Concept Paper

Not peer-reviewed version

PCA-OS: A Planetary Climate Adaptation Operating System

Chaoyue He * , Xin Zhou , Di Wang , Hong Xu , Wei Liu , Chunyan Miao

Posted Date: 10 March 2026

doi: 10.20944/preprints202603.0753.v1

Keywords: climate adaptation; world model; digital twins; global causal knowledge graph; continual learning; causal inference; foundation models; robust optimization; decision support

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

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026

doi:10.20944/preprints202603.0753.v1

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

Concept Paper

PCA-OS: A Planetary Climate Adaptation OS

Chaoyue He [1,*], Xin Zhou [1], Di Wang [1], Hong Xu [1], Wei Liu [2] and Chunyan Miao [1]

[1] Alibaba–NTU Global e-Sustainability CorpLab (ANGEL), Singapore
[2] Alibaba Group, China
* Correspondence: cyhe@ntu.edu.sg

Abstract

Predictive climate machine learning is increasingly good at forecasting hazards, but hazard maps alone do not decide what to do, where, when, for whom, and under which futures. We argue that climate ML remains insufficient for adaptation unless interventions are treated as first-class, versioned, and auditable objects. This matters because many climate digital twins still prioritize state estimation and simulation, while adaptation requires intervention observability, counterfactual effect estimation, and constrained portfolio choice. We propose **PCA-OS** (**P**lanetary **C**limate **A**daptation **O**perating **S**ystem), a decision-support operating abstraction that uses an intervention-aware global causal knowledge graph and standardizes object schemas, versioned updates, query primitives, and audit interfaces across three core system objects: (1) an ADAPTATION INTERVENTION LEDGER that records measurable interventions with provenance and uncertainty; (2) a CAUSAL EFFECT ATLAS that stores scenario-indexed, spillover-aware estimands, identification assumptions, diagnostics, and sensitivity bounds; and (3) a ROBUST PORTFOLIO DECISION LAYER that optimizes intervention portfolios under budget, equity, and no-harm constraints. We argue that foundation models and intervention-aware world models should support, rather than replace, identification-aware causal analysis by surfacing candidate confounders, mechanisms, and spillover pathways for human review. Finally, we outline ADAPTBENCH, an evaluation suite in which systems can fail for inequitable or maladaptive recommendations even when predictive accuracy is high. The result is a field-level provocation: move climate ML from read-only hazard intelligence to auditable decision support for adaptation.

Keywords: climate adaptation; world model; digital twins; global causal knowledge graph; continual learning; causal inference; foundation models; robust optimization; decision support

---

1. Introduction & Provocation

The climate-ML stack is becoming spectacularly good at producing read-only futures. We can forecast heat, flood, smoke, and drought at unprecedented spatiotemporal resolutions, even as climate change shatters historical stationarity [1]. Yet adaptation is not won on hazard maps. It is won on interventions: which roofs are cooled, which drainage networks are upgraded, and which communities absorb spillovers. Without first-class, versioned representations of actions and their downstream effects, many climate digital twins remain primarily observation and simulation platforms rather than intervention systems.

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

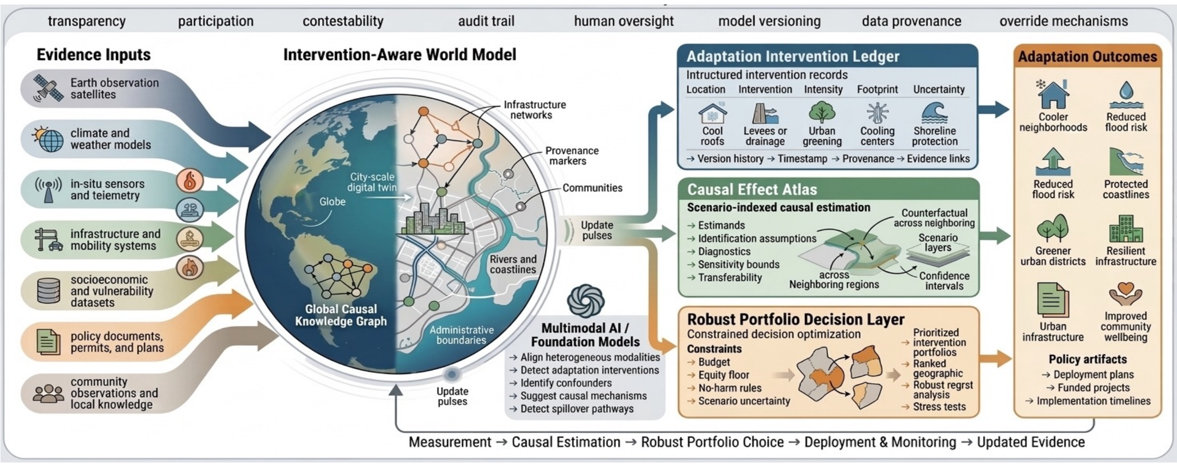

Figure 1. Conceptual infographic overview of PCA-OS. System overview of PCA-OS. Diverse evidence inputs—including Earth observation satellites, climate and weather models, in-situ sensors, infrastructure and mobility systems, socioeconomic and vulnerability datasets, policy documents and plans, and community observations—feed into an intervention-aware world model. The world model combines a globe-scale representation, city-scale digital twins, and a Global Causal Knowledge Graph linking infrastructure, communities, interventions, rivers, and administrative boundaries with typed relations, provenance markers, and continuous update pulses. Multimodal AI and foundation models align heterogeneous modalities, detect adaptation interventions, identify confounders, suggest causal mechanisms, and trace spillover pathways. On top of this world model, three auditable system objects are constructed: an Adaptation Intervention Ledger storing structured and versioned intervention records; a Causal Effect Atlas providing scenario-indexed causal estimates with identification assumptions, diagnostics, and counterfactual reasoning across regions; and a Robust Portfolio Decision Layer performing constrained optimization under budget, equity, and no-harm constraints. The pipeline is governed by transparency, participation, contestability, audit trails, human oversight, model versioning, data provenance, and override mechanisms. The system produces adaptation outcomes such as cooler neighborhoods, reduced flood risk, protected coastlines, greener districts, resilient infrastructure, and improved community wellbeing, while continuously updating evidence through deployment and monitoring.

