JOR BAYESIAN FUSION FRAMEWORK

Technical Manual v3.1 | Decoupled MCMC Edition

Methodology	James Orion
Implementation	Jake James
Calibration Standard	K = 0.20 (AARO 2024 Uncertainty Standards)
Default Prior (NH)	0.20

1. Framework Introduction & Core Philosophy

The JOR-V3.1 (James Orion Report) Bayesian Fusion Framework is a quantitative diagnostic tool designed to evaluate Unidentified Aerial Phenomena (UAP) sightings. It transitions analysis from qualitative anecdotes to probabilistic assessments by fusing witness credibility, environmental data, and physical sensor evidence.

Decoupled Architecture: v3.1 separates human heuristic scoring from mathematical inference. Human analysts provide initial scoring (Phase 1), while a Markov Chain Monte Carlo (MCMC) engine performs high-fidelity stochastic sampling to generate a final result (Phase 2).

2. Phase 1: Interactive Scoring Protocol

Scoring is performed in **jor_fusion.py**. Each case is evaluated across three weighted pillars to produce two distinct probability metrics:

SOP (Solid Object Probability) uses the *raw* Physical Evidence score (P_raw) — before any flight behaviour modifier — and represents a baseline probability that the observed object is a real, solid object.

NHP (Non-Human Probability) uses the *flight-adjusted* Physical Evidence score (P_final) and represents the probability of non-human origin. Both SOP and NHP share the same Credibility and Environmental inputs; only the Physical component differs.

Pillar	Variable	Default Weight
Witness Credibility	C	40%
Environmental Conditions	E	30%
Physical / Sensor Evidence	P	30%

```
SOP = (Weight_C x C) + (Weight_E x E) + (Weight_P x P_raw)
```

```
NHP = (Weight_C x C) + (Weight_E x E) + (Weight_P x P_final)
```

2.1 Witness Credibility (C) — 40% Weight

Measures the reliability, training, and independence of human observers.

Category	Score Range	Criteria
Weak	0.30 – 0.50	Single untrained civilian; anonymous report; no corroboration.
Moderate	0.55 – 0.65	1 trained observer (Pilot/LEO) OR 2–3 civilian witnesses.
Strong	0.70 – 0.80	Multiple trained observers OR multiple independent corroborating civilians.
Very Strong	0.81 – 0.85	Military/Scientific personnel + formal independent documentation.

Credibility Modifiers

Condition	Modifier
Independent written reports or time-stamped sensor logs available	**+0.03**
Witnesses viewed from >2 independent geographic positions (triangulation)	**+0.02**
Notable witness inconsistencies or conflicting reports	**-0.03**
Known high-risk misidentification potential or historically unreliable witness	**-0.05**

Credibility Hard Caps

Hard caps enforce absolute **ceilings** on the Credibility score regardless of any modifiers applied. When multiple caps apply simultaneously, the most restrictive (lowest value) cap governs.

Condition	Max Score	Effect
Anonymous witness	0.45	Score cannot exceed 0.45
Single untrained civilian	0.50	Score cannot exceed 0.50
No trained observer present	0.70	Score cannot exceed 0.70

2.2 Environmental Conditions (E) — 30% Weight

Evaluates observation quality and environmental noise during the sighting.

Category	Score Range	Criteria
Weak	0.30 – 0.45	Fog, heavy cloud, night (no illumination), duration <10s.
Moderate	0.50 – 0.60	Light cloud, partially obstructed, night (some illumination), 10–30s.
Strong	0.65 – 0.85	Broad daylight, clear sky, multiple viewing angles, duration >30s.

Environmental Modifiers

Condition	Modifier
Object observed from multiple vantage points (Air/Land/Sea)	**+0.03**
Local weather conditions officially documented at the timestamp	**+0.02**

Condition	Modifier
Estimated target distance exceeds 1 km (increased geometric uncertainty)	**-0.03**
Total observation duration under 5 seconds	**-0.05**

Environmental Hard Caps

Environmental caps differ from Credibility caps in one critical way: the **Daytime Clear** cap enforces a **floor** (minimum score) rather than a ceiling. This prevents clear-daylight observations from being scored lower than 0.60 due to conservative modifier stacking. All other environmental caps enforce ceilings.

> **The Daytime Clear floor is the only minimum-type cap in the entire framework. Every other hard cap across all three pillars enforces a maximum.**

Condition	Cap Value	Type	Effect
Heavy fog	0.40	**CEILING**	Score cannot exceed 0.40
Nighttime OR single perspective	0.70	**CEILING**	Score cannot exceed 0.70
Daytime clear sky	0.60	**FLOOR**	Score cannot go below 0.60

2.3 Physical / Sensor Evidence (P) — 30% Weight

Serves as the objective anchor, prioritising hard instrumentation and physical interaction.

Category	Score Range	Criteria
Weak	0.30 – 0.45	Anecdotal report only; no photography or trace evidence.
Moderate	0.50 – 0.65	Single sensor type (e.g., cell video ONLY or radar return ONLY).
Strong	0.70 – 0.85	Two or more sensor types (Radar-Visual) or confirmed physical anomalies.
Very Strong	0.86 – 0.95	Multi-sensor fusion + active physical interaction/trace analysis.

