Layered Automation in Hospital Crisis Readiness: A Scoping Review and Framework for Deterministic, Intelligent, and Agentic Systems

Parag Bhatnagar
Digital Environment Research Institute
Queen Mary University of London
Parag.Bhatnagar@tewhatuora.govt.nz

ABSTRACT

Hospitals face recurring crisis conditions, including pandemics, mass casualty events, natural hazards, and infrastructure disruptions, which stress operational coordination, compress decision cycles, and amplify information fragmentation. Automation has increasingly been used to support crisis responses through workflow execution, decision support, and predictive modelling. However, crisis automation literature remains fragmented across system types, hazards, and organizational contexts, and lacks a governance-aware synthesis that distinguishes between levels of autonomy and corresponding accountability requirements. This study reports an scoping review using EBSCOhost databases (PRISMA-ScR) of hospital crisis automation. Across CINAHL Ultimate and MEDLINE Ultimate, two search families (core automation; AI/agent terminology) yielded 2,297 records. A supplementary AI-assisted discovery search yielded 1,000 additional records for review. After screening predefined inclusion criteria and deduplication, 116 studies were included in the analysis. The synthesis identifies two dominant strata of crisis automation: deterministic workflow automation and bounded intelligent decision augmentation. In contrast, policy-constrained agentic automation, defined as goal-directed orchestration across heterogeneous systems, appears largely absent as a mature operational form in literature. Building on this synthesis, we propose a layered framework that links automation strata to integration patterns, failure modes, and escalating governance controls. A concrete use case illustrates how deterministic, intelligent, and agentic layers would differ in a surge scenario, highlighting practical constraints regarding data access, interoperability, and human authorization. This study contributes (i) a consolidated map of hospital crisis automation research, (ii) a governance-aware layered framework for disaster public health and healthcare informatics, and (iii) implications for designing and evaluating emerging agencies in crisis settings.

Keywords

Crisis informatics; hospital operations; disaster public health; robotic process automation; intelligent automation; agentic automation; governance; PRISMA-ScR

INTRODUCTION

Hospital systems operate under chronic operational constraints, even under routine conditions. During crisis events—pandemics, mass casualty incidents, extreme weather, cyber disruptions, infrastructure failure, and other surges—the mismatch between demand and available capacity becomes acute. Crisis conditions also degrade the coordination mechanisms required to respond: staffing becomes volatile, physical space is reconfigured, supply chains are destabilised, and information flow becomes fragmented across departments and tools. Crisis response thus depends not only on clinical capability but also on the operational capacity to coordinate, allocate scarce resources, and sustain safe throughput.

Digital technologies have become central to hospital crisis responses. Across contexts, hospitals have deployed automated routing, bed and capacity dashboards, surveillance reporting, triage support, and predictive models to stabilise operations and reduce reliance on manual coordination. These tools are often described as "automation", "decision support", "AI", or "digital response", but the literature remains conceptually and empirically fragmented on the topic. Studies frequently focus on narrow system deployments, specific hazards, or operational functions, making it difficult to draw generalisable lessons about what automation can realistically do in crisis contexts and what governance structures are required to do it safely.

Simultaneously, the rapid diffusion of large language model (LLM) tooling and "agentic" architectures has reopened questions about automation autonomy. In principle, agentic systems can coordinate across multiple operational systems by synthesising heterogeneous data and recommending multi-step actions. In practice, crisis settings impose heightened requirements: data protection constraints intensify, auditability becomes non-negotiable, model drift risk increases, and high-impact actions demand explicit authorisation. Without governance-aware framing, "agentic automation" risks being either oversold or prematurely dismissed.

This study shows that hospital crisis automation is mature at deterministic and intelligent levels, but systems that coordinate actions across workflows are not yet in use. More importantly, it shows that as automation becomes more advanced, governance requirements increase.. By empirically mapping automation strata and articulating their governance implications, the paper advances crisis informatics scholarship and helps evaluate more advanced automation systems in hospitals.

Research question

This study asks: How are automation systems deployed in hospital crisis contexts, and what changes as automation becomes more advanced?

Contributions

This paper contributes:

1. A map of hospital crisis automation studies
2. A framework with three levels of automation
3. A discussion of governance requirements for each level

Roadmap

Section 2 locks definitions and motivates the layered perspective. Section 3 describes the scoping review method and PRISMA-ScR flow. Section 4 reports results and synthesised themes. Section 5 presents the layered framework. Section 6 provides a concrete use case walkthrough. Section 7 discusses implications for crisis informatics, governance, and research design. Section 8 concludes.

CONCEPTUAL FRAMING AND DEFINITIONS

Why "automation strata" matter in crisis informatics

Crisis informatics has long emphasised socio-technical coordination under uncertainty: information systems are not neutral conduits but shape what is seen, by whom, and when; they allocate attention and structure collaboration. In hospital settings, this is amplified by regulatory constraints, fragmented IT landscapes, and heterogeneous data sources. Automation therefore cannot be treated as a single category. Systems that execute predefined workflows differ qualitatively from systems that predict future demand, and both differ from systems that plan and coordinate actions across multiple tools.

A layered framing helps address two recurrent problems in this literature:

1. Conceptual ambiguity: "AI," "automation," "decision support," and "digital response" are often conflated.
2. Governance mismatch: as autonomy increases, governance and assurance demands escalate. Treating all automation as equivalent obscures this.

Definitions

RPA (Robotic Process Automation): rule-based automation that executes predefined digital workflow steps in structured environments (Lacity & Willcocks, 2018; Van Der Aalst et al., 2018).

Intelligent Automation: automation augmented by predictive, classification, extraction, or modelling capabilities (e.g., ML, simulation), typically supporting decisions rather than autonomously executing high-impact actions

(Jordan & Mitchell, 2015).

Agentic Automation: automation that can plan and carry out multi-step actions across systems (Yao et al., 2022), invoking tools across systems (Schick et al., 2023), and adapting strategies within explicit policy constraints and human approval gates (Tabassi, 2023).

Hospital crisis readiness: the operational capacity of hospitals to maintain safe throughput, coordination, and service continuity under crisis/surge conditions (Kruk et al., 2015; Luke et al., 2023).

Information systems for crisis response: systems that support coordination (Comfort et al., 2004), triage (Bazyar et al., 2019), resource allocation (Romanowski et al., 2015), monitoring (Endsley, 1995), and decision-making under emergency conditions, including cross-department and cross-agency information exchange (Bazyar et al., 2019; Van De Walle & Turoff, 2007, 2008).

