FAB CITY FOUNDATION · RESEARCH INITIATIVE · APRIL 2026

PLANETAI

An instrument for closing the loop between Earth observation and community-scale response.

From observation to action, at community speed. PLANETAI is a five-tier open AI stack that binds Earth observation to the Fab City network — ~2,700 fab labs, 56 signatory cities, a 15-year community-sensor archive — as a standing response infrastructure. Two pre-registered hypotheses, four bioregions, 36 months. The network exists. This funds the wire.

Core Ideas Paper · April 2026

Authors: Tomas Diez, Tomas Vivanco

THE BET

The claim that distributed production reduces urban material throughput has carried a decade of Fab City policy without being measured. PLANETAI is the instrument that tests it — alongside a second claim that has never been tested at all: that observation coupled to a distributed manufacturing network closes the response loop faster than top-down policy. Both will be pre-registered on the Open Science Framework. A null result gets published. Either way, the instrument is built and open.

planetai.fab.city · github.com/fabcity/planetai · Apache 2.0 / OpenRAIL / CC-BY 4.0

1 The premise

Earth already has a lot of eyes on it. Copernicus flies, Google Earth Engine indexes a global tile pyramid, Aurora and GraphCast forecast the atmosphere at 25-km resolution. None of that infrastructure is the problem. The problem is what happens *after* the observation.

A satellite pass tells a European ministry that PM2.5 in a valley crossed a threshold. The ministry feeds it into a policy cycle that reports in 2028, legislates in 2030, permits industrial retrofits in 2032, measures impact in 2035. The observation was excellent. The response cycle took a decade. More critically, even when localised interventions do occur, cities lack the scientific instruments to determine whether they *causally* reduce ecological impact at territorial scales. Existing Earth observation is observation-rich, action-poor, and disconnected from robust causal attribution.

PLANETAI is built for a different loop: community-scale sensing → bioregional decision → distributed production response, cycled in days to months instead of years to decades. The sensing layer is two-tier. The primary signal — load-bearing for H_0-T causal claims — is municipal and district open data (Open Data BCN, Analyze Boston, Santiago Metropolitan Region portals, Bali Satu Data) harvested by four community hubs (one per pilot — Barcelona, Boston, Santiago, Bali) into per-pilot bioregion servers, with sensor campaigns filling gaps where municipal data is silent. Complementary to this, the Fab Lab network's existing crowdsensing instrumentation — SmartCitizen kits with a 15-year archive in Barcelona, OpenAQ's reference-grade station federation, 2,700+ labs on Fablabs.io, and the Fab City registry of 56 signatory cities — operates as a spatial-density layer at ±30–80% uncertainty (Castell 2017), useful for engagement and Agent-1 trigger logic but not load-bearing on H_0-T causal claims. The decision layer is a set of bioregional community councils with CARE sovereignty over the data that rises from their territory. The action layer — the part no Earth observation platform has — is a standing distributed manufacturing network that can actually respond. A fab lab can retool an air-filtration stack, an open-source water-test rig, or a Precious Plastic recycling line in weeks, bypassing municipal procurement cycles. While product-level life-cycle assessments suggest distributed interventions can achieve 40–60% lower cumulative energy demand, systematic city-scale causal evidence remains scarce.

This is the core differentiator. PLANETAI is not another dashboard. It is a **federated AI causal inference instrument**. Its architecture assumes that observation only matters if there is something at the other end of the wire that can act on it, and an AI framework capable of fusing multimodal, hyperlocal data to evaluate that action. By keeping data processing local through bioregional community councils, it solves the privacy bottlenecks of traditional assessments.

With six planetary boundaries already transgressed [Richardson et al. 2023], closing this gap is urgent. The AI-driven framework connecting local actions to bioregional sustainability outcomes has not been built. Moreno and Charnley [2016] argued that re-distributed manufacturing coupled with digital intelligence is the specific mechanism by which throughput reductions become plausible at bioregional scale. A decade later, digital intelligence has arrived. PLANETAI builds the coupling.

---

2 Two pre-registered hypotheses

PLANETAI is anchored by two falsifiable hypotheses, both pre-registered on the Open Science Framework before data collection begins. Three outcomes — strong positive, partial positive, null — are possible for each. All three are committed to publication. The pre-registration is the instrument's own peer review.

2.1 H_0-T — the causal throughput hypothesis

> **Null.** Distributed production at community and city tiers *does not* causally reduce urban material throughput or energy demand by more than **15%** over a 36-month intervention window, measured against a synthetic control baseline.

Three outcomes are possible: a strong positive in the **40–60%** range suggested by recent product-level LCAs; a partial positive in the **15–20%** range; or a null. We commit publicly to publishing all three. As empirical validation remains limited and context-dependent [Angstmann et al. 2025], a null would be one of the most consequential findings in circular-economy policy in the last decade. We are prepared to publish it if that is what the instrument returns.

2.2 H_0-A — the action-latency hypothesis

> **Null.** Coupling multimodal observation to a distributed manufacturing network *does not* measurably shorten the response cycle time — from community-scale detection of an environmental anomaly to a deployed intervention — relative to a policy-only baseline.

No existing Earth observation platform tests this, because no existing platform has an action layer at the other end. We measure it at **four community hubs — one per bioregion** (Barcelona, Boston, Santiago, Bali) — each anchored to a defined district, against matched non-intervention districts using difference-in-differences estimators. The scientific novelty is the causal measurement; the impact is that the measurement is attached to a response network that acts at bioregional speed.

Priors from the literature. The action-latency claim does not arrive without evidence. The COVID-19 maker mobilisation produced documented detection-to-prototype cycles of hours — Isinnova in Brescia reverse-engineered and 3D-printed a missing Venturi-valve prototype within six hours of a hospital running out of supplier stock on 13 March 2020, with 100 working units in 24 hours and patients on the device the following day — and detection-to-territory-deployment cycles of weeks — Barcelona's Ateneus de Fabricació, declared an essential service in late March 2020, jointly met PPE demand for all hospitals in the city and metropolitan area by the third week of April [CIDOB 2020; Tabarés et al. 2021]; Maker's Asylum's M-19 collective shipped 1,073,704 face shields across 42 Indian cities and villages in 49 days through 21 design iterations and a switch from 3D-printing to laser-cut sheets [Open Source Medical Supplies 2020; Corsini et al. 2021]. The peer-reviewed cross-case literature [Corsini, Dammicco & Moultrie 2021; Kieslinger et al. 2021; Tabarés et al. 2021; Hamdi-Cherif et al. 2021] reads these as innovation-community mobilisations whose speed was cashed out from years of fab-lab social infrastructure rather than built under crisis pressure. Outside crisis time, the distributed-prosthetics literature [Hofmann et al. 2016, 2017; Buehler et al. 2015; Zuniga et al. 2015; Yaqin et al. 2024] shows the same pattern at lower stakes: e-NABLE has delivered an estimated 10,000–15,000 upper-limb devices through ~40,000 volunteers across 100+ countries; the PPDFI Surabaya 3D-printing community-service programme demonstrates the institutional surface in Indonesia; Yayasan Peduli Tuna Daksa runs a national-scale prosthetic-distribution operation across the archipelago. What the literature establishes is detection-to-prototype and detection-to-territory-deployment speed for low-regulatory-burden artefacts; what it does not establish is sustained-production-at-scale or outcome against a controlled policy baseline. No study in this body of work measures action-latency reduction against a matched policy-only counterfactual. That is the gap H_0-A is built to fill.

The hypothesis is operationalised through a **response coefficient $\rho \in [0,1]$**, computed per tier as Coverage × Speed — Coverage being the fraction of pre-registered observations that produce a fitted, human-approved response within the tier's latency budget, Speed being the median Δt of those fitted responses normalised against a pre-registered policy-baseline latency. ρ is reported separately at every tier, never averaged without showing components, and a null result on ρ is publishable. Per-tier protocol — fitted-response criterion, approver authority, latency budgets, sub-latency decomposition — is documented in the methodology annex.

Pre-registration locks the project's reporting commitments before data collection begins and rules out post-hoc selection on the response coefficient.

