

YUCLAW v5.1

User Guide — Evidence-First Financial AI. Every signal traces to a real SEC filing.

An open research laboratory where market hypotheses are tested, challenged, and preserved — with every conclusion carrying its own audit trail.

The YUCLAW Team · August 2026 · yuclaw.ca · github.com/YuClawLab/yuclaw-brain · `pip install yuclaw` · Apache-2.0 License

Disclaimer. YUCLAW is a research and education platform. Nothing in this guide or on any YUCLAW surface is investment advice, a recommendation, or an offer of any product. Signal labels are research classifications, not buy/sell recommendations. Past results — in-sample or forward-tracked — do not predict future performance. *Investment implication: none established — no buy, sell, or alpha conclusion is supported by this material.*

Why YUCLAW exists

Financial markets generate millions of claims every day — headlines, scores, ratings, predictions. Almost none of them can be checked. The claim arrives; the evidence, the assumptions, the track record of similar claims, and the list of things that were tried and failed — all of that stays hidden. Claims without evidence become noise, and noise with confidence becomes dangerous.

YUCLAW was built on a simple refusal: **never publish a number a stranger cannot recompute**. Every signal traces to a real filing. Every statistic carries its protocol, its sample anatomy, and its count of how many other things were tried. Every failed hypothesis stays on the record, because knowing what does not work is half of knowing anything at all.

The product is not a prediction. The product is a record you can check — and a set of instruments that check it before you do.

Contents

- | | |
|---|--|
| 1 · What YUCLAW is (and is not) | 6 · Reading the statistics honestly |
| 2 · Quick start in three minutes | 7 · Verifying everything yourself |
| 3 · A map of the website | 8 · The research engines, in plain language |
| 4 · Signal labels and published thresholds | 9 · The protocol registry and pre-registration |
| 5 · Command-line reference | 10 · Evidence packets, exports & the AI layer |
| 11 · Bring Your Own Signal (research service) | 13 · Status: what is and is not established |
| 12 · Known limitations | 14 · FAQ · 15 · Glossary · 16 · Citing YUCLAW |

1 · What YUCLAW is (and is not)


YUCLAW ingests real SEC filings (8-K, 10-K, Form 4, and the foreign-private-issuer track: 6-K, 20-F, 40-F), extracts typed evidence events with locally hosted language models, classifies them under a fixed public vocabulary, and studies whether that evidence stream carries information — publishing every result **exactly as measured, including where it is adverse to the hypothesis**. All inference runs on local hardware; the full pipeline, data derivations, and methodology are open source under Apache-2.0.

Three design commitments define the platform:

- **Evidence first.** Every accepted event links to a primary filing with an accession number and a verified excerpt. Every research memo cites its sources sentence by sentence, and citations are machine-checked before anything renders.
- **Pre-registration.** Statistical protocols are locked into a tamper-evident, append-only registry *before* they touch data. Editing a locked specification is impossible without a visible supersession. Test counts are ledgered, and the expected number of chance findings is printed beside the results.
- **Self-audit.** The platform runs falsification tests, baseline comparisons, and robustness grids against its own headline results — and publishes the outcomes. Several of YUCLAW's own hypotheses have been retired by its own instruments; those retirements remain on the record.

What YUCLAW is not: it is not a stock picker, not a trading system, not an advisory service, and not a performance product. It never publishes sell or short recommendations. If you are looking for something to tell you what to buy, this is the wrong tool — deliberately.

1.1 · How it fits together



Filings become cited events; events feed either the scored universe or the never-scored evidence tier; all research runs under pre-registered protocols; gates stand between every result and the public record; and the dashed loop is the point — the public record feeds independent verification back against the source.

2 · Quick start in three minutes

```
pip install yuclaw
yuclaw demo          # guided 3-minute tour — works on any machine, no setup
yuclaw replay-lab    # reproduce the published Lab statistics from the public bundle
yuclaw verify AMD --date 2026-05-20 # check a signal against the public ledger
```

Every command in the read/verify family works from any computer with internet access — they consume only the published public bundle. Commands that read the research backend (`events`, `lens`, `memo`, `export`) require access to a YUCLAW research node; on other machines they print a clear one-paragraph explanation and exit with code 3 — never a stack trace.

First thing worth doing: run `yuc law replay-lab`. It recomputes every statistic and ledger root on the Validation Lab from published data. Exit code 0 means your machine just reproduced the site. That single command is the platform's thesis in miniature.