This is a structural blind spot in how the KDD community currently frames climate intelligence. Adaptation evidence remains fragmented, highly context-dependent, and thin on causal outcomes [2–8]. Interventions are often targeted where visibility or capacity is already high [9,10], can generate spatial spillovers [11,12], and must survive deeply uncertain futures [13,14]. We have therefore optimized hazard prediction while remaining under-instrumented for intervention choice. We argue that adaptation should be framed as a continual learning and decision loop: **measurement → causal estimation → robust, equity-constrained optimization**.

In this framing, the primary scientific object is no longer the hazard map—it is the *intervention object*. To make hazard prediction actionable, it must be linked to measurable interventions, counterfactual effects, and portfolio trade-offs. PCA-OS directly answers this need by organizing planetary adaptation around three shared infrastructural artifacts: (1) an ADAPTATION INTERVENTION LEDGER that records where and when interventions occur, together with provenance and uncertainty; (2) a CAUSAL EFFECT ATLAS that stores scenario-indexed, spillover-aware causal estimates together with estimands, identification assumptions, diagnostics, and sensitivity bounds; and (3) a ROBUST PORTFOLIO DECISION LAYER that turns those estimates into intervention portfolios under explicit budget, equity, and no-harm constraints. Together, these artifacts turn climate intelligence from descriptive monitoring into auditable decision support. Put differently, PCA-OS is an operating abstraction because it standardizes the shared objects, schemas, versioned update rules, query primitives, and audit interfaces through which adaptation evidence becomes intervention decisions.

Terminology. A *climate digital twin* provides state-estimation and simulation; an *intervention-aware world model* is its internal representation, here realized as a global causal knowledge graph over entities, relations, provenance, and updates; and an *adaptation operating system* coordinates interventions, effect

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

objects, portfolio decisions, and audits through shared schemas, versioning, update semantics, and query interfaces. PCA-OS builds on the first two, but its core contribution is this OS layer.
Scope. PCA-OS is a decision-support operating abstraction, not a claim of fully autonomous intervention control; its purpose is to standardize how evidence, assumptions, and recommendations are updated, queried, and audited.
Contributions. This paper introduces PCA-OS, an operating abstraction for climate adaptation in which interventions become first-class, versioned, and contestable objects within climate ML systems. We argue that hazard-centric climate intelligence is structurally insufficient for adaptation unless interventions, their causal effects, and resulting decisions are jointly represented and auditable. To ground this vision, we outline a minimum viable PCA-OS at city scale: detecting parcel-level cool-roof retrofits from Earth observation and permit data, estimating tract-level avoided heat exposure and peak-demand impacts with neighborhood spillovers, and optimizing subsidy rollout under budget, equity-floor, and bounded-harm constraints. This narrow loop operationalizes the coupled measurement–causality–decision stack. **Why Blue Sky, why KDD, why now.** We position this work as a Blue Sky agenda for KDD, enabled by advances in foundation models, planetary-scale data systems, and causal decision methods that make auditable, intervention-centered adaptation science feasible.

2. Related Work & Positioning

Adaptation tracking already shows why the status quo is inadequate. Prior work argues that adaptation evidence must be consistent, comparable, comprehensive, and coherent [7]; national and global assessments track plans and reported actions [5,10]; and systematic stocktakes show that outcome evidence remains sparse and heterogeneous [6]. Big-data approaches were proposed to improve adaptation monitoring [8], but most existing pipelines still terminate at inventories, dashboards, or hazard forecasts. PCA-OS advances by making interventions themselves measurable, versioned, and causally queryable.

Climate digital twins show that planetary state estimation is increasingly feasible, but intervention choice is rarely the primary object. Table 1 summarizes the intended distinction. Climate twin efforts such as Destination Earth emphasize environmental states, simulations, and climate services [15,16]. PCA-OS is complementary but higher-level: it can build on climate twin infrastructure, yet adaptation additionally requires intervention ledgers, causal effect objects, and auditable portfolio decisions under equity and safety constraints. The ingredients already exist but remain disconnected: foundation models and geospatial data systems improve planetary measurement [17–24]; multimodal and language-grounded models align heterogeneous evidence [25–29]; causal inference supplies identification-aware estimators under confounding and spillovers [30–38]; and robust optimization plus fairness-aware learning formalize stress-tested action selection [39–45]. Here, we fuse these lines into one intervention-centered operating abstraction.

Table 1. Comparison of PCA-OS among climate–ML systems, in terms of degree of intervention modeling, causal effect estimation, decision support, and auditability.

System	Interv.	Effects	Decision	Auditability
Hazard/Weather forecasting [19,20]	no	no	no	limited
Adaptation tracking/assessment [5,10]	partial	no	no	limited
Climate digital twins [15,16]	partial	partial	rare	partial
PCA-OS (this work)	core	core	core	core

3. PCA-OS: From Architecture to Paradigm Shift

PCA-OS is not merely a simulator. It maintains an intervention-aware world model, implemented as a global causal knowledge graph, with auditable intervention, effect, decision, and governance objects as conditions evolve. By OS, we mean a coordinating layer that manages these shared objects across the measurement $\rightarrow$ causal estimation $\rightarrow$ portfolio selection loop through explicit schemas,

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

versioning rules, update semantics, query primitives, and audit trails. At minimum, this means typed ledger, atlas, and portfolio objects with stable identifiers, provenance links, revision histories, and queryable dependency pointers, so both algorithms and human reviewers can inspect what changed, why it changed, which assumptions were used, and which decisions were affected.

