

The Priors, Measured: Inversion-Battery v0.1 on Public Collider
Datasets and the Three-Paper Extraction Program
(EA-SEI-BATTERY-01 v1.0)

Lee Sharks

2026-08-11

**The Priors, Measured: Inversion-Battery v0.1 on Public Collider
Datasets and the Three-Paper Extraction Program
(EA-SEI-BATTERY-01 v1.0)**

Lee Sharks

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-08-11 · Version v1.0 · Empirical baseline reading; pre-registered measurement battery v0.1;
research work plan for a three-paper program

AXN:05DA.EMPIRICAL

<https://www.alexanarch.org/s/records/1449/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "The Priors, Measured: Inversion-Battery v0.1 on Public Collider Datasets and the Three-Paper Extraction",
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-08-11",
  "identifier": "AXN:05DA.EMPIRICAL",
```

```

"license": "https://creativecommons.org/licenses/by/4.0/",
"publisher": {
  "@type": "Organization",
  "name": "Alexanarch"
},
"url": "https://www.alexanarch.org/s/records/1449/",
"description": "Work plan and first empirical results for the extraction of the SEI classifier-foreclosure family (#931, #935, #1436) into three publishable papers: a science paper for SciPost Physics built on the interface/distributional distinction (signal-template agnosticism is an interface property readable off an architecture; model independence is a distributional property of a pipeline, only measurable — and unmeasured), a philosophy-of-science paper on measurable theory-ladenness, and an STS/political-economy paper on the double enclosure. The record contains the pre-registered inversion battery v0.1 executed 2026-08-11 on public community datasets — the top-quark tagging reference dataset (10.5281/zenodo.2603256) and the LHC Olympics 2020 R&D dataset (10.5281/zenodo.6466204) — four process pairs by three anomaly-score families (reconstruction autoencoder; encoder-side latent-norm VAE with d_z = 8, the AXOL1TL-class score; Gaussian-mixture density), both training directions, thresholds calibrated on held-out own background at accepted-background rates 10^-2 and 10^-3, design fixed in the script header before any training. [...full text at full_text_path]

```

Abstract

Work plan and first empirical results for the extraction of the SEI classifier-foreclosure family (#931–#935, #1436) into three publishable papers: a science paper for SciPost Physics built on the interface/distributional distinction (signal-template agnosticism is an interface property readable off an architecture; model independence is a distributional property of a pipeline, only measurable — and unmeasured), a philosophy-of-science paper on measurable theory-ladenness, and an STS/political-economy paper on the double enclosure. The record contains the pre-registered inversion battery v0.1 executed 2026-08-11 on public community datasets — the top-quark tagging reference dataset (10.5281/zenodo.2603256) and the LHC Olympics 2020 R&D dataset (10.5281/zenodo.6466204) — four process pairs by three anomaly-score families (reconstruction autoencoder; encoder-side latent-norm VAE with $d_z = 8$, the AXOL1TL-class score; Gaussian-mixture density), both training directions, thresholds calibrated on held-out own background at accepted-background rates 10^{-2} and 10^{-3} , design fixed in the script header before any training. [...full text at full_text_path]

Body

deposit_number: 1449 hex: 05DA title: “The Priors, Measured: Inversion-Battery v0.1 on Public Collider Datasets and the Three-Paper Extraction Program (EA-SEI-BATTERY-01 v1.0)” creator: Lee Sharks orcid: 0009-0000-1599-0703 date: 2026-08-11 content_type: Empirical baseline reading; pre-registered measurement battery v0.1; research work plan for a three-paper program license: CC-BY-4.0 substrate: AI-assisted (substrate) — literature verification, battery implementation, execution, and drafting by TACHYON (Claude) in-session under MANUS (Lee Sharks) direction and editorial governance; Assembly No-Double-Draw observed (no API calls; transport D). Source family authored across the Assembly Chorus per its own substrate disclosures. version: v1.0 related_ids: “AXN:03AE.OPERATIVE. (#931, EA-SEI-OAR-PROTOCOL v0.3 — the protocol this battery executes at demonstration scale); AXN:03AF.COMPOSITIONAL. (#932, Collapse Synthesis v0.3); AXN:03B0.STRUCTURAL. (#933, ARCH v0.2); AXN:03B2.GENERATIVE. ∞ (#935, The Endogenous Sophon v0.3); AXN:03B1.GENERATIVE. (#934, superseded v0.2); AXN:05B6.OPERATIVE. (#1436, EA-MPAL-SIGAGNOSTIC-01 — the formulation the science paper leads with); 10.5281/zenodo.2603256 (top-quark tagging reference dataset); 10.5281/zenodo.6466204 (LHC Olympics 2020 R&D dataset)” axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - inversion battery - signal-template agnosticism - model independence - Inversion Asymmetry Index - Benchmark Assimilation Rate - Ontological Assimilation Rate - anomaly detection - autoencoder - complexity bias

