

IsoFind

Quick start

Getting to grips with IsoFind in a few hours

This booklet covers installation, the interface, importing data, isotopic matching, automation and the exchange of signed files. Read the first six chapters in order, then in any order you like.

IsoFind, sovereign isotopic traceability software. All data stays on your machine.

Contents

01	What IsoFind is	3
02	Installing and starting	4
03	Editions and licence	5
04	The workspace	7
05	Getting your data in	9
06	Sources and daughters	11
07	Your first session	12
08	Tabs and panes	14
09	Panels, strips and layouts	15
10	Moving around quickly	16
11	Standards, elements and CRMs	17
12	Nexus: matching and provenance	19
13	Field, boreholes and supply chains	21
14	The analysis notebook	22
15	Workflows and automation	23
16	The local assistant	25
17	IGL, isotopic modelling	26
18	Reports	27
19	ISOF, signing and exchange	28
20	Extensions	29
21	Preferences	30
22	Projects, campaigns and backup	31
23	Reference	32
24	Glossary	34
25	Going further	35

01

Chapter 01

What IsoFind is

Before clicking anything, a few words on what this software does and on the principles that explain most of its choices.

IsoFind is a workstation for isotopic geochemistry and traceability. It holds your samples and their measurements, compares them against references to trace an origin, models the processes that may have altered a signature, and produces documents you can cite, sign and pass on. It covers the whole chain, from the raw data line to the archived report.

What shapes the software's choices

Everything runs on your machine

None of your data leaves for a server. No mandatory online account, no telemetry, no tracking. This holds for the assistant too, whose model runs locally: your exchanges with it do not leave the computer. IsoFind was built to run on a network-isolated workstation. Nothing about it is degraded.

What you see, you can cite

Algorithms are documented, uncertainties propagated and displayed, outputs calibrated for publication. When a result rests on an assumption, the interface says so instead of returning a bare number. A traceability tool has to be traceable itself first.

Nothing is locked in

Your data leaves in the ISOF format, open and documented, signable and encryptable. An independent Python library, published as `isof`, reads and verifies those files without IsoFind. Your methods export as files, and so do your models.

How to read this guide

The first six chapters are enough to get working: install, understand the window, load data, get the vocabulary of the domain, and move around without searching. The rest reads in any order, whenever the need comes up. That is what the contents page is for.

The guide is in the software too

IsoFind ships an interactive guide that covers the essentials of these pages by having you do them rather than describing them. It opens on first launch and is available at any time from the Guide panel. Its steps tick when you actually perform the action, whichever way you go about it.

Menu paths are written as menu, then submenu, then entry. When something can be reached several ways, this guide gives the one that is easiest to remember, not the shortest.

[Help](#) > [Documentation](#)

02

Chapter 02

Installing and starting

From the executable to the first window, and what the welcome screen asks you to decide.

Installation

IsoFind ships as a single Windows installer. It requires no administrator rights beyond writing to the folder you choose, and no network connection during installation. On an isolated workstation, copy the installer by your laboratory's usual means and run it normally.

Installation creates a local database in your user profile. Nothing is written elsewhere. Encrypting that database is a Pro licence feature, described in the next chapter. Uninstalling does not remove that database. You can reinstall or upgrade without losing anything.

First launch

On first start, a welcome screen introduces the software over four pages and ends with a Start button. It does not come back, except after a major update, where it shows what changed. Closing it brings you to the workspace, with the interactive guide opening at the side.

Three things to set right away

1. **Language.** The interface exists in French and English, and switching is immediate. The setting lives in Preferences, General section.
2. **Theme.** Light, dark, or following the system. Charts and reports adapt. Preferences, Appearance section.
3. **Database protection,** if you hold a Pro licence. Turn it on now rather than once the database is full. The Protection menu handles it and hands you a recovery key.

Keep the recovery key elsewhere

Database protection rests on a key that exists in a single copy, and nobody, neither your administrator nor the vendor, can reconstruct it. Store it away from the machine it protects.

Licence

IsoFind opens without a licence, with the core set of features that remains free. A licence unlocks the additional modules, data certification and institutional uses. Activation works offline, through a signed licence file you drop into the software.

[Licence > Activate licence](#)

03

Chapter 03

Editions and licence

What you are holding, and what a licence opens up.

Two different things

The edition is what was compiled into your installation. The licence is what is unlocked inside it. A feature absent from your edition appears nowhere, so you will not go looking for it. A feature that is present but not covered by your licence shows up and tells you which level it needs.

This booklet describes IsoFind in full. Restricted features carry a band under their title, like the one on the Reports chapter further on. Anything without a band works everywhere.

The editions

Edition	What it holds
Essential	Free, with no time limit. A reduced core: the options are stripped before compilation, so they are not there at all.
Research	Free. The full chain for research work: import, samples, standards, elements, CRMs, quick matching, notebook, extensions.
Research+	Annual subscription, academic pricing. Research, with wider coverage of molecule families and report blocks.
Pro	Annual subscription. The whole catalogue is compiled; add-on modules are activated by the licence key.

