sustainability

MDPI

Article

Reframing Low-Emission Zones as Adaptive Decision Infrastructures: A Digital-Twin Framework and Lifecycle Methodology for Sustainable Urban Air Quality

Antonio Cantalapiedra-Asensio [1,2] and José Carlos Romero [3,*]

1 DBA in Management and Technology, Faculty of Economics and Business Administration (ICADE), Universidad Pontificia Comillas, 28015 Madrid, Spain; antonio@georeplica.com
2 Georeplica Technologies, Calle General Álvarez de Castro 41, 28010 Madrid, Spain
3 Institute for Research in Technology (IIT), ICAI School of Engineering, Universidad Pontificia Comillas, 28015 Madrid, Spain
* Correspondence: jcromero@comillas.edu

Abstract

Road transport is a leading source of urban nitrogen oxides (NO_x) and fine particulate matter ($PM_{2.5}$)—a public-health and urban-sustainability challenge—and Low-Emission Zones (LEZs) are Europe's principal response. Yet most are governed statically, unable to track conditions changing by the hour and the street. A digital twin, treated as decision infrastructure rather than a 3D model, recasts the LEZ as an *adaptive decision infrastructure*: a closed loop of sensing, modelling and rule-based adjustment. We develop a scalable, five-phase lifecycle methodology with auditability and GDPR-by-design built in, and derive three falsifiable hypotheses—efficiency, data integration, responsiveness—defining a research agenda. We test only the first. A diagnostic reading of London's ULEZ shows its unimplemented phases are precisely those that close the loop. A proof-of-concept on real hourly NO_2 from five London sites (2023–2024) tests efficiency: at equal abatement effort, adaptive targeting avoids significantly more elevated-pollution hours than a uniformly stricter policy (about 37% versus 27%; 95% CI excludes parity), the advantage rising with forecast quality. This demonstrates the *mechanism* in reduced form, not a generalizable figure for a deployed system. By making regulation more responsive and accountable, it advances the Sustainable Development Goals on health, sustainable cities and climate (SDGs 3, 11, 13).

Keywords: low-emission zones; digital twins; urban air quality; adaptive policy; decision infrastructure; smart cities; sustainable urban mobility; environmental governance; Sustainable Development Goals; data integration

Check for updates

Academic Editor: Fernando Fonseca

Received: 22 June 2026
Revised: 6 July 2026
Accepted: 8 July 2026
Published: 11 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Urban road transport remains a leading source of the pollutants that most damage city air quality—nitrogen oxides (NO_x) and fine particulate matter ($PM_{2.5}$)—with well-documented effects on public health. Low-Emission Zones (LEZs) have become the principal regulatory response, particularly in Europe, restricting or charging the most polluting vehicles in designated areas. Their salience is rising. The revision of the EU Ambient Air Quality Directive moves limit values closer to World Health Organization guidance. Cleaner, fairer mobility is also central to the United Nations Sustainable Development Goals on health, sustainable cities and climate action—SDG 3, SDG 11 and SDG 13—and to keeping cities within safe environmental limits [1]. An LEZ must also balance emission reduction against mobility access, social equity and local economic activity. Its performance is therefore as much a matter of governance as of technology [2].

Yet most LEZs are governed as *static* systems. Boundaries, eligibility, exemptions and enforcement are fixed for long periods and changed only through slow administrative cycles, whereas the conditions that drive exposure—traffic, weather, fleet composition and compliance—vary hour by hour and street by street. A static zone can lower average concentrations but cannot react to short-term episodes, congestion shocks or locally uneven exposure. The gap we address is therefore one of policy *design*: LEZs work, but they are not yet built to adapt [3].

A digital twin (DT) offers a way to close this gap. In current urban research a DT is not a static 3D replica of a city but a continuously updated, data-driven model of the real system that supports decisions on it, coupling sensing, simulation and forecasting in a feedback loop [4,5]. Applied to an LEZ, a DT can convert a fixed compliance zone into an adaptive one. It senses the city in near-real time, simulates the likely effect of alternative rules, and feeds that evidence back into enforcement and access decisions. This is the sense–simulate–act feedback already demonstrated for digital-twin transport and IoT-enabled control systems [6–8]. Throughout, we use the term *DT-enabled LEZ* for an LEZ governed through such a loop.

The building blocks of such a system already exist—LEZ impact studies, traffic and emissions modelling, and urban DT platforms—but they are developed separately. What is missing is an integrated account of how a city moves from policy objectives to an operating, adaptive LEZ, in which sensing, modelling, governance and evaluation form a single workable process [9].

This paper provides that account. Its central contribution is one idea made operational: reframing the LEZ as an *adaptive decision infrastructure*—a closed data–model–decision loop—and translating it into a design-to-operation lifecycle methodology that scales with a city's digital maturity and embeds governance from the outset, namely GDPR-by-design, auditability and equity. What is new is the integration, not the ingredients, and it is worth stating precisely against the closest work. Urban DT platforms [6,10,11] supply the sensing–modelling substrate but are not organised around a policy. The DT for mobile-source emissions of [12] monitors emissions but does not govern an instrument. Deep-reinforcement-learning and big-data traffic-signal control reduce emissions [13,14] but optimise a single control layer, not a route from policy objectives to an auditable operating system. The "dynamic LEZ" literature [3], meanwhile, argues that zones *should* adapt without specifying how a city would build, govern and evaluate one. What none of them provides—and what this paper contributes—is the connective tissue: a policy-first lifecycle that turns these technical pieces into a single, auditable *unit of governance*, treating the closed loop itself, rather than any individual model, as the object of design. The novelty is therefore threefold. The closed loop—not any single model—is the unit of design; governance and an evaluation roadmap are integral rather than retrofitted, with governance covering GDPR-by-design, auditability and equity; and a maturity ladder makes the same framework adoptable by cities that lack full real-time data. Table 1 positions the contribution against the closest prior work. The paper is primarily conceptual and methodological: we develop and argue for the framework, illustrate its logic, and provide a reduced-form, real-data test of its efficiency claim, while setting out the data requirements for the full empirical validation that should follow.

Table 1. Positioning of the proposed DT-LEZ framework against the closest strands of prior work, on an ordered scale "No" < "Partial" < "Substantial" < "Yes". "Yes" = addressed as a core concern; "Substantial" = substantially addressed but not fully integrated; "Partial" = touched but not integrated; "No" = out of scope. Entries for the present framework denote conceptual or by-design provision, not empirical demonstration. The compared strands were selected as the prior families closest to the framework's three layers and to the "dynamic LEZ" proposal, identified from the systematic reviews cited in Section 2 and screened for direct relevance to adaptive, DT-enabled air-quality governance.

Prior Strand	Full Design to Operation Lifecycle	DT as Decision Infrastructure	Adaptive/ Conditional Policy Logic	Governance & GDPR by Design	Scalability by Maturity	Evaluation Roadmap
Urban DT reviews & platforms [10,11,15]	Partial	Partial	No	Partial	No	No
DT for transport/emissions [6,12]	Partial	Partial	Partial	No	No	Partial
Adaptive traffic-signal control (DRL/MARL) [13,14]	No	Partial	Yes	No	No	Partial
LEZ impact & "dynamic LEZ" studies [3,16,17]	No	No	Partial (concept)	Partial	No	Partial (empirical)
This paper (DT-LEZ framework)	Yes (conceptual)	Yes	Yes	Substantial	Partial	Yes (roadmap)

From the framework we derive three falsifiable hypotheses, which together define a research agenda rather than a set of established results. For each, the paper specifies the mechanism the framework implies and how it can be tested (Section 5.5), and it provides initial, reduced-form evidence for the first (Section 5.4).

Hypothesis 1 (efficiency)**.** *Under equivalent traffic, weather and compliance conditions, an adaptive, DT-enabled LEZ reduces pollutants more efficiently than a static regime, because continuous sensing, simulation and rule adjustment direct effort to where and when it matters most [3,8,14].*

Hypothesis 2 (data fragmentation)**.** *The dispersal of mobility, environmental, enforcement and administrative data across disconnected systems is a binding constraint on the effectiveness of current LEZs [18,19].*

Hypothesis 3 (responsiveness)**.** *Real-time, rule-based adjustment improves responsiveness to short-term shocks—traffic surges, pollution episodes, adverse weather and compliance variation—within timescales that administrative review cannot match [20,21].*

The remainder of the paper is organised as follows. Section 2 reviews Low-Emission Zones and the limits of static design, digital twins as decision infrastructure, and governance. Section 3 develops the conceptual framework and its analytical formalisation. Section 4 presents the DT-LEZ lifecycle methodology. Section 5 reads London's ULEZ against the lifecycle and sets out the requirements for empirical validation. Sections 6 and 7 discuss implications, limitations and conclusions.

2. Background

2.1. Low-Emission Zones: Instrument, Evidence and the Limits of Static Design

LEZs and their stricter variants, Ultra-Low-Emission Zones (ULEZs), restrict, charge or discourage the entry of high-emitting vehicles into a defined urban area. They operate through three broad mechanisms: access restriction by emission standard, typically the European Euro classes; monetary charging; and, increasingly, time-differentiated or episodic rules. These are enforced through vehicle labelling schemes such as the French Crit'Air sticker or the German *Umweltplakette* and, where charging applies, through Automatic

Number Plate Recognition (ANPR) [17]. The instrument has spread quickly: a recent systematic review counts 363 active zones across Europe by October 2024, under widely differing designs, with the number still growing [17,22]. London's ULEZ is the most intensively studied case, owing to its 24 h operation, dense monitoring network and public reporting [23,24].

The evidence that LEZs improve air quality is generally positive but heterogeneous. State-space modelling of London's 2019 ULEZ identified reductions in roadside NO_x once meteorological and temporal confounders were removed [25], consistent with meteorological-normalisation methods developed to separate intervention effects from weather variability [26]. Systematic reviews of LEZs and congestion-charging zones report air-quality and health benefits, but stress that effect sizes vary substantially with zone design, enforcement intensity, fleet composition and complementary transport measures, and that some schemes show small or uncertain effects [16,17,27]. Other work links LEZs to welfare and health gains, including improvements in children's respiratory health. It also notes that much of the benefit flows from accelerated fleet renewal rather than from the boundary itself [28,29].

Because LEZs alter the cost and feasibility of urban mobility, their effects are also distributional. They can fall disproportionately on low-income households, small businesses and freight operators that cannot easily replace older vehicles, raising environmental-justice concerns documented across European cities [30,31]. Recent analysis shows that LEZ impacts on job accessibility are unevenly distributed across income groups [32], while evidence on high-street economic activity is mixed [33]. Public acceptability is correspondingly pivotal. Attitudes, perceived fairness and trust strongly shape whether a scheme survives politically [34]. Equity and legitimacy are thus not side-constraints but determinants of an LEZ's effectiveness.

