



NOVASYNNE

EEG Workbench

User Guide

Import a recording in any of five formats, filter it, take the artifacts out with ICA, epoch it around its own markers, and measure what happened — with every step recorded and every refusal explained.

Contents

Before you start	5
What you need	5
The recording used in this guide	5
The workspace	6
Before a recording is imported	6
Once a recording is loaded	7
What the file says about itself	8
Moving the window	9
The channel list	10
Importing a recording	12
The Import dialog	12
Formats that are more than one file	12
The montage decides whether you get scalp maps	13
Looking at the recording	14
One channel at a time	14
The spectrum	15
Scalp topographies	16
Time and frequency together	18
What is not available yet	19
Preprocessing	20
Where it lives	20
The presets	20
Choosing the band	21
What it reports	22
The panel follows the data	23
The spectrum afterwards	24
Taking artifacts out with ICA	25

Fitting ICA on its own	25
Inspecting the components	26
Cropping	27
Events, annotations and epochs	28
Annotations you add yourself	28
The markers the file already carries	30
Cutting epochs	30
Event-related potentials	32
The average at one channel	32
Every channel at once	33
Trial by trial	34
The scalp at chosen latencies	35
Measuring a peak	36
ERSP and ITC	37
Statistics	39
Exporting	41
Exporting to the Analytics Discovery Workbench	42
Version and support	46
Appendix A — Control reference	47
The rail	47
Toolbar	47
The right-hand panel	47
Preprocess	48
Views	48
Export	48
Appendix B — Supported formats	50
Appendix C — What the workbench refuses to do	51

Appendix D — The recording used in this guide 52

Before you start

The EEG Workbench takes a recording off a clinical or research amplifier and gives you somewhere to work on it: filter it, re-reference it, take the artifacts out, cut it into epochs around its own markers, and then look at it as waveforms, spectra, scalp maps, ERPs and time-frequency plots. It is built on MNE-Python, and every step it applies is one you could name in a methods section.

Two things are true of this workbench and worth knowing before the first page.

Nothing leaves the machine. There is no assistant here, no API key, no model call, no telemetry. The recording is read, held in memory and written back out; that is the whole data path.

One recording is loaded at a time. The workbench holds a single session, and importing again replaces it. Reset clears it.

This guide is task-oriented. Each chapter answers a question you might arrive with — *why will it not draw me an ERP?*, *what did preprocessing actually do?*, *which of these numbers is safe to quote?* The appendices are the control reference, the supported formats, and the list of things the workbench declines to do.

What you need

The application	Running and reachable in a browser. See the README shipped with the package for installation.
Your data	A recording in EDF/EDF+, BDF, FIF, BrainVision or EEGLAB format.
Nothing else	No API key, no account, no network.

SET SECRET_KEY BEFORE YOU START USING IT IN EARNEST

Left unset, the application generates a new signing key every time it starts, so restarting it invalidates the page already open in your browser. The next thing you click reports that the security token is no longer valid and asks you to reload — nothing is lost, but it is avoidable. Set `SECRET_KEY` in `.env` once and restarts stop mattering. The README says how.

The recording used in this guide

Every figure in this guide comes from one file: `rc01.edf`, the sample recording that ships in `DemoData/`. Seventy channels at 2048 Hz, forty-six seconds long, carrying 243 stimulus markers of fifteen kinds in its EDF+ annotations.

It is short, which is deliberate — it is long enough to show every feature and short enough that ICA and time-frequency analysis finish while you watch. Your own recordings will be longer, and the chapters on ICA and export say where that changes what you should do.

The workspace

Before a recording is imported




Figure 1 — The workbench on first open.

Import and About are the only controls enabled. Export, Preprocess, Channels, Events and Statistics switch on when a recording arrives, because none of them mean anything before one does.

The rail down the left side is the whole application: Reset, Import, Export, the preprocessing menu, Channels, Events, Statistics, About.

Once a recording is loaded




Figure 2 — Seventy channels on one time base.

Three regions.

The chart, filling the middle, is the recording itself — a stacked multi-channel view by default, one trace per channel, sharing one time axis. The markers the file carries appear as ticks along the bottom.

The right-hand panel has three tabs: Info, Channels and History.

The bar along the bottom states the sampling rate, how many channels are being drawn out of how many exist, the duration, how many preprocessing steps have been applied, and where in the recording the window sits.

THE STACKED VIEW DRAWS 32 CHANNELS

A seventy-trace stack is unreadable, so the multi-channel view caps at 32 and says so in the bottom bar — **Channels: 32/70**. The Channels tab decides which ones. This is a display cap and nothing else: every channel is loaded, filtered, referenced and exported whether or not it is on screen.