---

3 Five tiers, one instrument

The scientific scaffolding underneath the two hypotheses is the Fab City Full Stack [Diez, Niaros & Ferro 2024], building on the PI's earlier civic-platform articulation [Diez 2012] and operationalised through the **Full-Stack Metrics 4-Pillar × 5-Scale framework** — the five scales at which distributed production and its measurement operate: **Planet, Bioregion, Region, City, Community**. These tiers are not administrative conveniences. They are the granularity at which different actors have authority over different decisions. A Community observation sits with the council. A City observation sits with the mayor. A Region observation is the supply-chain and inter-municipal governance layer. A Bioregional observation is governed by the bioregional consortium. The topology of the five tiers *is* the governance model.

Planet tier (2,048-d). Federated aggregation across open-weight Earth foundation models — Aurora [Bodnar et al. 2025] for atmospheric physics, GraphCast [Lam et al. 2023] for 10-day deterministic forecasts, GenCast [Price et al. 2024] for probabilistic tail-event risk. Surface-tier complement: Sentinel-2 + NASA-IBM Prithvi (open weights, 10–30 m land-cover embeddings) as the always-on capability; Google Earth AI products (Earth Engine, AlphaEarth Foundations, WeatherNext, Population Dynamics) extend capability where accessible. We are not training a planet-scale model. We are specialising existing open models to the four-pillar measurement task. The architecture is fully open-weight throughout — permissionless to replicate, with the multi-vendor topology as additional sovereignty: PLANETAI cannot be held hostage to a single foundation-model provider.

Bioregion tier (1,024-d). Four partner-hosted servers — IAAC Barcelona, MIT Center for Bits and Atoms Boston, Universidad Católica de Chile Santiago, IT Del Bali. Each runs Weaviate (vector retrieval), Neo4j (production-actor knowledge graph), TimescaleDB (time series). The bioregion server is where the community council's governance veto is technically enforced. Nothing leaves upward without a signed approval on that server. Bioregion is defined ecologically (watershed, ecoregion) rather than administratively — the framework's tier with the longest time horizon.

Region tier (512-d). Multi-city / state / province scale. The tier where multi-municipal governance and supply-chain coordination operate, and where most regional open-data publishers sit — Generalitat de Catalunya for Barcelona, Massachusetts state-level data for Boston, Metropolitan Region of Santiago for the Chilean pilot, Bali Province for Bali. The Full-Stack Metrics framework places the strongest signals here for inter-city material exchange, regional energy independence, supply-chain resilience, joint procurement, and regional policy alignment. PLANETAI's federated open-data discovery layer (e.g., the OpenData.cat MCP for Catalunya, with the same pattern targeted for the other three

pilots) operates at this tier.

City tier (256-d). Single-municipality dashboards, Metroverse Economic Complexity Index [Hidalgo & Hausmann 2009; Hausmann et al. 2014], Boeing's [2024] Fab City Index sectoral self-sufficiency score (Hamburg piloted at 37/100), and the Amsterdam Circular Monitor pattern [City of Amsterdam 2022]. The city tier is where a mayor or planner interacts with the instrument. For Bali, where Metroverse has no Denpasar coverage, the City-tier throughput signal is reconstructed from Bali Satu Data (provincial open data published under a Tri Hita Karana taxonomy) — same PITO/DIDO/ρ structure as every other pilot, different ingest path.

Community tier (64-d). Neighborhood scale (1,000–50,000 people per the Full-Stack Metrics framework definition). **Four community hubs — one per pilot** (Barcelona, Boston, Santiago, Bali), each a fab lab or equivalent, each running a PLANETAI node — a compact general-purpose computing surface whose specifics are deferred to the hardware annex. Hubs are not measurement targets; they are (i) **harvesters** of municipal/district open data (Open Data BCN, Analyze Boston, Santiago Metropolitan Region portals, Bali Satu Data), (ii) **sensor-campaign operators** filling gaps where municipal data is silent — ≥30-day colocation calibrations against reference instruments plus longer-term campaign deployments — and (iii) **local-intelligence instantiators** running quantised SLMs over federated weights with Matryoshka embeddings. The treatment unit for community-tier H_0-T is the **district anchored by the hub**, with synthetic controls built per pilot from similar non-PLANETAI districts. Community scientists trained as Fab Academy Data Champions, integrated into the existing curriculum. This is where the instrument makes contact with the ground, where community sensors (±30–80% uncertainty against reference instruments [Castell et al. 2017]) require colocation calibration, and where the Plaça del Sol pattern of campaign-bounded measurement-to-action runs [Balestrini, Diez et al. 2015]. The PKC framework [Koo 2025] sits adjacent as the *individual-scale antecedent* PLANETAI extends upward — it is not a separate measurement tier; PLANETAI's smallest measurement unit is the community.

The binding mechanism across all five tiers is **Matryoshka Representation Learning** [Kusupati et al. 2022], combined with quantised Small Language Models (e.g., Llama 3B) and Flower-based federated learning with causal discovery. MRL produces embeddings where the first m dimensions are as representative as an independently trained m-dimensional model — the same trained embedding space runs on cloud-tier compute, a bioregion-tier server, and a community-tier node, each tier using a smaller slice, all slices aligned. Raw observational data stays at its source; model updates flow upward. Privacy-preserving federated causal inference across four distinct bioregions becomes computationally feasible.

4 The node — PLANETAI as deployable open hardware

The five-tier stack is not a web service. It is an instrument, and the unit of deployment is a **node**: a standardised piece of open hardware that a fab lab, bioregional partner, or community hub plugs in and runs. The observatory is the node's cloud face. The node itself is what materialises the instrument in the territory.

A PLANETAI node does four things at once. It subscribes to upstream planetary-tier inferences — Aurora, GraphCast, GenCast atmospheric; Sentinel-2 + Prithvi surface; Earth AI products where accessible. It fuses them with the locally-curated sensor stream on the LAN — community air-quality kits, environmental HAT readings, pilot-specific add-ons deployed per campaign rather than as a continuous fleet. It serves the observatory interface locally to the lab and its bioregion. And it enforces data sovereignty at the box: raw observational data does not leave the node unless the community council explicitly federates it. The node is the technical instantiation of sovereignty-by-architecture [PKC, Koo 2025] at Community scale, one tier up from the individual-scale antecedent.

Node v0 specification.

Layer	Component
Compute	Compact general-purpose node · storage-rich · optimised for asynchronous federated workflows, not latency-sensitive edge inference · spec'd in the hardware annex
Sensors	Campaign sensor kit · environmental HAT · pilot-specific add-ons · deployed per campaign, not as continuous fleet
Runtime	Docker Compose — local observatory UI, MRL embedding layer (64–512-d community / city / region slices), message broker for sensor ingest, cached tier outputs, agent runtime
Federation	LAN-local discovery; bioregion-server federation over an opt-in overlay

Layer	Component
Enclosure	Fab-labable case, BOM published under Apache 2.0 / CERN-OHL, OSHWA-track certification

Four anchor nodes ship to the four bioregion pilots in year one. The BOM is published open so additional labs can build their own — the Fab Lab network, not the project team, is the distribution mechanism. Specific compute and sensor selections are documented in the hardware annex, with the spec calibrated to the federated-data workload (storage-bound, asynchronous, locally-curated) rather than to latency-sensitive edge inference.

Planet-tier foundation models are cloud-native; the node does not replicate them. It subscribes to tier outputs and performs local fusion and embedding at the Community (64-d), City (256-d), and Region (512-d) tiers. The cloud is where Aurora runs. The node is where Aurora's output becomes actionable at community speed. **The node is a federated-data layer, not a datacentre and not an edge-AI box** — it retains local data, caches what the bioregion needs, and pairs upstream Earth-foundation embeddings with locally-curated observations.

Federation protocol (v0). Nodes announce themselves to the bioregion server via a Fab City node-registry endpoint. Each node writes MRL embeddings (not raw data) to a weekly diff that the bioregion server ingests *after* community-council approval. A node that goes offline retains all local work; reconnection triggers a delta sync with no data loss. The protocol is thin by design — a node should remain useful even if the upstream federation layer fails.

The node is the project's most concrete unit of leverage into the Fab City network. It turns a 56-signatory-city pledge into a 56-potential-node deployment surface.

