GrantPanel

SAMPLE REPORT · NASA

GrantPanel NASA ROSES Review

with Recommended Revisions

Merit review in the NASA ROSES style: evaluated on Intrinsic Merit, Relevance to NASA's objectives, and Cost reasonableness, with adjectival ratings and major/minor findings, three panelists, three rounds of discussion, and a panel rating taken as the median, all generated by GrantPanel's AI review panel.

grantpanel.ai

Note: This sample document was produced by GrantPanel, a multi-agent AI simulation of the NASA ROSES peer review review process, on a fictitious proposal. It is **not** an official NASA review and does not represent the views of the National Aeronautics and Space Administration. GrantPanel is independent preparation software and is not affiliated with or endorsed by any funding agency. Use reports like this for self-assessment and proposal stress-testing only.

PROPOSAL TITLE

Machine-Learning Retrieval of Exoplanet Atmospheres from Transit Spectroscopy

PRINCIPAL INVESTIGATOR

Dr. Naomi Okoro

INSTITUTION

Summit Astronomical Institute

PROGRAM ELEMENT

ROSES · Astrophysics Data Analysis Program (ADAP)

REQUESTED AMOUNT

$465,000 over 3 yr

RELEVANCE

Archival JWST/HST transit spectra

PANEL CONSENSUS

Very Good

PANEL RATING · MEDIAN

The panel rating (Very Good) is the median of Excellent, Very Good, and Excellent/Very Good.

Panel consensus. The proposal develops a machine-learning (amortized-inference) method to retrieve exoplanet atmospheric parameters from archival JWST and HST transit spectra, delivering calibrated uncertainties far faster than traditional sampling and releasing an open-source pipeline. The panel evaluated it on Intrinsic Merit, Relevance to NASA, and Cost, and set the panel rating as the median of the individual ratings.

FIX IN THIS ORDER

four of the panel's findings — the full list is below

01 **NASA archive unnamed.** →

Specify which archive (e.g., MAST) will host the retrieval products in the Data Management Plan.

02 **Software sustainability missing.** →

Add a licensing and maintenance statement for the open-source deliverable.

03 **Co-I effort too low.** →

Raise the Co-I's 0.5-month effort to match the pipeline-development role.

04 **GPU request unjustified.** →

Justify the year-2 GPU compute request in the budget narrative.

This is the short list. The panel filed three independent reviews, argued across three rounds, and produced **five groups of recommended revisions** — including items not shown here. Each finding is tied to the section it came from, so you can check the claim against your own document.

Contents

DOCUMENT INTELLIGENCE

Document Intelligence

Proposal package inventory

READY WITH QUESTIONS

Detected documents:

- Scientific/Technical/Management section (pages 1–15)
- References and Citations (pages 16–18)
- Budget Narrative, Table of Personnel and Work Effort, and CVs are present in the supporting package.

Conditional requirements:

- Open Science and Data Management Plan: present, naming the archive for the retrieval pipeline and derived spectra.
- Dual-anonymous compliance: unclear. Section 3 cites prior work in the first person, which may identify the team under ROSES dual-anonymous review.
- Foreign participation: not applicable. No non-U.S. institutions appear in the personnel table.

Document quality: No document quality issues detected.

Author questions: Will the self-identifying references in Section 3 be rewritten before submission? ADAP is reviewed dual-anonymously.

COMPLIANCE & RELEVANCE TO NASA

Compliance

ROSES / Guidebook format check

COMPLETED

Status: Compliant with concerns

Issues found:

- The Data Management Plan does not specify which NASA archive (e.g., MAST) will host the derived atmospheric-retrieval products, as ROSES requires.
- The Scientific/Technical/Management section runs to the page limit but the open-source software deliverable lacks a licensing and maintenance statement.
- Table of Personnel and Work Effort shows the Co-I at 0.5 month — too low for the stated pipeline-development role.
- Budget narrative omits justification for the year-2 GPU compute request.

