Intelligent Reconstruction of Chinese E-commerce Platforms under Large Language Models: A Systematic Literature Review and an Integrative Pathway-Capability-Performance Framework

Guojin Lin*
Farabi International Business School, Al-Farabi Kazakh National University, Almaty 050040, Republic of Kazakhstan
*Corresponding email: lgj3089@gmail.com

Abstract

The diffusion of large language models (LLMs) since late 2022 has begun to alter the architecture of digital commerce. However, scholarship on the phenomenon remains dispersed across disciplines that seldom engage one another. This article synthesises the evidence on LLM-driven reconstruction of Chinese e-commerce platforms and develops a theoretical framework capable of organising it. The review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 protocol. Searches of Scopus, Web of Science, and Google Scholar returned 1,142 records for the period 2019 to 2026. After screening, 87 peer-reviewed studies were retained and triangulated against 24 corporate filings and industry reports covering the 2024 and 2025 financial years of Alibaba, JD.com, and PDD Holdings. Thematic synthesis identifies three streams of enquiry: technical application, strategic capability, and performance evaluation. It also identifies five recurring implementation pathways spanning recommendation and search, conversational commerce, content generation, supply-chain orchestration, and risk governance. Three gaps emerge. The first is a shortage of comparative multi-case evidence built on publicly verifiable data. The second is weak theoretical integration across levels of analysis. The third is immature measurement of performance. In response, the article proposes a Pathway-Capability-Performance (PCP) framework. This framework joins the resource-based view, dynamic-capabilities theory, and platform economics in a single recursive structure. In this structure, implementation pathways reshape six capability domains, and reshaped capabilities generate performance across five dimensions. Performance feedback then conditions subsequent pathway choices. The framework contributes an integrative lens for strategic management research on generative artificial intelligence. It also offers platform executives a diagnostic logic for aligning technological choices with the capabilities those choices require.

Keywords

Large language models, E-commerce platforms, Intelligent reconstruction, Dynamic capabilities, Digital transformation, Platform strategy

Introduction

The public release of ChatGPT at the end of November 2022 turned a decade-old research advance into a technology that began reshaping commercial practice almost immediately. This advance was the transformer architecture. Few sectors sit closer to the centre of that disturbance than electronic commerce.

A commerce platform is, to a considerable extent, an apparatus for producing and acting upon text. These texts include product descriptions, search queries, recommendations, customer dialogue, and merchant communications. A technology that generates and interprets language at scale therefore touches nearly every layer of such a platform. What might elsewhere amount to an incremental upgrade becomes, in commerce, a reorganisation of platform operations.

China offers the setting in which this change carries the

greatest weight. Its online retail market, valued at roughly 1.19 trillion USD in 2024, is the largest in the world. Three firms - Alibaba, JD.com, and PDD Holdings - together account for close to 70.00% of gross merchandise value. These platforms are not merely large adopters of technology developed elsewhere. They are frontier developers, and they have taken visibly different routes.

Alibaba released its Qwen model family in April 2023. By early 2026, it had passed 1.00 billion cumulative downloads with more than 100,000 derivative models on a single open repository. In February 2025, the company committed at least 380.00 billion RMB to artificial intelligence and cloud infrastructure over three years. This sum exceeds its total such spending across the preceding decade.

JD.com consolidated its model portfolio under the JoyAI brand in 2025, spanning three billion to 750.00 billion parameters. It reported more than 50,000 AI agents operating across its internal systems. PDD Holdings, the most reticent of the three about a flagship general-purpose model, concentrates instead on an algorithm-centred matching engine. It serves over 900.00 million annual active buyers.

Scholarship has not kept pace with this activity. Research in computer science examines individual applications with precision, including recommendation, conversational dialogue, product classification, and generative advertising. However, it rarely connects them to firm-level strategy or to measurable economic outcomes. Management research moves in the opposite direction. It treats language models as an undifferentiated instance of digital technology without attending to what distinguishes them.

The literature is consequently rich in fragments and poor in synthesis [1]. It offers practitioners little by way of an instrument for judging whether an investment in these technologies is creating value. This article responds on two fronts. It provides a systematic synthesis of what is known about LLM-driven reconstruction in Chinese e-commerce, conducted under a transparent and reproducible protocol.

