

FROM GRAPH THEORY TO EXPLAINABLE DECISION INTELLIGENCE: INTEGRATING GRAPH AI, EXPLAINABILITY AND MANAGERIAL DECISION-MAKING. A CRITICAL, INTEGRATIVE REVIEW AND RESEARCH AGENDA FOR MANAGEMENT RESEARCH

Shashikumar R

Associate Professor
Department of Mathematics, Government First Grade College, Yelahanka

Keerthi H K

Assistant Professor
Department of Business Administration, Nitte Meenakshi Institute of Technology, Nitte(Deemed to be
University, Bengaluru

ABSTRACT

Organizations increasingly make decisions in settings where employees, customers, suppliers, technologies, alliances and other stakeholders are densely interconnected, so that an action affecting one entity can propagate consequences across a wider relational system. Graph theory offers a natural language for representing such interdependence, and over recent decades its principles have diffused into network analysis, business analytics, graph databases, knowledge graphs, graph machine learning and graph neural networks (GNNs). More recently, the emergence of explainable artificial intelligence (XAI) and decision intelligence has opened a further possibility: moving graph-based analysis beyond describing relationships and predicting outcomes toward explaining, contextualising and informing managerial choice. The literature supporting this trajectory, however, remains fragmented across mathematics, computer science, information systems, operations research and management, with graph prediction and graph explanation developing considerably faster than research connecting these outputs to constrained, context-sensitive managerial decisions. This review synthesises that fragmented landscape through a critical, integrative methodology and proposes Graph-Based Explainable Decision Intelligence (GEDI) as an organizational capability comprising relational sensing, graph analytics, predictive intelligence, explainable intelligence, decision integration and adaptive reconfiguration. Rather than treating graph theory, GNNs or XAI as novel in themselves, the review's contribution lies in integrating these previously separate streams into a management-oriented architecture and in identifying where the evidence base is mature and where it remains conceptual. The review closes with a research agenda organised around dynamic, causal, human-governed and prospectively agentic graph-based decision intelligence, together with guidance for management, information systems and decision-support scholarship.

Keywords: graph theory; graph neural networks; explainable artificial intelligence; decision intelligence; organizational networks; dynamic capabilities; management research

1. INTRODUCTION

Contemporary organizations operate through webs of interdependence: employees collaborate across teams, firms compete and cooperate within alliance ecosystems, suppliers and customers form extended value chains, and financial institutions are linked through overlapping exposures. In each of these settings, the outcome experienced by one entity is frequently shaped less by its own attributes

than by its position within a wider structure of relationships. Conventional analytical approaches, which typically treat observations as independent and identically distributed records, can therefore underrepresent a defining feature of organizational life: relational dependence.

Graph theory supplies a formal vocabulary for representing entities as nodes and relationships as edges, and its core constructs — connectivity, centrality, paths, clusters and topology — have long informed operations research and social network analysis. Over the past two decades this foundation has been extended substantially. Network analysis and organizational network analysis brought relational thinking into management scholarship; graph analytics and graph databases enabled the systematic discovery of structural patterns in organizational and business data; knowledge graphs added semantic structure by relating entities and concepts; and, most recently, graph machine learning and graph neural networks have made it possible to learn predictive representations directly from relational structure. In parallel, explainable artificial intelligence has emerged to address the interpretability of increasingly complex model outputs, while decision intelligence has emerged as a discipline concerned with converting analytical insight into actionable choice.

These streams, however, have developed largely in isolation from one another. Classical graph-theoretic work emphasises structural description; graph-learning research emphasises predictive accuracy; XAI research emphasises the fidelity and comprehensibility of explanations; and decision-intelligence research emphasises the translation of insight into action. Few studies connect all four. The result is a literature that is comparatively mature in representing and predicting relational phenomena, but considerably less developed in explaining, deciding and adapting on the basis of that intelligence. This review addresses that fragmentation directly. Its central research problem is: how has graph-based analytical thinking evolved from structural relationship modelling into explainable, AI-enabled decision intelligence, and what does this evolution imply for management research and managerial decision-making?

The review does not claim originality for graph theory, graph neural networks or explainable AI considered individually; each is well established in its own literature. Its contribution lies instead in integrating these previously fragmented streams into a management-oriented capability architecture — Graph-Based Explainable Decision Intelligence (GEDI) — that connects relational sensing, graph analytics, prediction, explanation, decision integration and adaptive reconfiguration. In doing so, the review shifts the guiding question from "which graph algorithm performs best?" toward "how can organizations convert relational data into explainable and actionable intelligence?" — a question of organizational capability rather than algorithmic performance alone.

The remainder of the paper is organised as follows. Section 2 outlines the review methodology and research questions. Section 3 traces the evolution of relational intelligence from classical graph theory to graph representation learning. Section 4 examines applications of graph intelligence across management domains. Section 5 reviews the transition from graph learning to explainable graph AI. Section 6 addresses the translation of explanation into decision intelligence. Section 7 integrates the theoretical lenses used throughout the review. Section 8 presents the GEDI framework. Section 9 identifies critical research gaps, Section 10 sets out a future research agenda, Section 11 discusses emerging technological frontiers, and Section 12 concludes.

