

Management control meets AI: From data-driven literature review to research gap*

Omar Feria Alonso

Universidade de Vigo - UVIGO, Spain

Elena Shakina

Universidade de Vigo - UVIGO, Spain

Maria Beatriz Gonzalez-Sanchez

Universidade de Vigo - UVIGO, Spain

Jose Berbel-Vera

Universitat de València - UV, Spain

El control de gestión se encuentra con la IA: De la revisión de literatura basada en datos a la identificación de vacíos de investigación

La integración de la inteligencia artificial (IA) en la investigación académica constituye un instrumento de alto impacto para la gestión del conocimiento científico. A partir de ello, este artículo presenta una revisión de la literatura basada en datos que explora la intersección entre los sistemas de control de gestión (SCG) y la IA, mapea los principales clústeres temáticos, e identifica vacíos de investigación. Sobre la base de un corpus curado de artículos revisados por pares, publicados entre 2010 y 2025, se identificaron cinco clústeres temáticos principales; además,

* We are grateful for the support of the following institutions: Spanish Ministry of Science and Innovation (PID2022-139974NB – I00), project title “Accounting and management control in small firms and startups: Implications in a context of high uncertainty”; Brazilian Ministry of Science, Technology and Innovation (Ref. MCTI/CNPq nº 14/2023) project title “Management control systems in small businesses and startups: Implications in a context of high uncertainty and their impacts on labour well-being, digital transformation and sustainability.”



se evaluó en qué medida cada uno aborda la transparencia y la explicabilidad, preocupaciones centrales en la implementación de IA en contextos de SCG. Los hallazgos revelan que solo dos clústeres abordan explícitamente la IA explicable (XAI), lo que evidencia una importante laguna en la literatura. Este estudio ofrece una doble contribución: proporciona un mapeo sistemático de las investigaciones actuales sobre sistemas de control habilitados por IA, y propone una agenda de investigación que destaca la necesidad de un enfoque más integrado y transparente hacia la explicabilidad en contextos de toma de decisiones impulsados por IA. Además, el estudio demuestra la capacidad de las técnicas basadas en datos para orientar futuras investigaciones, subrayando el papel indispensable de la lectura crítica y del juicio humano en la aplicación de métodos de IA a la investigación académica.

Palabras clave: revisión de la literatura, control de gestión, inteligencia artificial explicable

Management control meets AI: From data-driven literature review to research gap

The integration of artificial intelligence (AI) into academic research constitutes a highimpact instrument for managing scientific knowledge. Building on these capabilities, this paper presents a datadriven literature review that explores the intersection of management control systems (MCS) and AI, maps key thematic clusters, and identifies research gaps. Drawing on curated corpus of peer-reviewed articles published between 2010 and 2025, we identify five major thematic clusters and assess the extent to which each addresses transparency and explainability, core concerns in implementing AI within MCS contexts. Our findings reveal that only two clusters explicitly engage with explainable AI (XAI), revealing a significant research gap. This study offers a twofold contribution: it provides a systematic mapping of current research on AI-enabled control systems and proposes a research agenda that emphasizes the need for a more integrated and transparent approach to explainability in AI-driven decision-making contexts. The study further demonstrates the capacity of datadriven techniques to steer future inquiry, while simultaneously underscoring the indispensable role of critical reading and human judgment in the application of AI methods to scholarly research.

Keywords: literature review, management control, explainable artificial intelligence

O controle de gestão se encontra com a IA: Da revisão da literatura baseada em dados à identificação de lacunas de pesquisa

A integração da inteligência artificial (IA) na pesquisa acadêmica constitui uma ferramenta de alto impacto para a gestão do conhecimento científico. Com base nessas capacidades, este artigo apresenta uma revisão da literatura baseada em dados que explora a interseção entre sistemas de controle gerencial (SCG) e IA, mapeia os principais clusteres temáticos e identifica lacunas de pesquisa. Com base em um corpus selecionado de artigos revisados por pares, publicados entre

2010 e 2025, identificamos cinco clusters temáticos principais e avaliamos em que medida cada um aborda a transparência e a explicabilidade, preocupações centrais na implementação de IA em contextos de SCG. Nossos resultados revelam que apenas dois clusters abordam explicitamente a IA explicável (IAX), o que destaca uma importante lacuna significativa na literatura. Este estudo oferece uma dupla contribuição: fornece um mapeamento sistemático das pesquisas atuais sobre sistemas de controle habilitados por IA e propõe uma agenda de pesquisa que destaca a necessidade de uma abordagem mais integrada e transparente para a explicabilidade em contextos de tomada de decisão orientada por IA. Além disso, o estudo demonstra a capacidade das técnicas baseadas em dados para orientar futuras pesquisas, ao tempo em que destaca o papel indispensável da leitura crítica e do julgamento humano na aplicação de métodos de IA à pesquisa acadêmica.

Palavras-chave: revisão de literatura, controle gerencial, inteligência artificial explicável

1. INTRODUCTION

Artificial intelligence (AI) is transforming the way academic research is conducted, offering tools for processing and analysing large volumes of text quickly and accurately, exceeding the capabilities of traditional manual methods. One of the most relevant developments of this transformation is natural language processing (NLP), which refers to the ability of machines to understand, interpret, and generate human language. This makes NLP especially valuable for analysing large volumes of textual data such as academic publications (Teerasoponpong et al., 2025).

