

VarianceLabs
User guide

Mineral Resource and Reserve Estimation Assistant: app and tools

Version 0.3.6

Contents

1. **Projects and the Data page**
    - Which project
    - Loading data
    - The Data page
    - Pages remember your view
    - Saving figures
    - Tools
2. **Data processing: desurvey and composites**
    - Use
    - Results
    - Tools
3. **Data quality (QAQC)**
    - Load the QC files
    - Use
    - The QAQC page
    - Reviewing flags
    - Settings
    - Automatic choices
    - Tools
4. **Exploratory data analysis (EDA)**
    - Use
    - The EDA page
    - Automatic choices
    - Tools
5. **Geometry: wireframes and block models**
    - Use
    - Results
    - Block models
    - Tools
6. **Property modelling**
    - Set up
    - Use
    - Results
    - Tools

Ask the AI Agent in your own words: the quoted requests in each chapter are examples, not commands to learn. The AI Agent picks the tools.

Projects and the Data page

Everything you load or build — drillholes, composites, wireframes, block models, estimates — lives in a **project**, in `variancelabs-output/projects/<project-id>/` next to where Claude was started. The project keeps a copy of your files, the parameters of every object and what it was built from, so any object can be changed later and everything built from it follows.

Which project

You don't need to name one: with a single project in the folder, Claude uses it (the first time, one is created and named after the folder). With several, Claude asks which one. Ask for a new project to keep work separate.

Loading data

> Load `collar.csv`, `survey.csv` and `assays.csv`.
> Load `composites.csv` — it is already composited.
> Load the QC files `standards.csv` and `duplicates.csv` for these drillholes.

Raw tables stay editable (desurvey parameters, compositing). Tables that were desurveyed or composited elsewhere (X, Y, Z per interval) can be shown and used for wireframes and block models, but their processing cannot be changed: the raw data is not in the project.

The Data page

Open it from the database tile in the app, or ask Claude to open the project. It lists every object: drillholes with their composites underneath, wireframes, block models with their estimates underneath, each with a one-line summary and its status.

Click an object for its menu:

Action	What it does
Edit parameters…	shows every parameter with a short explanation; values detected from your data say what was detected and why. Change any (↺ **Use detected** goes back), then **Validate & rebuild**: the object and everything built from it are rebuilt in order. NEXUS estimates: **Validate & re-run**
Rename	names are yours; they don't have to be unique
Remove	also removes what was built from it (you are asked first)
Retry	for an object whose rebuild failed, or an estimate that is out of date (re-runs it)
View details…	LVA-DKNN estimates and QC data: held-out validation and the parameters used, read-only
Cancel run	stops a running estimate (a first run disappears; a re-run keeps the previous result)

Main commodity (top right): pick the column, e.g. AU1, that EDA, QAQC and the 3D viewer open on, wherever an object has it (upper/lower case ignored): drillholes are coloured by it and block models by their estimate of it. Each plot can still switch to another column. *Not set* goes back to the first column. You can also ask Claude to set it.

Element groups (under the list): each assay column as a chip in its geochemical group, used by the correlation matrix. Claude assigns them from geochemical knowledge when you load drillholes; common names (Cu, AU1,

Fe2O3_pct) start from a preset. Drag a chip to another group; click one to change its element or group, or to mark it as not an element (RQD, density). Claude does not overwrite what you set unless you ask.

Below detection and missing values: each assay value is read when the data is used, and your file is never changed. Click a chip to see what was found and how it was read.

In the file	Read as
`<0.01`, `BDL`, `ND`	below detection, set to half the limit
`NS`, `N/A`, `LNR`, blank	missing
`-9`, `-99`, `-999`	missing
`-0.01`	below detection with a limit of 0.01, unless too large to be a limit (over the column's median): then missing
`0` filling a whole hole of 5 or more samples	missing (not assayed)
other `0`	below detection
`>100`	above limit, kept at 100
a repeated lowest value, no limit stated	below detection, value kept (it may already be half the limit)
any other number	kept as the lab reported it, even under a limit stated elsewhere in the column (you get a warning)

Only columns read as assays are treated: the name is an element, a value carries a `<` or `>` code, or you say so. Lithology, sample ids and domain codes stay as they are. In the chip you can set the limit (it does reclassify the numbers at or under it), the codes that mean missing, how zeros read, and whether it is an assay column; this rebuilds the data and what is built from it. Ask Claude to do it with `project_update` and `assay_values`. QAQC audits the file as loaded, so it still reports these codes.