2. REVIEW METHODOLOGY

This review adopts a systematic, bibliometric-informed and integrative methodology, combining three complementary elements: a structured, reproducible literature search; an evidence-classification process that distinguishes mature from emerging findings; and a critical conceptual synthesis that integrates the resulting evidence into a novel management-oriented framework. This combination is

deliberate. A purely systematic review would document what has been studied without necessarily explaining why the field remains fragmented; a purely bibliometric review would describe publication patterns without producing a conceptual contribution; and a purely narrative review risks descriptive summary rather than critical synthesis. The integrative approach used here is intended to satisfy the methodological expectations of high-impact management, information-systems and decision-support outlets while preserving space for the conceptual contribution — the GEDI framework — that constitutes the review's principal output.

The review's conceptual boundary is explicitly management-oriented. It includes graph theory and graph algorithms relevant to decision problems, network analysis, organizational network analysis, business network analytics, knowledge graphs, graph machine learning, graph neural networks, explainable graph AI, explainable AI more broadly where it bears on relational or graph-based decisions, decision intelligence, decision support, human–AI decision-making, and managerial applications of graph-based methods. It excludes, or treats only briefly, purely mathematical graph-theory proofs without management relevance, low-level computational optimisation unrelated to decision-making, generic AI literature without a graph or network dimension, generic XAI literature that does not inform relational or graph-based decisions, and applications in which graphs function solely as a visualisation device.

Six research questions structure the review, moving deliberately from evolution to application, technological transformation, explanation, gaps and future architecture:

- RQ1. How has graph-based analytical thinking evolved from classical graph theory to contemporary decision intelligence in management research?
- RQ2. How have graph analytics and network-based methods contributed to understanding complex managerial relationships and decision problems?
- RQ3. How have graph machine learning and graph neural networks extended the predictive capabilities of traditional graph-based approaches?
- RQ4. How is explainable artificial intelligence transforming graph-based analytical outputs into interpretable and managerially meaningful insights?
- RQ5. What theoretical, methodological and contextual gaps constrain the development of explainable graph-based decision intelligence in management research?
- RQ6. What future research architecture can integrate graph theory, graph AI, explainability and managerial decision-making into a coherent, and potentially adaptive and agentic, decision-intelligence paradigm?

A full reproducible search protocol — including Scopus and Web of Science search strings, PRISMA screening criteria, and a bibliometric science-mapping plan covering co-citation, bibliographic coupling, keyword co-occurrence and thematic evolution analysis — is a necessary next step beyond the scope of the present conceptual synthesis and is proposed as a companion methodological exercise in Section 10. Consistent with that boundary, this review reports evidence maturity qualitatively (established, developing, emerging or frontier) rather than presenting invented publication counts, citation rankings or bibliometric clusters, which should be generated only once a verified database export is available.

2.1 Bibliometric and Scientometric Analysis Protocol

To meet the reproducibility expectations of a Q1 review, the systematic search results are additionally organised through a bibliometric-informed integrative methodology that converts the search output into a scientific map of the research landscape using descriptive bibliometrics, network scientometrics, science mapping and thematic-evolution analysis. Consistent with the principle stated in Section 2, this protocol does not fabricate bibliometric findings; every figure in this subsection is an illustrative template, populated with placeholder values for layout and interpretive-framework purposes only, and is explicitly marked for replacement once a verified Scopus and Web of Science export is available.

The protocol follows an eight-stage workflow: (1) export of Scopus and Web of Science metadata; (2) merging of datasets and removal of duplicates; (3) standardisation of author, affiliation and keyword names; (4) descriptive bibliometric analysis; (5) network analysis of authors, institutions and countries; (6) science mapping through co-citation, bibliographic coupling and co-word analysis; (7) thematic evolution analysis; and (8) interpretive synthesis integrated with the management literature reviewed in Sections 3 through 8. Metadata standardisation (stage 3) includes keyword harmonisation — for example, consolidating "graph neural network," "GNN" and "graph neural nets" into a single controlled term — which materially improves the quality of the resulting keyword clusters.

Reporting is organised around descriptive indicators (publication counts, citation counts, h-index, average citations per publication), publication-evolution trends, a journal analysis following Bradford's Law of scattering, an author-productivity analysis following Lotka's Law, and country- and institution-level collaboration mapping. Table 2.1 summarises the journal productivity zones expected under Bradford's Law across the four disciplinary clusters relevant to this review; the specific journals listed are illustrative of each cluster's scope rather than a confirmed core-zone ranking, which can only be established once citation data are retrieved.

