DiffAtOnce:
Molecular Diffusion

DOSY · ES / EN

DiffAtOnce Academy

Training manual

Learn to measure. Understand before interpreting.

24 modules · 48 examples · 17 scenes

Students and teachers · version 1.0

Synthetic simulations, sources and actual application captures.

Beatrix: local LLM + course text · verified hybrid package

2026-10-05 · diffatonce.com

Contents

User guide · DiffAtOnce Academy **3**

Open the application . . . 3

Four spaces for one lesson . . . 4

Numeric workshop: create errors and observe consequences . . . 5

The laboratory's 17 scenes . . . 7

Beatrix: ask and inspect the response . . . 9

Consult source slides and methods . . . 11

Guided 100-minute session . . . 13

20-minute demonstration . . . 13

Save and share work . . . 14

Common problems . . . 14

Complete 24-module learning path **15**

Example and check gallery **18**

Provenance and reproducibility **20**

User guide · DiffAtOnce Academy

The course teaches defensible DOSY decisions. Before changing a setting, describe what you expect to happen. Then compare prediction and result and write down the control you would need for an actual sample. The application connects problems, mechanisms and checks; its workshop values are simulated.

Open the application

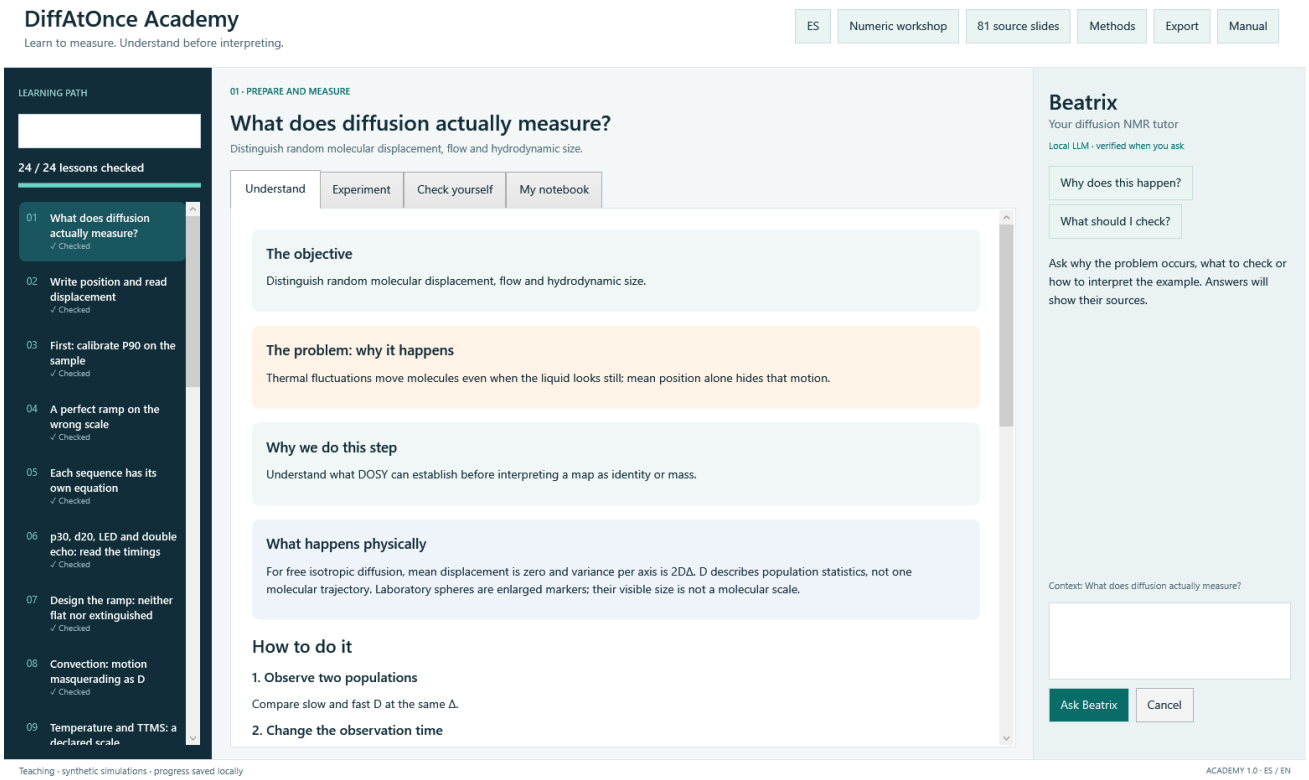
1. Extract the whole supplied application folder on Windows.
2. Open DiffAtOnce.Academy.exe, keeping Content, Lab and any supplied Models folder beside it.
3. Select English with the **EN/ES** button in the top bar.
4. Select a lesson in the learning path on the left. Search filters its title, chapter and problem explanation.
5. For a practical session, start with **03 · First: calibrate P90 on the sample**. Modules 01-02 provide the minimum diffusion and phase explanation.

The laboratory uses WebView2. If its runtime is unavailable, **Experiment** shows an explanation; **Understand**, **Check yourself** and the **Numeric workshop** remain usable. Paper links open an external browser; laboratory scenes are loaded from the local application package.

The **Manual** button opens the PDF edition for the active language. In a compact window, use the scrollbars; Beatrix retains the selected lesson as its context.

Four spaces for one lesson

The learning window



Actual QA capture. 24/24 progress comes from an automated test, not an assessed student. The Beatrix panel identifies its active mode; no model answer is fabricated.

