American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Context-Aware Artificial Intelligence for Adaptive Decision Support in Resource-Constrained Institutions

Dr. Liam Anderson

Associate Professor, Northern Ontario Institute of Digital Technology, Toronto, Canada

Abstract

Artificial intelligence has become increasingly relevant to institutional decision support; however, many contemporary systems implicitly assume reliable connectivity, adequate computational capacity, sufficient staffing, stable workloads, and consistently available data. These assumptions are frequently unrealistic in resource-constrained institutions, where decisions must be made under fluctuating bandwidth, limited computing infrastructure, personnel shortages, incomplete information, and sudden changes in service demand. This study proposes a Context-Aware Artificial Intelligence Decision Support System (CAAI-DSS) designed to adapt its decision logic, computational behavior, confidence thresholds, and level of human involvement according to prevailing institutional conditions.

The proposed architecture integrates contextual sensing, lightweight predictive inference, adaptive resource-aware reasoning, confidence-based human escalation, and governance controls within a unified decision-support framework. A reproducible simulation-based evaluation was conducted using 1,500 synthetic decision episodes distributed across five operational scenarios: low bandwidth, restricted computing capacity, staff shortage, demand surge, and mixed-resource constraints. The proposed system was compared with a conventional static artificial intelligence decision-support configuration and a rules-based decision approach. Simulation results indicated that the context-aware model achieved an overall decision accuracy of 78.6%, compared with 61.3% for static artificial intelligence and 51.5% for the rules-based configuration.

The proposed architecture also reduced simulated mean decision latency from approximately 5.18 seconds for static artificial intelligence to 2.05 seconds while selectively transferring uncertain cases to human review. Across individual scenarios, context-aware decision accuracy ranged from 75.0% to 82.0%, demonstrating greater resilience when resource conditions changed. The findings suggest that institutional artificial intelligence should not be treated merely as a predictive model but as an adaptive socio-technical decision infrastructure capable of recognizing the conditions under which a recommendation is generated. The proposed framework contributes a resource-sensitive, human-centered, and governance-aware model for institutions seeking responsible artificial intelligence adoption without depending on permanently high levels of digital infrastructure.

Keywords: context-aware artificial intelligence; adaptive decision support; resource-constrained institutions; human–AI collaboration; edge intelligence; institutional decision-making; resource optimization; explainable AI; responsible AI; adaptive computing

1. Introduction

Artificial intelligence is progressively influencing organizational decision-making by supporting prediction, classification, prioritization, resource allocation, service planning, and operational monitoring. In public institutions and other service-oriented organizations, AI can potentially improve responsiveness and administrative productivity, but its effective use depends heavily on institutional capabilities, including accessible data, workforce competence, technical infrastructure, governance arrangements, and dependable information systems. OECD analyses of public-sector AI emphasize that improved decision-making and service delivery are possible only when enabling conditions such as data management, organizational capability, accountability, and appropriate governance are established.

A central limitation in conventional AI deployment is the assumption that the operational environment surrounding the algorithm remains sufficiently stable. In reality, institutions may experience temporary network disruption, reduced processing capability, staff shortages, incomplete data, equipment limitations, sudden demand surges, or competing demands for scarce resources. These contextual conditions influence not only whether an AI model can operate efficiently but also whether a recommendation is practically feasible. Dey's foundational treatment of context-awareness established that information becomes contextually relevant when it helps characterize the situation surrounding an interaction, while later context-aware computing research emphasized the importance of interpreting computing behavior relative to dynamic environmental conditions rather than treating context as a static background variable.

Resource limitations are particularly important as increasingly sophisticated AI architectures are transferred from centralized computing environments to operational settings. Edge-computing research has demonstrated that placing computation closer to the point where data are produced can reduce dependence on distant infrastructure and improve responsiveness, although constrained memory, processing capacity, network quality, and energy availability continue to influence performance. Contemporary Edge AI studies therefore emphasize model compression, efficient inference, adaptive execution, hardware-aware optimization, and dynamic resource management.

The problem is not exclusively computational. Decisions generated under uncertainty can affect institutional employees, service recipients, budgets, and the allocation of limited resources. Consequently, trustworthy decision support requires appropriate transparency, traceability, risk management, and human oversight. The NIST AI Risk Management Framework explicitly identifies governance, contextual mapping, measurement, and risk management as interconnected functions and recognizes that human roles in AI-supported decision-making should be clearly defined. UNESCO similarly stresses transparency and opportunities for explanation when AI contributes to consequential decisions, while human-centered AI research argues that high automation need not eliminate meaningful human control.

