Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Who can learn from whom? The geography of urban climate evidence

Andrew Sudmant

andrew.sudmant@ed.ac.uk

University of Edinburgh

Egide Kalisa

University of Western Ontario

Simon Montfort

École Polytechnique Fédérale de Lausanne

Wenjia Cai

Tsinghua University

Shaurya Patel

Ahmedabad University

Felix Creutzig

Potsdam Institute for Climate Impact Research

Research Article

Keywords: climate change, learning, cities

Posted Date: March 31st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9268986/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Who can learn from whom? The geography of urban climate evidence

Sudmant, A; Kalisa, E; Montfort, S; Cai, W; Patel, S; Creutzig, F

Abstract

Published research on climate change in cities is highly uneven, but cities are not equally isolated. We map the geography of urban climate evidence by combining a city-by-domain evidence matrix with domain-specific similarity spaces, allowing us to distinguish direct evidence from the evidence cities can access through comparable analogues. Across vulnerability, adaptation, mitigation, and governance, case study evidence is sparse and concentrated in a limited set of region × type cells, with adaptation consistently the thinnest domain. Yet accessible evidence is broader than direct evidence alone would suggest. Forty-five percent of cities have no direct evidence, but 95.5% of those cities can access analogue experience by drawing from peers and all can do so when learning pools widen to regional or global peers. The largest gains come from widening within regions across city types rather than moving immediately to global pools of evidence. At the same time, the supplier side of the system is highly concentrated: evidence-rich hubs account for a disproportionate share of published and accessible climate action knowledge. These findings show that urban climate learning is structured, directional, and only partly open. By moving from evidence geography to learning geography, the paper identifies where cities can credibly learn from others and where new primary evidence remains most urgently needed.

Introduction

What does a city do when it has local knowledge of a climate problem, but little documented experience from comparable places to draw on? Cities rarely begin from zero knowledge, responses to climate change are shaped by lived experience, professional judgement, community expertise, institutional practice, and political contestation. The narrower question taken up here is how much climate knowledge is available in forms that are documented, comparable, and portable enough to support learning across cities. Research on climate change in cities now provides rich accounts of hazards, vulnerability, adaptation, mitigation, experimentation, and the politics of response. What remains less clear is the structure of transferable climate knowledge: where it is concentrated, how far it travels, and under what conditions it retains contextual relevance (Castán Broto and Bulkeley, 2013; Anguelovski et al., 2014; Bulkeley et al., 2013, 2014).

This gap sits at the intersection of three literatures. First, research on climate change in cities has shown that urban responses emerge through experiments, adaptation planning, justice claims, and increasingly diverse actor coalitions, rather than as simple reactions to climate risk alone (Castán Broto and Bulkeley, 2013; Anguelovski et al., 2014; Bulkeley et al., 2013; Klein et al., 2018). Second, work on policy transfer and policy mobility shows that policies do not move as fixed packages; they are translated, reassembled, and selectively borrowed across place (Benson and Jordan, 2011; McCann and Ward, 2012). Third, comparative urbanism has argued that cities should be compared relationally rather than treated as sealed categories whose differences foreclose learning in advance (Robinson, 2011). Recent work in urban studies has also argued that replication and generalizability need to be treated more explicitly if urban research is to cumulate coherently across cases and disciplines (Sudmant et al., 2025). Taken together, these literatures imply that city-to-city learning is real, but context-dependent. What they do not yet provide is a systematic account of the evidence infrastructure through which such learning becomes possible.

This missing piece matters for practice because documented examples often do more than illustrate what cities have done; they help make action legible to planners, funders, and decision-makers. This is not because documented research is superior to local or community knowledge. It is because some forms of knowledge circulate more easily through planning guidance, economic appraisal, funding applications, and claims about feasibility. The economics and co-benefits literature has been influential for precisely this reason. Gouldson et al. (2015), for example, frame the economic case for climate action as a way to strengthen political commitment, institutional capacity, finance, and implementation and Sudmant et al. (2016) make a similarly direct argument that ambitious urban mitigation can be understood as an investment rather than a net burden. But the force of such arguments depends on where the supporting examples come from and whether they are seen as relevant. If the cases that anchor the economic or institutional case for action come from a narrow set of cities, the question is not simply whether the case is compelling; it is whether it is portable.

This paper addresses that problem by conceptualising published research on climate change in cities as an architecture of urban climate learning. Rather than asking only which cities appear in climate research, it distinguishes between direct evidence, accessible evidence, replication readiness, scaling access, supplier roles, and priority deficits. In this paper, "evidence" is used in a narrow analytical sense to refer to documented urban climate experience in forms that can circulate through research, planning, and policy processes; the term does not imply superiority over local, professional, or community-based knowledge. The analysis covers five domains, hazards, vulnerability, adaptation, mitigation, and governance, because similarity is domain-specific rather than generic. The peers that matter for hazards are not necessarily the peers that matter for mitigation or governance. Reachable evidence is therefore defined not as any evidence from anywhere, but as the stock of documented experience that sits near a city in the similarity space most relevant to the problem at hand.

We operationalise this idea by building a city-by-domain evidence matrix from the published literature and combining it with domain-specific similarity spaces to identify analogue cities. This allows us to compare the direct supply of documented experience with the experience cities can access from neighbouring analogues, and to distinguish among three learning pools: within the same region × type cell, across types within a region, and globally. From there, we assess replication readiness under strict similarity, the gains from widening the learning pool, the supplier roles of different city-types, the concentration of hubs in the global similarity network, and the cells that appear to be highest priorities for future evidence generation. The central claim is simple: urban climate learning is not flat, universal, or frictionless. It is structured, uneven, and directional.

Our results show that published research on climate change in cities contains a substantial documented record, but that this record is unevenly distributed across problems and places. Evidence supply is concentrated in a limited number of region × type cells and hub cities, with hazards and vulnerability consistently stronger than adaptation. Yet access is broader than direct evidence alone would suggest. At the city level, 95.5% of zero-direct cities still have accessible evidence under strict analogues, and that figure rises to 100% when the pool widens to within-cell, regional, or global neighbours. The largest distributive gains come not from moving immediately to the global pool, but from widening from same-type peers to broader cross-type pools within regions. Higher-evidence city-types generally emerge as the principal suppliers of transferable knowledge, while lower-evidence types are more often the beneficiaries of borrowing. Cross-domain learning exists, but it is presently much less differentiated than the cross-type story.

This argument matters for theory because it connects research on climate change in cities to broader debates on policy mobility, comparability, and the geographies of knowledge production. It matters for practice because it distinguishes cities that can borrow, replicate, or scale existing knowledge from cities that still require new primary evidence. The paper therefore contributes conceptually by reframing the published literature on climate change in cities as an architecture of learning; methodologically by combining city-domain evidence mapping with domain-specific similarity spaces; empirically by identifying supplier, borrower, hub, and frontier positions in the global

evidence base; and practically by clarifying where new research is likely to add the most value for urban climate action.

The remainder of the paper proceeds as follows. The next section explains the construction of the evidence matrix and similarity spaces before presenting results on supply, readiness, scaling, supplier hierarchies, hubs, and priorities. The paper concludes by considering what a more explicit politics of urban climate learning implies for comparative urban research and for the design of more usable climate knowledge infrastructures.

2. Methods

2.1 Analytical design

This study treats published urban climate research as an architecture of transferable learning rather than as a simple inventory of case studies. In line with work on policy transfer, policy mobility, and comparative urbanism, the central question is not only which cities appear in the literature, but which kinds of documented experience may travel credibly across cities under conditions of contextual similarity. Empirically, the analysis combines two linked objects: a city-by-domain evidence matrix and a set of domain-specific similarity spaces. The first captures where documented urban climate experience is directly available. The second identifies which cities are plausible analogues for a given problem. Bringing these objects together makes it possible to distinguish direct evidence from accessible evidence and, from there, to estimate replication readiness, the gains from widening the learning pool, supplier roles, hub positions, and priority deficits.

A central design choice is to separate the evidence layer from the context layer. The evidence layer is built from the content of published case studies: what domains each study speaks to, and which cities it covers. The context layer is built from city characteristics: the variables used to decide which cities are sufficiently similar for learning in a particular domain. The former therefore identifies what kind of climate evidence exists; the latter identifies where that evidence may be relevant and portable.