3 · A map of the website (yuclaw.ca)

Page	The question it answers
Home	Current research classifications for the 79-ticker scoring universe, with composite scores. Caption on the table: <i>research classifications — not recommendations</i> .
Validation Lab	Does the composite signal contain measurable forward information? Deciles, information coefficients, clustered confidence intervals, baseline comparisons, calibration, neutralized ICs — with negative and inconclusive results shown as measured.
SMH Covered-Constituent Evidence Lens	Event-study analysis of the YUCLAW-covered share of disclosed SMH weight (coverage percentage stated dynamically on the page). Explicitly not a full-fund inference.
XLK Evidence Lens	The second theme lens — admitted because it passed the published lens admission standard, not by popularity.
Canada Resources Evidence	A source-traceable evidence view of four Canadian resource lenses (XEG, ZEO, GDX, URM) built on the 6-K/40-F filing substrate. Evidence-tier names are never scored.
Forward Tracking	The forward out-of-sample ledger: one period per trading day since Day 0, plus the label-calibration panel measuring what each classification has historically preceded.
Ledger	Daily SHA-256 roots committed to a public git repository before pages update — the anchor that makes retroactive edits visible.
Methodology	Backfill disclosure, estimator ladder, score → label thresholds, corporate-action and price-adjustment policies, trading-day conventions.
Weekly Evidence Note	A Friday note generated from the registry and evidence store directly; a build gate fails the site if its counts ever disagree with the chain.
Today's Evidence Changes	What changed in the evidence substrate — last completed UTC day plus the day in progress.
Replication	How to reproduce the record, the vocabulary of verification, and the (honestly empty until earned) log of independent replications.

4 · Signal labels and published thresholds

YUCLAW uses a fixed eight-label vocabulary. Labels are **sentiment-style research classifications**, chosen deliberately over directive verbs. The exact score → label mapping is published so classification is reproducible:

Composite score	Label	Reading
$\geq +0.55$	STRONG_BULLISH	strongly positive evidence posture
$\geq +0.40$	BULLISH	positive evidence posture
$\geq +0.20$	NEUTRAL	balanced evidence
≥ 0.00	WATCH	mixed; monitoring
≥ -0.20	WEAKENING	deteriorating evidence posture
≥ -0.40	NEGATIVE_EVENT	a materially negative evidence event dominates
< -0.40	BEARISH_WATCH	strongly negative evidence posture
—	RISK_ALERT	risk-state overlay; explicitly <i>not</i> score-mapped

Two honesty notes travel with this table. First, **the thresholds carry no outcome promise; the calibration panel measures the outcomes** — as of this guide's writing, pooled directional consistency has a confidence interval that includes 0.5, so directional meaning is not yet demonstrated, and the panel says so. Second, the evidence-impact component (C6) carries its own standing finding: insider-sale evidence may function more plausibly as a risk-state input than a near-term directional signal — which is why a name can carry heavy insider-sale evidence and a cautionary label without any directional claim being made.

5 · Command-line reference

Command	What it does	Runs where
<code>yuclaw demo</code>	Guided three-minute tour of the platform's ideas.	anywhere
<code>yuclaw why TICKER --as-of DATE</code>	Signal anatomy: which components and events produced a score, point-in-time.	anywhere*
<code>yuclaw verify TICKER --date DATE</code>	Checks a signal against the public ledger; prints VERIFIED or a plain-language explanation of what could not be checked on this machine.	anywhere
<code>yuclaw replay-lab</code>	Recomputes the published Lab statistics and ledger roots from the public bundle. Exit 0 = full reproduction.	anywhere
<code>yuclaw events --ticker T [--since D]</code>	Lists accepted evidence events with filing citations.	backend
<code>yuclaw lens NAME</code>	Lens summary: coverage, admission verdict, current posture.	backend
<code>yuclaw memo --ticker T</code>	Generates an evidence memo with per-sentence, machine-verified citations under a restricted conclusion vocabulary.	backend
<code>yuclaw export ...</code>	Exports YUCLAW-derived data (events, signals, statistics). Raw vendor price/options data is never exported.	backend

*why works anywhere for dates covered by the bundled demo data; full history requires the backend. **Exit codes:** 0 = success · 2 = usage error (message shows correct syntax) · 3 = environment cannot support the command (clear explanation, no traceback). If you ever see a raw Python traceback, that is a bug — please report it.