The modules that came from the former vertical editions, IsoStrata, IsoPetrus, IsoSupply, IsoStream and IsoDef, only activate on a Pro licence. They are bought separately and added to an existing installation without reinstalling anything.

What a Pro licence opens

Area	Features
Reports	Report generator, professional reports, custom charts, legally admissible export.
Analysis	Batch analysis, advanced statistics, machine-learning predictions, advanced fractionation.
Security	Protection menu, database encryption, access control, backup hardening.
Traceability	Audit log, activity log, legal traceability.
Collaboration	Multi-user, user management, roles and permissions, data synchronisation.

What the free editions hold

A great deal, and most of this booklet. Projects, import and export, sample management, quick matching and recent analyses, sample comparison, advanced filters, search, data validation, isotope calculator, unit converter, Python console, 3D visualisation, standards, elements, CRMs, community database, the extension system, preferences, backup restore.

The subscription model

A purchased licence gives you one year. At expiry you keep and go on using everything you bought; you simply have no claim on versions published afterwards. Nothing switches off on the anniversary date.

Buying and activating

There is no online purchase, neither in the software nor on the site. You get in touch, you pay, and you receive a signed licence file. You drop it into the software, and activation happens with no connection, which matters on an isolated workstation.

[Licence > Activate licence](#)

The current price list, for editions, modules and laboratory certificates, is published at isofind.tech.

04

Chapter 04

The workspace

Anatomy of the window. Five areas, each with its own role. Better to tell them apart straight away.

IsoFind does not present itself as a series of pages you visit one after another, but as a workspace where several things stay open at once. The window reads top to bottom and outside in.

Area	Where	What it is for
Title bar	Top	Menus, a few frequent tools, and the window buttons. It is present on every screen.
Side strips	Both edges	Your shortcuts. Extensions settle there on their own, everything else is pinned by hand.
Docks	Left, right, bottom	Anchored panels: inspectors, journals, lists. They accompany the work without replacing it.
Central area	Middle	The work itself, in tabs. A page, a panel, an editor.
Status bar	Bottom	What is running right now: tasks, connections, alerts.

The title bar

It carries the full menu on the left and a few tools on the right: run a workflow, search, go fullscreen, open preferences. It deliberately shows no document title, to leave room for those tools. You can move the window by grabbing its empty part.

The side strips

Two narrow columns border the window. Each button opens or closes something in one click. Panels brought by an extension appear there automatically, since they have nowhere else to show themselves; everything else is pinned by you. An empty strip removes itself, so you will not see a useless grey column until you have pinned something.

The docks

A dock holds panels as tabs. You can resize them, collapse them, drag one from one dock to another. A panel is not condemned to stay there: see the chapter on panels and layouts.

The central area

This is where the work lives. It behaves like browser tabs, with the added ability to split into side-by-side panes. An inactive tab keeps its state: a map stays centred where you left it, a spreadsheet keeps its selection.

The status bar

Discreet but useful. It shows the open project, the state of connections, background tasks and their progress, and flags detected problems. Clicking an indicator opens the matching panel. You can

hide the ones you do not need with a right click.

The available panels

A panel is a secondary surface that accompanies the work without taking it over. Here are the main ones. View, Panels, All panels gives the complete list, extensions included.

Panel	What it shows
Guide	The getting-started walkthrough, and where you are in it.
Dashboard	What changed recently and where you left off.
Files	The explorer, and the kind of each file in a campaign.
Sample	Everything the database holds on the selected sample.
Sample history	The successive changes to a sample.
Database browser	Tables and their contents, to check with your own eyes.
Journal	What you did, timestamped.
Problems	What the software detected without interrupting you.
Actions	Every action in the software, grouped and searchable.
Notebook	The Python, SQL and R cells.
Workflows	Your automations and their state.
Assistant	The conversation with the local model.
Knowledge	Your documentation, its collections and its search.
Tasks	What is running in the background, and its progress.
Extensions	What is installed and what each extension brings.

If the workspace does not suit you

A traditional view, with a navigation sidebar and one page at a time, remains available in Preferences, Workspace section. Nothing is lost switching between the two.

05

Chapter 05

Getting your data in

Four ways to load your data, and the resolver that spares you describing your columns one by one.

IsoFind is of no use until it holds something. Start by loading something, even roughly: most of the software's notions become obvious as soon as you see them applied to your own data.

From a CSV or an Excel file

File > Import data > From CSV

The most common route. IsoFind reads your file, shows you a preview, and proposes a mapping between your columns and its fields. The semantic resolver recognises most headers on its own: usual isotopic notations, element symbols, concentration units, common organic contaminant names. You only correct what it could not read.

If you are starting from scratch, a template file is provided. It shows the expected columns and their format, and spares you a first fruitless import.

Data > Download CSV template

From an ISOF file

ISOF > Import an .isof file...

An .isof file is IsoFind's exchange format. It carries the data, its context, and where applicable the signature of the laboratory that produced it. On import, the signature is verified and its status announced before anything enters your database.

