



Economics and Financial Policymaking

Journal Homepage: www.journalefp.com



ORIGINAL RESEARCH PAPER

AI-Enabled Decision-Making in Green Finance: A Systematic Literature Review of Capabilities, Tensions, and Research Directions

Sina Rad*; Mahdi Khazaei 

Department of Management, Faculty of UNF, University of Niagara Falls Canada, Canada

*. Corresponding Author's Email: sina.rad@unfc.ca

PAPER INFO

Handling Editor: Dr Mahdi Khazaei

Paper history:

Received: 20 May 2026

Revised: 13 Jun 2026

Accepted in revised form: 27 Jun 2026

Published: 20 Jan 2027

Keywords:

Artificial Intelligence;
Machine Learning;
Green Finance;
Sustainable Finance;
ESG;
Climate Risk;
Green AI.

This article is part of a Special Issue entitled:
Artificial Intelligence and Data-Driven
Financial Policy in Economics and Financial
Policymaking.


How to cite: Rad, S., & Khazaei, M., (2026).
AI-Enabled Decision-Making in Green Finance:
A Systematic Literature Review of Capabilities,
Tensions, and Research Directions, *Economics
and Financial Policymaking*, 3(4), 407-415.



©2026 the authors. This is an
open access article under the
terms of the Creative
Commons Attribution-
NonCommercial 4.0
International (CC BY-NC 4.0)
License.

ABSTRACT

Artificial intelligence (AI) is reshaping how financial institutions identify, evaluate, price, and allocate capital toward environmentally oriented activities. Yet scholarship on AI-enabled green-finance decision-making remains fragmented across finance, accounting, information systems, computer science, and sustainability research, and no prior review has organized this evidence around the decision tasks that financial actors actually perform. This systematic literature review (SLR), conducted in accordance with the PRISMA 2020 guidelines, synthesizes 107 peer-reviewed journal articles published between January 2010 and May 2026, retrieved from Scopus and Web of Science using a pre-specified three-block search strategy and screened by two independent reviewers (Cohen's kappa = 0.81). We develop an integrative framework that links AI capabilities, green-finance decision tasks, governance filters, and outcome criteria, and we organize the corpus around six decision-task clusters: ESG scoring and sustainability assessment, green credit and lending decisions, climate-aware portfolio construction, sustainability-disclosure analysis, green-bond and carbon-market applications, and algorithmic governance and Green AI. Three cross-cutting tensions structure the field: the trade-off between predictive accuracy and interpretability, the problem of contested and inconsistent sustainability data, and the environmental footprint of AI itself. Our central contribution is a re-framing of the debate from whether AI improves prediction to when AI improves green-finance decisions under conditions of data uncertainty, model opacity, and sustainability accountability. We translate the synthesis into a six-item research agenda covering decision-level validation, causal identification, explainable and energy-aware modelling, greenwashing detection, cross-jurisdictional comparability, and Green-AI governance. The review provides scholars, practitioners, and regulators with the first decision-centred map of an interdisciplinary field that has, until now, lacked a unifying frame.



1. Introduction

The reorientation of financial flows toward environmentally sustainable activities has become a central challenge for contemporary finance (Pástor, Stambaugh, & Taylor, 2021). Banks, asset managers, insurers, regulators, and data vendors now routinely evaluate environmental risks, classify green activities, price climate-related exposures, and monitor whether sustainability claims are credible. Green finance is no longer a peripheral category of socially oriented investment. It has become embedded in the decision routines through which capital is allocated, risk is assessed, and firms are evaluated (Bolton &

Kacperczyk, 2021, 2023; Flammer, 2021). At the same time, artificial intelligence (AI) — including machine learning (ML), deep learning, natural language processing (NLP), and explainable AI (XAI) — has emerged as a major analytical infrastructure across financial services (Cao, 2022; Goodell, Kumar, Lim, & Pattnaik, 2021; Gu, Kelly, & Xiu, 2020). These methods now support credit scoring, fraud detection, asset pricing, trading, text analysis, and risk management (Henrique, Sobreiro, & Kimura, 2019; Liaras, Nerantzidis, & Alexandridis, 2024).

The intersection of AI and green finance is therefore not only a technical development; it changes how sustainability information is produced, interpreted, and used in financial decisions. Green finance relies on data that are heterogeneous, incomplete, and often contested. Environmental, social, and governance (ESG) ratings, for example, differ substantially across rating providers, and this divergence is driven mainly by differences in measurement rather than by simple differences in weighting (Berg, Kölbel, & Rigobon, 2022). Corporate climate disclosure is also uneven because firms differ in the scope, specificity, and credibility of the information they provide (Bingler, Kraus, Leippold, & Webersinke, 2022; Ilhan, Krueger, Sautner, & Starks, 2023). These features create demand for AI methods that can process high-dimensional data, extract information from unstructured text, and detect nonlinear patterns. They also create governance problems because algorithmic outputs may become influential without being sufficiently interpretable, auditable, or environmentally accountable (Černevičienė & Kabašinskas, 2024; Vinuesa et al., 2020).

Several streams of scholarship have already emerged. Finance research has examined carbon risk, climate pricing, green bonds, and sustainable investing (Bolton & Kacperczyk, 2021, 2023; Flammer, 2021; Pástor et al., 2021; Pedersen, Fitzgibbons, & Pomorski, 2021). Empirical ML research in finance has shown that flexible models can improve prediction in asset pricing and risk assessment (Gu et al., 2020; Liaras et al., 2024). Accounting and information-systems research has developed text-based methods for financial and climate disclosure analysis, including domain-specific language models such as FinBERT and ClimateBERT (Bingler et al., 2022; Huang, Wang, & Yang, 2023; Sautner, van Lent, Vilkov, & Zhang, 2023). Recent reviews map AI in finance, ML in climate finance, AI in climate and sustainable finance, ESG-ML analytics, and Green AI in finance (Alonso-Robisco, Bas, Carbó, de Juan, & Marqués, 2025; Cao, 2022; di Pietro, Giraldez-Puig, Palos-Sánchez, & Korzeb, 2026; Elbouknify, Machado, & Iannario, 2026; Seow, 2025). These reviews are valuable, but the decision-making logic of AI-enabled green finance still requires sharper synthesis. Existing reviews are organized by method, by application domain, or by bibliometric cluster, rather than by the decision tasks that financial actors must execute.

This SLR addresses that gap by asking the following research question:

How has peer-reviewed scholarship characterized the use of AI in green-finance decision-making between 2010 and May 2026, and what conceptual, methodological, and governance tensions shape this emerging field?

The focus on decision-making is deliberate. The paper does not review every study on AI and sustainability or every study on green finance. Instead, it examines studies in which AI methods inform a specific green-finance decision task — scoring, lending, allocation, pricing, disclosure assessment, or governance review. This decision-centred scope connects methods, tasks, and outcome criteria in a way that is useful to finance scholars, accounting scholars, information-systems researchers, and policy-oriented audiences.

