Agentic AI in Public Sector Applications: A Multi-agent AI Planning Framework

Illugi T. Hjaltalín
www.stirna.is

October 1, 2025

1. Introduction: Analyzing Agentic AI through Public Sector Case Studies

This paper analyzes nine comprehensive case studies of agentic AI implementation in public and quasi-public sector contexts to examine how multi-agent systems transform service delivery while addressing critical challenges of public value creation. The cases span healthcare administration (Mass General Brigham, 2024), social services automation (UiPath, 2022), government research enhancement (World Bank DIME, 2024), NGO impact measurement (AI for Impact, 2025), academic research acceleration, public health surveillance (Boston University, 2025), financial risk management (Xu et al., 2024), transportation optimization (Chen & Ding, 2024), and legal document processing.

Building on the public values framework by Hjaltalin and Sigurdarson (2024), this analysis addresses three research questions: (1) How do agentic AI systems conceptualize and measure public value creation across different domains? (2) What contextual factors most significantly influence the social impact of AI initiatives? (3) How do implementation strategies align with public sector objectives and ethical considerations? The case studies provide empirical evidence for examining conceptual convergence in AI-driven public value creation.

2. Analytical Framework: Integrating Case Studies with AI Planning Architecture

The AI planning framework depicted in Figure 1 provides a structured approach for analyzing the nine case studies, with each implementation mapping to specific components of the three-module architecture: data acquisition, algorithmic modeling, and policy implementation with feedback mechanisms.

The framework's architecture, which emphasizes balancing computational efficiency with democratic accountability (Hjaltalin & Sigurdarson, 2024), is exemplified across the case studies. Mass General Brigham's implementation (Mass General Brigham, 2024) demonstrates the data acquisition module through ambient listening agents and the policy implementation component through validation agents that address hallucination prevention-a critical concern raised in the reflection prompts. Their clinical-trial approach to deployment reflects the framework's emphasis on iterative validation and gradual scaling.

The New York Foundling case (UiPath, 2022) illustrates complete framework integration: robotic-agent crews handle data processing, NLP triage agents perform algorithmic modeling, and sequencing agents ensure compliance in policy implementation. This addresses the reflection prompt regarding additional KPIs beyond hours-saved, showing how comprehensive framework implementation enables tracking of data accuracy (95%) and staff retention improvements (turnover reduced from 42% to 17%).

The World Bank's ImpactAI system (World Bank DIME, 2024) exemplifies sophisticated algorithmic modeling with retrieval agents processing embedded reports, reasoning agents extracting causal estimates, and critique agents flagging bias-directly addressing reflection prompts about assurances for confidential data extraction and FedRAMP compliance. This implementation reduced evidence-synthesis time from 6 weeks to 48 hours while maintaining rigorous data protection standards.

3. Comparative Analysis: Thematic Findings from Multi-Domain Implementation

The comparative analysis of nine agentic AI implementations reveals consistent patterns in public value creation while highlighting domain-specific challenges and solutions. Table 1 provides a structured

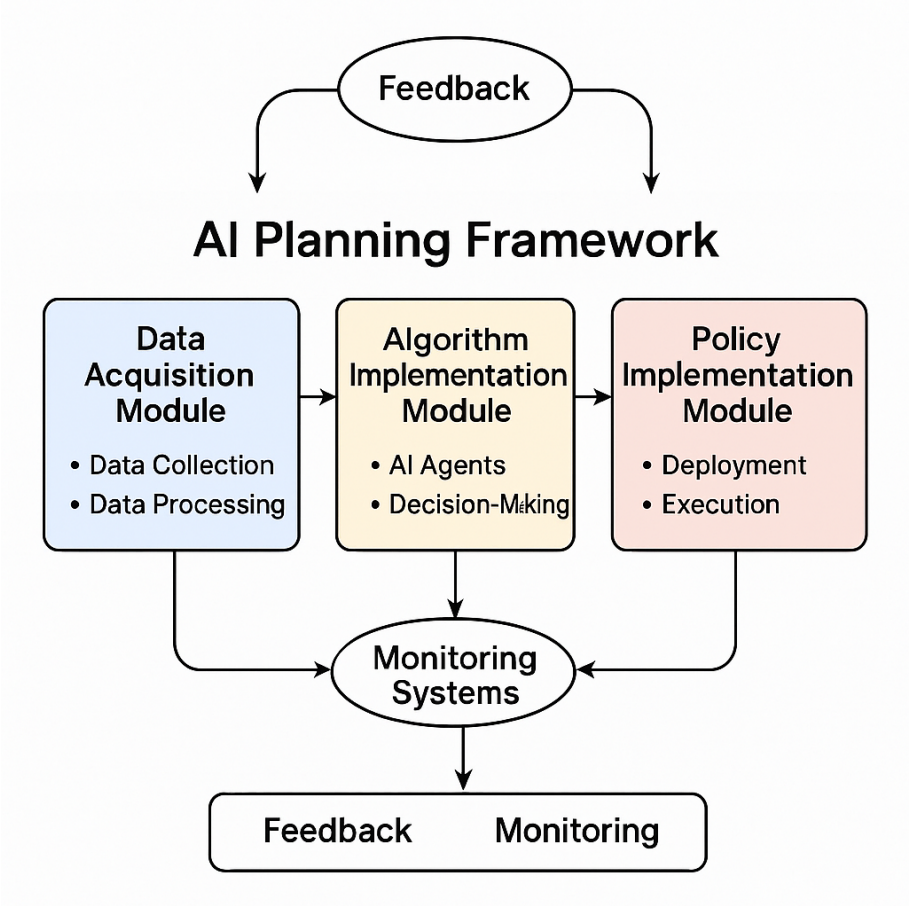

Figure 1: Multi-agent AI Planning Framework for Public Sector Applications

overview of architectural components, outcomes, and critical considerations across all cases.