---

5 Four-pillar metrics framework

PLANETAI measures the Fab City Full Stack through the **Full-Stack Metrics 4-Pillar × 5-Scale framework**: **Environmental, Social, Economic, Governance**. These four pillars apply at every one of the five scales (Community, City, Region, Bioregion, Planet) and are the axes of the Production Constellation — a radar chart that operates at any tier. Every indicator collected is placed in a pillar and a scale. The framework's working-paper articulation [Vivanco 2024] and the bioregional material-mapping fieldwork across four Chilean macrozones [Vivanco 2025] are the two reference artifacts PLANETAI builds on.

The Fab City Index lineage — three generations. PLANETAI is the third generation of the Fab City Index. **Generation 1** [Florentin, Chabanel & Guimas 2018, Utopies & FabCity Paris] ran a self-sufficiency index across ~600 French urban areas using 257 sectors aggregated to 12 macro-sectors, scored on a priority × self-sufficiency lattice and aggregated 0–100 via Utopies' proprietary LOCAL SHIFT® urban-economy simulator. Paris topped the ranking at **37.58/100**; the average French city produced 3.1% of what its population consumed (€3 of every €100); 95% of French cities scored under 10. The accompanying LOCAL FOOTPRINT® analysis showed that 85% of the average French metro's raw-material footprint is extracted abroad. **Generation 2** [Boeing 2024, Springer Open Access] re-implemented the same construct under European-statistical-system rigour — 16 macro-sectors with a NACE Rev. 2 × COICOP concordance, public open data only — and pilot-tested it in Hamburg, returning **37/100**. The convergence between Paris (37.58) and Hamburg (37) — six years apart, two statistical systems, the same number — is not a city-specific finding. It is the public-data ceiling for any major Western metropolis. **Generation 3** (PLANETAI) extends that single City × Economic cell into the four-pillar × five-scale matrix described below, federated across four bioregions, with a coupled action layer that makes response cycle time a measurable quantity for the first time. The expert-baseline panel for the H_0-T synthetic-control arm combines the Utopies 2018 priority × self-sufficiency methodology applied to peer cities, the Chilean-macrozone bioregional material-mapping baseline [Vivanco 2025], and a Fab Academy reviewer panel. Where formal NACE/COICOP data does not apply (Bali / Indonesian KBLI), the City × Economic cell is reconstructed from Bali Satu Data through the Tri Hita Karana ↔ four-pillar mapping described below — same PITO/DIDO/ρ formulation, different ingest path.

The narrative layer — PITO, DIDO, and the response coefficient ρ. The four-pillar × five-scale matrix is the underlying instrument. The narrative layer that translates twenty cells into Fab City vernacular returns to the founding state-diagnosis vocabulary [Diez 2016]: **PITO — *Products In, Trash Out*** (the linear-extractive metabolism — what the city receives and rejects: imports, energy, food, raw materials in; waste, emissions, externalities out) and **DIDO — *Data In, Data Out*** (the regenerative-distributed metabolism — what the city generates and circulates: open data, fab-lab capacity, distributed manufacturing, recycling infrastructure, transparency, community sensing). Each cell of the 4×5 matrix carries a documented PITO weight and DIDO weight summing to 1 — throughput cells (most Environmental, Economic-export-side) lean PITO; capacity cells (Governance, Social, Economic-making-side) lean DIDO; a few cells split. The Fab City Index is then defined as $\mathbf{FCI_t = DIDO_t \cdot (1 - PITO_t) \cdot \rho_t}$ in snapshot form, where ρ_t is the response coefficient that operationalises H_0-A. The trajectory $\mathbf{\Delta FCI/\Delta t}$ is what PLANETAI is built to measure across its first 36-month measurement window. Boeing's 37/100 and Utopies' 37.58/100 fall out as the single-cell, ρ-implicit, snapshot projection

at *Economic* × *Region* — same instrument, less resolution. The product form (not sum) is deliberate: high DIDO with high PITO is performative (fab labs without metabolic shift); fast ρ on a small set of pre-registered observations is selection bias controlled by pre-registration discipline. The full aggregation table, recovery derivation, and per-tier ρ protocol are in the methodology annex. Every reported ΔFCI/Δt carries propagated CI bounds — DIDO/PITO from official-statistics revision tolerances (~5–10%), ρ from agent-count-bounded variance (~10–20%); crowdsensing noise does not enter FCI computation.

Environmental. Biophysical baseline, material flows, and ecological limits. Data hierarchy is explicit: official statistics and reference-grade instruments (Eurostat EW-MFA, OpenAQ FRM/FEM stations, Sentinel-5P, AMB PREMET25, MassDEP, Bali Satu Data, *What a Waste 2.0*) are the **primary** signal layer for causal-attribution claims; crowdsensing networks (SmartCitizen, OpenAQ AirGradient) sit as a **complementary** spatial-density layer with ±30–80% uncertainty (Castell 2017) — useful for engagement and Agent-1 trigger logic, not load-bearing on H_0-T causal claims. Operationalised through this layered air-quality stack, climate risk (GenCast probabilistic tail events), biodiversity density, urban heat, and the imports→waste Sankey triptych at three scales simultaneously (Planet, Bioregion, Local). Three methodological scaffolds anchor this pillar: Hoornweg et al.'s [2016] city-scale planetary boundaries (the direct ancestor of the radar-chart form), Richardson et al.'s [2023] six transgressed planetary boundaries (why the pillar is non-negotiable), and the Amsterdam Circular Monitor pattern [City of Amsterdam 2022] for material mass × ecological impacts across food, textiles, electronics, construction, plastics, metals (whose finding — material use 15–61× higher than previously estimated, Scope 3 = 86% of city emissions — sets the scale of the observability gap PLANETAI addresses). The Fab City Index lineage [Florentin, Chabanel & Guimas 2018; Boeing 2024] is operationalised here as the city-tier sectoral self-sufficiency proxy for material consumption, complemented at the bioregion tier by Mealy & Teytelboym's [2022] Green Complexity Index and the distributed-manufacturing / urban-circulation-hub simulation thresholds [Martin et al. 2024] — distributed manufacturing outperforms centralised above 6% overproduction, and outperforms with a local materials-recovery hub above a 12.5% recycling rate. These are the quantitative thresholds behind the throughput-reduction hypothesis (H_0-T).

Social. Health, equity, employment, education, community cohesion. This is the pillar PLANETAI's Community tier is built for. Operationalised at community scale through fab-lab participation logs (workshop attendance, skills training delivered, vulnerable groups included as % of participants), Fab Academy archive analytics (skills + diversity), and PM2.5 health-outcome attribution from the Environmental pillar's air-quality data. At city and bioregion scale: citizen engagement / electoral participation, life expectancy, affordable housing share, digital inclusion. At Community tier in Santiago and Bali, ~75%

of urban employment is informal [ILO 2023] and invisible to formal-sector accounting; PLANETAI fills that gap by measuring participation through fab-lab and community-council channels rather than tax-registry-based labour statistics. The PKC framework [Koo 2025] enters this pillar as the individual-scale sovereignty antecedent — a Social pillar concern that PLANETAI extends upward through architecture.

Economic. Value creation, circular-economy jobs, productive capability, supply-chain localisation, R&D, innovation. This is the cell the Fab City Index lineage above directly populates. At City tier, PLANETAI scores the Economic pillar Boeing-style: 0–100 self-sufficiency on a NACE × COICOP-concordant sector breakdown using public open data, with the Utopies priority × self-sufficiency 2×2 logic supplying intervention targeting (high priority + low self-sufficiency = red-zone sectors — the methodology that identified 40 strategic priority sectors across the 50 largest French cities in 2018). At City tier PLANETAI also draws on Harvard Growth Lab Metroverse's Economic Complexity Index and Industry Space [Hidalgo & Hausmann 2009; Hausmann et al. 2014], which covers Barcelona, Boston, and Santiago directly and provides peer-comparison via cosine distance. At Community tier, the pillar is operationalised through distributed-manufacturing indicators that neither the Utopies nor the Boeing index can see: fab-lab catchment, materials processed locally, active skills density, project documentation output. The Community-tier indicator family is PLANETAI's original contribution at this pillar.

