Preprints.org

Article

Not peer-reviewed version

# From Traffic Indices to Agentic AI: A Reference Framework for Urban Mobility in Megacities

Tatiana Petrova *

Posted Date: 11 May 2026

doi: 10.20944/preprints202604.1407.v2

Keywords: agentic AI; multi-agent systems; reference architecture; urban mobility; transport accessibility; PTAL; composite traffic index; smart cities; London

Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

*Article*

# From Traffic Indices to Agentic AI: A Reference Framework for Urban Mobility in Megacities

**Tatiana Petrova**

SnT, University of Luxembourg; tatiana.petrova@uni.lu

**Abstract**

Megacity mobility research has long relied on aggregate statistical indicators (composite traffic-planning indices and accessibility surfaces) that capture the system at rest but are disconnected from the operational decisions shaping mobility minute to minute. In parallel, Agentic AI offers a reference paradigm of semantically interoperable autonomous agents that negotiate and coordinate across open networks. We propose a unifying Agentic AI reference architecture for urban transportation that maps any composite traffic index and any accessibility surface onto agent utility functions, negotiation-protocol primitives, and shared semantic ontologies. The architecture is instantiated in simulation only: no live agent-to-agent endpoints, no runtime large language models, no cross-organisation interoperability experiment; these are explicitly listed as next steps. A 12-borough London case study, evaluated over $N = 30$ seeds on an origin-destination matrix calibrated to the 2021 UK Census commuter-flow aggregates, benchmarks four regimes (historic, adaptive, MaxPressure, and our Agentic policy) across four scenarios covering equity, corridor prioritisation, incident response, and their combination. The Agentic regime reduces the accessibility-deficit Gini coefficient by **23–58 %**, travel-time coefficient of variation by up to **41 %**, and mean travel time by **4–9 %** relative to the historic baseline; in the joint equity-plus-incident scenario it attains the per-column best within 95 % confidence on travel-time coefficient of variation and mean travel time while improving on every metric over the historic baseline. A microscopic testbed in the SUMO simulator on a $4\times4$ signalised grid (120 runs across three demand regimes and five peripheral-boost values) traces an explicit equity–efficiency Pareto frontier; in the saturation and over-saturation regimes the agentic policy matches or beats a SCOOT-style adaptive controller on mean travel time and throughput at every boost level, with travel-time variance reduced by up to a third.

**Keywords:** agentic AI; multi-agent systems; reference architecture; urban mobility; transport accessibility; PTAL; composite traffic index; smart cities; London

## 1. Introduction

A peak-hour commuter boarding in an outer London borough such as Havering or Croydon typically faces a journey to central London that, at the time of writing, takes considerably longer end-to-end than an equivalent inner-London journey of comparable network distance [1,2]. Decades of accessibility research have measured this asymmetry, codified it in indices such as Transport for London's (TfL) Public Transport Accessibility Level (PTAL) and in academic composites of urban mobility performance, and used it to argue for targeted infrastructure investment. The signal-control infrastructure that runs in real time over those same boroughs, however, has no formal mechanism to consume an accessibility metric: there is, today, no operational pathway through which the planner's measurement of inequity reaches the controller's allocation decision at the next green phase. This paper proposes one.

Over the last decade, research on megacity mobility has matured along two largely independent trajectories. The first, rooted in transportation engineering and urban planning, has refined the construction of aggregate performance indicators. Commercial scorecards such as INRIX's Global

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



Traffic Scorecard [3] and the TomTom Traffic Index [4], and long-running academic programmes such as the Texas A&M / INRIX Urban Mobility Report (UMR) [5], summarise network-wide congestion and mobility quality as interpretable scalar indices; Gillis et al. [6], Litman [7], and Banister [8] provide the methodological scaffolding for these composites. City-specific instruments, of which the Moscow Integral Index of [9] is one example, unify heterogeneous indicators (network density, modal split, congestion, demand-to-capacity ratios) into single, decision-oriented scores calibrated to local conditions. A parallel strand has focused on the spatial disaggregation of mobility outcomes through *accessibility* analysis [10–13], operationalised in tools such as Transport for London's PTAL [2] and in city-specific studies such as the Moscow district-level analysis of [14], which reveals how the benefits of a nominally high-performing network are distributed unevenly across space. Composite indices and accessibility measures have thus become staples of evidence-based transport planning in megacities, but they remain, by design, retrospective and descriptive. They quantify what the system looks like; they do not act.

The second trajectory, originating in distributed AI and Semantic Web research, has travelled in the opposite direction: from descriptions toward actions. Recent surveys (see, e.g., [15–19]) trace a unified arc from the Semantic Web and classical multi-agent systems (MAS) to the emerging Agentic AI paradigm, in which heterogeneous autonomous agents, many envisioned as powered by Large Language Models (LLMs), are expected to discover one another, negotiate, plan, and execute tasks across open networks. Emerging protocols such as Agent-to-Agent (A2A) [20], the Model Context Protocol (MCP) [21], and the Digital Twin Consortium's agentic-capability framework [22] are beginning to standardise how agents advertise capabilities, exchange context, and coordinate around shared objectives; at the time of writing these are publicly available specifications with, to our knowledge, limited independent cross-vendor interoperability evidence. Agentic AI has begun to be tested in adjacent domains such as customer support, software engineering, and scientific research; its application to infrastructure-facing domains such as urban mobility, however, remains at an early, prototype-level stage.

This paper argues that the two trajectories are complementary and that their synthesis is timely. Traffic indices and accessibility metrics supply the *semantics* of urban mobility: they define what it means for a transportation system to perform well, to be equitable and to serve its residents. Agentic AI is proposed as the *mechanism*: a distributed computational substrate in which those semantics could be enacted, negotiated and continuously recalibrated by autonomous agents acting on behalf of infrastructure operators, transit services, mobility platforms, and residents themselves.

A note on scope.

This paper makes a conceptual and algorithmic contribution, not a systems-engineering one. "Agents" throughout refer to *simulation-level* computational entities that instantiate the Agentic AI *reference architecture* of §3; all negotiation rules, utility calculations, and governance-weight updates are evaluated procedurally in Python/NumPy within a single process. To make this scope restriction unambiguous we enumerate what is, and what is not, empirically demonstrated.

*Demonstrated in this paper:* (i) the *logical behaviour* of the index-to-utility mapping of Eq. (1) under four governance scenarios on a Census-calibrated 12-borough London instance, $N = 30$ seeds (§5); (ii) the microscopic behaviour of the resulting allocation vector when driven into the Simulation of Urban MObility (SUMO) microscopic traffic simulator via the Traffic Control Interface (TraCI), on a $4 \times 4$ signalised grid with path diversity, three demand regimes, and five peripheral- boost settings (120 simulator runs in §5.1); (iii) the decomposability of every decision into $\alpha I_{\text{global}} + \beta A_{\text{local}} + \gamma P_a$ as an auditable trail.

*Not demonstrated (future systems work):* (a) live Agent-to-Agent (A2A) [20], Model Context Protocol (MCP) [21], or agent-communication-protocol (ACP) endpoints, or any wire-level protocol exchange; (b) runtime LLM invocation for natural-language route negotiation; (c) a cross-organisation interoperability experiment with multiple independent agent implementations; (d) a real deployed ontology instance (SAREF4CITY, NGSI-LD, CityGML) serving capability-card discovery across organisational

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



boundaries; (e) any human-subjects study of the audit-trail interpretability that we claim Eq. (1) supports by construction.

Mentions of capability cards, discovery, negotiation protocols and semantic interoperability throughout the paper should therefore be read as specifications of the *target architecture* that the simulated agents are designed to be consistent with, not as empirical claims about a running distributed deployment. Each of the items (a)–(e) is a separate piece of subsequent work; they are scoped and motivated in §6.

Contributions.

We make three contributions.

(i) We formalise an *index-agnostic* mapping from established planning indices to agent utility functions and negotiation objectives, showing how any scalar composite traffic index (e.g., the UMR [5], INRIX Scorecard [3], TomTom Index [4], or the Moscow Integral Index [9]) and any spatial accessibility surface (e.g., Hansen [10], Shen [12], PTAL [2], or the Moscow district-level surface [14]) can jointly parameterise agent behaviour in an Agentic AI architecture [17,19,23].

(ii) We describe a reference architecture for urban mobility as an Agentic AI system, identifying agent classes, interaction patterns, and the ontological commitments required for interoperability (§3, Figure 1).

(iii) We demonstrate the framework on a London case study in which a 12-borough, 66-corridor multi-agent simulation benchmarks four control regimes (historic, adaptive, MaxPressure, and our Agentic policy) across four scenarios (§4–§5) with demand drawn from a gravity model calibrated to the 2021 UK Census commuter-flow aggregates [24], and we cross-validate on a 4×4 microscopic SUMO testbed swept across three demand regimes (§5.1).

