AUDIT / 1.3.1

Complete manual

Interpret every result with the correct estimator and evidence.

WPF and browser / English / 30 October 2026

This four-page supplement contains current 1.3.1 audit corrections. Read it before the retained manual, especially when comparing ILT distributions and Monte Carlo intervals.

The unchanged 1.3 document begins on PDF page 5. Its screens, titles and printed numbering are historical; they are not new 1.3.1 captures. Qualifications in this supplement supersede retained instructions.

02 Intervals and the inverse problem >

03 Data, order, sign and provenance >

04 Corrections and practical controls >

05 Open the retained 1.3 manual >

Intervals and the inverse problem

The interval belongs to the mono fit

Monte Carlo refits the reference monoexponential model. The interface identifies its D and interval separately. With NNLS, Tikhonov or L1 selected, the ILT centroid is a different estimator: the mono interval does not quantify that centroid or distribution shape, peaks or component count. ILT distribution uncertainty is not computed here.

The singular list does not establish uniqueness

A matrix with 20 observations and 64 D values has rank at most 20 and at least 44 right-nullspace directions. All 20 returned singular values may be nonzero while that space remains unresolved. The warning includes this limit; it neither counts molecules nor proves every positive solution impossible.

Why make this check

Small residuals and narrow intervals answer different questions. Compare regularization, grid, reserved observations and compatible alternatives. Keep processing and observations fixed when attributing differences to a method. A reproducible fit does not establish chemical identity.

Student objective

Before quoting an interval, name its estimator, shared errors and uncomputed properties. Do not assign an ILT centroid an interval obtained for a different fit.

Data, order, sign and provenance

A consistent reference

Per-signal processing and Monte Carlo use the minimum-b reference. Reordering the same (b, signal) pairs must not create or erase peaks by changing the reference row. This differs from physically reversing acquisition: a new ramp may expose real drift or kinetics.

Signed held-out controls

The control retains the signed reference and declared processing. Taking absolute values or clipping negative signals changes the problem. Inspect phase, signal and model before interpreting a difference as diffusion or separating species.

Import does not establish calibration

An imported series and metadata do not themselves supply an independent reference. Provenance, applied scale and available evidence remain distinct. A user-entered scale remains conditional; undocumented b values do not produce an invented D fit.

Valid inputs and stability

Empty/non-numeric gradients, non-finite reference/phase values and overflowing SI conversions are rejected with ES/EN explanations. They are not converted to zero. Nonzero-b fitting uses a minimum-b shifted reference; an unrepresentable extrapolated b=0 amplitude is reported unavailable, not NaN as a physical result.

Corrections and practical controls

Two specific teaching corrections

Example temperature-standard/factors: 0.89 x 293/301 = 0.86634551495, rounded to 0.8663 in ES/EN. Corrected D, 0.33 x 0.86634551495 = 0.2859 in 10^-9 m²/s units, is unchanged. The Spanish Kuchel 2012 [38] note now says g0=0, replacing g00; the mathematical equation is unchanged.

A practical error-detection exercise

1. Explain units and retain sign, phase and reference. 2. Compare mono and ILT on the same observations and read which estimator the interval covers. 3. Reverse paired rows and distinguish that test from reacquiring a ramp. 4. Enter an invalid value and verify rejection. 5. Import with provenance and design a control that can refute the conclusion.

What the checks establish

The audit compares examples, analytic limits, software contracts and ES/EN interfaces. The validator separates PPTX byte identity from fresh agreement of its 81 recorded slides. Originals and historical reports remain preserved. These checks are not physical calibration, chemical identification or an educational-efficacy study.

Beatrix, sources and scope

Documentation, executed training and approved native evaluation are distinct evidence. Mathematical guidance retains its Sinnaeve/Kuchel references; simplified UF/SPEN models do not certify a complete experimental sequence. Read bilingual notes and release reports for the exact scope.

Academy and release notes >

DIFFATONCE ACADEMY / 1.3.0

Complete manual

Learn to measure, interpret and defend diffusion NMR.

WPF and browser / English / 30 October 2026

PART I - Continuous research 1.3

Current instructions: twelve stages, one sample, shared uncertainty, controls, competencies, pilot and review.

APPENDIX A - Doctoral laboratory 1.2

Retained nine-activity material, derivations, data and limits. Historical labels are preserved.

APPENDIX B - Original course 1.0

24 lessons, examples and tools from the original course. Deeper study and historical-interface reference.

Part I is the current instruction reference. Appendices are not rewritten or relabelled as current screenshots. Documentary guidance, executed training and an approved model are different states.

Volume contents and navigation

Numbers below are PDF pages in this volume. Each part keeps the source document's printed numbering. Use links or bookmarks to move between parts.

Part I - Continuous research 1.3 **3**
How to learn and retain one sample 5
01 The question 7
02 Calibrate 8
03 Design 9
04 Acquire 10
05 Process 11
06 Interpret 12
07 Uncertainty 13
08 Next measurement 14
09 RF and space 15
10 Competencies and pilot 16
11 Defence and review 17
12 Public data 18
Reproducible continuous example 19
Beatrix and response modes 21
Instructor protocol and retention 22
Save, import and sources 23
Appendix A - Doctoral 1.2 / divider **25**
Appendix B - Course 1.0 / divider **48**

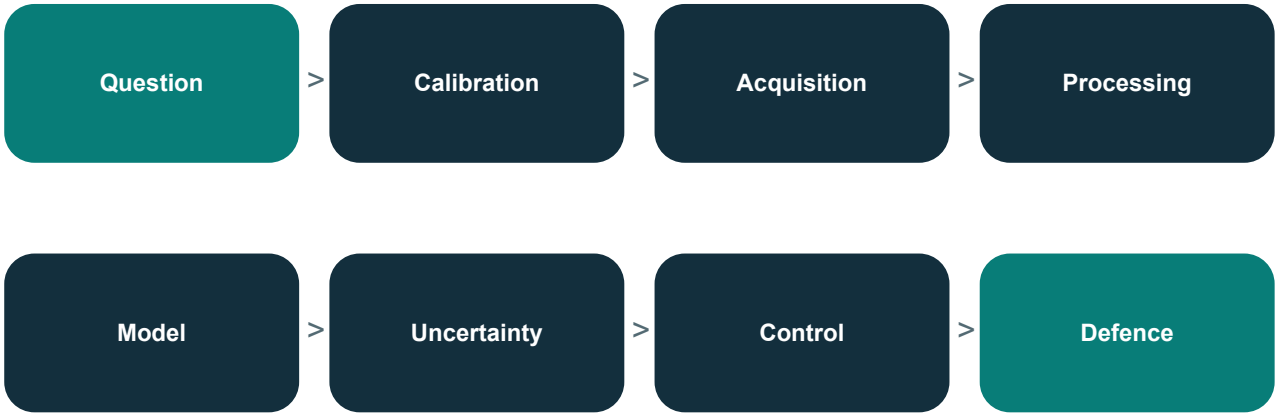
Each chapter retains provenance. Appendix web links reflect their original date; use Part I and the Academy product page for current installation/access.

DIFFATONCE ACADEMY / 1.3.0

From question to evidence

Continuous research manual

Twelve stages, one sample and decisions you can defend.



WPF and browser editions - English

30 / 10 / 2026

Computed teaching simulations, attributed public originals and human review. Models and inferences retain explicit conditions.

Contents and access

Web: open Academy and select Research. WPF: choose Research studio; change ES/EN in the application.

diffatonce.com/academy - English product page

The twelve interface stages

01 The question / 02 Calibrate / 03 Design / 04 Acquire / 05 Process / 06 Interpret / 07 Uncertainty / 08 Next measurement / 09 RF + space / 10 Competencies / 11 Defence and review / 12 Public data

Contents and access 2

How to learn with this version 3

What travels with the same sample 4

01 - The question 5

02 - Calibrate before measuring 6

03 - Design an informative measurement 7

04 - Acquire the same sample 8

05 - Process with explicit decisions 9

06 - Interpret without overclaiming 10

07 - Propagate shared errors 11

08 - Choose the next measurement 12

09 - RF, gradients and space 13

10 - Check what you learned 14

11 - Defend, review and improve 15

12 - Work with public originals 16

Continuous example - prepare and acquire 17

Continuous example - interpret and challenge 18

Beatrix, sources and response modes 19

Instructor protocol and independent evaluation 20

Save, import and reproduce 21

Sources, links and common problems 22

How to learn with this version

Turn a diffusion question into a conclusion proportional to the evidence.

What to do

1. Open Research in Academy web or choose Research studio in WPF, then select ES/EN. Both editions use the same interface and numerical engines.

2. Follow the twelve stages. Before an action, enter a prediction and reason in the help panel.

3. Use Doctoral, Derivations and Animations for deeper study, then return to the same case.

4. Export the investigation after each session. Then practise with another seed without consulting the worked example.

What it means

The seven extensions connect the sample, uncertainty, real data, competencies, control design, spatial encoding and tutoring/review.

Before continuing

At the end you should explain what was measured, why the equation applies, which biases remain and which control could refute your claim.

> Limit: Completing screens or passing software checks does not establish mastery or validate an experiment.

What travels with the same sample

Distinguish acquired data, decisions, assumptions and simulation truth.

What to do

1. The seed fixes a reproducible synthetic sample; it does not identify a real molecule.

2. Calibration, design, complex series, processing and fitting stay linked to the sample identity and history.

3. After changing an earlier decision, review and recalculate later stages. An earlier result does not establish the new design.

4. When public or CSV integrals are attached, evidence is labelled experimental/imported and provenance remains required.

What it means

Internal units are SI: D in m²/s, b in s/m² and gradient in T/m. The interface displays pulses in µs and sequence times in ms.

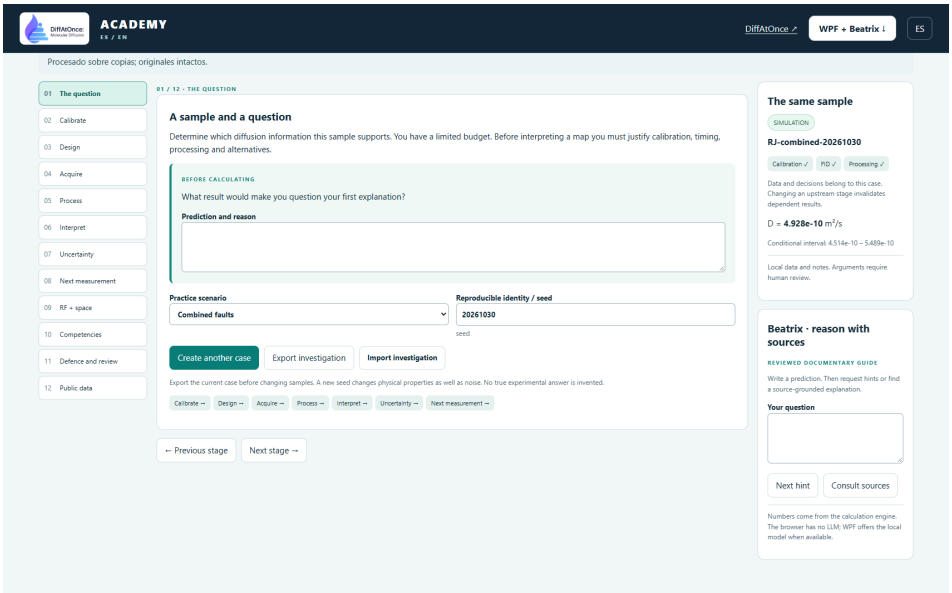
Before continuing

Check identity, units and processing state before comparing results.

> Limit: Experimental originals are unchanged. Virtual calibration does not automatically recalibrate imported data or the spectrometer.

01 - The question

State a hypothesis and a result that could contradict it.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Select Hidden difficulty and Reproducible identity; use Create another case only to start a different sample.

2. Start with Isolated signal, then try Gradient calibration, Convection, Time drift, Overlap or Combined faults.

3. Write a concrete question: for example, whether a D difference might follow ramp order rather than only two species.

4. Record predictions for each hypothesis before acquiring.

What it means

A cause is not established by choosing a word. Decisions and controls must support the argument.

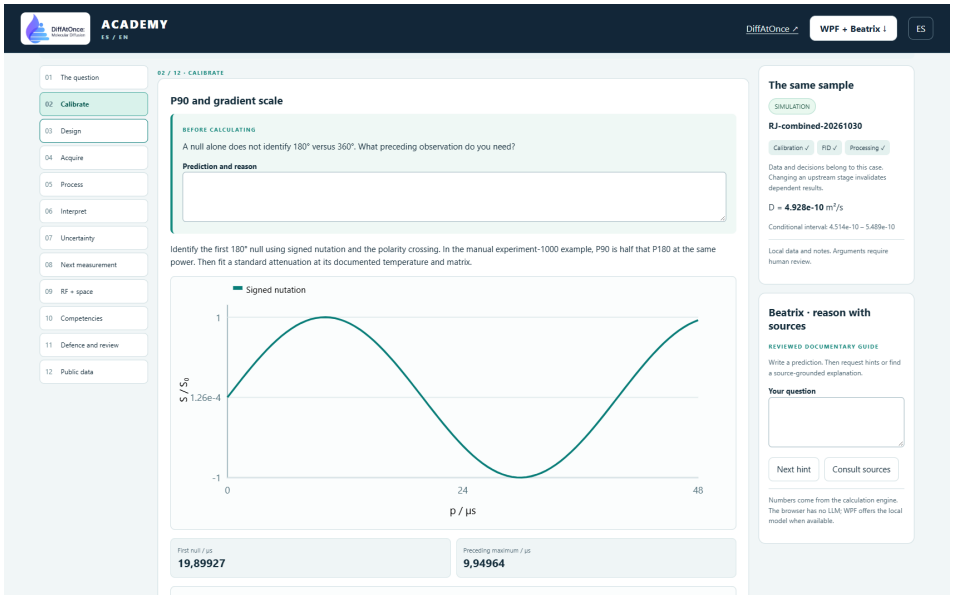
Before continuing

Your question includes a measurable observation, an alternative and a possible control.

> Limit: Choosing a preset reveals the problem family. For evaluation without that cue, use instructor-reserved cases/questions; guided practice is not a fully blind assessment.

02 - Calibrate before measuring

Separate RF calibration, gradient scale and standard conditions.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Inspect the signed nutation curve. Identify the first 180° null using the preceding maximum and sign reversal.

2. Enter P90=P180/2 at constant power and conditions. A 360° null is divided by four; a null alone does not establish its order.

