HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

RESPONSIBLE AI IN URBAN DESIGN: A FRAMEWORK FOR INCLUSIVE AND INTEGRATED PLANNING

Diyor Ergashev
Project Manager Institute for Macroeconomic and Regional Studies
E-mail: d.ergashev@imrs.uz

F-Doi: https://doi.org/10.5281/zenodo.23737074

Abstract: Urban artificial intelligence increasingly influences which neighbourhoods receive investment, whose needs enter design briefs and how competing spatial claims are evaluated. Yet technical performance offers an incomplete basis for public decisions. This article develops a governance framework that connects AI-assisted urban design with distributive assessment, meaningful participation and accountable implementation. A targeted narrative review of eleven research and institutional sources published within the period from 27 October 2024 to 27 October 2026 is combined with a transparent numerical illustration.

The synthesis distinguishes evidence of technical capability from evidence of public benefit and identifies four recurring problems: uneven representation, uncertain transferability, fragmented responsibility and evaluation through narrow performance indicators. The proposed framework links problem definition, data review, scenario generation, public deliberation, authorisation and monitoring through a documented decision record.

An illustrative allocation exercise shows why aggregate accessibility gains can conceal small benefits for disadvantaged residents: imposing a minimum benefit requirement changes the preferred alternative even when the aggregate score falls. The article specifies indicators, institutional responsibilities and an evaluation design for testing this proposition in municipal practice.

Its contribution is a procedure for making normative choices visible and contestable, it does not establish that the framework has already improved outcomes. Responsible adoption requires evidence that affected groups benefit, that officials can explain and revise decisions, and that the cost of oversight is included in implementation budgets.

Keywords: responsible artificial intelligence, urban design, integrated planning, spatial justice, participatory governance, algorithmic accountability

Introduction

An urban design decision distributes opportunities over space. The location of a school, a shaded walking route or a bus interchange changes who can reach essential services and at what cost. When an AI system helps rank sites or generate redevelopment options, it enters this distributive process. Its recommendations depend on a definition of the problem, a selection of data and a judgement about acceptable trade-offs. These decisions remain political and professional responsibilities even when software performs the calculations.

Recent public-sector evidence explains the attraction of AI but also the limits of an efficiency-led agenda. The OECD's 2025 review examined 200 government use cases, 57% supported the

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

automation, streamlining or tailoring of processes and services. This describes the orientation of the reviewed applications, rather than the prevalence of AI across all governments. For planning authorities, faster processing is useful only insofar as it supports defensible decisions. A shorter assessment period cannot compensate for an inaccessible consultation process or an investment programme that systematically overlooks poorly documented neighbourhoods.

The recent research agenda has become more explicit about institutional power. Kerimi examines how agentic AI may acquire influence over urban decisions without adequate accountability. Mashhadi Moghaddam and Cao analyse the divergence between technical indicators and social outcomes. Xu, Seto and Weng connect algorithmic accountability with the security of urban infrastructure. These contributions support a practical question: what evidence should a municipality require before allowing an AI output to influence a spatial decision?

This article addresses that question through three related objectives. It identifies the main governance implications of recent evidence, develops a procedure linking ethical commitments to planning decisions, and demonstrates how a distributive safeguard can change the ranking of project alternatives. Inclusion is defined here as both an opportunity to influence decisions and a defensible distribution of their benefits and burdens. Integrated planning refers to coordination across land use, mobility, housing, environmental management and public finance, including the connections between neighbourhood and metropolitan scales.

The central proposition is that a responsible planning system must preserve a traceable relationship between a public objective, the evidence used to pursue it and the authority that accepts the resulting trade-off. This relationship should survive model replacement, staff turnover and changes in political leadership. A technically sophisticated system without that record leaves residents unable to distinguish an administrative judgement from a software recommendation.

Research design and evidence boundaries

The study combines a targeted narrative review with conceptual framework development and a worked numerical example. Searches were undertaken on 27 October 2026 using combinations of “responsible AI”, “urban planning”, “urban design”, “participation”, “algorithmic accountability” and related terms. Publisher pages, arXiv records, OECD publications, ITU statistics, UK government guidance and IEA analysis were consulted. The publication window was the preceding twenty-four months. Searches also used 2025 and 2026 date terms to identify recent developments.