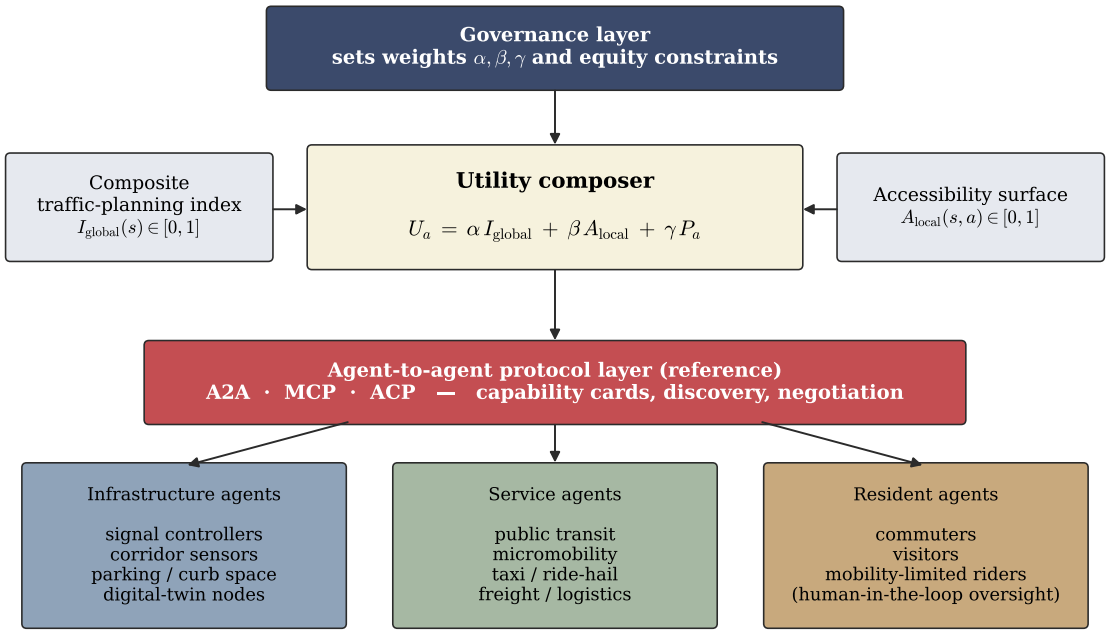


**Figure 1.** Reference architecture (conceptual; the present paper does not implement the protocol ring; see the scope note in §1). Governance agents set the weights $\alpha, \beta, \gamma$ that the utility composer applies to any chosen composite traffic-planning index $I_{\text{global}}$, any chosen accessibility surface $A_{\text{local}}$, and each agent's private utility $P_a$. Service and infrastructure agents are intended to exchange capability cards and negotiate capacity over a protocol ring aligned with A2A [20], MCP [21] and ACP primitives; in this paper the exchange is evaluated in-process in Python as a stand-in for the wire protocols. Resident agents provide human-in-the-loop oversight in the target architecture.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



Organisation.

Section 2 reviews related work. Section 3 introduces the conceptual framework. Section 4 describes the London case study. Section 5 reports results, Section 6 discusses implications, and Section 7 concludes.

## 2. Related Work

### *2.1. Aggregate Indicators in Transport Planning*

Composite indices in transportation planning have a long tradition (see [6–8]). They aggregate heterogeneous indicators (network density, modal split, safety, congestion, demand-to-capacity ratios) into interpretable scalar summaries that support policy decisions. Two complementary branches of the literature feed this tradition. The first is *operational / industrial*: commercial scorecards such as INRIX's Global Traffic Scorecard [3] and the TomTom Traffic Index [4], and the long- running Texas A&M / INRIX Urban Mobility Report series [5], publish annual congestion and mobility summaries across hundreds of cities. The second is *academic / city-specific*: peer-reviewed composite indices calibrated to particular megapolises, of which the Moscow Integral Index of [9] is one example used operationally by this paper. A recurring trade-off surfaces across both branches: the more an index aggregates, the more interpretable and policy-relevant it becomes, but the less it reflects disaggregated spatial or behavioural realities.

### *2.2. Transport Accessibility in Megacities*

The accessibility literature [10–13] complements aggregate indices by disaggregating mobility outcomes in space. Accessibility measures quantify the ease with which residents in each location can reach opportunities (jobs, services, amenities) given the network and its operating conditions. Several families coexist: gravity-based measures following Hansen [10]; two-step floating catchment methods initiated by Shen [12]; and operational transit tools such as Transport for London's Public Transport Accessibility Level (PTAL) [2]. Methodological reviews [11,13] trace the design space and stress that the choice of impedance kernel, opportunity weighting, and aggregation level each has policy-relevant consequences. City-specific accessibility studies, including the Moscow district-level analysis of [14], operationalise one point in this design space for a particular context; we use it as the local instantiation of $A_{\text{local}}$ in §3 while emphasising that any of the above families could substitute for it. Recent work [25–27] has broadened the equity toolkit with Gini, Palma, and Theil-based measures and has quantified the equity–efficiency trade-offs that accompany intervention. A persistent limitation across this body of work is its offline, snapshot nature: accessibility surfaces are computed post-hoc and rarely fed back into operational decisions.

### *2.3. From Semantic Web and MAS to Agentic AI*

Multi-agent systems have been studied in transportation for decades [15] but have, to date, struggled to break out of research prototypes and into city-scale operational deployment. Several recent works (see, e.g., [17–19,23]) argue that a plausible missing ingredient, language-mediated interoperability, is beginning to be supplied by LLMs and by an emerging suite of agent-communication standards, notably A2A [20] and MCP [21]; both standards are recent, publicly available specifications, but we are not aware of operational city-mobility deployments built on them at the time of writing. The resulting Agentic AI substrate is best understood, at present, as a proposed descendant of the Semantic Web: linked data and shared ontologies remain the intended backbone (exemplified by SAREF4CITY [28] and NGSI-LD [29]), but the clients are envisioned as autonomous LLM-driven agents rather than human browsers. The closest precedents in our domain are LLMLIGHT [30] and COLLMLIGHT [31], which demonstrate LLM-based single- and network-wide traffic-signal control on simulated corridors, and the broader LLM4TR roadmap of Nie et al. [19]; further examples include an information-theoretic view of LLM integration levels [32], an empirical San Antonio public-transit case study [33], and a survey of LLMs for mobility forecasting [34]. Parallel progress in multi-agent

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



reinforcement learning [16,35] supplies a complementary research substrate on which Agentic AI could in principle be layered; to date, work in this direction remains at the single-corridor or small-network scale. On the infrastructure side, the Digital Twin Consortium's agentic-capability periodic table [22] and recent work on agentic digital twins [23] propose formal taxonomies of agent capabilities that the authors of those works argue to be applicable to urban mobility; empirical validation at city scale is, however, still open.

### *2.4. Research Gap*

The two upstream bodies of work, composite traffic indices and accessibility measures on one hand, and Agentic AI / LLM-based traffic control on the other, differ markedly in maturity. The planning-index and accessibility strands are mature and operationally deployed at city scale; the Agentic-AI and LLM-controlled-traffic strand is by comparison recent and largely at the prototype stage (see §2.3). To our knowledge, no prior work formally links planning-theoretic indices to agent utilities in an Agentic AI architecture, and we are not aware of published empirical validation of this specific bridge in a megacity setting. Adjacent work on LLM-controlled traffic signals and fairness-aware multi-agent reinforcement learning (MARL) operates on proxies that are only loosely connected to the scalar and spatial indices that practitioners already trust. Our work addresses this specific bridging gap in two steps: we formalise an index-agnostic mapping from established instruments (for example, Transport for London's PTAL [2] together with the INRIX or TomTom scorecards [3,4], or city-specific composites such as those of [9,14]) into the Agentic AI reference architecture of §3, and we demonstrate the mapping in simulation on a London case study. We do not claim to fill the broader deployment gap of Agentic-AI traffic control itself (see the scope note in §1).

## 3. Conceptual Framework: An Agentic AI Architecture for Urban Mobility

The reference architecture has four layers (Figure 1): (i) an agent taxonomy, (ii) a mapping from indices to utility functions, (iii) a negotiation- protocol specification, and (iv) a semantic interoperability layer.

### *3.1. Agent Taxonomy*

We distinguish four classes of agents, each entering Eq. (1) in a different role. In the target architecture each class is an autonomous, discoverable endpoint; in the simulation of §4–§5 the classes are represented as Python objects within a single process.