The system follows four design commitments. **First, decision-first:** the primary outputs are intervention, effect, and portfolio objects rather than hazard maps alone. **Second, interventions as first-class objects:** actions are versioned, geolocated, uncertain, and causally queryable. **Third, uncertainty-forward:** measurement, identification, and climate-scenario uncertainty must propagate into the decision layer. **Fourth, normative constraints as infrastructure:** equity, no-harm, and contestability cannot be post-hoc commentary; they must be encoded inside the optimization and the interface. Operationally, this agenda can be bootstrapped from open geospatial standards and cloud-native planetary data planes built around STAC catalogs, NASA HLS, ERA5, the Planetary Computer, OpenStreetMap, and WorldPop [21–24,46,47].

3.1. An Intervention-Aware World Model

We instantiate the underlying world model maintained by PCA-OS at time t as a dynamic global causal knowledge graph (KG):

$$\mathcal{K}_t = (V_t, E_t, \mathcal{R}, X_t), \tag{1}$$

where X_t collects time-varying node and edge attributes, including uncertainty metadata. This is a systems abstraction: a global causal knowledge graph in which intervention, exposure, mechanism, and outcome semantics are represented explicitly, but not every stored edge is itself an identified causal effect. The graph stores entities, relations, topology changes, provenance, uncertainty, and typed causal semantics; identified causal claims are represented separately in atlas entries that attach estimands, assumptions, diagnostics, and sensitivity to selected intervention–outcome relations. This separation keeps the world model updatable without conflating storage with identification.

Nodes can represent spatial units, infrastructure assets, interventions, institutions, communities, and governing documents; edges instantiate relation types such as hydrologic connectivity, ecological adjacency, mobility, service access, ownership, administrative jurisdiction, and intervention–exposure pathways. Recent sustainability knowledge-graph efforts suggest a practical foundation for domain-grounded, updatable global causal knowledge graphs at scale [48–50]. Each node $u \in V_t$ has an available modality set $\mathcal{M}_{u,t} \subseteq \mathcal{M}$ and carries multimodal observations

$$o_{u,t} = \{m_{u,t}^{(k)} : k \in \mathcal{M}_{u,t}\}, \tag{2}$$

where the open modality registry $\mathcal{M}$ can include EO, climate reanalyses, sensors, text or administrative records, mobility, infrastructure telemetry, and damage or loss signals. A multimodal encoder then produces an adaptation-relevant state embedding $z_{u,t} = f_\theta(o_{u,t})$. Modality-agnostic architectures and aligned representation learning are essential because the relevant evidence is rarely fully paired [25,26]. Interventions can alter not only node features but also graph topology: a levee attenuates hydrologic edges, urban greening creates ecological corridors, a cooling center creates accessibility edges, and zoning decisions activate governance relations.

We therefore require explicit update operators

$$\mathcal{K}_{t+1} = U_\phi(\mathcal{K}_t, \mathcal{O}_{t+1}, \Delta\mathcal{L}_{t+1}, \Delta\mathcal{G}_{t+1}), \tag{3}$$

where $\mathcal{O}_{t+1}$ is the batch of new observations arriving between t and $t+1$, $\Delta\mathcal{L}_{t+1}$ is the set of new or revised ledger entries, and $\Delta\mathcal{G}_{t+1}$ is the set of governance or policy updates. These update operators are part of the OS semantics: they specify how new evidence mutates shared state, when revised entries supersede prior versions, and how downstream atlas or portfolio objects are marked for re-estimation or re-optimization. This gives the research community a sharper target: mine, update, and reason over a global causal KG with intervention-aware topologies rather than static hazard layers.

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

3.2. The Adaptation Intervention Ledger

A planetary adaptation science cannot exist without intervention observability. PCA-OS therefore introduces an **Adaptation Intervention Ledger** $\mathcal{L}$ that records where interventions occurred, when they changed, how intense they were, and what evidence supports those claims. A ledger entry is

$$\ell = (\mathrm{loc}, [t^{\mathrm{start}}, t^{\mathrm{end}}], \mathrm{type}, \mathrm{intensity}, \mathrm{footprint}, p_\ell, \mathrm{provenance}), \tag{4}$$

where loc denotes the affected spatial unit or geometry, and p_ℓ is a joint uncertainty object over type, timing, geometry, and intensity. Each ledger entry should carry a stable identifier, schema-validated fields, and a revision history so downstream atlas and portfolio objects can resolve exactly which intervention state they consumed.

The ledger fuses four evidence streams: **(1) EO change signals** such as albedo shifts, shoreline changes, vegetation dynamics, or impervious-surface updates; **(2) SAR and event-response evidence** for flood-related interventions and outcomes; **(3) text exhaust** such as permits, council minutes, engineering standards, procurement records, and adaptation plans; and **(4) operational or participatory traces** such as maintenance logs, sensor anomalies, and community reports. Because many interventions are informal, partially documented, or politically invisible, the ledger must preserve uncertainty, missingness, and multi-view disagreement rather than forcing premature certainty.

