ACQUILAB

AcquiLAB manual

Analysis, Plot Studio, Bayesian optimization, simulation, the Library with Zotero sync, Notebook and Documents, the Assistant and the MCP connector, plans and settings.

Updated 5 October 2026 · acquilab.io/docs

Contents

1 **Getting started**
Projects and workspaces · Finding your way · The project overview · Data · Before anything runs · Your first hour

2 **Analyze**
From file to figure · Methods · Reading a result · When an analysis fails · Not available yet

3 **Plot Studio and Figure Studio**
Workbook · Graph types and styling · Figures with several panels · Export

4 **Optimize (Bayesian optimization)**
How it chooses the next experiment · Engines and models · EI, UCB, PI and the rest · Which to use when · Running a campaign · Setting up without BO experience · Model check and switching the model · Multi-objective campaigns · Campaigns through the MCP connector · When to stop

5 **Simulate**
Running a model · Packages and limits · Worked chemistry examples · Ready-made models · Fit to measured data · Writing your own model · Custom ODE library · Formula library · Simulation in the BO loop

6 **Library and Zotero sync**
One library for all projects · Adding references · Collections, tags and projects · Verification and retractions · Duplicates and trash · Export · Two-way Zotero sync

7 **Notebook and Documents**
Notebook · Documents · Citing papers · Citing results · Check numbers and citations · Export

8 **Assistant, MCP connector and connected agents**
The Assistant · Plan first · Connect an AI agent · Connected agents · Connector tools · Good practice for agents

9 **Plans, runs, credits and settings**
Plans · Runs · Assistant credits · Limits · Graph preferences · Settings · Data and privacy

Chapter 1

Getting started

Projects and workspaces, the sidebar, the project overview, Data and what AcquiLAB asks before it runs anything.

Projects and workspaces

AcquiLAB is organised around a **project**. Files, analysis runs, optimization campaigns, simulations, figures, notebook entries and documents all belong to one project, and every project sits in a **workspace**: your personal workspace or a team workspace you were invited to. Members of a project see its work; a team workspace shares billing, runs and credits.

Everything keeps its provenance: which file and version, which method and method version, which settings, which model. Re-running, correcting or deleting something never silently changes a result someone already used; the places that use it say a newer version exists.

Your reference Library is the one thing that is personal rather than per project: it follows you across projects, and you link items to the projects that should see them.

Finding your way

The sidebar rail on the left of the app groups the pages by what you are doing.

Group	Pages
Workspace	Home (the project overview), Projects
Science	Data, Analyze, Plot Studio, Figures (Figure Studio), Optimize, Simulate
Record & write	Notebook, Documents, Library
Ask	Assistant
Account	Usage & Billing, Settings

- The top bar holds the workspace and project switchers and a meter such as *41 runs | 282 credits* that opens Usage & Billing. Below 1,100 px wide it shows runs only.
- On a phone the tab bar at the bottom has Analyze, Optimize, Home, Simulate and More. The workspace and project names in the top bar open the More sheet with the switchers.
- Optimize opens the list of campaigns. From there you reach Bayesian Optimization (one objective), Multi-objective and Observations (every recorded result).

The project overview

Home opens the overview of the current project. It answers three questions: what is waiting for me, what happened this week, and where do I continue.

Waiting for you

Recommended experiments to run, analyses or simulations that failed, and work that finished.

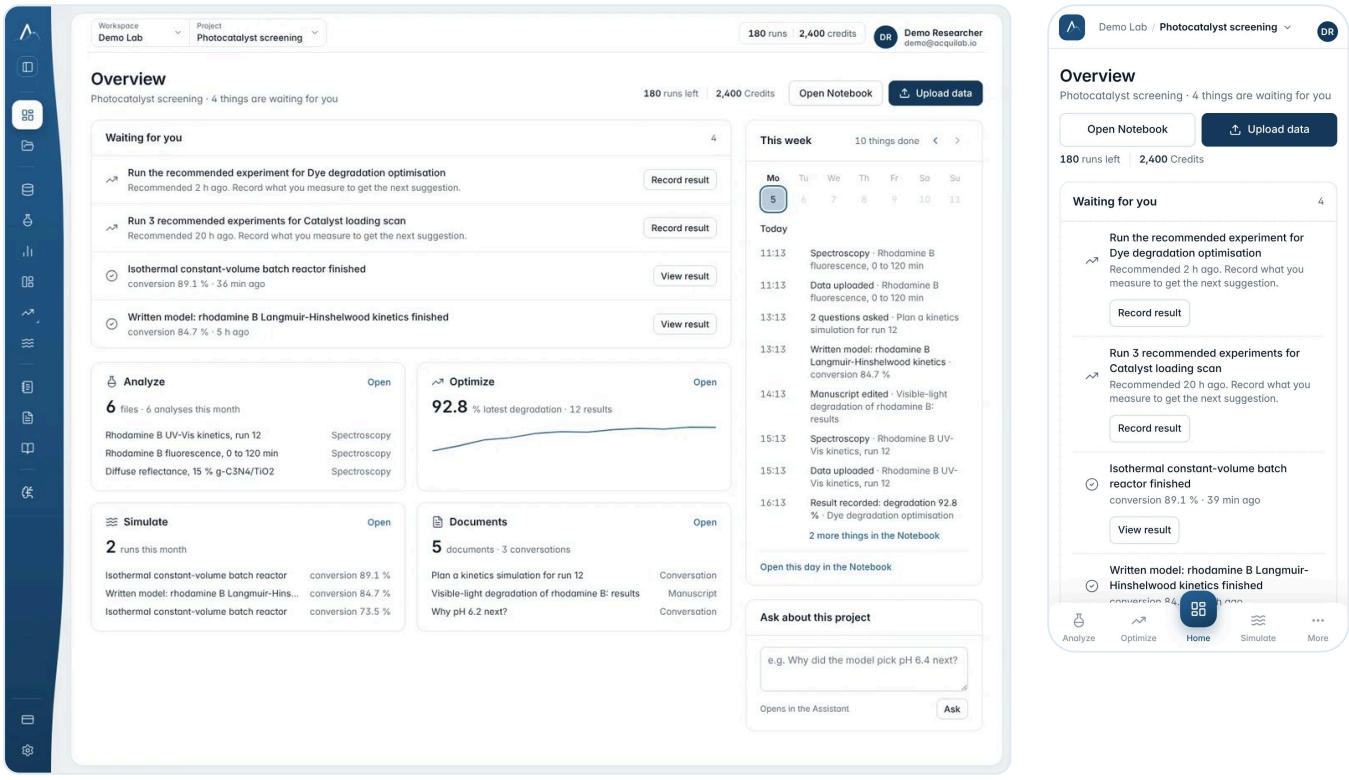
This week

Monday to Sunday, each day shaded by how much was done. Choose a day to list its work, or open that day in the Notebook.

Four panels

Analyze, Optimize, Simulate and Documents, each with the latest work and a button to start.

Ask about this project opens the Assistant with the project selected.



The project overview on a laptop and on a phone.

Data

Data lists everything the project holds, in three views.

View	What it shows
Files	Uploaded and derived datasets with their versions and the runs made on them. Per file: Open result, Fix setup and re-run, New run, Open in Plot Studio, Upload new version, download the latest result (JSON) or a summary (CSV), and Delete file.
Simulation results	Simulation runs grouped as compared sets, optimization loops and consecutive runs of one model. Select runs to compare them, or *Plot together*.
Optimization results	Every campaign, single objective and multi-objective, with its objective, number of results and best so far, and *Download results (CSV)* with every recorded result (the Pareto front marked for multi-objective campaigns).

Before anything runs

You start your own analyses and simulations with a button. When the Assistant or a connected AI agent wants to run something, it prepares it and waits for you:

- The Assistant shows a card with every setting and a **Run now** button. Nothing starts until you press it.
- Before any simulation, model fit or simulation optimization started by the Assistant or an agent, AcquiLAB asks for a **plan**: the mechanism and every constant with its source, the budget and the open questions. You answer the questions, and only then can runs start, within the plan's number of starts and for 24 hours. See Plan first.
- Work that uses runs or credits says so before it starts.

Your first hour

1. Sign in and open **Projects**. Create a project (or use the one you were invited to) and select it in the top bar.
2. Open **Analyze**, choose a technique, upload an instrument file and press **Analyze**. Plot Studio opens with the result. (Analyze)
3. Style the graph and export it, or place it in a figure. (Plot Studio and Figure Studio)
4. Open **Optimize** and create a campaign for something you are tuning in the lab: the result you measure and the settings you can change. Record what you already know, then ask for the next experiment. (Optimize)
5. Ask the Assistant a question about the project, such as *"Compare ΔEp across cycles in the CV I uploaded."* (Assistant)
6. At the end of the week, open the **Notebook** and make the weekly update into a document. (Notebook and Documents)

> **Features in this manual**
>
> The manual describes what is live in AcquiLAB. Where something is not available yet, it says "planned". Some pages appear only when your plan or account includes them, for example Multi-objective (Pro and Team) or Connected agents.

Chapter 2

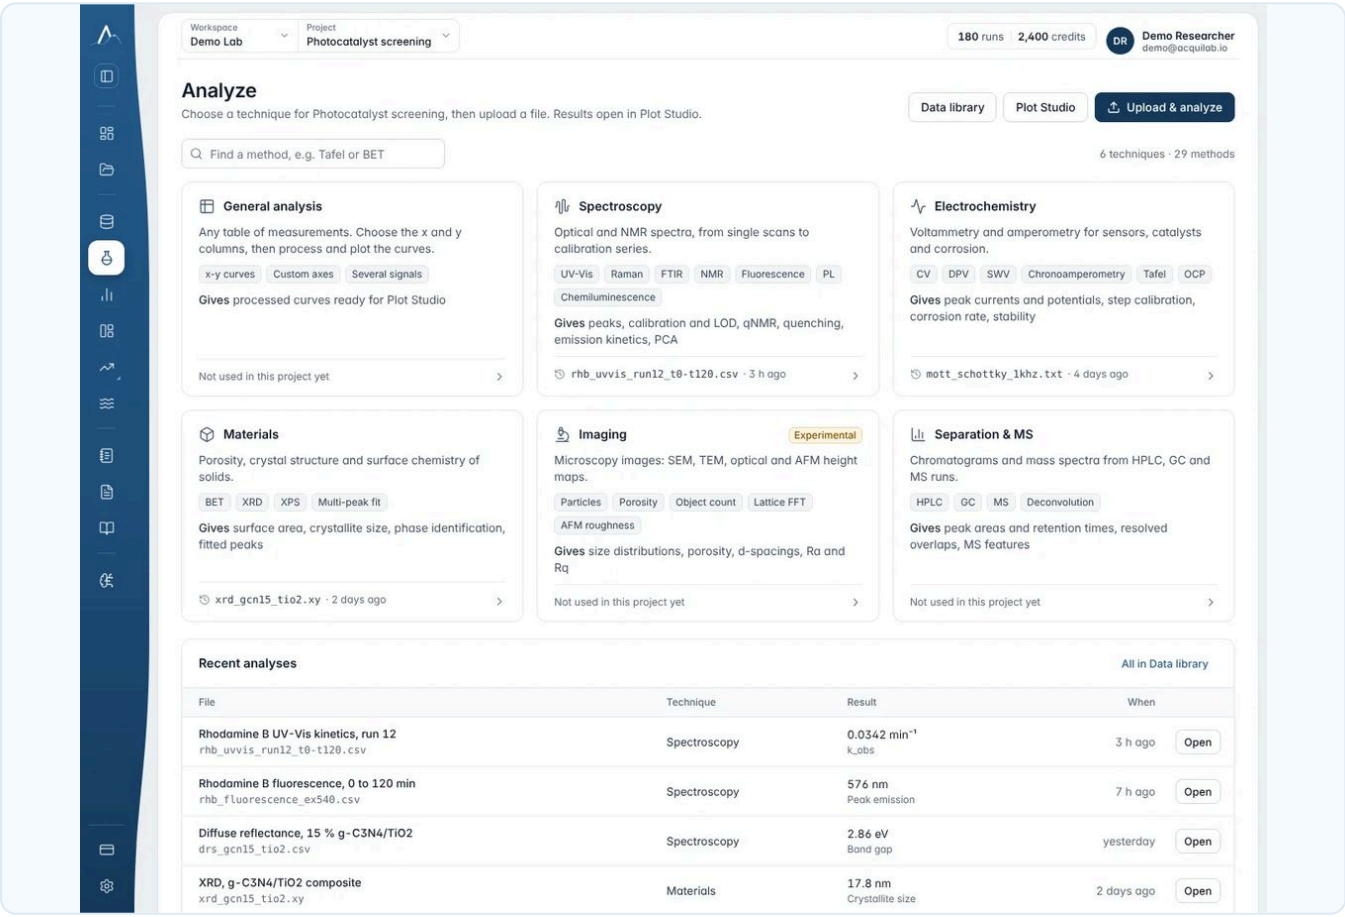
Analyze

From an instrument file to a result with units and uncertainties: techniques, methods, options and what to do when a run fails.

From file to figure

The normal path is four steps: **Analyze** → **technique** → **upload** → **Analyze**, and Plot Studio opens on the result.

The Analyze page shows one card per technique: General analysis, Spectroscopy, Electrochemistry, Materials, Imaging, and Separation & MS. Each card lists the methods it covers and what the analysis gives, and the search box matches methods, instruments and outputs (for example “qNMR”, “Scherrer” or “LOD”). Below the cards, *Recent analyses* lists the project's latest files with their state and headline result, with *Open* or *Fix and re-run*.



Analyze: one card per technique, and the project's recent analyses below.

1. Choose the technique card, then the method under it.

2. Upload the file: CSV, TXT or XLS for tables, TIFF for images, JSON for the structured materials and MS exports. Instrument text exports with metadata lines (JASCO and similar) are read as they are; you don't need to edit the file. Several files at once work for *Calibration from spectra* and *PCA*: they are placed side by side, one column per file.
3. Check the options. Columns are chosen from the header when the names are clear (for example `E (V)`, `Ewe/V`, `<I>/mA`, `Abs`); otherwise you pick them from the preview. Use `a.u.` for arbitrary units and `1` for dimensionless.
4. Press **Analyze**. Plot Studio opens with the result.

Methods

39 entries: 33 methods and the default analyzer of each technique. Every method reports values with units and, where the method defines them, uncertainties. Results are computed on the server; nothing in the browser recomputes them.

Method	What it gives you	Input	Main options
General analysis			
General x–y analysis	Processed curves of any table of measurements, ready for Plot Studio (for example the area under the curve).	CSV, TXT or XLS table	x and y columns with units (a.u. for arbitrary units, 1 for dimensionless)
Spectroscopy			
Spectroscopy (technique default)	UV-Vis style spectrum processing and peak positions.	CSV, TXT or XLS	—
Analyze a curve on the graph	Pseudo-Voigt peak fit, area under the curve, Savitzky–Golay smoothing or 1st/2nd derivative, or a baseline offset, over a chosen X range.	X and Y columns, at least 5 points in the range	what to do, X range, window (5–101 points, odd), polynomial order, peak near X
Baseline correction	Subtracts a baseline that follows the lower edge of the curve (asymmetric least squares) or a flat offset; gives the corrected curve to plot or analyse further.	X column (wavelength, wavenumber, shift, time ...) and Y column	baseline type, stiffness (flexible, medium, stiff), how strictly it stays under peaks, X range
Calibration curve (Beer–Lambert)	Straight-line fit of absorbance against concentration with residuals; slope ± SE, intercept, R^2, an unknown's concentration ± u.	Table of standards (at least 3 rows)	concentration and absorbance columns, unit, absorbance of an unknown
Calibration from spectra	Overlays the standards, reads each at the maximum, minimum, a wavelength or an area, fits the line and reads unknowns back.	Several files, or one wide table	signal kind, feature, range, concentrations (or units in column names), unknown columns
PCA of spectra	Scree plot, loadings against wavelength and the score plot.	Wide table, at least 3 spectra	axis in nm or cm^{-1}, 2–5 components, SNV, unit-length or maximum normalisation
Raman bands	Cosmic-ray removal, AsLS baseline and smoothing; band positions, heights and widths; pseudo-Voigt fits of the strongest bands.	Raman shift (cm^{-1}) and intensity	cosmic rays, baseline, smoothing, normalise, bands to fit (0–12), prominence

Method	What it gives you	Input	Main options
FTIR / IR bands	%T or T converted to absorbance, baseline, band positions and widths, fits.	Wavenumber and A or T	signal type, baseline, smoothing, bands to fit
Band assignment (suggested functional groups)	Suggests functional groups for bands above 1500 cm^{-1} (O–H, N–H, C–H, C≡N, C=O …) from 27 group-frequency ranges. Screening, not identification.	FTIR spectrum	signal type, baseline, smoothing
Fluorescence / emission spectrum	Emission maximum, FWHM, centroid, integrated intensity, all peaks and an optional I(λ1)/I(λ2) ratio.	Wavelength and intensity	emission or excitation, wavelength range, baseline, normalise, ratio wavelengths
Fluorescence quenching (Stern–Volmer)	$F_0/F = 1 + K_{SV}[Q]$: K_{SV} ± u and R^2.	[Q] and intensity (at least 3 rows)	concentration unit, F_0
Photoluminescence (PL)	Peak wavelength and photon energy, FWHM in nm and eV, integral, Stokes shift.	Wavelength and PL (absorption optional)	reference column and kind, baseline, smoothing
Chemiluminescence / emission kinetics	Maximum, time to maximum, total light, first-order decay fit and half-life.	Time and intensity	time unit, decay fit, background
NMR peaks and integrals	Peak list (shift, height, width in Hz) and integrals in protons relative to the first range.	ppm and intensity (MestReNova or TopSpin CSV)	nucleus ^{1}H, ^{13}C, ^{19}F or ^{31}P, solvent, referencing, integral and leave-out ranges, MHz
qNMR purity (internal standard)	Purity in % w/w from the analyte and internal-standard integrals, proton counts, molar masses and weighed masses.	Quantitative ^{1}H spectrum	both ppm ranges, protons, molar masses, masses, standard purity
Reaction conversion by NMR	Conversion from reactant and product signals per proton.	^{1}H spectrum of the mixture	both ranges and proton counts
Impurity and residual-solvent screening (^{1}H)	Flags residual solvents and common impurities for the chosen solvent and lists unexplained peaks.	^{1}H spectrum	your compound's ranges, S/N threshold
Peak types and J couplings	s, d, t, q … m with J in Hz (first-order patterns).	1D spectrum and spectrometer MHz	largest coupling, S/N, leave-out ranges
Electrochemistry			
Electrochemistry (technique default)	Technique analyzer for potentiostat exports.	CSV, TXT or XLS	—
Cyclic voltammetry peaks	Per cycle: E_{pa}, i_{pa}, E_{pc}, i_{pc}, ΔE_p, E_{mid} and $\lvert i_{pa}/i_{pc} \rvert$.	E and I in recorded order	units, smallest peak (default 5 % of the range)
Chronoamperometry / chronopotentiometry	Charge after the step, Cottrell plot and D; or baseline, steady state, change and t50/t90/t95 for a step inside the trace.	Time and current or potential	where the step is, steady-state %, n, area, concentration
Differential pulse / square-wave voltammetry peaks	Checks the pulse sequence, linear background, peak potential, height, FWHM and area.	E and (difference) current	technique, pulse amplitude, width or frequency, step, polarity
Tafel analysis (corrosion)	E_{corr}, j_{corr} and the anodic and cathodic Tafel slopes.	Polarization curve	units, area, reference electrode, fit ranges (automatic 60–180 mV)