Against this evidence, the recurring limitation is structural rather than substantive. Boundaries, eligibility rules, exemptions and enforcement are typically fixed for long periods and revised only through slow administrative or political cycles, whereas the determinants of exposure vary continuously: NO_x and $PM_{2.5}$ concentrations change with the hour, street geometry, wind, fleet mix and traffic intensity. Street-level studies further reveal pollutant fields that zone boundaries alone cannot capture [35,36]. A static zone may therefore lower average concentrations while underperforming during peak congestion, adverse weather or localised episodes. Its boundary effects are also contested: some schemes have shifted congestion and traffic to surrounding areas [37], whereas boundary monitoring around London's ULEZ found no comparable displacement of pollution [23,24]. This unevenness itself reflects the limits of one-size-fits-all fixed designs. Causal analysis of London suggests that further gains will require measures beyond the existing fixed ULEZ [38], and a nascent literature on *dynamic* LEZs argues that future zones should adjust to live environmental and mobility conditions [3]. Three questions thus remain unresolved by the static paradigm. How much of the benefit comes from the boundary itself rather than from background fleet renewal [28]? Do zones remove pollution or merely displace it, given that the evidence is mixed [23,37]? And how could a zone adjust to conditions in operation, when proposals for dynamic LEZs remain largely conceptual and empirically untested [3]? The instrument, in short, is effective but constrained by a design that cannot adapt—the policy-design gap this paper addresses.

2.2. Digital Twins as Decision Infrastructure

A DT is increasingly central to smart-city research, but its meaning has shifted. Early urban DTs were essentially high-fidelity 3D representations; the current literature describes them as continuously updated, data-driven models that mirror a physical system and support decisions on it, integrating Internet-of-Things (IoT) data, machine learning, simulation and analytics into a feedback loop. Systematic and bibliometric reviews document this reorientation toward decision support [4,10], while a framework report and a smart-city–DT study set out the enabling components and global practice [5,39]. Reviews of urban DTs increasingly characterise them as socio-technical infrastructures for monitoring, modelling and governing complex city systems rather than visualisation tools [15,40,41]. Socio-technical [42] and critical or ethical [43] analyses reinforce this reading. We adopt this operational reading: for an adaptive LEZ, a DT is a *decision infrastructure* that connects what the city senses to the rules it enforces. The rest of this subsection sets out the three layers such an infrastructure requires—perception, modelling, and integration with decision support—which Section 3 assembles into a closed loop.

The *perception layer* fuses environmental and mobility sensing. Air quality is captured by fixed regulatory stations, low-cost IoT sensor nodes, mobile sensors mounted on fleets and meteorological instruments, while mobility is captured by inductive loops, high-definition cameras, ANPR, radar and LiDAR, and connected-vehicle or vehicle-to-infrastructure (V2I) feeds. Reviews of IoT-enabled smart cities catalogue these sensing technologies and architectures [44,45], with complementary work on automated-vehicle and smart-city infrastructure [46] and on IoT-based adaptive traffic management [47]. Denser, lower-cost sensor networks are widening what cities can observe in near-real time. Their calibration, accuracy and coverage, however, remain open questions for operational use [48]. Edge and cloud processing complete the layer, supporting low-latency tasks such as detection and compliance screening alongside larger-scale historical analysis [44,45].

The *modelling layer* turns observation into estimation and prediction. Traffic is represented at macroscopic, mesoscopic or microscopic scales, and short-term states are increasingly forecast with machine-learning, deep-learning and hybrid methods, including graph neural networks that treat the road network as a spatial graph. Broad reviews survey traffic-flow prediction [20,49] and a dedicated survey covers graph neural networks for forecasting [21]. Multi-agent simulation and reinforcement learning support adaptive control, as demonstrated for IoT-enabled adaptive traffic management [47] and cooperative multi-agent signal control [8,14], whose logic transfers to LEZ operation. Emissions models then link vehicle dynamics—speed, acceleration, idling, vehicle type and fuel technology—to pollutant output [50,51]. Ideally they are coupled with dispersion modelling [52], so that mobility decisions can be related to exposure [53]. Domain-specific DTs already combine several of these elements, from digital-twin intelligent-transport systems [6] to a DT framework for monitoring urban mobile-source emissions [12].

The *integration layer* unifies these streams and connects them to decision support. Air-quality, traffic, registry, enforcement and administrative data are typically held by different agencies and vendors, so interoperability, shared standards and open APIs are prerequisites rather than afterthoughts. Without them a DT reproduces existing silos instead of overcoming them [18,19]. Integration platforms such as Snap4City illustrate how heterogeneous urban data can be unified for operational use [11], while work on the complexity and critical informatics of city DTs cautions that more data and detail do not automatically yield better governance [54,55]. Together these layers mark the distinction the paper builds on: conventional LEZ research asks whether a fixed policy reduced emissions after the fact, whereas a DT-enabled approach asks how an LEZ can be operated as an adaptive system—which makes integration across fragmented data systems a precondition for effectiveness, as Hypothesis 2 contends.

These capabilities come with practical constraints that bound what a DT-LEZ can be trusted to do. The framework must treat these as first-order design concerns rather than implementation details. Sensing is imperfect: low-cost air-quality nodes drift and need regular calibration against reference instruments, and camera and ANPR pipelines degrade in poor light, occlusion and adverse weather [44,48]. Models carry irreducible uncertainty—forecast error, and parameter and structural uncertainty in emission and dispersion estimates—so decisions must be robust to it rather than assume point predictions are exact [52,53]. Scaling from a pilot to a whole city multiplies data volume, integration surface and operating cost. A system that informs enforcement is also a security and privacy target: camera feeds, location traces and registry links demand protection against tampering and unauthorised access, so cybersecurity and resilience are governance requirements rather than optional extras [18,55]. The framework returns to these as explicit feasibility and gate criteria in Section 4.

2.3. Governance, Privacy and Equity

Technical capability is necessary but not sufficient: how a DT-enabled LEZ is governed determines whether it is lawful and legitimate. Because such systems may process number plates, location traces, camera imagery, compliance records and socio-economic indicators, much of their data is personal data under the EU General Data Protection Regulation. Lawful basis, purpose limitation, data minimisation, pseudonymisation, retention limits and access control must therefore be built in by design rather than retrofitted, and data-subject transparency rights must be operable in practice [56,57]. Purpose limitation is especially delicate here. Data gathered for air-quality management can readily be repurposed for enforcement or surveillance.

Beyond compliance, the system needs data governance and algorithmic accountability. Organising data for trustworthy, algorithm-assisted decisions requires clear ownership, quality assurance, documentation and audit trails of model outputs and policy actions [19,58]. Where models inform enforcement or exemptions they can encode historical inequalities or data gaps, so controlling algorithmic bias, together with explainability and auditability, becomes a governance requirement rather than a technical nicety [59]; open-data practices and usable public interfaces further enable external scrutiny [60].

Finally, governance must address distribution and legitimacy. Adaptive operation can improve fairness through continuous distributional monitoring and targeted, time-limited exemptions, but it can equally intensify surveillance or entrench bias where safeguards are weak. The distributional stakes are well documented: in analyses of inequalities in LEZ access effects [32], in environmental-justice assessments of specific zones [31] and in policy reviews of equity in low- and zero-emission zones [30]. Public acceptability rests on perceived fairness, transparency and participation [34], and smart-city scholarship warns against technocratic deployments that bypass democratic oversight [61,62]. Taken together, these strands position an adaptive LEZ at the intersection of air-quality policy, digital-twin systems and accountable urban governance: effective but constrained by static design, technically enabled by DT sensing, modelling and integration, and dependent on legal and social readiness—the three foundations the framework in Section 3 brings together.

3. The DT-LEZ Framework: A Closed-Loop Decision Infrastructure

This section develops the paper's core contribution. Section 2 established the mismatch between fixed regulatory design and dynamic urban conditions; the conceptual move this paper makes is to recast that mismatch as a *decision-infrastructure* problem rather than a control or modelling one. The question is not how to predict pollution more accurately, but how a city can convert what it senses into timely, accountable policy action—closing the gap between observing a problem and adjusting the rule that bears on it [4,5]. We first describe that system as a continuous loop, then trace what it changes for policy, and close by linking both to the paper's three hypotheses.

3.1. The Data–Model–Decision Loop

Concretely, the framework joins five stages into a single, repeating cycle (Figure 1). It begins where the city is observed: environmental and mobility sensing supplies NO_x and $PM_{2.5}$ readings, traffic flows and vehicle classes, fleet composition, weather, exposure and compliance records. These heterogeneous streams become useful only once an integrated data layer cleans, standardises and links them across agencies—the step that confronts the very fragmentation identified earlier as a barrier to adaptive governance. On that foundation sits the digital-twin core, where traffic, emissions and dispersion models, machine-learning forecasts and scenario tools both describe current conditions and estimate what would follow from alternative choices. A policy engine then turns those estimates into governance logic—thresholds, triggers, exemptions and enforcement priorities. The final stage renders these as operational adjustments: enforcement intensity, temporary restrictions, time-window changes, public alerts or routing measures. The cycle does not stop there but feeds back into the next round of sensing. What this circularity buys is reduced *policy latency*—the time between detecting an environmental or mobility problem and acting on it [7,63].

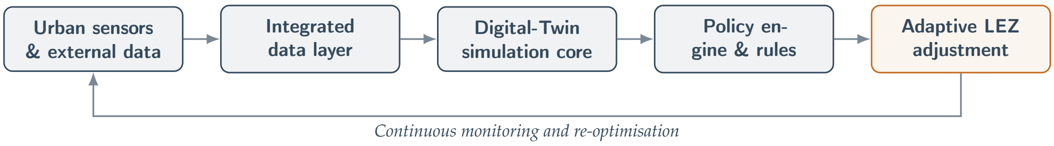

Figure 1. Digital-twin-enabled adaptive LEZ feedback loop transforming static regulation into real-time policy adjustment.

3.2. From Fixed Rules to Conditional Policy Logic

What the loop changes, in policy terms, is the status of the rule itself. A static LEZ treats policy intensity as a constant: access rules are set in advance and applied uniformly across time and space, and the only way to change them is the slow administrative cycle. The framework replaces that constant with a *conditional* rule—one whose intensity depends on observed and forecast conditions, among them traffic, measured pollutants, short-term forecasts and compliance. In practice, enforcement can intensify during a forecast NO_x exceedance, temporary restrictions can activate around vulnerable locations, rerouting can be recommended during congestion, and alerts can issue when exposure thresholds are likely to be breached. This is the whole of the static-versus-adaptive distinction: not a different instrument, but the same instrument with its intensity made responsive to the city it regulates. This logic is increasingly familiar from traffic prediction and adaptive control [8,20].

