



GEM4D DOCUMENTATION

Wedge Stability at Minodes

User Manual

PRODUCT	GEM4D
DOCUMENT	Wedge Stability at Minodes — User Manual
VERSION	1.9.x
DATE	14 June 2026
STATUS	First issue

Contents

1. What this tool does.....	2
2. Background theory	2
2.1 Why roof wedges matter.....	2
2.2 The ubiquitous-joint assumption.....	3
2.3 Wedge geometry: three joints plus the back	3
2.4 Failure modes and Factor of Safety	4
2.5 Strength model	4
2.6 Orientation uncertainty: Fisher Monte-Carlo	4
3. What you need before running.....	5
4. Step-by-step	6
5. The settings dialog	7
5.1 Joint strength & rock weight	8
5.2 Analysis options.....	8
5.3 Joint sets table.....	8
5.4 Wedge meshes (scene + DXF)	9
6. Result columns reference	9
7. Interpreting the results	10
7.1 The deterministic triplet: WdgFoS, WdgMode, WdgDepth, WdgWeight	10
7.2 WdgPfail — probability of failure	10
7.3 WdgDepth95 / WdgWeight95 — how big, when it fails	11
7.4 A suggested screening workflow	11
8. Wedge meshes	11
9. How the local span is measured	12
10. Assumptions and limitations	13
11. Messages and troubleshooting	14
12. References.....	14
Revision history.....	14


1. What this tool does

Wedge stability at Minodes (Calc dropdown → *Wedge stability at Minodes*) screens an entire underground development network for gravity-driven roof wedges in a single run.

At every Minode Marker (a point on the development with a known local tunnel orientation), GEM4D:

1. measures the **local back span** by ray-casting against the visible tunnel meshes,
2. forms every possible **tetrahedral roof wedge** from the joint sets you supply (three joint planes plus the excavation back),
3. solves each wedge for its **failure mode** and **Factor of Safety (FoS)** using classic gravity limit-equilibrium theory (Hoek & Brown 1980; Goodman 1989 — the same family of methods as UNWEDGE),
4. optionally runs a **Monte-Carlo simulation** that perturbs the joint orientations to estimate a **probability of failure** and percentile wedge sizes, and
5. writes nine result columns back onto the Marker table, where they can be filtered, colour-mapped, and exported like any other Marker data. Optionally it also builds the governing wedges as colour-coded scene meshes.

The intent is **screening, not final design**: the tool highlights *where* along kilometres of development the wedge risk concentrates, which locations deserve a detailed single-location analysis (e.g. in UNWEDGE), and how sensitive the answer is to orientation scatter.



2. Background theory

This section is written for users who have not run wedge analyses before. If you are familiar with UNWEDGE-style analysis you can skim to §2.5, which covers the conventions specific to GEM4D.

2.1 Why roof wedges matter

Rock masses are cut by families of roughly parallel discontinuities — **joint sets**. Where three joint planes intersect above an excavation back (roof), they can isolate a tetrahedral block of rock: three

sides formed by joints, the fourth side being the excavation surface itself. If gravity can move that block into the opening — straight down or sliding along one or two of the joint planes — it is a potential **roof wedge**, and it is one of the most common hazards in jointed-rock underground excavations.

2.2 The ubiquitous-joint assumption

GEM4D does not know where each individual joint actually is. Instead, it makes the standard conservative screening assumption that joints of every set can occur **anywhere** ("ubiquitous joints"). Under this assumption, every combination of three joint sets forms the *largest wedge that can geometrically fit* at every location. The question the tool answers is therefore not "is there a wedge here?" but "if the worst-positioned wedge formed here, what would it look like and would it be stable?"

This is the same philosophy as UNWEDGE: it deliberately overstates how often wedges occur, in exchange for never missing a kinematically possible one.

2.3 Wedge geometry: three joints plus the back

For each combination (triplet) of three joint sets, the three joint planes cut traces across the local back. The three traces bound a triangle on the back — the **face triangle** — and the three planes meet at a point above it — the **apex**. The tetrahedron between them is the wedge.