Manual literature reviews, by contrast, present notable limitations in both scope and objectivity. Traditional bibliographic review usually requires a considerable amount of time -sometimes months or even years- especially in broad or interdisciplinary fields, thereby restricting the volume of literature that can be rigorously analysed. Additionally, cognitive fatigue accumulated during the process may compromise the consistency and rigor of analytical judgments.

Despite these limitations, bibliographic reviews remain essential for gaining a comprehensive understanding of key topics and recent advances within a discipline, as well as for identifying significant gaps in existing body of knowledge. However, they also reflect a certain degree of subjectivity (Pugliese et al., 2023) as they rely heavily on researchers' exposure, experience, and critical assessment (Kraus et al., 2022). For early-career researchers in particular, it is crucial to develop a solid understanding

of the literature they engage with recognizing well-established findings, identifying underexplored areas and understanding ongoing debates.

To overcome these challenges, automated methods such as topic modelling using latent Dirichlet allocation (LDA) have emerged as valuable tools. LDA identifies latent topics in large collections of documents, making it essential for structuring automated literature review. In parallel, the growing capabilities of large-scale language models (LLM) offer new opportunities for generating thematic summaries complementing LDA-based structure analysis,

Building on these capabilities, this paper conducts an automated literature review at the intersection of AI and management control systems (MCS). The convergence of these two domains is rapidly evolving and emerging as a strategically area of organizational research. As AI technologies are increasingly embedded in decision-making and performance monitoring, scholars have begun to investigate the implications, challenges, and transformative potential.

This study presents a data-driven literature review that maps the emerging intellectual landscape of this convergence. It identifies thematic areas and evaluates how well current research addresses foundational AI concerns, particularly those related to the transparency and explainability, collectively termed explainable ai (XAI), which constitute primordial requirements from an MCS perspective.

To identify prevailing trends and research gaps, we apply a structured text mining implementing advanced NLP techniques on a curated corpus of peer-reviewed academic publications. Our analysis identifies five major thematic clusters that shape current discourse. Notably, concerns around XAI, crucial for mitigating algorithmic opacity and preventing hallucinations, are addressed explicitly in only two clusters, while the remaining three, including the widely discussed role of MCS in decision-making, largely neglect this dimension. This gap is particularly significant, as interpretability and trust are prerequisites for the responsible deployment of AI in organizations, especially when the technology aims to enhance, or even automate, decisionmaking processes.

Our findings suggest that while scholarly interest in the intersection of AI and MCS is growing rapidly, literature remains fragmented, and critical concerns such as XAI are unevenly addressed. This review not only synthesizes the state of the field but also highlights a pressing research agenda: the need for a more systematic integration of explainability into theoretical and empirical work on AI-enabled MCS.

2. THEORETICAL BACKGROUND

2.1. MCS: Overcoming black-box opacity of AI

MCS have been recognized as essential mechanisms through which managers guide organizational members to implement strategies and align actions with organizational goals (Simons, 1995). Over time, this traditional perspective has evolved toward a more dynamic and context-sensitive understanding of control. Recent studies emphasize that MCS are not static structures but flexible, adaptive configurations that enable organizations to respond to environmental uncertainty, foster learning, and encourage creativity and innovation (Bedford, 2015; Speklé et al., 2017).

Contemporary frameworks emphasise the multi-faceted nature of MCS: they combine formal mechanisms as budgets, performance reports, incentive contracts with informal elements grounded in culture and social norms (Anthony & Govindarajan 2007). These mechanisms may be implemented diagnostically to monitor variance or interactively to stimulate strategic debate (Simons, 1995; Tessier & Otley 2012), and they can fulfil both enabling and coercive functions depending on how managers deploy them (Ahrens & Chapman 2004). This plurality of purposes allows managers to balance efficiency, adaptability, and strategic responsiveness. Control instruments provide managers with quantitative and qualitative information that supports decision-making, mitigates risks, and uncovers opportunities (Chenhall & Moers, 2015).

In recent years, however, the digital transformation of organizations has profoundly reshaped how MCS are designed and used. AI is becoming increasingly embedded in managerial processes, automating tasks traditionally reliant on human judgment and enhancing the scope and depth of managerial analytics. This trend has opened an exciting field of inquiry, as researchers explore the implications of AI-driven decision-making (Dwivedi et al., 2021).

While AI holds significant promise for improving decision quality and organizational performance (Masialetti et al., 2024) its inherent “black-box” nature raises new challenges for accountability, transparency, and managerial understanding. As advanced machine learning models are deployed in high-stakes decisions, managers and stakeholders alike demand greater interpretability and control over AI-generated insights (Jarrahi, 2018).

In response, the field of explainable AI (XAI) has emerged as a critical approach to making AI models more transparent and accessible without substantially reducing their predictive power, a dilemma labelled as the “black-box” problem (Arrieta et al., 2020; Meske et al., 2022). While computer scientists devise explanation techniques, researchers grounded in management control emphasise how those explanations

must be integrated with existing accountability routines and decision processes inside organisations, highlighting that technical explanations are not sufficient. So, these explanations must be embedded within existing controls (Elbashir et al., 2021) as review meetings, variance-analysis protocols and ethical boundary systems if they are to sustain meaningful dialogue and informed judgement. When explanations are aligned with the task domain, experimental evidence shows that transparency can enhance user trust (Yu & Li, 2022); when mis-aligned, it risks information overload or misplaced reliance. XAI has therefore emerged as a bridge between algorithmic technical capability and managerial accountability.