Disciplinary cluster	Illustrative journals
Management & strategy	Strategic Management Journal; Organization Science; Journal of Business Research; Decision Support Systems
Information systems & AI	Information Systems Research; MIS Quarterly; Expert Systems with Applications; Information Sciences
Operations research & analytics	European Journal of Operational Research; Omega; Annals of Operations Research; Networks
Graph AI (technical core)	IEEE Transactions (various); ACM journals; machine-learning and artificial-intelligence journals

Figures A1 through A6 present the corresponding illustrative visualisation templates: a publication-evolution framework across six-year bands (Figure A1); an author-productivity template consistent with the concentration predicted by Lotka's Law, in which a small number of authors account for a disproportionate share of output (Figure A2); a management-domain publication-mapping template spanning the domains introduced in Section 4 (Figure A3); a country-collaboration network template (Figure A4); a co-citation cluster map illustrating how foundational social-network, graph-neural-network and graph-explainability sources might cluster once real co-citation data are analysed (Figure A5); and a keyword co-occurrence network illustrating the conceptual adjacency of terms central to this review's argument (Figure A6). Bibliographic coupling, institutional collaboration, a full thematic-evolution Sankey map, a centrality–density thematic map, trend-topic analysis and a three-fields plot

linking authors, keywords and journals follow the same template logic and are proposed as companion outputs once the underlying dataset is retrieved.




Figure A1. Illustrative publication-evolution framework (placeholder values; template for layout only).




Figure A2. Illustrative author-productivity template following Lotka's Law (placeholder values).





Figure A3. Illustrative management-domain publication mapping (placeholder values).


Figure A4. Illustrative Country Collaboration Network (schematic template)



Illustrative template — placeholder values for manuscript layout only.
To be replaced with verified Scopus / Web of Science data prior to submission.

Figure A4. Illustrative country-collaboration network (schematic template; not derived from retrieved co-authorship data).

Figure A5. Illustrative Co-Citation Cluster Map (schematic template)
Clusters: social-network foundations / GNN methods / graph-XAI methods



Illustrative template — placeholder values for manuscript layout only.
To be replaced with verified Scopus / Web of Science data prior to submission.

Figure A5. Illustrative co-citation cluster map (schematic template) spanning social-network foundations, GNN methods and graph-explainability methods.

Figure A6. Illustrative Keyword Co-Occurrence Network (schematic template)
Node size = illustrative keyword frequency (not real frequency data)



Illustrative template — placeholder values for manuscript layout only.
To be replaced with verified Scopus / Web of Science data prior to submission.

Figure A6. Illustrative keyword co-occurrence network (schematic template; node size is illustrative, not a real frequency count).

The interpretive value of this protocol lies not in the figures themselves but in the questions they are designed to answer once populated with real data: which sub-fields are intellectually mature versus emergent; which authors, journals and countries anchor the field's core; and how the field's thematic emphasis has moved over time relative to the evolutionary trajectory proposed in Section 3. Distinguishing mature research fronts from emerging ones in this way — rather than merely describing cluster colours or node positions — is what differentiates a Q1-level bibliometric synthesis from a purely descriptive bibliometric exercise, and this distinction is carried forward into the reproducible search protocol proposed in Section 10.6.

3. The Evolution of Relational Intelligence

The intellectual trajectory addressed by this review can be summarised as a progressive expansion in the purpose of graph-based analysis, moving from representing relationships, to analysing their structure, to predicting outcomes from them, to explaining those predictions, to informing decisions, to enabling action, and ultimately to supporting organizational learning: represent, analyse, predict, explain, decide, act, learn. Each stage in this progression is associated with a distinct dominant question and analytical capability, summarised in Table 1.

Evolutionary stage	Dominant question	Dominant capability
Classical graph theory	How are entities connected?	Structural representation
Network analysis	What patterns exist in relationships?	Structural analysis
Organizational network analysis	Who influences whom?	Relational intelligence
Graph analytics	Which structures and relationships matter?	Pattern discovery
Knowledge graphs	What entities and concepts are related?	Semantic intelligence
Graph machine learning	What is likely to happen?	Prediction
Graph neural networks	Can complex relational patterns be learned?	Representation learning
Explainable graph AI	Why did the model produce this result?	Interpretability
Decision intelligence	What should be done?	Decision support
Explainable decision intelligence	Why this decision, under which conditions, and with what consequences?	Actionable, accountable intelligence

The transition from explicit graph structure to learned relational representation is a pivotal, and often underappreciated, step in this trajectory. Perozzi, Al-Rfou and Skiena's DeepWalk demonstrated that latent representations of nodes could be learned from truncated random walks over a graph, translating ideas from language modelling into the relational domain and showing that a network could itself become a source of learned representation rather than merely an object of structural description. Kipf and Welling's graph convolutional network extended this logic by combining node attributes with neighbourhood structure within a scalable, semi-supervised learning framework — a combination of direct managerial relevance, since organizational outcomes (for example, an employee's performance) plausibly depend on both individual characteristics and relational context. Hamilton, Ying and Leskovec's GraphSAGE addressed a further limitation of early embedding methods by introducing an inductive framework capable of generating representations for previously unseen nodes through neighbourhood sampling and aggregation, a property of particular importance for organizational networks that are rarely static, as employees join and leave, suppliers change, and new relationships

form continuously. Veličković and colleagues' graph attention networks introduced masked self-attention over graph neighbourhoods, allowing models to weight neighbouring nodes differently — a mechanism that later becomes relevant to explainability, with the important caveat, developed in Section 5, that attention weights should not automatically be treated as faithful explanations.