For Bali, where Metroverse does not cover Denpasar, the City-tier Economic pillar is built on two coherent sources feeding the same PITO/DIDO/ρ taxonomy every other pilot uses: **Bali Satu Data** (4,700+ sovereign datasets published by Tim SPBE Diskominfos under Pergub Bali 53/2021, organised through Tri Hita Karana / Balinese cosmological categories — Jana, Atma, Jagat, Wana, Danu, Segara) for throughput at City and Region tiers, and the **WIPO Global Innovation Index 2025** [WIPO 2025] Indonesia profile (rank 55/139, climbing 30 positions since 2020) as the country-level peer-comparison anchor at the Bioregion tier. Engineering ingests data via the national Satu Data Indonesia portal (data.go.id, CKAN 3.0) and scrapes Bali Satu Data directly; partner collaboration with Diskominfos for a sovereign API endpoint is a year-one Bali deliverable. Community-tier DIDO inputs (fab-lab catchment, Fab Academy throughput, Precious Plastic activity density, SmartCitizen sensor density, banjar council activity) feed the same DIDO computation as every other pilot — no Bali-specific index. The **Tri Hita Karana ↔ four-pillar mapping** is a **Bali pilot only co-design experiment** between CAST Foundation, MDG, IT Del and a named Balinese customary authority (banjar adat, desa adat, or PHDI scholar — confirmed in v2); **publication is conditional on customary-authority sign-off**, not a methodological export. The Bali pilot is the operational test of PLANETAI's pluralistic measurement claim: the same indices, computed over sovereign provincial data published under a non-Western cosmological taxonomy.

Governance. Institutional capacity, transparency, participation, coordination — measured directly, not deferred to a sovereignty narrative. The pillar's operationalisation in PLANETAI is **sovereignty-by-architecture**: CARE principles (Collective Benefit, Authority to Control, Responsibility, Ethics) enforced *by code at the bioregion-server boundary*, not by procedure. Concrete signals: ODIN-style open-data maturity score, % of observatory sensors governed under CARE-aligned protocols, active Fab City Pledge programs, Fab City OS Sovereignty Index, presence and exercise of the community-council veto (a binary instrumented at the bioregion server). The PKC framework [Koo 2025] provides the individual-scale antecedent that PLANETAI extends upward through five tiers; the community council's veto is not a consultation procedure but a technical constraint at the bioregion server.

---

6 The 2054 Pledge — temporal benchmark

The Fab City Index without a deadline is a measurement; with one, it is a check on a commitment the network has been carrying for over a decade. At FAB10 in Barcelona, **6 July 2014**, the mayor of Barcelona invited cities to join the **Barcelona Pledge**: a commitment to become at least **50% self-sufficient by 2054** [Diez 2016]. Forty years from launch to target. Twelve years elapsed by April 2026; twenty-eight years remaining. The pledge is what gives the trajectory **ΔFCI/Δt** a benchmark — without it, the index is a free-floating measurement; with it, FCI(t) is a curve compared against a network-wide commitment dated, signed, and publicly held since 2014.

The pledge as Index-target. The 50% self-sufficiency threshold maps onto the Fab City Index as the *crossover state* at City × Economic, where Boeing-style self-sufficiency (the ratio that recovered Hamburg at 37/100 and Paris at 37.58/100 in 2024 and 2018) reaches 0.5. In PITO/DIDO terms, this is the point at which DIDO_t and (1 − PITO_t) at City × Economic become balanced — the regenerative metabolism is at least equal to the extractive one. Aggregating across the four-pillar × five-scale matrix with documented weights, the equivalent **FCI_pledge** target is in the methodology annex; the headline framing is the Boeing 0.5 crossover at the Index's signature cell.

Pledge progress and schedule slip. Two derived quantities fall out and become reportable per pilot. **Pledge Progress %** = (FCI_t − FCI_2014) / (FCI_pledge − FCI_2014), where FCI_2014 is the pre-PLANETAI baseline reconstructed retrospectively from 2014 public data — the year the pledge was launched and the start of the temporal benchmark. **Schedule slip** = (years remaining at current ΔFCI/Δt to reach FCI_pledge) − (years actually remaining to 2054). Negative slip means the pilot is ahead of pace; positive slip

means behind. At April 2026, the linear-progress benchmark is 30% of the 40-year horizon elapsed; a pilot reporting Pledge Progress < 30% is behind the linear pace, > 30% is ahead. Both quantities are reported with the same provenance discipline as the pillar scores: live where data supports it, and clearly labelled as indicative reconstruction where it does not.

What the first 36 months contribute. This window covers months 144–180 of the 480-month pledge horizon. ΔFCI/Δt measured across the four pilots over this window is the first systematic, federated, multi-bioregion test of *whether coupled observation-to-action moves the trajectory* against the policy-only baseline.

H_0-T is staged across two delivery tracks. (i) **Community-tier H_0-T** — operational throughput at neighbourhood/district scope (1,000–50,000 people per Vivanco) measured on hub catchment data: municipal waste-collection records, distributed-production material flows, localised energy draw. Fully powered on weekly/monthly cadence. This is the in-window deliverable. (ii) **Metropolitan-scale H_0-T** — synthetic-control on official statistics (Eurostat EW-MFA, Atlas of Economic Complexity, Bali Satu Data, *What a Waste*). Pre-registered on OSF before any data collection; the attribution paper is drafted in-window and published post-window on month 60+ data when both data lag (1–2 years) and statistical power (n=4 treated bioregions × donor-pool sparsity) improve. Both tracks committed to publication including null outcomes, with explicit power analysis for either result. A null on H_0-T or H_0-A within this window is consistent with the pledge requiring a different mechanism; a positive signal makes the coupling architecture the most empirically supported strategy for closing the remaining 28-year horizon. The pledge frames the question; PLANETAI is the instrument that measures the answer.

---

7 Action agents — closing the observation-to-action loop

The Action pillar is where H_0-A is instrumented and where the response coefficient ρ is computed in the field. An **action agent** is a purpose-built, human-in-the-loop software workflow that runs on the node, consumes multi-tier observational data, and proposes a specific intervention to a designated human approver at the community hub. Each agent is scoped to a single decision loop, pre-registered on OSF with its own observation threshold, fitted-response definition, approver authority, and latency budget, and logged step-by-step so that Coverage and Speed (and from them ρ) become falsifiable quantities at quarterly resolution.

Agent-1 — Bali air-quality response. SmartCitizen sensors within a 5 km radius of Fab Lab Bali record PM2.5 above a community-agreed threshold for more than 60 minutes.

The agent pulls the regional Aurora PM2.5 forecast, drafts a bilingual community advisory in Bahasa Indonesia and English, surfaces a shortlist of OSHWA / OKH open-hardware air-filter designs fabricable on the lab's current inventory, and queues the package for human approval at the Fab Lab Bali admin interface. Every step's timestamp is logged. The year-one target is to establish a baseline median latency against the existing provincial air-quality policy response. Success is not a particular latency number — success is that the measurement exists.

One agent per pilot.

Agent	Pilot	Trigger	Output	ρ_community sub-latency clock
Agent-1	Bali	PM2.5 threshold crossed in 5 km radius (≥60 min confirmation window)	Bilingual advisory + OSHWA air-filter shortlist	Sensor-detection → approved advisory
Agent-2	Barcelona	Imports→waste Sankey anomaly (material-category spike)	OSHWA / OKH alternative designs from local fab-lab inventory	Import-spike → procurement-alternative queued
Agent-3	Santiago	Weekly cadence	Municipal-policy brief with informal-economy inclusion flag [ILO 21st ICLS 2023]	Data-stream update → municipal-brief draft
Agent-4	Boston	Sensor-drift signal from calibration tracker	Replacement-part fab file + maintenance ticket	Drift-signal → fab-file delivered

Each agent maps to its pilot's specific instrumentation strength and contributes a distinct

ρ_community sub-latency to the H_0-A test. For the first 36 months of measurement, the four Community-tier agents are the primary ρ instruments; City, Region, Bioregion, and Planet ρ instruments come online in subsequent waves with pre-registered tier-weighting protocols.

Responsible-AI protocol. Agents never act autonomously on the physical world. Every agent output is draft → queue → approve. The human-in-the-loop gate is a technical requirement of the agent runtime, not a recommended practice. Audit logs record the full observation-to-approval chain and are published under the same open licence as the rest of the instrument. Community-council authority over agent deployment is enforced at the bioregion-server boundary — a council can disable any agent in its territory at any time without consulting the central team.

