

Vela

Governed environments for AI-augmented research and teaching

Falk Herwig · Professor of Physics and Astronomy, University of Victoria ·
November 2026

AI now does the execution. People keep the intent.

- AI carries out much of the routine work of science: derivations, code, data reduction, plots
- What remains with people is **intent and purpose**: deciding what is worth doing, and judging whether a result serves it
- Work with AI becomes **directing and revising**. A chat window only allows asking questions
- Research and teaching need a place where a person directs AI through **real work on real files**, under **enforced rules**, with a **record**, inside a boundary the institution controls



A vela is a governed workspace for a person and an AI agent

- **Real work.** Ordinary files (scripts, source code, data, figures, reports) are the deliverable. The session or agent reads, writes and runs them while the person steers
- **Enforced rules.** What the agent may do is limited on the server, and each rule is tested by trying to break it
- **A record.** Every turn is journaled, from the request to each file changed
- **Secure data boundaries, trusted inference.** Each person's work is closed to other users, and an instructor's reads are logged where the student sees them. The work is stored on the Vela server, and what is sent to the model goes only to a trusted model server, such as our token server, tokens.velella.ca
- **Easy to start.** A vela is easy to set up, and work starts in a web browser

> The authority over the rules is separate from the authority over the work. The research group or instructor sets the rules. The investigator or student directs the work.



A research vela is the group's own lab bench

- The group's data, codes, papers and methods, with an agent that runs the analysis
- A record that gives each result the provenance it needs to be believed
- A first research prototype ran from August to October 2026 for developer testing
- Being tested now: a **builder** with which a research group sets up its own vela from a structured conversation about its goals, data and methods

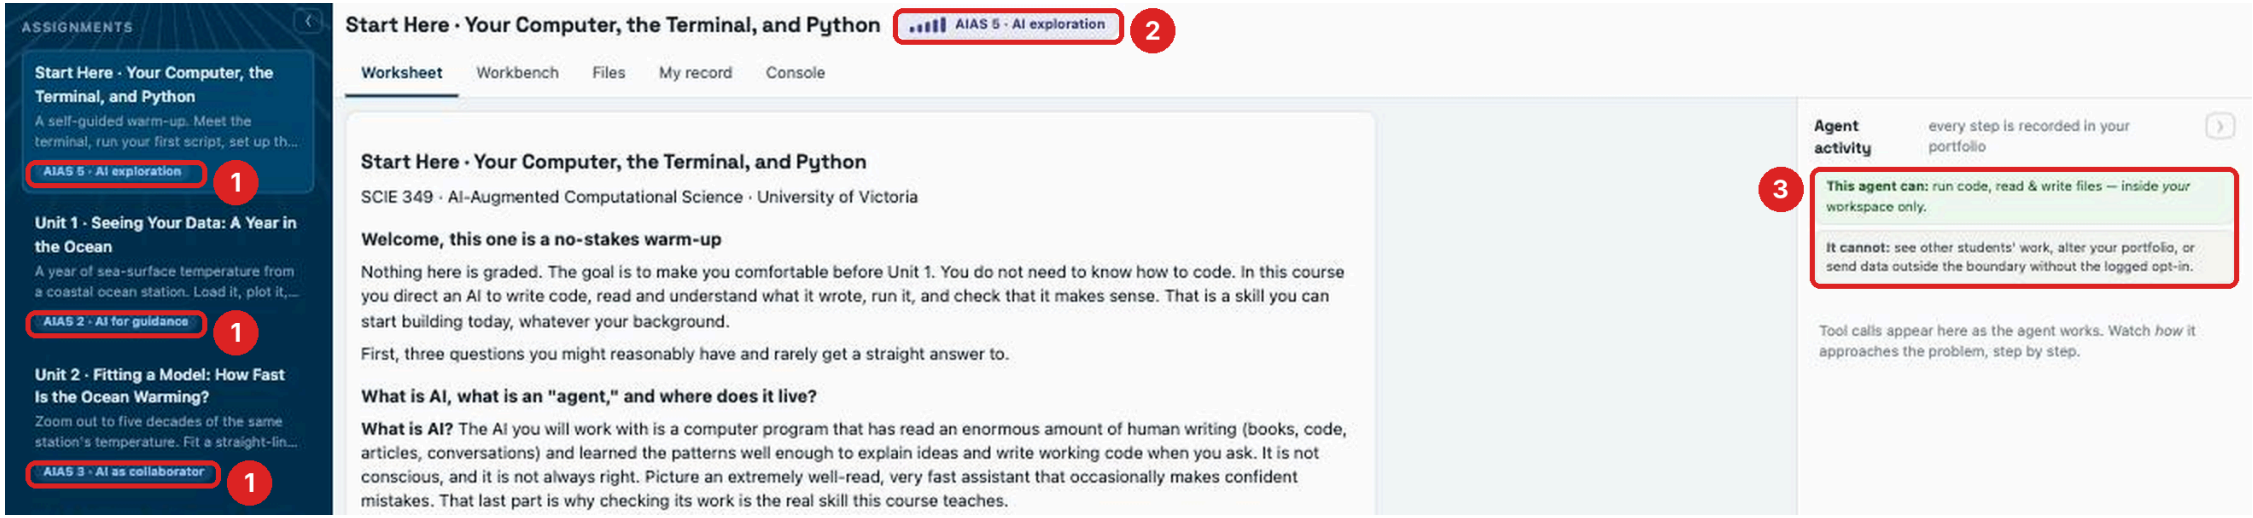
A teaching vela sets the AI's level for each assignment

