



Algorithmic Resilience Memory: Designing Agentic AI Systems for Organizational Learning and Climate-Crisis Adaptation

Dr. Harsha Sammangi¹, Aditya Jagatha², Navyasri Maddukuri³

¹Assistant Professor, MIS & Business Analytics, Augustana University

^{2,3}Graduate Researcher, Dakota State University, Madison, South Dakota

Abstract: Climate disruption has become a persistent organizational condition rather than an episodic event, yet most information systems designed to support organizational resilience treat each disruption as an isolated incident. Existing digital resilience platforms, disaster recovery systems, and AI-driven decision support tools lack the capacity to accumulate, encode, and reuse organizational knowledge across successive climate-related crises. This paper introduces Algorithmic Resilience Memory (ARM), a novel IS construct defined as an AI-enabled organizational capability through which agentic AI systems sense climate-related disruptions, encode prior organizational responses, preserve decision rationale, generate contextually adaptive recommendations, and reconfigure future actions through structured outcome feedback. Drawing on Design Science Research (DSR), we propose and develop an Agentic AI-Based Algorithmic Resilience Memory Framework as the primary artifact. The framework integrates six interdependent functional layers—environmental sensing, knowledge encoding, agentic AI reasoning, explainable decision support, human governance, and adaptive learning—grounded in organizational memory theory, dynamic capabilities theory, sociotechnical systems theory, and responsible AI governance principles. We demonstrate the framework through a detailed scenario involving a regional flood disrupting a manufacturing firm's supply chain operations and evaluate its utility using scenario-based assessment and expert panel validation. The paper makes three primary contributions: it introduces ARM as a theoretically grounded IS construct that advances digital resilience research; it offers a design-science artifact that organizations can adopt for AI-driven climate-crisis adaptation; and it establishes design principles for building agentic AI systems capable of institutional learning across repeated climate disruptions.

Keywords: Algorithmic Resilience Memory; Agentic AI; Digital Resilience; Organizational Learning; Design Science Research; Climate-Crisis Adaptation; Explainable AI Governance

I. INTRODUCTION

The acceleration of climate-related disruptions—from prolonged droughts and catastrophic floods to cascading supply-chain failures and infrastructure collapses—has fundamentally altered the operating environment for organizations worldwide. Between 2000 and 2023, the economic losses attributable to extreme weather events exceeded \$3.8 trillion globally (UNDRR, 2023), and these losses disproportionately affect organizations that lack the informational infrastructure to sense, interpret, and respond

to compounding threats. What distinguishes the climate crisis from other forms of organizational risk is not its magnitude alone but its iterative, non-stationary character: floods will recur, heat waves will intensify, and supply disruptions driven by carbon-regulation pressure or resource scarcity will compound over time. Organizations are, in this sense, not preparing for a single catastrophe but managing an ongoing adaptive challenge.

Information systems research has long engaged with questions of organizational resilience, disaster recovery, and business continuity (Boin & McConnell, 2007;

Herbane et al., 2004). More recently, scholarship on digital resilience has examined how information technologies help organizations detect, absorb, and recover from disruptive shocks (Linkov & Trump, 2019; Bhamra et al., 2011). Concurrently, advances in artificial intelligence—particularly the emergence of agentic AI architectures capable of autonomous goal-directed behavior, multi-step reasoning, and adaptive learning—have opened new possibilities for AI-mediated organizational decision support (Buxmann & Schmidt, 2021; Wang et al., 2023). However, the intersection of these developments remains undertheorized in IS research. Current systems tend to respond to disruptions reactively and episodically; they do not accumulate structured organizational memory across successive disruptions, and they do not learn in ways that improve the quality of future climate-crisis response.

This gap is not merely technical. It reflects a deeper conceptual limitation: existing IS constructs—organizational memory systems, business intelligence platforms, disaster recovery frameworks, and AI decision support tools—were not designed with the iterative, learning-oriented demands of repeated climate-crisis management in mind. Organizational memory research (Walsh & Ungson, 1991; Stein, 1995) addresses how organizations store and retrieve knowledge but does not theorize how this memory should be structured to support AI-mediated learning across sequential disruptions. Dynamic capabilities theory (Teece et al., 1997) provides insight into organizational reconfiguration but does not specify the IS artifacts through which such reconfiguration occurs. The result is a structural research gap: no integrated IS construct currently captures the memory, learning, and reconfiguration logic that organizations need to manage the climate crisis as a sustained, evolving condition.

This paper responds to that gap by introducing Algorithmic Resilience Memory (ARM) as a new IS construct. ARM is defined as an AI-enabled organizational capability through which agentic AI systems sense climate-related disruptions, encode the structure and context of those disruptions, preserve the rationale underlying prior organizational decisions, generate adaptive recommendations for current events, and reconfigure future response logic through outcome-based feedback. ARM is not a metaphor for system resilience; it is a theoretically grounded, design-oriented construct that specifies the functional requirements, architectural components, and governance conditions under which AI-enabled climate-

crisis memory and learning can be instantiated in organizational settings.

The paper addresses four research questions. RQ1: How can agentic AI systems be designed to capture, retain, and reuse organizational knowledge from repeated climate disruptions? RQ2: What design principles govern the development of Algorithmic Resilience Memory in organizations? RQ3: How does ARM enhance organizational sensing, learning, and adaptive response capability during climate-crisis conditions? RQ4: How can explainability and human accountability be embedded into agentic AI-based resilience systems without undermining their adaptive learning capacity?

We employ Design Science Research (DSR) as our guiding methodology (Hevner et al., 2004; Peffers et al., 2007), developing an Agentic AI-Based Algorithmic Resilience Memory Framework as the primary artifact. The framework is evaluated through scenario-based demonstration and structured expert panel validation. The paper's theoretical contributions include the conceptualization of ARM as a new IS construct, a set of design principles for agentic AI resilience systems, and an extension of digital resilience theory to encompass AI-mediated organizational memory. Practical contributions include a reusable framework architecture and governance guidelines for organizations deploying AI in climate-crisis settings.

The remainder of the paper is organized as follows. Section 2 reviews relevant literature across six thematic streams and identifies the research gap. Section 3 develops the ARM construct and compares it with related IS constructs. Section 4 presents the DSR methodology. Section 5 details the artifact design. Section 6 demonstrates the framework through a realistic scenario. Section 7 presents the evaluation plan. Section 8 discusses theoretical and practical implications. Sections 9 and 10 address contributions, limitations, and future research directions. Section 11 concludes the paper.

II. LITERATURE REVIEW

Digital Resilience And The Climate Crisis

Digital resilience refers to an organization's capacity to use information technologies to anticipate, absorb, adapt to, and transform in the face of disruptive events (Linkov & Trump, 2019; Bhamra et al., 2011). Early IS research in this

domain focused primarily on disaster recovery planning, business continuity management, and information system recovery (Herbane et al., 2004). More recent scholarship has broadened this focus to address systemic and recurring disruptions, including those attributable to climate change, geopolitical instability, and pandemic-related shocks (Wieland & Wallenburg, 2013; Hohenstein et al., 2015).

However, existing work on digital resilience remains largely reactive in orientation. Systems are designed to restore prior operational states rather than to learn from disruption sequences. The climate crisis complicates this model fundamentally. Recurring floods, extreme heat events, and chronic resource scarcity are not discrete recoverable shocks; they are persistent conditions that require organizations to develop cumulative adaptive capacity over time. Lengnick-Hall and Beck (2005) argue that resilience involves cognitive and behavioral transformation, not merely operational recovery. IS research has not yet produced an artifact or framework that enables this transformation to be mediated by AI systems capable of institutional memory.