What the file says about itself

File:	rc01.edf
Format:	EDF/EDF+ (.edf)
Sampling Rate:	2048 Hz
Channels:	70
EEG Channels:	70
Duration:	45s
Samples:	94,207
Bad Channels:	None
High-pass:	0.1 Hz
Low-pass:	417 Hz

Quick Actions

 **Apply Standard Filter**

 **Average Reference**

Figure 3 — What the file declares.

The Info tab is the recording's own account of itself: file, format, sampling rate, channel count, duration, sample count, bad channels, and the high-pass and low-pass bounds already applied — 0.1 Hz and 417 Hz here, which is what the file's own header declares rather than anything you chose.

Watch those two rows. They are the fastest way to confirm that a filter you just applied really took.

Quick Actions at the foot of the panel run the two most common steps directly — a standard 0.1–40 Hz filter, and an average reference — for when you do not need the full preprocessing dialog.

Moving the window



Figure 4 — The window: how much to show, and where.

Duration is a choice of 5, 10, 20, 30 or 60 seconds. The arrows step a window at a time, the outer buttons jump to the start and the end, and the Time box takes a position in seconds if you know where you are going. The left and right arrow keys step the window, and Home and End jump to either end, whenever the focus is not in a text box.




Figure 5 — A twenty-second window.

Every view that shows time — the stacked view, the single channel, the spectrogram, the Morlet map — reads this same window. Changing it changes all of them, which is the point: a spectrogram of a stretch you are not looking at is not much use.

The channel list




Figure 6 — Every channel, coloured by region.

The Channels tab lists every channel with its type, coloured by scalp region — frontal, central, parietal, occipital, temporal, auricular. Tick the ones you want drawn; Select All and Select None are there for when you want to start over.

CHANNEL TYPE IS WHAT THE FILE SAID, NOT WHAT THE NAME IMPLIES

In this recording all seventy channels are typed as EEG, including the four EOG derivations (VE0GL , VE0GR , HE0GL , HE0GR) and the two ear references. The analyses that should not include an EOG trace — the ERP plots, the scalp maps — leave them out by name. The stacked view does not: it draws what you tick.

Importing a recording

The Import dialog




Figure 7 — The Import dialog.

Three things to set.

The format. Five are supported, and picking a file sets the radio button to match its extension. It is stated rather than sniffed because two of the five are multi-file formats where the extension you select is the one that matters.

The file, or files. Drag them in or browse. The picker takes several at once, which is what BrainVision and EEGLAB need.

The montage. This is where electrode positions come from.

Formats that are more than one file

BrainVision is three files — `.vhdr`, `.vmrk`, `.eeg` — and EEGLAB is often two, `.set` with `.fdt`. Select them together. The dialog says so when you choose either of those formats, because a `.vhdr` on its own opens nowhere.

Uploads are capped at 1 GB per request.

The montage decides whether you get scalp maps

The workbench does not know where an electrode sat. The file names the channels; the montage turns those names into positions. Choose one at import and every topographic view works. Skip it, or choose one that does not match your cap, and the topographic view says **Montage Required** instead of drawing.

The dialog offers nine: Standard 10-20 and 10-05, four BioSemi sizes, EasyCap M1, and two EGI HydroCel arrays. This guide uses Standard 10-05, which places this recording's 64 scalp channels; the ear and EOG derivations are not in any standard montage and simply have no position.

THE MONTAGE IS NOT A DISPLAY PREFERENCE

It is applied at import and it is what every scalp map, every ERP topography and ICLabel itself are computed against. Changing your mind means importing again.

Looking at the recording

Six views share the toolbar at the top right. The first two are the signal itself; the rest are ways of summarising it.

One channel at a time




Figure 8 — One channel, full height.

The single-channel view draws whichever channel is selected in the toolbar over the current window, at a scale that uses the whole pane. This is the view for judging whether a channel is any good — the drift over the first few seconds here is obvious at this scale and invisible in a seventy-trace stack.

The spectrum




Figure 9 — The spectrum, with band power beside it.

The power spectral density of the visible window, drawn to 50 Hz, with relative power in six bands beside it. Both are computed over the channels you have selected.

THE SPECTRUM STOPS AT 50 HZ

That is a fixed display range, not the data's limit. It is the band that carries almost everything clinical EEG cares about, but it does mean the mains peak sits at the very edge of the plot and that the effect of a notch filter is hard to see here. The Info panel's low-pass row is the honest check on what a filter did.

Scalp topographies




Figure 10 — Voltage across the scalp at one instant.

Voltage mode maps one instant. The slider chooses which, the link button pins it to the main view's window, and the time is printed above the map so a figure taken from it can be labelled honestly.




Figure 11 — Band power instead of voltage.

Band Power mode maps average power in one band — six to choose from, delta through high gamma — over the whole recording rather than one instant.




Figure 12 — Every band at once.

PSD mode maps all bands side by side, normalised so the maps show where power sat rather than how much of it there was.