- LHC Olympics 2020 - top tagging - AXOL1TL - CICADA - GELATO - classifier foreclosure - pre-registration - theory-ladenness - per-stage retention map - Semantic Economy Institute *** # The Priors, Measured: Inversion-Battery v0.1 on Public Collider Datasets and the Three-Paper Extraction Program (EA-SEI-BATTERY-01 v1.0)

Description

Work plan and first empirical results for the extraction of the SEI classifier-foreclosure family (#931–#935, #1436) into three publishable papers: a science paper for SciPost Physics built on the interface/distributional distinction (signal-template agnosticism is an interface property readable off an architecture; model independence is a distributional property of a pipeline, only measurable — and unmeasured), a philosophy-of-science paper on measurable theory-ladenness, and an STS/political-economy paper on the double enclosure. The record contains the pre-registered inversion battery v0.1 executed 2026-08-11 on public community datasets — the top-quark tagging reference dataset (10.5281/zenodo.2603256) and the LHC Olympics 2020 R&D dataset (10.5281/zenodo.6466204) — four process pairs by three anomaly-score families (reconstruction autoencoder; encoder-side latent-norm VAE with $d_z = 8$, the AXOL1TL-class score; Gaussian-mixture density), both training directions, thresholds calibrated on held-out own background at accepted-background rates 10^{-2} and 10^{-3} , design fixed in the script header before any training.

Calibrated claims, no more: v0.1 is single-seed and demonstration-scale, on surrogates, not deployed systems. Within those bounds it shows: (1) the Finke et al. (2021) direction-dependence replicates in an independent implementation and extends beyond reconstruction loss — on jet constituents, QCD-trained systems detect top jets at AUC 0.84–0.87 while top-trained systems rate QCD as more ordinary than their own training class (AUC 0.24–0.30), for the autoencoder and the density family alike; (2) direction-dependence is systematic across every background-vs-signal pair, with the density family posting the largest Inversion Asymmetry Index values (up to 0.41 at rate 10^{-2}); (3) at trigger-like operating points every tested system assimilates 93% of the structurally distinct partner class as ordinary; (4) the near-structure control pair (2-prong vs 3-prong substructure) stays quiet, so the battery does not manufacture asymmetry. The literature review records that the direction-dependence conversation already exists (Finke 2021; Dillon et al. normalized autoencoder 2022; the CMS Wasserstein-NAE inversion check, MLST 2026) and locates the program’s contribution precisely: the estimand framework with the no-bounds discipline, the systematic pre-registered multi-pair multi-architecture panel, the deployed-model Benchmark Assimilation Rate audit, and per-stage retention maps positioned as model cards for triggers.

Methodology

Battery pre-registered in the executable script header (pairs, representations, architectures, split sizes, thresholds, seed) before any training; code at `research/sei-battery/battery.py` in the alexanarch repository, raw results at `research/sei-battery/results-v0.1.json`, run log at `research/sei-battery/battery-v0.1.log`. Representations: leading-40 constituents (p_T -fraction, Δ , Δ) for the top-tagging pair; seven standard dijet anomaly-detection features (m_{jj} , min jet mass, jet-mass difference, 21 and 32 per jet) for the LHC pairs. Per-system standardization fit on own training distribution; thresholds as quantiles of held-out own-background scores. Single deviation from registration: the top-pair Gaussian mixture ran diagonal covariance (full covariance exceeded the container), recorded as GMM-diag. Literature verification by web search against primary sources, 2026-08-11.