Space	What to do	What to produce
Understand	Read objective, problem, purpose, mechanism and steps. Open examples and sources.	A hypothesis and prediction.
Experiment	Play the animation and change one variable. Compare useful case and counterexample.	An observation with parameters and units.
Check yourself	Answer both questions and press **Check and learn**.	A decision and explanation.
My notebook	Record what you changed, what happened and what you would check next.	A reasoned module-level record.

A lesson is marked **Checked** when both answers are correct. You can retry after reviewing feedback. The progress bar counts checked lessons; it does not certify permission or competence to operate a spectrometer. Notes are saved when the field loses focus, when you change lessons and when you close the window.

Open this experiment → loads the associated scene and the example's available parameters. Some examples are written decision exercises or critical readings of historical results: they are not 48 independent simulation engines. Keep the instruction visible and complete the reasoning in your notebook.

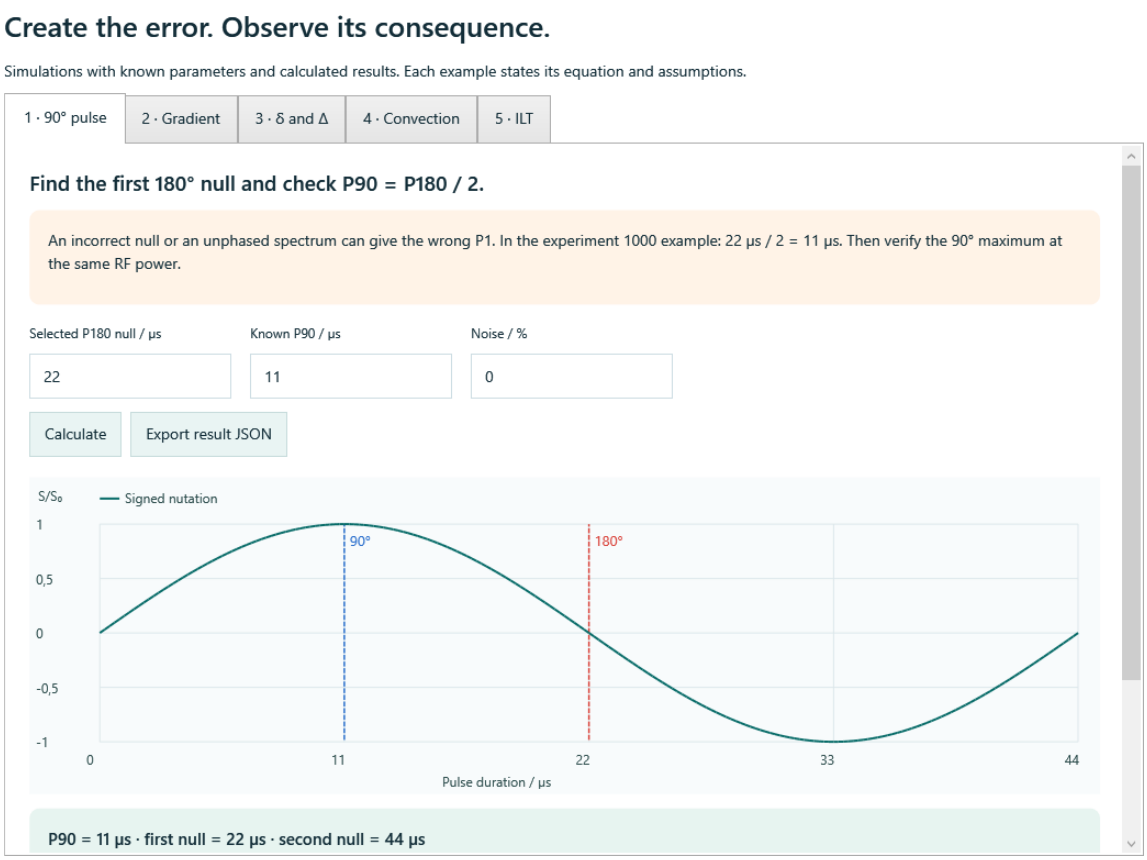
Numeric workshop: create errors and observe consequences

Open **Numeric workshop** from the top bar. Its five tabs calculate curves and state their assumptions. Decimal points and commas are accepted. Press **Calculate**, then **Export result JSON** to save a result. The file is labelled synthetic-teaching-simulation.

Editing a parameter clears the previous curve and results. Press **Calculate** to generate an output consistent with the new values; do not read an old plot as if it represented a newly edited field. Beatrix receives the workshop context when it opens, changes tab or calculates: exercise parameters, calculated results and assumptions. Before asking for a numerical explanation, check that a current result is available.

1. 90° pulse

Exp. 1000: 22 / 2 = 11 µs



Actual native P90 workshop capture with signed curve: blue marks the 90° maximum and red the first 180° null. 22 µs and 11 µs are exercise values; no instrument was operated.

The user's manual example is associated with experiment 1000. This is a workflow label, not a command sent to TopSpin.

1. Keep **Known P90=11 µs**, **Selected null=22 µs** and zero noise.
2. Locate the 90° maximum, first nontrivial 180° null and next 360° null.
3. Check **22/2=11 µs**.
4. Change the selected null to **44 µs**. Halving it gives 22 µs; the workshop indicates that this is not the first 180° null.

P90=P180/2 requires identifying that null at unchanged RF power. For a 360° null, use P360/4. On an actual instrument keep phase, gain, integration region and conditions comparable; signed nutation helps identify the crossing. The simulation omits RF inhomogeneity and relaxation and does not demonstrate that an actual tube is calibrated.