Agentic Ai In Organizational Decision-Making

Agentic AI refers to artificial intelligence systems that pursue multi-step goals, make sequential decisions, and adapt their behavior based on environmental feedback with minimal human intervention (Wang et al., 2023; Xi et al., 2023). Unlike classical decision support systems, which require explicit human queries and return single-step outputs, agentic AI architectures execute plans, use tools autonomously, and refine their strategies based on observed outcomes. Recent advances in large language model-based agents—including ReAct-style reasoning and multi-agent coordination systems—have demonstrated that agentic AI can perform complex organizational tasks such as supply-chain monitoring, anomaly detection, and multi-source data synthesis (Yao et al., 2023; Shinn et al., 2023).

In IS research, the emergence of agentic AI introduces new questions about the design of sociotechnical systems. Buxmann and Schmidt (2021) argue that AI agents alter the traditional human-in-the-loop model of decision support, creating new governance requirements. Baptista et al. (2020) examine how algorithmic systems reconfigure organizational decision-making processes and identify risks of opacity and accountability loss. Despite this growing body of work, agentic AI has not been theorized in relation to organizational resilience or climate-crisis

adaptation. The conditions under which agentic AI systems should operate in high-stakes, high-uncertainty environments—and the mechanisms through which they should store and reuse institutional learning—remain unspecified in the IS literature.

Organizational Memory And Learning In Information Systems

Organizational memory (OM) refers to stored information from an organization's history that can be retrieved and applied to present decisions (Walsh & Ungson, 1991). In IS research, organizational memory systems (OMS) are defined as information systems designed to capture, store, and retrieve organizational knowledge (Stein & Zwass, 1995). These systems encompass a range of formats—case-based reasoning repositories, lessons-learned databases, enterprise knowledge management systems—and have been studied extensively in contexts ranging from project management to healthcare (Jennex & Olfman, 2005; Markus, 2001).

The organizational learning literature, building on Argyris and Schön's (1978) foundational work, distinguishes between single-loop learning—adaptive responses within existing frameworks—and double-loop learning—reconceptualizing the frameworks themselves. Crossan et al.'s (1999) 4I model captures how learning involves intuiting, interpreting, integrating, and institutionalizing knowledge across individual and organizational levels. These frameworks provide valuable conceptual grounding for ARM, but they were not developed with AI-mediated memory or climate-disruption sequences in mind. The technical and design requirements for an AI system that enacts organizational memory through agentic reasoning—rather than passive retrieval—remain unaddressed.

Decision Support Systems For Crisis And Climate

Adaptation

Decision support systems (DSS) have been applied to crisis management and disaster response for several decades (Power, 2002; Turban et al., 2011). Climate-specific DSS applications include early warning systems for floods and droughts, carbon accounting dashboards, climate risk scoring engines, and supply-chain vulnerability tools. Gallopín (2006) examines how climate-adaptive decision support must integrate temporal dynamics, uncertainty, and multi-stakeholder values in ways that conventional DSS

architectures do not accommodate. Arnott and Pervan (2014) identify a gap in DSS research around adaptive systems that learn from use rather than being statically configured.

Recent work in IS has begun to address AI-enhanced decision support for operational risk (Shim et al., 2002; Bhatt & Zaveri, 2002), but climate crisis adaptation as an ongoing organizational condition remains marginal in this literature. Existing climate DSS tools operate at the level of environmental monitoring or scenario planning; they do not function as learning agents that remember past organizational decisions, evaluate their outcomes, and update their recommendation logic accordingly. ARM addresses this gap directly by redefining the DSS architecture for a world in which climate disruption is iterative rather than exceptional.

Explainable And Responsible Ai Governance

Explainable AI (XAI) refers to AI systems designed to make their reasoning processes interpretable and auditable by human users (Gunning et al., 2019; Arrieta et al., 2020). In high-stakes organizational settings, explainability is not merely a technical feature but a governance requirement: AI-generated recommendations that cannot be interrogated by human decision-makers undermine accountability and introduce systemic risks (Doshi-Velez & Kim, 2017). Responsible AI governance frameworks—including those of the OECD (2019), IEEE (2019), and the EU Artificial

Intelligence Act (2024)—specify requirements for transparency, fairness, auditability, and human oversight.

In IS research, responsible AI is an emerging field that examines how AI systems should be designed, governed, and institutionalized in organizational contexts (Dignum, 2019; Floridi et al., 2018). For climate-crisis resilience, the stakes of AI-generated recommendations are high: an AI system that recommends an ineffective supply-chain rerouting or an inappropriate resource allocation during a flood could cause significant harm. ARM explicitly incorporates XAI and governance principles into its design, requiring that every AI recommendation be accompanied by an explanation of its reasoning, a reference to the precedent decisions that informed it, and a mechanism for human override and accountability logging.

Research Gap And Theoretical Positioning

The preceding review reveals that while substantial bodies of scholarship exist on organizational resilience, agentic AI, organizational memory, decision support, and responsible AI governance, no prior research has integrated these streams into a unified IS construct that addresses the iterative, AI-mediated, memory-dependent demands of climate-crisis adaptation. Table 1 summarizes the key literature streams, their contributions, and their limitations relative to the ARM construct.

Table 1. Summary of Key Literature Streams and Research Gaps

Literature Stream	Key Contributions	Representative Sources	Limitations Relative to ARM
Digital Resilience	Recovery; continuity; adaptive capacity; technology inventories	Linkov & Trump (2019); Bhamra et al. (2011); Herbane et al. (2004)	Reactive and episodic; no AI-mediated iterative learning; no cross-disruption memory accumulation
Agentic AI	Autonomous multi-step reasoning; tool use; adaptive planning	Wang et al. (2023); Yao et al. (2023); Shinn et al. (2023)	Not applied to organizational resilience; no institutional memory architecture theorized
Organizational Memory	Knowledge storage; OMS design; lessons-learned systems	Walsh & Ungson (1991); Stein & Zwass (1995); Markus (2001)	Passive retrieval; not AI-mediated; not designed for climate-disruption specificity

Literature Stream	Key Contributions	Representative Sources	Limitations Relative to ARM
Decision Support Systems	Crisis DSS; climate early warning; scenario planning tools	Power (2002); Arnott & Pervan (2014); Gallopín (2006)	Static configuration; no learning from successive events; no rationale preservation
Responsible / XAI	Governance; auditability; transparency; human oversight	Arrieta et al. (2020); Floridi et al. (2018); IEEE (2019)	Governance principles not operationalized within resilience framework designs
Dynamic Capabilities	Sensing; seizing; reconfiguring organizational resources	Teece et al. (1997); Lengnick-Hall & Beck (2005)	Highly abstract; no IS artifact specified; no AI operationalization provided

ARM is positioned precisely at the intersection of these streams: it operationalizes dynamic capabilities through agentic AI; it extends organizational memory theory into AI-mediated, climate-specific settings; and it embeds responsible AI governance into a design artifact intended for real-world organizational deployment. This positioning makes ARM a genuinely new IS construct rather than a recombination of existing ones.

III. THEORETICAL BACKGROUND AND CONCEPT DEVELOPMENT

A. Defining Algorithmic Resilience Memory

Algorithmic Resilience Memory (ARM) is defined as an AI-enabled organizational capability through which agentic AI systems sense climate-related disruptions, encode the context and structure of organizational responses, preserve the rationale underlying managerial decisions, generate contextually adaptive action recommendations, and reconfigure future response logic through structured outcome feedback. The term 'algorithmic' denotes that memory is processed, indexed, and retrieved through AI-driven computation rather than static database queries. The term 'resilience' signifies functional orientation toward improving an organization's capacity to absorb disruption and adapt over time. The term 'memory' emphasizes continuity: ARM is designed to connect current disruptions with prior events, enabling

cross-event learning that conventional IT systems do not support.