2.2 Global urban typology used a baseline to identify learning possibilities

We build on an existing urban typology covering 11,000 urban centres across the world. Developed using 12 structural indicators (population, population growth, population density, population density growth, ratio between old and young population, Human Development index, Gender equality index, GDP pc, GDP pc. Growth, critical infrastructure, heating and cooling degree days), it consists of four types of urban centres (Montfort et al. 2026).

Type 1 cities are smaller, lower-income settlements with limited infrastructure and high cooling needs, growing at intermediate to fast rates and concentrated in Central America, sub-Saharan Africa, and South and Southeast Asia. Their most pressing climate-action priorities are foundational: strengthening basic services, expanding resilient water and sanitation systems, reducing exposure in informal settlements, and building preparedness for heat-related health risks. Mitigation capacity is constrained by development needs, making adaptation the near-term priority.
Type 2 cities are economically growing, mid-tier urban centres characterised by rapid horizontal expansion, low population density, and intermediate infrastructure, found primarily across South and Central America, Africa, and South and Southeast Asia. The central challenge is to upgrade infrastructure ahead of climate impacts while mainstreaming climate action within broader development planning. Spatially, the priority is steering expansion toward more compact, transit-oriented forms before low-density development becomes locked in.

Type 3 cities are wealthier, slower-growing urban centres with strong infrastructure and high per-capita emissions, concentrated in North America, Europe, East Asia, and southern South America. Having largely passed the infrastructure build-out phase, these cities face a deep mitigation agenda: building retrofits, electrified transport, and densification of low-density suburbs. Adaptation needs centre on extreme heat, flood risk, and the renewal of ageing infrastructure.

Type 4 cities combine high density and fast growth with relatively advanced infrastructure, but face acute cross-sectoral pressures distributed across Africa, Asia, Europe, and the Americas. Their climate challenge is the most complex: simultaneously managing rapid development, addressing poverty and informal settlements, reducing congestion and pollution, scaling nature-based cooling, and building the governance capacity needed to coordinate across these competing demands. For this type, mitigation and adaptation cannot be sequenced — they must be integrated from the outset.

2.4 From case studies to domain-specific evidence

Each article is assigned a domain profile across five domains: vulnerability, adaptation, mitigation, and governance. The pipeline uses dictionary-based procedures applied to article titles and abstracts. Domain-specific terms and compound phrases associated with each of the five domains are applied to identify article-level weights that are then normalised across domains so that a single paper may contribute to one domain or to several at once. Where an article contains no detectable domain signal under the fallback procedure, all domain weights are set to zero.

Vulnerability dictionary	impact*; risk*; expos*; vulnerab*; susceptib*; mortality; morbidity; health; poverty; informal; slum*; inequal*.
Adaptation dictionary	adapt*; resilien*; prepared*; early_warning; risk_assessment; green_infrastructure; nature_based; adaptation_plan; disaster_risk; relocation.
Mitigation dictionary	mitigat*; emission*; decarbon*; net_zero; carbon; co2; ghg; renewable_energy; energy_efficiency; transport; public_transport; building*; retrofit.
Governance dictionary	govern*; policy; institution*; planning; budget; finance; investment; climate_finance; participat*; equity; justice; gender; capacity_building; accountab*.

This method therefore does not infer whether a study is about mitigation, adaptation, governance, hazards, or vulnerability from the characteristics of the cities alone. It estimates domain relevance at the level of the study itself. City-level covariates enter later, when the analysis asks where that evidence may travel.

Region	Number of case studies				Cases per city				Median cases per city			
	Type 1	Type 2	Type 3	Type 4	Type 1	Type 2	Type 3	Type 4	Type 1	Type 2	Type 3	Type 4
Africa	11204	3480	3210	2112	6	12	84	64	0	1	36	17
Asia	12270	16660	18936	21643	4	7	29	305	0	1	4	94
Australasia	0	0	4815	0			109				15	
Europe	1	2798	47882	7105	1	15	45	1421	1	3	8	1020

North America	449	2488	35835	5001	17	17	90	1667	2	2	16	1249
Small Islands	741	640	2160	0	7	13	98		2	7	15	
South America	2632	8612	5977	3649	10	18	74	608	2	4	29	468

2.5 Article-city allocation and the five domains

The literature table is expanded into one row per article-city pair. If article a mentions n_a cities, each linked city i receives the fractional city weight shown below. Multi-city studies therefore contribute fractionally to each constituent city while retaining the same article-level domain profile. This prevents a multi-city paper from counting as a full case study for every city it mentions, while still allowing it to contribute evidence to each of them.

$$w^{(city)}_{ai} = 1 / n_a$$

The five domains play a dual role in the analysis. First, they classify the topical content of the literature. Second, they define the city-level comparison spaces used to identify analogue cities. Table 1 summarises that second role by listing the city-level data types and variables associated with each domain in the similarity analysis. The table should therefore be read as an operational description of how the comparison spaces are built, not as a claim that these variables exhaust the substantive meaning of each domain.

Table 1. Domains, associated city-level data types, and full variable sets used in the baseline similarity spaces

Domain	Article-level evidence focus	Associated city-level data types in the similarity space	Variables used in the baseline similarity space
Vulnerability	Published evidence on exposure, sensitivity, susceptibility, and coping capacity	Demographic structure; infrastructure capacity; human development; gender inclusion	old_pop_relative_to_young; critical_infrastructure; human_development_index; female_gender_index
Adaptation	Published evidence on preparedness, resilience, adaptation planning, and adaptive response	Demographic change; urban form change; infrastructure capacity; climatic exposure	population_growth; population_density_growth; critical_infrastructure; heating_degree_days; cooling_degree_days
Mitigation	Published evidence on emissions reduction, decarbonisation, energy demand, and low-carbon transition	Urban form and density; economic scale and change; climate-related energy-demand context	population_density; gdp_ppp; gdp_ppp_growth; heating_degree_days; cooling_degree_days
Conditions for Governance	Published evidence on policy, planning, institutions, finance, participation, justice, and enabling capacity	Economic capacity; enabling infrastructure; human development; gender inclusion	gdp_ppp; critical_infrastructure; human_development_index; female_gender_index

Note. The evidence-focus column refers to the coding of the literature. The data-types and variables columns refer to the separate city-level comparison architecture used to identify analogue cities. The city-level variables do not classify the articles into domains.

The constructs in table 1 are intentionally parsimonious. Vulnerability combines demographic, infrastructural, and development proxies but does not observe informal settlements, housing quality, or fine-grained inequality directly. Adaptation and governance are especially difficult to represent using globally harmonised city data, so the variables used here should be read as proxies for implementation context and enabling conditions rather than direct measures of institutional quality or policy effectiveness.

The constructs in Table 1 are intentionally parsimonious, and each carries clear limitations. Hazards is the narrowest construct in the baseline specification, capturing broad climatic stress through heating and cooling degree days rather than a full multi-hazard profile. Vulnerability combines demographic structure, infrastructure, and development as coarse markers of susceptibility and coping capacity, but it does not directly observe informal settlements, housing quality, distributive inequality, or fine-grained exposure. Adaptation and governance are harder still to represent with globally harmonised city-level variables, so the covariates used here should be read as proxies for implementation context and enabling conditions rather than as direct measures of institutional quality, political leadership, fiscal capacity, or policy effectiveness. Mitigation captures broad structural conditions relevant to urban energy demand and decarbonisation opportunity, but not sector-specific emissions profiles, energy systems, or regulatory design.

These limitations do not invalidate the exercise, but they do set its scope. The domain spaces are best understood as problem-oriented similarity devices that help identify where documented evidence may travel more plausibly across cities. They are not complete ontologies of hazards, vulnerability, adaptation, mitigation, or governance, and they should not be interpreted as definitive rankings of cities' underlying conditions in those domains. A richer representation of each domain would certainly capture additional nuance, but it would also increase the number of missing cities, weaken global comparability, and make the analogue logic less transparent. We therefore prioritise a smaller set of auditable covariates that can be applied consistently across the full city sample, while keeping the interpretation of results correspondingly modest.