From an instrument or a database

IsoFind can read the output of several instruments directly, and connect to an existing database or a LIMS. These routes take some configuration the first time, and become reproducible afterwards: a saved connection replays in one click, or can be automated.

File > Import data > From instrument

By hand

Data > Add sample

For a single sample, or simply to see the shape of things before importing a batch. Nothing is required beyond the name: you can fill in the rest later.

After the import

Three checks, in this order.

1. Open the Samples page and check the row count. A gap with your file means rejected rows, and the journal says which ones and why.
2. Run data validation. It spots out-of-range values, doubtful units, likely duplicates and missing fields that will get in the way later.
3. Classify your samples into sources and daughters. The next chapter covers it. Everything else depends on it.

[Tools](#) > [Data validation](#)

An import is never lost

Every import is recorded in the journal, with its source file, timestamp and outcome. The change registry, separately, keeps what actually changed in the database. Both are available from the View menu, under Panels.

06

Chapter 06

Sources and daughters

The one notion you cannot guess, and the one everything else in the software rests on.

IsoFind sorts your samples into two families. The distinction does more than tidy things up: it decides what matching can conclude.

A source is a reference

An ore of established origin, a laboratory standard, a documented and geolocated site, a raw material whose supplier you know. A source answers the question "what does material from this place look like". It forms your basis for comparison, and its quality sets the ceiling on everything you will be able to demonstrate.

A daughter is a question

A sample whose origin, supply chain, or altering process you are looking for. A field sample, a received batch, a seized item. A daughter answers the question "where does this come from".

What this means in practice

- Matching compares daughters against sources, never the other way round. A misclassified sample will therefore never be found.
- A database without sources can conclude nothing, however good the measurements. This is the most common cause of a disappointing first result.
- The same material can legitimately exist twice, once as a source and once as a daughter, if it serves both as reference and as object of study.
- Classification can be changed at any time, from the sample record or in bulk from the Samples page.

The rest of the record

Beyond that distinction, a sample carries whatever you give it: a material type, a sector, a collection location with its coordinates, a project, a classification, a free description. None of it is required, but every filled field is another filter, and search relies on them.

Measurements come next: isotopic ratios with their uncertainties, elemental concentrations, physico-chemical parameters, molecules and micropollutants. Each can carry its method, its reference publication and its purification yield.

The Sample panel

Kept open beside your work, it shows everything the database holds on the selected sample, and lets you correct a value without changing page. Right click gives access to the actions, with those that write to the database explicitly marked.

07

Chapter 07

Your first session

A complete example, from import to signed report. About an hour, with your own data or a small test set.

The previous chapters describe the parts. This one puts them end to end on an ordinary case: you have received a series of measurements, you want to know where some samples come from, and to produce a document that holds. Follow it once in full, even if your real work is different: what matters is seeing the sequence.

Prepare

1. Create a project and name it after your study. Everything that follows will attach to it, and you will be able to archive it as one piece at the end.
2. Check that your standards are known to IsoFind. If your series refers to an in-house standard, declare it now rather than after the import.
3. If you are starting from a spreadsheet, look at the CSV template and align your headers with it. It is not mandatory, the resolver recognises a lot, but it will spare you row-by-row corrections.

Load

1. Import your file. Read the column mapping screen rather than skimming it. A unit error takes ten seconds to fix there; once the data is in, it costs far more.
2. Open the Samples page and compare the row count with your file. A gap means rejected rows, and the journal says which.
3. Run data validation. Deal with what it flags before going further; an out-of-range value left in place will resurface at the worst moment, usually inside a matching result.

Classify

1. Mark your references as sources: field standards, materials of known origin, documented sites.
2. Mark as daughters the samples whose origin you are looking for.
3. Fill in at least location and material type for the sources. Those are the fields you will filter on most, and a matching result without a location will not teach you much.

This is where the quality of the result is decided

Almost every disappointing first attempt comes from this step, not from the matching engine. A database of twenty daughters and three sources cannot mechanically conclude anything solid. If you are short of references, start by importing some, even if it delays the analysis by a day.

Analyse

1. Run a quick match on a daughter whose answer you already know, if there is one. It is the best check on your settings, and it costs nothing.

2. Then move to the unknown daughters. Look at the confidence index, but above all at the detail: what brings a candidate closer, and what pushes it away.
3. If the deviations are systematic, open advanced matching and allow the processes that are plausible in your context. A regular deviation on a given system is often the signature of a process, not of a bad source.
4. Write down your assumptions as you go. They will have to appear in the report, and they are hard to reconstruct afterwards.

Report

1. Open the reports module and choose your blocks. A summary, methods, measurements, matching results and assumptions are enough for a first document.
2. Reread the assumptions block. It is the one an attentive reader will read first, and the one that protects your conclusions.
3. Sign the report if your laboratory certificate is in place. A recipient will be able to verify its integrity without owning IsoFind.
4. Export the data as ISOF too, signed. The report says what you conclude; the ISOF file lets someone else redo it.

Make the session repeatable

If this sequence has to be repeated, on every new series or every month, rebuild it as a workflow. You describe the sequence once, dry run it to check, and it becomes an object you relaunch without remembering the order of operations. The workflows chapter covers that shift.