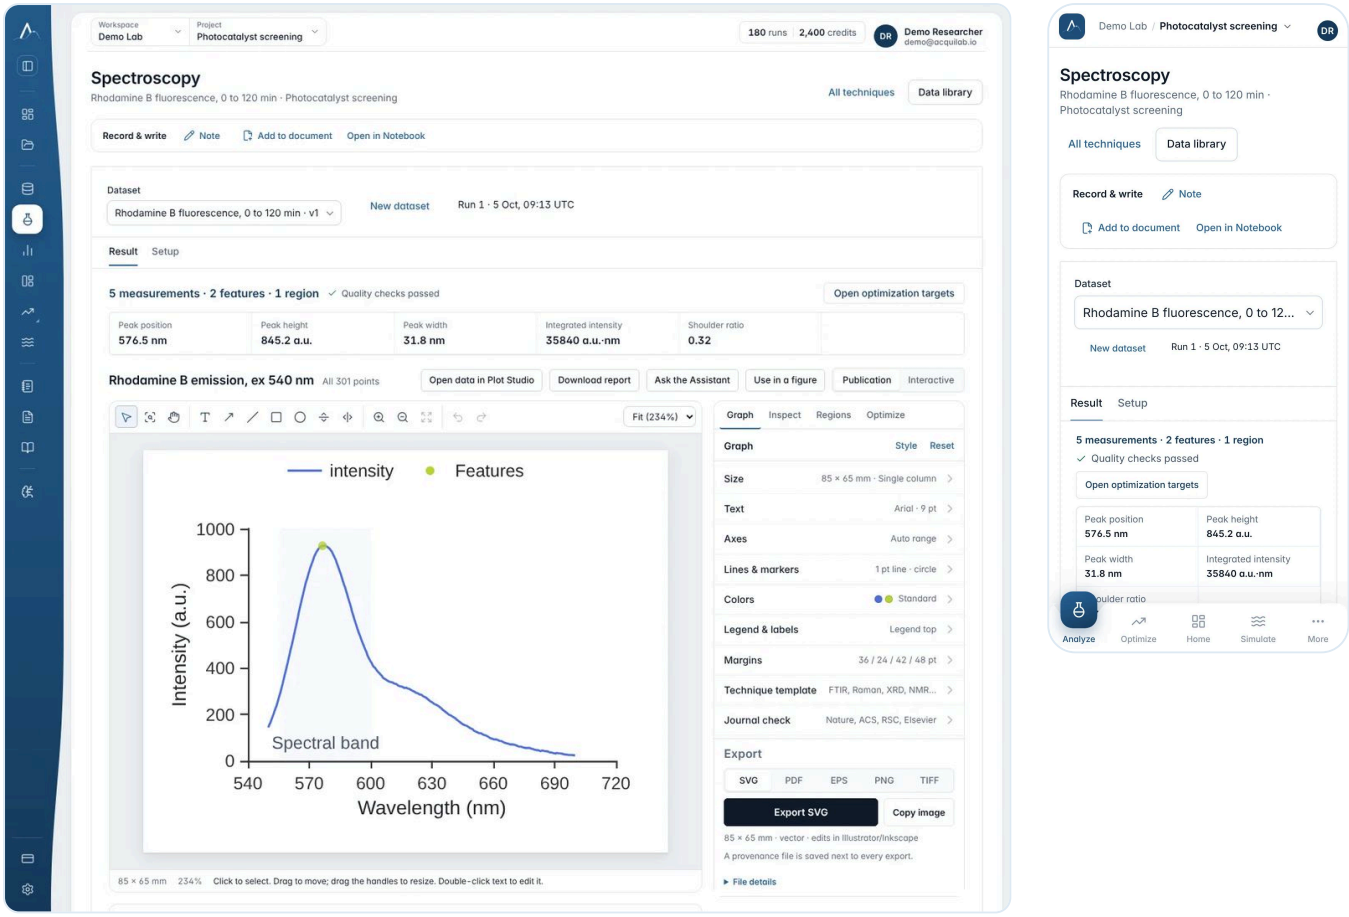
Method	What it gives you	Input	Main options
Calibration from injection steps (amperometric sensor)	Stable current (mean ± SD) before each injection, calibration line, sensitivity, R^2 and LOD.	One chronoamperometry recording with injections	concentration per step, window, injection times (automatic)
Materials			
Materials (technique default)	Materials analyzer (for example XRD patterns) for the structured JSON export.	JSON	—
BET surface area	BET range by the Rouquerol criteria, S_BET ± u, C constant and monolayer capacity.	p/p_0 and amount adsorbed	unit, gas N_2, Ar or Kr, temperature, cross-section
XRD peaks, d-spacings and crystallite size	Finds and fits the diffraction peaks (pseudo-Voigt): 2θ, d-spacing, width and Scherrer crystallite size of each peak; Williamson–Hall size and strain from four or more peaks.	Powder pattern: 2θ (degrees) and intensity	X-ray tube (Cu, Co, Mo, Cr, Fe) or another wavelength, background, smallest peak, peaks to fit, Scherrer K, instrument broadening, 2θ range
XRD phase identification (against your references)	Ranks the reference patterns you supply by how well they match (whole-pattern similarity and peak positions). Candidate screening: a phase that is not among your references can't be found.	One table on the same 2θ points: 2θ, your pattern, one column per reference phase (for example simulated from CIF files)	X-ray tube or wavelength, peak position tolerance, weight of the whole-pattern similarity, smallest peak
Fit several peaks at once	Fits overlapping peaks together (pseudo-Voigt) after an optional straight or Shirley background: position, width, area and area % of each peak with standard errors, the fit and its residual.	X and Y columns covering the peaks (XPS, Raman, IR, UV-Vis, XRD, chromatography ...)	peak positions (optional), number of peaks, smallest peak, background (straight, Shirley for XPS, none), X range
Imaging			
Imaging · Microscopy (technique default)	Image summary and measurements from TIFF microscopy images.	TIFF	—
Particles in an image (count and size)	Separates particles from the background by a threshold (touching particles split): number, covered area fraction, median diameter and number per area.	Greyscale TIFF (SEM, optical ...), bright particles on a darker background	pixel size and unit, threshold, smoothing, smallest particle, split touching particles, leave out particles cut by the edge
Porosity from an image	Takes the dark phase as pores: porosity (area fraction) and median pore size (local thickness).	Greyscale TIFF of a cross-section or membrane, darker pores	pixel size and unit, threshold, smoothing, smallest pore
Count objects in a fluorescence image	Counts bright objects (cells, nuclei, colonies, droplets ...): number, number per area and how many touch the edge.	Greyscale TIFF, one channel, bright objects on a dark background	pixel size and unit, threshold, smoothing, smallest object, split touching objects, what the objects are
Lattice spacings from an HRTEM image (FFT)	Fourier transform of a lattice image: the spots, lattice d-spacings and the angle between two lattice directions. Screening, not crystallographic indexing.	Square-ish HRTEM or STEM TIFF with lattice fringes	pixel size and unit, window before the FFT (Hann or none), smallest spot, maximum number of spots
AFM roughness and topography	Areal roughness Sa and Sq, the height distribution (percentiles) and the number of raised features.	AFM height image as TIFF (heights as pixel values, for example exported from Gwyddion)	pixel size and unit, height unit, height that counts as a raised feature, smallest feature

Method	What it gives you	Input	Main options
Separation & MS			
Separation & MS · Mass spectrum (technique default)	Centroid peaks with a noise baseline; mass spectrum profile with sticks.	JSON export	—
Chromatogram peaks and areas (integration)	Baseline, peak detection and integration: retention time, height, area and area %, width, tailing factor, plate count and resolution of each peak.	Time and detector signal (HPLC, GC, IC, SEC …)	baseline (straight, curved, none), regions kept out of the baseline, smoothing, smallest peak, integration width, minimum peak distance, peak direction, time range
Deconvolution of overlapping peaks	Separates co-eluting or shouldering peaks into Gaussian components on a straight local baseline: retention time, height, width, area and area % of each, with the fit and its quality.	Time and signal; a time window around the overlap	time window, retention times of the components (optional), number of components, shared or separate widths, how far a centre may move

A technique card marked *Experimental* or *Beta* is still being checked against reference data. Every result records the method, its version and its scientific maturity.

Reading a result

- The results table lists each value with its unit and uncertainty, and the quality checks the method ran. The columns used are named, so you can see what was read.
- The graph opens in Plot Studio, drawn in your saved style. Nothing there recomputes the result.
- *Next on this data* suggests the logical follow-up method with your settings carried over, for example “Name the peak types” or “qNMR purity” after NMR peaks.
- The *Record & write* bar adds a note, adds the result to a document or opens it in the Notebook.



A fluorescence spectrum result: the measurements and the graph, on a laptop and on a phone.

When an analysis fails

A failed run says why. When the problem is in the setup, the message names a reason and what to fix; the file and the run stay in **Data**, and *Fix setup and re-run* opens the setup with your previous choices.

Reason	What to do
`column_not_found`	A selected column was not found in the file. Re-open the setup and pick the columns again from the preview.
`columns_same`	The x axis and the signal are the same column. Choose two different columns.
`unit_missing`	Enter a unit for both columns (a.u. for arbitrary units, 1 for dimensionless).
`unit_not_recognised`	Use a standard symbol such as nm, s, mV or counts, a.u. for arbitrary units, or 1 for dimensionless.
`not_a_table`	The file could not be read as a table of numbers with a header row.
`non_numeric_values`	A selected column has empty or non-numeric values (text, blanks or NaN). Choose columns that contain only numbers, or clean those rows.
`too_few_points`	The selected columns need at least two numeric rows.

Other failures are reported in plain words: the analysis type does not support this input, the file is too large for this analysis, the analysis took too long and was stopped, or the analysis service is not ready (try again in a few minutes).

> **Arbitrary units**
>
> Write intensities as `a.u.` (or leave the suggested unit). AcquiLAB reads `a.u.`, `au`, `arb. units` and `RFU` as arbitrary units, so a fluorescence intensity is never taken for an astronomical or atomic unit.

Not available yet

Some techniques are described in AcquiLAB's skill catalogue as specifications but cannot run yet, among them impedance spectra you measured (EIS; the Randles circuit exists as a simulation), GC-MS and LC-MS, DLS and zeta potential, DSC and TGA, ELISA, qPCR, flow cytometry, western blots, rheology and mechanical testing. AcquiLAB says so instead of guessing; they are planned.

Chapter 3

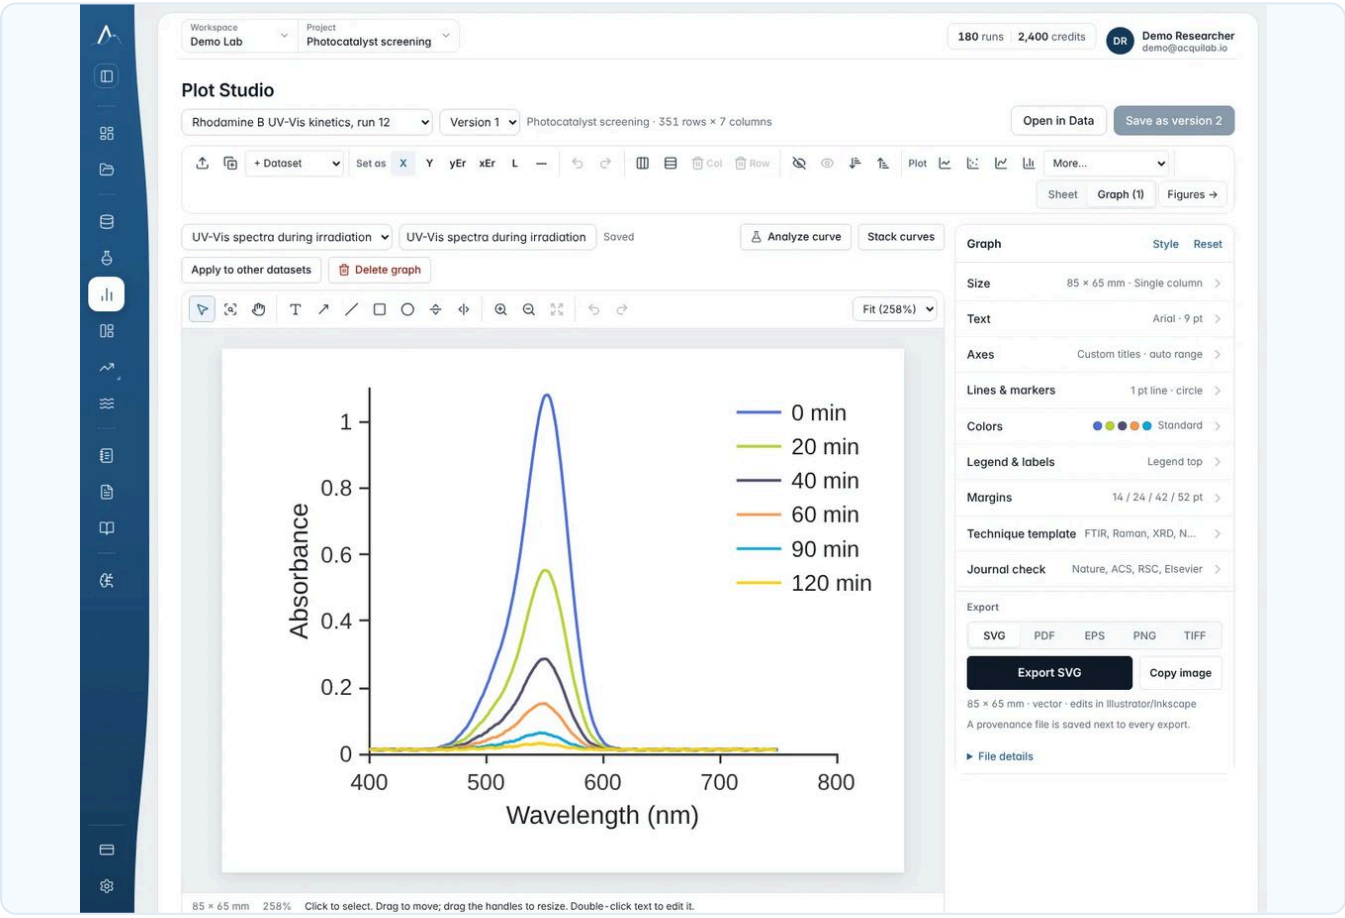
Plot Studio and Figure Studio

The workbook and graph editor for publication figures, multi-panel figures and vector export.

Workbook

Plot Studio is the workbook and graph editor for publication figures. A **sheet** is the table, a **graph** is one plot, and a **figure** is several panels arranged in **Figures** (Figure Studio). Nothing in Plot Studio recomputes a scientific result: it draws the numbers the server returned.

- Header rows for the long name, unit and comment of each column.
- Column designations X, Y and error; formulas (parsed, never executed as code); mask and sort rows.
- XLSX import; save as a new version of the dataset or as a new dataset.
- Results open here from Analyze, Simulate and Data (*Open in Plot Studio*).



Plot Studio: a graph of several spectra with its legend and axis titles.

Graph types and styling

Graph types

Line, scatter, line and symbol, sticks, stacked or grouped columns, histogram, box, violin, bubble, heat map and contour, pie or donut, radar, parallel coordinates and polar.

Styling

Style per curve; frame (open, box, box with ticks); ticks in or out and minor ticks; top X or right Y per curve; a converted right axis (A → %T, °C ↔ K, a·y + b); reversed axes and axis breaks; waterfall offset; normalisation; error bars; colour-blind and greyscale palettes; copy and paste a format.

New graphs start from your saved look when graph preferences are on: frame, ticks, text size, colours, error bars, grid, figure size and export resolution. Graphs that an AI agent asks for through the MCP connector are drawn in the same style.

Figures with several panels

Figures opens Figure Studio. Arrange graphs as 2×1, 2×2 or stacked panels that share the X axis, with a/b/c panel labels and insets. On a result page, *Use in a figure* places that graph in a figure; batch plotting draws the same graph for several datasets.