6 · Reading the statistics honestly

Every statistic on YUCLAW carries machinery designed to keep you from over-reading it. Learn these five habits and the site reads correctly:

6.1 · Badges are the grammar

Badge	Meaning
PRELIMINARY	The pre-registered interval excludes the null <i>and</i> sample floors are met. The strongest label the system awards — deliberately, nothing stronger exists.
DESCRIPTIVE	Measured and reported, but the interval includes the null. A description, not a conclusion.
UNDERPOWERED	Below pre-registered sample floors. Shown for completeness; statistically uninformative.
DATA-LIMITED	The input is frozen, missing, or constant; no statistic is computed on it.
INCONCLUSIVE	A gated read whose arms did not meet pre-committed minimums. Reported, never decorated.

6.2 · Clustered intervals outrank naive ones

Events sharing an issuer or a calendar day are not independent. Wherever both appear, the **cluster-aware interval is primary** and the naive one is labeled naive. The conservative envelope is the wider of the issuer-clustered and date-

clustered intervals. An interval that narrowly excludes zero should always be read alongside its falsification context — the pages place them together deliberately.

6.3 · The false-positive ledger is printed on purpose

The registry counts every secondary statistical cell ever computed and prints the expected number of chance findings at $\alpha=0.05$ beside the panels. If the site shows hundreds of exploratory cells, some will look significant by luck — the ledger tells you how many, before you are impressed.

6.4 · Falsification context travels with headlines

Headline event-study results are accompanied by placebo tests: date-shuffle nulls (does the timing matter?), direction randomization, and pre-event windows. Example of the house style, from the SMH lens: an adverse pooled result whose date-shuffle percentile is unremarkable is described as *era-generic direction alignment, not event-timed information* — the result stands, and its most deflationary explanation stands next to it.

6.5 · Sample anatomy beats sample size

Beside every panel: unique issuers, median events per issuer, top-issuer share, window overlap. Five hundred filings can be three stories (see §8). The anatomy tells you which.

7 · Verifying everything yourself

YUCLAW's core claim is that you do not have to trust it. Four verification workflows, in increasing depth:

7.1 · Reproduce the Lab (one command)

```
pip install yuclaw && yuclaw replay-lab
# exit 0: every statistic and daily ledger root recomputed from published data
```

7.2 · Full replication in a clean environment

```
git clone https://github.com/YuClawLab/yuclaw-brain && cd yuclaw-brain
make replicate # scrubbed-env fetch → verifier → diff against published outputs
```

The replication page defines the verifier vocabulary — an *exact day* (recomputed root matches byte-for-byte) versus an *anchored-subset day* (an intraday anchoring artifact, disclosed, in which every anchored hash still recomputes unchanged). Independent replications can be reported via the repository's issue template; the public log lists them as they are earned.

7.3 · Verify the protocol registry chain

```
curl -s0 https://raw.githubusercontent.com/YuClawLab/yuclaw-brain/main/registry/
protocols.jsonl
curl -s0 https://raw.githubusercontent.com/YuClawLab/yuclaw-brain/main/tools/
yuclaw_protocol_registry.py
python3 -c "import yuclaw_protocol_registry as r; \
    reg=r.Registry('protocols.jsonl'); print('chain OK:', reg.verify_chain())"
```

The registry is an append-only hash chain: every line carries the hash of the line before it. Any edit or deletion anywhere in history breaks the chain for every verifier on Earth. You can further recompute a protocol's `method_hash` from the committed specification text and compare — confirming specifications were not altered after registration.

7.4 · Trace one piece of evidence end to end

The site's evidence-trace walk-through follows one real filing from EDGAR through exhibit extraction, event typing, grading, and posture — every step linked. Ten minutes, and you will know exactly what an "accepted event" is.

8 · The research engines, in plain language

Three engines run routinely over every lens and every registered result. Their panels appear across the site under plain names:

8.1 · Evidence structure — "how many independent facts is this, really?"

Events that share an issuer and a time window are usually one story, not many signals. The engine clusters events into stories under a pre-committed linkage rule and computes an *effective* evidence count via design effects. A real example from the SMH lens: **513 filings** → **81 distinct events** → **3 stories**, with one story carrying 95% of the mass. Statistics on the page use cluster-aware inference accordingly — and now you know why.

8.2 · Context robustness — "where does this conclusion survive, and where does it break?"

Every registered primary result is re-estimated across a pre-declared grid — horizons, volatility regimes, trend regimes where computable, populations, eras — and summarized honestly: *sign held in X of Y cells; breaks in: [list]*. No pass/fail verdicts exist; cells that cannot be computed say so rather than being estimated around. A standing rule makes the grid part of what "done" means: every new registered result receives its robustness profile within one cycle.