Eleven sources were retained for a defined analytical purpose: two institutional reports on public-sector and municipal AI, three journal contributions on governance and social outcomes, two preprints on participatory and reasoning-based planning, two statistical releases on connectivity, one government transparency instrument and one energy assessment. Publication dates and data reference periods were treated separately. In particular, a 2026 paper examining deployments from 2015-2024 provides recent analysis of earlier experience, not evidence of deployments occurring in 2026.

Inclusion required identifiable authorship or institutional responsibility, a verifiable publication record and a direct contribution to at least one of five categories: representation, participation, decision authority, outcome evaluation or infrastructure impacts.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

Promotional material and unsupported claims of municipal performance were excluded from the synthesis. Preprints were retained where they described emerging approaches, but their status is identified and their claims are not treated as established policy effectiveness. Full-text material was used where accessible the WeDesign and reasoning-framework discussions rely on their accessible abstracts and publication records. This uneven depth of access limits comparison.

The review is purposive and non-exhaustive. It does not provide a PRISMA-style systematic review, a bibliometric census or a meta-analysis. No claim is made that the selected sources represent all cities, all languages or all relevant publications. The framework was derived by translating recurrent concerns into requirements for an identifiable planning decision: what must be recorded, who must review it and what would justify revision. These requirements are proposals developed in this article rather than standards validated by the reviewed studies.

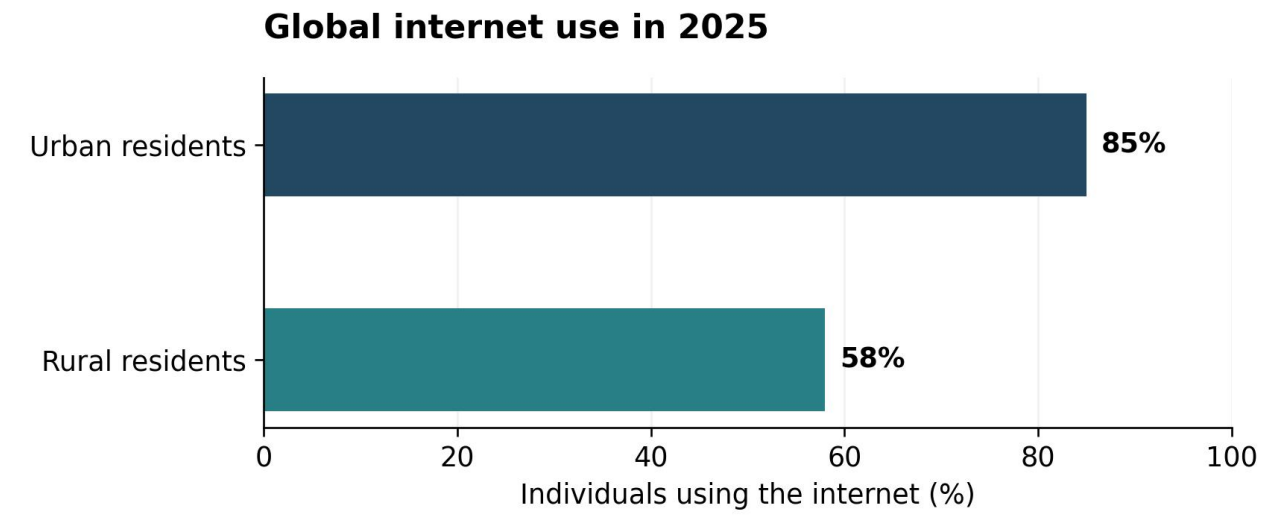
The numerical example uses three hypothetical project alternatives and two population groups. All inputs, weights and thresholds are stated. Its purpose is to expose a decision mechanism, not estimate an intervention effect. No residents were recruited, no personal data were processed and no municipal deployment was conducted for this article. Consequently, the results concern the internal behaviour of the proposed decision rule and should not be interpreted as causal findings about AI adoption.

Recent evidence and its implications

1.1 Participation requires more than a digital interface

The WeDesign preprint reports a Montreal workshop involving five focus groups and interviews with six planning professionals. Participants found immediate image generation useful for discussion, while identifying difficulties with accessibility needs, local architectural details and bilingual prompts. The study illustrates the value of visual dialogue and the importance of facilitation. Its limited setting does not establish that generated imagery produces representative participation or more equitable completed projects.