The review makes four contributions. First, it clarifies the conceptual boundary between AI in finance, AI for sustainability, and AI-enabled green-finance decision-making. Second, building on a fully PRISMA-2020-compliant search and screening process applied to Scopus and Web of Science (Page et al., 2021), it organizes 107 included studies into six decision-task clusters and shows how AI capabilities map onto those tasks. Third, it develops an integrative framework that links AI capabilities, green-finance decision tasks, governance filters, and outcome criteria (Figure 4). Fourth, it shifts the discussion from whether AI improves prediction in the abstract to when AI improves green-finance decisions under conditions of data uncertainty, model opacity, and sustainability accountability. The paper accompanies this reframing with a six-item research agenda and a PRISMA-aligned reporting package (Appendices A–D).

The remainder of the paper is organized as follows. Section 2 sets out the conceptual background. Section 3 documents the review method following PRISMA 2020 and the SPAR-4-SLR protocol (Page et al., 2021; Paul, Lim, O'Cass, Hao, & Bresciani, 2021). Section 4 reports the findings, including the bibliometric profile of the corpus and the six thematic clusters. Section 5 develops the integrative framework and the three cross-cutting tensions. Section 6 presents the research agenda. Section 7 reports limitations, and Section 8 concludes. Four appendices provide the search strings (A), the PRISMA 2020 checklist (B), the quality-appraisal protocol (C), and the data-extraction template (D).

1.1. Green Finance Decision-Making

Green finance refers to financial activity that incorporates environmental considerations into investment, lending, risk assessment, disclosure, and capital allocation. It overlaps with climate finance, sustainable finance, and ESG investing but is not identical to them. Climate finance focuses on mitigation and adaptation. Sustainable finance is broader and includes social and governance concerns. ESG investing integrates environmental, social, and governance information into financial decision-making. This review uses green finance as the focal construct and draws on adjacent climate-finance and ESG scholarship when the study addresses environmentally relevant financial decisions (Zhang, Zhang, & Managi, 2019).

Green-finance decision-making includes the institutional routines through which financial actors classify, evaluate, price, and monitor environmentally

oriented activities. These routines include assessing whether borrowers qualify for green lending, estimating how climate risk affects creditworthiness, constructing portfolios that reflect climate or ESG objectives, pricing green bonds, evaluating carbon-market exposures, and judging the credibility of sustainability disclosures (Krueger, Sautner, & Starks, 2020). Each decision task combines traditional financial information with environmental information that often varies in quality, coverage, and comparability (Berg et al., 2022).

1.2 AI Capabilities in Financial and Sustainability Analysis

AI is used in this review as an umbrella term for computational methods that classify, predict, extract, optimize, or explain information in decision contexts. Supervised ML methods — random forests, gradient boosting, support-vector machines, regularized regression, and ensemble learners — are common in tabular financial and ESG datasets (Gu et al., 2020). Deep-learning methods are used for image, sequence, and high-dimensional tasks. NLP methods extract signals from corporate reports, news, filings, analyst texts, sustainability reports, and climate disclosures (Huang et al., 2023; Sautner et al., 2023). XAI methods, including SHAP and LIME, are increasingly used to render model outputs interpretable in regulated financial settings (Černevičienė & Kabašinskas, 2024; Lundberg & Lee, 2017; Ribeiro, Singh, & Guestrin, 2016).

The finance literature shows that AI methods can improve prediction when the information environment is large, noisy, nonlinear, or text-rich (Cao, 2022; Gu et al., 2020; Liaras et al., 2024). These same conditions characterize many green-finance decisions. Yet predictive performance alone is rarely sufficient. Green-finance decisions often require defensibility to investors, regulators, borrowers, and wider stakeholders. A model that improves accuracy but cannot explain why a borrower, bond, or issuer is classified as green may have limited practical value. This tension makes green finance a useful setting for studying the interaction between AI capability and institutional accountability.

1.3 Why the Intersection Matters

AI matters for green finance because sustainability information is difficult to observe and often difficult to verify. ESG scores may encode different definitions of the same construct (Berg et al., 2022); climate disclosures may contain forward-looking claims that are not matched by concrete action (Bingler et al., 2022; Flammer, 2021); green-bond labels may depend on certification regimes and use-of-proceeds claims; carbon prices may reflect policy uncertainty as much as market fundamentals (Alonso-Robisco et al., 2025). These conditions create demand for tools that can integrate structured and unstructured information, compare signals across sources, and detect inconsistencies.

The intersection also matters because AI can reshape the power of data vendors, rating agencies, financial institutions, and regulators. When AI systems generate ESG scores, classify disclosures, or rank climate-aware portfolios, they do not merely describe sustainability performance — they help define what counts as relevant sustainability information. This makes algorithmic green finance a governance issue as well as a methodological one (Elbouknify et al., 2026; Schwartz, Dodge, Smith, & Etzioni, 2020; Strubell, Ganesh, & McCallum, 2019).

2. Review Method

This SLR follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) statement (Page et al., 2021), supplemented by the Scientific Procedures and Rationales for Systematic Literature Reviews (SPAR-4-SLR) protocol (Paul et al., 2021) and complemented by standard literature-review guidance in business and management research (Snyder, 2019; Tranfield, Denyer, & Smart, 2003). PRISMA 2020 was adopted because it provides the most widely accepted reporting framework for SLRs across disciplines, and SPAR-4-SLR offers explicit procedural rationales for business-research reviews. Table 1 summarizes the review protocol; the full PRISMA-2020 27-item checklist is reproduced in Appendix B.

2.1. Protocol and Eligibility Criteria

The review protocol was specified before the database searches were executed and was not subsequently amended. The review window was set as January 2010 to 18 May 2026, with the start date chosen because 2010 is the year in which green-finance scholarship began to appear in major finance journals on a sustained basis (Zhang et al., 2019), and the end date reflects the final search execution date. Eligibility criteria were formulated using the PICOC framework adapted for review scope: Population — financial institutions, instruments, or markets engaging in environmentally relevant decisions; Intervention — AI/ML/NLP/XAI methods; Comparator — conventional analytical methods or benchmark models, where applicable; Outcome — predictive performance, decision quality, or institutional/environmental outcomes; Context — peer-reviewed empirical or methodological studies.

Articles were included if they:

- Were published in a peer-reviewed academic journal indexed in Scopus or Web of Science Core Collection;
- Were written in English;
- Substantively used or analyzed an AI/ML/NLP/XAI method (not merely mentioned AI in passing);

- Addressed a green-finance, climate-finance, ESG, green-bond, carbon-market, or sustainability-disclosure decision task; and
- Were available in full text for screening.