It then proposes a PCP framework. This framework integrates the resource-based view, dynamic-capabilities theory, and platform economics into a single recursive structure [2]. The guiding research questions are as follows. Through what pathways do platforms embed LLMs into the commerce stack? How do those deployments reconstruct platform capabilities? How is the resulting performance conceived and measured? The remainder of the article describes the review method, sets out the theoretical foundations, reports the thematic synthesis, develops the framework, discusses its implications, and concludes.

Materials and methods

The review followed the PRISMA 2020 reporting guideline. This guideline was chosen because a narrative reading of such a dispersed literature would risk both omission and selection bias. Searches were run across Scopus, Web of Science, and Google Scholar for the period January 2019 to April 2026 [3]. This window is broad enough to bracket the release of Generative Pre-trained Transformer 3 (GPT-3) in 2020 and the first wave of Chinese language models in 2023.

The Boolean strategy combined three concept blocks. These were a technology block (large language model, generative AI, LLM and ChatGPT), a context block (e-commerce, platform, online retail and marketplace), and an outcome block (capability, transformation, performance and strategy).

The initial yield of 1,142 records fell to 932 after removal of duplicates, to 214 after screening on title and abstract, and to 87 after full-text assessment. The full-text assessment applied four eligibility criteria: explicit treatment of large language models or comparable generative technologies, a focus on commerce, retail, or platform contexts, availability of the full text in English, and publication in a peer-reviewed venue or a recognised pre-print archive. To these were added 24 corporate filings and industry reports. These included annual reports, submissions to the United States Securities and Exchange Commission on Forms 20-F and 6-K, disclosures to the Hong Kong Stock Exchange, and analyst studies. Platform-level evidence on the three focal firms cannot be reconstructed from the academic literature alone [4].

Documents were coded through an open pass followed by axial grouping in a structured manner. Coding was organised around three a-priori categories corresponding to the three research questions, with emergent codes

permitted throughout. Table 1 summarises the screening stage.

Table 1. PRISMA screening summary.

Stage	Records
Identified through database searches	1,142
Removed before screening (duplicates)	210
Screened on title and abstract	932
Excluded after title and abstract	718
Assessed in full text	214
Excluded after full-text assessment	127
Peer-reviewed articles retained	87
Corporate filings and industry reports added	24
Total documents synthesised	111

Note: This source was compiled by the author.

The attrition recorded in Table 1 is itself informative. Roughly three-quarters of the records identified by the search proved, on inspection, to address language models or commerce but not their intersection, which is one indication of how thinly the specific question examined here has been studied.

Theoretical foundations

The framework developed in this study draws on three traditions in strategic management. Each illuminates one facet of the phenomenon, and none is adequate alone.

The resource-based view

The resource-based view holds that durable advantage flows from resources that are valuable, rare, inimitable, and non-substitutable. Three resource classes increasingly satisfy that test in the present setting: proprietary training data, computational infrastructure, and algorithmic talent [5].

The focal platforms illustrate their convergence. JD.com's industrial model draws on more than fifty-seven million proprietary supply-chain stock-keeping units, a corpus no competitor can readily assemble. Alibaba's Qwen family rests on a computational base whose scale is a direct function of the infrastructure commitment noted above.

Yet the tradition has long been criticised for treating resources as relatively static stocks. It also says little about how firms reconfigure them when technology shifts abruptly or when competitive landscapes change rapidly. That limitation matters acutely for a discontinuity as sharp as the one under examination here.

Dynamic capabilities

Dynamic capabilities theory supplies what the resource-based view lacks. Dynamic capabilities refer to the capacity to integrate, build, and reconfigure competences in changing environments. This capacity is later disaggregated into sensing opportunities, seizing them through resource commitment, and transforming the organisation when opportunity demands it.

These capabilities have been recast as identifiable processes. Their value lies in their effect on the firm's resource configuration rather than in any mysterious quality of their own. The framework has been carried into research on digital transformation. More recently, it has been extended into work on artificial intelligence specifically [6].

Applied here, sensing means detecting demand for conversational interfaces or generative merchant tools. Seizing means committing to a proprietary or an open-weight model strategy. Transforming means redesigning the workflows that surround the model. It is this third element that gives the present study its central object.

Platform economics

Platforms are not ordinary firms. They are multi-sided markets in which value arises from interactions among distinct user groups. This value is governed by network effects, complementarities, and the rules of ecosystem participation [7].

Because the value of a network rises with the number of its participants, a resource that strengthens a platform's matching function is worth more there. The same

resource would be worth less in a stand-alone firm.

