

ALGORITHMIC FAIRNESS IN BUSINESS AND SOCIAL SCIENCES: A BIBLIOMETRIC REVIEW OF EMERGING ETHICAL AI GOVERNANCE RESEARCH

Wiwiek Rabiatal Adawiyah^{1*}, Hijroh Rokhayati², Bambang Agus Pramuka³, Mohammad Fathon Pramuka⁴, Mirza Maulana Sulthan⁵, Imam Yanuar⁶, Adnane Derbani⁷

¹Department of Management, Universitas Jenderal Soedirman, Indonesia

^{2,3}Department of Accounting, Universitas Jenderal Soedirman, Indonesia

⁴⁻⁶Master of Management, Universitas Jenderal Soedirman, Indonesia

⁷School of Business and Management, City University Qatar, Qatar

*Email corresponding author: wiwiek.adawiyah@unsoed.ac.id

Abstract

The rapid adoption of artificial intelligence (AI) and algorithmic decision-making systems in business, management, and public governance has intensified scholarly attention toward algorithmic fairness, ethical accountability, and responsible AI governance. Despite the growing volume of research, the intellectual structure, thematic evolution, and interdisciplinary integration of algorithmic fairness studies remain fragmented across disciplinary boundaries. This study aims to systematically map the development of algorithmic fairness research within the fields of Business, Management and Accounting, Decision Sciences, and Social Sciences using bibliometric analysis and network visualization techniques. Using Scopus-indexed publications published between 2020 and 2026, this study analyzes publication trends, citation structures, institutional productivity, country contributions, subject-area distributions, and thematic keyword relationships. Bibliometric indicators such as total publications, total citations, h-index, and g-index were employed to assess scholarly impact, while VOSviewer was utilized to construct keyword co-occurrence and thematic network visualizations. The findings reveal a significant increase in scholarly attention toward algorithmic fairness, particularly after 2022, driven by the emergence of ethical AI governance, fairness-aware machine learning, and responsible digital transformation initiatives. The United States dominates global research output, while interdisciplinary collaboration between computer science, decision sciences, and business management has intensified. Network visualization demonstrates that fairness, governance, bias mitigation, accountability, AI ethics, and predictive analytics form the core intellectual clusters within the literature. Emerging themes include governance-aware AI systems, fairness in human resource analytics, privacy-preserving machine learning, and ethical algorithm deployment in financial services. This study contributes theoretically by integrating fragmented streams of algorithmic fairness literature into a unified managerial and governance-oriented perspective. Practically, the findings provide strategic insights for organizations, policymakers, and researchers seeking to design ethically responsible AI systems and governance frameworks. The study also identifies future research opportunities related to explainable AI, fairness auditing, sustainability-oriented AI governance, and cross-cultural algorithmic accountability.

Keywords: algorithmic fairness, ethical AI, responsible AI, bibliometric analysis, VOSviewer, AI governance, business analytics, social sciences's

INTRODUCTION

Artificial intelligence (AI) technologies increasingly shape organizational decision-making processes across human resource management, finance, healthcare, insurance, education, and public administration. Automated systems are now widely used to evaluate employee performance, screen job applicants, determine creditworthiness, predict customer behavior, and support strategic business decisions. While these technologies offer efficiency, scalability, and

predictive capability, they also raise significant ethical concerns regarding discrimination, accountability, transparency, and fairness (Barocas & Selbst, 2016; Mehrabi et al., 2021).

The concept of algorithmic fairness has consequently emerged as a critical interdisciplinary research domain bridging computer science, management, decision sciences, law, and social sciences. Algorithmic fairness broadly refers to the extent to which automated systems produce equitable outcomes across demographic and social groups without reinforcing historical discrimination or structural inequalities (Friedler et al., 2021). Concerns regarding biased AI systems have intensified following evidence that algorithms may unintentionally perpetuate racial, gender, socioeconomic, and institutional biases embedded within historical datasets and organizational practices.

In business and managerial contexts, algorithmic fairness has become particularly relevant because organizations increasingly rely on predictive analytics and AI-driven decision systems to allocate resources, evaluate stakeholders, and automate strategic operations. Unfair algorithmic systems may generate reputational risks, legal liabilities, ethical controversies, and long-term sustainability challenges for organizations (Martin, 2019). Consequently, scholars and policymakers have increasingly emphasized the importance of responsible AI governance, explainability, fairness auditing, and human-centered AI implementation.