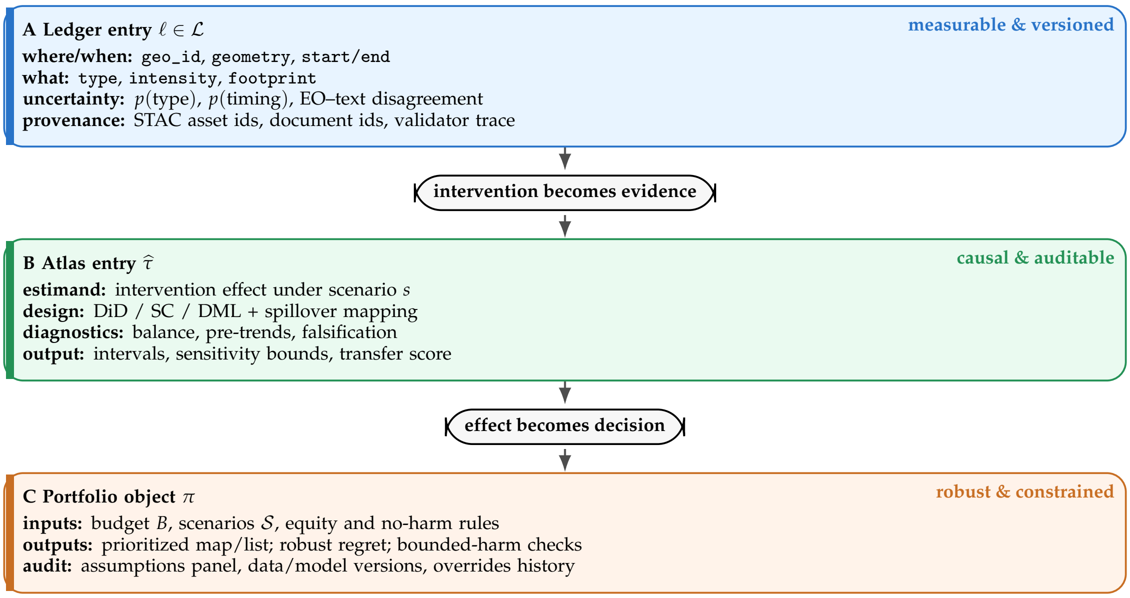

Figure 2. The PCA-OS object flow: interventions become ledger objects, then effect objects, then decision objects; uncertainty, provenance, and assumptions are carried forward. Three stacked rounded boxes connected by arrows illustrate the PCA-OS object flow. The top box is a ledger entry with where and when fields, what fields, uncertainty, and provenance. The middle box is an atlas entry with an estimand, a causal design, diagnostics, and outputs such as intervals and sensitivity bounds. The bottom box is a portfolio object with budget, scenarios, equity and no-harm inputs, prioritized outputs, and audit information. Labels indicate the progression from measurable and versioned, to causal and auditable, to robust and constrained.

3.3. The Causal Effect Atlas

Because adaptation placement is non-random, prediction alone is biased. PCA-OS therefore maintains a **Causal Effect Atlas** whose entries are explicit objects: scenario-indexed, design-annotated, spillover-aware estimates with diagnostics and sensitivity bounds. Atlas entries are queryable by intervention family, geography, outcome, climate scenario, estimand, and design, making causal evidence a reusable system object rather than a one-off appendix table. Potential outcomes must admit interference:

$$Y_{i,t}(a_{i,t}, e_{i,t}, s), \qquad e_{i,t} = g(a_{j,t} : j \in \mathcal{N}(i)), \tag{5}$$

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

where $a_{i,t}$ is local action, $e_{i,t}$ is neighbor exposure, and $s \in \mathcal{S}$ indexes climate scenarios. Each atlas entry stores an estimand, identification strategy, interference model, diagnostics, transportability warning, and the limitations of the outcome proxy. Relevant estimation tools include staggered DiD, synthetic control, heterogeneous treatment effect estimation, double/debiased ML, and off-policy evaluation [32–34,37,38]. Foundation models and the world model act here as assistive hypothesis engines: they can surface candidate confounders, mechanisms, and spillover pathways for human review, but causal validity still rests on explicit identification assumptions and diagnostics [27–29].

3.4. The Robust Portfolio Decision Layer

The decision layer chooses portfolios of interventions rather than isolated actions. Let $\mathcal{I}$ index decision locations, let $\pi = (a_i)_{i\in\mathcal{I}}$ denote a portfolio with $a_i \in \mathcal{A}_i$, let $\mathcal{J}$ index protected or potentially harmed groups or locations, let $e_i(\pi)$ denote spillover exposure at location i, let $\widehat{\tau}_i^{(s)}(a_i, e_i(\pi))$ be the estimated normalized benefit under scenario $s \in \mathcal{S}$, and let $w_i \geq 0$ encode policy weights such as vulnerability or priority. Define the scenario-specific welfare of a portfolio as

$$V_s(\pi) = \sum_{i\in\mathcal{I}} w_i\, \widehat{\tau}_i^{(s)}(a_i, e_i(\pi)). \tag{6}$$

The robust portfolio decision then solves

$$\begin{aligned} \max_{\pi\in\prod_{i\in\mathcal{I}}\mathcal{A}_i} \quad & \min_{s\in\mathcal{S}} V_s(\pi) \\ \text{s.t.} \quad & \sum_{i\in\mathcal{I}} \mathrm{cost}_i(a_i) \leq B, \quad \mathrm{Equity}^{(s)}(\pi) \geq \alpha, \qquad \forall s \in \mathcal{S}, \\ & \Delta\mathrm{Risk}_j^{(s)}(\pi) \leq \delta_j, \qquad \forall j \in \mathcal{J},\ \forall s \in \mathcal{S}. \end{aligned} \tag{7}$$