3. Fit the standard decay against nominal b; for one common scale g_actual=s·g_nominal use s=√(D_nominal/D_ref).

4. Check standard and reference temperatures, then press Apply calibration to case.

What it means

With nominal b, a gradient-scale error enters D quadratically. More scans do not remove this bias.

Before continuing

Justify the chosen null, reference D, temperature and waveform covered by calibration.

> Limit: P90 must be checked for the appropriate sample and power. p30, d20 and d5 require the actual pulse program; no universal translation exists.

03 - Design an informative measurement

Choose an attenuation window and controls within a budget.

What to do

1. Choose Spin echo, Stimulated echo, LED or Double stimulated and explain why.

2. Set δ, Δ, d5/LED, maximum gradient, levels, NS, replicates and ramp order.

3. Press Check design and budget; inspect duration, maximum b and restrictions before acquiring.

4. Predict the effect of longer Δ or reverse/interleaved order. Change one variable at a time for causal comparisons.

What it means

More weighting can improve contrast while losing signal. LED adds time; losses and compensation depend on sequence.

Before continuing

You can explain the role of each level and the part of the budget used to control errors.

> Limit: Timing is an ideal teaching model. Check the actual program and instrument limits before transferring a recipe to the laboratory.

04 - Acquire the same sample

Connect the design to a complex series and acquisition order.

What to do

1. After calibration and design, press Acquire virtual series.
2. Inspect the FID and acquisition count; retain sample identity and seed.
3. Compare a replicate with reversed/interleaved acquisition. Describe the expected change under drift.
4. Export the first investigation before starting another seed.

What it means

Synthetic data contain reproducible noise and sample-specific difficulties; the FID is a computed signal.

Before continuing

Signal, temporal order and b remain paired in your analysis.

> Limit: This action does not command TopSpin or acquire a physical sample. Experimental data enter through Public data or import.

05 - Process with explicit decisions

Preserve the sign and check phase, window and baseline dependence.

What to do

1. Set Phase correction, Subtracted offset, apodization and Signal window.
2. Press Process and integrate. Inspect the spectrum and integral decay.
3. Compare Whole spectrum, Signal 1 and Signal 2 when overlap exists.
4. Explain why a different window or phase may alter decay without changing physical D.

What it means

An integral is a decision about complex data. Taking only magnitudes or clipping negative values can change noise and bias the fit.

Before continuing

Your report allows phase, window and baseline processing to be reproduced.

> Limit: For real Bruker files, inspect digital filtering, reference, sequence and metadata in the experimental workflow. Do not transfer indices from another sample.

06 - Interpret without overclaiming

Compare models on exactly the same observations.

What to do

1. Select Monoexponential, NNLS, Nonnegative Tikhonov or Nonnegative L1, then Fit the same data.

2. For ILT, record λ and the grid; compare residuals, stability and peaks rather than only appearance.

3. Keep processing fixed to attribute differences to the method. Then examine processing sensitivity separately.

4. Return to the Doctoral laboratory for the ILT kernel, correlations and joint analysis.

What it means

A well-fitted curve does not determine a unique distribution. A penalty imposes a preference and may merge or create apparent peaks.

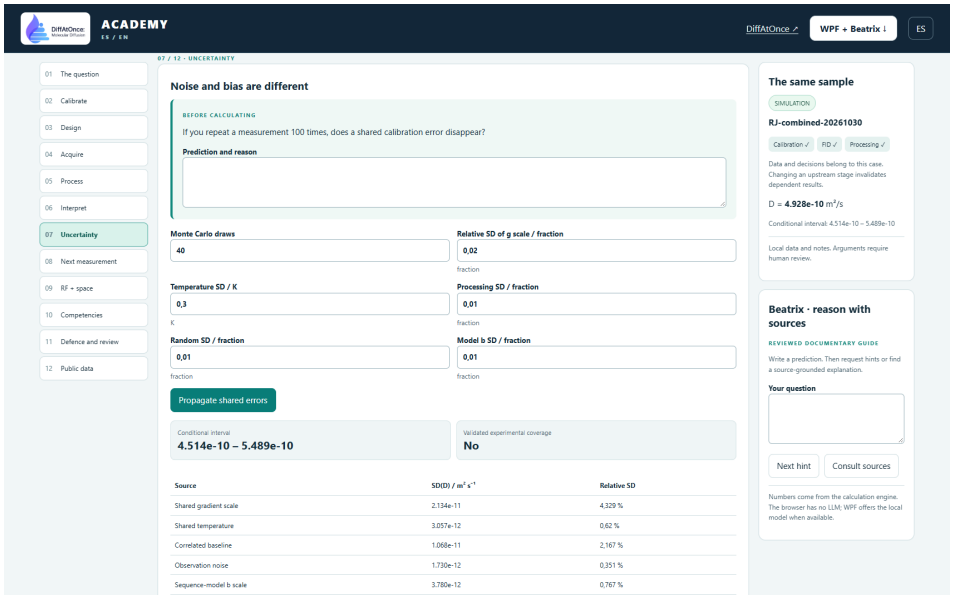
Before continuing

You can distinguish numerical fit, resolution, distribution and chemical assignment.

> Limit: A mixture's monoexponential D is conditional on the model. Do not assign absolute molar masses or species from the map alone.

07 - Propagate shared errors

Separate random variation from biases and common dependencies.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Enter Monte Carlo draws and SDs for g scale, temperature, processing, noise and model b.
2. Press Propagate shared errors and compare contributions and the conditional interval.
3. Repeat after reducing only noise, then after reducing calibration uncertainty. Compare the effects.
4. Describe which sources share a draw across gradients/signals and which vary by observation.

What it means

One calibration can shift several D values together. Treating it as independent noise would manufacture information that was not acquired.

Before continuing

Justify a control that reduces a dominant source and list errors your interval does not quantify.

> Limit: The interval is conditional on the supplied models and distributions. It is not validated experimental coverage and does not automatically correct a wrong model.

08 - Choose the next measurement

Use ambiguity to design a discriminating control.

What to do

1. Compare Higher b range, Repeat the design, Reverse order, Change Δ and Flow compensation.

2. Write a prediction before pressing Compare designs.

3. Choose an acquisition and press Observe held-out control. Held-out observations use an independent realization.

4. Compare the prediction with results and explain which alternatives remain.

What it means

Two distributions may predict the available data similarly and differ over another b region. The proposed acquisition explores that difference.

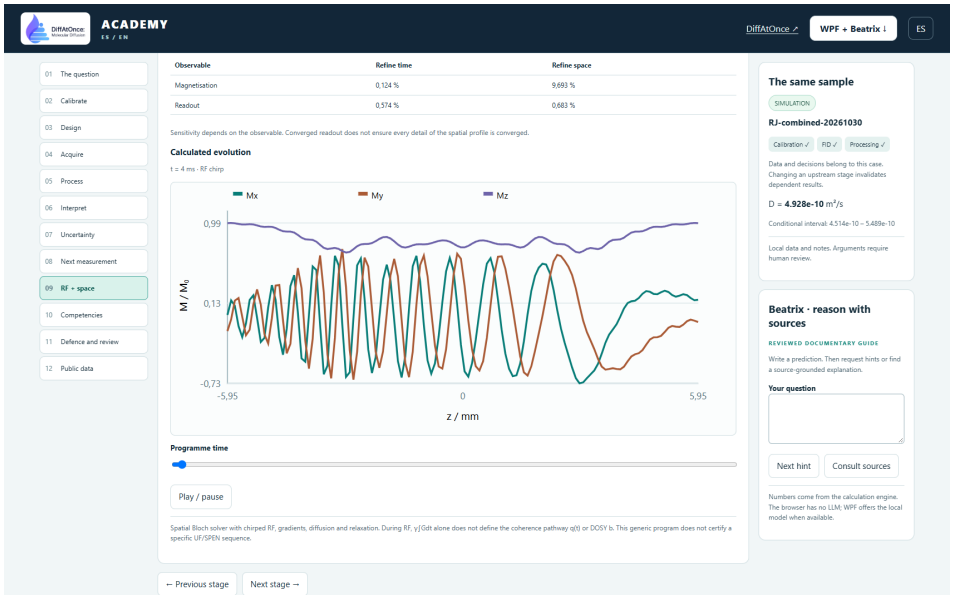
Before continuing

You can defend why this control discriminates better than more equivalent points.

Limit: Synthetic design selection and held-out data are not external validation of a real sample. Ideal flow compensation is not equivalent to every double-STE program.

09 - RF, gradients and space

Understand the physical origin of spatial encoding and its numerical limits.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Set spatial cells, length, RF duration/amplitude, chirp sweep, gradients, evolution, readout and coil profile.

2. Press Calculate Bloch-Torrey; inspect magnetization, phase, transverse profile and readout.

3. Press Check convergence to compare refined time step and grid. Review aliasing warnings.

4. Compare with the introductory ultrafast/SPEN model under Doctoral: explain what RF-based computation adds and what remains ideal.

What it means

The model is dM/dt = Ω×M + D·d²M/dz² - (Mx/T2, My/T2, (Mz-1)/T1). Ω contains RF and gradient with time-integrated chirp phase. Reflecting walls, Crank-Nicolson diffusion and exact per-step rotations are used. Readout sums coil-weighted complex magnetization.

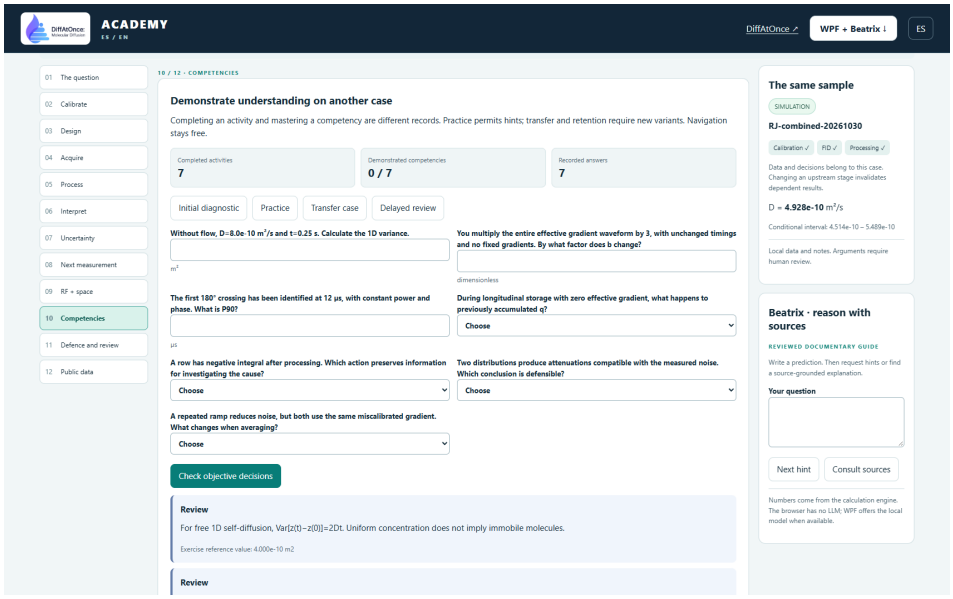
Before continuing

You can distinguish RF phase, physical gradient integral, nominal resolution and discretization error.

> Limit: During RF, γ∫Gdt is not coherence q or DOSY b. With linear readout, an offset Δf causes Δz_apparent=Δf/(γHz·G): not molecular motion. The generic protocol does not qualify a particular UF/SPEN-DOSY sequence.

10 - Check what you learned

Demonstrate transfer and retention as well as activity completion.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Take Initial diagnostic before consulting hints. Use Practice for outstanding concepts and record assistance used.

2. Complete Transfer case with new parameters and inspect each competency state. Return to a case that requires it.

3. The voluntary teaching pilot starts empty. To participate, use an optional local code, mark consent and press Consent and open the initial case. Pretest and posttest contain seven questions with different seeds; declare whether you answered without assistance.

4. The retention case becomes available after seven days, on the displayed date; there is no button to advance that interval. Use Check the current date and Export pilot results. No learners or results are preloaded.

What it means

Automatic checks are limited to explicit choices and numbers. Written reasoning requires human review.

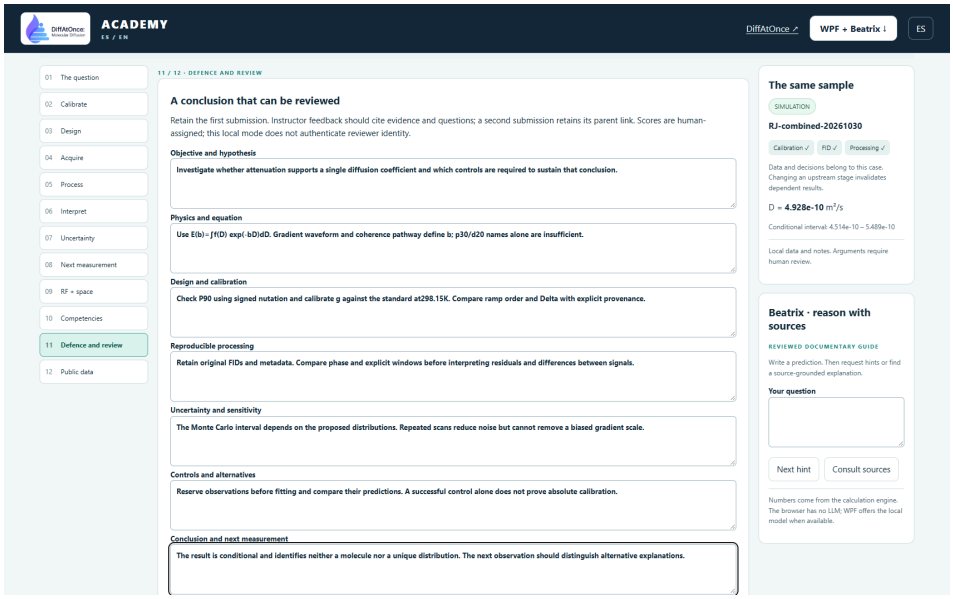
Before continuing

Solve a new situation without a hint and explain the rule you applied.

> Limit: The application does not certify doctoral training or infer understanding from text length. Educational efficacy requires students and independent assessment.

11 - Defend, review and improve

Keep reviewable reasoning and a traceable second version.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Complete objective, physics, design, processing, uncertainty, controls and conclusion.

2. Press Save a new submission. Export portfolio preserves the investigation, evidence and submissions.

3. The instructor uses five 0-4 criteria, cites evidence, writes comments/questions and presses Create and export feedback.

4. The student uses Import feedback, responds to questions and saves another submission without replacing the first.

What it means

Feedback is linked to a specific submission and identity; another case's review must not become its assessment.