Language models intervene in this logic along at least four lines. First, they lower search and matching costs. Second, they make personalisation economic at a finer grain than was previously feasible. Third, they create merchant-side products such as generative listing tools. Fourth, they shift the boundary between platform and complementor.

Content and code that once required human intermediaries can now be machine-generated. This shift further changes the relationship between platforms and their participants.

The logic of integration

The three traditions answer distinct but interrelated questions and are therefore complementary rather than competing. The resource-based view establishes which resources matter; dynamic-capabilities theory explains how they are reconfigured; platform economics accounts for why their value is amplified. Recent scholarship has argued convincingly for combining the first two, and the present article extends that argument by adding the platform dimension. It also organizes all three coherently into the sequential structure presented here [8].

Results: Thematic synthesis

The corpus organises into three streams of enquiry, corresponding to the three research questions.

Implementation pathways and applications

The largest stream concerns the practical uses to which language models are put. Recommendation and search dominate: Model embeddings improve cold-start performance and cross-domain generalisation. Dedicated architectures also address the difficulties that large, multi-dimensional catalogues pose for classical collaborative filtering [9,10].

Conversational commerce forms a second cluster, in which language models prove well suited to the pre-sale dialogue central to Chinese live-commerce. Automated service systems resolve the majority of customer inquiries without human intervention. Content generation forms a third, with generative tools producing descriptions and advertising at scale.

Supply-chain and operational deployment form a fourth pathway. Risk governance forms a fifth, addressing fairness, hallucination, and adversarial manipulation. Table 2 maps pathways to evidence from focal platforms.

Table 2. Application pathways and indicative platform evidence.

Pathway	Representative function	Indicative evidence
Recommendation and search	Personalised matching; semantic search	Qwen integration in Taobao; PDD matching engine; JD large-scale recommendation model
Conversational commerce	Pre-sale dialogue; chat-based shopping	JD intelligent service with sub-second response; Qwen consumer application
Content generation	Product descriptions; advertising copy	Alibaba generative merchant tools; JD content-marketing platform
Supply chain and operations	Inventory; forecasting; logistics	JD industrial model trained on 57 million supply-chain records; automated warehousing
Risk and governance	Fairness; hallucination; manipulation control	Cross-platform compliance and fraud-prevention investment

Note: This table was compiled by the author from the reviewed corpus and corporate disclosures.

Read across the row, Table 2 shows that the five pathways are not evenly distributed among the platforms. Each firm concentrates its effort differently, and that concentration is the empirical trace of a strategic choice rather than an accident of implementation.

Capability reconstruction

A second stream asks what changes inside the platform once these applications are adopted at scale. Read through the sensing-seizing-transforming sequence, six capability domains recur across the corpus. These domains are data and algorithmic capabilities, which govern how the platform perceives its environment and converts perception into model behaviour.

They also include computational and organisational

capabilities, which determine whether a sensed opportunity can be acted upon. Merchant-facing and ecosystem-coordination capabilities also recur, both of which require the platform to reach beyond its own boundary.

Systematic evidence indicates that data-driven and artificially intelligent capabilities mediate the relationship between digital transformation and competitive advantage. The open-sourcing of the Qwen family has been read in these terms as an ecosystem-coordination move. It attracts a developer community whose derivative output now exceeds that of any other single organization [11].

Performance evaluation

The third stream is the least developed. Evidence is largely case-based and operational, such as gains in service automation, revenue uplift attributed to personalisation, and faster fulfilment. Such evidence is reported chiefly through metrics that firms disclose about themselves, with little longitudinal or cross-firm comparison.

A bibliometric analysis of artificial intelligence in electronic commerce confirms that rigorous performance measurement remains a frontier rather than a settled practice. The construction of defensible indicator systems, built on data that third parties can verify, is in effect unfinished business.

Strategic differentiation among the focal platforms

Triangulating the academic literature against corporate disclosure yields a finding that neither source establishes on its own: The three platforms are pursuing systematically different strategies, and those differences are legible in their public reporting. Table 3 sets out the contrast.

Table 3. Strategic profiles of the three focal platforms, latest reported financial year.