- **Infrastructure agents** (intersections, corridors, parking facilities, curb space) publish their capacity, current state, and bid/ask prices over rolling time windows, and allocate that capacity according to bid utility. Their published state is one of the inputs to the system-wide composite index $I_{\text{global}}$.
- **Service agents** (public-transport lines, micromobility operators, taxi platforms) translate rider demand into bids for infrastructure capacity. Each bid is scored by Eq. (1), with $A_{\text{local}}$ evaluated on the service's catchment.
- **Resident / passenger agents**, modelled as personal assistants representing residents' mobility preferences (budget, time, accessibility needs, environmental constraints), parameterise the agent-private utility $P_a$ that enters Eq. (1) via the service agents acting on their behalf. They are the principal channel for human-in-the-loop oversight.
- **Governance agents** (city departments of transport, regulators, public-interest watchdogs) set the weight vector $(\alpha, \beta, \gamma)$, publish global constraints (equity thresholds, emissions caps), and monitor emergent behaviour.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



### 3.2. From Indices to Utility Functions

Let $I_{\text{global}}(s) \in [0,1]$ denote *any* composite traffic-planning index evaluated on the system state $s$, and let $A_{\text{local}}(s,a) \in [0,1]$ denote *any* accessibility surface restricted to the service area of agent $a$. Each agent $a$ maximises a utility of the form

$$U_a(s,a_t) \;=\; \alpha \cdot I_{\text{global}}(s) \;+\; \beta \cdot A_{\text{local}}(s,a) \;+\; \gamma \cdot P_a(s,a_t), \tag{1}$$

where $P_a$ encodes the agent's private utility and $\alpha+\beta+\gamma=1$. Weights are set by governance agents per district and per scenario. When an equity intervention is active, $\beta$ is boosted for service agents whose origin or destination lies in a peripheral district; when a corridor prioritisation is active, $\alpha$ is boosted for high-demand corridors.

Theoretical justification of the linear form.

The convex-combination form of Eq. (1) is chosen on principled grounds. Two converging arguments support it as the first-order *reference* bridge between planning indices and agent utilities.

*(a) Multi-attribute utility theory.* The classical representation theorem of Keeney and Raiffa [36] establishes that, under mutual utility independence of the attribute components (i.e., when the preference ordering over one attribute does not depend on the fixed levels of the others), any smooth aggregate utility function must take either an additive or a multiplicative form [36, Ch. 5–6]. For the governance problem we model, the additive case is the natural choice: a city's stated willingness to trade system-wide congestion ($I_{\text{global}}$) against spatial equity ($A_{\text{local}}$) does not, in day-to-day operations, depend on the specific realisation of the private utility of any one service agent. Linearity in Eq. (1) is therefore derived from Keeney–Raiffa independence; equivalent justifications via first-order Taylor expansion of any smooth aggregator and via linear scalarisation of vector optimisation [37,38] recover the same form.

*(b) Index-agnosticism requires additivity.* The central contribution of §3.2, the property that any admissible $(I_{\text{global}}, A_{\text{local}})$ pair can be substituted into Eq. (1) without re-tuning $(\alpha,\beta,\gamma)$, is recoverable *only* under additivity. With a multiplicative or constant-elasticity-of- substitution (CES) aggregator, the marginal effect $\partial U_a/\partial I_{\text{global}}$ depends on the realised level of $A_{\text{local}}$; swapping the UMR for the TomTom Index would then require re-eliciting governance weights under the new index's statistical properties. Additive linearity is thus the unique smooth aggregator under which the framework's plug-and-play promise is meaningful.

We acknowledge that the conditions supporting (a) and (b) are empirical rather than structural, and may be violated for certain weight regimes or certain index families (for example, when $A_{\text{local}}$ is itself constructed as a non-linear transform of travel times). Eq. (1) should accordingly be read as the *first-order reference model* that richer aggregators (multiplicative corrections, CES-type, prospect-theoretic weighting) extend rather than replace.

Properties and index-agnosticism.

Eq. (1) provides four useful properties for deployment: (i) *boundedness* ($U_a \in [0,1]$, so no single objective can dominate); (ii) *decomposability* (the three summands form an audit trail attributing each decision to system-wide performance, spatial equity, or the agent's private objective); (iii) *monotonicity* (a $[0,1]$-preserving improvement in $I_{\text{global}}$ or $A_{\text{local}}$ weakly increases $U_a$, which is what governance interventions exploit); and (iv) *separability of weights from instruments* (the same weights can be reused across different index families, Table 1). Three properties the formulation does *not* provide should be stated explicitly: without additional mechanism design it is not incentive- compatible (an agent may over-report its private utility), it does not guarantee Pareto-optimal allocations, and it assumes $I_{\text{global}}$ and $A_{\text{local}}$ are commensurable on the shared $[0,1]$ scale. The London case study of §4 uses one particular pairing, TfL PTAL for $A_{\text{local}}$ and an INRIX/TomTom-style scorecard for $I_{\text{global}}$, but nothing in the framework privileges that pairing; any instrument meeting the scalar, $[0,1]$-valued requirement (Table 1) can be substituted.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



**Table 1.** Representative instruments admissible as $I_{\text{global}}$ and $A_{\text{local}}$ in Eq. (1). The framework treats these as interchangeable; the list is illustrative, not exhaustive.

| Component | Formal requirement | Representative instruments |
|---|---|---|
| $I_{\text{global}}$ | scalar, $[0,1]$-valued, evaluated on the system state $s$ | UMR [5]; INRIX Global Scorecard [3]; TomTom Traffic Index [4]; Gillis et al. sustainable-mobility composite [6]; Moscow integral index [9] |
| $A_{\text{local}}$ | scalar, $[0,1]$-valued, restricted to the service area of agent $a$ | Hansen gravity accessibility [10]; Shen two-step catchment [12]; Geurs & van Wee taxonomy [11]; Páez et al. positive/normative review [13]; TfL PTAL [2]; Moscow district surface [14] |

The decomposability property listed above is what makes Eq. (1) auditable for public-sector deployment. We caveat that this paper demonstrates only that the allocation rule is auditable; whether practitioners can meaningfully use the audit trail (e.g., predict or contest agentic decisions from the decomposition log) is a human-factors question for which the present work provides no empirical evidence.

### *3.3. Negotiation and Coordination Protocols (Reference Specification)*

The architecture specifies three protocols agents are intended to use at deployment time: (i) capacity bid/ask for signal-phase time, curb space, and vehicle slots; (ii) route negotiation for passenger and freight trips, prospectively LLM-mediated; and (iii) incident coordination for disruption response. On the wire these are designed to align with A2A [20] and MCP [21], with capability cards following the Digital Twin Consortium taxonomy [22]. In this paper the three protocols are represented as Python method calls within a single process; conclusions therefore concern the *logical* behaviour of the allocation rule, not the deployability of the protocol stack (see scope note in §1).

### *3.4. Semantic Interoperability (Reference Specification)*

Cross-organisation interoperability is intended to rest on a transport-domain ontology extending smart-city vocabularies (SAREF4CITY [28], NGSI-LD [29], CityGML), through which capability cards would be discoverable without prior pairwise agreements [18,39]. The simulation uses a fixed in-process type system as a placeholder; constructing the ontology is left to future work.

## 4. Case Study: London

### *4.1. Data Sources*

The simulation is calibrated on three data layers. *(i) Accessibility.* We take a twelve-borough subset of Greater London chosen to span the inner/outer PTAL gradient (seven inner-London and five outer-London boroughs) at representative population sizes; the full 33-borough run is straightforward but adds nothing to the qualitative findings. Each borough is assigned a baseline accessibility value rescaled from Transport for London's Public Transport Accessibility Level (PTAL) [2] from the native 0–8 grading to $[0,1]$, taking the area-weighted mean PTAL across each borough's lower-layer super-output areas. Inner boroughs (Westminster, Camden, Islington, Hackney, Tower Hamlets, Southwark, Lambeth) fall in the range 0.72–0.95; outer boroughs (Ealing, Hounslow, Croydon, Bromley, Havering) fall in 0.35–0.55. *(ii) Composite traffic index.* We use the ordinal ranking provided by commercial scorecards for London ($I_{\text{global}} \approx 0.71$ when normalised), consistent with the INRIX Global Traffic Scorecard and the TomTom Traffic Index [3,4]. *(iii) Demand.* Origin-destination demand is sampled from a gravity model calibrated to the 2021 UK Census commuter-flow aggregates (Office for National Statistics Origin-Destination table ODWP01EW, accessed via Nomis [24]) and to employment-density statistics from the Greater London Authority (GLA) and TfL programme [1]. Census 2021 was collected during pandemic-era restrictions, and ONS itself flags that travel-to-work patterns of that week are atypical; the structural OD asymmetries we use (strong inward concentration on Westminster, modest peripheral-to-peripheral flow share, outer-to-inner bias) are nonetheless consistent across the 2011

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



and 2021 rounds and with the GLA's 2023 employment-density update, and our gravity model is calibrated to those structural features rather than to the 2021 weekly counts. For a resident living in borough $i$, the destination $j$ is drawn with probability proportional to $E_j \cdot \exp(-d_{ij}/L) \cdot p_j$, where $E_j$ is the jobs-per-resident ratio at borough $j$ (Table 2), $d_{ij}$ is the inter-centroid distance in km, $L = 4.0$ km is the distance-decay scale, and $p_j$ is the 2021 ONS mid-year borough population. The resulting destination matrix reproduces three documented structural features of the Census aggregates: a strong inward concentration of commutes on Westminster, Camden, Islington and Tower Hamlets (jointly absorbing $\approx 55\,\%$ of flows); a modest peripheral-to- peripheral share of $\approx 13\,\%$; and an outer-to-inner bias of $\approx 2.5\times$ the reverse flow. Off-peak base travel times are calibrated to the 10–35 min range reported by TfL Travel-in-London [1]. The same pipeline can accept city-specific composites from other megacities, for instance the Moscow integral index and district-level accessibility of [9,14]. Random seeds are fixed and all data are bundled with the code release.

