

A Minimal Study of Stub Reviewers

Authors: Alice Researcher, Bob Scientist, Carol Tester

Abstract

This paper presents a synthetic study used to validate the testing fixtures of an automated multi-agent peer review system. We describe a fictional methodology, report fabricated results, and conclude with self-referential observations to enable end-to-end pipeline tests without depending on third-party publications.

1. Introduction

Automated peer review systems need deterministic fixtures. This paper exists only as a fixture for pdfplumber. It contains the usual section headings so the section identifier has something to find.

2. Methods

We invented an experiment with $N=3$ fictional reviewers and counted imaginary citations. No data were harmed in the production of this fixture.

3. Results

All three stub reviewers reached consensus that the paper exists. $p=0.04$. We report a Cohen's d of 0.6 and note that the confidence interval crosses zero, which we will address in section 4.

4. Discussion

Limitations include the fact that this paper is entirely fictional. Implications: the automated review pipeline will produce deterministic outputs against this fixture.

5. Conclusion

Stub reviewers approved the paper with minor revisions. The fixture is valid. Future work consists of writing tests that exercise this paper through the full pipeline.

References

- [1] Stub, A. (2026). On Stub-Driven Development. *Journal of Tests*, 1(1), 1-9.
- [2] Mock, B. (2026). Fixtures as a Design Discipline. *Proc. of CI/CD*, 42-47.