Comprehensive surveys of graph neural networks — including taxonomies organised around recurrent, convolutional, autoencoder and spatial-temporal architectures, and reviews framing GNNs as models that capture dependency through message passing across structural and non-structural application scenarios — consolidate this foundation and confirm that representation learning and predictive graph modelling now constitute a mature area of the literature. What this literature does not by itself establish is that a graph neural network constitutes decision intelligence: prediction and decision are related but distinct constructs, a distinction that becomes central to the argument developed later in this review.

4. GRAPH INTELLIGENCE IN MANAGEMENT RESEARCH

Once relational structure can be represented, analysed and predicted, a wide range of managerial applications become tractable. Table 2 summarises the principal management domains in which graph-based intelligence has been applied, together with the typical nodes, edges, problems addressed and the character of the decisions such intelligence supports. These domains are presented as an organising taxonomy for the field rather than as a set of empirically verified claims; a verified domain-by-domain evidence matrix, populated from a systematic search of Scopus and Web of Science, is proposed as subsequent work.

Domain	Typical nodes / edges	Typical problem	Decision supported
Human resources	Employees / collaboration, communication	Influence, turnover, informal leadership	Talent management, intervention design
Strategy	Firms, business units / alliances, competition	Ecosystem and competitive positioning	Partnering and strategic choice
Supply chain	Firms, suppliers / transactions, dependencies	Disruption propagation, resilience	Sourcing and routing decisions
Marketing	Customers / interactions, influence	Community formation, recommendation	Targeting and segmentation
Finance	Firms, assets / ownership, exposure	Systemic risk, contagion	Portfolio and risk management
Innovation	Firms, researchers / collaboration	Knowledge diffusion	Collaboration and R&D decisions
Governance / risk	Actors, institutions / influence, control	Fraud and interconnected vulnerability	Governance and risk mitigation
Sustainability	Stakeholders / dependencies	ESG interdependence, supply-chain impact	Stakeholder intervention

A methodological issue that cuts across every domain in Table 2 deserves explicit emphasis: the graph is constructed, not simply found, in organizational data. An HR network, for instance, can be defined through email communication, meeting participation, project collaboration, formal reporting lines, advice-seeking ties or friendship, and these choices produce materially different graphs — a graph built from email exchange is not equivalent to one built from advice-seeking or collaboration ties,

even when drawn from the same organization. Graph construction therefore functions as a management-theory decision, not merely a technical pre-processing step, because the choice of nodes and edges encodes assumptions about which relationships are theoretically and practically relevant to the managerial question at hand. This issue should be treated as a first-order methodological concern in any study applying graph intelligence to organizational data, and it re-emerges below as a constraint on the transferability of graph foundation models across contexts.

Across these domains, a consistent pattern emerges: graph-based approaches consistently improve the visibility of relational dependencies that conventional, attribute-based analytics tend to obscure, yet the large majority of documented applications remain concentrated on structural description or predictive tasks — identifying influential employees, forecasting supplier disruption, predicting customer churn within a network — rather than on the subsequent step of managerial intervention. This concentration is not a criticism of the underlying technical work, which has genuine value in its own right; it is, however, the empirical basis for the translation gap that motivates the remainder of this review.

5. FROM GRAPH LEARNING TO EXPLAINABLE GRAPH AI

As graph neural networks have grown in sophistication, so has concern about their interpretability. Message-passing architectures aggregate information across multiple hops of a graph, making it difficult for a human decision-maker to trace why a particular node, edge or subgraph produced a particular prediction. This opacity is not a peripheral concern for management research: a supplier-risk score, a turnover-risk flag, or a fraud-detection alert generated by an uninterpretable graph model offers limited support for a manager who must justify, contest or act on that output.

Explainable graph AI has consequently become a substantial research field in its own right, encompassing post-hoc explanation techniques that vary in what they attribute predictions to. Feature-attribution methods identify influential node or edge features but often carry limited managerial meaning in isolation. Node- and edge-level explanations, and subgraph explanations that identify influential relational substructures, are more directly relational and therefore more interpretable to a manager reasoning about organizational structure. Counterfactual explanation techniques go further, exploring how a prediction would change if a specific relationship were strengthened, removed or reconfigured — for example, indicating that a supplier's predicted risk would fall if a particular dependency were diversified. This counterfactual framing moves the explanation closer to actionable managerial intervention than attribution-based methods, and for that reason occupies a privileged position within the GEDI architecture developed in Section 8. It is nonetheless essential to note that a counterfactual explanation is not automatically a causal explanation: demonstrating that altering a modelled relationship would change a prediction is not equivalent to demonstrating that altering the real relationship would change the real-world outcome, and the review treats this distinction as a standing caveat rather than a resolved matter.