Not a product suite. The agents are scoped as measurement instruments for H_0-A, not as a standalone product line. If the first four agents demonstrate measurable latency reduction against the policy baseline, the instrument earns a platform framing; if they do not, the project publishes the null with the same pre-registration discipline that binds the throughput hypothesis.

---

8 Methodological foundations

Six methodological pillars ground PLANETAI in existing research. Each names a body of prior work the instrument extends rather than replaces.

8.1 1. Economic complexity (Hausmann; Harvard Growth Lab)

Metroverse [Harvard Growth Lab 2021–] covers 1,000+ metropolitan regions in 79 countries, including Barcelona, Boston, and Santiago directly. The "Similar Cities" cosine-distance methodology is directly usable for peer-comparison at the City tier. PLANETAI's extension is anchored to two gaps in Metroverse: Santiago and Bali's informal economies (~75% of urban employment) are invisible to formal-sector data; Bali's anchor city has no Metroverse coverage at all. PLANETAI names both gaps explicitly and instruments around them. The bioregional peer-matching extension — combining industry-space, material-flow, and ecological vectors in a single 1,024-d embedding — is PLANETAI's original methodological contribution at this layer, to be published as a companion paper.

8.2 2. City Anatomy (IAAC / City Protocol Society)

Taxonomic backbone of the coverage matrix: the City Anatomy framework [Guallart et al. 2015], originated at IAAC and published as an IEEE-compatible open standard. Struc-

ture × Interactions × Society with 155 indicators extending ISO 37120. PLANETAI's ten phenomena map cleanly onto nine City Anatomy subsystems. Where City Anatomy has zero ISO indicators — Built Domain at all sub-city scales, distributed production, Scope 3 consumption, AI-governance observability — PLANETAI extends the taxonomy rather than silently assuming coverage. The methodological inheritance is named explicitly so PLANETAI's contributions sit on top of an existing standard rather than alongside it.

8.3 3. City Science and CityScope (MIT Media Lab, Larson)

Agent-based urban simulation with tangible interfaces. Two City Science Network labs sit inside PLANETAI pilot geographies. CityScope validates the UX principle for city-tier deployment: municipal partners need spatial simulations they can interact with, not charts they read.

8.4 4. Supply-chain empirics (Amsterdam, Barcelona, ILO)

The Amsterdam Circular Monitor [City of Amsterdam 2022] is the operational template for city-scale material-flow observability done rigorously. Amsterdam established the 15–61× under-estimation factor and the 86% Scope 3 finding that PLANETAI takes as motivation rather than as findings of its own. Barcelona's Zero Waste Plan 2021–2027 provides the Mediterranean-bioregion policy frame. ILO's 21st ICLS [2023] provides the Santiago and Bali informal-economy indicator backbone.

8.5 5. Hyperlocal scale and measurement-to-action

Seven empirical studies establish the grounding for the Community tier. The key insight across them: observation only matters if it produces decisions. The measurement-to-action pipeline:

Step	Reference	Tier	Mechanism
Observe bio-physical limits	Hoornweg et al. 2016	City	Planetary-boundary radar → safe operating space
Observe causal structure	Wu 2026	City / Bioregion	Causal discovery → which upstream drivers actually matter

Step	Reference	Tier	Mechanism
Identify evidence gap	Angstmann et al. 2025	Community	GUM SLR: 60% of claims lack empirical grounding
Model production system	Martin et al. 2024	Community	Distributed-manufacturing thresholds at 6% and 12.5%
Model circular logistics	Marmiroli et al. 2020	City / Community	EV LCA: electricity mix determines logistics gains
Optimise local production	Luo & Ball 2025	Community	DT + Q-learning: 78.5% demand fulfillment, 15% electricity savings
Policy lever	He et al. 2023	City / Bioregion	Manufacturing–services co-location drives 28.8% carbon-productivity gain

This pipeline instantiates H_0-A at each tier.

8.6 6. Sovereignty-by-architecture (PKC, Koo 2025)

The sovereignty-by-architecture pattern that PLANETAI operationalises at the Community-and-bioregion tiers has a direct antecedent in the **Personal Knowledge Container** framework [Koo 2025], which establishes a scale-free, self-sovereign data infrastructure spanning individual to community to public-network scales. PKC's architecture is **triadic**: raw data is stored as immutable, content-addressed records (locally controlled, cryptographically verified); a mediation layer transforms and routes that data without exposing the underlying records; and an authorisation layer enforces governance access control before any data is released upward.

PLANETAI extends this triadic pattern through the five Vivanco tiers:

- Raw observational data stays at its source — sensor stream, fab-lab telemetry, council records — as an immutable community-controlled record at the node.

- Matryoshka Representation Learning embeddings (64–2,048 d, sliced per tier) constitute the **mediated transformation layer** — the data is reshaped without exposing the underlying records.
- The community council's federation approval, technically enforced at the bioregion server, functions as the **governance authorisation gate** before any upward flow.

PKC's use of content-addressable storage, decentralised identity, and locally-operable open-source tooling directly parallels the approach at PLANETAI's Community and City tiers. **PKC is operational at IT Del** — PLANETAI's Bali bioregion anchor — making this architectural continuity operational rather than rhetorical, and deployable across Indonesia's 17,000+ islands. Ben Koo (IT Del), PKC's architect, co-anchors the Indonesian bioregional consortium and is co-author on the sovereignty-architecture sections of the methodology. PKC sits in the methodological-foundations stack at the same level as the five pillars above: not as a passing reference but as the scale-free sovereignty pattern PLANETAI scales upward through architecture.

8.7 The methodological convergence: why this instrument is possible now

Peuckert et al. [2025] reviewed approximately 1,000 quantitative impact studies on Fab Labs using the Fab City Full Stack as the analytic frame. Their finding is precise: quantitative evidence is heavily concentrated in Layer 2 (skills and learning) and Layer 3 (entrepreneurship). **Layer 6 — bioregional strategies — recorded zero rigorous quantitative impact studies in the strictest classification.** They identify why: bioregional and systemic impacts involve multiple stakeholders, emergent outcomes, and long timescales that resist the individual-level surveys and pre/post tests that dominate the evidence base. Their proposed remedies are explicit — "participatory evaluation, geospatial analysis, and longitudinal mixed-methods research."

Liu et al. [2024] describe exactly what AI now makes possible in urban observing: foundation models fusing satellite imagery, street-level sensing, and community-generated data at city and sub-city scales; real-time geospatial analysis without centralised data aggregation; multimodal interpretation that converts raw sensor streams into actionable indicators. These are not speculative capabilities — Aurora, GraphCast, and SmartCitizen are already operational. What has been missing is their coupling to a governance architecture that allows community data to participate without surrendering sovereignty.

PLANETAI is the intersection. Peuckert names the gap at Layer 6. Liu names the toolkit that makes that layer measurable. PLANETAI builds the instrument that deploys the toolkit specifically at the tiers Peuckert's review confirms have never been rigorously measured, across four bioregions, with federated AI ensuring sovereignty at the tier that matters most.

9 Four bioregional pilots

Barcelona — Mediterranean. IAAC, Fab Lab Barcelona, with regional HPC and municipal partners under negotiation. Climate: Mediterranean heat stress, dense urban tissue, mature circular-economy policy backbone. Instrumentation reality: the strongest in the network — 15-year SmartCitizen archive, ~600 active sensors, dense OpenAQ FRM/FEM, full Eurostat EW-MFA, direct Metroverse ECI. This is where the stack is proven before it is exported.

Boston — North Atlantic. MIT Center for Bits and Atoms, Fab Foundation, Fab Hub Kendall Square. Climate: cold-climate urban-heat-island dynamics; 2026 Fab City Summit host. Metroverse covers Boston directly. The community sensor fleet remains to be built; this is named as a year-one deliverable rather than an existing capability. Boston tests whether the Matryoshka embedding family survives a statistically different bioregion without losing alignment.

Santiago — Southern Cone. Universidad Católica de Chile, Núcleo Milenio FAIR, with Chilean AI research partners under negotiation. Climate: Andean / coastal / central-valley bands with Spanish-language and Andean-indigenous-knowledge governance dimensions that change the CARE protocol. ~75% informal urban employment invisible to ECI. ILO 21st ICLS 2023 is the dedicated Community-tier instrument. Santiago demonstrates that the stack generalises to a context where formal-sector data alone produces a misleading picture — and that the Community tier is where the informal economy becomes observable.