In consequence, within the context of MCS, XAI is increasingly viewed as a necessary enabler for ensuring that AI not only improves decision-making processes but does so in a way that managers can comprehend its outputs and justify them. Without sufficient explainability, the adoption of AI-enabled control systems will continue to face critical scepticism, as opaque outputs may undermine trust and hinder integration into established accountability structures and governance routines.

The emerging intersection of MCS, AI, and XAI represents a promising yet underexplored avenue for both research and practice. Recent studies—many still in the form of systematic literature reviews (e.g., Abdel-Karim et al. 2021; Weber et al. 2023) highlight that this field remains in its formative stages. Going forward, it is essential to develop control systems that not only harness the analytical power of AI but also ensure that decision-makers retain sufficient understanding to exercise informed and responsible judgment. This evolving interplay between technological sophistication and managerial comprehension will likely shape the future design of management control, positioning XAI as a vital link to unlock the potential of AI in increasingly data-driven organizations.

3. METHODOLOGY

To achieve the stated objectives, a five-stage, data-driven methodological framework was developed, combining advanced computational techniques with qualitative validation strategies. The analysis was conducted using the R software environment, chosen for its user-friendly interface and its ability to run specialized text analysis packages. This tool facilitated the extraction of key information such as representative bigrams and enabled the clear and replicable visualization of results. R is particularly useful for researchers as it doesn't require advanced knowledge of NLP or complex programming skills.

The core of the approach is grounded in LDA (Blei et al., 2003), a probabilistic topic modelling technique widely used to reveal latent thematic structures within large corpora of unstructured text. These AI techniques have already been used in accounting (Murphy et al., 2024), but as far as we know, not specifically in our field of research, which focuses more on the internal accounting tools than on disclosure of accounting information. LDA was chosen due to its capacity to model the distribution of topics across documents and the distribution of terms across topics, making it particularly suitable for identifying clusters of research themes in a large, heterogeneous set of academic publications. Each document in the corpus is treated as a mixture of multiple topics, and each topic as a distribution over words. This allows for the discovery of underlying thematic patterns that may not be immediately observable through manual review. The model parameters were tuned to identify five major clusters, which were determined to best balance interpretability and thematic distinctiveness based on coherence metrics and expert judgment.

While LDA typically operates on unigrams (individual words), this can lead to overly generic or ambiguous topic descriptors. To address this, we implemented bigram extraction—the identification of frequently co-occurring word pairs within the corpus. Bigrams offer a higher degree of semantic precision, as they often correspond to domain-specific expressions (e.g., “project_management”, “stability_strategy”, or “management_practices”). By associating each LDA-generated topic with its most representative bigrams, we enhanced the clarity and contextual richness of the topic labels. The five-steps (see figure 1) methodological framework are as follows:

1) Corpus construction via Scopus

A targeted search was conducted in the Scopus database using refined Boolean expressions containing the keywords: “*management control systems*”, “*performance measurement systems*”, and “*artificial intelligence*”. The temporal scope was restricted to publications from 2010 to 2025. The search excluded 18 subject areas deemed peripheral to the research focus and was limited to peer-reviewed journal articles written in English. This filtering ensured disciplinary relevance to the domains of business management, accounting, economics, econometrics, and finance.

2) Metadata extraction and preprocessing

Bibliographic metadata -including authorship, publication year, title, abstract, keywords, and institutional affiliation- was extracted and organized in a structured spreadsheet format. This served as the foundation for subsequent text analysis. The analysis was conducted running R programming using a set of specialized pack-

ages for each stage of the work. Initially, libraries for data loading and manipulation (readxl, dplyr) were installed to allow the selection of relevant columns, filtering records, and structuring the information. For corpus preprocessing, tools such as tidytext, tm, and others were used, aimed at removing stop words and tokenization. In the following stages, libraries designed to create term matrices and run topic models such as LDA were applied, as well as others to generate visual representations (word clouds, charts, bigram visualization, and metrics), with the aim of facilitating the interpretation of the results and identifying thematic patterns.

Particular emphasis was placed on the content of titles, abstracts, and author-defined keywords, as these represent the most concise and intentional expressions of a paper's thematic focus. The full text of the articles was not included in the analysis, as the objective of the study is to identify research trends—information that can be effectively extracted from these three key elements.

Limiting the corpus to specific components such as abstracts or key words is due to both methodological and practical considerations. First, these elements are specifically designed to concisely synthesize the fundamental content of the research, including the research questions, methodologies employed, and main findings. Additionally, licensing restrictions constitute a significant barrier to complete access to full texts. Finally, this methodological approach facilitates the replicability and scalability of the analysis, allowing for its consistent application across diverse disciplines and data sets, resulting in a more robust and transferable framework.

3) Topic modelling via LDA

The cleaned corpus was analysed using LDA to uncover latent thematic structures within the data. Five dominant topics were identified as the most representative within the literature. To enhance interpretability, the analysis prioritized the extraction of bigrams - particularly those relevant in the context of MCS and AI-related terminology. Each cluster was described using its top 10 bigrams, with the top 20 used in calculating representativeness for statistical robustness. We tested multiple cluster numbers ($k = 2$ to 20) and assessed coherence using the *ldatuning* package, applying metrics implemented by Cao et al. (2009) and Deveaud et al. (2014).