ARM operates across six theoretically grounded dimensions that together constitute its functional architecture. Figure 1 presents the conceptual model.



Figure 1. Conceptual Model of Algorithmic Resilience Memory

Note. The six dimensions of ARM are arranged as a closed learning cycle. Each stage produces structured information outputs consumed by the next stage, with feedback arrows from Outcome Assessment back to all prior stages enabling continuous institutional learning. Theoretical anchors are shown at each dimension.

Dimension 1 — Climate-Risk Sensing: ARM continuously monitors environmental, operational, geopolitical, and regulatory data streams to detect early warning signals of climate-related disruptions. This dimension draws on information processing theory (Galbraith, 1974), which frames organizational adaptation as dependent on the quality and timeliness of environmental information intake.

Dimension 2 — Disruption-Event Encoding: When a disruption signal crosses a threshold, ARM encodes the event's structural properties—location, severity, affected assets, temporal sequence, and organizational context—into a structured knowledge representation. This encoding logic draws on case-based reasoning (Kolodner, 1993) and organizational memory theory (Walsh & Ungson, 1991).

Dimension 3 — Decision-Rationale Preservation: ARM captures not only what decisions were made during a disruption but why—the managerial reasoning, contextual constraints, available alternatives, and trade-offs considered. Preserving rationale is essential because contextual differences between disruptions mean that the same decision may be appropriate in one setting and counterproductive in another.

Dimension 4 — AI-Supported Response Recommendation: Drawing on accumulated encoded memory, ARM's agentic AI layer compares current disruption characteristics with precedent events and generates ranked recommendations tailored to the current organizational context. Recommendations are accompanied by confidence levels, precedent citations, and explanatory narratives.

Dimension 5 — Feedback-Based Organizational Learning: After a disruption resolves, ARM collects structured feedback on the outcomes of recommended and implemented actions, updating its knowledge base and refining its recommendation algorithms. This dimension operationalizes double-loop learning (Argyris & Schön, 1978) through AI-mediated feedback cycles.

Dimension 6 — Adaptive Reconfiguration: ARM adjusts its sensing thresholds, encoding schemas, and recommendation weights based on cumulative learning, progressively improving its fit with the evolving risk landscape. This dimension draws on dynamic capabilities theory (Teece et al., 1997) as operationalized through IS artifact design.

B. Comparison With Related IS Constructs

Table 2 presents a systematic comparison of ARM with five related IS constructs along eight analytical dimensions. The comparison demonstrates that ARM is not a repackaging of prior constructs but a genuinely new integration that addresses the limitations of each antecedent construct.

Table 2. Comparison of Algorithmic Resilience Memory with Related IS Constructs

Dimension	Org. Memory Systems	Business Intelligence	Digital Resilience Platforms	AI Decision Support	ARM — This Study
Primary Function	Store And Retrieve Organizational Knowledge	Analyze Structured Historical Data	Restore Operations After Disruption	Support Individual Decision Episodes	Sense, Encode, Reason, Recommend, Learn, Reconfigure Across Successive Events
Climate-Crisis Focus	None	Indirect (Reporting)	Partial — Recovery Only	Emerging	Central Design And Functional Goal

Dimension	Org. Memory Systems	Business Intelligence	Digital Resilience Platforms	AI Decision Support	ARM — This Study
Cross-Event Learning	Limited — Static Retrieval	None	None	Limited	Core Architectural Mechanism
Agentic AI Integration	None	None	None	Emerging In Literature	Fully Integrated Across All Layers
Rationale Preservation	Partial — Outcomes Only	None	None	Rare	Required Design Feature — Structured Decision Logs
Explainability / Governance	None	Limited Dashboards	None	Emerging	Embedded Dedicated Governance Layer
Adaptive Reconfiguration	None	None	None	Partial	Outcome-Driven Parameter Updating
Theoretical Grounding	OM Theory; KM Theory	Data Warehousing; Analytics	Crisis Management; Continuity	DSS Theory; Human Factors	OM Theory + Dynamic Capabilities + Sociotechnical + Responsible AI

IV. RESEARCH METHODOLOGY: DESIGN SCIENCE RESEARCH APPROACH

Justification For Design Science Research

Design Science Research is adopted as the governing methodology because ARM is fundamentally a prescriptive construct: its value lies not in explaining a phenomenon but in enabling organizational action through a well-designed artifact. Hevner et al. (2004) established that DSR addresses IS problems through the creation and evaluation of innovative artifacts—constructs, models, methods, or instantiations—that extend the knowledge base of a field. Peffers et al.'s (2007) DSR methodology framework provides a structured six-activity process guiding artifact development from problem identification through design, demonstration, evaluation, and theoretical communication. Gregor and Hevner (2013) further articulate that DSR contributions can take the form of new constructs, design theories, or design principles—all three of which this paper

produces. Figure 3 presents the DSR process followed in this study.



Figure 3. Design Science Research Process Followed in This Study

Note. The six-stage process follows Peffers et al. (2007). A design refinement feedback arrow runs from Stage 5 (Evaluation) back to Stage 3 (Artifact Design), reflecting the iterative character of DSR. Stages 1–6 are sequentially executed, with Stage 4 providing the demonstration evidence for Stage 5 evaluation.

The choice of DSR over interpretive or positivist methodologies is deliberate. ARM does not emerge from observed phenomena in existing organizations; it is a

designed response to a theorized need. The climate-crisis adaptation context is sufficiently novel that empirical observation of organizations deploying ARM-like systems is not currently feasible at scale. DSR permits rigorous scientific contribution without requiring longitudinal empirical observation, provided the artifact is theoretically grounded, rigorously designed, and systematically evaluated (Hevner et al., 2004). The scenario-based evaluation and expert panel validation employed here are well-established DSR evaluation strategies (Venable et al., 2016).

Problem Identification And Motivation

The problem motivating this research can be stated precisely: organizations facing repeated climate-related disruptions lack an AI-mediated information system capable of accumulating structured memory from prior disruptions, preserving decision rationale, and generating learning-informed recommendations for future events. Existing systems—DSS platforms, ERP modules, crisis management tools—address disruptions episodically. They do not connect prior organizational responses to current events; they do not preserve the reasoning underlying past decisions in machine-interpretable form; and they do not update their recommendation logic based on observed outcomes. The result is that organizational learning from climate disruptions occurs informally through human recollection, with all the vulnerability to organizational turnover, cognitive bias, and incompleteness that this implies.

Objectives Of The Solution

The ARM artifact should achieve seven specific objectives: (1) provide real-time climate-risk sensing through integration with multiple environmental and operational data streams; (2) encode disruption events and organizational responses in machine-interpretable knowledge formats that support future retrieval and comparison; (3) preserve managerial decision rationale in a structured, queryable form; (4) generate ranked, context-sensitive recommendations drawing on precedent events in the ARM knowledge base; (5) present recommendations with transparent, auditable explanations supporting human oversight; (6) collect and process post-disruption outcome data to update recommendation logic; and (7) comply with responsible AI governance requirements including

explainability, auditability, human override capacity, and bias monitoring.

Artifact Design And Development

The ARM artifact is designed as a six-layer framework architecture integrating sensing infrastructure, knowledge representation, agentic AI reasoning, explainable decision interfaces, governance controls, and adaptive learning mechanisms. Each layer corresponds to one of ARM's six functional dimensions and is grounded in specific IS theories and design principles detailed in Section 5. The framework is designed at the conceptual and logical levels, drawing on established system architecture principles and design science conventions (March & Smith, 1995; Gregor & Hevner, 2013).

