ChemGraphs

Piper, Durov, Stiff, Schoeller and SAR diagrams from your own analyses – in-browser app • Hydrogeological Repository

What it does

ChemGraphs draws the six diagrams a hydrochemist actually uses – Piper, Durov, expanded Durov, Stiff, Schoeller and the USSL sodium–salinity plot – from a spreadsheet of your own water analyses. Download the template, fill in what you have, upload it, and every diagram is drawn from the same selection and the same date setting, beside a map of where the samples came from.

It is built for a monitoring record rather than a single round of sampling. A date slider runs over the sampling visits, so a network can be stepped through one visit at a time, watched accumulate up to a date, or drawn whole with older analyses fading behind newer ones. What a water is doing over ten years is then visible on the diagram itself, not only in a table.

Everything runs in your browser

No installation, no account, and no upload of your analyses to a server. The workbook is read, the chemistry computed and the diagrams drawn by the page itself, on your own machine – which matters when the data is a client's record under a compliance licence.

Getting your data in

The template is generated by the app rather than shipped as a file, so it cannot drift out of step with the reader that parses it back. It has three sheets: a Read me, a **Sites** sheet with one row per monitoring point, and a **Samples** sheet with one row per analysis. Only four cells are compulsory – a Site ID on each sheet, and a sample date. Everything else is optional, and a blank cell means “not determined”, which is not the same as zero.

The file	What is accepted
Formats	.xlsx, .xls or .csv. A laboratory's own export can usually be dropped in unmodified: sheets are identified by the shape of their header row, not only by their name.
Column headings	Matched loosely. Case, spacing, punctuation and the unit in brackets are all ignored, and each column accepts several names – “Borehole ID”, “Site”, “Monitoring point” and “Station” all read as the site.
Alkalinity	Give it as $CaCO_3$ or as HCO_3, in either of two columns. Bicarbonate is converted on the way in at 61.017 mg per milliequivalent.
Nitrate	Give it as N or as the nitrate ion, again in either of two columns. NO_3 is converted to N at 62.004 mg per milliequivalent.
Conductivity	mS/m or µS/cm, in either of two columns. A column headed only “EC” is read in whichever unit the page is set to; a column that names its own unit always overrides that setting.
Below detection	Write it as “<0.05” or as “-0.05”. Either is read as a value of 0.05 and used at that magnitude, so a half-limit chloride still counts in the ion balance instead of throwing the analysis out of it.
Dates	Best written YYYY-MM-DD, or entered as a real spreadsheet date. A date written DD/MM/YYYY is read day first.
Coordinates	Decimal degrees on WGS-84, on the Sites sheet. A southern latitude is negative. A site without a coordinate still gets its diagrams; it simply does not appear on the map.

Nitrate as N, or nitrate as NO_3? Getting these the wrong way round changes a result by 4.43 times, and nothing downstream can detect it – the number is chemically plausible either way. Fill one column, leave the other empty; the export records which one every analysis came in as.

A row that cannot be read is reported rather than silently dropped. Each one is collected with its sheet, its row number and the reason – a missing site, an unreadable date, a coordinate that is not a decimal-degree pair – and listed after the upload, so a partly loaded file is visible as one rather than as a quietly short dataset.

The six diagrams

Diagram	Built from	What it is for
Piper	Two ternary triangles and the rhombus between them; each analysis appears three times.	Classifying a water, and seeing groups of waters. Mixing lines are straight in the rhombus, so a water made of two end members plots on the line between them.
Durov	The same two ternaries projected onto a square, with a pH panel below and a logarithmic EC panel to the right.	Telling processes apart. Mixing, dissolution and ion exchange separate on the pH and EC panels in a way they do not in the square alone.
Expanded Durov	The Durov square in nine fields, with a small triangle, a rhombus and a small triangle over the three thirds of each edge.	Naming the hydrochemical facies. Position advances in whole thirds, so a water crossing a boundary steps rather than slides.
Stiff	Three horizontal axes in meq/L – Na+K against Cl+NO_3, Ca against alkalinity, Mg against SO_4.	A shape you can recognise at a glance, and carry onto a map. Panels share one scale, so width means concentration.
Schoeller	Six determinands on one logarithmic meq/L axis, one line per analysis.	Comparing waters of very different strength. The same water type at different strengths gives parallel lines.
SAR (USSL)	Sodium adsorption ratio against conductivity, on the US Salinity Laboratory classification.	Whether a water is fit to irrigate with. It needs Na, Ca, Mg and an EC, so fewer analyses reach it than reach the others.

