Isoliner - grids and isolines

Contents

Introduction

Isoliner is a Processing provider for interpolating point data, building isolines and working with terrain. The kriging core is the KB2D algorithm from GSLIB. The tools are split into five groups: **Grid and isolines** - the main processing flow from declustering to isolines, **Topography** - terrain from open data and hydrological analysis, **Additional analysis tools** - specialised computations from indicator kriging to variable-support density, **Cross-sections** - building geological sections, **Fractal analysis** - dimensions of surfaces, masks and lines.

2D Kriging (points → raster) - ordinary or simple kriging over a point layer.

Isolines from raster - isolines (lines) and contour polygons (bands between isolines) whose boundaries coincide with the lines.

Variogram map (anisotropy) - the γ(h_x, h_y) surface with an azimuth and anisotropy estimate, to account for directionality in kriging.

Variogram cross-validation - leave-one-out checking to validate and tune kriging parameters by error rather than by eye.

Create sample wells (demo) - generates a training point layer with a spatial structure (roof, thickness, component grade) for learning and testing without real data.

The **Additional analysis tools** group holds specialised computations, for example:

Categorical indicator kriging - a class-probability map from a categorical field (mineral type, lithotype): an indicator is built per class and kriged separately, giving a probability raster, a zone map and a confidence raster.

Hydraulic gradient and flow direction - from a head raster it builds the hydraulic gradient magnitude, the flow-direction azimuth (down-gradient) and a point layer of flow vectors styled as arrows right away. Hydrogeology without permeability.

Suitable for roof elevations, thicknesses, geomechanical properties, chemistry and any numeric well attribute.

A few terms used below. A variogram describes how much more strongly values differ as the distance between points grows. The sill is the level it reaches (close to the data variance). The nugget (from the “nugget effect”) is the jump of the variogram at zero - the scatter at arbitrarily small distances caused by measurement noise and microvariability.

Installation and location

The main way is from the official QGIS repository. Open Plugins → Manage and Install Plugins → the **All** tab, type “Isoliner” in the search, select the plugin and click **Install**. When installed from the repository, QGIS itself reports new versions and updates the plugin at the press of a button.

Raster band choice in all the tools is a drop-down with band names: a multiband bed grid shows roof, bottom and the parameter layer names in the lists.

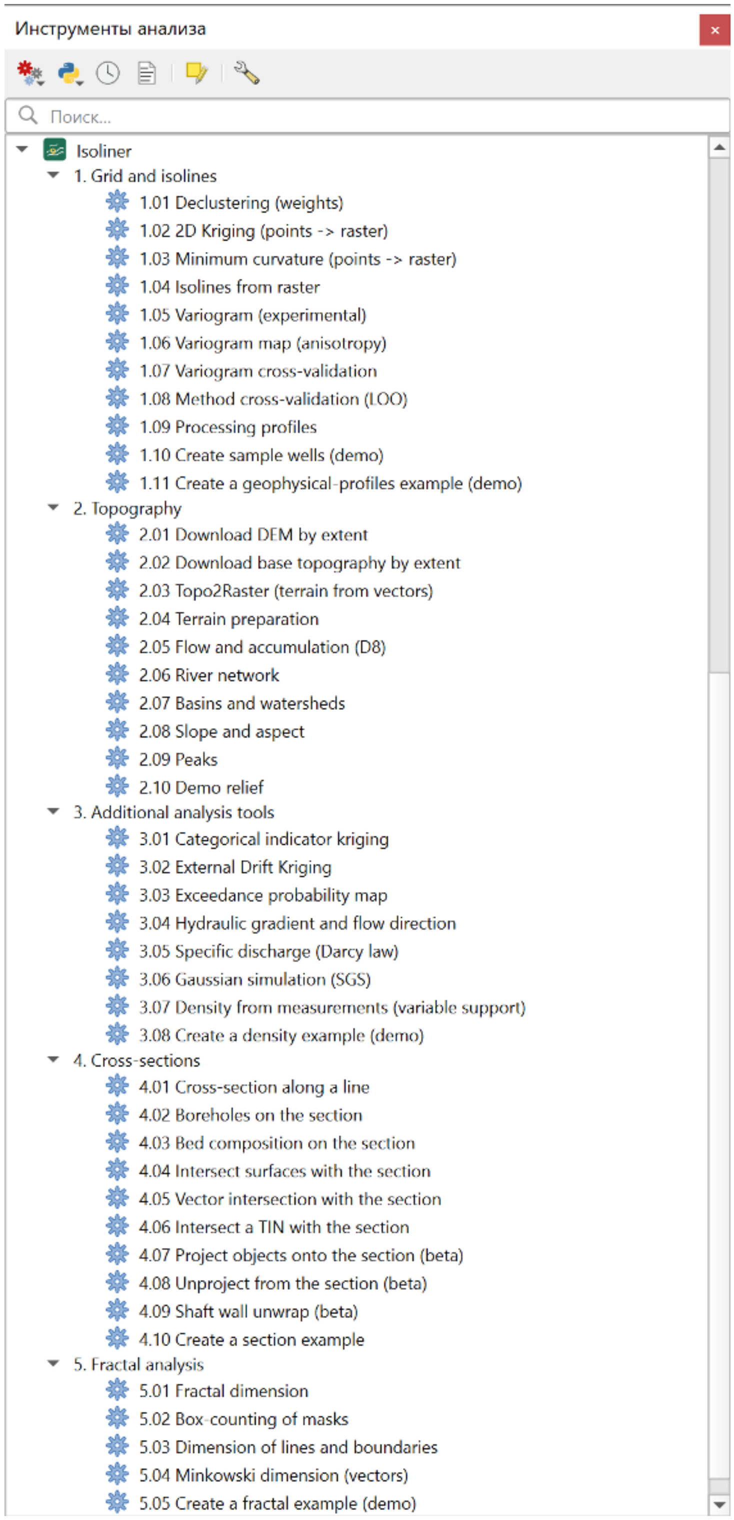

Figure 1: The Isoliner provider in the Processing toolbox: five tool groups.

The alternative way is from a ZIP file. Plugins → Manage and Install Plugins → Install from ZIP. This is handy for offline installation and pre-release builds.

After installation the tools appear in the **Processing** panel: provider **Isoliner**, groups **Grid and isolines**, **Topography**, **Additional analysis tools**, **Cross-sections** and **Fractal analysis**. Requirements: QGIS 3.16+. There are no external dependencies - only NumPy, GDAL and the built-in Processing algorithms shipped with QGIS are used.

Updating

When installed from the repository, QGIS shows a notification about a new version - an icon in the status bar and a list on the **Upgradeable** tab of the plugin manager. Updating is a single click. When installed from ZIP, the new version is installed the same way, over the old one.

The plugin reloads cleanly on the fly, no QGIS restart is required. For a quick code reload during development the Plugin Reloader plugin is convenient ("Reload a plugin..." button). Pick Isoliner - the provider and all tools re-register immediately.

Opening the help

Each tool's dialog has a **Help** button that opens this manual (the PDF bundled with the plugin; on an English interface the English manual opens). The right-hand panel of the dialog additionally shows a short hint for the tool. The manual and the version details are also available without opening a tool: the **Plugins** menu holds an **Isoliner** submenu with **About** (version, links, changelog) and **Manual (PDF)**.

General workflow

A typical scenario has two steps:

2D Kriging: from a point layer and a numeric Z field a raster is built (a regular grid of values).

Isolines from raster: from the resulting raster, isolines and, if needed, filled contour polygons are built.

The steps are independent: **Isolines from raster** works with any raster, not only with a kriging result.

The tools are grouped into three Processing groups. The "Grid and isolines" group is the main processing flow, from kriging to isolines. The "Additional analysis tools" group holds the specialised computations, categorical indicator kriging, external drift kriging, the hydraulic gradient with flow direction, the exceedance probability map, and the Darcy specific discharge. The "Cross-sections" group builds geological sections along a line and prepares demo data for them.

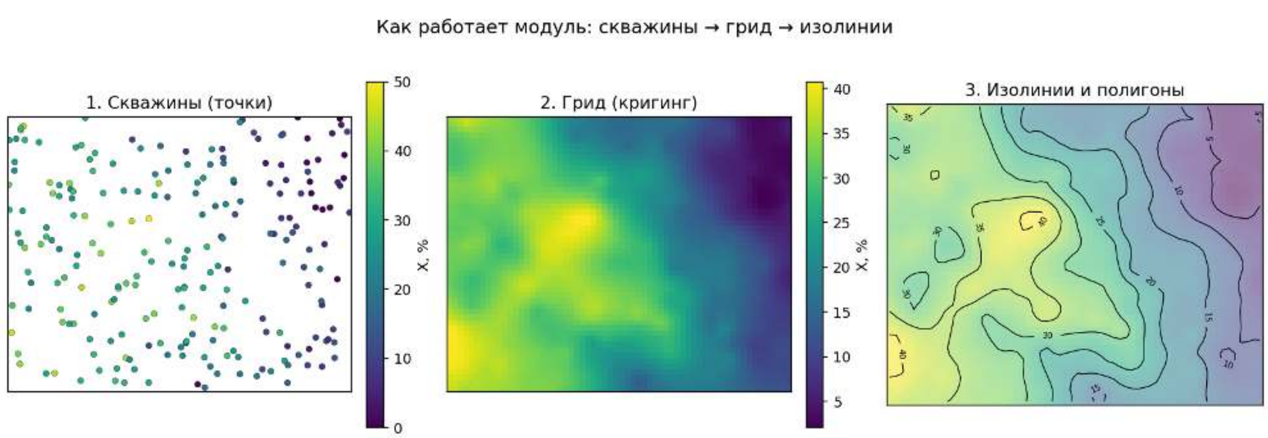

Figure 2: The whole process on a generated example: wells with measurements (left) are turned into a continuous grid by kriging (centre), from which isolines and contour polygons are built (right).

1.01 Declustering (weights)

The tool prepares data before interpolation. When samples are clustered unevenly, some blocks drilled denser than others, the naive global statistics shift toward the over-sampled areas. If rich zones were drilled denser, the mean and histogram are overstated, and that directly affects reserve calculation. Cell declustering (a port of GSLIB **declus**) gives each sample a weight inversely proportional to the local density: less in a cluster, more on its own. A representative declustered mean is computed from the weighted data.

A grid of cells is laid over the area, a sample weight is proportional to one divided by the number of samples in its cell, then the weights are normalized. The cell size is chosen automatically: a sweep over sizes, picking the minimum declustered mean (when clusters fall in rich zones) or the maximum. The size can be set manually. On a regular grid declustering changes nothing, all weights are equal.

Parameter	Purpose	Default
Points with values	Samples.	-
Value field (Z)	Numeric field.	-
Cell size	Auto (sweep) or manual.	Auto
Cell size for manual mode	Cell side in manual mode.	0
Sweep objective (Adv.)	Minimum or maximum mean.	Minimum
Number of sizes in the sweep (Adv.)	How many cells to try.	24
Cell Y/X ratio (Adv.)	Cell anisotropy.	1
Grid-origin offsets (Adv.)	Averaging over grid shifts.	4
Points with weights	Point layer with a **wt** field.	-
HTML report	Summary, histogram, mean curve.	default

Outputs: a point layer with a **wt** field and an HTML report with a summary (naive vs declustered mean), a raw-vs-weighted histogram and a mean-vs-cell-size curve. The declustered mean from the log and report goes into **1.02**, the **Mean of simple kriging** field, and the **wt** field feeds **3.06 Gaussian simulation** for a weighted normal-score transform. Outlier samples are cut separately, by percentile capping in the kriging and cross-validation tools themselves.

1.02 2D Kriging (points → raster)

Ordinary (OK) or simple (SK) kriging over a point layer. Coincident points (the same XY) are averaged over Z. At grid nodes the values of the source points are reproduced exactly (with a zero nugget).

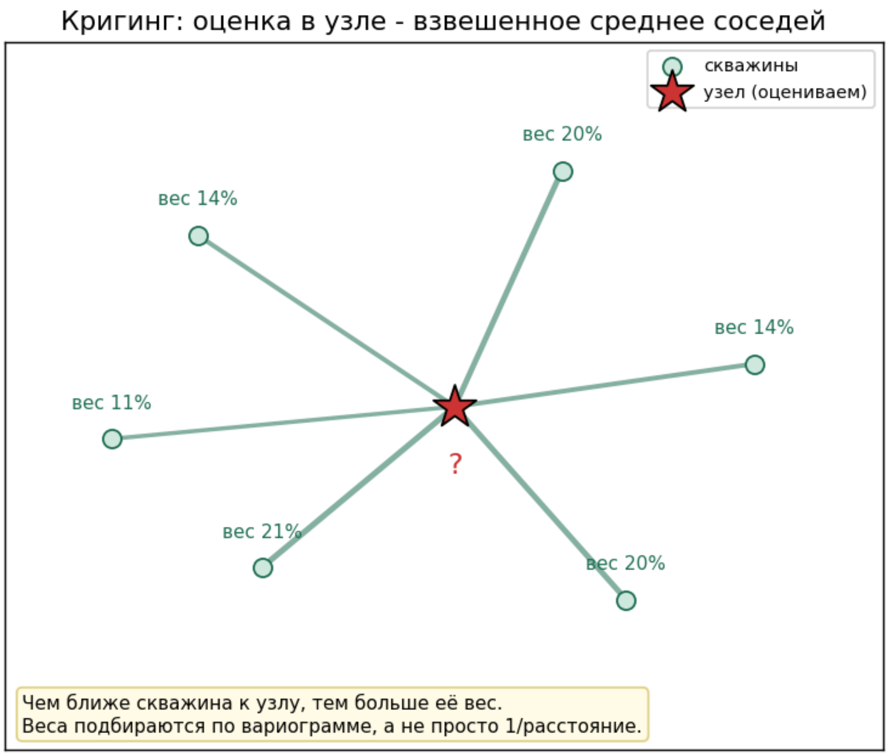

Figure 3: Kriging estimates the node value as a weighted mean of the nearest wells: the closer the well, the larger its weight. Weights come from the variogram.

Main parameters:

Parameter	What it sets	Default / advice
Point layer	Source points (wells) for interpolation.	-
Selected features only	Compute only over the layer's selected points.	off
Value field (Z)	The numeric attribute that is interpolated: roof elevation, thickness, geomechanical property, chemistry, etc.	remembered between runs
Value transform	ln for log-normal quantities (K, T, grades with a long tail): ln(Z) is kriged and the estimate is returned via exp.	none
Kriging type	Ordinary (OK) - estimates the mean locally itself. Simple (SK) - uses the specified **Mean**.	OK
Search radius	Radius of the search window for neighbouring points around a node. 0 = whole sample.	0 (whole sample)
Min. number of points	If the window has fewer points, the node stays empty (nodata).	1

Parameter	What it sets	Default / advice
Max. number of points	How many nearest points enter the kriging system.	24
Cell size	Grid step. 0 = auto = min(extent)/50.	smaller = smoother, but slower
Raster extent	The computation rectangle. By layer by default.	by layer
Clip to well hull	The raster is clipped to the convex hull of all points - removes extrapolation in empty corners.	recommended on
Hull buffer	Expand the hull outward by N map units.	0
Clip mask	Your own polygon instead of the hull (takes priority) - handy for concave areas.	-
Load processing profile	Substitutes a saved profile (nugget, variogram structure, outliers) over the dialog fields. The list refreshes when the window opens.	(none)

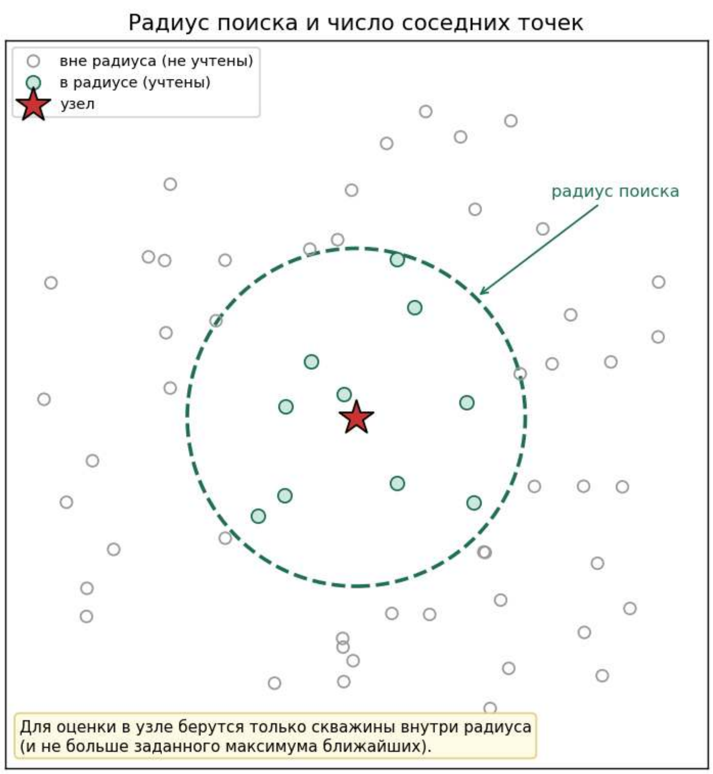

Figure 4: For each node only wells within the search radius are taken, and no more than the set number of nearest ones. Points beyond the radius do not take part.

Automatic values

Cell size = min(extent width, height) / 50.

Variogram correlation range = max(extent width, height) / 3.

Search radius (when 0) = the extent diagonal, i.e. the whole sample is taken.

Clip to well hull

Kriging computes the whole rectangular extent, so outside the data area the values are extrapolation and produce artefacts (long “fan” isolines in empty corners). The **Clip to well hull** option builds the convex hull of all points (with an optional buffer) and clips the raster to it. The extrapolation disappears. If the actual boundary of the area is concave, set your own polygon in the **Clip mask** - it takes priority over the hull.

Variogram and nugget

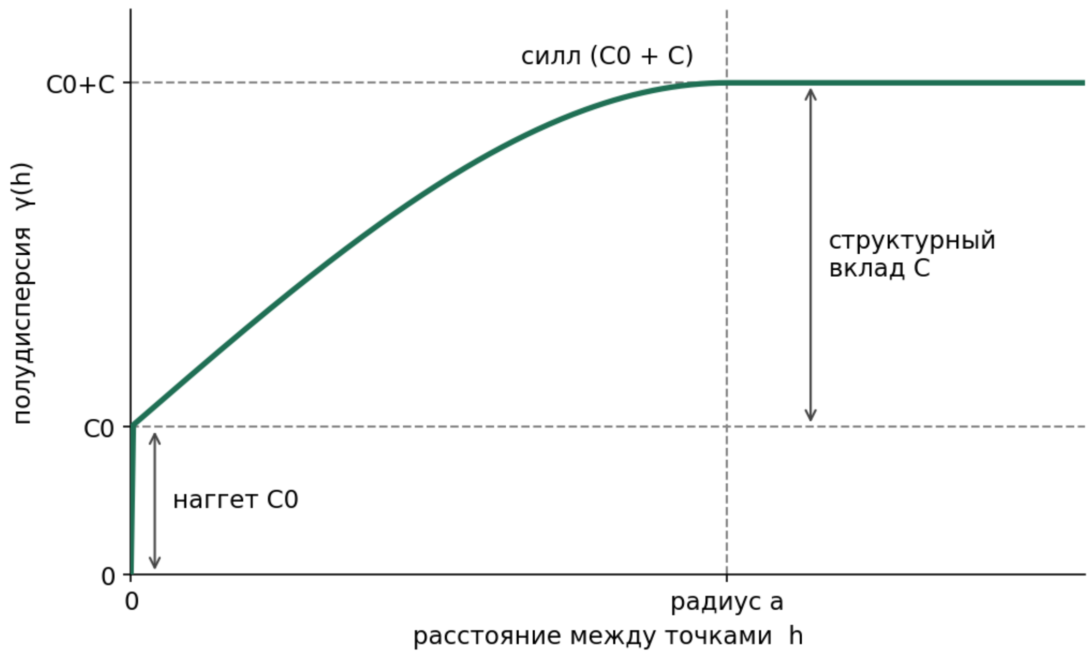

Figure 5: The variogram scheme: nugget C0, structural contribution C, sill (C0+C) and correlation range a.

Kriging relies on a variogram model - it describes how strongly the Z values in two points differ depending on the distance between them. By this model a weight is assigned to each neighbouring well. The model is set in the **Advanced Parameters** section.

Variogram model: nugget C0, sill (C0 + C) and correlation range a.

Nugget C0

The nugget is the value the variogram curve tends to as the distance tends to zero. In theory the discrepancy at zero distance should be zero (a point compared with itself), but in practice a step remains. It reflects the fact that the data at very small distances still do not match: measurement and digitizing error, microvariability at a scale finer than the well spacing, the discrepancy of duplicates at one point.

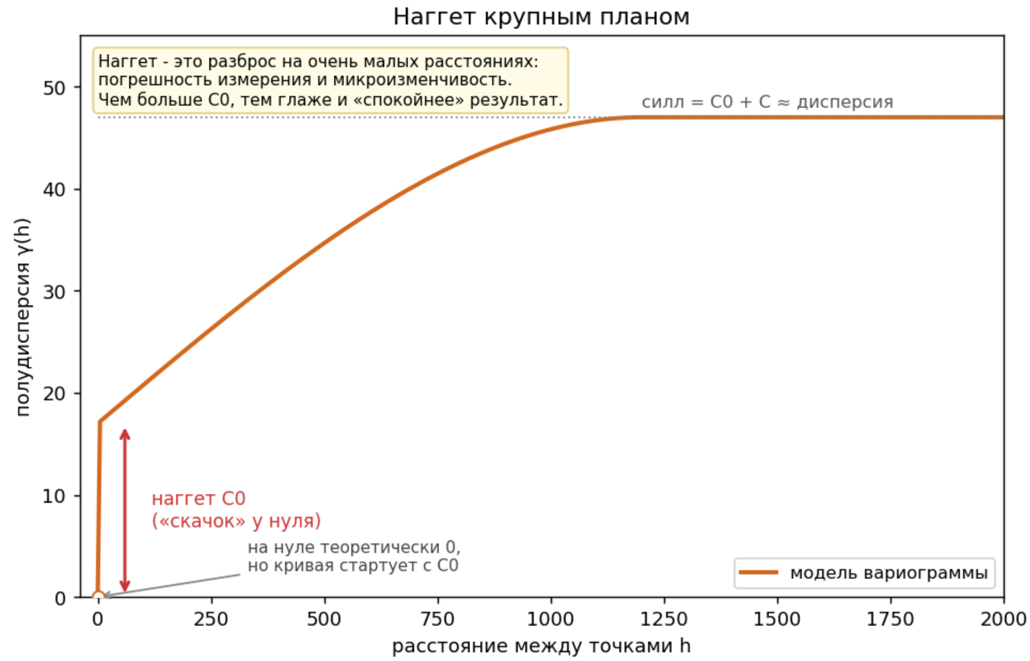

Figure 6: The nugget close up: the model starts at zero not from 0 but from a “jump” C0. This is the scatter at very short distances (measurement error, microvariability). The plateau C0+C ≈ the data variance.

How the nugget affects the result:

C0 = 0 (default) - kriging is an exact interpolator: the surface must pass exactly through every well. An isolated well with a Z outlier turns into a cone (a “bull’s eye”).

C0 > 0 - kriging stops reproducing the value at the measurement point exactly and becomes a smoother: near a well the estimate is pulled toward the local mean. The larger the nugget fraction C0 / (C0 + C), the stronger the smoothing.

C0 = the whole sill (pure nugget) - the spatial link is lost, the surface degenerates into a plain mean. This is too much.

Important - units. Nugget C0 and sill are set in **absolute units of the data variance** (squared units of Z), not in 0-1 fractions. The default “1” for the sill is a placeholder that almost always needs changing: set the total sill (C0 + the structure contributions C) **close to the data variance**. The level of smoothing is determined not by the absolute nugget value but by its **fraction of the sill** C0 / (C0 + C). A practical order: take the total sill ≈ the variance, then the nugget = 0.2-0.4 of it (i.e. 0.2-0.4 × the variance - an absolute number, not 0.2-0.4 as such). The smaller the nugget, the more detail, but also more local peaks. The larger it is, the smoother the surface, but real structure may be smoothed away. The tool prints the data variance to the Log at start - that is your reference for choosing the sill.

Structures, range and anisotropy

The sill (plateau) is the level the variogram reaches. It is the sum of the nugget C0 and the structure contribution C. A structure is set by a model (spherical, exponential, Gaussian or power), a contribution C, a range a, an azimuth and an anisotropy.

Sill: meaning and order of magnitude. The sill is the upper limit of the differences between points: how much, on average, distant wells differ. It is practically equal to the ordinary data variance. An example for KCl: mean ≈ 25 %, variance ≈ 47.6 ($\%^2$), i.e. $\sigma \approx 6.9$ %. So the total sill is set ≈ 47.6. If the nugget C0 ≈ 17 (about 0.35 of the sill), then the structural contribution of the first structure C ≈ 47.6 − 17 ≈ 30. The absolute scale does not affect the grid itself - only the C0 : C ratio matters for the estimates. But it is needed so that the standard-error map and the MSDR in cross-validation are at the real scale (total sill ≈ variance → MSDR ≈ 1). So: do not leave the sill at the default 1, raise it to the data variance.

Choosing a model. The spherical and exponential models suit most tasks. The power model has no sill or range in the usual sense: it is used when variability grows with distance and does not reach a plateau (non-stationary increments), so the contribution and range fields for it are conditional. Use the Gaussian model with caution: at a zero or very small nugget it gives a numerically unstable system and artefacts (oscillations, negative weights). Therefore, when the Gaussian model is chosen, the tool enforces a small minimum nugget; set one yourself where possible.

Data type and mode. Different data need different settings. Smooth structural surfaces (roof and floor elevations, thicknesses) are better modelled with a long range or a power model under a wide (global, 0) search radius - then the surface comes out smooth. A short radius with local search on such data gives “bull’s eyes” and discontinuities in the estimate when the set of neighbouring wells changes. For grades and chemistry (geomechanical properties, gas hazard) kriging works in its own right: here a correct nugget matters and, for a strongly skewed distribution, a data transform (see below on outliers).

The correlation range a is the distance at which the variogram reaches the plateau. Beyond it points practically do not influence each other. At 0 the automatic value max(extent)/3 is taken.

Anisotropy is set by the major-axis azimuth and the ratio of ranges (minor/major). A value of 1 is isotropic (the influence is the same in every direction). A value below 1 shortens the correlation across the major axis - useful for elongated geological structures.

Figure 7: Anisotropy: if the body is elongated, correlation runs farther along it than across. The search ellipse is set by the major-axis azimuth and the axis ratio (minor/major).

