CURIOSITYBITS LIVING LAB BOSTON · STUDY #04 PRE-REGISTRATION & STUDY DESIGN CARD (V1.2)

Pre-Registration & Observatory Design Card (`STUDY_DESIGN_CARD.md`)

- **Study ID:** `study-04`
- **Title:** *Where the Cloud Hits the Ground: Monitoring Local, State, and National Pushback and Bargaining Over AI Infrastructure*
- **Edition:** `v1.0` (Pre-Data Locked Pre-Registration)
- **Locked Timestamp (UTC):** `2026-11-05T02:58:00Z`
- **Lab:** CuriosityBits Living Lab Boston (`By CuriosityBits Living Lab Boston`)
- **Study Genre:** Living Measurement, Description & Monitoring Observatory (with Pre-Registered Comparative Propositions `P1` – `P3`)

---

1. Plain-English Research Objective & Observatory Scope

While national debates over artificial intelligence often revolve around abstract existential risk or frontier model benchmarks, the most concrete political fault line over AI in the United States is emerging where physical AI infrastructure—hyperscale data centers, high-voltage substations, transmission corridors, water-cooling systems, and gigawatt power generation—meets local communities, ratepayers, city councils, county boards, public utility commissions, and state legislatures.

This Living Paper establishes a continuously updated, reproducible, and auditable measurement infrastructure that monitors public, media, and legislative responses to AI infrastructure across **local municipalities and counties (Tier 1)**, **key state legislatures and utility commissions (Tier 2)**, and **national federal regulatory/congressional dockets (Tier 3)**. Rather than reducing public reaction to a simplistic `"pro-AI"` vs. `"anti-AI"` binary, we measure the specific points of contention, the institutional posture of government actions, and a composite **Public Pressure & Institutional Friction Index (`public_pressure_index`, `0–100`)**.

---

2. Q1 — Operational Measurement, Multi-Tier Data Sources & Powered Sample Sizing (`SKILL-012` & `SKILL-013`)

A. Systematic Public Data Sources (`2023-01-01` to `2026-11-05`, Zero Synthetic Fallback)

1. **Tier 1A — Municipal & County Legislative Records via Granicus Legistar Web API (`https://webapi.legistar.com/v1/{client}/matters`):** - Systematically audits 28 U.S. municipal and county Legistar endpoints (`phoenix`, `mesa`, `goodyear`, `santaclara`, `sanjose`, `mountainview`, `oakland`, `sacramento`, `fresno`, `longbeach`, `seattle`, `kingcounty`, `columbus`, `naperville`, `madison`, `milwaukee`, `alexandria`, `hampton`, `richmondva`, `greensboro`, `nashville`, `sanantonio`, `plano`, `denver`, `boston`, `pittsburgh`, `baltimore`, `newark`) for ordinances, resolutions, zoning applications, utility substation agreements, water assessments, and municipal AI governance actions introduced between `2023-01-01` and `2026-11-05`.
2. **Tier 1B & Tier 2 — County Board, State Legislative, Public Utility Commission (PUC) & Public Radio/Hearing Records (`https://api.search.brave.com/res/v1/web/search`):** - Systematically retrieves dated, URL-verified local county actions (including Virginia's Data Center Alley in Loudoun, Prince William, and Fairfax Counties; Metro Atlanta/Georgia; Central Ohio; Arizona; Oregon; Texas; Illinois; North Carolina; Pennsylvania), state legislative bills (sales-tax exemption audits/repeals, water/energy disclosure mandates), state PUC large-load tariff dockets (e.g., Virginia SCC, Ohio PUCO, Georgia PSC, Arizona Corporation Commission, ERCOT/Texas), and NPR / local public radio / municipal beat coverage of public hearings and community mobilization.
3. **Tier 3 — National Regulatory & Executive Dockets via Federal Register API (`https://www.federalregister.gov/api/v1/documents.json`):** - Systematically retrieves official U.S. federal agency notices, proposed rules, final rules, and presidential documents from FERC, DOE, EPA, Commerce/NTIA, and OMB addressing data-center electricity load, grid interconnection, nuclear co-location, cooling/environmental standards, and national AI infrastructure policy (`2023-01-01` to `2026-11-05`).