All other technical terms are treated as sub-capabilities within intelligent or agentic automation. Intelligent systems generate predictions but do not initiate actions, whereas agentic systems are defined by their ability to propose or execute multi-step actions across systems under explicit constraints.

Crisis Informatics and Socio-Technical Coordination Under Disruption

Crisis informatics scholarship examines how information systems support coordination, sensemaking, and decision-making under conditions of uncertainty and time compression (Palen et al., 2009). In disaster contexts, information flows are disrupted, organisational roles become fluid, and digital infrastructures must adapt rapidly to evolving situational demands. Hospitals face extreme pressure during crises. Coordination becomes difficult as demand rises and information changes rapidly across systems.

Prior work in crisis informatics has emphasised situational awareness platforms, information sharing networks, and collaborative technologies that support distributed coordination. However, less attention has been paid to how varying levels of automation alter the distribution of coordination authority between human and technical actors. From a socio-technical systems perspective (Orlikowski, 1992; Trist, 1981), technologies do not merely transmit information but reshape patterns of control, responsibility, and decision latency.

The layered automation framework proposed in this study extends crisis informatics by conceptualising automation as an escalation of socio-technical agency. Deterministic systems stabilise predefined workflows. Intelligent systems reduces uncertainty by introducing probabilistic abstraction. Agentic systems—if realised—would further redistribute coordination authority across heterogeneous systems under policy constraints. In crisis settings, such redistribution carries heightened governance implications because authority, accountability, and data protection are amplified under emergency conditions.

By situating automation strata within crisis informatics and socio-technical theory, this study reframes hospital crisis automation not only as technical deployment but as structured reallocation of coordination power under disruption.

METHOD: SCOPING REVIEW AND FRAMEWORK SYNTHESIS

Design

A scoping review was conducted following PRISMA-ScR guidance to map the landscape of automation systems deployed in hospital crisis contexts. The objective was to identify patterns in automation capability and governance articulation rather than to conduct quality-weighted effect synthesis. The review process included structured identification, screening, eligibility assessment, and data extraction.

Information Sources and Search Strategy

We searched two databases using EBSCOhost:

- **CINAHL Ultimate**
- **MEDLINE Ultimate**

Two structured search families were developed to capture both operational automation and AI-driven systems:

1. **Search Family 1** (Operational Automation Focus)
    a. (crisis OR disaster OR emergency OR pandemic OR surge) AND
    b. (hospital OR healthcare) AND

   c. (workflow automation OR process automation OR RPA OR decision support OR intelligent automation)
2. **Search Family 2** (AI / Agent Focus)
   a. (crisis OR disaster OR emergency OR pandemic OR surge) AND
   b. (hospital OR healthcare) AND
   c. (AI OR machine learning OR LLM OR agent OR autonomous)

Searches were conducted for studies published between 2010 and 2025.

Inclusion and Exclusion Criteria

Studies were included if they:

1. Examined hospital or healthcare system settings.
2. Addressed crisis, disaster, surge, or emergency contexts.
3. Described deployment of automation affecting operational workflow or decision support.
4. Were empirical studies, system evaluations, or structured technical deployments.

Studies were excluded if they:

1. Focused solely on clinical treatment efficacy without operational automation.
2. Examined public health surveillance outside hospital operational contexts.
3. Were commentary without system description.

Screening and Data Handling

Search results were exported and consolidated in structured spreadsheets.

Screening proceeded in two stages:

1. Title and abstract review against eligibility criteria.
2. Full-text review for records meeting inclusion thresholds.

We also used Elicit to identify additional studies using the same core query terms. Results were limited to English-language publications and screened using the same inclusion criteria as database results. All results were screened using the same criteria. This was used to supplement database searches and capture emerging AI-related terminology. A total of 3,297 records were identified. Following title and abstract screening, 3,170 records were excluded. A total of 127 records were assessed at full text, of which 11 were excluded due to duplication or ineligibility. This resulted in a final corpus of 116 studies. The PRISMA-ScR flow diagram summarises this process (Figure 1).

Data Extraction

Deduplication was performed on the included set using DOI matching as the primary key and title–year as a secondary key. A structured extraction sheet was developed to capture crisis type, operational domain, automation characteristics, governance controls described, integration patterns, and publication year. Automation systems were then categorised into one of three strata: deterministic automation, intelligent automation, and agentic automation (conceptual or advisory). Stratum classification was guided by predefined criteria relating to execution authority and coordination scope. Studies were classified as candidate agentic only if they described multi-step coordination across systems; purely predictive or advisory systems were excluded.

Classification Validation

To enhance classification reliability, all records initially categorised as "agentic" or "unclear" were manually reviewed. A random subset of deterministic and intelligent classifications was also validated to confirm consistency with the predefined stratum definitions. Any discrepancies were resolved through re-examination of full texts. This validation procedure strengthens confidence in the interpretive synthesis while remaining appropriate to the exploratory scope of a scoping review.

Limitations of Method

As a scoping review, this study maps conceptual and functional patterns rather than conducting formal quality appraisal. Screening and classification were conducted by a single researcher; while structured validation was performed, formal inter-coder reliability was not calculated. Future work may incorporate multi-reviewer coding to enhance reproducibility.

