



AI Coding Is a Force Multiplier for Biotech R&D

Your scientists have ideas for tools they need. AI-assisted development turns those ideas into production-grade software in days, with an auditable quality trail for scientific R&D, clinical trials, and regulated contexts.

The opportunity

AI coding agents have fundamentally changed the economics of software development. A scientist with no programming background can now describe a workflow in plain language and receive a working application in hours. A computational chemist can build in a day what previously took a team weeks. This is not incremental improvement — it is a structural shift in how quickly ideas become working software.

For biotech and biopharma R&D teams, this shift has specific consequences. The backlog of custom tools that scientists need — assay analysis pipelines, protocol trackers, molecular property calculators, screening dashboards — no longer requires months of engineering queue time. The gap between a scientific hypothesis and a computational tool to test it can shrink from quarters to days. R&D teams that adopt AI-assisted development gain a direct capacity advantage: more experiments, faster iteration, less dependence on scarce engineering resource.


But speed without quality is dangerous in this domain. The opportunity is real only if the code that AI produces is trustworthy enough for scientific and regulatory contexts. That is the problem this paper addresses.

The problem with AI-generated code

AI coding agents can produce a working proof of concept in minutes. But when the same agent writes both the code and the tests in the same pass, the tests only confirm what the AI *thought* was correct — not what *should* be correct. A defect in the reasoning produces a matching defect in the tests. Everything passes. The bug ships.

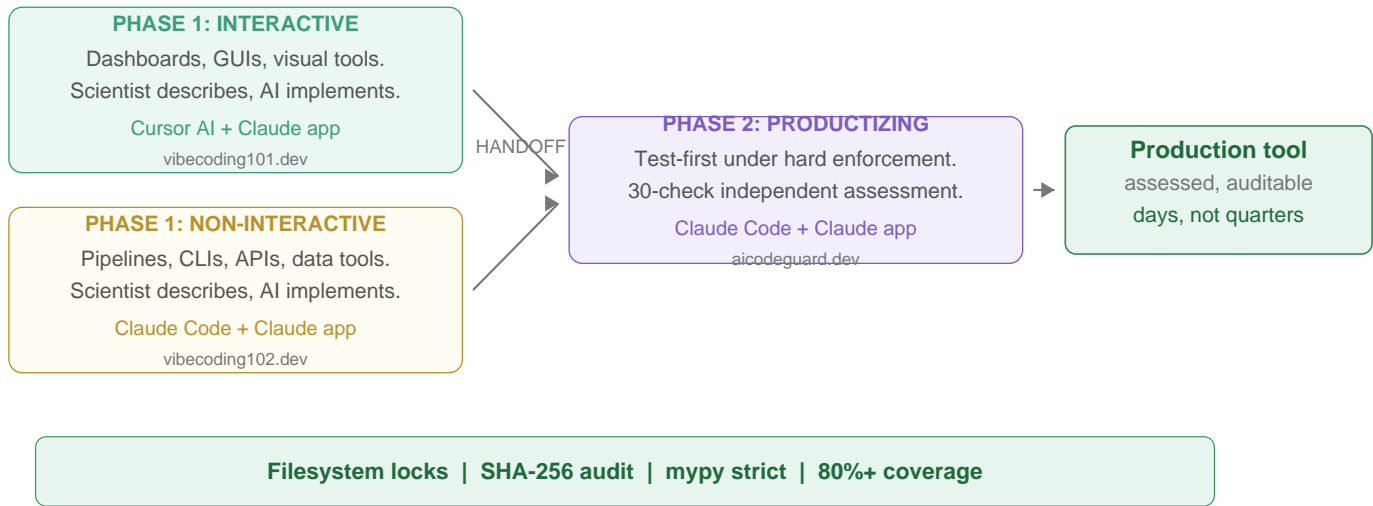
In a web application, these failures cause bugs and downtime. In scientific R&D software, they cause wrong conclusions. A software quality bug in your domain is a scientific integrity bug.

Real example: During conversion of a cell confluency assessment tool used in drug discovery, a boundary condition test discovered that `mask[-0:]` selects an entire NumPy array instead of an empty slice — causing incorrect image processing output in assay analysis. It was caught only because the test was written before the implementation existed.