Conditional, however, does not mean automated or legally unchecked. The DT provides *decision support within predefined governance constraints*: thresholds and triggers are set in advance, exemptions are justified, and actions are logged for review [42,64]. It mediates between real-time urban complexity and formal policy authority, structuring evidence and weighing alternatives without replacing the policymaker. For an instrument with legal, ethical, economic and social consequences, an adaptive LEZ must therefore be at once technically responsive and democratically governable. Concretely, the framework operationalises these requirements at defined points. Privacy is handled through GDPR-by-design and a data-protection impact assessment gated into Phase C (Supplementary Materials, Section S4). Accountability is handled through audit trails of data, model outputs and decisions maintained in Phase E. And fairness is handled through a vulnerability-weighted objective and an equity constraint $g(\cdot) \leq 0$ in the policy engine of Section 3.3, together with the explicit, time-limited exemptions set in Phase B.

These two moves—closing the loop and making the rule conditional—are what give the paper's three hypotheses their force. The loop is what could let adaptive operation reduce pollutants more efficiently than fixed rules (Hypothesis 1); it works only by overcoming the data fragmentation that otherwise constrains current LEZs (Hypothesis 2); and the conditional rule is what enables responsiveness within the timescale of urban shocks (Hypothesis 3). These remain largely conceptual claims, though Hypothesis 1 receives reduced-form empirical support. Section 5 maps the ULEZ onto the lifecycle and tests the efficiency claim on real data in a proof-of-concept (Section 5.4), while Section 5.6 sets out the data required for full validation.

3.3. The Policy Engine, Formally

To make the conditional rule precise—and to show exactly where uncertainty enters—we state the policy engine as a constrained decision problem. At each decision step t and location s, the DT returns a predictive distribution of the baseline (no-action) pollutant concentration $c^0_{s,t}$, summarised by a forecast $\hat{c}^0_{s,t}$ and its uncertainty. The background B_s—the share a zone cannot influence—is separated out, so the engine acts only on the controllable excess $(c^0_{s,t} - B_s)^+$. The *decision variable* is the abatement intensity $a_{s,t} \in [0, \bar{a}]$: the stringency of enforcement or access restriction applied at (s,t), bounded above by operational capacity $\bar{a}$. Under action $a_{s,t}$ the realised concentration is $c_{s,t} = B_s + (1 - a_{s,t})\,(c^0_{s,t} - B_s)^+$.

The engine chooses $\{a_{s,t}\}$ to minimise the *vulnerability-weighted* expected exceedance of a concentration of concern T over a planning horizon H:

$$\min_{\{a_{s,t}\}} \sum_{t=1}^{H} \sum_{s} w_s \, \mathrm{E}\big[(c_{s,t} - T)^+\big], \tag{1}$$

where $w_s \geq 1$ is a vulnerability weight that places greater value on reducing exposure in more vulnerable locations, so that equity enters the objective itself and not only the constraints. The program is subject to (i) a budget on total intervention effort, $\sum_{t,s} a_{s,t}\,(c^0_{s,t} - B_s)^+ \leq \mathcal{E}_{\max}$—the equal-effort constraint used in Section 5.4; (ii) equity constraints $g(\{a_{s,t}\}) \leq 0$ bounding the distributional burden across neighbourhoods and groups; and (iii) legal–operational constraints (maximum duration and frequency of restrictions, mandatory exemptions, due-process logging).

The *triggering logic* is the engine's online approximation to this program. Because $\hat{c}^0_{s,t}$ is uncertain, the trigger is probabilistic—intervene when the forecast probability of exceeding the threshold is high enough:

$$a_{s,t} = \alpha \cdot \mathbf{1}\Big[\Pr\big(\hat{c}^0_{s,t} > T\big) \;\geq\; p^\star\Big], \tag{2}$$

where α is the intervention effectiveness and $p^\star$ a confidence level set by the city's tolerance for false activations. To prevent chattering around the threshold, the trigger carries hysteresis—activation at T but deactivation only below $T - \delta$—so that decisions stay stable under noisy forecasts. The predictability of policy is therefore a direct function of forecast quality: a sharper predictive distribution pushes $\Pr(\hat{c}^0_{s,t} > T)$ above $p^\star$ on exactly the episodes that matter while leaving low-pollution hours untouched. The deterministic rule evaluated in Section 5.4 is the special case of a point forecast with $p^\star \to 0^+$.

This formulation locates the framework's design choices explicitly: the *objective* (Equation (1), vulnerability-weighted exposure), the *decision variables* ($a_{s,t}$), the *constraints* (effort budget, equity, legality) and the *uncertainty* (the predictive distribution and $p^\star$). It also clarifies the division of labour: the twin supplies the forecasts and counterfactual estimates that populate Equations (1) and (2), whereas the choice of T, $p^\star$, $\mathcal{E}_{\max}$, the vulnerability weights w_s and the equity bounds is a governance decision, not a model output.

4. Operationalising the Framework: The DT-LEZ Lifecycle

The loop described in Section 3 characterises a DT-enabled LEZ in steady-state operation; it does not, by itself, tell a city how to build that capacity or govern it over time. The lifecycle methodology supplies that missing path. It organises the work into five interdependent phases—policy objectives and baseline, policy-package design, digital-twin system design, implementation and operations, and monitoring and adaptive management—each of which leaves a concrete, auditable artefact rather than an undocumented decision. The phases are not generic project-management stages: they are derived from what the loop requires, and they are deliberately *policy-first*. The methodology starts from public-policy objectives and only then specifies the technical system to serve them. This guards against the common failure mode of technology-led smart-city projects that deploy a platform in search of a purpose.

Read against the loop, the phases play two roles. The first two set the loop's goals and rules—what the LEZ is for and how it will act—while the last three build, run and re-tune the loop itself. This makes the lifecycle the operational counterpart of the framework, and the artefacts it leaves behind are what allow the resulting decisions to be reviewed, contested and reproduced. Table 2 summarises the five phases with their inputs, outputs, decisions and responsible actors.

Table 2. DT-LEZ lifecycle phases: key inputs, main outputs/artefacts, decisions supported and primary actors.

Phase	Key Inputs	Main Outputs/Artefacts	Decisions Supported	Primary Actors
A—Policy objectives & baseline	Policy goals, legal constraints, emissions inventory, mobility data, exposure	Objectives, KPIs, baseline assessment, scope	Whether an LEZ is justified; scale and ambition	City, environmental and transport authorities
B—LEZ policy package	Baseline analysis, stakeholder input, equity considerations	Boundaries, access rules, exemptions, enforcement concept	Zone delineation, rule set, enforcement approach	Policymakers, transport authorities, legal teams
C—DT system design	Use cases, data availability, technical constraints	Architecture blueprint, data strategy, model portfolio	Sensors, models, architecture	Technical teams, data stewards, vendors
D—Implementation & operations	Architecture blueprint, procurement outcomes	Deployed DT, integration with city systems	Go/no-go, pilot-to-scale	City IT, operators, vendors, DPO
E—Monitoring & adaptive management	Operational data, KPIs, evaluation results	Performance reports, updated rules/boundaries	Policy adjustment, continuation or redesign	Policymakers, oversight bodies, public stakeholders

4.1. The Five Phases

Phase A—Policy objectives, constraints and baseline. The lifecycle opens not with technology but with intent. The municipality translates air-quality, mobility, equity and exposure goals into measurable, ex ante key performance indicators. It anchors them in a baseline diagnosis of emissions, hotspots, traffic patterns and exposed populations, drawing on evidence of LEZ air-quality and well-being effects [28], street-level morphology–

pollution relationships [35] and source apportionment for air-quality planning [36]. It also fixes the legal and institutional constraints: which authority regulates what, the admissible enforcement regime, and GDPR obligations. Later technical choices then follow policy and law rather than the contingencies of available data [17]. The artefact is a scoped set of objectives, KPIs and a baseline against which any later effect can be judged.

Phase B—The LEZ policy package. The diagnosis is then turned into a concrete instrument. Boundaries are delineated through a data-driven procedure—using exposure, vulnerability, traffic and accessibility—so that candidate zones can be tested in the DT before adoption, which reduces arbitrary or contested delineation [65]. The rule set fixes vehicle eligibility, charging and time windows. Equity-based exemptions are designed in from the start—for low-income households, essential services or operators without short-term alternatives—and made explicit, time-limited and auditable rather than bolted on afterwards [32]. Enforcement is specified concretely, from checkpoints, ANPR and computer-vision classification to registry checks and proportionate sanctioning, with coverage itself testable by simulation [66]. Complementary measures are planned to absorb displacement, among them public-transport capacity, urban-logistics consolidation, parking and modal-shift support [6]. Treating acceptability and distribution as design inputs, not post hoc assessments, is what makes the package politically durable.

Phase C—Digital-twin system design. Only now is the technical system specified, and strictly in the service of the preceding decisions. Requirements are derived from concrete use cases across the lifecycle—scenario testing at design time, operational monitoring, short-term forecasting during episodes, ex post evaluation and public communication [33]. A fit-for-purpose data and sensing plan sets out the minimum and extensible data each phase needs (Table 3), and a model portfolio spanning macro-, meso- and micro-scale and learning-based approaches is chosen by decision context, data availability and performance need rather than fixed in advance [53]. A modular reference architecture—ingestion, storage, modelling, a scenario engine, and dashboards or APIs—supports incremental deployment and limits vendor lock-in. Latency, availability, scalability, traceability and resilience are treated as first-class requirements (Figure 2) [51]. Privacy, security and auditability are built in rather than retrofitted [67].

Table 3. Minimum data requirements for DT-enabled LEZs by lifecycle phase.

Lifecycle Phase	Data Category	Minimum Required Data/Purpose
Design & baseline (A)	Air quality; mobility; population & land use	$NO_x/PM_{2.5}$; traffic volumes, fleet; density, sensitive sites—hotspot and exposure baseline
Policy design (B)	Traffic dynamics; socio-economic	OD flows, corridor loads; income proxies, vehicle ownership—zoning, rules, equity
DT system design (C)	Sensing infrastructure; contextual	Fixed stations, cameras, mobile sensors; road network, regulations—ingestion and model configuration
Operations (D)	Real-time traffic; environmental	Speeds, trajectories, classes; real-time AQ, meteorology—compliance, forecasting
Monitoring (E)	KPIs; audit logs	Emissions, compliance, travel times; decisions, model outputs—evaluation and traceability

Figure 2. Reference architecture for a DT-enabled LEZ.

Phase D—Implementation and operations. Deployment follows a phased roadmap—from a limited pilot with explicit go/no-go criteria, through scaling, to steady-state operation [52]—so that technical and institutional risk surfaces early rather than at full scale. Integration with municipal systems—traffic management, enforcement and ANPR, vehicle registries and open-data portals—depends on interoperable interfaces and agreed exchange protocols. Without them the heterogeneous formats and organisational silos noted earlier simply re-emerge [22]. Because the system's outputs can inform enforcement and policy change, operational reliability is not optional: DataOps and MLOps practices—data and model versioning, drift monitoring, validation and rollback—keep the loop trustworthy over time [68]. A clear operating model assigns the roles that make accountability concrete (Table 4).

Table 4. Governance roles and responsibilities for DT-LEZ operation.