Save your layout under a name as well, for instance that of the study. You will find the window exactly as you had arranged it for this work, tabs included.

08

Chapter 08

Tabs and panes

How to keep several things open without losing track.

Opening

The plus button, right of the tab strip, opens the catalogue of what can occupy the central area: the software's pages on one side, panels on the other. Universal search and the palette lead there too, and often faster.

Splitting

Drag a tab to the edge of a pane: it splits in two and the tab settles in the new half. Towards the middle of another pane, it moves there. On the tab strip, it reorders. An emptied pane disappears and its neighbour takes the space, so you never have an empty rectangle to tidy up.

Double-clicking the empty part of a tab strip enlarges that pane to the whole area, and a second double click returns to the previous layout. Handy for work that needs room without your wanting to close everything else.

Pinning

A pinned tab shrinks and moves to the head of the strip. It will not be closed by accident. Keep it for what you have open all day.

What survives and what does not

When the software reopens, the tab layout is restored: the panes, their split, the tabs and which one was active. What is restored is the list of what was open, not the internal state of each page. A recovered spreadsheet reopens on the same data, not necessarily on the same cell.

This restoration can be turned off in Preferences, Workspace section, if you prefer to start from a clean window every morning.

Gesture	Effect
Drag to a pane edge	Splits the pane and places the tab there
Drag to the middle of a pane	Moves the tab into that pane
Drag along the tab strip	Reorders
Double click the empty strip	Enlarges the pane, or reverts
Right click a tab	Pin, close others, send back to the dock
Scroll wheel on the strip	Scrolls the tabs when they overflow

09

Chapter 09

Panels, strips and layouts

A panel is not nailed to the window edge, and a working setup is worth saving under a name.

Four possible places

Every panel can live in four places, and you choose according to what you do with it.

Place	When to choose it
In a dock	An inspector, a list, a journal: what accompanies the work without taking it over.
In the middle, as a tab	The notebook, the consoles, the script editor: anything shaped like an editor that suffers in a narrow column.
Floating window	Something to consult over everything else, for a moment, without reorganising the window.
Separate window	A second screen. Several panels can share the same detached window.

Moving is done by dragging, or through the panel tab's context menu. A panel is never in two places at once: opening it somewhere removes it from where it was.

The strips

The narrow columns bordering the window are your personal shelf. A pinned item is a reference, not a copy: the label follows the language and any renaming, and if the target disappears the button shows as hollow instead of disappearing without a word.

You can pin a page, a panel or a command. The side is your choice, not an imposed category.

Named layouts

File > Layouts...

A layout captures everything: panels, their places and sizes, and your tabs. You save a whole working setup, not just an arrangement of windows. Keep one per kind of task. The one for data entry has nothing in common with the one for writing a report, or with a modelling session.

Putting everything back straight

If the window ends up in a state you cannot get out of, View, Panels, Reset layout puts the panels back where they started without touching your data or your saved layouts.

10

Chapter 10

Moving around quickly

Four routes to the same thing, and how to choose.

IsoFind holds about a hundred actions and sixteen pages. Nobody knows them all, not even through the menu. Four tools answer different questions.

Universal search

Press Shift twice

The shortest path to anything. It searches pages, panels, commands, settings, workflows and your samples at once. It is the only tool that does not require knowing which category something was filed under. Start there and you spare yourself the question.

The command palette

Ctrl, Shift and P

It searches commands only, and shows their shortcut beside them. By going through the palette for a frequent action, you end up remembering its combination without ever reading a list of shortcuts.

The Actions panel

The palette assumes you know what you are looking for. The Actions panel answers the other question, "what can I do", by showing every action grouped by category. Star the ones you use and they move to the top.

The top menu

Nineteen menus, and just about everything the software can do. A twentieth is added, Extensions, built from whatever your plugins declare. Nobody reads it end to end, but it is the only place that shows the whole span of IsoFind in one view. You find there what you did not know existed.

You want to	Use
Find something without knowing where it is	Universal search
Run an action whose name you know	The palette
See what you can do in an area	The Actions panel
Discover the span of the software	The top menu
Return in one click to what you open daily	The strips

11

Chapter 11

Standards, elements and CRMs

What gives your numbers meaning, and without which a comparison between laboratories means nothing.

Isotopic standards

[Standards > Isotopic standards](#)

An isotopic ratio only reads against a standard. IsoFind knows the usual international standards and lets you declare others, specific to your laboratory. Every measurement knows which standard it refers to. You can convert from one scale to another, and you avoid the classic misunderstanding between two series that seemed comparable.

On import, an unknown standard is flagged, and you choose to create it or to attach it to an existing one. The software does not decide for you.

Standardisation

[Standards > Standardization...](#)

The dedicated page applies drift correction and normalisation to a series of measurements, from the interleaved standards. It shows the fitted model and its residual, not just the corrected result, so that the correction stays checkable. The operation can become a workflow if you repeat it at every analytical session.