QC data (see QAQC) is listed under the drillholes it controls, with its counts (CRM results, blanks, duplicates); it is removed with them.

A drillholes panel groups the desurvey settings (method, azimuth correction, dip sign, coordinates) and the columns used, and lists its composites: **Edit ›** opens one, **+ Add composites** creates new ones from an interval length and a method. A block-model panel lists its estimates the same way, with **+ Add estimate** (see Property modelling).

Imported objects (a wireframe brought in from other software, or a table loaded already desurveyed or composited) show **View details…** instead of Edit: their parameters are not editable.

A failed rebuild keeps the previous version and shows the error. If the object you edited is the one that fails, what depends on it keeps its own current version untouched; only a failure further downstream marks what depends on *that* object *out of date*. An estimate is never rebuilt automatically: when what it was trained on changes, it is marked *out of date* until you retry it. The 3D viewer updates by itself within 2 seconds. Claude relays the warnings of everything that was rebuilt.

Setting `VARIANCELABS_OUTPUT_DIR` makes every project folder live under that one path, regardless of the working folder — project resolution then spans every folder that used it, not just the current one.

On Windows, a project's files add about 140 characters under the working folder path: if loads start failing with a path-length error, keep the working folder path short, or turn on `LongPathsEnabled`.

Pages remember your view

Each page reopens the way you left it, also after closing the browser or in a new Claude session: in the 3D viewer the camera, what is shown or hidden, the colours and the selection; in EDA and QAQC your data choices, the panels you switched off or collapsed and their settings. The view is saved in the project, so a copy of the project opens on it too. New objects always appear, shown. A link from Claude opens on what it describes.

Saving figures

Every chart of the EDA and QAQC pages, and the 3D viewer, has a **save as image** button (a picture with an arrow, beside the zoom-reset and maximize buttons; in the viewer, at the end of the toolbar). It asks for a title and a subtitle (filled in for you), a folder of the project (`figures` the first time, then the last one you used) and a file name, and saves a PNG on a white background, laid out as a report figure: the title, the subtitle, the framed chart or 3D view, a legend row with counts, and a footnote naming the data, the project, the app version and the date. Maps get a north arrow and a scale bar. A small `.json` file beside it says what the figure shows.

Claude saves figures too, for a report it writes, and writes their title: it takes the figure from the page you have open, as you see it; when no page shows it, the server draws it itself (this needs Edge or Chrome; the 3D view is then seen from the default angle). Ask for a subtitle, a footnote (e.g. the coordinate system), a dark image, the on-screen layout, or another folder if you want one.

Tools

Claude picks the tools; name one, or a parameter, to be specific. **Bold** parameters are required; the others are optional and, when left out, detected from your data. Every tool that works on a project also takes `project`: its name or id (by default the only one in the folder). Ask Claude for this guide at any time.

Tool	Use it to	Parameters
`project_open`	open the project: every object, its parameters and status	`open`: a page to open (data, viewer, eda or qaqc)
`project_create`	start a new, empty project	**`name`**; `open`
`project_load_drillholes`	load drillhole tables	raw tables `collar`, `intervals`, `survey`; or `desurveyed`, one table that already has X, Y, Z, with `load_as` drillholes or composites; `mapping`: which column is the hole id, from, to…; `desurvey`: the desurvey settings of Data processing; `name`
`project_load_qaqc`	load QC files for drillholes (see QAQC)	**`drillholes`**; `standards`, `reference`, `duplicates`, `blanks`; `mapping`: columns, blank ids, duplicate types, element names; `name`
`project_load_wireframe`	import an OBJ wireframe from other software	**`path`**; `name`
`project_update`	change an object's parameters, then rebuild it and everything built from it (Edit on the Data page)	**`object_id`**; **`params`**: the values to change, `auto` for the detected one; none re-runs an out-of-date estimate
`project_remove`	remove objects	**`object_ids`**; `cascade`: also remove what was built from them (Claude asks you first)
`project_set_main_commodity`	set the main commodity	**`column`**, e.g. AU1; empty clears it