Demonstration Strategy

The artifact is demonstrated through a detailed scenario depicting a regional flood that disrupts the supply chain of a mid-sized manufacturing organization. The scenario traces ARM's operation from initial sensing through knowledge encoding, AI-generated recommendations, human review, decision implementation, and post-disruption learning. The scenario is constructed to be industry-grounded and sufficiently detailed to illustrate how each layer of the framework functions under real-world conditions. Scenario-based demonstration is an accepted DSR strategy for artifacts not yet instantiated in live organizational settings (Peffer et al., 2007).

Evaluation Strategy

The artifact is evaluated using a two-stage approach. In the first stage, the flood disruption scenario and a complementary extreme heat scenario are used as structured evaluation instruments. Expert participants—drawn from IS research, operations management, climate risk management, and AI governance—assess each scenario along seven dimensions: usefulness, explainability, decision quality, organizational learning support, governance accountability, system completeness, and organizational fit. In the second stage, structured expert panel interviews provide qualitative elaboration on evaluation ratings and surface design implications not captured in the quantitative instrument. Evaluation criteria and measurement items are detailed in Section 7.

Research Rigor, Validity, And Ethical Considerations

DSR rigor is maintained through three mechanisms: theoretical grounding of all design decisions in established IS theory; systematic alignment between research questions, design objectives, artifact layers, and evaluation criteria; and transparent reporting of design choices and their justifications (Hevner et al., 2004; Venable et al., 2016). Construct validity is addressed by anchoring ARM's dimensions in established organizational memory and dynamic capabilities theory. External validity is addressed by designing the artifact at a level of abstraction supporting application across multiple industries and disruption types. Ethical considerations include privacy protections for organizational data encoded in the ARM knowledge base, bias monitoring requirements embedded in the artifact design, and explicit human override mechanisms preventing unchecked autonomous AI action in high-stakes contexts.

V. ARTIFACT DESIGN: AGENTIC AI-BASED ALGORITHMIC RESILIENCE MEMORY FRAMEWORK

The Agentic AI-Based Algorithmic Resilience Memory Framework (ARM Framework) is the primary artifact produced by this research. It is a six-layer architecture in which each layer performs distinct computational and organizational roles and interfaces with adjacent layers through defined information flows. Figure 2 presents the full framework architecture.



Figure 2. Agentic AI-Based Algorithmic Resilience Memory Framework

Note. Each layer is numbered 1 (bottom) through 6 (top), reflecting the upward flow of information from raw sensing data to adaptive feedback. Bidirectional arrows between adjacent layers indicate that information flows in both directions. The dashed adaptive feedback loop on the right side represents the return path from Layer 6 (learning outcomes) to all lower layers for system reconfiguration.

Layer 1: Environmental And Organizational Data

Sensing

The sensing layer is responsible for continuous ingestion and pre-processing of heterogeneous data streams relevant to climate-related risk. These include satellite-based weather and environmental monitoring feeds, IoT sensor data from physical assets (factory equipment, logistics infrastructure, energy systems), supply-chain monitoring platforms, regulatory and carbon-reporting databases, hydrological and meteorological services, and macroeconomic indices sensitive to climate-driven disruption. Data fusion algorithms harmonize these streams into a unified organizational risk signal, applying anomaly detection and threshold-based alerting to identify emergent disruption conditions.

The layer implements edge-computing preprocessing to reduce latency in early warning detection, consistent with information processing theory's emphasis on reducing environmental uncertainty through rapid signal acquisition (Galbraith, 1974). Multi-modal ingestion supports both structured numerical data and unstructured text—regulatory announcements, news feeds—processed through natural language understanding modules.

Layer 2: Event Memory And Knowledge Encoding

When a disruption signal is validated, the encoding layer transforms it into a structured knowledge record anchored to a domain-specific disruption ontology. The ontology defines entity classes (flood event, heat wave, supply disruption), relationship types (caused-by, affects, mitigated-by), and attribute schemas enabling consistent cross-event comparison. Each event record captures the disruption's physical characteristics, organizational impacts, response decisions, decision rationale, resources deployed, and timeline.

Decision-rationale preservation is implemented through a structured decision log capturing the goal being pursued, alternatives considered, constraints in force, information

available at decision time, and the managerial authority responsible. This log is inspired by issue-based information systems (Conklin & Begeman, 1988) and organizational learning frameworks distinguishing between what was decided and why (Argyris & Schön, 1978). The knowledge base is stored in a graph database optimized for relationship-based retrieval, enabling the AI reasoning layer to perform semantic similarity matching across disruption histories.

Layer 3: Agentic AI Reasoning And Recommendation

The agentic AI layer is the computational core of ARM, deploying a multi-agent architecture in which specialized agents perform distinct functions: a monitoring agent continuously interrogates the sensing layer; a similarity agent identifies precedent events structurally analogous to the current disruption; a planning agent generates ranked action recommendations combining precedent-derived strategies with current organizational constraints; and a coordination agent manages inter-agent communication and resolves conflicting recommendations. Agents operate through a reasoning loop informed by retrieval-augmented generation and chain-of-thought reasoning, enabling nuanced contextualization rather than simple template matching.

Recommendations are produced as structured outputs comprising: a ranked list of candidate responses; a confidence score reflecting precedent match quality and evidential reliability; a set of precedent citations identifying the specific prior events informing the recommendation; and a natural language explanation of the reasoning chain. The planning agent explicitly factors organizational resource availability, current disruption severity, and strategic priorities—information drawn from an organizational context module updated in real time by the sensing layer.

Layer 4: Explainable Decision Support

The explainable decision support layer translates agentic AI recommendation outputs into formats interpretable and actionable for human decision-makers. SHAP-based feature attribution identifies which aspects of the current disruption and which precedent characteristics drove each recommendation. Explanations are presented at three levels: an executive summary for senior managers requiring rapid situational awareness; an analytical narrative for operational managers evaluating specific

response options; and a technical audit report for IT governance functions and post-incident review.

Risk visualization dashboards present disruption maps, impact projections, resource allocation scenarios, and comparison matrices between the current event and retrieved precedents. These tools are designed per evidence-based decision support principles (Montibeller & Winterfeldt, 2015) and XAI literature requirements for transparency and comprehensibility (Arrieta et al., 2020). Scenario comparison tools allow managers to explore counterfactuals derived from the historical outcome archive.

Layer 5: Human Oversight And Governance

ARM is designed as a human-in-the-loop system. The governance layer implements approval workflows requiring human authorization before any AI-recommended action with significant resource, financial, or safety implications is executed. Override controls allow managers to reject or modify recommendations, with modification rationale captured for learning. Bias monitoring modules assess whether recommendation patterns exhibit systematic disparities across event types, geographical regions, or organizational units.

Accountability logging creates a tamper-evident audit trail connecting each recommendation to the AI model state, knowledge base state, reasoning artifacts, and human decision outcome. This log satisfies responsible AI governance framework requirements (IEEE, 2019; EU AI Act, 2024) and supports post-incident review, regulatory compliance, and organizational accountability. The governance layer is grounded in sociotechnical systems theory (Mumford, 2006), which emphasizes that technology design must incorporate human values and accountability norms rather than treating automation as an unconditional goal.

Layer 6: Feedback, Learning, And Reconfiguration

The feedback layer closes the ARM learning loop by systematically collecting and processing post-disruption outcome data. Outcome metrics include operational recovery time, financial impact, effectiveness of recommended versus implemented actions, resource efficiency, and stakeholder satisfaction. These metrics are codified and stored in the outcome archive, linked to corresponding event records and decision logs.