A design image can make a proposal easier to discuss while also narrowing the discussion. If the first image depicts a commercial promenade, participants may debate paving and seating before asking whether the site should accommodate affordable housing. For this reason, the proposed framework records the initial brief, the alternatives rejected before visualisation and the changes requested by residents. Facilitators should also accept non-visual contributions, including accounts of unsafe routes, informal uses and maintenance problems. These observations may be more consequential than



HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

preferences between rendered scenes.

Connectivity remains a material constraint. ITU's 2024 release reported internet use of 83% in urban areas and 48% in rural areas worldwide. Its 2025 release reported 85% and 58%, respectively. Figure 1 uses the latest of these releases
as a cross-sectional comparison. The figures imply a 27 percentage-point difference in 2025 and a 15% urban non-user share. They do not measure planning participation, affordability, digital skills or the representation of any particular city. Earlier and later release estimates should not be treated as a harmonised panel without checking revisions and methods.

The planning implication is a sampling problem. An online consultation observes the preferences of respondents who can access the process and choose to participate. Neither internet access nor a large response count makes this group representative. A municipal team should compare participants with the affected population, document missing groups and provide staffed alternatives. The inclusion assessment should examine influence over the eventual decision as well as attendance.

1.2 Technical capability and public benefit are different forms of evidence

Mashhadi Moghaddam and Cao examined 28 cases selected from 157 documented deployments across six domains, covering 2015-2024. Their analysis identifies situations in which technical performance and social outcomes diverge. Purposeful selection and documentation requirements limit statistical generalisation. The study is nevertheless relevant to evaluation design because it questions the use of a model's benchmark score as a proxy for successful policy.

For a heat-mitigation project, a model may classify exposed surfaces accurately while recommending investment mainly where land is available. The resulting intervention could improve the measured citywide average and leave residents with the greatest exposure largely unaffected. Evaluation therefore needs two distinct tests. The first examines the model's ability to represent the relevant conditions. The second examines the consequences of the decision made using that representation. Passing the first test does not logically entail passing the second.

The 2025 preprint by Yang, Li and Biljecki proposes an architecture for reasoning-capable planning agents with explicit verification and evaluation functions. It is useful as an account of an emerging technical direction, not as proof of reliable autonomous planning. In the framework developed here, machine-generated explanations remain claims to be checked against inputs, constraints and actual computations. An articulate justification is insufficient if the system cannot reproduce the relevant result or identify the evidence supporting it.

1.3 Institutional examples and the limits of transfer

The OECD's 2025 municipal issues note describes Eslöv's use of the Hektar platform for early-stage zoning and design exploration. It reports the combination of geospatial and project data with scenario generation and professional review. The account draws on an August 2025 interview and should be understood as an implementation description rather than an independent causal evaluation. Its practical significance is the connection between scenario exploration and an interdisciplinary review process.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

The same note describes Barcelona's RESPIRA ventilation system, which combines operational measurements with predictive control. This infrastructure application differs from allocating redevelopment sites: its operating variables can be monitored frequently, while displacement or access to public space may take years to evaluate. A governance approach must therefore reflect the observability and reversibility of each decision. Evidence from an operational control system cannot be transferred directly to land-use decisions with long-lived consequences.