8.3 · Evidence lifecycle — "how does filed information emerge, peak, and fade?"

Per event type: the average magnitude path after the event, time-to-peak, and half-life where reached. The current honest reading across qualifying types is *no decay distinguishable from diffusion* within the studied window — an informative null, published like any other result.

9 · The protocol registry and pre-registration

The registry is the platform's constitution. How to read it:

- **A protocol** is a locked specification — question, population, exactly one primary endpoint, secondary endpoints, analysis plan, sample floors, badge rules — registered *before* computation, identified by a hash of its own text.
- **A run** records a computation under a protocol: data window, cell counts (which feed the false-positive ledger), and a hash of the results artifact.
- **A supersession** is the only way a specification changes. The old version is preserved forever, mechanically refuses further runs, and a notice links old to new. Version lineages are public — including, in one case, a supersession that permitted late computation of a process-failure gap strictly from as-of inputs, with per-artifact lag disclosed and a mechanical verifier enforcing the constraint.
- **Questions** track open research hypotheses — including retired ones. YUCLAW's registry currently records, among others, the retirement of the simple insider-selling directional hypothesis on three pre-registered grounds. Failed hypotheses are preserved, not deleted: they are part of the record's value.

Why this matters to you: when a YUCLAW page shows a result, the criteria for that result were fixed before the data arrived, the number of other things tried is printed, and the specification's text can be hash-checked by anyone. Those three properties together are what "harder to fake" means in practice.

10 · Evidence packets, exports & the AI layer

10.1 · Evidence packets

Each research surface ships a downloadable packet: coverage statement, event-study summary, events CSV with event IDs, engine outputs, METADATA (data-through date, build commit, ledger root, methodology version, known limitations) and a ready CITATION.txt. Packets regenerate in the daily chain.

10.2 · Export rules — stated plainly

YUCLAW-derived data is exportable: signals, components, classifications, extracted events from public SEC filings, hashes, statistics, the track record. Raw vendor market data (price series, options data) is never redistributed — that is a licensing boundary, disclosed rather than blurred.

10.3 · For AI agents

YUCLAW positions itself as an open evidence layer underneath AI research tools. Machine surfaces: `yuc law . ca / l l m s . t x t` (identity, vocabulary, endpoints, citation format, the not-advice frame) · `evidence_index . j s o n` (a machine-readable map of every page, packet, and protocol with stable URLs and data-through dates) · MCP server tools mirroring the CLI with the same friendly no-backend behavior. Agents consuming YUCLAW inherit its citation discipline: event-level claims cite event IDs from the packet CSVs.

11 · Bring Your Own Signal (research service)

YUCLAW's validation machinery can be applied to *your* signal. You provide a point-in-time CSV (`date, ticker, signal_value`); the platform runs the full diagnostic suite — information coefficients with cluster-aware intervals, quantile monotonicity, churn, horizon decay, placebo tests — under a client-namespace protocol, and returns a hash-stamped memo plus a reproduction bundle whose standalone script re-derives every number.

- **Intake validation:** wrong columns, unparseable dates, NaN/infinite values, duplicate ticker-days, future-dated or lookahead-suspicious rows are each rejected with a row-numbered, plain-language message.
- **Honest ceilings:** user-defined lenses cap at EXPLORATORY (CLIENT) — higher maturity labels are reserved for the canonical public record and cannot be purchased.
- **Custody:** client files never enter the public repository or any backup path; machine-enforced, with deletion on request.
- **Sample honesty:** baskets under 15 names will carry power caveats — that is a property of sample size, not a defect in your signal, and the memo says so.

This service runs as a concierge research engagement during the current pilot phase. Terms, delivery, and scope are provided on inquiry.

12 · Known limitations

- **The forward record is young.** The out-of-sample ledger accrues one period per trading day; most forward statistics remain DESCRIPTIVE or UNDERPOWERED at current sample sizes, and the site says so wherever it applies.
- **The scoring universe is 79 tickers, frozen.** Evidence-tier names (Canada Resources; foreign SMH constituents) are ingested and analyzed but never scored — enforced by positive gating and a standing negative check.