Tool	Use it to	Parameters
`project_set_element_groups`	say which element and group each assay column is	**`columns`**; `group_order`, `overwrite`
`figure_save`	save a figure of the app (EDA or QAQC chart, 3D view) as an image for a report	**`page`**: eda, qaqc or viewer; `figure`: its title as shown (the start is enough); `view`: what to show instead of your view; `title`, `subtitle`, `note`: the words on the image; `folder` (default `figures` in the project); `name`; `theme`: light or dark; `layout`: report or screen; `width` in pixels; `side_notes`: Claude's interpretation beside the correlation matrix
`doc_open`	open this guide	`open`: off only gives the file's location

Data processing: desurvey and composites

Turns drillhole tables (collar, survey, assay intervals) into located samples: X/Y/Z down the hole (desurvey) and regular-length composites. Available through the MCP; the 3D viewer and the project use the same hole geometry.

Use

Ask Claude, for example:

> Composite the assays in `C:\data\assay.csv` with `collar.csv` and `survey.csv`.

That's enough: the composite length, the columns to average, the coordinate system (UTM or local) and the dip sign are chosen from your data. Claude reports what it used, for example "*length 1 m: median sample length 1.05 m*". To change anything, say so:

> Composite at 2 m, restart at LITH changes, and drop short end pieces.

You can set	Default
Composite length (`length`)	most common sample length, rounded
Columns to average (`columns`)	all numeric assay columns (not sample ids or interval length)
Domain column (`domain`, e.g. LITH)	none — Claude will suggest candidates
Composite method (`residual`, `min_fraction`)	fixed length, a short end piece (under half the length) merged into the previous composite; or keep / drop short ends, or *best fit* (length adjusted so each run splits evenly)
Start (`start`)	at the first sample; or at the collar
Unsampled gaps (`gaps`, `min_coverage`)	ignored in averages (or counted as zero); composites less than half covered are left blank
Coordinates (`coordinate_system`)	UTM if present, else local; never mixed between collar and survey
Dip sign (`dip_convention`)	negative down, unless every dip is positive
Desurvey method (`method`)	minimum curvature; or balanced tangent, or tangent (straight lines)
Azimuth correction (`azimuth_correction`)	0°; added to every azimuth (e.g. magnetic declination)

To add X/Y/Z to your samples without compositing:

> Desurvey `assay.csv` with `collar.csv` and `survey.csv`.

Results

Composites are added to the project under their drillholes (Data page), where their parameters can be edited; the CSV is in `objects/<object-id>/v<version>/composites.csv`. Desurvey (not part of a project load) writes to `variancelabs-output/processing/<run-id>/`:

File	Content
`composites.csv` / `desurveyed.csv`	your rows with X, Y, Z (composites also LENGTH, COVERAGE)
`report.json`	parameters used and why, skipped holes, warnings, before/after means

Check the report's skipped holes (e.g. holes with no assays) and the before/after means. Composites can go straight to property modelling or into the 3D viewer.

Tools

Tool	Use it to	Parameters
`processing_composite`	composite drillholes, with X, Y, Z at the midpoints; added to the project	**`drillholes`**: a drillholes object, or the collar, intervals and survey files (with their `mapping` and `desurvey` settings); `options`: the compositing settings above; `name`
`processing_desurvey`	add X, Y, Z to a table without compositing (a file, not a project object)	**`collar`**; **`table`**: intervals (hole, from, to) or points (hole, depth); `survey`; `at`: interval midpoints (`mid`) or both ends (`from_to`); `options`: `mapping` and the desurvey settings above

Data quality (QAQC)

Checks that the samples can be trusted before estimation: was the laboratory accurate (certified reference materials, CRMs), clean (blanks) and precise (duplicates and repeats), and are the drillhole data sensible (assay values, collars, surveys, intervals, grades far above their neighbours). Claude runs the checks through the MCP and reads their numbers; the **QAQC page** of the app (rail tile after Data) shows the same checks as interactive charts. Every check raises **flags** on suspect samples, holes or batches, which you review.