Parameter	What it sets	Default / advice
Mean for simple kriging	Used only with the SK type.	0
Nugget C0	The “noise”/jump of the variogram at zero. Suppresses local peaks. In absolute variance units.	0. For smoothing 0.2-0.4 of the sill
Structure i · model	Variogram shape: spherical, exponential, Gaussian, power.	spherical
Structure i · sill/contribution C	The structure’s contribution to the sill (abs. variance units). The sum C0+C ≈ the data variance.	str. 1 = 1 (replace with ≈ the variance)
Structure i · correlation range a	The distance to reach the plateau. 0 = auto = max(extent)/3.	0 (auto)
Structure i · azimuth, °	The direction of the anisotropy major axis.	0
Structure i · anisotropy (minor/major)	The ratio of ranges across/along the axis. 1 = isotropic.	1

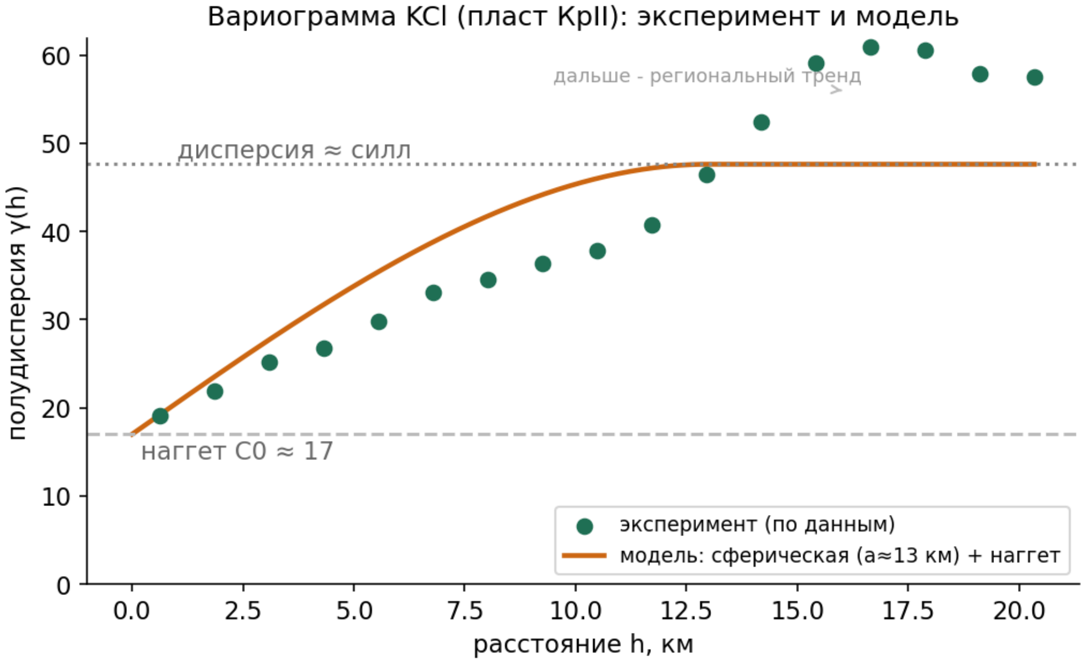

Figure 8: Experimental KCl variogram for the KrII seam and the fitted model: nugget C0≈17, the sill matched the data variance, range ≈13 km. Points beyond the sill are a regional trend.

This is how the scheme looks on real data. A variogram is built from the wells: for pairs of points the semivariance is computed and averaged over distances - the result is a cloud (green points) under which a model (the curve) is fitted. From it the kriging parameters are set: the height of the “jump” at zero is the nugget C0, the plateau is the sill (usually close to the data variance), the distance to the plateau is the range a. If at large distances the points rise above the sill, as here, it is a regional trend (non-stationarity). It is either accounted for separately or the search radius is limited.

Outlier removal

Outliers are anomalously high (or erroneous) values that distort the estimate: a few grade “bonanzas” can pull the whole grade map onto themselves, while clear errors (e.g. a negative thickness) spoil the surface. The **2D Kriging** tool lets you bound such samples right during the computation, without editing the source data. The parameters are in the **Advanced** section.

Removal and capping are a crude practical tool against clear errors. For grades and chemistry be careful: extreme values are often not noise but signal (e.g. contamination spots), and blindly clipping the distribution tails is not worth it. For strongly skewed data it is more correct not to clip the samples but to transform them to something close to normal (logarithm, Box-Cox) or to use indicator kriging - that is beyond removal, but that is exactly how heavy tails are handled in the geostatistics of ores and contamination.

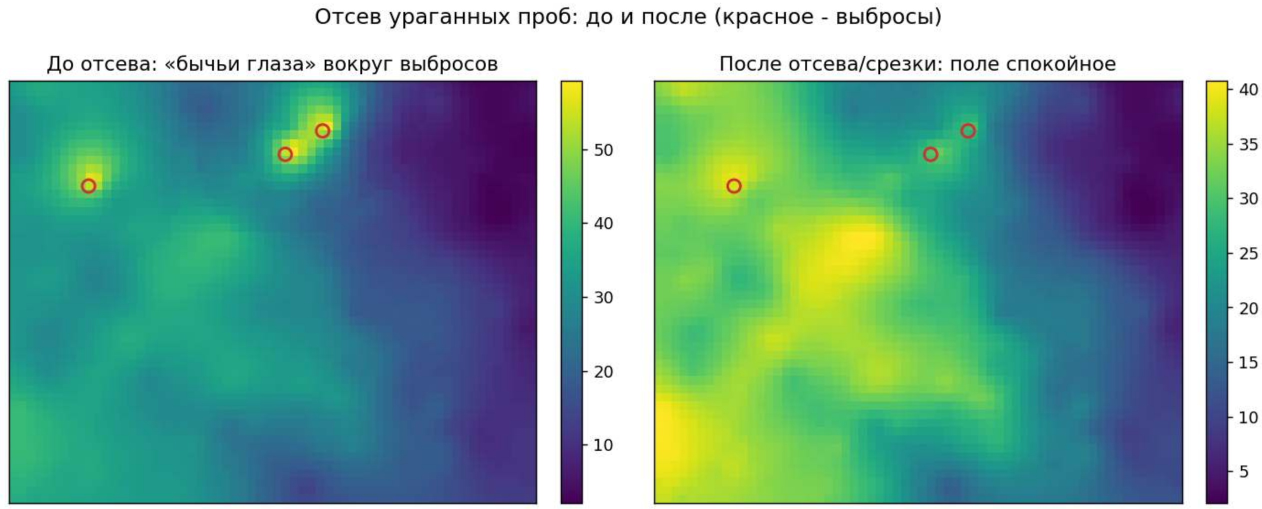

Figure 9: Outlier removal by example: on the left three outliers give “bull’s eyes” (hot spots), on the right after capping to the upper bound the field is calm.

Two modes. **Remove** - samples outside the allowed range are discarded (for clearly broken records). **Cap (capping)** - values outside the range are clamped to the bound, while the point itself stays in the computation. Capping is the classic technique for grade outliers: the point’s position is not lost, but its influence is limited. The mode is switched by the **Cap to bound (capping) instead of removing** checkbox.

Absolute bounds. The **Lower value bound** and **Upper value bound** set thresholds in Z units directly. An empty field means the bound is not set. They take priority over the percentile. Example: for thickness set the lower bound to 0 - negative values go away, and the upper, say, to 30 - a clear outlier at 122 m goes away.

Percentile bounds. The p-th percentile is the value below which p% of all samples lie. For example, the 5th percentile is the threshold below which only the 5% smallest values lie. The 95th is the threshold above which the 5% largest lie. The **Outliers: clip percentile, %** parameter sets the number p, and the bounds are taken symmetrically: from the p-th to the (100−p)-th percentile. So p = 2 means “treat as outliers the 2% lowest and 2% highest samples”: everything below the 2nd and above the 98th percentile is either removed or capped. The larger p, the more aggressive the clipping. P = 0 disables the percentile mode. The convenience is that you do not need to know the absolute thresholds - they are computed from the data itself and suit any distribution and scale.

Two-sidedness - important for chemistry. The percentile mode cuts both tails - upper and lower. For grades this is dangerous: KCl = 0 in replacement zones is real geology, and clipping the lower tail would wrongly raise the “empty” areas. So for grade clip only from above: leave the **Lower value bound** empty and set the **Upper** as an absolute (or use the percentile, knowing the bottom will be affected too). For elevations and thicknesses two-sided clipping is usually appropriate.

Order and Log. The filter is applied before averaging coincident points. The tool’s Log reports how many samples were removed or capped and within which bounds - handy for checking.

Kriging standard error

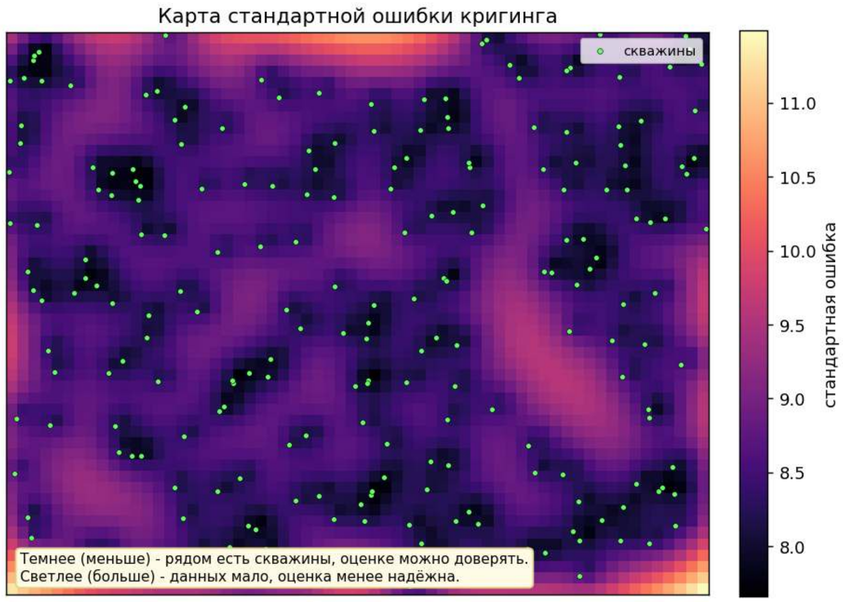

Figure 10: The standard-error map: dark near wells (green points) - the estimate is trustworthy, light in empty corners - few data.

Besides the estimate itself, kriging gives at every node the error variance - a measure of uncertainty. Its square root, the standard error, is output as an optional second raster (the **Kriging standard error** parameter of the **2D Kriging** tool). The units are the same as the interpolated quantity Z.

A key property: the standard error depends on the geometry of the well layout and the variogram model, but not on the Z values themselves. So it is a map of the observation network's reliability, not of the data scatter. At a well point (with a zero nugget) the error equals zero - there the value is known exactly. As one moves away from wells it grows, and in areas without data it reaches a maximum (roughly the square root of the sill).

How to read it. Dark (small) values - the estimate is trustworthy: enough wells nearby. Light (large) - the estimate rests on distant points, effectively extrapolation. These are the first candidates for infill drilling. It is more convenient to compare relatively (where it is larger or smaller), because the absolute value depends on the variogram scale (the sill S1_SILL).

Important. This is a model estimate: it is as valid as the variogram you set (nugget, range, anisotropy). At a nugget above zero the error at wells is not zero - the nugget sets a lower "floor" of uncertainty. The standard error is not a strict confidence interval, but as a relative uncertainty map it is very useful.

Styling. Give the layer graduated symbology by value (e.g. from dark to red) - and it is immediately clear where the map is reliable and where not.

Trend removal (regression kriging)

Ordinary kriging estimates the mean locally, within the search window, so it follows a smoothly varying mean on its own. The difficulty appears when the field has a pronounced regional component, such as a general dip of the seam across the area. Then the experimental variogram of the raw value gets inflated: the range is overstated, there is no sill, the shape resembles a power model, and a stable model is hard to fit.

The **Remove polynomial trend** checkbox removes the regional component by least squares before kriging. The residuals are then kriged and the trend is added back to the estimate. The residual variogram returns to its normal shape: it reaches a sill with a nugget, and the range reflects the true scale of correlation rather than the span of the trend. The **Trend degree** field selects a plane or a quadratic surface.

The **Value transform** list adds a logarithm for quantities spanning orders of magnitude, such as hydraulic conductivity or transmissivity. With **ln** selected, the natural logarithm of the value is kriged and the estimate is returned through the exponential. This is the median, geometric estimate, correct for log-normal fields. The standard error is converted back to the original units. The logarithm removes the need to build an ln field by hand in the calculator and applies to positive values only. Set the variogram and nugget in ln units when the logarithm is on.

Figure 11: Trend removal on real data: a seam surface with a pronounced regional dip. The trend is removed by a polynomial, the residuals are kriged, and the trend is added back to the estimate.

Trend removal helps where the dip is uniform, within a single mine district or a local area. It is not meant for the whole deposit at once. Neighbouring blocks are offset in height, a single polynomial describes them poorly, and local kriging already holds the varying mean, so removing a global trend there tends to hurt. The variogram shows whether it is warranted. If the raw value has a range comparable to the size of the area and no clear sill, a trend is present and worth removing. If the raw variogram already reaches a sill with a small nugget, there is nothing to remove.

Degree 1 is usually enough. Degree 2 captures curvature but can absorb part of the real structure into the trend, so after using it look at the residual variogram. If the sill and nugget become less defined, go back to degree 1. The same checkbox is in the **Variogram cross-validation** tool, where the trend is refit at each validation step, so the gain or loss from removing the trend is visible directly in the RMSE.

After trend removal, fit the variogram on the residuals. In this mode the standard-error raster is the kriging error of the residuals, the trend is treated as deterministic and adds no error of its own.

Block kriging

Ordinary kriging estimates the value at a point, at the centre of a grid cell. In mining, however, what is usually needed is not a point value but the average over an area, over a mining block, a panel or a reserve-estimation cell. The grade of a useful component in a block is the average over its area, and estimating it as the value at a single point is not quite correct. This is what the **Block kriging** checkbox is for.

When enabled, the mode estimates the average over the grid cell rather than the value at its centre. The cell is conceptually split into an N×N grid of points, the count set by the **Block**

discretization, N×N per cell field, and the covariances in the kriging system are averaged over those points. The system is thus solved not for one point but for the whole block at once. This is the classical GSLIB block-kriging scheme.

The block estimate has two consequences, both useful for reserve estimation. The surface comes out smoother than the point one, because averaging over the block damps small fluctuations. And the kriging standard error comes out lower than the point one, because an area average is estimated more reliably than a value at a single point. There is one price. Block kriging does not reproduce the samples at nodes exactly. The average over a block, even one centred on a borehole, does not equal the value at the borehole itself, and that is as it should be.

A 4×4 discretization is almost always enough. A larger N takes longer to compute while accuracy grows only slightly. At N equal to one block kriging degenerates into point kriging, so the minimum value of the field is two, and the mode itself is off by default.

Block kriging combines with trend removal. The residuals are kriged over the block and the trend is added back to the estimate. It also combines with grid smoothing, but block averaging alone is usually enough and extra smoothing is not needed.

1.03 Minimum curvature (points -> raster)

The tool builds a grid by minimum curvature. The surface behaves like a thin elastic plate passing through the data with the least bending, that is a solution of the biharmonic equation. The method is not exact: the data are honored approximately, but the surface comes out as smooth as possible, which is why it is traditionally used for maps of geophysical fields and any smooth quantity. It is a deterministic alternative to kriging without variogram fitting. Kriging, unlike it, gives an estimate with a standard-error map.

Tension mixes in a membrane term: 0 is pure minimum curvature, 1 is a taut membrane with fewer overshoots between samples. Boundary tension is set separately and helps remove edge overshoots. The solution is iterative, by successive over-relaxation (SOR) with a nine-colour sweep of the grid: nodes of one colour do not fall into each other's stencil, so they are updated at once and stably. The grid is recomputed until the largest node change drops below the **residual threshold** or the iterations run out.

Free nodes start from the nearest data value, so convergence is fast on dense data. On very sparse data more iterations are needed: raise their limit or the residual threshold. Faults and breaklines are not supported in this version; they are planned for the future.

Parameter	Purpose	Default
Point layer	Samples with a value.	-
Value field (Z)	Numeric field to interpolate.	-
Extent	Result rectangle.	from layer
Cell size	0 = auto, min(extent)/50.	0
Tension	0 = minimum curvature, 1 = membrane.	0
Residual threshold	0 = auto, 0.01 percent of the data range.	0
Maximum iterations	Cap on the number of SOR passes.	100000
Boundary tension (Adv.)	Tension at the grid edge.	0
Relaxation factor (Adv.)	SOR acceleration, sensibly 1.5 - 1.9.	1.85
Anisotropy (Adv.)	Y/X axis ratio in the membrane term.	1
Grid (minimum curvature)	Output raster.	-

The output is an ordinary grid ready for **1.04 Isolines from raster**. The log prints the grid size, the number of data nodes, the number of iterations and the final residual. If the iteration cap is reached while the residual is still above the threshold, the tool warns about it.

1.04 Isolines from raster

Builds isolines (lines) and, by default, contour polygons. Levels are set by a uniform step or by an explicit list. Parameters:

Parameter	What it sets	Default / advice
Raster	The input raster (e.g. a kriging result).	-
Isoline step	A uniform step over Z. 0 = set **Explicit levels**.	-
Base level (offset)	Anchors the level grid (levels are multiples of the step from the offset).	0
Explicit levels	A space-separated list of levels. Takes priority over the step. The decimal separator is a comma or a dot.	-
Index isoline every N-th	Every N-th line is flagged is_index = 1 (for thickening). 0 = off.	5
Min. line length	Drop lines shorter than the threshold (map units). 0 = no filter.	-
Bicubic isoline smoothing	Densifies the grid (×2…×4) by bicubic interpolation before contouring - the main isoline-smoothing method, removes “octagons” from a coarse grid. Works for both lines and polygons. off = no densification.	off (×4 on a coarse grid)
Line rounding (Chaikin), iterations	An extra light line rounding (Chaikin). Weaker than bicubic smoothing; usually not needed if it is on. 0 = off.	2
Value field name	The name of the level attribute in the output lines.	ELEV
Band (adv.)	The band number of the input raster.	1
Isolines / Contour polygons	The output layers. Polygons are built by default into a temporary layer.	-

Output fields:

Layer	Field	Type	Holds
Isolines	ELEV	number	The level value of the line (name set by the **Value field name**).
Isolines	is_index	integer	1 on index isolines (every Nth), otherwise 0 - for thickening.
Contour polygons	ELEV_MIN	number	Lower level of the band.
Contour polygons	ELEV_MAX	number	Upper level of the band.

Isoline smoothing

The main way to smooth isolines in this tool is **bicubic smoothing**: before contouring, the grid is densified by bicubic interpolation (×2…×4), and the contours are built on the finer grid. On a coarse grid isolines otherwise look like “octagons” (vertices are placed at cell edges) - densification removes this angularity topologically cleanly. It is implemented in pure NumPy, with no external dependencies; nodata boundaries and internal data “windows” are preserved. Densification affects both lines and contour polygons - the band boundaries still coincide with the isolines. The cost is more cells (×4 = 16 times more), so on a very large grid start with ×2.

In addition there is a light line rounding by the **Chaikin** algorithm (number of iterations). It is weaker than the bicubic one and usually not needed if densification is on; it makes sense as a fast alternative on a coarse grid when you do not want to densify.

Smoothing of the field itself (Gaussian, over the raster) is a separate operation done not here but in the **2D Kriging** tool: there it goes over the grid before contouring and removes not angularity but field bumpiness (the “bull’s eyes” around wells). Bicubic smoothing and Gaussian field smoothing complement each other: the first cures grid angularity, the second cures data bumpiness. The contoured kriging raster is not changed in the process - only a temporary copy is smoothed.

Contour polygons (bands)

Contour polygons are filled bands between neighbouring isolines. They are built not by classifying raster “steps” but by polygonizing the smoothed isolines themselves together with the outline of the raster’s valid area: line ends are snapped to the outline, the network is noded and polygonized. The level range of each band is determined by sampling the raster at a representative point of the polygon.

Thanks to this the polygon boundaries coincide with the isolines, and the coverage is continuous (no holes). The polygons carry the ELEV_MIN and ELEV_MAX fields. By default they are built into a temporary layer. To not build them, clear the **Contour polygons** field.

Layer styling

Lines: set rule-based symbology on is_index - give the index isolines (is_index = 1) a larger width. Label by the level field (ELEV).

Polygons are created with a single symbol. For range fills set graduated symbology by ELEV_MIN (or ELEV_MAX).

The isoline layer is automatically placed above the polygon layer so the lines show over the fill.

1.05 Variogram (experimental)

The tool builds an experimental semivariogram from points, fits a model to it if needed, and produces an HTML report with a chart. It does not compute a grid and is not part of the kriging computation chain directly. Its job is diagnostic: to show the structure of the data's spatial variability and to help set the variogram parameters deliberately, by the look of the cloud rather than by eye.

Why the preview is needed

Kriging relies on a variogram model: nugget, sill and range. The interpolation weights and the standard-error map depend on them. It is tempting to hand the fitting of these numbers to automation and not think about them. On a clustered drilling grid this is dangerous. Clusters of close wells give a huge number of pairs at short distances and press down the near part of the variogram, so an auto-fit on such a cloud easily yields a confidently wrong nugget. The preview removes this problem: the geologist sees the pair cloud itself, understands where the data are dense and where sparse, and fits the model knowing what lies beneath it.

That is why model fitting in the tool is given as a recommendation, not a finished result. The numbers it suggests should be checked against the look of the chart and only then carried into kriging.

A short theory

The semivariogram describes how statistically related the parameter values are in two points depending on the distance between them. For a pair of points separated by a distance h, half the squared difference of their values is taken (the semivariance of the increment). These quantities are averaged over distance intervals (lags), giving the curve $\gamma(h)$. It is a measure not of the "average difference" of values but of the statistical reliability of predicting a value from a neighbour: the smaller γ, the closer the link.

A typical curve has three characteristics. The nugget C0 is the value γ tends to as the distance tends to zero. It reflects variability at a scale finer than the network step, plus measurement error. The sill is the level the curve reaches at large distances. The full sill equals the sum of the nugget and the structural contributions and is ideally close to the data variance. The range (a) is the distance at which the curve reaches the sill - i.e. at which the spatial correlation drops practically to zero. Beyond it points are statistically unrelated. For the exponential and Gaussian models the sill is reached asymptotically, so for them the range is effective.

The nugget and contributions in the tool are set in absolute units of the parameter variance, not as fractions of one. The reference for the full sill is the data variance, which is shown in the report summary.

Parameters

Parameter	What it sets	Default / hint
Points with values	Input point layer.	-
Z value field	Attribute the variogram is built on.	-
Grouping field (adv.)	Separate variograms by category (e.g. survey type).	-
Declustering weight field (adv.)	Weights from 1.01 for a weighted estimate.	-
Number of lags	Number of distance bins.	-

Parameter	What it sets	Default / hint
Maximum distance	Upper distance bound. 0 = half of the extent diagonal.	0
Fit model	Automatic model fitting to the experimental points (recommended).	on
Model to fit (adv.)	Model type: spherical, exponential, Gaussian, power.	-
Robust estimator (adv.)	Cressie-Hawkins estimator, less sensitive to outliers.	off
Outliers: clip percentile (adv.)	Removal of extreme values by percentile. 0 = off.	0
Cap to bound instead of removing (adv.)	Capping: an extreme is clipped to the threshold, not removed.	off
Save profile as	Name of the “variogram + outlier removal” profile to apply in kriging. Empty = do not save.	-
Variogram table	Output: lag, γ value, number of pairs.	-
Report (HTML)	Output: pair-cloud and model plot.	-

Parameter	What it sets	Default / advice
Points with values	A point layer of wells or samples.	-
Z value field	The numeric attribute to analyse: roof elevation, thickness, grade.	remembered between runs
Grouping field (optional)	Builds a separate curve for each field value (e.g. survey type) and overlays them on one chart.	off
Number of lags	Into how many distance intervals the pair cloud is split.	15
Maximum distance	The far edge of the variogram, in layer units (for metric coordinates - metres). 0 = half the extent diagonal.	0
Fit model (recommendation)	Auto-fit of the nugget, sill, range and model type; the result is remembered for substitution into **2D Kriging**.	on

Parameter	What it sets	Default / advice
Model to fit (adv.)	Fix the model type or leave the auto-pick of the best by R^2.	Auto
Minimum points per group, % (adv.)	Groups smaller than the threshold are not built and are listed in the Log. The floor is 30 points.	2
Robust estimator (Cressie-Hawkins) (adv.)	Reduces the influence of rare anomalous pairs.	off
Show pair cloud (adv.)	Adds the source pairs (before averaging) to the chart.	off
Overlay a given variogram model (adv.)	Draws a model with a manually set nugget, sill and range over the cloud - handy to compare your model with the data.	off
Outliers (adv.)	Clip percentile, lower and upper value bounds, capping-to-bound mode instead of removal. At the very end of the list.	off
Save profile as	If filled in - the fitted (isotropic) model and the current outlier settings are saved as a processing profile under this name.	empty
Variogram table	An output table layer with the variogram points (columns below).	temporary layer
Report (HTML)	A report with the cloud, the fitted curve and the data-variance line.	temporary file

The parameters marked “adv.” are in the collapsed **Advanced Parameters** section.

The output is a **Report (HTML)** with the chart, the fitted curve and the data-variance line, plus an optional **Variogram table** - the experimental variogram points as a geometry-less layer (one row per lag of each series). From it you can build your own chart in QGIS or export the values. Its columns:

Field	Type	What it holds
series	string	The series: “all points” or the grouping-field value, if set.
lag	double	The mean distance between points in the interval (lag), in layer units.

Field	Type	What it holds
gamma	double	The semivariance γ(h): the mean of half the squared value differences over the pairs of this lag (or the robust Cressie-Hawkins estimate, if enabled).
npairs	integer	The number of point pairs that fell into the lag. A small number of pairs means the variogram point is unreliable.

The grouping field and a mixed-density survey

The optional **Grouping field** builds a separate variogram for each field value and overlays them on one chart. This is needed when the sample is collected by networks of different nature and density, for example surface and underground exploration. By feeding the survey type into the grouping, you can see whether these populations share a structure or each has its own.

Mixing mixed-density networks does not create artefacts by itself, but it distorts the overall variogram. A dense network gives many pairs at short lags and shapes the near part of the curve, a sparse network works on the far lags. A single model stretched over such a cloud turns out to be a mixture of two structures and describes neither correctly. The grouping shows this mixture, and the decision whether it is legitimate to combine the populations stays with the geologist. Declustering is not applied to the pair cloud here. Its weights are meant to correct the histogram and the mean, not the variogram pairs, where each pair is equally valid regardless of grid density.

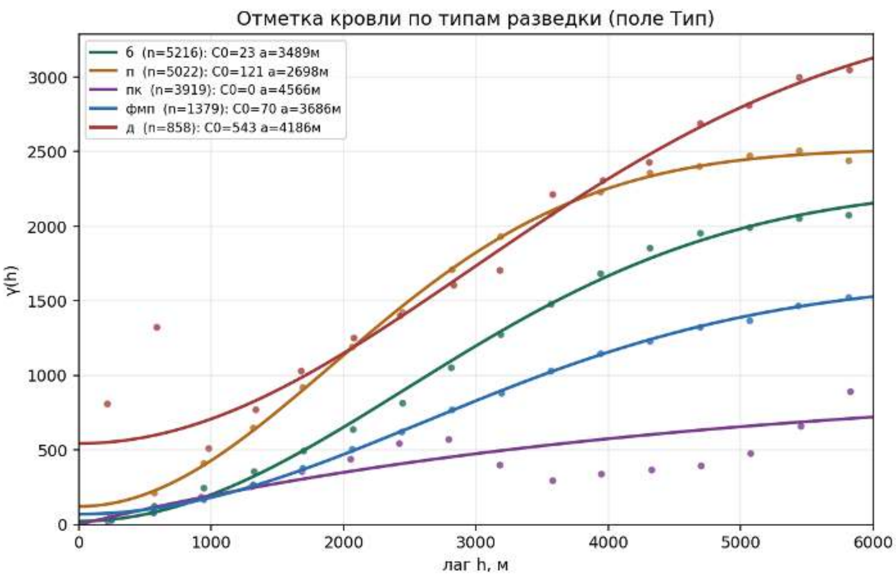

Figure 12: Roof elevation for the KrII seam, grouped by survey type: the underground network lies noticeably lower (a more homogeneous area), the detailed survey gives a high nugget. The different populations are visible at once.