2.6 Direct evidence

Direct evidence for city i in domain d is defined as the weighted sum of the article-city-domain contributions attached to that city. The baseline specification leaves the article-level multiplier λ_a equal to 1. The scripts implement optional recency weights, citation weights, and open-access filters, but these are switched off in the main run so that the baseline evidence measure is driven by city allocation and domain weights alone.

$$E_{id} = \Sigma_a \lambda_a \cdot w^{(city)}_{ai} \cdot w^{(domain)}_{ad}$$

This measure yields a city-by-domain evidence matrix. In substantive terms, E_id is the effective stock of documented experience available directly for city i in domain d, after allowing both multi-city allocation and multi-domain attribution. The matrix is summarised in two ways. First, it is treated as a continuous measure of evidence mass. Second, it is converted into thresholded coverage measures, particularly at $E \geq 1$ and $E \geq 5$, so that the analysis can distinguish sparse from more developed documented coverage. In the grouped heatmaps, the quantity shown is the share of cities in each region × type × domain cell that exceed the relevant threshold, not the total evidence mass in that cell.

Alongside evidence mass, the pipeline also records support-aware counts: the number of article-city pairs and unique articles contributing soft support, active support, and primary-domain support to each city-domain combination. These measures are used mainly for diagnostics and robustness checks, while the main empirical results rely on evidence mass and its thresholded coverage.

2.7 Domain-specific similarity spaces and accessible evidence

A separate similarity space is estimated for each domain using the city-level variables listed in Table 1. Before neighbour search, continuous variables are winsorised at the 1st and 99th percentiles. Population, population density, and GDP PPP are log-transformed where included. Missing values in the human development index and female gender index are first imputed using region medians; any remaining missing values are then imputed using global medians. All domain-specific feature matrices are standardised before distances are computed.

For each domain d, nearest-neighbour distances in the standardised feature space are converted to similarities using a Gaussian kernel. In the baseline specification, the domain-specific bandwidth σd is estimated as the median non-zero distance among the 50 nearest neighbours in that domain. The main analysis uses k = 50 nearest neighbours, with k = 25 and k = 100 retained for sensitivity analysis. High similarity is defined at $S \geq 0.70$, with additional checks at 0.60 and 0.80.

$$S^{(d)}_{ij} = \exp\left(- (Dist^{(d)}_{ij})^2 / (2\sigma_d^2) \right)$$

Accessible evidence is the stock of direct evidence that lies within a city's analogue neighbourhood in the relevant domain. In the main specification, accessible evidence is measured as an unweighted sum of the direct evidence held by reachable neighbours. The baseline measure is therefore extensive rather than similarity-weighted: it estimates the volume of documented experience available within a city's admissible comparison set. Similarity-weighted variants are retained as robustness checks.

$$A^{(P,k)}_{id} = \Sigma_{j \in N^{(P,k)}_{id}} E_{jd}$$

The analysis distinguishes four learning pools. The first is the strict analogue pool, which counts only neighbours in the same region × type cell whose similarity exceeds the high-similarity threshold. The second is the within-cell pool, which uses all k neighbours in the same region × type cell regardless of whether they exceed the strict threshold. The third is the regional pool, which widens comparison to all city types within the same region. The fourth is the global pool, which widens further to all cities. This nested structure is substantively important because it separates learning from close same-type peers, learning from broader cross-type regional peers, and learning from the global stock of evidence.

Two scaling metrics are then constructed. Horizontal gain is the increment in accessible evidence obtained by moving from the within-cell pool to the regional pool. Global gain is the further increment obtained by moving from the regional pool to the global pool. Replication readiness is assessed under the strict analogue pool, using both the amount of high-similarity reachable evidence and indicator measures of whether at least one qualifying neighbour exists.

3. Results

At the core of the results is a tension between comparability and reach. Cities can usually expand the stock of documented experience available to them, but those gains are often secured only by widening the comparison set beyond close analogues, and the terms of that trade-off vary sharply across region x type cells and domains. The analysis therefore treats urban climate learning as a problem of structured access rather than simple coverage. It asks not only where evidence is directly available, but where it remains usable under stricter or looser analogue conditions, and how those differences sort cities into distinct positions within the wider architecture of urban climate learning.

3.1 Direct evidence is sparse and uneven, with adaptation the thinnest domain

Figure 1 establishes the direct-evidence baseline using E >= 5 as a deliberately conservative threshold. At that level, direct coverage is concentrated in a limited set of Type 3 and Type 4 cells, while many Type 1 and Type 2 cells remain weakly documented or fall below the visibility threshold altogether.

The domain ordering is also clear. Vulnerability and mitigation retain the broadest direct coverage, governance is intermediate, and adaptation is consistently the thinnest domain.

Panel B shows that this is not only a story of a few exceptional cells. Across visible cells, median coverage remains low in every domain, and adaptation sits closest to the floor. Cities therefore enter the learning architecture from sharply unequal evidentiary starting points. Some begin with a modest but usable stock of directly documented experience, whereas others begin with very little that is locally anchored in the published record.

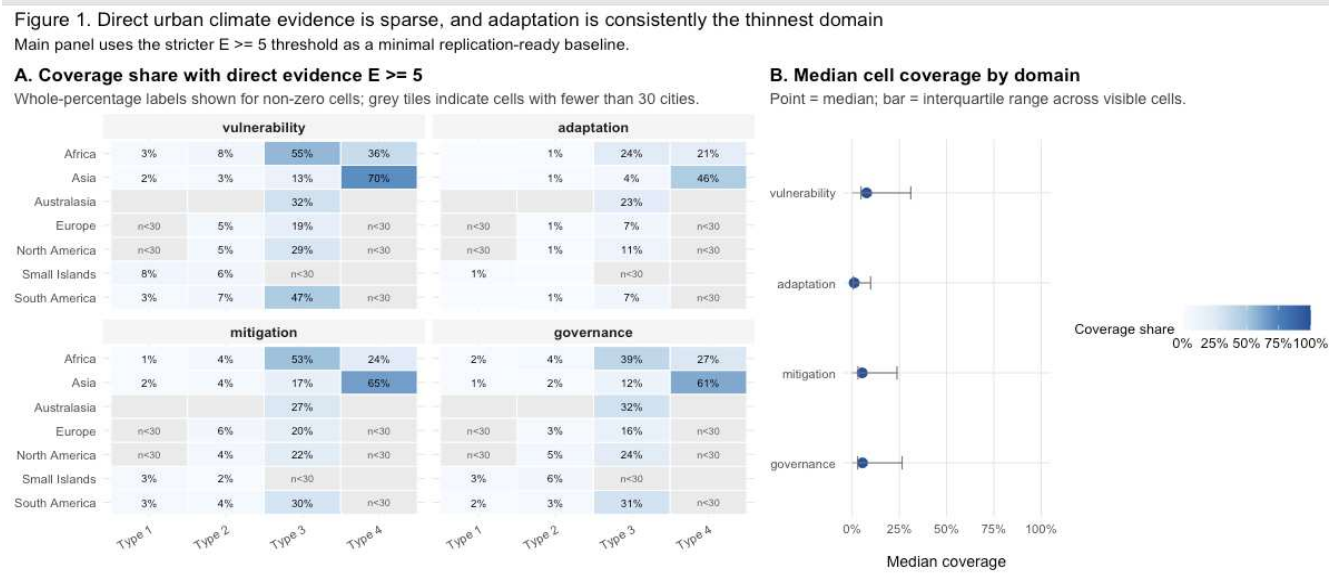


Figure 1. Direct urban climate evidence is sparse and uneven across the four non-hazard domains. Panel A shows the share of cities in each region x type cell with direct evidence mass E >= 5; grey tiles mark cells with fewer than 30 cities. Panel B shows median cell coverage by domain across visible cells.

3.2 The trade-off between analogue quality and evidence volume varies sharply across cells

Figure 2 shows, for six selected region x type x domain cells, how median cumulative reachable evidence grows as the admissible frontier widens from very high to progressively lower similarity under within-cell, regional, and global learning pools. The top row preserves the absolute evidence scale while the bottom row normalizes each cell within each pool to its own maximum, which makes frontier shape easier to compare independently of volume.