B. Powered Initial Sample Target & Minimum Detectable Effect (`SKILL-013`)

- **Sample Target:** N ≥ 240 unique, dated, URL-verified public/legislative/regulatory records across G ≥ 35 distinct jurisdictional clusters (`jurisdiction_id`: local municipalities/counties, state jurisdictions, and federal agencies) spanning `2023–2026`.
- **Ex-Ante Power & Minimum Detectable Effect (MDE):** At N = 240 across G ≥ 35 clusters (assuming a conservative intra-cluster correlation \rho = 0.08, average cluster size \bar{m} ≈ 6.5, design effect DEFF = 1 + (\bar{m}-1)\rho ≈ 1.44, effective sample size N_{eff} ≈ 167), a two-sided test at α = 0.05 achieves $80% statistical power (z_{0.025} + z_{0.20} = 1.96 + 0.8416 = 2.8016) to detect a standardized effect size of

d_{MDE} = 2.8016 × \sqrt{4 / N_{eff}} ≈ 0.43 SD (or a standardized partial regression slope r_{MDE} ≈ 0.21$).

C. Nuanced Multi-Dimensional Construct Definitions

Every record i is scored on **7 Contention & Issue Frames (`[0.0, 1.0]`)**, **1 Non-Binary Policy Posture Classification (5 categories + continuous directional support vs. restriction scores)**, and the **Public Pressure & Institutional Friction Index (`public_pressure_index` , `[0.0, 100.0]`)**: 1. **The 7 Contention & Issue Frames (`[0.0, 1.0]` continuous salience):** - `frame_grid_ratepayer` : Salience of electric grid capacity, gigawatt load growth, substation/transmission siting, and residential ratepayer cross-subsidization concerns. - `frame_water_land_env` : Salience of water consumption/cooling, diesel backup generator emissions, farmland/forest loss, residential proximity, and acoustic noise. - `frame_fiscal_tax` : Salience of sales/property tax exemptions, fiscal giveaways vs. local tax revenue, and cost-per-permanent-job debates. - `frame_transparency_nda` : Salience of non-disclosure agreements (NDAs), secretive site selection, fast-tracked zoning bypassing public hearings, and FOIA/public-record demands. - `frame_econ_growth_jobs` : Salience of local property-tax windfalls for schools/infrastructure, construction union jobs, and regional economic modernization. - `frame_labor_civil_rights` : Salience of algorithmic displacement, workplace surveillance, public-sector AI accountability, and civil liberties. - `frame_natsec_competitiveness` : Salience of U.S. compute leadership, domestic energy abundance, and national security. - **Composite Physical Resource & Ratepayer Contention Index (`resource_ratepayer_salience` , `[0.0, 1.0]`):** Equal-weighted mean of `frame_grid_ratepayer` , `frame_water_land_env` , and `frame_fiscal_tax` . 2. **Non-Binary Policy & Institutional Posture Typology (`policy_posture_category`):** - `1_unconditional_acceleration` : Unconditional subsidies, tax-exemption extensions, fast-track permitting, or by-right zoning expansion. - `2_conditional_bargaining` : Conditional approval ("Yes, *if*"): large-load utility tariffs protecting residential ratepayers, clean-power/closed-loop cooling mandates, noise buffers, or Community Benefit Agreements (CBAs). - `3_procedural_study_or_pause` : Formal impact study mandate, commission audit (e.g., JLARC), or temporary moratorium on new permits/rezoning. - `4_targeted_restriction_or_repeal` : Ending by-right zoning near residences, tightening noise/water caps, or repealing/capping tax exemptions. - `5_direct_denial_or_blockade` : Direct rezoning/permit denial, grassroots referendum/recall blockade, or litigation. 3. **Public Pressure & Institutional Friction Index (`public_pressure_index` , `[0.0, 100.0]`):** - Continuous `[0, 100]` composite index measuring the intensity of public, media, and institutional pressure on AI infrastructure developers, combining: (a) `mobilization_intensity` (`[0, 1]` , organized resident/coalition opposition, public comment surges, bipartisan pushback), (b) `institutional_friction_level` (`[0, 1]` , binding zoning/tariff/tax friction vs. routine administrative approval), and (c) `concession_or_guardrail_demand` (`[0, 1]` , demands for infrastructure self-funding, tariff ring-fencing, environmental limits, or site withdrawal):

> *PPI_i = 100 × \left(0.40 × mobilization_intensity_i + 0.35 × institutional_friction_level_i + 0.25 × concession_or_guardrail_demand_i\right)*