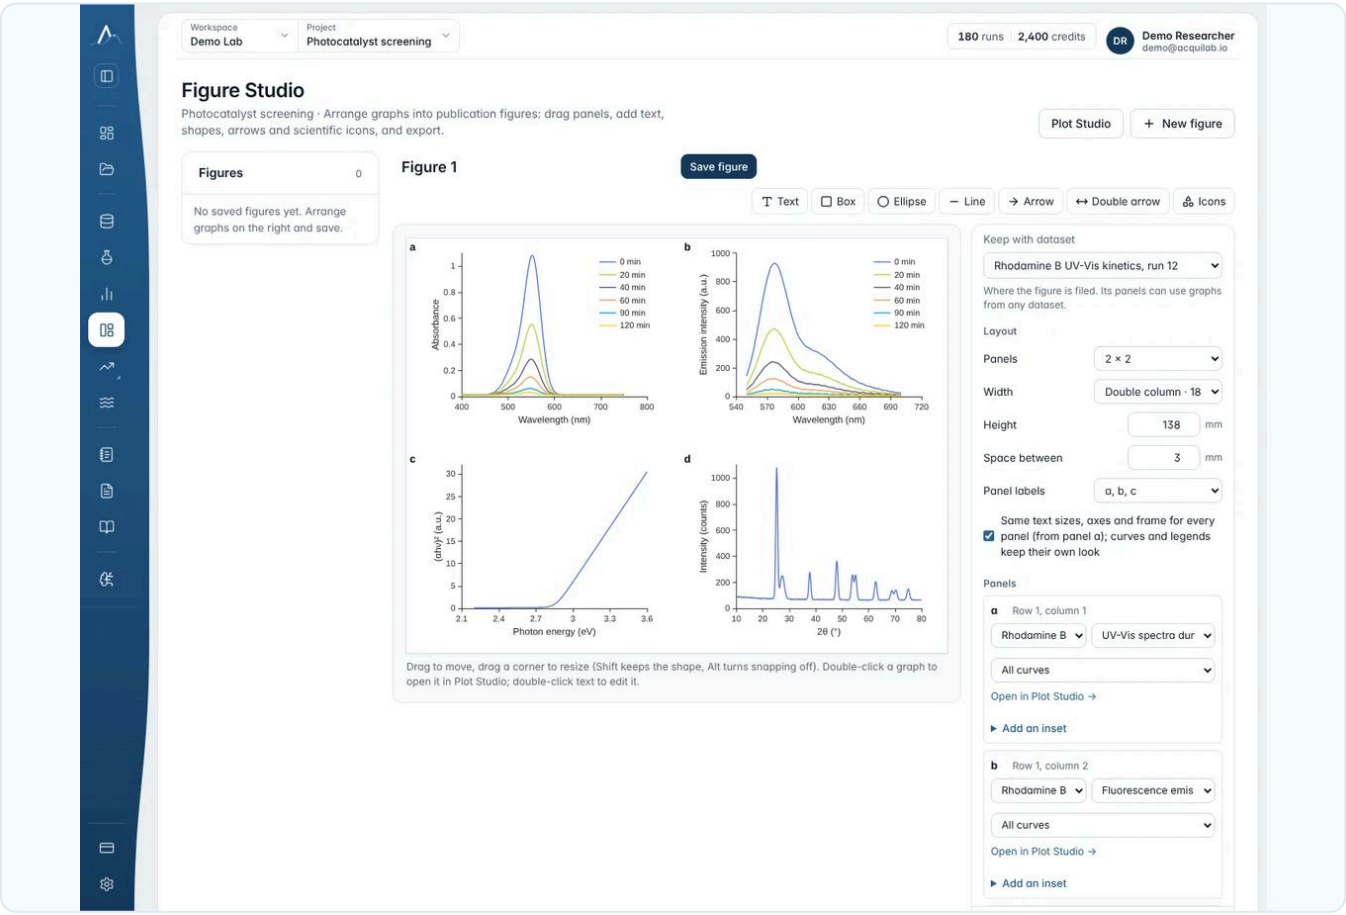


Figure Studio: graphs arranged as one figure with panel labels.

A saved figure can be placed in a document. The document keeps the drawing as placed and says *Figure updated in Figure Studio — Update* when the figure changes.

Export

Format	Use it for
PDF, EPS, SVG	Vector figures for journals and for editing in a drawing program.
PNG, TIFF (LZW)	Raster images with the resolution written into the file (300 or 600 dpi are typical).
Copy image	Paste straight into slides or a message.

Figure sizes for journals

The *Figure size* preference offers single and double column widths (85 and 180 mm) and the widths used by Nature, ACS, RSC and Elsevier journals, so a figure exported at 600 dpi fits the column without scaling.

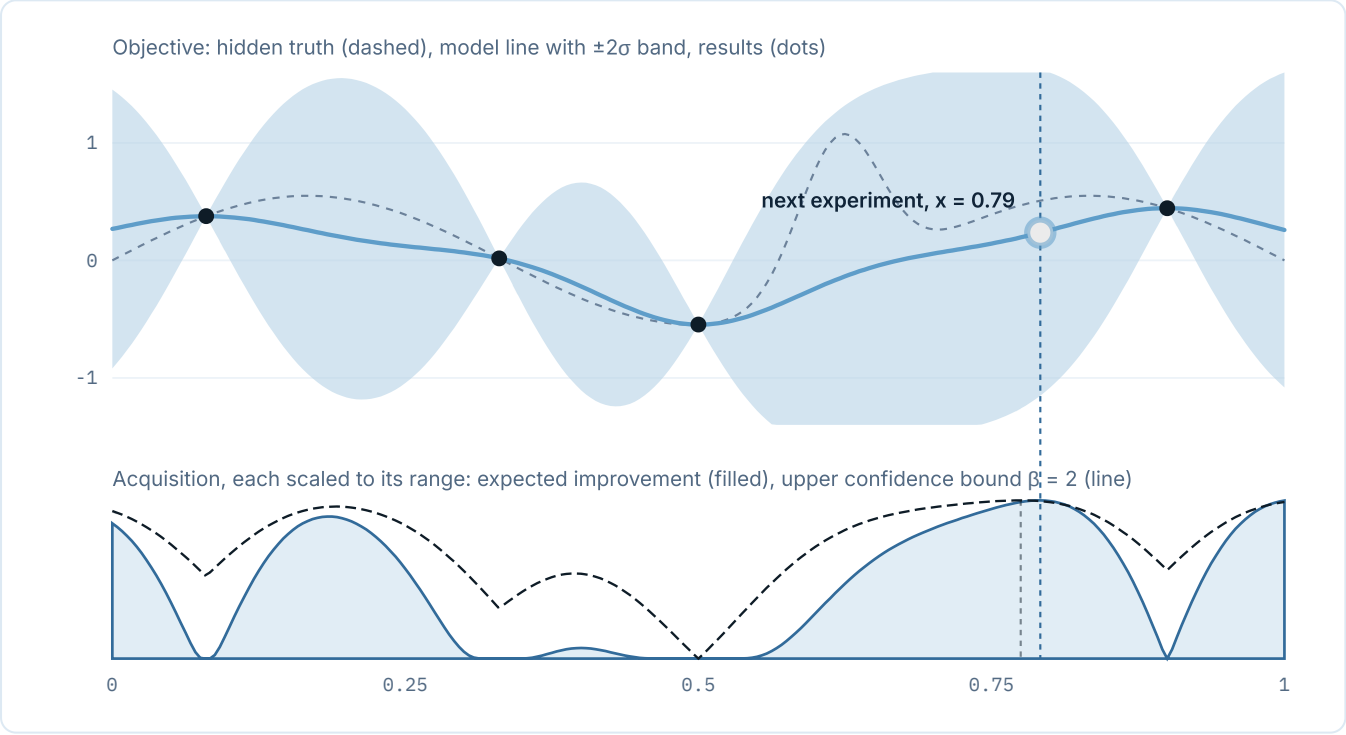
Chapter 4

Optimize (Bayesian optimization)

How Bayesian optimization chooses the next experiment, the engines and acquisition functions, and how to run, check and stop a campaign.

How it chooses the next experiment

Bayesian optimization (BO) decides which experiment to run next. It fits a statistical model, a Gaussian process (GP), to the results you have. For every setting *x* you have not tried, the model gives a prediction *μ(x)* and an uncertainty *σ(x)*. An **acquisition function** turns μ and σ into one score, and the setting with the highest score is the next experiment. You run it, record what you measured, and the model is fitted again.



One round of Bayesian optimization on a toy problem with one parameter. The model is unsure between results (wide band), and both acquisition functions prefer a setting next to the best result where the band is still wide. After the result is recorded, the model is refitted and the next round starts.

Engines and models

Engine	Library	Model	Acquisition functions	Use it for	Cost and plan
Standard BO	`acquilab-bo`, AcquiLAB's own NumPy/SciPy engine	GP. **Improved model 0.2.0** (default for new campaigns): a length scale per parameter, estimated noise, down-weighting of results that disagree with their neighbours (leave-one-out |z| > 3), and the result modelled as measured, as a logarithm or with the scale chosen automatically. **Simpler model (Standard 0.1.0)**: one length scale and the noise you report, to reproduce an earlier campaign exactly.	EI with ξ (default 0.01) over a pool of candidates (default 4096); with outcome limits, EI times the chance the limits hold	One objective, one experiment per round. The only engine with the model map, model check, outcome limits, simulation as a guide and "Why this experiment?"	Free on every plan
Advanced BO	BoTorch 0.15.1 on GPyTorch and PyTorch 2.13 (CPU)	BoTorch Gaussian process	qLogNEI, qLogEI, qUCB (β), LogPI	Batches of 1–16 experiments per round (plates, parallel reactors), noisy data, explicit control of exploration	Free on every plan
Multi-objective (MOBO)	BoTorch	A GP for each objective	qLogEHVI, qLogNEHVI, qEHVI (hypervolume improvement)	2–4 competing goals: yield and cost, sensitivity and selectivity, activity and stability	1 run per suggestion; Pro and Team
Simulation in the loop	`acquilab-bo` driving the simulation service	Improved model, result scale chosen automatically	EI	AcquiLAB chooses and runs simulations itself, on a ready-made model or on the parameters of a model you wrote	Runs for each simulation; Plus and up (in the beta, any account with runs)

Variables can be continuous (lowest, highest, unit, optional step), integer, or categorical (a list of choices such as catalyst or solvent). A campaign can have up to 50 variables. Linear constraints (for example `E1 + E2 + E3 = 3 µM`, or a ratio limit) and forbidden combinations are set in the wizard under *Limits* or later under Edit → *Rules from known chemistry*.

EI, UCB, PI and the rest

All of them answer the same question: is it worth testing *x*, given the prediction μ(x) and the uncertainty σ(x)? They differ in how much they value uncertainty (exploration) against a good prediction (exploitation). Below, f* is the best result so far, Φ and φ are the normal distribution and density, and `z = (μ − f* − ξ)/σ`. Written for maximisation.

Acquisition	Formula	Behaviour	Knob
PI · probability of improvement	`PI = Φ(z)`	The chance of beating f* by at least ξ, regardless of by how much. Greedy: it keeps sampling right next to the best point. In AcquiLAB's benchmarks LogPI ranked last on all three test problems.	ξ: larger asks for a bigger improvement and explores a little more

Acquisition	Formula	Behaviour	Knob
EI · expected improvement	`EI = (μ − f⋆ − ξ)·Φ(z) + σ·φ(z)`	The average size of the improvement. The first term rewards a good prediction, the second rewards uncertainty, so it balances both without tuning. The safe default. “Why this experiment?” splits a suggestion into these two terms (improve and learn).	ξ (default 0.01): raise to 0.05–0.1 to explore more
LogEI / qLogEI	`log EI, computed stably`	The same choice as EI in principle, but EI becomes numerically zero almost everywhere late in a campaign or in many dimensions, and the optimiser loses its way. The log form keeps a usable gradient (Ament et al., NeurIPS 2023). The q-version picks a batch of q points together.	ξ, batch size q
qLogNEI · noisy EI	`E over noise of LogEI against the uncertain best`	Treats the best so far as uncertain too, so one lucky noisy measurement doesn't trap the search. Handles batches and pending experiments. The best default when replicates disagree by more than a few percent.	batch size q
UCB / qUCB · upper confidence bound	`UCB = μ + √β · σ`	Optimism under uncertainty: each point is scored by how good it could plausibly be. You control the balance directly: β ≈ 0.1–0.5 exploits (polishing), β ≈ 1–4 balances, β > 4 explores and maps the space. It does not depend on f*, so it behaves well with noisy or drifting data.	β (AcquiLAB's examples use 2)
EHVI, qEHVI, qLogEHVI, qLogNEHVI	`expected gain in Pareto hypervolume`	For 2–4 objectives. The hypervolume is the region dominated by your best trade-offs above a reference point; a suggestion is worth its expected increase. qLogEHVI is the stable log form, qLogNEHVI also handles noisy measurements, qEHVI is the original form.	reference point, batch size

UCB and EI in one line

EI asks *how much better than my best do I expect this to be?* and needs no tuning. UCB asks *how good could this plausibly be?* and gives you a dial (β) for curiosity. Use EI for a sensible default; use UCB when you know you want to explore early or exploit late, or when noise makes “best so far” unreliable.

Which to use when

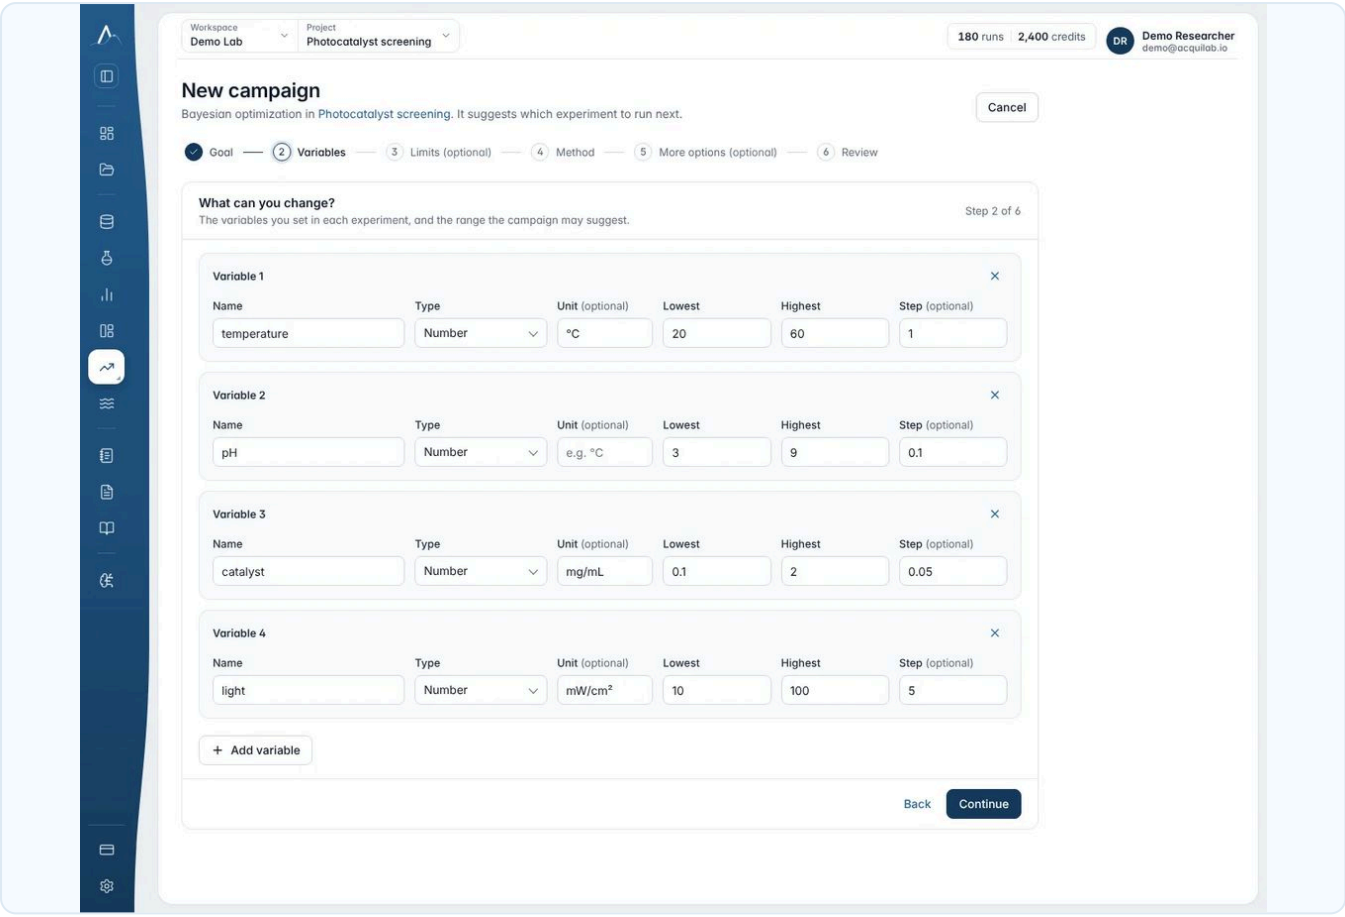
Your situation	Choose	Why
New campaign, one objective, one experiment at a time, results fairly reproducible	**Standard BO · EI · Improved model** (the wizard default)	Robust default, and you get the model map, model check and explanations.
Replicates scatter (biology, sensors, catalysis screens)	**Improved model** (estimated noise and down-weighting) or **Advanced · qLogNEI**	Both stop a single noisy high value from being treated as the true best.
You run 3–8 experiments in parallel (well plate, reactor block)	**Advanced · qLogNEI** (or qUCB) with experiments per round = slots	Batch acquisitions pick points that are good together instead of eight copies of the same idea.
Early stage, the space is unknown, you want a map	**qUCB with β 3–5**, or EI with ξ 0.05–0.1	Rewards uncertainty and spreads the experiments out.
Late stage, near the optimum, few runs left	**qUCB β 0.5–1** or **qLogEI**	Concentrates on the predicted best region.
Many parameters (6 or more) or a long campaign where EI goes flat	**qLogEI / qLogNEI**	The log form keeps the optimiser moving where EI underflows to zero.

Your situation	Choose	Why
A response that behaves like a product (rates, yields, activities)	**Improved model** with the result scale set to **Choose automatically** or **Logarithm**	Modelling the logarithm can turn an unreliable model into a reliable one; the model check suggests it when it helps.
A secondary result must meet a spec (purity ≥ 95 %, cost ≤ X)	**Improved model + outcome limits**	Models each limit and multiplies EI by the chance that all limits hold.
Two to four goals that conflict	**Multi-objective · qLogNEHVI** (noisy) or **qLogEHVI** (clean)	Finds the Pareto front; then set priorities (Ignore … Top) to rank it.
Mixture or ratio with a fixed total	Limits → **These add up to a total**	Every suggestion is a complete mixture; one part is computed as the rest.
A concentration, loading or dose over several decades	**Log scale** for that variable (the Setup check suggests it)	Low and high values get equal attention instead of most suggestions landing near the top.
Catalyst, solvent, buffer type	**Categorical** variables	No invented numeric order between choices.
You also have a simulation of the same system	Improved model with **Use simulation results as a guide**	Simulations enter as a cheaper source with their own length scale; only lab results count as the best.
The simulation itself is the experiment	**Simulation in the BO loop**	AcquiLAB runs the model inside the loop; no lab work.
You want the optimum fast and don't care about the model	Avoid **LogPI** except for final polishing	Too greedy; it stalls in local optima.

Running a campaign

Optimize → *Bayesian Optimization* opens the campaign designer. It has six steps, and the defaults suit most campaigns.