Three tracks, one lifecycle

The methodology separates the work into three tracks. **Phase 1** has two paths depending on the type of application: interactive apps (dashboards, GUIs) built with Cursor AI, or non-interactive apps (pipelines, CLIs, APIs) built with Claude Code. Both use the Claude app for strategic planning. **Phase 2** then hardens any prototype into production-grade code under strict enforcement, also using Claude Code with the Claude app.



Phase 1 (Interactive) uses Cursor AI with the Claude app: the scientist describes a workflow in domain language, Cursor implements visual tools with live previews. Working prototype in hours. Full guide: vibecoding101.dev

Phase 1 (Non-interactive) uses Claude Code with the Claude app: the same rapid prototyping, but for pipelines, CLIs, and data processing tools that don't need a GUI. Full guide: vibecoding102.dev

Phase 2 uses Claude Code with the Claude app to productize: test-first development enforced by filesystem locks and SHA-256 hash audits, with a 30-check independent production readiness assessment. The handoff artefact is running code, not a document. Full specification: aicodeguard.dev

How this works for your R&D team

Biotech R&D teams naturally contain both halves of this workflow. The scientist who needs the tool and the data scientist or developer who could harden it are often in the same group — sometimes the same person. The handoff between Phase 1 and Phase 2 happens within your team, not across an organisational boundary.

What changes

Phase 1 addresses the backlog. The scientist — who understands the domain better than any developer — builds a working prototype using AI-assisted development. This takes hours to days. **Phase 2** addresses the quality gap. The prototype is converted to production code under hard enforcement. Your data scientists or developers review architecture and test quality instead of writing code from scratch. Their capacity is used for judgement, not implementation.



| Scientists | Data scientists / devs | R&D leadership |
|---|---|---|
| Describe the tool you need in domain language. AI builds a working prototype. Test it on real data. | Review architecture and test quality during hardening. Your skills are used for judgement. Every commit is traceable. | Tools built in days, not quarters. Auditable quality trail. Production-assessed code your QA team can evaluate. |

A 30-person biotech with 2 computational scientists and no dedicated software engineers is currently limited by how many tools those 2 people can build and maintain. If each of them can produce production-assessed tools at 5–10x the rate, that is not a marginal improvement — it is a structural change in what the team can attempt.

Lab and in-silico workflows

The methodology works for any R&D tool your team needs — whether it touches physical instruments in the lab or runs computations on molecular data. The same three-track lifecycle applies.

| Lab workflows | In-silico workflows |
|---|---|
| Assay protocol tracking — version control with audit trail | Molecular property calculator — ADMET, Lipinski, custom scoring |
| Plate reader analysis — dose-response curves, IC50 values | Virtual screening pipeline — docking, ranking, hit list export |
| Screening QC dashboard — Z-prime, drift detection | Genomics annotation tool — variants, pathways, impact scoring |
| Instrument integration — readers, handlers, imagers to database | Knowledge graph builder — papers and transcripts to networks |

Scope profiles adapt to the domain: *Scientific R&D (reproducibility, numerical regression, provenance)* · *Clinical Trials (audit trail, ALCOA+, edit checks)* · *Regulated Medical (IEC 62304, traceability)* · *Pre-clinical baseline for general R&D tools.*

Learn more

| | |
|-----------------------------------|--|
| aidrug.dev | Overview of the complete methodology for life sciences teams |
| vibecoding101.dev | Phase 1: AI-accelerated prototyping guide for interactive apps (Cursor AI) |
| vibecoding102.dev | Phase 1: AI-accelerated prototyping guide for non-interactive apps (Claude Code) |
| aicodeguard.dev | Phase 2: VP-model orchestrator, enforcement, scope profiles, 30-check assessment |

Ready to explore this for your R&D team?

Whether you need lab workflow tools, in-silico pipelines, or both — the methodology is documented, validated, and ready to discuss. We work with biotech and biopharma R&D teams to turn scientific ideas into production-assessed software.

Dr Raminderpal Singh

raminderpal@hitchhikersai.org

20/15 Visioneers · raminderpalsingh.com