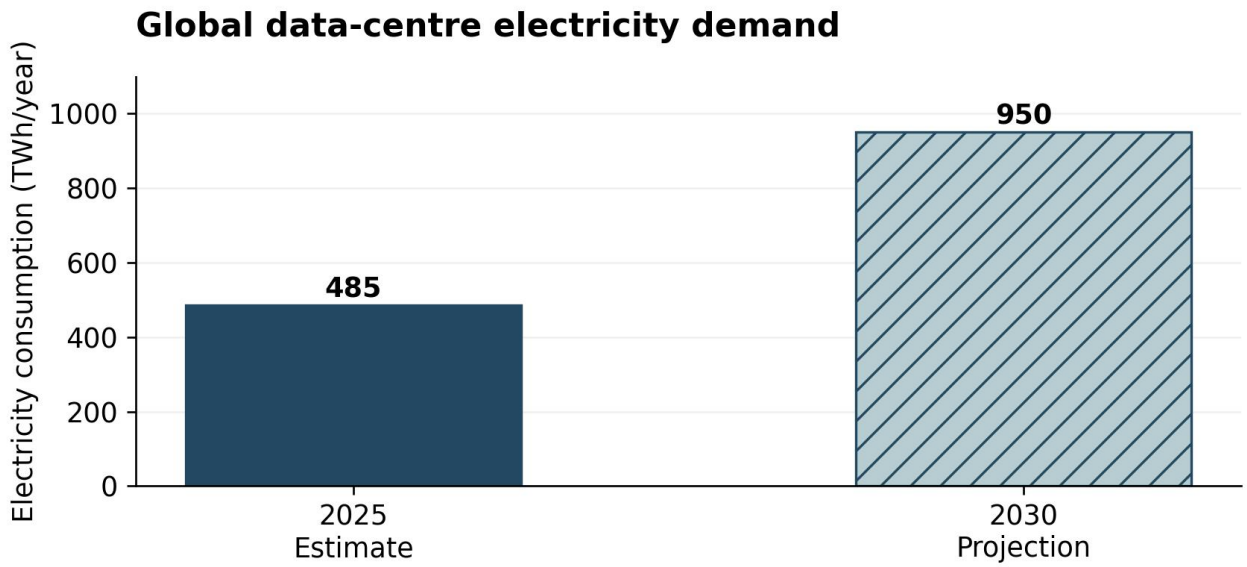
The UK Algorithmic Transparency Recording Standard provides another institutional reference. Guidance updated in May 2025 distinguishes mandatory central-government scope from voluntary use elsewhere. This article draws on its documentation approach, without claiming that the requirement applies automatically to every municipality or another jurisdiction. A planning register should disclose the purpose, responsible authority, data categories, role of the algorithm, evaluation and avenues for challenge in language residents can understand.

1.4 Urban AI has physical and organisational dependencies

Xu, Seto and Weng argue that urban AI governance must address both accountability and infrastructure security. Their 2026 Comment highlights the difficulty of assigning responsibility where municipalities are simultaneously users and overseers of systems embedded in public services. For planning practice, the implication developed here is that procurement must specify who can investigate an incident, suspend a recommendation service and restore a workable manual process.

The IEA's 2026 assessment estimates that global data-centre electricity consumption reached 485 TWh in 2025 and projects approximately 950 TWh in 2030. It reports 17% growth in total data-centre demand during 2025 and faster growth in AI-focused facilities. Figure 2 distinguishes the estimate from the projection. These totals cover data centres broadly they cannot be attributed to municipal AI or used to calculate the footprint of a particular planning application.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr



A municipality should instead request application-level resource information where available and disclose the limits of measurement. Repeated generation of detailed images or large scenario ensembles has to be justified by its contribution to a decision. For proposed data-centre development, integrated planning must examine the local grid, water arrangements and competing land uses. The fact that a facility serves digital activities does not remove its physical demands from spatial planning.

A framework for responsible planning decisions

1.1 Define the decision before selecting the model

The proposed framework begins with a decision charter. The charter states the planning problem, the affected area, the responsible public authority, the alternatives to AI and the decisions that software may support. It identifies the population entitled to participate, including people who use the area without owning property there. It also states what the project will not decide. A tool commissioned to compare walking routes should not silently become a mechanism for prioritising redevelopment parcels.

Risk is assessed through the severity of possible harm, the number and vulnerability of affected people, the reversibility of the decision and the consequences of error. These dimensions guide the intensity of review rather than produce a universal numerical risk score. A draft illustration used in a workshop may require source checks and clear labelling. An allocation affecting essential services requires independent assessment, documented distributional tests and an accessible challenge procedure. These are proposed governance categories, not a statement of legal classifications.

1.2 Establish what the data can support

The data review records provenance, collection dates, geographic coverage, missingness, permissions and known limitations. It distinguishes observed demand from administratively recorded

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

demand. Complaint records, for example, capture both a problem and the ability or willingness to report it. Treating low complaint volume as low need may direct resources away from residents whose concerns are least visible to the authority.

Spatial validation should withhold areas or time periods that approximate the intended deployment conditions. Randomly dividing neighbouring observations can exaggerate performance when training and test records share local characteristics. The appropriate division depends on the task: predicting heat exposure across unfamiliar districts requires a different test from updating conditions in a well-observed district. The validation report should explain why its test resembles the future use case.

Uncertainty must follow the data into the decision. Where disability access, tenure or informal activity is poorly recorded, the output should identify that limitation rather than substitute a confident ranking. Additional data collection must remain proportionate. An inclusion objective does not justify collecting identifiable sensitive information when aggregated observation or voluntary participation can answer the planning question.