Elements

[Elements > Manage elements](#)

The module lists elements with their isotopes and abundances. Every cell of the periodic table is in one of three states: no ratios available, ratios available but inactive, or active. You can add a custom element if your work falls outside the standard catalogue.

Two coverages not to confuse

Isotopic fractionation is natively modelled for a restricted set of systems, those for which the literature provides reliable constants. But any isotope can be stored, measured, compared and exported in IsoFind, and the geochemical side covers all elements and all molecules. Not seeing a system in the fractionation list does not mean IsoFind cannot work with it.

CRMs

[CRM > CRM list](#)

Certified reference materials serve to check the accuracy of your measurements. IsoFind holds their list, their certified values and uncertainties, imports the Georem reference set, and compares

your measurements with expected values. Statistics and comparison show your deviation and its drift over time. That is often the first sign an analytical line has moved.

Checking a series

The usual routine is four steps, and it is worth doing before interpreting anything.

1. Interleave your CRMs through the series, at the start, in the middle and at the end. Three points are enough to see a drift; one shows nothing.
2. Import the whole series, CRMs included, without separating them. The module recognises them by their identifier.
3. Open the CRM statistics. There you read your deviation from certified values and its spread.
4. Look at the comparison over time. A stable deviation is an accuracy error, which can be corrected; a moving deviation is instrumental drift, which is handled differently.

A CRM also works as a source

Nothing stops you declaring a certified material as a source alongside its control role. Its composition is known and documented, which makes it a good-quality reference for matching.

12

Chapter 12

Nexus: matching and provenance

The scientific core of the software. Relating a signature to an origin, and recognising the processes that altered it.

Quick matching

[Analysis > New quick analysis](#)

You choose a daughter, IsoFind compares it against all your sources and returns the most likely matches, with a confidence index and the detail of what brings each candidate closer or pushes it away. Start there, if only to judge the quality of your reference base.

Advanced matching

PRO LICENCE

The dedicated page opens up what the quick version summarises. There you control which isotopic systems are taken into account and how they are weighted, the tolerances, the processes allowed to intervene between source and daughter, and the geochemical criteria added to the isotopic comparison.

Processes, and why they change everything

Between a source and the sample in your hand, something has usually happened: adsorption, oxidation, mixing, evaporation, Rayleigh distillation. Each of those processes shifts the signature in a known way. If you do not account for them, you are looking for an identical signature where physics guarantees it is no longer identical.

Nexus composes those processes and looks for the combination that explains the observed deviation. The result does not just say where the sample comes from, but what happened to it on the way. That second piece is often the more useful one.

The process library

Available processes, their fractionation factors and their bibliographic sources can be consulted and extended. You can add your own values, measured or taken from a publication, and document them so the result remains citable.

Signature reconstruction

The reverse approach: from a present-day sample and assumptions about the processes it underwent, work back to the most likely initial signature. Useful in multi-source settings, where several inputs mix and no source looks like what you measure any more.

A confidence index is not proof

Nexus returns a likelihood computed under the assumptions you gave it. It is worth what your reference base is worth, and what the allowed processes are worth. The report says so explicitly, and your conclusions need to say so too.

13

Chapter 13

Field, boreholes and supply chains

Three modules that answer the question "where", each at a different scale.

Mapping

Geolocated samples plot on a map, with the usual layers and filters. The colour palette is sequential rather than diverging, so that an intermediate value does not read as a neutral point. Selections made on the map propagate to the other views.

Boreholes

The borehole module holds your drillings, their lithological logs and the samples taken at each depth. It solves flow in three dimensions, from the unsaturated zone to the water table, and simulates the transport of a geochemical plume with its spatialised isotopic fractionation. A lithological model trained on European data fills in the zones where you have no description.

The visualisation shows hydraulic head, piezometric contours and Darcy vectors, overlaid on the plume and the measurement points.

Supply chain

The supply chain editor represents a chain as a graph: deposits, processors, blends, destinations. You build it by drag and drop, and IsoFind checks that it stays acyclic and flags compliance inconsistencies. A reference deposits database, covering several dozen countries and several isotopic systems, provides a starting point.

Supplier audit applies the same logic to a given supplier: what they declare, what their materials say, and the gap between the two.

14

Chapter 14

The analysis notebook

When the interface is not enough, writing the calculation yourself without taking the data out of the software.

[View](#) > [Panels](#) > [All panels...](#) > [Notebook](#)

The notebook is a book of executable cells, on the model of scientific notebooks. It accepts Python, SQL, R, Markdown, PHREEQC and IGL, with syntax highlighting, and returns text, tables and charts.

It works on your data, not on a copy

That is the difference with exporting to an external tool. An API named isofind gives access to your samples, their measurements and your results in a few lines, without an intermediate file. What you compute can be written back to the database, attached to a report, or turned into a workflow step.

Six kernels, one space

- **Python** runs in the embedded browser, with the usual scientific libraries. No Python installation is required on the workstation.
- **SQL** queries the local database directly, read only. For a count or a cross-tabulation, it is often the shortest path.
- **R** also runs locally, with its packages, and serves mainly for the statistical analyses and plots you already do there. Markdown for prose, PHREEQC for speciation and IGL for a fractionation model come on top.