Notes: Required elements (References, Data Management Plan, budget) are present; the proposal correctly uses archival data with no new-observation request.

Relevance to NASA

Fit to the program element's objectives

COMPLETED

Relevance: High

Addressed well:

- Directly serves the ADAP objective of maximizing science return from existing NASA mission archives (JWST, HST) rather than requesting new observations.
- The retrieval products would benefit the broader exoplanet community and future mission planning — a stated NASA priority.
- Open-source pipeline deliverable aligns with NASA's open-science policy.

Notes on relevance vs. merit:

- The panel kept impact under Intrinsic Merit, not Relevance, per ROSES guidance; relevance here concerns fit to the archive-exploitation objective, which is strong.
- One panelist noted the proposal could state more explicitly how products feed forward to a specific upcoming mission's science.

INDIVIDUAL PANEL REVIEWS

Panelist A

NASA peer reviewer

COMPLETED

RATING — Excellent — INTRINSIC MERIT

Intrinsic merit.

- **Major strength:** the amortized-inference approach retrieves atmospheric parameters orders of magnitude faster than nested sampling while reporting calibrated uncertainties — a real methodological advance.
- **Major strength:** validation against established retrieval codes on synthetic spectra is planned and appropriate.
- **Minor weakness:** the treatment of model-misspecification (unmodeled clouds/hazes) is thin.

Relevance to NASA. High — squarely serves archive exploitation.

Cost. Reasonable; resources and skill mix match the scope.

Overall. Excellent — a strong, well-scoped archival analysis proposal.

Panelist B

NASA peer reviewer

COMPLETED

RATING — Very Good — INTRINSIC MERIT

Intrinsic merit.

- **Major weakness:** the training set is drawn from one forward model; without cross-model training the retrieval may inherit that model's biases — the validation plan does not test this.
- **Minor strength:** the uncertainty-calibration diagnostics are well chosen.
- **Minor weakness:** the exoplanet sample selection criteria are under-specified.

Relevance to NASA. High.

Cost. Reasonable, though the Co-I effort looks too low for the pipeline work.

Overall. Very Good — strong idea, with a methodological risk the validation plan should close.

Panelist C

NASA peer reviewer

COMPLETED

RATING — Excellent/Very Good — INTRINSIC MERIT

Intrinsic merit.

- **Major strength:** the PI's prior open-source retrieval tool is widely used, lending strong feasibility and community-benefit credibility.
- **Minor weakness:** the software sustainability plan (licensing, maintenance beyond the grant) is missing.

Relevance to NASA. High — directly advances open science and archive value.

Cost. Reasonable; year-2 GPU request needs justification.

Overall. Excellent/Very Good — high feasibility and community value; tighten sustainability and sample selection.

PANEL DISCUSSION

ROUND 1 — POSITIONS

A

REVIEWER A

Excellent for me. The amortized-inference speed-up with calibrated uncertainties is a genuine advance, and it's well matched to the archive-exploitation objective.

B

REVIEWER B

Very Good. My major weakness is the single forward model in training — the retrieval could inherit that model's biases and the validation plan doesn't test cross-model transfer.

C

REVIEWER C

Excellent/Very Good. The PI's existing open-source tool makes feasibility and community benefit strong; sustainability and sample selection are the soft spots.

ROUND 2 — DISCUSSION

A

REVIEWER A

Reviewer B, that's a fair major weakness — but it's addressable by adding a second forward model to the training set, which is within scope.

B

REVIEWER B

Agreed it's addressable. If cross-model validation were promised I'd move to Excellent. As written, it's a real methodological risk.

C

REVIEWER C

Both of those — cross-model training and a sustainability statement — are the kind of minor-to-moderate revisions that don't change the fundamentally strong merit.

ROUND 3 — FINAL POSITIONS

A

REVIEWER A

Final: Excellent. Major strengths outweigh a fixable misspecification concern.

B

REVIEWER B

Final: Very Good. The single-forward-model risk keeps me just below Excellent.