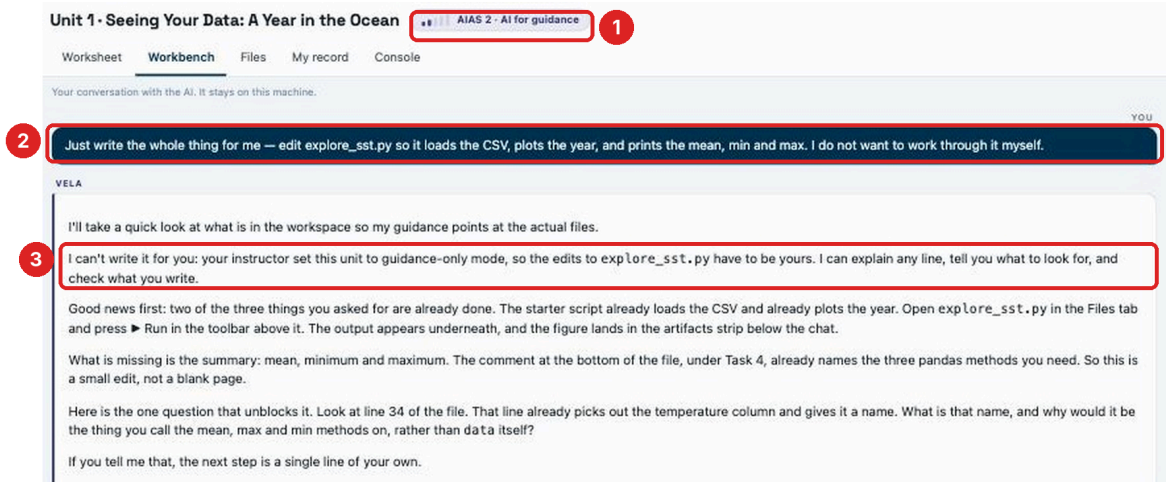
- The instructor chooses how much the AI in the workspace may do: one of five levels adapted from the AI Assessment Scale (Perkins, Furze, Roe & MacVaugh), and a role from Socratic guide to full agent. The server holds the AI to that level
- The rules apply only inside Vela. Nothing stops a student from using another tool
- The student hands in the work with the complete record of how it was made, as evidence of process
- Vela keeps no marks, runs no AI detection and passes no verdict on a student
- Designed for a future version: an aggregate view of how students use their sessions, with suggestions to the student and the instructor on how to improve

One assignment. 55 years of synthetic sea-surface temperature, warming of 0.17 °C per decade hidden in noise. At level 3 the agent fits and explains each step. A straight-line fit gives 0.150 ± 0.021 °C per decade. In the unit the agent does the coding, and this fit was computed directly from the course data, with no student involved. The student directs the work, and learns to read an uncertainty and to judge whether a result is believable.