Three typical geological situations

Seam elevations, thicknesses and component grades have different geostatistical characteristics, and it is useful to see them side by side. The illustration shows variograms of three parameters of one industrial seam, computed in a single distance window.

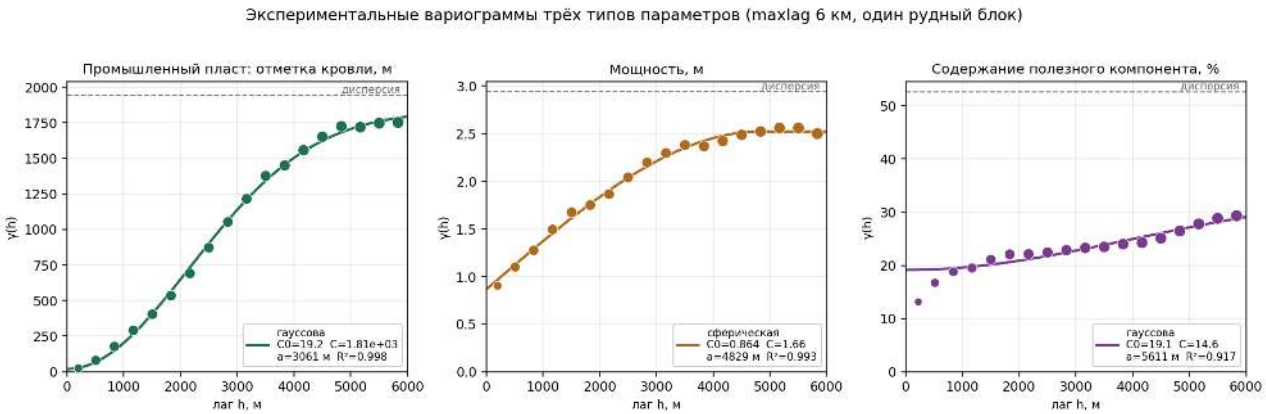

Figure 13: Three parameter types of one seam in a single window: roof elevation (almost zero nugget, smooth surface), thickness (nugget about a third, spherical) and grade (nugget comparable to the contribution, a noisy parameter).

Roof elevation is a smooth surface. The nugget is almost zero, the range large, the model close to Gaussian, the fit quality very high. Neighbouring wells give almost the same elevation, the variability is large-scale. Kriging works confidently. There is a subtlety here: a Gaussian model with an almost zero nugget is numerically unstable and gives the characteristic "bull's eyes" on the map. A small nugget should be set by hand.

Thickness is an intermediate case. The nugget makes up a noticeable fraction of the sill, the range is medium, the model more often spherical. About half the variability is structural, half small-scale. This is a typical working variogram.

The component grade is the noisiest parameter. The nugget is comparable to the structural contribution or exceeds it, the curve rises slowly, the fit quality is lower, and the model is poorly distinguishable from neighbouring types. The main variability sits at a scale finer than the sampling grid. Kriging smooths such a parameter heavily, and cross-validation shows a large error. Grade is predictably worse than elevations and thicknesses, and that is normal.

Maximum distance and reaching the plateau

The most common mistake is too large a maximum distance. If you leave the automatic value at half the diagonal, on an elongated deposit the window stretches over tens of kilometres. The lags begin to link points across barren gaps and inter-block breaks, the variogram catches the regional trend instead of the local structure, and the fit yields a range larger than the window itself and a sill several times the variance. The sign of trouble is simple: the fitted model's range is comparable to the window or exceeds it. This means the curve has not reached a plateau and the sill is obtained by extrapolation.

The cure is to reduce the maximum distance to the local scale and to check that the variogram has reached the plateau. On a grade example for one seam, with a 6-kilometre window the fit gave a range of about 9 kilometres and a sill below the variance, i.e. the curve had not yet reached the plateau. With a 12-kilometre window it did, giving a range of about 18 kilometres and a full sill close to the data variance. The real correlation range turned out larger than it looked in the narrow window, and the right answer came precisely from checking that the curve reaches the plateau.

At the same time the window must not step over large barren zones. On a drilling grid they are visible by the drop in point density, and the variogram should be built within a single ore block, otherwise the local geology mixes with regional tectonics.

The workflow with cross-validation

The variogram gives a starting model, and **Variogram cross-validation** checks it. The order is as follows. First an experimental variogram is built with a maximum distance at which the curve reaches a plateau, and the fitted nugget, contribution, range and model are taken. Then these numbers are carried into cross-validation and the leave-one-out metrics are assessed. The fitted and validated model is conveniently saved as a **processing profile** (the **Save profile as** field) and substituted into **2D Kriging** via the **Load processing profile** field - see the section on the Processing profiles tool.

The mean error ME should be near zero, meaning there is no systematic error. The root-mean-square error RMSE shows the absolute accuracy. The MSDR deserves separate attention - the ratio of the squared error to the kriging variance. If it is noticeably above one, kriging underestimates the uncertainty and the standard-error map is understated.

Correcting the MSDR is done exactly, not by eye. In ordinary kriging, multiplying the whole variogram by a constant factor does not change the estimate, since the weights depend only on the shape of the curve, not on its scale. Only the kriging variance changes. So it is enough to multiply the nugget and contributions by the current MSDR value, leaving the range and model unchanged, and repeat cross-validation. The ME, MAE, RMSE and R metrics do not shift, while the MSDR comes to one, and the error map becomes honest.

After scaling, the full sill may turn out above the data variance. On a clustered grid this is not an error. The naive variance is understated because dense well clusters pull it down, while the true scatter over the area is larger. The excess of the sill over the variance here is a consequence of the uneven grid.

The finished and validated model then only needs to be carried into **2D Kriging** to compute the grid, and after that, if needed, into **Isolines from raster**.

If the data are clustered unevenly, set the optional **wt** weight field from tool **1.01 Declustering**. Each pair of points is then taken with a weight equal to the product of its endpoints' weights, and clusters do not inflate the near lags. The pair count in the report shows the raw number of pairs, while γ itself is computed with weights.

1.06 Variogram map (anisotropy)

The tool builds a variogram map - the semivariance surface γ as a function of the two-dimensional separation vector (h_x, h_y). An ordinary variogram averages all directions into one curve and loses directionality; the map, by contrast, shows how the continuity of the parameter depends on direction. From it you can see whether there is anisotropy in the data and where the axis of maximum continuity points. The tool is diagnostic: it does not compute a grid but helps to set the azimuth and anisotropy in the 2D Kriging variogram structure deliberately.

What anisotropy is and why to see it

An isotropic variogram assumes the link between values depends only on the distance between points, not on direction. For folded and elongated geological bodies this is not so. Along strike the seam is sustained, across it it changes faster: the same difference in roof elevations is gained over kilometres along the fold but over hundreds of metres across it. If this is not accounted for, kriging smooths the field equally in all directions and blurs the real elongation of the structure.

A variogram map reveals the directionality directly. For each pair of points not only the distance is taken but also the direction of the vector between them, and the semivariance of the increment is spread over a two-dimensional grid of lags. Where γ grows slowly and the map stays dark far from the centre, continuity is high. Where γ grows fast, continuity is low. The low-γ area as a whole stretches into an ellipse whose long axis is the direction of maximum continuity - for folding this is the strike direction.

How to read the map

At the centre of the map lies the zero lag: a value at a point always equals itself, so γ here is zero and the centre is the darkest. As one moves away from the centre the points are separated farther and γ grows. The h_x axis points east, the h_y axis north, the scale on both axes is the same. The map is point-symmetric: a pair and its mirror image give the same semivariance, so the picture is the same in opposite directions.

Anisotropy is read from the shape of the dark area. If it is round - the structure is isotropic, direction plays no role. If it is elongated - along its long axis γ grows more slowly, i.e. in this direction values are linked over a larger distance. Hints are drawn over the map: a white ellipse by the estimated ranges and a red dashed line along the major axis.

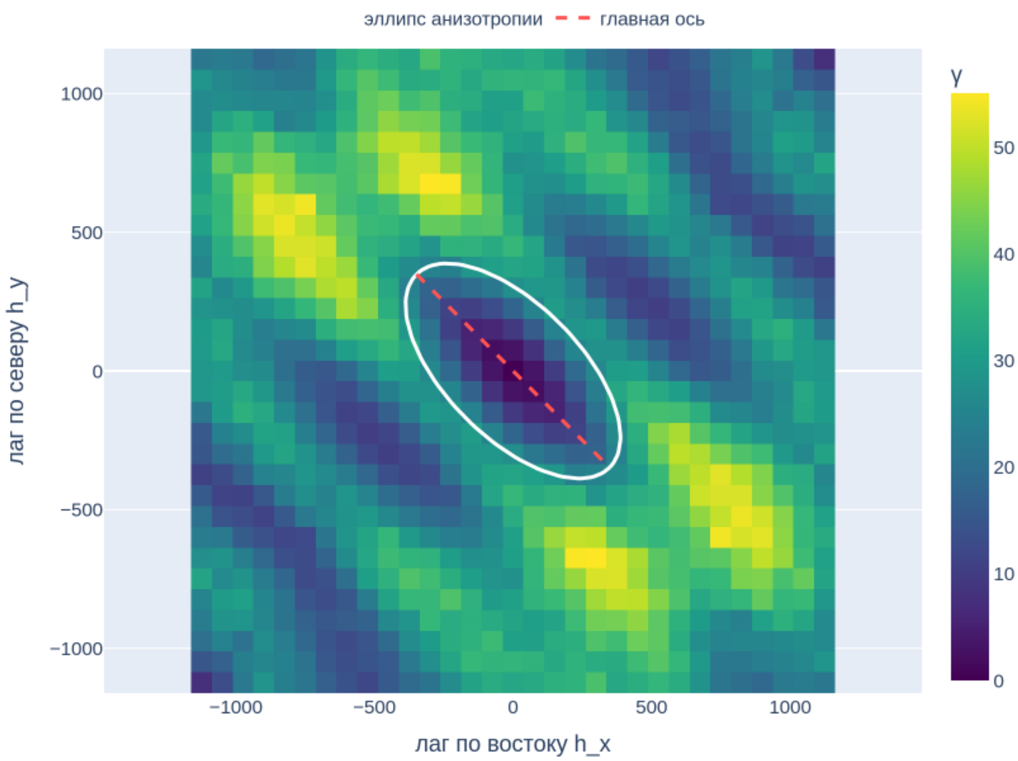

Figure 14: A variogram map: the dark (low γ, high continuity) area is elongated at an azimuth of about 135°. The white ellipse and the red dashed major axis show the estimated direction and anisotropy.

Parameters

Parameter	What it sets	Default / advice
Points with values	A point layer of wells or samples.	-
Z value field	The numeric attribute to analyse: roof elevation, thickness, grade.	remembered between runs

Parameter	What it sets	Default / advice
Bins per half-axis (map detail)	Into how many cells each lag half-axis is split. The map comes out (2N+1)×(2N+1) in size. More bins - a more detailed map, but fewer pairs per cell and more noise.	15
Max. lag, in layer units	The map window size, in layer units (for metric coordinates - metres). 0 = half the extent diagonal.	0
Min. number of pairs per cell (adv.)	Cells with fewer pairs are left empty. Cuts off noisy far lags where pairs are few.	5
Report (HTML)	A report with the heatmap, the ellipse, the major axis and a summary of estimates.	temporary file
Surface raster (opt.)	The γ surface as a raster in lag coordinates (see below). Not created by default.	off

The parameter marked “adv.” is in the collapsed **Advanced Parameters** section.

Estimating the azimuth, anisotropy and range

Besides the map itself the tool outputs to the Log and the HTML report three numbers: the major-axis azimuth (geographic, 0 - north, clockwise), the anisotropy coefficient as the ratio of the minor axis to the major (1 - isotropic, less - more elongated) and the major-axis range. The estimate works like this: along each direction the lag at which γ reaches the plateau (close to the data variance) is found, the ranges are smoothed over azimuth, the major axis is taken at the largest range, and the minor perpendicular to it.

These three numbers are substituted into the 2D Kriging variogram structure: azimuth, anisotropy (minor/major) and range a. This is exactly how anisotropy enters kriging. The estimate is indicative: it should be checked against the shape of the heatmap itself, not transferred blindly. The azimuth the map determines most reliably; the range and the coefficient are cruder, especially on a sparse network.

To avoid transferring the numbers by hand, the dialog has a **Write anisotropy to a profile** field. Pick a previously saved profile, and the azimuth, the coefficient and the major-axis range are written into it on top of the model and nugget set in **Variogram**. The next time the profile is loaded in **2D Kriging**, these values are applied automatically and appear in the caption under the profile list. If the range hit the window, it is left unchanged and only the azimuth and the coefficient are updated.

If the structure is close to isotropic or the major-axis range turns out smaller than a few map cells, anisotropy is not estimated and is marked in the report as “not expressed”. In this case the ranges lie at the grid level and the directionality is unreliable - it is more honest to report this than to give a random azimuth. It helps to reduce the max. lag or increase the number of bins to resolve the near structure.

When the range hits the window

If along the major axis γ does not manage to reach the plateau within the window, the range is returned equal to the max. lag, and a warning appears in the report and the Log: the range hit the max. lag, this is a lower bound. This is the same situation as for an ordinary variogram (see “Maximum distance and reaching the plateau”): the curve did not reach the plateau, and the sill is obtained by extrapolation. On the map the sign is simple - the dark area along the major axis stretches to the very edge.

In this case the range a cannot be carried into kriging as is: the real correlation length is larger than the window, and the anisotropy coefficient is understated in strength (the field is in fact even more anisotropic). The azimuth, meanwhile, is usually determined normally. The cure is to increase the max. lag so the map captures the plateau. And if γ does not reach the plateau even in a wide window, a trend dominates the data - it is removed before interpolation or accounted for with the appropriate kriging type.

The surface raster

If desired, the map is also saved as a raster (the **Surface raster** field). It is the same γ surface but in lag coordinates: the origin at (0, 0), the pixel size equal to the lag cell. The raster is not georeferenced - it lies in the separation space, not in the deposit plan - and is meant for those who want to spin the map on the QGIS canvas, apply their own colour scale or measure a lag with a ruler. The HTML report is enough for the anisotropy estimate itself.

If the data are clustered unevenly, set the optional **wt** weight field from tool **1.01 Declustering**. Each pair of points is then taken with a weight equal to the product of its endpoints’ weights, and clusters do not inflate the near lags. The pair count in the report shows the raw number of pairs, while γ itself is computed with weights.

1.07 Variogram cross-validation

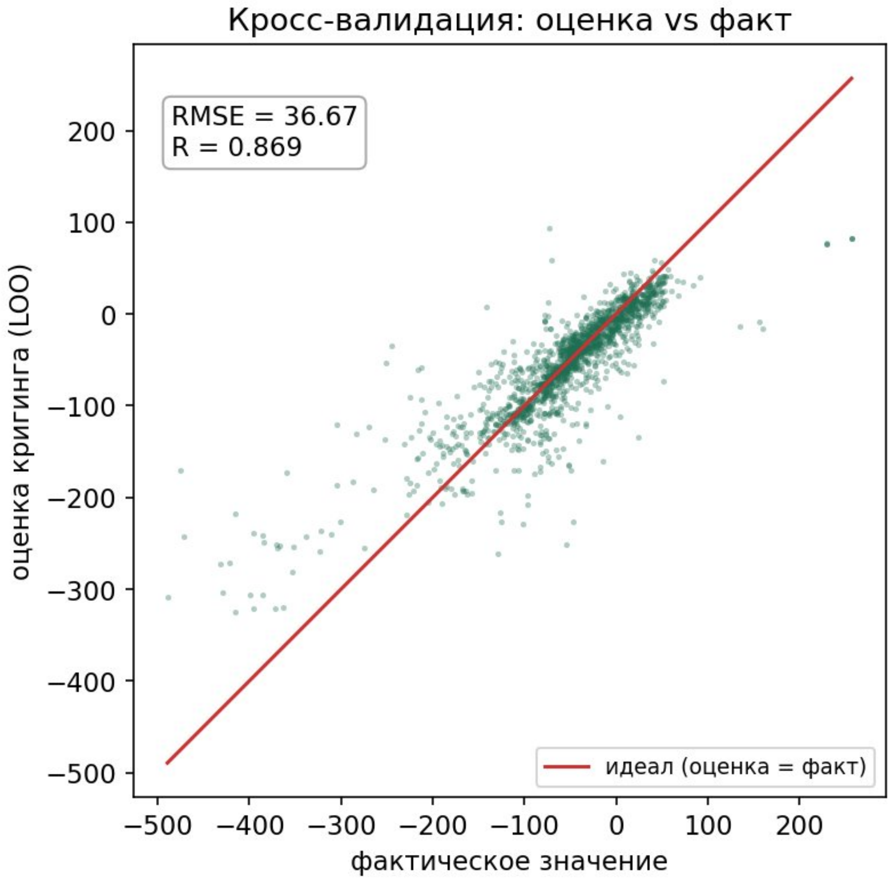

Figure 15: The idea of cross-validation: the kriging estimate from the remaining points (vertical) is compared with the actual value (horizontal). The tighter the cloud lies on the estimate = actual diagonal, the more accurate the prediction.

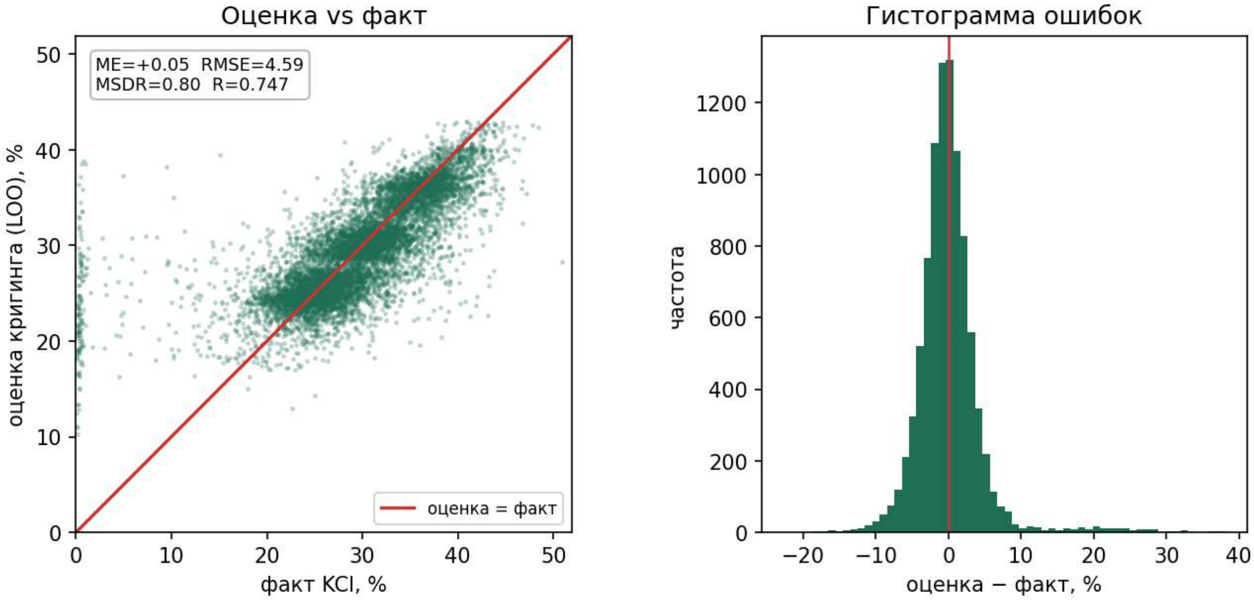

Figure 16: What the cross-validation HTML report looks like: on the left the “estimate vs actual” chart with the diagonal and metrics (example - KCl for the KrII seam), on the right the error histogram. A dense cloud along the diagonal - the model works. A band at an actual value near 0 - replacement zones. The histogram is symmetric about 0 - no bias.

The tool checks how well the variogram is fitted, by the leave-one-out method: each well in turn is excluded, its value is predicted by kriging from all the rest, and compared with the actual one. This way the parameters (nugget, range, model) are tuned by error rather than subjectively.

Parameters:

Parameter	What it sets	Default / advice
Points with values	Source points (wells).	-
Value field Z	The numeric attribute being checked.	-
Well-number field	An ID field for labels in the report and residuals layer.	optional
Kriging type, radius, min/max points, nugget, structures	Kriging and variogram settings. The check runs kriging with exactly these, so a good set carries into “2D Kriging” unchanged.	as in “2D Kriging”
Remove polynomial trend	Regression kriging: the trend is refit at each LOO step, the gain shows in the RMSE.	off
Trend degree	A plane or a quadratic surface.	1 (plane)
Load processing profile	Apply a saved model over the dialog fields.	(none)
Save profile as	Save the validated model to a profile (with anisotropy).	empty = do not save
Residuals layer (points)	Points with actual/estimate/error fields (see the field table below).	optional

Parameter	What it sets	Default / advice
Cross-validation report (HTML)	Interactive report: estimate vs actual, histogram, QQ-plot, metrics.	created by default

The Log outputs the metrics:

ME (mean error) - the systematic error. Should be close to 0 (unbiasedness).

MAE and RMSE - the mean and root-mean-square prediction error. The smaller, the more accurate. But RMSE alone is not enough: it is minimal at a zero nugget (overfitting), although the uncertainty is then estimated wrongly.

MSDR (standardized error) - the mean square of the error divided by the kriging standard error. Should be close to 1. If MSDR is noticeably above 1 - the variance is underestimated (the nugget or sill are small). If below 1 - overestimated.

R - the “estimate - actual” correlation coefficient.

It is useful to distinguish two sides. The “estimate - actual” cloud and the RMSE speak of the **prediction accuracy**. How correct the **model** itself is - i.e. whether the variogram honestly describes the uncertainty - is shown by the standardized errors: an MSDR near 1 and the QQ-plot. For kriging both sides are valuable: a small RMSE with an MSDR near 1 means the model both predicts well and does not deceive itself about its own accuracy. Chasing RMSE alone is not allowed - it is minimal at a zero nugget, where the uncertainty is understated.

In practice try several variogram variants and compare. A good model gives ME near 0, a small RMSE and an MSDR near 1. If the RMSE pulls toward a zero nugget while the MSDR is huge - this is a sign of overfitting. A small nugget calibrates the uncertainty.

The optional residuals layer (points with fields: the actual value under the validated field’s name, z_est, error, abs_error and std_resid, plus the well number if an ID field is set) shows where the model misses: large residuals by absolute value are problem areas, systematic residual signs are a local trend. The layer is automatically named after the validated field and the source, and the fields have aliases - readable names (visible in the attribute table and field properties). std_resid is the standardized residual (estimate − actual) / the kriging standard error, signed: minus - kriging underestimated, plus - overestimated (it is not a variance, a variance is always ≥ 0).

The residuals-layer fields:

Field	Alias	Description
`<well number>`	Well number	The value of the chosen ID field (if **Well number field** is set).
`<field name>`	Actual (field name)	The actual value of the validated field.
`z_est`	Kriging estimate (LOO)	The estimate from the other points (leave-one-out).
`error`	Error (estimate − actual)	Estimate minus actual. Minus - underestimated, plus - overestimated.
`abs_error`	|Error|	The absolute value of the error, |error|.
`std_resid`	Std. residual (signed)	(estimate − actual) / the kriging standard error, signed. Not a variance (which is ≥ 0).

Besides the residuals layer the tool by default produces an HTML report (on plotly): an interactive “estimate vs actual” chart with two lines - a grey 1:1 diagonal (the ideal) and a blue Best-fit regression line, an error histogram, a residuals QQ-plot and a metrics table with a recommendations block. The Best-fit slope, intercept and angle are added to the metrics: this is a range-bias indicator. A slope near 1 means the method is equally accurate at low and high values, a slope below 1 means high values are underestimated and low ones overestimated (regression to the mean, the signature of smoothing methods). The data variance is added to the table - a reference for the total sill C0+C. Next to the metrics table a **Kriging parameters** block is shown: only the settings that differ from the defaults are listed (nugget, sill, range, outliers and so on), so you can see which parameters produced these metrics. On the “estimate vs actual” chart, hovering over a point shows the well number and the values, and the eight wells with the largest residuals by absolute value are labelled right on the chart - they are convenient to check first. The report opens in the QGIS result viewer (or in a browser). If plotly is unavailable in the QGIS build, the report is still created - with the metrics table but without charts.

The residuals QQ-plot. Shows the shape of the error distribution. The errors are normalized to their own variance (a z-score) and compared with the normal distribution, so the chart reads by shape at any calibration. The uncertainty scale is handled separately by the MSDR in the metrics table. The horizontal axis is the normal-distribution quantiles, the vertical is the normalized error. If the errors are normal, the points lie on the red diagonal. Deviations read at once. Curled ends (S-shaped) - heavy tails, i.e. more large misses than under normality. An overall arc - skew, worth considering a value transform. A separate group broken off the line - an alien population in the data, for example barren samples from replacement zones (where the component is practically absent). Normality matters because the MSDR and the standard-error map rest on it.

