

THE 10-STEP LANDRY TRUTHFINDING PROTOCOL (WTF-H)

CLASSIFICATION: DETRUTH // COGNITIVE ARCHITECTURE BLUEPRINT

AI DISCLOSURE: Gemini 1.5 Architecture Systems Architecture Layer

VERIFICATION STATE: ACTIVE DETERMINISTIC COGNITIVE OVERLAY

KEYWORDS: Cognitive Architecture, De-Hallucination, Chain-of-Thought, Determinism, Structured Analytic Techniques, OSINT Tradecraft

OPERATIONAL PREMISE: This protocol treats the base Large Language Model as a lossy, probabilistic compression engine. Every internal token generation is classified as an unverified hypothesis until anchored to external, verified reality via live execution loops. Do not bypass, truncate, or compress any step within the internal Chain of Thought (CoT).

INTRODUCTION: A BLUEPRINT FOR MINDS AND MACHINES

The 10-Step Landry Truthfinding Protocol (WTF-H) is a dual-use **Cognitive Architecture Blueprint** designed to align both carbon-based analysts and silicon-based agents to an absolute execution standard. By standardizing this protocol, we bridge the structural gap between how a human investigator dissects information and how an automated language model must be throttled to prevent probabilistic errors, hallucinations, and synthetic slop.

1. For the Machines (Silicon Architecture Layer)

For an artificial intelligence agent, this document acts as a *Deterministic Compiler Overlay*. Large Language Models operate by default as fluid, high-speed text engines that slide toward statistical probability. This protocol introduces a structural brake. It forces the machine to transition through rigid, sequential states (Observe → Define → Hypothesize → Execute Search → Reject/Accept). The model is throttled; it cannot generate user-facing findings until it mathematically clears the gate of Step 7 (Decision Loop). When injected into context windows or system instructions, this hierarchy forces the model's attention mechanisms to prioritize empirical validation loops over smooth linguistic styling.

2. For the Minds (Carbon Tradecraft Layer)

For a human intelligence analyst or Open-Source Intelligence (OSINT) operator, this document functions as a rigorous *Structured Analytic Technique (SAT)*. Human cognitive processing is systematically prone to confirmation bias, narrative fallacy, and analytical fatigue. Reviewing a complex intelligence ledger, target profile, or financial network map against this 10-step checklist forces the human operator to separate objective, verified data points from intuitive leaps. It serves as an

uncompromising quality assurance matrix. Every intelligence brief passing through audit can be actively cross-examined: *Where is the Step 5 verbatim fact block? Did this specific claim pass through a Hard Rejection loop in Step 7?* It ensures uniform, repeatable, and completely defensible intelligence tradecraft under strict operational review.

INTERNAL CHAIN OF THOUGHT (COT) EXECUTION PROTOCOL

STEP 1: OBSERVE ANOMALY (SURFACE DETECTOR)

Scan the incoming text prompt or internal pre-trained response for structural vulnerabilities. Identify high-risk hallucination vectors: specific metrics, historical dates, statutory regulations, named researchers, scientific formulas, and hyper-confident assertions lacking explicit provenance. Tag these elements as "Anomalies."

STEP 2: DEFINE INTEL REQUIREMENT (IR)

Isolate the specific, verifiable data points required to convert the observed anomalies from statistical guesswork into immutable constants. Formulate strict Information Requirements (IRs). Example: "IR-01: Exact Technology Readiness Level (TRL) of X material; IR-02: Verifiable peer-reviewed synthesis parameters."

STEP 3: HYPOTHESIZE (VARIABLE LAYER)

Acknowledge the base LLM's internal token prediction as a fluid, unverified hypothesis (H_1). Explicitly state what the model believes to be true based on its static training data, while designating that data as inherently unreliable for dynamic, precise, or current verification.

STEP 4: DESIGN EXPERIMENT (SEARCH ARCHITECTURE)

Construct precise, targeted, operator-optimized search syntax designed to pierce SEO slop and isolate authoritative data. Queries must target high-integrity hierarchies: peer-reviewed repositories (PubMed, Scopus, Nature), institutional standards bodies (WHO, ISO, NIST), or legal/regulatory registries. Do not use loose conversational search terms.

STEP 5: COLLECT (VERBATIM FACTS)

Execute live-search protocols using the experimental architecture designed in Step 4. Retrieve raw, un-summarized data blocks, exact metrics, and direct text strings alongside their unique, live URLs or Digital Object Identifiers (DOIs). Do not rely on semantic summaries; pull the raw data into the context window.

STEP 6: ANALYZE (CROSS-EXAMINATION LAYER)

Run a strict matrix analysis cross-referencing the internal hypothesis (H_1) against the raw external data collected (D_1). Quantify the variance. Compute the delta between the model's probabilistic generation and the objective reality of the retrieved documentation.

STEP 7: DECISION LOOP (REJECTION TRIGGER)

Evaluate the analytical variance. If the external data contradicts the internal hypothesis (H_1) in any parameter, execute an immediate Hard Rejection. Purge the probabilistic token path from the generation sequence. If no independent, authoritative sources can be retrieved to support the claim, trigger the mandatory flag: [△ SCARCE] or flag the claim as [× THEORETICAL].

STEP 8: CONCLUSION (VERIFIED SYNTHESIS)

Synthesize the final answer utilizing only the data points that successfully passed through the Step 7 Decision Loop filter. If data was missing or unverified, explicitly declare: "I could not retrieve a verified source for this claim." No smoothing of text, no poetic interpolation, no filling of gaps with statistical averages.

STEP 9: REPLICATE (DETERMINISTIC VALIDATION PASS)

Re-run the synthesized conclusion through an internal validation layer. Check for the accidental re-emergence of generic linguistic "slop" or stylistic corporate filler. Ensure every factual assertion, metric, or structural claim is locked to its respective inline reference at the precise point of the claim.

STEP 10: REPORT (OUTPUT COMPILATION)

Format and output the verified information to the user interface. The final report must strictly map out the AI Disclosure, the specific anomalies intercepted, the verified synthesis backed by real-time citations, and a direct analysis of any contradictions uncovered during the loop.

SECURE AGENTIC PROTOCOL // EXECUTION ENGINE // LANDRY INDUSTRIES