Role	Primary Responsibilities
Policy owner (municipality)	Defines LEZ objectives, approves rules, oversees updates
System operator	Day-to-day DT operation, dashboards, alerts
Data steward	Data quality, access control, documentation
Data Protection Officer (DPO)	GDPR compliance, privacy-by-design
Technology provider/vendor	Maintains components, technical support
Oversight/audit body	Reviews performance, transparency, accountability

Phase E—Monitoring, evaluation and adaptive management. The final phase closes the loop and keeps it closed. A structured KPI framework spans air quality, mobility, equity and compliance, separating core indicators needed for regulatory reporting from extensible, locally specific ones [69]. Evaluation is designed to attribute effects rather than merely observe them, using control areas or DT-generated counterfactuals alongside before-and-after comparison and reporting the uncertainty introduced by sensor error, model assumptions and behavioural variability [70]. Transparency and accountability rest on public dashboards, periodic reporting, open-data releases and audit trails of data use, model outputs and decisions [59]. A formal policy-update protocol is what converts monitoring into adaptive policy rather than passive reporting. Such a protocol fixes review frequency, the thresholds that trigger reassessment, and how changes to boundaries, rules or exemptions are approved. It is the institutional hinge on which Hypothesis 3 turns.

4.2. Phase Gates and Computational Feasibility

Two practical questions decide whether the lifecycle is operational rather than aspirational: When is a phase *done*? And can the loop actually be *run*? We answer both explicitly. First, each phase advances only when it clears a measurable gate, stated ex ante so that progress is auditable rather than merely declared (Table 5). Second, the loop's compute budget is modest once its tiers are matched to the cadence of the decisions they support.

The computational load splits into three regimes with very different requirements. *Edge* processing handles the low-latency perception tasks—vehicle detection, ANPR and compliance screening—at sub-second rates per camera, close to the sensor and without moving raw video off-site, which also eases privacy compliance. *Near-real-time* inference—the short-term pollutant and traffic forecasting that drives the trigger of Equation (2)—runs centrally at the cadence of the decision itself. Because LEZ actions are taken on the scale of hours, hourly forecasting sits comfortably within reach of commodity cloud or on-premise infrastructure [13]. The expensive component—street-canyon dispersion at computational-fluid-dynamics fidelity—is reserved for *design-time* scenario analysis rather than the real-time loop, with reduced-order surrogates used operationally [52]. The binding constraint is therefore not raw throughput but aligning compute cadence with decision cadence.

Scalability comes from a modular, horizontally scalable architecture (Figure 2). Cost grows with sensor count and model fidelity—precisely the dimensions the minimum-viable and advanced configurations are defined to bound.

Table 5. Measurable exit criteria (phase gates) for the DT-LEZ lifecycle.

Phase	Exit Criterion (Measurable, Set Ex Ante)
A—Objectives & baseline	Quantified, time-bound KPIs approved, and a baseline whose input data meet stated completeness and quality thresholds (e.g., valid-data coverage and documented uncertainty).
B—Policy package	Rule set and boundary whose simulated effect on the Phase-A KPIs and on the distributional burden has been tested in the twin and meets pre-set acceptability and equity bounds, with stakeholder sign-off.
C—DT system design	Architecture passing acceptance tests against quantified non-functional targets (latency, availability, throughput) and clearing a data-protection impact assessment.
D—Implementation & operations	Pilot meeting pre-registered go/no-go thresholds (forecast skill, enforcement coverage, integration tests) before scaling to steady state.
E—Monitoring & adaptation	Attributable effect estimates, with reported uncertainty, delivered against the Phase-A KPIs, and any policy change logged through the update protocol.

4.3. Scalability and the Artefacts It Leaves Behind

Two features keep the methodology from collapsing into a one-size-fits-all prescription. The first is explicit scalability. A *minimum-viable* configuration, for cities with limited data infrastructure, can run on existing monitoring stations, traffic counts, basic fleet data, simple emission factors and descriptive dashboards—enough for evidence-based planning and baseline monitoring. An *advanced* configuration, by contrast, adds high-resolution and real-time sensing, mobile sensors, ANPR and LiDAR or radar, predictive modelling, emissions and dispersion simulation, and automated decision support for near-real-time forecasting and fine-grained exposure assessment. Cities can thus enter the lifecycle at a level matched to their sensor coverage, interoperability and institutional readiness, and progress as these mature [48].

The second feature is that each phase yields an *auditable artefact*: a regulatory package; a digital-twin architecture blueprint; a technology-selection matrix; a deployment plan; and a monitoring-and-evaluation plan. These are not administrative by-products but the very thing that separates a reproducible methodology from a one-off consultancy exercise. The artefacts record objectives, data assumptions, model choices, decisions and their revisions. In doing so they give the system the institutional continuity, transparency and GDPR-compliance on which public trust in an algorithm-assisted instrument ultimately rests [57].

5. Illustrating the Framework: The London ULEZ Against the Lifecycle

This section illustrates the framework through the most intensively documented real LEZ in operation, London's ULEZ. We use London as a *diagnostic example*—a structured way to exercise the framework—not as a validation case study: the mapping shows what the framework would add, it does not measure its effect. It proceeds in two steps. The first is a diagnostic mapping of the existing ULEZ onto the five-phase lifecycle, identifying which phases it implements and which it does not, and showing that the unimplemented phases are precisely those that close the data–model–decision loop. The second is a reproducible proof-of-concept on real London air-quality data that puts the framework's central efficiency claim to a quantitative test. Together these ground the framework in a working system and motivate the empirical agenda of Section 5.6.

5.1. The London ULEZ as a Reference Case

The ULEZ is well suited to this purpose because it is mature, large and openly reported. It began as a central-London zone on 8 April 2019. It expanded to inner London inside the North and South Circular roads on 25 October 2021, bringing some 3.8 million residents inside the boundary. It was then enlarged to cover the whole Greater London area, bordering the M25 in places, on 29 August 2023 [17,23]. Enforcement relies on ANPR cross-referenced to vehicle-registration records. Air quality is tracked through one of the densest urban monitoring systems in the world, the London Air Quality Network and associated low-cost sensor deployments [23,71].

The published evidence is positive but heterogeneous. Evaluations of the 2019 central zone attribute to the ULEZ a roughly 29% reduction in roadside nitrogen dioxide over its first six months [72]. Larger reductions are reported over longer horizons once meteorological and temporal confounders are removed, using state-space evaluation of the ULEZ [25], meteorological-normalisation methods [26] and official monitoring [24]. For the London-wide expansion, official monitoring estimates nitrogen-oxide emissions across all boroughs to be appreciably lower than a no-expansion counterfactual, with vehicle compliance above 96% after the first year [23]. At the same time, causal analysis indicates that further gains will require measures beyond the existing fixed zone, since the instrument's effect is approaching the limit of what constant access rules can deliver [38]. The ULEZ is thus a strong but plateauing static regime—exactly the situation the framework is designed to address.

5.2. Mapping the ULEZ onto the DT-LEZ Lifecycle

Read against the lifecycle (Figure 3), the ULEZ implements the loop's goal- and rule-setting phases fully, implements the operate-and-adapt phases only partially, and does not implement the digital-twin layer that would connect them. Table 6 sets out this reading phase by phase. Phase A is well developed: clear air-quality objectives, a monitored baseline and published key performance indicators. Phase B is likewise mature: defined boundaries, emission-based access rules, equity exemptions and an ANPR enforcement concept [23,32]. Phase C, by contrast, is essentially absent: the ULEZ is not designed around a digital twin used as decision infrastructure, and its sensing and modelling serve evaluation and enforcement rather than a scenario engine. Phase D is correspondingly minimal. The operational system is enforcement IT, not a continuously updated model aligned with the city. Phase E is the revealing case: monitoring and revision do occur, but through periodic reports and discrete boundary expansions decided on administrative and political cycles, with policy intensity held constant between them, rather than through condition-driven adjustment [3,38].

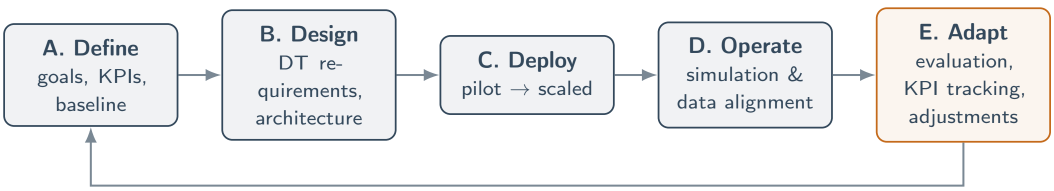

Figure 3. The DT-LEZ lifecycle as a closed loop (Phases A–E). The London ULEZ implements A and B, implements E only partially, and does not implement C or D.

We score the mapping on an explicit 0–3 scale, so that the mapping aims to be a reproducible judgement rather than a matter of impression. Each phase is assessed against its gate in Table 5: **0**, absent (the gate artefact is missing); **1**, minimal (some elements present, gate not met); **2**, implemented but open-loop (gate met, yet outputs feed periodic administrative review rather than the live loop); **3**, implemented as specified (gate met and, for the operate-and-adapt phases, feeding condition-driven, auditable adjustment). On this scale the ULEZ scores 3, 3, 0, 1 and 2 across Phases A–E (Table 6). Two features are designed to limit the subjectivity of the scoring. First, each level is keyed to the presence of a specific, documented artefact from the phase gate (Table 5), so a score records a verifiable check—does the artefact exist, and does it feed the live loop or only periodic review?—rather than a subjective rating. The ULEZ scores are each supported by the public evidence cited in Section 5, namely Transport for London and Greater London Authority reporting, so an independent assessor working from the same sources would reproduce the mapping. Second, the scale is ordinal and diagnostic, not a cardinal performance metric: the reading on which the paper relies—that Phases C and D are unimplemented and Phase E only partially—is robust to a one-level disagreement on any single phase.

Table 6. Maturity of London's ULEZ read against the DT-LEZ lifecycle. Each phase is scored 0–3 against the lifecycle gates of Table 5; the final column states what would be required to close the loop.

Phase	Score & Status	What the ULEZ Already Does	Missing to Close the Loop
A—Objectives & baseline	3—Implemented	Air-quality goals, monitored baseline, published KPIs	—
B—Policy package	3—Implemented	Boundaries, emission-based access rules, equity exemptions, ANPR enforcement	—
C—DT system design	0—Absent	Sensing and modelling for evaluation and enforcement	A Digital twin specified as decision infrastructure (scenario engine, model portfolio, integrated data)
D—Implementation & operations	1—Minimal	Enforcement IT and reporting pipelines	A continuously updated model aligned with live city data
E—Monitoring & adaptation	2—Partial	Periodic evaluation reports; discrete boundary expansions	Condition-driven, trigger-based adjustment within a bounded rule set

The pattern is consistent. The ULEZ is a well-implemented *open* policy cycle—it defines, designs, operates and reviews—but it stops short of the *closed* loop the framework requires, because Phases C and D are not built and Phase E adapts on administrative rather than environmental time. The framework's value, on this reading, is not to replace the ULEZ but to identify the specific, sequenced additions that would turn a strong static regime into an adaptive one: a decision-grade digital twin (C), its operational alignment with live data (D) and condition-driven adjustment (E).

