AI and Sustainability: From Mining the Past to Inventing the Future

Co-Creating the Next Digital Tools with Cities and Regions

Co-Authored by

ClimateView · ICLEI Local Governments for Sustainability · Global Covenant of Mayors for Climate & Energy · Scottish Climate Intelligence Service · UNU EHS · technische universität dortmund · sfs

With Case Studies from

高雄市政府 Kaohsiung City Government · City of Dortmund

How to Cite

Shalit, T., Dixon, M., Corbett, D., Salehi, P., Astolfo, J. P., Emer, D. E., Wharmby, C., Alshimmari, A., Rodriguez, A., Colodro Gotthelf, U., Sandholz, S., Hölsgens, H., Fassbender, L., Schultze, J., & Sadlowski, C. (2026). *AI and sustainability: From mining the past to inventing the future*. ClimateView AB. https://events.climateview.global/ai-whitepaper

Launched at Innovate4Cities 2026

Elements Stockholm
Jakobsbergsgatan 24,
111 44 Stockholm,
Sweden

ClimateView.global

Index

Introduction

Cities worldwide have set ambitious climate goals. However, the gap between these goals and the actual progress remains a major challenge of the climate transition [1], [2]. Bridging this gap requires not just better analysis of what has already been done, but also new ideas about what needs to happen next.

Standards and Reporting as the Foundation for AI

The potential of AI in climate governance depends upstream on the quality, consistency, and comparability of the data available. Without common standards for measuring and reporting urban emissions and climate action, the knowledge base that AI systems analyze remains fragmented, inconsistent, and difficult to compare across cities and geographies.

Initiatives such as the WWF One Planet City Challenge, which draws on data reported through the CDP-ICLEI Track, and the Global Covenant of Mayors Common Reporting Framework have spent years building exactly this foundation: shared methodologies, common metrics, and structured databases of city-level climate action. Thousands of cities now report through these frameworks annually, creating a growing and increasingly comparable record of what cities are doing, where progress is being made, and where gaps remain.

The One Planet City Challenge (OPCC) exemplifies how voluntary city disclosure can be operationalized into a robust data infrastructure. Redefined by WWF in 2017 to align with the Paris Agreement, the OPCC has evolved into a global initiative engaging 389 cities across 51 countries in its most recent cycle, growing more than fivefold from 64 cities in 2013. By leveraging the CDP-ICLEI Track unified reporting system, where, in 2025 alone, 507 cities disclosed 2,164 climate projects seeking a record US$105 billion in finance, the initiative ensures that information on emissions, risks, and action plans is independently evaluated and validated. Through tools such as its Strategic Feedback Report (SFR), as well as technical assistance and one-to-one training with cities, the OPCC translates these standardized inputs into structured, context-specific guidance, creating high-quality, verifiable datasets that support robust analysis and informed decision-making in urban climate governance.

This infrastructure matters for AI in two ways. First, it provides the structured, standardized data that AI systems need to identify patterns, surface relevant precedents, and generate credible suggestions. Second, it helps ensure that AI outputs can be traced back to real-world evidence rather than relying on unstructured, inconsistent, or unverifiable sources. The reporting work that cities, networks, and organizations have invested in over decades is a precondition for leveraging it effectively.

Artificial Intelligence (AI) is quickly becoming a part of sustainability and climate governance, and most current uses focus on organizing and analyzing existing knowledge: scanning climate plans, classifying policy measures, and comparing interventions across different areas. This initial wave of AI for climate can be useful but has fundamental limitations. A general concern from practitioners is emerging:

> ***"What if we are searching for gold but only find sludge?"***

Although a valid approach, mining the past is limited in its ability to develop the interventions cities need for the next phase of their transition. Existing interventions may offer valuable insights for future action, but this approach alone is not sufficient. It must be complemented by guided creativity that leverages AI, human expertise, modeling, and context-specific translation into action to drive innovation that accelerates progress.

This white paper advocates a different paradigm in which AI not only retrieves and organizes existing knowledge but also actively helps cities create new climate action pathways. This role for AI is only possible if humans are part of the process, not just in the loop. Practitioners who understand local context, stakeholder dynamics, and political feasibility must be part of both the design and the use of AI-enabled tools to support human judgment rather than bypass it.

Realizing this potential requires structure. ClimateView's AI Intervention Library is one example of what this can look like: generative AI grounded in the Transition Element Framework and causal Outcome Logic (a German national standard for impact assessment under DIN SPEC 91637) that can propose new intervention strategies, map their impact pathways, and connect them to real-world evidence where it exists - but is not limited to what has already been documented. A city planner wanting to increase cycling and walking, for instance, can move from that outcome to a structured set of interventions: familiar options like cycling infrastructure and public transport subsidies, but also newly generated approaches such as parking policy reforms, social norm campaigns, or employer incentive schemes that may not yet appear in any existing climate action plan. This prototype points toward something larger: a standard for how AI tools in cities should be built, grounding generative capability in open causal frameworks rather than only recycling the past. Instead of replacing current climate plans, tools built on this model are designed to activate them, transforming static lists of actions into dynamic portfolios.

The prototype has been tested through collaborative workshops in Dortmund and Kaohsiung; the next version will be developed using insights from city officials, researchers, private-sector representatives, and civil society actors from these workshops. These sessions confirmed the potential of AI-supported climate planning while identifying a clear set of requirements for the next generation of tools: improved navigation and contextualization, actor-specific perspectives, more comprehensive decision-support information, and the ability to understand how interventions combine and interact.

Used responsibly and built in genuine partnership with practitioners, AI can become a new form of shared transition intelligence. This living knowledge infrastructure helps cities shift from static plans to dynamic, fundable climate action. Every city that contributes makes the next city smarter.

The Limits of Retrospective AI

Cities, regions, and national governments are creating climate plans, policy documents, and implementation reports at an unprecedented pace. AI has allowed us to understand this expanding body of knowledge much faster than manual analysis ever could.

The most advanced current applications focus on extracting and organizing insights from existing documents and datasets. AI systems are now being used to scan thousands of climate policies and city climate action plans to identify recurring interventions, classify policy measures by sector, target, or implementation mechanism, analyze trends across jurisdictions, and identify commonly adopted approaches [3]. They also structure large volumes of climate-related information into searchable databases that support benchmarking and comparison.

These capabilities are important. Climate governance has long faced challenges with fragmented knowledge spread across reports, plans, and institutional silos. AI is already helping turn this scattered information into accessible intelligence, allowing policymakers to see, for example, how many cities have adopted congestion pricing or building retrofit programs, and how these policies are being implemented in different settings[1].

But this first generation of AI applications remains essentially retrospective. The main workflow is: documents → extraction → classification → comparison. It answers the question of what has already been tried, but doesn’t effectively address what should be tried next.