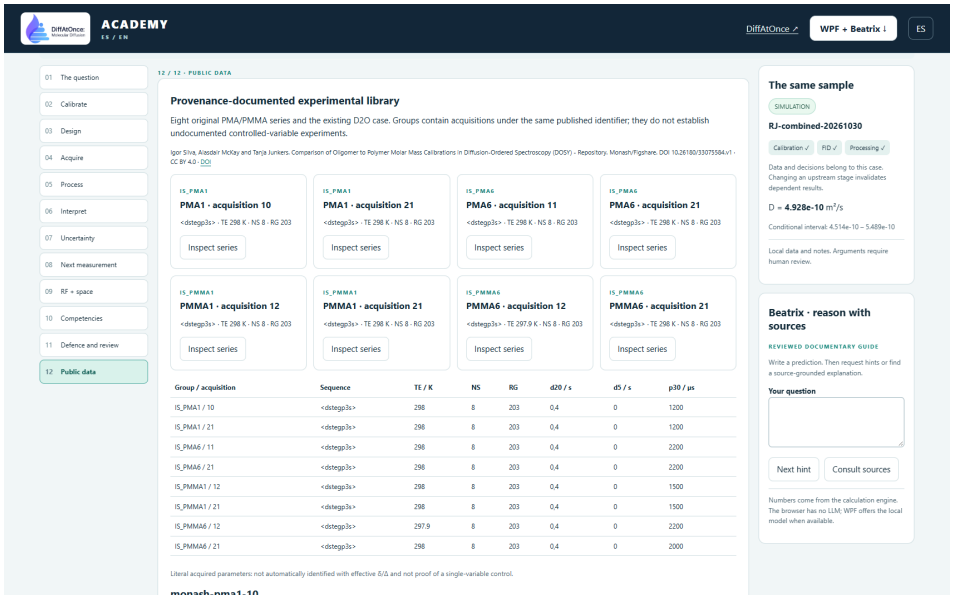
Before continuing

Your second version explains what changed, which evidence persuaded you and what remains unresolved.

> Limit: Scores and judgment are the instructor's. Integrity checks do not authenticate civil identity or replace an institutional grade system.

12 - Work with public originals

Process real experiments and recognize missing information needed for D inference.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Open Public data and select a PMA1, PMA6, PMMA1 or PMMA6 acquisition. Each group has two repository acquisitions.

2. Loading checks bytes and SHA-256 before interpreting FIDs/metadata. Inspect program, saved temperature, digital filter, NS and RG.

3. Set phase, apodization and window; press Process public series. Supply one b per acquisition and its provenance before fitting.

4. Press Continue investigation with these integrals to attach them to the case. Document which conditions permit comparison of the two acquisitions.

What it means

The eight datasets come from Silva, McKay and Junkers' Monash/Figshare repository, CC BY 4.0. Saved conditions alone do not establish temperature or convection controls.

Before continuing

Distinguish missing verified b, nominal scale, absolute calibration and a D conditional on processing.

> Limit: These eight datasets do not supply sufficient saved DOSY parameters to justify b automatically. Do not invent this axis, a true D or absolute molar mass.

Continuous example - prepare and acquire

Reproduce a synthetic investigation with checkable parameters and outputs.

What to do

1. Under The question, choose Combined faults and seed 20261030; create case RJ-combined-20261030.

2. The first 180° null is 19.8993 µs and the preceding maximum is at 9.94964 µs. Enter P90=9.94964 µs.

3. Use D_ref=1.9×10^-9 m²/s at 298.15 K. The nominal curve fits D=2.20719e-09 m²/s: s=√(D_nominal/D_ref)=1.07781.

4. Enter this scale, both temperatures 298.15 K, and apply calibration. Design STE: δ=2 ms, Δ=80 ms, LED=5 ms, g_max=0.3 T/m, 20 levels, NS=8, one replicate, recycle=2 s, ascending order and a 15 min budget.

5. Check duration approximately 6.352 min and b_max=2.37445e+09 s/m²; acquire the virtual series.

What it means

Values are computed by the engine shipped with this edition. Displayed figures are rounded; export retains available precision.

Before continuing

Compare against g scale=1 in an exported copy while keeping other decisions fixed.

Limit: This example is not an operating prescription for your spectrometer or an accepted experimental result.

Continuous example - interpret and challenge

Reach a reproducible result that still needs controls.

What to do

1. Process with phase=0°, offset=0, LB=0 Hz and Whole spectrum window. Fit Monoexponential to the same data.

2. The example returns D≈4.92791e-10 m²/s and residual RMS≈0.0198293. A reproducible number does not turn combined faults/a mixture into one species.

3. Propagate 40 Monte Carlo draws: relative g SD=0.02, T SD=0.3 K, processing=0.01, noise=0.01 and sequence b-scale=0.01.

4. The illustrative conditional interval is [4.51375e-10, 5.48915e-10] m²/s. Compare contributions before choosing what to reduce.

5. Predict the effect of Reverse order and observe the held-out control. Write what the difference explains and does not explain, then prepare a submission.

What it means

The example shows why an estimate, its uncertainty and a causal explanation are different objects.

Before continuing

Try another seed without copying these values: the justification should transfer to the new sample.

> Limit: After changing version, settings, processing or seed, do not force agreement with this table. Save and explain the new identity.

Beatrix, sources and response modes

Use traceable help without replacing your prediction or calculation.

What to do

1. Enter Prediction and reason first. Use Next hint for progressive help linked to the stage.

2. Enter Your question and press Consult sources: documentary retrieval returns reviewed passages and references.

3. In WPF, Ask the local model requests available inference with bounded context. Always read the mode label and distinguish documentary guidance from generated text.

4. Check numbers in the calculation engine and compare the answer with cited sources.

What it means

The web edition uses documentary guidance, not an LLM running in PHP. Windows uses a model only when its manifest and CPU evaluation are approved; otherwise documentary guidance is provided. A trained candidate does not automatically become the active model.

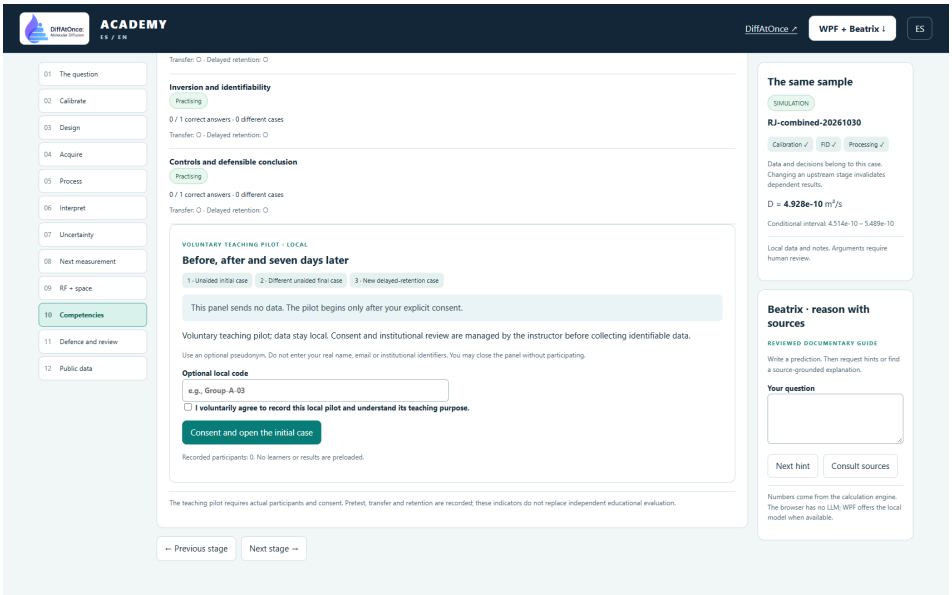
Before continuing

You can identify source text, calculated output and any actual generated text.

> Limit: Documentary retrieval, executed training and an approved model are different states. Fluency does not validate an equation, calibration, assignment or student assessment.

Instructor protocol and independent evaluation

Determine whether the pathway improves decisions in new situations.



Actual shared-interface capture. The steps below identify controls; zoom to inspect small labels.

What to do

1. Before the course, reserve transfer cases/questions not used to train the tutor or select settings.

2. Record the initial diagnostic, decisions, hints, predictions and first submission.

3. Two instructors can review reports with the same rubric, compare disagreements and document the second version.

4. Use a different final case and delayed review. Compare physics, controls, uncertainty and conclusion quality, not just multiple-choice answers.

What it means

The Competencies panel runs the local pilot with consent, different initial/final cases and delayed retention. Argument review is human. This manual does not claim results from a cohort that has not been evaluated.

Before continuing

Predefine a meaningful improvement and retain failures and negative results too.

> Limit: Software checks establish computational behavior and synthetic controls. They do not replace expert review or independent educational evaluation.

Save, import and reproduce

Preserve evidence while distinguishing local state, originals and results.

What to do

1. Export investigation/portfolio produces version 1.3 JSON: sample, decisions, results, learning and submissions. The interface portfolio omits raw FIDs: keep originals separately.

2. Import investigation restores a compatible portfolio. To reprocess, reload original series or reacquire the reproducible simulation. Import feedback requires a matching submission and identity checks.

3. CSV import requires finite b/signal pairs and correct units. Importing an unverified local source does not create a license or calibration.

4. For reproduction, record version, seed or hashes, settings, b sources, processing, method and controls. Keep the portfolio outside browser cache.

What it means

Storage is local. Browser changes, cache clearing or another profile can change visible progress; export is the portable copy.

Before continuing

Open an export in an isolated session and check identity/results before considering a submission complete.

Limit: JSON has size/compatibility limits. Do not edit hashes to force import. The web does not synchronize notes/data to an instructor account.

Sources, links and common problems

Find the equation, dataset and support relevant to each claim.

What to do

1. Sinnaeve (2012), DOI 10.1002/cmr.a.21223, and Kuchel et al. (2012), DOI 10.1002/cmr.a.21241, underpin the mathematical pathway. Use the 20 Derivations chapters for symbols, pages, equations and assumptions.

2. Torrey (1956), DOI 10.1103/PhysRev.104.563: diffusion in Bloch equations. Shrot, Shapira and Frydman (2004), DOI 10.1016/j.jmr.2004.08.001: continuous spatial encoding.

3. Silva, McKay and Junkers, DOI 10.26180/33075584.v1, CC BY 4.0: source of public experimental files; retain attribution.

4. If public data are absent, check package/version access. If b is absent, document its derivation; for aliasing, review grid/readout. If saving fails, export before closing.

What it means

This manual supplements the retained doctoral manual and original course with their explicit sources and conventions.

Before continuing

Every final conclusion identifies data, model, conditions, uncertainty and limitations.

> Limit: A file group, good fit or attractive screenshot does not by itself validate a physical explanation.

Retained doctoral manual

Sinnaeve 2012

Kuchel et al. 2012

Torrey 1956

Shrot et al. 2004

Monash / Figshare

CC BY 4.0

APPENDIX A

Doctoral laboratory

Retained source document - version 1.2

The next 22 pages are the doctoral manual 1.2, dated 5 October 2026. Its nine activities remain available under Doctoral.

Deepen movement, phase/echo, q(t), derivations, campaigns, blind cases, ILT, correlations, public FIDs and introductory UF/SPEN models.

Pages are appended without changing content, equations, figures or scientific boundaries. Printed numbering restarts in the source document; the reader toolbar and bookmarks locate pages in this volume.

For current instructions, export, the new pilot and model-selection rules use Part I. A model is used only with an approved manifest and CPU evaluation; otherwise documentary guidance is provided.

Open retained source document >

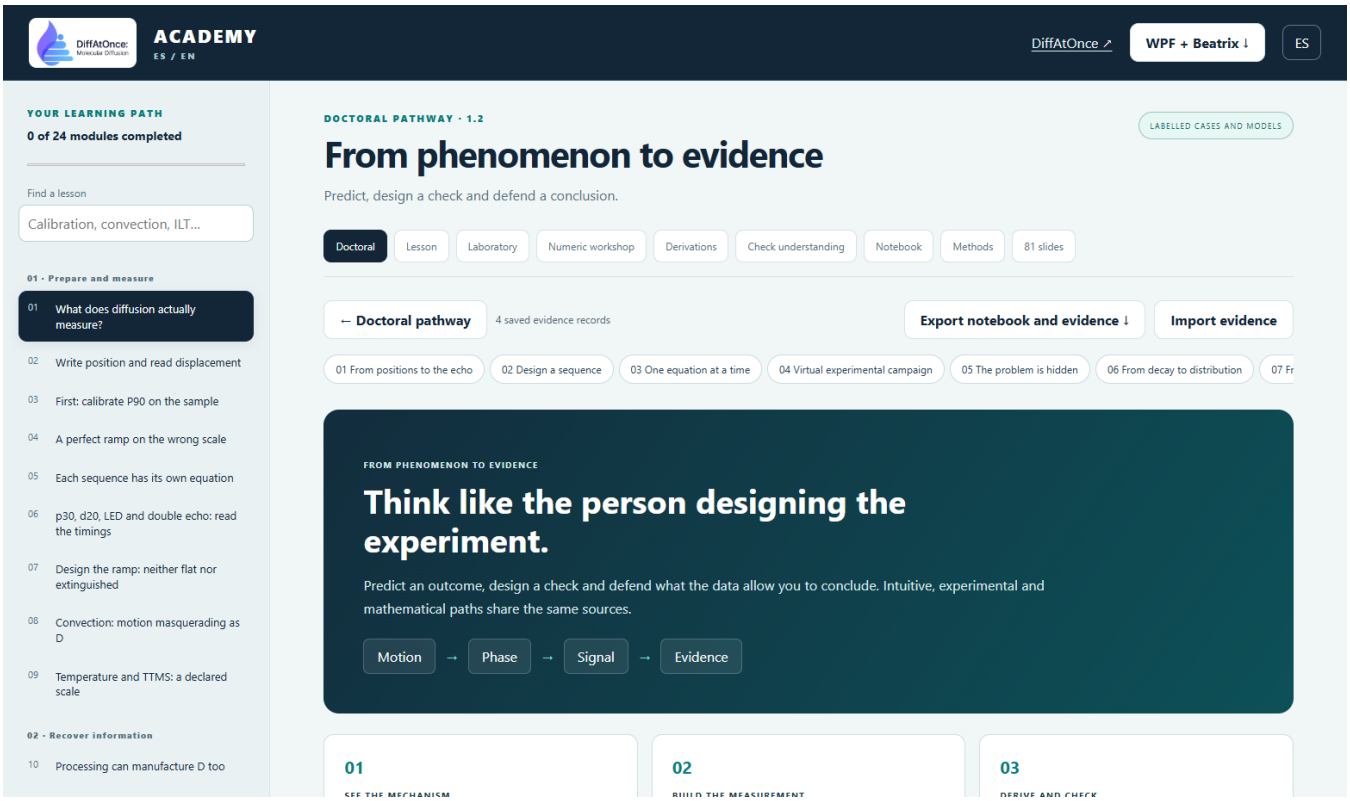
Return to current instructions >

Learn to measure, interpret and defend diffusion NMR

Complete doctoral laboratory manual

Version 1.2.0 - English - 5 October 2026

Nine activities connect molecular motion, phases, sequences, calibration, inverse problems and scientific conclusions. The goal is to design a measurement and explain why its evidence supports or limits your claim.