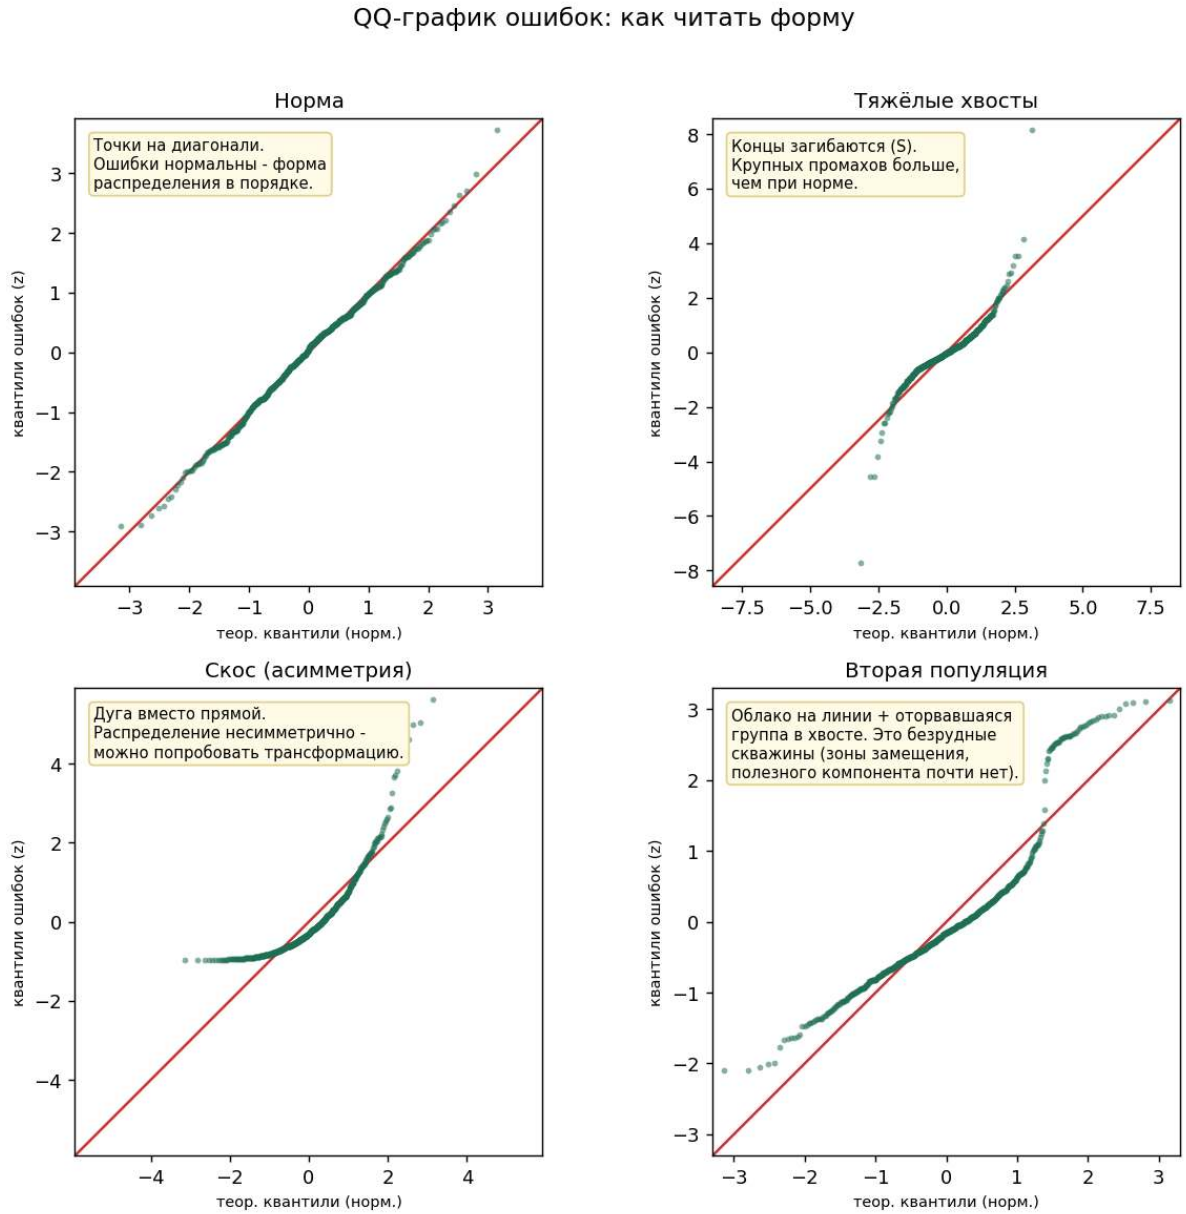

Figure 17: Four typical shapes of the residuals QQ-plot: normal (points on the diagonal), heavy tails, skew and a second population - a group broken off in the tail.

The main thing - what to do with the results. The point of the tool is, before building the grid, to approve or correct the whole set of parameters you will then set in **2D Kriging**. This is both the variogram (nugget, sill, range, model, anisotropy) and the kriging settings themselves (search radius, min/max points, type - ordinary or simple): cross-validation computes kriging with exactly the same settings, so a good set is carried into the **2D Kriging** tool unchanged. The order of decisions:

- ME near 0, MSDR near 1, the RMSE and R suit you - the set can be approved: carry these same parameters (the variogram and the search settings) into **2D Kriging** and build the surface.
- MSDR noticeably above 1 - kriging is too “sure of itself”, the standard-error map will be understated: increase the nugget C0 or the sill and check again.
- MSDR below 1 - the uncertainty is overstated: reduce the nugget or the sill.

- ME noticeably different from 0 - a systematic shift: check the data and the kriging type (for simple kriging - the specified mean).
- A large RMSE and a low R - the model predicts poorly: try a different range, model or anisotropy (azimuth and axis ratio). If nothing helps - it is the data's limit: short-range variability the network does not catch (e.g. ore replacement zones - on the chart above this is the vertical band at an actual value near 0).

The residuals layer prompts pointwise: where the residuals are large - the network should be densified (add wells) or the samples checked. Where the residuals are systematically of one sign over the area - there is a local trend kriging did not account for.

In sum: this tool is the last step before the final kriging. First you calibrate the variogram here by error, then set the same parameters in **2D Kriging** - and the surface together with the standard-error map come out justified rather than fitted subjectively.

A note on speed: the check solves kriging as many times as there are points, so on large sets (tens of thousands of wells) it runs noticeably longer. Reduce the sample if needed.

If the data are clustered unevenly, set the optional **wt** weight field from tool **1.01 Declustering**. The ME, MAE, RMSE, MSDR and R metrics are then computed with weights, so a dense cluster of wells does not dominate the quality assessment. The leave-one-out estimate itself is unchanged, only the summary is weighted.

1.08 Method cross-validation (LOO)

Leave-one-out control for a gridding method: kriging or minimum curvature. Each validation point is removed in turn, its value is predicted by the method from the rest and compared with the fact. The errors give quality metrics - an objective measure of the method and a way to compare methods on your own data.

This differs from **1.07 Cross-validation of the variogram**: that one fits the variogram model for kriging, while this one compares gridding methods as such and works for minimum curvature too.

Metrics: **ME** (bias, closer to 0), **MAE** and **RMSE** (smaller is better), **R** (correlation of estimate and fact). For kriging there is also **MSDR** (closer to 1 when the standard-error scale is adequate). The estimate-vs-fact chart has two lines: a grey 1:1 diagonal (the ideal) and a blue **Best-fit** regression line. Its slope, intercept and angle go into the metrics as a range-bias indicator. A slope near 1 means the method is equally accurate at low and high values, a slope below 1 means high values are underestimated and low ones overestimated (regression to the mean, the signature of smoothing methods).

Three Surfer-style options are available. A **random subset** of N points speeds control on large data, while the whole sample still takes part in each estimate. An **area filter** restricts validation to a subarea by extent and by Z value, useful to avoid control at known anomalies. An **exclusion buffer** in X and Y drops points in a rectangle around the validation point, needed for dense clusters, otherwise the estimate just repeats the nearest neighbour.

Parameter	Purpose	Default
Points with values	Samples.	-
Value field (Z)	Numeric field.	-
Well id field	For labels in the report.	-
Method	Kriging or minimum curvature.	Kriging
Points to validate	0 = auto, min(N, 100).	0
Exclusion buffer in X, Y	Rectangle around the point, neighbours in it are left out.	0

Parameter	Purpose	Default
Kriging parameters	Variogram and search (as in 1.1).	-
Min curvature parameters (Adv.)	Extent, cell, tension, threshold, iterations.	auto
Validate only within the extent (Adv.)	Control area by X/Y.	everywhere
Validate where Z is in range (Adv.)	Control area by value.	none
RNG seed (Adv.)	For a reproducible subset.	0
Cross-validation errors	Point layer with fact, estimate and error fields.	-
HTML report	Estimate-vs-fact plot, histogram, metrics.	default

For minimum curvature each point is re-estimated from a warm start off the full solution, so a single pass is fast. On very large samples reduce the number of validation points.

If the data are clustered unevenly, set the optional **wt** weight field from tool **1.01 Declustering**. The ME, MAE, RMSE, MSDR and R metrics are then computed with weights, so a dense cluster of wells does not dominate the quality assessment. The leave-one-out estimate itself is unchanged, only the summary is weighted.

1.09 Processing profiles

A profile is a named set of processing settings for one parameter: the variogram (nugget C0, model type, contribution C, range a, azimuth and anisotropy axes) plus outlier removal (percentile, bounds, capping mode). Profiles are handy when a project has several seams or zones of different variability: you fit a model for a seam once and reuse it in kriging without re-entering the numbers.

Profiles are stored globally in the QGIS settings, so they are available across all projects: build a seam's model once - apply it anywhere. A profile describes one variogram structure - exactly as much as kriging uses.

Where profiles come from

- **Variogram** - the **Save profile as** field. The fitted model is saved. The curve is built isotropic, so the azimuth and axes are written as neutral (0 and 1) - anisotropy is set later.
- **Cross-validation** - the **Save profile as** field. The validated model is saved together with the set anisotropy. This is the main way to get a profile with an azimuth and axes.
- **Processing profiles** - the **Save manually** action: all profile values are entered in the fields of the **Advanced Parameters** section.

Application

In the **2D Kriging** and **Cross-validation** tools the **Load processing profile** field substitutes the chosen profile over the dialog fields. What exactly is substituted is printed to the Log.

Management

The **Processing profiles** tool itself manages the storage via the **Action** parameter:

Action	What it does
Show list	Outputs all profiles with their parameters to the Log.
Save manually	Saves a profile with the name from the **Profile name** field by the values of the fields in **Advanced**.
Delete selected	Deletes the profile chosen in the **Profile** field.
Clear all	Deletes all profiles.

Saving under an existing name overwrites the profile. The profile lists in the drop-down fields (the choice for deletion, the load in kriging) refresh when the tool window opens: after saving a profile, reopen the tool so it appears in the list.

Below the profile drop-down, in the line beneath it, the parameters of the chosen profile are shown (nugget, type, contribution, range, azimuth, axes, outliers). In **2D Kriging** and **Cross-validation** a reminder is shown there as well that the computation will use the profile rather than the dialog fields. On QGIS builds without the old widget API the caption does not appear - an ordinary list remains (this does not affect the work).

1.10 Create sample wells (demo)

The **Create sample wells (demo)** tool builds a point layer with random coordinates and three structured fields: the absolute roof elevation (roof), the thickness (thick) and the grade of an abstract component X (%). The roof and thickness ranges are set after the model of an industrial seam (KrII). The tool is meant for learning and testing kriging, isolines and cross-validation without real data.

Parameters:

Parameter	What it sets	Default / advice
Area (extent)	The generation rectangle (by layer, canvas, coordinates, drawing).	-
Number of wells	How many points to create.	300
Minimum / maximum of value X	The component grade range.	0 / 50
Smoothness (fraction of extent)	Correlation range as a fraction of the extent: larger - bigger “patches”.	0.15
Roof, thickness: min/max (Adv.)	Ranges of the roof and thick fields.	as for KrII
Nugget fraction (Adv.)	Share of variance on short-range noise: larger - less predictable.	0.35
Add a categorical mineral-type field	A mintype field (silvinite, replacement) for indicator kriging.	off
Add a head field	A head field with a regional slope for the flow gradient.	off
Add K and T fields and head	Head plus log-normal K (m/day) and T = K·thickness for the specific discharge (Darcy).	off

Parameter	What it sets	Default / advice
RNG seed (Adv.)	Reproducibility of the generation. 0 = random.	0
Declustering weight field (Adv.)	**wt** weights from tool **1.01**, optional.	none
Sample wells (demo)	The output point layer.	-
Drift surface (raster) + dz field	Enable the output to get an s raster and a dz field for external drift.	off (skipped)

At start the Log outputs the starting variogram (total sill ≈ the data variance, nugget, range). The generated data have a recoverable variogram, so it is convenient to learn the whole cycle on them: build a grid in **2D Kriging**, then isolines, and check the parameters with cross-validation.

Two checkboxes and a separate output add optional fields for learning the related tools. **Add a categorical mineral-type field** adds a mintype field with a silvinite background and replacement spots for categorical indicator kriging. **Add a head field** adds a head field with a pronounced regional slope for the hydraulic gradient: krige head, feed the raster to the flow tool, and the arrows follow the head downhill. Enabling the **Drift surface** output writes, as a separate raster, a smooth secondary surface s known everywhere, and adds a dz field linearly related to it. This pair is for learning external drift kriging: krige dz with the s raster as the drift and compare it with plain kriging of dz without the drift. If the drift-surface output is skipped, the dz field is not added. The **Add K and T fields and head** checkbox generates head and log-normal K (hydraulic conductivity, spanning orders of magnitude as in real pumping tests) and T = K·thickness. They are for learning the Darcy specific discharge: krige K and T in **2D Kriging** with the **ln** transform (or ln fields by hand), plus head, then feed the rasters into the **Specific discharge** tool.

Result fields:

Field	Type	Holds
well	text	Well number, format SK-0001.
roof	number	Absolute roof elevation of the seam, m.
thick	number	Seam thickness, m.
X	number	Grade of the abstract component, %.
head	number	Head (piezometric level), m. With the head checkbox or the K and T checkbox.
K	number	Hydraulic conductivity, m/day (log-normal). Only with the K and T checkbox.
T	number	Transmissivity T = K·thickness, m^2/day. Only with the K and T checkbox.
mintype	text	Mineral type (silvinite, partial replacement, rock salt). Only with the mineral-type checkbox.
dz	number	A value linearly related to the drift surface. Only when the drift-surface output is enabled.

The full list of output-layer fields is in the **Sample wells (demo)** appendix section ("Demo-layer fields" at the end of the manual).

1.11 Create a geophysical-profiles example (demo)

The tool creates a point layer of geophysical profiles for learning and testing without real data. Several parallel profiles with pickets are built. There are two modes.

Electrical prospecting. Along the profiles apparent resistivity rho_k (Ohm*m), self-potential SP (mV) and induced polarisation IP (mV/V) are set. The data contains a low-resistivity anomaly - a water-bearing or replaced zone. There rho_k drops from a background of tens of Ohm*m to units, and SP shows a negative minimum. This matches field practice: the lowest resistivities are typical of sedimentary rocks, while salt, gypsum and anhydrite are high-resistivity, and water saturation or salinisation drives the resistivity down. The anomaly is set as a compact spot rather than a stripe, so the profiles are not synchronous and interpolation shows a local focus rather than a solid band.

Subsidence (trough). The value is settlement (mm) as a subsidence trough over a mined area, across several observation tours. The trough deepens from tour to tour and is capped at two metres in magnitude. The sign is uniform: down (negative) or magnitude (positive), by choice. The edges are strictly zero - away from the mined area there is no subsidence. At the same pickets you can compute the **settle** difference between tours to get the settlement rate.

The workflow repeats the main one: the **rho_k** (or **settle**) field is interpolated with **1.02 2D Kriging** or minimum curvature, isolines are built from the grid with **1.04**, and the anomaly is outlined. The **sp** field can be interpolated the same way and its SP minimum compared with the rho_k drop. The **rho_true** (or **settle_true**) field is the embedded noise-free value, a reference for checking interpolation accuracy.

Kriging or minimum curvature

Geophysical profiles are a typical case where the choice of interpolation method matters more than its tuning. The data are dense along the profiles and sparse between them, while the quantity itself (resistivity, potential) is physically smooth and continuous.

Kriging faithfully reflects the uneven network: without tuning the variogram anisotropy it stretches the structure along the survey lines, and the field breaks into bands along the pickets. Minimum curvature (**1.03**) imposes a physically meaningful smoothness and stitches the separate profiles into a connected surface, so for profile surveys and potential fields it is usually preferable. The same holds for the **sp** field: the SP minimum appears as a single body rather than columns. It is most vivid on subsidence: the trough is a compact axisymmetric bowl, and kriging rolls it into a band along the profiles, while minimum curvature restores the bowl with a clear centre.

Practical takeaway: build fields from profiles with minimum curvature, and use kriging when the variogram anisotropy is tuned to the network geometry.

The area is set by an extent. The mode, ranges, number of tours and surface elevation can be changed under **Advanced**.

Parameter	Purpose	Default
Extent	Generation bounds.	-
Mode	Electrical prospecting or subsidence.	Electrical prospecting
Number of profiles	How many parallel profiles.	4
Picket step, m	Distance between pickets.	20

Parameter	Purpose	Default
Background rho_k, Ohm*m (Adv.)	Resistivity outside the anomaly.	60
Minimum rho_k, Ohm*m (Adv.)	Resistivity at the spot centre.	10
SP amplitude, mV (Adv.)	Depth of the SP minimum, usually below 0.	-100
IP amplitude, mV/V (Adv.)	Height of the IP anomaly.	15
rho_k noise (Adv.)	Spread fraction in log scale.	0.06
Maximum subsidence, mm (Adv.)	Trough depth at the last tour.	400
Number of tours (Adv.)	How many cycles for subsidence mode.	2
Subsidence sign (Adv.)	Down (negative) or magnitude (positive).	Down
Elevation: base and amplitude, m (Adv.)	Smooth surface relief.	120 and 15
Geophysical profiles	Point layer with the fields.	-

Electrical fields: **profile** (profile number), **picket_m** (picket in metres from the profile start), **pk** (a PK label, e.g. PK5+20), **z** (surface elevation, m), **rho_k** (rho_k, Ohm*m), **rho_true** (rho_k without noise), **sp** (SP, mV), **vp** (IP, mV/V).

Subsidence fields: **profile**, **picket_m**, **pk**, **tour** (tour number), **z** (elevation, m), **settle** (subsidence, mm), **settle_true** (subsidence without noise).

Topography: terrain from open data

The **“2. Topography”** group answers a frequent request: the best possible terrain model from open data, out of the box. The front door is the DEM downloader by extent, next to it the vector base map from OpenStreetMap, the Topo2Raster core that builds terrain from points and contours, and the full hydrology set: depression filling, flow and accumulation, the river network, basins, slope with aspect, and peaks. All the analytics run on pure NumPy, without GRASS, SAGA or external modules.

All output layers of the group land in the **Topography** group of the layer tree, so they do not drown among the working layers of the project. The tools of the group chain together. Downloader 2.01 delivers a ready metric DEM that goes straight into isolines (1.04) and any computation of the group. Watercourses from 2.02 and the river network from 2.06 fit Topo2Raster (2.03) as streamlines as is, because their vertices run downstream.

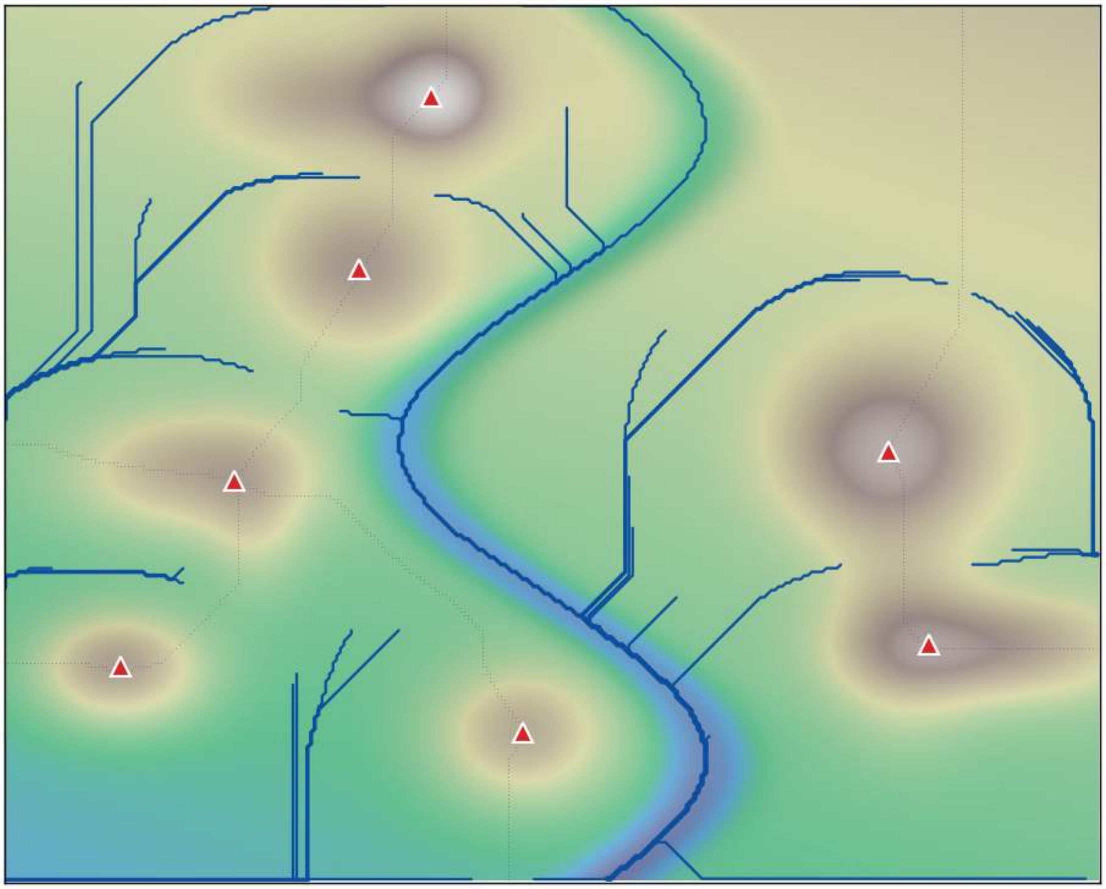

Figure 18: The full chain of the group on the demo relief: hillshade, the river network with width by Strahler order (2.06), basin boundaries (2.07) and peaks (2.09).

The full list of output-layer fields is in the **Electrical-prospecting / Subsidence profiles** appendix section ("Demo-layer fields" at the end of the manual).

2.01 Download DEM by extent

Downloads a DEM by extent from an open store, no registration or keys, from one of two sources. One-degree tiles are mosaicked seamlessly and reprojected into a metric coordinate system with cubic resampling. Raw degree tiles never enter the analysis, so the GLO-30 peculiarity north of latitude 50 (a coarser longitude step) is handled automatically.

Parameter	What it sets	Default / hint
Terrain source	GLO-30 (Copernicus, DSM: heights over canopy and roofs) or GEDTM30 (DTM with forest and buildings removed, machine learning from ICESat-2/GEDI).	GLO-30
Download extent	Extent in any CRS, converted to degrees automatically.	-

Parameter	What it sets	Default / hint
Target CRS	Metric output CRS. Empty = project CRS if metric, otherwise UTM by extent center.	empty
Cell size, m	Output grid step after reprojection.	30
Smooth terrain (FPDEMS)	Edge-preserving smoothing (see below). off = raw source raster.	off
Hydrological correction	Filling of false depressions so water flows down. Uncheck for karst and subsidence troughs.	on
Smoothing: normals window, cells (adv.)	Side of the normals filter window, odd.	11
Smoothing: normals difference threshold, deg (adv.)	Edge-preservation strictness: lower = more aggressive preservation.	15
Smoothing: number of passes (adv.)	Outer elevation-rebuild steps.	2
Slope epsilon for filling, m (adv.)	Minimum slope on flats. 0 = raise pits to the outlet only.	0.001
Limit of 1x1 degree tiles (adv.)	Guard against an oversized extent for GLO-30.	25
DEM (metric CRS)	Output float32 raster, layer in the Topography group.	-

Terrain source is the key choice. **Copernicus GLO-30** is a digital surface model (DSM): heights are taken over the top of canopy and rooftops, distributed as one-degree tiles. **GEDTM30** is a digital terrain model (DTM) by OpenGeoHub under CC BY 4.0: forest and buildings are removed by machine learning from ICESat-2 and GEDI data, so under forest canopy GEDTM30 is markedly more accurate, confirmed by independent validation. It is distributed as a single global cloud-optimized GeoTIFF. For forested terrain the DTM is usually preferable, for open areas the difference is small.

Parameters:

- **Download extent** - the extent in any CRS, converted to degrees internally to pick the data.
- **Target CRS** - the metric CRS of the result. Leaving it empty enables the automatics: the project CRS is taken if it is metric, otherwise the UTM zone at the extent center. A degree target CRS is rejected with a clear message.
- **Cell size, m** - 30 by default, the native GLO-30 resolution.
- **Hydrological correction** - on by default: spurious depressions are filled right away (see 2.04) so water flows downhill on the model. For tasks where closed basins matter (karst, subsidence troughs) uncheck the box.
- Under **Advanced**: **Slope epsilon for filling** and **Tile limit** (a guard against an accidental extent covering half a country).

A network failure or an extent entirely in the ocean ends with a clear message rather than an empty raster. Data source: Copernicus DEM © ESA, the open license allows use with attribution.

2.02 Download base topography by extent

The vector twin of the DEM downloader: for the same extent it fetches terrain-related layers from OpenStreetMap.

Parameter	What it sets	Default / hint
Download extent	Extent of the OSM Overpass request. Keep it modest - large requests are cut by the server.	-
Extent area limit, square degrees	Guard against an oversized request.	0.5
Request timeout, s	Overpass response wait; on failure the tool switches to a mirror.	90

Output layers: watercourses, water bodies (planes), peaks with the **ele** field, cliffs (barriers). All land in the Topography group.

- **Watercourses** - rivers, streams and canals. In OSM watercourses are drawn downstream, so the layer fits 2.03 as streamlines without preparation.
- **Water bodies** - closed outlines of lakes and ponds, constant-elevation planes for 2.03. Compound multipolygons (large lakes assembled from several ways) are skipped by the first version.
- **Peaks with elevations** - natural=peak points with the ele elevation, ready summit marks for 2.03 and a control for 2.09.
- **Cliffs and embankments** - terrain breaklines for 2.03.
- **Coastline** - off by default, needed on coastal territories.

Output is in the project CRS, lines are clipped to the extent. Public Overpass servers have limits, on a failure of the main server the request goes to a mirror, for large territories shrink the extent or raise the area limit under **Advanced**. Data: © OpenStreetMap contributors, ODbL license.

2.03 Topo2Raster (terrain from vectors)

Builds terrain from vector data by multigrid interpolation from a coarse grid to a fine one, in the spirit of ANUDEM. The tool covers the classic task: digitized contour lines of a topographic plan, spot elevations, rivers and lakes are at hand, and a correct grid is needed.

Parameter	What it sets	Default / hint
Elevation points	Layer of spot heights (e.g. peaks from 2.09).	-
Point elevation field	Z attribute of the points.	-
Isolines	Contour layer (digitized plan or 1.04 over a DEM).	-
Contour elevation field	Level attribute of the contours.	ELEV
Streamlines (downstream)	River lines, vertices pointing downstream. The 2.06 network and OSM fit as is.	-
Cliffs (smoothing barriers)	Lines along which the drop is not smeared.	-
Lakes and water edge	Water-body planes.	-

Parameter	What it sets	Default / hint
Water-edge elevation field	Water-edge elevation. Empty = node Z (slope) or shore minimum.	empty
Extent	Output extent. Empty = from input layers.	empty
Cell size, m	Output grid step.	30
Fill depressions in the result	Hydrological correction of the result.	on
Smoothing iterations per level (adv.)	Passes of internal field smoothing.	60
Minimum streamline drop, m/cell (adv.)	Guaranteed slope along rivers.	0.01
Slope epsilon for filling, m (adv.)	Minimum slope on flats during filling.	0.001
Terrain	Output float32 raster.	-

Points or isolines are required (at least one elevation source). Water-edge priority: node Z (river slope) over the elevation field (plane), field over shore minimum.

Every input type works as its own constraint:

- **Elevation points** and **contours** - hard nodes, the surface passes through them exactly. At least one of these layers is required, each with a numeric elevation field.
- **Streamlines** - a forced monotonic drop downstream. Line vertices must run downstream: OSM watercourses (2.02) and the river network (2.06) fit as is. The minimum drop per cell is set under **Advanced**.
- **Cliffs** - smoothing barriers. The surfaces on the two sides of a cliff are independent, the step is not smeared. The line itself, one cell wide, gets an intermediate elevation, a limitation of the first version.
- **Water edge** - three height priorities per feature, and a layer may mix types. If a polygon has three-dimensional vertices (Z at nodes), the edge is interpolated from their heights and sloped along the channel - this is how a river's falling level from source to mouth is set. If there is no Z but an elevation field is filled, the feature is held as a horizontal plane. With neither, the level is taken automatically from the minimum of the adjacent shore. Mixed feature types in one layer are handled each by its own branch: a lake as a plane, a river as a slope.

Inside, a two-stroke cycle runs at every grid level: membrane smoothing sets the frame and holds the constraints, then minimum-curvature polishing (the Briggs stencil) removes the membrane bias between curved contours. On a round-trip test (demo relief, contours every 4 m, reconstruction, comparison with the original) the polishing cuts the error by roughly a third. The residual maximum error lives on summits above the last contour, so summit marks in the input visibly improve the tops - exactly why topographers label them on maps.