This approach has significant limitations, including hallucinations, issues with accuracy, and inability to read tables and diagrams. As such, a new approach that focuses on the strength of the AI, which is parsing together information from different sources to find optimal solutions for the unique problems in the area to create an ‘intervention’, is a stronger and better application of AI as a tool. Transformational climate action often requires interventions that aren't yet documented in any plan or policy database. Searching existing documents reveals familiar, sometimes outdated solutions and reinforces common patterns, precisely when cities need to think beyond them. Relying only on retrospective AI risks solidifying yesterday's thinking at the scale of tomorrow's ambitions.

The infrastructure for organizing climate knowledge is already substantial and is developing at speed. The IEA's Policies and Measures Database[2], for example, catalogs over 5,000 policy records across 85 countries, searchable by sector, type, and geography, and represents one of the most comprehensive repositories of

[1] https://openkfw.github.io/d4dtools/docs/climate-resilience/
[2] https://www.iea.org/policies/about

documented climate action in existence. Although valuable, Rosenzweig et al. [4] highlight one challenge with this type of database to be that existing case study sources are fragmented, biased towards large cities in the Global North, and largely inaccessible to practitioners in the regions that need them most, meaning the knowledge base that retrospective AI draws on systematically underrepresents the urban contexts where climate action is most urgent. AI is now being layered on top of knowledge bases like this one: tools such as ChatClimate[3] and ClimateGPT[4], for example, apply large language models to climate science literature, enabling practitioners to query the IPCC's Sixth Assessment Report and a growing corpus of climate research in natural language, with referenced answers returned in seconds [5]. Meanwhile, a growing number of city-facing transition planning platforms are integrating AI into their workflows, moving beyond data management and reporting support to analyze climate risk, generate insights, and suggest relevant interventions from existing catalogs. Taken together, these developments represent real and meaningful progress in making climate knowledge more accessible and actionable. What remains missing, however, is the combination of a structured framework and a causal logic that together allow AI to reason about how interventions actually produce change, rather than simply retrieve examples of where they have been tried. Without both, even the most sophisticated AI tools are constrained to recombining what already exists. With them, the possibilities are fundamentally different.

The next progression is to move beyond just retrieving and organizing past interventions toward creating new ones. This requires a different approach: merging the generative power of modern AI with structured frameworks that allow it to reason about how systems change, causal chains, and which combinations of actions are necessary to achieve real-world results.

[3] https://www.chatclimate.ai/
[4] https://climategpt.ai/

A New Approach

Generative AI models are powerful pattern-recognition systems. However, without a structured representation of how transitions occur, they tend to produce fragmented, superficial, or implausible suggestions. To support an effective climate strategy, AI must be able to reason about how interventions generate systemic change: how policies and actions create enabling conditions, how those conditions influence behavior, how behavior changes activities, and how activity shifts ultimately lead to measurable outcomes such as emissions reductions.

This requires both a causal model of how transitions unfold and a structured representation of the available transition pathways. In essence, for AI to be leveraged to its full potential, both an ontology (Outcome Logic), which establishes the relationships between key transition concepts, and a taxonomy (Transition Element Framework), which classifies those concepts into meaningful categories, are essential.

Ontology: Outcome Logic

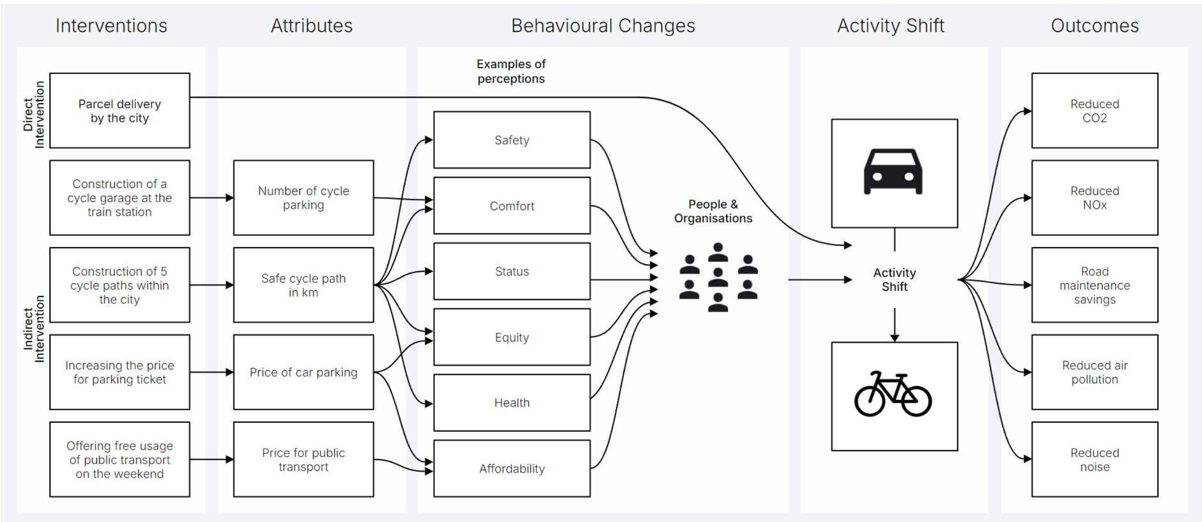

The Outcome Logic provides the causal ontology for climate transitions. It defines the fundamental relationships between interventions, enabling conditions, behavior, activity shifts, and outcomes. Within this framework, activity shifts are the primary mechanism through which interventions ultimately produce measurable impacts.

This distinction is important because interventions do not directly reduce emissions. Instead, they create the conditions that influence the decisions and behaviors of people and organizations [6]. Those behavioral changes alter activities, and those activity shifts generate outcomes.