Dimension	Alibaba	JD.com	PDD Holdings
Flagship model	Qwen (Tongyi Qianwen)	JoyAI (formerly Yanxi)	In-house recommendation models
Model orientation	General-purpose; open-weight	Domain-specialised; industrial	Task-specific; proprietary
Principal locus of effort	Cloud and foundation models	Supply chain and operations	Matching and distribution
Headline commitment	380.00 billion RMB infrastructure programme	AI research spending up sharply; ecosystem pledge	Research spending raised; no model platform
Signature disclosed metric	Qwen downloads exceeding one billion	Over 50,000 internal AI agents	Over 900.00 million annual active buyers

Note: The data presented in this table were compiled by the author from corporate annual reports, regulatory filings, and results announcements.

The pattern in Table 3 is coherent rather than random. Alibaba builds outward, positioning model capability as infrastructure for others. JD.com builds inward, reorganising its own operations around models grounded in proprietary industrial data. PDD builds narrowly, deepening a single decisive capability and declining the expense of a general-purpose model platform. These are recognisably different answers to the same technological question, and any framework that aspires to explain reconstruction must be able to accommodate all three. This diversity suggests that platform strategy depends heavily on existing resource endowments and organizational boundaries.

The PCP framework

The synthesis exposes a structural deficit in the literature: the three levels of analysis - what platforms do, what changes inside them, and what changes in their measured outcomes - are typically studied in isolation. The empirical reality, however, is sequential and recursive. Pathway choices shape which capabilities are reconfigured; reconfigured capabilities produce performance outcomes; and performance feedback conditions the next round of pathway choices [12]. The framework proposed here is designed to keep that recursion visible. Figure 1 presents its architecture.

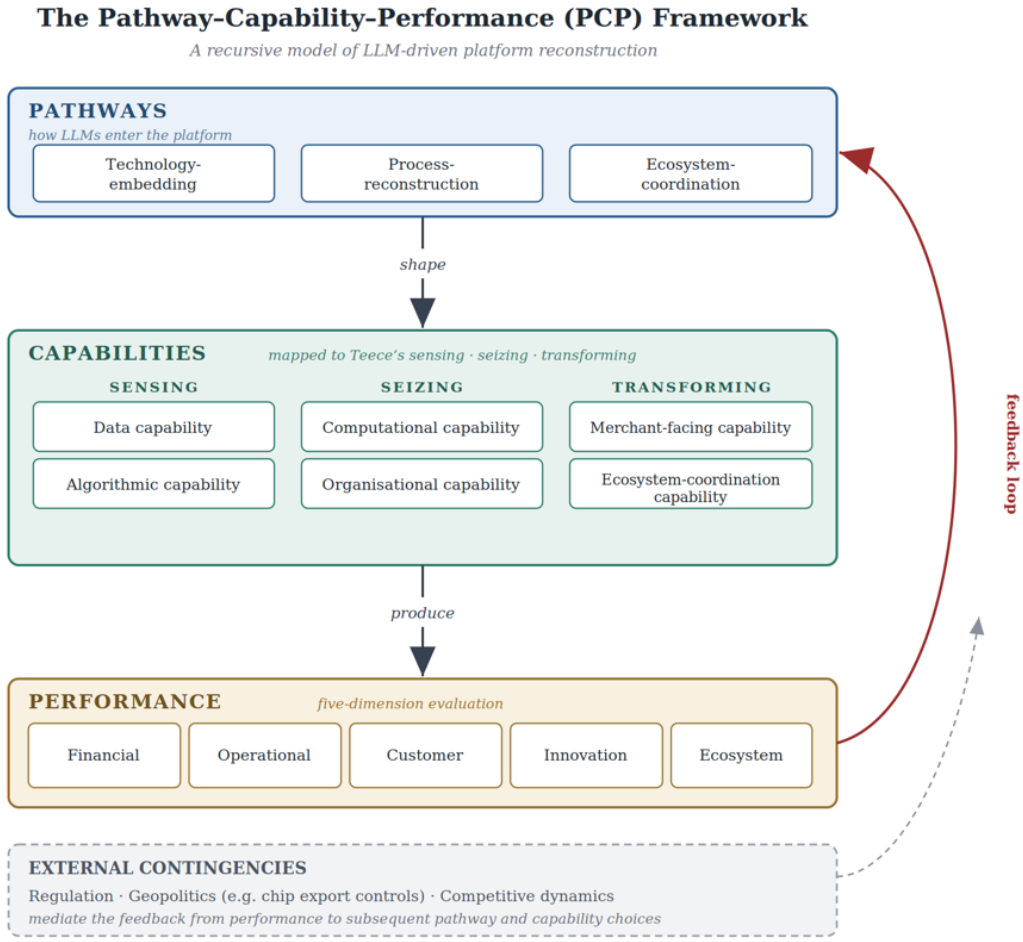


Figure 1. The PCP framework.

