

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.

The screenshot displays the 'DiffAtOnce ACADEMY' interface. On the left, a sidebar shows 'YOUR LEARNING PATH' with '0 of 24 modules completed'. Below this is a search bar and a list of activities under '01 - Prepare and measure'. The main area is titled 'DOCTORAL PATHWAY · 1.2' and 'From phenomenon to evidence'. It includes a sub-header 'Predict, design a check and defend a conclusion.' and a row of tabs: Doctoral, Lesson, Laboratory, Numeric workshop, Derivations, Check understanding, Notebook, Methods, and 81 slides. Below the tabs, there's a section for 'Doctoral pathway' with '4 saved evidence records' and buttons for 'Export notebook and evidence' and 'Import evidence'. A horizontal list of activities follows: 01 From positions to the echo, 02 Design a sequence, 03 One equation at a time, 04 Virtual experimental campaign, 05 The problem is hidden, 06 From decay to distribution, and 07 From... The bottom section features a large dark blue box with the text 'Think like the person designing the experiment.' and a flowchart: Motion → Phase → Signal → Evidence. Below this, three numbered boxes are visible: 01 SEE THE MECHANISM, 02 BUILD THE MEASUREMENT, and 03 DERIVE AND CHECK.

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.

Self-diffusion without net transport

At equilibrium each molecule changes position although the ensemble mean displacement is approximately zero. A gradient labels position with phase; the echo tests how well that label survives. In this model the gradient does not exert a force driving molecules: it encodes their position.

Ideal sequence
PGSE

D (10⁻⁹ m²/s)
1

Initial slab (μm)
240

Translation motion
No flow

Mean velocity (μm/s)
200

Uniform half-width (μm/s)
400

Ensemble molecules
512

Reproducible seed
20261005


Presets: δ=1.5 ms, g=200 mT/m and leading-edge Δ=60 ms per block; edit those timings and amplitudes in the table. In double STE, Δ=60 ms in each of two blocks: it is not the article figure's total Δ.

01 Positions

02 Individual phase

03 Complex sum

04 Observed echo

01 • Motion is random
Each dot is a trajectory; its color follows its phase.


02 • The phase label spreads
Shown: φ modulo 2π; accumulated phase in the particle card.


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.

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

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

Mean sequence
PGSE

p
1

Initial size (μm)
240

Transfer method
No flow

Mean velocity (μm/s)
200

Uniform half-width (μm/s)
400

Ensemble molecules
512

Reproducible seed
20261005

t / T
0 / 61.5 ms

p
1

q(t)
Predict → reveal

b(T)
Integrate q²

u(T)
Integrate signed q

Final echo
q(T)=0 · balanced

Ideal RF transfer
1 × S₀

▼ Builder: physical gradient × coherence pathway

Edit rectangular intervals. Every boundary with changed p represents an ideal RF transfer of the selected pathway; the diagram marks illustrative 90° or 180° pulses. p=0 stores a modulation: q does not return to zero.

#	Duration (ms)	g (mT/m)	p	Order
1	1,5	200	+1	↑ ↓ ×
2	29,25	0	+1	↑ ↓ ×
3	29,25	0	-1	↑ ↓ ×
4	1,5	200	-1	↑ ↓ ×

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 q² with the first moment ∫ 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

Pathway
Chapter

Physical intuition
Units, randomness and displacement

▼ Prerequisite diagnostic: four essential ideas

In an equilibrium solution, do molecules still move?

☐ Yes: self-diffusion without macroscopic net flux

☐ No: uniform concentration means no movement

☐ Only when a concentration gradient is applied

What units must b have for $\exp(-bD)$ to be dimensionless?

☐ s/m^2

☐ m^2/s

☐ T/m

If the whole effective gradient waveform is doubled, with fixed timing and no fixed gradients, how does b change?

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

Calibrate before choosing the ramp

Selected null / μs

Selected null order

Applied P90 / μs

Applied gradient scale

Standard D / $\text{m}^2 \text{s}^{-1}$

Reference temperature / K

Measured temperature / K

Apply selected null → P90

Apply standard calibration

A reference must match temperature and matrix.
 No universal temperature rescaling is applied.