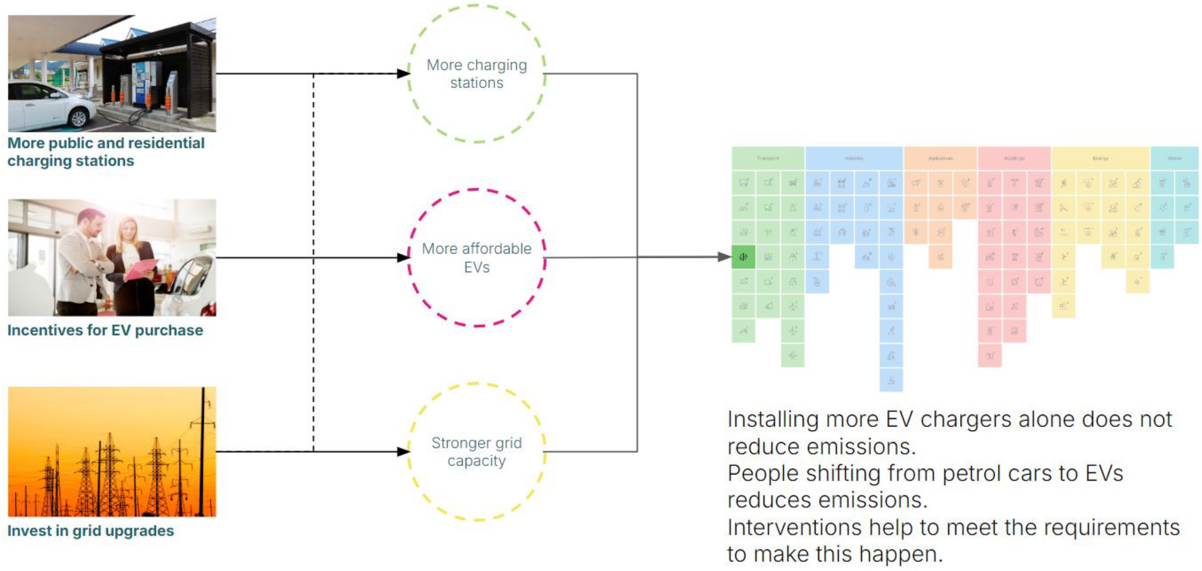

Consider the transition from petrol vehicles to electric vehicles. Installing charging infrastructure alone does not reduce transport emissions. Emissions decrease when people switch from petrol vehicles to electric vehicles. That behavioral shift depends on multiple enabling conditions, including charging availability, vehicle affordability, grid capacity, supportive policy, and broader social acceptance. Multiple interventions may therefore be required to achieve a single activity shift.

This causal structure has been formalized as Outcome Logic and codified through a DIN specification[5]. It provides a shared representation of how climate transitions generate outcomes, independent of any particular mitigation pathway or technology choice.

[5] https://www.dinmedia.de/en/technical-rule/din-spec-91637/393817040

Taxonomy: The Transition Element Framework

While Outcome Logic defines the structure through which transitions occur, it does not specify the concrete transitions available to cities. The Transition Element Framework (TEF) builds upon this ontology by providing a taxonomy of climate transition pathways.

The TEF defines and classifies specific activity shifts and mitigation options. Each transition element represents a distinct pathway through which activities can change and is linked to classifications and evidence from IPCC mitigation literature. Rather than organizing climate action solely through sectors such as energy, buildings, or transport, the framework organizes mitigation options around the underlying activity shifts and system transformations required to achieve climate goals.

In this sense, Outcome Logic defines the causal structure of transitions, while the TEF provides a systematic classification of the transition options that exist within that structure. The TEF is grounded in IPCC mitigation science and published as open source, providing a transparent and reusable foundation for climate planning and analysis.

Together, Outcome Logic and the TEF allow interventions to be understood not as isolated actions, but as components within a broader transition system. This structure is particularly valuable for AI because it enables reasoning about how interventions create enabling conditions, how those conditions contribute to shifts in activity, and how combinations of interventions interact to produce systemic outcomes.

A well-prompted general-purpose AI can generate a plausible list of climate actions. However, without a causal ontology and a structured taxonomy of transition options, it cannot reliably reason about which enabling conditions are missing, how interventions may need to be sequenced, or which combinations of measures are mutually reinforcing rather than redundant. By combining Outcome Logic and the TEF, AI gains a framework for reasoning about how interventions interact to produce systemic outcomes and real-world climate impact.

Why Structure Enables Inventive AI

Together, the TEF and Outcome Logic create the structured environment that generative AI needs to develop meaningful climate strategies. Within this framework, AI can generate new intervention ideas aligned with specific transition elements; reverse-engineer desired outcomes to find out which enabling conditions are missing; combine policies, infrastructure investments, and behavioral incentives into cohesive strategies; and explore different pathways toward the same climate goal.

Instead of just retrieving policies from existing plans, AI used in this way can reason about how systems change within a specified context. This shifts it from a tool for organizing knowledge to one for creating climate action pathways.

Without grounding in causal frameworks, generative AI operating in the climate domain carries significant risks. Richards and Worden [7] document how GenAI tools, capable of producing content indistinguishable from expert analysis, are already being deployed by actors across the political spectrum to influence climate decisions at every scale, from individual behavior to international diplomacy. The same generative capacity that allows AI to propose novel interventions can equally produce plausible-sounding but causally unfounded recommendations, amplify misinformation about specific climate measures, or manufacture synthetic consensus where none exists.

When AI suggestions cannot be traced to a causal logic, showing *why* an intervention would produce a given outcome through which mechanisms, practitioners have no basis for distinguishing a well-reasoned strategy from a confident hallucination. The shift from correlational to causal AI is a prerequisite for trustworthy decision-support. Systems that can only identify patterns in past data will systematically mislead when applied to novel or contested situations. In the context of climate governance, this risk is compounded by the political stakes involved.

Manivannan et al. [8] show that even well-intentioned AI policy tools can generate recommendations that are technically sound but publicly unacceptable, undermining implementation before it begins, because the models lack a structured representation of the social and institutional conditions that determine feasibility. The TEF and Outcome Logic address this directly: by requiring every generated intervention to be anchored to a specific activity shift, a defined set of enabling conditions, and an evidence-grounded causal chain, the framework makes AI outputs auditable and allows practitioners to interrogate the reasoning and the outcome.