Load the QC files

QC files are optional; the grade, geometry and spatial checks need only your drillholes.

> Load the QC files standards.csv, reference.csv and duplicates.csv for these drillholes.

File	Holds
standards	CRM (and blank) results: standard id, one column per element, batch and date if you have them
reference	certified values per standard, optionally an `<ELEMENT>_SD` column
duplicates	duplicate and repeat results: sample id, the original's sample id, a type (field, coarse, pulp, repeat, umpire)
blanks	blank results, if kept apart from the standards

At least one of standards, duplicates or blanks. The files are copied into the project and appear on the Data page under their drillholes. Columns, blank ids and duplicate types are detected and Claude reports what it found; ask for a different choice if one is wrong. If the QC files name an element unlike your assays (Au_ppm for AU), ask Claude to map it. In your drillholes, the sample id, batch and recovery columns are detected by name. Re-assay columns in your assays (AU1, AU2…) are compared automatically as pairs.

Use

> Run the QAQC of the project and tell me what needs attention.
> Check the CRMs for gold: any drift or bias?
> List the failed flags of the drillhole geometry checks.

Claude starts with an overview (which panels have data, a verdict for each, the open flags), then goes through the panels and gives you a link that opens the QAQC page on each result. Open the page yourself with "*open the QAQC page*".

The QAQC page

The left column picks the drillholes, their QC data and the element, and switches panels on or off. A panel appears only when the project has what it needs; otherwise its switch is greyed with the reason.

Panel	Shows	Needs
Summary	verdict and open flags per panel, the automatic settings	any drillholes
CRM control chart	results against the certified value and SD, the 8 Nelson rules, bias per standard	standards with certified values

Panel	Shows	Needs
Blanks	results against the contamination threshold, share exceeding	blanks
Duplicates & repeats	HARD plot, cross plot, precision and bias per pair set	duplicates, or re-assay columns
Grade checks	text or missing-data values, duplicate or missing sample ids, spikes, unassayed runs, core recovery	drillholes
Drillhole geometry	collar, survey and interval checks, a summary per hole	drillholes (survey checks need a survey file)
Spatial outliers	samples far above their neighbours in other holes (primary elements only)	located samples
Issues	every flag, filterable by panel, check, severity, status, element and hole	any flag

Charts: wheel to zoom (over an axis: that axis only), drag a box to zoom in, right-drag to pan, double-click or the circular arrow at a chart's top right to reset, hover for values; the arrows next to it maximize the chart (Esc to close). The page address keeps every setting, so a link reopens the same view.

Reviewing flags

A flag has a severity (fail, warn or info). On the page, or by asking Claude, mark a flag **Resolve** (dealt with) or **Dismiss** (not an issue), with a comment; **Reopen** undoes it. The Issues panel does the same for all the open flags shown at once.

> Dismiss the zero-value info flags, comment "below detection".

Reviews are kept in the project and survive rebuilds; if the data changes so that a flag disappears, its review is listed as unmatched rather than lost. Claude and the page see the same flags and reviews.

Settings

A few thresholds can be set on the page (the field next to the chart they affect) or by asking Claude; clearing a field, or "*auto*", restores the automatic value.

> Set the gold detection limit to 0.01 and the HARD threshold of field duplicates to 25 %.

Automatic choices

Choice	Default
Detection limit (per element)	lowest positive assay if at least 5 % of the assays carry it, else the smallest positive value
Blank threshold	3 × the detection limit
CRM standard deviation	the certified SD, else the robust SD of the results
CRM verdict	fail beyond 3 SD or bias > 10 %; warn at bias > 5 % or a Nelson rule 2–6 or 8; needs 5 results
HARD threshold	10 % pulp, repeat, umpire and re-assay; 20 % coarse and other; 30 % field
Pair set verdict	pass when at least 90 % of pairs are within the threshold, warn at 80 %

Choice	Default
Minimum recovery	80 %
Spike	at least 10 × both neighbours
Spatial outlier	at least 10 × the upper quartile of the neighbours in other holes, within 2 × the median collar spacing
Run order	by date when the files have one, else by file row

Each analysis lists the choices it made and why, and its warnings.

