

# STEM BIO-AI Evidence-Surface Scan v1.7.5

Repository: yorkeccak/bio | Commit: 100a0bf7497e | Branch: main | Audit Date: 2026-05-15 | Mode: LOCAL\_ANALYSIS | Policy: default (mirror\_only)

# 48 / 100

Final Score

## T1 Quarantine

### Use Scope:

Exploratory review only; no patient-adjacent use.

Weighted model: Stage 1 x 0.40 + Stage 2R x 0.20 + Stage 3 x 0.40 - Risk Penalty = 48

Stage 1  
README Evidence

# 75 / 100

Stage 2R  
Repo-Local Consistency

# 40 / 100

Stage 3  
Code / Bio Responsibility

# 25 / 100

Stage 4  
Replication Evidence

# 30 / 100

## Code Integrity

PASS

### C1 Credentials

No direct credential patterns detected by local CLI scan.

WARN

### C2 Dependency Pinning

External operational dependency signal surfaced in code-integrity lane.

PASS

### C3 Deprecated Paths

No deprecated patient-adjacent metadata patterns detected.

WARN

### C4 Exception Handling

Unsupported legal/compliance claim surfaced in boundary-integrity lane.

## Remediation Targets

- Clinical-adjacent surfaces exist without an explicit non-diagnostic/non-clinical boundary.
- Self-asserted compliance or privacy-governance claim requires independent verification.
- Legal, privacy, or compliance claim appears without supporting governance or security-grounding evidence in re
- Core workflow appears materially dependent on named external service providers; local or self-host claims may

## Positive Evidence

- Package metadata was available for repo-local consistency checks.

## Bio Deterministic Diagnostics

### SMILES Surface Integrity

not\_detected=1

### SMILES RDKit Validation

not\_applicable=1

### SMILES Parser Guard

not\_detected=1

### Silent Mock Fallback

not\_detected=1

### Traceability Manifest Surface

not\_detected=1

### Bio Subprocess Run Trace

not\_detected=1

## Regulatory basis note

Aligned to current official source classes as of May 2026: EU AI Act (Regulation (EU) 2024/1689), FDA QMSR, FDA AI-enabled device guidance themes, and IMDRF SaMD/GMLP frameworks.

This is a traceability aid, not a compliance or clearance determination.

**Traceability summary:** Structural signals partially align with traceability scaffolding. This remains a pre-audit traceability aid, not a compliance determination.

## Stage 1 — README Evidence Signal | Weight: 0.40

75 / 100

S1 Score

Check	Points	Evidence / Finding
Baseline	+60	Non-nascent README evidence baseline.
BIO/medical terms in README	+10	README exposes bio/medical domain vocabulary.
BIO/medical terms in package	+5	Package metadata exposes bio/medical domain vocabulary.
R2: Regulatory Framework	+5	Self-asserted privacy/compliance language detected without stronger regulatory-framework evidence.
R3: Clinical Boundary	-5	CA-INDIRECT surface lacks explicit non-clinical or non-diagnostic boundary.

Calculation: 60 plus Stage 1 evidence additions/deductions = 75

• Clinical-Adjacent: **YES** (CA-INDIRECT) • Explicit Disclaimer: **ABSENT** • T0 Hard Floor: **Clear**

## Stage 2R — Repo-Local Consistency | Weight: 0.20

40 / 100

S2R Score

Check	Points	Evidence / Finding
Baseline	+60	Non-nascent local repository baseline. — Every repository that is not nascent starts at 60. This baseline accounts for basic structural maturity.
R2R-1: README / Package Alignment	+15	README has domain overlap with package metadata or entry points. — README and package metadata share bio-domain vocabulary, indicating claim-to-implementation alignment.
R2R-D2: Missing Clinical Boundary (PENALTY)	-20	Clinical-adjacent surfaces exist without an explicit non-diagnostic/non-clinical boundary. — Clinical-adjacent repository lacks an explicit 'research use only' or 'not for...

60 plus local consistency additions/deductions = 40 **Local Contradiction / Insufficient Consistency**

## Stage 3 — Code & Bio Responsibility | Weight: 0.40

25 / 100

S3 Score

### Engineering Accountability (T-series)

Check	Points	Evidence / Finding
T1: CI/CD Workflow	0 / 15	No workflow files detected. — CI/CD workflows (GitHub Actions, GitLab CI, CircleCI) verify that commits do not silently break the pipeline. Full credit (15) requires...
T2: Domain-Specific Tests	0 / 15	No tests detected. — Domain-specific tests verify biological outputs — e.g., sequencing pipeline correctness, variant call validation, or genomic data integrity. Full credit...
T3: Changelog & Release Hygiene	0 / 15	No changelog detected. — A CHANGELOG tracks which version fixed which defect — essential for regulatory traceability and reproducibility audits. CHANGELOG.md, CHANGELOG, or...

### Biological Integrity (B-series)

Check	Points	Evidence / Finding
B1: Data Provenance Controls	15 / 15	Dependency manifest detected with data source, IRB, or dataset citation language. — Dependency manifests (requirements.txt, pyproject.toml, environment.yml) establish...
B2: Bias / Limitations Documentation	0 / 15 [Manual review required]	No bias/limitations language detected by local CLI scan. — Documentation of algorithmic bias, limitations, or model boundary conditions. Score 8 for boundary language; max...
B3: COI & Funding Disclosure	5 / 5	COI, funding, sponsor, or acknowledgement language detected. — Conflict of interest and funding disclosure in README or FUNDING.md. Required for institutional review context...