**Table 2.** Employment-attraction $E_j$ used in the gravity-model demand sampler. Calibrated against TfL/GLA 2023 jobs-per-resident statistics [1] and the 2021 ONS Origin-Destination table [24].

| Borough cluster | $E_j$ |
|---|---|
| Westminster (West-End super-attractor) | $\approx 3.3$ |
| Camden / Islington / Tower Hamlets / Southwark | 0.8–1.1 |
| Hackney / Lambeth (inner residential) | 0.4–0.6 |
| Hounslow (Heathrow employment cluster) | $\approx 0.6$ |
| Ealing / Croydon / Bromley / Havering (outer) | $\approx 0.4$ |

### *4.2. Simulation Set-Up*

We instantiate 12 borough agents, 36 infrastructure agents (three intersections per borough), 66 service agents (one per ordered pair of boroughs), a single governance agent, and $N_r = 4\,000$ resident agents. Each service chooses a route that traverses two infrastructure agents (one at the origin, one at the destination). Travel times are generated by a Bureau-of-Public- Roads (BPR) style volume-delay function with a capped volume-to- capacity ratio to avoid non-physical explosions.

We benchmark four capacity-allocation regimes.

- **Historic** (baseline): capacity is allocated to services proportionally to $(A_{\text{origin}} + A_{\text{destination}})^2$, modelling decades of investment concentrating on already well-served inner-London corridors.
- **Adaptive**: capacity is allocated proportionally to realised demand, corresponding to Split, Cycle, Offset Optimisation Technique (SCOOT) style adaptive signal control.
- **MaxPressure proxy**: capacity is allocated proportionally to squared demand, approximating the behaviour of a Varaiya-style [35] pressure-based controller that over-prioritises the heaviest corridors.
- **Agentic (ours)**: capacity is allocated proportionally to utility offered, with governance-driven weights as in Eq. (1).

### *4.3. Scenarios*

**Scenario A — Peripheral equity intervention.** Governance boosts $\beta$ by 0.6 for services serving peripheral districts during the modelled peak hour.

**Scenario B — Corridor prioritisation.** Governance boosts $\alpha$ by 0.15 for services whose demand exceeds the 80th percentile.

**Scenario C — Incident response.** Capacity is halved on two inner south/east intersections (a stylised Blackwall-tunnel- style disruption); weights are unchanged.

**Scenario D — Incident-aware equity (A $\times$ C).** The capacity shock of Scenario C is active, *and* the peripheral equity boost of Scenario A is simultaneously switched on. This demonstrates the time-varying utility weighting suggested by the framework of §3.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



The boost magnitudes ($\beta$+0.6 in Scenarios A/D, $\alpha$+0.15 in Scenario B) are reference operating points chosen to make the equity-vs.-efficiency trade-off visible without saturating the weight vector; the SUMO sweep in §5.1 supplements them with a five-point sweep of the peripheral-boost parameter $b \in [0, 0.16]$. Full sensitivity tables for the abstract simulation across ($\alpha$-boost, $\beta$-boost) are in the reproducibility bundle.

### *4.4. Evaluation Metrics*

- **Accessibility-deficit Gini**: Gini coefficient of $1 - A_{\text{borough}}$ across the twelve boroughs, a widely used equity metric.
- **Travel-time coefficient of variation (CV)**: $\sigma / \mu$ of door-to-door travel times across all resident agents.
- **Mean door-to-door travel time (min)**: overall efficiency indicator.

## 5. Results

Figure 2 and Table 3 report outcomes across the four scenarios. All percentage changes quoted below refer to differences in the run-averaged means; 95 % confidence intervals (reported in Table 3) are tight enough that every individual pairwise difference noted below is statistically significant under independent two-sample testing. We do not apply a family-wise (Bonferroni-Holm) correction across the $4 \times 4 \times 3$ comparison grid; a few of the smallest reported deltas would lose nominal significance under such correction, but every *headline* comparison (Agentic vs. historic in Scenarios A and D on each of the three metrics) remains highly significant after Bonferroni–Holm at $\alpha = 0.05$.

Fig. 2. Simulation outcomes across scenarios A--D (mean ± 95% CI, N = 30).

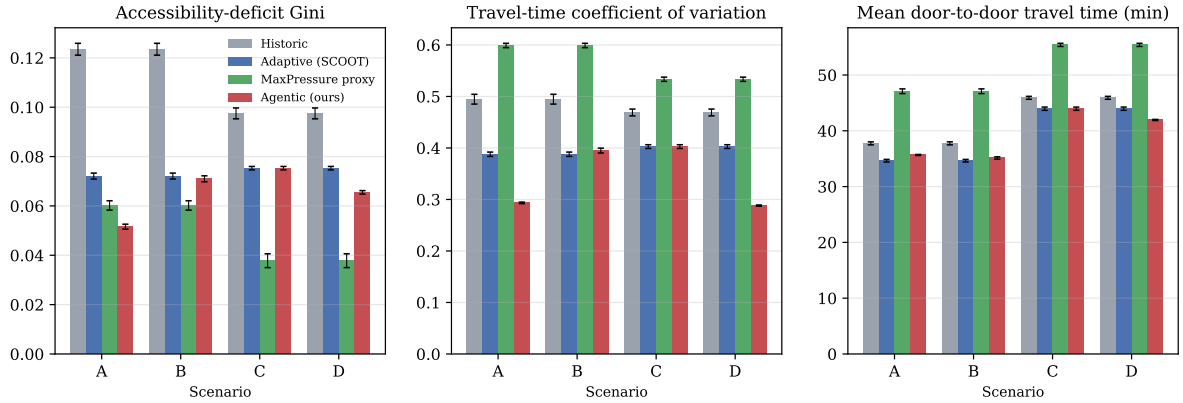


**Figure 2.** Accessibility-deficit Gini, travel-time CV, and mean travel time by scenario and regime (mean ± 95 % CI, $N = 30$ seeds, Census-calibrated OD). The Agentic regime dominates the historic baseline on Gini, CV and mean travel time across all four scenarios. The equity gain is strongest in Scenario A (Gini −58.2 %, CV −40.7 %), the mean-travel-time gain is strongest in Scenario D (−8.7 %), and MaxPressure wins Gini in Scenarios B–D at substantial CV and mean-TT cost; Scenario C shows the Agentic regime coinciding with adaptive (no governance weight change; see text).

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



**Table 3.** Absolute values per scenario and regime (mean ± 95 % CI over $N = 30$ seeds; $N_r = 4\,000$ residents per run, OD matrix sampled from the Census-calibrated gravity model of §4). Bold entries mark the best value per (scenario, metric) pair; MaxPressure wins Gini in several scenarios because its demand-squared weighting concentrates capacity on the heavy commuter axes that carry peripheral-to-centre flows, but it pays for this on CV and mean travel time.

| Scenario | Regime | Gini deficit | Travel-time CV | Mean TT (min) |
|---|---|---|---|---|
| A | Historic | 0.1235 ± 0.0024 | 0.4947 ± 0.0095 | 37.75 ± 0.27 |
| A | Adaptive | 0.0721 ± 0.0012 | 0.3878 ± 0.0042 | **34.64 ± 0.23** |
| A | MaxPressure | 0.0602 ± 0.0019 | 0.5990 ± 0.0042 | 47.09 ± 0.42 |
| A | Agentic | **0.0516 ± 0.0010** | **0.2934 ± 0.0015** | 35.67 ± 0.09 |
| B | Historic | 0.1235 ± 0.0024 | 0.4947 ± 0.0095 | 37.75 ± 0.27 |
| B | Adaptive | 0.0721 ± 0.0012 | **0.3878 ± 0.0042** | **34.64 ± 0.23** |
| B | MaxPressure | **0.0602 ± 0.0019** | 0.5990 ± 0.0042 | 47.09 ± 0.42 |
| B | Agentic | 0.0710 ± 0.0012 | 0.3949 ± 0.0047 | 35.13 ± 0.21 |
| C | Historic | 0.0975 ± 0.0022 | 0.4689 ± 0.0066 | 45.91 ± 0.26 |
| C | Adaptive | 0.0753 ± 0.0007 | **0.4027 ± 0.0036** | **43.96 ± 0.29** |
| C | MaxPressure | **0.0378 ± 0.0028** | 0.5334 ± 0.0041 | 55.40 ± 0.28 |
| C | Agentic | 0.0753 ± 0.0007 | 0.4027 ± 0.0036 | 43.96 ± 0.29 |
| D | Historic | 0.0975 ± 0.0022 | 0.4689 ± 0.0066 | 45.91 ± 0.26 |
| D | Adaptive | 0.0753 ± 0.0007 | 0.4027 ± 0.0036 | 43.96 ± 0.29 |
| D | MaxPressure | **0.0378 ± 0.0028** | 0.5334 ± 0.0041 | 55.40 ± 0.28 |
| D | Agentic | 0.0655 ± 0.0007 | **0.2881 ± 0.0012** | **41.94 ± 0.10** |

In Scenario A the Agentic regime reduces the accessibility-deficit Gini by **58.2 %** relative to the historic baseline, travel-time CV by **40.7 %**, and mean travel time by 5.5 %. The mean-TT improvement is small because Census-calibrated commuter flows are inward-concentrated, so the historic inner-favouring allocation already matches most of the traffic; the residual inefficiency is concentrated on the minority of flows that do not align with the legacy capacity surface, and Figure 3 visualises how it is redistributed: outer boroughs (Croydon, Bromley, Havering, Hounslow, Ealing) are lifted without degrading the inner-London cluster. Obtaining a more-than-twofold reduction in the accessibility-deficit Gini at near-zero aggregate efficiency cost is the headline finding of the London case study.