---

3. Q2 — Text Annotation & 4-Judge Open-Weight Ollama Cloud + JEV Validation Plan (`SKILL-012`, `SKILL-018`, `SKILL-019`)

To prevent single-model bias, every record in the analysis corpus is scored independently at `temperature = 0.0` across a **4-Judge Cross-Family Panel** combining **Ollama Cloud Open-Weight Models** and **JEV**: 1. **Judge 1 (`ollama-cloud/deepseek-v4.1-flash`):** Open-weight `deepseek-v4.1-flash` accessed via Ollama Cloud API (`https://ollama.com/v1/chat/completions`). 2. **Judge 2 (`ollama-cloud/glm-5.3-flash`):** Open-weight `glm-5.3-flash` accessed via Ollama Cloud API (`https://ollama.com/v1/chat/completions`). 3. **Judge 3 (`ollama-cloud/minimax-m3.1-flash`):** Open-weight `minimax-m3.1-flash` accessed via Ollama Cloud API (`https://ollama.com/v1/chat/completions`). 4. **Judge 4 (`jev/typesafe-jev-router`):** JEV Decision Verifier (`typesafe/jev-router` / `meta-llama/llama-3.3-70b-instruct` at `temperature = 0.0`).

- **Reliability & Ballot Transparency Requirements:**
- All raw per-judge ballots are shipped in `data/raw_judge_ballots.json` (`SKILL-018`).
- We report both single-rater `ICC(1,1)`, `ICC(2,1)`, `ICC(3,1)` and 4-judge average-measures `ICC(1,k)`, `ICC(2,k)`, `ICC(3,k)` (Shrout & Fleiss 1979) plus mean pairwise Pearson r and per-judge standalone regression slopes (`SKILL-019`).

---

4. Q3 — Pre-Registered Empirical Propositions, Confounder Controls & Explicit Falsification Criteria (`SKILL-005` & `SKILL-024`)

Alongside descriptive monitoring across jurisdictions and time, we pre-register three testable empirical propositions (`P1`, `P2`, `P3`), each estimated via OLS with standard errors clustered by jurisdiction (`jurisdiction_id`):

1. **Proposition `P1` (Local/State vs. National Resource-and-Ratepayer Divergence):** - **Specification:** `resource_ratepayer_salience ~ tier_local_or_state + word_tone + excerpt_char_len_z + C(event_year)` - **Prediction:** `tier_local_or_state` (`1` for Municipal/County/State records, `0` for Federal/National records) has a positive and statistically significant coefficient ($\beta > 0$, $t \geq 1.96$, clustered by `jurisdiction_id`), showing that sub-national AI politics centers significantly more on physical grid, water/land, and tax/ratepayer contention than national regulatory discourse.

2. **Proposition `P2` (Physical Resource Contention Predicts Public Pressure & Institutional Friction Beyond Tone and Confounders): - Specification:** `public_pressure_index ~ resource_ratepayer_salience + frame_labor_civil_rights + tier_local_or_state + word_tone + excerpt_char_len_z + C(region) + C(event_year)` - **Prediction:** `resource_ratepayer_salience` positively and significantly predicts `public_pressure_index` ($\beta > 0$, $t \geq 1.96$, clustered by `jurisdiction_id`), and within high-pressure records (`public_pressure_index >= 50`), **Conditional Bargaining + Procedural Pause/Moratorium (`posture 2 + 3`)** outnumber **Direct Denials/Bans (`posture 5`)**, demonstrating that local pushback predominantly operates as institutional bargaining and tariff/zoning conditioning rather than binary technology prohibition.
3. **Proposition `P3` (Post-2024 Escalation of Public Pressure & Grid-Ratepayer Contention): - Specification:** `public_pressure_index ~ period_2025_2026 + tier_local_or_state + word_tone + excerpt_char_len_z + C(region)` - **Prediction:** Records from the `2025–2026` cohort (`period_2025_2026 == 1`) exhibit significantly higher `public_pressure_index` ($\beta > 0$, $t \geq 1.96$, clustered by `jurisdiction_id`) than records from the early `2023–2024` buildout cohort (`period_2025_2026 == 0`).

Pre-Stated Falsification / Null Criteria

- If any pre-registered proposition (`P1`, `P2`, or `P3`) yields a jurisdiction-clustered $|t| < 1.96$ or an opposite sign after controlling for `word_tone`, `excerpt_char_len_z`, region, and year/tier controls, we commit to reporting that proposition as a **null or falsified result** in the abstract, main text, and public Feature without altering the specification ex post.