Two rules size the wedge:

- **Span scaling.** The face triangle is scaled so that its minimal enclosing circle equals the locally measured back span (§9). A wedge cannot daylight across more rock than the opening exposes — this is the span-limited, conservative sizing convention.
- **Persistence capping (optional).** Real joints have finite size. If you give a set a persistence (trace length, metres), no trace of that set on the face triangle may exceed it, which can shrink the wedge below the span-limited size. This follows UNWEDGE's "scale wedge by joint length" convention.




Figure 2: A single intersection wedge viewed from the side, with the tunnel semi-transparent, and the Minodes coloured and scaled on WdgWeight.

2.4 Failure modes and Factor of Safety

With the wedge geometry known, gravity limit equilibrium classifies each wedge into one of four modes:

Mode code	Mode	Meaning
1	Fall	The block can translate straight down without engaging any joint plane — there is no frictional resistance at all. FoS = 0 by definition. The geometric picture: a wedge falls when the vertical (plumb) line through its apex passes inside the face triangle.
2	Plane slide	Gravity presses the block onto exactly one joint plane and it slides down-dip on that plane. FoS = resisting forces (friction + cohesion on the sliding plane) ÷ driving force.
3	Edge slide	The block stays in contact with two planes and slides along their line of intersection. FoS computed with the normal forces resolved onto both planes.
4	Stable	No kinematically admissible gravity-driven mode exists (e.g. the block is wider at the top than its outlet). FoS is reported as 99 (a "stable" sentinel, not a real number).
0	None	The triplet is geometrically degenerate (e.g. a joint nearly parallel to the back) — no wedge could be solved.

The **Factor of Safety** is the classical ratio of available shear resistance to the gravity driving force. FoS < 1 means the wedge cannot support itself without ground support; FoS slightly above 1 means it is nominally stable but with little margin. Mines typically use a **design FoS** of 1.3–2.0 for screening (the dialog default is 1.0, i.e. bare limit equilibrium — raise it to your site standard).

A counter-intuitive point worth knowing: a wedge formed by one shallow-dipping plane and two steep planes does **not** necessarily slide on the shallow plane — if the apex plumb line is inside the face triangle, the block free-falls regardless of how flat one plane is. Sliding requires the geometry to be skewed enough that the apex overhangs the face triangle's edge.

2.5 Strength model

Shear resistance on each joint plane uses the **Mohr–Coulomb** model:

$$\text{resistance} = \text{cohesion} \times \text{contact area} + \text{normal force} \times \tan(\text{friction angle})$$

- **Friction angle (°)** — applies to all joint sets (one global value).
- **Cohesion (kPa)** — also global. With cohesion = 0 (the default, and a common conservative choice for screening), FoS becomes *independent of wedge size* — only the angles matter. Persistence and span then affect the reported depth/weight but not the FoS. With cohesion > 0, smaller wedges get proportionally more help from cohesion, so FoS rises as size shrinks.
- **Rock unit weight (t/m³)** — converts wedge volume to weight (default 2.7 t/m³).

The analysis is **dry and gravity-only**: no water pressure, no clamping from the in-situ stress field, no seismic loading (see §10, Assumptions).

2.6 Orientation uncertainty: Fisher Monte-Carlo

Joint-set orientations are not single values — mapped poles scatter around the set mean. The standard statistical model for that scatter is the **Fisher distribution**, whose concentration parameter κ (**kappa**) plays the role of an inverse variance: large κ (say 100+) means tightly clustered poles; small κ (below ~20) means a diffuse set.

When you set **Monte-Carlo trials** > 0, the tool repeats the whole wedge analysis many times per location, each trial drawing a random orientation for every set from its Fisher distribution. Each trial

finds its own worst wedge, and the trial **fails** if that wedge's FoS is below your design FoS. From the trials the tool reports:

- **WdgPfail** — the percentage of trials that failed: an estimate of the probability that, given the orientation scatter, the worst credible wedge at this location does not meet your design FoS.
- **WdgDepth95 / WdgWeight95** — the 95th-percentile apex depth and weight **among the failing trials only** (§7.3 explains how to read these).