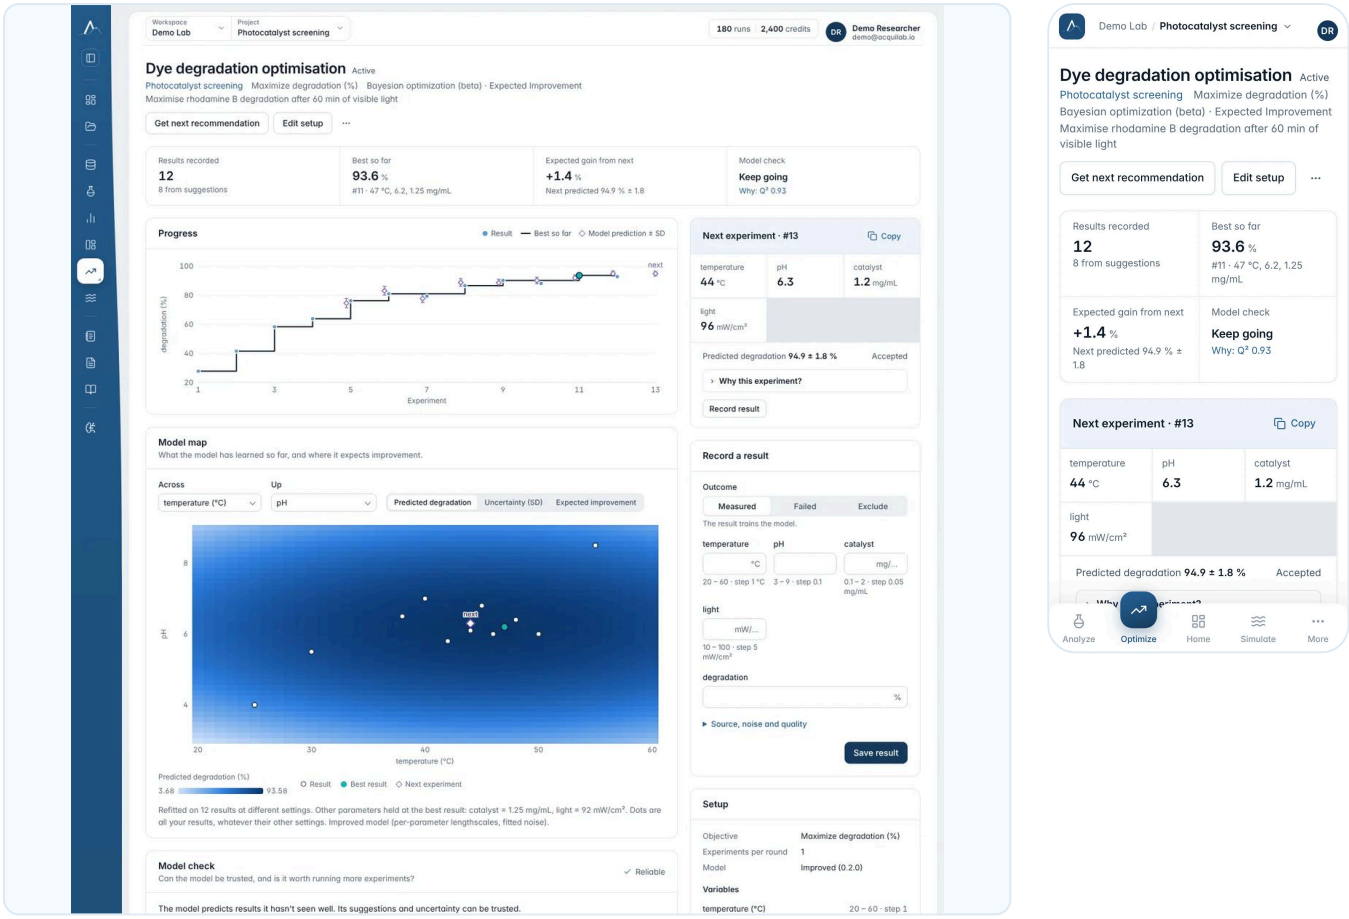
Step	What you set
Goal	Campaign name, the result you measure (objective name and unit), maximise or minimise, a description.
Variables	What you can change: name, type (continuous, integer, categorical), unit, lowest and highest value or the choices, and the **scale** of each number (Linear or Log). A log scale suggestion appears under the variable it is about.
Limits (optional)	Parts of a mixture (*These add up to a total*), sums and ratios such as total volume ≤ 10 mL, and combinations that must never be suggested; leave empty if every setting in the ranges is safe.
Method	Engine and acquisition function, experiments per round (Advanced BO: 1–16), the model (Improved, recommended, or Standard), outcome limits, and *Use simulation results as a guide*.
More options (optional)	Random seed, starting experiments, candidates scored per suggestion, exploration ξ, noise floor, largest grid.
Review	The **Setup check**, then create. The method and model can change later; variables are fixed once results exist.



The campaign designer, Variables step: what you can change and the range the campaign may suggest.

On the campaign page

1. Start from what you know: *Import past results* takes a CSV, a pasted table or an earlier campaign. Otherwise run the starting experiments the campaign spreads across the ranges.
2. Press **Get next recommendation**, run the experiment in the lab, and record the result (or mark it failed or excluded). A typo can be fixed once with *Correct*.
3. Watch the **model check** and the **model map** (prediction, uncertainty and EI over one or two variables). *Why this experiment?* explains a suggestion.
4. Pause, edit ranges or rules, and resume at any time. *Add to document* puts the campaign's best result and recent results into a report.



A single-objective campaign. The model check, the suggestion history and the results table follow further down the page.

Setting up without BO experience

The campaign designer checks your setup while you fill it in and says in plain words what to change. Each piece of advice has **Accept**, which changes the setup for you, and **Keep as is**. The check only recommends: *Create* always works, whatever it says.

Scale per variable

Every number variable has **Scale: Linear | Log**. Linear spreads the search evenly over the range. Log spreads it evenly over each factor of ten, so 0.001–0.01 mg/mL gets as much attention as 1–10 mg/mL. With a linear 0.001–10 range, nine in ten suggestions would land above 1.

- Log needs a lowest value above 0, and 1 or more for whole numbers. Choices have no scale.
- Log is **suggested** for a number that spans 100× or more when its name or unit marks it as a quantity that behaves on a log scale: a concentration, loading, dose, amount, flow, pressure, time, molecular weight, enzyme activity, light intensity or frequency (M, mM, µM, mg/mL, mol%, wt%, mL/min, bar, min, U/mL and similar). Any other range of 1,000× or more gets a softer suggestion: most suggestions would land near the top of a linear range that wide.
- It is never suggested for temperature, pH, pKa, potential, angles or the parts of a mixture. If you choose log there, the check says why linear fits better.
- *Keep linear* hides the suggestion for this setup. Values, steps and limits always stay in your units; on a log variable the model map draws a log axis (0.001, 0.01, 0.1 …).

Mixtures: parts that add up to a total

In *Limits*, **These add up to a total** turns two or more number variables with the same unit into a mixture, such as TiO2 + g-C3N4 = 100 wt%. Choose the total (100 for a percentage) and which part is **the rest** (the last one unless you pick another).

- The rest is computed: total minus the other parts. Suggestions show it marked *computed*, the record form fills it in as you type the other parts, and the results table shows it.
- The lowest values must add up to no more than the total, and the highest values to at least the total. The rest has no step and stays linear. A variable belongs to one mixture; a campaign has up to five.
- The check suggests a mixture when variables in %, wt%, vol% or as fractions can add up to 100 % (or 1). Accept adds it.

The Setup check

The *Review* step lists what looks fine (a tick per variable) and what to consider. In Edit the same check is a one-line summary, for example “Setup check: 2 suggestions”, that opens on a click.

Advice	When
Use a log scale	As above. Accept switches the variable to Log.
These look like parts of a mixture	Parts in %, wt%, vol% or fractions that can add up to the total.
Few experiments for this many settings	Enter *Experiments you plan to run* (used only by the check). Fewer than 4 per variable (6 for multi-objective) means the first suggestions mostly explore; fix the settings you already know.
Two settings tied together	Amount, concentration and volume of the same thing: optimise two of them, not all three.
Many choices	A choice with more than 8 options; start with the most promising ones.
Only a few whole numbers	A whole number with 3 or fewer values; Accept makes it a choice.
A step on a log scale	A fixed step removes the benefit of the log scale at the low end; Accept removes the step.
Outcome limits	A limit such as purity ≥ 95 % needs that value recorded with each result.

> **When the Assistant or an agent builds the setup**
>
> The Assistant and connected agents run the same check. When they build a setup, they switch on a log scale only for a recognised quantity (a concentration, loading, dose, flow, pressure and the others above) that spans 100× or more, and tell you which, for example “Log scale switched on for E1 (it spans 200×)”. A range that is merely very wide stays a suggestion for you to decide, as do mixtures and the other advice. In the designer you decide every one yourself.

In *Edit*, changing a scale or adding a mixture makes a new revision; recorded results are kept, and a mixture can be added only when the recorded parts already add up. Simulation loops (*Simulate → Optimize*) have the same scale control under each search range.

Model check and switching the model

The model check tells you whether to trust the suggestions. It refits the model leaving each result out in turn and reports Q² (how well left-out results are predicted), how many results fall inside the 95 % band, and a verdict: **reliable**, **caution** or **unreliable**, with a *keep going?* hint and *converged* when the suggestions stop improving.

A better model, one click away

When another model predicts your results clearly better (Q² higher by at least 0.1), the check names it, for example the Improved model or the Improved model on a log scale, and offers **Switch model…**. It opens Edit with those options set; save to apply.

Scale of the result

With the Improved model, Edit → *Scale of the result*: *As measured*, *Logarithm* (results must all be positive) or *Choose automatically from the results*. Automatic uses the log when every result is positive, there are at least five distinct settings and the log model predicts better. The check then says "log scale (chosen automatically)".

Outcome limits

Other things you measure that must stay within a limit, such as purity ≥ 95 %. The model learns each one and suggests settings likely to meet them all. Adding a limit switches to the Improved model.

Simulation as a guide

Use simulation results as a guide: cheap simulation runs from a Simulate loop on this campaign show where to look. The model learns how far they are from the lab, and suggestions aim at the best lab result.

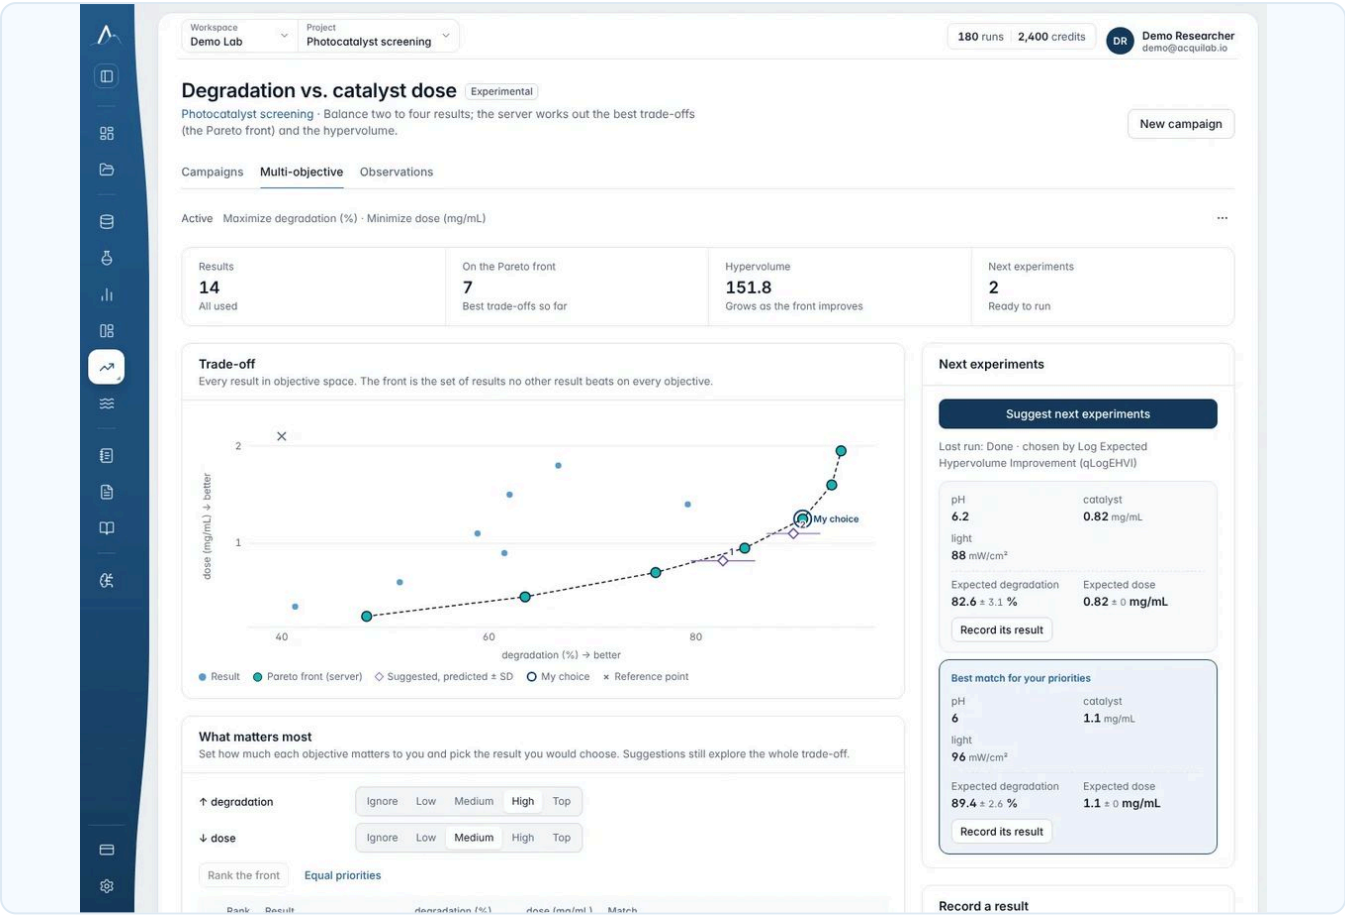
New campaigns, simulation loops and campaigns created by agents use the Improved model. Campaigns created earlier keep the model they were created with until you switch.

Multi-objective campaigns

Optimize → *Multi-objective* (Pro and Team) handles two to four objectives that conflict. Instead of one best point it finds the **Pareto front**: the settings where no objective can improve without another getting worse.

In the app	Method	When
Push the trade-off (recommended)	qLogEHVI	Picks experiments that should improve the best trade-offs the most. 1–4 per round.
For noisy results	qLogNEHVI	Treats recorded values as uncertain, so one lucky reading doesn't steer it. Slower. 1–4 per round.
Original method	qEHVI	Kept for comparison; the recommended method is usually better. Experimental.

Set how much each objective matters (Ignore, Low, Medium, High, Top) to rank the front and pick a candidate. Each suggestion uses 1 run.



A multi-objective campaign and its Pareto front.

Campaigns through the MCP connector

An AI agent connected through the MCP connector creates and runs campaigns with the same engines:

```
create_optimization(
  name="Cu-catalysed coupling",
  parameters=[
    {"name":"temperature","minimum":40,"maximum":110,"unit":"°C"},
    {"name":"catalyst_loading","minimum":0.5,"maximum":5,"unit":"mol%"},
    {"name":"solvent","kind":"categorical",
     "choices":["DMF","DMSO","MeCN"]}],
  objective={"name":"yield","direction":"maximize","unit":"%"},
  batch_size=1)
check_optimization_setup(...)  # free: log scales, mixtures, budget
add_results(...)               # what you measured
suggest_next_experiments()  # free for one objective
explain_recommendation()    # improve vs learn
```

- Two to four objectives in `objectives=[…]` make a multi-objective campaign.
- `batch_size` greater than 1 uses Advanced BO batches.
- A single-objective campaign created this way uses the Improved model, as in the designer.

- `check_optimization_setup` reviews a setup first (free, nothing is created) and returns the parameters with log scales switched on for recognised quantities spanning 100× or more; other advice is returned for the agent to show you. `create_optimization` takes `scale` per parameter and `composition` for parts that add up to a total.
- Agents record only measured values, never predictions.

When to stop

- The model check says *converged*: recent suggestions no longer improve the predicted best.
- You reached the target you needed, or the experiment budget is spent.
- The model check stays *unreliable* after switching model: the variables may not explain the result. Revisit ranges, add a variable, or check the measurement.

A converged campaign has found a good region under its model; it does not prove a global optimum.

Chapter 5

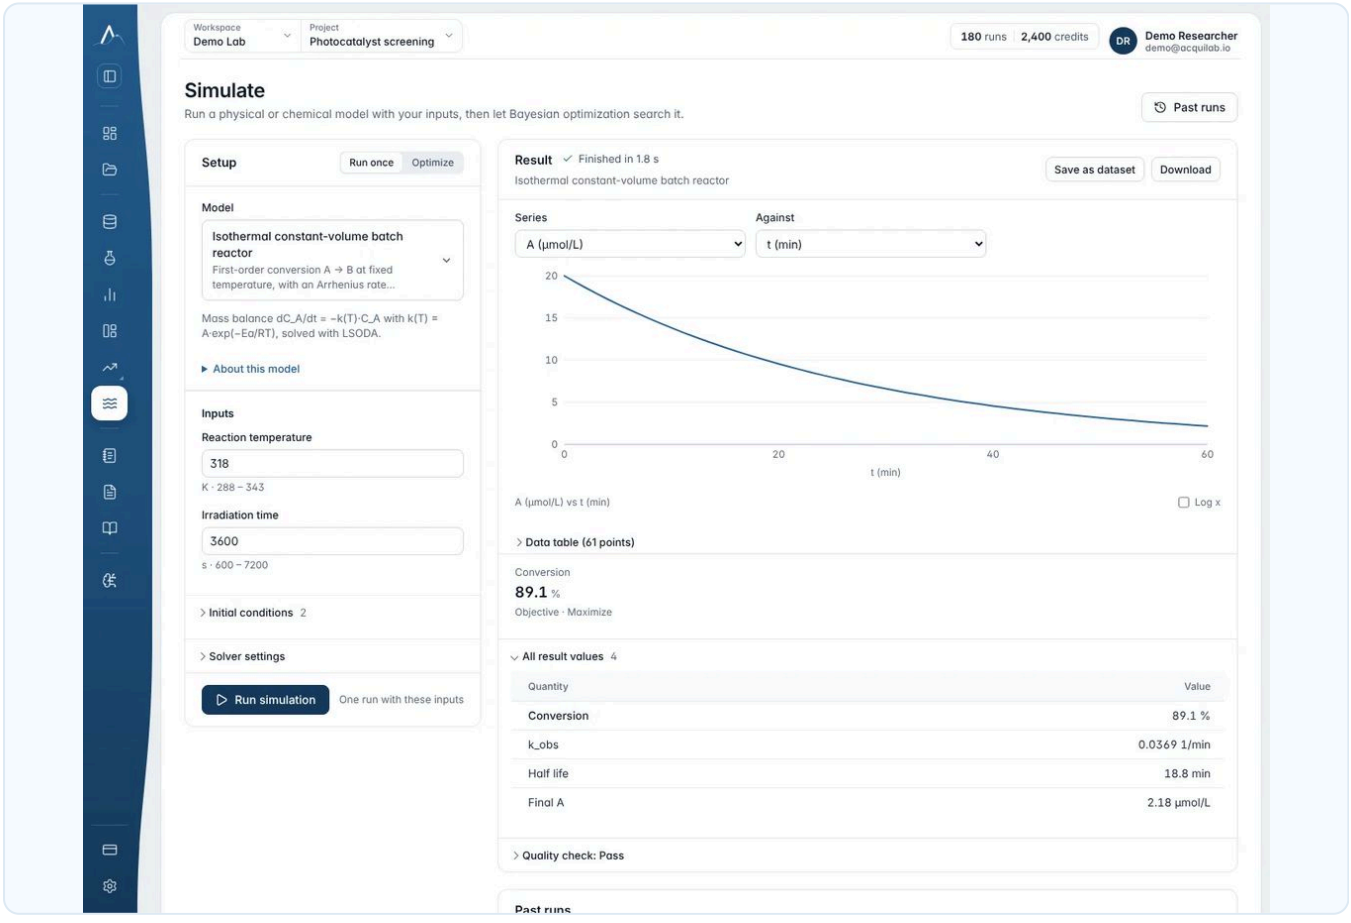
Simulate

Ready-made physics and chemistry models, models you write as rate equations or formulas, fitting to data and simulation in the BO loop.