1.3 Generate alternatives under explicit constraints

Scenario generation should begin with a conventional baseline, such as an existing planning proposal or a transparent rule-based allocation. The AI-assisted alternatives are then evaluated against the same budget, service requirements and physical constraints. This comparison makes the claimed benefit of AI testable. A complex system that reproduces the baseline without improving decisions may not justify its maintenance cost.

Some requirements should operate as constraints rather than compensable preferences. A proposal that fails an adopted accessibility requirement should not recover its acceptability merely by scoring well on architectural appearance. Other considerations, such as travel-time benefits and construction costs, can be evaluated through explicit trade-offs. The responsible authority must state which category each criterion belongs to and explain the basis for that choice.

The model should retain multiple feasible alternatives when no option performs best on every objective. An alternative is dominated if another feasible option is at least as good on all stated objectives and better on at least one. Removing dominated alternatives can simplify deliberation, but the definition of the objectives remains open to public scrutiny. Criteria omitted from the model can change which options deserve consideration.

1.4 Connect deliberation with decision authority

Figure 3 presents the proposed workflow. Residents and independent reviewers interact with the process before authorisation, while monitoring can return a project to earlier stages. Each stage produces evidence attached to the same decision record. This record is the unit of accountability: it connects a dataset and model version with an actual public decision, rather than documenting software in isolation.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr



Public deliberation should test whether the model represents the relevant problem. Participants must be able to challenge the chosen indicators, propose an additional alternative and identify burdens missing from the analysis. The authority should publish a response explaining which submissions changed the design, which did not and why. A record of attendance alone cannot show that participation influenced the outcome.

The authorising official remains responsible for selecting an alternative. The decision statement should identify the model's contribution, unresolved uncertainties, distributional effects and reasons for accepting particular trade-offs. Professional judgement may depart from the highest model score, but the departure should be explained. Otherwise, neither adherence to the model nor rejection of it can be evaluated consistently.

1.5 Monitor consequences and preserve the ability to intervene

Monitoring compares realised outcomes with the decision's stated objectives. It includes operational performance, benefit distribution and the handling of complaints. Each indicator needs an owner, reporting frequency and an action triggered by a material deviation. A model update that changes the ranking of neighbourhoods should require renewed review when it affects the basis of an authorised allocation.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

Contracts should provide access to the information needed for independent investigation, including model versions, relevant logs and evaluation outputs. Where disclosure would expose personal information or security vulnerabilities, controlled audit access may be appropriate. A public summary should still explain what was tested and what remains uncertain. Contract termination and data export arrangements matter because dependence on an inaccessible supplier can make a formal right to review ineffective.

Table 1. Proposed responsibilities and evidence required before progression

Stage	Accountable role	Required evidence
Problem definition	Planning authority	Decision charter, affected population, non-AI baseline
Data review	Data steward	Provenance, coverage, permissions, missingness
Alternative appraisal	Planning and technical team	Comparable scenarios, validation, distributional tests
Deliberation	Participation lead	Accessible channels, submissions, reasoned responses
Authorisation	Designated public official	Decision reasons, trade-offs, unresolved uncertainty
Monitoring	Service owner and independent reviewer	Outcome reports, incidents, appeals, revision record

Operational indicators and an illustrative allocation

1.1 Measuring representation and distribution

Representation can be examined with a participation ratio. For group g, let p_g be its share of the affected population and s_g its share of consultation participants. Define $R_g = s_g / p_g$, where $p_g > 0$. A value below one indicates underrepresentation relative to population share. The ratio is a diagnostic, not a certificate of legitimacy: it does not measure speaking time, influence, accessibility or diversity within a group. When group information is incomplete or voluntarily withheld, the uncertainty should be reported.

For a planning alternative x, let $b_g(x)$ denote the change in an agreed outcome for group g relative to a common baseline. The population-weighted benefit is $B(x) = \Sigma_g p_g b_g(x)$. This

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

equation requires compatible units and an explicit population denominator. Travel-time reductions should not be added directly to temperatures or income changes without a justified transformation. Even when all outcomes share a unit, the average conceals their distribution.