Signed nutation: the first maximum and two nulls

The 90° maximum occurs at 9 μs . The 180° null crosses polarity at 18 μs ; divide by two for P90. The 360° null at 36 μs requires division by four. These values belong to this virtual fixture.

Observed data

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

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

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 $b_{\max} \cdot D$. 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.

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

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

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: $b = 2b_{\text{loop}}$ 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

identify the cause.

Submit only after gathering evidence.

Experiment controls

Case

CASE-A

The decay is nearly monoexponential, but D differs between two instruments.

Reproducible seed

5

Working hypothesis

—

Choose independent controls

- ☐ Reference standard at known temperature
- ☐ Change Δ while matching b range
- ☐ Compare mirrored q loops at matched b
- ☐ Compare signed and magnitude integrals
- ☐ Rephase and check baseline
- ☐ Reverse gradient order
- ☐ Zero-gradient reference versus time

Observed data

b / 10³ s m⁻²

— Observed data

Choose independent controls

Your controls do not yet distinguish the competing explanations. Add an independent control.

Explain which observations distinguish your hypothesis from an alternative


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.

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

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

Direct intensity: additive signed noise, all points

b for each row (s/m², comma-separated; never percentages)


```
653509.1344914145,11530181.401660925,35784151.52365601,73431163.71996517,124446787.74256998,188864139.8680188,266668833.81751567,357829109.84582174,462404143.7244584,580326073.6533743,711671447.4611506,856404163.0929744,1014470745.027592,1185973800.5885558,1370802036.4237823,1569075431.9138863
```

Fit D and explore uncertainty

conditional-monoexponential-fit

$D = 1.92716 \times 10^{-9} \text{ m}^2/\text{s} \cdot R^2 = 0.95471 \cdot n=16$

Conditional 95% interval: $[1.56126, 2.20371] \times 10^{-9} \text{ m}^2/\text{s}$



Log-linear fitting needs positive integrals. Direct fitting retains every signed point and assumes additive intensity noise: compare results to recognise

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 Γ , G_{len} , G_{dist} and $diff_{list}$ to build b . A nominal estimate near $1.93 \times 10^{-9} \text{ m}^2/\text{s}$ 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

Experiment controls

Sequential series duration / s	Reaction k / s^{-1}
120	0,012
Start time / s	Gradient points
40	24
Maximum $b / s\ m^{-2}$	SNR of unattenuated signal
4000000000	150
Spatial blur / bins	Coil profile variation
1,4	0,2
Readout loss proxy	Reproducible seed
0,15	22

Write your prediction before running

A → B during the acquisition window

Same reaction, different acquisition histories

Forward/reverse pairing takes twice as many transients and is not an instantaneous measurement.

An apparent mono D in a reacting mixture is not the D of either pure component.

The advanced activity: conversion $A \rightarrow B$ 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 $\sqrt{N}$ 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 ^1H 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

Scientific question and falsifiable prediction
Sequence and why it measures that quantity
Calibration, temperature and units
Reproducible processing: parameters and windows
Random uncertainty and systematic biases

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

$$\mathbf{J} = -D \nabla c ; \partial c / \partial t = -\nabla \cdot \mathbf{J} = D \nabla^2 c$$

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 $\langle |r-r_0|^2 \rangle = 6Dt$.

$$\partial M / \partial t = D \nabla^2 M + i \gamma p(t) g(t) x M - R_2 M$$

$$q(t) = \gamma \int_0^t p(s) g(s) ds ; E = \exp[-D \int_0^T q^2(t) dt]$$

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.

$$E_v = \exp(-Db) \exp(-ivM_1) ; M_1 = \int_0^T q(t) dt$$

$$E_{\text{uniforme}} = \exp(-Db) \text{sinc}(v_{\text{max}} M_1)$$

With positive encoding phase and $q(T)=0$, integration by parts gives velocity phase $-vM_1$. The last expression averages a uniform velocity distribution between $-v_{\text{max}}$ and $+v_{\text{max}}$. 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 $\exp(+ivM_1)$. 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 $\text{rad s}^{-1} \text{T}^{-1}$; g in T m^{-1} ; δ and Δ in s ; q in rad m^{-1} ; b in s m^{-2} ; D in $\text{m}^2 \text{s}^{-1}$. 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.