All selected cells gain reachable evidence as similarity constraints loosen, but the source and timing of those gains. Europe / Type 3 / mitigation and Europe / Type 3 / adaptation show the tightest overlap across pools, which suggests that same-cell analogues already recover much of the stock that broader searches would provide. Africa / Type 2 / governance and Europe / Type 2 / adaptation sit closer to the opposite end. Their within-cell frontiers remain thin through much of the curve, and sizeable gains appear only once borrowing widens to broader regional or global pools. Africa / Type 1 / mitigation and Africa / Type 1 / vulnerability occupy an intermediate position, where widening the pool helps but does not completely transform the reachable stock.

The normalized row shows that these contrasts are not only about magnitude. Some cells accumulate a large share of reachable evidence relatively early, while others remain flat until similarity has already been relaxed substantially. That distinction matters because it separates cells that can learn mainly through closer replication from cells that can learn only by moving quickly toward weaker

analogues and heavier contextual translation.

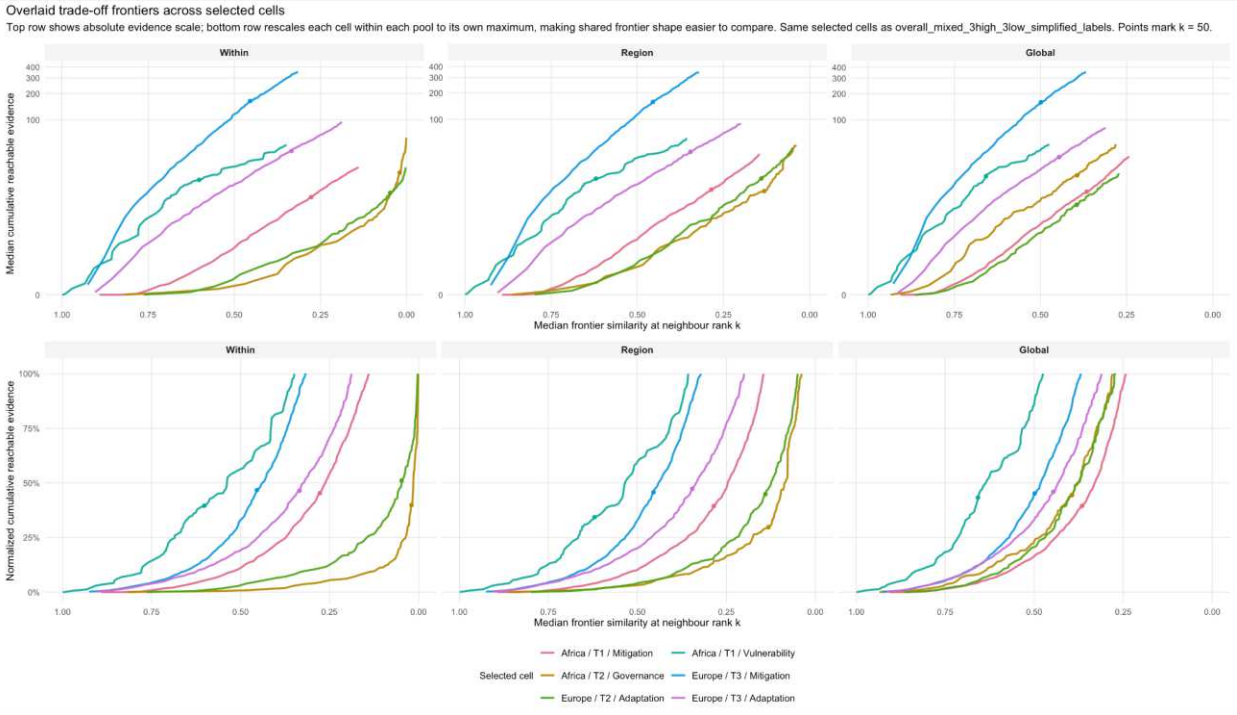


Figure 2. Overlaid trade-off frontiers across selected cells. Columns show within-cell, regional, and global learning pools. In each panel, the reversed x-axis moves from higher to lower similarity as the admissible neighbour frontier widens. The top row retains the absolute evidence scale; the bottom row normalizes each cell within each pool to its own maximum. Points mark k = 50. Selected cells are Africa / T1 / mitigation, Africa / T1 / vulnerability, Africa / T2 / governance, Europe / T2 / adaptation, Europe / T3 / mitigation, and Europe / T3 / adaptation.

3.3 Cells occupy distinct learning conditions once close-peer and broader reach are considered jointly

The frontier profiles in Figure 2 show that borrowing conditions differ sharply across contexts. Figure 3 generalizes that pattern across the full displayed cell set by combining two dimensions of reach: median within-cell high-similarity reachable evidence on the x-axis and median global reachable evidence on the y-axis, benchmarked within each domain. Cells occupy all four quadrants rather than collapsing into a single gradient of advantage. Some remain weak on both axes and therefore sit at the frontier. Some have limited close-peer reach but stronger access to broader pools and are therefore borrowable. Others are replication-ready, with stronger close-peer reach but less reliance on global pools. The most advantaged cells combine both strengths and operate as scaling platforms.

A second point is methodological but substantively important. The plotted view is a positive-within-cell subset: cells with zero median within-cell high-similarity reach are not displayed. Even within that restricted set, the heterogeneity is pronounced. The full architecture is therefore, if anything, more uneven than the figure alone suggests. Read against Figure 2, the message is consistent. Broader

borrowing is widespread, but close-peer readiness is much more selectively distributed.

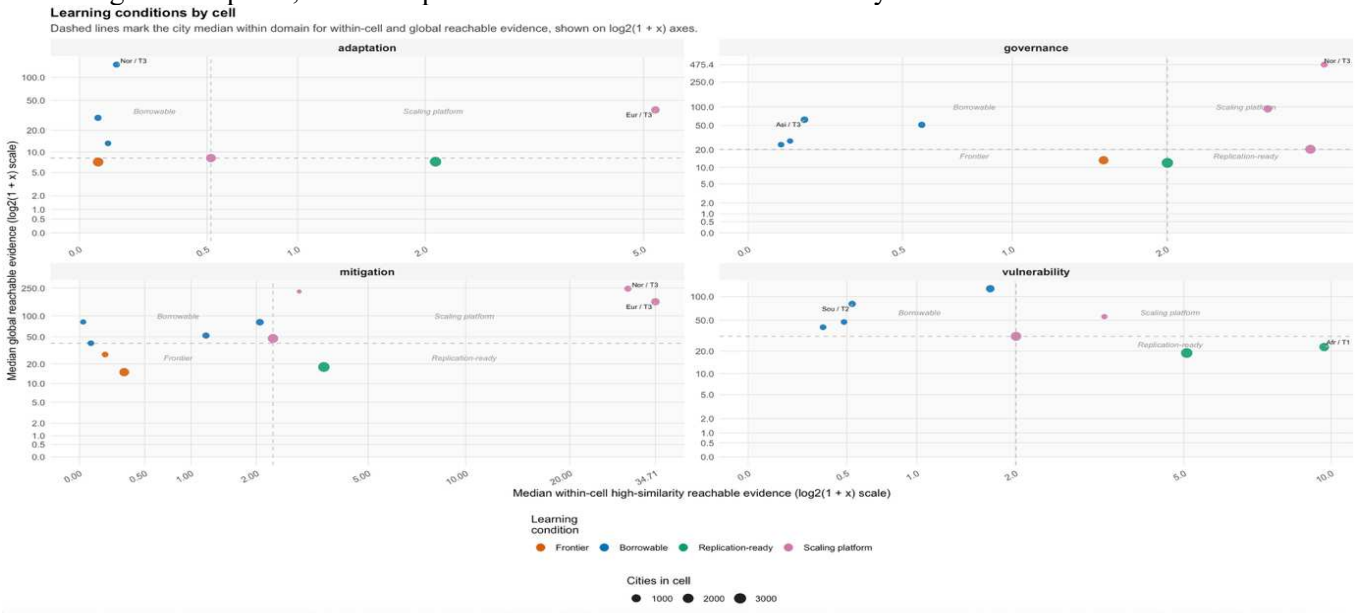


Figure 3. Learning conditions by cell. The x-axis shows median within-cell high-similarity reachable evidence and the y-axis shows median global reachable evidence, both on log2(1 + x) scales. Dashed lines mark the city median within each domain; only cells with positive median within-cell high-similarity reach are shown.