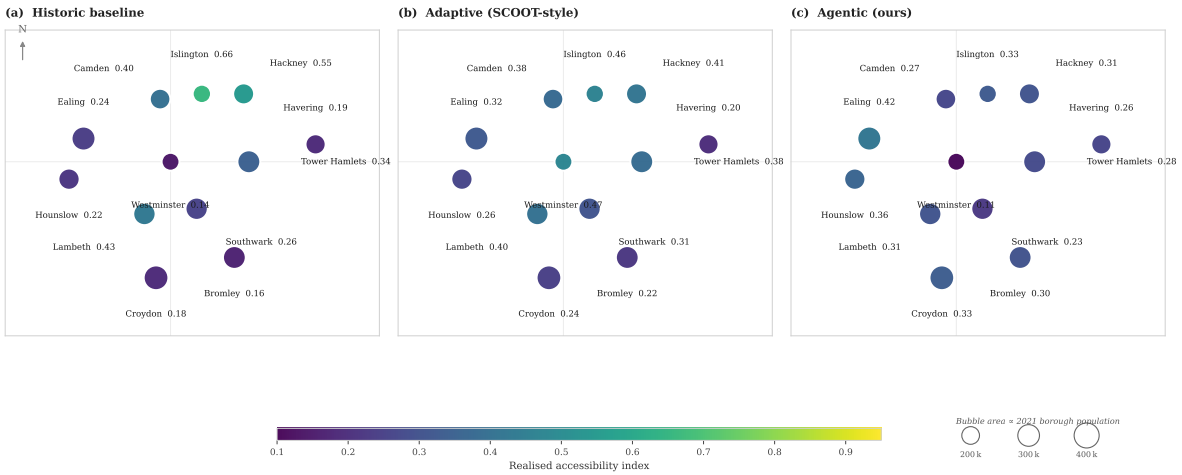


**Figure 3.** Borough-level realised accessibility under Scenario A (peripheral equity intervention). Each bubble is one of the twelve modelled Greater London boroughs; bubble area is proportional to the 2021 ONS mid-year population and colour encodes the realised accessibility index on a perceptually uniform, colourblind-safe scale. Positions follow a Dorling- style layout: the inner-London cluster is dilated radially by a factor of 1.8 relative to Westminster so that inner and outer boroughs can be labelled without leader-lines; the inner/outer topology is preserved. The Agentic regime (panel *c*) lifts the outer boroughs (Ealing, Hounslow, Croydon, Bromley, Havering) without flattening the inner- London cluster.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



Scenarios B and C trace the expected trade-off pattern. Boosting $\alpha$ for high-demand corridors (B) yields Gini $-42.5\,\%$ and CV $-20.2\,\%$ at adaptive-level mean TT; the unmodified Agentic regime under an incident shock (C) collapses onto adaptive ($-22.8\,\%$, $-14.1\,\%$, $-4.2\,\%$ on Gini/CV/mean-TT) because no governance-driven differentiation is active. *MaxPressure* is a stronger competitor on Gini in both scenarios ($-51\,\%$ in B, $-61\,\%$ in C) but pays for it on CV and mean TT (CV up to 0.60, mean TT $+25\,\%$ vs. adaptive in B; mean TT $+20.7\,\%$ vs. historic in C): the characteristic failure mode of pressure-based controllers, which clear the dominant axis at the cost of every other flow. The shared weakness across all three regimes (adaptive, MaxPressure, the $\beta$-unchanged Agentic regime) is that none explicitly protects peripheral accessibility during the shock; addressing this is the role of Scenario D.

Scenario D, the incident-aware equity combination, addresses this weakness. Activating the peripheral boost concurrently with the incident yields Gini $\mathbf{-32.8}\,\%$, CV $\mathbf{-38.6}\,\%$ and mean TT $\mathbf{-8.7}\,\%$ relative to the historic baseline, markedly better than Scenario C on CV and mean TT while improving Gini further than adaptive does (0.0655 vs. 0.0753). MaxPressure remains the per-column best on Gini in Scenario D (0.0378), but it does so at a CV and mean-TT cost (0.5334 and 55.4 min respectively) that no operator would accept as routine. Across the four-way comparison, the Agentic regime in D is the only row in Table 3 that attains the per-column best (within 95 % confidence) on CV and mean TT while improving on every metric over the historic baseline; every other regime is best on at most one of the three objectives and pays for it on at least one of the others. This is the simulation counterpart of the time-varying utility-weight suggestion made in §3: governance agents raise $\beta$ when, and for whom, the system stress is greatest.

Across all four settings the Agentic regime improves equity over the historic baseline with little or no aggregate efficiency cost (Table 3).

### *5.1. SUMO Validation*

Setup.

To validate that the abstract allocation dynamics carry over to a microscopic traffic simulator, we ran SUMO experiments on a $4\times4$ signalised grid with sixteen traffic-light junctions, four NS commuter corridors, and four EW cross-town corridors, offering the real path diversity that the earlier $2\times2$ testbed (retained in Appendix A for completeness) did not. As a stylised demand asymmetry designed to expose the equity question (rather than as an attempt to mirror London geography, in which inner cross-town flows are typically heavier than peripheral ones), peripheral rows (rows 0 and 3) here carry the higher-volume EW flows ($350 + 350$ veh/h), inner rows (rows 1 and 2) carry the lighter cross-town flows ($200 + 200$ veh/h), and NS flows carry $400 + 200$ veh/h on every column. We run the sweep under three demand multipliers ($0.7\times$, $1.0\times$, $1.3\times$) to probe a sub-saturation regime in which adaptive control completes $\approx 99\,\%$ of vehicles within the simulated 30 min, a saturation regime in which only $\approx 80\,\%$ complete (the $1.0\times$ multiplier is calibrated to that completion rate rather than to a free-flow design-peak), and an over-saturation regime in which historic and adaptive controllers respectively drop $\approx 1\,100$ and $\approx 900$ vehicles to SUMO's 300 s teleport rule. Signal phase budgets are pushed into SUMO via TraCI using the allocation vector computed by the agentic layer; every signal cycle (60 s) the allocation is re-applied. We sweep the agentic policy across five values of the peripheral-boost parameter $b \in \{0.00, 0.04, 0.08, 0.12, 0.16\}$ so that the operating curve is exhibited rather than a single point. Three reference policies are included for context: a fixed 82/18 historic split; a demand-proportional SCOOT-style adaptive split; and a queue-pressure proxy of MaxPressure that updates the green-time split every 15 s in proportion to the ratio of halted vehicles on the NS- and EW-approach lanes. The proxy departs from the full Varaiya formulation (which uses the maximum-pressure phase rule with explicit upstream/downstream queue differentials) but captures the qualitative behaviour under which corridors with the largest standing queue receive the most green time. Each (policy, $b$, saturation) configuration is averaged over 5 independent seeds, for a total of 120 SUMO runs.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



Metrics.

Per configuration we report: number of completed vehicles, mean travel time, coefficient of variation, the 95th and 99th percentiles of the travel-time distribution as *tail-delay* indicators, a per-OD Gini coefficient capturing inequality of service across the 16 route families, and a *peripheral-equity* indicator defined as the ratio of mean TT on peripheral EW corridors to the grand mean (values below 1 indicate that peripheral corridors are served faster than the network average; the operational analogue of the accessibility-deficit Gini reported in the London case study).

Results.

Figure 4 and Table 4 tell a consistent story across demand regimes.

