# BayesGWL

Bayesian groundwater level estimation – in-browser app • Hydrogeological Repository

## What it does

BayesGWL estimates the groundwater level across a study area from a scatter of borehole measurements. Rather than interpolating water levels on their own, it co-kriges them against surface elevation: because the water table generally follows topography, the terrain acts as a ‘qualified guess’ that guides the estimate between boreholes and keeps the result hydrogeologically plausible where data are sparse.

The method is a direct port of the Bayesian estimation routine from *Tripol*, a Fortran geostatistics package, reimplemented to run entirely in the browser. The numerical core was validated against the original program’s published outputs.

## Everything runs in your browser

No installation, no account, and no upload of your data to a server. Your observation file and study-area boundary are read and processed locally by the page itself. The only outbound request is an elevation lookup for the interpolation grid.

## What you need

| Input | Format | Notes |
|---|---|---|
| **Study area boundary** | .zip containing .shp, .dbf, .shx (and ideally .prj) | A polygon shapefile. Projected coordinate systems are supported and reprojected automatically when a .prj is present. |
| **Borehole observations** | CSV: Name, Longitude, Latitude, Elevation (mamsl), Water Level (mbgl) | Water level is depth below ground – the standard dip-meter reading. A header row is expected. Can also be typed straight into the table. |

A minimum of three points is required; results improve markedly with more, and with points spread across the area rather than clustered.

## Data it supplies for you

| Source | Used for | Detail |
|---|---|---|
| **SRTM elevation** | Surface elevation at every grid point, and the ‘guess’ surface the estimate is built on | 30 m or 90 m product, selected automatically from the size of your study area. Retrieved on demand via OpenTopoData. |
| **South African National Groundwater Archive (2019)** | Optional extra water-level control points | Static water levels for archive sites falling inside your study area, merged with your own observations. Each site can be included or excluded individually. |

## How a run works

| Step | What happens |
|---|---|
| **1. Study Area & Grid** | Upload the boundary shapefile. The app sizes the area, picks a suitable SRTM product and grid spacing, and fetches elevations for the interpolation grid. Optional de-sampling thins dense grids. |
| **2. Observation Data** | Load your CSV or type the boreholes in. Archive sites are added automatically if enabled. Both tables can be sorted, and any site can be excluded from the analysis with a tick box. |
| **3. Semivariogram** | The spatial correlation structure of the water levels is computed. Sensible lag settings are suggested from the data; the sliders let you explore alternatives, and the plot updates live. |
| **4. Fit Variogram Model** | A theoretical model (Gaussian by default; Spherical, Exponential, Linear/Power and De Wijsian are also available) is fitted to the experimental points by bounded non-linear least squares. |
| **5. Predict Water Levels** | Bayesian estimation runs at every grid point. Leave-one-out cross-validation reports the expected error. Predictions can be clipped so the water table never sits above ground. |
| **6. Cross-sections** | South–north and west–east sections through the result, positioned with sliders, showing the water level beneath the ground surface. |
| **7. Report** | A single self-contained HTML report: settings, statistics, and every map and chart, ready to print or save as PDF. |
| **8. Export** | CSV export of the interpolation grid, the water-level dataset actually used, and the predicted levels at every grid point. Every map can also be exported as a PNG. |

## Reading the results

- **Leave-one-out (LOO) cross-validation** holds each observation out in turn, re-estimates it from the others, and compares. The RMS error is the headline number: the typical difference between estimate and measurement, in metres.
- **Standard error** in the error panel is a calibration check rather than a distance. It should be near 1.0 if the model's stated uncertainty matches its actual errors; well below 1.0 means the uncertainty is overstated.
- **Estimates follow the terrain.** Because elevation guides the interpolation, predicted levels can extend beyond the range of the measured ones where the ground rises or falls past any borehole. The depth below surface is usually the more revealing quantity.

## Limitations worth knowing

- Results are only as good as the input: sparse, clustered or inconsistent water levels give a correspondingly uncertain surface.
- Archive water levels are historic and of mixed vintage and quality. They add spatial coverage, not currency – review them before relying on them, and exclude any that look wrong.
- SRTM is a surface model of roughly 30–90 m resolution; it will not resolve local topographic detail, and elevations may differ from a surveyed collar by several metres.
- The method assumes a water table that broadly follows topography. Confined systems, steep hydraulic gradients near abstraction, and perched water are not represented by that assumption.

**Beta.** BayesGWL is published for evaluation and feedback. Check results against your own understanding of the site before using them in a report or a decision.