The AI Intervention Library: From Framework to Practice

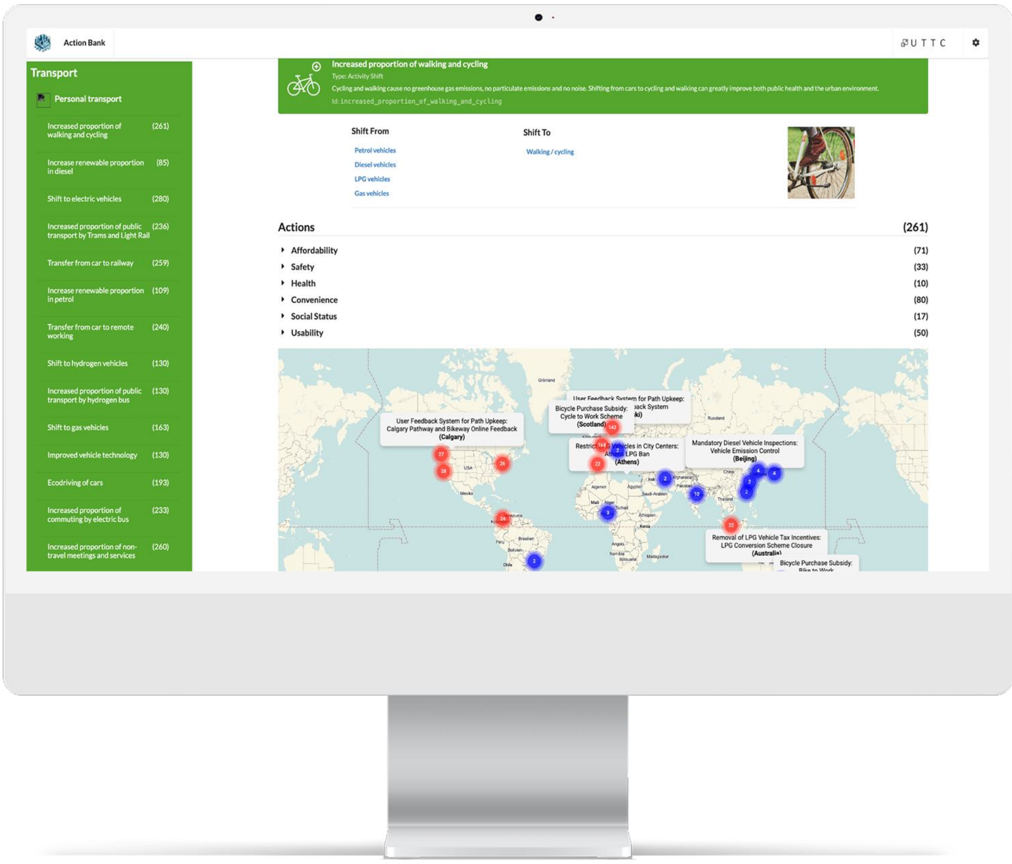

ClimateView's AI Intervention Library prototype was created to test this method in practice. The system integrates three main components.

First, all interventions are mapped using the Transition Element Framework and causal Outcome Logic, ensuring that AI understands not only what an intervention is but also how it contributes to systemic change. Second, within this structured environment, generative AI capabilities enable the system to propose new intervention ideas aligned with climate goals, combine interventions into coherent strategies, reverse-engineer outcomes to identify missing actions, and connect newly generated ideas to existing real-world examples. Third, the prototype includes hundreds of documented climate interventions, each linked to multiple real-world implementations across cities worldwide. These examples provide credibility, context, and practical grounding.

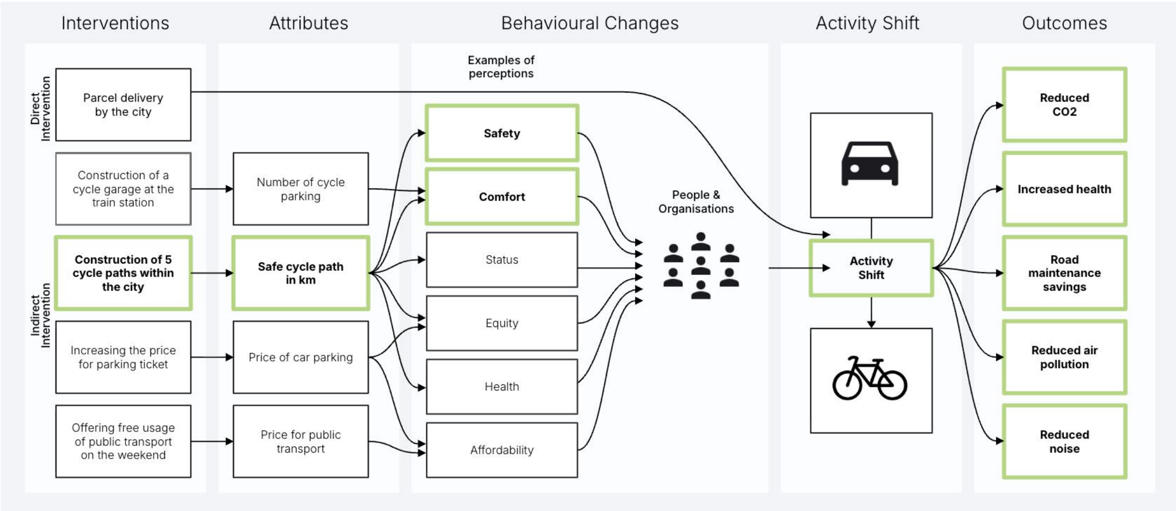

To understand how this works in practice, consider a city planner working on transport decarbonization. Instead of starting with a blank slate or digging through unorganized policy databases, they begin by identifying a desired activity shift: increasing the share of trips made by cycling and walking. Working backward through the Outcome Logic, the system identifies the enabling conditions that would need to change to make that shift happen, in this case, perceptions of safety, comfort, and affordability among the people and organizations whose behavior needs to shift. From those enabling conditions, it identifies the interventions that could create them: not only familiar measures such as building cycle infrastructure or offering free weekend public transport, but also less obvious approaches such as raising parking prices or social norm campaigns targeting perceptions of status and equity. As the image above illustrates, these interventions work not in isolation but in combination, each targeting a different part of the causal chain. Each suggestion is linked to causal reasoning explaining why it could be effective, and where real-world implementations exist, those are surfaced too. Where no documented examples exist yet, the system still generates the intervention, grounded in the causal logic rather than constrained by what has already been tried.

The result is a prototype for a creative tool for climate planning that helps cities discover, test, and refine action pathways that might not surface in traditional planning methods. Instead of replacing existing climate plans, the AI Intervention Library aims to complement them by enabling a portfolio approach and suggesting ways to fill gaps where existing interventions are insufficient to create the enabling conditions for the desired change.

The Potential of this Approach for Unlocking Finance

Structured demand plays a significant role in unlocking finance for sustainability projects. Cities frequently lack the capacity to develop bankable projects and investment pipelines that meet the requirements of funders, suppliers, and procurement processes [9], [10]. Structured demand, a clear picture of what is needed, why, and how it connects to measurable outcomes, is a prerequisite for accessing finance that many cities have yet to build.

By working backward from a desired activity shift to identify missing enabling conditions and the interventions that address them, the AI Intervention Library helps cities build an investment case grounded in causal logic rather than a list of actions [11], [12]. This matters because individual interventions rarely create the conditions for systemic change on their own. Research on urban decarbonization consistently finds that coordinated, cross-sectoral portfolios outperform isolated measures in both impact and investability [13]. The stronger case for finance lies in showing how a portfolio of mutually reinforcing measures works together, and why that combination will produce the intended outcome.

When multiple cities work within the same framework, a further opportunity emerges. Recurring needs across cities can be aggregated into demand signals large enough to support joint procurement, supplier matchmaking, and national funding programs that city-by-city planning struggles to unlock [10]. A next step for the library is to connect AI-suggested portfolios directly to relevant funding sources and financing mechanisms, helping cities move from structuring their demand to securing the resources to act on it.

Co-Creating AI Solutions with Cities

For AI tools to be truly useful in climate governance, they need to be developed with practitioners, not only for them [14]. Cities operate within complex environments shaped by political constraints, local economies, existing infrastructure, institutional cultures, and social dynamics that no external tool can fully capture. AI systems that overlook this reality risk generating recommendations that are technically sound but ultimately ineffective.

The AI Intervention Library was developed through an iterative co-creation process involving direct collaboration with cities, networks, and researchers. Workshops enable stakeholders to interact with the tool, explore intervention ideas, and provide structured feedback on what is helpful, what is missing, and what could be misleading. These sessions serve two main purposes: they offer immediate value to participants by supporting creative policy exploration, and they generate practitioner insights that inform the development of the next version of the tool.