2. Objectives of the Study

2.1 Development of a Context-Sensitive Decision Architecture

The first objective is to develop artificial intelligence decision-support architecture capable of interpreting operational circumstances before generating or implementing recommendations. The framework is designed to treat contextual information such as available computing power, network connectivity, staffing, data completeness, workload intensity, urgency, and resource pressure as active decision variables rather than peripheral system conditions.

2.2 Evaluation of Adaptive Decision Performance

The second objective is to compare the proposed context-aware system with a conventional static AI decision configuration under different institutional constraints. Particular attention is given to decision accuracy and response latency because resource-constrained organizations require systems that remain useful even when the conditions under which they were originally configured deteriorate.

2.3 Integration of Human Oversight and Responsible AI Controls

The third objective is to incorporate human review into adaptive decision-making without requiring manual intervention for every routine recommendation. Low-confidence decisions are selectively deferred, while higher-confidence recommendations remain machine-assisted. This objective is consistent with contemporary human-centered AI principles that favor complementary human–machine relationships and clearly defined responsibility rather than indiscriminate automation. Amershi and colleagues emphasized the importance of supporting appropriate human–AI interaction throughout the lifecycle of intelligent systems, while Shneiderman argued that meaningful human control can coexist with high levels of automation.

3. Review of Literature

3.1 Context-Aware Computing and Adaptive Decision Systems

The conceptual foundation of context-aware computing predates contemporary generative and predictive AI systems. Dey defined context as information capable of characterizing the situation of an entity and established a foundation for systems that change their behavior according to contextual conditions. Dourish subsequently argued that context should not be regarded simply as a predetermined collection of variables; instead, its meaning emerges through interaction and activity. These perspectives remain highly relevant to modern decision-support applications because institutional decisions are rarely independent of environmental conditions, organizational capacity, user roles, urgency, and available resources.

Recent work has broadened the context-aware concept into more complex decision environments. Violos and colleagues reviewed cognition- and context-aware decision-making systems across sustainability-oriented applications and identified their relevance to areas including resource utilization, healthcare, urban systems, and environmental management. Their review demonstrates that contextual reasoning is increasingly connected with intelligent decision infrastructures rather than limited to user-interface personalization.

3.2 Resource-Efficient and Edge Artificial Intelligence

Shi and colleagues described edge computing as a paradigm in which computing capabilities are moved closer to data sources, helping address latency and network-related limitations. Satyanarayanan similarly characterized edge computing as an important development for applications in which responsiveness and proximity to data are critical. These foundations have become increasingly significant as AI inference is moved from centralized infrastructure to local or distributed environments.

Contemporary Edge AI research focuses on the difficulty of executing increasingly complex models under restricted memory, processing, communication, and energy budgets. Sander and colleagues reviewed acceleration techniques for resource-constrained Edge AI and highlighted approaches including model pruning, quantization, knowledge distillation, neural architecture optimization, and deployment-level acceleration. Wang and Jia similarly proposed a data–model–system optimization perspective for reliable and efficient edge deployment. Dong, Mao, and Zhang demonstrated that adaptive early-exit mechanisms can improve the trade-off between inference performance and computational overhead under changing resource and bandwidth conditions.

Resource awareness.

3.3 Human-Centered, Explainable, and Governance-Aware Decision Support

Human–AI collaboration has become a central consideration in decision-support design. Amershi and colleagues proposed interaction guidelines addressing user expectations, correction mechanisms, system uncertainty, and the changing relationship between the user and intelligent system. Shneiderman's human-centered AI framework argues that dependable systems can combine substantial automation with substantial human control rather than treating the two as opposites.

Kovari further emphasized that decision-support effectiveness involves a balance between predictive performance, transparency, and user trust rather than accuracy alone. Recent systematic-review evidence by Fanfani, Ipsaro Palesi, and Nesi similarly reflects increasing attention to architectures that support human-centered, context-sensitive, and interpretable decision assistance.

The governance dimension has become equally important. NIST's AI RMF organizes responsible AI risk activity through Govern, Map, Measure, and Manage functions. UNESCO emphasizes accountability, transparency, human rights, and access to explanations in AI-informed decisions. WHO's AI governance work, although developed for healthcare, likewise emphasizes human autonomy, transparency, accountability, inclusiveness, and responsible governance. More recent adaptive allocation research also demonstrates the value of incorporating governance constraints directly into learning systems rather than applying governance only after model development.

Identified Research Gap

Existing research provides substantial knowledge about context-awareness, Edge AI, model efficiency, human–AI interaction, and trustworthy AI. However, these areas are frequently examined independently. Limited research integrates institutional context detection, resource-sensitive AI inference, operational capacity, confidence-based human escalation, and explicit governance controls into a unified adaptive decision-support architecture for resource-constrained institutions. The present study addresses this intersection.