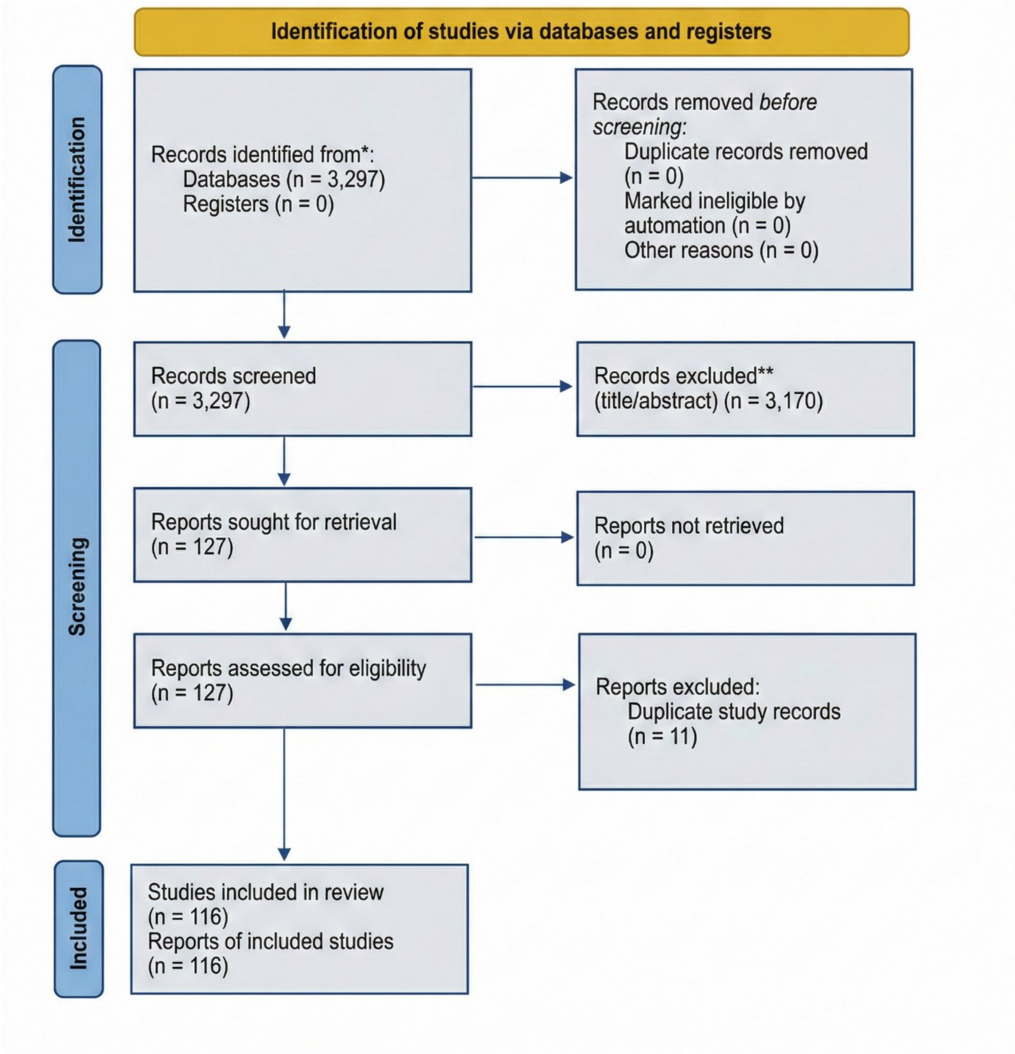

Figure 1. PRISMA-ScR Flow Diagram

RESULTS

Corpus overview

A total of 116 studies met inclusion criteria. As shown in Table 1, the literature is concentrated in pandemic and surge-related contexts, reflecting the acceleration of digital deployment during COVID-19. Deterministic and intelligent automation dominate the corpus, while agentic systems appear primarily in conceptual or advisory form. Operationally, most systems are deployed at the emergency department, ICU, or hospital-wide coordination level, with fewer studies examining cross-institutional orchestration. Publication volume increases markedly post-2020, indicating rapid crisis-driven digital transformation. Of the included studies, over half addressed COVID-19 contexts (53.4%), with surge and emergency capacity forming the second largest category (27.6%). Intelligent automation constituted the dominant stratum (56.0%), followed by deterministic systems (25.9% when merged), while only 3.4% of studies approached agentic automation in conceptual or advisory form. Publication volume increased markedly post-2020, reflecting crisis-driven digital acceleration.

Table 1. Characteristics of Included Studies (N = 116)

Crisis Type	n	%
COVID-19	62	53.4
General emergency / Surge	32	27.6

Natural hazard / Fire	5	4.3
Pandemic / Outbreak (non-COVID)	5	4.3
Unspecified / General	5	4.3
Mass casualty	4	3.4
Mixed categories (combined labels)	3	2.6

Table 2. Automation Stratum

Automation Stratum	n	%
Intelligent	65	56.0
Deterministic / Decision-support	29	25.0
Deterministic (explicit)	1	0.9
Unclear	17	14.7
Agentic (candidate)	4	3.4

Table 3. Distribution of operational domains across included studies

Operational Domain	n	%
Unspecified / General	20	17.2
Triage / Routing	8	6.9
ICU / Critical Care Ops	8	6.9
Capacity / Bed Management	5	4.3
Diagnostics Coordination	5	4.3
Staffing / Workforce	4	3.4
Supply / Logistics	4	3.4
ED Operations	3	2.6
Surveillance / Reporting	3	2.6
Other (remaining domains)		