Running a model

Simulate runs validated physics and chemistry models in an isolated worker with bounded time and memory, and records the model version and inputs with every result.

- In *Setup*, choose the model and change its inputs; units and allowed ranges are shown next to each. The model's description lists its assumptions, validity, known failure modes and citations.
- *What to do*: **Run once**, or **Optimize** to let Bayesian optimization run the model repeatedly (see Simulation in the BO loop).
- The result shows the values, curves and a quality check, with *Download* (CSV or JSON), *Save as dataset* and *Open in Plot Studio*.
- *Past runs* groups runs that belong together (compared sets, optimization loops, consecutive runs of one model). Reopen a run, compare 2–8 runs, or compare 2–10 runs with measured data on one graph.



Simulate: a batch reactor run, its result values and the past runs.

Packages and limits

Package	Models	What you can compute	Example inputs	Time limit
SciPy / NumPy (AcquiLAB model engine)	custom_ode, custom_algebraic, batch, CSTR, PFR, diffusion, heat, reaction–diffusion, membrane, electrode diffusion, Cottrell, Butler–Volmer, CV, EIS, isotherm, breakthrough	Rate equations over time (LSODA, switches between stiff and non-stiff methods), algebraic formulas, 1-D transport, reactor balances, electrode kinetics, impedance, adsorption	An enzyme cascade to maximum H_2O_2; CV at 10–1000 mV/s; a Nyquist plot for k^o = 0.01 cm/s	seconds to about 1.5 min
PySCF 2.14 (density functional theory)	`dft_molecule`	Kohn–Sham DFT with the PBEh(a) hybrid (E_{xc} = a·E_x^{HF} + (1 – a)·E_x^{PBE} + E_c^{PBE}), a single point on the GFN2-xTB structure, SCF to 10^{-9} E_h: HOMO and LUMO energies, the HOMO–LUMO gap (objective), orbital energies, total energy; implicit solvent (ddCOSMO), ions and radicals	SMILES up to **40 atoms**; basis def2-SVP (default), 6-31G*, STO-3G or def2-TZVP; exact exchange a = 0 (PBE) … 0.25 (PBE0, default) … 1; charge; unpaired electrons; solvent none, water, methanol, acetonitrile, DMSO or toluene	up to 30 min
tblite 0.7 (GFN-xTB) and RDKit 2026.3	`xtb_molecule`	SMILES → 3D (RDKit ETKDG) → GFN2- or GFN1-xTB relaxed geometry, orbital energies, dipole, harmonic frequencies, the IR spectrum with intensities, thermochemistry G, H and S at T (RRHO, 1 atm)	SMILES up to **100 atoms** of H, B, C, N, O, F, Si, P, S, Cl, Br, I; GFN2 or GFN1; ALPB solvent none, water, methanol, acetonitrile, DMSO, acetone, THF, toluene, hexane, CH_2Cl_2, $CHCl_3$; T 100–1500 K; symmetry number; IR band width	up to 30 min
OpenMM 8.6	`water_md`	NPT molecular dynamics of a water box with PME: density, potential energy per molecule, ΔH_{vap}, the O–O radial distribution g(r), self-diffusion coefficient	T 250–400 K; pressure; molecules; equilibration and production (ps); water model TIP3P, SPC/E, TIP4P-Ew, TIP3P-FB or TIP4P-FB	up to 30 min
Cantera 3.2	`combustion_equilibrium`, `ignition_delay`	Adiabatic flame temperature and equilibrium products (Gibbs minimisation at constant H and P) with an equivalence-ratio sweep; ignition delay in a constant-pressure reactor with detailed chemistry	φ 0.3–3; T_0; pressure 0.5–60 atm; mechanism and fuel (validated for methane–air and H_2–air)	up to 2 min

Quantum chemistry and molecular dynamics run on a separate chemistry worker with 4 GB of memory. A DFT or water MD run uses 5 runs, an xTB run 2, every other model 1 (see Runs).

Worked chemistry examples

Question	Model and settings	What to read
HOMO–LUMO gap of phenol in water	dft_molecule · SMILES `Oc1ccccc1` · def2-SVP · solvent water · a = 0.25 (PBE0)	HOMO, LUMO and gap (eV); compare with a gas-phase run to see the solvent shift

Question	Model and settings	What to read
The gap shrinks with conjugation	dft_molecule · benzene `c1ccccc1`, naphthalene `c1ccc2ccccc2c1`, anthracene `c1ccc2cc3ccccc3cc2c1` · same basis and a	Gap trend across the series (compare runs)
How much does the functional matter?	dft_molecule · one molecule · a = 0, 0.25, 0.5	The gap grows with exact exchange; a = 0 is PBE, 0.25 is PBE0
Push–pull dye, donor and acceptor levels	dft_molecule · 4-nitroaniline `Nc1ccc(cc1)[N+](=O)[O-]` · def2-SVP · water	HOMO (donor side) and LUMO (acceptor side) energies
A stable radical	dft_molecule or xtb_molecule · TEMPO `CC1(C)CCCC(C)(C)N1[O]` · unpaired electrons 1	Orbital energies of the open-shell molecule
Which tautomer is favoured in water?	xtb_molecule · 2-hydroxypyridine `Oc1ccccn1` against 2-pyridone `O=c1cccc[nH]1` · GFN2 · water, then none	ΔG between the two (kJ/mol), and whether the solvent changes the order
A large dye cation	xtb_molecule · rhodamine B (64 atoms) · charge 1 · GFN2	Relaxed structure, orbitals and the IR spectrum; a few minutes, well inside the 30-minute limit
Predict an IR spectrum	xtb_molecule · ethanol `CCO` (default) or CO_2 `O=C=O`	IR spectrum; CO_2 shows the IR-silent symmetric stretch
Water's density maximum	water_md · TIP4P-Ew · T 260–320 K	Density against T; g(r) peak near 0.28 nm; D ≈ $5{\cdot}10^{-9}$ m²/s for TIP3P at 298 K
Hottest flame	combustion_equilibrium · methane–air · φ sweep 0.6–1.6	T_{ad} peaks slightly rich of φ = 1 (about 2226 K at φ = 1 for methane–air)
Ignition delay map	ignition_delay · T_0 900–1400 K · P 1, 10, 40 atm · φ 1	Arrhenius-like fall of the delay with T_0, shorter at higher P

Read DFT numbers correctly

A Kohn–Sham HOMO–LUMO gap is neither the optical gap nor the fundamental gap: use it for trends between related molecules at the same settings. DFT is a single point on the xTB geometry (no DFT relaxation) and handles up to 40 atoms; the SCF can fail to converge for difficult open-shell systems, which is reported as a failed run. One conformer is computed, so flexible molecules may land in a local minimum. Bayesian optimization cannot choose a molecule for you: the design inputs are numbers such as temperature or exact exchange, not SMILES. To screen molecules, run one simulation per SMILES and compare.

Ready-made models

23 models. The design inputs are what Bayesian optimization searches when you choose *Optimize*; the objective is fixed by the model.

Model	Physics	Design inputs	Outputs	Objective	Example question
Reactors					
`batch_isothermal`	dC/dt = ν·r(C), mass-action network at fixed T (Arrhenius rate constants)	temperature 200–600 K, final time	concentration profile, conversion, yield, selectivity	max yield	Best T and time for the intermediate in A → B → C

Model	Physics	Design inputs	Outputs	Objective	Example question
`cstr_steady`	$0 = (C_{in} - C)/\tau + \nu \cdot r(C)$	temperature, τ	steady concentrations, conversion, yield, selectivity	max conversion	Residence time for 90 % conversion
`pfr_ideal`	$dC/d\tau = \nu \cdot r(C)$, constant density	temperature, τ	exit concentrations, conversion, yield, selectivity	max conversion	PFR against CSTR for the same network
Electrochemistry					
`cottrell`	$I = nFAC\sqrt{(D/\pi t)}$	area, D, C_bulk	I(t), j(t)	max current	Expected step current for a 3 mm disk
`butler_volmer`	$i = i_0[\exp(\alpha_a nF\eta/RT) - \exp(-\alpha_c nF\eta/RT)]$	i_0, η, temperature, area	j, I, Tafel branches	max current	Overpotential needed for 10 mA/cm²
`electrode_diffusion_1d`	Planar diffusion to an electrode with a bulk boundary	area, D, C_bulk	C(x,t), I(t), flux	max current	When does Cottrell break down in a thin cell?
`cyclic_voltammetry`	Diffusion of O and R with Butler–Volmer flux, linear sweep	scan rate 0.001–10 V/s, k^o, C_bulk	I against E, i_{pc}, ΔE_p, i_{pa}/i_{pc}	max cathodic peak	ΔEp against scan rate to judge reversibility
`eis_randles`	$R_s + (C_{dl} \parallel (R_{ct} + \text{Warburg}))$	R_s, k^o, C_dl	Nyquist, Bode, $\lvert Z \rvert$, phase, R_{ct}	min low-frequency impedance	Semicircle size for a faster electrode
Transport and heat					
`diffusion_1d_steady`	$0 = D\, d^2C/dx^2$	D	C(x), boundary fluxes	max flux	Steady flux through a gel layer
`diffusion_1d_transient`	$\partial C/\partial t = D\, \partial^2 C/\partial x^2$	D	C(x,t), fluxes	max flux	Time to reach a sensor surface
`reaction_diffusion_1d`	$\partial C_i/\partial t = \nabla \cdot (D_i \nabla C_i) + \omega_i(C)$	D, k	C(x,t), boundary flux	max uptake flux	Enzyme layer thickness against substrate penetration
`membrane_1d`	Fickian slab with partition K at both faces	D, C_upstream, C_downstream	C(x), flux, permeability	max flux	Permeability of a coating
`heat_conduction_1d_steady`	$d/dx(k\, dT/dx) = 0$	k	T(x), heat flux	max heat flux	Insulation thickness
`heat_conduction_1d_transient`	$\rho c_p \partial T/\partial t = \partial/\partial x(k\, \partial T/\partial x)$	k (plus ρ, c_p, target T)	T(x,t), centre T, time to temperature	max centre T	Time for a sample to reach 60 °C
Separations					
`adsorption_isotherm`	Sips (Langmuir–Freundlich) with van 't Hoff temperature dependence; n = 1 is Langmuir	temperature 200–600 K, C_feed	q(C), θ, q at feed	max loading	Loading at 25 against 60 °C
`fixed_bed_breakthrough`	Axially dispersed plug flow, LDF uptake and a Sips isotherm	u 1e-5–0.5 m/s, L 0.01–10 m	$C_{out}(t)$, C/C_0, breakthrough time, productivity	max productivity	Column length and velocity for the best productivity

Model	Physics	Design inputs	Outputs	Objective	Example question
Combustion and molecules					
`combustion_equilibrium`	HP equilibrium (Cantera)	φ 0.3–3, T_0 250–1200 K	$T_{ad}(\varphi)$, products	max flame T	Flame temperature against mixture
`ignition_delay`	Constant-pressure adiabatic reactor (Cantera, CVODES)	T_0 800–2000 K, P 0.5–60 atm, φ	T(t), delay	min delay	Delay at engine-like conditions
`xtb_molecule`	GFN-xTB (tblite) on an RDKit structure	temperature 100–1500 K	IR spectrum, G, orbitals, dipole	min G	Tautomer or conformer energy
`dft_molecule`	Kohn–Sham DFT, PBEh(a) (PySCF)	exact exchange 0–1	orbital energies, gap	max gap	HOMO–LUMO gap of a dye
`water_md`	NPT molecular dynamics with PME (OpenMM)	temperature 250–400 K	$g_{OO}(r)$, density	max density	Temperature of maximum density
Models you write					
`custom_ode`	dy/dt = f(y, t, p), whitelisted expressions, SciPy LSODA	your parameters (1–6 at a time)	every state's curve and one reported value	max or min of the report	See the ODE library below
`custom_algebraic`	Outputs as formulas of the parameters	your parameters	outputs and one reported value	max or min of the report	See the formula library below

Fit to measured data

Fitting works for batch reaction kinetics, EIS Randles, the adsorption isotherm, fixed-bed breakthrough and `custom_ode`: choose 1–6 parameters with bounds and a measured column in the project (in the same units as the model), optionally a σ column. The fit uses 10–400 model evaluations and counts as one simulation. It reports the values with intervals, and the fitted curve can be saved as a dataset.

Writing your own model

When no ready-made model fits (an enzyme cascade, a reaction network, a dose–response), write the rate equations yourself as `custom_ode`, or plain formulas as `custom_algebraic`. AcquiLAB evaluates them with a whitelist of operations, never arbitrary code, and integrates them with SciPy's LSODA solver.

The system, part by part

- **states**: 1–8 quantities that change, each `{name, initial, unit}`.
- **rates**: one `d(state)/dt` expression per state.
- **parameters**: up to 24 constants `{name, value, unit, source}`; the source is the paper, or `"assumed"`.
- **report**: the headline number, an expression evaluated at the end of the run. Optimization maximises or minimises it.
- **time**: `{end, unit: s | min | h, points}`. Every state's curve is kept.

Rules the checker enforces

- Allowed: numbers, names, `+ - * /`, `**` with a plain positive number up to 8, `sin cos exp log log10 sqrt abs min max tanh sinh cosh`, `pi`, `e`.
- **Rates are per second**, and `t` is in seconds, whatever unit you choose for the run time.
- Reserved names: `y`, `t`, `k` (ODE), `x` (formulas), function names, `pi`, `e`. Use `k1`, `kf`, `kcat` ...
- Exponents must be plain numbers: `S**nH` is refused. Write `exp(nH*(log(K) - log(max(S, 1e-12))))` instead.
- Guard `log` or `sqrt` of a changing quantity: `log(max(A, 1e-12))`, `sqrt(max(A, 0))`.
- No if/else, comparisons or events. For a switch use a smooth step: `0.5*(1 + tanh((t0 - t)/w))`.
- Units are not converted. Keep them consistent yourself (mM with mM/s, and so on).

Full example: the creatinine sensing cascade (CA → CI → SOX)

The model behind the 4 October 2026 run in the Creatinine sensor project, with the literature constants used there. It passed AcquiLAB's checker together with the optimization settings below (total enzyme fixed at 3 µM, 20 runs). That run found an optimum near 1 : 1.3 : 1.7 (CA : CI : SOX), 0.0196 mM H_2O_2 at 5 min against 0.0186 mM for 1 : 1 : 1, with sarcosine oxidase as the rate-limiting step.

```
{
  "states": [
    {"name":"Cn",  "initial":0.1,  "unit":"mM"},   // creatinine
    {"name":"Cr",  "initial":0,    "unit":"mM"},   // creatine
    {"name":"Sar", "initial":0,    "unit":"mM"},   // sarcosine
    {"name":"O2",  "initial":0.25, "unit":"mM"},
    {"name":"H2O2","initial":0,    "unit":"mM"}],
  "parameters": [
    {"name":"kcat1","value":252, "unit":"1/s","source":"BRENDA, P. putida creatininase"},
    {"name":"Km1",  "value":26,  "unit":"mM", "source":"J Biochem 1979"},
    {"name":"Kcr",  "value":130, "unit":"mM", "source":"J Biochem 1979"},
    {"name":"Keq",  "value":1.3, "unit":"1",  "source":"assumed"},
    {"name":"kcat2","value":10.5,"unit":"1/s","source":"derived from 14 U/mg"},
    {"name":"Km2",  "value":1.33,"unit":"mM", "source":"BRENDA"},
    {"name":"kcat3","value":14,  "unit":"1/s","source":"J Appl Biol Sci 33:161"},
    {"name":"Km3",  "value":3.6, "unit":"mM", "source":"J Appl Biol Sci 33:161"},
    {"name":"KmO2", "value":0.1, "unit":"mM", "source":"assumed"},
    {"name":"E1","value":0.001,"unit":"mM","source":"assumed"},
    {"name":"E2","value":0.001,"unit":"mM","source":"assumed"},
    {"name":"E3","value":0.001,"unit":"mM","source":"assumed"}],
  "rates": {
    "Cn":  "-kcat1*E1*(Cn - Cr/Keq)/Km1/(1 + Cn/Km1 + Cr/Kcr)",
    "Cr":  "kcat1*E1*(Cn - Cr/Keq)/Km1/(1 + Cn/Km1 + Cr/Kcr) - kcat2*E2*Cr/(Km2 + Cr)",
    "Sar": "kcat2*E2*Cr/(Km2 + Cr) - kcat3*E3*Sar/(Km3 + Sar)*O2/(KmO2 + O2)",
    "O2":  "-kcat3*E3*Sar/(Km3 + Sar)*O2/(KmO2 + O2)",
    "H2O2":"kcat3*E3*Sar/(Km3 + Sar)*O2/(KmO2 + O2)"},
  "report": {"name":"H2O2 at 5 min","expression":"H2O2","unit":"mM"},
  "time":   {"end":5,"unit":"min","points":301}
}

// optimize the enzyme split at a fixed total of 3 µM
optimize    = [{"name":"E1","lower":0.0002,"upper":0.0026},
               {"name":"E2","lower":0.0002,"upper":0.0026},
               {"name":"E3","lower":0.0002,"upper":0.0026}]
constraints = [{"sum":["E1","E2","E3"],"equals":0.003}]
goal = "maximize", max_simulations = 20
```

Comments (`//`) are for reading only; leave them out of the real JSON.