3.4 The supplier side of the system is concentrated in evidence-rich hubs more than in analogue centrality alone

Learning is not only a question of who can borrow; it is also a question of which cities supply the transferable stock. Figure 4 decomposes concentration in the global learning network into global in-degree, evidence mass, and their product as the hub score. Global in-degree is comparatively diffuse: the top 5% of cities account for only a small share of total in-degree in each domain. Evidence mass and hub score, by contrast, are sharply concentrated. In the plotted runs, the top 5% of cities account for roughly three-quarters of total evidence mass and of total hub score across the four domains.

The proximity between the evidence-mass and hub-score curves suggests that urban climate hubness is driven at least as much by concentration in documented experience as by analogue centrality alone. Broad nominal access to analogue cities therefore should not be mistaken for a broadly distributed supplier ecology. The reachable stock remains anchored disproportionately in a limited set of

evidence-rich hubs.

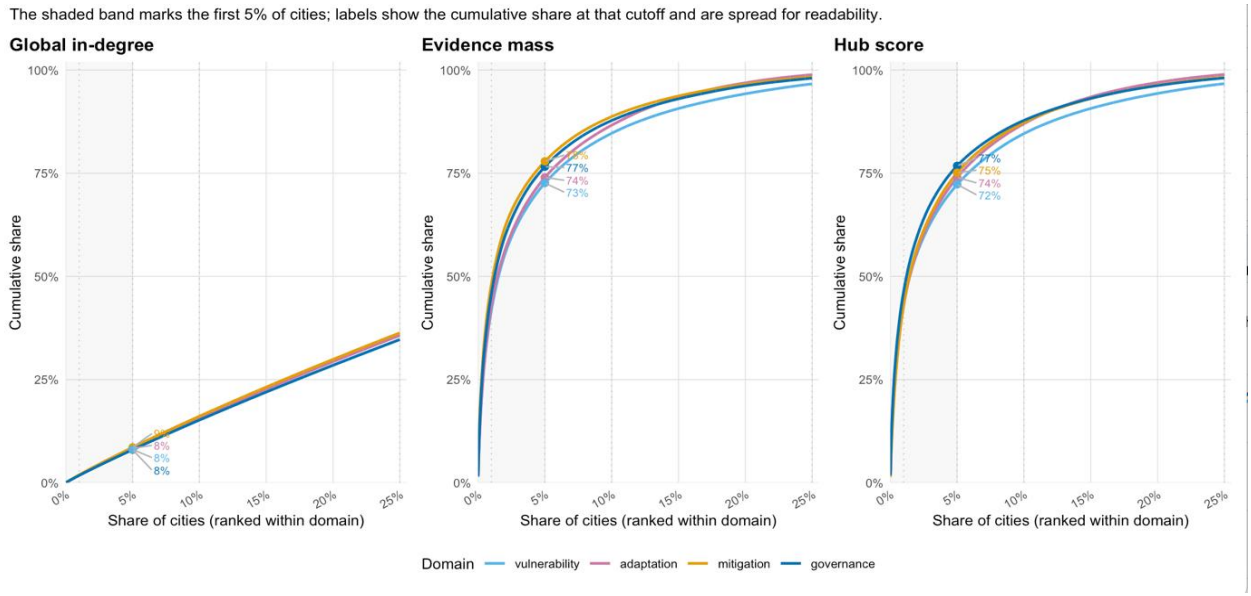


Figure 4. Concentration in the global urban climate learning network. Cities are ranked within each non-hazard domain by global in-degree, evidence mass, and hub score. Curves show the cumulative share accounted for by the top x% of cities, with the shaded band marking the top 5%.

3.5 The highest-priority agenda is broader than the frontier alone

The final step is to translate descriptive positions in the learning architecture into a practical agenda for new evidence generation. Figure 5 shows that the highest-priority set is not confined to a single learning condition. Frontier, replication-ready, and scaling-platform cells all appear in the ranked agenda, which means that priority is not simply another name for frontier status. Cells rise to the top when they combine population scale with deficits in direct evidence and reachable evidence.

The left panel clarifies why particular cells rank highly. Because the score is multiplicative, the weighted log decomposition is best read as a constraint profile rather than as additive shares. Among the top-ranked set, within-cell deficits are often among the strongest penalties, while direct-evidence deficits are consistently negative across most rows, indicating that thin direct documentation remains pervasive even among otherwise different learning positions. The right panel then translates those profiles back into an overall ranking and shows that several cells remain near the top under alternative weighting scenarios.

In the run shown here, Africa / Type 1 / mitigation, Africa / Type 2 / governance, Small Islands / Type 1 / mitigation, and Europe / Type 2 / adaptation appear among the most stable high-priority cells. The broader implication is that the agenda for new evidence generation should not be defined only by where cities currently sit at the frontier of learning, but by where scale, under-documentation, and weak borrowing opportunities coincide.

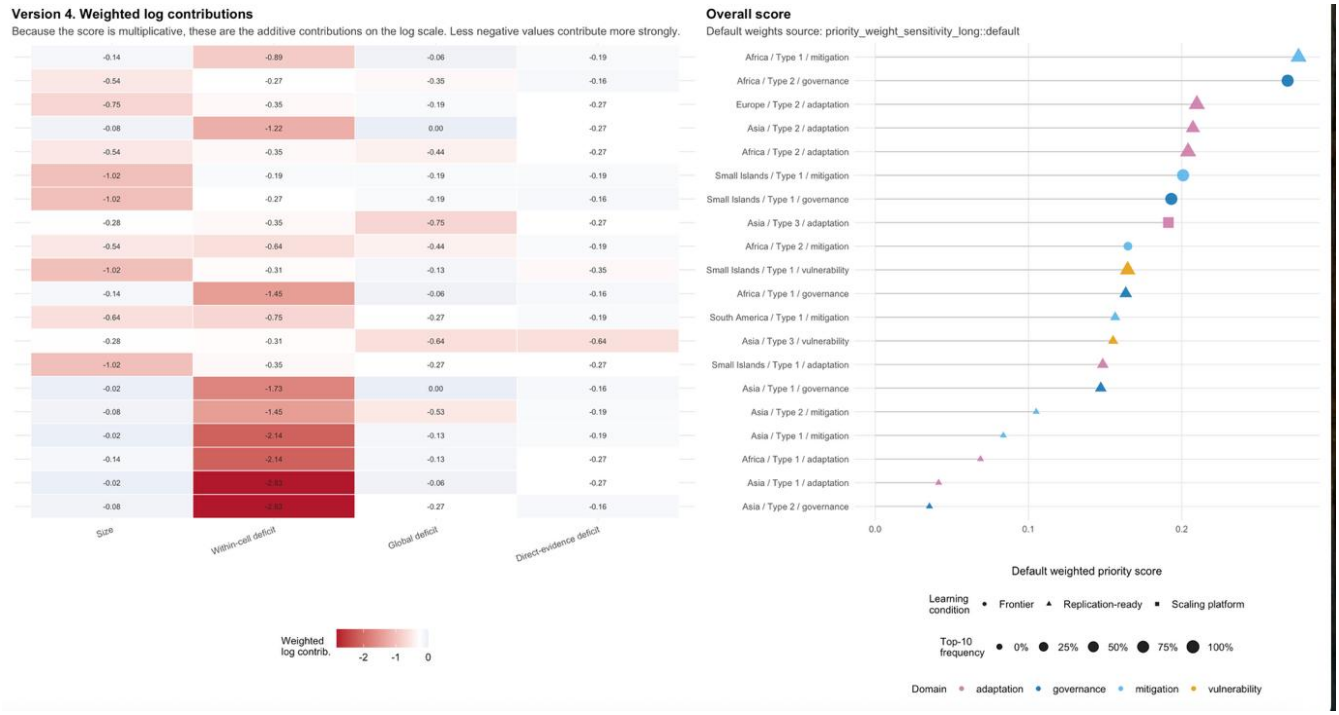


Figure 5. Priority agenda for new evidence generation. The left panel decomposes the default priority score into weighted log contributions from size and three deficits. The right panel shows the corresponding overall scores for the top-ranked cells, with symbols indicating learning condition and point size indicating top-10 frequency across weighting scenarios.