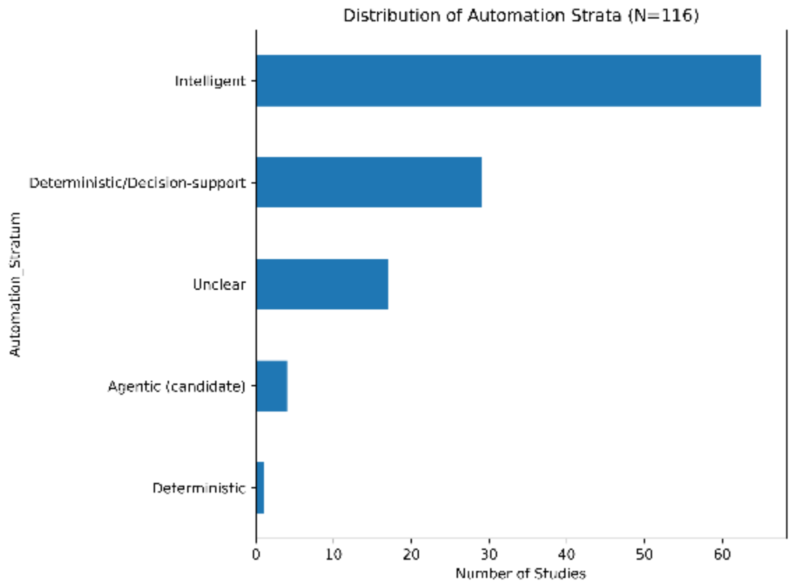

Figure 2. Automation Strata Distribution

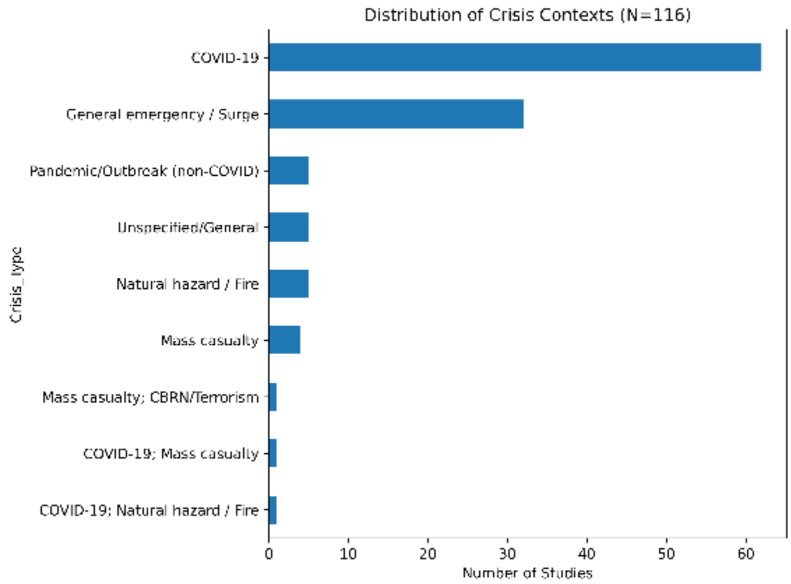

Figure 3. Crisis Context Distribution

Operational Domains Represented

Across the corpus, crisis automation systems cluster into operational domains:

1. Triage and routing: prioritisation support, routing logic, queue management, and escalation workflows.
2. Capacity and throughput: bed management, occupancy forecasting, surge dashboards, discharge coordination.
3. Diagnostics and clinical operations coordination: coordination of imaging/lab workflows under surge conditions; prioritisation support.
4. Command-and-control information integration: dashboards and shared situational awareness tools, often aggregating multiple sources.
5. Resource allocation and logistics: staffing allocation, supply constraints, scheduling under disruption.
6. Reporting and public health interface: surveillance reporting, cross-agency information exchange under emergency response.

Thematic Synthesis: Three Cross-cutting Clusters

Three synthesis clusters emerged:

Cluster A: Operational capability enablement

Most studies focus on stabilising operational throughput during crisis conditions by reducing manual coordination overhead. Deterministic automation and structured decision support dominate: routing logic, triggers, dashboards, and queue management are frequently used to make processes legible and tractable under pressure. These systems tend to be narrow in scope but high in operational relevance.

Cluster B: Governance and safety/assurance controls

Governance discussion is uneven. Many deterministic systems incorporate basic controls (role-based access, logging, manual override), but explicit articulation of accountability, escalation protocols, and auditability under crisis degradation is inconsistent. Intelligent systems often retain human confirmation gates for high-impact decisions, but detailed treatment of model drift and distribution shift under crisis conditions is variable.

A critical observation is that governance is often treated as implementation detail rather than a first-class design requirement, despite crisis contexts amplifying the consequences of automation error.

Cluster C: Integration patterns and constraints

Integration is a structural constraint across the corpus. Many systems are deployed as standalone platforms or lightly integrated dashboards, often using bespoke interfaces rather than API-native interoperability. This limits cross-system coordination, restricts automation scope, and inhibits higher-autonomy orchestration. Fragmented architectures also amplify governance risk by increasing opacity, reducing end-to-end traceability, and complicating incident response when automation misbehaves.

Automation Strata: What is Present

Deterministic Automation (dominant)

Deterministic automation constitutes the most prevalent form of crisis automation in the reviewed corpus. These systems execute predefined rules to maintain throughput and visibility under surge conditions. Examples include emergency department decision-support tools (Boltin et al., 2018), rapid informatics deployments during COVID-19 (Reeves et al., 2020), and digitally adapted operational workflows (Pankhurst et al., 2021). Dashboard-based command systems further illustrate the role of structured information consolidation during crisis operations (Parker et al., 2024). Such systems operate within bounded authority: they execute structured logic but do not independently redefine goals. Their primary function is friction reduction within known workflows rather than adaptive coordination across systems. Importantly, governance at this layer is relatively mature. Studies commonly report audit logging, role-based access, and manual override. However, brittleness emerges when crisis protocols evolve rapidly, requiring rule reconfiguration.

Typical patterns include: routing based on stable criteria, automated status updates and task creation, threshold-based triggers and alerts, and queue management and escalation.

Typical failure modes: rule brittleness, misalignment under rapidly shifting protocols, brittle dependencies on upstream data quality, and silent failure when interfaces change.

Typical controls: manual override, audit logs, role-based access, and exception handling. Crisis readiness at this layer depends heavily on monitoring and operational fallback processes.

Intelligent Automation (substantial; bounded)

Intelligent automation appears primarily as decision augmentation: predictive models, simulation-based planning tools, and classification/risk scoring systems. These systems aim to compress uncertainty and support faster decision cycles, especially for capacity planning and triage. A substantial subset of studies deploy predictive modelling and simulation to anticipate capacity constraints or clinical deterioration. For example, near-real-time occupancy modelling supports surge anticipation (Preiss et al., 2022). Machine learning models predict hospitalisation risk among COVID-positive patients (Song et al., 2022) and mortality risk within EHR systems (Sottile et al., 2021). Simulation models further inform crisis planning under epidemiological uncertainty (Groves-Kirkby et al., 2023). These systems shift coordination from reactive to anticipatory modes. However, they remain embedded within human decision structures. The literature consistently describes model outputs as advisory rather than autonomously executable. This bounded authority reflects recognition of distribution shift and model drift risks in crisis contexts. Governance mechanisms, i.e. validation procedures, calibration, and human confirmation gates are more visible at this layer than in deterministic systems. Thus, intelligent automation augments, but does not supplant, human-led coordination.

Typical failure modes: distribution shift (crisis conditions differ from training conditions), overconfidence in model outputs, and feedback loops (model predictions influence the system in ways that invalidate assumptions).

Typical controls: validation checkpoints, human confirmation gates, monitoring for drift, and conservative thresholds.

Agentic Automation (largely absent as mature operational form)

Across the included corpus, mature agentic automation—goal-directed multi-step planning across heterogeneous hospital systems—appears largely absent as an operationally deployed and evaluated form. Some studies describe distributed coordination or advanced decision-support concepts, and recent papers referencing LLMs introduce advisory possibilities, but these do not amount to policy-constrained tool-using orchestration operating across multiple clinical/operational systems with auditability and human authorisation.

This absence is a key result: it suggests that the “frontier” challenge for crisis automation is not simply adding intelligence, but achieving governance-safe coordination across fragmented systems.