Two workshops demonstrate this approach, one with stakeholders in Dortmund and another with international city representatives in Kaohsiung.

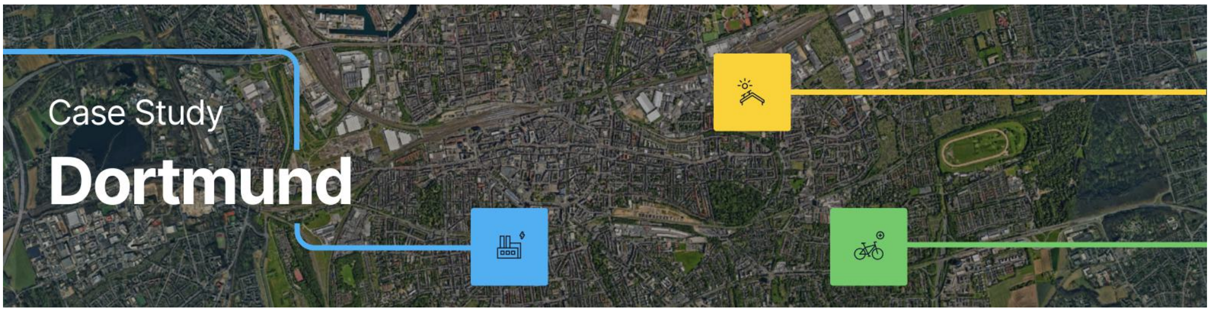

Workshop Setting

The workshop in Dortmund, held on March 4, 2026, brought together city officials, researchers from the Social Research Center at TU Dortmund University, private-sector representatives, and civil society actors. The session took place within the context of Dortmund's participation in the EU Net Zero Cities mission, as one of 112 European cities working toward climate neutrality by 2035, and the workshop was organized as part of the Climate City Dashboard 2.0 project.[6]

Participants were briefly introduced to the Transition Element Framework ClimateView platform and the AI Intervention Library prototype. They then participated in structured interactive exercises. Using the Transition Element Framework, groups identified focus transition areas, examined the enabling conditions needed for those transitions, generated new intervention ideas based on the AI Intervention Library's suggestions, and mapped the actor landscape necessary for implementation.

[6] https://netzerocities.eu/germanys-pilot-activity-climate-city-dash-2-0

Outcomes and Insights

The workshop produced five categories of insight that together establish a clear agenda for the next generation of the tool.

Navigation and usability at scale. Participants responded positively to the library's wide range of interventions, seeing it as a real catalyst for new ideas. However, the volume also posed a usability challenge. Without effective filtering and search options, the library's richness can become overwhelming rather than helpful. Participants requested a multi-layer filter system that allows navigation by sector, actor, impact, cost, and timeframe, as well as a keyword search to help users quickly find relevant areas. The library's value as a discovery tool relies on the ability to navigate it purposefully.

Context matching and local relevance. A recurring theme was the gap between what the library can provide and what a specific city can actually implement. Participants emphasized that the tool currently lacks a mechanism to assess how well a particular intervention aligns with a city's policy environment, infrastructure, and economic conditions. A feasibility assessment system, possibly using a red, amber, green (RAG) or traffic-light rating system to show high suitability, conditional fit, or poor alignment, and including indicative timelines, basic requirements, and an assessment of the biggest levers, was suggested to address this. This is not just a minor feature request; it raises a fundamental question about whether AI-generated suggestions can be trusted by practitioners responsible for implementing them.

Expanding a multiple-actor perspective. The current library design primarily focuses on municipal actions, steps cities can take directly. Participants noted that this view is too limited. Many of the most effective climate actions require coordinated efforts among various actors: companies, utilities, property owners, civil society organizations, and regulators. The tool should do more than highlight what cities can do; it should also help map the entire landscape of involved actors, making clear who else needs to act, what each actor's role should be, and how the city fits within a larger system of stakeholders. Crucially, this includes surfacing whose goal is to drive a given intervention, so that when an intervention makes sense but cannot be led by the city, the library can help identify which actor is best placed to take it forward. Participants also requested actor-specific views, so that a company or community organization could filter the library to find interventions relevant to their particular role and capacity.

Richer decision-support information. To move from exploration to action, practitioners need more than just a list of intervention ideas. Participants requested estimates of the potential for emissions reductions for each intervention, indicative costs and implementation timelines, co-benefits such as economic, health, and social impacts, and, importantly, possible negative effects or trade-offs. This last point was illustrated directly: an intervention that reduces emissions but harms a particular economic sector or community group must be fully understood in its complexity if it is to be responsibly argued for and implemented. Participants noted that the library already has significant value as an argument tool, helping cities find and build the case for interventions, and that an integrated cost-benefit analysis would substantially strengthen this function. The workshop also revealed interest in integrating budget and funding contexts, so that the library could assist cities in identifying not just what is possible in theory, but what is feasible given available resources.

Strategic insight and intervention bundling. Perhaps the most ambitious insight from the workshop was the recognition that the library's greatest potential value may not lie in individual intervention suggestions. Still, at the higher level of intelligence, it could produce insights into how interventions relate to one another. Participants expressed a desire to understand which interventions are most influential, how groups of mutually reinforcing interventions can be combined into coherent strategies, and how those combinations could be aligned with a city's available budget so that a portfolio of actions, rather than isolated measures, can be planned and justified together. This points toward a vision of the tool as a strategic intelligence system, rather than just an idea generator.

Taken together, these insights confirm that AI can genuinely inspire new ideas in climate planning, and that the library is already seen as a powerful tool for thinking beyond familiar experience and surfacing interventions that might not otherwise enter a city's field of view. At the same time, participants were clear that the value of that inspiration depends on context, transparency, and a clear understanding of the stakeholder landscape in which interventions must be carried out.

Workshop Setting

The Kaohsiung workshop was held in May 2026 as part of the Kaohsiung–ICLEI Climate Neutral and Smart Cities Community of Practice, a thematic program bringing together cities across the Asia-Pacific region to advance urban climate action. The session was a Climate Action workshop designed to help cities translate their climate plans into concrete implementation pathways using ClimateView's Transition Element Framework. Ten cities participated: Kaohsiung (host city), Newcastle, Adelaide, Penang Island, Pasig, Quezon, Oakland, Kyoto City, Goyang City, and Gwangmyeong City.

> *"Opportunities for exchange between city practitioners are more important than ever. AI holds real promise in sustainable urban development, but it must work hand-in-hand with human knowledge and intelligence, not replace them."*
>
> **Jui-Hun Chang**
>
> **Director General of the Environmental Protection Bureau, Kaohsiung City**

As part of the workshop, participants used the AI Intervention Library prototype to identify and fill gaps in their city's intervention portfolios, working through the Transition Element Framework to map existing actions, surface missing pieces, and explore new possibilities. Following this hands-on exercise, participants were invited to reflect on the tool's usefulness and share ideas for features they would like to see in a future version. The session was conducted in a multilingual setting, with contributions in English, Mandarin Chinese, Japanese, and Korean, reflecting the regional diversity of the participant group.