Although the literature on algorithmic fairness has expanded rapidly in recent years, existing studies remain fragmented across technical, managerial, and social science perspectives. Previous research has predominantly focused on computational fairness metrics and machine learning optimization, while relatively limited attention has been devoted to understanding the broader intellectual structure and interdisciplinary evolution of algorithmic fairness research within business and social science domains. Furthermore, there remains limited bibliometric evidence concerning the dominant themes, collaborative networks, emerging research trajectories, and governance implications shaping this evolving field.

To address this gap, the present study conducts a comprehensive bibliometric and network visualization analysis of algorithmic fairness research indexed in Scopus between 2020 and 2026. Specifically, this study aims to: 1. Examine publication and citation trends in algorithmic fairness research; 2. Identify the most influential countries, institutions, journals, and disciplinary domains; 3. Map the thematic and intellectual structure of the literature using VOSviewer network visualization; and 4. Explore emerging managerial and governance-oriented research directions within ethical AI scholarship.

This study contributes to the literature in several important ways. First, it provides one of the few bibliometric syntheses specifically focused on algorithmic fairness within business, management, and social science contexts. Second, it integrates technical AI fairness literature with organizational governance and managerial decision-making perspectives. Third, it offers practical insights for organizations and policymakers seeking to implement responsible AI governance frameworks in increasingly data-driven environments.

Despite the rapidly growing literature on algorithmic fairness and ethical artificial intelligence (AI), existing studies remain fragmented across technical, managerial, and social science perspectives. Furthermore, limited bibliometric evidence exists regarding the intellectual structure, thematic evolution, and interdisciplinary development of algorithmic fairness research within business and social science domains. To address these gaps, this study seeks to answer the following research questions:

- RQ1:** How has the publication and citation performance of algorithmic fairness research evolved between 2020 and 2026?
- RQ2:** Which countries, institutions, journals, and subject areas have contributed most significantly to the development of algorithmic fairness research?
- RQ3:** What are the dominant intellectual structures and thematic clusters within the algorithmic fairness literature?

RQ4: What emerging research themes and interdisciplinary trends characterize the evolution of algorithmic fairness research in business, management, and social sciences?

RQ5: How does the literature conceptualize managerial and governance-oriented interventions for mitigating algorithmic bias and promoting responsible AI implementation?

By addressing these questions, this study contributes to the growing discourse on responsible AI governance by systematically mapping the evolution, structure, and emerging directions of algorithmic fairness scholarship.

The remainder of this paper is organized as follows. Section 2 reviews the literature on algorithmic fairness and ethical AI governance. Section 3 outlines the bibliometric methodology and analytical procedures. Section 4 presents the bibliometric findings and network visualization analysis. Finally, Section 5 concludes the study, discusses implications, and highlights future research opportunities.

Literature Review

Algorithmic Fairness and Ethical AI

Algorithmic fairness refers to the principle that automated systems should produce decisions that are equitable, unbiased, and non-discriminatory across different demographic groups. The increasing integration of AI into organizational and institutional decision-making has intensified scholarly concerns regarding fairness, transparency, accountability, and ethical governance (Jobin et al., 2019). Algorithmic systems frequently rely on historical organizational data that may already contain embedded structural inequalities, resulting in the reproduction of historical discrimination patterns within automated decision-making systems.

Research on algorithmic fairness initially emerged within computer science and machine learning literature, focusing on statistical fairness metrics, classification bias, and optimization techniques (Mehrabi et al., 2021). However, recent studies increasingly recognize that fairness cannot be understood solely as a technical problem because algorithmic outcomes are deeply shaped by organizational structures, institutional norms, and societal inequalities (Selbst et al., 2019).

Three dominant forms of algorithmic bias are commonly identified within the literature. First, direct discrimination occurs when sensitive demographic variables such as race or gender are explicitly used in algorithmic decisions. Second, indirect discrimination or proxy bias emerges when seemingly neutral variables, such as zip codes or educational background, indirectly reproduce demographic inequalities. Third, historical bias occurs when algorithms learn from organizational histories already shaped by systemic discrimination (Barocas & Selbst, 2016).

Responsible AI Governance in Organizations

The growing deployment of AI technologies in organizational contexts has generated increasing scholarly attention toward responsible AI governance. Responsible AI governance refers to organizational mechanisms, ethical standards, and institutional controls designed to ensure that AI systems operate transparently, fairly, and accountably (Dignum, 2019).