These metrics are based on studies cited in the topic modelling literature. The index proposed by Cao et al. (2009) measures the similarity between topics based on their word distributions: the lower the value, the greater the thematic differentiation. For example, if two topics share many words like “money,” “bank,” and “credit,” the metric penalizes them. In contrast, the index proposed by Deveaud et al. (2014) evalu-

ates the informational separation between topics, selecting models where each topic contributes distinctive content (higher values indicate better performance). Thus, if one topic focuses on “financial education” and another on “artificial intelligence,” the metric considers it favourable due to its clear distinction. Based on the trade-off between coherence and thematic interpretability, $k = 5$ was selected as the optimal number of clusters.

4) Cluster labelling and thematic interpretation via LLMs

To enhance the interpretability of the topics generated by the LDA model, we employed a LLM to support the labelling and thematic refinement of each cluster. This step involved a two-phase interaction with the LLM, using carefully crafted prompts aligned with the goals of semantic clarity and analytical depth.

The model was tasked with proposing accurate and contextually appropriate labels for each topic based on their top 20 bigrams. The following prompt was used:

Prompt 1: “I am conducting topic modelling as part of a research project at the intersection of management control and artificial intelligence. Below, I present the output for one of the topics identified, which includes the top 20 most relevant bigrams, each accompanied by its beta value and relative weight. Based on this list, please propose the most thematically appropriate and contextually accurate name for the topic. The suggested name should concisely reflect the core theme revealed by the bigrams, capturing both the conceptual and terminological focus within the domain of management control and AI. It is essential that the proposed label takes into consideration all the bigrams provided, not just the most dominant ones.”

5) Assessment of explainability coverage

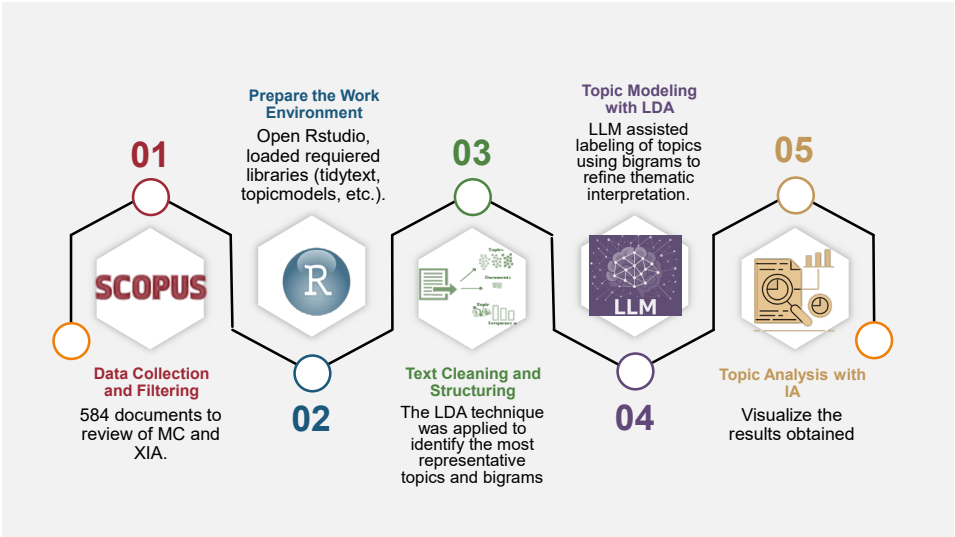
A second intervention using LLMs was performed to analyse whether—and how—each of the identified clusters explicitly or implicitly addressed the issue of AI explainability. This step was critical to assess whether transparency-related concerns constitute a discernible trend across the thematic clusters or remain marginal in the current discourse. The following prompt was employed:

Prompt 2: “As part of a topic modelling analysis on research at the intersection of management control and artificial intelligence, I am seeking to identify whether the issues on AI explainability and interpretability are present within the discovered topics. For each topic, please review the list of most relevant bigrams (along with their beta values and weights) and determine whether any of the terms - either directly or indirectly - relate to the concepts of transparency, explainability, or interpretability

of AI systems. This includes both explicit keywords (e.g., explainable AI, model transparency) and more implicit references (e.g., black box, decision rationale, user trust). Highlight any such connection.”

By integrating LLM-assisted interpretation into the workflow, we ensured a more coherent thematic labelling and a focused assessment of explainability-related dimensions within each topic, contributing to both the interpretive richness and the analytical rigor of the study.

Figure 1. Summary of the work



4. RESULTS

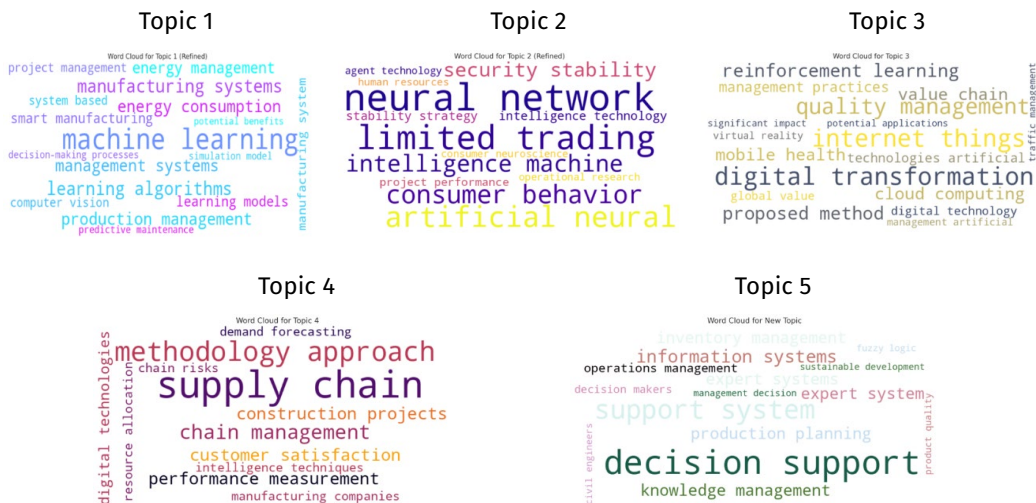
Figure 2 presents the most significant terms extracted from the entire document corpus using frequency-based analysis. These terms reflect the most recurrent and contextually significant lexical items found in the titles, abstracts, and keywords of peer-reviewed publications at the intersection of management control and AI. While this initial lexical mapping offers only a high-level overview of the discourse, it serves as a crucial starting point for identifying the foundational vocabulary around which scholarly debates are organized.