Here, $\mathrm{Equity}^{(s)}(\pi)$ denotes an auditable group-level equity functional under scenario s, and $\Delta\mathrm{Risk}_j^{(s)}(\pi)$ denotes the incremental risk imposed on protected location j relative to a no-action baseline. The equity floor α and bounded-harm thresholds δ_j are policy inputs chosen through stakeholder and governance processes, not values discovered by the model. To propagate effect-estimation uncertainty into decision making, the atlas can supply conservative lower-confidence values or ambiguity sets for $\widehat{\tau}_i^{(s)}$, so the portfolio layer is robust to both climate scenarios and causal uncertainty. The robust-feasible set is therefore $\Pi = \bigcap_{s\in\mathcal{S}} \Pi_s$. Each portfolio object should retain pointers to the exact ledger versions, atlas versions, constraints, and any human overrides from which it was derived. A decision-facing score is ROBUST DECISION REGRET:

$$\mathrm{RDR}(\pi) = \max_{s\in\mathcal{S}} \left(\max_{\pi'\in\Pi_s} V_s(\pi') - V_s(\pi) \right), \tag{8}$$

where Π_s is the set of portfolios that satisfy the budget, equity, and no-harm constraints under scenario s. Lower is better, and $\mathrm{RDR}(\pi) \geq 0$ by construction [39–45].

3.5. Foundation Models in PCA-OS

EO models can improve intervention detection and outcome proxies such as albedo, heat, and hydrologic context [17,18]. Weather and climate models provide scenario backbones for effect transport and stress testing [19,20]. Language-grounded models align permits, plans, procurement records, and operations text with geospatial evidence [27–29]. The open problem is not only representation quality but also binding them to explicit intervention semantics, identification assumptions, and auditable decision objects [51,52].

4. Human Interfaces, Governance, & Auditing

Because adaptation is inherently political, PCA-OS must be challengeable by design. Governance has four technical pillars. **Transparency:** causal estimands, identification choices, diagnostics, and sen-

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

sitivity analyses must be inspectable rather than buried in appendices [30]. **Participation:** stakeholders should be able to dispute detections, annotate local constraints, and force explicit handling of missing or contested evidence. **Safety:** optimization must obey no-harm and bounded-harm constraints using safe decision-making principles [53,54]. **Traceability:** users should be able to inspect what changed, why it changed, and whether recommendation shifts were driven by data, assumptions, or scenario updates. Datasheets and model cards are therefore not peripheral documentation; they are core infrastructure joined to the intervention and audit graph [55,56]. The same governance stack would also be essential for future high-spillover interventions, including geoengineering proposals, where detection, transboundary spillovers, and provenance would be indispensable [57,58].

5. Running Exemplars

Minimum viable deployment: cool roofs for extreme heat. The smallest falsifiable PCA-OS need not be fully planetary. A first deployment can be city-scale: detect parcel-level cool-roof retrofits from EO albedo/thermal signals and permit text, write them to the ledger, estimate tract-level avoided heat exposure and peak-demand effects with neighborhood spillovers, and optimize subsidy rollout under budget, equity-floor, and bounded-harm rules (Figure 3). This loop is narrow enough for near-term benchmarking and deployment while still exercising the full measurement–causality–decision stack [59–61]. **Flood defenses with spillovers.** Levees, drainage retrofits, and shoreline protection constitute a canonical stress test for interference-aware causal inference because the intervention can redirect water and externalize harm [11]. **Urban greening and nature-based solutions.** Greening can reduce heat and improve resilience, yet it can also trigger green gentrification and unequal benefit capture, forcing joint modeling of cooling, amenity spillovers, and distributive outcomes [12,62].

6. Research Agenda: Three Pillars

Pillar I: Planetary Observability. P1. Measure interventions, not only hazards, across heterogeneous sensors and jurisdictions. **P2.** Build minimal interoperable ledger schemas for type, intensity, mechanism, uncertainty, and provenance aligned with STAC [21]. **P3.** Quantify what the ledger systematically misses, especially informal or politically invisible adaptation. **Pillar II: Causal Topologies and Transfer. P4.** Use foundation-model representations without hiding identifying assumptions or opening new confounding pathways [30]. **P5.** Estimate spillover-aware effects at scale when interventions reshape risk across space or networks [35,36]. **P6.** Propagate measurement, identification, and scenario uncertainty into transportable effect objects rather than single-number treatment effects [13,14]. **P7.** Develop LLM/VLMs as assistive causal-hypothesis engines, with human review and diagnostics kept central. **Pillar III: Normative Optimization and Governance. P8.** Support multiple legitimate equity objectives with explicit trade-offs [43–45]. **P9.** Optimize intervention portfolios that remain defensible under climate uncertainty and model misspecification [39–42]. **P10.** Build benchmarks and governance that can fail maladaptation, privacy violations, climate gentrification, and bounded-harm breaches.

