INTI JOURNAL | Vol.2025, Issue 3, No.1
eISSN:2600-7320

Optimizing Urban Mobility in Hangzhou: A Case Study of the City Brain's AI-Driven Traffic Management

Guo Hanxiang[1*], Leong Wai Yie[2]

[1]Faculty of Business and Communications, INTI International University, Malaysia.
[2]Faculty of Engineering and Quantity Surveying, INTI International University, Malaysia.

[*]**Email**: i25034188@student.newinti.edu.my

Abstract

This study examines Hangzhou's City Brain as an AI-enabled traffic governance platform. Using Leong & Kumar (2023) four-dimensional ITS framework—data acquisition, connectivity, intelligence, and responsiveness, the paper evaluates operational outcomes, governance conditions, and transferability. We find that (i) average traffic efficiency improved in pilot corridors and (ii) emergency response times shortened markedly, with (iii) gains shaped by a public–private partnership that couples municipal mandates with cloud-scale analytics. However, challenges persist around data governance and public trust, interoperability with legacy ITS, and context-dependent scalability. Comparative references to Singapore and Amsterdam underscore how institutional design conditions technological payoffs. The case contributes practice-oriented insights for cities seeking reproducible, ethically governed AI in transport.

Keywords

Smart City, Urban Traffic, City Brain; Traffic Optimization, Intelligent Transportation Systems.

Introduction

Hangzhou's "City Brain" integrates multi-source urban data with real-time analytics to orchestrate traffic management at city scale. This study applies the Intelligent Transportation System (ITS) framework of Leong and Kumar (2023) as an analytical lens, examining how sensing, connectivity, machine intelligence, and operational responsiveness translate into measurable mobility outcomes. The research is guided by two questions: (Q1) To what extent does City Brain improve traffic efficiency and emergency responsiveness? (Q2) Under what governance conditions are these improvements replicable? In addition, the study outlines implications for citizen trust, transparency, and long-term sustainability.

Submission: 6 July 2025; **Acceptance:** 28 August 2025; **Available online:** September 2025

Methodology

This study employs a qualitative case study approach based on triangulated secondary sources, including annual reports from the Hangzhou Municipal Bureau of Traffic, Alibaba Cloud's City Brain documentation, and peer-reviewed literature. To address the limitations of secondary data, key indicators are cross-validated across at least two independent sources, consistency checks are conducted over time using pre/post baselines, and evidence is coded into an ITS analytic matrix. The matrix maps findings onto four dimensions: data acquisition (sensor, camera, and GPS feeds), connectivity (platform integration and data pipelines), intelligence (AI models for detection and forecasting), and responsiveness (signal priority and incident clearance). This process demonstrates explicitly how Leong's model structures the analysis, rather than serving merely as a citation.

Results and Discussion

Hangzhou's outcomes are interpreted through Leong & Kumar (2023) four-dimensional ITS framework—data acquisition, connectivity, intelligence, and responsiveness—to address two research questions: (Q1) To what extent has City Brain improved operational performance? (Q2) Under what governance conditions are such gains reproducible?

As summarized in Table 1, post-deployment gains are evident in average corridor speed and emergency response, reflecting the combined effects of acquisition, connectivity, intelligence, and responsiveness.

Table 1. Traffic Performance Before and After City Brain Implementation. Sources: Hangzhou Municipal Bureau of Traffic (2019); Alibaba Cloud (2021).

Indicator	Before	After	Improvement
Avg Speed (Main Roads)	22.1 km/h	25.5 km/h	+15%
Emergency Response Time	12 min	6 min	-50%
Incident Clearance Time	30 min	<20 min	-33%
Violation Detection Accuracy	<40%	>90%	Significant gain

Residents report more predictable commutes and improved access to emergency services. The City Brain platform also supports mobile app access for real-time updates, enhancing public decision-making.