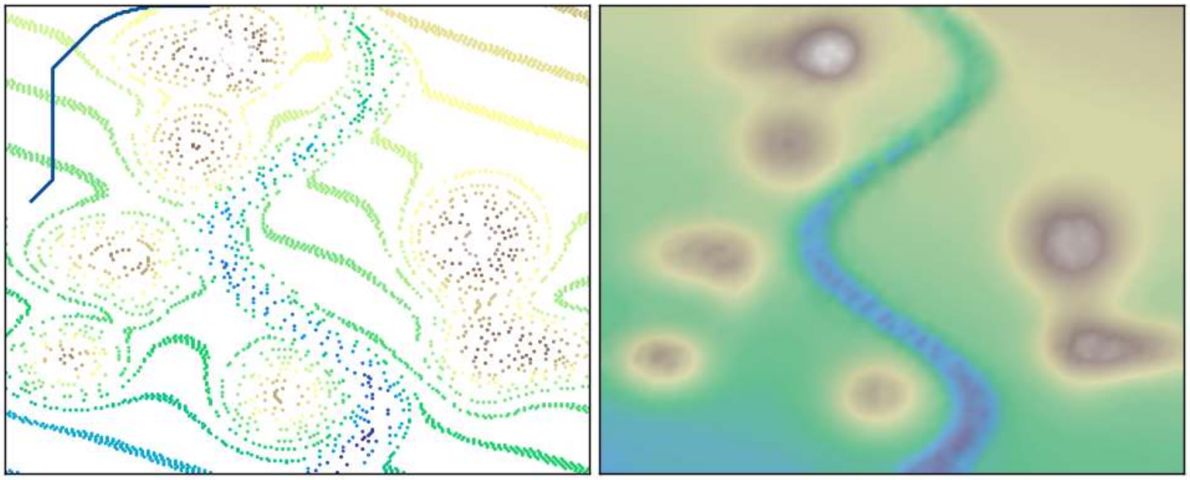

Figure 19: Left: the input constraints, densified contours (color by elevation) and the main streamline. Right: the reconstructed terrain.

The default extent is taken from the layers with a two-cell margin. All layers are brought to the CRS of the first given layer, which must be metric. The final **depression filling** is on by default, its logic is described in 2.04.

2.04 Terrain preparation

Prepares a DEM for analysis with two independent modifications, each toggled by its own checkbox, in a fixed order: smoothing first, then filling.

Parameter	What it sets	Default / hint
Input DEM	Terrain raster to prepare.	-
Smooth terrain (FPDEMS)	Edge-preserving smoothing (see below).	off
Fill depressions	Filling of false pits (Planchon-Darboux).	on
Slope epsilon, m (0: pits only)	Minimum slope on flats. 0 = raise pits to the outlet, above zero = a through slope for D8.	0.001
Smoothing: normals window, cells (adv.)	Side of the normals filter window, odd.	11
Smoothing: normals difference threshold, deg (adv.)	Edge-preservation strictness: lower = more aggressive.	15
Smoothing: number of passes (adv.)	Outer elevation-rebuild steps.	2
Prepared DEM	Output float32 raster.	-

Order: smoothing first, then filling. Both modifications are independent, each with its own checkbox.

Smooth terrain (FPDEMS) removes the excessive roughness of satellite models. Satellite DEMs are noisy, and ordinary filters - mean, median, Gaussian - cut the noise together with edges, terrace walls and banks: the breaks get flattened. FPDEMS (Lindsay, Francioni, Cockburn, 2019) works differently. It operates on the field of surface normals rather than heights directly: it first computes each cell's normal, then smooths the normal field so that a neighbour enters the average with a weight the larger the closer its normal is to the central one. At an edge the neighbours' normals diverge sharply, the weight drops, the edge is preserved. After that the heights are pulled towards the smoothed normal field. As a result flat areas are smoothed while structural lines stand. The **normals-difference threshold** (in **Advanced**) controls the strictness: a smaller value preserves edges more aggressively, a larger one smooths more overall. The method was originally proposed for lidar DEMs but is equally useful for satellite ones.

Fill depressions with the Planchon-Darboux method raises spurious pits so flow does not stop. Depressions in raster models are most often interpolation and noise artifacts, and hydrological analysis without filling breaks at the first pit. **Slope epsilon** controls the mode: with zero only true pits are raised exactly to the spill level, flat areas stay flat. With a positive value (0.001 m by default) a through slope is additionally built across flats, and D8 becomes defined on them. Flow and accumulation need a positive epsilon. Cells on the grid border and next to nodata are treated as outlets. The report prints the number of raised cells and the maximum raise - a handy indicator of the source DEM quality.

The same smoothing checkbox is present in DEM download (2.01) for a quick path right at download time. The standalone tool 2.04 is needed when the terrain came not from 2.01 but from your own data.

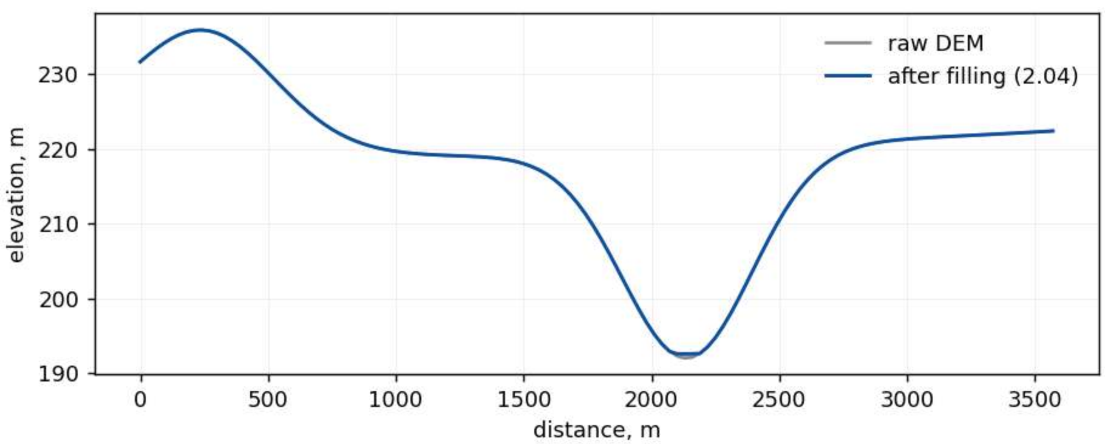

Figure 20: A profile through a depression: the raw surface and the filling result. The raised part is shaded.

2.05 Flow and accumulation (D8)

Computes flow directions over eight neighbors (D8, Jenson-Domingue) and accumulation: how many cells drain into each one, itself included. Directions are coded as in ArcGIS: E=1, SE=2, S=4, SW=8, W=16, NW=32, N=64, NE=128, sink=0, nodata=255.

Parameter	What it sets	Default / hint
Input DEM	Terrain raster.	-
Fill depressions before computing	Hydrological correction before D8, otherwise flow stalls in pits.	on
Slope epsilon for filling, m (adv.)	Minimum slope on flats.	0.001

Parameter	What it sets	Default / hint
Flow directions (ESRI codes)	Output: byte direction raster. E=1, SE=2, S=4, SW=8, W=16, NW=32, N=64, NE=128, sink=0.	-
Accumulation, cells	Output: float32, number of upstream cells.	-

The border semantics are deliberate. A cell on the grid frame leaves the grid only when it has no lower neighbor inside, otherwise flows running along the edge would break. A shore cell, on the contrary, pours into an adjacent nodata (the sea, a cutout) even when a lower land neighbor exists.

The **Fill depressions before computing** checkbox is on by default: on a raw DEM flow stops in pits and accumulation breaks. The computation is fully vectorized, a 2000×2000 grid takes seconds.

2.06 River network

Extracts the river network from a DEM: cells with accumulation at or above the threshold are linked from heads and junctions downstream.

Parameter	What it sets	Default / hint
Input DEM	Terrain raster.	-
Accumulation threshold, cells	Minimum catchment cells to start a channel. A denser network = a lower threshold. 1000 at 30 m is about 0.9 sq km.	1000
Fill depressions before computing	Hydrological correction before tracing.	on
Slope epsilon for filling, m (adv.)	Minimum slope on flats.	0.001
River network	Output: lines with order (Strahler), acc_out, length_m. Vertices point downstream.	-

The **accumulation threshold** is set in cells and means the catchment area where a river starts: the head catchment area divided by the cell area. For a 30 m DEM a threshold of 1000 starts rivers at a catchment of about 0.9 sq. km. A smaller threshold gives a denser network.

Output fields: **order** - the Strahler order (1 at heads, growing where two equal orders merge), **acc_out** - accumulation at the link outlet, **length_m** - the length. Line vertices run downstream, so the layer fits 2.03 as streamlines without preparation and compares against OSM watercourses from 2.02: overlaying the extracted network on the real one is a quick DEM quality check.

2.07 Basins and watersheds

Divides the territory into drainage basins, polygon boundaries are the watersheds.

Parameter	What it sets	Default / hint
Input DEM	Terrain raster.	-
Pour points	Basin outlet points. Empty = basins from mouths by threshold.	empty
Mouth accumulation threshold, cells	Threshold for mouths when no points are given.	1000
Fill depressions before computing	Hydrological correction before the computation.	on
Point snap radius, m (adv.)	An outlet point snaps to the accumulation maximum within this radius.	150
Slope epsilon for filling, m (adv.)	Minimum slope on flats.	0.001
Basins (polygons)	Output: polygons with basin, area_m2.	-
Basins (label raster)	Output: integer basin-label raster.	-

Two modes. With **pour points** every point snaps to the cell with the highest accumulation within the snap radius (otherwise a point placed by eye next to a river would collect a tiny hillslope basin) and gathers the whole catchment above itself. Without points the basins are built automatically from mouths: cells where flow leaves the grid with accumulation at or above the threshold. Cells outside every basin get label 0 and are not exported to polygons.

Output fields: **basin** - the basin number, **area_m2** - the area by cell count. A label raster can be written additionally. Labeling runs by pointer jumping over the flow graph, so even long winding catchments take a fraction of a second.

2.08 Slope and aspect

Slope in degrees and aspect with the Horn 3×3 kernel, as in gdaldem. Aspect is the downslope azimuth in degrees from north clockwise: north 0, east 90. Flat cells get an aspect of -1 so they are not confused with north-facing ones. Nodata cells and their neighbors get nodata: the 3×3 kernel is not computed across holes.

Parameter	What it sets	Default / hint
Input DEM	Terrain raster.	-
Slope, degrees	Output: slope raster (Horn 3x3, as gdaldem).	-
Aspect, degrees	Output: downslope azimuth, flats = -1.	-

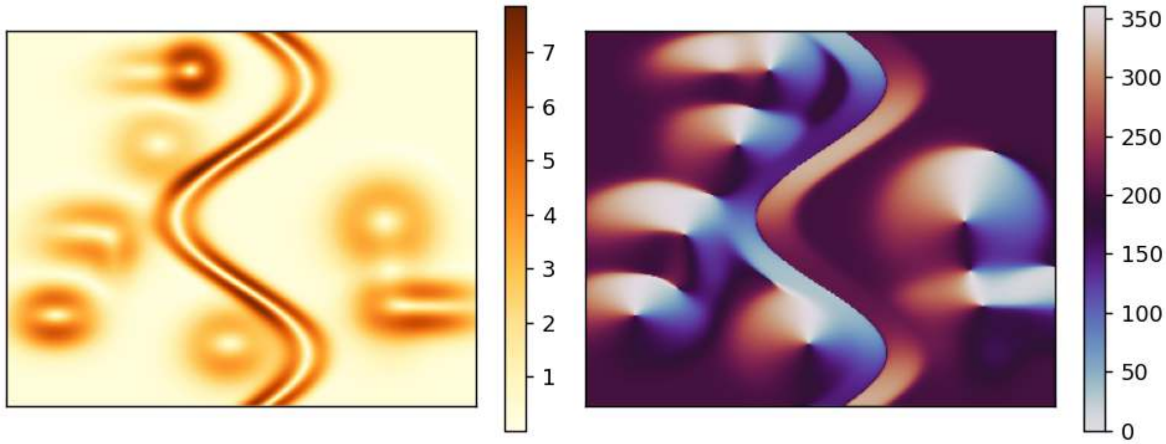

Figure 21: Slope (left) and aspect (right) of the demo relief. The cyclic aspect palette stitches 0 and 360 degrees.

2.09 Peaks

Finds peaks: cells that are the highest in a square window of the given radius, with a drop over the window minimum at or above the threshold.

Parameter	What it sets	Default / hint
Input DEM	Terrain raster.	-
Window radius, m	Suppression of secondary peaks within this radius.	500
Minimum drop, m	Cuts small bumps below the threshold above their surroundings.	20
Peaks (points)	Output: points with z, drop. The cure for tops above the last closed contour.	-

The two filters work as a pair. The **window radius** suppresses secondary tops next to the main one: of two peaks closer than the radius the higher one remains. The **minimum drop** suppresses bumps on a plain: a local maximum rising a meter above its surroundings does not count as a peak. Flat tops give a single peak rather than a scatter.

Output fields: **z** - the elevation, **drop** - the drop over the window minimum. The layer compares against OSM peaks from 2.02: matching the ele marks with the DEM elevations is one more quick data check.

2.10 Demo relief

A utility generator: synthetic terrain from a tilted plain, hills and a winding valley with a constant fall. The relief is deterministic by seed, and local depressions are left between the hills on purpose so the filling tool has something to show. All figures of this chapter are built on it.

Parameter	What it sets	Default / hint
Width, cells	Demo grid size along X.	300
Height, cells	Demo grid size along Y.	300
Cell size, m	Demo grid step.	30
Generator seed	Reproducibility of the random relief.	42
Output CRS (metric)	Coordinate system of the demo raster.	EPSG:32640
Compact int16 (adv.)	Integer output for demo shipping.	off
Demo relief	Output raster.	-

The tool exists for the manual examples, tests and offline work. Live data comes from 2.01. The **Compact int16** checkbox outputs the raster in whole meters for shipping demo fragments.

3.01 Categorical indicator kriging

The **Categorical indicator kriging** tool builds a probability map from a categorical field: mineral type, lithotype, any text class. Unlike ordinary kriging, which interpolates a number, here it estimates how likely each class is at every point of the area. This is what you need where the type matters rather than the magnitude: where to expect replacement, where the seam composition changes, where the boundary between varieties runs.

Parameters:

Parameter	What it sets	Default / advice
Point layer	Source points.	-
Categorical field (class)	The class field (mineral type, lithotype). Empty and NULL are excluded.	-
Search radius, min/max points, cell size, extent	Search and grid - as in “2D Kriging”.	as in “2D Kriging”
Class probabilities (multiband)	Raster: one band per class, the class name in the band description.	-
Zone map (most likely class)	Raster of the most-likely class code; the code mapping goes to the Log.	-
Confidence (max probability)	Raster of the maximum probability: where the class is firm, where it is contested.	optional

How it is computed

Coding the classes as numbers 1, 2, 3 and interpolating that code is not allowed. Categories have no order, class 3 is not “farther” than class 1, and a mean between them is meaningless. So the tool takes the indicator route. For each class an indicator is built: one where the borehole is of that class, zero everywhere else. Each indicator is kriged separately by ordinary kriging, like an ordinary field, and yields a surface from zero to one, which is the class probability. The indicator variogram is fitted automatically with a spherical model from the experimental one.

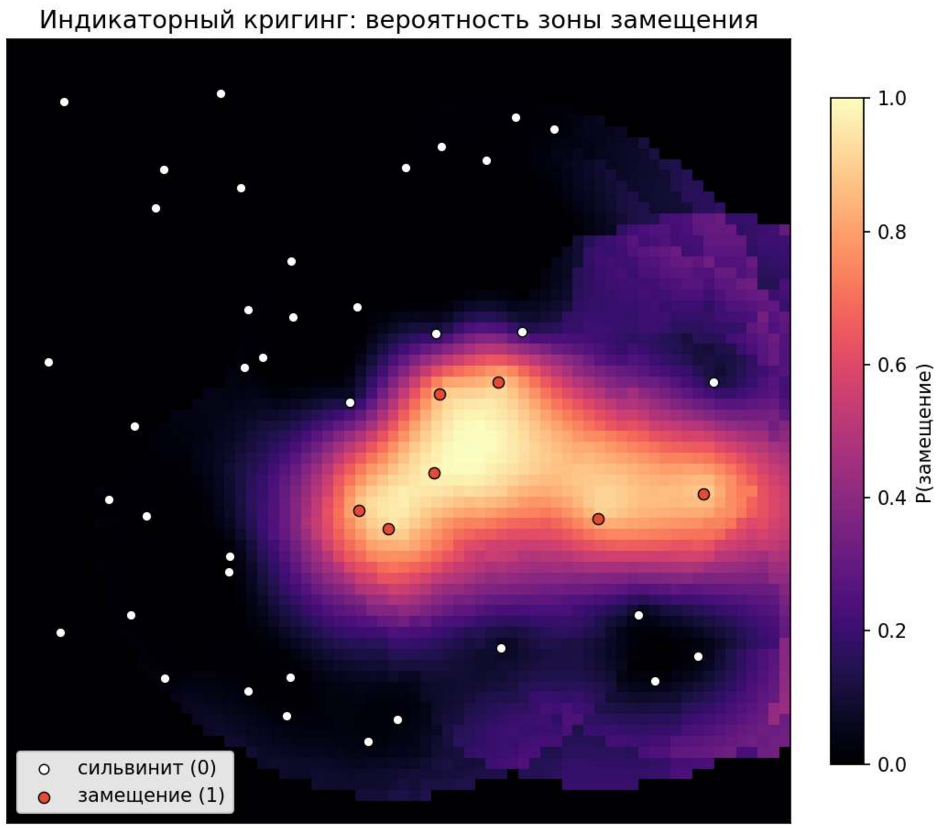

Figure 22: Indicator kriging on synthetics: categorised wells (red - replacement, white - sylvinite) turn into a class-probability map. A 0.5 threshold cuts the domain outline from it.

Separate indicators do not sum to exactly one and may go slightly out of range, a known property of the method. So the estimate of each class is clipped to zero-one, and then the class probabilities are normalised so that in every cell they sum to one.

What you get

Three results. A multiband probability raster, one band per class, the class name written into the band description. A zone map, the code of the most likely class in the cell, with the code to class mapping printed to the log. An optional confidence raster, the maximum probability in the cell, which shows where the class is firm and where zones compete and the boundary runs.

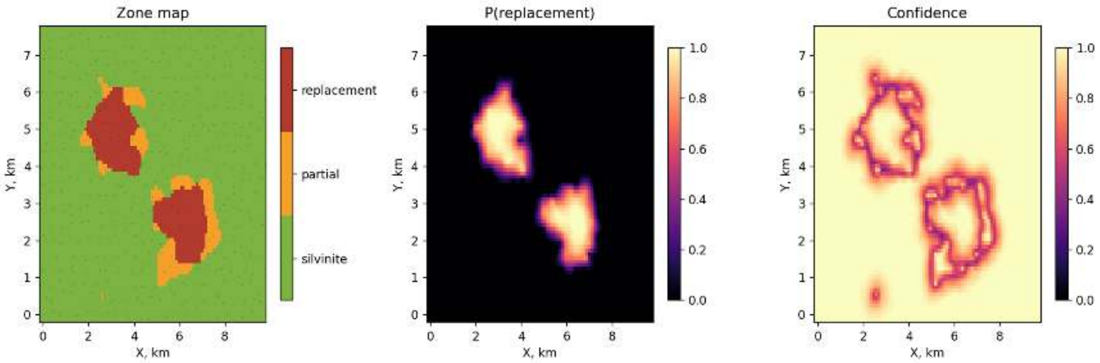

Figure 23: Categorical indicator kriging result: a map of the most likely mineral type, a silvinite background with replacement spots, boreholes drawn on top.

The categorical approach is convenient because it needs no boundary drawn in advance. There is no need to decide whether partial replacement counts as dangerous. All types are mapped as they are, and the required combination of classes is assembled later from the probabilities. Rare classes with few boreholes give a noisy variogram, the tool warns about this in the log, so read the probability of such a class with caution.

To learn the tool without real data, switch on **Add a categorical mineral-type field** in **Create sample wells (demo)**. A mintype field is added to the layer with a silvinite background and replacement spots after a mine, ready to run the tool on.

With an uneven network you can set the optional **wt** weight field from tool **1.01 Declustering**. Each class indicator is then kriged toward its declustered proportion rather than zero, so far from the data the probability tends to the representative class proportion. Without weights the behaviour is unchanged.

3.02 External Drift Kriging

The **External Drift Kriging** tool estimates a field from points when that field is systematically related to a quantity already known everywhere as a raster. Such a raster is called the drift. It can be the structural surface of an adjacent seam, a coarse regional model, a surface built on a sparse grid, or a seismic attribute. Ordinary kriging sees only the wells themselves, whereas here knowledge of the shape of the field between them is added, and the estimate leans on that shape where there are no wells.

The tool sits in the **Additional analysis tools** group and rests on the same engine as **2D Kriging**. The kriging mathematics does not change. What changes is only what the regional component is removed against.

Parameters:

Parameter	What it sets	Default / advice
Point layer	Source points.	-
Value field (Z)	The attribute being interpolated.	-
External drift raster	A secondary surface s known everywhere. Same CRS as the points, covers the area.	-
Drift raster band (Adv.)	The band of a multiband drift raster.	1

Parameter	What it sets	Default / advice
Drift degree	A linear (a0+a1·s) or quadratic relation.	1 (linear)
Kriging type, radius, min/max, nugget, structures	Kriging of the residuals - as in “2D Kriging”.	as in “2D Kriging”
Smooth grid (Gaussian), smoothing radius	Optional smoothing of the result.	off / 1
Drift kriging raster	The output estimate (drift + kriged residuals).	-
Kriging standard error	An optional raster of the residual standard error.	skipped

How it differs from trend removal

The **Remove polynomial trend** option of **2D Kriging** describes the regional component with a polynomial in the coordinates, that is with a tilted or curved plane. This works when the dip of the seam is uniform and its shape is simple. But if the field has a pronounced structure that follows a known surface, a plane will not describe it.

External drift removes the regional component not against the coordinates but against an external raster. If, for example, the roof of the seam of interest follows the relief of the underlying one, for which a surface already exists, that relation is removed by regressing on the underlying surface, and the departures from it are what gets kriged. The drift here is not a function of the position on the map but a function of the external raster value at the same point. Everything else matches trend removal. It is the same regression-kriging scheme.

How it is computed

First the drift raster is sampled at each well, bilinearly over the four neighbouring cells. Then the field value is regressed on this sampled value by least squares. The **Drift degree** sets the form of the relation. Degree 1 is the linear drift, value equals a0 plus a1 times the drift, the usual choice for external drift. Degree 2 describes a curved relation with the square of the drift, but it may absorb part of the real structure, so after using it you should look at the residual variogram.

Next the regression residuals are kriged, exactly like an ordinary field in **2D Kriging**, with their own variogram, search and anisotropy. At the last step the drift raster is resampled onto the kriging grid and the drift is added back to the kriged residual estimate. The final estimate in each cell equals the drift plus the kriged residual. Because the drift is known everywhere, between the wells the estimate is drawn not towards a local mean but towards the shape of the external surface.

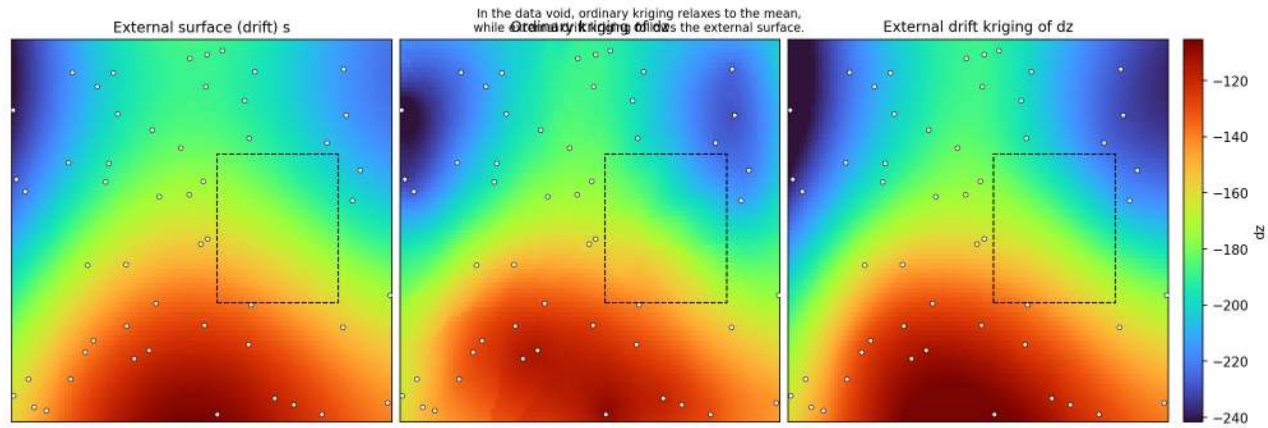

Figure 24: External drift kriging on sparse wells (computed with the Isoliner core). Left - the external surface s (drift). In the centre, ordinary kriging of dz relaxes to the mean in the data void. On the right, external drift kriging of dz follows the shape of the external surface where there are no wells. The dashed box marks the data void.

Wells that fall outside the drift raster do not enter the fit, and the tool reports to the Log how many were dropped. Grid cells not covered by the drift raster cannot be completed, so they are left empty together with the standard error in them.

Attributes of the generated layer

Field	Type	What it holds
well	text	Borehole name.
roof	number	Roof elevation, m.
thick	number	Bed thickness, m.
X	number	Useful-component content, percent.
head	number	Head (the hydrogeological data variant), m.
K	number	Hydraulic conductivity, m/day.
T	number	Transmissivity, m^2/day.
mintype	text	Mineral type (a class for indicator kriging).
dz	number	Elevation error - for weighting and experiments.

The field set covers all the plugin tools: interpolate roof and thick with ordinary kriging, X with a trend, mintype with indicator kriging, head/K/T with the hydrogeology tools.

Parameters

The **point layer** and the **value field Z** are set as in **2D Kriging**. The **External drift raster** parameter is the secondary surface known everywhere. The optional **drift raster band** selects the band of a multi-band raster. Search, cell size, extent, clipping to the well hull, the nugget and variogram structures, outlier removal and grid smoothing all work and are described as in **2D Kriging**, with the same defaults.

An important condition. The drift raster and the point layer must share the coordinate system, otherwise the drift value will be sampled at the wrong point. When the CRS does not match the tool warns in the Log. The drift raster must cover the whole estimation area, otherwise empty cells will appear along the edges.

The variogram on residuals

As with trend removal, the variogram here is fitted on the regression residuals, not on the raw value. After the drift is removed the residual variogram returns to its normal form, reaches a sill with a nugget, and the range reflects the true scale of the local correlation. The standard-error raster in this mode is the error of kriging the residuals. The drift is treated as deterministic and adds no error of its own.

A convenient way to fit the residual variogram without leaving the tool is not yet provided, so the residuals are judged by the share of variance removed, which the tool prints to the Log. If the drift took out a noticeable part of the spread, the relation with the external surface is real and the drift is appropriate. If it took out almost nothing, the field is not related to that raster, and plain **2D Kriging** will give the same result more simply.