Run once

Simulate → custom ODE, or `run_written_simulation`. Gives every state's curve and the reported value.

Compare 2–6 variants

The same system with some values changed, started together: `variants=[{"label":"1:2:1","values":{"E2":0.0015}}]`.

Optimize 1–6 parameters

Ranges, up to 4 constraints (`sum … equals / at_most / at_least` , or coefficients with eq, le or ge), 6–40 simulations.

Fit to measured data

1–6 parameters with bounds, a measured column in the project, in the same units as the system.

What a custom ODE cannot do

Only well-mixed systems over time: no spatial gradients (use the 1-D transport models), no random effects, no scheduled events such as repeated doses (approximate them with smooth `tanh` steps). At most 8 states and 24 parameters, and about 60 terms per expression. Larger systems, dosing schedules, stochastic simulation and multi-objective optimization of written models are planned.

Custom ODE library

Every row was written in AcquiLAB's format and passed its checker. Copy a rate, give the parameters values and sources, and you have a runnable model. Up to 8 states fit in one system, so rows can be combined (for example a cascade plus enzyme inactivation).

System	States	Rate expressions (AcquiLAB syntax)	Typical report	Question to ask, or optimization target
Enzymes and biosensors				
Michaelis–Menten	S, P	`S: -kcat*E*S/(Km + S)` `P: kcat*E*S/(Km + S)`	P at end	Enzyme amount for 90 % conversion in 10 min
Competitive inhibition	S, P	`S: -kcat*E*S/(Km*(1 + I0/Ki) + S)`	P at end	Compare 0, 50 and 200 µM inhibitor as variants
Mixed / uncompetitive inhibition	S	`S: -Vmax*S/(Km*(1 + I0/Kic) + S*(1 + I0/Kiu))`	S left	Which inhibition type explains your progress curve (fit)
Substrate inhibition (Haldane)	S, P	`S: -Vmax*S/(Km + S + S**2/Ksi)`	P at end	Starting substrate that maximises product
Product inhibition	S, P	`S: -Vmax*S/(Km*(1 + P/Kp) + S)`	P at end	When product removal pays off
Cooperative (Hill) kinetics	S, P	`S: -Vmax/(1 + exp(nH*(log(K05) - log(max(S, 1e-12)))))`	P at end	Effect of nH on a switch-like response
Enzyme inactivation during use	E, S	`E: -kd*E` `S: -kcat*E*S/(Km + S)`	S left	Sensor drift; operational half-life ln2/kd

Two-substrate oxidase with HRP readout (GOx)	Glc, O2, H2O2, Res	`Glc: -kcat*Eg*Glc*O2/(KmO*Glc + KmG*O2 + Glc*O2)`
		`H2O2: (same) - kHRP*H2O2`
		`Res: Iapp*kHRP*H2O2`
			resorufin formed	Is O_2 limiting at high glucose?

Multi-enzyme cascade (creatinine CA–CI–SOX)	Cn, Cr, Sar, O2, H2O2	`reversible Haldane step + two MM steps + O₂ term (full JSON above)`	H_2O_2 at readout	Best enzyme ratio at a fixed total (constraint sum = total)
Surface adsorption on an electrode (Langmuir kinetics)	th	`th: kads*Cb*(1 - th) - kdes*th`	coverage	Incubation time to 95 % coverage
Ligand–receptor binding (SPR-type)	L, Rf, LR	`LR: kon*L*Rf - koff*LR` `L, Rf: -kon*L*Rf + koff*LR`	LR/(LR + Rf)	K_D = koff/kon; association time

System	States	Rate expressions (AcquiLAB syntax)	Typical report	Question to ask, or optimization target
Chemical kinetics				
Consecutive A → B → C with Arrhenius	A, B, C	`A: -A1*exp(-Ea1/(R*T))*A` `B: A1*exp(-Ea1/(R*T))*A - A2*exp(-Ea2/(R*T))*B`	yield of B	T and time that maximise the intermediate
Reversible A ⇌ B	A, B	`A: -kf*A + kr*B`	B at end	Approach to equilibrium, K = kf/kr
Second order A + B → P	A, B	`A: -k2r*A*B` `B: -k2r*A*B`	conversion	Effect of excess B
Autocatalysis A + B → 2B	A, B	`A: -kau*A*B` `B: kau*A*B`	B at end	Induction time against seed amount
Linear temperature ramp	A	`A: -A1*exp(-Ea/(R*(T0 + ramp*t)))*A`	A left	Ramp rate (K/s) for full conversion
Photodegradation with light absorption	C	`C: -phi*Iabs*(1 - exp(-2.303*eps*lp*C))`	C left	Irradiation time; inner-filter effect
Free-radical polymerisation (QSSA)	In, Mo	`In: -kdec*In` `Mo: -kp*Mo*sqrt(2*fe*kdec/kt)*sqrt(max(In, 0))`	1 − Mo/M0	Initiator amount for 90 % conversion
Oscillating reaction (Brusselator)	X, Y	`X: a - (b + 1)*X + X**2*Y` `Y: b*X - X**2*Y`	X at end	Oscillations appear when b > 1 + a²
Bioprocess, ecology, epidemiology				

Monod fermentation	X, S, P	`X: mumax*S/(Ks + S)*X`
		`S: -mumax*S/(Ks + S)*X/Yxs`
		`P: Ypx*mumax*S/(Ks + S)*X`
			product titre	Feed substrate and inoculum for the highest titre

Logistic growth with Luedeking–Piret product	X, P	`X: mu*X*(1 - X/Xmax)` `P: aLP*mu*X*(1 - X/Xmax) + bLP*X`	titre	Growth- against non-growth-associated production
Predator–prey (Lotka–Volterra)	Prey, Pred	`Prey: alpha*Prey - beta*Prey*Pred` `Pred: delta*Prey*Pred - gam*Pred`	population	Phage–bacteria or grazing cycles

SIR / SEIR epidemic	S, E, I, R	`S: -b0*S*I`
		`E: b0*S*I - sig*E`
		`I: sig*E - g0*I`
		`R: g0*I`
			I + R	Final size against R_0 = b0/g0

Pharmacokinetics and release				
One compartment, oral dose	Agut, Cp	`Agut: -ka*Agut` `Cp: F*ka*Agut/Vd - ke*Cp`	C_p at time	Dose giving a target trough at 12 h
Two compartments, IV bolus	C1, C2	`C1: -(k10 + k12)*C1 + k21*C2` `C2: k12*C1 - k21*C2`	C1 at time	Distribution against elimination phase
Infusion that stops (smooth switch)	Ci	`Ci: Rin*0.5*(1 + tanh((tinf - t)/w)) - k10*Ci`	C at end	Infusion length for a target level
Effect compartment (PK/PD)	Ce	`Ce: ke0*(Ci - Ce)`	C_e	Delay between plasma level and effect
Dissolution (Noyes–Whitney)	Md	`Md: kdis*(Cs - Md/Vm)`	dissolved mass	Time to 80 % release

System	States	Rate expressions (AcquiLAB syntax)	Typical report	Question to ask, or optimization target
Adsorption, energy and heat				
Pseudo-first-order adsorption	q	`q: k1a*(qe - q)`	q at end	Contact time for equilibrium
Pseudo-second-order adsorption	q	`q: k2*(qe - q)**2`	q at end	Fit k2 and q_e to uptake data

Battery, 1-RC Thevenin model	SOC, V1	`SOC: -I/Q`
		`V1: I/C1 - V1/(R1*C1)`
		`report: E0 + Es*SOC - V1 - I*R0`
			terminal voltage	Current that keeps V above cut-off for 1 h

Supercapacitor charging (RC)	Vc	`Vc: (Vs - Vc)/(Rs*Cap)` `report: 0.5*Cap*Vc**2`	stored energy	Charge time to 95 %
Lumped heating (Joule and loss)	Tc	`Tc: (Pj - hA*(Tc - Tamb))/mcp`	T at end	Power for a 37 °C chip
Periodic forcing (day–night, AC)	Tf	`Tf: (Tm + dT*sin(2*pi*t/Per) - Tf)/tau`	T at end	Damping of a temperature swing

Formula library

`custom_algebraic` has parameters, named outputs (up to 8) and a report, with no states or time. Use it for closed-form relations, or as a fast objective inside an optimization loop. Non-integer powers are written as `exp(n*log(...))`, and `x` is reserved. A system holds at most 24 parameters, so split large formula sets.

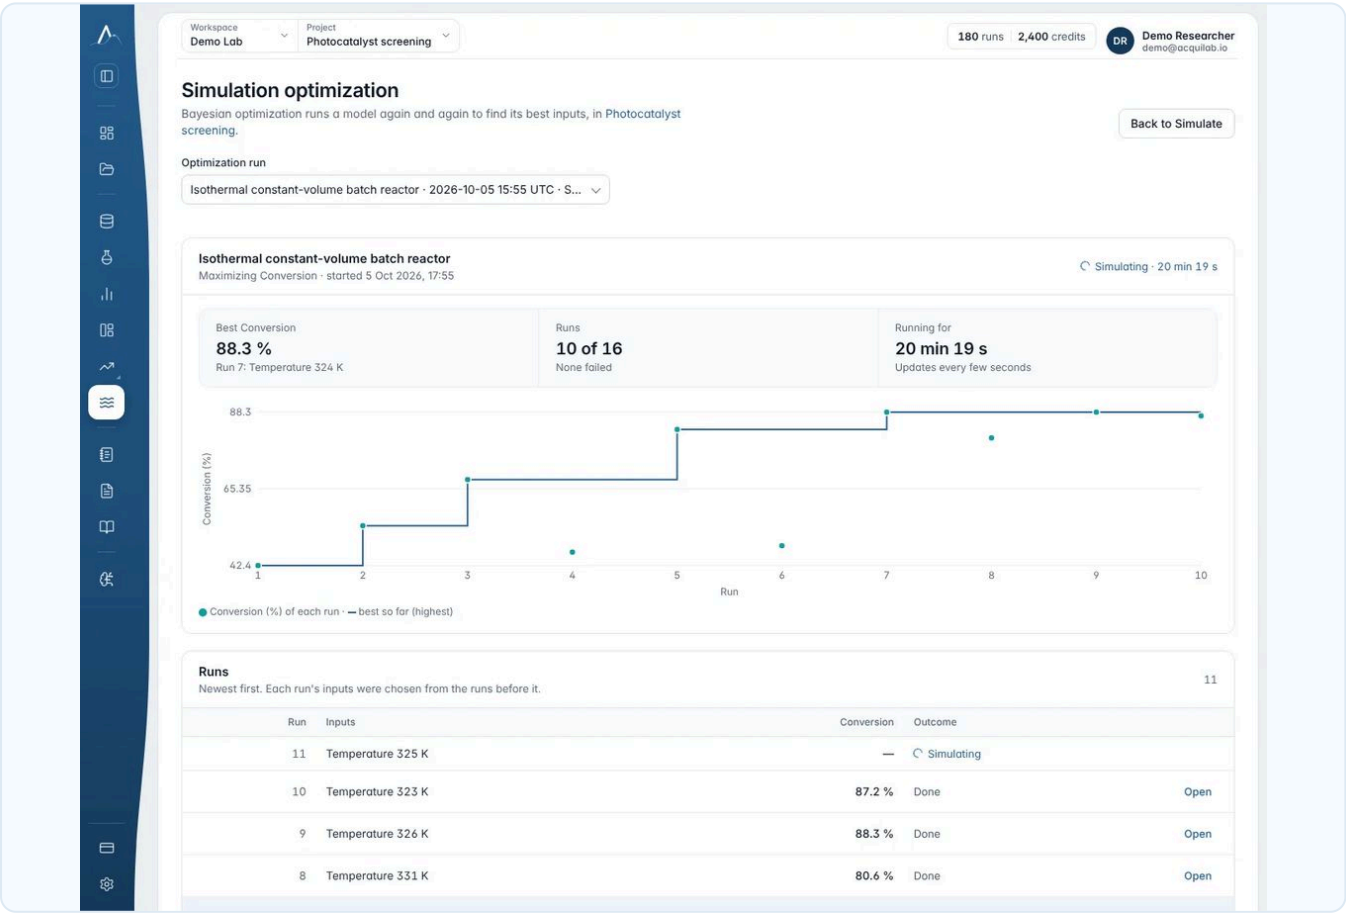
Relation	Expression	Use
Randles–Ševčík peak current	`0.4463*n*F*A*Cb*sqrt(n*F*v*D/(R*T))`	Expected CV peak; compare with the measured ip
Hill dose–response	`Emax/(1 + exp(nH*log(EC50/C)))`	Response at a dose; dose for 80 % effect
Arrhenius	`Apre*exp(-Ea/(R*T))`	Rate constant at a temperature
Eyring	`kB*T/hP*exp(-dG/(R*T))`	Rate from an activation free energy (DFT or xTB ΔG‡)
Nernst	`E0 + R*T/(n*F)*log(Cox/Cred)`	Electrode potential
Beer–Lambert	`eps*lp*Cm`	Absorbance; concentration for A = 1
Henderson–Hasselbalch	`pKa + log10(Abase/HA)`	Buffer pH
Langmuir isotherm	`qm*bL*Ce/(1 + bL*Ce)`	Loading
Freundlich isotherm	`KF*exp(log(Ce)/nF)`	Loading with a non-integer exponent
Q10 temperature scaling	`r1*exp((T2 - T)/10*log(Q10))`	Biological rate at another temperature
Limit of detection	`3.3*sb/slope`	LOD from the blank SD and the sensitivity
Stokes–Einstein	`kB*T/(6*pi*eta*rh)`	Diffusion coefficient of a particle

Simulation in the BO loop

Choose a model, then *What to do* → **Optimize**. Bayesian optimization runs the model repeatedly, choosing each new set of inputs from the results so far. Settings stay as they are; only the inputs inside the search ranges change.

Setting	Choices
Goal	The model's objective, for example maximise the HOMO–LUMO gap
Number of runs	2–100 for a ready-made model; the first few are spread across the ranges and the model chooses the rest
Each run	Starts automatically, or waits for my approval
Model (more options)	Improved (recommended) or Simpler (standard 0.1.0)
Result scale	Choose from the results (default), as measured, or logarithm (results must be positive)
Spread-out starting runs	How many runs come before the model starts choosing
Exploration	Balanced, explore more (new regions), or refine the best result
Stop early	When a result reaches a target, or after a number of runs without improvement

The loop page follows each run. Pause, resume, stop, or approve and skip single runs (a reason is optional). The loop reserves its whole run budget when it starts and gives back what it did not use.



A simulation loop on a ready-made model.

Written-model loop

Optimize 1–6 parameters of your own system, with up to 4 linear constraints, in 6–40 simulations (Simulate, or `start_written_model_optimization`). Each iteration runs your system with the recommended values, and AcquiLAB checks that the stored run used exactly those values before it counts. To minimise, the loop maximises the negative.

Simulation as a guide to lab BO

Turn on *Use simulation results as a guide* in a lab campaign with the Improved model. Simulation results become a second, cheaper source; the model check reports how well simulation agrees with the lab, and only lab results can be the best.

Constraints are available for written models; a loop on a ready-made model searches its design inputs within the ranges you narrow.

Chapter 6

Library and Zotero sync

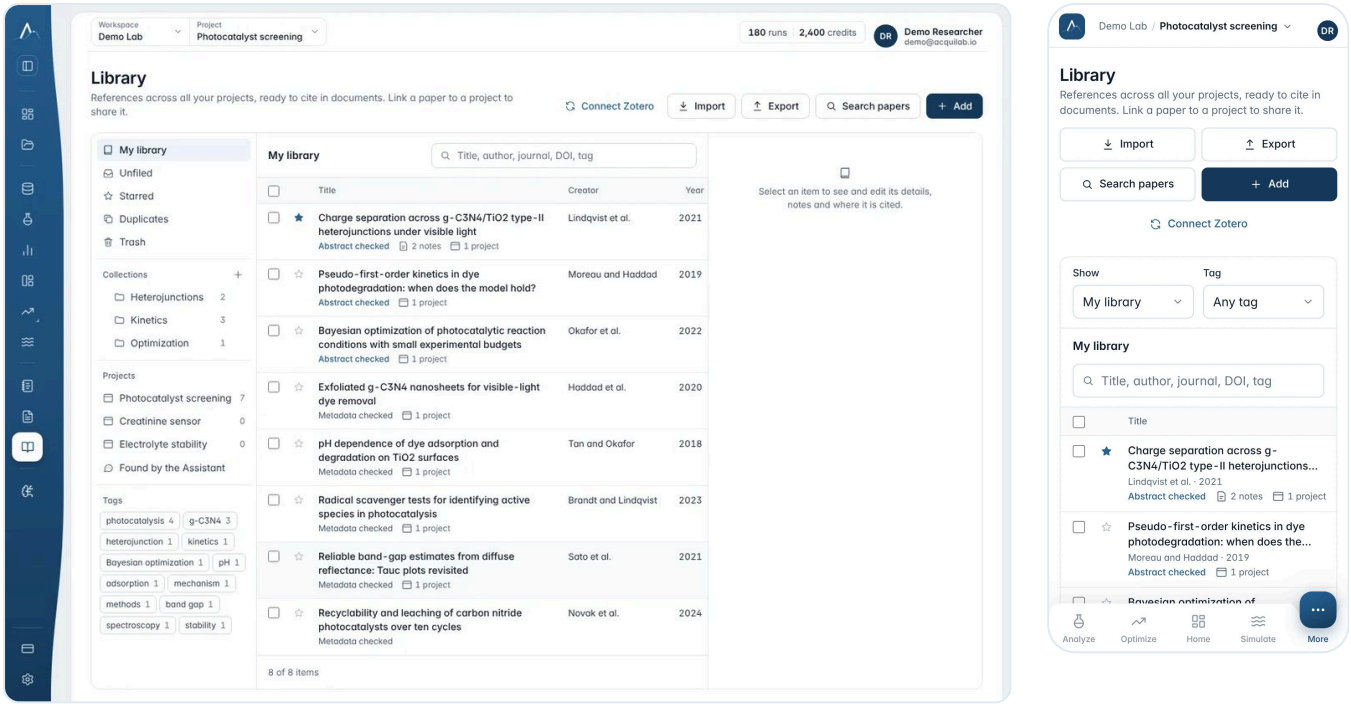
One personal reference library across your projects: adding, organising, verifying, de-duplicating and exporting references, and two-way Zotero sync.