While the system excels in operational responsiveness, its success is contingent upon strong institutional support and technological readiness. Hangzhou's model, driven by Alibaba Cloud, benefits from an advanced digital infrastructure and centralized policy environment—conditions that may not exist in other cities. Figure 1 visualizes the control dashboard for real-time monitoring and alerts, providing operators with actionable cues for priority signaling and incident dispatch.

Figure 1. City Brain Traffic Dashboard (schematic, adapted from Alibaba Cloud, 2021). Caption: Real-time flows, incident alerts, and operator prompts are co-displayed to inform interventions. Source: Alibaba Cloud (2021).

The observed ~15% uplift in main-road speed and ~50% reduction in emergency response time aligned with four reinforcing mechanisms. First, acquisition expands coverage and granularity (cameras, GPS, IoT probes), reducing blind spots. Second, connectivity consolidates heterogeneous data via platform pipelines, enabling low-latency fusion. Third, intelligence (anomaly detection, short-horizon forecasting) converts into actionable insights. Fourth, responsiveness operationalizes those insights through green waves, pre-emption for ambulances, and incident dispatch. Gains appear strong on signalized arterials with stable demand patterns, while saturated bottlenecks display diminishing marginal returns—a reminder that algorithmic coordination complements but does not replace capacity management.

Where pre/post windows cover multiple quarters, effects persist beyond initial tuning, though peak-hour volatility remains higher than off-peak. Corridor-level heterogeneity suggests that data quality and controller actuation limits condition the achievable ceiling.

City Brain shortens the sense–decide–act loop. On the sense side, multi-source feeds raise recall for non-recurrent incidents (stalls, minor crashes). On the decide side, learned policies prioritize phase splits and offsets, balancing throughput with delay variance. On the act side, operator prompts and automated strategies (e.g., emergency vehicle priority and detour advisories) narrow the gap between detection and field execution. The operations room and visualization

INTI JOURNAL | Vol.2025, Issue 3, No.1
eISSN:2600-7320

wall—shown in Figure 2—enable cross-agency situational awareness, which shortens the sense–decide–act loop.
Figure 2 outlines the operations room setting, enabling cross-agency coordination and network-wide situational awareness.

Figure 2. Central Control Room & Visualization Wall (schematic)
Source*:* Hangzhou Municipal Bureau of Traffic (2019).

Crucially, these mechanisms are co-dependent: intelligence without reliable acquisition degrades; responsiveness without organizational readiness stalls.

Table 2 contrasts institutional setups and citizen interfaces in Hangzhou, Singapore, and Amsterdam, highlighting how governance design shapes technological payoffs and transferability.

Table 2. Comparative Features of Smart Traffic Systems: Hangzhou, Singapore, Amsterdam.
Sources: Alibaba Cloud (2021); Government Technology Agency (2018); van Winden & van den Buuse (2017).

Feature	Hangzhou	Singapore	Amsterdam
Leading Agency / Operator	Public–Private Partnership (Hangzhou Municipal Govt + Alibaba Cloud)	Government-led (LTA + GovTech)	Municipality–University consortium（pilot-oriented）
Technical Focus	Real-time optimization & incident response	Predictive coordination & centralized planning	Sustainable mobility & living-lab pilots
Citizen Interface	City services via Alipay/municipal apps	myTransport.SG official app	Co-creation platforms & pilot participation

Key Strength	Speed & AI scaling	Institutional stability & integration	Citizen engagement & experimentation

Hangzhou's public–private model (municipal mandate + cloud-scale analytics) optimizes for speed and scaling; Singapore's government-led integration favors predictability and whole-of-government coordination; Amsterdam's living-lab approach maximizes experimentation and citizen co-creation. Consequently, what transfers is not a turnkey "system" but a design pattern: align institutional capacity, data rights, and procurement with the chosen optimization objective.