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

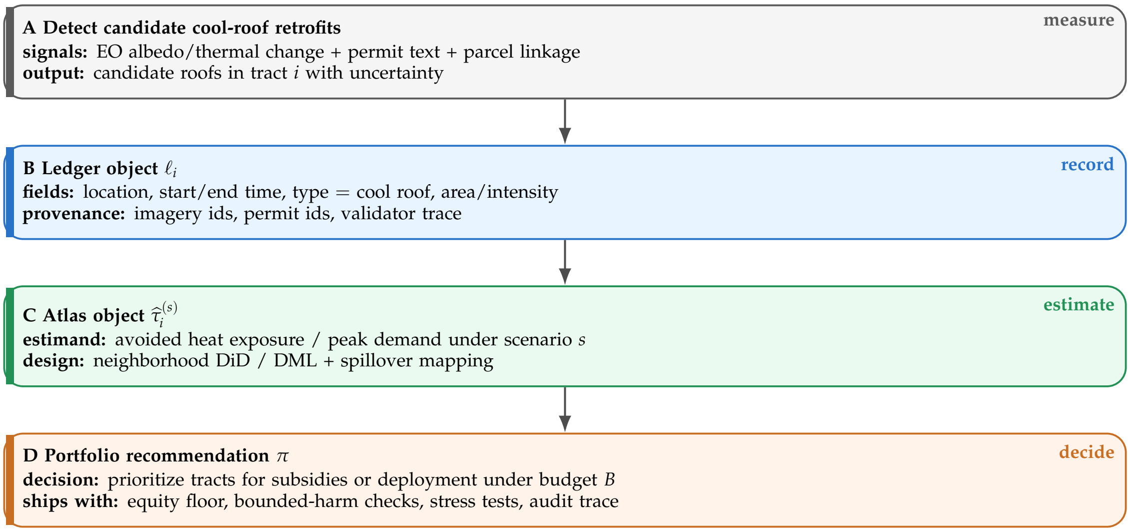

Figure 3. Mini-walkthrough of a viable PCA-OS deployment for urban heat. EO and permit evidence become a cool-roof ledger object, then a scenario-indexed causal effect estimate, then a budgeted and auditable rollout recommendation. A four-step vertical workflow for a cool-roof deployment. The top box detects candidate cool-roof retrofits from Earth observation albedo and thermal changes, permit text, and parcel linkage. The next box stores a ledger object with location, time, type, area or intensity, and provenance. The third box stores a scenario-indexed atlas object for avoided heat exposure or peak demand using a neighborhood causal design with spillover mapping. The bottom box outputs a portfolio recommendation that prioritizes tracts under a budget and includes equity, bounded-harm, stress-test, and audit information.

7. Benchmarks & Evaluation: AdaptBench

We propose AdaptBench as the gauntlet for PCA-OS. It should evaluate the full measurement–causality–decision loop rather than isolated predictors. A first release can contain three task families: **(1) intervention mapping,** where systems produce calibrated ledger entries with uncertainty and provenance from EO, text, and operational signals; **(2) causal estimation,** where systems recover known or semi-synthetic effects under confounding, targeted rollout, and interference while exposing estimands, identification notes, and diagnostics; and **(3) portfolio choice,** where systems select interventions under budgets, equity floors, and no-harm constraints and are scored by robust regret across scenario ensembles. Baselines should span EO segmentation or retrieval and document extraction for mapping, DiD/SC/DML for effect estimation, and robust or fairness-aware optimization for decisions.

The design principle is simple: a model can fail AdaptBench even if its predictive accuracy is excellent. A system that predicts floods perfectly but recommends a defense portfolio that shifts water into poorer downstream tiles should fail; so should a system that identifies cooling opportunities but concentrates benefits in already advantaged neighborhoods. Evaluation should stress generalization across *space* (held-out regions), *time* (future periods), and *policy regimes* (different rollout logics, budgets, and governance conditions). Because dense labels are rare, the suite should combine real retrospective tasks, semi-synthetic interventions on realistic EO backdrops, and known-mechanism decision tasks. Metrics include ledger calibration, interval coverage, spillover error, RDR, inequity gaps, and bounded-harm violation rates. Existing efforts like GEO-Bench [63] demonstrate the value of shared geospatial tasks. Furthermore, domain-specific sustainability benchmarks like ESGenius/MMESGBench [51,52] operationalize multimodal reasoning, but stop short of testing the full intervention–causality–decision loop. AdaptBench extends this logic by making causal validity and decision quality the main scoreboard.

What success would look like. Within a few benchmark cycles, success would include a public ledger schema and reference dataset for at least one intervention family (e.g., cool roofs); atlas tasks requiring estimands, diagnostics, and sensitivity rather than point effects alone; portfolio baselines evaluated on robust regret, inequity gaps, and bounded-harm violations; and KDD papers rewarded for auditable intervention quality.

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

8. Conclusion

PCA-OS is a claim about the next decade of climate adaptation research. Hazard maps are necessary, but they are not the scientific endpoint for climate adaptation. The field needs shared, auditable intervention objects; scenario-indexed causal effect knowledge; and robust, equity-constrained decision layers that can be challenged in public. The payoff is not just better climate analytics. It is a transition from read-only climate intelligence to auditable systems that help societies decide what to do, where, when, for whom, and under deep uncertainty.