Bali — Indonesian Archipelago. IT Del, CAST Foundation, Meaningful Design Group, Fab Lab Bali. Climate: tropical-humid, cyclone-exposed, archipelagic, intermittent connectivity, multiple indigenous-language communities. Bali falls into the same PITO/DIDO/ρ taxonomy as every other pilot — same indices, same observatory surfaces — with two sovereign data sources feeding the framework rather than a Bali-specific index. **Bali Satu Data** (balisatudata.baliprov.go.id) — the official Bali provincial open data portal published by Tim SPBE Diskominfos Provinsi Bali under Pergub Bali 53/2021, with 4,700+ datasets organised through six Balinese cosmological categories: **Jana** (society/people), **Atma** (wellbeing/spirit), **Jagat** (territory/world), **Wana** (forest), **Danu** (water), **Segara** (sea/coast) — provides the City- and Region-tier Environmental and Economic data with full data sovereignty. **WIPO Global Innovation Index 2025** [WIPO 2025] places Indonesia at rank 55/139 (8th of 36 upper-middle-income economies, 12th of 17 in South-East-Asia / East-Asia / Oceania), climbing 30 positions since 2020 — the country-level peer-

comparison anchor for the Bioregion tier, with strengths in entrepreneurship policies (rank 10), state of cluster development (rank 11), and university-industry R&D collaboration (rank 13). Engineering ingest path: data.go.id (CKAN 3.0 national portal, contains Bali subset via federation) plus a Bali Satu Data scraper for province-specific datasets; direct API access through partner collaboration with Diskominfos is a year-one Bali deliverable. Community-tier DIDO inputs (fab-lab catchment, Fab Academy throughput, Precious Plastic activity density, SmartCitizen sensor density, banjar council activity) feed the same DIDO computation as Barcelona, Boston, and Santiago — same metric, same formula, different ingest sources. The lack of Metroverse coverage for Denpasar is handled by reconstructing the City × Economic cell from sovereign provincial open data published under Bali's own cosmological taxonomy. The sovereignty architecture is stress-tested most directly at this pilot, and the **Tri Hita Karana ↔ Full-Stack Metrics pillars** mapping is run as a **Bali pilot only co-design experiment** between CAST, MDG, IT Del and a named Balinese customary authority (banjar adat, desa adat, or PHDI scholar — confirmed in v2); **publication is conditional on customary-authority sign-off**, not a methodological export (see *What's inherited, what's novel*).

---

10 What's inherited, what's novel

Inherited. Economic Complexity Index [Hidalgo & Hausmann 2009; Metroverse]. City-scale ontological taxonomy [City Anatomy, Guallart et al. 2015]. Agent-based simulation with tangible interfaces [CityScope, MIT Media Lab]. Sovereignty-by-architecture at individual scale [PKC, Koo 2025]. Planetary-boundaries radar form [Hoornweg et al. 2016]. Community-tier geographic units + quintile scoring [Dark Matter Labs & TreesAI 2025]. **Fab City Index, Generation 1** — 257-sector / 12-macro-sector self-sufficiency index across ~600 French urban areas via LOCAL SHIFT® / LOCAL FOOTPRINT® [Florentin, Chabanel & Guimas 2018, Utopies & FabCity Paris]. **Fab City Index, Generation 2** — NACE Rev. 2 × COICOP concordance, 16 macro-sectors, Hamburg pilot 37/100 [Boeing 2024]. **Full-Stack Metrics 4-Pillar × 5-Scale framework** — the measurement scaffolding underneath the Fab City Full Stack [Vivanco 2024], with bioregional material-mapping fieldwork across four Chilean macrozones [Vivanco 2025]. Re-distributed manufacturing coupled with digital intelligence as the causal mechanism [Moreno & Charnley 2016]. Measurement-to-action at local scale [Luo & Ball 2025]. AI transforming urban observing [Liu et al. 2024].

Novel. Federated AI causal inference across five governance tiers simultaneously — not in any prior framework. **H_0-A**, testing whether observation-to-action cycle time is mea-

surably shorter than a policy-only baseline — no existing platform tests this. **The Fab City Index Generation 3 formulation**: PITO and DIDO as the founding state-diagnosis vocabulary [Diez 2016] reactivated as continuous indices over the 4×5 matrix, the response coefficient ρ as the per-tier operational form of H_0-A, and the trajectory **ΔFCI/Δt** as the metric the instrument is built to measure. Bioregional peer-matching combining industry-space, material-flow, and ecological vectors in a single 1,024-d embedding. A distributed-production indicator stack at Community tier: fab-lab catchment, materials processed locally, active skills density, project documentation count — none exist in municipal datasets today. Production capacity as a Community-tier adaptation measure alongside climate and socioeconomic vulnerability. CARE-governed sovereignty enforced *by architecture* at the bioregion-server boundary. An HS→material crosswalk connecting Atlas of Economic Complexity import categories to Eurostat EW-MFA material-flow taxonomy. **One taxonomy across all pilots** — every pilot computes PITO, DIDO and ρ from the same formula on the same matrix; per-pilot data sources differ, the indices do not. **The Tri Hita Karana ↔ four-pillar mapping (Bali pilot only · Path B).** Bali's provincial open data is published through a Balinese cosmological taxonomy (Jana, Atma, Jagat, Wana, Danu, Segara). PLANETAI runs the mapping from this taxonomy to the Full-Stack Metrics 4-Pillar × 5-Scale framework as a **Bali pilot only co-design experiment** between CAST, MDG, IT Del and a named Balinese customary authority (banjar adat, desa adat, or PHDI scholar — confirmed in v2). **Publication is conditional on customary-authority sign-off** — a methodological-companion paper if and only if the customary authority co-authors and approves it. This is not a methodological export and not a generic alignment claim; it is a CARE-governed local experiment whose publication path is decided by the people whose cosmology is being mapped. AI-enabled causal discovery across urban subsystems as city-tier policy intelligence, following Wu [2026]. **The PLANETAI Node** — the five-tier stack instantiated as OSHWA-certifiable open hardware (BOM, Docker runtime, federation protocol, maintenance model). **The PLANETAI Agent runtime** — pre-registered action-agent framework with one reference agent per pilot, human-in-the-loop by architecture, and community-council disable-authority at the bioregion-server boundary.

---

11 Platform state

The observatory is a working demonstration, not a production system. Live, synthetic, and to-be-built capabilities are labelled distinctly throughout the interface and reproduced here.

Live. Global satellite map. 56 Fab Cities from the signatory registry. OpenSky flights. NASA FIRMS fires. Open-Meteo weather. SmartCitizen Barcelona. OpenAQ reference stations. Atlas of Economic Complexity imports (Planet Sankey). *What a Waste* MSW composition. AMB PREMET25 and MassDEP bioregion Sankeys for Barcelona and Boston. WIPO Global Innovation Index 2025 Indonesia profile (Bali Bioregion-tier peer anchor).

Synthetic. Production Constellation radar scores per pilot. Bali and Santiago Bioregion Sankeys (national *What a Waste* per-capita adjusted). Local Sankey (410 kg/cap fab-lab-shed budget, shape real, weights indicative). Community-tier production panels. MRL embedding values.

To be built. Cross-network sensor normalisation layer — named as a research deliverable, not a solved primitive. ≥30-day colocation campaigns at every community hub, without which causal-attribution claims cannot be made. Bioregion-server infrastructure hosting calibration-status tracking. Federated aggregation with non-IID baselines. Community-council protocols with a real CARE veto, operationalised technically at the bioregion-server boundary rather than written into governance policy alone. **Bali Satu Data ingestion pipeline** — provincial open data portal (4,700+ datasets) ingested via the national Satu Data Indonesia (data.go.id) CKAN-3.0 API and the provincial portal scraper. The Tri Hita Karana → four-pillar mapping is a **Bali pilot only co-design experiment** (Path B) gated on CAST + MDG + IT Del + a named Balinese customary authority (banjar adat, desa adat, or PHDI scholar) sign-off; mapping is not applied at ingest by default and publication is conditional on customary-authority approval.

---

12 Open scientific questions

The seven items below are unresolved at the time of writing and define the live boundary of the project.