4. Research Methodology

4.1 Research Design

The study adopted a design-science and simulation-based evaluation approach. The research artifact was the proposed Context-Aware Artificial Intelligence Decision Support System. Because no real institutional dataset, human participants, or organization-specific decision records were supplied for this study, empirical field results were not fabricated. Instead, the evaluation used a reproducible synthetic benchmark containing 1,500 decision episodes.

Five resource-constrained scenarios were modeled, with 300 decision episodes assigned to each scenario. A fixed random seed of 42 was used to improve reproducibility. Each synthetic episode contained normalized contextual variables representing case severity, urgency, data quality, bandwidth availability, computational capacity, staff availability, and institutional demand pressure.

4.2 Proposed Context-Aware AI Framework

The proposed system operates through six logically connected stages:

Context Acquisition → Context Interpretation → Resource Assessment → Adaptive AI Inference → Confidence and Governance Check → Decision or Human Escalation

During context acquisition, the system identifies available operational signals. Context interpretation transforms these inputs into a representation of the current institutional condition. The resource-assessment layer determines whether computing, staff, connectivity, and data conditions are sufficient for normal AI execution.

The adaptive-inference component then modifies decision behavior according to institutional circumstances. In relatively stable situations, the system can perform routine machine-assisted decision support. During resource degradation, it can prioritize lightweight inference and local processing. When prediction confidence falls within an uncertainty region, the system transfers the case to simulated human review.

Table 1: Context Variables and Adaptive Functions in the Proposed CAAI-DSS

Context Variable	Operational Meaning	Adaptive System Response
Bandwidth availability	Reliability of network communication	Preference for local/lightweight processing during weak connectivity
Compute availability	Processing resources currently accessible	Reduction in computational complexity when capacity is limited
Staff availability	Human capacity for review or implementation	Modification of escalation and resource-allocation behavior
Data quality	Completeness and reliability of decision input	Increased caution and uncertainty handling when data quality declines
Demand pressure	Current institutional workload	Dynamic prioritization of higher-value or higher-urgency cases
Severity	Consequence associated with a	Increased priority for potentially consequential

	case	decisions
Urgency	Time sensitivity of the required response	Higher weighting for time-critical decisions

4.3 Comparative Simulation Procedure

Three decision approaches were modelled:

The rules-only DSS applied a fixed heuristic based primarily on case severity and urgency. It represented a low-complexity institutional decision tool with minimal adaptation.

The static AI DSS used severity, urgency, and data-quality information but did not explicitly adapt its decision boundary to staff availability, demand pressure, or changing computational capacity. Simulation noise increased as bandwidth deteriorated, representing the vulnerability of a less context-sensitive architecture to infrastructure instability.

The proposed CAAI-DSS incorporated severity, urgency, institutional demand, data quality, staffing conditions, and computational capacity into the decision function. Its simulated predictions included lower environmental degradation than the static configuration because the architecture was modeled as locally adaptive.

An uncertainty interval was defined around the decision boundary. Predictions close to that boundary were transferred to simulated human review. Human-review accuracy was set at 90% for benchmark purposes. This is a modeling assumption rather than a measured estimate of actual institutional staff performance and should be recalibrated in future field studies.

5. Analysis and Results

5.1 Overall Comparative Performance

Across all 1,500 simulated decision episodes, the proposed CAAI-DSS achieved 78.6% decision accuracy, compared with 61.3% for static AI and 51.5% for the rules-only decision configuration. The performance difference indicates that including institutional context in the decision function can be advantageous when the underlying environment is intentionally designed to fluctuate.

The proposed system also produced an average simulated latency of approximately 2.05 seconds, whereas the static AI configuration averaged approximately 5.18 seconds. The rules-only approach was computationally simpler and therefore faster, but its lower decision accuracy illustrated the limitation of optimizing system simplicity without sufficient contextual reasoning.

The confidence-sensitive mechanism transferred approximately 26.5% of proposed-model episodes to simulated human review. This finding should not be interpreted as an optimal real-world deferral rate. Rather, it demonstrates how selective escalation can be incorporated as an adjustable architectural parameter.