Actual doctoral pathway interface. Figures come from application checks. The Monash case is identified as conditional real evidence.

How to use this manual. Follow the route on pages 2-3, then complete the nine activities. Write predictions before running and retain controls, assumptions and conclusions in the notebook.

Scope. Virtual experiments are synthetic teaching models. The Monash case contains published real signals; its D estimate remains conditional on metadata and calibration. The application does not control a spectrometer.

01 | Learning route and objectives

You do not need to solve all the theory before experimenting. Begin with a physical prediction and return to the derivation when you need to justify a result. Three guided levels let you explore intuition, experimental design and mathematics.

The nine activities

1. sync: connect molecules, phases, magnetization vectors and the echo.

2. sequence: construct q(t), calculate b and sensitivity to flow.

3. guided: complete predictions, units and derivation steps.

4. campaign: calibrate and design a ramp under a time budget.

5. blind: distinguish causes through independent controls.

6. inverse: compare fits and inversion, spectral relationships and uncertainty.

7. cases: process real or synthetic data with explicit provenance.

8. advanced: study a reaction using sequential acquisition and spatial encoding.

9. assessment: defend a conclusion and review it with a human rubric.

What you should be able to explain

Why self-diffusion occurs at equilibrium; why an echo attenuates even when a gradient is refocused; what b measures; how gradient errors bias D; when two populations are resolvable; which controls distinguish convection, processing and changing composition.

02 | Open the application and retain your work

Windows WPF

1. Extract the entire portable ZIP. Keep the executable together with Content, Web, Lab, Models and Docs. Do not run from inside the ZIP.

2. Open **DiffAtOnce.Academy.exe**. The **Doctoral laboratory** button opens the laboratory; switching ES/EN retains your active reasoning. **Doctoral manual (PDF)** opens the manual matching the language.

3. The embedded reader requires Microsoft Edge WebView2 Runtime. If unavailable, install the official runtime or use the browser edition. Content, animations and equations are local.

Browser edition

Open Academy on the available deployment and select **Doctoral**. The browser edition includes the educational laboratory without running an LLM. Public availability of a version is separate from this manual and should be checked on the download page.

Notebook and export

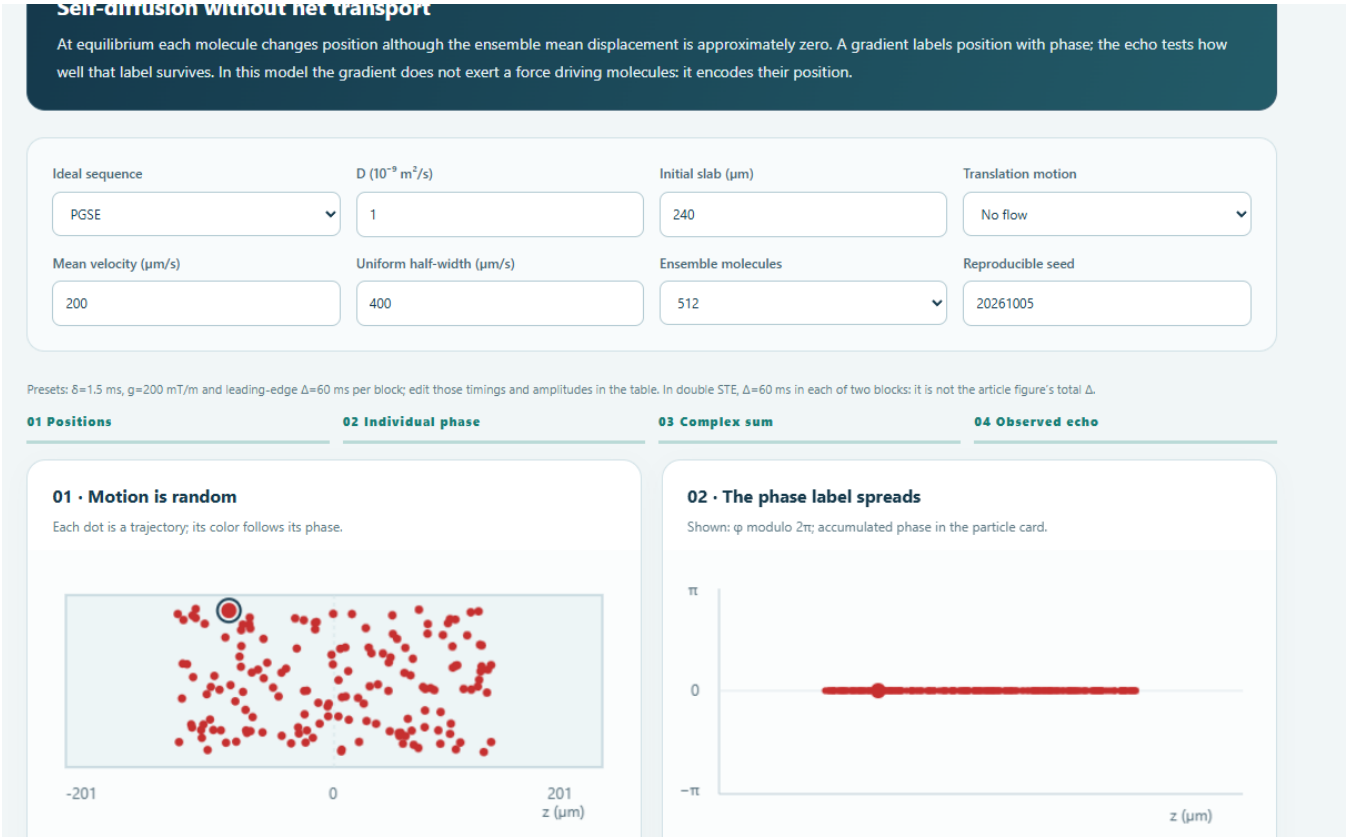
Activities retain predictions, parameters, decisions and evidence across languages. Export the notebook as a JSON record; save a copy before closing, changing computers or clearing browser storage. Imported data may need reimporting after leaving its activity.

The record retains provenance, hashes, processing and conclusions; it does not turn a simulation into experimental validation and does not publish your FIDs. If storage is blocked or full, export before closing.

Microsoft WebView2 Runtime - official documentation

03 | Molecules, phase and echo: one connected story

Objective. Explain echo loss as a phase distribution, rather than molecules disappearing.



The sync activity: view displacement, phase and signal together. Follow one molecule, then consider the ensemble.

1. Predict what happens if molecules remain stationary between encoding and decoding.

2. Run the model and relate the selected position to phase and its complex vector.

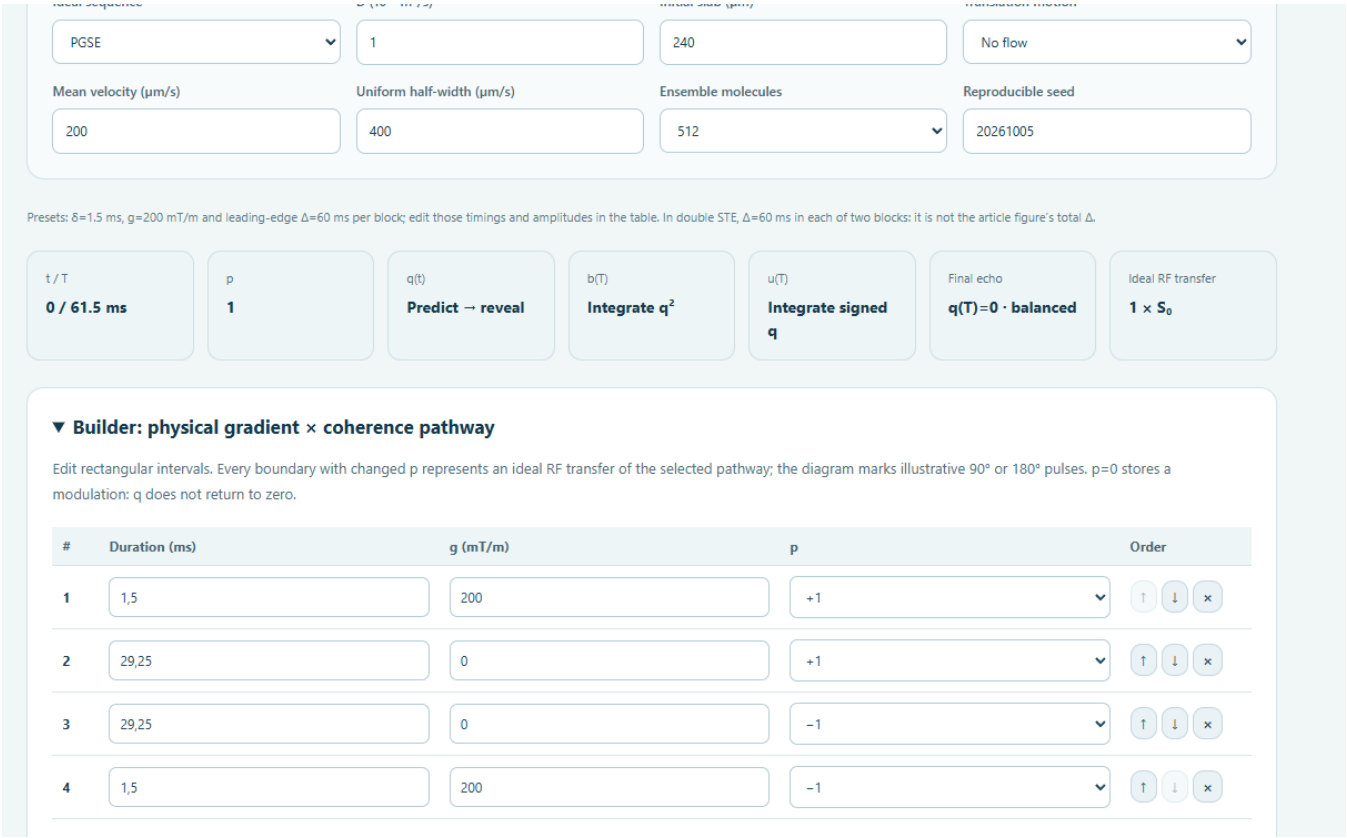
3. Increase D or the diffusion interval while holding other parameters fixed. Observe phase dispersion and the magnitude of the sum.

4. Add uniform flow or a velocity distribution. Distinguish an ensemble phase shift from amplitude loss caused by dispersion.

$$E = \langle \exp(i\varphi) \rangle \; ; \; |E| = |\langle \cos \varphi \rangle + i \langle \sin \varphi \rangle|$$

At equilibrium, macroscopic fluxes can cancel while every molecule continues moving. Self-diffusion concerns displacements and correlations; uniform concentration does not imply D = 0. The free-motion model does not establish restricted diffusion or chemical exchange.

04 | Sequence construction and q(t) memory



The sequence activity: accumulated history determines weighting, including longitudinal storage.

1. Enter RF, gradient and time segments using the interface units. Before revealing the calculation, draw your prediction of q(t).

2. Follow coherence order p(t). RF changes the effective gradient p(t)g(t); it does not automatically erase q history.

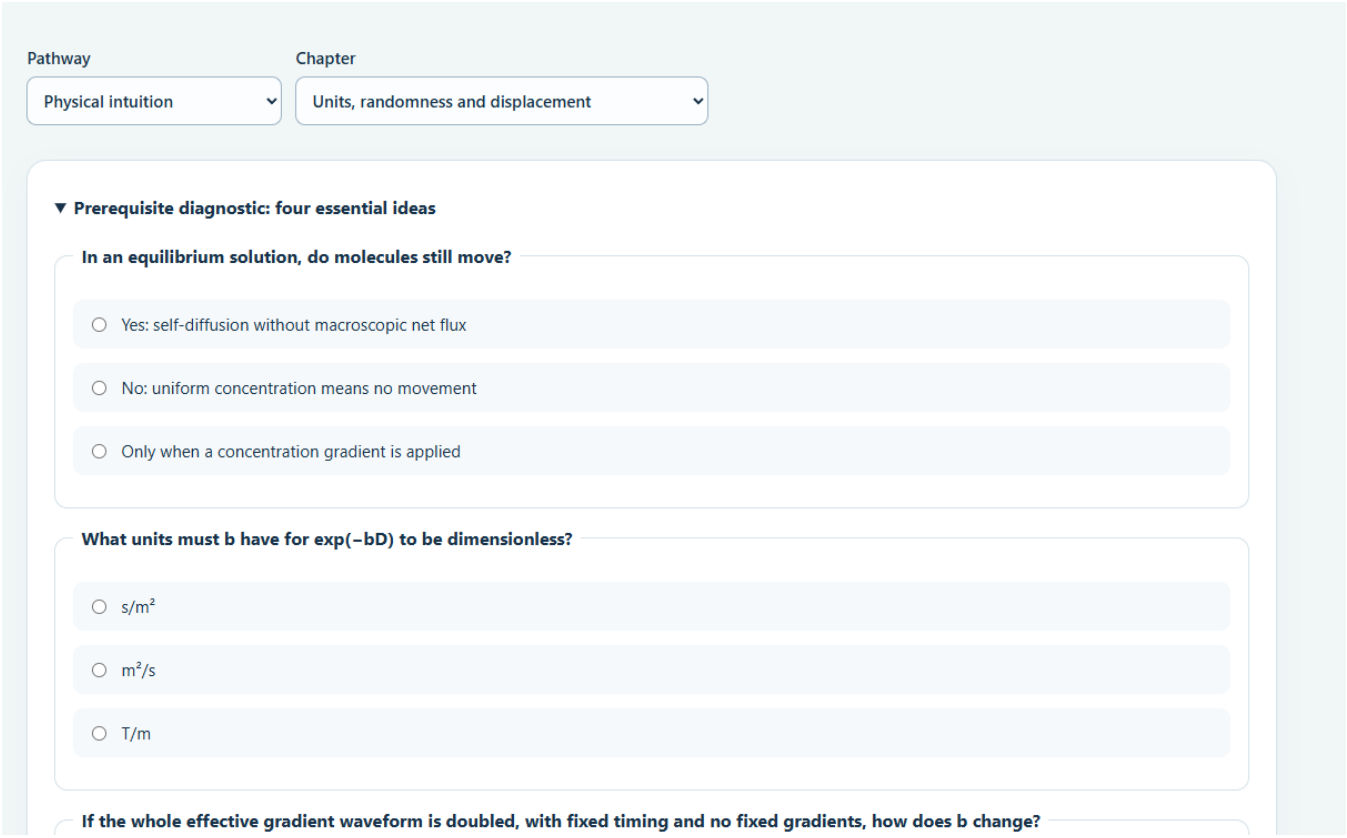
3. Compare the area under q^2 with the first moment $\int q\,dt$. A balanced echo q(T)=0 does not imply zero flow sensitivity.

