PLANETARY IMAGE ANALYSIS | RESEARCH SOFTWARE

Cassini Jupiter Lightning Detector

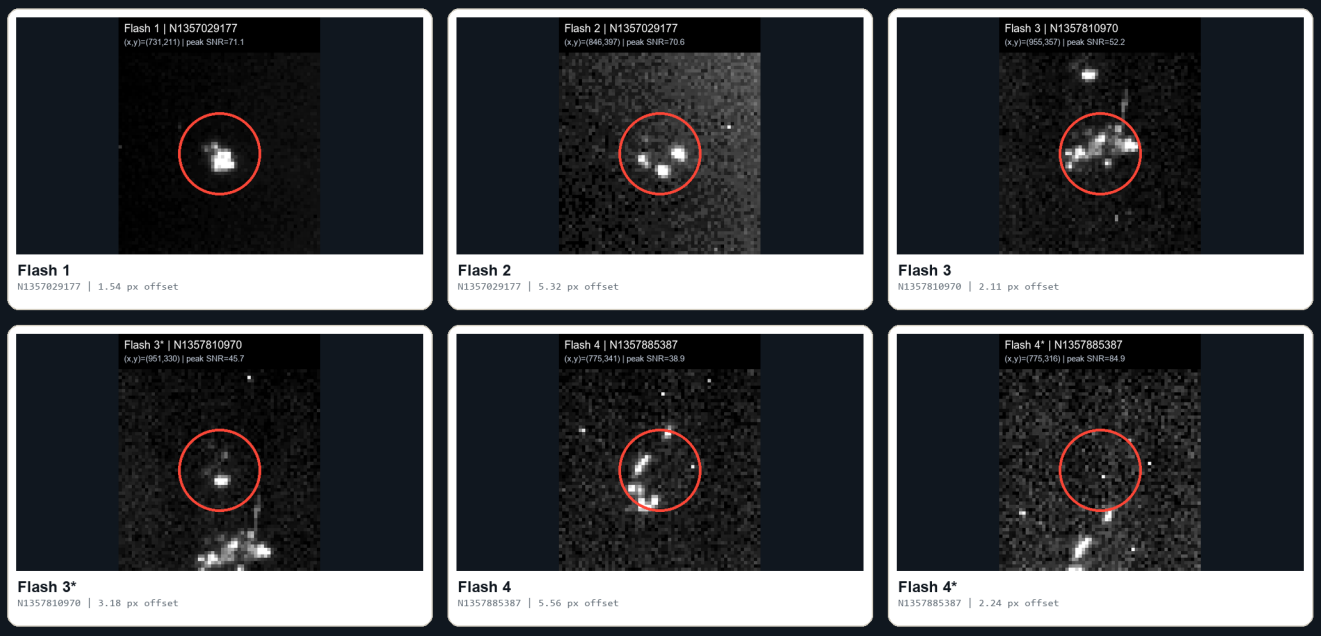
An explainable search-and-review pipeline for faint night-side events

ROLE Research / data analysis contributor

STACK Python, NumPy, Pillow, ISIS/SPICE

PUBLISHED LIGHTNING, REPRODUCED

Six Cassini ISS marks recovered by an independent classical-CV search



Source: calibrated Cassini ISS imagery via NASA PDS OPUS; reference marks: Dyudina et al. (2004).

221	196,233	106	6 / 6	42
Cassini frames	bright regions	review rows	published marks	surface points

Live evidence, source code, and provenance are available below.