The value of this preliminary output lies not only in its ability to highlight core concepts, such as “artificial”, “performance”, “intelligence”, or “support”, but also in its role as a diagnostic tool for guiding more focused lines of inquiry. In particular, this overview raises key analytical questions that shape the direction of the study:

- To what extent are these core concepts thematically co-located in scholarly texts, and how do they cluster into identifiable domains of knowledge?
- Which of these frequently used terms reflect established areas of inquiry, and which may be indicative of emerging research frontiers?
- Are critical concerns, such as AI explainability, sufficiently visible in the vocabulary of the field, or do their relative absence signal a conceptual blind spot?

The keywords generated by the LDA model represent the most characteristic terms for each topic, selected for their high probability of matching within a thematic group. These keywords allow for content interpretation without reviewing the entire corpus and act as a bridge between the original text and a more structured analysis. Visualizing them as in figure 2 validates thematic coherence and facilitates the exploration of semantic relationships and gaps in the literature.

Figure 2. *The 100 most representative terms identified through topic-based word modelling*



LDA clusters bigrams according to their frequency of co-occurrence in texts and assigns them a probability within each topic, indicating how representative they are for that thematic group. Figure 2 presents the 100 most representative bigrams - pairs of words that frequently co-occur within the corpus - which significantly enhance the semantic granularity and interpretive depth of the thematic analysis. Unlike single-word frequencies, which may yield overly generic or contextually ambiguous findings, bigrams allow for the detection of domain-specific conceptual units that offer more precise and meaningful interpretations. A notable example is the expression “artificial

intelligence”, whose combined meaning constitutes a well-established and distinct research domain, whereas its individual components as “artificial” and “intelligence” lack such clarity when considered in isolation.

By revealing these collocations, the bigram analysis contributes to reconstruct the underlying conceptual architecture of the literature. It illustrates how terms are operationalized in context and how they contribute to the formulation of key academic debates. For instance, bigrams such as “machine_learning”, “neural_network” and “digital_transformation” not only reflect dominant research themes but also anchor broader discussions about the evolving role of AI in managerial decision-making and control systems.

Furthermore, the identification of these compound expressions serves as a crucial intermediate step in the topic modelling process. Bigrams were used both to label the LDA-generated topics and to interpret their content with greater specificity. Their frequency and distribution offer arguments for how scholars articulate complex ideas, facilitating the construction of more coherent thematic clusters (see table 1) and enabling more accurate cross-topic comparisons.

Although both models (figure 2 and table 1) use LDA, they are analysing different units of language and therefore identify different words or combinations as representative.

Following the implementation of LDA, five distinct thematic clusters were generated based on the distribution of terms across the corpus. To enhance the semantic interpretability of these machine-generated clusters, a two-step process involving LLM – GPT-4o model interventions was employed, allowing for a more accurate thematic classification and evaluation of content relevance.

In the first step, the top 20 bigrams with computed betas associated with each cluster - representing the most frequent and meaningful word pairings within each cluster - were extracted and used as input for the LLM. The model was prompted to generate a concise and conceptually coherent title for each cluster based on these bigrams and their representativeness within the overall distribution. This procedure enabled the transformation of abstract probabilistic topic structures into human-readable thematic labels, thereby facilitating meaningful interpretation and subsequent analysis.

Table 1. The 50 most representative bigrams in MCS- AI intersection

Word rank	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5
1	<i>machine</i>	<i>neural</i>	<i>digital</i>	<i>supply</i>	<i>decision</i>
	<i>learning</i>	<i>network</i>	<i>transformation</i>	<i>chain</i>	<i>support</i>
2	<i>learning</i>	<i>information</i>	<i>internet</i>	<i>design</i>	<i>support</i>
	<i>algorithms</i>	<i>limited</i>	<i>things</i>	<i>methodology</i>	<i>system</i>
3	<i>manufacturing</i>	<i>limited</i>	<i>quality</i>	<i>methodology</i>	<i>support</i>
	<i>systems</i>	<i>trading</i>	<i>management</i>	<i>approach</i>	<i>systems</i>
4	<i>management</i>	<i>artificial</i>	<i>reinforcement</i>	<i>originality</i>	<i>proposed</i>
	<i>systems</i>	<i>neural</i>	<i>learning</i>	<i>value</i>	<i>system</i>
5	<i>energy</i>	<i>consumer</i>	<i>proposed</i>	<i>purpose</i>	<i>information</i>
	<i>consumption</i>	<i>behaviour</i>	<i>method</i>	<i>paper</i>	<i>systems</i>
6	<i>production</i>	<i>intelligence</i>	<i>exclusive</i>	<i>chain</i>	<i>knowledge</i>
	<i>management</i>	<i>machine</i>	<i>licence</i>	<i>management</i>	<i>management</i>
7	<i>energy</i>	<i>security</i>	<i>cloud</i>	<i>limitations</i>	<i>expert</i>
	<i>management</i>	<i>stability</i>	<i>computing</i>	<i>implications</i>	<i>systems</i>
8	<i>learning</i>	<i>using</i>	<i>value</i>	<i>performance</i>	<i>inventory</i>
	<i>models</i>	<i>artificial</i>	<i>chain</i>	<i>measurement</i>	<i>management</i>
9	<i>smart</i>	<i>intelligence</i>	<i>mobile</i>	<i>literature</i>	<i>production</i>
	<i>manufacturing</i>	<i>technology</i>	<i>health</i>	<i>review</i>	<i>planning</i>
10	<i>project</i>	<i>stability</i>	<i>management</i>	<i>customer</i>	<i>expert</i>
	<i>management</i>	<i>strategy</i>	<i>practices</i>	<i>satisfaction</i>	<i>system</i>