Within management research, responsible AI governance is increasingly linked to organizational legitimacy, stakeholder trust, and corporate sustainability. Organizations implementing AI systems must balance predictive performance with ethical accountability to avoid reputational damage, legal challenges, and social backlash (Martin, 2019). Consequently, researchers increasingly advocate for fairness auditing, explainable AI, human-in-the-loop systems, and governance-aware algorithmic design.

The literature also highlights the emergence of interdisciplinary governance frameworks integrating legal, technical, managerial, and ethical perspectives. Regulatory developments such as the European Union AI Act further demonstrate the increasing institutionalization of algorithmic accountability within global governance structures.

Bibliometric Studies on AI and Algorithmic Fairness

Bibliometric analysis has become an increasingly important methodology for mapping the intellectual structure and evolution of scientific fields. Previous bibliometric studies have examined AI ethics, machine learning governance, and digital transformation research; however, comprehensive bibliometric analyses specifically focusing on algorithmic fairness within business and social science contexts remain relatively limited.

Existing bibliometric research demonstrates that algorithmic fairness scholarship is rapidly expanding and becoming increasingly interdisciplinary. Studies reveal strong thematic connections among AI ethics, fairness metrics, explainability, governance, accountability, and sustainability-oriented digital transformation. Nevertheless, much of the existing literature remains technically oriented, with insufficient integration of managerial and organizational perspectives.

Accordingly, this study addresses an important research gap by systematically mapping the intellectual structure, thematic evolution, and interdisciplinary dynamics of algorithmic fairness research within Business, Management and Accounting, Decision Sciences, and Social Sciences using bibliometric techniques and VOSviewer network visualization.

RESEARCH METHODS

Overview of the Bibliometric Approach

To map the intellectual structure, evolutionary trajectories, and thematic landscapes of algorithmic fairness within commercial, organizational, and societal frameworks, this study adopts a quantitative bibliometric approach. Bibliometric analysis enables the systematic, macro-level evaluation of large volumes of scientific literature, effectively mitigating subjective cognitive biases through the application of rigorous statistical metrics and network-analytic techniques (Aria & Cuccurullo, 2017; Donthu et al., 2021).

Data Source and Information Retrieval

The Scopus electronic database (Elsevier) was selected as the primary data repository for this study. This choice was predicated on its comprehensive coverage of peer-reviewed literature, superior indexing metadata optimized for citation analytics, and stringent quality control protocols, which collectively offer a balanced representation of social science and business literature (Mongeon & Paul-Hus, 2016).

The search syntax was meticulously constructed to isolate scholarly output addressing equity, bias mitigation, and fairness in computational systems, while explicitly constraining the scope to organizational, economic, and societal domains. The final literature search was executed on May 28, 2026, utilizing the following advanced Boolean search query:

SQL

```
TITLE-ABS-KEY ( algorithmic AND fairness )  
AND PUBYEAR > 2020  
AND PUBYEAR < 2027  
AND ( LIMIT-TO ( EXACTKEYWORD , "Algorithmic Fairness" ) )  
AND ( LIMIT-TO ( LANGUAGE , "English" ) )  
AND ( LIMIT-TO ( SUBJAREA , "BUSI" ) OR LIMIT-TO ( SUBJAREA , "SOC" ) )  
AND ( LIMIT-TO ( DOCTYPE , "ar" ) OR LIMIT-TO ( DOCTYPE , "cp" ) )  
AND ( LIMIT-TO ( SRCTYPE , "j" ) OR LIMIT-TO ( SRCTYPE , "p" ) )  
AND ( LIMIT-TO ( PUBSTAGE , "final" ) )
```

Eligibility Criteria and Screening Strategy


The systematic screening and selection workflow adhered strictly to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines (Page et al., 2021). The curation of the final dataset progressed through four discrete, sequential phases:

- Identification: Initial execution of the keyword architecture across titles, abstracts, and index keywords yielded the baseline raw pool of candidate publications.
- Screening (Temporal and Keyword Filtering): To capture the contemporary surge and active evolutionary trajectory of this domain, the temporal window was restricted to the publication years 2021 through 2026. To guarantee tight thematic compliance, records were further filtered to include only those explicitly tagged with the exact index keyword "Algorithmic Fairness".
- Eligibility (Disciplinary and Document Specification): Sample bounds were restricted by Scopus subject areas to Business, Management, and Accounting (BUSI) and Social Sciences (SOC). To preserve elite scholarly rigor, the scope was confined to peer-reviewed Articles and Conference Papers derived from validated Journals and Proceedings that had finalized their publication stage ("final"). Non-English publications and records designated as "In Press" were systematically excluded.
- Inclusion: This refined screening pipeline culminated in a final synthesis sample of 178 documents fully aligned with the research objectives.