An equity condition can require $b_g(x) \geq m_g$ for each designated priority group, where m_g is a publicly adopted minimum benefit. The threshold is a normative choice that should reflect local priorities and feasibility. It is not generated by the predictive model. A wider appraisal should also report burdens such as displacement, service interruption and construction impacts, which may require separate constraints rather than incorporation into the same benefit score.

Table 2. Proposed evaluation indicators and their interpretation

Dimension	Indicator	Interpretation and limitation
Representation	R_g = participant share / population share	Below 1 indicates underrepresentation, does not measure influence.
Distribution	$B(x)$ and group-specific $b_g(x)$	Report average and group outcomes in the same stated units.
Minimum benefit	$b_g(x) \geq m_g$	Threshold requires public justification and uncertainty assessment.
Traceability	Decisions with complete records / reviewed decisions	Specify the audit sample and required fields.
Contestability	Resolution time and substantive changes after review	Low complaint counts may reflect barriers to complaints.
Resource use	Measured energy and total cost per assessed alternative	Include audit and participation costs, disclose allocation assumptions.

1.2 A transparent example

Consider a hypothetical district selecting among three equal-cost accessibility projects. Sixty per cent of resident's form group A and 40% form a priority group B with poorer baseline access. Benefits are measured as additional essential destinations reachable within an adopted travel-time threshold. The values represent assumed average gains per resident they are not observed data, model predictions

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

or claims about any city. Feasibility, maintenance and other impacts are assumed equal solely to isolate the distributional mechanism.

Alternative X yields gains of 12 destinations for group A and 1 for group B. Alternative Y yields 8 and 6, while alternative Z yields 6 and 8. The population-weighted gains are therefore 7.6, 7.2 and 6.8 destinations, respectively. Maximising the average selects X. Requiring a gain of at least four destinations for the priority group excludes X and selects Y among the remaining options. Relative to X, Y reduces the aggregate gain by 0.4 destinations, or approximately 5.3%, while increasing the priority-group gain from one to six destinations.

Table 3. Hypothetical gains in reachable destinations per resident with equal project costs

Alternative	Group A 60%	Group B 40%	Weighted benefit	B gain ≥ 4
X	12	1	7.6	No
Y	8	6	7.2	Yes
Z	6	8	6.8	Yes

This example does not establish that Y is universally fair. Residents may dispute the destination definition, the travel-time threshold or the grouping. A destination that is physically reachable may remain unaffordable or inaccessible to a wheelchair user. The example shows a narrower result: an explicit minimum-benefit condition changes the selected alternative, and the magnitude of the resulting trade-off can be reported instead of hidden inside a composite score.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

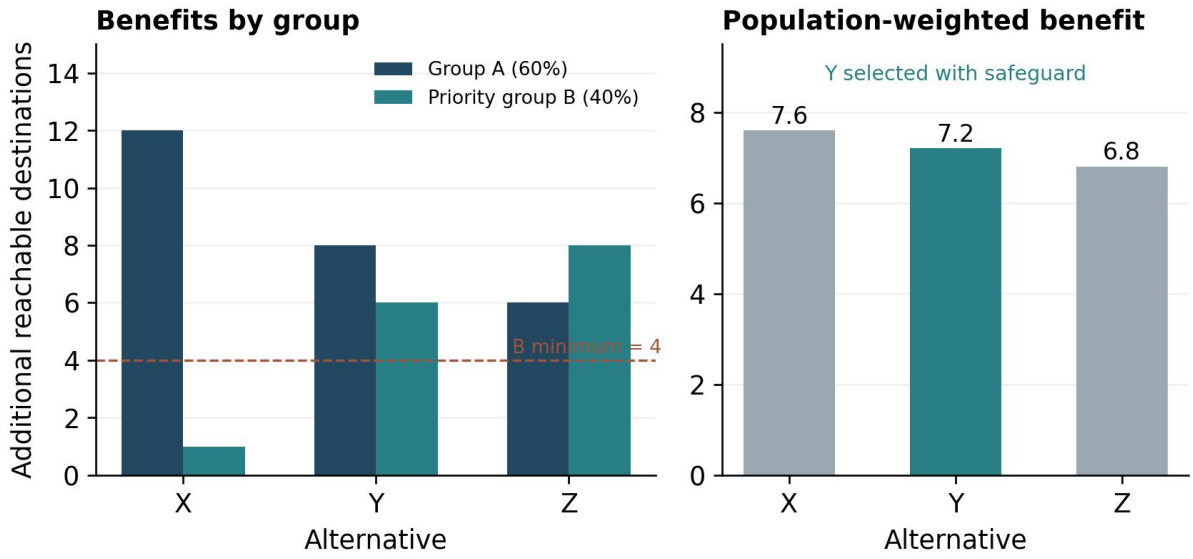


Figure 4. Distributional trade-off in the hypothetical accessibility example. Original calculations using Table 3. The minimum applies to group B, Y maximises the population-weighted benefit among qualifying alternatives. All values are illustrative.