Outcomes and insights

The workshop produced five categories of insight that together point to a regionally distinct set of priorities for the tool's continued development.

Regional relevance and contextual filtering. A recurring theme across the session was that the library's value depends on its ability to surface interventions that are actually applicable in a given regional context. Participants from across the Asia-Pacific group specifically noted that many examples in the current library appear to reflect conditions in European or North American cities and requested the ability to filter by region, for instance, to see interventions implemented by Southeast Asian or

East Asian cities. This is not merely a convenience feature: participants explained that when seeking to justify a policy to local decision-makers, being able to point to a precedent from a comparable city in the same region carries substantially more persuasive weight than a case from a distant context.

City typology and sector-specific sorting. Participants observed that cities vary significantly in their economic profiles, environmental and physical characteristics, with some primarily industrial, some coastal, some agricultural, and some commerce-focused, and that the library's usefulness would be greatly enhanced if it could sort and present interventions by city type. The suggestion was that an AI-powered platform should help a user characterize their city and then provide a curated set of interventions matched to that profile, rather than navigating a general catalog. The City of Kaohsiung noted that the current library appears to have limited coverage of industrial transition pathways and expressed interest in expanding the library to include more examples from industrial cities and sectors.

Multilingual access. One of the most emphatic and widely shared pieces of feedback was the need for the library to operate in languages other than English. Participants noted that all current content is in English, which creates a significant access barrier for city officials and practitioners whose primary working language is Mandarin, Japanese, Korean, or another Asian language. This was identified not as a secondary enhancement but as a prerequisite for the tool to be genuinely useful across the region. Participants connected this to a broader point about the tool's potential: an AI system that can surface relevant examples and translate them into a practitioner's own language would represent a qualitatively different kind of resource from anything currently available.

AI as a decision-support and simulation tool. Several participants engaged with the question of what distinguishes the AI Intervention Library from a general-purpose AI assistant such as a chatbot or search engine. The discussion revealed a desire for the tool to go beyond information retrieval and provide simulation or scenario-modelling capabilities, specifically, the ability to estimate the cross-sectoral co-benefits of an intervention (for example, the health, economic, and social effects of an energy efficiency subsidy), to compare those benefits across different policy levers, and to present the results in a form compelling enough to support funding requests and budget justifications. A participant from Kyoto described this directly in the context of subsidy allocation: local government officials often struggle to make the case for climate interventions when compared against spending in other sectors such as health or education, and a tool that could model and communicate comparative impact would address a genuine institutional weakness.

Responsible use and the role of human judgment. Participants from Pasig raised a point that applied broadly across the session: while the library is a genuinely useful tool for expanding the range of ideas a city considers, and for thinking outside the constraints of familiar frameworks, it should be accompanied by clear guidance on its limitations. Interventions drawn from the library must still be assessed for local feasibility, stakeholder fit, potential gaps and challenges, and community impact before being pursued. Participants suggested that a future version of the tool could support this process more actively, for instance, through a chatbot interface that helps a user trace an intervention back to its real-world implementers, or through connectivity features that allow cities to share experience and make direct contact with practitioners who have implemented a given approach. This points toward a vision of the library not as a static catalog but as a living network connecting cities with one another's knowledge and experience.

Taken together, the insights from the Kaohsiung workshop underscore that the AI Intervention Library's potential in the Asia-Pacific region is substantial, but its realization depends on addressing significant gaps in language access, regional coverage, and contextual matching. Participants were enthusiastic about the tool's capacity to inspire new thinking and surface examples that would otherwise remain invisible, and equally clear that this inspiration must be grounded in local realities to translate into effective action.

Synthesis of Workshop Insights

Feature	Description	User need addressed	Source
Navigation & discovery			
Keyword search & multi-layer filters	Filter interventions by sector, actor type, impact, cost, and timeframe. Include keyword search to quickly find relevant areas.	Prevent overload; enable purposeful navigation of a large library	**Both workshops**
City typology matching	Allow users to characterize their city (industrial, coastal, agricultural, etc.) and receive a curated set of interventions matched to that profile.	Surface relevant interventions without manual browsing	**Kaohsiung**
Regional filtering	Filter examples by geography so practitioners can find precedents from cities in comparable regional contexts.	Increase persuasive and political credibility of suggestions	**Kaohsiung**
Strategic intervention bundling	Enable users to select multiple interventions and explore how they could be strategically combined, surfacing synergies, dependencies, and the biggest levers within a portfolio.	Support coherent strategy-building, not just individual action selection	**Dortmund**
Context & feasibility			
Context-fit indicator (traffic light)	A traffic-light or scoring system assessing how well each intervention fits a given city's policy environment, infrastructure, and economic conditions, including basic requirements and key enablers.	Build practitioner trust; flag mismatches before effort is invested	**Dortmund**
Budget & timeline estimates	Provide indicative cost ranges and implementation timelines per intervention. Where possible, link to city budget context so users can see what is feasible given available resources.	Move from inspiration to actionable, resource-aware planning	**Dortmund**
Funding discovery & signposting	For each intervention, surface relevant funding sources, including grants, programs, and financing mechanisms, from national and regional bodies, and	Help local authorities proactively identify and apply for available funding individually or	**Scottish Climate Intelligence Service**

Feature	Description	User need addressed	Source
	link them by sector/intervention type.	through aggregation.	
Obstacles & constraints flagging	Identify common barriers to implementation for each intervention, political, financial, technical, or social, so that cities can anticipate and plan for them.	Set realistic expectations; support risk-aware decision-making	**Dortmund**
Sector gap expansion	Expand library coverage beyond transport and energy to include industrial transition, agriculture, and other underrepresented sectors.	Serve cities whose primary challenges lie outside current coverage	**Kaohsiung**
Impact & decision support			
Emissions reduction estimates	Provide an indicative quantification of potential emissions reductions per intervention, to support prioritization and argument-building.	Enable evidence-based prioritization and stakeholder buy-in	**Both workshops**
Co-benefits display	Surface economic, health, social, and other co-benefits alongside climate outcomes. Strengthen the case for interventions by connecting them to a city's broader goals.	Support cross-departmental and political buy-in	**Both workshops**
Negative effects & trade-off warnings	Flag potential negative effects or unintended consequences of interventions for specific stakeholder groups, making trade-offs visible and enabling more honest, responsible advocacy.	Prevent harm; support credible and nuanced policy arguments	**Dortmund**
Integrated cost-benefit analysis	Provide a structured cost-benefit framework for each intervention, including direct and indirect effects, to serve as a tool for securing funding and political approval.	Strengthen the case for investment in specific interventions	**Dortmund**
Cross-sector impact simulation	Model the comparative impact of an intervention across sectors (e.g. climate vs. health vs. education) to support subsidy justification and budget negotiations with central government.	Help officials make the case for climate spending against competing priorities	**Kaohsiung**