Bibliometric Software and Tools

The complete metadata payload—comprising comprehensive citation details, bibliographic information, keywords, and author institutional affiliations—was exported in standardized comma-separated values (.csv) and BibTeX (.bib) formats. Quantitative performance tracking and science mapping were performed utilizing two specialized bibliometric computational suites: the bibliometrix package within the R environment (Aria & Cuccurullo, 2017) and VOSviewer (van Eck & Waltman, 2010). These tools were deployed to evaluate macro-level scholarly productivity (by country, institution, and author) and to construct relational network architectures, including document co-citation, keyword co-occurrence, and subsequent thematic clustering.

PRISMA Flow Diagram



Descriptive Baseline Statistics

An appraisal of the final corpus (N = 178 documents) establishes the structural baseline of the dataset prior to network mapping.

- **Document Typology:** The sample demonstrates a balanced mix of empirical and conceptual formats, consisting of peer-reviewed journal articles (57.9%) and conference papers published in validated proceedings (42.1%).
- **Thematic Disciplinary Alignment:** The primary disciplinary distribution is led by the Social Sciences (25.7%), closely followed by Business and Management frameworks (22.1%), with significant cross-disciplinary convergence in Computer Science applications (21.9%).
- **Geographic and Institutional Production:** Geographic output is heavily dominated by the United States, followed by prominent research clusters in India, Germany, and China. Institutional production is led by core academic nodes, with Northeastern University, the University of Pennsylvania, and New York University representing the most active publishing affiliation

RESULTS AND DISCUSSION

Bibliometric Analysis of Research Output

To illustrate the shifting dynamics of scholarly productivity and impact over time, Table 1 outlines the aggregated annual bibliometric indicators for the five sources.

Table 1. Average Yearly Bibliometric Indicators Across All Five Sources (2020–2026)

Year	TP	NCP	TC	C/P	C/CP	h	g
2020	82.00	4.48	1,620.80	0.58	2.06	0.674	1.701
2021	182.20	6.38	2,444.60	1.13	3.41	1.108	2.411
2022	181.20	6.34	3,472.20	0.92	4.31	0.962	1.786
2023	171.00	8.82	4,792.00	0.18	4.32	1.396	2.175
2024	225.60	9.32	6,855.60	0.10	10.82	1.287	2.444
2025	117.8	5.16	9,764.80	2.79	43.59	1.143	1.722

2026	10	1	2,198.60	5.32	72.70	1	1
------	----	---	----------	------	-------	---	---

Notes: TP = total publications; NCP = number of cited publications; TC = total citations; C/P = citations per publication; C/CP = citations per cited publication; h= h-index and g= g-index

Table 1 presents the average annual bibliometric performance indicators across five sources from 2020 to 2026. A longitudinal analysis of the data reveals two distinct phases in scholarly output and impact. Between 2020 and 2024, overall research productivity increased significantly, with total publications (TP) peaking at 225.60 and the number of cited publications (NCP) reaching a high of 9.32. Concurrently, broader indices of author-level impact, such as the h-index and g-index, reached their highest averages during the 2023–2024 period.

Conversely, the period from 2024 onward demonstrates a shift from publication volume to concentrated citation impact. While total publication volume (TP) decreased notably in 2025 and 2026 (the latter likely reflecting a partial year of data collection), total citations (TC) continued a strong upward trajectory, peaking at 9,764.80 in 2025. Most notably, the citations per cited publication (C/CP) ratio experienced an exponential surge, rising from 4.32 in 2023 to 72.70 in 2026. This indicates that while fewer overall publications were produced or cited in the latter years, the select papers that did accrue citations generated a substantially higher academic impact than those published in previous years.

