



ORIENT · BUILD · AUTOMATE

LLM WORKSHOPS FOR ORGANIZATIONS

Biotech AI Enablement Program

Turning AI curiosity into safe, everyday capability for life-science teams.

Instructor led · In person first · Hybrid and remote options

Aeneas Bio · aeneas.bio · 2026

WHY THIS MATTERS

Training Decides the Outcome

LEADERS

Move faster, stay ahead

Leaders want AI to make the organization faster and more profitable, and they fear being left behind.

EMPLOYEES

Stretched thin

Employees feel stretched thin and fear the same outcome from a different position.

THE DIFFERENCE

Risk or advantage

Untrained use creates risk.
Trained use can become a competitive advantage.

Participants leave understanding what AI can do within their own role. Preserving users' judgment is the core principle.

WHY NOW

AI Is Scaling Faster Than Teams Can Absorb It

\$8.8B

Projected AI-in-biopharma market by 2030, up from roughly \$2.1B in 2025.

~95%

Estimated share of GenAI pilots in pharma that stall for lack of trained staff to run them.

~60%

Share of the workforce expected to need reskilling by 2027, per World Economic Forum estimates.

THE GAP

Enterprise AI platforms are scaling fast. Hands-on, role-specific training for scientists and non-coders is undersupplied.

The Current State

01

Adoption without instruction

Staff already use AI tools on their own, which produces uneven output quality.

02

Uneven skill, uneven results

Online content overstates capabilities, and tool exposure varies across teams.

03

Real exposure, quietly leaked

Sensitive data, intellectual property, and regulated workflows can end up in the wrong tool, creating personal and organizational risk.

04

Grow faster, shared experience

Workshops and chat groups support idea sharing, and participants connect quickly with resources and experts.

Orient, Build, Automate

01 · 5 TOPICS

Orient

Build the fundamentals and understand the tools.

- LLM Foundations and Mechanics
- Platform and Model Selection
- Data Security and Governance
- Structured Prompt Engineering
- Development Environment Bridge

02 · 5 TOPICS

Build

Construct the tools, skills, and plugins that perform tasks.

- Work Product Definition
- Spec-Driven Development
- Knowledgebases and RAG Architecture
- Verification, Testing, and Refinement
- Compliance and Audit Guardrails

03 · 5 TOPICS

Automate

Bring it together and express your agency with confidence in your current job.

- Scheduling and Routine Execution
- Agentic Loops and Decision Logic
- Life-Science Platform Integration
- Agent-to-Agent Orchestration
- Human-in-the-Loop Oversight

Build the Fundamentals

1.1	LLM Foundations and Mechanics	How models generate, reason, and fail, plus the core vocabulary.
1.2	Platform and Model Selection	Compare providers on cost, capability, and data handling.
1.3	Data Security and Governance	What can and cannot go into a prompt under your policies.
1.4	Structured Prompt Engineering	Frameworks like ROSE and TAG for consistent, reliable output.
1.5	Development Environment Bridge	The tools and access that carry training into daily work.

HANDS-ON
EXERCISE

Each participant takes one real task from their own role, runs it on two platforms, and scores cost, capability, and data handling. A data-sorting drill then separates what can and cannot go into a prompt, and the task is rewritten with ROSE or TAG.

Construct Tools, Skills, and Plugins

- | | | |
|-----|--|---|
| 2.1 | Work Product Definition | Name the deliverable first: what done looks like. |
| 2.2 | Spec-Driven Development | Write the spec the AI executes against, like an assay protocol. |
| 2.3 | Knowledgebases and RAG Architecture | Ground outputs in your own documents and data. |
| 2.4 | Verification, Testing, and Refinement | Check output against ground truth and iterate until it holds. |
| 2.5 | Compliance and Audit Guardrails | Build review and traceability into the workflow from the start. |

HANDS-ON EXERCISE

Choose a real deliverable from your team, such as a weekly QC summary. Define what done looks like, write the spec, ground the AI in your own SOPs, then test the output against ground truth and keep the review trail.