In the second step, the same bigrams were analysed by the LLM to determine whether each cluster explicitly or implicitly addressed issues related to AI explainability and transparency. The LLM was asked to assess the presence of concepts associated with algorithmic interpretability, prevention of hallucinations, and trust-building in AI-enabled systems.

The resulting cluster titles are as follows:

- Cluster 1 AI Integration in MCS for smart manufacturing and energy efficiency
- Cluster 2 Neural models and performance control: Toward transparent and explainable ai in consumer analytics

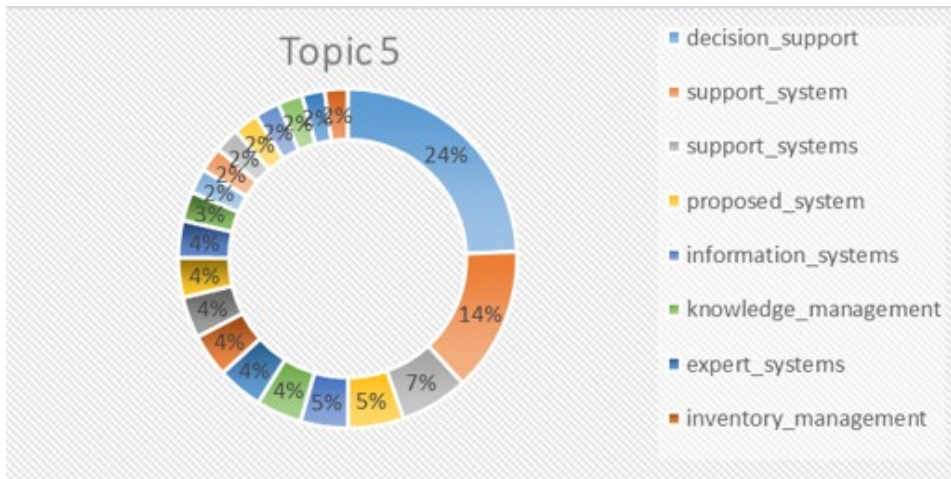
- Cluster 3 Management control in digital transformation: Aligning AI tools with quality and value chain objectives
- Cluster 4 Explainable AI in supply chain forecasting: Enhancing risk control through management systems
- Cluster 5 Decision support systems and sustainable control: Building trust in AI for managerial decision-making

A quick analysis of the clusters leads us to the following insights: the first cluster is related to *smart manufacturing* and is more closely associated with the engineering field; the second focuses on *consumer analytics*, within the area of marketing; the third concerns *quality and value chains*, aligning with quality management research; the fourth clearly centers on *supply chain* topic, mainly linked to logistics; and the fifth is associated with *managerial decision-making*, which is especially relevant to managerial control.

The results of this targeted analysis revealed that only two clusters - cluster 2 and cluster 4 - contained discernible thematic elements connected to explainability and transparency. Cluster 2 focused on neural models and their application in consumer analytics, highlighting the need for transparency in algorithmic decision-making. Cluster 4 dealt with supply chain forecasting, where explainability appeared as a mechanism to enhance risk management through interpretable AI tools. In contrast, the remaining three clusters - although thematically relevant to the broader integration of AI in management control - lacked any significant references to explainability-related constructs.

Finally, figure 3 presents a detailed view of cluster 5, which focuses on decision support and sustainable control. The figure shows the 20 most representative bigrams within this cluster, along with their relative probabilities (β), highlighting the key terms that define this specific area of the literature. The higher the β value for a term, the greater its relevance within the cluster.

Figure 3. Detailed analysis of cluster 5, including the 20 main bigrams and their respective percentages of representation



4.1. Discussion

This study offers a novel, data-driven mapping of the intersection between MCS and AI, revealing a fragmented yet rapidly evolving body of literature. Through LDA topic modelling complemented by LLM-assisted interpretation, five distinct thematic clusters were identified. Although explicability and interpretability are widely recognised as fundamental principles for the responsible use of AI, our analysis shows that these concerns remain peripheral in current MCS-AI research. Only two clusters contain elements explicitly related to explainability, while the other three, despite being highly relevant to the integration of AI in management control, lack significant references to this critical dimension.