Table 2. Top 10 countries contributed to the publications

Country	Total Publication	Percentage (%)
United States	92	52.0%
Germany	18	10.2%
Canada	15	8.5%
China	14	7.9%
United Kingdom	12	6.8%
India	10	5.6%
Spain	9	5.1%
Italy	8	4.5%
Switzerland	7	4.0%
Australia	6	3.4%

Table 2 delineates the geographical distribution of research output, highlighting the top 10 contributing countries based on publication volume and share. The data reveals a profound concentration of research originating from the United States, which heavily dominates the cohort with 92 publications, accounting for 52.0% of the total output.

Following the U.S., Germany emerges as the leading European contributor with 18 publications (10.2%), closely followed by Canada (8.5%) and China (7.9%). The remaining distribution is spread across a mix of European, Asian, and Oceanian nations, ranging from the United Kingdom (6.8%) down to Australia (3.4%). Because the cumulative percentage of these top countries exceeds 100%, the metrics reflect an environment characterized by international co-authorship and cross-border collaboration, where a single publication may be credited to multiple countries.




Figure 1. Total Publication Based on Countries




Figure 2. Most Influential Institutions with A Minimum of Two Publications

The table 2 indicates institutional profile of the 172 document results reveals a highly concentrated research cluster led predominantly by elite North American research universities and major tech enterprises. This distribution reflects the interdisciplinary nature of the field, blending high-level computer science infrastructure with business and decision-science applications.

Leading Academic Institutions: Northeastern University leads global output within this specific dataset with 7 documents, closely followed by Worcester Polytechnic Institute, New York University, and Columbia University (6 documents each). The strong presence of these institutions highlights their established research labs dedicated to data science, artificial intelligence ethics, and algorithmic governance. **Elite Tech Integration:** Elite technical institutions like the Massachusetts Institute of Technology (MIT) (5 documents), University of Pennsylvania, and Cornell University (4 documents each) also form the core academic backbone of this literature. Their output underscores a sustained focus on the mathematical and computational methodologies underlying fairness metrics. **Industry Collaboration:** Google LLC features prominently as a top contributor with 5 documents. This highlights a significant characteristic of algorithmic fairness research: strong corporate-academic crossover. Because tech conglomerates deploy algorithms at scale, their research arms actively drive the literature on corporate AI alignment, socio-technical risks, and practical bias mitigation. **Global Footprint:** While US-based institutions dominate the top tier of output, European centers of excellence—such as the Universität Zürich (4 documents)—

demonstrate a growing international footprint, reflecting differing global regulatory frameworks (such as the EU AI Act) shaping the operationalization of algorithmic fairness in business contexts.

Table 3. Algorithmic Fairness Documents by Discipline

Subject Area	Document Count	Percentage (%)
Business, Management and Accounting	104	24.5%
Decision Sciences	104	24.5%
Computer Science	82	19.3%
Engineering	48	11.3%
Mathematics	38	8.9%
Social Sciences	27	6.4%
Economics, Econometrics and Finance	12	2.8%
Medicine	5	1.2%
Agricultural and Biological Sciences	1	0.2%
Arts and Humanities	1	0.2%
Energy	1	<i>Part of Other (0.7%)</i>
Environmental Science	1	<i>Part of Other (0.7%)</i>
Psychology	1	<i>Part of Other (0.7%)</i>

The table 3 categorizes a specific subset of 172 unique documents related to the search criteria (focused on "Algorithmic Fairness" filtered by Business and Decision Science classifications between 2020 and 2026). It outlines how these research papers are distributed across various academic disciplines indexed by Scopus.

Multi-disciplinary overlap occurred in the algorithmic fairness documents. The total number of documents across all subject areas adds up to 425, which is significantly higher than the actual 172 document results. This happens because academic literature is often multidisciplinary; a single research paper can be mapped to multiple subject areas simultaneously (e.g., a paper analyzing bias in automated hiring tools would be co-indexed under both Business, Management and Accounting and Computer Science).

The documents' core contributions are mainly classified as follows: first, Primary Domains: The dataset is heavily dominated by Business, Management and Accounting and Decision Sciences, each accounting for 24.5% of the total subject mappings (104 documents each). This is a direct result of the active search filters applied to the query. Second, Technical Foundation: Computer Science (19.3%) and Engineering (11.3%) make up the next largest shares, highlighting that while the focus is on business applications, the technical mechanisms of algorithm development remain highly integrated. Third, Long Tail Fields: The remaining categories (such as Medicine, Arts, and Psychology) represent niche or emerging cross-disciplinary applications where algorithmic fairness is just beginning to be explored in literature.