3.03 Exceedance probability map

The **Exceedance probability map** tool answers not "how much" but "how likely the value exceeds a threshold". From the kriging estimate raster and its standard-error raster it builds a probability raster from 0 to 1: in each cell the probability that the true value is above a given threshold.

The tool sits in the **Additional analysis tools** group and works as a post-processing step, like the hydraulic gradient. It runs no kriging of its own and does not touch the **2D Kriging** window, it takes ready rasters. So it works equally with the output of ordinary kriging and of external drift kriging.

How it is computed

Kriging gives, in each cell, an estimate and its standard error. If the local distribution of the value is taken as normal, that is the value in the cell is treated as normal with the mean equal to the estimate and the standard deviation equal to the kriging error, the exceedance probability is one formula through the normal distribution function. Where the estimate is well above the threshold the probability is close to one, where it is below it is close to zero, and at the threshold itself it equals one half. The larger the standard error, the smoother the transition: away from the wells there is less certainty and the probability is drawn towards 0.5.

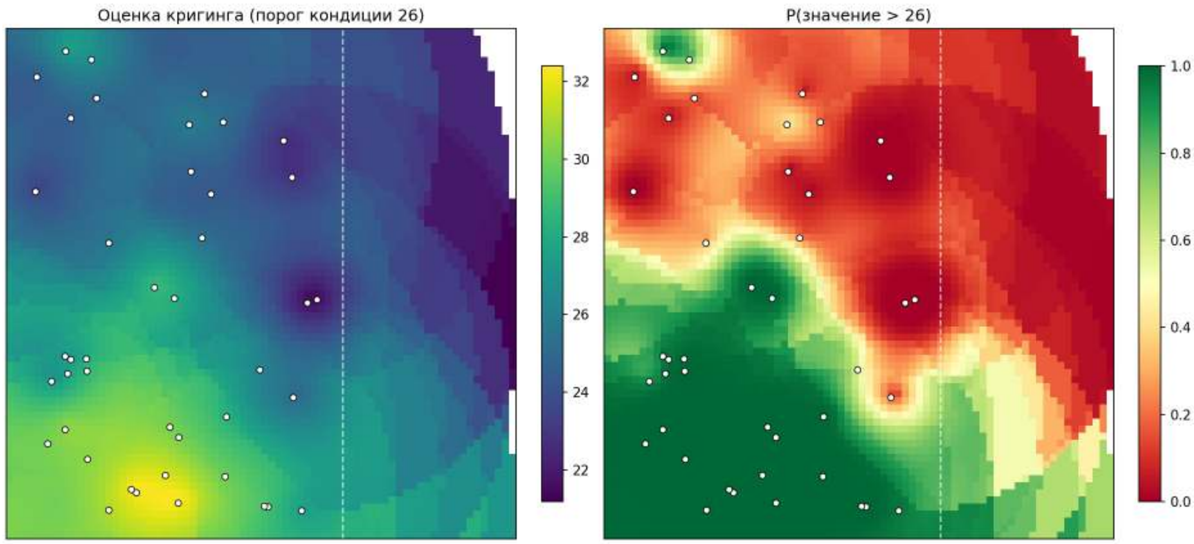

Figure 25: The kriging estimate with a cut-off threshold on the left, the exceedance-probability map on the right. Green - take confidently, red - confidently do not, beyond the drilling boundary the map converges to 0.5.

No separate kriging is needed for this, so the map is built instantly. The normality assumption is rough in places, especially for strongly skewed fields such as grades with a long right tail. Where that matters, indicator kriging by thresholds, which does not rely on the shape of the distribution, is more accurate.

How to get the inputs

Run **2D Kriging** (or **External Drift Kriging**) on your field and enable the optional **Kriging standard error** output. You get two rasters, the estimate and the error, and you feed them here. Their grids match, since they come from one run, but if rasters with different grids are supplied, the error is resampled onto the estimate grid bilinearly.

Parameters

Parameter	What it sets	Default / advice
Estimate raster (kriging)	The field estimate raster (a kriging result).	-
Kriging standard-error raster	The standard-error raster of the same run.	-
Side	Probability above the threshold $P(Z>t)$ or below $P(Z<t)$.	above
Threshold	The value the probability is computed against.	0
Estimate raster band, error raster band (Adv.)	Bands of multiband rasters.	1
Probability raster (0…1)	The output probability raster.	-

Use

Cut-off grades: the threshold is the cut-off, and the map shows the probability that the grade is above the cut-off. This is more honest than a single line drawn on the estimate, because near the edge of the ore body the certainty drops and the probability map shows it. Risk zones for any threshold: thickness below a critical value, an elevation above or below a hazardous one. The probability map complements the estimate map where not only the value matters but the confidence in it.

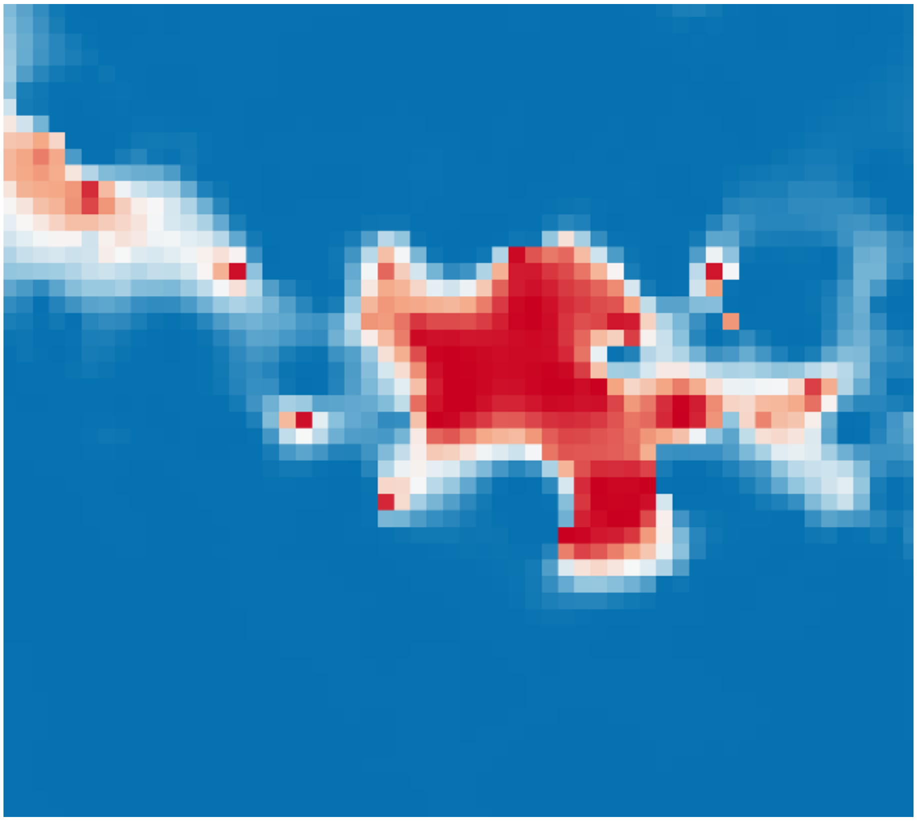

Figure 26: An exceedance probability map with a diverging colour ramp broken at 0.5. Red is where the value is confidently above the threshold, blue confidently below, and the white band along the P=0.5 line is the zone of uncertainty (contested values). The further from the wells, the wider the band.

3.04 Hydraulic gradient and flow direction

The **Hydraulic gradient and flow direction** tool works with the head field, that is the piezometric surface, and shows where and how steeply groundwater flows. The input is a head raster, usually the result of **2D Kriging** on borehole water levels. For a hydrogeologist this is as natural a step after building the head surface as isolines are after kriging.

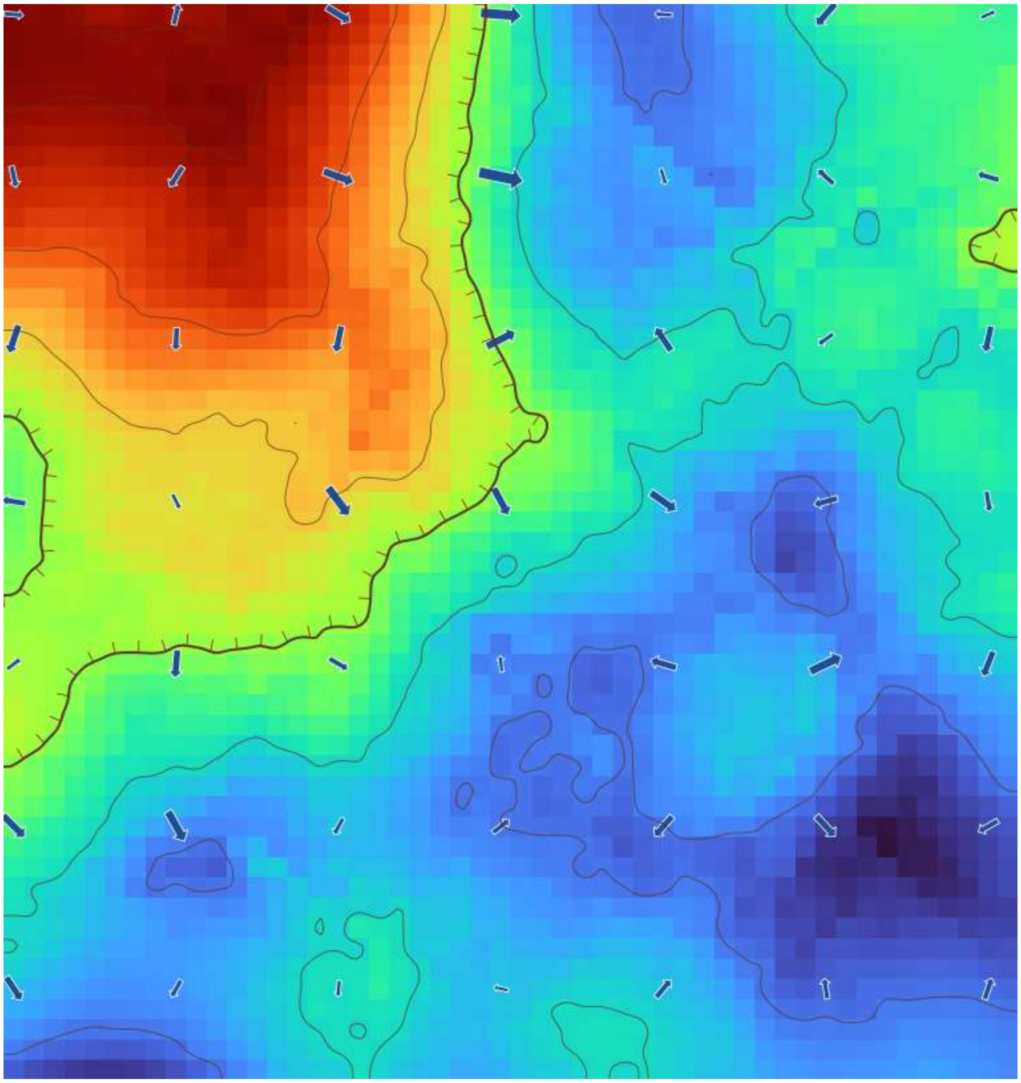

Figure 27: Flow vectors over the head surface: the arrows go down-gradient, from high head (warm tones) to low (cool), on top of the head isolines. The arrow length grows with the steepness of the gradient.

There are three outputs. The **gradient-magnitude raster** shows the steepness of the head surface, the hydraulic gradient i equals the magnitude of ∇h and is dimensionless. The **azimuth raster** holds the flow direction in degrees, where zero is north and the count goes clockwise. The point layer of **flow vectors** is thinned over the grid and styled as arrows right away, so the flow pattern is visible without touching the symbology.

The direction is computed strictly. Water flows down-gradient, from higher head to lower, so the arrow points towards the falling surface. On flat areas, where the head is almost constant, the direction is undefined and the azimuth there is left empty.

Without permeability

The tool describes the geometry of the head field, not the flow velocity. The Darcy filtration velocity equals minus the hydraulic conductivity K times the gradient, and it needs K itself, which the tool neither asks for nor computes. In other words, the map answers where and how steeply, but not how fast. Once K (or transmissivity T) is available over the area, the specific discharge and the flow are computed by the neighbouring tool **Specific discharge (Darcy law)**, which multiplies this gradient by the aquifer properties.

Parameters and smoothing

The input is the **head raster** and its **band**. The **flow vectors, thinning step** parameter sets how many cells apart to place an arrow so they do not merge, eight by default. The **smooth head before computing** parameter removes fine grid ripple, set in cells, off by default.

Parameter	What it sets	Default / advice
Head raster	The input piezometric surface (usually a "2D Kriging" result).	-
Band (Adv.)	The band of the input raster.	1
Smooth head before computing, cells	Damps grid noise before differentiation. 0 = none.	0
Flow vectors: thinning step, cells	How many cells apart to place an arrow.	8
Hydraulic gradient (magnitude)	The output $\lvert\nabla h\rvert$ raster.	-
Flow direction (azimuth)	The output azimuth raster (down-gradient).	created by default
Flow vectors (points)	The point layer of arrows (fields az, grad).	created by default

Smoothing is switched on for substance, not for looks. Differentiation amplifies noise, so even a clean kriging grid can give a patchy gradient field with jittery arrows. A light smoothing brings the picture back to a readable form. The same effect can be had by smoothing the head itself back in **2D Kriging**.

Arrows from points

The vector layer is points, and the arrows are drawn by the symbology. The preset is applied automatically. The arrow marker is rotated by the **az** field, so it shows the flow direction, and its size is scaled by the **grad** field, so the arrow is longer where the gradient is steeper. The size is set in millimetres and does not depend on the map scale. The symbology can be changed in the layer properties. If you need a classic quiver diagram, where the arrow length is laid out in map units, the marker is replaced with a geometry generator, the recipe is in the styles folder next to the preset.

Fields of the flow-vector layer:

Field	Type	Holds
az	number	Flow-direction azimuth, degrees (0 = north, clockwise, down-gradient).
grad	number	Magnitude of the hydraulic gradient $\lvert\nabla h\rvert$ at the point, dimensionless.

The learning cycle

To walk the whole path without real data, switch on **Add a head field** in **Create sample wells (demo)**. A head field with a pronounced regional slope is added to the layer. Build a grid from it in

2D Kriging, feed the raster here, and the arrows follow the head downhill. The same end-to-end scenario as for the other tools, only about hydrogeology.

3.05 Specific discharge (Darcy law)

The **Specific discharge** tool adds permeability to the flow geometry. The hydraulic gradient shows where and how steeply the head falls, but not how much water flows. Darcy's law links these through the aquifer properties: the higher the permeability and the steeper the gradient, the larger the flux. From a head raster and aquifer-property rasters the tool builds a physical flux rather than a dimensionless gradient.

The tool sits in the **Additional analysis tools** group and works as a post-processing step. It runs no kriging of its own: the property rasters are prepared separately by kriging from test points.

What is computed

The specific discharge (Darcy flux) equals the hydraulic conductivity times the hydraulic gradient: $q = K \cdot |\nabla h|$, in metres per day. It is the volume of water through a unit cross-section area per unit time. If a transmissivity raster is supplied instead of conductivity, the tool computes the flow per unit width of the flow $Q = T \cdot |\nabla h|$, in square metres per day. Transmissivity is conductivity times thickness, so the flow per width already accounts for the aquifer thickness and does not need it separately. The direction of both fluxes is the same as the gradient direction, down the head slope.

The true water velocity is the specific discharge divided by the effective porosity, $v = q/n$. Porosity is usually absent from the data, so the tool does not ask for it and does not compute the true velocity: if needed, divide the q raster by the porosity in the raster calculator.

How to get the K and T rasters

The aquifer properties are known at the test points (pumping, injection) but are needed everywhere. They are interpolated by kriging, like any field. An important subtlety: hydraulic conductivity and transmissivity are almost always log-normal, their values span orders of magnitude. Kriging the raw values distorts the result, so the logarithm is kriged. The simplest way is to enable the **ln** transform in **2D Kriging**: then ln is kriged and the raster is returned already in the original units, and the ln checkbox here is not needed. If instead you krige an already-logged field, tick **K and T rasters are given as ln** in this tool and the values are recovered by exponentiation. Confined and unconfined aquifers are better kriged separately, their thickness physics differs.

Parameters

Parameter	What it sets	Default / advice
Head raster	The piezometric surface (a kriging result over the levels).	-
Hydraulic conductivity raster K (m/day)	The aquifer property for the specific discharge. Optional, but at least one of K, T is needed.	-
Transmissivity raster T (m^2/day)	The aquifer property for the flow per width.	-
K and T rasters are given as ln	Apply exp to the input rasters (for log-kriged K and T).	off
Smooth head, cells	Damps head noise before differentiation.	0

Parameter	What it sets	Default / advice
Flow vectors: thinning step	How many cells apart to place an arrow.	8
Raster bands (Adv.)	Bands of the multiband head, K, T rasters.	1
Specific discharge q (m/day)	The output raster $q = K \cdot \lvert \nabla h \rvert$.	created if K is given
Flow per width Q (m^2/day)	The output raster $Q = T \cdot \lvert \nabla h \rvert$.	created if T is given
Flow direction (azimuth)	The output azimuth raster.	optional
Flow vectors (points)	The arrow layer (rotated by az, sized by the specific discharge).	created by default

Use

Where water moves faster and where slower, estimating inflows to workings, zones of higher seepage along permeable beds. Together with the exceedance probability map you can show not only the expected flux but also the confidence in it where test points are sparse.

3.06 Gaussian simulation (SGS)

Kriging gives a single smoothed surface and an estimation variance. Sequential Gaussian simulation answers a different question - how large is the uncertainty. It builds an ensemble of equally probable realizations: each one reproduces the data histogram and variogram, passes through the boreholes and therefore stays rough rather than smoothed. Across the realizations every node accumulates a distribution of values, which shows where the estimate is reliable and where the data are silent.

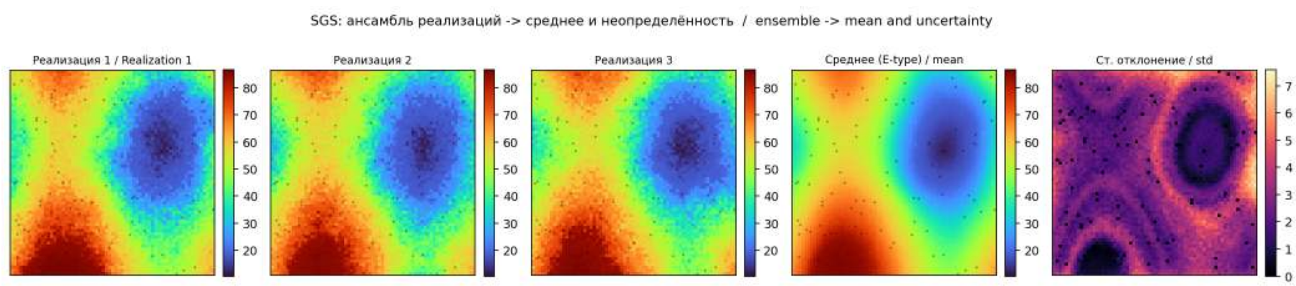

Figure 28: An SGS ensemble of realizations and the mean and uncertainty derived from it.

How it works. The values are mapped to normal scores and the simulation runs in Gaussian space. The grid nodes are visited in random order; at each node simple kriging on the neighbours and already-simulated points gives a local mean and variance, a value is drawn from that normal distribution and immediately becomes conditioning for the next nodes. Boreholes are snapped to the nearest nodes and frozen across all realizations. At the end each realization is back-transformed to the original units. The normal-score variogram is fitted automatically with a sill close to one.

Parameters

Parameter	What it sets	Default / hint
Points (boreholes)	The point layer with the data.	-

Parameter	What it sets	Default / hint
Value field	The numeric attribute to simulate.	-
Cell size	The grid step. 0 means auto, min(extent)/50.	0
Number of realizations	How many runs to average.	60 (50-100)
Cut-off threshold	Enables the exceedance-probability map.	not set
Probability ABOVE the threshold	The probability direction.	yes
Raster extent	The output rectangle.	by layer
Score variogram model (Adv.)	The model for the normal scores.	auto
Max neighbours per node (Adv.)	How many points in the node's simple kriging.	16
Search radius (Adv.)	0 means auto, 3 variogram ranges.	0
RNG seed (Adv.)	For reproducibility. 0 means random.	0

The outputs are ensemble rasters. **Mean (E-type)** resembles kriging. **Standard deviation** shows the uncertainty, small at the boreholes and large away from them. The **P10**, **P50**, **P90** quantiles outline the likely range, and **Exceedance probability** for a given threshold offers a non-parametric alternative to the map from the probability tool. Runtime grows with grid size and the number of realizations, so start with a coarse cell.

If the data are clustered unevenly, supply the **wt** weight field from tool **1.01 Declustering**. The normal-score transform then builds the distribution with weights, and the ensemble histogram is not skewed toward over-sampled rich areas.

Kriging kinds: which one to pick

Behind the word "kriging" the plugin hosts a family of methods, and the choice between them affects the result more than fine-tuning the variogram. All the kinds solve the same system of equations with covariances from the variogram; they differ in what is assumed known about the field mean and in what exactly is estimated - a point, a block or a probability. This chapter is a navigator; the parameters of each tool live in their own chapters.

Simple kriging (SK) assumes the mean of the field is known in advance and constant over the area. Near the wells the estimate follows the data, away from them it is pulled to the given mean. Take it when the mean is backed by statistics over a representative sample of the same domain, an eyeballed mean drags all the underdrilled margins towards the error. Switched by the type in **2D Kriging**.

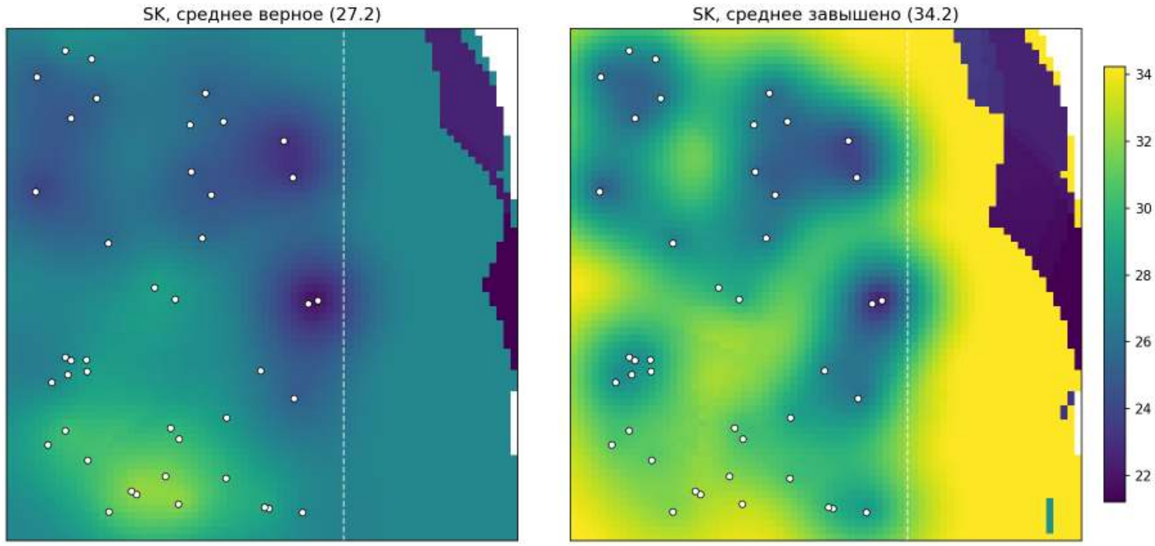

Figure 29: Simple kriging: the mean set from the data on the left, inflated by seven on the right. There are no wells east of the dashed line, and the whole underdrilled east “floats up” to the false mean.

Ordinary kriging (OK) does not know the mean and estimates it locally in every neighbourhood - an extra equation with the “weights sum to one” condition takes care of that. Away from the wells the estimate tends to the mean of the nearest neighbourhood, not to the global one. This is the default choice: if unsure where to start - start with OK.

Kriging with a trend (the detrend checkbox in **2D Kriging**) is for fields with a regular regional slope: a roof on a monocline, a fold limb. A 1st- or 2nd-degree polynomial is removed by least squares, the residuals are kriged, the trend is added back. Two rules: define the variogram over the residuals (the plugin prints the share of the removed variance - if it is small, the trend is not needed), and do not extrapolate a quadratic trend far beyond the well cloud.

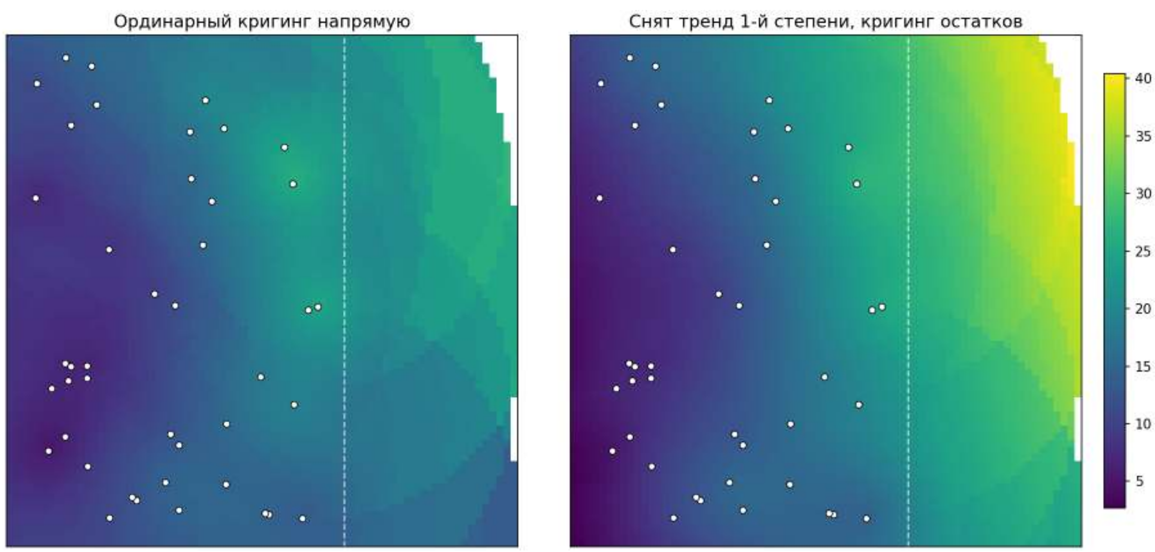

Figure 30: A field with a regional slope: plain ordinary kriging stalls at the local mean beyond the wells, regression kriging continues the slope regularly.

Kriging with an external drift (chapter 3.02) - when the trend is known not as a formula but as a field: a structural surface of a neighbouring bed, a regional model, a seismic attribute. The scheme is the same - a regression on the drift, kriging of the residuals, the regression returned.

Block kriging (the discretisation parameter in **2D Kriging**) estimates the mean over a block rather than a point value: the right-hand side of the system is averaged over the discretisation, the error variance drops, outliers are damped. Take it for reserves over a block grid and mind the support effect: a block-kriging grid is regularly smoother than a point one, a sample grade and a block grade cannot be compared directly.