## Stage 3 Gap Analysis — Path to Next Tier

- **T-series (engineering) attained:** 0 / 45    **B-series (bio integrity) attained:** 20 / 35
- **Local CLI scan maximum:** 55 / 100 (T1+T2+T3 max 15 each; B1 max 10; B2/B3 require manual review)
- **Gap to T3 (final score >= 70):** 22 points needed across all stages
- **Gap to T4 (final score >= 85):** 37 points needed across all stages
- **B2 Bias/Limitations:** Not detectable — requires manual audit of README, model card, or supplementary documentation for validation boundaries and algorithmic limitations
- **B3 COI/Funding:** Not detectable — requires inspection of README or FUNDING.md for conflict of interest and funding source disclosure

## Code Integrity — Deep Analysis

Check	Points	Evidence / Finding
C1: Hardcoded Credentials	PASS	No direct credential patterns detected by local CLI scan.   Scan: Scans for AWS access keys (AKIA*), OpenAI keys (sk-*), GitHub tokens (ghp_*), and api_key = '...' patterns...
C2: Dependency Pinning	WARN	External operational dependency signal surfaced in code-integrity lane.   Scan: Checks whether requirements.txt / pyproject.toml / environment.yml use exact version pins...
C3: Deprecated Patient Paths	PASS	No deprecated patient-adjacent metadata patterns detected.   Scan: Scans deprecated/directories for patient metadata patterns: patient_id, patient_age, patient_sex...
C4: Fail-Open Exceptions	WARN	Unsupported legal/compliance claim surfaced in boundary-integrity lane.   Scan: Detects fail-open exception patterns: 'except Exception: pass' or 'except: return True' in...

## Remediation Guidance

### [WARN] C2: Dependency Pinning:

→ Pin all dependencies to exact versions (== for pip, hash-pinning for conda). Run pip-audit or safety regularly. Consider pip-compile for reproducible lock files. Unpinned ranges in clinical-adjacent pipelines create silent regression risk.

### [WARN] C4: Fail-Open Exceptions:

→ Replace broad 'except Exception: pass' or 'except: return True' patterns with specific error types and explicit failure logging. In clinical-adjacent code paths, any silent failure is a patient safety risk. Fail closed, not open.

## Classification & Repository Analysis

Check	Points	Evidence / Finding
Clinical Adjacent	YES	Severity: CA-INDIRECT. Triggered by BIO/CLINICAL_OUTPUT term regex match across README, docs, and code.
T0 Hard Floor	Clear	No T0_HARD_FLOOR condition detected.
Explicit Disclaimer	ABSENT	Disclaimer pattern not found in README or docs. High impact on Stage 1 and Stage 2R scores.
Files Scanned	166	Total files indexed by recursive walk. Text files only for content analysis; binary files counted but not read.
Execution Mode	LOCAL_ANALYSIS	No LLM calls. No network access. No runtime execution. Deterministic regex + file-system scan only.

## File Integrity (SHA-256)

File	SHA-256 Hash
README.md	199862D708D85AF0B126FD4129E5F134D6E9E804F6F8249F940F3DA16DC190AA

## Priority Improvement Roadmap

### Priority 1: Clinical-adjacent surfaces exist without an explicit non-diagnostic/non-clinical boundary.

→ Add a prominent 'Research Use Only — Not for Clinical or Diagnostic Use' disclaimer to README H1 or H2 section. Reference applicable frameworks: FDA SaMD guidance, EU AI Act Article 6, or IRB oversight requirements for your deployment context.

### Priority 2: Self-asserted compliance or privacy-governance claim requires independent verification.

→ Review this finding and implement appropriate controls before supervised or clinical-adjacent deployment.

### Priority 3: Legal, privacy, or compliance claim appears without supporting governance or security-grounding evidence in reviewed repository sources.

→ Review this finding and implement appropriate controls before supervised or clinical-adjacent deployment.

### Priority 4: Core workflow appears materially dependent on named external service providers; local or self-host claims may overstate operational independence.

→ Review this finding and implement appropriate controls before supervised or clinical-adjacent deployment.

### Priority 5: C2\_dependency\_pinning: WARN

→ Pin all production dependencies to exact versions (== for pip). Add pip-audit or safety to CI pipeline for vulnerability scanning. Consider pip-compile for deterministic lock files.

### Priority 6: C4\_exception\_handling\_clinical\_adjacent\_paths: WARN

→ Replace broad exception handlers with specific error types and explicit logging. In any clinical-adjacent code path: fail closed, not open. Never silently return True or pass on exception.

## Positive Evidence Summary

- Package metadata was available for repo-local consistency checks.

## Method Boundary

Deterministic local CLI scan. No LLM, network, or runtime test execution is required.

**Scope boundary:** Runtime behavior, model output correctness, dynamic validation, wet-lab reproducibility, and clinical validation are outside the scope of this local CLI scan. This report assesses structural signals only.

## Report Metadata

Field	Value
Schema Version	stem-ai-local-cli-result-v1.6
STEM BIO-AI Version	1.7.5
Generated (local date)	2026-05-15
Report Validity	180 days from audit date
Execution Mode	LOCAL_ANALYSIS
Repository	yorkeccak/bio
Remote URL	<a href="https://github.com/yorkeccak/bio.git">https://github.com/yorkeccak/bio.git</a>
Branch	main
Commit (HEAD)	100a0bf7497e62ead024df34d8c2e00ae74b8d99
Files Scanned	166
Final Score / Tier	48 / 100 — T1 Quarantine