Why–Repair: Mathematically Constrained Search for Local Proof Repair

Abstract

Repairing a mathematical proof requires more than correcting an isolated statement: the replacement must follow from admissible premises and remain sufficient for subsequent arguments. We propose a mathematically constrained search extension to Why–Repair, a dependency-guided system for auditing and locally repairing natural-language proofs. Each repair region is governed by a two-sided contract that specifies admissible upstream premises and the conclusions required by retained downstream steps. With the original theorem and assumptions fixed, a budgeted controller explores alternative patches, prunes candidates using validated counterexamples within their applicable scope, and expands the editable region when local candidates are exhausted. Verification feedback thus constrains subsequent candidates and guides changes to repair scope, rather than merely accepting or rejecting individual patches. Versioned dependencies track unresolved assumptions and invalidate affected evidence for revalidation before acceptance. Historical evaluation of the underlying auditor on a frozen engineering set of 50 algebraic proofs yields 94% proof-validity classification accuracy. Separately, an exact-arithmetic prototype constructs a Lyapunov certificate in a controlled example, demonstrating counterexample-guided pruning and repair-region expansion. Earlier repair pilots using historical diagnoses include human-confirmed successes and an unresolved case in which revalidation detected a remaining downstream error. These preliminary results support an executable formulation of proof repair in which mathematical constraints guide both candidate selection and the scope of revision.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.