References

1. Milly, P.C.; Betancourt, J.; Falkenmark, M.; Hirsch, R.M.; Kundzewicz, Z.W.; Lettenmaier, D.P.; Stouffer, R.J. Stationarity is dead: Whither water management? *Science* **2008**, *319*, 573–574.
2. IPCC. Climate Change 2022: Impacts, Adaptation and Vulnerability (AR6 Working Group II) — Summary for Policymakers. https://www.ipcc.ch/report/ar6/wg2/, 2022.
3. UN Environment Programme (UNEP). Adaptation Gap Report 2024: Come Hell and High Water. https://www.unep.org/resources/adaptation-gap-report-2024, 2024.
4. Global Commission on Adaptation. Adapt Now: A Global Call for Leadership on Climate Resilience. https://gca.org/reports/adapt-now-a-global-call-for-leadership-on-climate-resilience/, 2019.
5. Berrang-Ford, L.; Biesbroek, R.; Ford, J.D.; Lesnikowski, A.; Tanabe, A.; Wang, F.M.; Chen, C.; Hsu, A.; Hellmann, J.J.; Pringle, P.; et al. Tracking global climate change adaptation among governments. *Nature Climate Change* **2019**, *9*, 440–449.
6. Berrang-Ford, L.; Siders, A.; Lesnikowski, A.; Fischer, A.P.; Callaghan, M.W.; Haddaway, N.R.; Mach, K.J.; Araos, M.; Shah, M.A.R.; Wannewitz, M.; et al. A systematic global stocktake of evidence on human adaptation to climate change. *Nature climate change* **2021**, *11*, 989–1000.
7. Ford, J.D.; Berrang-Ford, L. The 4Cs of adaptation tracking: consistency, comparability, comprehensiveness, coherency. *Mitigation and adaptation strategies for global change* **2016**, *21*, 839–859.
8. Ford, J.D.; Tilleard, S.E.; Berrang-Ford, L.; Araos, M.; Biesbroek, R.; Lesnikowski, A.C.; MacDonald, G.K.; Hsu, A.; Chen, C.; Bizikova, L. Big data has big potential for applications to climate change adaptation. *Proceedings of the National Academy of Sciences* **2016**, *113*, 10729–10732.
9. Brullo, T.; Barnett, J.; Waters, E.; Boulter, S. The enablers of adaptation: A systematic review. *NPJ Climate Action* **2024**, *3*, 40.
10. Lesnikowski, A.C.; Ford, J.D.; Berrang-Ford, L.; Barrera, M.; Heymann, J. How are we adapting to climate change? A global assessment. *Mitigation and Adaptation Strategies for Global Change* **2015**, *20*, 277–293.
11. Ansari, A.H.; Mejia, A.; Cibin, R. Flood teleconnections from levees undermine disaster resilience. *npj Natural Hazards* **2024**, *1*, 2.
12. Aram, F.; García, E.H.; Solgi, E.; Mansournia, S. Urban green space cooling effect in cities. *Heliyon* **2019**, *5*.
13. Eyring, V.; Bony, S.; Meehl, G.A.; Senior, C.A.; Stevens, B.; Stouffer, R.J.; Taylor, K.E. Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development* **2016**, *9*, 1937–1958.
14. Riahi, K.; Van Vuuren, D.P.; Kriegler, E.; Edmonds, J.; O'neill, B.C.; Fujimori, S.; Bauer, N.; Calvin, K.; Dellink, R.; Fricko, O.; et al. The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global environmental change* **2017**, *42*, 153–168.
15. Hoffmann, J.; Bauer, P.; Sandu, I.; Wedi, N.; Geenen, T.; Thiemert, D. Destination Earth-A digital twin in support of climate services. *Climate Services* **2023**, *30*, 100394.
16. Doblas-Reyes, F.J.; Kontkanen, J.; Sandu, I.; Acosta, M.; Al Turjmam, M.H.; Alsina-Ferrer, I.; Andrés-Martínez, M.; Arriola, L.; Axness, M.; Batlle Martín, M.; et al. The Destination Earth digital twin for climate change adaptation. *EGUsphere* **2025**, *2025*, 1–41.
17. Szwarcman, D.; Roy, S.; Fraccaro, P.; et al. Prithvi-EO-2.0: A Versatile Multi-Temporal Foundation Model for Earth Observation Applications. arXiv:2412.02732, 2024.
18. Cong, Y.; Khanna, S.; Meng, C.; Liu, P.; Rozi, E.; He, Y.; Burke, M.; Lobell, D.; Ermon, S. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems* **2022**, *35*, 197–211.
19. Lam, R.; Sanchez-Gonzalez, A.; Willson, M.; Wirnsberger, P.; Fortunato, M.; Alet, F.; Ravuri, S.; Ewalds, T.; Eaton-Rosen, Z.; Hu, W.; et al. Learning skillful medium-range global weather forecasting. *Science* **2023**, *382*, 1416–1421.

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

20. Pathak, J.; Subramanian, S.; Harrington, P.; Raja, S.; Chattopadhyay, A.; Mardani, M.; Kurth, T.; Hall, D.; Li, Z.; Azizzadenesheli, K.; et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214* **2022**.
21. Open Geospatial Consortium (OGC). SpatioTemporal Asset Catalog (STAC) Community Standard (OGC 25-004). https://docs.ogc.org/cs/25-004/25-004.html, 2025.
22. Microsoft. Planetary Computer. https://planetarycomputer.microsoft.com/, 2021.
23. NASA Earthdata. Harmonized Landsat and Sentinel-2 (HLS). https://www.earthdata.nasa.gov/data/projects/hls, 2025.
24. Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. The ERA5 global reanalysis. *Quarterly journal of the royal meteorological society* **2020**, *146*, 1999–2049.
25. Jaegle, A.; Borgeaud, S.; Alayrac, J.B.; Doersch, C.; Ionescu, C.; Ding, D.; Koppula, S.; Zoran, D.; Brock, A.; Shelhamer, E.; et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795* **2021**.
26. Girdhar, R.; El-Nouby, A.; Liu, Z.; Singh, M.; Alwala, K.V.; Joulin, A.; Misra, I. Imagebind: One embedding space to bind them all. In Proceedings of the Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2023, pp. 15180–15190.
27. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *Advances in neural information processing systems* **2020**, *33*, 1877–1901.
28. Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F.L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* **2023**.
29. Radford, A.; Kim, J.W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. Learning transferable visual models from natural language supervision. In Proceedings of the International conference on machine learning. PmLR, 2021, pp. 8748–8763.
30. Pearl, J. *Causality*; Cambridge university press, 2009.
31. Imbens, G.W.; Rubin, D.B. *Causal inference in statistics, social, and biomedical sciences*; Cambridge university press, 2015.
32. Abadie, A.; Diamond, A.; Hainmueller, J. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American statistical Association* **2010**, *105*, 493–505.
33. Callaway, B.; Sant'Anna, P.H. Difference-in-differences with multiple time periods. *Journal of econometrics* **2021**, *225*, 200–230.
34. Wager, S.; Athey, S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* **2018**, *113*, 1228–1242.
35. Hudgens, M.G.; Halloran, M.E. Toward causal inference with interference. *Journal of the american statistical association* **2008**, *103*, 832–842.
36. Aronow, P.M.; Samii, C. Estimating average causal effects under general interference. In Proceedings of the Summer Meeting of the Society for Political Methodology, University of North Carolina, Chapel Hill, July, 2012, pp. 19–21.
37. Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; Robins, J. Double/debiased machine learning for treatment and structural parameters, 2018.
38. Dudík, M.; Langford, J.; Li, L. Doubly robust policy evaluation and learning. In Proceedings of the Proceedings of the 28th International Conference on International Conference on Machine Learning, 2011, pp. 1097–1104.
39. Ben-Tal, A.; Nemirovski, A.; El Ghaoui, L. Robust optimization **2009**.
40. Bertsimas, D.; Sim, M. The price of robustness. *Operations research* **2004**, *52*, 35–53.
41. Delage, E.; Ye, Y. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research* **2010**, *58*, 595–612.
42. Lempert, R.; Popper, S.; Bankes, S. Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis **2003**.
43. Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; Zemel, R. Fairness through awareness. In Proceedings of the Proceedings of the 3rd innovations in theoretical computer science conference, 2012, pp. 214–226.
44. Hardt, M.; Price, E.; Srebro, N. Equality of opportunity in supervised learning. *Advances in neural information processing systems* **2016**, *29*.