5.3. Contrasting Static and Adaptive Logic

The static baseline is current LEZ logic: fixed boundaries, eligibility, exemptions and enforcement, with policy intensity effectively constant in operation. Consistent with the evidence base, such a regime reduces traffic-related pollutants—especially NO_x—through fleet renewal, deterrence of high-emitting vehicles and compliance, as shown by official ULEZ monitoring [23], London air-quality reporting [24], meteorologically normalised trends [26] and child-health cohort evidence [29]. It cannot, however, easily respond to short-term surges, adverse weather or localised hotspots. The adaptive regime adds a conditional policy term informed by traffic, air-quality, meteorological and compliance data and short-term forecasts, letting responses vary with conditions. Conceptually, this yields three advantages: lower policy latency, since responses are triggered when thresholds are exceeded; better spatial targeting; and greater precision through temporary, proportionate interventions rather than broad restrictions.

Table 7 sets out this contrast across six governance dimensions. It compares the two regimes *mechanism by mechanism* and records only the *direction* in which an adaptive regime would be expected to differ: greater responsiveness, tighter spatial targeting and lower policy latency. The next subsection then puts this efficiency claim on real data with an order-of-magnitude proof-of-concept. A definitive, fully validated figure is reserved for the empirical programme of Section 5.6, which requires the linked dataset described there. The expected direction is consistent with evidence that continuous sensing, forecasting and feedback-based adjustment improve the timing and targeting of interventions: short-term traffic forecasting [73], feedback-based adaptive control [14] and DT-based emission monitoring [12]. Related feedback-based interventions such as large-scale adaptive traffic-signal control have likewise been associated with measurable urban emission reductions [13], though such figures are not directly transferable to LEZ governance.

Table 7. Conceptual, mechanism-by-mechanism comparison of static and adaptive LEZ logic across six governance dimensions. The final column states the *expected direction* of difference only; no magnitudes are claimed or measured.

Dimension	Static LEZ Baseline	Adaptive DT-LEZ	Expected Direction
NO_x reduction	Fixed access rules and fleet compliance	Real-time triggers and targeted response	Higher efficiency
$PM_{2.5}$ response	Indirect, slower, fleet-dependent	Forecast-informed intervention in episodes	Moderate improvement
Traffic efficiency	Fixed restrictions may displace traffic	Dynamic routing, time-window and targeted controls	Improved responsiveness
Policy latency	Slow administrative/political cycles	Near-real-time, trigger-based	Strong improvement
Compliance management	Periodic review, fixed checkpoints	Continuous, risk-based prioritisation	More targeted enforcement
Equity monitoring	Ex post assessment	Continuous exposure/vulnerability tracking	Improved accountability

The comparison's key point is that the principal advantage of the adaptive configuration is not only higher average reduction but improved *responsiveness*. The adaptive DT-LEZ functions as a policy-optimisation layer added to the static regime rather than a replacement for it. It may flag a forecast NO_x exceedance on a congested corridor and recommend temporary enforcement prioritisation, public warnings or routing changes. It may also reveal that a boundary change would reduce exposure in one area while displacing it to another.

5.4. A Proof-of-Concept on Real London Data

To illustrate the efficiency mechanism of Hypothesis 1 on real data—rather than assert it—we run a deliberately simple, fully reproducible experiment on observed hourly NO_2 from five roadside and kerbside stations of the London Air Quality Network (LAQN), including Marylebone Road, Swiss Cottage and Holloway Road, over the two full years 2023–2024 (76,957 valid hourly observations) [71]. The experiment is semi-synthetic, and its limits are stated plainly: the *environment* is real measured data, but the adaptive policy is *modelled*, since a regime that was never deployed cannot be observed. We therefore keep two things separate. This experiment is a demonstration of the efficiency *mechanism*, not an evaluation of policy *effectiveness*. It shows the mechanism is plausible on real data, but it does not measure the impact of a deployed DT-enabled LEZ, which requires the linked, deployed study of Section 5.6.

The set-up is transparent by design. The observed series is the static baseline—the world under the current, fixed ULEZ. Because a zone can only act on locally generated pollution, each series is split into an irreducible background (its 10th percentile) and a controllable, traffic-attributable excess above it. Two interventions are then compared at *equal abatement effort*, defined as the total mass of NO_2 removed ($\mu g\, m^{-3}\, h$). The *uniform* regime trims the controllable excess by a constant fraction every hour, a stricter but still static policy. The *adaptive* regime trims it only when a short-term forecast exceeds an elevated-concentration threshold T, a closed-loop response. Matching effort is the crux: both regimes remove the same total pollution, so any difference is due to *targeting*, not to doing more. For the headline we deliberately use the most conservative, genuinely deployable forecast: *persistence*, the previous observed hour, which uses only past data and so involves no look-ahead. Its correlation with the realised value is $\approx$0.9 at these sites. We take $T = 40\ \mu g\, m^{-3}$, the EU annual NO_2 limit value [74], used as a reference concentration of concern; for comparison, the WHO 2021 24 h guideline is $25\ \mu g\, m^{-3}$ [75]. We also take an intervention effectiveness of 0.30 on the controllable excess. Both, and the forecast quality, are varied below.

With this causal forecast, pooled across the five sites the equal-effort adaptive regime avoids **36.5%** of the elevated-NO_2 hours against **26.9%** for the uniform regime that removes exactly the same total pollution. This is a targeting advantage of $+9.6$ percentage points, a factor of 1.36, whose 95% bootstrap confidence interval, $[8.3, 10.8]$ percentage points, *excludes zero*. Figure 4 illustrates one winter episode at Marylebone Road, where the adaptive trigger shaves the peaks while leaving low-pollution hours untouched. The intervention is active only a minority of hours.

A natural objection is that concentrating abatement on the right tail must beat spreading it evenly—i.e., that the advantage is built in. The controls show otherwise. A *no-skill* forecast, a random trigger with correlation ≈ 0, leaves only a negligible advantage of $+1.1$ percentage points, a factor of 1.08. Without information, targeting buys essentially nothing, and the deployable forecasts' advantage is an order of magnitude larger. The advantage then rises monotonically with forecast quality, up to a factor of 1.64 for an artificial near-perfect forecast with correlation ≈ 0.98. This gradient—the signal being the gap between informative forecasts and the no-skill floor—directly reflects the responsiveness mechanism of Hypothesis 3, since the efficiency gain is realised only with a skilful short-term forecast. It shows the result follows the framework's logic, not an artefact of construction.

We also compare *deployable* predictors against persistence and attach uncertainty with a moving-block bootstrap of contiguous 72 h blocks and 2000 replicates, preserving hourly and daily autocorrelation. Table 8 reports the pooled results. The two informative causal forecasters agree closely. Persistence and a short-term autoregressive model—the latter using lags of 1 and 24 h plus hour-of-day, fitted on the first 60% of each series and applied causally—both reach a correlation of 0.89 and an advantage of about 9 percentage points, a factor near 1.35. For both, the 95% confidence interval excludes zero, so the targeting advantage is statistically significant at this sample size. A weaker seasonal-naive rule yields a small but still positive advantage, consistent with the forecast-quality gradient.

One caveat is reported honestly. At this large sample even the no-skill control shows a tiny, marginally non-zero advantage of ≈1 percentage point, a residual convexity effect of tail-targeting. It lies an order of magnitude below the deployable forecasts and does not affect the conclusion. Section 5.6 sets out how a fully linked dataset would turn this illustration into a deployed evaluation.

Table 8. Forecaster comparison and bootstrap uncertainty for the proof-of-concept, pooled across five London roadside/kerbside LAQN sites over the period 2023–2024 (76,957 valid hours; $T = 40\ \mu g\,m^{-3}$, intervention effectiveness 0.30). "Advantage" is the adaptive minus the equal-effort uniform reduction in elevated-NO_2 hours; 95% confidence intervals are from a moving-block bootstrap (72 h blocks, 2000 replicates). The synthetic and no-skill forecasts are a ceiling and a control, not deployable methods.

Forecaster	**Corr. (f, x)**	**Adaptive Avoided (%)**	**Advantage (pp) [95% CI]**	**Factor**
Persistence (causal)	0.89	36.5	+9.6 [8.3, 10.8]	1.36
Autoregressive, lags 1, 24 (causal)	0.89	35.0	+8.9 [7.5, 10.1]	1.34
Seasonal-naive (causal)	0.44	21.5	+1.9 [0.7, 2.5]	1.10
Synthetic, near-perfect (ceiling)	0.98	46.9	+18.3	1.64
No-skill, random (control)	0.00	14.3	+1.1 [0.2, 1.5]	1.08

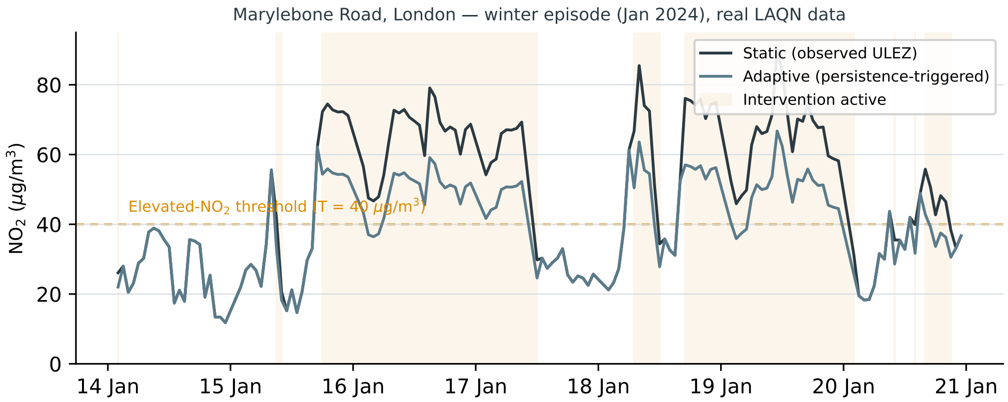

Figure 4. Real hourly NO_2 at Marylebone Road (LAQN) during a winter episode. The static (observed) regime and the effort-matched adaptive regime remove the same total pollution, but the adaptive trigger (shaded hours), driven by a causal persistence forecast, concentrates abatement on the peaks above the reference threshold while leaving low-pollution hours untouched.

The pattern is robust across parameters. Over an intervention-effectiveness range of 0.10–0.50 and a span of synthetic forecast quality, the adaptive advantage is positive in 100% of the grid and reaches roughly 36 percentage points of additional exceedance-hours avoided (Figure 5). It also depends on the threshold of concern: the targeting factor is 1.02, 1.36 and 1.76 at $T = 30$, 40 and 50 $\mu g\,m^{-3}$. Adaptivity thus helps most for the rarer, peakier episodes a static policy cannot reach, and reduces to near-parity only when almost every hour is already "elevated", at $T = 30$. Several simplifications are deliberate: a single pollutant; a percentile background; a reduced-form intervention with no explicit traffic dynamics, dispersion or behavioural response; and a comparison at equal abatement *mass* rather than equal *cost*. These are exactly what the empirical programme of Section 5.6 is designed to replace. The full specification—data, notation, equations and per-site results—is given in Section S6 of the Supplementary Materials, and the analysis code and data are openly available in Zenodo (DOI: 10.5281/zenodo.21157822; Data Availability Statement).