THREE VIEWS, THREE SETS OF BAND EDGES

The band-power map uses the workbench's own six bands (delta 0.5–4, theta 4–8, alpha 8–13, beta 13–30, low gamma 30–50, high gamma 50–100 Hz). The PSD map is drawn by MNE and uses MNE's five (0–4, 4–8, 8–12, 12–30, 30–45 Hz). The Statistics dialog reports a third set, five bands ending at 45 Hz. The numbers are therefore not interchangeable between the three, and an alpha figure quoted from one should say which one it came from.

Time and frequency together




Figure 13 — Time-frequency for the visible window.

The spectrogram is the fast route: a short-time Fourier transform of one channel over the visible window, 0.5–50 Hz.




Figure 14 — The same window through Morlet wavelets.

Morlet TFR is the slower and better-resolved route over the same window, 1–50 Hz across 40 log-spaced frequencies at seven cycles, with the cycle count reduced automatically at the low end where seven cycles would not fit in the window.

Both are per-channel. The channel selector sits in the same bar, and the link button re-reads the main view's window.

What is not available yet



Figure 15 — ERSP and ITC are disabled.

Two of the four time-frequency modes are greyed out, and the ERP view is unavailable altogether:




Figure 16 — The ERP view before any epochs exist.

ERSP, ITC and every ERP plot are averages *across trials*. Until the recording has been cut into epochs there are no trials to average, so the workbench disables the controls and says which step is missing rather than drawing something meaningless. Chapter 8 creates the epochs; the same views are alive from Chapter 9 onward.

Preprocessing

Where it lives




Figure 17 — Preprocessing and cropping share one menu.

The sliders button in the rail opens a two-item menu: Preprocess, which is filtering, referencing, interpolation, ICA and resampling; and Crop, which shortens the recording.

The presets

A screenshot of the 'Preprocess EEG Data' dialog box. The dialog has a title bar with a close button. Below the title bar, there are 'QUICK PRESETS' buttons: 'Auto', 'Minimal', 'Standard', and 'Thorough'. A descriptive text below these buttons explains the presets. The main area of the dialog contains three sections, each with a checkbox and a title: 'Bandpass Filter', 'Notch Filter (Line Noise)', and 'Re-reference'. The 'Bandpass Filter' section has input fields for 'HIGH-PASS CUTOFF (HZ)' (0.1) and 'LOW-PASS CUTOFF (HZ)' (40), with a note 'Removes slow drifts (0 = no high-pass)' and 'Removes high-frequency noise'. Below these are buttons for '0.1-40 Hz', '0.5-45 Hz', '1-30 Hz', and '0.1-100 Hz'. The 'Notch Filter (Line Noise)' section has a 'LINE FREQUENCY' input field set to '50 Hz (Europe, Asia, Australia)' and a checked checkbox for 'Include harmonics (100, 150 Hz, etc.)'. The 'Re-reference' section has a 'REFERENCE TYPE' input field. At the bottom right, there are 'Run Preprocessing' and 'Close' buttons.

Figure 18 — Preprocessing, with presets for the common cases.

Four presets set the checkboxes below them, and you can then adjust anything they set:

Preset	Sets
Auto	Filter, notch, average reference, and ICA with ICLabel choosing what to remove
Minimal	Filter and notch
Standard	Filter, notch and average reference
Thorough	Standard, plus bad-channel interpolation and ICA you inspect yourself

Everything below the presets is a card you can turn on or off independently: the bandpass filter, the notch and its harmonics, the reference, bad-channel interpolation, ICA, and resampling.

Choosing the band

Preprocess EEG Data

QUICK PRESETS

☐ Auto ☒ Minimal ☐ Standard ☐ Thorough

Auto: Filter, notch, reference and ICLabel, in one click | Minimal: Filter + notch | Standard: Filter + notch + reference | Thorough: Standard, plus bad-channel interpolation and manual ICA

☒ **Bandpass Filter**

HIGH-PASS CUTOFF (HZ) LOW-PASS CUTOFF (HZ)

0.1 100

Removes slow drifts (0 = no high-pass) Removes high-frequency noise

0.1-40 Hz 0.5-45 Hz 1-30 Hz 0.1-100 Hz

☒ **Notch Filter (Line Noise)**

LINE FREQUENCY

50 Hz (Europe, Asia, Australia) ☒ Include harmonics (100, 150 Hz, etc.)

☒ **Re-reference**

REFERENCE TYPE

Run Preprocessing Close

Figure 19 — 0.1–100 Hz, a 50 Hz notch, and an average reference.

The bandpass has four preset bands beside the cutoff boxes — 0.1–40, 0.5–45, 1–30 and 0.1–100 Hz — and takes any pair you type.

The notch defaults to 50 Hz with harmonics included; switch it to 60 Hz outside Europe, Asia and Australia. Getting this wrong does not fail, it simply leaves the line noise in place and removes a band that had nothing wrong with it.

The reference offers average, REST, Cz, linked ears, or channels you name. Average is the usual choice for a high-density array and is what this guide uses.