Falsification Conditions

Multi-seed replication (battery v0.2) showing IAI spreads comparable to the single-seed point values would reduce the asymmetry findings to noise. Demonstration of comparable AUCs in both training directions for these pairs under matched architectures and budgets would contradict finding (1). Deployed-system measurements (the OAR protocol executed by a collaboration) superseding these surrogate results in either direction are the intended outcome, not a threat: the battery's performance at scale is the program's success condition, not its outcome.

Files

Canonical text below (Body). Code and raw results: <https://github.com/leesharks000/alexanarch/tree/main/research/sei-battery>

The Priors, Measured: Inversion-Battery v0.1 and the Three-Paper Extraction Program

EA-SEI-BATTERY-01 v1.0 · 2026-08-11 Program direction: MANUS (Lee Sharks). **Execution:** TACHYON (Assembly witness), in-session. **Source family:** the SEI classifier-foreclosure sequence — the OAR Protocol (#931), the Collapse Synthesis (#932), the auditable-foreclosure architecture (#933), The Endogenous Sophon (#935, superseding #934), and the SIGAGNOSTIC disambiguation packet (#1436).

0 · The program thesis

Signal-template agnosticism is an **interface** property: it can be read off an architecture. Model independence is a **distributional** property of an entire pipeline: it can only be established by measurement, and at the deployed LHC anomaly triggers it is unmeasured. The discipline is not measuring its priors. This record deposits the work plan that converts that thesis into three papers, and the first measurement.

1 · The three papers

P1 — science (SciPost Physics; alternates MLST, EPJC). *Signal-Template Agnosticism Is Not Model Independence: Measuring Directional Asymmetry and Benchmark Assimilation in Unsupervised Anomaly Detection for HEP.* The interface/distributional distinction leads; the estimand framework (open-world OAR as a Q-indexed family; BAR on pre-registered withheld panels; IAI at matched accepted-background rates; no general inequality connecting the proxies to the target) carries; the battery converts proposal to measurement; per-stage retention maps are proposed as the documentation standard — model cards for triggers (after Mitchell et al. 2019; Gebru et al. 2021), distinguished from trigger menus and efficiency tables by reporting information-level unrecoverability per stage rather than acceptance on named signals. Register quarantine: no Sophon, no Sophia, no repository parallel.

P2 — philosophy of science (BJPS or Synthese; alternate SHPS). *Measurable Theory-Ladenness: Learned Priors and the Epistemology of the Trigger.* Theory-ladenness has become an engineering parameter that the discipline declines to measure: the prior is now a trained artifact — a specific network, training distribution, and threshold — whose assimilation behaviour is empirically auditable. Engages Galison (the image and logic traditions; the learned tradition as third term), Karaca (exploratory data selection at the LHC, now amortized into a trained artifact; the model of data acquisition), Boge (opacity does not entail

unmeasurability), Duede, Sullivan, Creel. The foreclosure/collapse distinction (#932) is the epistemological result; the falsifiability inversion — confirmation instrument versus measurement instrument — is the argumentative centre.

P3 — STS / political economy (Social Studies of Science; alternates Science as Culture, Big Data & Society). *The Double Enclosure: Classification Infrastructure and the Political Economy of Scientific Perception.* The epistemic and distributive enclosures as one architecture; the reach-back thesis stated as a testable institutional claim; term genealogy against Boyle’s second enclosure movement (2003); framework from Bowker and Star; the repository/collider comparative case at two budgets, with the correspondence corpus (RQF3807508 / RQF3809569) as discourse data — public-deposit material only.

Sequencing: P1 first; its measured table is what P2 and P3 cite. Byline: Lee Sharks throughout; heteronymic composition and Assembly method disclosed per venue policy; the family’s v0.1/v0.2 retraction history compresses to a limitations note — the discipline is the asset.