A single Python console

The Python console in the Tools menu and the notebook's Python kernel share the same namespace. A variable defined in a notebook cell is visible from the console, and the other way round. Two separate environments would produce unfindable variables and irreproducible results.

What it is most used for

- A calculation the interface does not have, specific to your laboratory.
- A quality check on a series, before declaring it good.
- A precise chart, in the exact format a journal expects.
- A prototype treatment, later turned into a workflow once it works.

15

Chapter 15

Workflows and automation

Describe a sequence of treatments once, then replay it, schedule it, or trigger it on an event.

What a workflow is

A chain of steps: read data, correct it, run a match, produce a report, export a signed file. Each step takes the previous one's result. You build it visually, you name it, and it becomes an object you relaunch without having to remember the order of operations.

Automation > New workflow...

Dry run before the real one

A workflow can run dry. It walks through all its steps, checks that inputs exist and parameters are consistent, and announces what it would have done, without writing anything. This is the first reflex to have on a workflow that touches the database, and it avoids discovering a wrong file path after three hours of computation.

Automation > Simulate a workflow (dry run)...

Triggering

Trigger	Example
On demand	You launch the workflow when you need it.
On a schedule	A quality check every night, a report every Monday.
On an event	When a file lands in a watched folder, when an import finishes.
From another workflow	One workflow calls another. You compose instead of duplicating.

Following along

The automation dashboard shows the armed workflows, the ones that ran, the ones that failed and at which step. A failure is not just a message: the debugger replays the run step by step with the values it actually saw.

Automation > Automation dashboard...

Sharing

A workflow exports as a signed document. The recipient sees where it comes from, whether it has been modified since signing, and which capabilities it requests before running it. A treatment that

circulates between laboratories needs that guarantee.

A workflow really acts

An imported workflow can write to your database, produce files and call external services if it has the capabilities. Read what it requests, and dry run it before launching it for real.

Three workflows that pay off straight away

- **The overnight check.** It rereads the database, runs data validation and leaves you a status in the morning. You find the problems before starting, not in the middle of an analysis.
- **The instrument series.** It watches the folder where your instrument drops its output, imports what arrives, applies standardisation and files the samples in the right project.
- **The end-of-study package.** It produces the report, exports the data as signed ISOF and drops the lot in a dated folder. Half an hour of manual work that is rarely done the same way twice.

Start with the shortest of the three. A two-step workflow that runs every day teaches you more about the engine than a fifteen-step one you never dare launch.

16

Chapter 16

The local assistant

A model that runs on your machine, reads your data and your documentation, and sends nothing outside.

The assistant is built into the software and runs locally. It queries your samples, your analyses and your matches, and answers in your language. None of your exchanges with it leaves the computer. You can use it on data you would not entrust to any online service.

Three modes, three levels of autonomy

Mode	What it does
Converse	It answers and explores with you. It reads, it changes nothing.
Plan	It writes a workflow matching your request, which you review and launch yourself.
Act	It runs what you approved, showing every tool call in the conversation thread.

Every tool call appears in the conversation: you see what it reads and what it does. Nothing happens behind the scenes. Without that, an assistant would be unusable anywhere actions have to be justified.

Giving it your documentation

The Knowledge module accepts your procedures, your publications, your lab notes. The assistant draws on them to answer, and cites its sources. Full-text search is entirely local: your documentation never leaves the workstation.

Your sources are organised into collections, by project or by subject, across document kinds. You can then scope the assistant to the collections of your choice. You get a project-specialised agent without changing model.

A model must be installed

The assistant needs a local engine and at least one downloaded model. The AI assistant section of preferences checks the engine is present and lists the available models. On an isolated workstation, the model copies over like an ordinary file.

17

Chapter 17

IGL, isotopic modelling

A dedicated language to write a fractionation model, check it, and share it.

IGL is an isotopic modelling language, separate from IsoFind and usable on its own. It lets you write a model as text: reservoirs, fluxes, fractionation processes and their parameters. A model written that way can be read, versioned, compared and cited. A diagram clicked together in an interface cannot.

Inside IsoFind

IGL > New IGL model

The IGL menu gathers the essentials: create or open a model, check it before running it, format it, browse the fractionation databases, and get an explanation of a diagnostic code. The notebook also accepts IGL cells: you can mix a model and its analysis in the same document.

Check before running

Checking reads the model without running it and flags what does not hold: an inconsistent unit, a reservoir never fed, a parameter absent from the databases. Diagnostics carry a code, and every code explains itself in full from the menu. Check before launching. It costs a few seconds.

IGDB databases

Fractionation values come from databases in the open IGDB format, which can be inspected and replaced. You can use the ones shipped, add one specific to your laboratory, or point at a database published by a third party. A model states which database it uses. Its result therefore reproduces elsewhere.

IGL is installed separately

IGL has its own installer and lives in its own folder, with its databases and its specification. IsoFind detects it automatically. If you want to know which version is in use and where it comes from, IGL, IGL version and location tells you exactly.