THE ORDER IS FIXED, WHATEVER ORDER YOU TICK THEM IN

The workbench applies the bandpass first, then the notch, then bad-channel interpolation, then the reference, then resampling, and ICA last. That matters most for ICA: it wants filtered, referenced data, and fitting it on an unfiltered recording wastes components on drift.

What it reports




Figure 20 — Each step listed as it completes.

Preprocessing does not simply succeed. It lists what it did, in order, with the parameters it used — and it is worth reading rather than dismissing, because a step that silently did nothing looks exactly like one that worked.

The panel follows the data

File:	rc01.edf
Format:	EDF/EDF+ (.edf)
Sampling Rate:	2048 Hz
Channels:	70
EEG Channels:	70
Duration:	45s
Samples:	94,207
Bad Channels:	None
High-pass:	0.1 Hz
Low-pass:	100 Hz

Quick Actions

 **Apply Standard Filter**

 **Average Reference**

Figure 21 — The filter bounds have changed.

The Info panel is read from the recording rather than from what you asked for. The low-pass row moved from 417 Hz to 100 Hz here because the filter really was applied. If a row does not move, the step did not do what you thought.

1. filter

```
{"h_freq":100,"l_freq":0.1}
```

2. notch

```
{"freq":50,"freqs_filtered":  
[50,100,150,200,250,300,350,400,450,  
500,550,600,650,700,750,800,850,900  
,950,1000],"harmonics":true}
```

3. reference

```
{"channels":[],"type":"average"}
```

Figure 22 — The history travels with the recording.

The History tab records every step with its parameters — including, for the notch, every harmonic it removed. This is the provenance record, and it leaves with the file on export.

The spectrum afterwards




Figure 23 — The same spectrum, filtered.

Compare this with the spectrum before filtering. Below about 1 Hz the drift is gone; the band-power bars beside it have shifted accordingly. The 50 Hz notch is not visible here for the reason given earlier — the plot ends where the notch begins.

Taking artifacts out with ICA

Independent component analysis separates the recording into components that mix to produce it, so that the ones carrying blinks, muscle and mains hum can be removed without touching the rest.

Fitting ICA on its own

Preprocess EEG Data

☐ Interpolate Bad Channels

CHANNELS MARKED AS BAD

Fp1
AF7
AF3
F1

Select channels to mark as bad, then enable interpolation

No channels marked as bad

☒ ICA Artifact Detection Show

ARTIFACT REMOVAL MODE

☒ Manual Inspection ☐ Auto (ICLabel)

Manual: Review and select components to remove

NUMBER OF COMPONENTS: 20
Typically 15-25 for standard EEG

METHOD: FastICA (fast)

① After ICA fitting, you'll review components and select which to remove.

Run Preprocessing Close

Figure 24 — ICA on its own.

ICA is a card in the same preprocessing dialog, marked **Slow** because it is. Untick the other cards to run it alone on data you have already filtered — which is what this guide does, and what you should do rather than re-filtering.

Number of components defaults to 20, which is a reasonable number for a standard array; 15–25 is the usual range. More components take longer and split the signal more finely, which is not automatically better.

Method is FastICA, Infomax or Picard.

Artifact removal mode is the choice that matters:

Manual Inspection fits the components, classifies them, and hands them to you to confirm. *Auto (ICLabel)* fits, classifies and removes in one go, with a brain-probability threshold you set.

AUTO MODE REMOVES COMPONENTS YOU NEVER SEE

It is the right choice for a batch of recordings processed the same way, and the wrong one for the first time you look at a dataset. Manual takes one extra click and shows you what is about to be thrown away.

Inspecting the components




Figure 25 — Every component, classified, with the artifacts pre-selected.

Each card is one component: its number, the class ICLabel gave it, the first five seconds of its own time course, and the probability that it is brain activity. Components ICLabel did not call brain are pre-selected for exclusion — ten of twenty here — and the count at the foot of the dialog tracks what you change.

Click a card to include or exclude it. Select All Artifacts and Clear Selection reset the selection wholesale.

WHAT THE CARD SHOWS, AND WHAT IT DOES NOT

The card shows the component's time course, not its scalp topography. A class and a brain probability are what you are judging on. Where that is not enough — a component you suspect is a single bad electrode rather than a blink — the honest move is to exclude it, apply, and compare the result against the unexcluded recording rather than to guess from the thumbnail.

Apply ICA & Remove Selected subtracts the selected components from the recording and adds two entries to the history: the fit and the exclusion, with the component numbers named. Cancel leaves the fit in place and removes nothing.

Cropping




Figure 26 — Cropping, with what it will keep stated before you commit.

Crop shortens the recording to a window. The overview at the top draws the whole file with the kept region highlighted, the start and end boxes take seconds, and the summary below states what will go, what will remain, and what percentage of the data that is.