The learning module uses outcome data to update three types of system parameters: similarity weights in the similarity agent (improving precedent matching); recommendation confidence thresholds (calibrating conditions for high- versus low-confidence recommendations); and sensing thresholds in the monitoring agent (improving early warning precision for

specific disruption types). This reconfiguration logic operationalizes Teece et al.'s (1997) dynamic capabilities construct at the IS artifact level, translating the abstract notion of organizational reconfiguration into specific algorithmic update mechanisms.

Table 3. Design Principles for Algorithmic Resilience Memory

Principle	Description	Theoretical Grounding	Practical Implementation
DP1: Continuous Sensing	ARM must monitor multiple heterogeneous data streams in real time without operator prompting	Information processing theory (Galbraith, 1974)	Multi-source data ingestion APIs; edge-computing preprocessing; anomaly detection alerts with configurable sensitivity thresholds
DP2: Structured Encoding	Every disruption event must be encoded against a consistent ontological schema enabling cross-event comparison	Organizational memory theory (Walsh & Ungson, 1991); Case-based reasoning (Kolodner, 1993)	Domain-specific disruption ontology; graph database storage; automated schema validation on data ingest; human review gate for edge cases
DP3: Rationale Preservation	Decision rationale must be captured explicitly at time of decision — not inferred retrospectively from outcomes	Organizational learning (Argyris & Schön, 1978); Issue-based IS (Conklin & Begeman, 1988)	Structured decision log with mandatory fields: goal, alternatives considered, constraints, available information, authorizing role, override rationale
DP4: Contextual Recommendation	AI recommendations must integrate current organizational context — resource availability, strategic constraints — not only historical precedent similarity	Dynamic capabilities (Teece et al., 1997); DSS theory (Power, 2002)	Real-time organizational context module updated by sensing layer; constraint-aware planning agent; recommendations flagged when precedent match is below 0.60 similarity
DP5: Layered Explainability	Explanations must be calibrated to the cognitive and authority needs of	XAI (Arrieta et al., 2020); Decision support design	Three-tier explanation: executive summary (1 page), analytical narrative (3–5 pages), technical audit report

Principle	Description	Theoretical Grounding	Practical Implementation
	different organizational roles	(Montibeller & Winterfeldt, 2015)	(full system trace); SHAP-based attribution
DP6: Human Primacy	No high-stakes AI recommendation may be executed without explicit human authorization	Sociotechnical systems (Mumford, 2006); Responsible AI (Floridi et al., 2018)	Mandatory approval workflows; override controls; authority-tiered authorization thresholds based on decision impact magnitude
DP7: Accountable Logging	All AI recommendations, human decisions, and system states must be logged in a tamper-evident, queryable audit trail	Responsible AI governance (IEEE, 2019; EU AI Act, 2024)	Immutable audit log with cryptographic integrity verification; linked recommendation-decision-outcome records; regulatory compliance export module
DP8: Outcome-Driven Learning	Post-disruption outcomes must trigger systematic updates to recommendation weights, similarity metrics, and sensing thresholds	Double-loop learning (Argyris & Schön, 1978); Dynamic capabilities (Teece et al., 1997)	Structured outcome collection protocol at 30, 60, and 90 days post-disruption; automated parameter update pipeline; human-review gate before major model updates

VI. DEMONSTRATION SCENARIO: REGIONAL FLOOD AND SUPPLY-CHAIN DISRUPTION

Scenario Context

Meridian Components Ltd. is a mid-sized manufacturing firm operating in the lower Midwest of the United States, producing precision components for the automotive sector. The firm's primary manufacturing plant is located adjacent to a river system that has experienced three significant flooding events in the past eight years. Following the second flood in 2021, which caused \$4.2 million in direct operational losses and a 12-week supply-chain interruption, Meridian's operations director commissioned a post-incident review. That review found that while the 2021 response was superior to the 2017 response in some respects—early supplier notification and logistics

rerouting—significant institutional knowledge from 2017 had been lost through staff turnover, and the decision rationale from 2021 was not systematically documented. Meridian subsequently adopted the ARM Framework as part of a broader digital resilience initiative.

Early Warning Detection — Arm Sensing Layer

In October 2024, the ARM sensing layer detects a convergence of signals: the National Weather Service issues a Stage 2 flood watch for the adjacent river system; satellite-based soil saturation data indicates the watershed is at 94% capacity following three consecutive weeks of above-average precipitation; and commodity shipping indices indicate an 18% increase in alternative logistics costs consistent with anticipated regional disruption. The monitoring agent's composite risk score crosses the 0.78 threshold—calibrated from the 2017 and 2021 events—that reliably precedes operational impact within 72 to 96 hours.

An automated alert is dispatched to Meridian's operations leadership.

Event Encoding — Arm Knowledge Encoding Layer

The encoding layer initiates a structured disruption record, mapping the current event against the disruption ontology and classifying it as a Category B River Flood—consistent with the 2021 event—with an estimated impact radius of 35 kilometers, primary risk to transport infrastructure and ground-level production assets, and secondary risk to tier-1 supplier delivery capacity. The encoding layer queries the knowledge base and retrieves the 2017 and 2021 event records, including decision logs and outcome archives. The similarity agent calculates a 0.83 similarity score between the current event profile and the 2021 event, noting key contextual differences: the current event occurs during a period of higher logistics market tightness and lower internal inventory buffers.

Ai Reasoning And Recommendation — Arm Agentic

Ai Layer

The planning agent analyzes similarity-matched precedents and generates three ranked recommendations. Recommendation 1—rated high confidence (0.87)—draws from the 2021 playbook: initiate supplier notification to the three tier-1 suppliers in the flood zone; redirect component deliveries through the Memphis logistics hub; and pre-authorize overtime production at the Indiana secondary facility. The precedent citation notes that hub rerouting reduced supply interruption from 12 weeks (2017) to 6.5 weeks (2021). Recommendation 2—rated moderate confidence (0.72)—adds a new element: given the current inventory buffer deficit, pre-purchase a 60-day component stock from an alternative Ohio supplier identified through the ARM supplier diversity registry, a capability indexed from the 2021 lessons-learned record. Recommendation 3—rated lower confidence (0.61)—proposes a temporary production-sharing agreement with a peer firm in Chicago, a strategy untested in Meridian's history but flagged as effective in analogous industry cases.

1. Human Review and Decision — Decision Support and Governance Layers

The explainable decision support layer presents the three recommendations to Meridian's operations director and supply-chain manager through a dashboard displaying the flood risk map, inventory buffer status, confidence scores, precedent citations, and a side-by-side event comparison.

The executive summary highlights the key contextual difference from 2021: the lower inventory buffer implies higher exposure if hub rerouting alone proves insufficient. After review, the operations director approves Recommendations 1 and 2 in full and declines Recommendation 3—the production-sharing agreement—citing competitive risk with the identified peer firm. The modification rationale is captured in the decision log. The approved alternative—a 15% overtime expansion at the Indiana facility—is also logged.

E. Post-Disruption Learning — Feedback And Reconfiguration Layer

The 2024 flood peaks on October 19th and causes a 4.5-week supply interruption—notably shorter than 2021, despite the more challenging inventory environment. The feedback layer collects outcome data 60 days post-incident: supply interruption of 4.5 weeks versus 6.5 weeks in 2021; total financial impact of \$1.9 million versus \$4.2 million; logistics rerouting cost 3% above standard; component pre-purchase cost-efficiency within budget. The learning module assesses that Recommendation 2 (the pre-purchase strategy) was particularly effective and updates its confidence weight from 0.72 to 0.81 for analogous low-inventory scenarios. The rejection of Recommendation 3 is recorded as a competitive-sensitivity override, and future recommendation logic for similar situations will flag this criterion more prominently in the explanatory output. Figure 4 illustrates the learning and feedback loop that drives this post-disruption reconfiguration.