The case of cluster 5, which focuses on decision support and sustainable control, is particularly illustrative. While it includes terms such as “trust” and “decision-making”, it does not adequately address the technical or ethical implications of using opaque AI systems. This highlights a conceptual gap: managerial decision contexts inherently demand accountability and auditability, yet the literature does not yet reflect a systematic engagement with the risks posed by black-box models. This may partly reflect the fact that managers are not fully aware of how AI-driven automated decision processes can impact organisational transparency and responsibility.

This finding underscores a broader issue: although the discourse on AI within management control is expanding, the integration of explainability as a core research

concern remains uneven and fragmented. The two-step LLM-assisted analysis proved especially valuable in identifying these latent conceptual gaps, pointing to the need for greater scholarly attention to the responsible and interpretable use of AI-enabled control systems.

From a disciplinary perspective, this study contributes to the emerging dialogue between accounting, management control research, and AI. While explainability is well established in AI studies, its systematic incorporation into MCS frameworks is still underdeveloped. Our results highlight the need for stronger theoretical integration, where concepts such as interpretability, algorithmic bias, and trustworthiness become embedded in the design and evaluation of AI-driven control practices.

Building on this, the literature review suggests several promising directions for future research. First, it is important to examine how different types of MCS (e.g., diagnostic versus interactive systems) may mediate or moderate the role and effectiveness of AI explainability mechanisms. Second, empirical studies, particularly those employing experimental or case-based designs, should investigate how managers interpret, trust, and respond to AI-generated outputs under varying levels of transparency. Third, conceptual work is needed to clarify what “explainability” should mean in the context of MCS: is it sufficient to offer post-hoc justifications, or should AI-driven control tools be inherently interpretable by design? Addressing this question will be crucial for theory development and practical guidance.

Finally, researchers should explore whether certain control environments, such as high-uncertainty or high-stakes contexts, place greater demands on explainable AI, and how organisations can adapt their control architectures to meet these demands effectively.

5. CONCLUSION

Given the increasing emphasis on efficiency also in research, it is essential to employ methods that provide timely and cost-effective insights while also promoting objectivity in the interpretation of complex literature. This paper demonstrates how the integration of NLP, advanced topic modelling techniques and interpretive support from LLMs facilitates the thematic organization of a fragmented body of literature and contribute to the strategic structuring of an academic field. By extracting latent patterns of co-occurrence and clustering them into coherent themes, this approach moves beyond traditional literature reviews, to offer a systematic mapping of intellectual landscape that highlights dominant conceptual cores as well as peripheral or underexplored areas of inquiry.

At the heart of this contribution lies its capacity to reveal research gaps with a relatively higher degree of objectivity and analytical precision. Unlike narrative or subjective judgements based on few numbers of academic contributions that often depend on expert selection and interpretation, this method identifies not only what is being studied, but, critically, what is not. Through comparative analysis of the thematic coverage across clusters, it becomes possible to pinpoint which foundational concepts such as explainability and transparency in AI systems, are systematically embedded in some areas yet conspicuously absent in others where they would be expected to play a central role, as it has been revealed in this work regarding the MCS field.

This evidence-based identification of omissions or imbalances in thematic coverage provides a clear and justifiable rationale for setting future research priorities. It enables scholars to engage in more strategic theorizing and research design by focusing on underexplored intersections, unaddressed ethical concerns, or conceptually inconsistent applications within the field. In our case, the finding that only two of five major clusters engage meaningfully with XAI signals not only a conceptual oversight but a potentially high-impact area for scholarly intervention.

Although in practice, MCS are increasingly supported by algorithmic tools, particularly in areas such as performance measurement, forecasting, or decision-making processes, cluster 5, titled “Decision support systems and sustainable control: Building trust in ai for managerial decision-making,” contains no explicit reference to XAI, despite the theoretical expectation that both concepts should be closely related. This absence may be explained by the inherent complexities of operationalizing XAI within organizational contexts, or by the limited availability of empirical studies that integrate both MCS and XAI as core constructs. Such a gap highlights the need for further research into how XAI can be effectively embedded into AI-supported control systems to enhance transparency, trust, and accountability in managerial decision-making.

From a broader research management perspective, particularly in relation to doctoral programmes and PhD supervision, —this methodology can offer valuable practical support. Supervisors and research programme coordinators can use it to assess whether PhD proposals or research projects align with current trends and gaps in the literature.

However, despite the growing capabilities of AI tools, researchers still need to engage deeply with the literature; there is no substitute for reading extensively to ensure critical understanding and nuanced interpretation. AI tools, as these presented in this paper, should be seen only as an additional aid for systematic literature reviews.

Author contributions:

Alonso, O. F.: Software, Investigation, Writing – original draft, Writing, review, and editing. **Shakina, E.:** Conceptualization, Methodology, Validation, Investigation, Data curation, Supervision, Project administration. **Gonzalez-Sanchez, M. B.:** Validation, Formal analysis, Investigation, Writing – original draft, Writing, review, and editing, Supervision, Project administration. **Berbel-Vera, J.:** Validation, Formal analysis, Investigation, Writing – original draft, Writing, review, and editing, Supervision.

Omar Feria Alonso (Alonso, O. F.)

Elena Shakina (Shakina, E.)

Maria Beatriz Gonzalez-Sanchez (Gonzalez-Sanchez, M. B.)

Jose Berbel-Vera (Berbel-Vera, J.)

Conflict of interest statement

Authors declare that, throughout the research process, there has not been any sort of personal, professional, or economic interest that may have influenced the researchers' judgement and/or actions during the elaboration and publication of this article.