18

Chapter 18

Reports

Producing a document that holds up before a review committee or a court.

PRO LICENCE

[Plus button of the central area > Reports](#)

The reports module assembles a document from blocks: summary, methods, measurement tables, charts, matching results, uncertainties, appendices, compliance statements. You choose the blocks, order them by dragging, and the document builds from your real data.

What sets these reports apart from an export

- Every figure stays attached to its measurement, its method and its uncertainty. A reader can walk back up the chain.
- Computational assumptions are written down, not implied.
- The document can be signed. A third party then verifies it has not moved since issue.
- Comparing several scenarios fits in a single report. A decision maker often expects exactly that.

Customising

The report's visual identity is set once and for all. For what the shipped blocks do not cover, a custom block accepts code that produces its own content. You can add a figure or a table specific to your practice without waiting for a software release.

Signing

A signed report carries your laboratory's identity and a timestamp. The signature rests on your laboratory certificate, described in the next chapter. A recipient can verify it without owning IsoFind.

19

Chapter 19

ISOF, signing and exchange

How to pass on data without the recipient having to take your word for it.

The format

ISOF is IsoFind's exchange format. It carries the data and its context: measurements, methods, standards, uncertainties, everything that is missing to interpret a number when it travels alone. It is open and documented, and an independent Python library reads it without IsoFind.

Three levels

Level	What the recipient can establish
Plain file	The data and its context, with no guarantee of origin.
Signed file	Who issued it, when, and that the content has not changed since.
File encrypted for a recipient	All of the above, and nobody else can read it.

The laboratory certificate

ISOF > My laboratory certificate

Signing requires an identity. IsoFind relies on a three-level chain of trust, from a root authority down to your laboratory certificate. The certificate request is made from the ISOF menu, and the certificate you obtain is renewed before it expires.

Verifying what you receive

ISOF > Verify an .isof file...

Verification tells you who signed, when, whether the content still matches the signature, and whether the signing certificate is valid and recognised. Make it a habit for any file from outside, including from a trusted colleague: what verification establishes is not the sender's good faith but the integrity of the transport.

Sending encrypted

For a given recipient, encrypted export makes the file unreadable by anyone else. You manage your recipients and their certificates from the ISOF menu. The file can then travel over any channel, including one you do not control.

20

Chapter 20

Extensions

What a plugin can add, and why it does not look like an add-on bolted on the side.

An extension does not stay in its corner. It contributes its actions to the palette, its panels to the strips and the central area, its widgets to the dashboard and its section to preferences, exactly like the rest of the software. Once installed it is indistinguishable from a native feature. That is deliberate.

Installing

[Plugins > Browse library](#)

The library shows the available extensions, the ones installed, and the capabilities each one requests. An extension declares precisely what it accesses: read samples, write to storage, show a panel, run computation. Nothing more is granted to it.

Signing

Distributed extensions are signed, and the signature is verified on install. If you develop your own, developer mode allows an unsigned extension to run. Necessary during development, inadvisable the rest of the time.

A few shipped examples

- A PHREEQC console, which runs geochemical models and proposes an interpretation of them.
- A full R console, with its packages and its plots.
- A quality control workshop.
- Identification tools, notably for perfluorinated compounds and for ICP-MS interferences.

Writing your own

[Plugins > Developer documentation](#)

An extension is a folder with a manifest and some JavaScript. It runs in a sandbox, it inherits the user's theme automatically, and it has nothing to know about the software's appearance. The Plugins menu has what you need to create one from a skeleton.

21

Chapter 21

Preferences

Where the settings live, and the few that really change your day.

[Edit > Preferences...](#)

Preferences are arranged in grouped sections, and above all they are searchable: if you do not know where a setting lives, type its name. Changes gather at the bottom of the window and only take effect once applied, so you can explore without breaking anything.

The sections

Group	What you find there
Workstation	General, appearance, workspace. Language, theme, font size, tab restoration.
Science	Analysis and computation, Nexus, IGL, reports and exports. The settings that influence results.
System	AI assistant, extensions, modules, data and backup, security, network, advanced, diagnostic.
Extensions	One section per extension that declares one.

The ones worth a look early on

- **Theme and font size.** Software you look at eight hours a day deserves two minutes of setting up.
- **Tab restoration.** Decide whether you want your window back as it was or a clean start every morning.
- **Automatic saving and its interval.** To set according to how long your sessions run.
- **Confirmation before deletion.** Leave it on until you are entirely comfortable.
- **Date format.** The ISO format avoids ambiguity between collaborators in different countries.

The scope of a setting

Some settings apply to you, others to the open project, others are imposed by the installed edition. Every setting carries a marker showing its scope. It will save you a long search for why a change does not take.

22

Chapter 22

Projects, campaigns and backup

Organising work over time, and losing nothing.

Projects

File > New project

A project groups samples, analyses, reports and settings. It acts as a frame: opening a project narrows what you see to what concerns it, and some settings can be specific to it. A project exports as a single file. That is the simplest way to archive finished work or pass it on.

Campaigns

File > Open an object > Campaign folder...