Figure 4. Climate-Disruption Learning and Feedback Loop

Note. The six-stage process runs left to right, beginning with the Disruption Event and concluding with Outcome Assessment. Three feedback arrows return information to earlier stages: (1) outcome data updates the Knowledge

Base (bottom arc, bold); (2) modification rationale updates AI recommendation logic (upper arc from Human Review to AI Recommendation); and (3) recalibrated sensing thresholds return from Outcome Assessment to the Disruption Event stage. Each completed cycle strengthens Algorithmic Resilience Memory.

VII. EVALUATION PLAN

Evaluation Strategy Overview

Consistent with DSR evaluation best practices (Venable et al., 2016), the ARM Framework is evaluated using a combination of scenario-based assessment and structured expert panel review. The evaluation is designed to assess both the utility of the artifact—whether it achieves the objectives specified in Section 4.3—and its quality: theoretical coherence, design completeness, and practical applicability. The evaluation does not provide longitudinal empirical validation, which the paper acknowledges as a future research requirement, but establishes the artifact's plausibility, internal consistency, and fitness for purpose within the DSR paradigm.

Scenario-Based Evaluation

Two scenarios serve as evaluation instruments. The primary scenario is the Meridian Components flood disruption described in Section 6. A secondary scenario depicts an extreme heat event affecting a food distribution company in the Southern United States, triggering cold-chain infrastructure failure and workforce productivity decline across three distribution centers. Expert participants engage with each scenario, reviewing the ARM Framework's response as documented in structured scenario artifacts, and rating the framework's performance on each evaluation dimension.

Expert Panel Composition

The expert panel consists of 12 to 15 participants drawn from four domains: IS researchers with expertise in digital resilience, AI, or organizational learning (4–5 participants); operations and supply-chain managers with direct experience managing climate-related disruptions (3–4 participants); AI governance specialists experienced in responsible AI deployment (2–3 participants); and climate risk officers from financial, logistics, or manufacturing organizations (2–3 participants). Participants are selected using purposive sampling to ensure domain coverage and practical experience relevance.

Evaluation Dimensions And Measurement Items

Table 4. Evaluation Criteria and Sample Measurement Items

Evaluation Dimension	Theoretical Basis	Sample Measurement Items (7-point Likert: 1 = Strongly Disagree, 7 = Strongly Agree)	Minimum Success Threshold
Usefulness	Technology acceptance; DSR utility (Hevner et al., 2004)	'The ARM Framework would help my organization respond more effectively to climate disruptions'; 'The framework addresses a genuine and important organizational need'	Panel mean ≥ 5.5 ; no item below 4.0
Explainability	XAI theory (Arrieta et al., 2020); Doshi-Velez & Kim (2017)	'The explanations provided for AI recommendations are sufficiently clear for managerial decision-making'; 'I could use the explanation to justify the decision to a senior stakeholder or regulator'	Panel mean ≥ 5.0
Decision Quality	DSS theory (Power, 2002); Arnott & Pervan (2014)	'The AI-generated recommendations in the scenario represent high-quality decision options for the disruption described'; 'The confidence scoring system supports well-calibrated managerial judgment'	Panel mean ≥ 5.0

Evaluation Dimension	Theoretical Basis	Sample Measurement Items (7-point Likert: 1 = Strongly Disagree, 7 = Strongly Agree)	Minimum Success Threshold
Organizational Learning	Argyris & Schön (1978); Crossan et al. (1999)	'The ARM Framework would enable systematic organizational learning from past climate disruptions'; 'The feedback and reconfiguration layer would meaningfully improve future response quality over time'	Panel mean ≥ 5.5
Governance Accountability	Responsible AI (IEEE, 2019; Floridi et al., 2018; EU AI Act, 2024)	'The human oversight mechanisms in the ARM Framework are sufficient to maintain organizational accountability for AI-influenced decisions'; 'The audit logging design would satisfy regulatory compliance requirements in my context'	Panel mean ≥ 5.5
System Completeness	DSR design quality (Gregor & Hevner, 2013; March & Smith, 1995)	'The six layers of the ARM Framework collectively address the full scope of organizational climate-disruption response'; 'I can identify no significant functional gap in the framework as presented'	Panel mean ≥ 5.0 ; no critical gap raised by >30% of panelists
Organizational Fit	Sociotechnical systems (Mumford, 2006)	'The ARM Framework could be integrated into existing organizational decision-making processes without requiring fundamental restructuring'; 'The governance layer is compatible with real-world organizational accountability structures'	Panel mean ≥ 4.5

Analysis Approach

Quantitative ratings from the Likert instrument will be summarized using descriptive statistics and analyzed for inter-rater consistency using Krippendorff's alpha. Qualitative data from structured expert panel interviews will be analyzed thematically (Braun & Clarke, 2006), with themes used to identify design refinements and future research priorities. The evaluation report will identify both strengths confirmed by the expert panel and limitations or design gaps surfaced by expert critique, consistent with the DSR principle of honest and transparent evaluation reporting (Hevner et al., 2004).

VII. DISCUSSION

Theoretical Implications For Information Systems Research

This paper makes a conceptual contribution to IS theory by introducing ARM as a new construct operating at the intersection of organizational memory, dynamic capabilities, digital resilience, agentic AI, and responsible AI governance. The introduction of ARM challenges IS researchers to reconceptualize what organizational memory means in an era of agentic AI. Walsh and Ungson's (1991) foundational model conceives of organizational memory as distributed across individuals, organizational roles, culture, structure, and ecology. ARM extends this model by

proposing that AI-mediated memory—structured, machine-interpretable, and algorithmically retrievable—constitutes a qualitatively distinct form of knowledge storage with different properties, capabilities, and risks than human-held memory.

Specifically, ARM introduces the concept of decision-rationale preservation, extending existing organizational memory theory by requiring that the reasoning underlying decisions—not merely their outcomes—be captured in a form supporting future AI-mediated retrieval and comparison. This is theoretically significant because it addresses a known limitation of case-based reasoning systems (Kolodner, 1993): contexts differ between past and present events in ways that make pure outcome-based precedent matching insufficient. Rationale preservation allows the AI to assess whether a past decision's reasoning remains valid in a new context, rather than blindly applying historical templates.

ARM also advances digital resilience research. Existing digital resilience frameworks are largely competency-based (Bhamra et al., 2011) or technology-inventory-based (Linkov & Trump, 2019); they describe what organizations should be capable of without specifying IS artifacts that operationalize those capabilities. ARM represents a step toward design-theoretic digital resilience research, grounding resilience capabilities in concrete artifact design decisions and linking those decisions explicitly to theoretical principles.

Practical Implications For Organizations

For organizations operating in climate-disruption-exposed sectors—manufacturing, logistics, agriculture, energy, and infrastructure—ARM offers a practical blueprint for transitioning from episodic, reactive crisis response toward systematic, AI-mediated institutional learning. The Meridian scenario demonstrates that even a modest ARM deployment, applied to a firm that had previously experienced two comparable disruptions, can generate measurably superior outcomes: a reduction in financial impact from \$4.2 million to \$1.9 million, and supply interruption duration from 6.5 to 4.5 weeks. Critically, ARM's value is cumulative: each disruption event enriches the knowledge base, improving recommendation quality for future events. Organizations that begin ARM adoption early in their climate-risk exposure trajectory will develop a compounding institutional advantage.

The framework's emphasis on explainability and human oversight addresses a practical barrier to AI adoption in organizational risk management: managerial reluctance to delegate consequential decisions to systems whose reasoning cannot be interrogated. ARM is designed to support rather than supplant managerial judgment, presenting AI recommendations as informed options rather than directives and preserving full managerial authority to override, modify, or decline them. This design philosophy is consistent with evidence that AI adoption in high-stakes organizational contexts is more successful when human decision-makers retain visible control and accountability (Logg et al., 2019).