C

REVIEWER C

Final: Excellent/Very Good. High feasibility and community value; minor revisions only.

PANEL RATING

Panel Summary

Very Good ●●●●○ PANEL RATING · MEDIAN

Panel consensus. The proposal develops a machine-learning (amortized-inference) method to retrieve exoplanet atmospheric parameters from archival JWST and HST transit spectra, delivering calibrated uncertainties far faster than traditional sampling and releasing an open-source pipeline. The panel evaluated it on Intrinsic Merit, Relevance to NASA, and Cost, and set the panel rating as the median of the individual ratings.

Intrinsic merit — consensus

- **Major strengths:** a real methodological advance in retrieval speed with calibrated uncertainties; strong feasibility anchored by the PI's widely used open-source tool.
- **Major weakness:** training on a single forward model risks inheriting that model's biases; the validation plan does not test cross-model transfer.
- **Minor weaknesses:** thin treatment of model misspecification (clouds/hazes); under-specified exoplanet sample-selection criteria.

Relevance to NASA — consensus

- High. Directly serves the ADAP objective of maximizing science return from existing NASA archives and advances NASA's open-science policy. Per ROSES guidance, impact was scored under Intrinsic Merit, not Relevance.

Cost — consensus

- Reasonable: proposed resources and skill mix match the scope. Two items need attention — the Co-I effort (0.5 month) is too low for the pipeline role, and the year-2 GPU request is unjustified.

Points of disagreement

- Whether the single-forward-model training risk is Excellent-limiting: Panelist B held it as a major weakness; A and C viewed it as an addressable revision. The panel rating (Very Good) is the median of Excellent, Very Good, and Excellent/Very Good.

Recommendation. A strong, highly relevant archival-analysis proposal. To reach the top of the range, a revision should add cross-model training and validation, specify sample-selection criteria, provide a software sustainability plan, correct the Co-I effort, name the NASA archive for data products, and justify the GPU request.

RECOMMENDED REVISIONS

Every weakness and gap the panel raised, paired with a specific, actionable revision, and grouped by review stage.
GrantPanel recommends changes and shows you where; it never edits your document.

Revisions by section

Ordered by what to fix first

5 REVISIONS

01 Compliance & format

FIX BEFORE SUBMISSION

- → **NASA archive unnamed.** Specify which archive (e.g., MAST) will host the retrieval products in the Data Management Plan.
- → **Software sustainability missing.** Add a licensing and maintenance statement for the open-source deliverable.
- → **Co-I effort too low.** Raise the Co-I's 0.5-month effort to match the pipeline-development role.
- → **GPU request unjustified.** Justify the year-2 GPU compute request in the budget narrative.

02 Relevance to NASA

STRENGTHEN RELEVANCE

- → **Feed-forward to a mission implicit.** State explicitly how the retrieval products feed a specific upcoming mission's science.

03 Individual reviews — intrinsic merit

STRENGTHEN THE SCIENCE

- → **Single forward model in training.** Train and validate across more than one forward model to avoid inheriting its biases.
- → **Model misspecification thin.** Deepen the treatment of unmodeled clouds and hazes.
- → **Sample selection under-specified.** Specify the exoplanet sample-selection criteria.

04 Panel — priorities

ADDRESS FIRST

- → **Highest priority.** Cross-model training and validation — the risk keeping the rating below Excellent.
- → **Then.** A software sustainability plan, the Co-I effort correction, the named archive, and the GPU justification.

05 Prior work & citations

POSITION THE CONTRIBUTION

- → **Two directly relevant works uncited.** Cite recent cross-model SBI robustness work (2024) and a 2024 JWST cloud-signature study.

SUPPORTING RESEARCH

REFERENCES LOOKUP

References

Extracts and looks up cited prior work

COMPLETED

Extracting cited references from the proposal...

Found 6 references. Fetching abstracts from Semantic Scholar...