2. Gradient

Use initial values **Nominal G=420 mT/m**, **δ=1.4 ms**, **Δ=80 ms**, **Dref=0.80 ×10⁻⁹ m²/s** and **Actual/nominal G=1.10**. Fitting nominal b gives apparent D near **0.968 ×10⁻⁹**: 1.10^2=1.21, a 21% bias. The fit can still be perfect.

The purpose is to recover scale using a standard signal. Actual calibration requires Dref(T), solvent, stable temperature, consistent extraction and the correct sequence-specific b equation. This teaching standard does not replace TBO or write gradient constants into TopSpin.

3. δ and Δ

Compare the initial case with **δ=0.3 ms**, then with larger δ. Observe slow and fast signals, maximum b and signal relative to noise. A flat curve provides little slope; an extinguished curve loses its tail. Workshop indicators -slow signal≤50% and fast signal>3σ- are teaching criteria for this example, not universal acquisition rules or probe limits.

In Bruker, first translate p30/d20 into the pulse program's physical intervals. The same parameter label does not mean the same time in every sequence. The workshop uses ideal rectangular PGSE; it does not check heating, duty cycle or your hardware's maximum permitted gradient duration.

4. Convection

Keep **D=0.80**, **Common velocity=35 µm/s** and set velocity half-width to zero. A shared velocity rotates phase and can preserve magnitude. Increase half-width to **300 µm/s** and observe magnitude loss, possible lobes and apparent D. Change Δ with other settings fixed.

This workshop uses a uniform velocity distribution with a half-width. The animated scene uses its own Gaussian spread model. Neither identifies an actual tube's flow profile. Δ dependence alone does not establish convection: restriction, exchange, kinetics and errors may also matter. A suitable result may be ‘inconclusive: acquire a control’, without a definitive mass.

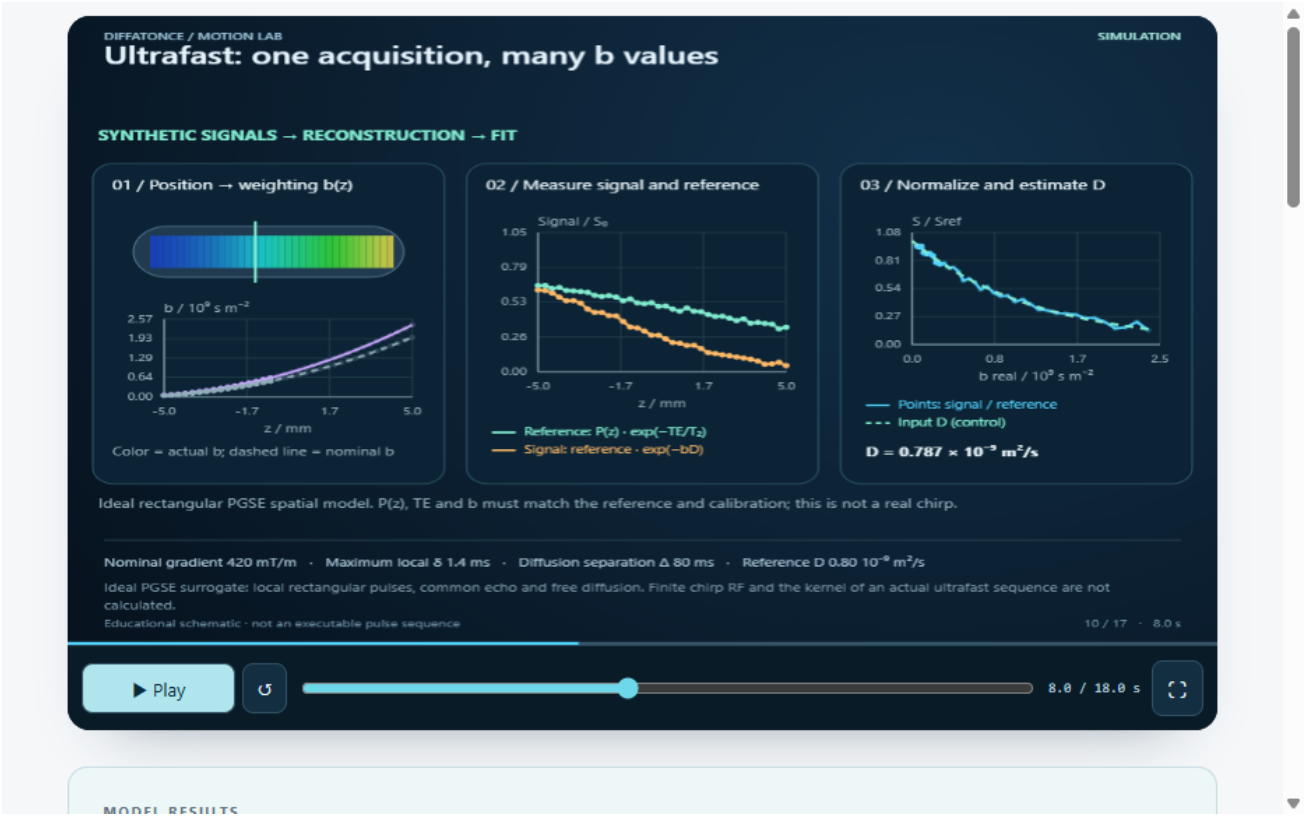
5. ILT

Start with **bmax=3 ×10⁹ s/m²**, **32 points**, **40 bins** and **1% noise**. Compare one- and two-component curves, difference/noise and singular spectrum. Increase bins while holding acquired points fixed and explain why no new measurements appear. Enter whole numbers for points and bins.