Two further qualifications matter for the argument developed in this review. First, attention weights in graph attention networks, although intuitively appealing as a source of explanation, should not be treated as inherently faithful explanations of model behaviour; attention indicates where a model allocates weight, not necessarily why a prediction was made. Second, and more fundamentally, the literature establishes a well-supported technical relationship between graph neural networks and explanation techniques, but a considerably less well-supported relationship between explanation and improved managerial decision-making. Work situating explainable AI in relation to business decision-making has begun to organise the field around data, method and application dimensions and to argue for explainability's relevance to more informed decisions, which supports the conceptual bridge from

explanation toward decision but does not establish a universal causal claim that explanation improves decision quality.

This review therefore proposes the construct of managerial XAI as a useful extension of technical XAI. Where technical explainability asks whether an explanation is faithful to the underlying model, managerial XAI additionally asks whether an explanation is understandable, contextually relevant, feasible to act upon, and genuinely useful for a specific managerial decision. Framed formally, technical XAI corresponds to model understanding alone, whereas managerial XAI corresponds to model understanding combined with organizational context, decision objectives, operative constraints and actionability. This construct is offered as a review-derived conceptual extension rather than an established universal standard, and its empirical validation — for instance, through studies of how managers actually use graph-based explanations — constitutes one of the clearer opportunities identified in Section 10.

6. FROM EXPLANATION TO DECISION INTELLIGENCE

Graph neural networks have substantially expanded the capacity to learn from relational structures, but predictive capability does not by itself resolve the managerial problem of deciding whether, why and under what constraints an intervention should be made. This distinction becomes particularly consequential when graph predictions concern organizational relationships that are themselves endogenous to managerial action — a supplier network reshaped by a sourcing decision, or a collaboration network reshaped by a restructuring, is not a fixed backdrop against which predictions are made but an object that the decision itself will alter. Accordingly, this review positions graph learning as a predictive layer within a broader decision-intelligence architecture rather than treating predictive performance as the endpoint of graph analytics.

A structured review of artificial intelligence in management decision-making, drawing on a substantial corpus of peer-reviewed articles analysed bibliometrically, characterises AI–management decision research as an emerging field concentrated around forecasting, classification, knowledge management and information extraction — a concentration that corroborates the pattern observed within the graph-specific literature reviewed in Sections 3 to 5. Prediction, explanation, recommendation and decision are frequently, and problematically, treated as interchangeable in applied literature; this review instead retains them as conceptually distinct: prediction is not equivalent to explanation, which is not equivalent to recommendation, which is not equivalent to decision. A model can predict accurately without explaining its prediction; it can explain a prediction without generating an actionable recommendation; and a recommendation can be generated without constituting an organizational decision, which additionally requires weighing objectives, constraints, alternatives and consequences within a specific governance context.

The translation from graph-derived intelligence to managerial decision can therefore be represented as a multi-step process — from prediction, through explanation, to contextual interpretation, consideration of alternatives, evaluation against constraints, and finally decision — rather than a single computational step. It is precisely this intermediate translation process, largely absent from both the graph-learning and the XAI literatures reviewed above, that the GEDI framework in Section 8 is designed to make explicit.

7. THEORETICAL INTEGRATION

Rather than introducing a single new grand theory, this review integrates six complementary theoretical lenses, each contributing a distinct explanatory function to the overall argument. Social network theory explains how relationships, network position, structural holes and cohesion shape

organizational outcomes, and provides the foundational logic for treating relational structure as analytically consequential rather than incidental. Information processing theory explains why graph representations can help organizations manage relational information of a complexity that conventional, attribute-based models struggle to capture. Complexity theory, arguably the most consequential lens for the review's future-oriented contribution, treats organizations as complex adaptive systems in which relationships evolve dynamically and managerial interventions can generate non-linear, difficult-to-anticipate consequences — a framing that motivates the adaptive component of GEDI developed below. Dynamic capabilities theory supplies the sensing–seizing–reconfiguring logic through which organizations translate graph-based intelligence into sustained advantage, moving from sensing relational change, to interpreting complex patterns, to deciding, to reconfiguring. Decision theory contributes the vocabulary of choice under constraint that distinguishes recommendation from decision. Socio-technical systems theory, finally, frames the governance relationship between human judgement and computational intelligence that runs throughout the review's treatment of managerial XAI and human-governed decision integration.

These six lenses form a connected chain rather than six independent theoretical pillars: social network theory establishes relational structure; information processing theory establishes relational information; complexity theory establishes interdependence and non-linearity; decision theory establishes choice; dynamic capabilities theory establishes adaptation; and socio-technical systems theory establishes human–AI governance. GEDI, presented next, functions as the integrating capability architecture that connects these separately-developed theoretical contributions to the empirical evidence synthesised in Sections 3 through 6.