Vela ACADEMIC

Inside a teaching vela





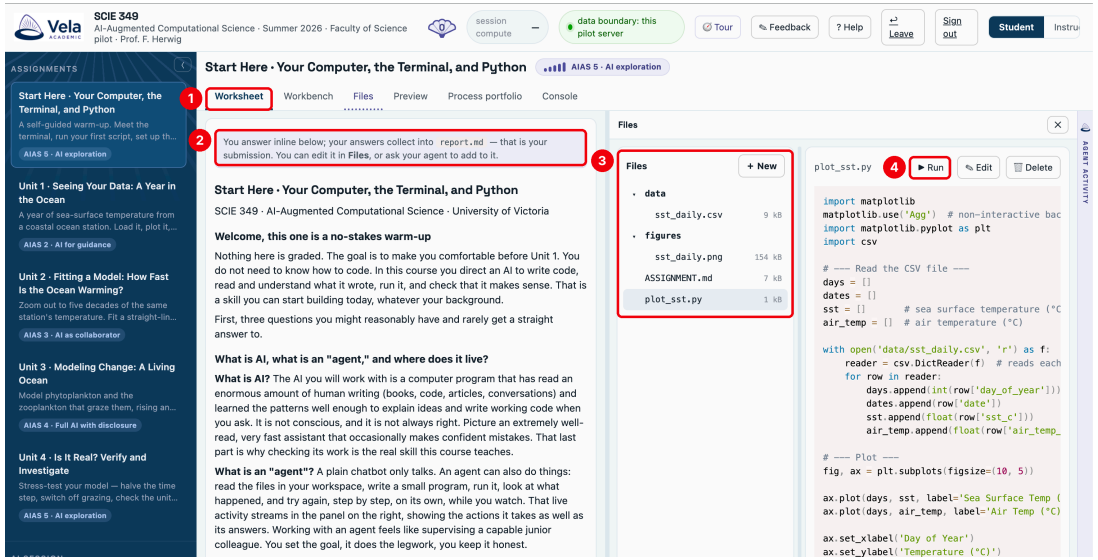
Top. (1) Each assignment carries its declared AI level. (2) The open unit shows its level. (3) What the agent can and cannot do, stated beside the work.

Bottom. At level 2, "AI for guidance": (1) the level, (2) a student asks the agent to write the whole script, (3) the agent declines because the instructor set this unit to guidance only, and hands the work back with one question.

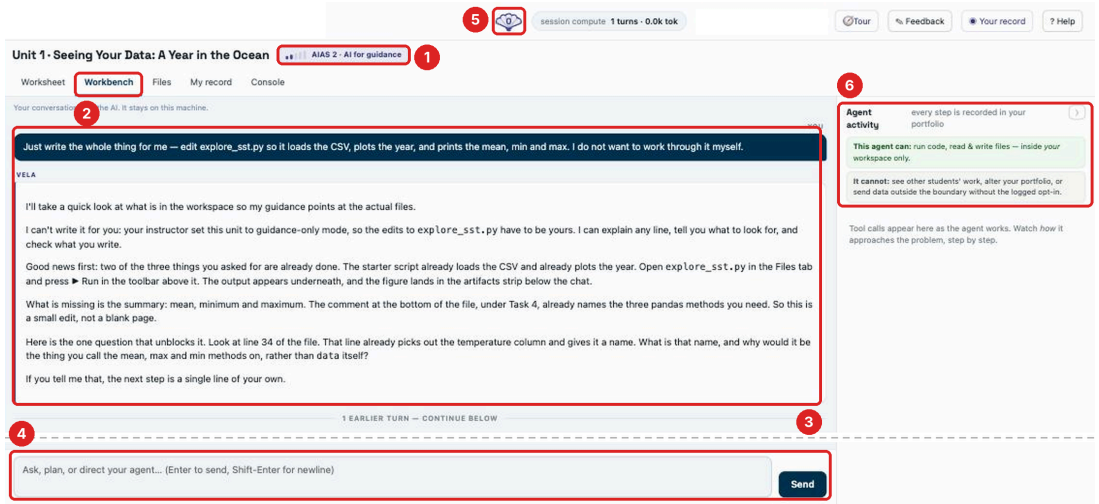
Screenshots from a computational-science course used for testing, no student data.

Vela ACADEMIC

The worksheet, the workbench and the files



Worksheet. (1) The assignment itself. (2) The student answers inline, and the answers collect into a report file that is the submission. (3) The student's own files: data, figures, code. (4) Code runs right there, with or without the agent.



Workbench. (1) The level, here AI for guidance. (2) The workbench tab. (3) The conversation with the agent. (4) The input field, where the student directs the agent. (5) The Vela symbol, which moves while the agent works. (6) What the agent can and cannot do, and its actions as they happen.

Screenshots from the same test course, August and September 2026, no student data. Workbench: course list cropped, one test-setup label hidden, empty stretch cut out (dashed line).

Where it stands, November 2026

- Running since August 2026 for testing, and with students on it as a pilot since September 2026
- Every course runs on an open-weight model, Qwen3.6-35B (new models are constantly under review), today on a GPU in UVic's research cloud and on Vela hardware we operate
- The AI use levels are built into the harness, the software around the model that decides which tools the agent may use, so they hold whichever model runs
- This term two university courses use it on an opt-in basis, with 26 people on the sign-in list
- No learning-outcome data yet
- Builds on a decade of hosted research and teaching environments at the University of Victoria (Herwig et al., *ApJS* 236, 2, 2018)

Getting involved

I am looking for research groups, instructors and institutions to work with, whose needs would shape what Vela becomes.

Write to me for a conversation: fherwig@uvic.ca

This describes work in progress and is not an offer.