Feature	Description	User need addressed	Source
Lever identification	Highlight which interventions represent the biggest levers in a given city context, and how they connect to or unlock other interventions.	Focus limited resources on highest-impact actions	**Dortmund**
Actor landscape & stakeholders			
Multi-actor role mapping	For each intervention, define the full set of actors required (e.g., companies, utilities, property owners, civil society, regulators) and specify each actor's role, including who should drive implementation if the city is not the right actor.	Reveal system-level requirements; clarify accountability beyond municipal government	**Dortmund**
Actor-specific filtered views	Allow non-municipal users (companies, community organizations, utilities) to filter the library for interventions where they can play a role, with context tailored to their capacity and perspective.	Expand the tool's usefulness beyond city government; engage the private sector	**Dortmund**
Conflicting goals visibility	Make visible where different actors' goals may conflict around a given intervention, so that trade-offs and political dynamics can be anticipated and addressed.	Support realistic planning and stakeholder negotiation	**Dortmund**
Real-world evidence & connectivity			
Richer real-world example linking	More clearly connect each intervention to real-world case examples, including outcome data, so the library serves as both an evidence base and an idea generator.	Strengthen argument-building; demonstrate that interventions have worked in practice	**Dortmund**
City-to-city connectivity	Enable cities that have implemented an intervention to be directly contactable, turning the library into a peer learning network and a knowledge base.	Facilitate direct knowledge exchange and trust-building between cities	**Kaohsiung**
Interface & accessibility			
Conversational / chatbot interface	An interactive AI assistant to help users navigate the library, validate ideas, handle complex queries, connect to real-world	Lower barrier to use; bridge gap between inspiration and action	**Both workshops**

Feature	Description	User need addressed	Source
	implementers, and guide users toward the right actor for a given intervention.		
Multilingual support	Make the library fully operable in languages beyond English.	Remove language as a barrier to access across regions	**Kaohsiung**
Usage guidance & responsible use disclaimers	Provide clear documentation on the library's limitations, recommended workflow, and the essential role of human judgment in validating and contextualizing AI-generated suggestions.	Ensure responsible use; maintain practitioner accountability	**Kaohsiung**

AI for Shared Transition Intelligence

The workshops described in this paper point toward a vision of AI in climate governance that extends beyond any individual tool or feature set. The greater opportunity lies in creating shared transition intelligence: a dynamic, collaborative knowledge infrastructure that becomes more valuable as more cities use, contribute to, and improve it.

The concept of shared transition intelligence is based on a simple yet powerful idea: every city working on climate action creates knowledge useful to other cities facing similar issues. Currently, most of that knowledge is scattered across local reports, institutional memory, and individual practitioners' experiences. AI helps to uncover, organize, and share that knowledge on a large scale.

This is the ultimate goal of the AI Intervention Library. Not just a static database of past interventions, but a dynamic system that learns from real-world use, updates its understanding of what works and in which situations, and constantly improves its ability to help cities identify and take effective action. Every city that contributes by documenting interventions, sharing implementation experiences, and giving feedback makes the next city smarter.

Realizing this ambition requires more than just technical development. It calls for governance structures that guarantee data quality and proper use; transparency about how AI recommendations are generated and their limitations; and ongoing collaboration with urban practitioners to keep the tool grounded in the realities of climate governance rather than drifting toward abstract optimization.

It also requires honesty about the environmental impact of AI itself. Generative AI systems use significant energy and computational resources [15]. Using AI irresponsibly in the name of sustainability would be counterproductive. The approach outlined in this paper is designed to intentionally use AI to deliver multiple benefits for scaling and accelerating climate action across regions with limited capacity by leveraging its generative capabilities to address the specific challenges of invention and strategy, where the significant potential benefit of AI assistance lies.

The workshops and the prototype described in this paper point to six principles that should guide the development and deployment of AI tools for urban sustainability.

Ground AI in causal frameworks, not just data. The value of AI for cities does not lie in the size of its training data but in the quality of its reasoning structure. A tool that can explain why and how an intervention might work, not just where it has been tried, will be fundamentally more useful than one that cannot. The TEF and Outcome Logic represent one such framework. The field needs more, and it needs them to be interoperable.

Build for the system, not just the city. Climate transitions require coordinated action across actors that extend far beyond municipal government [16], [17]. The next generation of AI tools should map the full actor landscape around each intervention, making clear who needs to act, in what sequence, and where the city's role ends and others' begin. An AI that only surfaces what the city can do unilaterally describes, at best, only a fraction of the transition.

Think in portfolios, act through causal chains. Individual interventions rarely deliver transformational change on their own. The enabling conditions for any meaningful activity shift, whether increasing cycling, accelerating building retrofits, or transitioning industrial processes, typically require a coordinated set of actions targeting different parts of the same causal chain simultaneously. This is where causal logic becomes not just analytically useful but practically essential. By mapping the enabling conditions that stand between a city's current state and its desired outcome, AI can identify not just individual interventions but the combination of measures that together create the conditions for change to occur. The goal is not a list of good ideas but a coherent portfolio, one where the pieces are chosen because of how they interact, not despite it.

Make context the filter, not the afterthought. Participants in both Dortmund and Kaohsiung pointed to the same gap: the distance between what AI suggests and what a given city can actually implement. Useful AI does not just identify and generate possibilities; it assesses their fit against a city's policy environment, infrastructure, economic conditions, regional context, and available resources. Feasibility is a prerequisite rather than a feature to be added later.

Treat every implementation as shared knowledge. The most powerful version of AI for cities is not a tool used in isolation but a system that becomes smarter with every use. Cities that document what they tried, what worked, and what did not contribute to a growing infrastructure of transition intelligence that benefits every city that follows. Realizing this requires shared standards, open data protocols, and governance structures that protect quality and ensure equitable access across geographies and languages.

Keep humans accountable, not just in the loop. AI-generated suggestions must be validated by practitioners who understand local context, stakeholder dynamics, and political feasibility. Responsible use is not a disclaimer to be appended at the end; it is a design requirement. The next generation of tools should actively support the human judgment process by making trade-offs visible, flagging uncertainty, and connecting practitioners to peers who have navigated similar decisions.

The immediate priority for tools like the AI Intervention Library is to translate structured input from cities into the next generation of development: improved contextual filtering, actor mapping, richer decision-support information, and multilingual access. Systematic co-creation with practitioners is not a one-off exercise but an ongoing requirement, and building the mechanisms to capture, synthesize, and act on city feedback at scale is itself a critical next step.

Conclusion

The cities that will lead in climate action over the coming decade are not those with the most extensive records of past efforts. They are the ones who can move most quickly from ambition to implementation, identifying the right combination of policies, infrastructure, and behavior changes for their specific situation, securing political and financial backing, and learning swiftly from their efforts. It will also be the ones recognized for sharing these experiences with others.