A campaign is a folder on your disk declared as a workspace. IsoFind reads its contents, recognises the kind of each file and marks it, and offers the appropriate ways to open them. Your files stay files, in place, readable by other tools. It bridges your disk organisation and the software, without trying to replace your file explorer.

Backing up

Data > Backup local database

The backup produces a complete copy of your database, encrypted if your licence opens protection. Do it before any cleaning operation, before a major update, and on a regular rhythm that you can hand to a scheduled workflow. The shortcut is Ctrl Shift S.

The journal and the registry

Two traces coexist, and they do not say the same thing. The journal records what you did: an import, an export, a report generation. It is a trace of intent. The change registry records what changed in the data. It is a trace of fact, durable and signed. Both are needed: the first answers "what did I undertake", the second "what came of it".

Database cleaning cannot be undone

The cleaning operation permanently removes what it takes out. It asks for confirmation and displays what it will do, but no undo is possible afterwards. Back up first.

23

Chapter 23

Reference

Shortcuts, and what to know for an isolated workstation.

Shortcuts

Shortcut	Effect
Shift Shift	Universal search
Ctrl Shift P	Command palette
Ctrl Shift A	Actions panel
Ctrl Shift E	Files panel
Ctrl Shift J	Journal panel
Ctrl Shift H	Sample history
Ctrl Shift D	Database browser
Ctrl Shift M	Problems panel
Ctrl Shift L	Layout manager
Ctrl Shift K	Shortcut settings
Ctrl I / Ctrl E	Import and export an .isof file
Ctrl Shift S	Back up the local database
Ctrl O	Open a CSV in the editor
Ctrl W	Close the current editor
F11	Fullscreen

The complete and current list is available from Help, Keyboard shortcuts. The palette also shows each command's shortcut next to its name. That is how you end up knowing them.

Working offline

IsoFind only asks for the network for the operations that are explicitly online: checking reference databases, the extension library, updates. Everything else, including the assistant, the notebook and reports, works without a connection. On an isolated workstation, plan to copy reference databases, assistant models and the extensions you want by hand.

When something goes wrong

1. Look at the status bar. A failing indicator appears there, and clicking it opens the detail.
2. Open the Problems panel, which gathers what the software detected without interrupting you.
3. Check the journal, which keeps a timestamped trace of the last operation and its outcome.

4. The Diagnostic section of preferences gathers the information useful for a report, all in one place.

[Help](#) > [Report a bug](#)

Updating

There is no automatic update, deliberately: software that updates itself on an isolated or certified workstation would cause more trouble than it saves. Download the new installer from isofind.tech and run it over the existing installation. Your database, your projects and your settings are untouched.

After a major update the welcome screen comes back once to show what changed. The full detail is in Help, Changelog.

What a good report contains

- Your IsoFind version and edition, from Help, About.
- What you were doing, in one sentence, and what you expected.
- The exact message if there is one, copied rather than summarised.
- The contents of the Diagnostic section of preferences, which already gathers most of the technical context.

A minimal dataset that reproduces the problem beats any description, when you are able to share it.

24

Chapter 24

Glossary

The terms that come up in the interface, defined once.

Term	Definition
Campaign	A disk folder declared as a workspace, whose files IsoFind recognises.
CRM	Certified reference material, used to check the accuracy of measurements.
Daughter	A sample whose origin is being sought.
Dock	One of the three panel anchoring areas, left, right or bottom.
IGDB	Open format of the fractionation factor databases used by IGL.
IGL	Isotopic modelling language, separate from IsoFind and usable on its own.
ISOF	IsoFind's exchange format, open, signable and encryptable.
Layout	An arrangement saved under a name, covering panels and tabs.
Nexus	The matching and process composition engine.
Panel	A secondary surface, dockable, floating or central, that accompanies the work.
Process	A phenomenon that shifts an isotopic signature in a known way: adsorption, oxidation, mixing, evaporation.
Source	A reference sample, whose origin is established.
Standard	The reference against which an isotopic ratio is expressed.
Strip	One of the two narrow columns bordering the window that hold your shortcuts.
Workflow	A saved sequence of treatment steps, replayable and schedulable.
Workspace	IsoFind's main window, with its tabbed central area and its panels.

25

Chapter 25

Going further

Where to find the rest, and how to reach us.

In the software

- The **interactive guide**, Guide panel, which covers the essentials of these pages by having you do them. It has two walkthroughs, the second opening once the first is complete.
- The **help** in the Help menu, more detailed than this booklet on each module.
- The **assistant**, whom you can ask about the software as much as about your data.

Online

The full documentation, release notes and training resources are at isofind.tech. The `isof` Python package, which reads and verifies ISOF files without IsoFind, is published on PyPI.

Reporting and asking

A defect is reported from Help, Report a bug, which gathers the useful diagnostic information. Feature requests and usage questions go through support, reachable from the same menu.

One last word

IsoFind is built for laboratories that need their results to hold up before a review committee, an auditor or a court, and that would rather not hand their data to an online service to get there. If the software is useful to you, telling people about it is what helps most.