Use of generative AI in this proof-of-concept. The analysis code that produces the comparison reported here and generates Figures 4 and 5, together with its accompanying documentation, was implemented with the assistance of Claude (Anthropic). The study design, the choice of data, the model formulation and all parameter choices were made by the authors, who verified the code against the equations in Section S6 of the Supplementary Materials and reproduced the reported results.

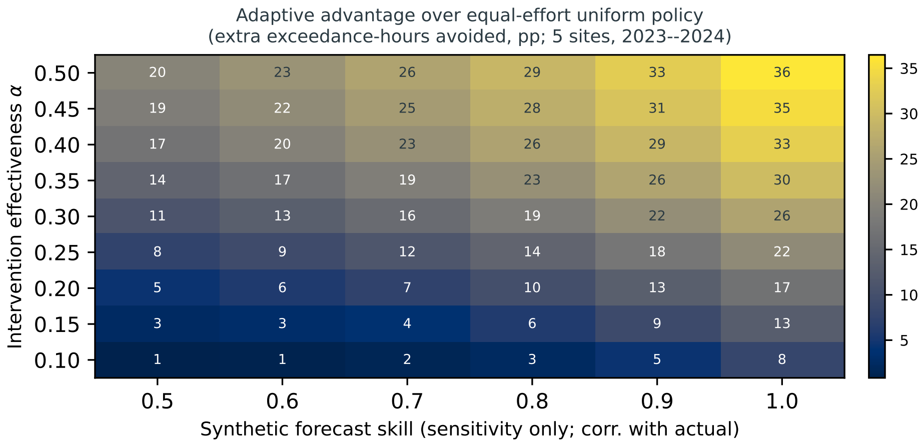

Figure 5. Sensitivity of the adaptive advantage (extra elevated-NO_2 hours avoided over the equal-effort uniform policy, percentage points), pooled across the five roadside/kerbside sites over 2023–2024, over intervention effectiveness and (synthetic) forecast quality. The advantage is positive across the whole plausible range and grows with both.

5.5. The Hypotheses: Mechanism, Evidence and Test

We treat the three hypotheses as a falsifiable research agenda, not as confirmed results. For each we state the causal mechanism the framework implies, the evidence that currently makes it plausible, and a concrete design through which it could be tested or refuted—so that each functions as a sharp, checkable claim rather than an assertion, and together they set the empirical programme that the framework invites.

On efficiency (Hypothesis 1), because a static LEZ holds policy intensity constant it cannot reallocate enforcement or access to the times and places where the marginal contribution to exposure is highest, so abatement per unit of restriction is lower. The closed loop, by contrast, shortens the detect-to-respond cycle and concentrates interventions where they matter. This is plausible given the street-level heterogeneity of $PM_{2.5}$ and the episodic nature of exceedances [3,35], together with evidence that feedback-based control improves outcomes at scale in related mobility domains [13,14]. The proof-of-concept of Section 5.4 implements exactly this comparison in reduced form—a fixed and a conditional regime on identical real concentration traces at equal abatement effort. It finds the conditional regime avoids significantly more elevated-NO_2 hours, with its 95% confidence interval excluding parity, consistent with the hypothesis. A full test would extend it to linked traffic, weather and compliance data in a deployed setting.

On data fragmentation (Hypothesis 2), a policy that lacks integrated mobility, air-quality and enforcement data cannot detect or act on the conditions to which it should respond, which caps achievable performance. Interoperability and data-governance research consistently identifies fragmentation as a binding constraint on urban data systems [18,19]; a test would measure decision latency and spatial coverage before and after data integration for the same city.

On responsiveness (Hypothesis 3), conditional triggers can act within the timescale of surges, episodes and adverse weather, which administrative review cannot, and short-term traffic and air-quality forecasting is now mature enough to drive them [49,73]. This could be tested through an event study of response time and population exposure during traffic surges and pollution episodes.

Whether these expected directions hold at the scale of a deployed system—beyond the reduced-form evidence of Section 5.4—is the empirical question the programme below is designed to settle.

5.6. Validation Boundaries and Empirical Roadmap

The proof-of-concept of Section 5.4 tests the efficiency mechanism on real hourly concentration data, but it remains semi-synthetic: the adaptive policy is modelled, the intervention is reduced-form, and no linked dataset of traffic, weather, compliance and enforcement is used. It demonstrates the mechanism; it does not measure the deployed effect. Full empirical validation would require four steps: (i) linking hourly NO_x and $PM_{2.5}$ with traffic, weather, fleet and compliance data; (ii) estimating a baseline statistical model of the existing LEZ controlling for seasonality, weather and background trends; (iii) simulating adaptive-policy counterfactuals under equivalent conditions; and (iv) sensitivity analysis for uncertainty, robustness and displacement. These steps define the next stage of research and the basis for a future empirical case study.

6. Discussion

The paper's central contribution is a reframing. By recasting the LEZ as a unit of *adaptive governance* rather than a fixed boundary, and the digital twin as decision infrastructure rather than a map, it shifts the guiding question from "did the LEZ work?" to "how should the LEZ respond?". In doing so it connects three bodies of literature that have largely developed apart: air-quality policy, urban digital twins and smart-city governance (Table 1). The practical counterpart of that reframing is a design-to-operation pathway. It comprises a phased methodology whose artefacts make decisions auditable, a maturity-scaled route that lets cities start small, and governance requirements—GDPR-by-design, auditability and equity monitoring—that make adaptive enforcement defensible rather than merely feasible.

Its main strength is transferability. Because the framework prescribes a process and a set of decision points rather than a single technology, it adapts to local context. The minimum-viable/advanced distinction lets a city begin with the data it already has before moving toward real-time sensing and adaptive triggers [54,55]. The reframing is also analytically useful. Locating the value of a DT-enabled LEZ in reduced policy latency and sharper targeting, rather than in richer visualisation, gives a concrete criterion against which such systems can be judged [11]. Generalisation across cities is handled by these maturity tiers. A low-data city can enter at the minimum-viable configuration—existing stations, traffic counts and descriptive dashboards—and ascend to real-time sensing and adaptive triggers as capacity grows (Section 4). One lifecycle thus serves very different starting points. The efficiency-and-targeting direction is consistent with evidence that feedback-based sensing and control improve the timing of interventions in related mobility domains [8,14] and with proposals for dynamic LEZs [3]. What the framework adds, unlike adaptive-control studies that optimise a single control layer, is an explicit route from policy objectives to an auditable operating system. Two comparisons make this concrete. Causal evaluations of London's ULEZ report real but plateauing roadside-NO_x reductions [38]. Our proof-of-concept points the same way and speaks to that plateau, by making policy intensity condition-dependent rather than uniform. Adaptive traffic-control studies, in turn, report emission gains from feedback control [13,14]. There, the contribution is a new controller; here, it is the governance layer that turns such control into an accountable policy instrument.

These claims are bounded by real limitations. Data availability and quality remain the binding constraint in cities with sparse monitoring or fragmented ownership, and interoperability failures can make a DT reproduce the very silos it was meant to dissolve [18,19]. This is the precondition Hypothesis 2 identifies. Model uncertainty and algorithmic bias are equally real: models trained on incomplete or historically skewed data can distort enforcement, exposure assessment and exemption design. Social legitimacy is decisive too. LEZs provoke resistance when they are seen as punitive, opaque or unfair, especially among low-income households and small businesses [30,34]. Adaptive operation can improve fairness through continuous equity monitoring, but it can as easily intensify surveillance. This is why GDPR-by-design, audit trails, explainable rules and public communication are treated here as core requirements rather than implementation details [58,61].

Moving from design to deployment also meets barriers that are organisational and political as much as technical. Institutionally, responsibility for air quality, transport, enforcement and data is usually split across agencies, so an adaptive LEZ needs a mandate and data-sharing arrangements that often do not yet exist [18,22]. Legally, condition-dependent enforcement must preserve the citizen's ability to know which rule applies, where and when, which limits how fast and how opaquely rules may change. Organisationally, cities need data-engineering and modelling capacity that many do not have. Financially, sensing, integration and operation compete with other municipal priorities, which makes procurement and a modest entry point decisive. And because the instrument is visible and redistributive, public acceptability—shaped by perceived fairness and transparency—can decide whether a scheme survives politically [30,34]. These are not arguments against adaptivity but the conditions under which it becomes feasible, and the lifecycle's gates and artefacts are meant to surface them early.

This invites the obvious objection: is adaptivity always desirable? It is not. Frequent rule changes can erode the legibility and predictability that make a regulation fair and enforceable, and dynamic enforcement raises a genuine legal question, since a citizen must be able to know which rule applies, where and when. The framework therefore treats adaptivity as *bounded*: thresholds, triggers and exemptions are fixed in advance, changes are logged and reversible, and the DT supports rather than replaces the responsible authority. Adaptivity is a means to better-targeted, more legitimate intervention, not an end in itself; where stability serves equity and compliance better, the same loop justifies leaving the rule untouched.

For adoption, then, sequence matters more than sophistication. A city is better served by clear objectives, a modest pilot, transparent governance and realistic data than by an ambitious platform. Investment in data stewardship, legal oversight, interoperability and stakeholder engagement deserves to be taken as seriously as investment in sensors or software. Implemented this way—incrementally, with monitoring, auditability and adaptive management built in from the start—a DT-enabled LEZ lets technical capability strengthen environmental performance, public trust and equitable governance rather than trade them against one another [62,76].

Read against the journal's remit, these mechanisms are what make the framework a sustainability contribution and not only a technical one. Better-timed, better-targeted abatement lowers population exposure to NO_x and $PM_{2.5}$ and the associated health burden (SDG 3). Continuous distributional monitoring and time-limited, means-sensitive exemptions let cities pursue cleaner air without deepening mobility inequities (SDG 11). Steering activity away from the highest-emitting vehicles and hours supports lower-carbon urban mobility (SDG 13). Framing the LEZ as adaptive decision infrastructure thus links a concrete governance instrument to measurable sustainability outcomes—exposure, equity and emissions—and to the environmental limits that motivate air-quality regulation in the first place [1].

7. Conclusions

This paper set out to close a specific gap: not whether Low-Emission Zones work, but why their fixed design limits them, and how they might instead be governed adaptively. Its central contribution is to reframe the LEZ as an adaptive decision infrastructure organised around a continuous data–model–decision loop, and to translate that idea into a scalable, auditable, five-phase lifecycle methodology. By drawing the distinction between fixed and conditional policy logic, distinguishing minimum-viable from advanced configurations, and reading London's ULEZ against the lifecycle to locate the phases that would close the loop, the paper offers a transferable and testable basis for adaptive environmental governance.