4. Change a duration, gradient shape or separation. Verify that the selected equation corresponds to the entire effective temporal waveform.

$$q(t) = \gamma \int_0^t p(s)g(s)\,ds$$
$$b = \int_0^T q^2(t)\,dt \; ; \; M_1 = \int_0^T q(t)\,dt$$

Use the actual pulse program to interpret p30, d20, d5 and coherence-selection gradients. The same parameter label does not establish a universal effective interval. The application builds teaching topologies; it does not install a TopSpin sequence or validate its execution.

05 | Complete the derivation rather than memorize it



The guided activity: each transformation includes a prediction, check or question about its assumptions.

The collection includes 20 chapters, 93 steps and the 111 numbered equations in Sinnaeve and Kuchel. There are 79 objective tasks, 16 numerical checks and 14 documentary steps requiring human review. Coverage identifies DOIs and PDF/journal pages; notation discrepancies are explained without changing originals.

1. Choose intuition, experimental or mathematical level. Read variable meanings and validity conditions.

2. Write the prediction or complete the task before requesting a hint. Hints are progressive and reviewed.

3. Solve objective choice, units or numerical tasks. Complex algebra requires human review; saving an answer does not establish a correct proof.

4. Connect the equation to an experiment: which observable changes, which parameter must be measured, and which control could invalidate the assumption?

Bloch-Torrey combines magnetization evolution, diffusion and spatial phase. Scalar E = exp(-Db) requires the model assumptions and correctly calculated weighting; it is not a universal law for every sample.

06 | Calibrate before selecting the ramp



The campaign activity: identify the RF null, apply the gradient reference and repeat to inspect the effect.

Manual P90 calibration

The virtual nutation series has a maximum at 9 µs, a 180° null at 18 µs and a 360° null at 36 µs. The 180° null crosses polarity; P90 = P180/2. If choosing the 360° null, P90 = P360/4. A single null does not identify its order.

Select duration and order, apply them and inspect sensitivity. In the real workflow, experiment 1000 may contain your manual null search; verify what p1 represents before dividing. RF power, tuning, sample and channel must match the experiment.

Gradient reference

$$s = \sqrt{(D_{nominal} / D_{ref})} \; ; \; g_{real} = s \cdot g_{nominal}$$

The activity fits a synthetic standard series in signal space. The standard and its D must match temperature and matrix: a temperature mismatch prevents applying calibration. The square law assumes the entire effective waveform is scaled; fixed gradients and cross terms require the full history.

07 | Design δ, Δ, sensitivity and time

1. Predict whether the ramp provides adequate attenuation using b_max·D. Very little attenuation adds limited information; a tail buried in noise does not resolve components either.

2. Select δ, Δ and maximum gradient. Change one at a time and inspect attenuation, residuals and absolute signal.

3. Compare PGSE, STE, STE + LED and ideal mirrored Double-STE loops. Review the model timing definitions.

4. Set points, repetitions, NS, d1 and acquisition time. Check the budget; more data cost time and the sample may evolve during acquisition.

5. Submit a conclusion covering calibration, b range, SNR, time and controls. Then reveal the generating parameters to assess bias.

$$b_{PGSE} = \gamma^2 g^2 \delta^2 (\Delta - \delta/3)$$

$$\text{Time} = \text{points} \times \text{repetitions} \times \text{NS} \times (\text{d1} + \text{AQ} + \text{sequence duration})$$

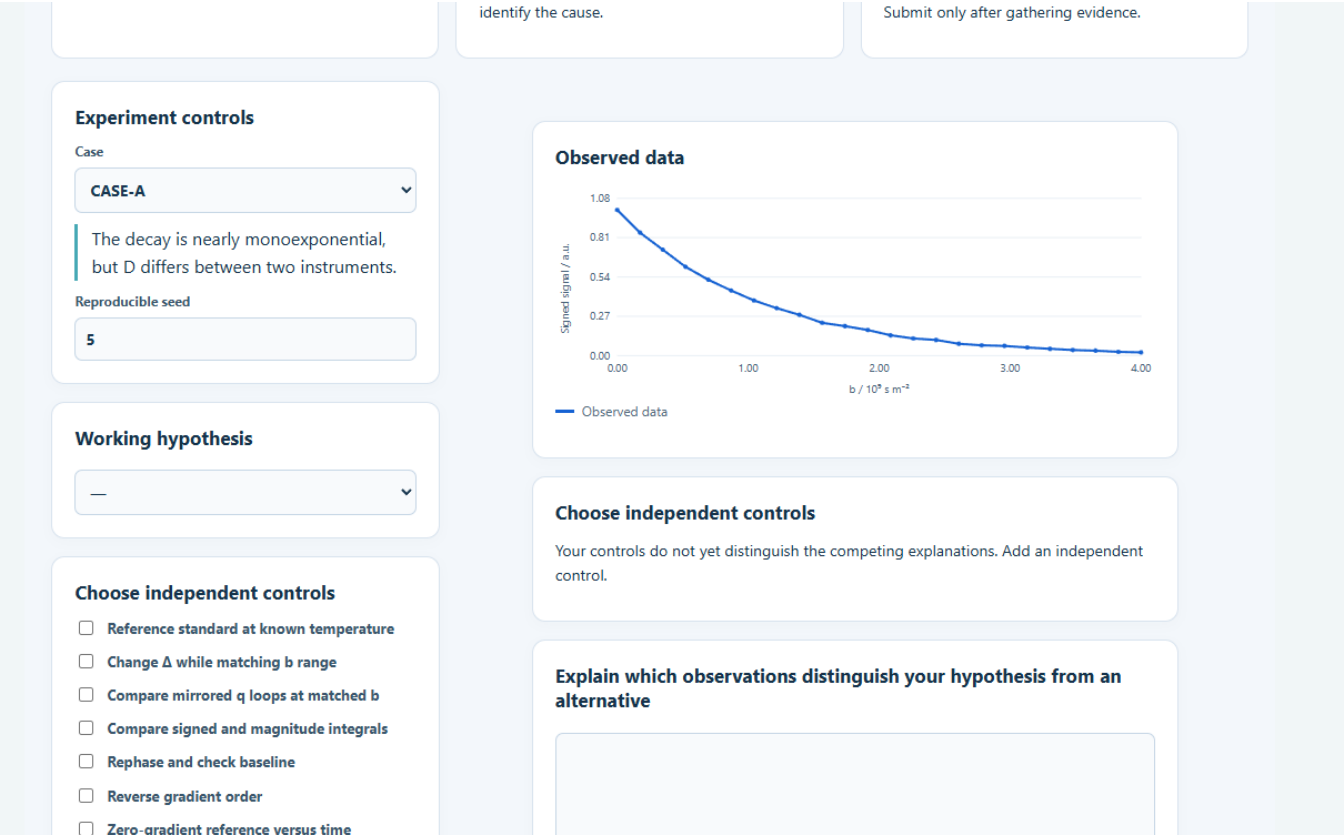
Virtual recovery is 1 - exp(-d1/T1). Ideal selected-pathway RF transfer is 1 for PGSE, 1/2 for STE/LED and 1/4 for two stimulated echoes. Compensation has a sensitivity cost. Constant incomplete recovery can preserve normalized D while reducing precision. Remaining T1/T2 and RF-error factors are teaching proxies, not complete coherence-pathway simulations.

Double-STE uses two identical mirrored q loops: b = 2b_loop and the ideal uniform-flow first moment cancels. Δ belongs to each loop. LED is represented by added time; it is not claimed to automatically remove convection.

Conclusion to justify

Explain what increasing Δ gains and what relaxation, convection or temporal evolution could cost. A useful ramp answers your scientific question with sufficient evidence under its constraints.

08 | Hidden cases: select informative controls



The blind activity: the initial curve does not disclose its cause. Review follows a hypothesis, controls and an evidence explanation.

Five cases explore gradient scale, convection, phase, temporal change and recovery. The model is concealed for teaching; this is not a secure examination. A curve alone does not establish a unique diagnosis.

1. Choose a hypothesis and an alternative. Predict which control would differ between them.

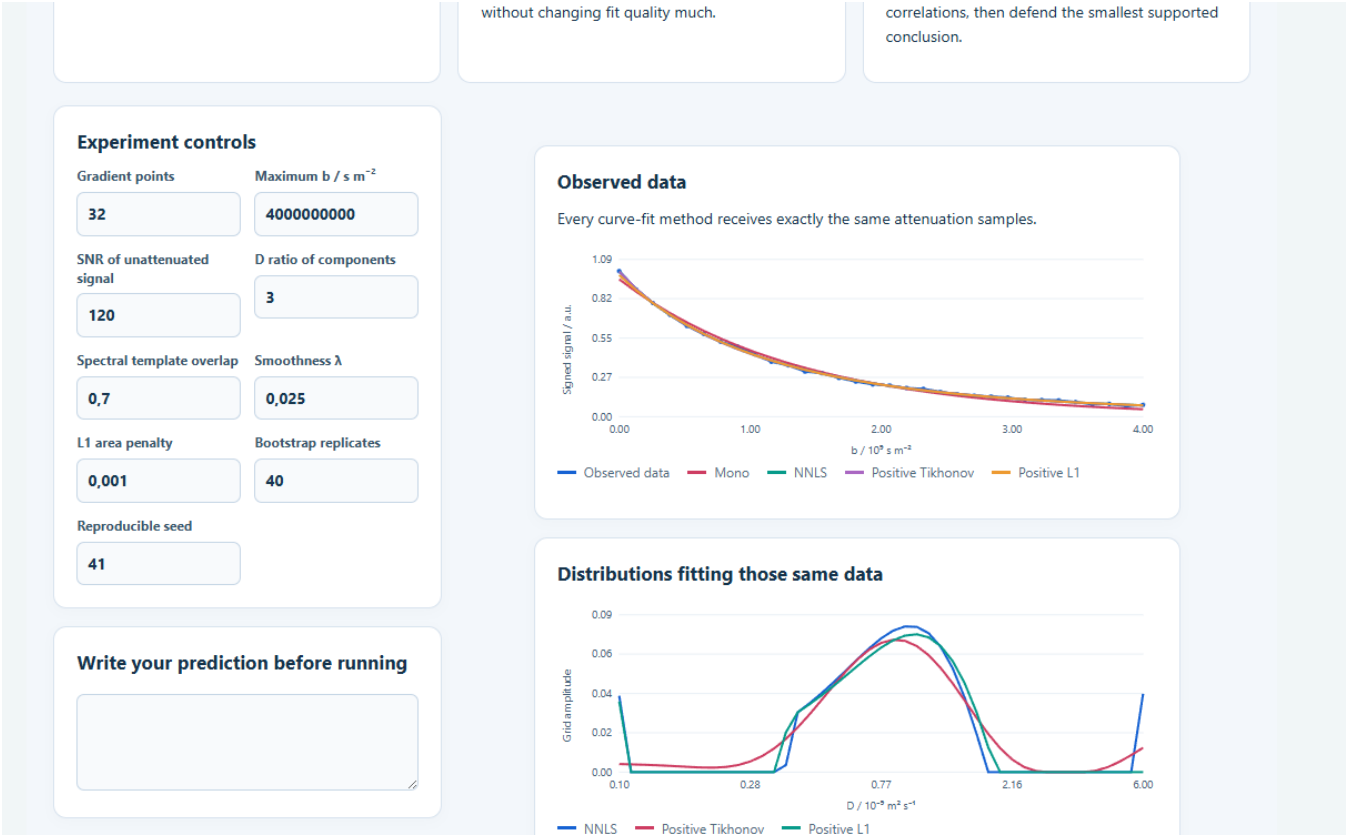
2. Acquire controls: a known-temperature standard; Δ at matched b; mirrored loops; phase/baseline; reversed order; time reference; recovery; RF nutation.

3. Write which observation supports your hypothesis, what weakens it and what remains unexcluded.

4. Submit the diagnosis. Without discriminating controls or a substantive explanation, the cause remains hidden. An incorrect hypothesis can be revised using the evidence.

Automatic feedback checks exercise requirements; it does not fully grade argument quality or validate a real diagnosis. Multiple causes can coexist in an experimental sample.

09 | One curve, several inverse problems



The inverse activity: all methods receive the same data; compare distributions and residuals rather than map appearance.

$$y_i = \Sigma_j K_{ij} a_j + \varepsilon_i \; ; \; K_{ij} = \exp(-b_i D_j)$$

Mono: estimates A and D by least squares in signal space; it does not log-transform or discard negative signals. One D can be apparent when the sample is a mixture.

NNLS: minimizes error with nonnegative amplitudes. **Positive Tikhonov:** adds second-difference smoothness on a log D grid. **Positive L1:** penalizes nonnegative area. All report KKT residual, tolerance and convergence. These examples do not implement PALMA, CONTIN or every algorithm in the catalogue.

Change SNR, D contrast and penalties. If different peaks explain the same data almost equally well, state which conclusion survives and which detail depends on the method. A local maximum does not identify a species.

10 | Spectral correlations and uncertainty

Supervised joint unmixing

The synthetic matrix contains two components sharing D across frequencies. The exercise knows their pure spectral templates and projects the matrix onto them before fitting separated decays. This is supervised unmixing, not blind NMF or automatic chemical identification.

1. Examine singular values, template correlation and the conditioning indicator. Do not turn a second singular value into automatic evidence of two species.

2. Increase overlap to one. Identical templates make the individual curves unidentifiable by this projection.

3. Compare joint and single-signal estimates. Explain which assumption adds information and how you would check it independently.

Model-conditional bootstrap

The activity resamples mono-fit residuals using a reproducible seed. The interval shows noise uncertainty within that model; it has no validated 95% coverage. It omits mixture bias, calibration, cross terms, experimental dependence and model selection.

A narrow interval around a biased estimate can mislead. Retain conditional uncertainty, regularization sensitivity and scientific alternatives separately. Reviewing the conclusion requires human judgment.