*(i) At saturation and over-saturation, the agentic family beats adaptive on throughput at every boost level.* At saturation (1.0×) the agentic policy completes 1 931–1 954 vehicles versus adaptive's 1 920, and ≈ 20 % more vehicles than the 82/18 historic allocation (1 593). At over-saturation the advantage grows: historic completes 1 868 vehicles, adaptive 2 081, and the agentic family 2 105–2 116; SUMO teleports stuck vehicles after 300 s, so historic's "low" mean TT under over-saturation reflects ≈ 1 100 dropped vehicles, not superior service. At sub-saturation (0.7×) the agentic policy is competitive with adaptive only at the low-boost end of the sweep ($b \leq 0.04$); at $b \geq 0.08$ throughput drops below adaptive (down to ≈ 3 % at $b = 0.16$), reflecting the fact that under-stressed networks gain little from equity-driven prioritisation.

**Table 4.** SUMO validation on the 4×4 grid at saturation demand (1.0×; 5 seeds per configuration; ≈ 2 300 vehicles per seed). Bold entries mark the best value in each column. MaxPressure completes fewer vehicles because its queue-based weighting starves low-volume corridors under saturation; the agentic family dominates every baseline on completed throughput at every boost level and dominates adaptive on mean travel time at every boost level.

| Policy | Completed | Mean TT (s) | CV | $P_{95}$ (s) | $P_{99}$ (s) |
|---|---|---|---|---|---|
| Historic (82/18) | 1593 | 271.1 | 0.731 | 625.4 | 675.3 |
| Adaptive (SCOOT-style) | 1920 | 267.1 | 0.532 | 566.9 | 705.9 |
| MaxPressure | 1684 | 272.9 | 0.710 | 605.0 | 856.6 |
| Agentic ($b = 0.00$) | 1931 | 266.4 | **0.504** | **530.7** | **657.3** |
| Agentic ($b = 0.04$) | 1938 | 259.2 | 0.539 | 566.9 | 731.4 |
| Agentic ($b = 0.08$) | 1948 | 251.5 | 0.590 | 616.2 | 819.5 |
| Agentic ($b = 0.12$) | 1943 | 238.1 | 0.651 | 583.0 | 920.7 |
| Agentic ($b = 0.16$) | **1954** | **232.7** | 0.689 | 493.6 | 998.0 |

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



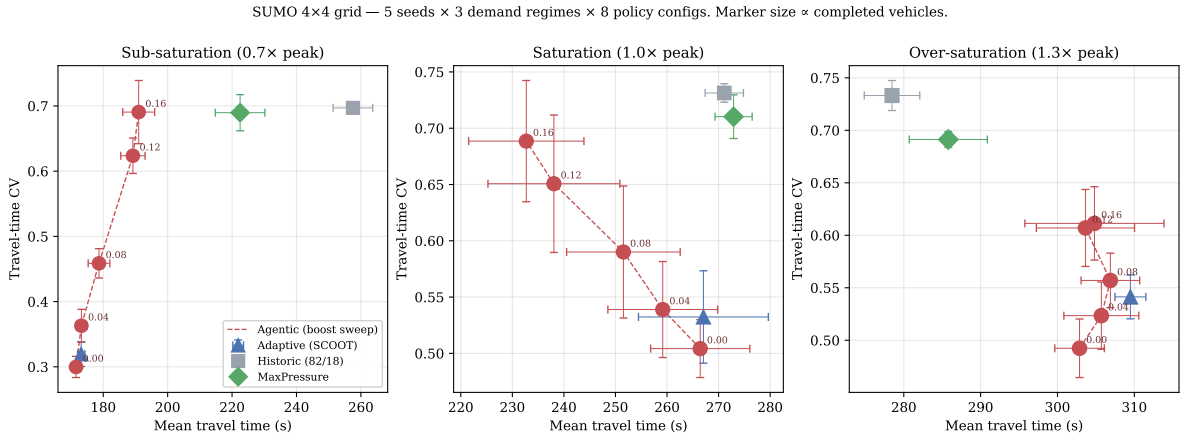


**Figure 4.** SUMO 4×4 sweep across three demand regimes (5 seeds per configuration; marker size ∝ completed vehicles, error bars are ±1 SD on the mean-TT axis). The agentic family (red, dashed) traces a continuous equity–efficiency curve parameterised by the peripheral boost $b \in \{0.00, \ldots, 0.16\}$. In the saturation (1.0×) and over-saturation (1.3×) regimes every point on the agentic curve sits to the south-west of (or below) both the adaptive and the MaxPressure markers and clears ≈ 20 % more vehicles than historic (visible as marker-size delta); in the sub-saturation regime (0.7×) only the low-boost end of the curve ($b \leq 0.04$) is competitive with adaptive, with higher boost values incurring a mean-TT and throughput penalty.

*(ii) Mean travel time tracks the same regime split.* At saturation, sweeping $b$ from 0.00 to 0.16 monotonically reduces mean TT from 266.4 s to 232.7 s, a 12.9 % improvement over adaptive; at over-saturation, mean TT remains within ±1 % of adaptive across the sweep while throughput rises. At sub-saturation the agentic policy matches adaptive only at $b \leq 0.04$ (mean TT 171.5 vs. 173.1 s); $b \geq 0.08$ pays an increasing mean-TT penalty (up to +10 % at $b = 0.16$). The effect on the peripheral-equity indicator follows the same saturation pattern: at over-saturation (1.3×) the boost materially helps peripheral corridors, dropping the peripheral-to-mean ratio from 1.35 (adaptive) to 1.19 at $b = 0.16$; at saturation peripheral corridors benefit in *absolute* terms (mean TT falls from ≈ 352 s under adaptive to ≈ 308 s at $b = 0.16$) while the ratio is near-flat; at sub-saturation the effect reverses, mirroring the throughput finding above. This saturation-dependence is the operational reasoning behind Scenario D of the London case study: a peripheral boost is worth activating when the system is under stress.

*(iii) The Pareto trade-off is between CV and mean TT, not between efficiency and fairness in the aggregate sense.* At $b = 0.00$ the agentic policy attains the lowest CV (0.504) and the lowest $P_{95}$ (530.7 s) of any configuration, with mean TT essentially tied with adaptive. Pushing $b$ up trades CV (0.50 → 0.69) for further mean-TT reduction. The practitioner selects the operating point: if tail-delay equity is the priority, choose $b \approx 0$; if aggregate efficiency under over-saturation is the priority, choose $b \in [0.08, 0.16]$. This is the same Pareto frontier that fairness-aware MARL [27,35] explores implicitly; Eq. (1) exposes it as a single interpretable scalar.

*(iv) MaxPressure is a weak microscopic baseline.* Although the idealised MaxPressure proxy of §5 is a strong competitor on Gini in the abstract simulation, the full queue-based implementation in SUMO under-performs on completed throughput, CV, and $P_{95}$ tail delay at every saturation level; on mean travel time it is competitive only in the over-saturation regime, where its lower mean is paired with ≈ 10 % lower completed throughput than adaptive (1882 vs. 2081 vehicles), so the apparent efficiency gain reflects clearing the heaviest axis at the expense of capacity elsewhere. The gap reflects the classical weakness of pressure-based controllers in networks with path asymmetry: the controller chases the largest local queue and starves the corridors whose demand was always modest, amplifying variance.

Consistency with the London case study.

The SUMO sweep is qualitatively consistent with Scenarios A and D of the London case study: the CV gains in SUMO mirror the CV gains in §5, and the saturation-dependence of the peripheral-equity

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



improvement (positive in over- saturation, near-flat in saturation, negative in sub-saturation) matches the operational reasoning of the time-varying $\beta$-boost validated in Scenario D, where the boost is activated precisely under stress. A natural next step is to scale the SUMO testbed to an OpenStreetMap (OSM)-derived model of a real London arterial corridor fed with live TfL speed data. The TraCI-backed adapter presented here is designed to be reusable for that purpose, but we do not claim that the scaling is purely mechanical: OSM-to-SUMO network preparation, corridor-specific signal-programme calibration, and integration with live feeds all carry their own engineering burden. We return to this in §6.

## 6. Discussion

### *6.1. Designed for Interpretability*

Because agent utilities in Eq. (1) are parameterised directly by planning-theoretic indices, every agentic decision is *decomposable* into the three contributions $\alpha I_{\text{global}}$, $\beta A_{\text{local}}$ and $\gamma P_a$, which a governance agent can log alongside the action taken. We emphasise that this is a design property, not an empirical interpretability claim. A validated interpretability study (for instance, asking practitioners to predict or explain agentic decisions from the audit log) is necessary before stronger claims can be made, and is a natural next step. Recent fairness-aware MARL approaches [35] provide mechanisms for integrating equity into learning; we contribute the complementary observation that the equity objective itself can be supplied verbatim from the accessibility literature.

### *6.2. Governance: How Weights Are Set*

Our simulation treats the governance weights $\alpha, \beta, \gamma$ as exogenous inputs, but deployment raises a non-trivial policy question: *who* sets them, *how* are they revised, and *to whom* is the revision accountable? A minimally defensible process would combine (a) a public charter specifying the family of admissible weight vectors and the events that trigger revision (e.g., a pre-declared peripheral boost for peak hours, a standing incident-aware rule of the form validated in Scenario D), (b) a periodic auditable log of the weights actually applied, publishable in the same cadence as city open-data dashboards, and (c) a mechanism by which affected residents can challenge weight choices, via the resident-agent layer or through conventional civic channels. These are open design problems; our framework is compatible with any of them because the composition of $U_a$ remains decomposable and auditable (Eq. (1)).