A ternary diagram is a diagram of proportions, and only of proportions. Two waters that plot on the same spot of a Piper can differ by a factor of ten in every ion. That is not a defect – it is what lets the diagram show that a water has kept its character while getting stronger – but it is the commonest misreading of one. If the question is “how much”, read the Stiff or the Schoeller.

Reading a record through time

Setting	What it draws
One visit	The sampling round nearest the slider, within a tolerance you set from a week to six months. Sampling rounds never land on the same day at every site, so an exact-date match would empty the diagram at almost every step.
Up to this date	Every analysis from the start of the record to the slider, older ones fainter.
The whole record	Every analysis for the selected sites at once.

Setting	What it draws
Play	Steps the slider through the sampling rounds at one of three speeds, redrawing as it goes.
Fading	Older samples are drawn more transparent than newer ones. The oldest is drawn at 20% of the newest one's strength by default, and 10, 35 and 60% are on offer if the record is sparser or busier than usual. The ramp is per site, so a borehole sampled twice and one sampled thirty times read the same way. On the Schoeller the old lines are thinner as well: transparency accumulates where lines overlap and width does not, so on a busy sheet width is what still carries the age.
Evolution path	On the SAR plot only, a line joining a site's analyses in time order. The three ternary diagrams do not offer it: both SAR axes are quantities, so a line between two analyses is a real path through concentration space, whereas on a diagram of proportions it would cross ground the water never occupied.

Up to ten sites may be drawn at once. Each takes a colour *and* a marker shape and keeps both everywhere. Ten colours cannot be told apart by a colour-blind reader on a scatter plot, so shape is the channel that carries identity and the legend shows both. Nothing on a diagram exists only as a colour: every mark has a hover label, and every value behind it is in the table.

One diagram at a time

The six sit behind a row of tabs and one is drawn at a time, with only the settings that belong to it beside it. Marker size is common to all of them; three settings belong to one diagram each:

Setting	What it does
Piper: size by EC	Marker area proportional to conductivity about the median, with a key. Offered on a single visit only: across a decade the marks would be sizing against each other in time as well as between sites, and neither reading would survive that.
Schoeller	A milligram scale per column as well as the shared meq/L axis, and an option to fold potassium into sodium and nitrate into chloride.
Stiff: site	Which site the sheet is drawn for. One analysis per site is a comparison between sites, and the whole selection fits on one sheet; more than one per site makes the sheet that site's own history, so it is drawn for one site at a time and you choose which.

What you get

Output	Detail
Map	Aerial imagery by default – a monitoring network is read against what is on the ground. Aerial with labels, a street map and no base map are a click away. Each site carries the symbol for its type; click one to add or remove it.
Diagrams	One at a time, chosen from the tabs, with its own PNG and SVG beside it. They are vector graphics, so the SVG export is the diagram itself and opens in Illustrator or Inkscape to be relabelled without losing anything.
Sample table	Every analysis on screen with what was derived from it – the plotting percentages, the water type, hardness, SAR and the ion balance. Sortable on any column; click a row to centre the map on that site.

Output	Detail
Report	A standalone printable page: the sites, the analyses, every row that could not be read and the arithmetic behind each derived figure, plus whichever figures you tick. Each is drawn fresh at report resolution, including diagrams not on screen, so a report is never limited to the tab you happen to be on.
Exports	PNG or SVG of any diagram, on screen or not; PNG of the map; CSV of the analyses with every derived value beside the reported one; CSV of the sites; CSV of the conversion factors; and the loaded dataset written back out as a filled-in template.

Site types

Twelve types, each with its own map symbol. Free text is matched onto them, so “Monitoring borehole”, “BH” and “Well” all read as a borehole; anything unrecognised becomes “Other” rather than being rejected.

Group	Types
Groundwater	Borehole, well point, spring or eye, piezometer, seep or decant
Surface water	River or stream, dam or reservoir, wetland or pan, canal or discharge
Other	Rainfall, tap or reticulation, other

The arithmetic behind every figure

Every concentration is divided by its equivalent weight to give milliequivalents per litre, which is the unit all six diagrams are built on. Potassium is then added to sodium and nitrate to chloride, and each side is expressed as a percentage of its own total.