Two refinements feed the simulation:

- **Persistence** is applied in every trial, exactly as in the deterministic solve.
- **Weight** (relative occurrence, 0–1] makes a set present in only that fraction of trials. A set seen in only 30 % of mapped windows can be given Weight = 0.3; trials where it is "absent" can only form wedges from the remaining sets. A trial with fewer than three present sets forms no wedge and counts as a *non-failure* (it stays in the Pfail denominator).

Results are **reproducible within a run** (each location gets a deterministic random seed) but a new run draws a new seed, so Pfail will wobble by sampling noise between runs — see §7.3 for how many trials you need.

3. What you need before running

Requirement	How to get it	Mandatory?
Minode Markers with Dip / DipDir columns	Create 2 → <i>Minodes from development meshes</i> , with the size calculations enabled. The Dip/DipDir columns carry the local tunnel-axis orientation used to orient the back at each point.	Yes
Visible tunnel meshes	Load the development triangulations and make them visible — spans are measured by ray-casting against whatever meshes are visible (in CLIPPING mode, against the visible clipped fragments).	Yes
Mapping data with stereonet clusters	Load structural Mapping data; the stereonet panel auto-detects joint clusters. These prepopulate the joint-set table (mean pole, κ , pole count, and mean trace length as persistence).	No — you can type joint sets manually

The analysis runs on the **visible** Marker rows only. This is a feature: filter the Marker table (e.g. to one mining area) to restrict the analysis. Note, however, that a run resets the *Wdg\ columns of all** rows to blank before writing the visible ones — a partial run does not preserve results from an earlier run on other rows (re-run with the wider filter if you need them together).

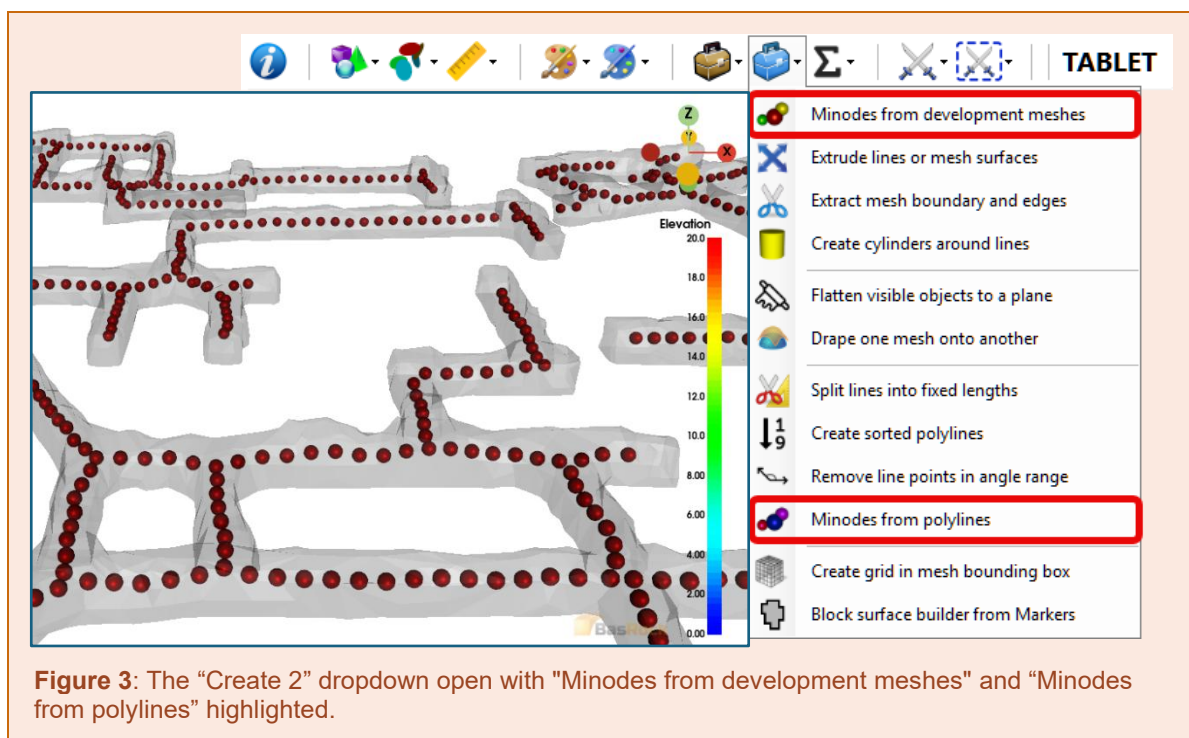


Figure 3: The “Create 2” dropdown open with “Minodes from development meshes” and “Minodes from polylines” highlighted.