Sensitivity analysis reveals where the decision depends on assumptions. If the minimum required gain for group B is at most one, all alternatives are feasible and X is selected under the stated weights. If it exceeds one but is at most six, Y is selected. If it exceeds six but is at most eight, Z is selected, above eight, none qualifies. With p denoting group A's population share, aggregate gains are $1 + 11p$ for X, $6 + 2p$ for Y and $8 - 2p$ for Z. X overtakes Y when p exceeds 5/9, while Y overtakes Z when p exceeds 1/2. These thresholds follow directly from the assumptions and can be checked without an AI model.

If predicted gains have substantial uncertainty, a point-estimate constraint is insufficient. A municipality could require an agreed lower uncertainty bound to exceed the minimum, provided that the interval construction is validated. Where this is impossible, the authority should describe the probability and consequences of failing the target. Presenting the range of plausible outcomes is more informative than adding decimal places to an uncertain score.

Implementation and empirical evaluation

An initial pilot should address a decision with a manageable scope, observable outcomes and a credible baseline. Public-space accessibility or the prioritisation of shade interventions may be suitable subjects when the necessary data and local partnerships exist. The choice should follow an assessment of feasibility and public need it should not be inferred from the popularity of a particular model. Pilots affecting individual entitlements or irreversible land acquisition require a different level of scrutiny.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

The first implementation stage should establish the charter, data inventory and baseline. A second stage should compare AI-assisted alternatives with the conventional approach while both remain advisory. Community review and an independent technical assessment should precede authorisation. The final stage should track implementation and outcomes, retaining a route to revision. These stages describe a proposed sequence rather than a fixed timetable. Data preparation or consultation may legitimately take longer than model configuration.

Evaluation should distinguish process effects from spatial outcomes. Process measures include staff time, correction rates, the ability to reproduce recommendations and the influence of public submissions. Outcome measures include accessibility, exposure reduction, affordability and the distribution of benefits. Satisfaction can inform the assessment, but it is not sufficient where residents have different expectations or lack an effective way to challenge the authority.

A prospective study could compare the proposed framework with conventional planning and with AI assistance lacking the additional governance procedure. Where allocation can be randomised ethically, comparable planning tasks or teams could be assigned to different procedures. Where randomisation is unsuitable, matched projects or phased implementation may support comparison. Any difference-in-differences design would require examination of pre-intervention trends and plausible control conditions. Spillovers between neighbouring areas and learning between staff could undermine otherwise credible comparisons.

The evaluation protocol should specify primary outcomes before results are available. It should also record unsuccessful trials, abandoned alternatives and the full costs of consultation and audit. Counting only software expenditure would bias the cost comparison ignoring those activities would underestimate responsible deployment costs. Conversely, documenting fewer legal disputes or less rework would require evidence over a sufficiently long follow-up period, rather than an assumption that better documentation automatically produces savings.

For municipalities with limited analytical capacity, a shared regional audit service may be more feasible than a specialist team in every authority. Such an arrangement should preserve local authority over priorities and public participation. Transfer to a new city requires checking language, settlement patterns, data coverage and institutional responsibilities. For an Uzbek municipal pilot, for example, language accessibility and the fit between cadastral records and actual service use would be questions to investigate this article does not assume that these data problems have been empirically demonstrated in a particular municipality.

Discussion and limitations

The framework treats inclusion as a property of the decision process and its consequences. This avoids making representativeness depend solely on the number of consultation responses or fairness depend solely on a model-level statistic. It also changes the role of explanation. A useful explanation must enable a resident or reviewer to identify an assumption, understand its effect and seek a correction. A general description of how a model works may provide transparency without offering that practical capacity.