Taken together, the results depict an unequal but partially navigable architecture of urban climate learning. Direct evidence is thin and adaptation remains especially under-documented, yet cities do not face identical learning constraints. Some cells can borrow credibly from nearby analogues, some depend primarily on broader regional or global pools, and some sit in positions from which evidence circulates outward as well as inward. Because the supplier side of the system is highly concentrated, closing the most consequential gaps will require more than adding cases in general. It will require targeted evidence generation in the cells where direct documentation is weak, close analogues are scarce, and the potential value of additional knowledge is high.

4. Discussion

This paper began with a practical but theoretically consequential question: what can cities do when climate problems are locally salient but documented experience from comparable places is thin? We suggest that the answer is best framed not in terms of whether urban climate knowledge exists in the abstract, but in terms of how it is organised, where it accumulates, and how far it can travel under bounded conditions of similarity. Read in that way, the paper contributes to a growing call for urban research that is more cumulative, more explicit about replication and generalisability, and more honest about the conditions under which findings may travel across heterogeneous urban contexts (Sudmant et al., 2025; Robinson, 2016; McCann, 2011).

The central result is that urban climate knowledge is neither universally available nor wholly absent. It is structured as an unequal but partially navigable learning system: direct evidence is concentrated, accessible evidence is broader, borrowing is conditional rather than frictionless, and some deficits remain stubborn even after wider pools of analogue experience are taken into account.

That claim matters because the manuscript should not be read as a simple mapping exercise. What is being mapped is not only the distribution of studies, but the architecture through which documented

urban climate experience becomes legible, portable, and potentially reusable. In that sense, the paper speaks to knowledge systems, learning, and justice as much as to urban climate planning. The discussion therefore proceeds in six steps. It first frames the paper through replication and conditional transfer, then moves from evidence geography to learning geography, before interpreting learning trajectories, supplier hierarchy, justice, and the still-limited integration of domains across the urban climate field.

4.1 Replication, generalisability, and conditional transfer

The first implication is conceptual. This paper attempts to moves the discussion from simple unevenness in research coverage to a more precise question about cumulative learning. In the language proposed by Sudmant et al. (2025), the results can be read as operationalising conditions for replication-building and generalisability-building within the urban climate field. Direct evidence matters because without a minimum stock of documented local or near-local cases, even basic replication remains difficult. High-similarity reachable evidence matters because it provides the nearest equivalent, in this framework, to bounded replication across comparable cities. Broader regional and global pools matter because they indicate whether knowledge may travel beyond close peers, albeit with looser comparability constraints. This is not a universalist model of transfer, and it does not assume a single benchmark city from which others should learn. Rather, it identifies patterned spaces of comparability in a way that is closer to comparative urbanism's call to think cities through elsewhere, while retaining the policy-mobility insight that what travels is always translated, reassembled, and politically reworked (Robinson, 2016; McCann, 2011).

Importantly, the paper should not be read as claiming that similarity scores prove transferability. At most, similarity scores identify evidence conditions under which transfer claims become more or less credible. This distinction is important in a field where municipal climate policy learning remains under-theorised with respect to context, transferability, and evidence quality, and where the IPCC continues to stress that there is no one-size-fits-all urban pathway even when cross-city learning is possible (Neij and Heiskanen, 2021; Dodman et al., 2022). One virtue of the present framework is that it occupies a middle ground between naive universalism and radical particularism. It rejects the idea that any case from anywhere can be lifted into policy guidance without friction, but it also rejects the opposite view that contextual difference forecloses meaningful comparison in advance.

5.2 From evidence geography to learning geography

The second implication is empirical and methodological. Existing stocktakes have already shown that urban climate research is heavily concentrated in a relatively narrow set of large, wealthy, and internationally visible cities, while many vulnerable cities remain thinly represented in the published record (Montfort et al., 2025). The results here confirm that concentration, but they also extend the argument in a more consequential direction. Direct evidence geography and accessible evidence geography are not the same thing. Direct coverage remains sparse across many region × type cells and is especially thin in adaptation, yet reachable evidence expands markedly once cities are allowed to borrow from analogue neighbours and widen their learning pools beyond strict within-cell peers. The object of analysis therefore shifts. The relevant question is no longer only where studies are produced, but where usable analogue experience is reachable and under what loss of comparability that reach is achieved.

That shift from evidence geography to learning geography is especially important for adaptation. The results show that adaptation is consistently the thinnest non-hazard domain in direct coverage and remains comparatively thin even after learning pools widen. This is consistent with broader assessments of adaptation evidence, which suggest that documented adaptation is often fragmented,

localised, and incremental, with limited evidence of demonstrated risk reduction or transformational change (Berrang-Ford et al., 2021). This also helps explain why the paper should be careful not to equate a thicker evidence base with stronger outcomes. A city or cell may be highly visible in the literature without being exemplary in effectiveness, justice, or long-term resilience. Published evidence is a record of legibility and circulation, not a direct ranking of success.

There is also a historical shift here. Earlier work on urban adaptation planning asked whether cities were planning at all and found that many were not yet doing so in any substantive way (Araos et al., 2016). The field has since matured into a much larger and more heterogeneous body of case material, as reflected in major stocktakes and syntheses of urban climate solutions (Lamb et al., 2019; Montfort et al., 2025). This paper represents a second-generation move within that trajectory. Rather than counting the spread of activity alone, it asks what kinds of cumulative learning become possible once the published literature is treated as an evidence infrastructure rather than as a loose collection of case studies.

5.3 From learning conditions to learning trajectories

The learning-condition typology is one of the manuscript's most promising theoretical devices because it classifies evidence states, not just evidence volume. Read narrowly, the x-axis captures median reachable evidence among close, high-similarity peers and the y-axis captures median reachable evidence in the broader global pool. Read more interpretively, those axes describe evidence conditions that make different forms of cumulative learning more or less possible. Rightward movement is closest to replication-building because it reflects denser evidence among close analogues. Upward movement is closest to generalisability-building because it reflects broader evidence that remains usable across more heterogeneous contexts. Direct evidence remains a third and distinct task, which is why weak cells cannot simply borrow their way out of deficit. On this reading, the figure does more than sort cells into four boxes. It becomes a map of uneven research pathways.

Seen in that light, the off-diagonal quadrants are analytically productive. Borrowable cells are not simply middling cases. They are settings in which the literature appears portable before it is densely replicated among close peers. Those cells may therefore depend more heavily on translation, brokerage, and selective reconstruction, because what is available is broad reach rather than thick near-peer confirmation. Replication-ready cells show the opposite pattern: evidence has accumulated among close analogues before broader portability is established. Those cases may be especially valuable for consolidating bounded lessons and then testing how far they travel. Scaling platforms, by contrast, are the clearest candidates for outward supply, while frontier cells remain the strongest case for new primary evidencing. The paper's contribution is therefore not only to identify unevenness, but to show that replication-building and generalisability-building do not necessarily unfold together or in a single sequence.

The quality-volume trade-off in reachable evidence strengthens that reading. Much of the additional accessible evidence appears only once cities move beyond the strict analogue frontier, which means that upward movement in the state space is often purchased by loosening comparability constraints. That is why broader access should not be mistaken for frictionless generalisation. It indicates wider opportunity for learning, but also greater need for translation across different institutions, political economies, and delivery conditions. In practical terms, adaptation seems more likely to occupy frontier or borrowable positions, whereas hazards and, in some cells, mitigation more often look replication-ready or scaling-oriented. Even that pattern remains uneven across cells, which is exactly why the discussion should resist collapsing back into a single North–South or large-city story.

5.4 Scaling, supplier hierarchy, and the politics of urban climate knowledge

A fourth implication concerns scaling. The results suggest that the central widening move is not a simple leap from local to global, but a more intermediate shift from same-type peers to broader cross-type pools within regions. That is an important finding because it implies that region is doing substantive work in the circulation of urban climate knowledge. Shared institutional repertoires, donor ecologies, planning traditions, development trajectories, and political histories may make regional widening more productive than an immediate turn to global exemplars. This reading aligns with emerging calls for a more systematic science of scaling in urban climate governance, and it also supports work distinguishing knowledge transfer from knowledge exchange and co-creation (van der Heijden, 2023; Bakker et al., 2024). What the analysis maps most clearly is the geography of plausible transfer and some exchange; it does not, by itself, demonstrate collaborative co-production or durable learning in practice.