Tools

The QC files are loaded with `project_load_qaqc` (see Projects). Most tools also take `drillholes` and `qc` (by default the first drillholes with QC data, and their first QC data) and `open` (opens the QAQC page on the result); `element` defaults to the first one.

Tool	Use it to	Parameters
`qaqc_overview`	start here: the panels with data, a verdict for each, open flags, settings	—
`qaqc_standards`	CRM control chart: bias and Nelson-rule events by batch	`element`; `standard`: one CRM (default the first)
`qaqc_blanks`	blank contamination	`element`
`qaqc_duplicates`	precision (HARD) and bias per pair set	`element`; `set`: one pair set, e.g. field (default the first)
`qaqc_checks`	grade, drillhole-geometry or spatial checks	**`family`**: grades, geometry or spatial; `element`
`qaqc_issues`	list the flags	filters `family`, `check`, `severity` (fail, warn, info), `status` (open, resolved, dismissed), `element`, `hole`; `limit` (50)
`qaqc_review`	resolve, dismiss or reopen flags	**`status`**: resolved, dismissed or open; `flag_ids`, or every flag of a `check` (optionally of one `element` or `hole`); `comment`; `only_open`: leave reviewed flags alone
`qaqc_settings`	read or change the thresholds	`values`: `detection_limits` (per element), `blank_factor`, `hard_thresholds` (% per pair type), `recovery_min` (%), `spike_ratio`, `spatial_ratio`; `auto` restores one

Exploratory data analysis (EDA)

Statistics and charts of the project's samples before estimation, and checks of an estimate after it. Claude runs the analyses through the MCP and reads their numbers; the **EDA page** of the app (rail tile under Data) shows the same analyses as interactive charts. EDA never changes the project.

Use

Ask Claude, for example:

> Do an EDA of the AU assays by lithology.
> Are there outliers in AU? Suggest a top cut.
> Draw a section through the holes along the main trend, coloured by LITH, with the AU estimate's blocks.
> Check the AU estimate: grade–tonnage curve at 2.7 t/m³ and swath plots.

Claude starts with an overview of what can be analysed, then runs the analyses one by one and gives you a link that opens the EDA page on each result. Open the page yourself with "*open the EDA page*".

The EDA page

The left column picks the **samples** (a drillholes or composites object) and an optional **group by** column (DOMAIN, LITH…), and switches panels on or off. A panel appears only when the project has what it needs; otherwise its switch is greyed with the reason.

Panel	Shows	Needs
Univariate statistics	histogram, probability plot, box plots by group, statistics	drillholes or composites
Bivariate analysis	correlation matrix grouped by element association (click a cell), below it the scatter of that pair	2 numeric columns
Outliers (top-cut)	decile analysis verdict, cap candidates with the metal each removes, cutting curve, plan of the samples above the suggested cap	drillholes or composites
Section	vertical section of the holes inside the corridor along A–A′, with the wireframes and blocks it cuts; the line's plan and display settings beside it	drillholes or composites
Grade–tonnage	tonnage (or volume) and grade above cut-off, table	a block-model estimate
Swath plots	sample mean against estimate mean along X, Y and Z	a block-model estimate

Charts: wheel to zoom (over an axis: that axis only), drag a box to zoom in, right-drag to pan, double-click or the circular arrow at a chart's top right to reset, hover for values; the arrows next to it maximize the chart (Esc to close). The page address keeps every setting, so a link reopens the same view; the page refreshes when objects are added or rebuilt.

Section

The column beside the section works like the 3D viewer's Display tab:

- **Line**: drag A, A′ or the corridor on the plan (holes inside it in white, the arrow shows which way you look); corridor width, **Flip** (look from the other side), **Reset**.
- **Samples**: property, colour by *Continuous* (range, log scale, *Auto* = 2nd–98th percentile), *Classed* (equal interval, quantile or custom breaks such as `0.3, 1, 3`) or *Categorical*; palette and reverse; interval width; hole labels. The legend gives the metres of samples in each class.
- **Wireframes** and **Block model**: switch them on or off; tick the wireframes to show and pick their colour; choose an estimate (coloured like the samples when it estimates the same property) or a block model (coloured by domain), its opacity and block edges.