Express Your Agency with Confidence

3.1	Scheduling and Routine Execution	Run a validated workflow on a recurring basis, no manual restart.
3.2	Agentic Loops and Decision Logic	Let an agent take multiple steps within set bounds.
3.3	Life-Science Platform Integration	Connect agents to LIMS, ELN, and other systems already in use.
3.4	Agent-to-Agent Orchestration	Coordinate multiple agents across a multi-step process.
3.5	Human-in-the-Loop Oversight	Define where a person must approve before the agent proceeds.

HANDS-ON EXERCISE

Take the validated workflow from Build and schedule it. Set the bounds the agent decides within, connect it to a system your team already uses, and mark the points where a person must approve before it continues.

DELIVERY

In Person First, Flexible by Design

OUR PREFERENCE

In person

Instructors observe each participant and confirm concepts stick before moving on.

MIXED TEAMS

Hybrid

On-site and remote participants share one session. Every seat works the same real task.

DISTRIBUTED TEAMS

Fully remote

Live facilitation with the same exercises, run in your own tools.

EXERCISES BUILT FOR LIFE-SCIENCE WORK

Every exercise is tailored to real work in life sciences, using your team's own tasks, documents, and data wherever policy allows.

Research · Lab operations · Commercial · Technical workflows

PROGRAM STRUCTURE

Scoped to Your Team

FORMAT	Instructor led, in person first, with hybrid and fully remote options
COHORT SIZE	Small groups with a hands-on ratio
DURATION	Modular half-day or full-day sessions
CADENCE	A single session or a multi-session engagement
AFTER THE WORKSHOP	An ongoing coach and accountability partner

BRIDGE SESSION	An optional lighter preview of the role-specific tracks for organizations planning multi-session engagements. Module depth, group size, and pacing are set during a discovery call.
----------------	---

Who Else Teaches AI to Biopharma

Provider type	What they offer	What is often missing
Academic exec ed	Strategic framing, certificates	Hands-on training, on-site delivery
Online platforms	Self-paced ML and GenAI courses	Team training, live instruction
Associations	Webinars, working groups, regulatory framing	Platform setup, agentic workflows
Large consultancies	Enterprise AI strategy, large rollouts	Boutique delivery, speed
Boutique trainers	Bioinformatics and data science bootcamps	LLM training, governance framing

THE OPEN SPACE

We have not found a provider that combines life-science depth, security and governance training, and a documented multi-session reference engagement.

Specialized, Live, and Hands-On

LARGE CONSULTANCIES

Broad strategy, higher
cost, waitlists.

GENERIC PLATFORMS

Self-paced video, no live
facilitation, no security or
governance focus.

AENEAS BIO

- Life sciences and biotech specialization with boutique delivery
- Live, hands-on facilitation, in person first, hybrid or remote when needed
- Field-specific labs and content
- A proven multi-session reference engagement

From Scattered Tools to One Baseline

CHALLENGE

Staff used different AI tools independently, skill levels varied across the organization, and leadership worried about data exfiltration.

ENGAGEMENT

Aeneas Bio helped select an LLM platform and delivered a multi-session program with live co-facilitation using the client's own staff and data.

OUTCOME

A consistent skill baseline, near 100% adoption within 3 to 4 weeks, more consistent and timely artifacts, and reusable data governance policies.

"This is the level of technical depth our staff needed, with each skill level accommodated."

Client feedback, biotech innovator

NEXT STEPS

Book a Discovery Call

01

Book a 30-minute discovery call.

02

We scope the workshop to your audience, frequency, goals, and delivery format: in person, hybrid, or fully remote.

03

Your team gets hands-on training with customized materials built around real work in life sciences.

training@aeneas.bio

aeneas.bio · Raleigh, NC