5.2 Scenario-Wise Performance

Table 2: Simulation-Based Performance across Resource-Constraint Scenarios

Institutional Scenario	Static AI Accuracy (%)	Proposed CAAI-DSS Accuracy (%)	Static AI Mean Latency (s)	Proposed Mean Latency (s)
Low bandwidth	63.7	75.0	6.12	1.74
Compute constrained	57.0	82.0	4.74	2.32
Staff shortage	52.7	76.0	4.29	2.05
Demand surge	68.7	80.3	4.30	1.75
Mixed constraints	64.3	79.7	6.45	2.38

5.3 Interpretation of the Performance Pattern

The simulation suggests that context-sensitive adaptation is most valuable when the decision environment differs from the assumptions embedded in a static system. Static AI performed comparatively better during the demand-surge scenario than during staffing or computational restrictions, but its behavior remained less consistent across conditions.

The proposed architecture reduced this sensitivity by explicitly incorporating operating conditions into decision logic. The result supports the broader Edge AI argument that systems deployed under resource constraints require adaptive execution rather than assuming permanently available infrastructure. Resource-aware Edge AI research similarly emphasizes that the optimal balance among performance, model complexity, latency, and available resources can change across deployment conditions.

Figure 1: Simulation-Based Decision Accuracy across Institutional Constraint Scenarios

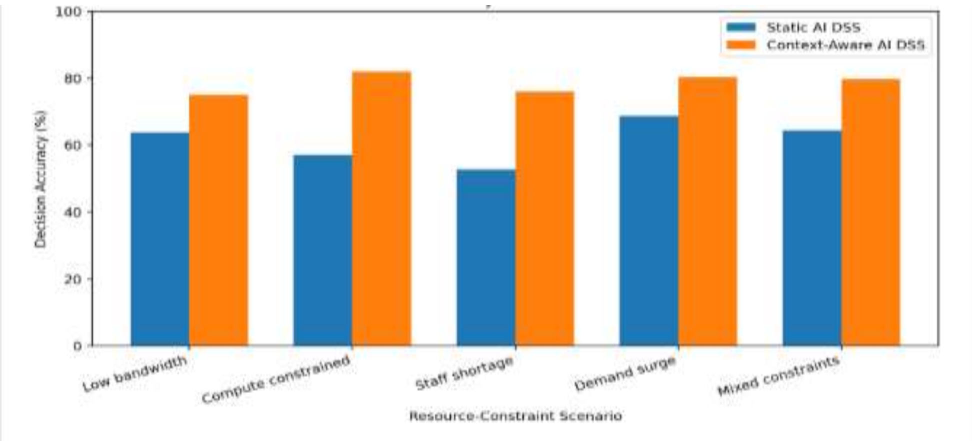

6. Discussion

6.1 Context Awareness as an Institutional Capability

Key Points:

- Context awareness should be treated as an essential institutional capability rather than only a technical feature.
- Effective decision support depends on understanding both the problem and the institution's actual capacity to respond.
- Resource-constrained institutions need systems that can function effectively within available infrastructure and operational limits.
- Decision-making tools should adjust to changes in connectivity, computing capacity, staffing, workload, and data availability.
- High computational complexity is not always necessary for achieving useful and timely institutional outcomes.
- System performance should be evaluated under realistic operating conditions, not only under ideal laboratory settings.
- Adaptability should be considered an important requirement during technology selection and implementation.
- Institutions should prioritize solutions that remain reliable during staff shortages, demand surges, weak connectivity, and incomplete data.
- Operational feasibility should receive the same attention as technical performance.
- A context-sensitive approach can improve institutional resilience, responsiveness, and resource utilization.

6.2 Human Oversight Without Universal Manual Review

One of the important characteristics of the proposed model is selective human escalation. Human oversight is often discussed as though every AI recommendation must either be completely automated or manually verified. In practice, such a binary design can create unnecessary workload and undermine the efficiency gains expected from AI.

A context-aware approach allows human involvement to be dynamically determined according to uncertainty, risk, decision consequence, and institutional capacity. NIST recognizes that AI systems may operate across a continuum ranging from autonomous processing to human decision-making supported by AI and states that responsibilities for oversight should be differentiated. Shneiderman's human-centered AI model similarly demonstrates that meaningful human control and advanced automation are not inherently incompatible.

6.3 Governance, Practical Deployment, and Limitations

Trustworthy institutional AI requires more than computational adaptation. UNESCO's AI ethics framework emphasizes transparency, accountability, human rights, and the ability to obtain explanatory information regarding AI-informed decisions. The NIST AI RMF similarly frames governance as a continuing organizational responsibility. OECD guidance concerning AI in public institutions also stresses that benefits depend on trustworthy implementation, organizational capacity, and appropriate accountability.

Accordingly, the proposed CAAI-DSS includes a governance layer that conceptually determines which decisions can remain automated, which require confirmation, and which must be transferred to an authorized human decision-maker. In future institutional implementation, this layer could include financial authorization thresholds, legal restrictions, professional competence requirements, privacy rules, safety conditions, and organization-specific accountability policies.