Wireframes and blocks are cut by the vertical plane through A–A′ (the corridor's middle); the samples are those inside the corridor.

Automatic choices

Choice	Default
Statistics	weighted by sample length (untick *Length-weighted* to weight samples equally)
Scale	log when all values are positive and skewed (skewness > 1)
Histogram bins	Freedman–Diaconis rule, 10 to 60 bins
Suggested cap	a break in the probability plot's top decile, else P99 (nothing is capped)
Section line and corridor	along the collars' long axis; 2 × the median collar spacing
Section colours	the samples' 2nd–98th percentile, log when that spans more than 100×
Density	none: the grade–tonnage curve is then in volume (m^3), not tonnes
Swath slices	a multiple of the parent block size, at most 25 slices

Each analysis lists the choices it made and why, and its warnings (for example an estimate that is out of date, or values left out of a log scale).

Tools

The sample analyses take `source` (a drillholes or composites object; default the first drillholes), `property` (default the first numeric column) and `group_by` (a category column such as DOMAIN or LITH: results per group). Every tool takes `open` (opens the EDA page on the result).

Tool	Use it to	Parameters
`eda_overview`	start here: what can be analysed, a summary per column, the suggested order	—
`eda_univariate`	statistics and histogram	`source`, `property`, `group_by`; `weight`: length or none; `log`: auto, on or off; `bins`
`eda_bivariate`	correlation of two columns, and the correlation matrix (Spearman by default)	`source`, `x`, `y`, `group_by`; `log`; `method`: pearson or spearman
`eda_outliers`	top-cut diagnostics and a suggested cap (nothing is capped)	`source`, `property`, `group_by`, `weight`

Tool	Use it to	Parameters
`eda_section`	vertical section along A–A′	`source`, `property`; `start`, `end`: X, Y of A and A′; `width` (m); `wireframes` (default all; an empty list for none); `block_model`: an estimate or block model to cut
`eda_grade_tonnage`	grade–tonnage curve of an estimate	**`estimate`**; `density` (t/m^3); `cutoffs`; `domain`
`eda_swath`	swath plots of an estimate against its samples	**`estimate`**; `axis`: all, x, y or z; `width` (m); `domain`

Geometry: wireframes and block models

Builds closed 3D wireframes from drillholes: **domain solids** from interval codes (lithology, domain) and **grade shells** around a cut-off. Available through the MCP; the result appears in the project (Data page) and the 3D viewer, and exports as OBJ.

Use

Ask Claude, for example:

> Wireframe the SAP lithology from `collar.csv`, `survey.csv` and `intervals.csv`.
> Build a 0.5 g/t AU grade shell and show it with the holes.

That's enough. The grid size, the model extent and the orientation of the body are chosen from your data, and Claude reports what it used and why, for example "*5 m cells: median SAP intercept 20 m / 3*". To change anything, say so:

> Redo it with 2 m cells, isotropic, and carry the shell further between holes.

You can set	Default
Grid cell size (`resolution`)	fine enough for the body (≈ 3 cells across) and the area, max 8 M cells
Model extent (`extent`)	the drilled volume + 5 %
Trend: orientation, anisotropy (`trend`)	follows the shape of the contacts, up to 10:1; or isotropic, or set azimuth, dip, pitch and ratios
Reach between holes (`max_distance`)	2 cells (raise it to carry thick intercepts further)
Grade shell method (`transform`)	indicator: samples above the cut-off are honoured exactly; or interpolate the grade (log or not)
Small parts (`min_part_volume`)	kept (set a minimum volume to drop them)

The wireframe appears in the open 3D viewer within 2 seconds (ask Claude to open the viewer otherwise). To change its parameters later, use Edit on the Data page. Any OBJ from other software can be loaded into the project the same way. Double-click a wireframe (in 3D or in the Data tree) to select it; the Display tab then sets its style (solid, solid + edges, edges only, points), colours and opacity.