8. GRAPH-BASED EXPLAINABLE DECISION INTELLIGENCE (GEDI)

The evidence reviewed above supports a specific and important conclusion: the maturity of the literature is not uniform across the relational-intelligence pipeline. Representation and prediction are comparatively mature; explanation is developing rapidly; decision integration is considerably less developed; and adaptive, feedback-driven reconfiguration is the least developed of all. This asymmetry, summarised in Table 3, is itself one of the review's principal findings and directly motivates the definition of GEDI proposed below.

Pipeline stage	Dominant question	Evidence maturity
Graph theory	What is connected?	Mature
Network analysis	How are entities positioned?	Mature
Graph analytics	What patterns/dependencies exist?	Mature
Graph representation learning	What representations can be learned?	Mature
Graph neural networks	What can be predicted?	Mature
Explainable graph AI	Why did the model predict it?	Developing
Counterfactual graph AI	What would change the prediction?	Emerging
Decision intelligence	What should be done?	Developing
Adaptive GEDI	What happened after the intervention?	Limited
Agentic GEDI	Can the system continuously sense, reason and coordinate?	Frontier

Graph-Based Explainable Decision Intelligence (GEDI) is defined here as the integration of relational data representation, graph analytics, graph-based artificial intelligence, explainability mechanisms and managerial judgement to generate transparent, context-sensitive and actionable decisions within complex organizational systems. This definition is deliberately constructed to distinguish GEDI from each of its constituent literatures. Graph analytics asks what the network looks like; graph machine learning asks what is likely to happen; explainable graph AI asks why the model produced a given result; decision intelligence asks what action should be considered; and explainable decision intelligence — the outcome GEDI is designed to support — asks what action should be considered, why, based on which relationships and evidence, under what assumptions, and with what potential consequences.

GEDI comprises six interdependent components. Relational sensing captures organizational, market and stakeholder data as an interconnected system rather than as independent records. Graph analytics identifies structural properties — centrality, connectivity, community structure — that reveal organizational and strategic patterns. Predictive intelligence, realised through graph machine learning and graph neural networks, learns representations capable of forecasting node-, edge- and network-level outcomes. Explainable intelligence renders those predictions interpretable through attribution, subgraph and counterfactual explanation, subject to the managerial-XAI standard introduced in Section 5. Decision integration combines predictions and explanations with organizational objectives, constraints and alternatives to produce a genuine managerial decision rather than a bare recommendation. Adaptive reconfiguration, finally, closes the loop: because managerial decisions frequently alter the very network from which intelligence was generated — restructuring changes communication networks, supplier substitution changes supply networks, acquisitions change ownership networks, hiring changes organizational networks, alliances change ecosystem networks — GEDI treats the network as an object that is itself in motion, generating a feedback relationship between network state, intelligence, decision, action and subsequent network state.

This can be represented as an evolutionary sequence from organizational data through graph representation, graph metrics, network analysis and knowledge graphs, into graph analytics, graph machine learning and graph neural networks, into explainable graph AI (comprising explanation, prediction and counterfactual analysis), into decision intelligence, managerial judgement and managerial action, into organizational outcomes, and finally into network transformation, which feeds back into the originating organizational data. The central conceptual transition captured by this sequence is: graph structure, then analytical intelligence, then predictive intelligence, then explainable intelligence, then decision intelligence — with the adaptive loop distinguishing GEDI from a linear pipeline model.

It is important to state plainly what GEDI is, and is not, intended to claim. It is a conceptual organizational-capability framework, grounded in the six theoretical lenses integrated in Section 7, rather than a fully validated empirical theory. Its relational-sensing, graph-analytics and predictive-intelligence components rest on a mature evidence base; its explainable-intelligence component rests on a rapidly developing but not yet settled evidence base; and its decision-integration and adaptive-reconfiguration components are substantially more conceptual, reflecting the genuine state of the underlying literature rather than an aspiration to overstate it. This graduated evidentiary status is intentional and is carried through into the ten review-derived propositions and the research agenda presented next.

ΠΕΡΙΕΧΟΜΕΝΑ PROPOSITIONS

The following propositions are offered as review-derived conceptual statements intended to guide future empirical work, not as validated hypotheses. P1–P5 are relatively well aligned with existing evidence; P6–P10 are progressively more conceptual and are stated accordingly.

- P1. Relational complexity increases the potential value of graph-based intelligence.
- P2. Graph representations improve organizational visibility of relational dependencies.
- P3. Graph learning extends relational analysis toward predictive intelligence.
- P4. Increasing graph-model complexity increases the need for effective explanation.
- P5. The managerial value of graph explanations depends on their fidelity, comprehensibility, contextual relevance and actionability.
- P6. Counterfactual graph explanations may offer greater intervention relevance than purely attribution-based explanations.
- P7. Integrating graph intelligence with organizational objectives, constraints and alternatives increases its potential decision value.
- P8. Human–AI complementarity is likely to be particularly important for complex managerial decisions.
- P9. Managerial interventions can alter the network structure from which subsequent intelligence is generated.
- P10. Feedback from intervention outcomes can support the adaptive improvement of graph-based decision intelligence over time.