7. Conclusion

This study proposed a Context-Aware Artificial Intelligence Decision Support System for institutions operating under variable and constrained resource conditions. Unlike static decision-support architectures that primarily optimize predictive output, the proposed framework considers bandwidth, computing capacity, staffing, data quality, demand pressure, urgency, severity, prediction confidence, and governance constraints as interconnected components of decision quality. A reproducible simulation involving 1,500 synthetic decision episodes demonstrated that the proposed system achieved 78.6% overall accuracy compared with 61.3% for static AI while also reducing modeled response latency. Performance remained comparatively stable across low-bandwidth, compute-constrained, staff-shortage, demand-surge, and mixed-constraint scenarios.

These findings do not substitute for institutional field validation, but they demonstrate the technical and conceptual value of designing AI around the conditions in which decisions actually occur. The principal contribution of the study is therefore a shift from resource-dependent artificial intelligence toward resource-aware artificial intelligence. For institutions with limited digital infrastructure, personnel, or operational flexibility, the future of effective decision support may depend less on maximizing model size and more on developing systems that recognize constraints, adapt intelligently, preserve meaningful human authority, and remain accountable under changing real-world conditions.

References

1. K. Dey, “Understanding and using context,” Personal and Ubiquitous Computing, vol. 5, pp. 4–7, 2001, doi: 10.1007/s007790170019.
2. P. Dourish, “What we talk about when we talk about context,” Personal and Ubiquitous Computing, vol. 8, no. 1, pp. 19–30, 2004, doi: 10.1007/s00779-003-0253-8.
3. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge computing: Vision and challenges,” IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, 2016.
4. M. Satyanarayanan, “The emergence of edge computing,” Computer, vol. 50, no. 1, pp. 30–39, 2017, doi: 10.1109/MC.2017.9.
5. S. Amershi et al., “Guidelines for human-AI interaction,” in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019, pp. 1–13, doi: 10.1145/3290605.3300233.
6. Shneiderman, “Human-centered artificial intelligence: Reliable, safe & trustworthy,” International Journal of Human–Computer Interaction, vol. 36, no. 6, pp. 495–504, 2020, doi: 10.1080/10447318.2020.1741118.
7. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023.

8. UNESCO, Recommendation on the Ethics of Artificial Intelligence. Paris, France: United Nations Educational, Scientific and Cultural Organization, 2021.
9. World Health Organization, Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. Geneva, Switzerland: WHO, 2021.
10. G. S. Tegenaw, D. Amenu, G. Ketema, F. Verbeke, J. Cornelis, et al., “Evaluating a clinical decision support point of care instrument in low resource setting,” BMC Medical Informatics and Decision Making, vol. 23, art. 51, 2023, doi: 10.1186/s12911-023-02144-0.
11. R. Dong, Y. Mao, and J. Zhang, “Resource-constrained Edge AI with early exit prediction,” Journal of Communications and Information Networks, vol. 7, no. 2, 2022.
12. S. S. Gill et al., “Edge AI: A taxonomy, systematic review and future directions,” arXiv:2407.04053, 2024.
13. Organisation for Economic Co-operation and Development, G7 Toolkit for Artificial Intelligence in the Public Sector. Paris, France: OECD Publishing, 2024.
14. J. Sander, A. Cohen, V. R. Dasari, B. Venable, and B. Jalaian, “On accelerating Edge AI: Optimizing resource-constrained environments,” arXiv:2501.15014, 2025.
15. X. Wang and W. Jia, “Optimizing Edge AI: A comprehensive survey on data, model, and system strategies,” arXiv:2501.03265, 2025.
16. J. Violos, G. Mamanis, I. Kompatsiaris, and S. Papadopoulos, “Cognition and context-aware decision-making systems for a sustainable planet: A survey on recent advancements, applications and open challenges,” Discover Sustainability, vol. 6, art. 235, 2025.
17. Organisation for Economic Co-operation and Development, Governing with Artificial Intelligence. Paris, France: OECD Publishing, 2025.
18. M. Fanfani, L. A. I. Palesi, and P. Nesi, “Human-centered AI for decision support systems: A systematic review of application domains, architecture designs, current trends and future directions,” Big Data and Cognitive Computing, vol. 10, no. 6, art. 186, 2026, doi: 10.3390/bdcc10060186.
19. V. Pelechano, A. Mestre, M. Albert, and M. Gil, “HAAS: A policy-aware framework for adaptive task allocation between humans and artificial intelligence systems,” arXiv:2605.02832, 2026.