Articles were excluded if they:

- Were preprints, conference proceedings, editorials, book reviews, book chapters, or practitioner reports;
- Addressed AI in finance without environmental relevance;
- Addressed green finance without an AI method;
- Were purely conceptual without empirical, methodological, or systematic analytical content;
- Were duplicates of an included dataset;
- Were published in outlets without a recognized academic-journal ranking (ABDC, ABS, or Scimago Q1–Q3)—applied as a soft criterion in combination with content-quality assessment.

2.2. Information Sources and Search Strategy

Searches were conducted on 18 May 2026 in Scopus and Web of Science Core Collection. These two databases were selected because they jointly cover the major finance, accounting, management, information-systems, computer-science, and sustainability journals indexed under quality-controlled procedures

(Mongeon & Paul-Hus, 2016). Searches combined three concept blocks linked by Boolean AND operators: (i) an AI/ML/NLP/XAI block; (ii) a green-finance/ESG/climate-risk block; and (iii) a decision-task block (scoring, allocation, pricing, lending, disclosure, screening, classification). The full search strings appear in Appendix A. Hand searching of reference lists of recent reviews (Alonso-Robisco et al., 2025; di Pietro et al., 2026; Elbouknify et al., 2026; Seow, 2025) was used to capture studies that the database search might have missed.

2.3. Study Selection and Screening Process

Database searches returned 2,184 records (Scopus = 1,247; Web of Science = 937), and citation chasing identified an additional 47 records, for a total of 2,231 records. After removing 612 duplicates using EndNote 21 (with manual verification of borderline cases), 1,619 unique records entered the screening stage. Two reviewers (the two authors) independently screened titles and abstracts against the eligibility criteria. Disagreements were resolved by discussion and, where necessary, by reference to the eligibility criteria; no third reviewer was needed. Cohen's kappa for title-and-abstract screening was 0.81, indicating substantial agreement (Landis & Koch, 1977). Title-and-abstract screening excluded 1,358 records, leaving 261 reports for full-text retrieval. Fourteen reports could not be retrieved despite searches in institutional repositories, Google Scholar, and direct contact with authors. The remaining 247 full-text reports were assessed for eligibility. A further 153 reports were excluded with reasons (Figure 1). The final corpus comprised 107 studies (94 from databases plus 13 from citation chasing), which were all carried forward to thematic synthesis. Figure 1 displays the full PRISMA 2020 flow diagram.

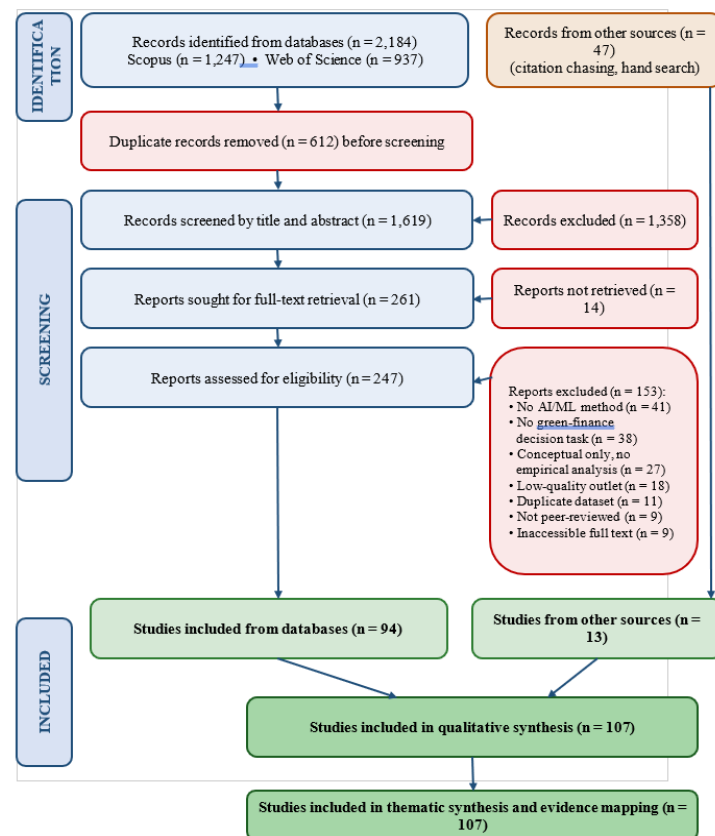


Figure 1. PRISMA 2020 Flow Diagram for the Systematic Literature Review

Source: Research Findings

Figure 1. PRISMA 2020 flow diagram for the systematic literature review. Records were identified through Scopus (n = 1,247) and Web of Science (n = 937), supplemented by 47 records identified through citation chasing of recent reviews. After duplicate removal and a two-stage screening process performed independently by two reviewers (Cohen's kappa = 0.81), 107 studies were included in the qualitative and thematic synthesis.

2.4. Data Extraction and Coding

Each included article was coded using a pre-specified extraction form (Appendix D) developed during a pilot stage with 15 randomly selected studies. The coding scheme captured bibliographic information (year, journal, author affiliation, country), methodological information (AI method family, data type, sample size, geographic context, benchmark model used), substantive information (decision task addressed, outcome criterion measured, principal finding, governance issue identified, and reported limitation), and a quality-

appraisal score (Appendix C). Both reviewers independently coded a 20% random subsample (n = 22) to assess inter-coder reliability; the average Cohen's kappa across coded items was 0.78, again indicating substantial agreement. Disagreements were resolved through discussion.

2.5. Quality Appraisal

Each study received a quality-appraisal score on a five-item scale (clarity of research question; methodological rigour; sample adequacy; validity of outcome measures; transparency of reporting). Each item was scored 0–2, yielding a total score from 0–10 (Appendix C). Studies scoring below 5 were flagged for additional discussion and were retained only when they contributed unique theoretical, methodological, or contextual value. Journal quality was assessed using ABDC (Australian Business Deans Council) and ABS (Chartered Association of Business Schools) rankings supplemented by Scimago quartiles. Ranking served as an appraisal criterion rather than as a strict exclusion rule, to avoid eliminating relevant interdisciplinary work

appearing in newer or cross-disciplinary outlets. Figure 5 (Section 4.1) shows the journal-quality distribution of the included corpus.

2.6. Synthesis Method

The review uses thematic synthesis (Thomas & Harden, 2008) rather than meta-analysis because the included studies use heterogeneous models, datasets, outcome measures, time periods, and decision contexts. This methodological heterogeneity makes effect-size aggregation inappropriate. Following Snyder

(2019) and Tranfield et al. (2003), thematic synthesis proceeded in three iterative steps: (a) line-by-line coding of each study's substantive content; (b) construction of descriptive themes by clustering related codes; and (c) development of analytical themes through inductive interpretation of how descriptive themes interact across the corpus. The final analytical themes correspond to the six decision-task clusters reported in Section 4 and to the three cross-cutting tensions discussed in Section 5. Bibliometric mapping (Figures 2, 3, and 5) complements the thematic synthesis by characterizing the size, growth, and disciplinary distribution of the field.