Governance, trust, and data stewardship

City-scale optimization heightens expectations for transparency, auditability, and data minimization. Publishing performance dashboards, clarifying data retention policies, and instituting third-party audits can mitigate "black box" concerns. Clear delineation of roles between municipal authorities (public interest, oversight) and platform providers (infrastructure, models) reduces ambiguity in accountability and incident response.

Scalability and interoperability

Scaling beyond pilot corridors requires: (i) interoperability with legacy controllers and heterogeneous SCATS-like deployments; (ii) lifecycle budgeting for data engineering and model re-training; (iii) vendor-neutral interfaces to limit lock-in and ease module replacement. Performance plateaus when sensing density, controller latency, or right-of-way constraints cap the feasible control envelope—pointing to the need for joint planning with physical network upgrades.

Equity and externalities

While average speeds improve, distributional effects merit scrutiny. Priority strategies should safeguard public transport reliability, non-motorized safety, and access for emergency services without displacing congestion to residential streets. Continuous monitoring for model drift and bias (e.g., under-reporting in low-sensor neighborhoods) is essential to sustain legitimacy.

Limitations and validity

Findings rely on triangulated secondary sources (municipal reports, provider documentation, peer-reviewed studies). We mitigate risks of measurement error and time misalignment through cross-source consistency checks and pre/post baselines. Nonetheless, unobserved confounders (construction, seasonal demand shifts) may influence effect sizes; city-wide randomized rollouts are infeasible, so causal claims remain cautious.

The policy implications are clear. Cities pursuing AI-enabled traffic management should lead with governance by design, codifying data rights, audit processes, and public communication alongside technical rollout. They should also procure for the lifecycle by budgeting for data plumbing, MLOps, and interoperability rather than focusing only on initial licenses. Equally important is the definition of KPIs and re-evaluation cadences to track speed variance, incident clearance, public transport reliability, and equity metrics. Finally, building organizational capacity is essential, embedding analytics and operations teams to close the loop between model insight and field practice.

The takeaway is that Hangzhou's case demonstrates measurable operational benefits when AI, data infrastructure, and institutional capacity co-evolve. Transferability hinges less on copying artifacts and more on reproducing these enabling conditions.

Conclusion

City Brain demonstrates that AI-assisted signal orchestration and incident management can deliver measurable gains in speed and emergency responsiveness under enabling governance. Yet durability depends on transparent data stewardship and public trust, interoperability with legacy ITS, and the capacity to adapt models as conditions evolve. For policy, cities should align governance design with technical ambition—not vice versa—and budget for continuous integration rather than one-off deployments. Future research should test transferability across diverse regulatory contexts and quantify long-term social equity effects.

Acknowledgements

The author wishes to thank Alibaba Cloud and the Hangzhou Municipal Government for providing access to publicly available data.

References

Alibaba Cloud. (2021). *Hangzhou City Brain: Project overview*. Alibaba Cloud. https://www.alibabacloud.com/solutions/smart-city

Government Technology Agency (GovTech). (2018). *Intelligent transport systems in Singapore*. Government Technology Agency. https://www.tech.gov.sg/products-and-services/ITS/

Guo, H., & Leong, W. Y. (2025). Comparative review of AI applications in urban transport: Insights from China's City Brain and Singapore's LTA Smart Mobility. *Journal of Information Technology*, *2025*(07). https://doi.org/10.61453/joit.v2025no07

Hangzhou Municipal Bureau of Traffic. (2019). *Annual report on traffic efficiency in Hangzhou*. http://hzjt.gov.cn

Leong, W. Y., & Kumar, R. (2023). Smart city 5G intelligent transportation system. In *Convergence of IoT, blockchain, and computational intelligence in smart cities* (pp. 1–22). CRC Press. https://doi.org/10.1201/9781003353034-1

van Winden, W., & van den Buuse, D. (2017). Smart city pilot projects: Exploring the dimensions and conditions of scaling up. *Journal of Urban Technology*, *24*(4), 51–72. https://doi.org/10.1080/10630732.2017.1348884