That is why institutional mediation has to stay explicit. Similarity scores do not make learning happen. Cities still require boundary-spanning organisations, open comparative frames, communicative infrastructures, monitoring, and sustained investment if evidence is to become actionable rather than merely available. Work on urban knowledge exchange institutions and learning pathways makes exactly this point: network participation or evidence availability alone does not guarantee that ideas are translated, embedded, or sustained (Dickey et al., 2022; Bellinson and Chu, 2019). The widening curves therefore reveal the structure of opportunity, not the social accomplishment of learning.

The same caution applies, in a different way, to supplier hierarchy and hubness. The results show that the supplier side of the network is concentrated much more strongly in evidence-rich hubs than in analogue centrality alone. Put differently, cities become influential suppliers not only because many others resemble them, but because their experiences have been repeatedly documented, circulated, and rendered comparable. That should be read as a form of epistemic concentration. Cities that are more deeply embedded in municipal networks and have stronger adaptation planning often also have greater fiscal and institutional capacity, stronger international visibility, and more repeated rounds of study, circulation, and citation (Heikkinen et al., 2020; Castán Broto and Bulkeley, 2013). The supplier core is therefore not a neutral bibliometric artefact. It is partly the outcome of how urban experimentation, networked visibility, and publication practices interact.

5.5 Frontier deficits, justice, and the priority agenda

The justice implications of this research follow directly from the concentration in the body of documented research we have documented. Accessible evidence is not automatically equitable evidence. A city may be able to borrow from elsewhere and still remain dependent on knowledge produced in settings with very different institutions, fiscal capacities, political settlements, or imaginaries of desirable urban futures. The priority results matter precisely because they show that the highest-value agenda is broader than the frontier alone. Cells rise to the top not simply because they are empty, but because they combine scale with persistent deficits in direct evidence and in usable learning pools. In the present ranking, stable priorities such as Africa / Type 1 / mitigation, Africa / Type 2 / governance, Small Islands / Type 1 / mitigation, and Europe / Type 2 / adaptation demonstrate that the most consequential deficits remain distributed away from the main hubs and cannot be resolved by borrowing alone.

That empirical pattern gives the justice argument much more traction than a generic call for inclusion. Justice-oriented urban adaptation research has long argued that low-capacity, rapidly changing, and politically marginalised urban contexts are often the least well served by mainstream evidence production (Shi et al., 2016). The findings here add a sharper epistemic dimension to that concern. When supplier cities dominate the documented record, they may also influence which problems,

solutions, and standards of feasibility appear authoritative. This is where epistemic justice becomes central. The issue is not only who is missing from the literature, but whose knowledge comes to count as portable, credible, and action-guiding in the first place (Mabon et al., 2022). The politics of learning therefore sit inside the evidence architecture itself.

This point is especially acute for adaptation, where the results show thinner direct and accessible evidence than for hazards or vulnerability. Borrowing under those conditions may transmit not only instruments and institutional scripts, but also dominant imaginaries about what adaptation should look like and for whom it is being designed. Work on just adaptation has stressed that urban responses are shaped by contestable visions of plausible and desirable futures, and that dominant imaginaries can reproduce exclusion or maladaptation when they travel uncritically (Castán Broto et al., 2024). The practical implication is not to reject cross-city learning. It is to treat frontier deficits and supplier concentration as reasons to diversify evidence production, widen who is recognised as a supplier, and scrutinise the justice risks that accompany supposedly portable models.

5.6 Weak vertical learning and the present limits of domain integration

A final implication concerns integration across domains. This subsection should remain slightly provisional until the final wording of the later figures is settled. Still, if the final results continue to show the weaker vertical-learning pattern described in the briefing notes, namely, that cross-domain borrowing is flatter and more row-invariant than the cross-type story, then that should be treated as a substantive finding rather than as a defect in the paper. It would suggest that the urban climate literature currently appears more connected across place-types than across thematic domains, which is consistent with broader diagnoses of fragmentation across urban research traditions (Wolman et al., 2024). Hazards, vulnerability, adaptation, mitigation, and governance may therefore still coexist more as adjacent literatures than as deeply integrated learning systems.

That interpretation would also align with existing evidence on urban climate planning. Integrated planning is certainly present, but empirical assessments suggest that the depth of integration between mitigation and adaptation remains moderate and uneven, with mitigation objectives often retaining greater prominence (Grafakos et al., 2020). IPCC assessments likewise emphasise the need for integrated, cross-sector, multi-scale action while also noting the governance and financial barriers that make such integration difficult in practice (Dodman et al., 2022). On that reading, a weak vertical-learning signal would be telling us something real about the current field: not that cross-domain learning is impossible, but that it is still comparatively under-developed, more supply-driven than target-specific, and therefore a priority area for future evidence-building.

The paper's limitations and future agenda follow directly from these findings. First, the analysis maps published and therefore legible evidence; it does not capture the full range of local, professional, tacit, or community-based knowledge that shape urban climate action. That distinction matters normatively as well as analytically, because what becomes visible in journals and assessments is already filtered by language, funding, institutional access, and publication conventions. Second, the governance domain should be read modestly. In the present operationalisation it captures documented evidence on governance-relevant conditions and enabling contexts more than governance quality or capability in any comprehensive sense. Third, the next generation of work could push beyond reachability alone by testing when analogue access actually leads to successful uptake, translation, or co-produced learning and when it does not. That would bring this architecture of learning into closer dialogue with the kinds of transferable yet place-sensitive public policy frameworks now being proposed for cities and human settlements (Creutzig et al., 2024).

Taken together, the discussion suggests that urban climate learning should not be understood as a flat global marketplace of best practice. It is a structured and directional knowledge system in which some cities and cells are far more visible as suppliers than others, in which access broadens faster than

direct documentation, and in which the largest gains often come from intermediate widening rather than immediate global abstraction. The practical consequence is twofold: widen access where bounded borrowing is plausible, and invest in new primary evidence where both direct documentation and reachable analogues remain thin. The theoretical consequences is equally important. What this paper offers is not simply another account of uneven urban research, but a way of thinking about how documented climate knowledge accumulates, travels, and fails to travel across the contemporary urban world.

5. Conclusion

This paper asked a simple but consequential question: who can learn from whom in the urban climate field? The answer cannot be read from the map of published case studies alone. The documented record is highly concentrated, and direct evidence remains thin across many region × type × domain intersections, with adaptation especially under-documented. Yet the picture changes once urban climate research is treated as an architecture of learning rather than as an inventory of cases. Although 45% of cities have no direct evidence, 95.5% of zero-direct cities can still reach analogue experience under strict high-similarity conditions, and all can do so once the learning pool widens. The largest distributive gains come less from an immediate move to global exemplars than from widening from same-type peers to broader cross-type pools within regions. Urban climate learning is therefore neither universal nor absent. It is structured, unequal, and only partly navigable.

The paper's main contribution lies in shifting the object of analysis from evidence geography to learning geography. Conceptually, it reframes published urban climate research as an evidence infrastructure through which replication, borrowing, and bounded generalisation become more or less possible. Methodologically, it combines a city-by-domain evidence matrix with domain-specific similarity spaces to distinguish direct evidence from accessible evidence and to locate frontier, borrowable, replication-ready, and scaling-platform positions. Empirically, it shows that the supplier side of this system is strongly directional: a narrow set of evidence-rich hubs exerts disproportionate influence over the stock of urban climate experience that becomes portable and legible. This matters for debates on comparative urbanism and policy mobility because it suggests that urban learning is not a matter of transporting universal best practice, but of translating knowledge across bounded spaces of similarity, institutional difference, and political contestation (Robinson, 2016; McCann, 2011).