- **Theme-lens coverage is partial and stated.** Lens pages analyze the covered constituent sleeve only, with the covered-weight percentage in the title area and an itemized account of what is uncovered and why.
- **Prices are unadjusted closes.** Dividends are not reinvested (a disclosed drag consideration for high-yield names); splits are handled by window exclusion, never retroactive re-pricing.
- **In-sample replay carries look-ahead caveats.** The backfill era is labeled as in-sample replay, with the language-model look-ahead disclosure in methodology; it is never blended with the forward ledger.
- **Some analyses await data or accrual.** Cells that cannot be computed honestly (insufficient history, thin arms, absent factor data) render as stated gaps rather than estimates.
- **One disclosed outage window** exists in the record; it is documented, not repaired.

13 · Status: what is and is not established

This mirrors the shared status block rendered on the site from a single source. Statuses are measured, not aspirational; the live pages are canonical.

Status	Item
✓ Verifiable today	Replay reproduces every Lab statistic and ledger root from published data · daily SHA-256 roots anchored to public git before pages update · every accepted event traces to a filing with accession number and verified excerpt · coverage stated as measured · snapshots are point-in-time with zero retroactive edits · evidence-tier names are never scored.
✗ Not established	Forward alpha — no spread, IC, or alpha significant at 5% with adequate power · directional meaning of the label set (calibration interval includes 0.5) · C6 risk-gate sign confirmation (first computable read scheduled under a locked protocol) · event-study lead over peer models.
· Accruing	The forward out-of-sample record, daily · matured event-study windows · the C6 elevated arm under live ingestion · robustness grids over new results · external replications, as reported.

14 · Frequently asked questions

Is this investment advice? No. Research and education only. No YUCLAW output is a recommendation, and the platform never publishes sell or short recommendations.

Why do you publish results that hurt your own hypothesis? Because a record that only contains favorable results is not a record. Several adverse findings on the site — including baseline comparisons the composite loses at current sample sizes — are published under the same pre-registered protocols as everything else. The platform's product is the trustworthiness of its process.

A label looked wrong in hindsight — is that a bug? No. Labels are point-in-time research classifications; the calibration panel measures, publicly, what each label has historically preceded — including where the answer is "nothing demonstrable yet."

Why is coverage of an index fund partial? Because YUCLAW only analyzes constituents inside its disclosed universe and filing substrate, and refuses to imply otherwise. The uncovered remainder is itemized with reasons on each lens page.

Can I use YUCLAW data in my own research or product? Yes — derived data is Apache-2.0-licensed and exportable with ready citations. Raw vendor market data is not redistributed.

How do I report a bug or a replication? Via GitHub issues on the repository; replication reports have a dedicated template and, once verified, enter the public log.

Who is the YUCLAW Team? An independent research operation based in Calgary, Alberta, running entirely on sovereign local hardware. The platform, methodology, and this guide are published under the YUCLAW Team name; the record itself — the chain, the ledger, the registry — is the authority, not any individual.

15 · Glossary

accepted event	A typed evidence item extracted from a filing that passed grounding checks; carries a source URL, accession number, and verified excerpt.
CAR	Cumulative abnormal return: return relative to a benchmark model, cumulated over event-time trading days.
cluster-aware CI	A confidence interval that respects dependence between events sharing an issuer or date; primary wherever shown.
conservative envelope	The wider of the issuer-clustered and date-clustered intervals.
estimand	A precisely defined quantity being estimated (e.g., capped-ETF-weighted mean CAR at +20 trading days).
evidence-tier name	A ticker that is ingested and analyzed but never scored; excluded from the 79-ticker scoring universe by machine-enforced gating.
IC	Information coefficient: cross-sectional rank correlation between a score and forward returns.
N _{eff} / design effect	The effective number of independent observations after accounting for clustering; $N_{\text{eff}} = N / \text{DEFF}$.
point-in-time	Recorded with the information available at that moment and never edited afterward.
protocol / run / supersession	See §9 — the registry's three record types.
story	A cluster of events linked by issuer and time under a pre-committed rule; the unit of independent evidence.

16 · Citing YUCLAW

```
YUCLAW <page name>, data through <date>, build <commit>,  
https://github.com/YuClawLab/yclaw-brain  
# or use the ready CITATION.txt inside any evidence packet;  
# event-level claims cite event IDs from the packet CSVs.
```

Links: yuclaw.ca · GitHub github.com/YuClawLab/yclaw-brain · PyPI pypi.org/project/yclaw · machine surface yuclaw.ca/llms.txt

Final disclaimer. Hypothetical research illustration. Not investment advice, not performance advertising, not an offer of any product. Research classifications are not recommendations. Past results — in-sample or forward-tracked — do not predict future performance. © 2026 The YUCLAW Team · Apache-2.0 License · research & education only.