While recent publications discuss distributed coordination and LLM applications, the corpus does not document operational, policy-constrained agentic automation within hospital crisis environments. Distributed multi-agent architectures for emergency coordination have been proposed (CINCAR & IVAŞCU, 2025; Lujak et al., 2016). However, these studies describe architectural coordination frameworks rather than integrated hospital IT orchestration under explicit governance controls. Similarly, recent studies examining LLM use in emergency medicine (Kuas et al., 2025; Shenoy et al., 2025) position such systems as advisory tools. They do not demonstrate autonomous multi-step planning across heterogeneous hospital systems with encoded policy constraints and auditable human approval. This empirical absence suggests that integration architecture and governance maturity, not model capability, are the primary constraints.

LAYERED AUTOMATION IN HOSPITAL CRISIS CONTEXTS

Empirical Consolidation of Automation Strata

The review shows three types of automation. Most systems fall into two groups: deterministic and intelligent. Systems that coordinate actions across multiple systems are not observed. Deterministic systems stabilise workflows. Intelligent systems support prediction and planning. By contrast, agentic automation—defined here as multi-step, cross-system planning and execution under explicit policy constraints—does not appear in sustained operational form within the corpus. Table 4 summarises these three types. These findings explain the structure of the framework in Figure 2. Most systems stabilise workflows or support decisions. Systems that coordinate actions across multiple systems are not observed. This pattern suggests that current hospital automation is constrained by system integration rather than algorithmic capability. Progression between layers depends on organisational and technical infrastructure rather than model sophistication alone.

Table 4. Automation strata observed in hospital crisis literature (N = 116)

Dimension	Deterministic Automation	Intelligent Automation	Agentic Automation (Conceptual / Advisory)
Primary function	Workflow execution and threshold control	Probabilistic prediction and planning support	Conceptual coordination or LLM-based advisory
Crisis deployment examples	ED routing tools (Boltin et al., 2018); COVID digital response workflows (Pankhurst et al., 2021; Reeves et al., 2020); command dashboards (Parker et al., 2024)	Surge prediction (Preiss et al., 2022); hospitalisation prediction (Song et al., 2022); EHR mortality prediction (Sottile et al., 2021); simulation planning (Groves-Kirkby et al., 2023)	Distributed coordination architectures (Lujak et al., 2016); multi-agent system proposals (CINCAR & IVAŞCU, 2025); LLM advisory use in emergency contexts (Kuas et al., 2025; Shenoy et al., 2025)
Execution authority	Executes predefined rules	Generates predictive outputs; human confirmation typical	Advisory or architectural; no documented autonomous cross-system execution
Integration pattern	Event-triggered workflow systems	Model + dashboard integration	Conceptual cross-system or advisory integration
Governance description	Logging and manual override	Validation, calibration, human oversight	Governance concerns discussed; operational safeguards rarely specified

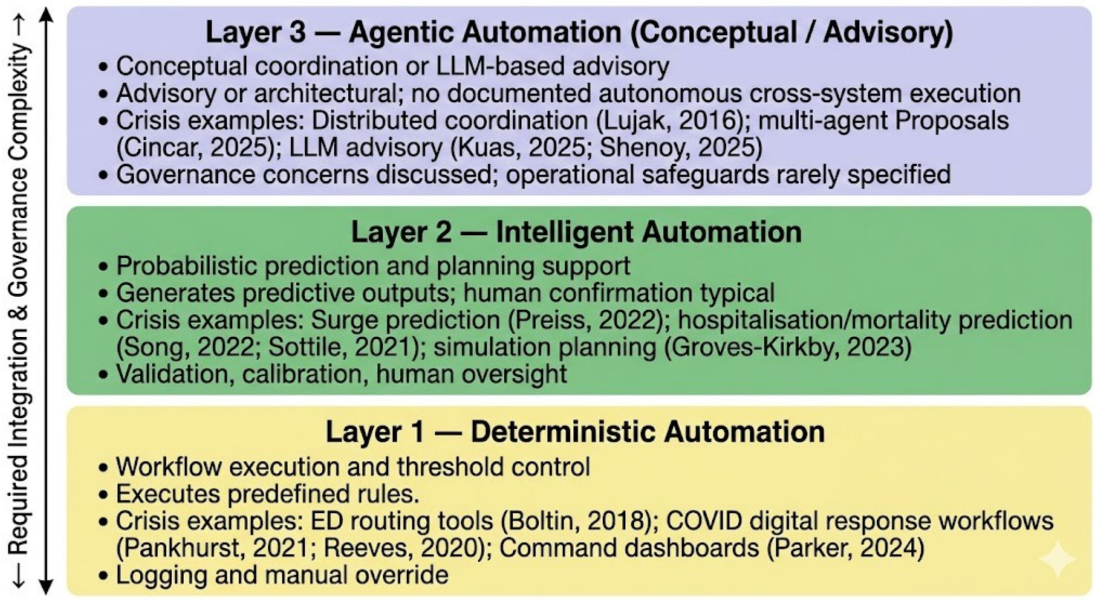

Figure 4. Layered Automation Framework

Deterministic Automation

Deterministic systems are widely used to stabilise workflows. They execute predefined rules such as routing and alerts. These systems are reliable but cannot adapt when conditions change.

Intelligent Automation as Bounded Decision Augmentation

Intelligent systems support prediction and planning. They are used to forecast demand and identify risk. These systems do not take action, and decisions remain with human operators.

Agentic Automation: A Structural Absence

No studies showed systems that can plan and act across multiple systems. Agentic automation remains conceptual in this literature. The main barrier is integration and governance.

Escalating Governance Burden Across Strata

Governance requirements increase as automation becomes more advanced. Deterministic systems require logging and manual override. Intelligent systems require validation and human review. Agentic systems require strict constraints, approval steps, and auditability. This is important for crisis settings, where safety and accountability are critical. Without explicit governance design, higher-autonomy systems risk operational fragility rather than resilience.