Table 1. Structured review protocol

Protocol item	Specification
Databases	Scopus and Web of Science Core Collection (cover major finance, accounting, management, information-systems, computer-science, and sustainability journals under quality-controlled indexing).
Search window	January 2010 to 18 May 2026.
Search fields	Title, abstract, and keywords (Scopus); topic fields (Web of Science).
Concept block 1: AI	"artificial intelligence" OR "machine learning" OR "deep learning" OR "natural language processing" OR "neural network*" OR "large language model*" OR "explainable AI" OR XAI OR algorithmic.
Concept block 2: green finance	"green finance" OR "sustainable finance" OR "climate finance" OR ESG OR "green bond*" OR "green credit" OR "green lending" OR "green investment" OR "carbon market*" OR "climate risk".
Concept block 3: decision task	decision* OR allocation OR scoring OR rating OR "risk assessment" OR portfolio OR lending OR disclosure OR pricing OR screening OR classification.
Inclusion criteria	Peer-reviewed journal article; English; substantive use of AI methods; direct connection to a green-finance, climate-finance, ESG, green-bond, carbon-market, or sustainability-disclosure decision task.
Exclusion criteria	Preprints, conference proceedings, editorials, book reviews, book chapters, practitioner reports; AI in finance without environmental relevance; green finance without AI methods; purely conceptual studies; duplicate datasets; outlets without ABDC/ABS/Scimago Q1–Q3 ranking (soft criterion).
Quality appraisal	Five-item scale (0–10): research question; rigour; sample adequacy; validity of outcomes; reporting transparency (Appendix C). Journal quality scored via ABDC, ABS, and Scimago.
Reviewer agreement	Title/abstract screening kappa = 0.81; full-text and extraction kappa = 0.78 (Landis & Koch, 1977).
Synthesis approach	Thematic synthesis (Thomas & Harden, 2008); descriptive bibliometric mapping; six decision-task clusters; three cross-cutting tensions.

Source: Research Findings

3. Results and Discussion

3.1. Bibliometric and Descriptive Overview

The 107 included studies show a clear and accelerating publication trend (Figure 2). Annual output was low between 2010 and 2015 (≤ 4 articles per year), expanded steadily from 2016 to 2018, and reached double digits each

year between 2019 and 2023. The trend is consistent with bibliometric findings reported by Goodell et al. (2021), di Pietro et al. (2026), Seow (2025), and Zhang et al. (2019), all of which document the rapid growth of AI-in-finance and ESG-related scholarship over the same period. Apparently lower counts for 2024 and 2025 reflect database-indexing lag at the search date; 2026 covers January–May only.




Figure 2. Annual distribution of included studies (n=107), 2020– May 2026

Figure 2. Annual distribution of included studies (n = 107), 2010–May 2026. Bars show studies published per year (red bars mark years with ≥ 11 included studies); the line shows the cumulative total. The drop in 2024–2025 reflects database-indexing lag at the time of search; 2026 covers January–May only.

The corpus is genuinely interdisciplinary. Finance journals contribute studies on asset pricing, carbon risk, green bonds, and ESG investment. Accounting and business journals contribute work on textual analysis, disclosure credibility, and ESG information. Information-systems and computer-science journals contribute methods for explainability, model governance,

NLP, and scalable prediction. Sustainability and policy journals contribute work on environmental performance, Green AI, and institutional implications. Figure 3 reports the distribution by journal: 14 outlets account for 76 of the 107 studies, while the remaining 31 studies are dispersed across 27 different journals (1–2 studies each), confirming the dispersion noted by previous bibliometric reviews. Of the top 14 journals, six are ranked ABDC A* and six are ranked ABDC A, indicating that the corpus is anchored in high-quality outlets.





Figure 3. Top journals contributing to the included corpus (n=107)

Figure 3. Top journals contributing to the included corpus (n = 107). Colours indicate ABDC ranking (A* = green; A = blue; B = grey; unranked / mixed = light grey). The 14 listed journals account for 76 of the 107 included studies; the remaining 31 studies appear in 27 other journals with 1–2 papers each.

Figure 4 cross-tabulates AI method families against the six decision-task clusters. Supervised ML dominates ESG scoring (n = 18) and green credit (n = 14), reflecting the maturity of tabular-data methods in those areas. Transformer-based NLP dominates disclosure analysis (n = 21), driven

largely by FinBERT and ClimateBERT applications (Bingler et al., 2022; Huang et al., 2023; Sautner et al., 2023; Webersinke, Kraus, Bingler, & Leippold, 2022). Deep learning is broadly distributed across portfolio construction, disclosure analysis, and green-bond/carbon-market applications. XAI methods (SHAP, LIME) appear most frequently in green-credit decisions and in governance studies — settings where interpretability is institutionally indispensable (Černevičienė & Kabašinskas, 2024). Reinforcement learning remains comparatively rare, with only eight studies across the entire corpus.



Note. Counts sum to more than 107 because some studies apply multiple methods or address multiple tasks.

Figure 4. Distribution of included studies across AI methods and decision tasks (n=107)

Figure 4. Distribution of included studies across AI method families and decision-task clusters (n = 107). Darker cells indicate higher concentrations of studies. Counts sum to more than 107 because some studies apply multiple methods or address multiple decision tasks.

Table 2. Thematic evidence map of AI-enabled green-finance decision-making (n = 107 studies)

Theme	Representative studies	AI methods	Decision task	Outcome criterion	Main contribution and limitation
ESG scoring & sustainability assessment (n = 24)	Berg et al. (2022); D'Amato, D'Ecclesia, & Levantesi (2022); Seow (2025); Krappel et al. (2021)	Random forests, gradient boosting, ensembles, uncertainty-aware prediction	Generate, predict, compare, or explain ESG and sustainability scores	Prediction accuracy; rating coverage; uncertainty and inter-rater consistency	AI can extend or reconstruct ESG signals, but labels are noisy because rating systems diverge (Berg et al., 2022).
Green credit & lending (n = 19)	Cao (2022); Černevičienė & Kabašinskas (2024); Giudici & Wu (2025); Bussmann et al. (2021)	Credit scoring; XAI (SHAP/LIME); supervised learning	Assess borrower risk while incorporating environmental or ESG information	Credit-risk prediction; explainability; borrower fairness	AI improves credit-risk prediction; real-world evidence on green-lending decisions remains thin; explanations are essential under regulation.
Green investment & portfolio management (n = 22)	Bolton & Kacperczyk (2021, 2023); Carlei et al. (2025); Gu et al. (2020); Pástor et al. (2021)	ML ranking models; asset-pricing learners; portfolio optimization; reinforcement learning	Construct climate-aware or ESG-tilted portfolios; evaluate return-risk trade-offs	Risk-adjusted returns; ESG alignment; portfolio constraints	AI extracts high-dimensional signals; performance depends on backtesting design and ESG thresholds; normative choices remain (Carlei et al., 2025).