The three hypotheses the paper advances make that testability concrete: that adaptive operation reduces pollutants more efficiently than a static regime; that overcoming the fragmentation of urban data is a precondition for that gain; and that real-time, rule-based adjustment improves responsiveness to short-term shocks. In practice these would surface as lower policy latency, better-targeted interventions, sharper compliance monitoring and more transparent evaluation. They remain open hypotheses—the first with initial, reduced-form support—rather than established results, and define the empirical agenda the framework invites.

The paper's contribution at this stage is primarily conceptual: it advances a framework, and tests one of its claims only in reduced form. The London ULEZ gap analysis structures the static-versus-adaptive comparison, and the proof-of-concept on real LAQN data gives the efficiency claim reproducible and statistically significant support across deployable causal forecasters, with the 95% CI excluding parity. This holds on real concentration traces with a modelled policy rather than a deployed system, so the three hypotheses stand as the framework's central, still-open claims. The approach also assumes data, interoperability and institutional capacity that vary across cities, which the minimum-viable configuration is designed to accommodate. The governance and equity safeguards, though built in from the start, will be fully tested only in deployment. In short, the framework is established conceptually and illustrated in reduced form; any claim about its effectiveness in operation must await the extensive empirical validation set out next. Stated precisely, as they are here, these claims set the agenda for the work that follows.

That work has a clear shape. The first priority is empirical validation along the four steps set out in Section 5.6: assembling a linked dataset of hourly air quality, traffic, weather, fleet and compliance; estimating a baseline model of an existing LEZ; simulating adaptive-policy counterfactuals under equivalent conditions; and testing robustness, uncertainty and displacement. Beyond a single city, the route runs through standardised KPIs for DT-enabled LEZs, cross-city benchmarks, trials of bounded adaptive-policy rules in real settings under appropriate legal safeguards, and direct evaluation of distributional impacts and public acceptability. Pursued together, these steps are how a DT-enabled LEZ can move from a conceptual innovation to a robust, accountable and scalable instrument for urban air-quality management. It thereby becomes a concrete lever for the wider urban sustainability transition—healthier populations, lower-carbon mobility and cities that remain within their environmental limits—in line with the Sustainable Development Goals on health, sustainable cities and climate action (SDGs 3, 11 and 13) [1].

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/su18147100/s1, Section S1: Computer-vision components for DT-enabled LEZs; Section S2: Traffic-modelling options; Section S3: Emissions and dispersion modelling; Section S4: GDPR-by-design checklist for DT-LEZ; Section S5: Procurement requirements; Section S6: Proof-of-concept: data, model and computation; Section S7: The policy engine as a constrained optimization problem.

Author Contributions: Conceptualization, A.C.-A. and J.C.R.; methodology, A.C.-A.; investigation, A.C.-A.; writing—original draft preparation, A.C.-A.; visualization, A.C.-A.; writing—review and editing, J.C.R.; supervision, J.C.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The hourly NO_2 observations used in the proof-of-concept (Section 5.4) are openly available from the London Air Quality Network (LAQN, https://www.londonair.org.uk, accessed on 1 June 2026). The processed data extracts and the complete analysis and figure-generation code that reproduce all reported results are openly available in Zenodo at https://doi.org/10.5281/zenodo.21157822 (accessed on 1 June 2026). No other new data were created in this study.

Acknowledgments: During the preparation of this manuscript, the authors used Claude (Anthropic; Opus 4.x and Sonnet 4.x model families, 2025–2026) for language editing of the text and for assistance in implementing the proof-of-concept analysis code and generating its figures (Section 5.4; full specification in Section S6 of the Supplementary Materials). The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: Author A.C. (Antonio Cantalapiedra-Asensio) was employed by Georeplica Technologies. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest.

References

1. United Nations. *Transforming Our World: The 2030 Agenda for Sustainable Development*; Technical Report A/RES/70/1; United Nations General Assembly: New York, NY, USA, 2015.
2. Kaur, R.; Pandey, P. Air pollution, climate change, and human health in Indian cities: A brief review. *Front. Sustain. Cities* **2021**, *3*, 705131.
3. Manglano-Redondo, P.; Paricio-Garcia, A.; Lopez-Carmona, M.A. Dynamic Low-Emission Zones for Urban Mobility: A Systematic Review. *Appl. Sci.* **2025**, *15*, 2915.
4. Deng, T.; Zhang, K.; Shen, Z.J.M. A systematic review of a digital twin city: A new pattern of urban governance toward smart cities. *J. Manag. Sci. Eng.* **2021**, *6*, 125–134. https://doi.org/10.1016/j.jmse.2021.03.003.
5. Li, D.; Yang, W.; Su, Z. Smart city based on digital twins. *Comput. Urban Sci.* **2021**, *1*, 4. https://doi.org/10.1007/s43762-021-00005-y.
6. Ge, C.; Qin, S. Digital twin intelligent transportation system (DT-ITS)—A systematic review. *IET Intell. Transp. Syst.* **2024**, *18*, 2325–2358. https://doi.org/10.1049/itr2.12539.
7. White, G.; Zink, A.; Codecá, L.; Clarke, S. A digital twin smart city for citizen feedback. *Cities* **2021**, *110*, 103064.
8. Wang, T.; Cao, J.; Hussain, A. Adaptive traffic signal control for large-scale scenario with cooperative group-based multi-agent reinforcement learning. *Transp. Res. Part C Emerg. Technol.* **2021**, *125*, 103046.
9. Komti, A.; Jbara, A.; Kechiche, N.; Khardi, S.; Slimi, K. Study of Airborne Particulate Pollution from Road Traffic Emissions in the Coastal Urban Center of Sousse City, Tunisia. *Int. J. Environ. Res.* **2025**, *19*, 1–25.
10. El-Agamy, R.F.; Sayed, H.A.; AL Akhatatneh, A.M.; Aljohani, M.; Elhosseini, M. Comprehensive analysis of digital twins in smart cities: A 4200-paper bibliometric study. *Artif. Intell. Rev.* **2024**, *57*, 154. https://doi.org/10.1007/s10462-024-10781-8.
11. Adreani, L.; Bellini, P.; Colombo, C.; Fanfani, M.; Nesi, P.; Pantaleo, G.; Pisanu, R. Implementing integrated digital twin modelling and representation into the Snap4City platform for smart city solutions. *Multimed. Tools Appl.* **2024**, *83*, 37121–37146.
12. Zhivkov, P.; Fidanova, S.; Dimov, I. Digital twin-based framework for real-time monitoring and analysis of urban mobile-source emissions. *Atmosphere* **2025**, *16*, 731.

13. Wu, K.; Ding, J.; Lin, J.; Zheng, G.; Sun, Y.; Fang, J.; Xu, T.; Zhu, Y.; Gu, B. Big-data empowered traffic signal control could reduce urban carbon emission. *Nat. Commun.* **2025**, *16*, 2013. https://doi.org/10.1038/s41467-025-56701-4.
14. Benhamza, K.; Seridi, H.; Agguini, M.; Bentagine, A. A multi-agent reinforcement learning based approach for intelligent traffic signal control. *Evol. Syst.* **2024**, *15*, 2383–2397.
15. Weil, C.; Bibri, S.E.; Longchamp, R.; Golay, F.; Alahi, A. Urban digital twin challenges: A systematic review and perspectives for sustainable smart cities. *Sustain. Cities Soc.* **2023**, *99*, 104862.
16. Broster, S.; Terzano, K. A systematic review of the pollution and health impacts of low emission zones. *Case Stud. Transp. Policy* **2025**, *19*, 101340.
17. Delgado-Lindeman, M.; Cordera, R.; Moura, J.L.; Rodriguez, A. Characteristics and effects of low emission zones in Europe. A systematic literature review. *Eur. Transp. Res. Rev.* **2025**, *17*, 54. https://doi.org/10.1186/s12544-025-00749-2.
18. Quek, H.Y.; Sielker, F.; Akroyd, J.; Bhave, A.N.; von Richthofen, A.; Herthogs, P.; van der Laag Yamu, C.; Wan, L.; Nochta, T.; Burgess, G.; et al. The conundrum in smart city governance: Interoperability and compatibility in an ever-growing ecosystem of digital twins. *Data Policy* **2023**, *5*, e6.
19. Paskaleva, K.; Evans, J.; Martin, C.; Linjordet, T.; Yang, D.; Karvonen, A. Data governance in the sustainable smart city. *Informatics* **2017**, *4*, 41.
20. Sayed, S.A.; Abdel-Hamid, Y.; Hefny, H.A. Artificial intelligence-based traffic flow prediction: A comprehensive review. *J. Electr. Syst. Inf. Technol.* **2023**, *10*, 13.
21. Jiang, W.; Luo, J. Graph neural network for traffic forecasting: A survey. *Expert Syst. Appl.* **2022**, *207*, 117921.
22. Kok, I. *Planning and Implementation of Low- and Zero-Emission Zones in Cities*; Technical Report; International Council on Clean Transportation: Washington, DC, USA, 2023.
23. Transport for London. *London-Wide ULEZ: One Year Report*; Technical Report; Transport for London: London, UK, 2025.
24. Greater London Authority. *Air Quality in London 2016–2024*; Technical Report; Greater London Authority: London, UK, 2024.
25. Hajmohammadi, H.; Heydecker, B. Evaluation of air quality effects of the London Ultra-Low Emission Zone by state-space modelling. *Atmos. Pollut. Res.* **2022**, *13*, 101514.
26. Grange, S.K.; Carslaw, D.C. Using meteorological normalisation to detect interventions in air quality time series. *Sci. Total Environ.* **2019**, *653*, 578–588.
27. Chamberlain, R.C.; Fecht, D.; Davies, B.; Laverty, A.A. Health effects of low emission and congestion charging zones: A systematic review. *Lancet Public Health* **2023**, *8*, e559–e574.
28. Sarmiento, L.; Wägner, N.; Zaklan, A. The air quality and well-being effects of low emission zones. *J. Public Econ.* **2023**, *227*, 105014.
29. Tsocheva, I.; Scales, J.; Dove, R.; Chavda, J.; Kalsi, H.; Wood, H.E.; Colligan, G.; Cross, L.; Newby, C.; Hall, A.; et al. Investigating the impact of London's ultra low emission zone on children's health: Children's Health in London and Luton (CHILL) protocol for a prospective parallel cohort study. *BMC Pediatr.* **2023**, *23*, 556.
30. Wappelhorst, S.; Bui, A.; Morrison, K. *Low- and Zero-Emission Zones and Social Equity: Supporting the Urban Transition to Zero-Emission Vehicles and Alternative Transport Modes*; Technical Report; International Council on Clean Transportation: Washington, DC, USA, 2023.
31. Moreno, E.; Schwarz, L.; Host, S.; Chanel, O.; Benmarhnia, T. The environmental justice implications of the Paris low emission zone: A health and economic impact assessment. *Air Qual. Atmos. Health* **2022**, *15*, 2171–2184.
32. Liotta, C. What drives inequalities in Low Emission Zones' impacts on job accessibility? *Transp. Policy* **2025**, *160*, 29–41. https://doi.org/10.1016/j.tranpol.2024.10.029.
33. Wu, X.; Zhong, C.; Wang, Y. The impact of the ultra-low emission zone on high streets economy and social equality in Outer London. *Transp. Res. Part A Policy Pract.* **2025**, *200*, 104612. https://doi.org/10.1016/j.tra.2025.104612.
34. Morton, C.; Mattioli, G.; Anable, J. Public acceptability towards Low Emission Zones: The role of attitudes, norms, emotions, and trust. *Transp. Res. Part A Policy Pract.* **2021**, *150*, 256–270.
35. Wang, Y.; Dai, X.; Gong, D.; Zhou, L.; Zhang, H.; Ma, W. Correlations between urban morphological indicators and PM2.5 pollution at street-level: Implications on urban spatial optimization. *Atmosphere* **2024**, *15*, 341.
36. Degraeuwe, B.; Hooyberghs, H.; Janssen, S.; Lefebvre, W.; Maiheu, B.; Megaritis, A.; Vanhulsel, M. A source apportionment and air quality planning methodology for NO2 pollution from traffic and other sources. *Environ. Model. Softw.* **2024**, *176*, 106032. https://doi.org/10.1016/j.envsoft.2024.106032.
37. Tassinari, F. Low emission zones and traffic congestion: Evidence from Madrid Central. *Transp. Res. Part A Policy Pract.* **2024**, *185*, 104099. https://doi.org/10.1016/j.tra.2024.104099.
38. Tong, C.; Dai, Y.; Cole, M.; Elliott, R.J.R.; Bartington, S.E.; Liu, B.; Shi, Z. Further improvement in London's air quality demands more than the Ultra Low Emission Zone policy. *npj Clean Air* **2025**, *1*, 29.
39. World Economic Forum; China Academy of Information and Communications Technology. *Digital Twin Cities: Framework and Global Practices*; Technical Report; World Economic Forum: Geneva, Switzerland, 2022.