Implications For Ai Governance

ARM's embedded governance layer addresses a growing concern in responsible AI research: the deployment of AI systems in high-stakes organizational environments without adequate accountability mechanisms (Floridi et al., 2018; Dignum, 2019). The ARM Framework demonstrates that governance is not an add-on to AI system design but an intrinsic architectural requirement. By making governance an explicit framework layer—with approval workflows, override controls, bias monitoring, and tamper-evident audit logging—ARM provides a design template that AI governance frameworks can reference when specifying architectural requirements for accountable AI in organizational resilience contexts. This is particularly timely as frameworks such as the EU AI Act (2024) begin to mandate specific human oversight requirements for high-risk AI applications in operational risk management.

Implications For Climate Resilience And Decision

Support Research

The paper's framing of climate disruption as a continuous organizational condition rather than an episodic event has important implications for both decision support research and climate resilience practice. Decision support research has long been organized around single-event architectures: a manager faces a decision, consults the DSS, and obtains a recommendation. ARM reframes this architecture around a learning loop: each event is both a decision challenge and a data-generating opportunity that improves future decision support quality. This reframing aligns DSS research more closely with the iterative, non-stationary character of climate risk, which is characterized by compounding severity and cross-event interdependence that static

architectures cannot accommodate. The ARM Framework thus extends Arnott and Pervan's (2014) call for adaptive DSS into the domain of climate-crisis management.

VIII. CONTRIBUTIONS

A. Theoretical Contributions

This paper contributes four theoretical advances to IS research. First, it introduces Algorithmic Resilience Memory as a new IS construct, providing a formal definition, a multi-dimensional theoretical structure, and clear differentiation from related constructs. ARM advances organizational memory theory by incorporating AI-mediated encoding and retrieval, and advances digital resilience theory by providing a design-theoretic account of how resilience capabilities can be instantiated in IS artifacts. Second, the paper extends dynamic capabilities theory into IS artifact design, operationalizing sensing, seizing, and reconfiguring at the level of specific system layers and design principles. Third, the paper introduces decision-rationale preservation as a distinct, design-relevant dimension of AI-mediated organizational memory, extending both organizational memory theory and decision support research. Fourth, the paper contributes eight theoretically grounded design principles that constitute an incipient design theory for ARM-class systems, consistent with Gregor and Hevner's (2013) framework for design theory contributions in IS.

B. Methodological Contributions

The paper demonstrates that DSR is a productive methodological approach for addressing the intersection of AI system design and organizational resilience theory. Specifically, it illustrates how scenario-based evaluation can be rigorously structured to assess multi-dimensional IS artifacts, providing a replicable methodological template for future DSR studies in related domains. The evaluation framework—with its seven dimensions, structured Likert-scale measurement items, and minimum success thresholds—is independently reusable for evaluating agentic AI system designs in organizational settings beyond the specific context of climate-crisis resilience.

C. Practical Contributions

Three practical contributions are offered. First, the ARM Framework provides a concrete architectural blueprint that IS professionals, AI engineers, and organizational risk managers can adapt for implementation. The six-layer

architecture is deliberately modular: organizations can begin with subset adoption—deploying Layers 1–3 initially—and expand as organizational capacity and AI maturity develop. Second, the eight design principles provide actionable guidance for organizations designing or procuring agentic AI systems for climate-resilience purposes, reducing the design uncertainty accompanying novel AI deployment contexts. Third, the Meridian demonstration scenario provides a worked example that practitioners can reference when communicating the ARM concept to organizational stakeholders, evaluating vendor proposals, or building a business case for adoption.

IX. LIMITATIONS AND FUTURE RESEARCH

A. Limitations

This paper has several important limitations. First, the ARM Framework is presented at the conceptual and logical design levels; no software prototype has been implemented or tested in a live organizational setting. The artifact's technical feasibility is supported by reference to existing technologies—graph databases, multi-agent AI architectures, XAI tools—but integration complexity, performance under real-world data volumes, and latency characteristics remain empirically untested. The evaluation is accordingly limited to expert judgment of design quality rather than measured system performance.

Second, the evaluation relies on expert panel assessment and scenario-based methods. While these are established DSR evaluation approaches (Venable et al., 2016), they do not provide the statistical generalizability of large-sample empirical studies or the ecological validity of longitudinal field research. Expert ratings reflect judgments about plausibility and potential utility, not measured organizational outcomes. Third, the ARM Framework is designed at a level of abstraction intended to support cross-sector applicability, but operational instantiation will necessarily be organization-specific. Sector-specific data integration requirements, regulatory constraints, and governance structures require customization not addressed in detail here.

Fourth, the paper does not address the technical challenges associated with encoding knowledge quality and completeness over extended time periods. Knowledge bases degrade if event records are incomplete, if decision

logs capture rationale inaccurately, or if feedback data is not collected systematically. ARM's effectiveness is contingent on organizational data discipline, which varies substantially across organizations and industries. Fifth, the agentic AI architectures referenced in this paper are evolving rapidly; specific technical choices embedded in the framework design may require revision as the field advances.

B. Future Research Directions

Several important future research directions emerge. First, the development and testing of an ARM prototype in a partner organization would provide empirical validation of the design principles and effectiveness, enabling design iteration across multiple disruption cycles. Second, longitudinal research studying organizations across successive climate disruptions would provide evidence for the claim that ARM generates cumulative learning benefits—testing whether knowledge base quality demonstrably improves recommendation accuracy over time. Third, cross-industry comparative research would examine how ARM design requirements differ across sectors with distinct climate-risk profiles: agriculture, coastal infrastructure, healthcare supply chains, and financial services. Fourth, sociotechnical adoption research examining the organizational behavior dimensions of human-AI collaboration in high-stakes crisis environments would complement this paper's design-theoretic contribution with empirical insights into implementation barriers and success factors. Fifth, development of ARM governance standards and regulatory alignment documentation would support adoption in jurisdictions where AI governance requirements are becoming legally binding. Sixth, future research should investigate AI fairness and equity dimensions of ARM, examining whether resilience memory systems systematically favor certain organizational units, regions, or stakeholder groups in their recommendation logic.

X. CONCLUSION

The climate crisis is an organizational reality, not a future risk. For many organizations, the question is no longer whether climate-related disruptions will occur but how quickly and effectively the organization can respond—and whether each successive disruption makes the organization better equipped for the next. This paper has argued that addressing this challenge requires a new category of IS capability: one that does not merely predict climate

disruption or restore prior operational states, but remembers, learns, explains, and adapts over time.

Algorithmic Resilience Memory is introduced as that capability: a theoretically grounded, design-oriented IS construct integrating agentic AI, organizational memory, dynamic capabilities, and responsible AI governance into a unified framework architecture. Through the design of the Agentic AI-Based Algorithmic Resilience Memory Framework and its demonstration in the Meridian Components flood scenario, the paper has shown that ARM is both conceptually coherent and practically plausible. The framework's six layers, eight design principles, and embedded governance architecture provide organizations with a structured pathway from episodic, reactive crisis response toward systematic, AI-mediated institutional learning.

The conceptual contribution of this paper extends beyond the artifact. By introducing ARM as a new IS construct, the paper opens a line of inquiry that connects organizational memory theory, dynamic capabilities research, digital resilience, and responsible AI governance in ways that none of these traditions has pursued individually. It reframes the IS challenge of climate-crisis adaptation from the question of what systems can predict to the question of what systems can remember, learn from, and improve upon. That reframing is, we argue, both theoretically necessary and practically urgent.