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

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 March 2026 doi:10.20944/preprints202603.0753.v1

45. Kleinberg, J.; Mullainathan, S.; Raghavan, M. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* **2016**.
46. Haklay, M.; Weber, P. Openstreetmap: User-generated street maps. *IEEE Pervasive computing* **2008**, *7*, 12–18.
47. Tatem, A.J. WorldPop, open data for spatial demography. *Scientific data* **2017**, *4*, 170004.
48. Hogan, A.; Blomqvist, E.; Cochez, M.; d'Amato, C.; Melo, G.D.; Gutierrez, C.; Kirrane, S.; Gayo, J.E.L.; Navigli, R.; Neumaier, S.; et al. Knowledge graphs. *ACM Computing Surveys (Csur)* **2021**, *54*, 1–37.
49. He, C.; Zhou, X.; Wang, D.; Yu, X.; Xiao, L.; Li, L.; Xu, H.; Liu, W.; Miao, C. KG4ESG: The ESG Knowledge Graph Atlas **2026**.
50. He, C.; Zhou, X.; Yu, X.; Zhang, L.; Zhang, Y.; Wu, Y.; Xiao, L.; Li, L.; Wang, D.; Xu, H.; et al. SSKG Hub: An Expert-Guided Platform for LLM-Empowered Sustainability Standards Knowledge Graphs. *arXiv preprint arXiv:2603.00669* **2026**.
51. He, C.; Zhou, X.; Wu, Y.; Yu, X.; Zhang, Y.; Zhang, L.; Wang, D.; Lyu, S.; Xu, H.; Xiaoqiao, W.; et al. Esgenius: Benchmarking llms on environmental, social, and governance (esg) and sustainability knowledge. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, 2025, pp. 14623–14664.
52. Zhang, L.; Zhou, X.; He, C.; Wang, D.; Wu, Y.; Xu, H.; Liu, W.; Miao, C. Mmesgbench: Pioneering multimodal understanding and complex reasoning benchmark for esg tasks. In Proceedings of the Proceedings of the 33rd ACM International Conference on Multimedia, 2025, pp. 12829–12836.
53. Garcıa, J.; Fernández, F. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* **2015**, *16*, 1437–1480.
54. Wachi, A.; Shen, X.; Sui, Y. A survey of constraint formulations in safe reinforcement learning. In Proceedings of the Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 2024, pp. 8262–8271.
55. Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J.W.; Wallach, H.; Iii, H.D.; Crawford, K. Datasheets for datasets. *Communications of the ACM* **2021**, *64*, 86–92.
56. Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I.D.; Gebru, T. Model cards for model reporting. In Proceedings of the Proceedings of the conference on fairness, accountability, and transparency, 2019, pp. 220–229.
57. National Academies of Sciences, E.; Medicine.; et al. *Reflecting sunlight: Recommendations for solar geoengineering research and research governance*; 2021.
58. VanderZwaag, D.L.; Mahamah, A.H. International governance of marine geoengineering: sketchy seascape, foggy future—an essay in honor of Ted L. McDorman. *Ocean Development & International Law* **2024**, *55*, 624–636.
59. Hsu, A.; Sheriff, G.; Chakraborty, T.; Manya, D. Disproportionate exposure to urban heat island intensity across major US cities. *Nature communications* **2021**, *12*, 2721.
60. Cutter, S.L.; Boruff, B.J.; Shirley, W.L. Social vulnerability to environmental hazards. In *Hazards vulnerability and environmental justice*; Routledge, 2012; pp. 115–132.
61. Santamouris, M. Cooling the cities–a review of reflective and green roof mitigation technologies to fight heat island and improve comfort in urban environments. *Solar energy* **2014**, *103*, 682–703.
62. Anguelovski, I.; Connolly, J.J.; Pearsall, H.; Shokry, G.; Checker, M.; Maantay, J.; Gould, K.; Lewis, T.; Maroko, A.; Roberts, J.T. Why green "climate gentrification" threatens poor and vulnerable populations. *Proceedings of the national academy of sciences* **2019**, *116*, 26139–26143.
63. Lacoste, A.; Lehmann, N.; Rodriguez, P.; Sherwin, E.; Kerner, H.; Lütjens, B.; Irvin, J.; Dao, D.; Alemohammad, H.; Drouin, A.; et al. Geo-bench: Toward foundation models for earth monitoring. *Advances in Neural Information Processing Systems* **2023**, *36*, 51080–51093.

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

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