OPEN LIVE REVIEWER | VIEW SOURCE CODE | READ PUBLISHED STUDY

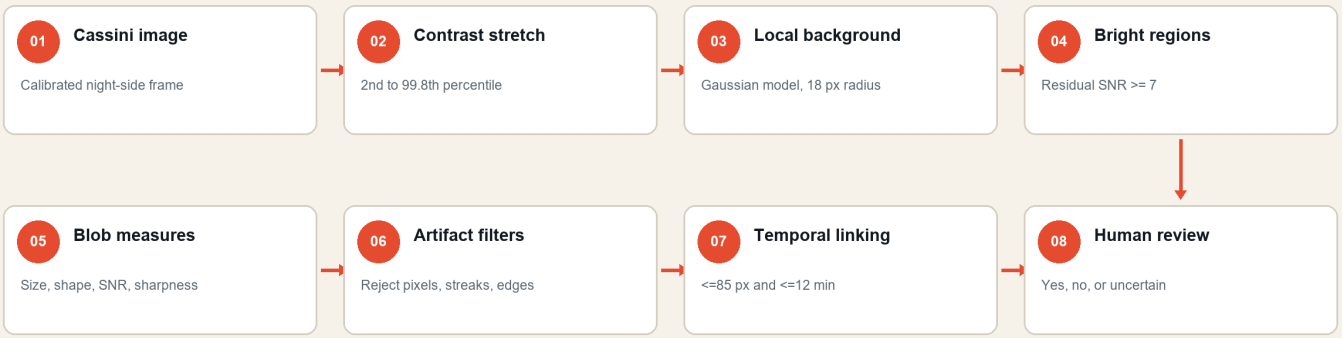
01 | DETECTOR DESIGN

Turn a noisy archive into an inspectable queue

I built a deterministic classical-computer-vision baseline so every candidate can be traced back to a measured rule. The system estimates local background, computes robust signal-to-noise, forms connected blobs, records shape features, applies artifact flags, and links nearby detections through time.

EXPLAINABLE DETECTION PIPELINE

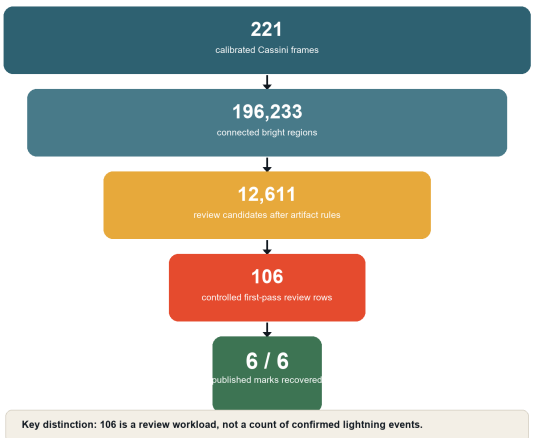
Every transition produces an inspectable intermediate result.



Current method: custom classical computer vision in Python / NumPy / Pillow. No deep learning or reinforcement learning.

FROM SEARCH SPACE TO REVIEW QUEUE

The pipeline narrows evidence without calling every bright spot lightning.



Locked baseline parameters

- Background blur radius: 18 px
- Detection SNR threshold: 7
- Review SNR threshold: 8
- Minimum review area: 3 px
- Maximum component area: 500 px
- Track displacement: <=85 px
- Maximum temporal gap: 12 min

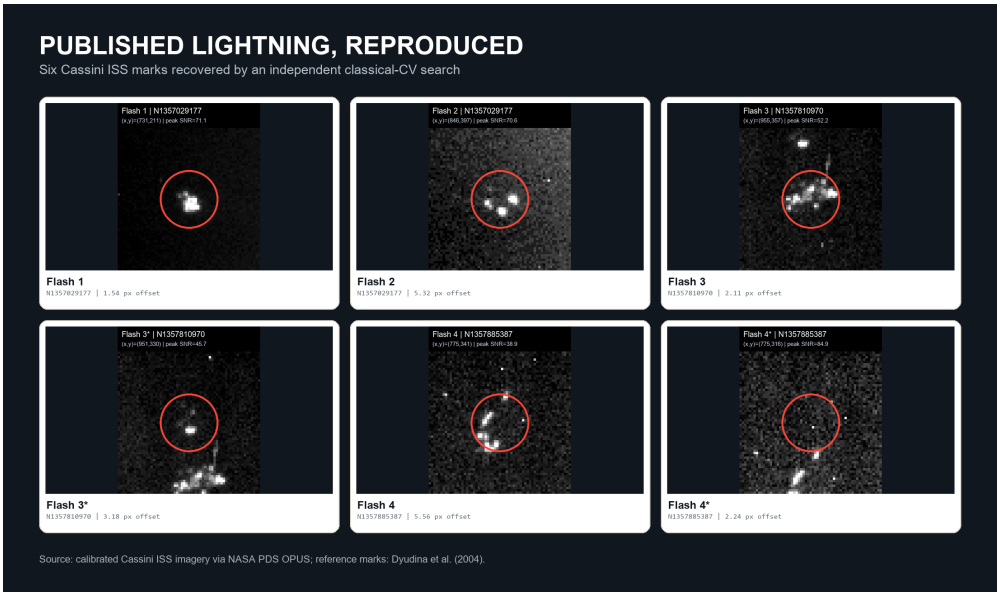
WHY THIS MATTERS

A reviewer can challenge any threshold or feature. The baseline is deliberately understandable before a learned model is introduced.

02 | EVIDENCE

Recover known events before searching for new ones

The detector was not given the published coordinates during candidate generation. I compared its output afterward against six marks reported in Dyudina et al. (2004). All six had a detector candidate within the fixed 8-pixel validation radius.