Physical Evidence Modifiers

Condition	Modifier
Documented EMP, interference, or electronic shutdown	**+0.05**
Long-duration multi-frame imagery or high-definition video	**+0.03**
Independent laboratory analysis of physical trace evidence	**+0.02**
Poor or ambiguous video quality (low SNR)	**-0.05**
Notably inconsistent sensor readings across different platforms	**-0.07**

Physical Evidence Hard Caps

All Physical Evidence hard caps enforce ceilings only. The multi-sensor maximum of 0.95 is an absolute limit that cannot be exceeded under any circumstance.

Condition	Max Score	Effect
No sensor data present	0.55	Score cannot exceed 0.55

Condition	Max Score	Effect
Only video evidence available	0.75	Score cannot exceed 0.75
Multi-sensor fusion (absolute maximum)	0.95	Score cannot exceed 0.95 under any circumstance

3. Flight Behavior & Kinematic Anomalies

Kinematic anomalies differentiate a Solid Object (SOP) from a Non-Human origin (NHP). The flight modifier is applied **exclusively to the Physical Evidence score (P)**, producing P_final. This ensures anomalous kinematics are moderated through the Physical Evidence weight (30%) rather than directly inflating the overall composite score.

```
P_final = MIN( P_raw + Flight_Mod, 0.95 )
```

Classification	Modifier	Description
Conventional Flight	+0.00	Behaves within expected aerodynamic parameters.
Minor Anomaly	**+0.02**	Slightly unusual maneuvers/speed; likely explainable.
Moderate Anomaly	**+0.04**	Clearly abnormal movement/trajectory; limited conventional explanation.
Major Anomaly	**+0.05**	Impossible maneuvers (hypersonic stop, instantaneous acceleration).

4. Phase 2: The Bayesian Inference Engine

After scoring, **jor_pymc_runner.py** initiates Phase 2. This transitions analysis from human heuristics to stochastic probability using PyMC's NUTS sampler.

4.1 Stochastic Modelling — Beta Distributions

The engine treats every analyst score (C, E, P) as the **mean** of a Beta Distribution with a standard deviation of **0.02**. This uncertainty cloud propagates scoring variance through the full Bayesian update. The prior (Prior_NH) is also treated as a Beta-distributed random variable with the same sigma.

4.2 Likelihood Equations & Calibration

The Calibration Constant **K = 0.20** defines the Skeptic's Buffer, ensuring the Human Hypothesis (H) remains statistically viable unless evidence is overwhelming.

```
P(Evidence | NH) = NHP
```

```
P(Evidence | H) = MIN( 1.0, 1.0 - NHP + (K x SOP) )
```

4.3 Bayesian Posterior Update

The posterior probability of Non-Human origin is computed via Bayes' theorem:

```
Posterior_NH = (Prior_NH x P(E|NH)) / [(Prior_NH x P(E|NH)) + (Prior_H x P(E|H))]
```

4.4 MCMC Sampling Configuration

Parameter	Value	Notes
Sampler	NUTS	No-U-Turn Sampler (PyMC default)
Chains	4	Parallel chains for convergence diagnostics
Draws per chain	1,000	4,000 total posterior samples per case
Tune steps	1,000	Burn-in / warm-up steps

Parameter	Value	Notes
Target accept	0.95	High acceptance rate for stable sampling
Random seed	42	Ensures reproducible results

5. Execution Workflow & Best Practices

Step 1 — Configuration	Verify **PRIOR_NH = 0.20** and **CALIBRATION_K = 0.20** are set in **jor_fusion.py**. These define the Conservative Standard. Edit values directly in the file for reproducible sensitivity runs. Do *not* rely on the interactive runtime tweak if you intend to run jor_pymc_runner.py afterward — runtime tweaks are memory-only and will not be seen by the runner.
Step 2 — Scoring	Execute **python jor_fusion.py**. Follow the prompts to score each pillar, apply modifiers, apply hard caps, and select a flight behaviour classification. Results are logged to **jor_scores.csv**.
Step 3 — MCMC Inference	Execute **python jor_pymc_runner.py**. The runner imports constants directly from jor_fusion.py and passes them to the PyMC engine. Each case receives 4,000 posterior samples (4 chains x 1,000 draws). jor_scores.csv is updated in-place.
Step 4 — Interpretation	Analyse the updated jor_scores.csv. Key output columns are described in the table below.

Output Column Reference

Column	Description
Posterior_Mean	Refined posterior probability of Non-Human origin after MCMC sampling.
CI_2.5%	Lower bound of the 95% credible interval.
CI_97.5%	Upper bound of the 95% credible interval.
SOP	Solid Object Probability — uses P_raw (no flight modifier applied).
NHP	Non-Human Probability — uses P_final (flight modifier applied to P).
Posterior_NH	Point-estimate posterior from the Phase 1 deterministic calculation.

Narrow credible interval bands *indicate high-confidence data with low scoring uncertainty.* ***Wide bands*** *suggest sensitivity to analyst input variance and warrant additional evidence before drawing conclusions.*