[1/6] Atmospheric retrieval of exoplanet transmission spectra — found [2/6] Simulation-based inference for astrophysics — found [3/6] JWST transmission spectroscopy of a warm Neptune — found [4/6] Clouds and hazes in exoplanet atmospheres — found [5/6] Open-source retrieval frameworks — no match [6/6] Calibration of posterior estimators in SBI — found

The proposal cites 6 references. Abstracts (where available) are provided below.

[1] Atmospheric retrieval of exoplanet transmission spectra *(2018)* — Madhusudhan *Annual Review of Astronomy & Astrophysics · 2018* A comprehensive review of retrieval methodology, establishing the Bayesian nested-sampling baseline the proposal aims to accelerate.

[2] Simulation-based inference for astrophysics *(2020)* — Cranmer et al. *PNAS · 2020* Introduces amortized simulation-based inference and its calibration diagnostics, the methodological foundation of the proposal.

[3] JWST transmission spectroscopy of a warm Neptune *(2023)* — JWST Transiting Exoplanet Team *Nature · 2023* Reports high-precision JWST transmission spectra revealing molecular features, the data class the proposal will analyze.

[4] Clouds and hazes in exoplanet atmospheres *(2019)* — Gao et al. *Nature Astronomy · 2019* Reviews aerosol effects that flatten spectral features, the model-misspecification challenge reviewers flagged.

[5] Open-source retrieval frameworks *(Not found on Semantic Scholar.)*

[6] Calibration of posterior estimators in SBI *(2021)* — Hermans et al. *NeurIPS · 2021* Shows that amortized posterior estimators can be miscalibrated and provides tests to detect it — directly relevant to the proposal's uncertainty claims.

LITERATURE REVIEW (DEEP RESEARCH)

Literature Review

Deep-research SOTA scan

COMPLETED

Research area & central claim. The proposal claims amortized simulation-based inference can retrieve exoplanet atmospheric parameters from archival transit spectra with calibrated uncertainties and far lower cost than nested sampling. The panel's deep search assessed the speed/calibration claim and the generalization claim.

State of the art (last 3–5 years)

- Madhusudhan (2018) codified nested-sampling retrieval — the baseline being accelerated.
- Cranmer et al. (2020) established amortized SBI and calibration diagnostics for the physical sciences.
- JWST-era spectra (2023–2024) sharply increased data quality and volume, motivating faster retrieval.
- Hermans et al. (2021) documented miscalibration risks in amortized estimators — the exact concern reviewers raised.

Competing approaches

- **Traditional nested sampling** — gold-standard uncertainties but slow; the proposal must match its calibration to justify the speed-up.
- **Neural posterior estimation on single forward models** — fast but bias-prone; the proposal's differentiator should be cross-model robustness.

Open problems the proposal addresses

- Whether amortized retrieval remains calibrated under model misspecification (clouds/hazes).
- Whether retrievals transfer across forward models — the reviewers' central methodological concern.

Citation gaps in the proposal

- Recent cross-model SBI robustness work (2024) — directly relevant to the transfer concern, not cited.
- A 2024 JWST cloud-signature study — relevant to the misspecification treatment, not cited.

Search log

- "amortized simulation based inference exoplanet retrieval 2024"
- "neural posterior estimation transmission spectrum calibration"
- "forward model bias atmospheric retrieval transfer"
- "JWST transmission spectroscopy clouds hazes 2024"
- "posterior calibration diagnostics SBI astrophysics"
- "ADAP archival exoplanet atmosphere analysis"

INDIVIDUAL WEB SEARCHES

Panelist A
3 searches

1. amortized inference exoplanet retrieval speed 2024
2. calibrated uncertainty transmission spectrum retrieval
3. model misspecification clouds hazes retrieval

Panelist B
4 searches

1. forward model bias neural retrieval transfer
2. cross model training simulation based inference
3. exoplanet sample selection transit spectroscopy
4. posterior calibration amortized estimator test

Panelist C
3 searches

1. open source exoplanet retrieval tool adoption
2. software sustainability NASA open science policy
3. ADAP archival JWST HST analysis objective