Realizing this potential depends in part on opportunities for cities to learn from one another. International city networks such as ICLEI and the Global Covenant of Mayors have become important platforms for sharing knowledge, building institutional capacity, and scaling promising local experiments to citywide adoption [18], [19], [20]. Initiatives such as WWF's One Planet City Challenge go further by providing common assessment frameworks that allow cities to benchmark progress against shared standards. These networks demonstrate that shared platforms and assessments are among the most effective mechanisms for turning isolated local experience into transferable transition intelligence.

AI has a role to play in all these challenges. But its most powerful contribution is not only retrospective, by cataloging what has already been tried. It lies in the generative ability to help cities imagine and design what needs to happen next, interventions that may not yet exist in any database, strategies that emerge from the combination of structured knowledge and creative thinking.

The AI Intervention Library prototype proves this vision is technically feasible. The workshops in Dortmund and Kaohsiung show that it is practically desired. What practitioners ask for is not a replacement for human judgment, but a tool that broadens what they can see and consider, and links their local thinking to a global body of knowledge and experience.

The next generation of AI for cities will not just analyze the past. It will help create the future of climate action, and, in doing so, bridge the gap between the climate goals cities have set and the changes those goals require.

References

[1] C. Diaz et al., “Addressing climate-action implementation gaps in cities: planning for meaningful and consistent implementation through the 5-UP framework,” Front. Sustain. Cities, vol. 7, no. 1649194, Oct. 2025, doi: 10.3389/frsc.2025.1649194.

[2] B. Kaplan Weinger, “Overcoming the implementation gap: everyday barriers and enablers of urban climate governance,” Front. Sustain. Cities, vol. 8, no. 1656615, Feb. 2026, doi: 10.3389/frsc.2026.1656615.

[3] F. Larosa, S. Hoyas, J. A. Conejero, J. Garcia-Martinez, F. Fuso-Nerini, and R. Vinuesa, “Large language models in climate and sustainability policy: limits and opportunities,” Environ. Res. Lett., vol. 20, no. 7, p. 074032, Jul. 2025.

[4] C. Rosenzweig et al., “Building and using the evidence base for urban climate action: the UCCRN City Solutions Case Study Atlas,” NPJ Urban Sustain., vol. 6, no. 1, Feb. 2026, doi: 10.1038/s42949-026-00342-z.

[5] S. A. Vaghefi et al., “ChatClimate: Grounding conversational AI in climate science,” Commun. Earth Environ., vol. 4, no. 1, Dec. 2023, doi: 10.1038/s43247-023-01084-x.

[6] D. de Ridder and S. Thomaes, “When and how behavior change can accelerate system change (and vice versa): mapping reciprocal processes for climate change mitigation,” Behav. Public Policy, pp. 1–18, Jun. 2025.

[7] D. Richards and D. Worden, “Applications of generative artificial intelligence to influence climate change decisions,” NPJ Clim. Action, vol. 3, no. 1, Dec. 2024, doi: 10.1038/s44168-024-00202-5.

[8] A. Manivannan et al., “Generative AI for climate governance and acceptability-constrained policy design,” NPJ Clim. Action, vol. 5, no. 1, p. 37, Mar. 2026.

[9] V. Todeschi, G. Hernandez-Moral, E. Clementi, P. Bertoldi, and G. Melica, “Financing climate change mitigation actions in cities: insights from the Covenant of Mayors initiative,” Sustain. Cities Soc., vol. 136, no. 107097, p. 107097, Jan. 2026.

[10] CCFLA, “The State of Cities Climate Finance 2024. Cities Climate Finance Leadership Alliance,” 2024. [Online]. Available: https://citiesclimatefinance.org/publications/2024-state-of-cities-climate-finance

Authors

Tomer Shalit
Mark Dixon
David Corbett
Pourya Salehi
Juan Pablo Astolfo
Deniz Esin Emer
Clare Wharmby
Alreem Alshimmari
Alvaro Rodriguez
Uri Colodro Gotthelf
Simone Sandholz
Henricus Hölsgens
Lana Fassbender
Jürgen Schultze
Christopher Sadlowski

[11] Shalit, T., Dixon, M. & Bergöö, M., “Standardising Climate Mitigation: The Transition Element Framework An Open-Source Initiative To Structure And Codify IPCC Knowledge For Implementation Of Climate Action,” ClimateView AB, 2024. [Online]. Available: https://www.transitionelements.org/downloads/white_paper_standardising_climate_mitigation.pdf

[12] Shalit, T., Dixon, M., Corbett, D., Salehi, P., Schultz, L., Manner, M., Andrén, C., Wharmby, C., Boertien, R., “Designing for delivery: How cities can structure fundable climate transitions,” ClimateView AB, 2025. [Online]. Available: https://www.climateview.global/en/whitepaper/designing-for-delivery

[13] M. Alméstar, S. Romero-Muñoz, and N. Mestre, “Breaking silos: A systemic portfolio approach and digital tool for collaborative urban decarbonisation,” Sustainability, vol. 17, no. 11, p. 5145, Jun. 2025.

[14] E. M. Leclercq and E. A. Rijshouwer, “Enabling citizens’ Right to the Smart City through the co-creation of digital platforms,” Urban Transform., vol. 4, no. 1, Dec. 2022, doi: 10.1186/s42854-022-00030-y.

[15] I. Alnafrah, “The Two Tales of AI: A Global assessment of the environmental impacts of artificial intelligence from a multidimensional policy perspective,” J. Environ. Manage., vol. 392, no. 126813, p. 126813, Sep. 2025.

[16] G. Doci, H. Dorst, S. Hillen, and T. Tjokrodikromo, “Urban transition governance in practice: exploring how European cities govern local transitions to achieve climate neutrality,” Front. Sustain. Cities, vol. 7, no. 1559356, Jun. 2025, doi: 10.3389/frsc.2025.1559356.

[17] K. Pereverza, H. Rohracher, and O. Kordas, “Fostering urban climate transition through innovative governance coordination,” Environ. Pol. Gov., vol. 35, no. 4, pp. 631–646, Aug. 2025.

[18] D. J. Gordon and C. A. Johnson, “City-networks, global climate governance, and the road to 1.5 °C,” Curr. Opin. Environ. Sustain., vol. 30, pp. 35–41, Feb. 2018.

[19] T. M. P. Nguyen, K. Davidson, and L. Coenen, “Understanding how city networks are leveraging climate action: experimentation through C40,” Urban Transform., vol. 2, no. 1, Dec. 2020, doi: 10.1186/s42854-020-00017-7.

[20] A. P. Jakobi, B. Loges, and R. Haenschen, “What do international city networks contribute to global governance? Towards a better conceptual and empirical assessment,” Glob. Soc., pp. 1–23, May 2024.