9. CRITICAL RESEARCH GAPS

Synthesising the evidence reviewed across Sections 3 through 8 yields a central research gap that is more specific, and more defensible, than a generic claim that graph theory in management has been insufficiently studied: although graph-based methods have matured from structural network analysis to representation learning and graph-neural prediction, and although explainability research has increasingly addressed the opacity of graph models, the management literature remains comparatively fragmented in explaining how relational predictions can be translated into context-sensitive, constraint-aware and actionable decisions. Even less developed is the feedback process through which managerial interventions alter the underlying network and generate new information for subsequent decisions.

Five subordinate gaps follow directly from this central proposition.

- Translation gap: the step connecting graph prediction and explanation to a bounded managerial decision — incorporating objectives, constraints and alternatives — is theoretically underspecified relative to the prediction and explanation stages that precede it.
- Construction gap: graph construction choices (which relationships to encode as edges, at what granularity, over what time window) are rarely treated as a management-theoretic decision, despite materially shaping downstream predictions and explanations.
- Dynamism gap: much of the literature retains a static view of organizational graphs, despite organizational relationships changing continuously; temporal and dynamic graph methods remain comparatively underused in management applications.

- Causality gap: counterfactual graph explanations demonstrate that altering a modelled relationship changes a prediction, but this is not equivalent to demonstrating that altering the underlying real-world relationship would change the real-world outcome; causal graph reasoning for management decisions remains an emerging rather than an established capability.
- Adaptation gap: research examining how managerial decisions reshape the networks from which they were derived, and how organizations might learn from that reshaping, is the least developed stage of the entire pipeline.

10. FUTURE RESEARCH AGENDA

Building on the gaps identified in Section 9, this review proposes a research agenda organised around five thrusts, together with the methodological groundwork required to pursue them.

10.1 Dynamic and temporal graph decision intelligence

Future work should extend GEDI beyond static network snapshots toward dynamic, temporal and multilayer graph representations capable of capturing how organizational relationships evolve. The governing research question shifts from whether a prediction was accurate at a given time to whether a decision remains valid as the underlying network subsequently changes — formally, from evaluating prediction at time t to evaluating the validity of a decision at time t against network state at time $t+1$.

10.2 Causal and counterfactual graph reasoning

Because counterfactual explanation is not automatically causal explanation, a priority for future research is the development and empirical validation of causal graph-learning methods capable of supporting genuine intervention reasoning in organizational settings, together with rigorous evaluation protocols for when a graph-based counterfactual can be treated as decision-relevant.

10.3 Managerial XAI and decision-quality evaluation

Given that the relationship between explanation and decision quality remains context-dependent and comparatively fragmented, this review recommends treating explanation-to-decision effectiveness as an empirical research agenda rather than an established causal relationship, and calls for behavioural and field studies that evaluate managerial-XAI explanations against the standard proposed in Section 5 — fidelity, comprehensibility, contextual relevance and actionability — rather than technical fidelity alone.

10.4 Human-governed and responsible GEDI

Socio-technical systems theory points toward a research stream examining how human governance should be structured around graph-based decision intelligence — including accountability, contestability and the appropriate allocation of judgement between managers and algorithmic systems — particularly in high-stakes domains such as workforce, financial-risk and supply-chain decisions.

10.5 Adaptive and prospectively agentic GEDI

The most forward-looking element of the agenda concerns adaptive GEDI, in which organizational decisions and their consequences for network structure are explicitly fed back into subsequent rounds of relational sensing and analysis, and the further, more speculative possibility of agentic GEDI — a future research architecture in which specialised computational agents continuously sense graph changes, analyse relational patterns, generate predictions, produce explanations, explore counterfactual interventions and support human-governed decisions. A plausible agentic sequence proceeds from a graph-data agent, through sensing, analytics and prediction agents, to explanation and counterfactual agents, to a decision agent operating under human governance, and on to action,

monitoring and graph update. This sequence is presented explicitly as a future research proposition rather than a claimed present-day management technology, and its organizational feasibility, governance requirements and risks are themselves a legitimate object of future study rather than an assumed benefit.

10.6 Methodological groundwork: toward a reproducible evidence base

Realising this agenda requires a reproducible systematic and bibliometric evidence base beyond the conceptual synthesis presented here. The recommended next step is a formal Scopus/Web of Science search protocol with separate search strings for graph theory and network analysis, graph analytics, graph representation learning, graph neural networks, explainable graph AI, decision intelligence, management applications, graph-large-language-model integration, graph foundation models and agentic graph intelligence; a PRISMA-compliant screening workflow; and a bibliometric analysis plan covering annual publication growth, source and author analysis, citation and co-citation analysis, bibliographic coupling, keyword co-occurrence, thematic mapping and a three-field plot. A useful working hypothesis for that analysis — to be tested against the data rather than assumed — is that the field's thematic emphasis has shifted progressively from network analysis, through graph analytics and graph learning, to graph neural networks, explainability and decision intelligence, with agentic intelligence as an emerging frontier; confirming, qualifying or disconfirming this trajectory empirically is left as the natural continuation of the present review.