1. H_0-T magnitude. LCAs suggest 40–60% cumulative energy-demand reduction at the product level. Bioregional-scale empirical measurement has never been done. Three outcomes are equally plausible. We publish all three.

2. Bali Satu Data ingestion + Tri Hita Karana mapping (Path B). Bali Satu Data (4,700+ provincial datasets) is the primary sovereign data source for the Bali pilot, but the per-cell mapping from the six-category Balinese cosmological taxonomy (Jana, Atma, Jagat, Wana, Danu, Segara) into the Full-Stack-Metrics four pillars is a v0 documented mapping, not yet validated by Bali partners (CAST, MDG, IT Del) and not yet co-designed with a Balinese customary authority (banjar adat, desa adat, or PHDI scholar). The mapping is a **Bali pi-**

lot only co-design experiment; **publication is conditional on customary-authority sign-off** and is not committed to as a project-wide deliverable. WIPO GII 2025 Indonesia profile (rank 55/139) provides the country-level Bioregion-tier peer-comparison anchor. Bioregion→City downscaling R^2 is still unmeasured for any pilot — separate research deliverable.

3. Cross-network sensor normalisation. Not solved. It is a research deliverable, not a completed primitive. Community-sensor ±30–80% uncertainty against reference instruments [Castell et al. 2017] is named in every Community-tier deployment.

4. Community-tier vulnerability re-fitting. Existing community vulnerability methodologies [e.g., Dark Matter Labs & TreesAI 2025, UK] are geography-specific. Bioregion-specific adaptation for Barcelona, Santiago, and Bali is year-one work.

5. Green Complexity Index at city tier. Mealy & Teytelboym's GCI is country-level in the published literature. A city-level GCI is an extension we propose to publish as a companion paper. Not a validated construct yet.

6. Node maintenance at scale. The Fab Lab Node-Operator role is newly defined. Whether the Fab Academy curriculum extension produces enough trained operators to scale beyond the four anchor nodes is a year-two empirical question.

7. Agent-governance protocol across jurisdictions. Human-in-the-loop approval is a technical constraint at the agent-runtime level, but the legitimacy of the human approver differs across the four jurisdictions (Spain, US, Chile, Indonesia). Community-council authority is technical; the constitutional weight of that authority is a governance question each pilot's community council resolves locally.

---

13 Scientific deliverables

Three or more peer-reviewed publications, including one explicitly reporting the H_0-T causal-falsification outcome with publication-grade error bounds. Three or more open datasets. A cross-network sensor normalisation layer as a standalone Apache 2.0 artefact. A federated-deployment protocol as reusable infrastructure. OSF pre-registration of both hypotheses. An HS→material crosswalk as a methodological contribution. A bioregional peer-matching methodology companion paper. **PLANETAI Node v1** — open-hardware specification (BOM, Docker runtime, federation protocol, node-operator training curriculum) released under Apache 2.0 with the enclosure design OSHWA-certifiable. **PLANETAI Agent runtime v1** — Apache 2.0 agent-runtime code with four reference agent implementations, pre-registered latency-measurement protocols, and a published audit-log

format.

---

14 The bet

The Fab City network has spent more than a decade on a claim it never tested. The claim is worth testing. It is also worth being ready for the answer to be *no*, or *not at this scale*, or *yes but smaller than hypothesised*. PLANETAI is the instrument that lets us hear whichever answer is true, and the action layer that lets us respond at community cadence rather than policy cadence.

The cases the literature surveyed above documents — Barcelona's Ateneus de Fabricació, Maker's Asylum's M-19 collective, the Tunisian Solidarity FabLabs, the PPDFI and Yayasan Peduli Tuna Daksa prosthetic-distribution surfaces in Indonesia — are not external case-studies. They are constituent nodes, or close cousins of nodes, inside the same distributed-fabrication network PLANETAI instruments. The question is whether a network that has already produced episodic action-latency compression under crisis conditions can produce measurable, attributable compression under everyday conditions, against a controlled baseline, repeatedly. That measurement has not been run. PLANETAI is the instrument that runs it.

If the answer is positive, the network has causal evidence it never had. If the answer is null, the network pivots with its honesty intact. Either way, the infrastructure — federated AI stack, calibrated sensor fleet, trained community scientists, sovereignty architecture — remains in place. A 56-city distributed scientific instrument, open, auditable, and attached to a response network that can act.

---

15 References

Organised thematically.

15.1 Earth observation and AI foundations

Bodnar, C., et al. (2025). Aurora: a foundation model for the Earth system. *Nature.*

Lam, R., et al. (2023). Learning skilful medium-range global weather forecasting. *Science* (GraphCast).

Liu, Z., et al. (2024). How will AI transform urban observing, sensing, imaging, and mapping? *npj Urban Sustainability,* 4, 50.

Price, I., et al. (2024). Probabilistic weather forecasting with machine learning. *Nature* (GenCast).

15.2 Representation learning and federated inference

Kusupati, A., et al. (2022). Matryoshka Representation Learning. *NeurIPS*. arXiv:2205.13147.

15.3 Distributed manufacturing and the Fab City Full Stack

Angstmann, A., et al. (2025). Green urban manufacturing: a systematic literature review. *Urban Planning*.

Balestrini, M., Diez, T., Marshall, P., Gluhak, A. & Rogers, Y. (2015). IoT community technologies: leaving users to their own devices or orchestration of engagement? *EAI Endorsed Transactions on Internet of Things*, 1(1), e1. https://doi.org/10.4108/eai.26-10-2015.150601. Empirical foundation for the community-sensing and engagement methodology that PLANETAI extends through the four-hub harvester architecture.

Boeing, N. (2024). The Fab City Index: a toolkit for measuring progress towards a circular economy. In Moritz, M., Redlich, T., Buxbaum-Conradi, S. & Wulfsberg, J. P. (eds), *Global Collaboration, Local Production.* SDG – Forschung, Konzepte, Lösungsansätze zur Nachhaltigkeit. Springer Gabler, Wiesbaden, 115–133. https://doi.org/10.1007/978-3-658-44114-2_9.

Diez, T. (2012). Personal Fabrication: Fab Labs as Platforms for Citizen-Based Innovation, from Microcontrollers to Cities. *Nexus Network Journal*, 14(3), 457–468. https://doi.org/10.1007/s00004-012-0131-7. PI's first peer-reviewed published articulation of the fab-lab-as-civic-platform thesis that PLANETAI operationalises at bioregional scale.

Diez, T. (2016). *Fab City Whitepaper: Locally productive, globally connected self-sufficient cities.* Fab City Global Initiative / IAAC / MIT Center for Bits and Atoms / Fab Foundation / Barcelona City Council, April 2016. Original published articulation of the *Products In, Trash Out → Data In, Data Out* transition.

Diez, T. (ed.) (2018). *Fab City: The Mass Distribution of (Almost) Everything.* IAAC / Actar Publishers, Barcelona. The PI-edited foundational volume on the Fab City model.

Diez, T., Niaros, V. & Ferro, J. (2024). The Fab City Full Stack. In Moritz, M., Redlich, T., Buxbaum-Conradi, S. & Wulfsberg, J. P. (eds), *Global Collaboration, Local Production.* SDG – Forschung, Konzepte, Lösungsansätze zur Nachhaltigkeit. Springer Gabler, Wiesbaden, pp. 13–26. Open Access. https://doi.org/10.1007/978-3-658-44114-2_2. The Full Stack chapter that the Full-Stack Metrics framework [Vivanco 2024] operationalises.

Florentin, A., Chabanel, B. & Guimas, V. (2018). *Toward Productive Cities: FabCity Index France.* Utopies & FabCity Paris, Étude N°13, June 2018.

Leal, R., et al. (2025). Circular economy adoption in manufacturing firms. *Business Strategy and the Environment.*

Luo, Y. & Ball, P. (2025). A digital twin and Q-learning-based approach for vertical-farm production scheduling. *Scientific Reports*, 15.

Marmiroli, B., et al. (2020). Electric-vehicle life-cycle assessment and logistics implications. *Journal of Cleaner Production.*

Martin, J., et al. (2024). Decentralised manufacturing systems and urban circulation hubs. *CIRP Annals.*

Moreno, M. & Charnley, F. (2016). Can re-distributed manufacturing and digital intelligence enable a regenerative economy? *Sustainable Design and Manufacturing* (*Springer IDT*), 52.