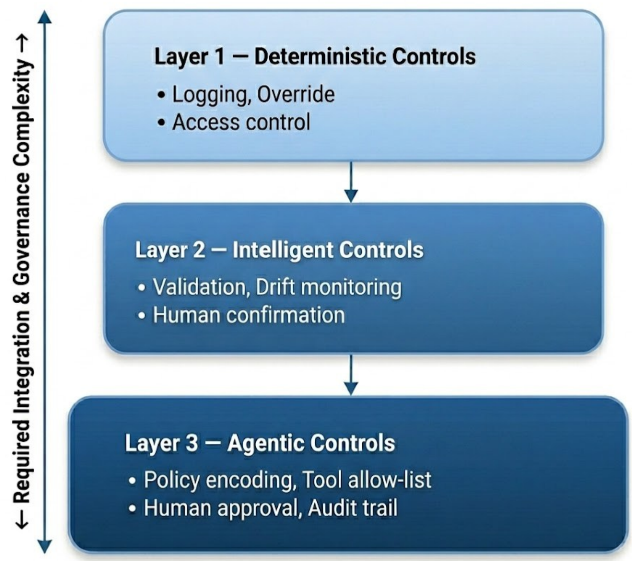

Figure 5. Governance Escalation Model

USE CASE: ICU SURGE MANAGEMENT DURING A RESPIRATORY OUTBREAK

This use case illustrates how the three automation layers identified in Section 5 operate in practice. In particular, this example shows how deterministic, intelligent, and agentic automation differs in a real hospital scenario.

Scenario:

A hospital experiences a rapid increase in patients presenting with respiratory illness. Emergency department (ED) wait times are rising, ICU beds are near capacity, and diagnostic turnaround times are delayed. Clinical and operational teams must coordinate patient flow, staffing, and bed allocation under time pressure.

Deterministic automation: stabilising workflow execution

Deterministic automation supports routine coordination tasks using predefined rules.

Inputs:

1. ED triage category (from patient administration system)
2. Current bed status (bed management system)
3. Threshold rules (e.g., ED wait time > 4 hours)

Actions:

1. Route patients into specialty queues based on triage category
2. Trigger alerts when ED or ICU thresholds are exceeded
3. Generate tasks for bed cleaning and patient transfer
4. Update operational dashboards

Outcome: Manual coordination effort is reduced and workflows become more consistent.

Limits: When multiple constraints conflict (e.g., no ICU beds and limited staff), the system cannot resolve trade-offs.

Intelligent automation: supporting decision-making

Intelligent automation adds predictive capability but does not act independently.

Inputs:

1. Historical admissions and discharge data
2. Current occupancy and ED inflow

3. Staffing availability indicators

Actions:

1. Forecast ICU demand for the next 48–72 hours.
2. Identify patients at high risk of deterioration
3. Generate capacity recommendations (e.g., open surge beds)

Outcome: Teams receive early warning signals and can plan.

Controls: All recommendations are reviewed by clinicians or operational leads before action.

Limits: Predictions may become unreliable as conditions change rapidly during a crisis.

Agentic automation: policy-constrained orchestration (candidate design)

Agentic automation extends beyond prediction to propose coordinated actions across systems.

Goal: Reduce ED congestion and maintain ICU capacity while meeting safety constraints.

System constraints:

The system can:

1. Read occupancy data (bed management API)
2. Read staffing summaries (rostering system)
3. Propose changes to bed allocation
4. Draft communications to teams

The system cannot:

1. Directly reassign beds
2. Change staffing rosters
3. Override clinical decisions

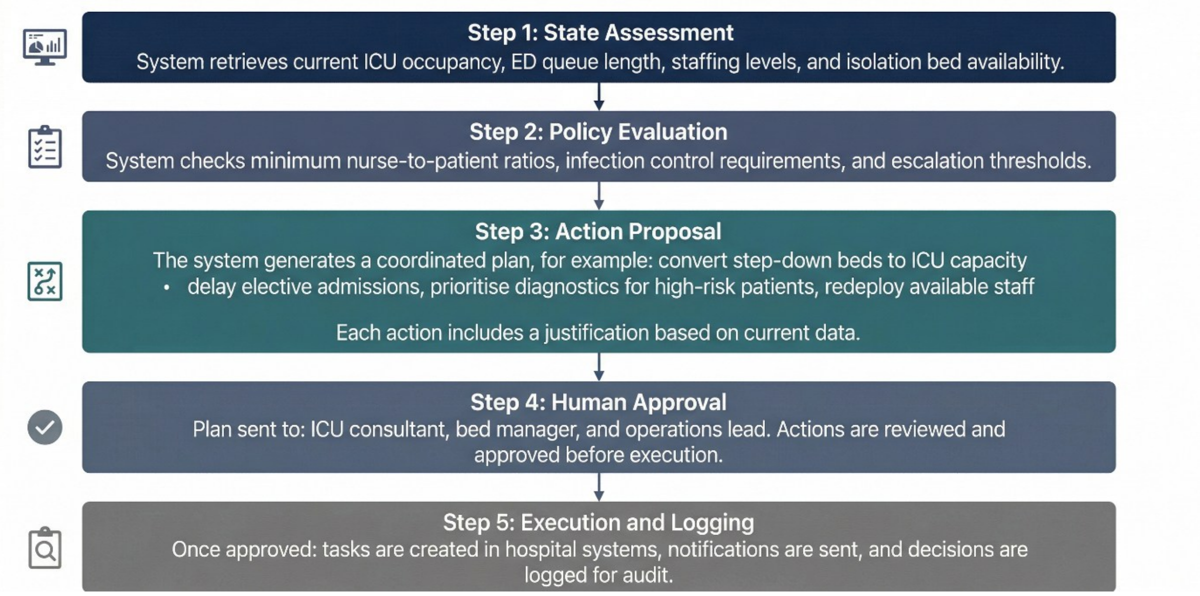

Figure 6. Policy-Constrained Agentic Automation in ICU Surge Management

Outcome: The system supports coordination across departments but does not act independently. All high-impact decisions remain under human control.

Failure containment: if system state is incomplete or contradictory, the agent must refuse to propose actions beyond safe bounds and escalate to human coordination.

This walkthrough makes clear that the agentic layer is not simply “more AI”; it is fundamentally a governance and integration problem. Figure 6 illustrates a five-step coordination workflow: state assessment, policy evaluation, action proposal, human approval, and execution. The figure shows how data inputs, constraints, and human oversight interact across stages. The example reflects the framework's core proposition: increasing coordination capability requires stronger governance and system integration.

DISCUSSION

What the Synthesis Implies

The central implication is that hospital crisis automation is already materially present, but predominantly at deterministic and bounded intelligent levels. The literature evidences substantial effort to stabilise throughput and augment crisis decision-making, but shows limited maturity in governance-aware cross-system orchestration.