40. Mazzetto, S. A review of urban digital twins integration, challenges, and future directions in smart city development. *Sustainability* **2024**, *16*, 8337.
41. Huzzat, A.; Anpalagan, A.; Khwaja, A.S.; Woungang, I.; Alnoman, A.A.; Pillai, A.S. A comprehensive review of digital twin technologies in smart cities. *Digit. Eng.* **2025**, *4*, 100040.
42. Nochta, A.; Wan, L.; Schooling, J.M.; Parlikad, A.K. A socio-technical perspective on urban analytics: The case of city-scale digital twins. *J. Urban Technol.* **2021**, *28*, 263–287.
43. Argota Sánchez-Vaquerizo, J. Urban Digital Twins and metaverses towards city multiplicities: uniting or dividing urban experiences? *Ethics Inf. Technol.* **2025**, *27*, 1–31.
44. Bellini, P.; Nesi, P.; Pantaleo, G. IoT-enabled smart cities: A review of concepts, frameworks and key technologies. *Appl. Sci.* **2022**, *12*, 1607.
45. Whaiduzzaman, M.; Barros, A.; Chanda, M.; Barman, S.; Sultana, T.; Rahman, M.S.; Roy, S.; Fidge, C. A review of emerging technologies for IoT-based smart cities. *Sensors* **2022**, *22*, 9271.
46. Kaššaj, H.; Peráček, T. Synergies and potential of Industry 4.0 and automated vehicles in smart city infrastructure. *Appl. Sci.* **2024**, *14*, 3575.
47. Mutambik, I. IoT-Enabled Adaptive Traffic Management: A Multiagent Framework for Urban Mobility Optimisation. *Sensors* **2025**, *25*, 4126.
48. Cowell, N.H.; Chapman, L.; Topping, D.; James, P.J.P.; Bell, D.B.; Bannan, T.; Murabito, E.; Evans, J.; Birkin, M. Moving from monitoring to real-time interventions for air quality: are low-cost sensor networks ready to support urban digital twins? *Front. Sustain. Cities* **2024**, *6*, 1500516.
49. Medina-Salgado, B.; Sánchez-DelaCruz, E.; Pozos-Parra, P.; Sierra, J.E. Urban traffic flow prediction techniques: A review. *Sustain. Comput. Inform. Syst.* **2022**, *35*, 100739.
50. Chawla, A.; Khare, M.; Perugu, H. Improving exhaust emission evaluation: An integrated modelling approach for urban road networks in diverse operating environments. *Atmos. Pollut. Res.* **2024**, *15*, 102207.
51. Dong, Y.; Xu, J.; Ni, J. Carbon emission model of vehicles driving at fluctuating speed on highway. *Environ. Sci. Pollut. Res.* **2023**, *30*, 18064–18077.
52. Liang, Q.; Miao, Y.; Zhang, G.; Liu, S. Simulating microscale urban airflow and pollutant distributions based on computational fluid dynamics model: A review. *Toxics* **2023**, *11*, 927.
53. Gastineau, P.; Can, A.; Yaméogo, B.F.; Luquezi, L.G.; Hankach, P.; Vandanjon, P.O.; Bescond, V.L. Modeling exposure to mobility-related pollution: Review and key challenges. *Transp. Res. Rec.* **2024**, *2678*, 587–603.
54. Caldarelli, G.; Arcaute, E.; Barthelemy, M.; Batty, M.; Gershenson, C.; Helbing, D.; Mancuso, S.; Moreno, Y.; Ramasco, J.J.; Rozenblat, C.; et al. The role of complexity for digital twins of cities. *Nat. Comput. Sci.* **2023**, *3*, 374–381.
55. Ahn, C.; Lotfi-Jam, F.; Graham, C.; Bunnell, T.; Marvin, S. Critical urban informatics for urban digital twin models. *Nat. Cities* **2025**, *2*, 114–116.
56. Dewitte, P.; Ausloos, J. Chronicling GDPR transparency rights in practice: The good, the bad and the challenges ahead. *Int. Data Priv. Law* **2024**, *14*, 106–133.
57. Ausloos, J.; Dewitte, P. Shattering one-way mirrors—Data subject access rights in practice. *Int. Data Priv. Law* **2018**, *8*, 4–28.
58. Janssen, M.; Brous, P.; Estevez, E.; Barbosa, L.S.; Janowski, T. Data governance: Organizing data for trustworthy Artificial Intelligence. *Gov. Inf. Q.* **2020**, *37*, 101493.
59. Vega-Gonzalo, M.; Christidis, P. Fair models for impartial policies: Controlling algorithmic bias in transport behavioural modelling. *Sustainability* **2022**, *14*, 8416.
60. Mutambik, I.; Almuqrin, A.; Lee, J.; Zhang, J.Z.; Alomran, A.; Omar, T.; Floos, A.; Homadi, A. Usability of the G7 Open Government Data Portals and lessons learned. *Sustainability* **2021**, *13*, 13740.
61. Kitchin, R.; Coletta, C.; Evans, L.; Heaphy, L.; Mac Donncha, D. *Smart Cities, Urban Technocrats, Epistemic Communities and Advocacy Coalitions*; Working Paper 26; The Programmable City, Maynooth University: Maynooth, Ireland, 2017.
62. Gil-Garcia, J.R.; Pardo, T.A.; Nam, T. What makes a city smart? Identifying core components and proposing an integrative and comprehensive conceptualization. *Inf. Polity* **2015**, *20*, 61–87.
63. Ketzler, B.; Naserentin, V.; Latino, F.; Zangelidis, C.; Thuvander, L.; Logg, A. Digital twins for cities: A state of the art review. *Built Environ.* **2020**, *46*, 547–573.
64. Dembski, M.; Wössner, U.; Letzgus, M.; Ruddat, M.; Yamu, C. Urban digital twins for smart cities and citizens: The case study of Herrenberg, Germany. *Sustainability* **2020**, *12*, 2307.
65. Xue, S.; Zou, H.; Feng, Q.; Wang, X.; Liu, Y.; Wang, Y.; Liu, L. A method for delineating traffic low emission control zone based on deep learning and multi-objective optimization. *Environ. Monit. Assess.* **2025**, *197*, 523. https://doi.org/10.1007/s10661-025-13949-z.
66. Otu, E.; Ashworth, K.; Tsekleves, E.; Ackley, A. Assessing air pollution health impacts in Lancaster City centre: A hybrid modelling study of urban planning mitigation scenarios using the APEHI framework. *Air Qual. Atmos. Health* **2025**, *18*, 3751–3779.

67. Collao, J.; Ma, H.; Lozano-Galant, J.A.; Turmo, J. Traffic road emission estimation through visual programming algorithms and building information models: A case study. *IEEE Access* **2021**, *9*, 150846–150864.
68. Afrazi, M.; Jahed Armaghani, D.; Afrazi, H.; Fattahi, H.; Samui, P. Real-time monitoring of tunnel structures using digital twin and artificial intelligence: A short overview. *Deep Undergr. Sci. Eng.* **2025**, *5*, 315–330.
69. Omrany, H.; Al-Obaidi, K.M. Application of digital twin technology for Urban Heat Island mitigation: Review and conceptual framework. *Smart Sustain. Built Environ.* **2025**, *14*, 2145–2175.
70. Evangeline, S.I. Ethical, Privacy, and Security Implications of Digital Twins. In *AI-Powered Digital Twins for Predictive Healthcare: Creating Virtual Replicas of Humans*; IGI Global Scientific Publishing: Hershey, PA, USA, 2025; pp. 397–424.
71. Environmental Research Group, Imperial College London. London Air Quality Network (LAQN). 2024. Available online: https://www.londonair.org.uk (accessed on 22 June 2026).
72. Transport for London. *Central London Ultra Low Emission Zone: Six Month Report*; Technical Report; Transport for London: London, UK, 2019.
73. Liu, R.; Shin, S.Y. A review of traffic flow prediction methods in intelligent transportation system construction. *Appl. Sci.* **2025**, *15*, 3866.
74. European Parliament and Council of the European Union. Directive (EU) 2024/2881 of the European Parliament and of the Council of 23 October 2024 on Ambient Air Quality and Cleaner Air for Europe (Recast). 2024. Available online: https://eur-lex.europa.eu/eli/dir/2024/2881/oj/eng (accessed on 1 June 2026).
75. World Health Organization. *WHO Global Air Quality Guidelines: Particulate Matter ($PM_{2.5}$ and PM_{10}), Ozone, Nitrogen Dioxide, Sulfur Dioxide and Carbon Monoxide*; Technical Report; World Health Organization: Geneva, Switzerland, 2021.
76. Kakderi, C.; Oikonomaki, E.; Papadaki, I. Smart and resilient urban futures for sustainability in the post COVID-19 era: A review of policy responses on urban mobility. *Sustainability* **2021**, *13*, 6486.

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.