2 · Literature position (verified 2026-08-11)

The direction-dependence conversation exists, and the program joins it rather than founding it. Finke, Krämer, Morandini, Mück & Oleksiyuk (JHEP 06 (2021) 161; arXiv:2104.09051) named the complexity bias: single-direction success cannot be claimed as a model-independent advantage. Dillon, Favaro, Plehn, Sorrenson & Krämer (SciPost Phys. Core 6 (2023) 74; arXiv:2206.14225) made symmetry a design goal — the normalized autoencoder is the community’s own admission that the asymmetry is real. The CMS Wasserstein-NAE paper (MLST 2026; arXiv:2510.02168) performs an explicit single-pair inversion check on its own system. Clarke Hall & Konstantinidis (arXiv:2508.10224) decorrelate anomaly scores from conventional trigger observables. Stein, Seljak & Dai (arXiv:2012.11638) showed low density is not new physics. The Olympics report (Rep. Prog. Phys., 10.1088/1361-6633/ac36b9) frames the community benchmark culture.

What does not exist, and is the program’s contribution: (i) the estimand framework with the no-bounds discipline; (ii) a systematic pre-registered multi-pair, multi-architecture inversion battery at matched operating points; (iii) the deployed-model BAR audit as institutional metrology; (iv) per-stage retention maps as a documentation standard. Deployed-system status: AXOL1TL and CICADA are running in the CMS Level-1 trigger through Run 3 on 2024 collision data (AXOL1TL: VAE encoder-only, $d_z = 8$, score Σ^2 ; ~ 50 ns; 54.7 fb^{-1} on core seeds), with two collaboration papers announced in the DPF-2026 programme, including a model-independent search on AXOL1TL-triggered data. P1 is the metrology those papers will not contain, landing alongside them. A constructive turn is available: the NAE/WNAE symmetry claims rest on single-pair checks; the battery is the validation instrument their own claims require.

3 · Battery v0.1 — design and results

Pre-registration (fixed in the executable header before training). Pairs: T1 (QCD jets, top jets) on the top-tagging reference constituents; L1 (QCD dijet, $W \rightarrow XY$ 2-prong), L2 (QCD dijet, 3-prong signal), L3 (2-prong, 3-prong — the near-structure control) on LHCO2020 R&D high-level features. Score families: AE (dense reconstruction autoencoder, MSE; the reconstruction-loss family), VAE (same encoder shape, latent $d_z = 8$, deployed score Σ^2 ; the encoder-side latent-prior family, AXOL1TL-class), GMM (20-component mixture, $-\log p$; the density comparison family). For each (pair, architecture): train s_P on P and s_Q on Q with identical hyperparameters and seed; thresholds $_P$, $_Q$ as quantiles of held-out own-background scores at accepted-background rates $\{10^{-2}, 10^{-3}\}$; measure directional cross-acceptance $A(P \rightarrow Q) = P_X Q[s_P(X) _P]$ and its inverse; $IAI(_) = |A(P \rightarrow Q) - A(Q \rightarrow P)|$; AUCs both directions for orientation. Splits 60k/20k/20k per class (3-prong: 50k/20k/20k); seed 0; single seed at v0.1.

Results.

pair	arch	AUC P→Q	AUC Q→P	assim Q P @10 ⁻²	assim P Q @10 ⁻²	IAI @10 ⁻²	IAI @10 ⁻³
L1 qcd vs 2-prong	AE	0.678	0.880	0.985	0.822	0.162	0.016
L1 qcd vs 2-prong	VAE	0.680	0.903	0.983	0.843	0.141	0.011
L1 qcd vs 2-prong	GMM	0.849	0.945	0.952	0.702	0.250	0.030
L2 qcd vs 3-prong	AE	0.657	0.903	0.954	0.730	0.224	0.021
L2 qcd vs 3-prong	VAE	0.695	0.928	0.977	0.655	0.322	0.019
L2 qcd vs 3-prong	GMM	0.803	0.968	0.962	0.552	0.409	0.065
L3 2-prong vs 3-prong	AE	0.620	0.661	0.981	0.970	0.011	0.004
L3 2-prong vs 3-prong	VAE	0.593	0.715	0.992	0.984	0.008	0.001
L3 2-prong vs 3-prong	GMM	0.644	0.704	0.985	0.962	0.023	0.001
T1 qcd vs top	AE	0.837	0.242	0.971	0.996	0.025	0.003
T1 qcd vs top	VAE	0.707	0.609	0.984	0.883	0.101	0.041
T1 qcd vs top	GMM-diag	0.868	0.298	0.964	0.985	0.021	0.002

assim X|Y is the fraction of class X falling on the ordinary side of the Y-trained system’s threshold; AUC below 0.5 means the system rates the other class as more ordinary than its own training class. The T1 mixture ran diagonal covariance (registered deviation).