REFERENCES

- Abdel-Karim, B. M., Pfeuffer, N., & Hinz, O. (2021). Machine learning in information systems—a bibliographic review and open research issues. *Electronic Markets*, 31, 643-670. <https://doi.org/10.1007/s12525-021-00459-2>
- Ahrens, T., & Chapman, C. S. (2004). Accounting for flexibility and efficiency: A field study of management control systems in a restaurant chain. *Contemporary Accounting Research*, 21(2), 271-301. <https://doi.org/10.1506/VJR6-RP75-7GUX-XH0X>
- Anthony, R. N., & Govindarajan, V. (2007). *Management control system* (12th ed.). McGraw Hill Companies Inc.
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Benetot, A., Tabik, S., Barbado, A., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Bedford, D. S. (2015). Management control systems across different modes of innovation: Implications for firm performance. *Management Accounting Research*, 28, 12-30. <https://doi.org/10.1016/j.mar.2015.04.003>

- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993-1022. <https://www.jmlr.org/papers/v3/blei03a.html>
- Cao, J., Xia, T., Li, J., Zhang, Y., & Tang, S. (2009). A density-based method for adaptive LDA model selection. *Neurocomputing: An International Journal*, 72(7-9), 1775-1781. <https://doi.org/10.1016/j.neucom.2008.06.011>
- Chenhall, R. H., & Moers, F. (2015). The role of innovation in the evolution of management accounting and its integration into management control. *Accounting, Organizations and Society*, 47, 1-13. <https://doi.org/10.1016/j.aos.2015.10.002>
- Deveaud, R., Sanjuan, E., & Bellot, P. (2014). Accurate and effective latent concept modelling for ad hoc information retrieval. *Document numérique*, 17(1), 61-84. <https://doi.org/10.3166/dn.17.1.61-84>
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., & Williams, M. D. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 1-49. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- Elbashir, M. Z., Sutton, S. G., Mahama, H., & Arnold, V. (2021). Unravelling the integrated information systems and management control paradox: Enhancing dynamic capability through business intelligence. *Accounting & Finance*, 61(S1), 1775-1814. <https://doi.org/10.1111/acfi.12644>
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organisational decision making. *Business Horizons*, 61(4), 577-586. <https://doi.org/10.1016/j.bushor.2018.03.007>
- Kraus, S., Breier, M., Lim, W. M., Dabić, M., Kumar, S., Kanbach, D., Mukherjee, D., Corvello, V., Piñeiro-Chousa, J., Liguori, E., Palacios-Marqués, D., Schiavone, F., Ferraris, A., Fernandes, C., & Ferreira, J. J. (2022). Literature reviews as independent studies: Guidelines for academic practice. *Review of Managerial Science*, 16(8), 2577-2595. <https://doi.org/10.1007/s11846-022-00588-8>
- Masialeti, M., Talaei-Khoei, A., & Yang, A. T. (2024). Revealing the role of explainable AI: How does updating AI applications generate agility-driven performance? *International Journal of Information Management*, 77, 1-17. <https://doi.org/10.1016/j.ijinfomgt.2024.102779>
- Meske, C., Bunde, E., Schneider, J., & Gersch, M. (2022). Explainable artificial intelligence: Objectives, stakeholders and future research opportunities. *Information Systems Management*, 39(1), 53-63. <https://doi.org/10.1080/10580530.2020.1849465>

Murphy, B., Feeney, O., Rosati, P., & Lynn, T. (2024). Exploring accounting and AI using topic modelling. *International Journal of Accounting Information Systems*, 55, 1-18. <https://doi.org/10.1016/j.accinf.2024.100709>

Pugliese, S., Giannetti, V., & Banerjee, S. (2023). How to conduct efficient and objective literature reviews using natural language processing: A step-by-step guide for marketing researchers. *Psychology & Marketing*, 41(2), 223-442. <https://doi.org/10.1002/mar.21931>

Simons, R. (1995). *Levers of control: How managers use innovative control systems to drive strategic renewal*. Harvard Business School Press. <https://doi.org/10.1002/smj.4250150301>

Speklé, R. F., van Elten, H. J., & Widener, S. K. (2017). Creativity and control: A paradox—evidence from the levers of control framework. *Behavioural Research in Accounting*, 29(2), 73-96. <https://doi.org/10.2308/bria-51759>

Teerasoponpong, S., Chernbumroong, S., & Suntornnond, V. (2025). Meta-analysis on sustainable development challenges: Lessons from Asia-Pacific industrial estates. *Business Strategy and the Environment*, (0), 1-23. <https://doi.org/10.1002/bse.70082>

Tessier, S., & Otley, D. (2012). A conceptual development of Simons' Levers of Control framework. *Management Accounting Research*, 23(3), 171-185. <https://doi.org/10.1016/j.mar.2012.04.003>

Weber, P., Carl, K. V., & Hinz, O. (2023). Applications of explainable artificial intelligence in finance—a systematic literature review of finance, information systems, and computer science literature. *Management Review Quarterly*, 74, 867-907. <https://doi.org/10.1007/s11301-023-00320-0>

Yu, L., & Li, Y. (2022). Artificial intelligence decision-making transparency and employees' trust: The parallel multiple mediating effect of effectiveness and discomfort. *Behavioural Sciences*, 12(5), 1-17. <https://doi.org/10.3390/bs12050127>

Reception date: 29/07/2025

Review date: 30/07/2025

Acceptance date: 01/08/2025

Contact: omar.feria@alumnado.uvigo.gal