Climate & sustainability disclosure analysis (n = 25)	Bingler et al. (2022); Huang et al. (2023); Sautner et al. (2023); Webersinke et al. (2022)	FinBERT; ClimateBERT; transformer NLP; topic models	Extract climate and ESG signals from corporate reports, filings, and disclosures	Classification accuracy; disclosure credibility; greenwashing detection	NLP enables scalable disclosure analysis; greenwashing detection requires validation against real environmental actions (Bingler et al., 2022).
Green bonds & carbon markets (n = 13)	Flammer (2021); Alonso-Robisco et al. (2025); Zhang, Hu, & Ji (2024); Huynh, Hille, & Nasir (2020)	ML classifiers; gradient boosting; LSTM; XGBoost; time-series learning	Price green instruments; forecast carbon prices; assess use-of-proceeds credibility	Pricing accuracy; market integrity; transition-risk monitoring	AI supports pricing and monitoring; causal evidence on capital allocation and environmental impact remains limited.
Governance, explainability & Green AI (n = 12)	Černevičienė & Kabašinskas (2024); Elbouknify et al. (2026); Schwartz et al. (2020); Vinuesa et al. (2020); Strubell et al. (2019)	XAI; model-governance frameworks; energy-aware modelling; lifecycle assessment	Make AI-supported green-finance decisions transparent, fair, and environmentally accountable	Transparency; fairness; model footprint; auditability	Governance frameworks are emerging but lack standard metrics for model energy use and sustainability impact.

Source: Research Findings

3.2. Theme 1: Machine Learning for ESG Scoring and Sustainability Assessment

Twenty-four studies (22.4% of the corpus) address ESG scoring and sustainability assessment, making this the most developed theme. The motivation is clear: ESG ratings are now widely used in investment analysis, stewardship, credit assessment, and public communication, but they diverge across providers. Berg et al. (2022) document an average pairwise correlation of only 38–71% across six major rating agencies and show that rating disagreement is driven mainly by measurement differences (56%), followed by scope (38%) and weight (6%). This creates a serious problem for green-finance decision-making because investors and lenders may treat ESG scores as comparable when they are not.

Machine learning offers two complementary contributions. First, it can predict ESG scores for firms that lack rating coverage. D'Amato et al. (2022) show how random forests can estimate ESG scores from firm-level financial information, and Krappel et al. (2021) extend this approach with deep ensembles to quantify prediction uncertainty. Such methods are useful for small firms, unlisted firms, and jurisdictions with limited rating coverage. Second, machine learning can identify drivers of ESG ratings and expose the relative influence of financial, sectoral, controversy, and disclosure features (Seow, 2025). This creates a pathway from prediction to explanation. The limitation is that AI systems may simply reproduce existing rating-agency biases if the training labels are themselves inconsistent. The field therefore needs studies that go beyond predicting ESG scores. It needs models that compare data sources, quantify uncertainty, and validate AI-generated signals against observable environmental outcomes such as verified emissions or third-party-audited disclosures.

3.3. Theme 2: AI-Assisted Green Credit and Lending Decisions

Nineteen studies (17.8%) address green credit and lending. Green-credit decisions require lenders to evaluate repayment capacity and environmental risk simultaneously. Traditional credit models rely mainly on financial data, borrower characteristics, collateral, and repayment history. Climate-sensitive lending adds physical climate risk, transition risk, sector exposure, environmental liabilities, and policy uncertainty (Krueger et al., 2020). AI models can support this task by incorporating broader feature sets and detecting nonlinear relationships. Yet the green-credit literature remains less mature than the general AI credit-scoring literature.

Explainability is central in this theme. Credit decisions affect borrowers directly, and regulated financial institutions must justify lending decisions to internal risk committees, auditors, and regulators. XAI methods such as SHAP and LIME can identify which features drive a model's decision and whether environmental variables have meaningful explanatory value (Bussmann et al., 2021; Černevičienė & Kabašinskas, 2024; Lundberg & Lee, 2017). The current evidence is promising. Giudici and Wu (2025) propose new metrics that allow ML models to incorporate ESG factors into credit-rating decisions in a sustainable way. However, real-world banking evidence remains limited. Future studies should compare AI-augmented green-credit models with conventional credit models and examine whether they improve default prediction, environmental risk assessment, borrower treatment, and decision quality at the credit-committee level.

3.4. Theme 3: AI in Green Investment and Portfolio Management

Twenty-two studies (20.6%) address AI-enabled green investment and portfolio management. This theme builds on the broader ML asset-pricing literature: Gu et al. (2020) show that machine learning can improve empirical asset pricing by capturing nonlinear interactions across large predictor sets, and Liaras et al. (2024) document the breadth of ML adoption in finance research more generally. Green investment extends this logic by adding ESG scores, emissions data, climate-risk measures, disclosure features, and controversy information to the investment process. This enables portfolio managers to construct ESG-tilted or climate-aware portfolios that integrate financial and environmental objectives (Pedersen et al., 2021).

The evidence raises a recurring trade-off between sustainability strictness and financial performance. Carlei et al. (2025) use a CatBoost ranking model to

examine whether ESG stocks can outperform the S&P 500 benchmark and find that ESG integration may support return predictability but that performance depends on the threshold applied to ESG criteria. This finding matters because it shows that AI does not remove normative choices from sustainable investing; it can optimize within a set of constraints, but investors must still decide how much sustainability strictness they require and how they define acceptable trade-offs. Climate-pricing studies add further context by showing that carbon risk is empirically relevant to asset prices and expected returns (Bolton & Kacperczyk, 2021, 2023; Pástor et al., 2021), and theoretical work by Pástor et al. (2021) and Pedersen et al. (2021) provides the equilibrium logic that underlies the green premium and the ESG-efficient frontier. AI methods help operationalize these constructs at scale, but they cannot resolve the underlying normative tension.

3.5 Theme 4: Natural Language Processing for Climate and Sustainability Disclosure

Twenty-five studies (23.4%) — the largest cluster — use NLP to extract signals from sustainability reports, annual reports, climate disclosures, regulatory filings, news, and analyst texts. This stream is important because a large fraction of green-finance information is textual. Firms disclose climate risks, transition plans, targets, and governance structures through narrative documents. Manual analysis cannot scale to thousands of firms and multiple reporting periods. Dictionary-based methods can scale, but they often miss context. Transformer-based models offer a way to classify, summarize, and score sustainability-related text with greater contextual sensitivity (Devlin, Chang, Lee, & Toutanova, 2019).