Peuckert, J., Wenzelmann, V. & Rüdebusch, G. (2025). Measuring the impacts of Fab Labs: a review of quantitative research and evidence. Presented at Fab25, Brno & Prague, 4–11 July 2025. Helmut Schmidt University Hamburg. Zenodo. https://doi.org/10.5281/zenodo.16211909.

Srai, J. S., et al. (2016). Distributed manufacturing: scope, challenges and opportunities. *International Journal of Production Research*, 54(23).

Vivanco, T. (2024). *Fab City Full Stack Metrics Framework: An Architectural Framework for Measuring the Impact of Design Interventions across five scales* (*Community, City, Region, Bioregion, Planet*) *through four pillars* (*Environmental, Social, Economic, Governance*). Fab City Foundation working paper.

Vivanco, T. (2025). *Research through Design for bioregional material mapping across four Chilean macrozones.* Doctoral thesis, Universidad Católica de Chile.

15.4 Civic-innovation response speed (priors for H_0-A)

Buehler, E., Branham, S., Ali, A., Chang, J. J., Hofmann, M. K., Hurst, A. & Kane, S. K. (2015). Sharing is caring: assistive technology designs on Thingiverse. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (*CHI '15*).

CIDOB (2020). *Barcelona's CoronavirusMakers: co-producing local solutions to a global pandemic.* Cities on the Frontline series, Barcelona Centre for International Affairs. https://www.cidob.org/en/publications/barcelonas-coronavirusmakers-co-producing-local-solutions-global-p

Corsini, L., Dammicco, V. & Moultrie, J. (2021). Frugal innovation in a crisis: the digital fabrication maker response to COVID-19. *R&D Management*, 51(2), 195–210. https://doi.org/10.1111/radm.12446.

Distributed Design Platform (2020). *Viral Design: the COVID-19 crisis as a global test bed for*

distributed design. https://distributeddesign.eu/wp-content/uploads/2020/11/DistributedDesignBook_2020-online.pdf.

Hamdi-Cherif, A., et al. (2021). Maker movement contribution to fighting COVID-19 pandemic: insights from Tunisian FabLabs. *Journal of Public Affairs*. PMC8661684. https://pmc.ncbi.nlm.nih.gov/articles/PMC8661684/.

Hofmann, M., Burke, J., Pearlman, J., Fiedler, G., Hess, A., Schull, J., Hudson, S. E. & Mankoff, J. (2016). Clinical and maker perspectives on the design of assistive technology with rapid prototyping technologies. *ASSETS '16: Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*, 251–256. https://doi.org/10.1145/2982142.2982181.

Hofmann, M., Harris, J., Hudson, S. E. & Mankoff, J. (2017). Understanding volunteer AT fabricators: opportunities and challenges in DIY-AT for others in e-NABLE. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 6184–6194. https://doi.org/10.1145/3025453.3026045.

Kieslinger, B., Schaefer, T., Fabian, C. M., Biasin, E., Bassi, E., Ruiz Freire, R., Mowoh, N., Arif, N. & Melis, P. (2021). Covid-19 response from global makers: the Careables cases of global design and local production. *Frontiers in Sociology*, 6, 629587. https://doi.org/10.3389/fsoc.2021.629587.

Manunza, M., et al. (2023). How makers responded to the personal protective equipment shortage during the COVID-19 pandemic: an analysis focused on the Hauts-de-France region. PMC10147541. https://pmc.ncbi.nlm.nih.gov/articles/PMC10147541/.

Open Source Medical Supplies (2020). *National case study: India's maker response against COVID-19*. https://opensourcemedicalsupplies.org/national-maker-response-case-studies/india/.

Pemerintah Provinsi Jawa Barat (2024). *Garut District prosthetic distribution programme reporting*. https://jabarprov.go.id/en/berita/pemdakab-garut-kolaborasi-salurkan-kaki-dan-tangan-pals

Tabarés, R., Cuhls, K., et al. (2021). The Coronavirus Makers Network: understanding the success of an innovation community facing COVID-19 in Spain. ResearchGate 352719934.

Tabarés, R., Cuhls, K., et al. (2021). Design, manufacture and save: Coronavirus Makers during the COVID-19 crisis in Spain. ResearchGate 354789721.

UK National Audit Office (2020). *Investigation into government procurement during the COVID-19 pandemic*. https://www.nao.org.uk/reports/government-procurement-during-the-covid-19-pand

Yaqin, M. A., et al. (2024). Design and implementation of a low-cost and functional prosthetic hand using 3D printing technology for a member of the Association of Physical Disabilities Indonesia (PPDFI Surabaya). *Frontiers in Community Service and Empowerment*.

https://ficse.ijahst.org/index.php/ficse/article/view/51.

Zuniga, J., Katsavelis, D., Peck, J., Stollberg, J., Petrykowski, M., Carson, A. & Fernandez, C. (2015). Cyborg Beast: a low-cost 3D-printed prosthetic hand for children with upper-limb differences. *BMC Research Notes*, 8, 10. https://doi.org/10.1186/s13104-015-0971-9.

15.5 Economic complexity and the city

Harvard Growth Lab (2021–). *Metroverse: Urban Economy Navigator.* https://metroverse.hks.harvard.edu.

Hausmann, R., et al. (2014). *The Atlas of Economic Complexity: Mapping Paths to Prosperity.* MIT Press.

Hidalgo, C. A. & Hausmann, R. (2009). The building blocks of economic complexity. *PNAS*, 106(26).

Mealy, P. & Teytelboym, A. (2022). Economic complexity and the green economy. *Research Policy*, 51(8).

Stojkoski, V., Koch, P. & Hidalgo, C. A. (2023). Multidimensional economic complexity and inclusive green growth. *Communications Earth & Environment*, 4, 130.

WIPO (2025). *Global Innovation Index 2025: Indonesia country profile.* World Intellectual Property Organization. https://www.wipo.int/global_innovation_index/en/2025/.

15.6 Indonesian sovereign open data and Tri Hita Karana

Pemerintah Provinsi Bali (2021). *Peraturan Gubernur Bali Nomor 53 Tahun 2021 tentang Satu Data Indonesia Tingkat Provinsi.* https://jdih.baliprov.go.id/uploads/produk-hukum/peraturan/2021/pergub/2021pergub0051053.pdf.

Tim SPBE Diskominfos Provinsi Bali. *Bali Satu Data — Portal resmi data terbuka Provinsi Bali.* https://balisatudata.baliprov.go.id. Integrated with Satu Data Indonesia (https://data.go.id) via CKAN API 3.0.

15.7 Urban metabolism and planetary boundaries

City of Amsterdam (2022). *Amsterdam Circular Monitor.* https://amsterdam.nl.

Dark Matter Labs & TreesAI (2025). *Building Climate Resilience in Croydon.* https://github.com/Dark-Matter-Labs/croydon_adaptation_plan.

Hoornweg, D., et al. (2016). An urban approach to planetary boundaries. *Ambio*, 45(5), 567–580.

Richardson, K., et al. (2023). Earth beyond six of nine planetary boundaries. *Science Advances.*

15.8 City-scale ontology and policy

Guallart, V., Meneses, N., Frogheri, F., Ibañez, O., Rubio, P. & Giralt, C. (2015). *City Anatomy* (CPA-I_001-v2). City Protocol Society / IAAC.

He, W., et al. (2023). Green technology, collaborative agglomeration, and carbon productivity. *Environment, Development and Sustainability.*

Wu, X. (2026). Data-driven system-dynamics modelling for urban ecosystem assessment. *Land*, 15.

15.9 Sensor and measurement reliability

Castell, N., et al. (2017). Can commercial low-cost sensor platforms contribute to air-quality monitoring and exposure estimates? *Environment International.*

15.10 Sovereignty, CARE, and governance architecture

International Labour Organization (2023). *Resolution concerning statistics on the informal economy (21st ICLS)*. ILO, Geneva.

Koo, B. (2025). *PKC: A Resource-Aware, and Self-Sovereign Computational Infrastructure for Community Learning.* pkc.pub; github.com/xlp0/PKC. Published June 2025; revised March 2026.

PLANETAI · Fab City · April 2026