Crop is destructive to the loaded session — the removed part is gone until you import again — and it is recorded in the history with both durations.

Events, annotations and epochs

The Events button opens one dialog with three tabs. They are easy to confuse, so: **EEG Events** are the markers your file already contains. **Annotations** are spans you add yourself. **Epochs** are the cuts made from either.

Annotations you add yourself




Figure 27 — Nothing marked yet.

An annotation is a labelled span: eyes closed, a task block, a stretch you want to remember was contaminated.

Figure 28 — One annotation.

Start and end in seconds from the beginning of the recording, a label, and an optional description.




Figure 29 — Two annotations over one recording.


They can be exported to JSON and imported back — the same shape the Wearable Workbench uses, so a protocol file written once can annotate both.



Figure 30 — Annotations shade every channel on the same clock.

They are not decoration. Each one shades the full height of the chart across its span, so a change at the start of a block is visible in every channel at once, and they follow the window as you move through the recording.

The markers the file already carries



The screenshot shows the 'Events & Annotations' window with the 'EEG Events' tab selected. A table lists five events with their sample numbers, times, IDs, and descriptions. To the right, an 'Event Info' panel provides summary statistics.

SAMPLE	TIME (S)	EVENT ID	DESCRIPTION
2840	1.387	1	#TgtM_RC01_SOA63
3528	1.723	3	#start
10511	5.132	13	MNOPQR
10818	5.282	14	STUVWX
11125	5.432	15	YZ1234

Event Info

Events are markers embedded in the EEG file (triggers, stimuli markers, etc.)


Total events: 243
Unique types: 15

Figure 31 — The markers this recording carries.

Events come from the file itself — the trigger channel or the EDF+ annotations — and are read at import. The table gives the sample, the time in seconds, the numeric ID and the description. The panel beside it counts them: 243 events of 15 distinct types in this recording.

You cannot edit these. They are what the acquisition system wrote, and the workbench treats them as fact.

Cutting epochs



The screenshot shows the 'Events & Annotations' window with the 'Epochs' tab selected. It features two main panels for creating epochs: 'Create from EEG Events' and 'Create from Annotations'. The 'Create from EEG Events' panel includes fields for 'EVENT TYPE' (set to 'All events'), 'TIME BEFORE (S)' (-0.2), 'TIME AFTER (S)' (0.8), and 'BASELINE CORRECTION' (Mean (tmin to 0)). The 'Create from Annotations' panel includes a 'LABEL FILTER' (All labels), a checked 'Use label duration as epoch length' option, 'PRE-LABEL (S)' (0), 'POST-LABEL (S)' (auto), and 'BASELINE CORRECTION' (None). A 'Current Epochs' section at the bottom states 'No epochs created yet. Create epochs from EEG events or from your annotations to enable ERP and time-frequency analysis.'

Figure 32 — Cutting an epoch around every marker.

An epoch is a window around a marker. Choose the event type — or all of them — then how much to keep before and after, and whether to baseline-correct.

-0.2 to 0.8 seconds with a mean baseline over the pre-stimulus period is the default and the usual starting point: enough baseline to correct against, and enough afterwards to contain a P300.

Events & Annotations

Annotations | EEG Events | **Epochs**

Create from EEG Events

EVENT TYPE
All events

TIME BEFORE (S)
-0.2

TIME AFTER (S)
0.8

BASELINE CORRECTION
Mean (tmin to 0)

Create from Events

Create from Annotations (when no EEG events)

LABEL FILTER
All labels

☒ Use label duration as epoch length

PRE-LABEL (S)
0

POST-LABEL (S)
auto
Leave empty for label duration

BASELINE CORRECTION
None

Create from Labels

ⓘ EEG events available - use events for epoch creation

Current Epochs

Epochs: 243

Time range: -0.200s to 0.800s

Source: EEG Events

By type:

- #TgtM_RC01_SOA63: 1
- #end: 1
- #start: 1
- 56789_: 20
- ABCDEF: 20
- AGMSYS: 20
- BHNTZ6: 20
- CIOU17: 20

Figure 33 — The epochs it produced.

The result states how many epochs it made, over what time range, from which source, and how many came from each event type. This recording yields 243 epochs, most types contributing 20 apiece.

EPOCHS FROM ANNOTATIONS, WHEN THERE ARE NO EVENTS

The right-hand card cuts epochs from your own annotations instead, using each label's own duration or a fixed window. It is disabled while the file has events of its own, and says why: events are the more precise source, and mixing the two would make the average mean two different things.

Event-related potentials

Everything in this part is an average across the 243 epochs, and every plot names the count in its title so a figure cannot be mistaken for a single trial.

The average at one channel



Figure 34 — The averaged response at one channel.

One channel, averaged across epochs, with a confidence band around it. Time is in milliseconds from the marker, and zero is the marker itself.

Every channel at once



Figure 35 — Every channel overlaid.