Those findings carry two practical implications. One is enabling. Cities and intermediaries should not read the absence of direct local evidence as proof that no usable experience exists; in many cases, relevant analogue evidence is available, especially through wider regional pools. The other is cautionary. Borrowing cannot substitute for new primary evidence where both direct documentation and close analogues remain thin, nor does accessibility guarantee effectiveness or justice. Because the supplier core is concentrated, accessible evidence is not automatically equitable evidence. The cases that circulate most easily through planning, assessment, and finance also help define which urban experiences count as credible, portable, and worth emulating. A more useful urban climate knowledge system would therefore do more than expand the volume of case studies in general. It would widen the supplier base, strengthen evidence in persistently under-documented cells, and direct new research towards the intersections where population scale, thin documentation, and weak reachability coincide.

The paper also sets clear limits for future work. It maps published and therefore legible evidence, not the full range of local, professional, and community-based knowledges that shape urban climate action. It identifies conditions for learning rather than policy quality or effectiveness, and the governance domain should be read as capturing governance-relevant contexts more than governance in any exhaustive institutional sense. The next step is therefore to examine when analogue access actually becomes uptake, adaptation, contestation, or durable change in practice. That agenda speaks

directly to the wider challenge of building climate knowledge systems that are cumulative without becoming universalising, and comparative without flattening difference (Dodman et al., 2022; Sudmant et al., 2025). The broader implication is modest but important: urban climate action does not proceed in a flat world of interchangeable examples. It proceeds through an uneven architecture of learning, and changing that architecture may matter as much as adding one more celebrated case.

References

Araos, M., Berrang-Ford, L., Ford, J. D., Austin, S. E., Biesbroek, R., & Lesnikowski, A. (2016). Climate change adaptation planning in large cities: A systematic global assessment. *Environmental Science & Policy, 66*, 375–382. https://doi.org/10.1016/j.envsci.2016.06.009

Bakker, J., Scholten, P., Fransen, J., & Minkman, E. (2024). A typology of urban knowledge sharing: From a systematic literature review to an integrated model. *Science and Public Policy, 51(4)*, 707–720. https://doi.org/10.1093/scipol/scae008

Bellinson, R., & Chu, E. K. (2019). Learning pathways and the governance of innovations in urban climate change resilience and adaptation. *Journal of Environmental Policy & Planning, 21(1)*, 76–89. https://doi.org/10.1080/1523908X.2018.1493916

Berrang-Ford, L., Siders, A. R., Lesnikowski, A., Fischer, A. P., Callaghan, M. W., Haddaway, N. R., Mach, K. J., Araos, M., Shah, M. A. R., Wannewitz, M., Doshi, D., Leiter, T., Matavel, C., Musah-Surugu, J. I., Wong-Parodi, G., Antwi-Agyei, P., Ajibade, I., Chauhan, N., Kakenmaster, W., ... Abu, T. Z. (2021). A systematic global stocktake of evidence on human adaptation to climate change. *Nature Climate Change, 11(11)*, 989–1000. https://doi.org/10.1038/s41558-021-01170-y

Castán Broto, V., & Bulkeley, H. (2013). A survey of urban climate change experiments in 100 cities. *Global Environmental Change, 23(1)*, 92–102. https://doi.org/10.1016/j.gloenvcha.2012.07.005

Castán Broto, V., Olazabal, M., & Ziervogel, G. (2024). Disrupting the imaginaries of urban action to deliver just adaptation. *Buildings and Cities, 5(1)*, 199–214. https://doi.org/10.5334/bc.456

Creutzig, F., Becker, S., Berrill, P., Bongs, C., Bussler, A., Cave, B., Constantino, S. M., Grant, M., Heeren, N., Heinen, E., Hintz, M. J., Ingen-Housz, T., Johnson, E., Kolleck, N., Liotta, C., Lorek, S., Mattioli, G., Niamir, L., McPhearson, T., ... Zekar, A. (2024). Towards a public policy of cities and human settlements in the 21st century. *npj Urban Sustainability, 4(1)*, Article 29. https://doi.org/10.1038/s42949-024-00168-7

Dickey, A., Kosovac, A., Fastenrath, S., Acuto, M., & Gleeson, B. (2022). Fragmentation and urban knowledge: An analysis of urban knowledge exchange institutions. *Cities, 131*, Article 103917. https://doi.org/10.1016/j.cities.2022.103917

Dodman, D., Hayward, B., Pelling, M., Castán Broto, V., Chow, W. T. L., Chu, E., Dawson, R., Khirfan, L., McPhearson, T., Prakash, A., Zheng, Y., & Ziervogel, G. (2022). Cities, settlements and key infrastructure. In H.-O. Pörtner, D. C. Roberts, M. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, & B. Rama (Eds.), *Climate change 2022: Impacts, adaptation and vulnerability: Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 907–1040). Cambridge University Press. https://doi.org/10.1017/9781009325844.008

Grafakos, S., Viero, G., Reckien, D., Trigg, K., Viguié, V., Sudmant, A., Graves, C., Foley, A., Heidrich, O., Mirailles, J. M., Carter, J., Chang, L. H., Nador, C., Liseri, M., Chelleri, L., Orru, H., Orru, K., Aelenei, R., Bilska, A., ... Dawson, R. (2020). Integration of mitigation and adaptation in urban climate change action plans in Europe: A systematic assessment. *Renewable and Sustainable Energy Reviews, 121*, Article 109623. https://doi.org/10.1016/j.rser.2019.109623

Heikkinen, M., Karimo, A., Klein, J., Juhola, S., & Ylä-Anttila, T. (2020). Transnational municipal networks and climate change adaptation: A study of 377 cities. *Journal of Cleaner Production, 257*, Article 120474. https://doi.org/10.1016/j.jclepro.2020.120474

Lamb, W. F., Creutzig, F., Callaghan, M. W., & Minx, J. C. (2019). Learning about urban climate solutions from case studies. *Nature Climate Change, 9(4)*, 279–287. https://doi.org/10.1038/s41558-019-0440-x

Mabon, L., Barkved, L., de Bruin, K., & Shih, W.-Y. (2022). Whose knowledge counts in nature-based solutions? Understanding epistemic justice for nature-based solutions through a multi-city comparison across Europe and Asia. *Environmental Science & Policy, 136*, 652–664. https://doi.org/10.1016/j.envsci.2022.07.025

McCann, E. J. (2011). Urban policy mobilities and global circuits of knowledge: Toward a research agenda. *Annals of the Association of American Geographers, 101(1)*, 107–130. https://doi.org/10.1080/00045608.2010.520219

Montfort, S., Callaghan, M., Creutzig, F., Lamb, W. F., Lu, C., Repke, T., Ge, K., & Minx, J. C. (2025). Systematic global stocktake of over 50,000 urban climate change studies. *Nature Cities, 2*, 613–625. https://doi.org/10.1038/s44284-025-00260-8

Neij, L., & Heiskanen, E. (2021). Municipal climate mitigation policy and policy learning - A review. *Journal of Cleaner Production, 317*, Article 128348. https://doi.org/10.1016/j.jclepro.2021.128348

Robinson, J. (2016). Thinking cities through elsewhere: Comparative tactics for a more global urban studies. *Progress in Human Geography, 40(1)*, 3–29. https://doi.org/10.1177/0309132515598025

Shi, L., Chu, E. K., Anguelovski, I., Aylett, A., Debats, J., Goh, K., Schenk, T., Seto, K. C., Dodman, D., Roberts, D., Roberts, J. T., & VanDeveer, S. D. (2016). Roadmap towards justice in urban climate adaptation research. *Nature Climate Change, 6(2)*, 131–137. https://doi.org/10.1038/nclimate2841

Sudmant, A., Creutzig, F., & Mi, Z. (2025). Replicate and generalize to make urban research coherent. *International Journal of Urban Sciences, 29(3)*, 582–603. https://doi.org/10.1080/12265934.2024.2382706

van der Heijden, J. (2023). Towards a science of scaling for urban climate action and governance. *European Journal of Risk Regulation, 14(3)*, 513–525. https://doi.org/10.1017/err.2022.13

Wolman, H., Barnes, W., Clark, J., Friedman, S., Harris, R., Lin, J., & Ogorzalek, T. (2024). The state of urban research: Views across the disciplines. *Journal of Urban Affairs, 46(3)*, 425–462. https://doi.org/10.1080/07352166.2022.2080073