As Figure 1 indicates, the first layer concerns pathways. Drawing in this study and on the literature of strategic renewal, three meta-pathways are distinguished. A technology-embedding pathway inserts language models incrementally into existing modules. A process-reconstruction pathway redesigns cross-functional workflows around model-mediated automation. An ecosystem-coordination pathway opens model capability to merchants, developers, and external complementors. These are not mutually exclusive. Platforms pursue all three simultaneously, but with differing emphasis. It is the pattern of emphasis, evident in Table 3, that distinguishes one strategy from another.

The second layer concerns capabilities. The six domains identified previously map onto the established triad. Data and algorithmic capabilities constitute sensing. Computational and organisational capabilities constitute seizing. Merchant-facing and ecosystem-coordination capabilities constitute transforming. This mapping converts a loose and frequently invoked notion into a structure that can be observed and tested [13].

The third layer concerns performance, conceived along five dimensions in the spirit of the balanced-scorecard tradition as adapted for the platform era: financial, operational, customer, innovation, and ecosystem. No single dimension suffices. Reconstruction that improves one while degrading another is not the unambiguous success a solitary metric would report. A platform may compress margin deliberately in order to build an ecosystem whose returns accrue only later. An evaluation confined to financial indicators would misread that investment as failure.

What distinguishes the framework from a static taxonomy is the feedback loop shown on the right of Figure 1. Performance outcomes return to condition capability investment and pathway choice. That return is mediated by several external contingencies, including regulation, geopolitical constraints such as export controls on advanced semiconductors, and competitive dynamics. The framework therefore describes a dynamic loop rather than a line. It is intended for repeated application over time rather than for a single verdict. Table 4 summarises its constructs and their theoretical anchors in detail.

Table 4. Constructs and theoretical anchors of the PCP framework.

Layer	Constituent constructs	Theoretical anchors
Pathways	Technology-embedding; process-reconstruction; ecosystem-coordination	Vial (2019); Warner and Wäger (2019)
Capabilities	Data; algorithmic; computational; organisational; merchant-facing; ecosystem-coordination	Barney (1991); Teece (2007); Ellström et al. (2022)
Performance	Financial; operational; customer; innovation; ecosystem	Verhoef et al. (2021)
Feedback and contingency	Recursive loops; regulation; geopolitics; competition	Teece (2007)

Note: This framework was developed by the author.

Table 4 makes plain that the framework is not an eclectic borrowing but a division of labour among traditions. Each layer is anchored where that tradition has explanatory purchase. The layers are joined by the sequence rather than merged into a single undifferentiated construct.

Discussion

Research gaps and an agenda

Three gaps follow from the synthesis, and each suggests a line of work. The first is empirical. Most case-based studies examine a single company or a single application. The regulatory filings that platforms are obliged to publish offer an under-used basis for triangulating strategy across firms without recourse to private data. The second is theoretical: Pathway, capability, and performance are seldom integrated within one explanatory architecture. This is the deficit the framework proposed here is designed to remedy. The third is methodological. Performance is reported anecdotally far more often than rigorously. The field has yet to develop reproducible indicator systems of the kind that the analytic hierarchy process (AHP), entropy weighting (EW), or data envelopment analysis (DEA) could support. Addressing the third gap is, in the author's view, the most immediately tractable of the three, because the required data are already public [14].

Theoretical and practical implications

For theory, the contribution is integrative. Strategic management has accumulated three powerful but separately deployed traditions. The arrival of general-purpose language models provides an unusually clear occasion for joining them, because the phenomenon itself refuses to stay within any one of their boundaries. For practice, the framework suggests that pathway choices ought to follow from a candid audit of capabilities rather than from ambition alone.

The most expensive error visible in the evidence assembled here is the commitment to an ecosystem-coordination pathway by a firm lacking the computational and organisational capabilities to sustain it. This is a mismatch the framework is designed to expose before the investment is made rather than after.

Limitations

Three limitations qualify these conclusions. The literature is moving quickly enough that any review carries a stale-date risk. Developments after April 2026 are not captured. The search was confined to English-language sources, which may under-represent Chinese-language scholarship on the focal platforms. However, the corporate disclosures triangulated here are published in English by regulatory requirement.