The butterfly plot overlays all 66 scalp channels — the four EOG derivations are left out — with global field power below it. GFP is the spread across channels at each instant: peaks in it are the moments when the scalp map had structure, which is where topographies are worth looking at.

Trial by trial




Figure 36 — Trial by trial, rather than averaged.

The ERP image is one row per epoch, colour for amplitude, with the average underneath. It answers the question an average cannot: whether a component is present in most trials or driven by a handful. Rows can be sorted by amplitude or latency.

The scalp at chosen latencies




Figure 37 — The scalp at four latencies after the marker.

Topographies of the average at the latencies you name — 100, 200, 300 and 400 ms by default, and any comma-separated list you type.

Measuring a peak



Figure 38 — Measuring a peak in a named window.

Name a search window in milliseconds, choose whether you are after the most positive, most negative or largest absolute deflection, and Measure reports the amplitude, the latency, and the mean over a 20 ms window centred on the peak.

STATE THE WINDOW, ALWAYS

A peak is only meaningful with the window it was found in. The workbench prints the window and the mode beside every result — **250 to 500 ms, pos** here — because a latency quoted without them cannot be checked or reproduced.

ERSP and ITC

With epochs in place, the two disabled time-frequency modes come alive. Both are per-channel, both run 1–50 Hz across 40 log-spaced frequencies, and both are computed with Morlet wavelets over the epoch window rather than the recording window.




Figure 39 — Power change against baseline, across trials.

ERSP is event-related spectral perturbation: power at each time-and-frequency point as a percentage change from the pre-stimulus baseline, averaged across trials. Red is more power than baseline, blue is less.




Figure 40 — How consistently phase lines up across trials.

ITC is inter-trial coherence: how consistently the phase at each time-and-frequency point lines up across trials, from 0 to 1. It is a measure of consistency, not of size — a response too small to see in the average can still be phase-locked, and ITC is where it shows.

Statistics




Figure 41 — Band power, spectral metrics, coherence and asymmetry.

The Statistics dialog computes a summary over the loaded recording:

Relative band power — delta through gamma, each as a percentage of the total power between 0.5 and 45 Hz. The five partition that range, so they sum to a hundred and a change in one is a change in the balance rather than in absolute power.

Spectral metrics — peak alpha frequency, the 1/f slope, spectral power variability and signal entropy.

Global and spatial — global field power, frontal alpha asymmetry, and average coherence.




Figure 42 — The event-related measures.

Derived indices — theta/beta ratio, a cognitive engagement index and a vigilance index, each a ratio of band powers.

Event-related potentials — P300 amplitude and latency, which require epochs to exist.

Microstate metrics — average microstate duration and coverage.

Save to File writes the whole set out as a report.

A MEASURE THAT COULD NOT BE COMPUTED IS BLANK, NOT A WORD

Several of these need something the recording may not have — P300 needs epochs, frontal alpha asymmetry needs frontal channels on both sides, coherence needs more than one channel. When a measure cannot be computed the panel shows an em dash and the saved report leaves the value blank. It never writes N/A : a word in a column of numbers turns the whole column into text the moment the report is read by anything else, and the measure then drops out of every calculation without an error. The recording used here produces all twenty, so none of them is blank above.

THESE ARE SUMMARIES OF ONE RECORDING, NOT CLINICAL FINDINGS

Every number here is computed over the whole loaded recording with the preprocessing you applied, and changes when either does. The derived indices in particular are ratios with a large literature behind them and no single agreed definition; quote them with the band edges and the reference you used, or not at all.

Exporting

Export EEG Data

Exporting: 70 channels, 46.0s duration, 2048 Hz, 4 preprocessing steps applied.

EXPORT FORMAT

☒ **FIF** (.fif) - MNE native format, preserves all metadata

☐ **EDF** (.edf) - European Data Format, widely compatible

☐ **BrainVision** (.vhdr/.vmrk/.eeg) - BrainVision format

☐ **EEGLAB** (.set) - EEGLAB format for MATLAB

☐ **CSV - epochs** (.csv) - one row per sample, labelled by epoch, for model training

☐ **CSV - raw** (.csv) - the continuous signal; large recordings are refused

☐ **Analytics Workbench** (.zip) - one tidy table of derived measures, with its manifest

FILENAME

rc01_processed

.fif

The file will be downloaded to your browser's download folder.

☐ Export epochs only (243 epochs instead of continuous data)

INCLUDE IN EXPORT

☒ Events/Annotations

☒ Bad channel markers

Cancel

Export

Figure 43 — Six formats: four recording formats, raw CSV, and a table for the Analytics Workbench.

The banner states what is about to be written — channels, duration, sampling rate and how many preprocessing steps have been applied to it.