---

5. Mandatory Data Completeness & Inventory Verification Protocol

Every run of the ingestion pipeline must generate `data/data_completeness_ledger.json`, auditing: 1. **Endpoint Availability & Freshness Status:** Across all 28 audited municipal Legistar APIs, explicitly separating active `2023–2026` municipal inventories from frozen legacy endpoints (`goodyear` frozen `2021-01-27`, `greensboro` frozen `2021-06-15`, `sanantonio` frozen `2021-09-24`, `longbeach` migrated `2023-11-10`) and HTTP 500 endpoints (`chicago`, `atlanta`). 2. **Timestamp Hygiene & Clerk-Typo Filtering (`SKILL-016`):** Logging and filtering municipal clerk date typos where `MatterIntroDate > 2026-11-05` (e.g., `richmondva` `9859-10-01`, `denver` `2030-09-30`, `milwaukee` `2028-02-16`) so zero future-year anachronisms enter the analysis panel. 3. **Inventory vs. Ingested Coverage Verification:** Comparing official source match counts (`N_official_inventory`) against retrieved and parsed records (`N_ingested_valid`) per source tier, region, and calendar year (`2023–2026`), documenting any gaps and recovery paths.

---

6. Endpoint-Probe Provenance Disclosure & Edition `v1.2` Audit Amendment Log (`SKILL-024`)

- **Immutable Pre-Data Locked Reference:** The original locked design card is preserved byte-for-byte in `STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md` (SHA-256: `bb6b1b709c5d99c1d9ebe6f954bfd99ca646879a332ec96b6d4fa29bf7b462b9`, locked `2026-11-05T02:58:00Z`).
- **How This Design Card Was Drafted Relative to the Endpoint Probe (`SG-AUDIT-S04-V1.1-2026-11-06`, Fix 6):** Prior to locking `STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md` at `2026-11-05T02:58:00Z`, the author ran an initial endpoint-health reconnaissance probe (`?top=3&orderby=MatterIntroDate desc`) against the candidate Granicus Legistar municipal OData APIs (`2026-11-05T02:49:00Z–02:56:00Z`) to determine which municipal endpoints were active versus frozen or returning HTTP 500 errors and to check date-field hygiene. That preliminary endpoint probe produced the specific frozen/migrated endpoint dates (`goodyear` `2021-01-27`, `greensboro` `2021-06-15`, `sanantonio` `2021-09-24`, `longbeach` `2023-11-10`), the two HTTP 500 endpoints (`chicago`, `atlanta`), and the three municipal clerk future-date typos (`richmondva` `9859-10-01`, `denver` `2030-09-30`, `milwaukee` `2028-02-16`) cited in Section 5 above. Full corpus ingestion (`collect_and_verify_corpus.py`, `retrieved_at_utc = 2026-11-05T03:04:18Z`, yielding N = 364 records across G = 48 jurisdictional clusters) and 4-judge scoring/regression estimation (`annotate_and_analyze_observatory.py`) were executed after locking the card, so the pre-registration preceded outcome scoring and hypothesis testing (`P1–P3`) but incorporated endpoint-availability and clerk-typo observations from the initial API health probe.
- **Post-Audit Amendments Logged in Edition `v1.2` (`2026-11-06`, following `SG-AUDIT-S04-V1.1-2026-11-06`):** While keeping the pre-registered `P1–P3` specifications on the baseline panel (N = 364, G = 48) as primary, Edition `v1.2` adds six post-audit verification and corpus-hygiene analyses: (1) reporting inter-rater reliability and construct scores strictly on live LLM ballots after dropping all `824 / 1,456` (`56.59%`) deterministic rule-fallback ballots; (2) deriving `policy_posture_code` from the median and mode of live ballots (plus standalone 100%-live JEV) rather than the rounded 4-ballot mean; (3) rebuilding the federal comparison group from 51 relevance-ranked and phrase-matched 2023–2026 Federal Register documents (including Executive Orders 14318, 14141, 14110, and 14365, the DOE AI-infrastructure RFI and *Speed to Power* notice, and the NTIA data-center RFC) scored 100% live (`0%` fallback) and testing `P1–P3` with Rademacher wild-cluster restricted bootstrap (B = 4,999); (4) filtering records for AI/data-center/large-load relevance prior to scoring (`N = 243, G = 44`), labeling the 97 non-`.gov`/`.us` records in `State_County_PUC_Primary_Search` as secondary sources (`183` primary government records + `60` secondary reportage/legal-analysis records in the cleaned `v1.2` panel), dating secondary pages by underlying event rather than web-page modification date, deduplicating multi-outlet secondary event records, and removing the FY1984–85 Richmond budget artifact (`AII-2023-013`); (5) reporting exact full-rank standard errors (`P2` SE = 9.8013, t = +2.6804; `P3` SE = 2.5359, t = +1.8975) and

grouping the §4.4 jurisdiction ranking strictly by `jurisdiction_id` (`US_PublicRadio_LocalBeat` N = 10, mean 31.80); and (6) documenting this endpoint-probe provenance disclosure.