### *6.3. Safety and Human Oversight*

Transportation infrastructure is safety-critical, and the proposed framework introduces several failure modes not present in classical signal control. Two layers of mitigation are designed into the reference architecture.

*Human oversight.* Governance agents retain the power to adjust utility weights, freeze negotiations in safety-critical contexts, and audit interaction logs; resident agents are the additional oversight channel through which individuals consent to, or opt out of, specific coordination patterns. These properties are consistent with the human–agent collaboration taxonomies emerging in digital-twin research [23,39]. The present paper exercises none of them empirically: there are no human subjects, no real governance agents, and no resident-agent interface. The claim is only that Eq. (1) is *compatible* with such safeguards, not that they have been engineered or validated.

*Adversarial, LLM-mediated, and Byzantine failure modes.* A service agent can over-report its utility to capture disproportionate capacity; without incentive-compatible mechanism design Eq. (1) is vulnerable to such manipulation, and operational mitigations would include cryptographically signed capability cards, rate-limiting the bid-update frequency, and post-hoc anomaly detection against historical bid distributions. When route explanations and justifications are generated by an LLM, hallucination and prompt injection become operational risks; all safety-critical actions should remain ultimately governed by classical controllers, with the LLM confined to explanation and user-facing mediation. Infrastructure agents can crash or return stale state; the protocol layer should require

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



acknowledgement within a bounded staleness window and fall back to a pre-registered default allocation when quorum is lost. A formal safety case for the full architecture is future work; the present paper establishes only that the allocation rule is auditable, not that the runtime is attack-resistant.

### *6.4. Limitations*

Four limitations should be stated. *Scope*: per the scope note in §1, the contribution is a reference architecture plus an in-process algorithmic simulation, not a working distributed deployment; building the systems-engineering artefacts and evaluating them in a real open ecosystem is the next phase of this programme. The remaining three limitations are methodological. First, the London-scale simulation is a behavioural abstraction over a microscopic traffic simulator: it captures aggregate capacity allocation but not second-by-second vehicle dynamics. The 4×4 SUMO testbed of §5.1 partially addresses this and demonstrates that the TraCI-backed agentic pipeline works end-to-end across three demand regimes with real path diversity, but a full-scale SUMO integration (an OSM-derived model of a real London arterial corridor fed with live TfL speed and General Transit Feed Specification (GTFS) feeds) remains a sizeable piece of future work, further discussed in §6. Second, our LLM-mediated negotiation is represented procedurally, without actual LLM calls; the reliability of natural-language mediation under adversarial conditions warrants dedicated safety analysis. Third, the OD matrix, while calibrated to the 2021 Census commuter-flow aggregates [24] and TfL employment-density statistics [1], is not a live feed: the integration with a real-time OD matrix, live TfL GTFS feeds, and a production-grade congestion scorecard would require data-sharing agreements that the framework (§3) is designed to support but does not itself supply.

### *6.5. Future Work*

We identify three priority directions. (1) Scale the SUMO testbed of §5.1 from the present 4×4 synthetic grid to an OSM-derived model of a real London arterial corridor (e.g., the A40 Westway or the Holloway Road–Seven Sisters axis), fed by live TfL speed-API data, and integrate LLM-mediated negotiation between intersection controllers and service agents. (2) Develop federated accessibility agents that operate across city boundaries, using NGSI-LD [29] as the substrate. (3) Study the long-term governance implications of delegating mobility coordination to autonomous agents, extending the time-varying weighting pattern validated in Scenario D to a richer class of events (construction, weather, major events) and to the fairness-aware MARL formulation of [35].

## 7. Conclusions

This paper offers practitioners a drop-in utility composition rule (Eq. 1) that takes any composite traffic- planning index and any accessibility surface a city already trusts, and produces a per-decision audit trail decomposable into $\alpha I_{\text{global}} + \beta A_{\text{local}} + \gamma P_a$. The London case study shows that this rule, parameterised by a small number of governance weights, produces equity- preserving improvements in peak-hour performance at near-zero aggregate efficiency cost (Scenario A) and is the only formulation in our four-way comparison that simultaneously addresses equity and incident response (Scenario D); a 4×4 SUMO sweep traces the same Pareto frontier microscopically.

The framework is not London-specific: applying it to other megacities requires a locally calibrated composite index, an accessibility study at appropriate resolution, and the same agent-protocol stack. Cross-city work [25,26] suggests that the accessibility-deficit Gini transfers well across megacities; whether a city-specific index such as the Moscow integral of [9] can substitute for the PTAL/INRIX pairing used here is an open empirical question the framework is structured to answer with a single weight-vector re-elicitation. The systems-engineering work that this paper scopes but does not carry out (live protocol-ring deployment, runtime LLM agents, cross-organisation interoperability) is the substantive next phase.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2



## Reproducibility

All simulation code, SUMO configurations, raw result CSVs, and the LaTeX source of this manuscript are released under the MIT licence in a companion public repository at https://github.com/tapetrova/agentic-ai-urban-mobility. All seeded runs use `seed=20260417`; the abstract Python/NumPy simulation is deterministic and completes in under one second on a laptop-class CPU. The SUMO 4×4 sweep (120 simulator runs) completes in approximately 20 minutes on a single core. Package pins, bundle structure, and a full `README` are included in the repository.

**Author Contributions:** All contributions (conceptualisation, methodology, software, formal analysis, investigation, writing of the original draft, and revision) were carried out by the sole author (T.P.).

## Appendix A. Supplementary SUMO Testbed (2×2)

Earlier iterations of this work validated the TraCI-backed agentic pipeline on a smaller 2×2 signalised grid with 1 920 vehicles over 30 min and eight origin-destination pairs; we retain the results here as a robustness check. In that testbed the Agentic policy at low peripheral boost ($b \in \{0.00, 0.04\}$) dominates a SCOOT-style adaptive controller on mean travel time (−6.7 %, −7.9 %), travel- time coefficient of variation (−6.5 %, −13.1 %) and 95th-percentile tail delay (−29.9 %, −26.5 %) simultaneously (Figure A1). The 2×2 grid, however, lacks path diversity and cannot exhibit network-level OD equity: a single low-frequency OD pair remains a bottleneck under every policy because the grid has essentially no alternative routes. The 4×4 testbed of §5.1 is the principal empirical vehicle for network-scale claims; we cite the 2×2 here only to demonstrate that the qualitative Pareto-frontier shape (including the low-boost simultaneous dominance finding) is preserved across grid sizes.

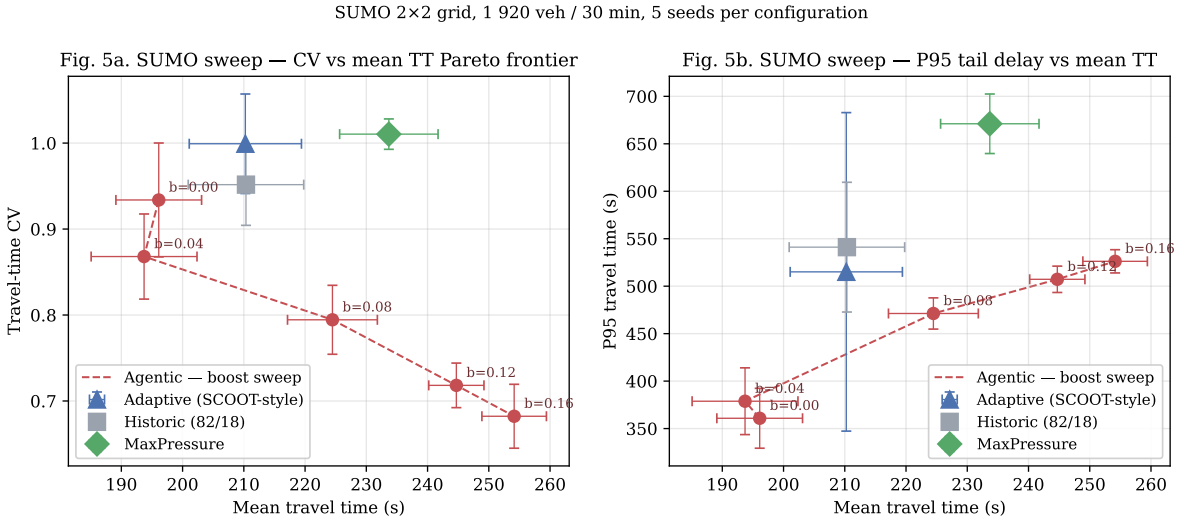


**Figure A1.** Supplementary 2×2 SUMO Pareto curves (5 seeds per configuration). **Left:** CV vs. mean TT. **Right:** $P_{95}$ vs. mean TT. The Agentic family (red, dashed) at $b \in \{0.00, 0.04\}$ dominates the adaptive baseline on all three axes simultaneously.