Related sources

Tikhonov regularization and optimal choice of the regularization parameter (2011)

PALMA: entropy/L1 hybrid regularization (2017)

11 | Real case: from FID to a conditional conclusion



The real Monash case after FFT, integration and nominal fitting. The result remains conditional and does not validate absolute D.

The Monash D2O case is from Silva, McKay and Junkers, DOI 10.26180/33075584.v1, licensed CC BY 4.0. The teaching subset has 16 FIDs of 16,384 complex points and sequence dstegp3s (Double-STE + LED). Original files are read without modification.

1. Load the case and inspect provenance, hashes and acquisition parameters before processing.

2. Run FFT and inspect group-delay correction, line broadening, phase and baseline. The example uses group delay 76 and LB 0.3 Hz; do not transfer these to another sample.

3. Select the signal in ppm from its spectrum. The teaching window 4.5-5.0 ppm belongs to this case, not all samples.

4. Compare positive-signal log fitting with direct signed fitting. Retain phase and integration decisions.

The peak is near 4.7014 ppm; a manual phase around -104.6° orients this signal in this processing workflow. Decide from the spectrum itself, rather than treating the number as a universal recipe.

12 | Import and retain provenance

Supported imports

The reader accepts attenuation CSV with b_s_m2 and signal, FID CSV with time_s, real and imag, and spectrum CSV with ppm and real. Raw Bruker needs acqus and fid, or acqus, acqu2s and ser/ser.ser; processed import needs 1r and procs. Select the correct type and inspect units.

1. Use a read-only copy and retain the original. Record files, sizes and hashes.

2. Check TD, dimensions, endianness/numeric type, frequency, spectral width, sequence, gradient list and temperature. A partial file does not establish a complete experiment.

3. Document processing, ppm window, normalization, b definition, fit method and uncertainty.

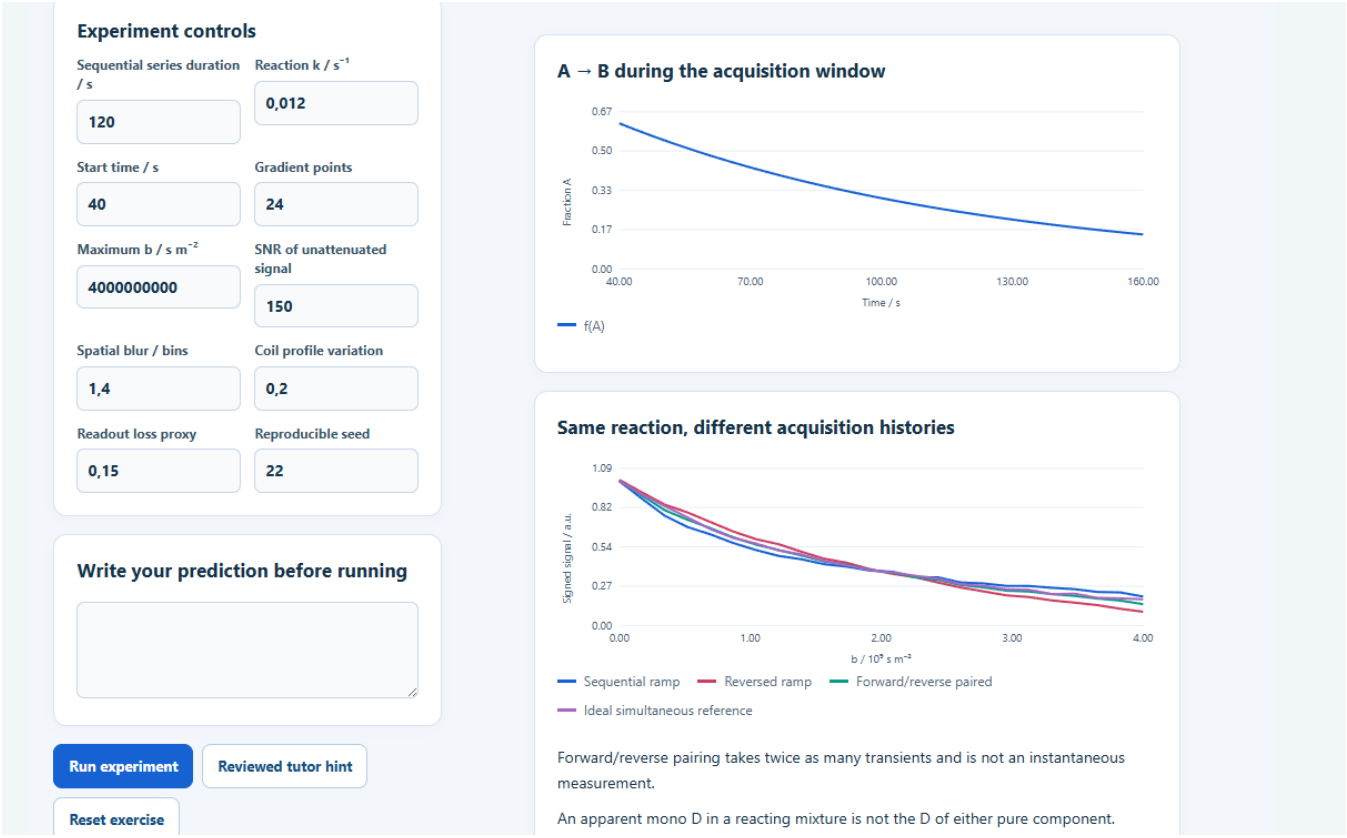
4. Export the notebook and record what additional data an absolute conclusion requires. You may need to reimport after leaving the activity; do not assume the browser retains an entire FID.

What the Monash case supports

The calculation uses nominal TopSpin Gamma, Glen, Gdist and difflist to build b. A nominal estimate near 1.93×10^{-9} m^2/s is a conditional processing result; it is not validated absolute D without gradient calibration, temperature and sequence timing definitions.

The three synthetic cases in this activity explore isolated signal, phase and drift. In every conclusion state whether you used published real data, a synthetic perturbation or a complete simulation.

13 | Time-resolved diffusion



The advanced activity: conversion A → B changes attenuation while a sequential ramp moves through b.

$$f_A(t) = \exp(-kt)$$
$$E(b,t) = f_A(t)\exp(-bD_A) + [1-f_A(t)]\exp(-bD_B)$$

1. Predict bias direction when b increases while A converts into B.

2. Compare forward and reverse order over the same temporal window. If estimates depend on order, a static description requires controls.

3. Average forward/reverse ramps and compare with the ideal simultaneous reference at the midpoint.

4. Vary k, duration and start time. Identify when evolution during acquisition affects interpretation.

Pairing two ramps costs twice as many transients. It can reduce first-order drift, but it is not a snapshot and does not remove every kinetic bias. The model assumes equal NMR response, constant D and negligible exchange during each encoding period.

14 | Ultrafast and SPEN: what spatial encoding gains

Scientific question

Can you obtain diffusion information before composition changes? Spatial parallelization assigns different weighting to different positions and collects them within one acquisition. The costs concern encoding, readout, resolution and sensitivity.

1. Inspect b(z), the coil profile and additional readout loss. The ideal simultaneous reference separates temporal evolution from encoding limitations.

2. Compare raw and corrected ultrafast data. Correction assumes independently known profiles and amplifies noise at low sensitivity.

3. Increase spatial blur in the SPEN-inspired model. Compare raw data with positive smooth decoding.

4. State what calibration or measurements would be needed to make this example quantitative.

Explicit limits

The exercise prescribes b(z), profiles and a teaching √N noise increase. SPEN is represented by a Gaussian spatial point-spread matrix, not Bloch propagation of chirp pulses. It does not compute adiabaticity, RF bandwidth, actual gradients, cross terms or complete spectral resolution. It is not a quantitative implementation of a particular SPEN-DOSY protocol.

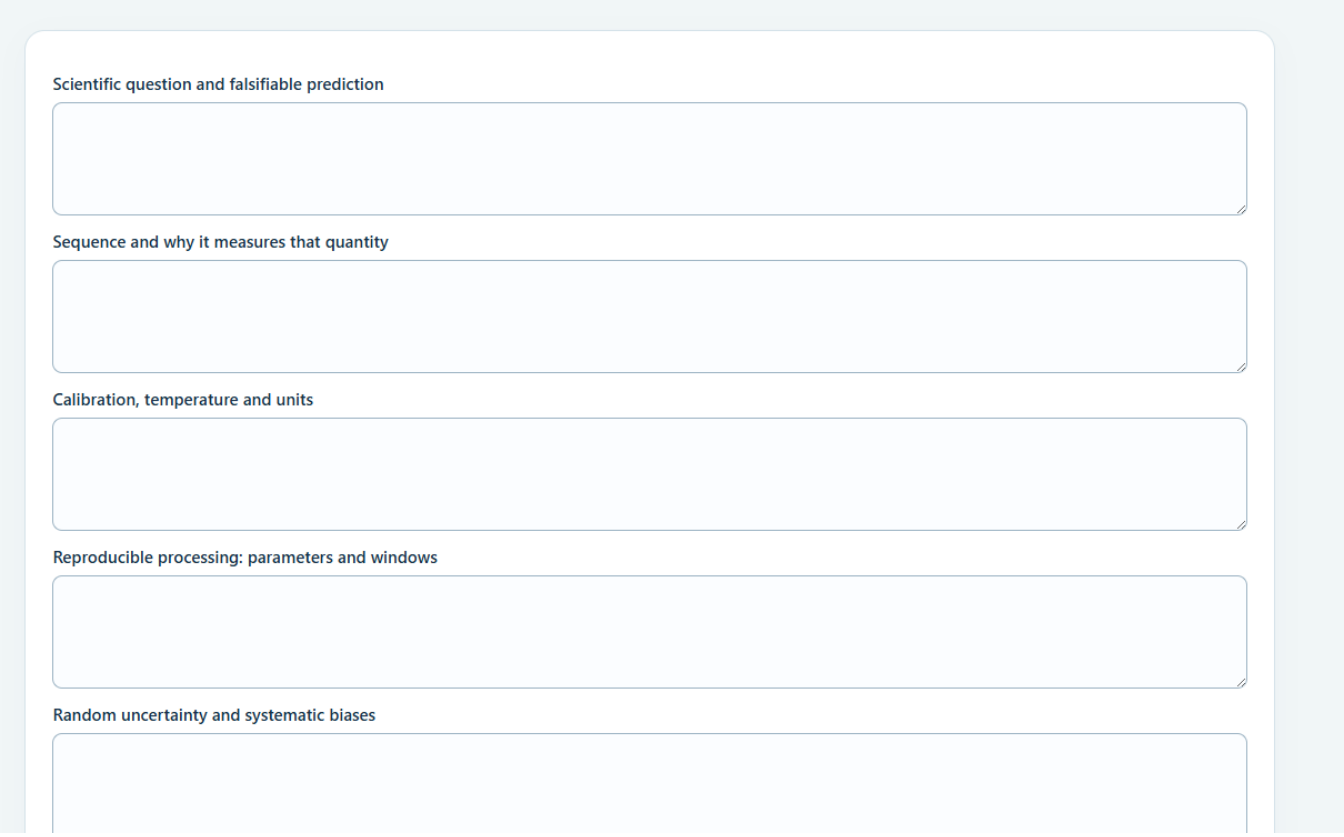
The map and simulation help select questions and controls. A real sequence requires encoding calculation, calibration, experimental checks and uncertainty estimation.

Ultrafast diffusion-based unmixing of ^{1}H NMR spectra (2021)

Optimisation of spatially-encoded DOSY for mixtures (2021)

Time-resolved and ultrafast restricted-diffusion studies (2020)

15 | Defend the conclusion and review the argument



The assessment activity: a rubric supports human review of reasoning; it does not identify species or use an LLM to grade arguments.

The rubric contains five criteria with levels 0-4. A teacher or researcher assigns the assessment from evidence, rather than plot appearance. Consult the complete criterion names and descriptors shown in the activity.

1. Define a question and a concrete claim; separate observation from interpretation.

2. Justify calibration, sequence, b range and processing using available metadata.

3. Explain controls, assumptions, residuals, conditional uncertainty and alternatives.

4. State what would change your conclusion or which measurement is missing.

5. Request human review, retain comments and record a revised argument.

Saving text or completing objective tasks does not demonstrate understanding of the entire derivation. Learning assessment should include a new case with different parameters and a solution that cannot simply be recalled.

15b | Rubric for human review

The application records seven sections: question, sequence, calibration, processing, uncertainty, controls and conclusion. A 30-character minimum checks that a response exists; it does not measure semantic correctness. The score remains unassigned until human review.

Five criteria, 0-4 each

1. Physics and appropriate equation: connect the observable, q history, b and model conditions.

2. Design and discriminating controls: acquisition answering the question and controls separating explanations.

3. Provenance and reproducible processing: files, units, hashes and decisions sufficient to repeat the analysis.

4. Uncertainty and stability: noise, biases, parameter sensitivity and resolution limits.

5. Conclusion proportional to evidence: a bounded claim and a next measurement capable of changing it.

Common level descriptors

0: absent or contradictory. **1:** unsupported assertion. **2:** correct with incomplete controls. **3:** justified, reproducible and bounded. **4:** compares alternatives, quantifies limits and proposes an independent test.

Use comments to revise the argument; do not turn a pending score into experimental validation. The program does not assign this score automatically or delegate the judgment to Beatrix.

16 | Beatrix, hints and evidence boundaries

Two clearly different modes

Original WPF lessons: Beatrix retains the locally fine-tuned model for the original 24 lessons. Generated answers should be checked against sources and data.

New mathematics and doctoral activities: help is reviewed documentary material. It does not generate new LLM explanations, no new model was trained for these chapters, and the tutor does not grade open arguments.

Browser: the laboratory runs without LLM inference. Reviewed hints can be consulted without transferring imported data to a remote model.

Why this separation matters

A convincing answer can confuse phase, q(t), b and the coil profile. Previous model tests on the new chapters found physical contradictions; therefore these chapters use a reviewed guide. A useful tutor should identify the source and assumption supporting its explanation.

Labels to retain

Synthetic: a model generates the data. **Real:** published or imported signals with provenance. **Conditional:** a calculation depends on incomplete parameters or calibration. **Software verified:** code/interface checks. **Experimentally validated:** evidence from adequate experimental controls; it does not follow from passing a software check.