Information systems research has a distinctive role to play in climate-crisis adaptation: it is the discipline best positioned to design, evaluate, and govern the digital infrastructures through which organizations sense, decide, and learn. By providing a rigorous design-science foundation for ARM's further development—theoretical, methodological, and practical—this paper contributes to that larger project. The organizations that invest in building algorithmic resilience memory today will not merely survive the disruptions of tomorrow; they will progressively master the conditions of a climate-altered organizational world.

REFERENCES

1. Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.

2. Arnott, D., & Pervan, G. (2014). A critical analysis of decision support systems research revisited: The rise of design science. *Journal of Information Technology*, 29(4), 269–293. <https://doi.org/10.1057/jit.2014.16>
3. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & 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>
3. Baptista, J., Stein, M. K., Klein, S., Watson-Manheim, M. B., & Lee, J. (2020). Electronic government and the information technology artefact. *Journal of Information Technology*, 35(3), 241–257. <https://doi.org/10.1177/0268396220921550>
4. Bhamra, R., Dani, S., & Burnard, K. (2011). Resilience: The concept, a literature review and future directions. *International Journal of Production Research*, 49(18), 5375–5393. <https://doi.org/10.1080/00207543.2011.563826>
5. Bhatt, G. D., & Zaveri, J. (2002). The enabling role of decision support systems in organizational learning. *Decision Support Systems*, 32(3), 297–309. [https://doi.org/10.1016/S0167-9236\(01\)00120-8](https://doi.org/10.1016/S0167-9236(01)00120-8)
6. Boin, A., & McConnell, A. (2007). Preparing for critical infrastructure breakdowns: The limits of crisis management and the need for resilience. *Journal of Contingencies and Crisis Management*, 15(1), 50–59. <https://doi.org/10.1111/j.1468-5973.2007.00504.x>
7. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
8. Buxmann, P., & Schmidt, H. (2021). Artificial intelligence. *Business & Information Systems Engineering*, 63(1), 1–4. <https://doi.org/10.1007/s12599-020-00683-0>
9. Conklin, J., & Begeman, M. L. (1988). gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Information Systems*, 6(4), 303–331. <https://doi.org/10.1145/58566.59297>
10. Crossan, M. M., Lane, H. W., & White, R. E. (1999). An organizational learning framework: From intuition to institution. *Academy of Management Review*, 24(3), 522–537. <https://doi.org/10.5465/amr.1999.2202135>
11. Dignum, V. (2019). Responsible artificial intelligence: How to develop and use AI in a responsible way. Springer. <https://doi.org/10.1007/978-3-030-30371-6>
12. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608. <https://arxiv.org/abs/1702.08608>
13. Artificial Intelligence Act. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council on artificial intelligence. Official Journal of the European Union, L 2024/1689. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People: An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
14. Galbraith, J. R. (1974). Organization design: An information processing view. *Interfaces*, 4(3), 28–36. <https://doi.org/10.1287/inte.4.3.28>
15. Gallopín, G. C. (2006). Linkages between vulnerability, resilience, and adaptive capacity. *Global Environmental Change*, 16(3), 293–303. <https://doi.org/10.1016/j.gloenvcha.2006.02.004>
16. Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337–355. <https://doi.org/10.25300/MISQ/2013/37.2.01>
17. Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G. Z. (2019). XAI—Explainable artificial intelligence. *Science Robotics*, 4(37), eaay7120. <https://doi.org/10.1126/scirobotics.aay7120>
18. Herbane, B., Elliott, D., & Swartz, E. M. (2004). Business continuity management: Time for a strategic role? *Long Range Planning*, 37(5), 435–457. <https://doi.org/10.1016/j.lrp.2004.07.002>
19. Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
20. Hohenstein, N. O., Feisel, E., Hartmann, E., & Giunipero, L. (2015). Research on the phenomenon of supply chain resilience: A systematic review and paths for further investigation. *International Journal of Physical Distribution & Logistics Management*, 45(1/2), 90–117. <https://doi.org/10.1108/IJPDLM-05->

- 2013-0128 IEEE. (2019). Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems (1st ed.). IEEE Standards Association.
21. Jennex, M. E., & Olfman, L. (2005). Assessing knowledge management success. *International Journal of Knowledge Management*, 1(2), 33–49. <https://doi.org/10.4018/jkm.2005040104>
 22. Kolodner, J. L. (1993). *Case-based reasoning*. Morgan Kaufmann Publishers.
 23. Lengnick-Hall, C. A., & Beck, T. E. (2005). Adaptive fit versus robust transformation: How organizations respond to environmental change. *Journal of Management*, 31(5), 738–757. <https://doi.org/10.1177/0149206305279367>
 24. Linkov, I., & Trump, B. D. (2019). *The science and practice of resilience*. Springer. <https://doi.org/10.1007/978-3-030-04565-4>
 25. Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
 26. March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems*, 15(4), 251–266. [https://doi.org/10.1016/0167-9236\(94\)00041-2](https://doi.org/10.1016/0167-9236(94)00041-2)
 27. Markus, M. L. (2001). Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success. *Journal of Management Information Systems*, 18(1), 57–93. <https://doi.org/10.1080/07421222.2001.11045671>
 28. Montibeller, G., & von Winterfeldt, D. (2015). Cognitive and motivational biases in decision and risk analysis. *Risk Analysis*, 35(7), 1230–1251. <https://doi.org/10.1111/risa.12360>
 29. Mumford, E. (2006). The story of socio-technical design: Reflections on its successes, failures and potential. *Information Systems Journal*, 16(4), 317–342. <https://doi.org/10.1111/j.1365-2575.2006.00221.x>
 30. Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
 31. Power, D. J. (2002). *Decision support systems: Concepts and resources for managers*. Quorum Books.
 32. Shim, J. P., Warkentin, M., Courtney, J. F., Power, D. J., Sharda, R., & Carlsson, C. (2002). Past, present, and future of decision support technology. *Decision Support Systems*, 33(2), 111–126. [https://doi.org/10.1016/S0167-9236\(01\)00139-7](https://doi.org/10.1016/S0167-9236(01)00139-7)
 33. Shinn, N., Cassano, F., Berman, E., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 8634–8652.
 34. Stein, E. W. (1995). Organizational memory: Review of concepts and recommendations for management. *International Journal of Information Management*, 15(1), 17–32. [https://doi.org/10.1016/0268-4012\(94\)00003-C](https://doi.org/10.1016/0268-4012(94)00003-C)
 35. Stein, E. W., & Zwass, V. (1995). Actualizing organizational memory with information systems. *Information Systems Research*, 6(2), 85–117. <https://doi.org/10.1287/isre.6.2.85>
 36. Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7<509::AID-SMJ882>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z)
 37. Turban, E., Sharda, R., & Delen, D. (2011). *Decision support and business intelligence systems* (9th ed.). Prentice Hall.
 38. Venable, J. R., Pries-Heje, J., & Baskerville, R. (2016). FEDS: A framework for evaluation in design science research. *European Journal of Information Systems*, 25(1), 77–89. <https://doi.org/10.1057/ejis.2014.36>
 39. Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57–91. <https://doi.org/10.5465/amr.1991.4278992>
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W., Wei, Z., & Wen, J.-R. (2023). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18, 186345. <https://doi.org/10.1007/s11704-024-40231-1>



40. Wieland, A., & Wallenburg, C. M. (2013). The influence of relational competencies on supply chain resilience: A relational view. *International Journal of Physical Distribution & Logistics Management*, 43(4), 300–320. <https://doi.org/10.1108/IJPDLM-08-2012-0243> Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, H., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., & Gui, T. (2023). The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*. <https://arxiv.org/abs/2309.07864>
41. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. *Proceedings of the International Conference on Learning Representations (ICLR 2023)*. <https://arxiv.org/abs/2210.03629>