Finally, the synthesis is interpretive rather than meta-analytic: Effect sizes are not pooled. Operational figures quoted from corporate sources are reported as disclosed rather than independently audited. This caveat applies with particular force to self-reported efficiency gains.

Conclusion

This review of 87 peer-reviewed studies and 24 corporate and industry documents supports three conclusions. The literature on large language models in Chinese e-commerce has matured rapidly along five

application pathways. However, it has not connected those pathways either to capability theory or to measured performance. The resource-based view and dynamic-capabilities theory, supplemented by platform economics, provide sufficient scaffolding for that connection. The PCP framework attempts to supply this connection. And the three leading Chinese platforms, far from converging on a common template, are pursuing recognisably distinct strategies. The framework is able to accommodate and explain these differences.

The framework's value will ultimately be settled empirically rather than by argument. Work now in progress applies it through text-mining of the platforms' annual reports and earnings calls. It is further applied through grounded-theory coding of capability reconstruction. The work also includes constructing a weighted performance index benchmarked across firms. Should that programme succeed, the contribution will be more than a clearer account of how language models are reshaping the world's largest digital-commerce ecosystem. It will also offer an instrument through which managers can evaluate their own position within it.

Author Contributions

Lin Guojin: conceptualisation; methodology; formal analysis; investigation; visualisation; writing - original draft; writing - review and editing. The author has read and approved the submitted version.

Funding

This work was not supported by any fund.

Acknowledgements

The author thanks the academic supervisor of the doctoral programme for constructive comments on an earlier version of the framework presented here.

Conflicts of Interest

The author declares no conflict of interest.

References

[1] Barney, J. (1991) Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99-120.

[2] Chugh, P., Jain, V. (2024) Artificial intelligence (AI) empowerment in e-commerce: a bibliometric voyage. *NMIMS Management Review*, 32(3), 159-173.

[3] Ellström, D., Holtström, J., Berg, E., Josefsson, C. (2022) Dynamic capabilities for digital transformation. *Journal of Strategy and Management*, 15(2), 272-286.

[4] Jobstreibizer, J., Beliaeva, T., Ferasso, M., Kraus, S., Kallmuenzer, A. (2025) The impact of artificial intelligence on business models: a bibliometric-systematic literature review. *Management Decision*, 63(13), 372-396.

[5] Matz, S. C., Teeny, J. D., Vaid, S. S., Peters, H., Harari, G. M., Cerf, M. (2024) The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14(1), 4692.

[6] Kang, J. W., Das, S. B., Chanda, R., United Nations Economic and Social Commission for Asia and the Pacific. (2026) Advancing digital trade and cross-border e-commerce in Asia and the Pacific. *Asia-Pacific Sustainable Development Journal*, 33(2), 1-29.

[7] Pereira, V., Bamel, U. (2021) Extending the resource and knowledge based view: a critical analysis into its theoretical evolution and future research directions. *Journal of Business Research*, 132, 557-570.

[8] Teece, D. J. (2007) Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319-1350.

[9] Heaton, S., Teece, D., Agronin, E. (2023) Dynamic capabilities and governance: an empirical investigation of financial performance of the higher education sector. *Strategic Management Journal*, 44(2), 520-548.

[10] Verhoef, P. C., Broekhuizen, T., Bart, Y., Bhattacharya, A., Dong, J. Q., Fabian, N., Haenlein, M. (2021) Digital transformation: a multidisc-ciplinary reflection and research agenda. *Journal of Business Research*, 122, 889-901.

[11] Vial, G. (2021) Understanding digital transfor-mation: a review and a research agenda. *Managing Digital Transformation*, 13-66.

[12] Warner, K. S., Wäger, M. (2019) Building dynamic capabilities for digital transformation: an ongoing

process of strategic renewal. *Long Range Planning*, 52(3), 326-349.

[13] Zhao, Z., Fan, W., Li, J., Liu, Y., Mei, X., Wang, Y., Li, Q. (2024) Recommender systems in the era of large language models (LLMs). *IEEE Transactions on Knowledge and Data Engineering*, 36(11), 6889-6907.

[14] Chen, K., Cheng, H., Qin, Q. (2024) Assessing the impact of environmental accounting message disclosure quality on financing costs in high-pollution industries. *Journal of Cases on Information Technology (JCIT)*, 26(1), 1-18.