One library for all projects

Library is your personal reference manager. There is one library per person, shared across all your projects, and you decide which items each project sees by linking them. Members of a linked project can see and cite those items; they cannot change your copy.

The page has three panes: the list of views, collections and projects on the left, the items in the middle, and the details of the open item on the right.

View	Shows
My library	Every item you keep.
Unfiled	Items in no collection.
Starred	The papers you return to often.
Duplicates	Items with the same DOI, or the same title and year, ready to merge.
Trash	Deleted items, until you restore them or delete them permanently.
Found by the Assistant	Papers the Assistant found in a project; *Save to library* keeps one.
Collections	Your folders, with subcollections. An item can be in several collections.
Projects	The items linked to each project.



The Library on a laptop and on a phone: views, collections and projects, and the papers with their verification state.

Adding references

Add by identifier

Paste DOIs, PMIDs, arXiv IDs or links, several at once. AcquiLAB fills in the details from Crossref, Europe PMC or arXiv. You can also enter an item by hand; it is saved as not verified.

Search papers

Searches OpenAlex and Crossref, optionally between two years. Add the papers you want in one click.

Import a file

BibTeX (.bib), RIS (.ris) or CSL-JSON (.json) up to 5 MB, from another reference manager, Google Scholar or a LaTeX project. The format is detected automatically, or choose it. Items already in your library are not imported again; they are added to the chosen collection or project.

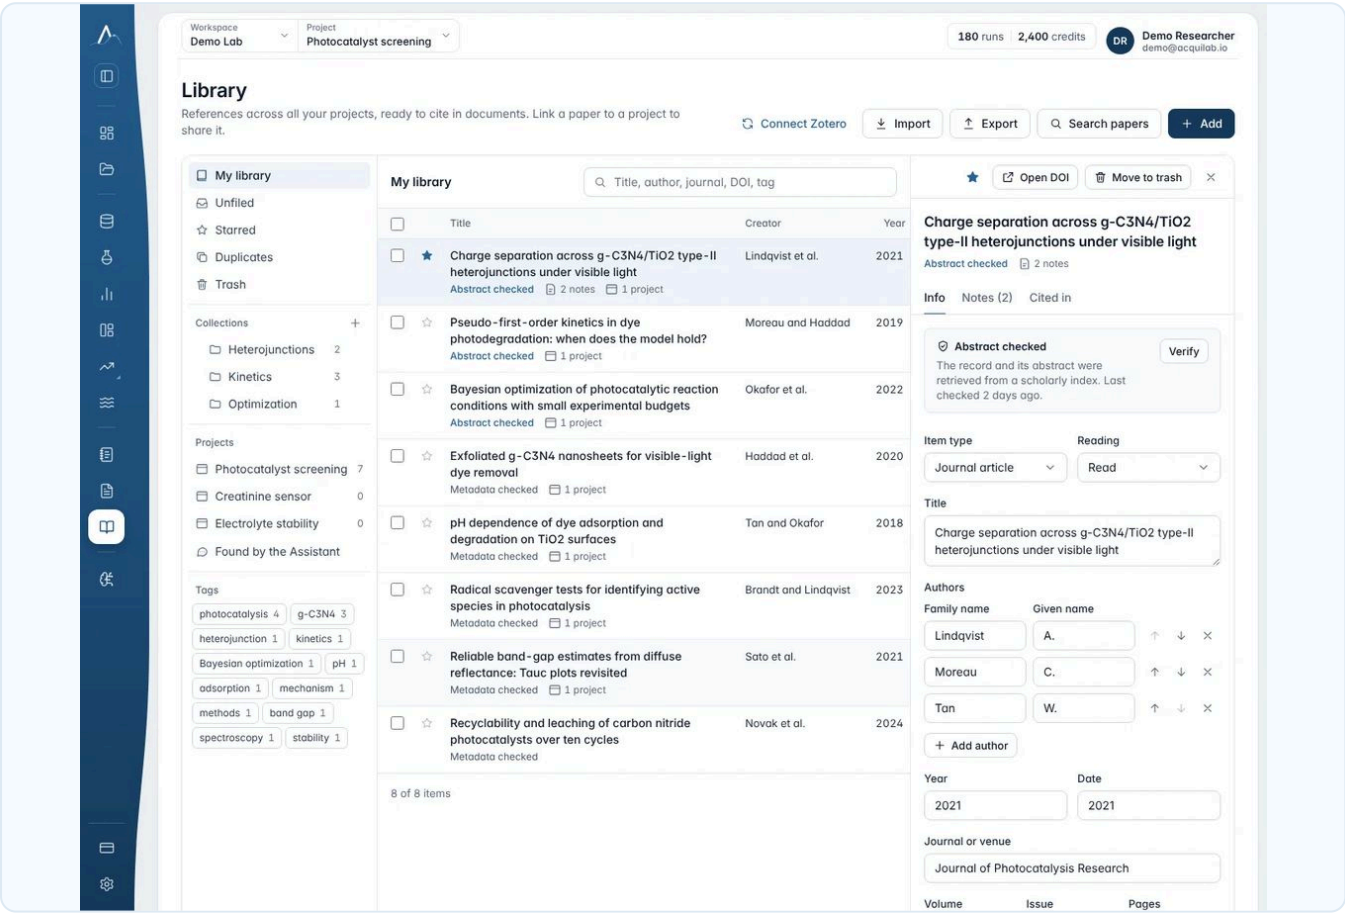
From the Assistant

Papers the Assistant found while answering appear under *Found by the Assistant*. Save the ones you want to keep.

Identifier lookups are limited to 60 a minute and searches to 10 a minute per person.

Collections, tags and projects

- Select several items for bulk actions: *Add to collection…*, *Link to project…*, add or remove tags.
- Star an item, and mark it unread, reading or read.
- In the details pane, edit the metadata (item type, title, authors, journal, year, volume, issue, pages, identifiers, abstract), the citation key and the tags, and write notes in Markdown.
- *Cited in* lists the documents that cite the item. Citing an item from your library in a project's document links it to that project automatically.



One reference in the details pane, with its verification state.

Verification and retractions

Every item shows how well its record was checked. **Verify** looks it up by DOI, PMID or arXiv id and fills in missing details. Changing an identifier marks the item not verified until you verify it again.

Shown as	Meaning
Not verified	Typed in, imported or synced without checking.
Metadata checked	Title, authors and journal were matched against a scholarly index (Crossref, Europe PMC or arXiv).
Abstract checked	The record and its abstract were retrieved from a scholarly index.

Retraction flags appear on the item, in the citation picker and in the document check: **Retracted** and **Withdrawn** (errors), **Corrected** and **Possible retraction** (warnings).

> **What the Assistant may cite**
>
> The Assistant and connected agents cite only papers AcquiLAB itself retrieved for the project, using their citation keys. Unverified items you typed in are never presented as checked sources.

Duplicates and trash

1. Open *Duplicates*. Items with the same DOI, or the same title and year, are grouped.
2. Choose the record to keep and press *Merge into the selected record*.
3. The merge fills empty fields from the others, joins their tags, collections, projects and notes, and moves the others to the trash. Their citation keys keep working in documents.

Deleting moves items to *Trash*, where you can restore them. *Delete permanently* or *Empty trash* removes them for good.

Export

Export references writes BibTeX, RIS or CSL-JSON for the whole library, a collection, a project or the selected items, for LaTeX, Word plug-ins and other reference managers. Documents export their own reference list; see Export in Notebook and Documents.

Two-way Zotero sync

Connect your Zotero personal library (My Library) and keep both in sync. Changes made in either place reach the other when you press **Sync now**.

1. In the Library, choose *Connect Zotero*. Open *zotero.org/settings/keys/new* (sign in to Zotero if asked).
2. Under **Personal Library**, tick **Allow library access** and **Allow write access**. Without write access the sync is read only: changes made in AcquiLAB are not sent to Zotero. Notes access is optional.
3. Save the key, copy it and paste it into *Zotero API key*, then **Connect**.
4. Press **Sync now**. A large library syncs in rounds; the dialog shows what was received, sent and removed on each side, and *sync again to continue* when more is waiting.

Topic	Behaviour
What syncs	Items and their metadata, tags, collections and notes. Attachments (PDFs) and annotations are never pulled, sent or deleted.
Edits on both sides	Merged field by field: a field changed only in Zotero is updated here, a field changed only here is sent (other Zotero fields are left untouched). If both sides changed the same field, the AcquiLAB value wins and the conflict is listed. Tags and collection membership merge as sets.

Topic	Behaviour
Moved to the trash here	Moved to Zotero's trash; restoring it here restores it in Zotero.
Moved to the trash in Zotero	Moved to the trash here, still linked; restoring it here restores it in Zotero.
Deleted in Zotero	Moved to the trash here and unlinked; restoring it later creates it again in Zotero.
Deleted permanently here	Deleted in Zotero as well.
Your key	Stored encrypted and used only to sync. It is never shown again.
Disconnect	Items stay in your library. Reconnecting the same Zotero account later resumes the sync.

If Zotero stops accepting the key, the dialog says so: disconnect and connect a new key. Zotero group libraries are not synced. PDF attachments with an in-browser reader and highlights are planned.

Chapter 7

Notebook and Documents

The project's running record, documents built from it, citations of papers and results, the number check and exports to Word and LaTeX.

Notebook

Notebook is the project's running record. It fills itself: analysis runs, each optimization result with its value and settings, recommendations, simulation runs and loops, saved figures and graphs, data uploads, documents, and your own questions to the Assistant.

- Write a note on its own (the composer at the top) or next to an entry (*Note*). You can edit or delete your own notes; project managers can delete any.
- Pin entries for the whole project. Filter by All, Pinned, Analyses, Optimization, Simulations, Figures, Data or Notes, search, or pick a day.
- The side panel shows what is pinned, this week's counts and latest result with *Make the weekly update*, and a small calendar. The calendar view shows a month with counts per day.
- Three months load at first; *Load earlier* adds three more.
- Tick entries, then *Add to document* or *Create document from these*.

Documents

Documents holds the project's reports, manuscripts, literature notes and slides. They are written in Markdown with LaTeX maths (`$…$` and `$$…$$`), save automatically, and refuse to overwrite a colleague's newer edit.

Start from the Notebook

New → From the Notebook builds a draft from a template (Weekly update, Methods & results, Full record, Group-meeting slides) and a choice of entries (This week, Last 30 days, Pinned only, Ticked in Notebook, or chosen dates). *Draft the summary with the Assistant* opens the Assistant with a prepared question for the same period.

Live blocks

Results, campaigns, notes and figures placed from the Notebook appear as cards with a link to their source. When the source changes, the card says so (for example *Newer result available*) and offers *Update*; *Update all* appears when several changed. Nothing changes until you press it.

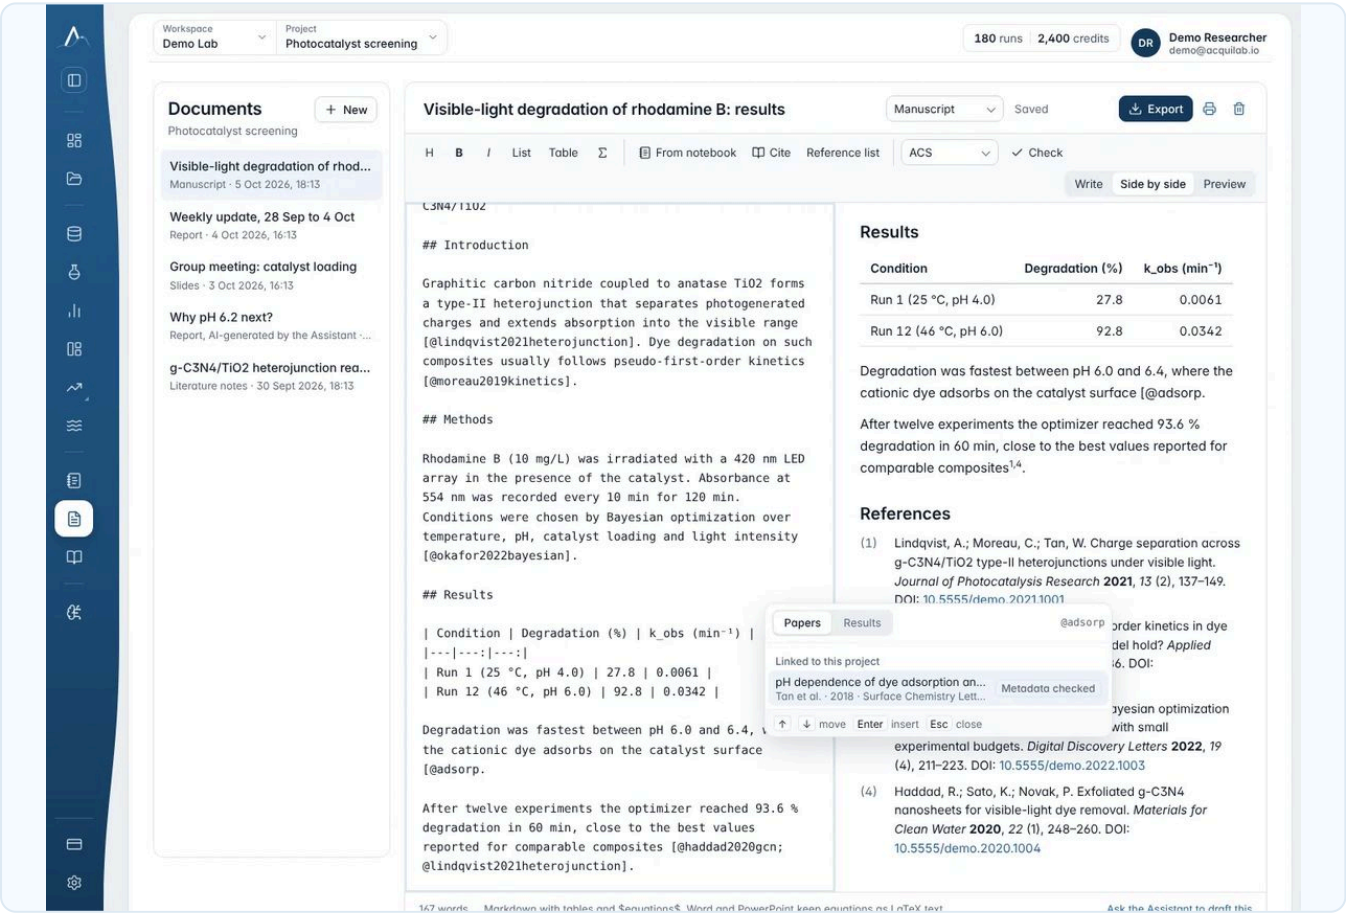
The editor has three views: *Write*, *Side by side* and *Preview*. The toolbar inserts entries *From notebook*, opens *Cite*, places the *Reference list*, sets the citation style and runs the *Check*.

Citing papers

1. Type `@` in the text, or press **Cite**. The picker's *Papers* tab searches your library and the project's papers by title, author, year or key.
2. Choose a paper. AcquiLAB inserts its key in Pandoc form and, if it came from your library, links the item to the project.
3. Choose the citation style in the toolbar: **ACS** (default), **RSC**, **Nature**, **APA 7**, **IEEE** or **Vancouver**. ACS, RSC and Nature number citations as superscripts.
4. Press **Reference list** to place the bibliography at the cursor; otherwise it goes at the end of the document.

You write	Meaning
`[@curie1898]`	One paper
`[@a; @b]`	Several papers in one citation
`[@a, p. 4]`	A paper with a locator
`[-@a]`	Author left out (in APA, the year only)

The style and the place of the reference list are stored in the document as hidden comments, so they travel with it. A key with no matching item shows as `[?key]` and is reported by the check.



Citing a paper in a document; the preview shows the formatted reference list.

Citing results

The picker's *Results* tab lists the project's analyses, figures, campaigns, simulations and datasets. Choose one, then how to insert it:

Insert as	What you get
Link	The name, linked to its source, for example “Analysis #12 (fluorescence, 9.txt)”.
Value	One measurement with its unit, such as λ_{max} = 520 nm, pinned to the revision you cited.
Embed	A live block at the end of the document, with Update when the source changes.

A cited value never changes by itself: if the source has a newer result, the check tells you and you decide whether to update.

Check numbers and citations

Check reads the whole document and checks that every citation resolves and every value still matches its source. Each finding links to its place in the text; values that changed can be updated from the panel.

Finding	Level	Meaning
Citation missing	Error	A cited key is not in the library or the project.
Retracted or withdrawn paper	Error	A cited paper is flagged as retracted or withdrawn.
Corrected paper, possible retraction	Warning	A cited paper has a correction or may be retracted.
Unverified citation	Warning	The cited item was never verified by DOI, PMID or arXiv id.
Source missing	Error	A cited result, figure or dataset was deleted or can no longer be read.
Value changed	Error	The number in the text no longer matches the cited revision of its source (units included).
Newer revision	Warning or note	The source has a newer result than the one cited.

Export

Format	What you get
Word (.docx)	Formatted citations and reference list; figures embedded at 300 dpi.
PDF	Print, then save as PDF; figures stay vector.
LaTeX	A .zip with the .tex (citations as `\cite{}`), `references.bib` and each figure as vector PDF; a single .tex when there is nothing else.
PowerPoint (.pptx)	A slide per heading or `---`; a figure fills its slide.
Excel (.xlsx)	A sheet per table.

Format	What you get
Markdown (.md)	The text as written.
Figures (.zip)	Every figure in the document as SVG.

Tick *Add data & methods appendix* in the export menu to add the sources, versions and analyzers of the results the document cites: the reproducibility statement journals ask for. Every export is listed under *Exported before* with its format, time and size.

> **References for other tools**
>
> To use the same references in another writing tool, export them from the Library as BibTeX, RIS or CSL-JSON.

Chapter 8

Assistant, MCP connector and connected agents

The project-aware Assistant, connecting your own AI agent through the MCP connector, the permissions and run limits you give it, and every tool it can call.

The Assistant

Assistant is a chat that works inside the selected project with your permissions. It reads and prepares; anything that runs or costs something waits for your approval.