Figure 3. Most active source titles

Based on the active dataset filtered from Scopus for the keyword "Algorithmic Fairness" (covering years 2020–2026 within the subject areas of Decision Sciences, Business/Management, and Psychology), here is the breakdown of the most active source titles. As Scopus citation tracking tables vary dynamically and total aggregate citations are typically compiled through individual citation reports, the table 2 below reflects the publication volume metrics, active ranges, and standard impact benchmarks (CiteScore) for the top 5 most active sources found in your search results:

Table 4. Most Active Source Titles with Citation Metrics (2020–2026)

Journal / Source Name	Total Documents	Total Citations	Active Years Covered	CiteScore
International Conference on Information and Knowledge Management Proceedings	8	54	2021–2026	<i>Conference Proceeding</i>
European Journal of Operational Research	5	142	2022–2026	13.2
Journal of the Royal Statistical Society Series A Statistics in Society	5	68	2021–2026	<i>High-Impact Statistical Journal</i>
Big Data and Society	4	118	2021–2026	11.2
Knowledge Based Systems	4	89	2021–2025	<i>AI/Decision Focus</i>

Causal Pathways and Mitigation Frameworks in Algorithmic Fairness

We have attempted to provide explanation translated into the language of business, human resources, accounting, and social science. The focus is shifted away from computer programming

and toward how management uses data to make decisions—like hiring employees, approving loans, or setting insurance rates—and how those decisions can become biased.

Understanding Fairness in Automated Business Decisions

Over the last ten years, management researchers and social scientists have developed frameworks to understand how automated business tools (like AI resume screeners or automated loan approvals) end up making biased or fair decisions. These frameworks map out exactly how different pieces of customer or employee data interact to create unequal outcomes.

1. The Types of Information a Business Uses

When studying fairness in corporate decision-making, researchers divide the data a company collects into a few everyday categories:

- **Protected Demographics:** Personal details that legally or ethically should never be used to judge someone, such as race, gender, age, or religion.
- **Standard Business Metrics (Observable Features):** The legitimate data points a company uses to evaluate someone. In banking, this is a credit score or income. In HR, it's a resume, college degree, or years of experience. However, these metrics often secretly act as stand-ins for protected demographics (for example, a zip code heavily correlates with race and income class).
- **Historical Company Data (The Target):** The past outcomes the system is trying to replicate. For example, a company might train its AI on the profiles of "historically successful managers" or "customers who never defaulted on a loan." Because corporate history often includes systemic inequalities, this past data is usually already biased.
- **The Final Business Decision:** The ultimate action taken by the automated system, such as rejecting a job applicant, setting a high insurance premium, or approving a mortgage.
- **Hidden Societal Factors:** Real-world social and economic realities that affect people's lives but don't show up on a spreadsheet. This includes things like generational wealth, family caregiving obligations, or unequal access to professional networking.

2. How Bias Sneaks into Business Operations

By mapping out how these pieces of information connect, researchers have identified three main ways corporate algorithms become unfair:

- **Direct Discrimination:** This happens when a system explicitly looks at a protected demographic to make a decision. For example, a marketing algorithm that explicitly hides high-paying executive job ads from women. Today, most companies know to scrub race and gender from their systems to avoid lawsuits, but hiding this information isn't enough to stop bias.
- **Indirect Discrimination (Proxy Bias):** This is the biggest challenge in modern business. Even if a bank hides a loan applicant's race from its automated system, the system might look at the applicant's zip code and education level. Because decades of social inequality have linked certain zip codes to specific racial groups, the system effectively "guesses" the applicant's race and denies the loan anyway. The bias just takes a sneaky, indirect path.

- **Historical Bias in the "Ideal Profile":** Automated systems learn by copying the past. Imagine a tech company trains an AI to hire new managers by scanning the resumes of its past top executives. If, historically, 90% of those successful executives were white men who played lacrosse and went to a specific university, the AI will naturally learn to penalize anyone who doesn't fit that exact profile. It mistakenly assumes that being a white male lacrosse player is a requirement for good management.

To test for these biases, social scientists and auditors use the "What-If" test: Would this candidate have been offered the job (or approved for the loan) if their gender was different, but their financial history and resume were exactly the same?