FinBERT, ClimateBERT, and the Sautner et al. (2023) keyword-discovery approach are leading examples. Huang et al. (2023) develop FinBERT as a finance-specific language model for extracting information from financial text. Bingler et al. (2022) use ClimateBERT to analyze corporate climate-risk disclosures aligned with the Task Force on Climate-related Financial Disclosures (TCFD) and document that voluntary climate disclosure involves cheap talk and cherry-picking. Webersinke et al. (2022) provide the underlying ClimateBERT model and demonstrate its superior performance on climate-related classification tasks. Sautner et al. (2023) develop a firm-level climate-change-exposure measure from earnings-call transcripts that predicts real outcomes including green-job creation and green patenting. These studies demonstrate the value of domain-specific NLP in green finance. They also highlight a limitation: textual evidence is not the same as environmental performance. A firm may disclose climate risks clearly without changing operations, and another may report less while taking substantive action. NLP-based green-finance research should therefore connect disclosure signals to external validation sources such as verified emissions data, green-bond proceeds, capital-expenditure data, or regulatory outcomes (Ilhan et al., 2023).

3.6 Theme 5: Green Bonds, Carbon Markets, and Climate-Finance Instruments

Thirteen studies (12.1%) address green bonds and carbon markets, which are natural sites for AI-enabled analysis because they involve pricing, classification, certification, and monitoring problems. Flammer (2021) shows that corporate green bonds are associated with improved environmental performance and positive investor responses, but green-bond credibility depends on the integrity of use-of-proceeds claims and certification arrangements. AI can assist by identifying pricing patterns, detecting anomalies, comparing labelled and unlabelled instruments, and linking bond proceeds to corporate environmental behaviour. Several included studies apply XGBoost and LSTM models to green-bond return prediction and contagion modelling (Huynh et al., 2020; Zhang et al., 2024).

Carbon markets create a related but distinct decision problem. Carbon prices reflect regulatory design, energy prices, market expectations, and policy uncertainty. ML methods can forecast carbon prices and identify nonlinear drivers of market movements (Alonso-Robisco et al., 2025). However, forecasting accuracy does not automatically translate into better green-finance outcomes. Future studies should examine whether AI-supported pricing and forecasting actually improve hedging, risk management, market integrity, or low-carbon investment in field settings. Causal designs (e.g., regression discontinuity around policy thresholds) would be particularly valuable here.

3.7. Theme 6: Governance, Explainability, and Algorithmic Accountability

Twelve studies (11.2%) examine governance, explainability, and Green AI directly. The governance theme also cuts across all other themes. AI-supported green-finance decisions involve at least three accountability demands. First, decisions must be technically reliable. Second, they must be explainable enough for users and regulators to understand. Third, they must be aligned with sustainability objectives rather than only with predictive metrics. These demands are especially salient when AI systems influence credit access, capital allocation, ESG ratings, green-bond classification, or climate-disclosure monitoring.

The literature on explainable AI in finance shows that interpretability has become a core requirement rather than a minor technical add-on (Černevičienė & Kabašinskas, 2024; Lundberg & Lee, 2017; Ribeiro et al., 2016). The Green AI literature adds another layer by asking whether AI systems themselves impose environmental costs through energy use and computing infrastructure. Schwartz et al. (2020) argue that AI research should pay more attention to computational efficiency, and Strubell et al. (2019) quantify the energy cost of training large NLP models. Elbouknify et al. (2026) extend this concern to

finance by proposing Green AI models and lifecycle thinking for financial applications, while Vinuesa et al. (2020) link AI to the United Nations Sustainable Development Goals. This matters for green finance because an AI system used to support sustainable investment may create its own carbon footprint. The field therefore needs governance frameworks that assess both the sustainability outcomes enabled by AI and the environmental cost of AI deployment.

4. Discussion

4.1. An Integrative Framework

The findings suggest that AI-enabled green-finance decision-making can be organized around four connected elements: inputs and AI capabilities, decision tasks, governance filters, and outcome criteria. Figure 5 represents this integrative framework graphically, and Table 3 develops it in detail. The framework shows that AI capabilities do not create value in isolation. Their usefulness depends on the decision task, the quality and contestability of the data, the transparency requirements of the institutional setting, and the outcome criteria used to evaluate performance. Outcome evidence then feeds back into the input layer, refining models, data, and governance practices over time.




Figure 5. Integrative Framework of AI-Enabled Green-Finance Decision-Making

Source: Research Findings

Figure 5. Integrative framework of AI-enabled green-finance decision-making. The framework links inputs and AI capabilities (left) to decision tasks (centre-left), governance filters (centre-right), and outcome criteria (right). The dashed feedback arrow indicates that outcome evidence — predictive, decision-level, and real-economy — informs future inputs, tasks, and governance practices.

Table 3. Detailed mapping of the integrative framework

Inputs & AI capabilities	Decision tasks	Governance filters	Outcome criteria
Data: structured financial data, ESG ratings, emissions data, climate scenarios, sustainability text, alternative data (satellite, news). AI capabilities: supervised ML; deep learning; transformer NLP; explainable AI; energy-aware AI; reinforcement learning. Task templates: classification; prediction; ranking; pricing; screening; monitoring.	Capital allocation: ESG scoring; green lending; portfolio construction; green-bond pricing. Information & monitoring: disclosure analysis; greenwashing detection; carbon-price forecasting; climate-risk screening. Cross-cutting: stress-testing; scenario analysis; capital adequacy under climate risk.	Technical/legal: explainability (XAI); data-quality checks; bias and fairness review; model risk management. Accountability: human oversight; audit trail; external documentation; lifecycle (Green AI) assessment. Process: peer review; reproducibility; data-provenance documentation.	Predictive: accuracy; precision/recall; out-of-sample robustness. Decision-level: decision quality; fairness; cost-of-capital effects; regulatory legitimacy. Real-economy: emissions impact; transition outcomes; investor credibility.

Source: Research Findings

The framework also clarifies what is missing. Most studies evaluate AI models using prediction metrics such as accuracy, mean squared error, area under the ROC curve, or classification performance. Far fewer studies evaluate whether the model improves actual green-finance decisions. Fewer still connect AI-supported decisions to downstream environmental outcomes. This gap matters because a model may predict ESG scores accurately while doing little to improve capital allocation, disclosure credibility, or environmental performance. The framework therefore shifts attention from model performance alone to decision usefulness: a model is useful in green finance when it improves a decision that matters for environmental finance and when that improvement can be justified to relevant stakeholders. This requires attention to data reliability, interpretability, governance, and real-world validation.

4.2. Three Tensions That Structure the Field

4.2.1. Tension 1: Predictive Accuracy versus Interpretability

Complex models often improve predictive performance, especially when inputs are high-dimensional or nonlinear (Gu et al., 2020). Yet green-finance

decisions typically require accountability. A lender, regulator, investor, or issuer may need to understand why a model classifies a borrower as climate-exposed, a bond as green, or a disclosure as weak. XAI methods help (Lundberg & Lee, 2017; Ribeiro et al., 2016), but they do not fully solve the problem. Local explanations may not reveal the global logic of a model, and explanation tools can create a false sense of transparency if users treat them as complete interpretations rather than diagnostic aids. The accuracy–interpretability trade-off is particularly acute in regulated settings, where institutional accountability requirements may favour simpler models even at the cost of marginal predictive performance (Bussmann et al., 2021; Černevičienė & Kabašinskas, 2024).