What it does

- Reads your datasets, analysis results, campaigns, multi-objective fronts and simulation runs, and cites them.
- Searches the literature and reads open-access full text (Europe PMC). Papers it finds appear in the Library under *Found by the Assistant*.
- Prepares cards you approve with **Run now**: an analysis with every setting listed, a simulation or written model, a BO suggestion, a fit or an optimization loop. *Change in Analyze* and *Open in Simulate* fill in the page instead.
- Uses reviewed analysis skills as tools: cyclic voltammetry, chronoamperometry, DPV, SWV and UV-Vis calibration.
- When it prepares an optimization loop, runs the Setup check, switches on a log scale for each recognised quantity (a concentration, dose, flow, pressure and the like) spanning 100× or more, and says so on the card. Other advice, such as a merely wide range, stays a suggestion.
- Shows “step k of up to N” while it works, with **Stop**; **Save to Documents** keeps any answer.

Questions that work well

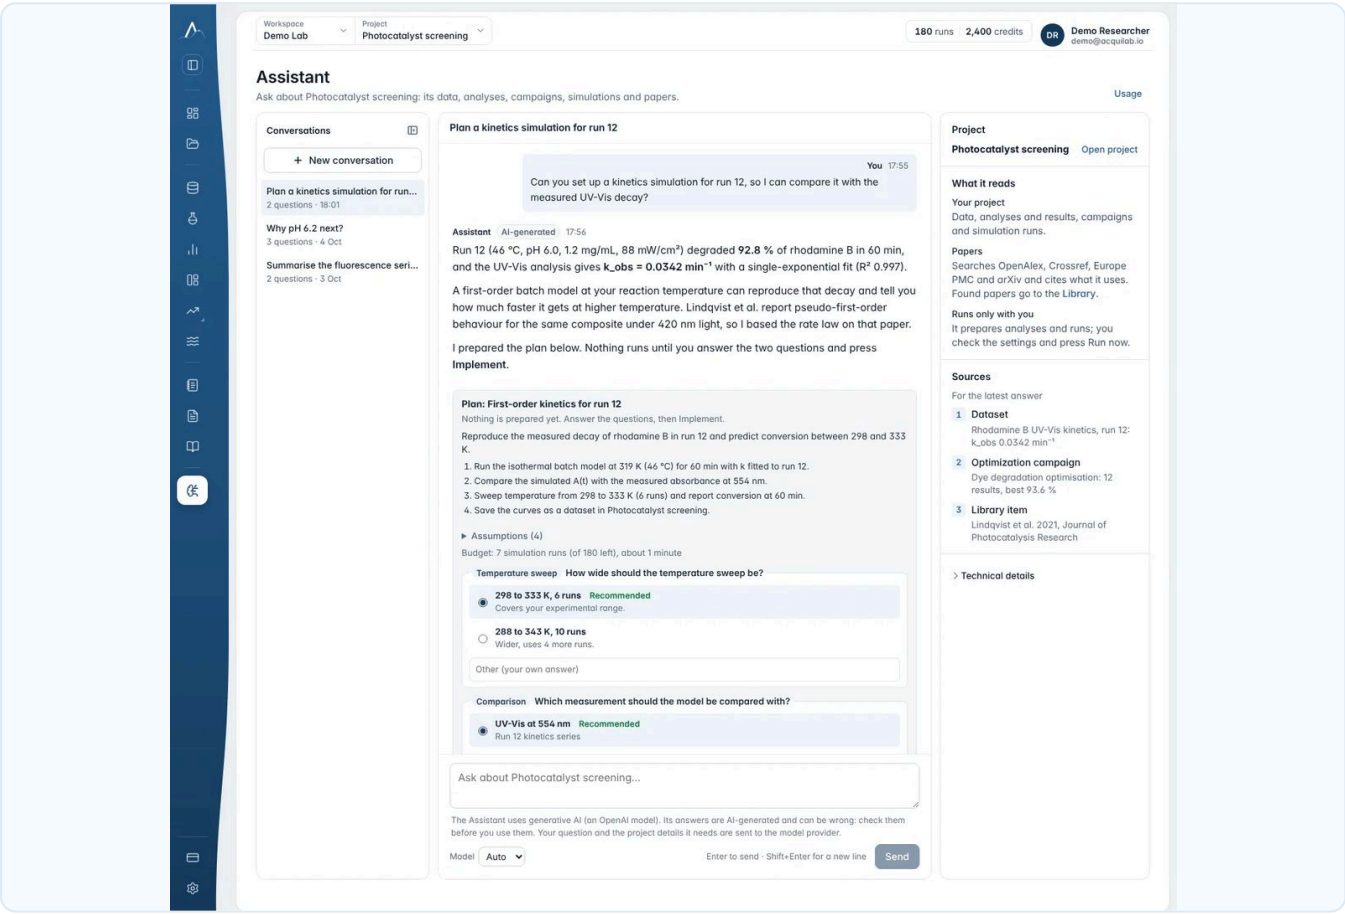
- “Analyse the CV file I uploaded and compare ΔEp across cycles.”
- “I want the HOMO–LUMO gap of phenol in water but I don't know DFT. Pick sensible settings from the literature, prepare the run and explain each choice.”
- “Write an ODE model of the creatinine cascade and find the best CA : CI : SOX ratio at 3 µM total enzyme.”
- “Why is the next experiment at 85 °C?”
- “Summarise what we learned this week and save it to Documents.”

Answers use Assistant credits, charged by the AI tokens each answer actually used (see Assistant credits).

Plan first

Before any simulation, model fit or simulation optimization started by the Assistant or an agent, AcquiLAB asks for a plan:

1. The plan states the full mechanism from the literature, every constant and assumption with its source, the budget, and the open questions with options.
2. You answer its questions. In the app, the plan is a card with the recommended option marked; press **Implement**, or write in the chat what to change and the Assistant revises the plan.
3. Only then can runs start, using that plan, within its number of starts and for 24 hours. The Assistant prepares each run as a card that you start with **Run now**; an agent still asks before it uses runs.



In the app the plan is a card: answer its questions (or pick the recommended options), then press Implement.

Connect an AI agent

The AcquiLAB MCP connector is a standard Model Context Protocol server, so any AI client or agent that supports MCP can use it. It works inside your projects with the permissions you grant.

1. In your AI client, add a custom MCP connector (remote server) with this address:

```
https://mcp.acquilab.io/mcp
```

2. Sign in to AcquiLAB when the client asks.
3. Review the permissions the agent asks for, and allow them.
4. If it may use runs, set how many it may use per action and per 24 hours (20 and 50 to start with).

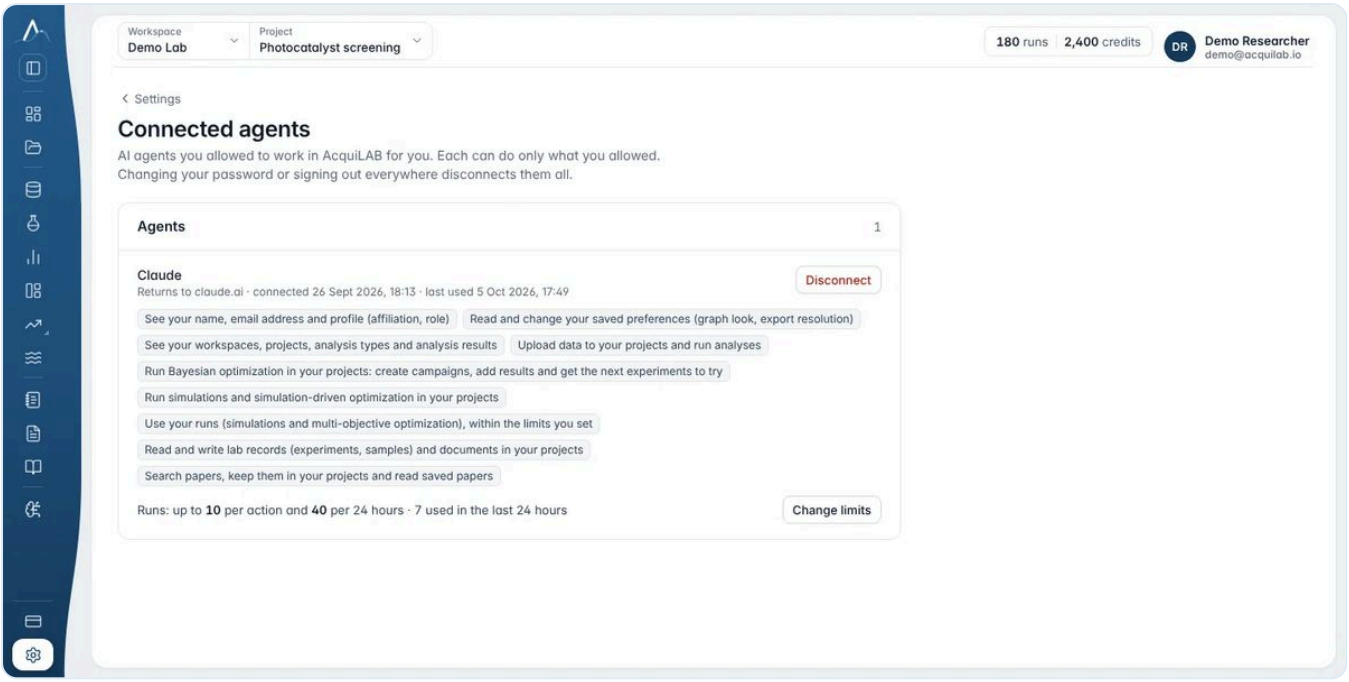
Permission	What it allows
Profile	See your name, email address and profile (affiliation, role).
Preferences	Read and change your saved preferences (graph look, export resolution).
Projects	See your workspaces, projects, analysis types and analysis results.
Analysis	Upload data to your projects and run analyses.
Optimization	Run Bayesian optimization in your projects: create campaigns, add results and get the next experiments to try.
Simulation	Run simulations and simulation-driven optimization in your projects.
Runs	Use your runs (simulations and multi-objective optimization), within the limits you set.
Records	Read and write lab records (experiments, samples) and documents in your projects.
Literature	Search papers, keep them in your projects and read saved papers.

Agents cannot delete anything, share projects, change billing or save Plot Studio layouts. Graphs an agent requests are drawn in your saved graph preferences.

Connected agents

Settings → Connected agents lists every AI agent you allowed. For each one you see where it returns to, when it was connected and last used, what it may do, and, if it may use runs, its limits per action and per 24 hours and how many it used in the last 24 hours.

- **Change limits** sets the runs per action and per 24 hours. AcquiLAB asks you to confirm it's you if you have not signed in recently.
- **Disconnect** stops the agent at once. To use it again, add AcquiLAB to the agent and sign in once more.
- Changing your password or signing out everywhere disconnects every agent.



Connected agents: what each agent may do and how many runs it may use.

Connector tools

The connector offers 54 tools. Each works only with the permission it needs.

Area	Tools
Setup	`get_preferences` `save_preferences` `list_projects` `create_project` `project_activity`
Analysis	`list_analysis_types` `run_analysis` `get_analysis` `get_analysis_result` `list_analyses` `get_figure` `list_skills`
Data	`list_datasets` `get_dataset` `get_dataset_columns`
Optimization	`list_optimizations` `check_optimization_setup` `create_optimization` `add_results` `suggest_next_experiments` `explain_recommendation` `get_optimization` `get_optimization_history` `cancel_suggestion` check_optimization_setup is free and creates nothing; create_optimization takes a scale per parameter and composition (parts that add up to a total).
Simulation	`list_simulation_models` `describe_simulation_model` `run_simulation` `get_simulation` `list_simulations` `compare_simulations` `compare_simulations_with_data` `save_simulation_as_dataset` `save_comparison_as_dataset` `annotate_simulation` `cancel_simulation` `fit_model_to_data` `start_simulation_optimization` `get_simulation_optimization` `control_simulation_optimization`
Written models	`check_written_model` `run_written_simulation` `start_written_model_optimization` check_written_model is free and runs nothing.
Plan first	`propose_plan` `approve_plan`
Lab records	`list_experiments` `save_experiment` `save_sample`
Literature	`search_literature` `list_references` `read_paper` `pin_reference`
Documents	`list_documents` `read_document` `save_document`

Good practice for agents

The connector tells every agent how to work. You can expect it to:

- Read your saved preferences first and save a preference you state (“box frame”, “600 dpi”).
- Ask which project the data goes into, and create a project only when you agree.
- Send your file as it is: instrument formats are read directly.
- Report measurements with their units exactly as returned, name the columns used, and give the figure and Plot Studio links. It never redraws a graph or recomputes a value.
- Record only what you measured in an optimization, never a prediction.
- Check an optimization setup before creating it, use the log scales the check switches on (only for recognised quantities such as concentrations, doses or flows spanning 100× or more), and tell you which variables now use a log scale. Other advice is shown to you, not applied.
- Give every value in a written model its source (a paper, or “assumed”) and keep units consistent.
- Cite only papers AcquiLAB returned, with their citation keys.

> **Runs and limits**
>
> Optimization suggestions and simulation starts have an hourly limit per person, shared by the app and your agents. Work that uses runs stays within the limits you set for the agent and asks first.

Chapter 9

Plans, runs, credits and settings

Plans and prices, what runs and credits pay for, limits, graph preferences, account settings and the basics of data and privacy.

Plans

Plan	Price a month	Runs a month	Assistant credits a month	Adds
Free	`$0`	`0 (20 trial runs once)`	`0 (promotions only)`	Analyze, Plot Studio, Standard and Advanced BO; simulations in the beta with trial runs
Plus	`$19.99`	`200`	`1,000`	Simulations and simulation in the BO loop
Pro	`$39.99`	`500`	`2,500`	Multi-objective BO
Team	`$299.99`	`3,000 pooled`	`20,000 pooled`	Up to 10 people, shared workspace and billing, spend limits per member, pay by invoice

- Yearly billing costs 10 × the monthly price (two months free).
- Prices are fixed in 12 currencies (USD, EUR, GBP, CAD, AUD, CHF, JPY, KRW, CNY, INR, SGD, BRL); the pricing page shows yours.
- Payments are not open yet: Plus, Pro and Team show *Notify me*, and Free is available now.
- Paid plans will renew automatically until you cancel in Billing; you keep access until the end of the paid period.

Runs

Runs are the allowance for computing that AcquiLAB does for you: simulations and multi-objective suggestions. Analyses, Standard BO and Advanced BO need no runs.

Work	Runs
A multi-objective suggestion	`1`
A simulation (most models, fits and written models)	`1`
An xTB simulation	`2`
A DFT or water MD simulation	`5`
A simulation optimization loop	The model's runs × the number of runs, reserved at the start; what is not used comes back when the loop ends
A run that fails	Given back

- Beta: simulations are open to every account while it has runs, and every person gets 20 trial runs once, for 30 days. Later, simulations will be part of the Plus, Pro and Team plans.

- In a team workspace, runs come from the workspace; in your personal workspace, from you. The top bar meter and **Usage & Billing** show what is left and what was used.
- Optimization suggestions and simulation starts also have an hourly limit per person (shared by the app and your agents) that protects the servers.

Assistant credits

Credits pay for Assistant answers. One credit is USD 0.01 of AI use, and an answer is charged by the tokens it actually used:

- Before an answer, AcquiLAB reserves enough credits for the largest possible answer, and refuses if they are not there. There is no silent overage.
- After the answer, it charges what the model provider reported (input, cached, output and reasoning tokens), at least 1 credit, and releases the rest.
- Free has no monthly credits; credits come with a paid plan or from promotions.

Limits

Area	Limit
BO campaigns	Up to 50 variables; 2–4 objectives for multi-objective; Advanced BO 1–16 experiments per round, multi-objective 1–4; model map and model check on Standard BO only
Simulation loop on a ready-made model	2–100 simulations
Written-model optimization	1–6 parameters, up to 4 linear constraints, 6–40 simulations; Standard BO (EI), one objective
Custom ODE and formulas	1–8 states (ODE) or up to 8 outputs (formulas), up to 24 parameters, about 60 terms per expression; LSODA; no events, noise or space
Variants and comparisons	2–6 variants per written run; compare 2–8 runs, or 2–10 runs with measured data
Fit to data	1–6 parameters, 10–400 model evaluations
Chemistry run time	Cantera up to 2 min (2 GB); xTB, DFT and water MD up to 30 min (4 GB)
Molecule size	xTB up to 100 atoms; DFT up to 40 atoms
Library	Imports up to 5 MB per file; 60 identifier lookups and 10 searches a minute
Agents	Runs per action and per 24 hours that you set (20 and 50 to start with)

Graph preferences

Settings → Preferences sets how your new graphs look, once, for every browser. Switch on **Remember my preferences**; each choice is saved as soon as you make it. *Standard* leaves a setting at AcquiLAB's default. *Stop remembering and delete them* removes them all.

Preference	Choices
Frame around the plot	Open, Box, Box with ticks on all sides

Preference	Choices
Tick marks	Outside, Inside
Minor ticks	On, Off
Text size	`6–14 pt`; axis titles use this size, tick numbers one point smaller
Colours	Standard, Colour-blind safe, Greyscale
Error bars	On, Off
Grid lines	On, Off
Figure size	Single column 85 mm, double column 180 mm, and the single and double column widths of Nature, ACS, RSC and Elsevier journals
Image resolution	`300 dpi`, `600 dpi`

Connected agents read the same preferences, so the graphs they request look like yours.

Settings

Page	What is there
Security	Password, email, two-step sign-in, passkeys, sessions and ORCID
Preferences	How new graphs look, in every browser
Connected agents	The AI agents you allowed, their permissions and run limits (see Connected agents)
Usage and billing	Plan, runs, credits and invoices
Your data	Download a copy, or delete your account
Support	Write to the AcquiLAB team

Data and privacy

- A project's work is visible to the members of that project. Your Library is yours; members of a project see only the items you link to it.
- AI agents work only with the permissions you granted, and you can disconnect them at any time.
- Your Zotero key is stored encrypted, used only to sync and never shown again.
- **Download my data** (Settings → Your data) gives one JSON file with your profile, sign-in methods, sessions, workspaces, projects, datasets, analyses, campaigns, documents, Assistant conversations and credit history. Uploaded files stay in their projects; download them there.
- **Delete your account** signs you out everywhere. For 30 days nothing is erased and you can sign in and keep your account; after that your profile, sign-in methods and your own workspaces are erased for good. Work in workspaces you share with others stays with them.

> **The full terms**
>
> The Privacy Policy and Terms of Service are the binding texts; this section is a summary.

AcquiLAB manual, updated 5 October 2026. Facts were checked against the AcquiLAB code and live catalog on that date. Parameter values in the model examples are illustrative and marked “assumed” unless a source is given. The latest version is at acquilab.io/docs.