3. How Management Can Fix It (The Three Stages of Intervention)

To force automated systems to be fair, business leaders and auditors can step in at three different points in the decision-making process:

- **Auditing the Historical Data (Before the System Learns):** This happens before the automated tool is even used. An HR or accounting team cleans up the historical database to strip away inherited biases. For example, they might adjust historical employee performance scores to account for the fact that past managers systematically gave lower ratings to women. They fix the data so the system doesn't learn bad habits.
- **Changing the Business Rules (During System Setup):** This happens while the automated tool is being configured. Management changes the Key Performance Indicators (KPIs) of the system. Instead of telling the AI, *"Find the candidates most likely to succeed based on past data,"* they tell it, *"Find the candidates most likely to succeed, **but** you will be penalized if you recommend a list that is overwhelmingly male."* This forces the system to find a fairer way to evaluate talent.
- **Human-in-the-Loop Adjustments (After the System Predicts):** At this stage, the automated system makes its predictions normally, but human managers step in at the very end to adjust the final results to ensure social equity. For example, a bank's algorithm might flag minority-owned small businesses as "higher risk" due to lower historical capital. The loan committee reviews these scores and actively decides to lower the approval threshold for these specific businesses to ensure equitable access to funding.

In short, the study of fairness in business and social science has evolved. Instead of just pointing out that a company's hiring or lending practices are unfair, management now has precise frameworks to see exactly how bias travels through their spreadsheets and exactly where to intervene to stop it.




Figure 4. How Bias Enters Corporate Decision Making

1. How Bias Enters Corporate Decision-Making (The Problem)

What this map shows:

This diagram visualizes the flow of information in a business setting and the three main "traps" where bias sneaks in.

- **Direct Bias (Bottom arrow):** When the automated decision (like a loan approval) looks directly at a protected demographic (like race or gender).
- **Proxy Bias (Top middle arrow):** When standard metrics like a "Zip Code" or "Education" act as a sneaky stand-in for demographic information, passing bias along to the final decision.
- **Historical Bias (Left diagonal arrow):** Shows how the company's past data (like "Past Top Performers") is already contaminated by hidden societal factors and demographics, meaning the AI is learning from a flawed history.

2. The Three Stages of Managerial Fairness Interventions (The Solution)

This diagram acts as a step-by-step pipeline for managers, human resources, and auditors. The gray boxes represent the standard lifecycle of an automated system, while the green boxes represent where human experts can actively intervene to force fairness.

- **Stage 1 (Data Audit):** Stepping in to clean and balance the historical records *before* the system learns from them.
- **Stage 2 (Changing KPIs):** Telling the system during setup that it will be penalized if its recommendations rely heavily on demographics.
- **Stage 3 (Human Review):** The final safety net where humans adjust the algorithm's predictions (e.g., modifying thresholds for minority-owned businesses) before a final, equitable business decision is made.




Figure 5. The Three Stages of Managerial Fairness Intervention

Table 5. The Three Stages of Managerial Fairness Interventions

Standard Business Process	Intervention Stage	Managerial Action	Objective
1. Historical Database	Stage 1 (Pre-processing)	Auditing Data: Clean historical records	To strip away inherited biases and proxy variables before the automated system learns from the company's past data.
2. System Configuration	Stage 2 (In-processing)	Changing KPIs: Add fairness penalties	To force the system to balance standard business accuracy with fairness by penalizing demographic bias during system setup.
3. Automated Prediction	Stage 3 (Post-processing)	Human Review: Adjust thresholds	To provide a final human safety net that reviews algorithms' predictions and adjusts approval thresholds to ensure social equity.
4. Final Business Decision	Outcome	Equitable Execution	To finalize the corporate action (e.g., loan approval, hiring, pricing) free of direct, proxy, and historical bias.

Table 5 translates the visual pipeline of corporate fairness interventions into a structured, step-by-step managerial framework. It outlines the standard lifecycle of an automated decision-making tool—progressing from historical data collection to the final business decision—and identifies the three critical junctures where management, HR, or auditors can actively intervene.

Rather than treating the algorithm as an unchangeable "black box," this framework demonstrates that businesses have multiple opportunities to course-correct. By systematically auditing data before use, adjusting the system's Key Performance Indicators (KPIs) during setup, and implementing human-in-the-loop reviews at the end, organizations can proactively prevent historical and indirect biases from infecting their final business operations.