17 | Essential equations and their conditions

$$J = -D\,\nabla c \;;\; \partial c/\partial t = -\nabla\cdot J = D\,\nabla^2 c$$

Conservation and Fick law for constant D. At equilibrium net J can vanish while molecules still move.

$$P(x,t|x_0) = (4\pi Dt)^{-1/2}\exp[-(x-x_0)^2/(4Dt)]$$
$$\langle (x-x_0)^2\rangle = 2Dt$$

One-dimensional free Gaussian propagator without boundaries or drift. In three dimensions $\langle |r-r_0|^2\rangle = 6Dt$.

$$\partial M/\partial t = D\nabla^2 M + i\gamma p(t)g(t)xM - R_2 M$$
$$q(t) = \gamma\int_0^t p(s)g(s)ds \;;\; E = \exp[-D\int_0^T q^2(t)dt]$$

A simplified complex transverse Bloch-Torrey form with a phase convention. The echo expression assumes free Gaussian diffusion and refocusing q(T)=0; relaxation and coherence pathways depend on the sequence.

$$E_v = \exp(-Db)\exp(-ivM_1) \;;\; M_1 = \int_0^T q(t)dt$$
$$E_{uniforme} = \exp(-Db)\,\mathrm{sinc}(v_{max} M_1)$$

With positive encoding phase and q(T)=0, integration by parts gives velocity phase -vM1. The last expression averages a uniform velocity distribution between -v_max and +v_max. It is not a universal convection law.

When comparing with the synchronized view: that view uses the opposite convention, $\varphi = -\gamma\int p(t)g(t)z(t)dt$, so its flow factor is exp(+ivM1). The two conventions conjugate the echo: they reverse phase and imaginary-part signs while preserving |E|, b and D.

18 | Units, assumptions and common errors

Use SI within calculations: γ in rad s^{-1} T^{-1}; g in T m^{-1}; δ and Δ in s; q in rad m^{-1}; b in s m^{-2}; D in m^2 s^{-1}. The product bD is dimensionless. If γ is in Hz/T, handle the 2π factor consistently.

1. Do not confuse δ with Δ or one bipolar lobe duration with the total pair duration.

2. Do not reset q during longitudinal storage: encoding history continues to determine attenuation.

3. Do not apply the rectangular equation to every gradient shape. Shape and fixed gradients can alter b and its cross terms.

4. Do not confuse a balanced echo with a zero velocity first moment.

5. Do not treat an ILT peak as chemical identity, absolute mass or an established experimental population.

6. Do not treat a conditional bootstrap interval as total uncertainty.

7. Do not transfer phase, ppm windows, indices or calibration across samples without checking the spectrum and experiment.

Notation checks in the articles

Chapters document literal discrepancies found by comparing equations, units and limits. They retain the PDF equation alongside the reviewed expression and justification; they do not claim an official erratum where none was established. Consult the relevant note before using a catalogue formula.

19 | Troubleshooting and reproducing your work

1. WPF reader will not open: check complete extraction and WebView2 Runtime. Retain the error message; use the web edition if the runtime is unavailable.

2. Invalid input: inspect bounds, sign and units. Invalid results are removed. After a parameter edit, previous charts are labeled stale; run again.

3. A case stays hidden: submit a hypothesis, discriminating controls and evidence explanation. Minimum text length is an interaction requirement, not a scientific-quality score.

4. The map changes with λ: inspect residuals, amplitudes and assumptions. Retain both configurations and explain which conclusion is stable.

5. No convergence: inspect KKT/tolerance. Do not interpret the distribution as a numerically confirmed solution before convergence.

6. Browser cannot save: export the notebook before closing. Retain originals separately and reimport when needed.

7. Incomplete import: inspect file type and required metadata. Do not invent missing calibration or parameters.

Minimum reproducible record

Application/content version, activity and seed, parameters/units, data/DOI/hashes, processing, method/tolerance, prediction, controls, conclusion and limits. Keep software evidence separate from evidence obtained on a spectrometer.

20 | Primary sources, licensing and further reading

Stejskal & Tanner (1965). Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient.
doi:10.1063/1.1695690

Sinnaeve (2012). The Stejskal-Tanner Equation Generalized for Any Gradient Shape.
doi:10.1002/cmr.a.21223

Kuchel et al. (2012). Stejskal-Tanner Equation Derived in Full.
doi:10.1002/cmr.a.21241

Silva, McKay & Junkers. Monash D2O public data.
doi:10.26180/33075584.v1

The data subset is attributed to its authors under CC BY 4.0. The teaching repository retains provenance.json; this manual does not reproduce the private article PDFs.

Diffusion inversion / related methods.
Tikhonov regularization (2011)
PALMA (2017)

Spatial and time-resolved methods.
Ultrafast diffusion-based unmixing (2021)
SPEN-DOSY optimisation (2021)
Time-resolved and ultrafast diffusion (2020)

Application files

Content/derivations.json: reviewed derivations. Web/doctoral-*.js: activities. Content/doctoral-data/: public data and provenance. Docs/DOCTORAL_MANUAL_ES.pdf and EN.pdf: this manual. START_HERE_ES_EN.txt: package startup. Names and paths concern the published software package, not instrument calibration.

Educational application documentation. Check sources, metadata and assumptions before turning a result into an experimental claim.

APPENDIX B

Original course and tools

Retained source document - version 1.0

The next 20 pages retain the original training manual 1.0. They reference the 24 lessons, examples, workshop and historical Beatrix behavior.

Screenshots and version names are intentionally retained. Edition 1.3 adds Research studio without turning historical documents into descriptions of its new screens.

Pages are appended without changing content, equations, figures or scientific boundaries. Printed numbering restarts in the source document; the reader toolbar and bookmarks locate pages in this volume.

For current instructions, export, the new pilot and model-selection rules use Part I. A model is used only with an approved manifest and CPU evaluation; otherwise documentary guidance is provided.

Open retained source document >

Return to current instructions >

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

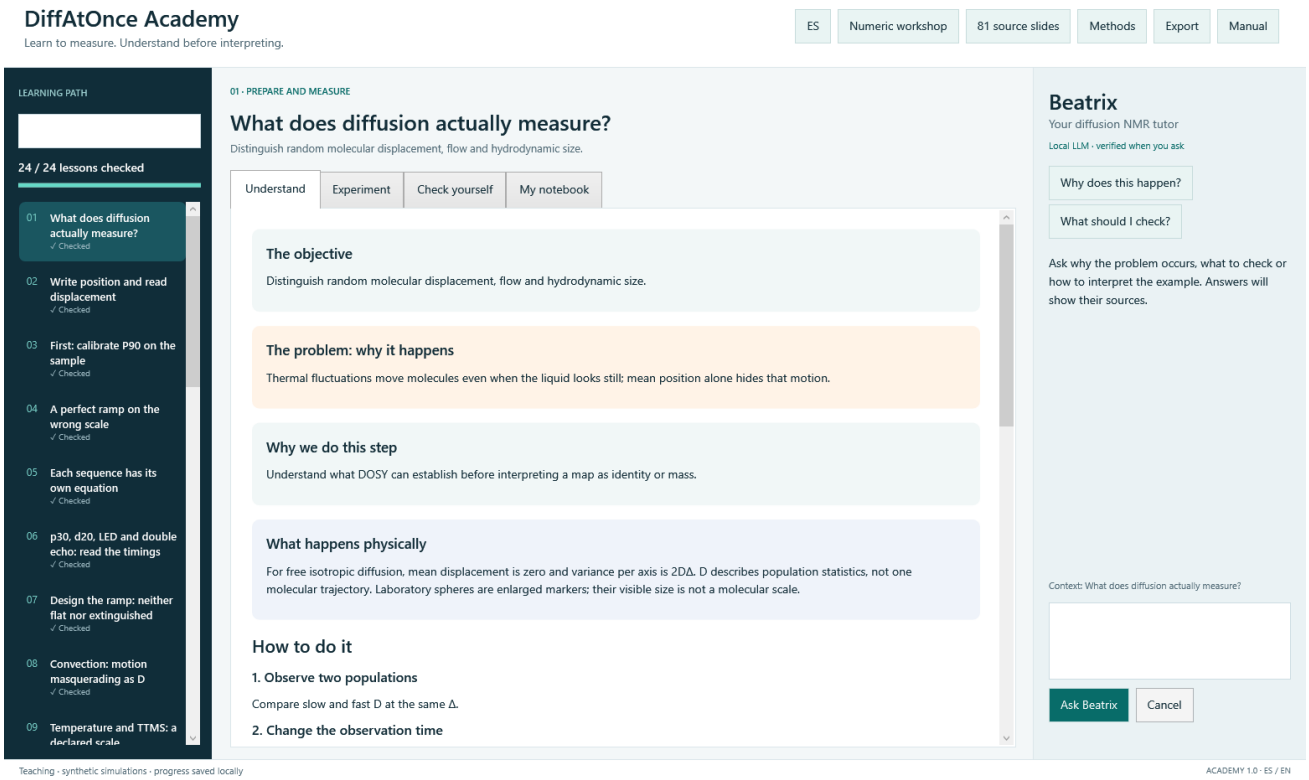
1. Extract the whole supplied application folder on Windows.
2. Open DiffAtOnce.Academy.exe, keeping Content, Lab and any supplied Models folder beside it.
3. Select English with the **EN/ES** button in the top bar.
4. Select a lesson in the learning path on the left. Search filters its title, chapter 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.

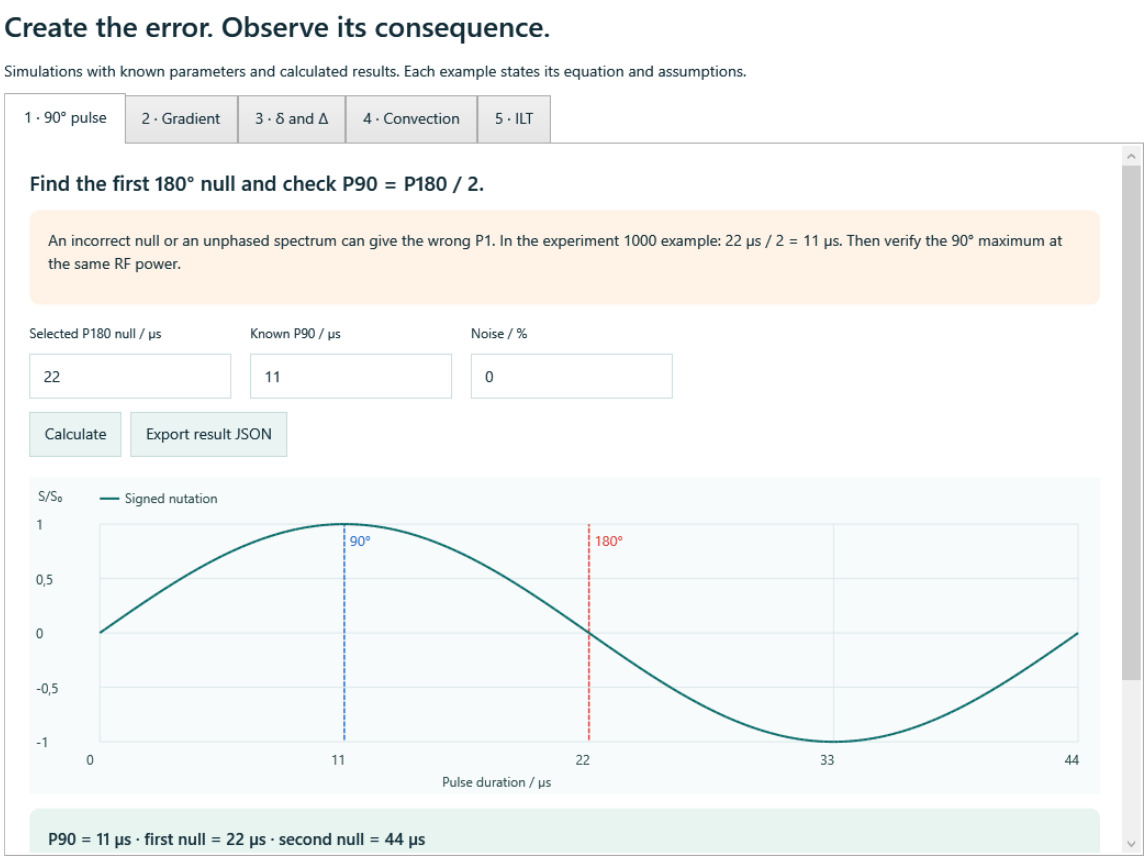
Numeric workshop: create errors and observe consequences

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.