Format	Use it for
FIF	Anything going back into MNE. Preserves the most metadata
EDF	Wide compatibility with clinical software
BrainVision	Brain Products tooling; writes three files
EEGLAB	MATLAB
CSV — epochs	One row per sample, labelled by epoch. The route into model training
CSV — raw	The continuous signal, for when nothing else will do
Analytics Workbench	Derived measures as a tidy table, for the Analytics Discovery Workbench

Export epochs only switches the export from the continuous recording to the 243 epochs. EDF and BrainVision cannot hold epochs and refuse rather than writing something misleading; FIF, EEGLAB and CSV can.

Version 1.0 · August 2026

Page 40 of 51

Include in export decides whether events and annotations, and bad-channel markers, travel with the data.

EVERY EXPORT IS A ZIP, AND THE HISTORY IS IN IT

Alongside the data file the workbench writes a plain-text preprocessing history — the file it came from, its properties, and every step with its parameters and timestamp — and zips the two together. That text file is the provenance record for whatever you do next, and it is the reason a CSV export is not the dead end it usually is.

CSV HAS A CEILING

A CSV export is refused past 40 million cells, with the row and column count in the message. Seventy channels at 2048 Hz reach that in about five minutes of recording. Export epochs, or crop first.

Exporting to the Analytics Discovery Workbench

The five formats above all write the **signal**. This one writes **derived measures**: one tidy table where rows are observations and columns are variables, which is the shape the Analytics Discovery Workbench reads. Band powers, spectral slope, entropy, asymmetry, coherence and the ERP measures — the same numbers the Statistics panel reports, from the same calculation, so the two cannot disagree.

Export EEG Data

① Exporting: 70 channels, 46.0s duration, 2048 Hz, 4 preprocessing steps applied.

EXPORT FORMAT

☐ FIF (.fif) - MNE native format, preserves all metadata

☐ EDF (.edf) - European Data Format, widely compatible

☐ BrainVision (.vhdr/.vmrk/.eeg) - BrainVision format

☐ EEGLAB (.set) - EEGLAB format for MATLAB

☐ CSV - epochs (.csv) - one row per sample, labelled by epoch, for model training

☐ CSV - raw (.csv) - the continuous signal; large recordings are refused

☒ Analytics Workbench (.zip) - one tidy table of derived measures, with its manifest

One row per observation, named so the Analytics Workbench reads the units and the roles without being told. Signal samples are not exported.

SUBJECT *

SESSION

rc01.edf

Rows cannot be told apart without it.

ONE ROW PER

Whole recording

70 channels, 243 epochs, 15 condition(s).

► Participant details (optional)

☒ EEG channels only (otherwise every channel in the file)

Downloads a .zip holding the table and its manifest, which the Analytics Workbench reads directly.

Enter a subject identifier to see what this export holds.

FILE NAME

Figure 44 — The export is refused until a subject identifies the rows.

A subject identifier is required and the Export button stays disabled until you give one. Once this table is loaded beside another recording — which is the point of writing it — rows with no subject cannot be told apart. **Session** defaults to the filename; name the visit if you are building a cohort.

One row per what

One row per observation, named so the Analytics Workbench reads the units and the roles without being told. Signal samples are not exported.

SUBJECT * **SESSION**

NS-101 rc01.edf

Rows cannot be told apart without it.

ONE ROW PER

Channel

70 channels, 243 epochs, 15 condition(s).

▼ Participant details (optional)

GROUP **AGE** **SEX**

A field left blank is left out of the table rather than written as an empty column.

☒ EEG channels only (otherwise every channel in the file)

Downloads a .zip holding the table and its manifest, which the Analytics Workbench reads directly.

64 rows × 31 columns, about 23.98 KB.

Figure 45 — One row per channel, with the shape of the table stated before it is written.

Granularity	Rows	Use it for
Whole recording	one	One row per subject per visit — the cohort table
Channel	one per channel	Topographic comparison; where on the scalp a measure differs
Epoch	one per epoch	Trial-by-trial analysis, with the condition on every row

Epoch is offered only once epochs exist. Until you cut them it is greyed out and says so; see *Cutting epochs*. The hint under the picker states what this recording holds — channels, epochs and conditions — so you can see what each choice would produce before choosing it.

EEG channels only is on by default. The demo recording has 70 channels of which 64 are EEG type, so a channel export is 64 rows rather than 70; turn it off to include the rest.

What comes out

The whole-recording table carries the same measures as the Statistics panel. The per-channel table replaces the global measures with per-channel ones — relative and absolute band power, peak alpha frequency, RMS amplitude, SD, peak-to-peak, spectral slope and entropy — and adds the channel's type and whether it is marked bad. The per-epoch table gives band powers, global field power, amplitude and, where the epoch window reaches 250–500 ms, a P300 amplitude and latency.

Identity and recording context — subject, session, sampling rate, channel count, reference, filter settings and the number of preprocessing steps — travel on every row, so the table still describes itself once it is sitting beside a table from a different instrument.