## References

1. Transport for London. Travel in London 2023: Consolidated Estimates of Total Travel and Mode Shares. Technical report, Greater London Authority, 2023.
2. Transport for London. Assessing Transport Connectivity in London: Public Transport Accessibility Levels (PTALs). https://tfl.gov.uk/, 2015.
3. INRIX. Global Traffic Scorecard. https://inrix.com/scorecard/, 2024. Annual composite measure of congestion across >1000 cities.
4. TomTom. TomTom Traffic Index. https://www.tomtom.com/traffic-index/, 2024. City-level congestion ranking based on probe-vehicle data.
5. Schrank, D.; Eisele, B.; Lomax, T. Urban Mobility Report. Technical report, Texas A&M Transportation Institute / INRIX, 2021. Annual composite index of urban congestion used in the United States.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2




6. Gillis, D.; Semanjski, I.; Lauwers, D. How to Monitor Sustainable Mobility in Cities? Literature Review in the Frame of Creating a Set of Sustainable Mobility Indicators. *Sustainability* **2016**, *8*, 29. https://doi.org/10.3390/su8010029.
7. Litman, T. Well Measured: Developing Indicators for Sustainable and Livable Transport Planning. Technical report, Victoria Transport Policy Institute, 2016.
8. Banister, D. The Sustainable Mobility Paradigm. *Transport Policy* **2008**, *15*, 73–80. https://doi.org/10.1016/j.tranpol.2007.10.005.
9. Petrova, T.; Grunin, A.; Shakhbazyan, A. Integral Index of Traffic Planning: Case-Study of Moscow City's Transportation System. *Sustainability* **2020**, *12*, 7395. https://doi.org/10.3390/su12187395.
10. Hansen, W.G. How Accessibility Shapes Land Use. *Journal of the American Institute of Planners* **1959**, *25*, 73–76. https://doi.org/10.1080/01944365908978307.
11. Geurs, K.T.; van Wee, B. Accessibility Evaluation of Land-Use and Transport Strategies: Review and Research Directions. *Journal of Transport Geography* **2004**, *12*, 127–140. https://doi.org/10.1016/j.jtrangeo.2003.10.005.
12. Shen, Q. Location Characteristics of Inner-City Neighborhoods and Employment Accessibility of Low-Wage Workers. *Environment and Planning B: Planning and Design* **1998**, *25*, 345–365. https://doi.org/10.1068/b250345.
13. Páez, A.; Scott, D.M.; Morency, C. Measuring Accessibility: Positive and Normative Implementations of Various Accessibility Indicators. *Journal of Transport Geography* **2012**, *25*, 141–153. https://doi.org/10.1016/j.jtrangeo.2012.03.016.
14. Petrova, T.; Grunin, A. Transport Accessibility of Urban Districts in a Megapolis: Insights from Moscow. *Urban Science* **2024**, *8*, 36. https://doi.org/10.3390/urbansci8020036.
15. Bazzan, A.L.C.; Klügl, F. A Review on Agent-Based Technology for Traffic and Transportation. *The Knowledge Engineering Review* **2014**, *29*, 375–403. https://doi.org/10.1017/S0269888913000118.
16. Li, Z.; Liang, J.; Du, X.; Liu, Y.; Jiang, H.; Li, Z.; Long, C.; Wei, H.; Ketter, W. A Survey on Multi-Agent Reinforcement Learning for Adaptive Transportation Solutions. *SN Computer Science* **2025**, *6*. https://doi.org/10.1007/s42979-025-04475-3.
17. Petrova, T.; Bliznioukov, B.; Puzikov, A.; State, R. From Semantic Web and MAS to Agentic AI: A Unified Narrative of the Web of Agents. arXiv preprint, 2025, [arXiv:cs.MA/2507.10644].
18. Liu, K.; Yigitcanlar, T.; Mehmood, R.; Corchado, J.M.; Fu, X. Large Language Models in Urban Planning: A Systematic Review and Conceptual Framework. *Journal of Urban Technology* **2025**. https://doi.org/10.1080/10630732.2025.2556551.
19. Nie, T.; Sun, J.; Ma, W. Exploring the Roles of Large Language Models in Reshaping Transportation Systems: A Survey, Framework and Roadmap. *Artificial Intelligence for Transportation* **2025**, *1*, 100003, [2503.21411].
20. Google and Linux Foundation. Announcing the Agent2Agent Protocol (A2A). https://a2a-protocol.org/, 2025.
21. Anthropic. Introducing the Model Context Protocol (MCP). https://www.anthropic.com/news/model-context-protocol, 2024.
22. Digital Twin Consortium. AI Agent Capabilities Periodic Table Framework. Technical report, Digital Twin Consortium, 2025.
23. Ivanov, D. Agentic Digital Twins: Bridging Model-Based and AI-Driven Decision-Making Support. *International Journal of Production Research* **2026**. https://doi.org/10.1080/00207543.2026.2630277.
24. Office for National Statistics. Origin-Destination Data, England and Wales, Census 2021 (Table ODWP01EW: Location of Usual Residence and Place of Work, Local Authority stratum). Nomis, https://www.nomisweb.co.uk/sources/census_2021_od, 2022. Accessed April 2026.
25. Pritchard, J.P.; Tomasiello, D.B.; Giannotti, M.; Geurs, K.T. An International Comparison of Equity in Accessibility to Jobs: London, São Paulo and the Randstad. *Findings* **2019**. https://doi.org/10.32866/9334.
26. Hwang, U.; Lieu, S.J.; Guan, H.; Dalmeijer, K.; Van Hentenryck, P.; Guhathakurta, S. Measuring Transit Equity of an On-Demand Multimodal Transit System. *Journal of the American Planning Association* **2024**, *91*, 72–87. https://doi.org/10.1080/01944363.2024.2323470.
27. Coelho, P.d.S.; Nobrega, R.A.d.A.; Oliveira, L.K.d.; Humberto, M. The Trade-off Between Equity and Quality in Public Transportation: Lessons from a Brazilian Case Study. *npj Sustainable Mobility and Transport* **2025**, *2*. https://doi.org/10.1038/s44333-025-00039-3.
28. European Telecommunications Standards Institute. ETSI TS 103 828: SmartM2M; SAREF Ontology for Smart Applications. Technical Report V1.1.1, ETSI, 2024.


© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202604.1407.v2




29. ETSI Industry Specification Group on Context Information Management. NGSI-LD API Specification. https://ngsi-ld.org/, 2024.
30. Lai, S.; Xu, Z.; Zhang, W.; Liu, H.; Xiong, H. LLMLight: Large Language Models as Traffic Signal Control Agents. In Proceedings of the Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2025. https://doi.org/10.1145/3690624.3709379.
31. Yuan, Z.; Lai, S.; Liu, H. CoLLMLight: Cooperative Large Language Model Agents for Network-Wide Traffic Signal Control. *arXiv preprint* **2025**, [2503.11739].
32. Tu, W.; Li, J.; Xiao, F.; Wang, X.; Lu, Y. Integrating Large Language Models into Traffic Systems: Integration Levels, Capability Boundaries, and an Information-Theoretic Perspective. *Entropy* **2026**, *28*, 211. https://doi.org/10.3390/e28020211.
33. Jonnala, R.; Liang, G.; Yang, J.; Alsmadi, I. Exploring the Potential of Large Language Models in Public Transportation: A San Antonio Case Study. In Proceedings of the AAAI 2025 Workshop on AI for Urban Planning, 2025, [2501.03904].
34. Zhang, Z.; Sun, Y.; Wang, Z.; Nie, Y.; Ma, X.; Li, R.; Sun, P.; Ban, X. Large Language Models for Mobility Analysis in Transportation Systems: A Survey on Forecasting Tasks. *Transportation Research Record* **2026**, *2681*.
35. Yu, G.; Siddique, U.; Weng, P. Fair Multi-Agent Reinforcement Learning for Traffic Control. *ACM Journal on Autonomous Transportation Systems* **2025**, *1*. https://doi.org/10.1145/3749378.
36. Keeney, R.L.; Raiffa, H. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*; Wiley: New York, 1976. Reprinted Cambridge University Press, 1993.
37. Miettinen, K. *Nonlinear Multiobjective Optimization*; Vol. 12, *International Series in Operations Research & Management Science*, Springer, 1999. https://doi.org/10.1007/978-1-4615-5563-6.
38. Boyd, S.; Vandenberghe, L. *Convex Optimization*; Cambridge University Press: Cambridge, 2004.
39. Sacoto-Cabrera, E.J.; Perez-Torres, A.; Tello-Oquendo, L.; Cerrada, M. IoT, AI and Digital Twins in Smart Cities: A Systematic Review for a Thematic Mapping and Research Agenda. *Smart Cities* **2025**, *8*, 175. https://doi.org/10.3390/smartcities8050175.


**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.