Determinand	mg per meq	Where it is used
Calcium	20.04	Every diagram, and the ion balance
Magnesium	12.156	Every diagram, and the ion balance
Sodium	22.989	Every diagram, the ion balance and SAR
Potassium	39.102	Carried with sodium; not part of SAR
Total alkalinity as $CaCO_3$	50.0	Every diagram, and the ion balance
Sulphate	48.03	Every diagram, and the ion balance
Chloride	35.453	Every diagram, and the ion balance
Nitrate as N	14.007	Carried with chloride
Fluoride	19.0	Ion balance only
Aluminium	13.49	Ion balance only
Iron	27.92	Ion balance only
Manganese	27.47	Ion balance only

Ion balance error is (sum of cations – sum of anions) ÷ (sum of cations + sum of anions) × 100, on milliequivalents. Fluoride is counted with the anions and aluminium, iron and manganese with the cations, so a ferruginous water

balances here where it would not on the six major ions alone. Within 5% is a good analysis, 5 to 10% is worth a query, and worse than 10% usually means a determinand is missing rather than wrong. Aluminium is carried at 13.49 mg per meq, i.e. as a divalent species; the trivalent figure is 8.99. It never reaches a diagram, only the balance.

Sodium adsorption ratio is Na ÷ √((Ca + Mg) ÷ 2) on milliequivalents; potassium is not part of it. The USSL salinity classes divide at 250, 750 and 2 250 µS/cm. The three sodium-class boundaries slope, because the SAR at which sodium begins to damage soil structure falls as the water gets more saline: from SAR 10, 18 and 26 at 100 µS/cm down to 1.5, 5 and 9 at 5 000 µS/cm. Those anchors are fixed, so rescaling the axes to fit a dataset moves the frame and never the classification.

Water type names the ion holding more than half the milliequivalents on each side; where none does, that side is “mixed”. It is the reading the Piper gives, not a Stuyfzand or Kurlov classification.

How a run works

Step	What happens
1. Template	Download it, fill in the Sites and Samples sheets, or load the demonstration dataset to see what a populated one looks like.
2. Upload	Drop the workbook on the page. It reports how many analyses it read, how many carry enough ions to plot, how they balance, and anything it could not use.
3. Choose sites	Up to ten, from the map or the list. The list can be filtered by group, by site type or by name, and shows how many of each site's analyses are plottable.
4. Set the date	One visit, up to a date, or the whole record – and play it through if the record is long enough to be worth watching. The date controls sit with the diagrams, so the picture moves as the slider does.
5. Pick a diagram	From the tabs. Its own settings appear beside it.
6. Read and export	Any diagram as PNG or SVG, the numbers as CSV, or a report of the figures you choose.

The demonstration dataset

Synthetic. Nothing in it comes from a real monitoring record, and nothing in it should be quoted as an observation. It is a generated ten-year record – 991 analyses from 29 sites over 39 quarterly rounds – built round an ordinary model with no facility and no pollution in it: a ridge recharging a shallow aquifer, groundwater moving toward a pan, and the water evolving from fresh Ca–HCO_3 through Ca–Mg–SO_4 to Na–Cl and Na–HCO_3 on the way. Rainfall, a spring, a stream, a farm dam and the pan are sampled beside the boreholes, so eight site types appear.

The analytical noise, the reporting limits and the occasional unreported determinand are generated too: 4% of the analyses fall outside a 5% ion balance, and a few carry no alkalinity and cannot be placed on a ternary diagram. Coping with that is part of what the app is for.

Three things the app does with imperfect data, rather than refusing it

- **A result below the detection limit is used at its magnitude.** Dropping those analyses instead would discard most of the fluoride and nitrate in a typical monitoring record, and would put the ion balance out by more than the value itself.
- **An analysis that does not balance is drawn anyway, and flagged.** The balance is in the table, in the hover label and in every export. Which analyses to set aside is a judgement about the laboratory and the sample, not one an app can make for you.
- **A missing determinand breaks a Schoeller line rather than being bridged across.** A straight segment drawn past an ion that was never measured would read as a measurement.

Limitations worth knowing

- **Ten sites is the limit, and it is a real one.** Past about ten series a Piper stops being readable whatever the colours and shapes do. Draw a subset, export it, and draw the next.
- **The Piper, Durov and expanded Durov say nothing about concentration.** The Stiff and the Schoeller carry the strength, and the sample table carries the numbers.
- **The Stiff sheet clips outliers by default.** One hypersaline borehole among fresh ones would leave every other panel a hairline, so the shared scale is set by the largest panel that is not an outlier and the few that overflow are clipped and counted. Full-range and per-panel scales are both a switch away, and the sheet says which is in force.
- **A sample needs Ca, Mg, SO_4 and alkalinity, plus one of Na or K and one of Cl or NO_3, to reach a ternary diagram.** The Durov panels also want pH and EC, and the SAR plot needs an EC as well as Na, Ca and Mg.
- **Nothing is validated against a laboratory certificate.** A transcription error in the spreadsheet is a transcription error in the diagram, and only an implausible ion balance will hint at it.

Beta. ChemGraphs is published for evaluation and feedback. Check any figure you rely on against the laboratory report it came from before using it in an assessment.