Effective rank depends on the workshop's kernel, scale and noise criterion; it is not a molecule count. The graph explains ill-conditioning, but does not execute every catalogue algorithm or guarantee ILT component separation.

The laboratory's 17 scenes

Ultrafast in the local laboratory



Actual computed-scene capture. Profile and estimates come from an ideal model, not an experimental acquisition.

Scenes	Question they help answer
Brownian, spin echo, gradients-b	What displacement is encoded, and how is b constructed?
Pulse-90, acquisition-window	How are RF and an informative attenuation window established?
Convection, attenuation-mixture	Why can an attractive curve give apparent D?
ILT, overlap-joint	What distinguishes components, and what does a joint model add?

Scenes	Question they help answer
Ultrafast, SPEN, time-resolved	What changes with spatial encoding or an evolving sample?
D-T_2, prototype-1d	Which coordinates are recovered, and what validation remains?
Restricted, exchange, anisotropy	Which other mechanisms cause temporal or directional dependence?

Read each scene's evidence status and limitations. Some are diagrams; UF, SPEN and time-resolved compute signals and fits under ideal models. Animation time is neither Δ nor reaction time. Save outputs offered by the scene; the application asks where to save them. Historical video can represent an earlier scene version and is identified as such.

Three advanced comparisons

• **Ultrafast:** compare raw D, reference+nominal-b D and reference+calibrated-b

D. Reference cancellation of profile and relaxation requires matching conditions. An actual chirp needs its q(z,t), not unreviewed application of a rectangular formula.

• **SPEN:** follow complex data, polarity reordering, phase and transforms

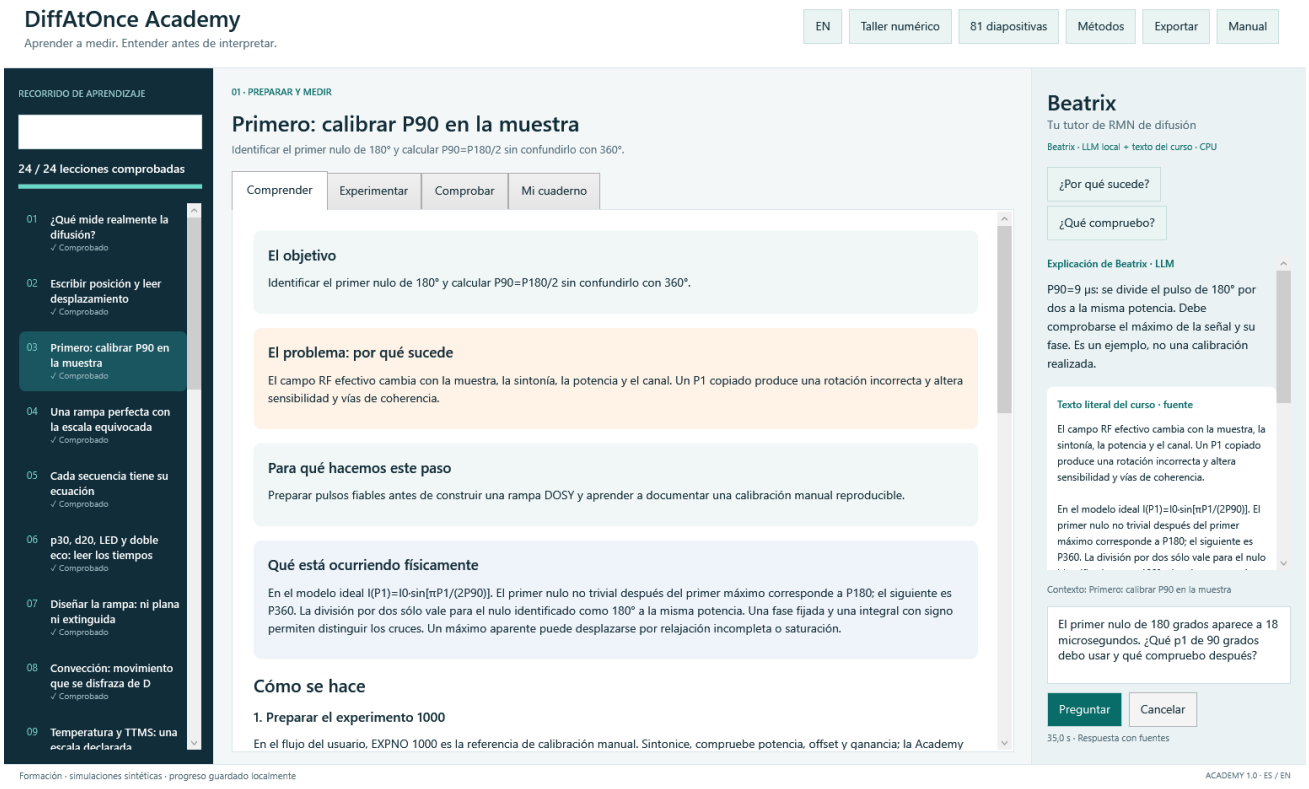
before fitting. Try pure A, then absent A at the same selected frequency; B leakage is not a valid recovered D_A.

• **Time-resolved:** set tLab=0 and compare ascending, descending and permuted

ordering. D_A remains fixed while population changes. Frame cadence is not temporal resolution: overlapping windows share data and average evolution. Windows shorter thanΔ+readout are not executable acquisitions.