Results

In the project folder, `objects/<object-id>/v<version>/`:

File	Content
`wireframe.obj`	the solid, in your coordinates (Leapfrog, Deswik, Micromine, Blender…)
`report.json`	parameters and why, volume, area, parts, fit checks, warnings
`points.csv`	the interpolation points with their fitted values

Check the warnings: a solid **cut by the model extent** (closed flat on that side), samples **on the wrong side** of the surface, or a fit that **hit the centre limit**. In the viewer, the Object tab shows volume, parts and whether the solid is cut by the extent.

Block models

Builds the block-model geometry; grades come from property modelling (see Property modelling). Ask Claude, for example:

> Build a block model inside the SAP wireframe.
> Make it 10 x 10 x 5 m, rotated 30° to follow the strike, and show it with the holes.

Blocks are cut into sub-blocks where a wireframe boundary crosses them, so the model follows the solid closely; each block gets the code of the wireframe that contains its centre (the first listed wins where wireframes overlap). Claude reports the block volume against the wireframe volume.

You can set	Default
Block size (`block_size`)	half the drill spacing in plan, half of that vertically
Sub-blocks per block (`subblocks`)	4 x 4 x 2 inside wireframes, none otherwise
Rotation (`rotation`)	none (azimuth, plunge and dip can be set)
Model extent (`extent`, or `origin` and `n_blocks`)	the wireframes, else the drilled volume + 5 %
Domain codes	1, 2, … in the order the wireframes are given
Blocks outside the wireframes (`keep_outside`)	dropped (can be kept as code 0)

In the viewer, select the block model in the Data tree: the Display tab colours it by domain (show/hide and recolour each one), by block size (parent vs sub-block) or by one of its estimates (a new estimate is applied by itself), with block edges. Blocks are opaque; to see through them, use the ray-traced mode below.

To look inside the blocks, use Display › **Section**: choose an Axis (Elevation, North–South, East–West), then drag the two handles to keep the rows between them. Drag one handle to cut the model on one side, or bring both together for a single bench or cross-section. Block models built before this feature need a rebuild first. Models over 2 million blocks keep their cut faces open.

To see inside a block model, tick **Ray-traced** under Display › Rendering: choose Voxels or Translucent, the opacity, and (with an estimate as Property) a cutoff such as "Show ≥ 0.5". Very large models open ray-traced only. Block models built before this feature need a rebuild first (Data page › Edit › Apply, or ask Claude to update the block model).

On a laptop with two graphics cards, Chrome draws on the integrated one, which makes the ray-traced view several times slower: in Windows Settings › System › Display › Graphics, set Google Chrome to **High performance**, then restart Chrome.

On the Data page, a block model's block size and rotation can be changed; it is rebuilt with its wireframes when they change.

In the project folder, `objects/<object-id>/v<version>/`:

File	Content
`block_model.csv`	one row per block: centre (XC, YC, ZC), size, I J K, sub-block index, DOMAIN

File	Content
`block_model.json`	origin, rotation, block size, counts, domains
`report.json`	parameters and why, block counts, volume per domain vs wireframe, warnings

Check the warnings: a wireframe that is **not closed**, one that **extends beyond the model**, one **closed by its own extent** (the blocks against those faces are extrapolated), or a block volume **far from the wireframe volume** (use smaller blocks).

Tools

`drillholes` is a drillholes or composites object, or the collar, intervals and survey files. `options` takes the settings in the tables above.

Tool	Use it to	Parameters
`geometry_domain_wireframe`	build a solid around intervals with given codes	**`drillholes`**; **`column`**, e.g. LITH; **`inside`**: the codes inside, e.g. SAP; `options` (advanced: `kernel`, `drift`, `smoothing`, `accuracy`, `max_centres`, `spacing`, `narrow_band`); `name`
`geometry_grade_shell`	build a shell around grades above a cut-off	**`drillholes`**; **`column`**, e.g. AU; **`cutoff`**; `options`, also `composite`: the compositing settings of Data processing; `name`
`geometry_block_model`	build a block model, without grades	`wireframes`, in priority order where they overlap; `drillholes`, only to size the model; `options`, also `max_blocks`; `name`

Property modelling

Estimates one property (e.g. a grade) at every block of a block model from drillhole samples. Two methods: **NEXUS** (default), a geostatistical estimator for deposits whose mean grade, spread or orientation change from place to place, and **LVA-DKNN**, a neural network that kriges each block from its nearest samples with a learned covariance. Start it from the Data page or ask Claude; results appear on the Data page and in the 3D viewer.