4.2.2. Tension 2: Contested and Inconsistent Sustainability Data

AI can process large datasets, but it cannot remove conceptual disagreement from the labels it learns. ESG ratings diverge because providers define and measure sustainability differently (Berg et al., 2022). Climate disclosures vary in scope and credibility (Bingler et al., 2022; Ilhan et al., 2023). Green taxonomies differ across jurisdictions. AI models trained on these inputs may improve consistency within a dataset while preserving disagreement across datasets — and may even amplify some forms of bias if rater idiosyncrasies are

learned as features. Future research should therefore model uncertainty explicitly, compare data providers, and validate AI-derived signals against independent environmental outcomes such as verified emissions, regulatory enforcement actions, or third-party audit results.

4.2.3. Tension 3: The Environmental Footprint of AI

The third tension concerns the environmental footprint of AI itself. Green-finance AI may support better allocation of capital, but model training, inference, data storage, and cloud infrastructure all consume energy. This creates a potential contradiction: AI may support sustainability while also increasing computational emissions. Green AI provides a way to address this problem by measuring and reducing the environmental cost of model development and deployment (Elbounkify et al., 2026; Schwartz et al., 2020; Strubell et al., 2019). Green-finance research should treat model footprint as part of model evaluation, not as a separate operational issue. Practical levers include energy-aware training schedules, model distillation, hardware selection, and selective use of large language models for high-value tasks.

4.3. Implications for Theory

For finance theory, the review shows that AI-enabled green finance is not only a matter of better prediction. It raises questions about how climate and sustainability information enters prices, lending decisions, and portfolio construction. The carbon-risk literature suggests that environmental exposures affect investor demand and expected returns (Bolton & Kacperczyk, 2021, 2023; Pástor et al., 2021), and the Sautner et al. (2023) firm-level climate-exposure measure shows that text-based signals are already priced in options and equity markets. AI changes this process by altering how environmental exposures are measured and ranked. Future finance research should examine whether AI-generated sustainability signals affect price formation, cost of capital, and market discipline — and whether they do so in ways that converge with or diverge from human-curated signals.

For accounting and disclosure research, the review highlights the importance of textual sustainability information. NLP tools can scale the analysis of climate disclosures, but disclosure quality must be evaluated against substantive action. Future accounting research can connect AI-based disclosure measures to assurance, audit risk, greenwashing detection, and regulatory enforcement (Bingler et al., 2022; Huang et al., 2023). For information-systems research, green finance is a rich setting for studying algorithmic accountability because model outputs are embedded in contested institutional environments. For sustainability and management theory, the review suggests that AI is not a neutral tool but a constitutive element of how sustainability is defined, measured, and enforced.

4.4. Implications for Practice and Policy

For financial institutions, the review suggests that model selection should depend on the decision task. Low-stakes screening may tolerate more complex models if they improve coverage and speed. High-stakes credit, rating, or investment decisions require stronger explainability, documentation, and human oversight. Institutions should also invest in data governance before investing in more complex models. In many green-finance settings, the binding constraint is not model sophistication but inconsistent labels, weak disclosure, and uncertain environmental measurement.

For policymakers and regulators, the review suggests that AI regulation and sustainable-finance regulation should not be treated as separate domains. Algorithmic ESG scoring, AI-supported climate-risk assessment, greenwashing detection, and sustainability-disclosure monitoring sit at the intersection of both domains. Regulatory guidance should address model transparency, data provenance, auditability, and environmental footprint. The European Union's Sustainable Finance Disclosure Regulation (SFDR), the International Sustainability Standards Board (ISSB) standards, and the EU AI Act create complementary requirements that, taken together, point toward an integrated framework for accountable, sustainable, and energy-aware AI in finance. This would help prevent AI from amplifying greenwashing while still allowing useful analytical methods to develop.

5. Conclusion

AI is becoming part of the informational infrastructure of green finance. It supports ESG scoring, green-credit assessment, climate-aware portfolio construction, disclosure analysis, green-bond and carbon-market monitoring, and algorithmic governance. The 107 studies synthesized in this PRISMA-2020-compliant SLR show clear progress, particularly in prediction and text analysis, and clear unresolved tensions around interpretability, data quality, and the environmental footprint of AI itself.

The central argument of this review is that AI-enabled green finance should be evaluated as a decision system, not only as a set of models. The relevant question is not whether AI improves prediction in the abstract. The relevant question is whether AI improves green-finance decisions in ways that are accurate, interpretable, fair, environmentally credible, and institutionally legitimate. The integrative framework developed in this paper, the six thematic decision-task clusters, the three structuring tensions, and the six-item research agenda together provide the conceptual scaffolding for that evaluation. They

also provide an actionable map for scholars in finance, accounting, information systems, sustainability, and policy who are now confronting questions that no single discipline can answer alone.

Author Contributions

Both authors contributed equally to the conceptualization, screening, coding, and writing of this manuscript. The authors declare no competing financial interests or personal relationships that could have influenced the work reported. No external funding was received for this research. The data-extraction workbook and screening log used in this review are available from the corresponding author on reasonable request.

Conflict of Interest

The authors, while observing publication ethics in the referencing, declare the absence of interest of conflict.

Acknowledgement

We sincerely thank the reviewers for their valuable and constructive comments, which helped us improve the quality of this manuscript.

Data availability

All supporting data and supplementary information for this article are available in the Supplementary Material section.

AI Used Statment

The authors retain full responsibility for the accuracy, originality, integrity, legality, and ethical compliance of all text, data, tables, figures, analyses, references, and conclusions, and have verified all AI-assisted outputs. Any fabrication, falsification, manipulation, plagiarism, copyright infringement, or other misuse of AI remains solely the responsibility of the authors and does not transfer or diminish their accountability for the published work.