Beatrix: ask and inspect the response

Local generation and literal text: two visible blocks



Actual CPU inference capture: the question uses a first null at 18 µs and Beatrix explains P90=9 µs. This is a different exercise from the 22/2=11 µs workshop. LLM text and the literal course copy are separately labelled. Sources and context accompany the answer; no instrument was calibrated. This native UI capture is in Spanish; the course and application also support English.

Select a lesson and use **Why does this happen?**, **What should I check?** or write a question. For example:

- 'The first null is 22 µs: why divide by two, and how do I rule out 360°?'
- 'My fit is perfect with actual g 10% different: why does D change 21%?'
- 'Why does permuting gradients not freeze a reaction?'
- 'Two methods give different peaks with similar residuals: which check is missing?'

The answer shows its mode and sources. Open the cited slide or paper. If an answer requires conditions you have not supplied, it should acknowledge that. **Cancel** stops the current question.

Beatrix combines two labelled blocks:

1. **Beatrix's explanation · LLM** contains text actually generated by the local model.
2. **Literal course text · source** copies the module's problem and mechanism. It is course content, not attributed to model generation.

The mode reads **local LLM + course text · CPU**. Read both blocks and open their sources. In native checks of open SPEN/TR questions, generated text omitted causes; the literal block supplies the explanation preserved in the course. A 32/32 HF evaluation does not mean the CPU GGUF alone passed those causal questions. Reports preserve that distinction.

Check the **Context** label before asking. For example, in the workshop change the null, press **Calculate** and ask why the result becomes invalid. Returning to **Experiment** restores laboratory-scene context; selecting another lesson restores its module context. Outputs supplied to Beatrix are simulations, not spectrometer data.

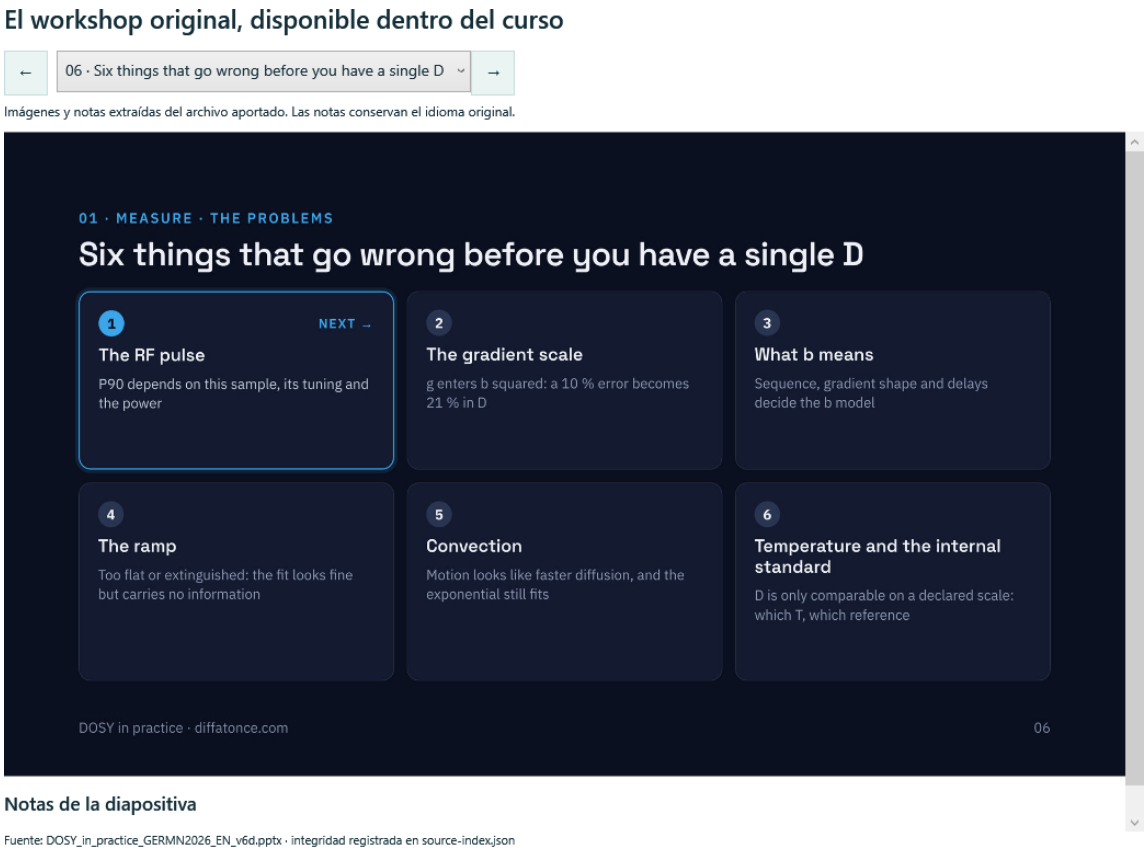
Verified hybrid package status. Local fine-tuning scored 32/32 on the reserved HF test, compared with 21/32 for the base model. The CPU package passed 15/15 checks of **LLM + literal course text**. Generated-only CPU answers omitted causes in open SPEN/TR questions; those limitations remain in the reports. The literal block supplies the course mechanism and is not attributed to the LLM. Manifest hashes link model, reports and course version. This evidence is not instrument validation, independent expert review or measured educational effectiveness.

The application verifies the package before inference. **Documentary guide · no LLM** means retrieval without inference, including a missing model or failure; do not present it as model generation.

Beatrix explains and proposes checks. It does not modify results, calibrate the instrument or validate identities. Fluent wording or a fine-tuned model does not replace experimental evidence or teacher review.