11. EMERGING TECHNOLOGICAL FRONTIER

Several technical developments sit at the frontier of the pipeline summarised in Table 3 and warrant brief discussion, with the explicit caveat that each is classified as emerging or frontier evidence rather than established management practice. Graph transformer architectures combine transformer-style attention mechanisms with graph inductive biases and have been identified as a recent direction in graph learning, alongside continuing challenges of scalability, generalisation, robustness, interpretability and handling of dynamic graphs. Graph foundation models represent an attempt to extend the foundation-model paradigm — large-scale pretraining followed by adaptation — to graph-structured data, with structural alignment, heterogeneity, scalability and evaluation identified as open challenges. For management research, graph foundation models raise a distinctive question: whether relational knowledge can transfer across organizational contexts, for instance from a human-resources network to a supply-chain network, without losing domain-specific meaning. Given the construction-gap argument developed in Section 4 — that nodes, edges and structural patterns encode context-specific theoretical assumptions — this transferability question is substantive rather than purely technical, and graph transferability across organizational contexts is proposed here as a promising area of intersection between graph AI research and organizational theory.

Retrieval-augmented approaches that ground large language model outputs in graph-structured evidence (sometimes referred to in the literature as GraphRAG), and the graph-agent and agentic-GEDI possibilities discussed in Section 10.5, extend this frontier further still. All of these developments are classified in this review as emerging or frontier evidence, appropriate for a forward-looking research agenda but not yet appropriate as a basis for claims about current organizational practice.

CONCLUSION

This review has traced the evolution of graph-based analytical thinking from classical graph theory through network analysis, graph analytics, graph machine learning and graph neural networks, to explainable graph AI and, most tentatively, decision intelligence. Four contributions follow from this

synthesis. The first is evolutionary: the review reconstructs the intellectual trajectory from graph theory to graph-based decision intelligence as a coherent progression rather than a set of disconnected applications. The second is integrative: the review connects previously fragmented streams of network science, graph AI, explainable AI and management decision-making. The third is conceptual: the review proposes Graph-Based Explainable Decision Intelligence (GEDI) as an organizational capability comprising relational sensing, graph analytics, predictive intelligence, explainable intelligence, decision integration and adaptive reconfiguration. The fourth is forward-looking: the review develops a research agenda around dynamic, causal, human-governed and prospectively agentic graph-based decision intelligence.

The review's principal finding is an asymmetry in evidence maturity: representation and prediction are comparatively mature, explanation is developing rapidly, and the translation of graph-derived intelligence into constrained managerial decisions — together with the adaptive feedback between decisions and the networks they reshape — remains substantially underdeveloped. This asymmetry, rather than the technical sophistication of any individual graph algorithm, constitutes the field's central research opportunity. Accordingly, this review does not claim novelty for graph theory, graph machine learning or explainable artificial intelligence considered individually. Its principal contribution lies in integrating these previously fragmented streams into a management-oriented architecture that explains how relational representations can evolve into predictive, explainable, actionable and adaptive decision intelligence. The GEDI framework proposed here is intended to redirect scholarly and managerial attention from the performance of individual graph algorithms toward the organizational capability required to convert relational intelligence into contextually grounded managerial action and sustained organizational learning.

REFERENCES

1. Coussement, K., Abedin, M. Z., Kraus, M., Maldonado, S., & Topuz, K. (2024). Explainable AI for enhanced decision-making. *Decision Support Systems*, 184, 114276. <https://doi.org/10.1016/j.dss.2024.114276>
2. Hamilton, W. L., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 1024–1034.
3. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*.
4. Oppioli, M., Sousa, M. J., Sousa, M., & de Nuccio, E. (2024). The role of artificial intelligence for management decision: A structured literature review. *Management Decision*, 63(10), 3262–3280. <https://doi.org/10.1108/MD-08-2023-1331>
5. Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). DeepWalk: Online learning of social representations. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 701–710.
6. Post-hoc explainability of graph neural networks: A comprehensive survey. (2026). *Information Sciences*, 740, 123202. <https://doi.org/10.1016/j.ins.2026.123202>
7. Shehzad, A., Xia, F., Abid, S., Peng, C., Yu, S., Zhang, D., & Verspoor, K. (2024). Graph transformers: A survey. *arXiv*. <https://arxiv.org/abs/2407.09777>
8. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. *International Conference on Learning Representations*.
9. Wang, Z., Liu, Z., Ma, T., et al. (2025). Graph foundation models: A comprehensive survey. *arXiv*. <https://arxiv.org/abs/2505.15116>
10. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
11. Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., & Sun, M. (2021). Graph neural networks: A review of methods and applications. *AI Open*, 1, 57–81. <https://doi.org/10.1016/j.aiopen.2021.01.001>