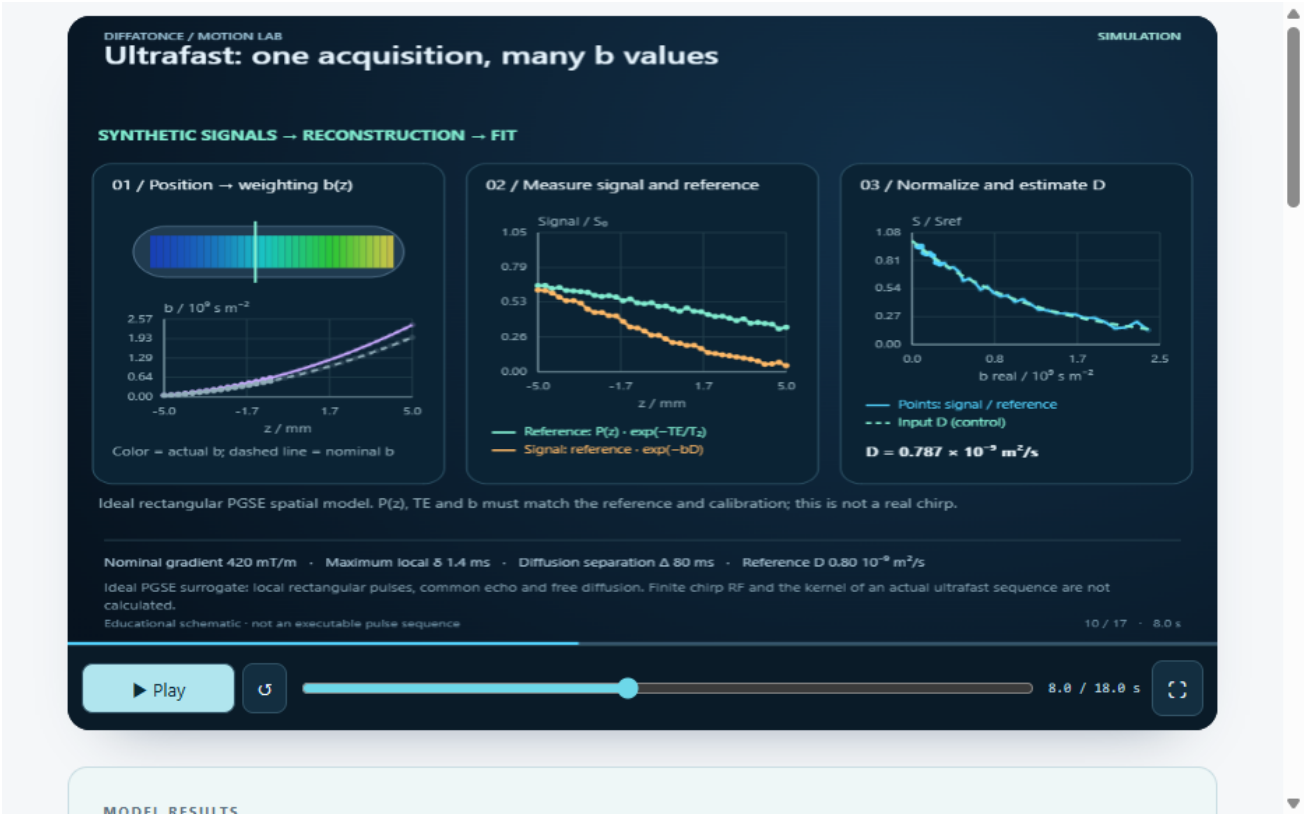
5. ILT

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

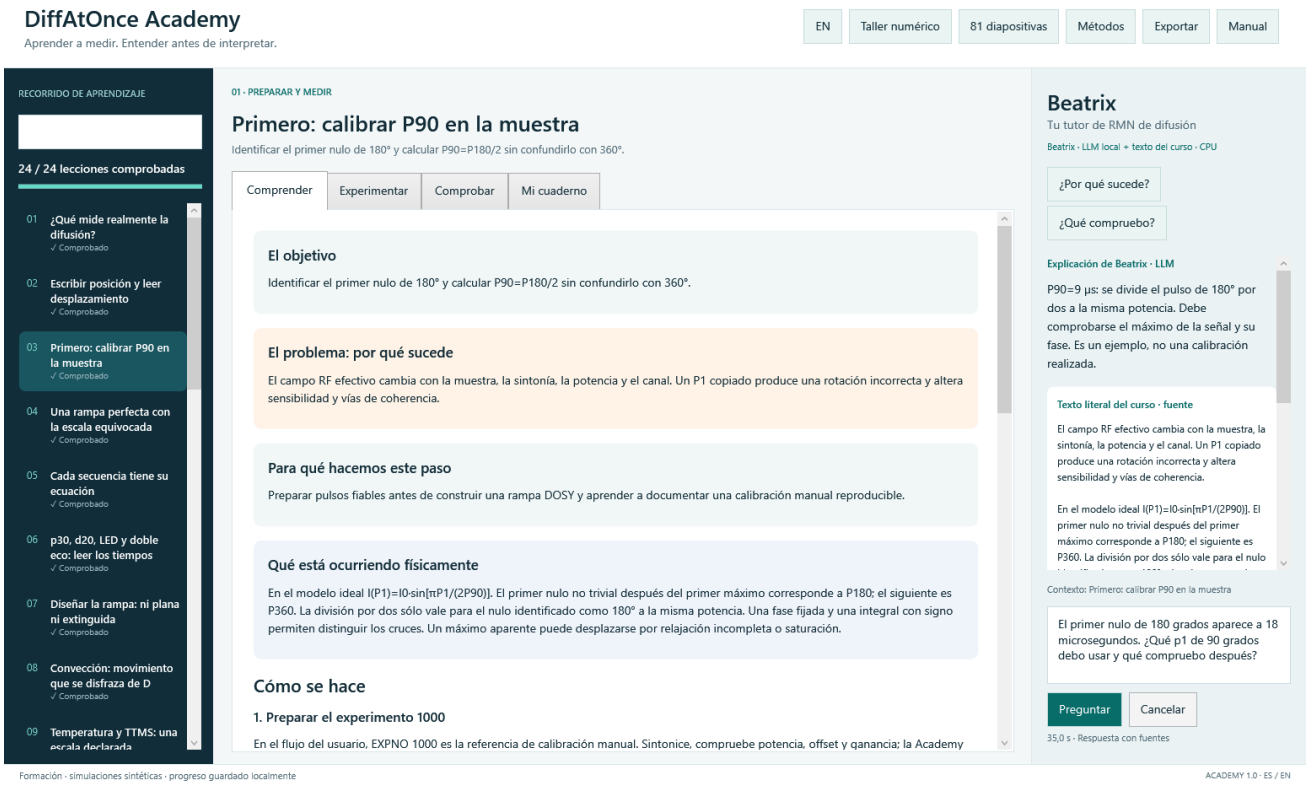
before fitting. Try pure A, then absent A at the same selected frequency; B leakage is not a valid recovered D_A.

• **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.

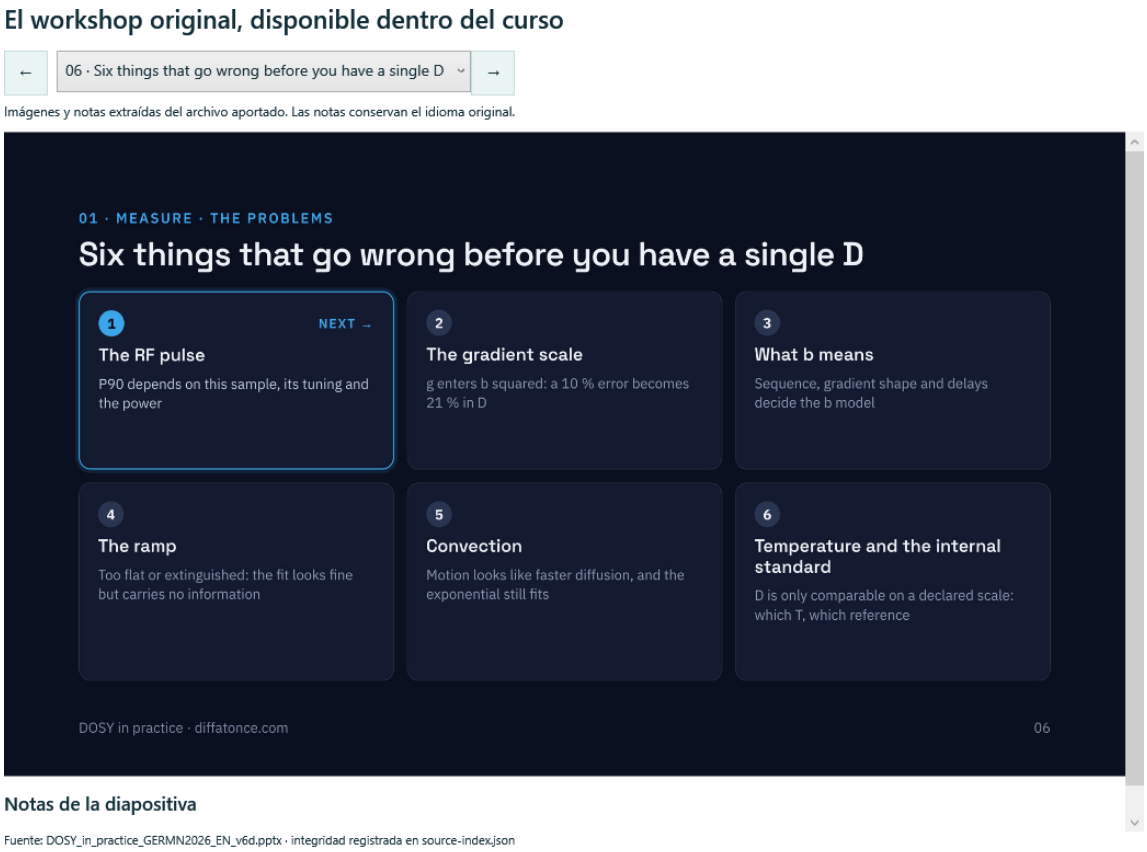
Verified hybrid package status. Local fine-tuning scored 32/32 on the reserved HF test, compared with 21/32 for the base model. The CPU package passed 15/15 checks of **LLM + literal course text**. Generated-only CPU answers omitted causes in open SPEN/TR questions; those limitations remain in the reports. 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.

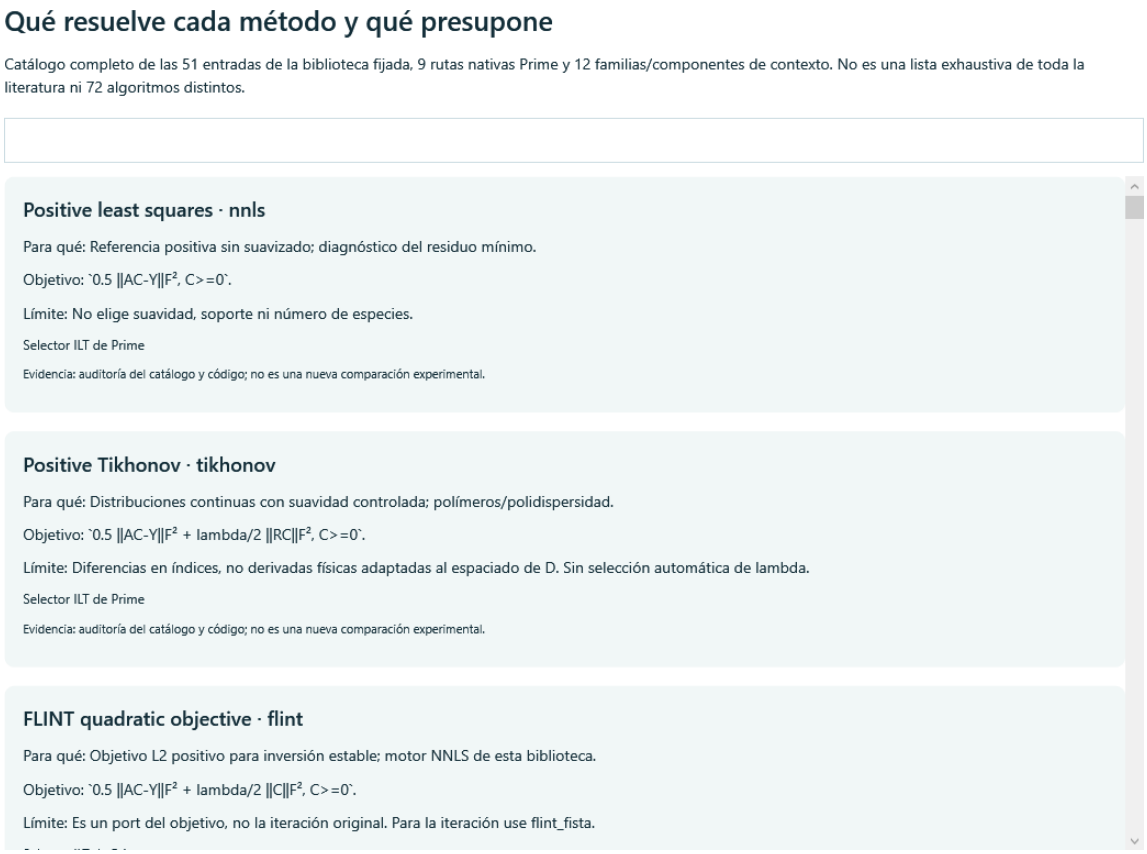
Consult source slides and methods

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

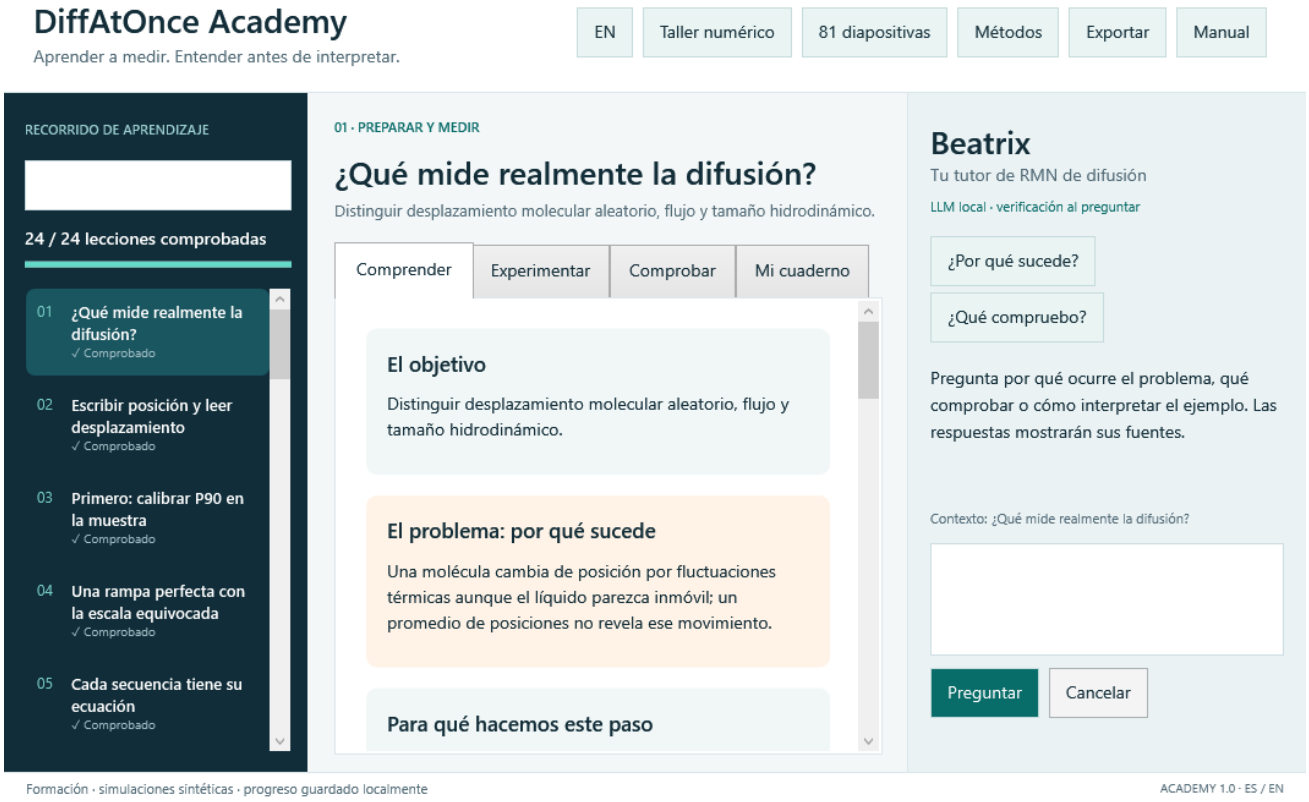
24 · Final case: decide what can be concluded

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.