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

There are unavoidable conflicts. Publishing detailed spatial evidence may expose sensitive households protecting confidentiality may restrict external scrutiny. More extensive participation may improve the brief while delaying an urgent intervention. An equity constraint may reduce aggregate benefit, as the illustration shows. These conflicts should be resolved through an explicit public mandate and proportionate review. They cannot be eliminated by declaring every objective equally important.

The framework also leaves several questions open. It does not establish universal minimum-benefit thresholds, determine which groups must receive priority or resolve disputes about land rights. Those questions require legitimate local processes and applicable law. Nor does it claim novelty for participation, auditing or distributional analysis considered separately. The contribution is their integration around a traceable planning decision, with criteria for authorisation and revision.

The evidence base is limited by purposive selection, English-language accessibility and the availability of well-documented cases. Some sources are conceptual contributions or preprints institutional case descriptions may rely on implementers' accounts. The review therefore supports analytical propositions more strongly than causal conclusions. The numerical example abstracts from heterogeneous costs, overlapping social categories and uncertain behaviour. Field testing is needed to establish whether the proposed procedure changes actual designs, benefits underrepresented residents and remains workable within municipal budgets.

Conclusion

Responsible urban AI requires a defensible chain from public objectives to evidence, alternatives and authorised action. The reviewed literature and institutional material show why technical performance, consultation technology and published documentation each provide only part of that chain. The proposed framework specifies how they can be connected through a decision charter, distributional assessment, deliberation, named responsibility and monitoring.

The worked example establishes a limited but important result: when an explicit minimum benefit for a priority group is introduced, the preferred project can change even though its aggregate benefit is lower. Reporting this trade-off gives public authorities and residents a basis for judgement. Concealing it within a single score does not. The next research task is a prospective municipal evaluation that measures both the additional governance costs and the resulting distribution of urban benefits.

References

1. OECD. (2025). Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions. OECD Publishing. https://doi.org/10.1787/795de142-en

2. Kerimi, D. (2025). From the power broker to a(i) power broker. npj Urban Sustainability, 5, 97. https://doi.org/10.1038/s42949-025-00290-0

3. Mashhadi Moghaddam, S. N., & Cao, H. (2026). The metrics trap: How technical sophistication masks social harm in urban AI systems. npj Urban Sustainability, 6, 94. https://doi.org/10.1038/s42949-026-00394-1

HSR (LONDON),
HOUGHTON STREET REVIEW JOURNAL
ISSN: 2058-5519|Impact factor-15.9
http://houghtonstreetreview.com/index.php/hsr

4. Xu, Y., Seto, K. C., & Weng, Q. (2026). The invisible gap: Urban AI security. npj Urban Sustainability, 6, 101. https://doi.org/10.1038/s42949-026-00446-6

5. Mushkani, R., Berard, H., & Koseki, S. (2025). WeDesign: Generative AI-Facilitated Community Consultations for Urban Public Space Design. arXiv preprint, version 3, 15 September 2025. https://doi.org/10.48550/arXiv.2508.19256

6. Yang, S., Li, J., & Biljecki, F. (2025). Reasoning Is All You Need for Urban Planning AI. arXiv preprint. https://doi.org/10.48550/arXiv.2511.05375

7. OECD. (2025). Artificial Intelligence for Advancing Smart Cities. Issues note, CFE/ RDPC/ URB (2025) 27. https://www.oecd.org/content/dam/oecd/en/about/ programmes/cfe/the-oecd-programme-on-smart-cities-and-inclusive-growth/Issues-Note-AI-for-advancing-smart-cities.pdf

8. International Telecommunication Union. (2024). Measuring Digital Development: Facts and Figures 2024. Internet use in urban and rural areas. https://www.itu.int/itu-d/reports/statistics/2024/11/10/ff24-internet-use-in-urban-and-rural-areas/

9. International Telecommunication Union. (2025). Measuring Digital Development: Facts and Figures 2025. https://www.itu.int/en/mediacentre/Pages/PR-2025-11-17-Facts-and-Figures.aspx

10. Government Digital Service. (2025). Algorithmic Transparency Recording Standard: Guidance for Public Sector Bodies. Updated 8 May 2025. https://www.gov.uk/government/publications/guidance-for-organisations-using-the-algorithmic-transparency-recording-standard/algorithmic-transparency-recording-standard-guidance-for-public-sector-bodies

11. International Energy Agency. (2026). Key Questions on Energy and AI. Executive summary. https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary