied subject area inclusion criteria. Only articles that fell under Economics, Econometrics and Finance or Business, Management and Accounting were retained, which narrowed down the sample to 101 items. Finally, a manual screening process was applied to the titles, abstracts, and key terms of the remaining items. What this step aimed at was to ensure that all the selected records actually addressed the phenomenon of XAI within a corporate finance context. The articles that survived all the previous screening steps were finally subjected to a manual assessment step, during which the papers purely studying technical peculiarities of machine learning algorithms or mentioning financial data sets only as an illustration were removed; thus, 83 articles were eventually retained and formed the corpus of this study. Figure 1 below presents the Workflow of bibliometric literature review based on the PRISMA Model.




Figure 1 - Workflow of bibliometric literature review

3.4 Analytical Tools

The analysis has been done using VOSviewer (version 1.6.20) and Bibliometrix package in R, used through the Biblioshiny interface.

4. Results of Bibliometric Analysis

4.1 Most Relevant Journals

Journal-level analysis identifies the publication venues that serve as the primary intellectual homes for XAI in corporate finance research. Table 2 shows the ten journals with the most articles in the corpus.

Table 2: Most Relevant Journals

Journal	Articles
Computational Economics	9
Journal of Risk and Financial Management	7
Finance Research Letters	6
Journal of Forecasting	4
Risks	4
Journal of Financial Data Science	3
Journal of Risk Model Validation	3
Data Science in Finance and Economics	2
Decision Support Systems	2
Economic Modelling	2

Computational Economics leads with nine articles, a concentration that reflects the journal's deliberate positioning at the intersection of algorithmic approaches and economic applications. Its willingness to publish technically demanding papers on tree-based machine learning, neural networks, and XAI in credit and bankruptcy modelling has made it a natural home for this research stream.

The Journal of Risk and Financial Management, an open-access publication with broad scope across risk quantification and financial management, has attracted seven articles. Its accessibility and rapid publication timelines have made it attractive to researchers seeking quick dissemination of applied XAI studies. Finance Research Letters, with six articles, reflects the appetite for shorter, high-impact communications that identify specific findings about XAI applications without the extended theoretical apparatus typical of full-length journal articles.

The Journal of Forecasting and Risks each account for four articles. This is not surprising given what each journal covers. Forecasting research naturally intersects with questions of model accuracy and interpretability, while Risks has become a home for quantitative risk methodology work where XAI techniques are increasingly applied. The Journal of Financial Data Science and the Journal of Risk Model Validation contribute three articles each. Both sit at the boundary between academic research and professional practice, and both attract readers who work directly with quantitative models in institutional settings making them natural outlets for applied XAI work. Taken together, these ten journals account for the bulk of the corpus.

4.2 Source Impact by H index

Finance Research Letters has an H-index of 4, meaning that 4 of its articles in this corpus have been cited at least 4 times each, a sign that its contributions have genuinely influenced subsequent work, rather than simply building up in the record. The Journal of Risk and Financial Management trails at H=3. Computational Economics, Data Science in Finance and Economics, Decision Support Systems, the Journal of Banking and Finance, the Journal

of Financial Data Science, the Journal of Forecasting, Risks all fall at H=2 as per Figure 2. That concentration is not insignificant; it suggests a fairly wide spread of influential work across a range of outlets that serve different readerships, from applied economists to quantitative practitioners, to risk modelers. No one journal outside Finance Research Letters has pulled decisively ahead. Applied Economics comes in at H=1, the lower edge of local impact within this corpus; its contributions to the XAI-corporate finance intersection specifically have not yet attracted the citation depth that others have.

Impact in this corpus is not monopolised by one or two outlets. Finance Research Letters leads, but the midrange is crowded; several journals sit at H=2 serve quite different audiences, suggesting that the field has not yet settled around a set of destination journals. Work is landing in multiple places and still being cited.




Figure 2 : Source Impact by H index

4.3 Most Cited Countries

Citation counts at the country level are unevenly distributed, which is itself informative. Table 3 lists the ten most-cited countries alongside their average citations per article which reflects whether the work being produced is actually shaping subsequent research and not just adding to the count.

Table 3: Most Cited Countries in XAI - Corporate Finance Research

Country	Total Citations (TC)	Average Citations Per Article
China	499	26.3
Italy	333	83.2
Germany	131	32.8
Norway	100	20.0
Spain	58	19.3
USA	34	3.8

Country	Total Citations (TC)	Average Citations Per Article
France	32	16.0
Turkey	27	5.4
United Kingdom	25	8.3
Saudi Arabia	9	9.0

China leads by total citation volume with 499 citations, reflecting the sheer number of Chinese researchers publishing in this space, particularly in credit risk modelling and financial distress prediction for listed companies on Chinese stock exchanges. However, Italy's average citation rate of 83.2 per article vastly exceeds China's 26.3, a disparity explained by the outsized influence of a small number of highly cited Italian researchers principally Paolo Giudici and colleagues at the University of Pavia whose foundational papers on explainable machine learning in credit risk management (Bussmann et al., 2021) have become the most-cited work in the entire corpus with 316 citations.

Germany's presence, centred on work from the University of Passau and University of Regensburg, reflects strength in quantile neural networks for loss given default modelling. Norway's contribution comes primarily from the NTNU Business School, which has produced influential papers on financial distress prediction using LightGBM with SHAP interpretations in Nordic corporate contexts. The United States, despite being the global centre of machine learning research generally, contributes relatively few articles to this specific corpus which is a reflection, perhaps, of the decentralised nature of US financial regulation and the relatively slower uptake of formal XAI requirements compared to the EU context.

4.4 Most Cited Documents

Citation counts identify the intellectual anchors of a field. It shows the papers that others build upon, contest, or extend. Table 4 presents the ten most highly cited documents in the corpus, including both raw citation counts and normalised citation metrics that adjust for publication age.

Table 4: Most Cited Documents in the Corpus

Paper	Total Citations	TC Per Year	Normalised TC
Kou et al. (2021) Bankruptcy prediction for SMEs, Decision Support Systems	345	57.5	2.05
Bussmann et al. (2021) Explainable Machine Learning in Credit Risk, Computational Economics	316	52.7	1.88
Zacharias et al. (2022) Explainable AI-based feature selection, Electronic Markets	85	17.0	1.61
Carmona et al. (2022) No more black boxes - XGBoost in business failure, RIBAF	75	15.0	1.42

Paper	Total Citations	TC Per Year	Normalised TC
De Lange et al. (2022) Explainable AI for credit assessment in banks, J Risk Fin Mgmt	65	13.0	1.23
Ghosh et al. (2023) NFT and DeFi price prediction via XAI, Int Rev Fin Analysis	53	13.3	2.49
Liu et al. (2023) Tree-based gradient boosting for financial distress, J Forecasting	50	12.5	2.35
Zhang et al. (2023) DS-XGBoost for financial risk early warning, Finance Research Letters	41	10.3	1.93
Kellner et al. (2022) Quantile neural networks for loss given default, J Banking & Finance	39	7.8	0.74
Nallakaruppan et al. (2024) XAI framework for credit evaluation, Risks	26	8.7	2.42

The predominance of the top two papers is quite spectacular. Kou et al. (2021) conducted research on bankruptcy prediction for small and medium-sized enterprises using transaction data and multi-objective feature selection, which has received 345 citations. This number is impressive and the methods' innovativeness applying ensemble learning to a practically relevant task while addressing the curse of dimensionality is a significant methodological contribution. Bussmann et al. (2021)'s article is equally groundbreaking as it was one of the first to apply SHAP values in a credit risk context and use correlation network analysis to group borrowers together by the similarity of their SHAP explanations, thus suggesting new research directions. The normalised citation metric, which adjusts for the age of publication, highlights more recent papers that are accruing citations at a rapid rate. Ghosh et al.'s (2023) analysis of NFT and DeFi price dynamics using XAI achieves the highest normalised score (2.49), suggesting it addresses a particularly timely application domain. Liu et al. (2023) and Nallakaruppan et al. (2024) similarly show high normalised scores, indicating that the most recent waves of publication are receiving strong scholarly attention.

4.5 Most Relevant Authors

Author-level productivity analysis identifies the scholars who have made the most consistent contributions to XAI in corporate finance. Table 5 presents the ten most prolific authors by number of articles in the corpus, alongside their fractionalised article count which is a metric that adjusts for co-authorship by dividing each article's contribution equally among its authors.

Table 5: Most Relevant Authors

Author	Articles	Articles (Fractionalised)
Ben Jabeur, Sami	3	0.92
Carmona, Pedro	2	0.58
Chen, Wei	2	0.75
De Lange, Petter Eilif	2	0.45
Li, Yubo	2	0.34
Liu, Jiaming	2	0.40
Liu, Wenxin	2	0.50
Teplova, Tamara	2	0.67
Wang, Yong	2	0.75
Wu, Chong	2	0.40

Sami Ben Jabeur leads the author rankings with three articles in the corpus and a fractionalised score of 0.92, reflecting a research programme that spans bankruptcy prediction, ESG modelling, and clean energy price forecasting, each employing XAI techniques for model interpretation. Based across French institutions, Ben Jabeur exemplifies the European academic tradition of combining rigorous machine learning with regulatory-sensitive financial applications.

Pedro Carmona, based at the University of Valencia, has contributed directly to the methodological conversation around opening black-box models in corporate failure prediction. His paper with Dwekat and Mardawi (2022), titled 'No more black boxes!', became one of the most-cited works in the corpus, articulating the XAI imperative for business failure research in terms that resonated beyond technical audiences.

The relatively modest number of articles per author (even the most prolific contributor) has only three papers which reflects the genuine novelty of this research area. The field has not yet developed the dense single-author research programmes characteristic of more mature disciplines. This fragmentation is also visible in the fractionalised counts: most highly cited papers are collaborative works, meaning the credit for any individual paper is distributed across multiple authors. This collaborative character is a feature, not a limitation. It reflects the inherently interdisciplinary nature of research at the intersection of computer science, finance, and regulatory policy.

4.6 Top Keywords by Frequency

Author keywords record how researchers choose to position their own work, which terms they select to signal what a paper is about and which conversations it belongs to. In order to understand the thematic work, the researchers looked at how frequently particular keywords appear across the corpus. The Table 6 below ranks the most common author keywords by frequency of occurrence and also maps the conceptual vocabulary that researchers in this space share.

Table 6: Top Author Keywords by Frequency

Words	Occurrences
machine learning	27
credit risk	11
explainable ai	11
explainable artificial intelligence	11
interpretable machine learning	10
Interpretability	7
bankruptcy prediction	6
explainable machine learning	6
feature selection	6
machine-learning	6

4.7 Keyword Co-Occurrence Map

The keyword co-occurrence map in Figure 3, produced using VOSviewer and a minimum keyword occurrence threshold of three, depicts the thematic structure of XAI research in corporate finance. The map contains 47 keywords that meet the threshold requirement, which were organised into 5 broad clusters by the software.

The central cluster encompasses the core XAI concepts of machine learning, explainable AI, SHAP values, interpretable machine learning, and XGBoost, alongside their typical application domain of credit risk and financial distress. The close association of these terms suggests strong coupling between the methodological development and particular application in credit risk management. In practical terms, a researcher publishing an XAI innovation in a financial paper before 2020 almost inevitably applied the new method to credit risk modelling, and a financial distress paper published after 2020 almost never appeared without some application of a novel explanation method.

The next largest cluster in the top right of the map includes bankruptcy prediction, financial distress prediction, and firm performance, associated with classifiers such as Random Forest, LightGBM, and neural networks. Within this cluster, connections are dense, which suggests emergence of a sub-literature with relatively established conventions. The methods characteristic of this sub-field are balanced/over-sampled training sets, cross-validated classification with multiple classifiers, and SHAP-based feature ranking and ranking visualisation. Comparisons between classification methods are common, with bar charts or beeswarm plots of SHAP values serving as a canonical summary of model behaviour. Another sub-field, represented by a smaller cluster in the map, links ESG, corporate governance, sustainability, and firm value to XAI. This cluster appears distantly linked to the credit risk cluster, reflecting the relative novelty of ESG applications in the XAI field. The few connections to the larger body of literature occur through the shared methodologies of SHAP and XGBoost, suggesting active adoption of credit risk innovations by ESG researchers.

Two more application-specific clusters represent the fintech space: one encompasses peer-to-peer lending and credit scoring associated with LIME, while the other covers fintech and digital finance. Both are characterised by larger and more granular data sets, alternative data sources, and focus on individual rather than corporate credit assessment. The smallest cluster in the figure encodes sovereign risk, macroeconomic variables, and country-level analysis, representing a separate sub-field of country-level financial risk assessment.




Figure 3: Keyword Co-Occurrence Map

4.8 Temporal Evolution

Each keyword's color reflects its peak period of appearance in terms of the literature's intensity as well as its general location within the study window. Dark blue represents earlier mean publication years, approximately 2023.8, and yellow-green signifies the more recent ones, nearing 2024.8. Explainable AI and bankruptcy prediction belong to the blue category, alongside feature selection and explainable artificial intelligence. Thus, they are older nodes in the system in terms of peak co-occurrence, and their value is closer to 2023. They still have good connections to the rest of the network, but they represent the settled area of this field of study rather than its frontier.

Credit risk and machine learning are central to this keyword network, which is reflected in their color green. They are situated near the middle of the color range, which implies their peak publication period is around the midpoint of the timeline used for the study. In turn, yellow-green keywords 5 and 6, interpretable machine learning and risk assessment, are the latest to appear. Interpretable machine learning in particular seems to gain more traction in contemporary research, as manifested in the shift from explainable AI and its umbrella term, explainable artificial intelligence. Forecasting is the outlier in this network, as its placement and color place it closer to the beginning of the timeline, but its connections reach further. It is a typical development for the research field to start with keywords that represent a wide scope of study and slowly narrow down to more specific terms. Forecasting belonged to the early adopters in terms of publication but failed to gain the same level of popularity as credit risk and machine learning. The study's visual representation of the keywords' evolution over time demonstrates how the trend evolved from encompassing a variety of approaches to focusing on a particular statistical method.




Figure 4: Temporal Evolution Map

4.9 Bibliographic Coupling

Bussmann (2021) and Kou (2021) are the two papers that stand out immediately. Their nodes are visibly larger on the map. Carmona (2022) and Zacharias (2022) grab the second position De Lange (2022). and Kellner (2022) occupy similar territory, mid-network, moderately connected, part of the consolidation layer that formed as the field picked up pace. The green and teal nodes representing Ghosh, Liu, Zhang, Szepannek from 2023, then Tang, Wang, Nazemi, Bone-Winkel from 2024 fill out the middle of the map. Their connections run both toward the central anchors and laterally across each other, which matters. It means more recent work is starting to reference other recent work, which is how a literature develops its own internal structure over time. Nallakaruppan (2024) seems to be an exception as it looks detached from the main body of the network. It is likely approaching the subject from a different disciplinary angle or drawing on a reference base that does not overlap much with the rest of the corpus.




Figure 5: Bibliographic Coupling Map

4.10 Document Co-citation Network

Guestrin's XGBoost paper carries the largest node on the map and connects across almost every cluster a reflection of how routinely this work gets cited alongside everything else in applied financial machine learning. Altman's financial ratios paper appears twice, once in the left cluster and again as a large green node on the right. That is not a mapping error; it captures the fact that Altman gets pulled into reference lists from two different angles, by distress prediction studies and by papers benchmarking against classical models.

The dense left cluster groups ensemble and credit risk methods together Krivorotov, Xia, Li, Papenbrock, Breiman. Papers working in this space tend to cite all of them in combination, which is why they sit so tightly. Berrada and Vandierendonck sit in the middle, bridging the applied left side and the interpretability-focused right both deal with opening up model internals, so their position makes sense.

The right side is looser. Altman, Breiman's Random Forests, and Caruana's InterpretML anchor a green cluster, while Athanasopoulos and McDonald sit further out in orange, representing a forecasting strand that intersects with XAI only partially.

Three intellectual traditions are visible overall, namely, classical financial distress modelling, applied ensemble methods, and interpretability methodology. Most papers draw on more than one, which explains why the cross-cluster connections are as dense as they are.




Figure 6: Document Co-citation Network

4.11 Annual Publication Trend

Publication numbers across the 83-article corpus shows that the output began low, and continued to be low for a while, but soon shot up. In 2015, the field of XAI was not coherent. The label was barely in circulation, work applying interpretable machine learning to corporate finance problems was a small cluster of credit scoring studies. This was largely true till 2018. Research was appearing though slowly, and the community that was producing was small. 2019 and 2020 ushered in a marked change. Bussmann et al.'s work on explainable machine learning in credit risk drew wider attention on what these methods could do in a financial context, not just technically, but in terms of regulatory relevance. GDPR's Article 22, which placed obligations on institutions using automated systems to make decisions affecting individuals, had come into force in 2018 and was beginning to create practical pressure on financial firms to demonstrate that their models could be explained. From 2021 onwards, output has accelerated considerably. In comparison to 2020, annual publication counts had more than doubled by 2022 and 2023. SHAP and LIME had been packaged into standard Python libraries by this point, making them accessible to researchers without deep machine learning expertise. The COVID-19 period had also shown a certain weakness of black-box models when structural breaks occur, models trained on historical patterns fail in ways that nobody can easily diagnose, which made the case for transparency more concrete. Large language models have introduced another possibility that model outputs could be accompanied by plain-language explanations rather than just numerical scores. Data of 2025 and early 2026 also suggest a rising trend. The trend is captured only till 27 March 2026, therefore, the trend is showing downwards post 20025. However, it is expected to be rising towards the end of 2026.




Figure 7 : Publication Trend

5. Thematic Synthesis and Future Research Directions

5.1 The Dominance of Credit Risk and Bankruptcy Prediction

The most surprising thing about the bibliometric evidence was just how much of the overall corpus seemed to be devoted to the task of credit risk analysis and bankruptcy prediction. Together these two application areas dominated the overall number of articles, citations, and keywords. This dominance can be explained by a combination of factors, the first being that this was the area where machine learning applications had been deployed and evaluated for the longest time. As such, this field provided a rich foundation of reference datasets and evaluation criteria to which the new algorithms could be compared. Similarly, the necessity of providing an explanation behind automated credit decisions created an additional impetus for researchers to engage with the domain, not only from a theoretical but also from a practical perspective. And finally, the relative availability of high-quality labelled data (default/non-default) provided fertile ground for applying and evaluating the classification algorithms, to which the XAI field is particularly concerned with understanding. This can be seen in Figure 8 below.




Figure 8 : Thematic Map

While exploring the credit risk cluster, a pattern appeared; XGBoost / LightGBM as the main classifier (with the addition of the SMOTE or similar techniques to combat the issue of class imbalance) and the utilization of SHAP values for global feature importance rankings as well as local individual explanations constitute the basis of methodology chosen by authors, working on a paper from different countries and institutions to such an extent, that it can be seen as a template for XAI research in credit risk. Numerous scholars have employed the XGBoost algorithm to investigate credit risk, plotting the resulting SHAP values and examining how they are distributed. Beyond what is already evident, what fresh insights can emerge from examining SHAP values when the same approach is applied to different data sets? Future investigations might explore how the same algorithm performs across different data sets, compare the explanations that SHAP and LIME offer for an identical model, or assess how stable the results remain during economic crises versus periods of stability. Moreover, examining how XAI analysis shapes the decision-making of credit officers would represent a valuable avenue for further research.

5.2 The Emerging ESG–XAI Nexus

The co-occurrence map suggests that ESG and sustainability is an emerging but still peripheral cluster. This reflects the current state of the field: the incorporation of ESG into machine learning models is still relatively new, and application of XAI to understand ESG-based predictions is even newer. Nevertheless, the number of papers within this intersection is growing at one of the highest rates. The systematic literature review of ESG analytics through machine learning by Seow (2025) and the study by Dossa et al. (2025) on ESG disclosure and performance of Chinese firms using Extra Trees and SHAP are two recent examples.

The intersection of ESG and XAI is an interesting one for three reasons. First, the multifaceted nature of ESG scores that combine environmental, social and governance factors into aggregated estimates invites the application of various XAI methods to understand what drives these predictions. There is also the problem of discrepancy between ESG scores provided by different ratings agencies. Within one company, significantly different scores can be found across ESG pillars. The use of SHAP decomposition might help explain the situation, highlighting variations between different ESG pillars. Finally, the requirement for companies to report on their sustainability performance is creating new data sets that did not previously exist, providing new inputs for analysis.

Several questions arise from this. Could SHAP values be used to evaluate ESG scores provided by for-profit ratings agencies, rather than simply taking them as given? If so, it would be interesting to see how the weights of different variables compare across agencies. Furthermore, the importance of individual ESG pillars is likely to vary both between industries and over time. Most interestingly, insights gained through the application of XAI to ESG scores could be used to inform changes to a company's strategy, transforming the analysis from a descriptive to a prescriptive exercise.

5.3 Regulatory AI and Fairness in Financial Decision-Making

A significant part of the readings in the Journal of Risk Model Validation and Journal of Financial Data Science relate to the regulatory aspects of the use of XAI in the finance field. Specifically, the works of García-Céspedes et al. (2025), Ramteke et al. (2022), Bharodia and Chen (2021), and Lin and Wang (2025) highlight the problems that arise in practice when trying to interpret the results of XGBoost models with SHAP explanations. The questions posed include whether the reason codes produced by XGBoost models are comparable to logistic regression in the context of credit risk management, whether there is stability in the SHAP values across different training runs, and how these explanations satisfy the legal requirements for adverse action alerts.

However, it is the question posed by Lin and Wang (2025) on the instability of SHAP values across different training runs of the same XGBoost model with the same data that seems the most pressing. The issue demonstrated in the study relates to the fact that for the majority of features with medium importance, SHAP values varied significantly (by more than 30%) across 100 different training runs. Therefore, the use of such explanations for credit scoring purposes seems questionable, as there is no guarantee of consistency and explainability of results, which is critical for legal compliance.

The topic of SHAP value stability across different machine learning runs needs further exploration, as well as the implications of such instability for the use of explanations in practice. It is necessary to formalize requirements for SHAP stability for credit scoring applications and examine how the instability issue could be addressed when communicating model evidence to regulators. Moreover, it is important to consider whether the initial use of inherently interpretable models that do not require post-hoc explanations could be a viable way to ensure algorithmic transparency and compliance with regulations.

5.4 Methodological Frontiers: Beyond SHAP

While SHAP dominates the corpus, the bibliometric evidence points to several methodological frontiers that are beginning to attract attention. Counterfactual explanations which answer the question 'what would have to be different about this loan application for the model to approve it?' appear in a small but growing set of papers (Wang et al., 2025) and represent a more actionable form of explanation for credit applicants. Graph neural networks with interpretability components are emerging as a promising approach for modelling financial networks, where the relationships between borrowers or firms carry risk information that tabular models cannot capture (Pourkshofigoftar et al., 2025; Hegde and Bhowmik, 2026). The integration of natural language processing with XAI is another emerging frontier. Tang et al. (2024) demonstrated that BERTopic, a transformer-based topic modelling approach, can be combined with SHAP to interpret the textual content of corporate bond news as a predictor of default. Hesse and Loy (2025) used BERT-based sentence embeddings from management discussion and analysis sections of annual reports to predict bankruptcy, with sentence-level analysis providing a form of textual interpretability. These approaches represent a confluence of quantitative approaches to XAI and qualitative analysis of financial text.

Federated learning has begun to appear in discussions of XAI, while the combination raises a problem that is worth considering. The basic principle of federated learning is that a model can be trained across several institutions without any of them having to provide their raw data. This is useful when privacy regulations are becoming more strict and financial institutions are reluctant to pool sensitive records together. However, this creates a real difficulty for explainability. If no institution has visibility over the full set of training data, it becomes significantly harder to explain why a model arrived at a particular prediction than it would be in a standard setting. How to generate meaningful explanations under these circumstances is an open methodological question.

5.5 Underexplored Corporate Finance Domains

The bibliometric evidence points towards several domains in corporate finance that have remained underserved by XAI research. Capital structure, the trade-off between debt and equity, does not appear to be a topic addressed by any of the keywords suggesting that XAI has failed to address the fundamental issues of optimal leverage, pecking order behaviour or market timing in financing decisions. Dividend policy, mergers and acquisitions and initial public offering pricing are completely absent in the corpus.

Corporate governance appears in the keyword co-occurrence map, but mostly as a secondary theme connected to ESG research and is not presented as a primary field of application for XAI. The application of XAI to understand the impact of board characteristics, ownership structures or governance indices on firm value is an unexplored opportunity. Working capital management, which is at the heart of corporate finance practice, appears in only one paper in the corpus (Akdogan, 2025) indicating a substantial disconnect between academic XAI research and the day-to-day decision making that occupies the time of many corporate finance professionals. Real estate finance, insurance pricing and municipal government finance have started showing up in the most recent papers in the corpus and reflect the expansion of XAI techniques beyond banking into other financial sectors. Pollestad and Oust (2025) applied SHAP to an automated valuation model for Oslo apartments; Radebe, Nomlala and Matenda (2026) used Random Forest with SHAP for South African municipal financial distress prediction. These early forays suggest that XAI's applicability extends well beyond banking and corporate credit, and that researchers in adjacent domains will increasingly adopt the methodological toolkit developed in credit risk research.

6. Conclusion

This bibliometric analysis has mapped, for the first time, the intellectual structure of the emerging field at the intersection of explainable artificial intelligence and corporate finance. The present article is based on the analysis of 83 peer-reviewed articles found in the Scopus database. The selection was made from screened results by applying inclusion criteria. The current research explores the trends of development of the topic by determining the growth of the number of publications over time, detecting the leading countries and organizations, assessing the most cited works, and analyzing the content structure.

To begin with, the overall scope of the topic can be identified as limited, given the number of publications since 2015 to 2026, which suggests that it is a relatively new area of research. However, since then, the number of articles has grown significantly. The reason for this trend can be linked to the regulator's interest in improving algorithmic transparency, the availability of gradient boosting machines and SHAP explanation techniques for non-machine learning experts, and the necessity to make more accountable AI-driven financial decisions.

It is noteworthy that the largest body of work focuses on credit risk assessment and bankruptcy prediction, including the use of a standard classification XAI framework that pairs ensemble classifier algorithms with SHAP explanation methods. Given the researchers' efforts, China and Italy lead the way in terms of quantity, whereas Italy dominates in quality with several influential papers.

Furthermore, the current scope of research allows detecting the frontier of knowledge in the area of XAI in corporate finance. First, ESG–XAI intersection is a promising direction given the rising importance of ESG integration into business practices. Second, there is a need for further research into the properties of the SHAP values, especially from the regulatory perspective, as well as an opportunity to explore such explanation approaches as counterfactuals, graph neural networks, and federated learning. Third, there are multiple areas within corporate finance that have not attracted the attention of researchers yet, including capital structure,

dividend policy, M&A, and working capital, among others. Thus, it can be concluded that the field of XAI in corporate finance has the potential for further development both in terms of the methodology and research topics.

In turn, the findings of the present study can be beneficial for practitioners and policymakers interested in the application of explainable AI in corporate finance. In particular, the bibliometric analysis highlights the topic's research potential, which can be of interest to both scholars and industry practitioners. At this stage of development, the area needs the input of researchers in order to address the concerns of the regulators and end-users of AI models. While the necessity to comply with the ESG-related regulations and explain one's decisions is growing, the opportunity to use the explainable AI models provides room for improvement. Therefore, the primary question for these stakeholders should be not whether they need to employ the methods of XAI but how seriously they need to consider their opportunities. In particular, the analysis of literature suggests that the use of such methods can facilitate more accurate economic forecasts, test long-standing theories, and improve the decision-making process in corporate finance, benefiting both its participants and the regulators.

CRedit author statement

Dr Deeksha Arora: Writing - Conceptualization, Data curation, Methodology, Writing – original draft, review & editing, Software, Formal analysis

Funding sources

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author used Claude (Anthropic) and ChatGPT in order to assist with language editing and improving the readability of the manuscript. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- [1] Abrahamsen, N., Nylén-Forthun, E., Møller, M., de, L.P., & Ristad, M. (2024). Financial Distress Prediction in the Nordics: Early Warnings from Machine Learning Models. *Journal of Risk and Financial Management*, 17(10), . <https://doi.org/10.3390/jrfm17100432>
- [2] Akdogan, Y. (2025). AI-driven insights into working capital strategies: An application on Borsa Istanbul. *Borsa Istanbul Review*, 25(6), 1359-1377. <https://doi.org/10.1016/j.bir.2025.07.017>
- [3] Akhmetova, A., Shayakhmetova, A., & Abdurakhmanov, N. (2026). Interpretable Multi-Model Framework for Early Warning of SME Loan Delinquency. *Risks*, 14(2), . <https://doi.org/10.3390/risks14020025>
- [4] Ali, S.S., & Patel, R. (2025). An Explainable ADASYN-Based Focal Loss Approach for Credit Assessment. *Journal of Forecasting*, 44(4), 1513-1530. <https://doi.org/10.1002/for.3252>
- [5] Anand, A., Baesens, B., & Vanpée, R. (2023). Sovereign credit risk modeling using machine learning: a novel approach to sovereign credit risk incorporating private sector and sustainability risks. *Journal of Credit Risk*, 19(1), 105-154. <https://doi.org/10.21314/JCR.2022.008>
- [6] Arias-Serna, M., Quiza-Montealegre, J., Montes-Gómez, L., Clavijo, L., & Orozco-Duque, A. (2025). Explainable Machine Learning Models for Credit Rating in Colombian Solidarity Sector Entities. *Journal of Risk and Financial Management*, 18(9), . <https://doi.org/10.3390/jrfm18090489>
- [7] Baklanova, V., Kurkin, A., & Teplova, T. (2024). Investor sentiment and the NFT hype index: to buy or not to buy?. *China Finance Review International*, 14(3), 522-548. <https://doi.org/10.1108/CFRI-06-2023-0175>
- [8] Batebi, S., & Elnahas, A. (2026). Are machines better predictors of insider trading?. *Global Finance Journal*, 69, . <https://doi.org/10.1016/j.gfj.2026.101237>
- [9] Beade, Á., Rodríguez, M., & Santos, J. (2024). Multiperiod Bankruptcy Prediction Models with Interpretable Single Models. *Computational Economics*, 64(3), 1357-1390. <https://doi.org/10.1007/s10614-023-10479-z>

- [10] Ben, J.S., Bakkar, Y., & Cepni, O. (2025). Do global COVOL and geopolitical risks affect clean energy prices? Evidence from explainable artificial intelligence models. *Energy Economics*, 141, . <https://doi.org/10.1016/j.eneco.2024.108112>
- [11] Bharodia, N., & Chen, W. (2021). What can we learn from what a machine has learned? Interpreting credit risk machine learning models. *Journal of Risk Model Validation*, 15(2), . <https://doi.org/10.21314/JRMV.2020.235>
- [12] Bitetto, A., Filomeni, S., & Modina, M. (2025). Machine Learning for the Unlisted: Enhancing MSME Default Prediction with Public Market Signals. *Journal of Corporate Finance*, 94, . <https://doi.org/10.1016/j.jcorpfin.2025.102830>
- [13] Bone-Winkel, G., & Reichenbach, F. (2024). Improving credit risk assessment in P2P lending with explainable machine learning survival analysis. *Digital Finance*, 6(3), 501-542. <https://doi.org/10.1007/s42521-024-00114-3>
- [14] Bozkurt, İ., & Kaya, M. (2023). Foremost features affecting financial distress and Bankruptcy in the acute stage of COVID-19 crisis. *Applied Economics Letters*, 30(8), 1112-1123. <https://doi.org/10.1080/13504851.2022.2036681>
- [15] Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable Machine Learning in Credit Risk Management. *Computational Economics*, 57(1), 203-216. <https://doi.org/10.1007/s10614-020-10042-0>
- [16] Carmona, P., Dwekat, A., & Mardawi, Z. (2022). No more black boxes! Explaining the predictions of a machine learning XGBoost classifier algorithm in business failure. *Research in International Business and Finance*, 61, . <https://doi.org/10.1016/j.ribaf.2022.101649>
- [17] Chan, C., Tang, F., & Yang, J. (2026). Financial Performance Prediction and Stability Analysis Using SHAP-Enhanced Machine Learning Models. *Computational Economics*, , . <https://doi.org/10.1007/s10614-025-11283-7>
- [18] Chen, R., Tong-Yu, L., Min, J., & Wenfu, X. (2025). A comprehensive explainable approach for imbalanced financial distress prediction. *Journal of Risk Model Validation*, 19(3), 1-12. <https://doi.org/10.21314/JRMV.2025.010>
- [19] Chen, W., Liu, W., Liu, W., & Tian, G. (2026). Network contagion and post-M&A default risk: Interpretable machine learning evidence from an emerging market. *British Accounting Review*, , . <https://doi.org/10.1016/j.bar.2026.101843>
- [20] Cil, A., & Yildiz, K. (2025). A systematic literature review on applications of explainable artificial intelligence in the financial sector. *Internet of Things (The Netherlands)*, 33, . <https://doi.org/10.1016/j.iot.2025.101696>
- [21] Conti, A., & Morelli, G. (2025). Bayesian probability of default models with Langevin dynamics. *Quantitative Finance*, 25(8), 1333-1341. <https://doi.org/10.1080/14697688.2025.2532025>
- [22] Davis, R., Lo, A., Mishra, S., Nourian, A., Singh, M., Wu, N., & Zhang, R. (2023). Explainable Machine Learning Models of Consumer Credit Risk. *Journal of Financial Data Science*, 5(4), 9-39. <https://doi.org/10.3905/jfds.2023.1.141>
- [23] de, L.P., Melsom, B., Vennerød, C., & Westgaard, S. (2022). Explainable AI for Credit Assessment in Banks. *Journal of Risk and Financial Management*, 15(12), . <https://doi.org/10.3390/jrfm15120556>
- [24] Deng, S., Li, Y., Zhu, Y., Wang, B., Ning, H., Yi, S., & Shimada, T. (2025). Financial distress warning and risk path analysis for Chinese listed companies: An interpretable machine learning approach. *Economic Modelling*, 152, . <https://doi.org/10.1016/j.econmod.2025.107288>
- [25] Dong, Q., & Li, L. (2026). Forecasting Corporate Default Risk Across Multiple Horizons With Interpretable Machine Learning. *Journal of Forecasting*, , . <https://doi.org/10.1002/for.70129>
- [26] Dossa, J., Ukwuoma, C., Thomas, D., Dossa, J., & Gopang, A. (2025). Prediction of nexus among ESG disclosure and firm Performance: Applicability, explainability and implications. *Innovation and Green Development*, 4(4), . <https://doi.org/10.1016/j.igd.2025.100261>
- [27] Dwivedi, P., Islam, B., & Kajal, M. (2026). HMDL-GFNR: A Hybrid Multi-Stage Deep Learning Model with Graph-Based Financial Network Representation for Credit Risk Assessment. *Computational Economics*, , . <https://doi.org/10.1007/s10614-025-11247-x>
- [28] Emin, D., & Emin, A. (2025). The role of AI in credit risk assessment: Evidence from OECD and BRICS via system GMM and random forest. *Finance Research Letters*, 81, . <https://doi.org/10.1016/j.frl.2025.107499>

- [29] Fasano, F., Adornetto, C., Zahid, I., La, R.M., Montaleone, L., Greco, G., & Cariola, A. (2024). The dilemma of accuracy in bankruptcy prediction: a new approach using explainable AI techniques to predict corporate crises. *European Journal of Innovation Management*, 28(11), 1-22. <https://doi.org/10.1108/EJIM-06-2024-0633>
- [30] Ferrara, M., & Ciano, T. (2024). Improving the Interpretability of Asset Pricing Models by Explainable AI: A Machine Learning-based Approach. *Economic Computation and Economic Cybernetics Studies and Research*, 58(4), 5-19. <https://doi.org/10.24818/18423264/58.4.24.01>
- [31] Franco, J., & Laurini, M. (2025). Integrating Choquet portfolios and machine learning interpretability for robust cryptocurrency investment strategies. *Journal of Finance and Data Science*, 11, . <https://doi.org/10.1016/j.jfds.2025.100172>
- [32] Gafsi, N. (2025). Machine Learning Approaches to Credit Risk: Comparative Evidence from Participation and Conventional Banks in the UK. *Journal of Risk and Financial Management*, 18(7), . <https://doi.org/10.3390/jrfm18070345>
- [33] García-Céspedes, R., Alias-Carrascosa, F., & Moreno, M. (2025). On Machine Learning models explainability in the banking sector: the case of SHAP. *Journal of the Operational Research Society*, 76(12), 2591-2603. <https://doi.org/10.1080/01605682.2025.2485263>
- [34] Ghosh, I., Alfaro-Cortés, E., Gámez, M., & García-Rubio, N. (2023). Prediction and interpretation of daily NFT and DeFi prices dynamics: Inspection through ensemble machine learning & XAI. *International Review of Financial Analysis*, 87, . <https://doi.org/10.1016/j.irfa.2023.102558>
- [35] Giraldo, C., Giraldo, I., Gomez-Gonzalez, J., & Uribe, J. (2023). An explained extreme gradient boosting approach for identifying the time-varying determinants of sovereign risk. *Finance Research Letters*, 57, . <https://doi.org/10.1016/j.frl.2023.104273>
- [36] Hegde, A., & Bhowmik, B. (2026). Network Centrality Feature Augmentation and SHAP Analysis of Inlier and Outlier Borrowers in Credit Scoring. *Computational Economics*, . <https://doi.org/10.1007/s10614-025-11292-6>
- [37] Helhel, Y., & Mengüç, E. (2025). Predicting corporate financial performance in the FMCG industry through machine learning: comparing mandatory and voluntary ESG disclosures. *European Journal of Finance*, . <https://doi.org/10.1080/1351847X.2025.2585962>
- [38] Hesse, M., & Loy, T. (2025). Unlocking bankruptcy clues: A novel sentence-based machine learning approach. *International Journal of Accounting Information Systems*, 56, . <https://doi.org/10.1016/j.accinf.2025.100751>
- [39] Jacobs, M. (2024). Benchmarking alternative interpretable machine learning models for corporate probability of default. *Data Science in Finance and Economics*, 4(1), 1-52. <https://doi.org/10.3934/DSFE.2024001>
- [40] Jang, J. (2025). Selective news selection model for explainable stock prediction via cross-attention integration. *Finance Research Letters*, 85, . <https://doi.org/10.1016/j.frl.2025.108146>
- [41] Jiang, S., Pan, X., Guo, Y., Liu, C., & Xiong, H. (2025). Transparent prediction of financial analyst recommendation quality using generalized additive model. *Electronic Commerce Research and Applications*, 73, . <https://doi.org/10.1016/j.elerap.2025.101524>
- [42] Kadkhoda, S., & Amiri, B. (2025). Enhanced Bankruptcy Prediction Model Based on Network Analysis and Explainable Machine Learning. *Computational Economics*, . <https://doi.org/10.1007/s10614-025-11016-w>
- [43] Kang, Y., Ryu, D., & Webb, R. (2025). How well do machine learning models in finance work?. *Financial Innovation*, 11(1), . <https://doi.org/10.1186/s40854-025-00870-0>
- [44] Kanta, K., Narkbunnum, W., & Hinthaw, K. (2026). Enhancing Credit Risk Management and Organizational Adaptation in Thai Agricultural Cooperatives. *Journal Of Organizational Behavior Research*, 11(1), 15-28. <https://doi.org/10.51847/pQQTtkWesl>
- [45] Kellner, R., Nagl, M., & Rösch, D. (2022). Opening the black box – Quantile neural networks for loss given default prediction. *Journal of Banking and Finance*, 134, . <https://doi.org/10.1016/j.jbankfin.2021.106334>
- [46] Kim, M., Benahderrahmane, S., & Rahwan, T. (2024). Interpretable machine learning model for predicting activist investment targets. *Journal of Finance and Data Science*, 10, . <https://doi.org/10.1016/j.jfds.2024.100146>
- [47] Kou, G., Xu, Y., Peng, Y., Shen, F., Chen, Y., Chang, K., & Kou, S. (2021). Bankruptcy prediction for SMEs using transactional data and two-stage multiobjective feature selection. *Decision Support Systems*, 140, . <https://doi.org/10.1016/j.dss.2020.113429>

- [48] Li, Z., Zhu, W., Wu, Y., & Wu, Z. (2025). Reliability of analyst recommendations based on the DS-LightGBM model. *Pacific Basin Finance Journal*, 94, . <https://doi.org/10.1016/j.pacfin.2025.102961>
- [49] Lim, J., & Jeong, B. (2025). The Impact of Cash Holding Decisions on Firm Performance in the IT Industry. *Journal of Risk and Financial Management*, 18(11), . <https://doi.org/10.3390/jrfm18110625>
- [50] Lin, C., Jian, M., Lin, S., & Kuang, X. (2025). Optimal structure of an expected loss credit rating model. *Applied Economics*, 57(31), 4537-4559. <https://doi.org/10.1080/00036846.2024.2364079>
- [51] Lin, L., & Wang, Y. (2025). SHAP Stability in Credit Risk Management: A Case Study in Credit Card Default Model. *Risks*, 13(12), . <https://doi.org/10.3390/risks13120238>
- [52] Lin, Y., Padliansyah, R., Lu, Y., & Liu, W. (2025). Bankruptcy prediction: Integration of convolutional neural networks and explainable artificial intelligence techniques. *International Journal of Accounting Information Systems*, 56, . <https://doi.org/10.1016/j.accinf.2025.100744>
- [53] Liu, J., Li, C., Ouyang, P., Liu, J., & Wu, C. (2023). Interpreting the prediction results of the tree-based gradient boosting models for financial distress prediction with an explainable machine learning approach. *Journal of Forecasting*, 42(5), 1112-1137. <https://doi.org/10.1002/for.2931>
- [54] Mu, Y., Fu, B., & Hu, Q. (2025). The Determinants of Limited Household Participation in Risky Financial Markets: Evidence from China Using Explainable Machine Learning. *Journal of Risk and Financial Management*, 18(12), . <https://doi.org/10.3390/jrfm18120686>
- [55] Nallakaruppan, M., Chaturvedi, H., Grover, V., Balusamy, B., Jaraut, P., Bahadur, J., Meena, V., & Hameed, I. (2024). Credit Risk Assessment and Financial Decision Support Using Explainable Artificial Intelligence. *Risks*, 12(10), . <https://doi.org/10.3390/risks12100164>
- [56] Nazemi, A., & Fabozzi, F. (2024). Interpretable machine learning for creditor recovery rates. *Journal of Banking and Finance*, 164, . <https://doi.org/10.1016/j.jbankfin.2024.107187>
- [57] Nguyen, H., Viviani, J., & Ben, J.S. (2025). Bankruptcy prediction using machine learning and Shapley additive explanations. *Review of Quantitative Finance and Accounting*, 65(1), 107-148. <https://doi.org/10.1007/s11156-023-01192-x>
- [58] Nwafor, C., & Nwafor, O. (2023). Determinants of non-performing loans: An explainable ensemble and deep neural network approach. *Finance Research Letters*, 56, . <https://doi.org/10.1016/j.frl.2023.104084>
- [59] Ozturkkal, B., & Wahlström, R. (2025). Explaining Mortgage Defaults Using SHAP and LASSO. *Computational Economics*, 66(4), 3291-3325. <https://doi.org/10.1007/s10614-024-10763-6>
- [60] Peng, M., Stern, E., & Hu, H. (2025). Forecasting China bond default with severe class-imbalanced data: A simple learning model with causal inference. *Economic Modelling*, 144, . <https://doi.org/10.1016/j.econmod.2024.106985>
- [61] Philips, D., Tilles, D., & Law, T. (2021). Interpretable, Transparent, and Auditable Machine Learning: An Alternative to Factor Investing. *Journal of Financial Data Science*, 3(4), 84-100. <https://doi.org/10.3905/jfds.2021.1.077>
- [62] Phothong, L., Sukprasert, A., Boonlua, S., Chubsuwan, P., Seetha, N., & Kunsrison, R. (2026). An Explainable Voting Ensemble Framework for Early-Warning Forecasting of Corporate Financial Distress. *Forecasting*, 8(1), . <https://doi.org/10.3390/forecast8010010>
- [63] Pollestad, A., & Oust, A. (2025). Harnessing uncertainty: a new approach to real estate investment decision support. *Quantitative Finance*, 25(11), 1771-1788. <https://doi.org/10.1080/14697688.2025.2468269>
- [64] Pourkhoshgoftar, S., Shahbahrani, A., & Esmi, N. (2025). Graph-Based Inductive Learning for Credit Risk Prediction with Imbalance Mitigation. *Computational Economics*, , . <https://doi.org/10.1007/s10614-025-11114-9>
- [65] Radebe, N., Nomlala, B., & Matenda, F. (2026). Forecasting Municipal Financial Distress in South Africa: A Machine Learning Approach. *Forecasting*, 8(1), . <https://doi.org/10.3390/forecast8010018>
- [66] Ramteke, A., Wadhwa, P., & Yan, M. (2022). Interpretability of Machine Learning versus Statistical Credit Risk Models. *Journal of Financial Data Science*, 4(2), 111-124. <https://doi.org/10.3905/jfds.2022.1.089>
- [67] Seow, R. (2025). Transforming ESG Analytics With Machine Learning: A Systematic Literature Review Using TCCM Framework. *Corporate Social Responsibility and Environmental Management*, 32(6), 7358-7389. <https://doi.org/10.1002/csr.70089>

- [68] Shoko, T., Verster, T., & Dube, L. (2025). Comparative analysis of classical and Bayesian optimisation techniques: Impact on model performance and interpretability in credit risk modelling using SHAP and PDPs. *Data Science in Finance and Economics*, 5(3), 320-354. <https://doi.org/10.3934/DSFE.2025014>
- [69] Stef, N., Ben, J.S., Carmona, P., & Scherer, R. (2025). Explainable artificial intelligence modeling for bankruptcy prediction under climate change risks. *Review of Quantitative Finance and Accounting*, , . <https://doi.org/10.1007/s11156-025-01466-6>
- [70] Sun, W., Zhang, X., Li, M., & Wang, Y. (2023). Interpretable high-stakes decision support system for credit default forecasting. *Technological Forecasting and Social Change*, 196, . <https://doi.org/10.1016/j.techfore.2023.122825>
- [71] Sunagatullin, A., & Bahrami, M. (2025). Modular Architectures for Interpretable Credit Scoring for Heterogeneous Borrower Data. *Journal of Risk and Financial Management*, 18(11), . <https://doi.org/10.3390/jrfm18110615>
- [72] Szepannek, G., & Lübke, K. (2023). How much do we see? On the explainability of partial dependence plots for credit risk scoring. *Argumenta Oeconomica*, 2023(1), 137-150. <https://doi.org/10.15611/aoe.2023.1.07>
- [73] Tang, P., Wang, X., Han, Q., & Dong, R. (2025). Predicting Credit Default Risk Crisis of Government Implicit Debt: An Interpretable Machine Learning Approach. *Computational Economics*, , . <https://doi.org/10.1007/s10614-025-11007-x>
- [74] Tang, W., Bu, H., Zuo, Y., & Wu, J. (2024). Unlocking the power of the topic content in news headlines: BERTopic for predicting Chinese corporate bond defaults. *Finance Research Letters*, 62, . <https://doi.org/10.1016/j.frl.2024.105062>
- [75] Teplova, T., Sokolova, T., & Kissa, D. (2023). Revealing stock liquidity determinants by means of explainable AI: The role of ESG before and during the COVID-19 pandemic. *Resources Policy*, 86, . <https://doi.org/10.1016/j.resourpol.2023.104253>
- [76] Wang, J., Jiang, C., Zhou, L., & Wang, Z. (2024). Assessing financial distress of SMEs through event propagation: An adaptive interpretable graph contrastive learning model. *Decision Support Systems*, 180, . <https://doi.org/10.1016/j.dss.2024.114195>
- [77] Wang, L., Yu, Z., Ma, J., Chen, X., & Wu, C. (2025). A Two-Stage Interpretable Model to Explain Classifier in Credit Risk Prediction. *Journal of Forecasting*, 44(7), 2132-2150. <https://doi.org/10.1002/for.3288>
- [78] Xie, S., & Shingadia, T. (2025). Explainable Machine Learning Framework for Predicting Auto Loan Defaults. *Risks*, 13(9), . <https://doi.org/10.3390/risks13090172>
- [79] Ying, C., Shi, A., & Li, X. (2025). Hybrid boosted attention-based LightGBM framework for enhanced credit risk assessment in digital finance. *Humanities and Social Sciences Communications*, 12(1), . <https://doi.org/10.1057/s41599-025-05230-y>
- [80] Zacharias, J., von, Z.M., Chen, J., & Hinz, O. (2022). Designing a feature selection method based on explainable artificial intelligence. *Electronic Markets*, 32(4), 2159-2184. <https://doi.org/10.1007/s12525-022-00608-1>
- [81] Zhang, C., Chen, H., Zhang, X., Chen, M., & Li, Y. (2026). Resolving Interpretability Conflicts of Gradient Boosted Decision Trees (XGBoost) in Credit Risk Assessment: Evidence From Internet Enterprise Finance Innovation A/B Testing. *Journal of Organizational and End User Computing*, 38(1), . <https://doi.org/10.4018/JOEUC.402038>
- [82] Zhang, T., Zhu, W., Wu, Y., Wu, Z., Zhang, C., & Hu, X. (2023). An explainable financial risk early warning model based on the DS-XGBoost model. *Finance Research Letters*, 56, . <https://doi.org/10.1016/j.frl.2023.104045>
- [83] Zhao, J., Zhou, N., & Wan, Z. (2026). Interpretable machine learning for default risk prediction in stress testing. *Journal of Risk Model Validation*, 20(1), . <https://doi.org/10.21314/JRMV.2025.015>