This aligns with a broader crisis informatics insight: coordination and accountability, not algorithmic cleverness, are often the binding constraints. In hospitals, fragmented architectures and strong privacy constraints make cross-system automation non-trivial even in non-crisis conditions; in crises, these constraints intensify.

Why Agentic Automation is Largely Absent (in this corpus)

The absence of mature agentic automation can be understood as the result of multiple coupled constraints. First, interoperability constraints remain significant, as many operational systems lack stable APIs and data remains siloed across platforms. Second, governance constraints limit deployment, as high-impact actions require explicit authority and auditability. Third, risk asymmetry in crisis settings means that failures carry disproportionate consequences, leading organisations to favour bounded forms of automation. Fourth, evaluation constraints persist, as crisis contexts are difficult to study prospectively and robust evaluation designs are rare. Finally, crisis response often relies on improvised coordination, which is difficult to formalise into explicit policy constraints required for safe agentic systems.

Design Implications for Crisis-ready Automation

The framework yields several practical design implications for crisis-ready automation. Systems should be designed with fallback in mind, ensuring that deterministic processes and manual runbooks remain available and that higher levels of automation can degrade gracefully under failure conditions. Governance must be treated as an explicit design component, with policies, approval mechanisms, and audit logging embedded as core system features rather than afterthoughts. Autonomy should be constrained by design through the use of tool allow-lists, scoped credentials, and mandatory human approval gates, particularly in high-risk crisis contexts. Finally, integration maturity should be explicitly assessed, as agentic systems depend on interoperable data access and stable operational interfaces; without these, they tend to revert to brittle workarounds such as screen-based automation or produce outputs that are difficult to trace and audit.

Research Implications

For the research community, the corpus suggests several priorities:

1. Evaluation under distribution shift: crisis-time performance varies; methods must address drift and emergent dynamics.
2. Governance frameworks as outcomes: studies should report auditability, override mechanisms, approval workflows, and accountability structures alongside operational metrics.
3. Cross-agency collaboration: in contexts that require inter-agency data sharing; privacy-preserving architectures become central.
4. Agentic safety in operational settings: rather than debating agentic autonomy in the abstract, research should investigate constrained orchestration designs, failure containment, and accountability mechanisms.

Positioning Relative to Prior Work

This paper's distinct contribution is not to claim that agentic automation is already deployed, but to show—through synthesis—that the literature is mature at deterministic and intelligent strata while lacking a coherent governance-aware framework for agentic automation. The layered model offers a way to connect operational crisis automation research to emerging architectures without collapsing governance considerations.

LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

This study has several limitations. First, the review was conducted primarily through EBSCOhost databases (CINAHL Ultimate and MEDLINE Ultimate), supplemented by an AI-assisted discovery search. Although this provides substantial coverage of healthcare and clinical informatics literature, additional databases and citation chaining may identify further relevant studies.

Second, screening decisions were logged within structured spreadsheets rather than dedicated systematic review software. While this does not alter inclusion logic, exclusion counts are reported in aggregate rather than separated into title/abstract and full-text stages.

Third, the included literature is heterogeneous in methodological rigor and reporting depth. As a scoping review, this study maps conceptual and functional patterns rather than conducting quality-weighted effect synthesis.

Fourth, the classification of automation strata involves interpretive synthesis. While guided by explicit definitions, future work should formalise a coding schema and report inter-coder reliability.

Future research should pursue three directions.

1. Empirical validation of the layered framework through case-based mapping of crisis workflows onto deterministic, intelligent, and candidate agentic architectures.
2. Formal governance design studies investigating policy encoding, tool scoping, and approval chain implementation for higher-autonomy systems.
3. Cross-agency interoperability research examining how data protection, API design, and auditability can be aligned in disaster public health contexts.

Such work would move beyond descriptive system reporting toward systematic design science for crisis-ready automation.

CONCLUSION

This paper synthesises hospital crisis automation literature through a scoping review using EBSCOhost databases (PRISMA-ScR) and proposes a layered framework distinguishing deterministic automation, intelligent decision augmentation, and policy-constrained agentic automation. Across 116 included studies, deterministic and intelligent systems dominate crisis response implementations, while mature agentic automation appears largely absent as an operationally evaluated form. The framework clarifies how autonomy escalation increases integration demands, failure risks, and governance requirements—core concerns for disaster public health and healthcare informatics. By making governance and integration constraints explicit, the paper provides a grounded foundation for evaluating emerging agentic automation proposals in crisis settings.

BIBLIOGRAPHY

Bazyar, J., Farrokhi, M., & Khankeh, H. (2019). Triage Systems in Mass Casualty Incidents and Disasters: A Review Study with A Worldwide Approach. *Open Access Macedonian Journal of Medical Sciences*, *7*(3), 482–494. https://doi.org/10.3889/oamjms.2019.119

Boltin, N., Valdes, D., Culley, J. M., & Valafar, H. (2018). Mobile Decision Support Tool for Emergency Departments and Mass Casualty Incidents (EDIT): Initial Study. *JMIR mHealth and uHealth*, *6*(6), e10727. MEDLINE (29934288). https://doi.org/10.2196/10727

CINCAR, K., & IVAŞCU, T. (2025). Artificial Intelligence-Driven Multi-Agent System for Enhanced Radiology Workflow and Decision Support. *Applied Medical Informatics*, *47*, S77–S77. Academic Search Index (188140142).

Comfort, L. K., Ko, K., & Zagorecki, A. (2004). Coordination in Rapidly Evolving Disaster Response Systems: The Role of Information. *American Behavioral Scientist*, *48*(3), 295–313. https://doi.org/10.1177/0002764204268987

Endsley, M. R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *37*(1), 32–64. https://doi.org/10.1518/001872095779049543

Groves-Kirkby, N., Wakeman, E., Patel, S., Hinch, R., Poot, T., Pearson, J., Tang, L., Kendall, E., Tang, M., Moore, K., Stevenson, S., Mathias, B., Feige, I., Nakach, S., Stevenson, L., O'Dwyer, P., Probert, W., Panovska-Griffiths, J., & Fraser, C. (2023). Large-scale calibration and simulation of COVID-19 epidemiologic scenarios to support healthcare planning. *Epidemics*, *42*, 100662. MEDLINE Ultimate (36563470). https://doi.org/10.1016/j.epidem.2022.100662

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, *349*(6245), 255–260. https://doi.org/10.1126/science.aaa8415