References

- Alonso-Robisco, A., Bas, J., Carbo, J. M., de Juan, A., & Marqués, J. M. (2025). Where and how machine learning plays a role in climate finance research. *Journal of Sustainable Finance & Investment*, 15(2), 456–497. <https://doi.org/10.1080/20430795.2024.2370325>.
- Berg, F., Kölbel, J. F., & Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6), 1315–1344. <https://doi.org/10.1093/rof/rfac033>.
- Beunza, D., & Stark, D. (2004). Tools of the trade: The socio-technology of arbitrage in a Wall Street trading room. *Industrial and Corporate Change*, 13(2), 369–400. <https://doi.org/10.1093/icc/dth015>.
- Bingler, J. A., Kraus, M., Leipold, M., & Webersinke, N. (2022). Cheap talk and cherry-picking: What ClimateBERT has to say on corporate climate risk disclosures. *Finance Research Letters*, 47, 102776. <https://doi.org/10.1016/j.frl.2022.102776>.
- Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2), 517–549. <https://doi.org/10.1016/j.jfineco.2021.05.008>.
- Bolton, P., & Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *Journal of Finance*, 78(6), 3677–3754. <https://doi.org/10.1111/jofi.13272>.
- 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>.
- Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys*, 55(3), Article 64. <https://doi.org/10.1145/3502289>.
- Carlei, V., Furi, D., Ceccarelli, A., & Cascioli, P. (2025). Outperforming ESG stocks portfolio: A machine learning ranking model with CatBoost regressor. *The North American Journal of Economics and Finance*, 80, 102517. <https://doi.org/10.1016/j.najef.2025.102517>.
- Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. *Artificial Intelligence Review*, 57, Article 216. <https://doi.org/10.1007/s10462-024-10854-8>.
- D'Amato, V., D'Ecclesia, R., & Levantesi, S. (2022). ESG score prediction through random forest algorithm. *Computational Management Science*, 19(2), 347–373. <https://doi.org/10.1007/s10287-021-00419-3>.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics* (pp. 4171–4186). Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-1423>.
- di Pietro, F., Giráldez-Puig, P., Palos-Sánchez, P. R., & Korzeb, Z. (2026). Artificial intelligence in climate and sustainable finance: A blessing or a curse? *Journal of Economic Surveys*. Advance online publication. <https://doi.org/10.1111/joes.70075>.
- Elbounkify, I., Machado, M. R., & Iannario, M. (2026). Designing green artificial intelligence (Green AI) models for finance: A novel approach for

- sustainable and responsible adoption. *Financial Innovation*, 12, Article 96. <https://doi.org/10.1186/s40854-026-00915-y>.
- Flammer, C. (2021). Corporate green bonds. *Journal of Financial Economics*, 142(2), 499–516. <https://doi.org/10.1016/j.jfineco.2021.01.010>.
- Giudici, P., & Wu, L. (2025). Sustainable artificial intelligence in finance: Impact of ESG factors. *Frontiers in Artificial Intelligence*, 8, Article 1566197. <https://doi.org/10.3389/frai.2025.1566197>.
- Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, 32, 100577. <https://doi.org/10.1016/j.jbef.2021.100577>.
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hha009>.
- Henrique, B. M., Sobreiro, V. A., & Kimura, H. (2019). Literature review: Machine learning techniques applied to financial market prediction. *Expert Systems with Applications*, 124, 226–251. <https://doi.org/10.1016/j.eswa.2019.01.012>.
- Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2), 806–841. <https://doi.org/10.1111/1911-3846.12832>.
- Huynh, T. L. D., Hille, E., & Nasir, M. A. (2020). Diversification in the age of the 4th industrial revolution: The role of artificial intelligence, green bonds and cryptocurrencies. *Technological Forecasting and Social Change*, 159, 120188. <https://doi.org/10.1016/j.techfore.2020.120188>.
- Ilhan, E., Krueger, P., Sautner, Z., & Starks, L. T. (2023). Climate risk disclosure and institutional investors. *Review of Financial Studies*, 36(7), 2617–2650. <https://doi.org/10.1093/rfs/hhad002>.
- Krappel, T., Bogun, A., Borth, D., Schimmelpfennig, T., & Stiehl, S. (2021). Heterogeneous ensemble for ESG ratings prediction. Proceedings of the KDD Workshop on Machine Learning in Finance. *arXiv preprint arXiv:2109.10085*.
- Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *Review of Financial Studies*, 33(3), 1067–1111. <https://doi.org/10.1093/rfs/hhz137>.
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174. <https://doi.org/10.2307/2529310>.
- Liaras, E., Nerantzidis, M., & Alexandridis, A. (2024). Machine learning in accounting and finance research: A literature review. *Review of Quantitative Finance and Accounting*, 63(4), 1431–1471. <https://doi.org/10.1007/s11156-024-01306-z>.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, 30, 4765–4774. Curran Associates.
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: A comparative analysis. *Scientometrics*, 106(1), 213–228. <https://doi.org/10.1007/s11192-015-1765-5>.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372(71). <https://doi.org/10.1136/bmj.n71>.
- Pástor, L., Stambaugh, R. F., & Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2), 550–571. <https://doi.org/10.1016/j.jfineco.2020.12.011>.
- Paul, J., Lim, W. M., O'Cass, A., Hao, A. W., & Bresciani, S. (2021). Scientific procedures and rationales for systematic literature reviews (SPAR-4-SLR). *International Journal of Consumer Studies*, 45(4), O1–O16. <https://doi.org/10.1111/ijcs.12695>.
- Pedersen, L. H., Fitzgibbons, S., & Pomorski, L. (2021). Responsible investing: The ESG-efficient frontier. *Journal of Financial Economics*, 142(2), 572–597. <https://doi.org/10.1016/j.jfineco.2020.11.001>.
- Pollock, N., & Williams, R. (2010). The business of expectations: How promissory organizations shape technology and innovation. *Social Studies of Science*, 40(4), 525–548. <https://doi.org/10.1177/0306312710362275>.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). ACM. <https://doi.org/10.1145/2939672.2939778>.
- Sautner, Z., van Lent, L., Vilkov, G., & Zhang, R. (2023). Firm-level climate change exposure. *Journal of Finance*, 78(3), 1449–1498. <https://doi.org/10.1111/jofi.13219>.
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). *Green AI. Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>.
- Seow, R. Y. C. (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>.
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>.
- Starks, L. T. (2023). Presidential address: Sustainable finance and ESG issues — Value versus values. *Journal of Finance*, 78(4), 1837–1872. <https://doi.org/10.1111/jofi.13255>.
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650). Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1355>.
- Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8, Article 45. <https://doi.org/10.1186/1471-2288-8-45>.
- Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management*, 14(3), 207–222. <https://doi.org/10.1111/1467-8551.00375>.
- Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Fuso Nerini, F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, Article 233. <https://doi.org/10.1038/s41467-019-14108-y>.
- Webersinke, N., Kraus, M., Bingler, J. A., & Leippold, M. (2022). ClimateBERT: A pretrained language model for climate-related text. SSRN Working Paper No. 4229146. <https://doi.org/10.2139/ssrn.4229146>.
- Zhang, D., Zhang, Z., & Managi, S. (2019). A bibliometric analysis on green finance: Current status, development, and future directions. *Finance Research Letters*, 29, 425–430. <https://doi.org/10.1016/j.frl.2019.02.003>.
- Zhang, H., Hu, Z., & Ji, Q. (2024). Green bond and yield spread: Evidence from machine learning. *Energy Economics*, 132, 107473. <https://doi.org/10.1016/j.eneco.2024.107473>.