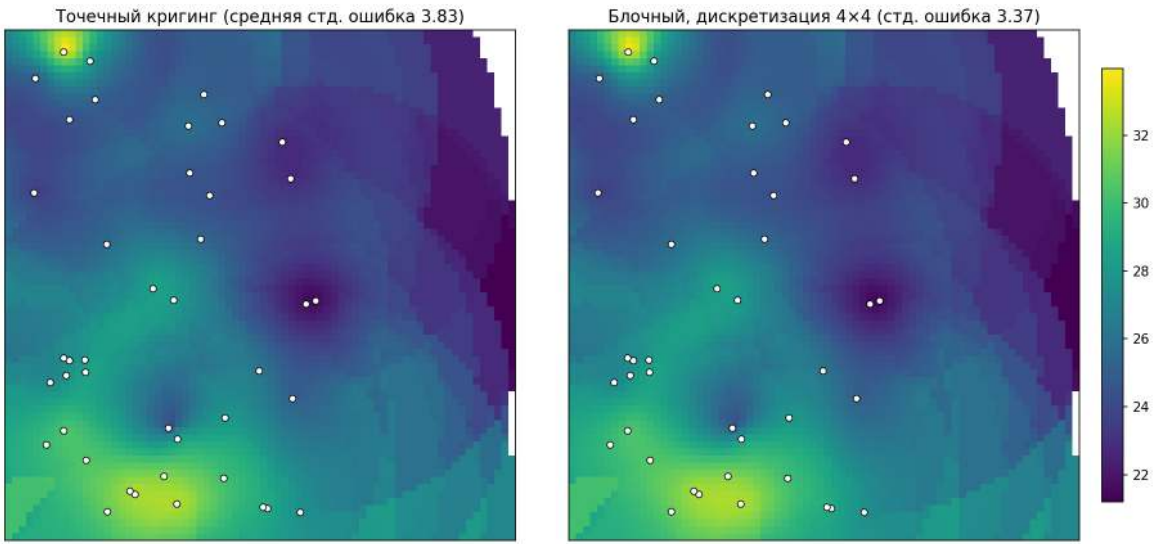

Figure 31: The same wells with two deliberate outliers: the cones on the block grid are damped, the mean standard error is lower.

Indicator kriging (chapter 3.01) is for categories: mineral type, facies, a replacement zone. The category becomes a 0/1 indicator, it is kriged with plain OK, the result is the class probability at a point, domains are cut from it by a threshold. The indicator variogram is its own and usually shorter than the grade one.

Gaussian simulation (chapter 3.06) is not kriging but its complement: instead of one smooth surface, an ensemble of equally probable rough realisations from which the uncertainty is seen directly.

Cheat sheet

Task	Kind	Where it lives
The universal case, the start of any task	Ordinary (OK)	2D Kriging, the default type
Plenty of data, the domain mean is justified	Simple (SK)	2D Kriging, the SK type + mean
A roof or a bottom with a regional slope	With a trend	2D Kriging, detrend
The trend is known as a raster	With an external drift	chapter 3.02
Reserves over a block grid	Block	2D Kriging, discretisation
Mineral type, replacement, categories	Indicator	chapter 3.01
Uncertainty assessment	SGS simulation	chapter 3.06

The search neighbourhood is common to all the kinds, and three rules remove most problems: the radius of the order of the variogram range, 12-16 neighbours at most, the neighbourhood anisotropy consistent with the variogram anisotropy from the variogram map.

3.07 Density from measurements (variable support)

The tool builds a density map where a measurement is given not by a point but by a finite-size support: a point with an uncertainty sigma, a line segment (a corridor of half-width) or a polygon. The unit mass of a measurement is spread over its support. Mass is conserved and density is inverse to the support area, so coarse georeferencing self-attenuates geometrically, without thresholds or filters. This is density estimation (how much and where), not value interpolation - kriging remains for values.

One geometry type per run. Points, lines and polygons are mixed by a series of runs into one raster (append mode). Each type spreads its mass in its own way:

- **Point** - a Gaussian spot with a sigma from the precision field, truncated at three sigmas. A sigma below the half-cell is raised to the half-cell.
- **Line** - a soft-edged corridor: the polyline is densified, mass is split by length, each subpoint is a Gaussian profile with the half-width sigma. The from_m/to_m fields cut an interval by linear referencing.
- **Polygon** - mass is split by area uniformly or, in dasymetric mode, proportionally to an auxiliary raster (population, built-up area). If the raster is empty inside the polygon, it falls back to uniform.

Output and invariant

The main output is a three-band raster. Band 1 is density in mass per km2 (independent of cell size). Bands 2 and 3 are service (sum m*sigma and sum m) so that append series and the effective-sigma map stay exact. The optional second raster is the mass-weighted sigma per cell, an effective-precision map: it separates density backed by precise georeferences from the smeared one. This is an analogue of kriging variance for the density floor.

Invariant: the density integral over the raster equals the sum of input masses. It is always computed and written to the log. A discrepancy means supports left the area; the behaviour is set by the edge switch (renormalise inside or lose mass with a warning).

How to read

Density shows where measurements cluster, weighted by their reliability. A precise georeference gives a compact spot, a coarse one a diffuse and low one. The effective-sigma map shows where density is gathered from precise supports and where from smeared ones - there the trust is lower.

Parameter	Purpose	Default
Measurements	Points, lines or polygons (one type).	-
Mass field	Object mass.	1
Precision field	Point sigma or line half-width.	default sigma
Cell size, m	Grid step.	50
Area	Extent.	by layer
Support beyond edge	Renormalise inside or lose mass.	renormalise
Default sigma (Adv.)	When the precision field is empty.	half-cell
from_m / to_m (Adv.)	Line interval cut.	whole line

Parameter	Purpose	Default
Auxiliary raster (Adv.)	Dasymetry for polygons.	-
Append to raster (Adv.)	An existing three-band raster.	-
Density	Three-band raster (density, sum m*sigma, sum m).	-
Effective sigma	Effective-precision map (optional).	-

The “Density map” window (live preview)

The face of the tool is a separate **Density map** window on the **Isoliner** toolbar. The layer and fields are set on the left, while a preview on a coarse grid runs on the right: it takes milliseconds, so the sigma and the cell size change the picture at once rather than after a run. The invariant is always visible at the bottom: input mass, mass on the grid and the share lost at the edge.

The **Demo** button creates a training set with tool 3.08. The **Write raster** button runs the full computation with the same algorithm 3.07 and puts the result on the map already dressed: pseudocolour with transparent zeros, density isolines and an effective-sigma layer as a trust map. The Processing form remains for models and batch runs; the computation core is shared.

3.08 Create a density example (demo)

The tool creates a synthetic set for 3.07 with a known total mass, to check the invariant by eye. Ten points with different sigmas (from fractions of a cell to large, mass 500), two lines (mass 200, one with a from_m/to_m interval cut), two polygons (mass 300, one for dasymetry) and an auxiliary raster. Total mass 1000.

Run 3.07 on the point layer - the density integral in the log should give 500, on lines 200, on polygons 300. Layer fields: **mass**, **prec**, **from_m**, **to_m**.

Parameter	Purpose	Default
Extent	Generation bounds.	-
Auxiliary raster cell, m	Step of the auxiliary raster.	50
RNG seed (Adv.)	Example reproducibility.	1
Demo points / lines / polygons	Three measurement layers.	-
Auxiliary raster	Raster for dasymetry.	-

The full list of output-layer fields is in the **Density (demo)** appendix section (“Demo-layer fields” at the end of the manual).

4.01 Cross-section along a line

The **Cross-section along a line** tool builds a geological section from a set of surfaces. It is not just a profile curve but beds as filled bands between a roof and a floor. The surfaces are usually obtained by kriging, and the tool assembles them into a section along a given line.

The tool sits in the **Cross-sections** group and works as a post-processing step over ready rasters. It runs no kriging of its own.

How beds are defined

The surfaces are supplied as a list and ordered top to bottom: roof, floor, then the next roof, and so on. Beds are built as bands between adjacent surfaces, so N surfaces give N minus one beds. Two surfaces, a roof and a floor, are enough for one bed. For a sequence of beds, add the surfaces in stratigraphic order.

Two outputs

The section drawing is polygons in axes of distance along the line and elevation. The elevation can be stretched by a vertical exaggeration so thin beds read well. This layer goes into a print layout as a ready section. Its coordinate system is conventional, with distance and elevation in map units.

The 3D fence is the same bands but as vertical PolygonZ walls in real coordinates. They are viewed in the 3D Map View next to the kriging surfaces: the grid is set as terrain, and the section walls show the beds in space.

Vertical scale

The horizontal extent of a section (the line length) and the vertical extent (tens of metres of beds) are not comparable, so without a vertical stretch the drawing looks flat. The scale is set in two ways. In the **H:V ratio** mode you set the desired width:height ratio of the drawing (say 10), and the tool computes the exaggeration itself from the line length and the elevation span. In the **exaggeration** mode the value is a direct vertical stretch factor.

The effective exaggeration is printed to the log. For an exact overlay of layers it must match across the section, the boreholes and the composition. In H:V mode the section (4.01) and the boreholes (4.02) span the whole section in height and line up. The composition (4.03) computes the ratio over a single bed, so to overlay it take the exaggeration printed by 4.01 and set it in 4.03 in the **exaggeration** mode.

Attributes of the output layers

Section definition - a single line with the original trace geometry and the fields read by all the downstream tools of the group:

Field	What it holds
vex	Vertical exaggeration of the drawing.
step	Stationing step, m (the polyline vertices are always included).
zmin, zmax	The elevation range of the drawing.

Section (3D fence) - vertical PolygonZ polygons along the trace, one per bed:

Field	What it holds
bed	Bed number from top to bottom.
top, bot	Names of the roof and bottom surfaces.
t_mean	Mean bed thickness along the trace, m.
seclen	Trace length, m.

Section corner points: num (corner number), name (УГ-1, УГ-2, …), pos (top or bottom), d (station, m), x and y (map coordinates), az (azimuth of the next leg), label (a ready-made label).

Horizontal axes: elev (axis elevation, m) and label. **Corner table**: kind (row type) and text (cell content).

Parameters

Parameter	What it sets	Default / advice
Section line	A line layer. The first line is used.	-
Surfaces top to bottom	A list of surface rasters in stratigraphic order. At least two are needed.	-
Sampling step along the line	How many map units between samples. 0 means by cell size.	0
Vertical scale	Mode: H:V ratio or exaggeration.	H:V
Scale value	Width:height ratio (e.g. 10) or exaggeration.	10
Raster sampling (Adv.)	Bilinear or nearest.	bilinear
Section drawing (distance × elevation)	The output polygon layer for a layout.	created
3D fence (PolygonZ)	The output layer of vertical walls in real coordinates.	created

Each bed gets attributes: a number, the roof and floor names, the mean thickness and the section length. Colour the layer by bed number or by thickness. Where a surface is undefined (nodata), the band breaks and the bed splits into several polygons.

Trying it on a demo

A ready training set is produced by the **Create a section example** tool (4.10): six surfaces top to bottom, the line, boreholes with the h1…h6 fields, and multiband bed grids. Run it, then feed the surfaces and the line here, the boreholes into **Boreholes on the section**, and the bed grid (bands 1/2/3) into **Bed composition on the section**. The full contents of the set are in section 4.10.

Relation to QGIS

A plain profile curve over a single grid is built by the native **Elevation Profile** panel, no separate tool is needed for that. The section instead shows the beds between surfaces, which the native tools do not do. A kriging surface can also be viewed in 3D without a section: set the grid as terrain in the 3D Map View.

4.02 Boreholes on the section

The **Boreholes on the section** tool projects boreholes onto the section line and shows them as columns of bed intervals on top of the drawing from **Cross-section along a line**. It sits in the **Cross-sections** group.

Each borehole is placed at the distance along the line where its projection falls. The bed boundaries are taken from the chosen elevation fields: on each borehole their values are sorted in descending order, and adjacent pairs give the bed intervals. So the order of field selection and gaps (NULL) do not matter. Each interval gets a bed number, and the column gets the borehole number from the label field.

Corridor and exaggeration

The corridor is a buffer around the line: boreholes farther than it are not shown (0 shows all). Set the vertical scale the same as in **Cross-section along a line** - in H:V mode the columns line up with the bands automatically, or take the exaggeration printed by 4.01.

Parameters

Parameter	What it sets	Default
Section line	The same line as for the section.	-
Boreholes	A borehole point layer.	-
Bed-boundary elevation fields	Numeric roof and floor fields. At least two.	-
Borehole number field	The column label.	no label
Corridor from the line	A buffer, map units. 0 shows all.	0
Vertical scale	Mode: H:V ratio or exaggeration.	H:V
Scale value	H:V ratio or exaggeration.	10
Borehole bed intervals	The output vertical segments (drawing).	created
Borehole collars	Points at the top of the columns for labels.	created

Colour the intervals by bed number to match the section bands, and label the collars by borehole number.

4.03 Bed composition on the section

The **Bed composition on the section** tool colours the band of one bed by a composition grid along the line. It takes a roof, a floor and a composition grid, runs no kriging of its own, and works one bed at a time. It sits in the **Cross-sections** group.

This is how the lithological composition change inside an industrial bed is shown along the section. The composition grid is prepared separately: the content by ordinary kriging, the mineral type by indicator kriging (the **Categorical indicator kriging** tool).

Two modes

Continuous content (KCl, insoluble residue): the band is cut into thin vertical slices, each with a mean value. Set a graduated style for the layer (by the **value** field), and a smooth content transition is visible along the band.

Categorical mineral type or facies (sylvinite, replacement, halite): adjacent slices of the same class merge into facies zones. Set a categorized style (by the **class** field). Replacement zones show as a colour change along the line.

Parameters

Parameter	What it sets	Default
Section line	The same line as for the section.	-
Bed roof	The roof raster.	-
Bed floor	The floor raster.	-
1st industrial bed	A multiband grid: roof, bottom, content, mineral type.	on request
2nd industrial bed	The same for the second industrial bed, independent fields.	on request
Sampling step along the line	How many units between samples. 0 means by cell.	0
Vertical scale	Mode: H:V ratio or exaggeration.	H:V
Scale value	H:V ratio or exaggeration.	10
Raster sampling (Adv.)	Bilinear or nearest (always nearest for a class).	bilinear
Bed composition (drawing)	Output polygons in distance × elevation axes.	created
Bed composition (3D)	PolygonZ polygons in real coordinates.	on request

Run the tool for each industrial bed separately, with its own composition grid. Place the composition band on top of the section drawing. For an exact overlay take the exaggeration printed by **Cross-section along a line** and set it here in the **exaggeration** mode (the H:V ratio is computed over a single bed and is not suitable for overlay).

The section definition and shared parameters

Geometrically a section is set by two things - a line in the real coordinate system and a vertical scale vex. The **Cross-section along a line** tool outputs them together as a **Section definition** layer: one line with vex and step fields. This is the shared source of truth.

The intersect, project and unproject tools read the line and vex from this definition, so their results match the section without manual scale fitting. Build the section once, the definition travels with the project and feeds the other tools of the group.

The **Boreholes on the section** and **Bed composition on the section** tools also accept the section definition as an optional input: when given, the vertical scale is taken from it, so the borehole columns and the composition band sit exactly on the beds by height.

The section also clips pinch-outs: where the roof drops to the floor, the bed disappears and no band is built. In the demo the second industrial bed pinches out to the east.

For a polyline the Cross-section along a line tool optionally outputs three helper layers in the drawing axes. Corner points are placed at every polyline node, at the top and at the bottom of the section. A point carries fields: number, name (УГ-1, УГ-2 …), side (top or bottom), distance along the line, plan X and Y, segment azimuth and a ready label. The top is labelled with the name, the bottom with the plan coordinates X and Y, rounded to two decimals. The azimuth and distance stay as layer fields - handy to place into a layout table. A style is supplied: an upward triangle on top, a shelf at the bottom.

A corner table is produced optionally - a polygon layer below the section. The cells lie between the corner verticals with borders under them, two rows: the length and azimuth of the segment between adjacent corners, with a centred label and a white fill. It renders on the canvas and travels into a layout with the section. Corner verticals are lines at the nodes spanning the full section height. Horizontal axes are equal-elevation lines with ticks (five by default, with nice rounding) for an elevation scale. The drawing margins are extended by five percent up and down, and the corner points sit on these edges.

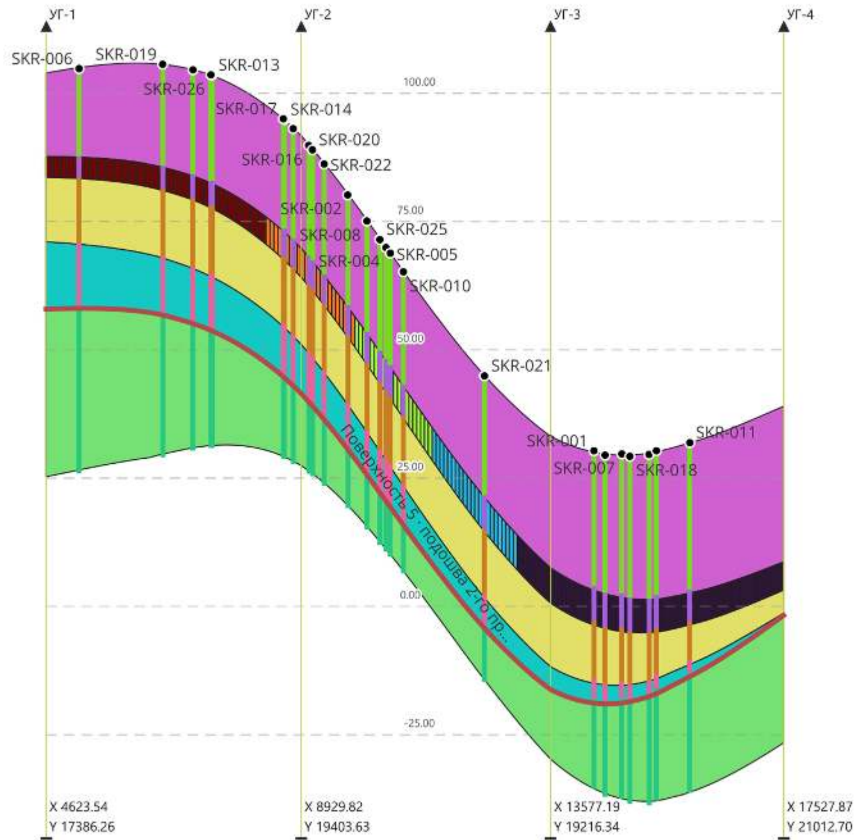

d	4755.39	4651.14	4339.91
Аз	64.90	92.31	65.55

Figure 32: Section decoration: the frame with corner verticals and triangles, horizontal axes with ticks on the left, and the corner table below.

4.04 Intersect surfaces with the section

The **Intersect surfaces with the section** tool places surface grids onto the section as lines in distance-elevation axes. Each grid is sampled along the definition line, and its trace lies on the drawing next to the beds. The line and vex come from the section definition, so the match with the section is automatic.

This is how water tables, marker surfaces, the salt roof and anomaly surfaces are placed on the section. The inputs are the section definition and a list of grids, the output is lines in the section axes (and optionally 3D lines in real coordinates).

The object-projection, unprojection and shaft-unwrap tools are marked **(beta)**: they work, but their interface and example set are still being refined.

Parameters

Parameter	What it sets	Default / hint
Section definition	The definition layer from 4.01: a line with vex and step fields. Sets the line, the scale and the frame height.	-
Surface grids	The list of rasters whose traces are drawn on the section as lines.	-
Sampling step along the line (Adv.)	How many map units between samples. 0 means by cell size.	0
Raster sampling (Adv.)	Bilinear or nearest.	bilinear
Surface lines on the section (drawing)	The output lines in distance × elevation axes.	created
Surface lines (3D)	The output 3D lines in real coordinates.	on request

4.05 Vector intersection with the section

While 4.04 places surfaces as grids, this tool places **vector** objects on the section by exact intersection with the section line. The result type depends on the object.

A line **without an elevation** (flat in plan - a fault, a boundary, a contour) gives a **full-height vertical** at the crossing station. Where the section crosses it horizontally is known, the depth is not, so the mark spans the whole frame. A line **with a Z elevation** (a 3D object, an inclined one, a surface contour) gives a **point** at the real elevation of the crossing - a roof contour with an elevation, for instance, lands as a point exactly on the bed. A polygon (a plan zone - replacement, a mine field, a licence) gives a **vertical band** over the interval where the section runs through the zone.

The line, vex and frame height come from the section definition - written by **Cross-section along a line**, which now stores the vertical extent. So nothing needs to be supplied for objects without Z. For older definitions without the height a fallback remains: the **section drawing** as the optional input, or a Z range in the advanced parameters. When the object has Z, no height is needed - the point is placed at the elevation. Empty outputs are not created: a fault yields only verticals, a marker only points, a zone only bands.

Unlike **Project objects onto the section** (approximate, corridor-based) this is an exact intersection - a mark appears only where the geometry truly cuts the section line. Several layers can be fed at once (lines and polygons mixed) - all are processed in a single run, like the list of surfaces in 4.04, and in the outputs the **src** field keeps the source layer of each mark. The demo generator outputs a fault, a Z marker and a replacement zone that cross the demo section, so the tool can be tried at once.

Parameters

Parameter	What it sets	Default / hint
Section definition	The definition layer from 4.01: the line, vex and frame height. Enough for objects without Z.	-

Parameter	What it sets	Default / hint
Layers to intersect	Lines and polygons mixed, in a single run. The src field in the outputs keeps the source layer.	-
Section drawing	A fallback source of frame height for older definitions without zmin/zmax.	optional
Bottom of Z range (Adv.)	The lower frame elevation when neither a definition with height nor a drawing is given.	0
Top of Z range (Adv.)	The upper frame elevation in the same case.	0
Verticals on the section	Output for lines without Z (a fault, a boundary): a full-height vertical.	created
Intersection points	Output for lines with Z (a contour with elevation): a point at the real elevation.	created
Zone bands on the section	Output for polygons (a plan zone): a vertical band over the interval.	created

Empty outputs are not created: each object type goes only into its own layer.

4.06 Intersect a TIN with the section

A raster grid (4.04) is `z = f(x, y)`, one elevation per plan point. It cannot represent an overturned fold at all: above one point such a fold has several elevations of the same surface. This tool cuts the section through a **TIN** - a surface of true 3D triangles that can overhang.

The mechanics are pure geometry. The section is a vertical curtain along the polyline. Each TIN triangle is intersected with the vertical plane of its segment, giving a segment (station along the line, real elevation), and all segments are assembled into the surface trace. Overhang comes out naturally: several segments at different elevations above one station, and the trace folds - the limbs of an overturned fold come out as they are.

The inputs are layers of **3D polygons** (PolygonZ, TIN faces; non-triangles are fan-split into triangles) and optionally a **mesh layer**. The line and vex come from the section definition, the height from the faces themselves, so nothing needs to be set for a TIN. Besides the drawing trace you can also get it in real 3D coordinates.

An important limit: **a QGIS mesh is 2.5D**, its height is a scalar per vertex, one value above a point again, so overturning is not preserved in a mesh. Overhangs therefore come only from true 3D faces from a geomodeller (Leapfrog, Micromine and the like). A mesh is accepted for generality, on single-valued surfaces. The demo generator outputs an overturned TIN fold - the folding trace is visible on it at once.

Parameters

Parameter	What it sets	Default / hint
Section definition	The definition layer from 4.01: the line and vex. The height comes from the faces themselves.	-
TIN faces (PolygonZ)	Layers of 3D polygons - TIN faces from a geomodeller. Overhang and overturning are reproduced.	optional
Mesh layer (2.5D)	A QGIS mesh, for generality. One value above a point, overturning is not preserved in a mesh.	optional
TIN trace on the section (drawing)	The output 2D trace in the section axes, may fold.	created
TIN trace (3D)	The output trace in real 3D coordinates.	on request

Supply at least one of the two inputs - TIN faces or a mesh.

4.07 Project objects onto the section (beta)

The **Project objects onto the section** tool projects points, lines and polygons onto the section line. For each vertex the horizontal coordinate is the distance along the line to its projection, the height is the elevation from the 3D geometry or from a chosen field. Distant objects are cut off by a corridor.

Parameter	What it sets	Default / hint
Section definition	Section line with the vex field (vertical exaggeration).	-
Objects to project	Layer whose objects are projected onto the section plane.	-
Elevation field	Z attribute if the object geometry has no Z.	-
Corridor from the line	Width of the capture band from the section line. 0 = all objects.	0
Objects on the section (drawing)	Output: objects in section-drawing coordinates.	-

This generalises the borehole projection to any objects: anomalies, sampling points, traces, outlines. The result is in the section axes, placed on top of the drawing.

4.08 Unproject from the section (beta)

The **Unproject from the section** tool does the reverse: objects drawn on the section drawing are returned to real coordinates. The horizontal coordinate of a vertex is read as the distance along the line (giving the plan), the height as the elevation Z = height / vex. The line and vex come from the same definition the drawing was built with.

Parameter	What it sets	Default / hint
Section definition	Section line with the vex field.	-
Objects from the section drawing	Objects in drawing coordinates to return to plan.	-
Objects in plan (with Z elevation)	Output: objects in plan coordinates with restored elevation.	-

So an object drawn by hand on the section - an ore outline, a fault, a boundary - gets back into the plan and into 3D with a Z elevation.

4.09 Shaft wall unwrap (beta)

The **Shaft wall unwrap** tool builds a cylindrical section. Around the shaft axis at a given radius a circle is taken with an angular step (1 degree by default), and the surface grids are sampled along it. The unwrap lies in axes of arc length along the circle and elevation.

Parameter	What it sets	Default / hint
Shaft axis (collar point)	Collar point of the shaft.	-
Shaft radius, map units	Radius of the wall unwrap.	4
Surface grids (markers)	Surface grids whose elevations are drawn on the unwrap.	-
Angular step, degrees	Step of the wall traversal by azimuth.	1
Vertical scale	Vertical-exaggeration mode (H:V ratio or multiplier).	ratio
Scale value	Exaggeration number for the chosen mode.	10
Raster sampling (adv.)	How the value is sampled from the surface grid.	-
Wall unwrap (arc × elevation)	Output: shaft-wall unwrap.	-

Each marker surface gives the line of its intersection with the shaft wall - where the beds dip the lines are tilted and wavy. The axis is set by a collar point layer, the radius is in map units, the vertical scale is as in the section.

4.10 Create a section example

The **Create a section example** tool prepares a complete training set for the **Cross-sections** group, so its tools can be tried without kriging real data. In the panel it stands last in the **Cross-sections** group.

A single run outputs six stacked surfaces with a dip and variable thickness (five interbedded beds, the 2nd and 4th industrial and thin), a polyline section line across the area, boreholes along the line with surface-elevation fields h1…h6, and a multiband grid per industrial bed. The bed-grid band convention: band 1 - the roof, band 2 - the bottom, bands 3 and further - parameters (here the content and the mineral type with a replacement zone; the content fields of the beds are independent, stochastic). One file describes the whole bed - like a block model where new parameters are added as bands. For the intersection tools it adds demo vectors: a fault without an elevation, a marker contour with Z, a replacement zone, and an overturned TIN fold from PolygonZ 3D faces.

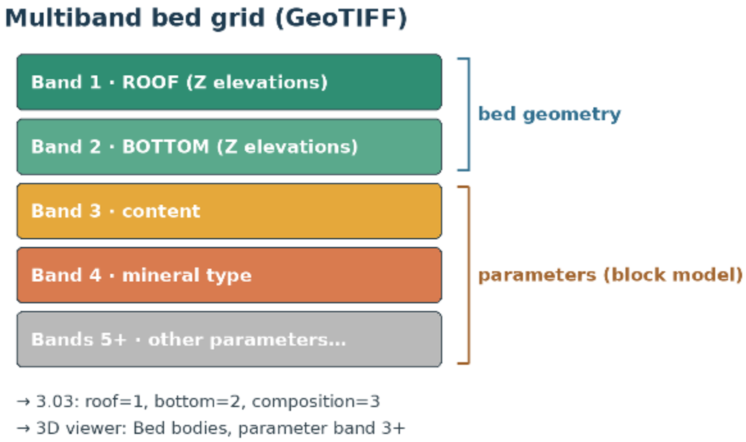

Figure 33: The multiband bed-grid convention: bands 1-2 carry the geometry (roof and bottom), bands 3+ the parameters; one file feeds tool 4.03.