THE MANIFEST IS NOT OPTIONAL

The download is a `.zip` holding the table and a `_analytics.json` beside it, carrying the units, the column roles and this session's preprocessing. It matters most for a channel export: 64 distinct channel names in 64 rows read as an identifier to anything guessing from the data alone, and only the manifest says `Channel` is a grouping factor. Without it the granularity is inert.

Version and support




Figure 46 — About.

About names the version you are running. Quote it in any bug report, along with the format of the file you imported — most reproducible problems are a format-and-version pair.

Appendix A — Control reference

The rail

Control	Effect
Reset	Clears the loaded recording
Import	Opens the import dialog; also replaces the current recording
Export	FIF, EDF, BrainVision, EEGLAB, CSV, or a table for the Analytics Workbench — always a zip
Preprocessing menu	Preprocess, and Crop
Channels	The channel list
Events	Annotations, EEG events and epochs
Statistics	The summary measures
About	Version

Toolbar

Control	Effect
Time	Window start, in seconds
Duration	5, 10, 20, 30 or 60 seconds
Step and jump	A window at a time, or to the start or end
Arrow keys, Home, End	The same four moves from the keyboard
View buttons	Multi-channel, single channel, PSD, topographic, time-frequency, ERP

The right-hand panel

Tab	Shows
Info	What the file declares, and the two Quick Actions
Channels	Every channel with its type and region colour
History	Every preprocessing step with its parameters

Preprocess

Card	Effect
Bandpass filter	High-pass and low-pass cutoffs, with four presets
Notch	50 or 60 Hz, harmonics optional
Re-reference	Average, REST, Cz, linked ears, or named channels
Interpolate bad channels	Rebuilds channels marked bad from their neighbours
ICA	5–100 components; FastICA, Infomax or Picard; manual or ICLabel
Resample	A new sampling rate, with three presets

Views

View	Needs	Shows
Multi-channel	—	Up to 32 channels stacked on one clock
Single channel	—	One channel at full scale
PSD	—	Spectrum to 50 Hz, with relative band power
Topographic	A montage	Voltage at an instant, band power, or all bands
Time-frequency	—	Spectrogram or Morlet TFR over the visible window
Time-frequency	Epochs	ERSP and ITC over the epoch window
ERP	Epochs	Waveform, butterfly, ERP image, topographies, peak measurement

Export

Control	Effect
Export Format	FIF, EDF, BrainVision, EEGLAB, CSV per epoch, CSV raw, or Analytics Workbench
Filename	The name inside the zip
Export epochs only	Writes the epochs rather than the continuous recording
Include in export	Whether events and bad-channel markers travel with the data

Shown only for an analytics export:

Control	Effect
Subject	Required. Identifies the rows once the table is loaded beside another
Session	Names the visit; defaults to the filename
One row per	Whole recording, channel, or epoch — epoch needs epochs to exist
EEG channels only	On by default; off includes every channel in the file
Participant details	Group, age and sex; a blank field is left out rather than written empty

Appendix B — Supported formats

Format	Extension	Also needs
EDF / EDF+	<code>.edf</code>	—
BDF	<code>.bdf</code>	—
MNE-Python FIF	<code>.fif</code>	—
BrainVision	<code>.vhdr</code>	<code>.vmrk</code> and <code>.eeg</code> , uploaded together
EEGLAB	<code>.set</code>	<code>.fdt</code> , when the data is in a separate file

Export offers the same five, plus CSV of either the continuous signal or the epochs.

Appendix C — What the workbench refuses to do

A refusal is working as designed. These are the ones you are most likely to meet.

An ERP, ERSP or ITC without epochs. All three are averages across trials. The controls are disabled and the view says which step is missing.

Epochs from annotations while the file has events of its own. Events are the more precise source. The card is disabled and says so, rather than letting you build an average from two different definitions of the same moment.

A topographic map without a montage. Electrode positions are not in the signal. The view says `Montage Required` and names the fix.

Epochs in EDF or BrainVision. Neither format can represent them. The export is refused rather than silently writing the continuous recording instead.

A CSV export past 40 million cells. The message gives the row and column counts and suggests the two ways out: crop, or export epochs.

An upload over 1 GB. A per-request cap, adjustable in `config.py` if your recordings need it.

Appendix D — The recording used in this guide

File	<code>rc01.edf</code> , with a companion <code>rc01.edf.event</code>
Format	EDF+, 12.6 MB
Shape	70 channels at 2048 Hz, 46 seconds, 94,207 samples per channel
Channels	64 scalp positions in the 10-05 layout, two ear references, four EOG derivations
Markers	243 events of 15 types, carried in the EDF+ annotations
Where	<code>DemoData/</code> in the development tree. It is not shipped with the packaged application

The events are read from the EDF+ annotations inside the recording itself; the `.event` companion beside it is not used by the workbench.

The two annotations in Chapter 8 — *Eyes closed* and *Task* — were added by hand to demonstrate annotating, and are not part of the file.