Set up

Both methods need PyTorch, installed with the `modelling` extra (CUDA build on Windows/Linux, about 2.5 GB). From the `variancelabs-mcp` folder:

```
uv tool install --python 3.11 --editable ".[modelling]" --force
```

Then restart Claude Code (or `/mcp` → `variancelabs` → reconnect). An NVIDIA GPU is used automatically when present; otherwise it runs on the CPU (slower).

Use

On the **Data page**, open a block model and click **+ Add estimate**: choose the samples (composites first), the column (your main commodity is preselected) and the domains (wireframe domains are ticked; each is estimated from its own samples only). The options can stay automatic: each is detected from the data and shown afterwards with its reason. **Create estimate** runs it in the background; the row shows its progress.

Or ask Claude, for example:

> Estimate AU1 in the PEG domain of the 10 × 10 × 5 m block model from the 2 m composites, with the probability above 0.5 g/t.

To change an option, open the estimate, edit it (↺ **Use detected** puts the detected value back) and click **Validate & re-run**. Main options: Transform, Trend levels, Anisotropy, Local scale, Cut-offs; under Advanced: model selection (a cross-validated search, or off for a faster run), smoothness, multipliers, neighbours, discretisation, physical bounds, device.

Results

`objects/<estimate id>/v<n>/estimate.csv`, one row per block in the order of `block_model.csv`:

Column	Meaning
`<value>_est`	expected grade of the block: use it for tonnage and grade
`<value>_p10`, `_p50`, `_p90`, `_sd`	spread at the block centre (P50 is below the mean for skewed grades)
`<value>_rel_unc`	block relative uncertainty (for classification, with the Competent Person)
`<value>_p_gt_<cut-off>`	probability above each cut-off
`<value>_trend`	the large-scale trend alone
`flag`	`trend_only`: no sample close enough, the trend alone; `extrapolated`: more than twice the drill spacing from any sample

Judge the estimate by its validation, on the estimate's panel: held-out holes scored against ordinary kriging (*vs OK*) and inverse distance (*vs IDW*) — positive means better; *80 % cover* near 80 % means honest uncertainty; *Max TRE* at most 1.5 means the swath means of the blocks follow those of the samples (swaths cutting a sharp contact, e.g. the edge of a thin body, can exceed it with a good estimate). *vs OK* compares the held-out distributions, not the block means: on sharp-edged bodies in barren ground the benchmark's NEXUS block means were less accurate than ordinary kriging's, so check such deposits against another estimate. Claude relays the warnings.

In the 3D viewer the block model is coloured by a new estimate as soon as it is ready; select the block model to change its *Property*, the palette or the colour range. `runs/<estimate id>/worker.log` has the details if a run fails.

When the block model, wireframes or composites change, the estimate is marked *out of date*; **Retry** re-runs it. LVA-DKNN estimates are started through Claude (`method: lva_dknn`) and are read-only on the Data page.

Tools

Tool	Use it to	Parameters
`modelling_start_run`	start an estimate in the background	**`block_model`**; **`samples`**: composites, drillholes, or a CSV with X, Y, Z and a hole id; **`value`**: the column to estimate; `method`: nexus (default) or lva_dknn; `domains`: the domain codes to estimate, each from its own samples; `columns`: the X, Y, Z, hole and length columns of a CSV; `options`; `name`
`modelling_run_status`	follow a run; its validation and files when done	**`object_id`**: the estimate
`modelling_cancel_run`	stop a running estimate	**`object_id`**

NEXUS `options`: `transform`, `trend_levels`, `anisotropy`, `scale` (local scale), `cutoffs`; advanced: `model_selection`, `smoothness`, `window_multiplier`, `range_multiplier`, `neighbours`, `per_hole`, `search_radius`, `discretisation`, `physical_min`, `physical_max`, `seed`, `device`. LVA-DKNN `options` (`epochs`, `learning_rate`, `device`…): ask Claude for the full list.