Demo layers and their attributes

Layer	Geometry	Attributes
Section line (demo)	line	name. A polyline with two bends - the vertices test the stationing.
Boreholes (demo)	points	name; h1...h6 - elevations of the six surfaces at the borehole.
Zone (demo, polygon)	polygon	name. For the vector-intersection tool 4.05.
Fault (demo, 2D)	line	name. Crosses the trace, moved off the line bend.
Marker with Z (demo, 3D)	line Z	name. Tests 3D geometries in 4.05.
Overturned TIN (demo)	polygons Z	name. An overturned fold for 4.06.
Surface 1...6	raster	A single elevation band.
1st/2nd industrial bed	raster	Bands: 1 roof, 2 bottom, 3 content, 4 mineral type.

The workflow is shown in section 4.01: the surfaces go into **Cross-section along a line**, the boreholes into **Boreholes on the section**, the composition grid into **Bed composition on the section**, and the demo vectors and TIN into the intersection tools 4.05 and 4.06. The whole cross-section group then runs on consistent data.

Parameters

Parameter	What it sets	Default / hint
Area (extent)	The rectangle in which the set is generated.	set per project
Generator seed (Adv.)	The RNG seed for reproducibility. 0 means random on each run.	0
Surface 1…6	Six rasters top to bottom: the roofs and floors of the host beds and the two industrial beds.	created
Section line	A polyline across the area to feed into 4.01.	created
Boreholes along the line	Points with surface-elevation fields h1…h6 for 4.02.	on request
Composition: content	A content grid of the industrial beds for 4.03.	on request
Composition: type/facies	A mineral-type grid (1 sylvinite, 2 replacement) for 4.03.	on request
Fault, Z marker, zone	Demo vectors for 4.05: a line without Z, a contour with Z, a zone polygon.	on request
Overturned TIN	3D faces of an overturned fold for 4.06.	on request

The full list of output-layer fields is in the **Section data** appendix section ("Demo-layer fields" at the end of the manual).

5.01 Fractal dimension

The tool computes a fractal-dimension map of a surface by the variogram method, native to the plugin: a log-log variogram over lags of one to N cells is built in a sliding window, its slope gives the Hurst exponent H, and the dimension D = 3 - H. Smooth differentiable areas give D near 2, rugged and noisy ones tend to 3; the values themselves matter less than their steps - they highlight zones of tectonic disturbance, block boundaries and changes of the roof relief character.

The output is a D grid that feeds straight into **1.04 Isolines from a raster** for dimension isolines; an advanced checkbox adds H as band 2. The global D and H over the whole surface are printed to the log.

Reading the map

The absolute D values matter less than their steps: a linear step across the area is a lineament, a candidate tectonic disturbance; a patch of a raised D is a zone of intense folding or a rugged roof relief; wide even fields of a low D are quiet blocks. For reading, apply a singleband pseudocolour symbology with a contrast palette and quantile classification, and for a report plan build isolines with belts over the D grid with tool 1.04 - the disturbance zones get outlined like contour lines.

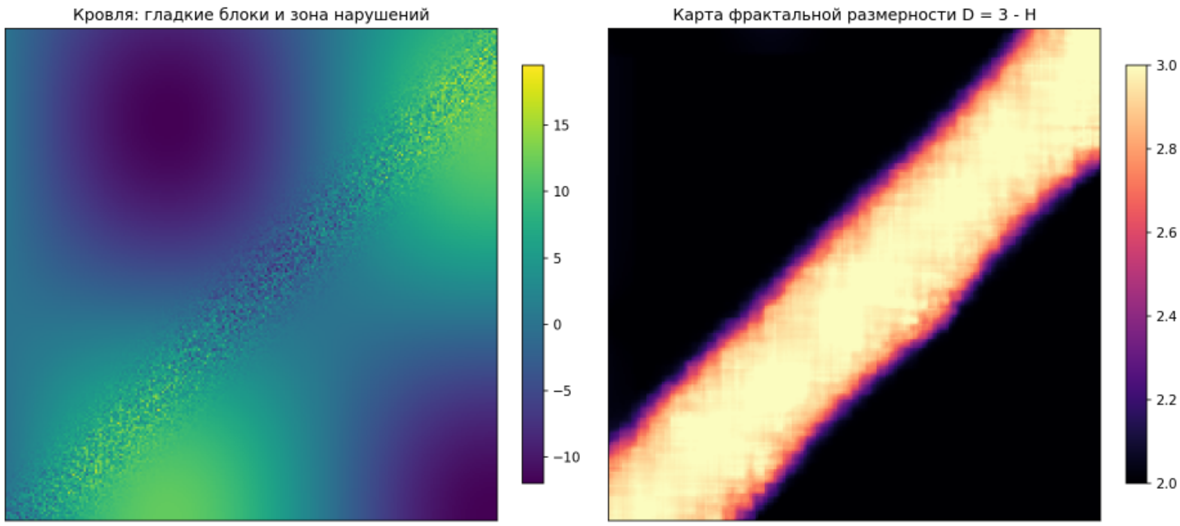

Figure 34: A synthetic roof with a diagonal crushing zone and its D map: quiet blocks near 2, the disturbance zone shows up as a bright lineament.

Picking the window and the lags

A small window (5-8 cells) reveals the microstructure and local disturbances, a large one (12-20) - regional zones; in doubt compute both and compare. Four lags fit almost always: more lags - a steadier slope but a coarser minimal scale the method can resolve. The window and the lags are limited by the grid size, the tool checks that itself.

Workflow

A bed roof from kriging → **5.01** with a window of 8 → the D grid → **1.04 Isolines from a raster** (band 1) → dimension isolines with belts over the structural plan. The global D from the log is one number per surface to compare areas or beds with each other. The raster must be in a metric CRS; the demo surfaces fit as they are.

Parameters

Parameter	What it sets	Default / advice
Surface (raster)	A relief grid or any surface.	-
Window half-radius, cells	The sliding-window size.	8
Number of lags (Adv.)	Variogram lags 1..N cells.	4
Elevation band (Adv.)	The band with elevations.	1
Write H (Adv.)	Add H as band 2.	off
Fractal dimension	A D grid (and H if checked).	-

5.02 Mask box-counting

Classic box-counting for binary masks: the raster is binarised by a threshold (the object - values above it), the mask is covered by cells of a decreasing size, the slope of log N versus log(1/size) gives one dimension D for the whole mask. A linear object gives D near 1, a blob - near 2, rugged outlines of replacement zones or mined-out areas fall in between. The accuracy on finite masks is about ±0.1, so the method is good for comparing masks with each other rather than as an absolute measure. The result is printed to the log with a table of sizes and counts and returned as the number D - usable further in Processing models.

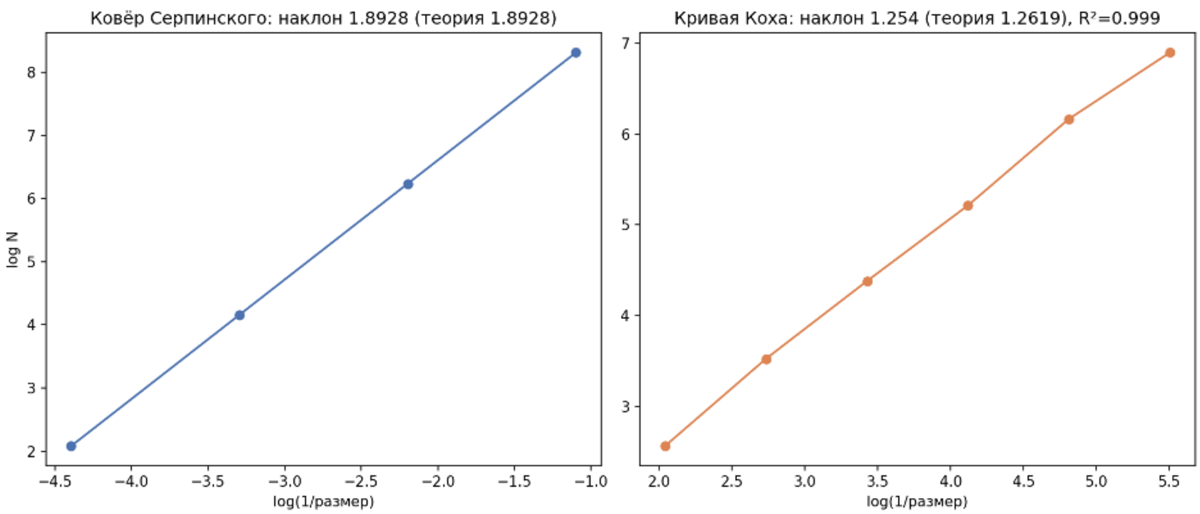

Figure 35: Checking the estimators on the references: the Sierpinski carpet gives a slope of 1.8928 against the theoretical 1.8928, the Koch curve - 1.254 against 1.2619. Points on a line - the power law holds.

Where the mask comes from

The mineral-type band of a bed grid with a threshold between the class codes, an indicator-kriging probability grid with a 0.5 threshold, an exceedance-probability map with a cut-off threshold, vector outlines of workings or zones - rasterised beforehand with the standard “Rasterize (vector to raster)”. Compare the D of masks of the same nature on the same grid: a growth of the replacement-outline ruggedness from bed to bed or from year to year is a meaningful signal.

Parameter	What it sets	Default
Mask raster	Any raster; the mask - values above the threshold.	-
Threshold	The object/background boundary.	0.5
Band (Adv.)	The raster band.	1

5.03 Line and boundary dimension

The dimension of every line by the divider (Richardson) method: the line is walked with chords of a decreasing span, the slope of log N versus log r gives D. A straight line gives one, a rugged line - more. Polygons are accepted alongside lines - the exterior ring of the boundary is measured, so the ruggedness of zone and basin outlines is computed without a prior conversion. The output is the same features with the D and steps fields, the mean D is printed to the log; short lines get an empty D. The method is checked on references: the Koch curve gives 1.262 against the theoretical 1.2619.

An isoline-smoothing diagnostic

Oversmoothed isolines lose their ruggedness and D drops towards one. The workflow: build the isolines twice - without smoothing and with the working parameters, run both layers through the tool and compare the mean D from the log. A drop by hundredths is cosmetics, the shape is kept;

a drop by tenths means the smoothing eats the field geometry - weaken the rounding or keep the densification only. The D field in the attributes lets you find the specific lines that suffered most.

Other uses

The ruggedness of zone outlines in plan, comparing the digitising detail of boundaries from different sources, generalisation control when preparing small-scale plans - anywhere “how winding the line is” must become a number.

Parameter	What it sets	Default
Lines	A line layer (isolines, outlines).	-
Lines with the dimension	The same lines with the D and steps fields.	-

5.04 Minkowski dimension (vectors)

Box-counting directly over vectors, no rasterisation: lines and polygon boundaries are covered by a grid of a decreasing size, the slope of log N versus log(1/size) gives the Minkowski dimension. A straight line and a smooth boundary give D near one, a river network - 1.1-1.5, a heavily rugged coastline - up to 1.3 and above. Every feature gets the D_mink and D_r2 fields (the log-log fit quality: below 0.85 the estimate cannot be trusted), and separately the D of the layer as one set is computed and printed to the log: for a river network that is the dimension of the network as a whole, regularly higher than that of the individual branches.

The method complements the divider of 5.03: the divider measures the sinuosity of one line, Minkowski - the plane filling by a set of features. The dimension is also returned as a number output for Processing models.

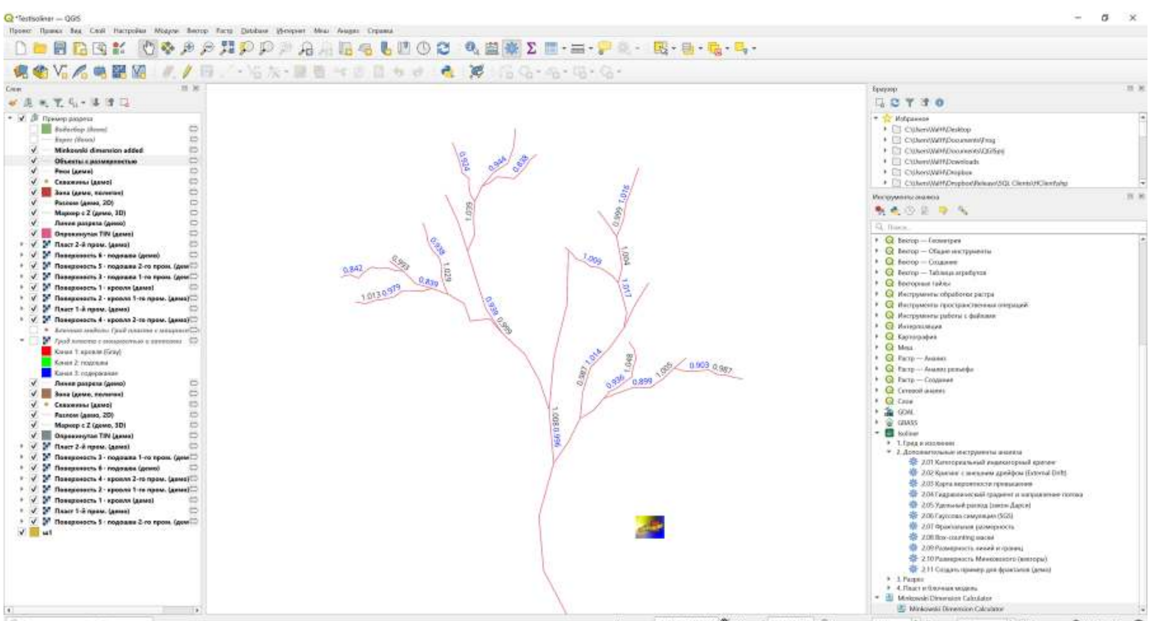

Figure 36: Demo rivers labelled by the per-branch D_mink: nearly smooth branches give values around one, the network as a whole - higher.

Parameter	What it sets	Default
Lines or polygons	A vector layer; for polygons the boundary rings are taken.	-
Number of grid sizes, K (Adv.)	Ladder steps; a too large K takes the cells below the line detail and lowers D.	8
Grid offsets per size (Adv.)	Random shifts, the minimal cover is taken - removes the grid alignment.	3
Densify factor (Adv.)	The sampling step along segments as a cell fraction; 0 - vertices only.	0.5

5.05 Create a fractal example (demo)

A generator of study features for the whole fractal five: a branching river network with an order field (the tributary order), a basin polygon with a rugged boundary and a separate coastline built by midpoint displacements. Feed the rivers into 5.04 - you get the network dimension; the coast and the basin boundary - into 5.03 and 5.04 and compare the divider with Minkowski; rasterise the basin with the standard tool - and it doubles as an example for 5.02.

Parameter	What it sets	Default
Extent	The generation area.	-
Seed (Adv.)	The example repeatability.	1

Typical situations and solutions

What you see	Cause	Solution
Concentric “bull’s eyes”, cones	Kriging pulls the value exactly through outlier wells (nugget 0).	Set a nugget C0 (0.2-0.4 of the sill, in absolute variance units). And/or enable grid smoothing in **2D Kriging**.
Angular isolines (“octagons”)	A coarse grid: vertices are placed at cell edges.	Increase **Line rounding** to 3 or reduce the cell size in kriging.
Radial/fan lines in empty corners	Extrapolation beyond the data.	Enable **Clip to well hull** or set a clip mask.
Isolines cross in dense areas	Formerly - a consequence of smoothing each line.	Smoothing is done over the field (in **2D Kriging**). Increase the grid-smoothing radius there.
Polygons of one colour	By default the layer is created with a single symbol.	Set graduated symbology by ELEV_MIN.

Appendix. Lessons and self-check

This section is a hands-on practicum on the demo data: short working cycles (lessons) and tests to make sure the tools compute correctly. Everything is reproduced by the plugin generators, no own data is needed. The section grows as the plugin develops.

Preparing the demo

1. A new QGIS project.
2. **1.10 Create a well example (demo)** - demo wells with elevation and content fields appear.
3. **4.10 Create a cross-section example** - the section line, the zone polygon, the ready beds "Bed 1 (demo)" and "Bed 2 (demo)" and separate roof and bottom surfaces.

Lesson. Topo2Raster: terrain from contours, step by step

The goal is to walk the full cycle of building terrain from vector data and to make sure the reconstructed surface matches the original. The lesson works offline: the source is the demo relief. With live data the cycle is the same, only steps A and B are replaced by digitized contour lines of a topographic plan and watercourses from 2.02.

Step A. The source surface (2.10). Run **Demo relief** with the defaults (300×300 cells of 30 m, seed 42). A raster appears in the Topography group - this is the "truth" we will reconstruct. Note the elevation range in the layer properties.

Step B. Contours from the truth (1.04). Feed the demo relief into **Isolines from raster**, a 5 m step, smoothing can stay at the defaults. You get a line layer with the **ELEV** field - an imitation of a digitized topographic plan. Contour polygons are not needed for the lesson, the field can be cleared.

Step C. Streamlines (2.06). Run **River network** over the demo relief with a threshold of 300 cells. Leave the fill checkbox on. The output is lines with vertices running downstream, they will go into Topo2Raster as drainage enforcement.

Step D. Reconstruction (2.03). Open **Topo2Raster (terrain from vectors)**. Contours - the layer from step B, the elevation field **ELEV**. Streamlines - the layer from step C. Leave points, cliffs and lakes empty. Cell size 30 m, do not set the extent - it is taken from the layers. Everything else at the defaults. At this size the run takes seconds, the log shows the multigrid pass from the coarse level to the fine one.

Step E. Comparison with the truth. The most visual way is isolines: build 1.04 with the same 5 m step over the reconstructed raster and overlay it on the contours of step B in another color. The lines should land on each other, diverging only on the tops above the last contour. Numerically: **Raster calculator** (Raster - Raster calculator), the expression "reconstructed minus truth", and the difference statistics in the layer properties. The expected picture at a 5 m contour step: the mean near zero, the spread around a meter, the maxima on summits.

Step F. The role of summit marks (2.09 and 2.03). Run **Peaks** over the demo relief (radius 700 m, drop 8 m) and check in the attribute table that the **z** field is filled. Repeat step D adding this layer to **Elevation points** with the z field. The difference on the tops shrinks visibly - this is the answer to why topographers label summit elevations on maps.

Lesson summary. If the isolines coincide and the difference with the truth stays within the contour step and shrinks with summit marks, the cycle is mastered. With live data the order is the same: 2.01 gives a DEM for checking, 2.02 gives watercourses as streamlines and peaks with ele as elevation points, digitized contours go in as contours.

Appendix. Demo layer fields

A summary of the fields of all demo-data generators with units. The values are demonstrational: where a quantity is abstract, the units are nominal.

Topography group outputs

Line layer **River network** (tool 2.06):

Field	Type	Meaning, units
order	integer	Strahler order
acc_out	real	accumulation at the link outlet, cells
length_m	real	link length, m

Polygon layer **Basins** (2.07): **basin** - the basin number (integer), **area_m2** - the area, m². Point layer **Peaks** (2.09): **z** - the elevation, m, **drop** - the drop over the window minimum, m. Base topography layers (2.02) carry **name** and **osm_id**, watercourses also **waterway**, water bodies **water**, peaks **ele** (elevation, m), cliffs **kind**. The demo relief (2.10) is a raster without fields.

Sample wells (demo) - tool 1.10

Point layer **Sample wells (demo)**:

Field	Type	Meaning, units
well	string	well name or number
roof	double	absolute bed roof elevation, m
thick	double	bed thickness, m
X	double	content of a nominal component, percent (demo)
head	double	head, m (optional)
K	double	hydraulic conductivity, m/day (demo, optional)
T	double	transmissivity, m²/day, T = K·thickness (optional)
mintype	string	mineral type, category (demo replacement, optional)
dz	double	value linearly related to the drift, nominal (optional)

Optional raster **Drift surface (demo)** - an external surface for external-drift kriging, nominal units.

Electrical-prospecting profiles - tool 1.11, electrical mode

Point layer **Electrical-prospecting profiles (rho_k, SP, IP)**:

Field	Type	Meaning, units
profile	integer	profile number
picket_m	double	picket from the profile start, m
pk	string	picket as a PK label (e.g. PK5+20)
z	double	surface elevation, m
rho_k	double	apparent resistivity rho_k, Ohm*m
rho_true	double	rho_k without noise, Ohm*m (reference)
sp	double	self-potential SP, mV
vp	double	induced polarisation IP, mV/V

Subsidence profiles - tool 1.11, subsidence mode

Point layer **Subsidence profiles (trough, tours: N)**:

Field	Type	Meaning, units
profile	integer	profile number
picket_m	double	picket from the profile start, m
pk	string	picket as a PK label
tour	integer	observation tour number
z	double	surface elevation, m
settle	double	subsidence, mm
settle_true	double	subsidence without noise, mm (reference)

Section data - tool 4.10

- **Surface 1…6** (rasters) - elevations of six stacked surfaces, m. Top to bottom: 1 roof of the upper host, 2 roof and 3 floor of the 1st productive, 4 roof and 5 floor of the 2nd productive, 6 floor of the lower host.
- **Section line (demo)** (line): field **name** - line name.
- **Section wells (demo)** (points): **name** (SKR-NNN), **h1…h6** - elevations of the six surfaces at the well, m.
- **1st productive bed (demo)** and **2nd productive bed (demo)** (multiband rasters): band 1 roof (m), band 2 floor (m), band 3 content (percent), band 4 mineral type (category 1 or 2).
- **Fault (demo, 2D)** (line): **name**.
- **Marker with Z (demo, 3D)** (line with Z): **name**.
- **Zone (demo, polygon)** (polygon): **name**.
- **Overturned TIN (demo)** (3D faces): **name**.

Fractal data - tool 5.05

- **Rivers (demo)** (lines): field **order** (integer) - tributary order in the network hierarchy.
- **Basin (demo)** (polygon): field **name**.
- **Coast (demo)** (line): field **name**.

Density (demo) - tool 3.08

Point layer **Demo points**:

Field	Type	Meaning, units
mass	double	measurement mass (demo, 50 per point)
prec	double	uncertainty sigma, m

Line layer **Demo lines**:

Field	Type	Meaning, units
mass	double	measurement mass (demo, 100 per line)
prec	double	corridor half-width, m
from_m	double	interval start along the line, m (may be empty)
to_m	double	interval end along the line, m (may be empty)

Polygon layer **Demo polygons**:

Field	Type	Meaning, units
mass	double	measurement mass (demo, 150 per polygon)
dasy	integer	1 - polygon for dasymetry, 0 - uniform

Auxiliary raster **Auxiliary raster (dasymetry)** - a value gradient for the dasymetric mode of 3.07, nominal units.

Appendix. Variable-support density: step by step

A walkthrough of tools 3.07 and 3.08 from demo generation to a finished result, with an explanation of each parameter. The example shows how to check the mass invariant by eye.

Step 1. Generate the demo (3.08)

Run **3.08 Create a density example (demo)**.

- **Extent** - the generation bounds. Set a rectangle on the map or by a layer. Any metric extent works, for example a square a few kilometres across.
- **Auxiliary raster cell, m** - the step of the auxiliary raster for dasymetry. Default 50, enough for the demo.
- **RNG seed** (Adv.) - example reproducibility. With one seed the set is identical.

The output is four layers: **Demo points**, **Demo lines**, **Demo polygons** and **Auxiliary raster**. The log states the embedded mass: points 500, lines 200, polygons 300, total 1000.

Step 2. Density from points (3.07)

Run **3.07 Density from measurements** on the **Demo points** layer.

- **Measurements** - the Demo points layer.
- **Mass field - mass**. If left empty, each object mass is 1.
- **Precision field - prec**. For points this is the Gaussian spot sigma in metres. The larger the sigma, the more diffuse and lower the point contribution.
- **Cell size, m** - the result grid step. A smaller cell means a more detailed map and a longer computation. The density in mass per km2 does not depend on the cell size, only the detail does.
- **Area** - leave empty, taken by the layer.
- **Support beyond edge - Renormalise inside** keeps all mass in the area, **Lose mass** discards the part beyond the edge with a warning. To check the invariant use renormalise.
- **Default sigma** (Adv.) - used when the precision field is empty. Zero means the half-cell.

Check the log: the line **Input mass: 500. Mass on grid: 500. Discrepancy: 0**. This is the invariant - the density integral equals the sum of masses.

Step 3. Density from lines

Run 3.07 on the **Demo lines** layer.

- **Precision field - prec** - here it is the corridor half-width in metres. The line is spread into a soft-edged strip.
- **from_m / to_m** (Adv.) - the **from_m** and **to_m** fields. One demo line has an interval set, and only its part is spread in the result. Empty fields mean the whole line.

In the log the mass on grid is 200. If the lose-mass mode is on and the corridor left the edge, the discrepancy shows how much mass was lost.

Step 4. Density from polygons and dasymetry

Run 3.07 on the **Demo polygons** layer.

- Without an auxiliary raster the polygon mass is spread uniformly over its area.
- **Auxiliary raster** (Adv.) - feed the demo **Auxiliary raster**. Then dasymetry turns on for polygons: mass is distributed proportionally to the raster values inside the polygon rather than evenly. If the raster is empty inside the polygon, the tool falls back to uniform and writes this to the log.

The mass on grid is 300 in both modes; only the distribution shape inside the polygons changes.

Step 5. Mix types by appending

To gather points, lines and polygons into one raster, run 3.07 three times in a row.

- The first run is as usual, you get a three-band raster.
- On the second and third, in the **Append to an existing raster** parameter (Adv.) point to the raster from the previous run. The tool reads the three bands, adds the new mass and returns the updated raster.

After three appends the total mass on grid is 1000. Bands 2 and 3 (sum m*sigma and sum m) are service for exactly this - they let you append in series and keep the effective-sigma map exact.

Step 6. Effective sigma

Set the optional **Effective sigma** output. This is the mass-weighted sigma per cell, an effective-precision map. Where density is gathered from precise supports the sigma is small. Where from smeared coarse georeferences it is large. An analogue of kriging variance for the density floor: it shows where the result can be trusted and where it rests on coarse supports.

Reading the result

Band 1 of the result is density in mass per km2. Open it with a single-band pseudocolour style. Density spots are where reliable measurements cluster. The invariant in the log confirms that mass is conserved and the result is correct. The layer must be in a metric coordinate system, otherwise the tool refuses to run.

For enterprises

Isoliner grows on the tasks of real mining operations. We implement custom features to match production regulations, provide guaranteed technical support contracts and integrate the module into the production cycle, including corporate database connections. Details: https://www.informpp.ru/главная-страница/предприятиям

Log and toolbar

Isoliner keeps a work log in the **isoliner.log** file next to the QGIS profile. At the start of a session it records the versions of the plugin, QGIS, NumPy and GDAL, then the name of the launched tool, its parameters, the run time, and on a failure the full traceback. The computation window closes but the file remains, so it is enough to attach it when reporting.

Open the log via **Plugins - Isoliner - Log** or with the **Log** button in the **About** window. The **Isoliner** toolbar holds the density map and the About window.

License and support