Conceptual Convergence in Value Measurement

The cases demonstrate remarkable consistency in how public value is conceptualized and measured. Efficiency metrics (time savings from 30-70%, cost reductions) are consistently paired with quality indicators (accuracy rates of 85-95%) and human benefits (burnout reduction, increased trust scores). This triangulation of metrics addresses reflection prompts about comprehensive impact measurement beyond simple efficiency gains.

The strategic approaches also converge around risk-aware implementation. Mass General Brigham's clinical-trial deployment model, rather than big-bang rollout, reflects widespread recognition that responsible AI adoption requires careful validation and gradual scaling. Similarly, the emphasis on guardrails against hallucinations, bias, and false positives across multiple cases indicates shared understanding that reliability and safety are fundamental to public value creation.

Contextual Factors in Social Impact

The social impact of AI initiatives shows strong dependence on organizational context and human factors. The New York Foundling case particularly illustrates how AI addressed profound human challenges-42% staff turnover resulting from repetitive data entry-by enabling more meaningful client engagement. This demonstrates that the most significant impacts often relate to human experience improvements rather than purely technical metrics.

Table 1: Comparative Analysis of Agentic AI Case Studies: Architecture, Outcomes, and Implementation Challenges

Application Domain	**Agent Architecture**	**Key Outcomes**	**Critical Considerations**
Healthcare Administration	Listening, summarizer, validation agents	30% time reduction, 20% burnout decrease	Hallucination prevention, clinical validation
Social Services Automation	Robotic crews, NLP triage, sequencing agents	16h/month saved, 42%→17% turnover	Data accuracy (95%), compliance ordering
Government Research	Retrieval, reasoning, critique agents	6w→48h synthesis, $3M savings	Confidential data handling, FedRAMP compliance
NGO Impact Measurement	Data collector, validation, dashboard agents	70% reporting time reduction	Bias mitigation in qualitative data collection
Academic Research	Retrieval, summarization, auto-drafting agents	40% preparation time reduction	Ethical concerns in AI-generated citations
Public Health Surveillance	NLP, forecasting, human-vetting agents	3-day early detection, 0.86 precision	False positive mitigation, privacy implications
Financial Risk Management	Anomaly detection, predictive, compliance agents	35% accuracy improvement	Model interpretability, Basel III compliance
Transportation Optimization	Sensor fusion, prediction, control agents	25% travel time reduction	Equity concerns, special event anticipation
Legal Document Processing	Ingestion, ranking, screening, summarizer agents	50% review speed increase	Privilege screening, court filing policies

Reflection prompts reveal critical contextual considerations: transportation optimization raises equity concerns about whose mobility needs are prioritized; NGO data collection must address potential biases in qualitative outcome stories; and social media monitoring for public health surveillance creates privacy implications that require careful ethical balancing.

Strategic Alignment with Public Objectives

The implementations consistently align with Hjaltalin and Sigurdarson (2024)'s framework through their emphasis on hybrid intelligence models that combine computational power with human oversight. Multi-agent architectures incorporate validation steps, human veto points, and continuous feedback loops-all reflecting the balance between innovation and accountability.

The cases also demonstrate strategic alignment with broader public objectives: service delivery improvement (healthcare, social services), decision-making enhancement (research, risk management), resource optimization (transportation, legal processing), and public trust strengthening (impact measurement, surveillance). However, challenges remain regarding model interpretability for regulatory compliance, ethical AI-generated content, and equitable service distribution-highlighting ongoing tension between technological capability and public sector values.

Bibliography

AI for Impact. (2025, May). *AI agents for NGOs: From intention to impact*. Retrieved August 23, 2025, from %5Curl%7Bhttps://mzninternational.com/ai-agents-for-ngos-from-intention-to-impact/%7D

Boston University. (2025, April). *Open source AI infectious diseases monitoring tool*. BU Today. Retrieved August 23, 2025, from %5Curl%7Bhttps://www.bu.edu/articles/2025/open-source-ai-infectious-diseases-monitoring-tool/%7D

Chen, H., & Ding, Y. (2024). Implementing traffic agent based on langgraph. *Proceedings of the Fourth International Conference on Intelligent Traffic Systems and Smart City (ITSSC 2024)*. https://doi.org/10.1117/12.3050639

Hjaltalin, I. T., & Sigurdarson, H. T. (2024). The strategic use of AI in the public sector: A public values analysis of national AI strategies. *Government Information Quarterly*, *41*(1), 101914. https://doi.org/10.1016/j.giq.2024.101914

Mass General Brigham. (2024, July). *How AI can help providers listen to their patients better*. Mass General Brigham. Retrieved August 23, 2025, from %5Curl%7Bhttps://www.massgeneralbrigham.org/en/about/newsroom/articles/ai-ambient-documentation%7D

UiPath. (2022). *NY foundling saves 16 hours per clinician* [Case study]. Retrieved August 23, 2025, from %5Curl%7Bhttps://www.uipath.com/resources/automation-case-studies/new-york-foundling-ngo-rpa%7D

World Bank DIME. (2024, May). *World bank DIME AI team overview*. Retrieved August 23, 2025, from %5Curl%7Bhttps://www.worldbank.org/en/about/unit/unit-dec/impactevaluation/ai%7D

Xu, H., Niu, K., Lu, T., & Li, S. (2024). Leveraging artificial intelligence for enhanced risk management in financial services: Current applications and future prospects [Open Access under CC BY-NC 4.0. Article received: 25-02-24, accepted: 06-06-24, published: 02-08-24]. *Engineering Science and Technology Journal*, *5*(8), 2402–2426. https://doi.org/10.51594/estj.v5i8.1363