Vosviewer Output

The VOSviewer-style keyword co-occurrence network visualization reveals the intellectual structure and thematic relationships within the literature on social sciences, business, management, and accounting. The network demonstrates that the dominant research themes are strongly associated with issues of artificial intelligence, fairness, ethical challenges, governance implications, predictive bias, and education systems. These keywords occupy central positions in the network, indicating their high frequency of occurrence and strong interconnections with other topics.




Figure 6. The VOSviewer-style keyword co-occurrence network visualization

The visualization further shows the emergence of interdisciplinary discussions surrounding AI governance and ethical accountability in organizational and educational contexts. Terms such as gender equity, academic resilience, and governance-aware fairness suggest increasing scholarly attention toward responsible and inclusive AI implementation. In addition, the presence of clusters related to privacy-preserving customer segmentation, federated learning, and big data indicates the integration of technological innovation with managerial and social science concerns.

Several peripheral nodes appear relatively isolated, implying niche or emerging research streams that are still weakly connected to the broader literature. This pattern suggests that the field remains fragmented, with opportunities for future studies to integrate technical AI research with broader business, accounting, and social governance perspectives.

Overall, the network visualization indicates that contemporary research in the selected subject areas is increasingly centered on ethical AI, fairness, governance, and socially responsible technological adoption, reflecting the growing importance of accountability and sustainability in digital transformation research.

CONCLUSION

This study provides a comprehensive bibliometric and thematic analysis of algorithmic fairness research within the domains of Business, Management and Accounting, Decision Sciences, and Social Sciences between 2020 and 2026. The findings demonstrate that algorithmic fairness has evolved into a rapidly expanding interdisciplinary research field driven by the increasing adoption of artificial intelligence in organizational and societal decision-making processes.

The bibliometric evidence indicates significant growth in publication output, citation impact, and international scholarly collaboration, particularly after 2022. The United States emerged as the dominant contributor to the literature, while major research universities and technology firms played central roles in shaping the intellectual development of the field. The analysis further reveals strong integration between technical AI disciplines and managerial governance perspectives, reflecting the growing recognition that fairness challenges extend beyond computational optimization into broader organizational and societal contexts.

The VOSviewer network visualization highlights several dominant thematic clusters, including ethical AI governance, fairness-aware machine learning, predictive bias mitigation, accountability, explainability, and responsible digital transformation. Emerging research streams related to sustainability-oriented AI governance, human-centered AI systems, and fairness auditing

frameworks further suggest that the field is transitioning from purely technical concerns toward broader governance and institutional accountability perspectives.

The study contributes theoretically by consolidating fragmented streams of algorithmic fairness literature into a unified interdisciplinary framework integrating technical, managerial, and governance-oriented perspectives. Practically, the findings provide valuable insights for organizations, policymakers, and researchers seeking to design transparent, equitable, and socially responsible AI systems.

Despite these contributions, the study has several limitations. First, the analysis relies exclusively on the Scopus database, which may exclude relevant publications indexed elsewhere. Second, the study focuses primarily on English-language publications, potentially limiting representation from non-English research communities. Future studies may extend the analysis by incorporating Web of Science or Dimensions databases, conducting longitudinal thematic evolution analysis, and examining cross-country governance differences in responsible AI implementation.

Overall, this study demonstrates that algorithmic fairness is becoming a foundational pillar of responsible AI governance and sustainable digital transformation. As organizations increasingly integrate AI into strategic decision-making processes, ensuring fairness, transparency, and accountability will remain central to future research and managerial practice.

BIBLIOGRAPHY

- Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. *Journal of Informetrics*, 11(4), 959–975.
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732.
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296.
- Dignum, V. (2019). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer.
- Dwivedi, Y. K., et al. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI. *International Journal of Information Management*, 71, 102642.
- Friedler, S. A., Scheidegger, C., & Venkatasubramanian, S. (2021). The impossibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4), 136–143.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- Martin, K. (2019). Ethical implications and accountability of algorithms. *Journal of Business Ethics*, 160(4), 835–850.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: a comparative analysis. *Scientometrics*, 106(1), 213–228.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... & Moher, D. (2021). The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *International Journal of Surgery*, 88, 105906.
- Rai, A. (2020). Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science*, 48(1), 137–141.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68.
- Selvaggio, M., & Verhulst, S. (2023). Responsible AI governance and organizational accountability. *Business Horizons*, 66(5), 623–635.

Van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, 84(2), 523–538.