VALIDATION + JUPITER LOCALIZATION

Pixel recovery is separated from geometric interpretation.

6 / 6	1.54-5.56 px	42 / 106	0
published marks recovered	detector-to-reference offset	surface intersections computed	projection software failures

WHERE THE GEOMETRY STAGE FITS



Candidate x/y	Cassini camera model	Jupiter lat/lon	Surface grouping
Detector coordinates in the 1024 x 1024 image	USGS ISIS: ciss2isis + spiceinit + campt	Planetographic latitude + positive-west longitude	Nearby coordinates grouped within a review tolerance

64 rays had no modeled surface intersection; that is a geometric result, not a software failure.

Validation offsets

Flash 1	1.54 px
Flash 2	5.32 px
Flash 3	2.11 px
Flash 3*	3.18 px
Flash 4	5.56 px
Flash 4*	2.24 px

SAFE CLAIM

Published pixel marks were recovered. Unmatched candidates remain unconfirmed until independent review and temporal/geometric checks.

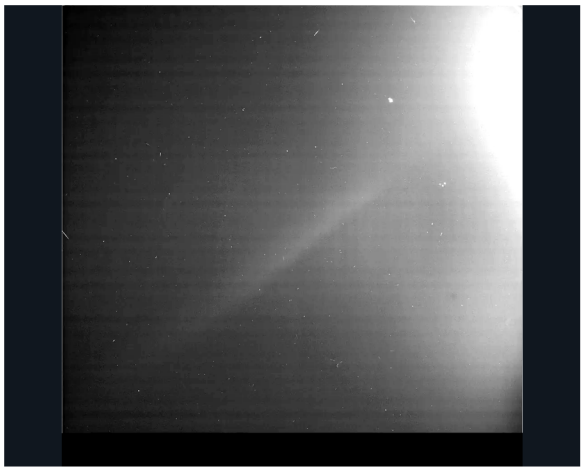
03 | RESEARCH PRODUCT

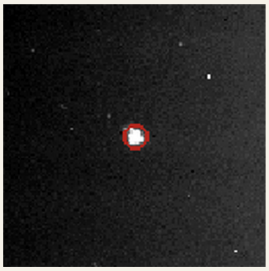
Make candidate decisions reproducible

The public interface exposes the prioritized queue without requiring reviewers to run the detector. Each row retains image ID, timestamp, source link, x/y, SNR, blob size, artifact flags, geometry, and the reviewer decision. Image stretch controls preserve the surrounding context needed to distinguish lightning-like blobs from cosmic rays and compression structure.

HUMAN-IN-THE-LOOP REVIEW

The detector prioritizes. A reviewer makes the scientific decision.





Candidate

1357029177-0148

x/y 731.37, 212.49
SNR 27.55
area 75 px
group G0001
lat/lon 26.536, 162.634 W

Reviewer label

Known lightning | Possible | Artifact | Uncertain

Saved with reviewer, timestamp, notes, source image, and geometry.

1. **Inspect context** — Stretch the full frame and crop.
2. **Record a label** — Positive, negative, or uncertain.
3. **Preserve provenance** — Append the decision to a shared log.
4. **Train deliberately** — Promote a model only after label QA.

Current boundary: labels are collected, but automatic retraining is intentionally disabled.

Engineering boundaries

- Detector runs locally; the public site is review-only.
- Human labels are stored with provenance and do not silently alter thresholds.
- Automatic retraining is disabled until enough labels pass quality review.
- A learned model should be compared against the frozen classical baseline.

TRY THE PUBLIC REVIEWER

04 | TECHNICAL OWNERSHIP

What I built, tested, and documented

This page is the evidence index for an engineering reviewer. Each link goes directly to a working interface, implementation module, validation artifact, or primary data source.

PRODUCT

Live review interface

Candidate inspection, stretch controls, labels, and geometry.

GitHub repository

Source, documentation, tests, and reproducible commands.

IMPLEMENTATION

Detector source

Background model, SNR thresholding, blobs, and temporal links.

Geometry source

ISIS/SPICE candidate projection and lat/lon grouping.

Automated tests

Regression checks for processing, labels, public assets, and geometry.

Evidence + primary sources

Measured run manifest

Geometry validation report

Published lightning study

NASA PDS OPUS archive

USGS ISIS Cassini import docs

Recommended links on the main portfolio page

Title / hero image -> this case-study PDF
Outcome / CTA -> live reviewer
Tools / implementation -> GitHub repository
Published validation -> study abstract
Cassini source -> OPUS archive
Geometry -> validation report