Kruk, M. E., Myers, M., Varpilah, S. T., & Dahn, B. T. (2015). What is a resilient health system? Lessons from Ebola. *The Lancet*, *385*(9980), 1910–1912. https://doi.org/10.1016/S0140-6736(15)60755-3

Kuas, C., Canakci, M. E., Acar, N., Kanbakan, A., Cetin, M., & Gunsoy, E. (2025). The Potential and Pitfalls of ChatGPT in Toxicological Emergencies. *Journal of Emergency Medicine*. ScienceDirect (S0736467925002513). https://doi.org/10.1016/j.jemermed.2025.07.002

Lacity, M. C., & Willcocks, L. P. (2018). A New Approach to Automating Services. In Mit Sloan Management Review, *How to Go Digital* (pp. 81–104). The MIT Press. https://doi.org/10.7551/mitpress/11633.003.0015

Lujak, M., Billhardt, H., & Ossowski, S. (2016). Distributed coordination of emergency medical service for angioplasty patients. *Annals of Mathematics and Artificial Intelligence*, *78*(1), 73–100. Springer Nature Journals (edssjs.8FB7A7ED). https://doi.org/10.1007/s10472-016-9507-9

Luke, J., Franklin, R. C., Dyson, J., & Aitken, P. (2023). Building Toward a Disaster Resilient Health System: A Study of Hospital Resilience. *Disaster Medicine and Public Health Preparedness*, *17*, e219. https://doi.org/10.1017/dmp.2022.204

Orlikowski, W. J. (1992). The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science*, *3*(3), 398–427. https://doi.org/10.1287/orsc.3.3.398

Palen, L., Vieweg, S., Liu, S. B., & Hughes, A. L. (2009). Crisis in a Networked World: Features of Computer-Mediated Communication in the April 16, 2007, Virginia Tech Event. *Social Science Computer Review*, *27*(4), 467–480. https://doi.org/10.1177/0894439309332302

Pankhurst, T., Atia, J., Evison, F., Gallier, S., Lewis, J. M., McKee, D., Ryan, S., Sapey, E., Ball, S., & Coleman, J. J. (2021). Rapid adaptation of a local healthcare digital system to COVID-19: The experience in Birmingham (UK). *Health Policy and Technology*, *10*(4). ScienceDirect (S2211883721000915). https://doi.org/10.1016/j.hlpt.2021.100568

Parker, F., Martínez, D. A., Scheulen, J., & Ghobadi, K. (2024). *An Interactive Decision-Support Dashboard for Optimal Hospital Capacity Management* (arXiv:2403.15634). arXiv. https://doi.org/10.48550/arXiv.2403.15634

Preiss, A., Hadley, E., Jones, K., Stoner, M. C. D., Kery, C., Baumgartner, P., Bobashev, G., Tenenbaum, J., Carter, C., Clement, K., & Rhea, S. (2022). Incorporation of near-real-time hospital occupancy data to improve hospitalization forecast accuracy during the COVID-19 pandemic. *Infectious Disease Modelling*, *7*(1), 277–285. MEDLINE Ultimate (35136849). https://doi.org/10.1016/j.idm.2022.01.003

Reeves, J. J., Hollandsworth, H. M., Torriani, F. J., Taplitz, R., Abeles, S., Tai-Seale, M., Millen, M., Clay, B. J., & Longhurst, C. A. (2020). Rapid response to COVID-19: Health informatics support for outbreak management in an academic health system. *Journal of the American Medical Informatics Association*, *27*(6), 853–859. CINAHL Ultimate (143652001). https://doi.org/10.1093/jamia/ocaa037

Romanowski, C., Raj, R., Schneider, J., Mishra, S., Shivshankar, V., Ayengar, S., & Cueva, F. (2015). Regional response to large-scale emergency events: Building on historical data. *International Journal of Critical Infrastructure Protection*, *11*, 12–21. https://doi.org/10.1016/j.ijcip.2015.07.003

Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). *Toolformer: Language Models Can Teach Themselves to Use Tools* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2302.04761

Shenoy, E. S., Branch-Elliman, W., & Wiens, J. (2025). The Search for the Golden Egg: Applied Large Language

Models to Detect Patients With Potential Exposure to Avian Influenza. *Clinical Infectious Diseases*, *81*(5), 904–906. CINAHL Ultimate (190481333). https://doi.org/10.1093/cid/ciaf371

Song, J., Hobensack, M., Bowles, K. H., McDonald, M. V., Cato, K., Rossetti, S. C., Chae, S., Kennedy, E., Barrón, Y., Sridharan, S., & Topaz, M. (2022). Clinical notes: An untapped opportunity for improving risk prediction for hospitalization and emergency department visit during home health care. *Journal of Biomedical Informatics*, *128*, 104039. https://doi.org/10.1016/j.jbi.2022.104039

Sottile, P. D., Albers, D., DeWitt, P. E., Russell, S., Stroh, J. N., Kao, D. P., Adrian, B., Levine, M. E., Mooney, R., Larchick, L., Kutner, J. S., Wynia, M. K., Glasheen, J. J., & Bennett, T. D. (2021). Real-time electronic health record mortality prediction during the COVID-19 pandemic: A prospective cohort study. *Journal of the American Medical Informatics Association*, *28*(11), 2354–2365. CINAHL Ultimate (152975083). https://doi.org/10.1093/jamia/ocab100

Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1; p. NIST AI 100-1). National Institute of Standards and Technology (U.S.). https://doi.org/10.6028/NIST.AI.100-1

Trist, E. L. (1981). *The evolution of socio-technical systems: A conceptual framework and an action research program*. Ontario Quality of Working Life Centre.

Van De Walle, B., & Turoff, M. (2007). Emergency response information systems: Emerging trends and technologies. *Communications of the ACM*, *50*(3), 29–31. https://doi.org/10.1145/1226736.1226760

Van De Walle, B., & Turoff, M. (2008). Decision Support for Emergency Situations. In F. Burstein & C. W. Holsapple, *Handbook on Decision Support Systems 2* (pp. 39–63). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-48716-6_3

Van Der Aalst, W. M. P., Bichler, M., & Heinzl, A. (2018). Robotic Process Automation. *Business & Information Systems Engineering*, *60*(4), 269–272. https://doi.org/10.1007/s12599-018-0542-4

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). *ReAct: Synergizing Reasoning and Acting in Language Models* (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2210.03629