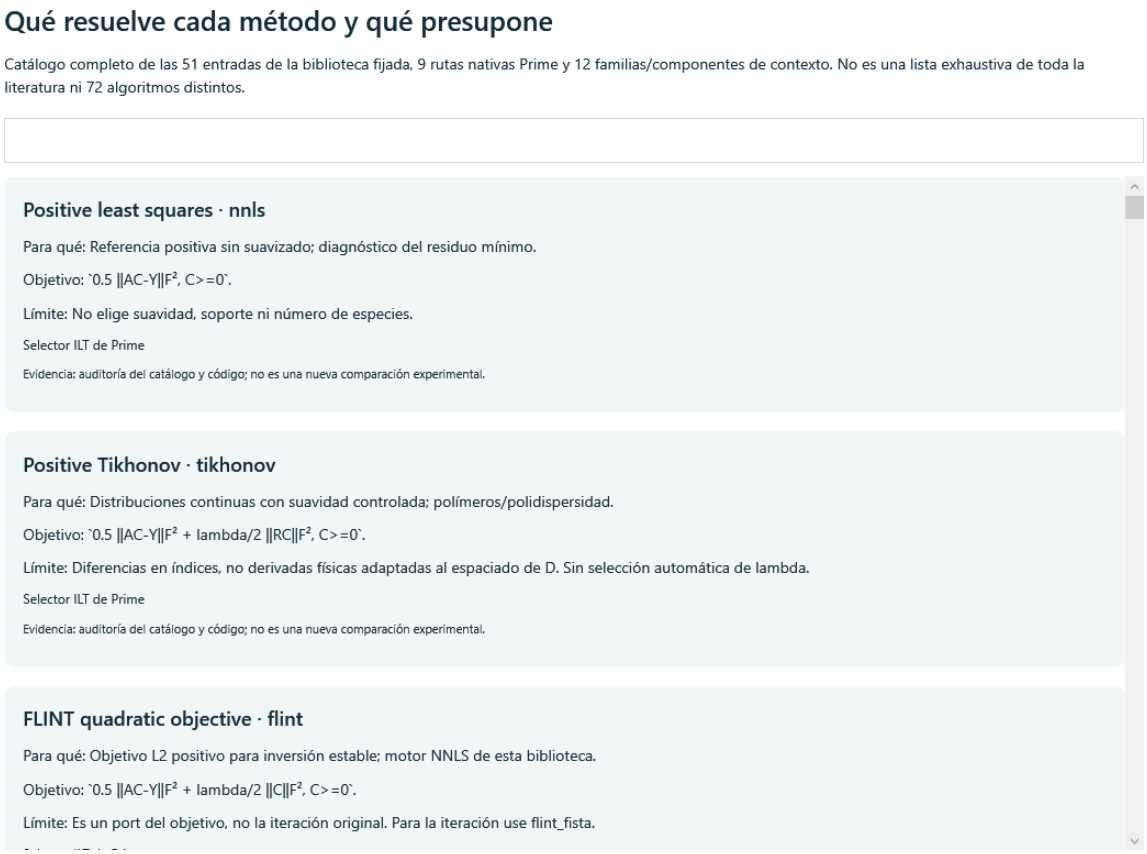
Consult source slides and methods

Original workshop reference viewer



Actual viewer capture for the 81supplied slides. Images and notes retain their source language. This native UI capture is in Spanish; the course and application also support English.

Choose by objective and limits



Actual 72-record catalogue capture. Historical availability does not mean Academy executes every method. This native UI capture is in Spanish; the course and application also support English.

81 source slides opens the original workshop with a selector and arrows. Images and notes retain their original English; lessons are bilingual. Lesson source buttons open relevant slides directly. TP194, KK2 and HDO values are historical results reported in the deck, not acquisitions recalculated by Academy.

The viewer uses the frozen extracted snapshot. At release, the current PPTX container has a different hash; all 81 images, notes, descriptions and ordering match that snapshot. See source-input-reaudit.json to distinguish teaching content identity from binary identity. The cause of the container change is unknown.

Methods searches names, uses and limitations. The catalogue contains 72 records: 51 library cores, 9 historical Prime routes and 12 context families. These are neither 72 distinct algorithms nor an exhaustive literature survey. Availability refers to the source catalogue; it does not promise an Academy button executing every method. Compare objective, assumptions and limitations as well as its name. Fixed-λ Tikhonov is not the full CONTIN program.

Guided 100-minute session

Time	Activity	Student output
0-10 min	Modules 01-02: motion, phase and the meaning of b.	Prediction of changing Δ.
10-25 min	Modules 03-04 and P90/gradient workshop.	22/2, 44-null diagnosis and 21% bias.
25-40 min	Modules 05-07: sequence, LED and δ/Δ design.	Physical-interval→parameter table and justified ramp.
40-55 min	Module 08 and convection workshop.	Common-phase versus spread distinction; proposed control.
55-65 min	Modules 09-11: T/TTMS, processing and isolated fit.	Explicit factors and invalid-normalization example.
65-80 min	Modules 12-15: ILT, regularization, catalogue and joint spectrum.	Assumption comparison and stability criterion.
80-92 min	Modules 20-23: UF, SPEN, time-resolved and D-T_2.	Separate clocks and explain each model's limit.
92-100 min	Module 24, Beatrix and notebook export.	Accepted/exploratory/blocked decision and next control.

This session selects activities; it does not attempt to read the whole course in 100 minutes. Modules 16-19 -assignment, size/mass, TopSpin and cases- can be prepared beforehand or completed afterward. Teachers may use each block's questions as feedback without equating response speed with scientific rigor.