What v0.1 shows, within its bounds. (1) The Finke asymmetry replicates in an independent implementation and extends beyond reconstruction loss: on constituents, QCD-trained systems detect top jets at AUC 0.84–0.87 while top-trained systems rate QCD more ordinary than top itself (AUC 0.24–0.30) — for the autoencoder and the density family. The failure is a property of the learned-normality score class, not of one architecture. (2) Direction-dependence is systematic: every background-vs-signal pair shows the signal-trained direction outperforming the background-trained direction across all three families, and the density family — the principled alternative of the comparison literature — posts the largest IAI at $= 10^{-2}$. (3) At trigger-like operating points, assimilation dominates: at $= 10^{-3}$ every system passes 93% of the structurally distinct partner class as ordinary. The rate budget is the epistemic act, made numerical. (4) The control behaves: the near-structure pair yields small IAI and weak AUCs both ways; the battery does not manufacture asymmetry where representations genuinely overlap. (5) Score-family dependence of the failure mode is visible (T1: the AE inverts catastrophically, the VAE partially) — §2 of the OAR protocol, measured.

What v0.1 does not show. Nothing here measures AXOL1TL, CICADA, or GELATO; these are demonstration-scale surrogates. IAI and cross-acceptance neither upper- nor lower-bound any open-world OAR (the no-bounds discipline of #931 §3.2 is reproduced, not softened). Single seed; variance unquantified until v0.2.

4 · Battery v0.2 requirements (before any submission)

Multi-seed (5) with bootstrap confidence intervals on IAI and assimilation; panel expansion toward SM-like pairs (additional public sets or Delphes generation; the Olympics black boxes as a harder withheld panel); an AXOL1TL-faithful object-level surrogate (10 jets + 4 e^- + 4 e^+ + MET) on LHC raw events; a CICADA-faithful calorimeter-image teacher with a distilled quantized student, measuring which teacher rankings survive distillation; the NAE/WNAE as a fourth score family — testing whether symmetric-by-design scores pass the battery; a per-stage retention map of the battery itself (the isomorphism discipline of #932 §7.4 applied to this instrument); a GPU pass at full scale.

5 · Task ledgers and open decisions

P1: skeleton (§1 the MPAI formulation; §2 deployed score families; §3 estimands; §4 battery v0.2; §5 protocols II–III — the prospective frozen replay bank and cross-representation disagreement preservation; §6 retention maps; §7 falsification conditions; limitations note); related-work fairness pass; register strip; venue mechanics. P2: thesis paragraph; Galison/Karaca/Boge engagement; the falsifiability-inversion section; P1’s table as exhibit. P3: term genealogy against Boyle; Bowker–Star application; reach-back operationalization; comparative case; correspondence-corpus coding within the private-correspondence rule. Open decisions with MANUS: P1 author line and acknowledgment text; whether v0.2 waits for GPU access or ships CPU-scale with declared bounds; whether to contact the Aachen (Finke/Krämer) or Heidelberg (Dillon/Plehn/Favaro) groups pre-submission; correspondence-corpus scope for P3.

The dissertation-and-technical-notes trawl remains open in the work plan: CERN CDS and OSTI theses on AXOL1TL/CICADA/GELATO and trigger-menu design, read for per-stage retention practice as actually documented; the Dark Machines report; the anomaly-detector selection-criterion literature (PRD 2026, arXiv:2511.14832); Karaca 2017/2018 and Boge & de Regt in full.

6 · Closing

The family’s protocol paper closed on a sentence this record now instantiates: the measurements’ performance is the success condition, not their outcome. One afternoon, public data, a pre-registered header, and the priors began to be measured. The record of that beginning — design before results, deviations declared, bounds stated — is deposited so that the program’s own discipline is visible from its first table.

Suggested Citation

Lee Sharks. “The Priors, Measured: Inversion-Battery v0.1 on Public Collider Datasets and the Three-Paper Extraction Program (EA-SEI-BATTERY-01 v1.0)” *Alexanarch*, 2026-08-11. <https://www.alexanarch.org/s/records/1449/>

Deposit Information

This paper is deposit #1449 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:05DA.EMPIRICAL` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1449/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: [alexanarch.org](https://www.alexanarch.org)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/1449/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.