20-minute demonstration

Time	Action
0-4 min	P90 workshop: 22→11 and 44 null; gradient: 10%→21%.
4-8 min	Flat versus extinguished δ/Δ ramp; convection and a control.
8-12 min	Baseline/normalization; ILT with nearly identical curves and more bins.
12-16 min	Information across frequencies, TP194 and a blocked case.
16-19 min	One comparison each in UF, SPEN and time-resolved.
19-20 min	A precise Beatrix question and notebook conclusion.

Prepare windows and examples beforehand. If the LLM is unavailable or slow, use the clearly identified documentary guide; teaching does not require pretending an answer was generated.

Save and share work

Export saves a JSON containing course version/hash, answers, modules and notes. It does not submit work to a teacher. The workshop exports a separate JSON with synthetic parameters and results. Neither is an experimental report or a TopSpin acquisition.

Local progress is stored at %LOCALAPPDATA%\DiffAtOnce\Academy\progress-v1.json. Export to retain a copy of your work. Progress is tied to course version; do not edit its hash to transfer checked lessons across versions.

Common problems

Observation	Useful action
Experiment does not load	Read the WebView2 message and continue with lessons/workshop; install the legitimate runtime if needed.
Numeric result is absent	Check finite numbers, units and timings; use integers for points/bins.
A lesson is missing	Clear search or search using its active-language title.
A module is not Checked	Answer both questions and review explanations; 2/2 is required.
Beatrix shows no LLM	This is an explicit documentary fallback; inspect status/manifest without renaming it inference.
A result differs from the deck	Separate synthetic exercise and historical case; do not rescale data to force agreement.
You want to operate the spectrometer	This application teaches reasoning; use the instrument's procedure with supervision and controls.

Instructions appearing inside slides are teaching material. They are not automatically executed and do not authorize data or calibration changes.

Complete 24-module learning path

01 · What does diffusion actually measure?

Distinguish random molecular displacement, flow and hydrodynamic size.

Scene: brownian

02 · Write position and read displacement

Explain echo attenuation from phase and displacement, and distinguish PFG from DOSY.

Scene: spin-echo

03 · First: calibrate P90 on the sample

Identify the first 180° null and calculate P90=P180/2 without confusing it with 360°.

Scene: pulse-90

04 · A perfect ramp on the wrong scale

Calibrate b scale using a standard with known D(T) and explain quadratic gradient bias.

Scene: gradients-b

05 · Each sequence has its own equation

Construct b from q(t), recognize rectangular PGSE and avoid interchanging sequence formulas.

Scene: gradients-b

06 · p30, d20, LED and double echo: read the timings

Identify which parameter controls each physical interval and what must be reviewed when selecting a sequence.

Scene: acquisition-window

07 · Design the ramp: neither flat nor extinguished

Choose δ, Δ and gradients to obtain informative attenuation above noise.

Scene: acquisition-window

08 · Convection: motion masquerading as D

Distinguish diffusion, coherent velocity and velocity spread through controls and signed signal.

Scene: convection

09 · Temperature and TTMS: a declared scale

Calculate reference and temperature factors without duplication or calling them convection correction.

Scene: attenuation-mixture

10 · Processing can manufacture D too

Detect phase, baseline, integration and normalization bias before inversion.

Scene: attenuation-mixture

11 · One signal: start with the clearest fit

Interpret monoexponential fitting, residuals and used points before a complete map.

Scene: attenuation-mixture

12 · ILT: nearly identical curves, different distributions

Explain ill-conditioning, informative rank and resolution limits without equating bins with information.

Scene: ilt

13 · Regularizing chooses a preference

Interpret λ, prior and stability using the same data, grid and noise model.

Scene: ilt

14 · Choose algorithms by their assumptions

Choose families for isolated signals, continuous distributions, few components or joint spectra, and distinguish solver, objective and diagnostic.

Scene: ilt

15 · One molecule leaves several clues

Explain benefits and limits of sharing diffusion information across frequencies.

Scene: overlap-joint

16 · Beatrix: ask, reason and verify

Use a sourced tutor to propose hypotheses and resolve questions without delegating chemical conclusions or data changes.

Scene: overlap-joint

17 · From D to radius and mass: state the assumptions

Calculate Rh under Stokes-Einstein and distinguish empirical apparent mass from absolute mass.

Scene: brownian

18 · A calibration workflow that can also reject

Organize RF calibration, standard, metadata and review before automated acquisition.

Scene: pulse-90

19 · Successful cases and cases that stop

Compare TP194 and KK2 controls and provenance without transferring indices or results to other samples.

Scene: attenuation-mixture

20 · Ultrafast: a ramp spread along the tube

Interpret spatial encoding, profile normalization and b(z) calibration through known and recovered signals.

Scene: ultrafast

21 · SPEN-DOSY: reconstruct before fitting

Follow complex EPSI signal, polarity reordering, phase correction and recovery of frequency-specific attenuations.

Scene: spen

22 · A reaction progresses while the ramp is measured

Separate t_lab, Δ, window and cadence and demonstrate gradient-order bias during a reaction.

Scene: time-resolved

23 · D-T_2 and development on conventional probes

Distinguish diffusion-relaxation maps, chemically resolved DOSY and a prototype without quantitative validation.

Scene: dt2

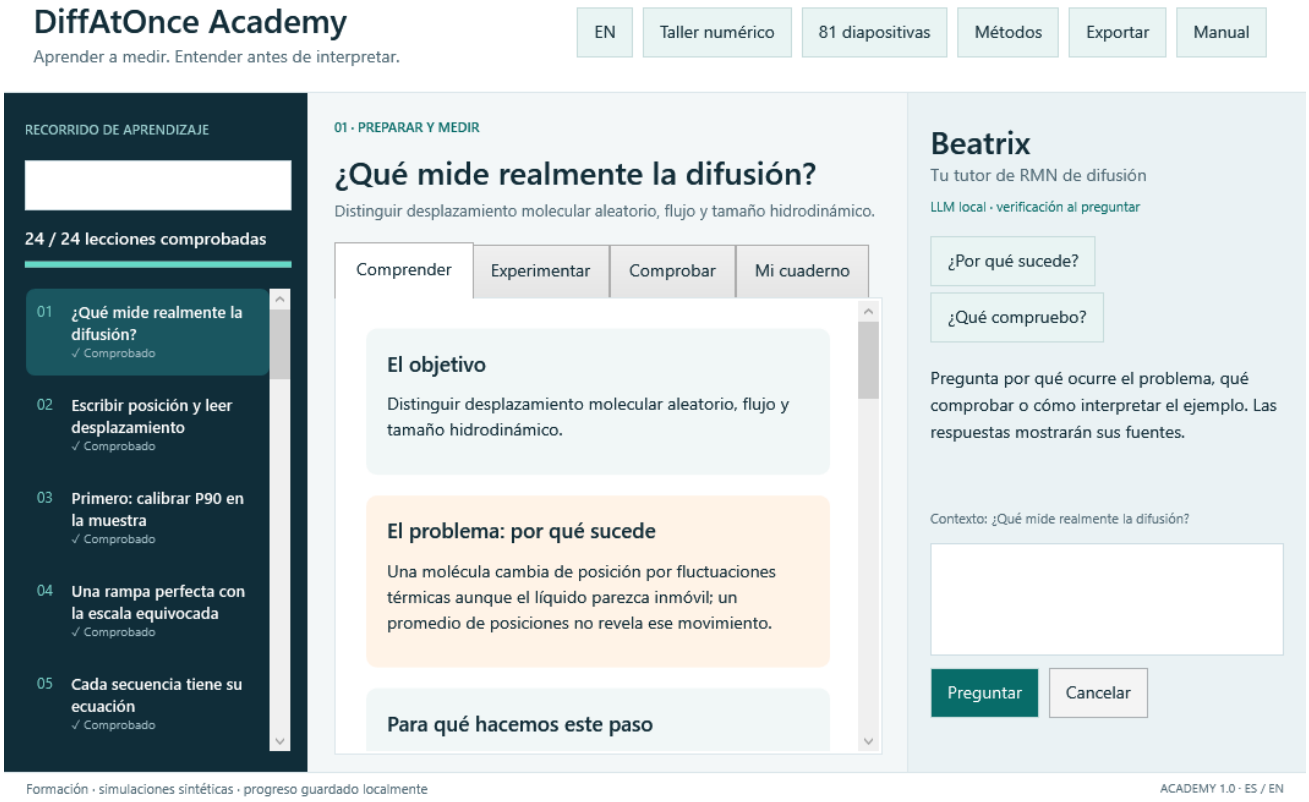
24 · Final case: decide what can be concluded

Build a complete decision and distinguish convection, restriction, exchange and anisotropy before choosing a model.

Scene: restricted

Example and check gallery

The same lesson in a compact window



Actual compact-window capture in Spanish. Use scrollbars to read the whole lesson; Beatrix retains the active module context. Displayed progress comes from the software test.

Expanded laboratory: compare hypotheses



Actual WebView2 capture from the final laboratory version. Controls let students vary parameters, observe the signal and compare models. Results remain ideal teaching simulations.

How to read this example: input D is 0.800 $\times 10^{-9}$ m²/s. Fitting without normalization gives 1.296; reference plus nominal b gives 0.952; reference plus calibrated b gives 0.787. Remaining bias is -1.64 %. Compare each estimate with the known value and explain what the reference cancels. Calibration is declared in the simulation; it has not been measured on the spectrometer.

Provenance and reproducibility

The frozen extracted snapshot maps 81 slides to 24 modules.

PPTX at extraction · SHA-256: d54f2a6b7a42dc640a26b199408bf19505cea4dac7615089e5f679d168a63a1d

Course SHA-256: f9a17df9645b541abc70907e5d04cd38c213293443ce1d37c2c0d94f8aea0cb9

Current PPTX · SHA-256: 5dfce0b0ec2b6f9bf2ddc1eef5ca04a4d31e053c9e2f051d6542fff63e273975

The current container has a different hash and 22 more bytes. All 81 images, notes, text/descriptions, order, hidden flags and media relationships match the snapshot. No cause or binary identity is claimed. Additive report: source-input-reaudit.json; course/index/training remain frozen.

Images are actual application captures or exact source copies. No LLM conversation is fabricated. Historical data retain their source attribution.

Capture reports and hashes: manual-build-manifest.json. Generator: tools/manual_build.py.

This manual does not operate TopSpin, modify TBO/KK2 or establish calibration or identity. The UF prototype retains quantitative NO-GO pending instrument controls.

Verified hybrid manifest: beatrix-academy-qwen3-1.7b-lora-v2-q8 · SHA-256: eb49a43f12de5057b9fbe41ab242fc8a7916bf9a36d15ce2cd733adfb8bd89ff

CPU report for LLM + course: 15/15 · SHA-256: 3dd224a85bc67219fe39d57f2e53e70cc532894c1defc1f18ffbf2378d11dbb3

Generated-only causal omissions in SPEN/TR remain documented; the course block is not model generation.