4. Step-by-step

1. Load the development meshes and create (or load) the Minode Markers.
2. Optionally load Mapping data and check the stereonet's auto-clusters look sensible (the cluster-membership context menu on the stereonet lets you tighten or loosen what counts as a cluster).
3. Optionally filter the Marker table to the rows you want analysed.
4. Open **Calc** → **Wedge stability at Minodes**.
5. Review/edit the settings dialog (§5) and press **Analyse**.
6. Wait — the analysis runs in the background with the loading animation; a few thousand Minodes with 1 000 trials each typically takes seconds to a couple of minutes.
7. Read the summary dialog (how many rows were analysed vs skipped, the worst deterministic FoS, what columns were written).
8. Colour the Markers by a result column (e.g. `WdgFoS` or `WdgPfail`) and interpret (§7).

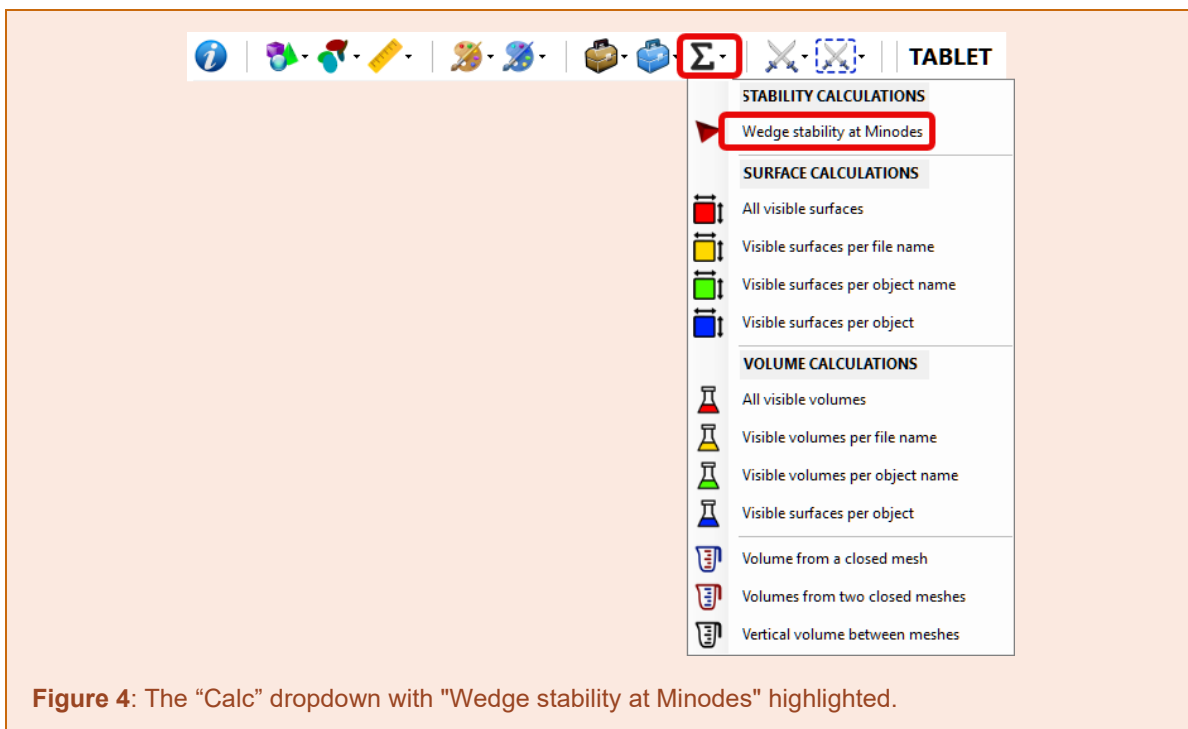


Figure 4: The “Calc” dropdown with "Wedge stability at Minodes" highlighted.

5. The settings dialog

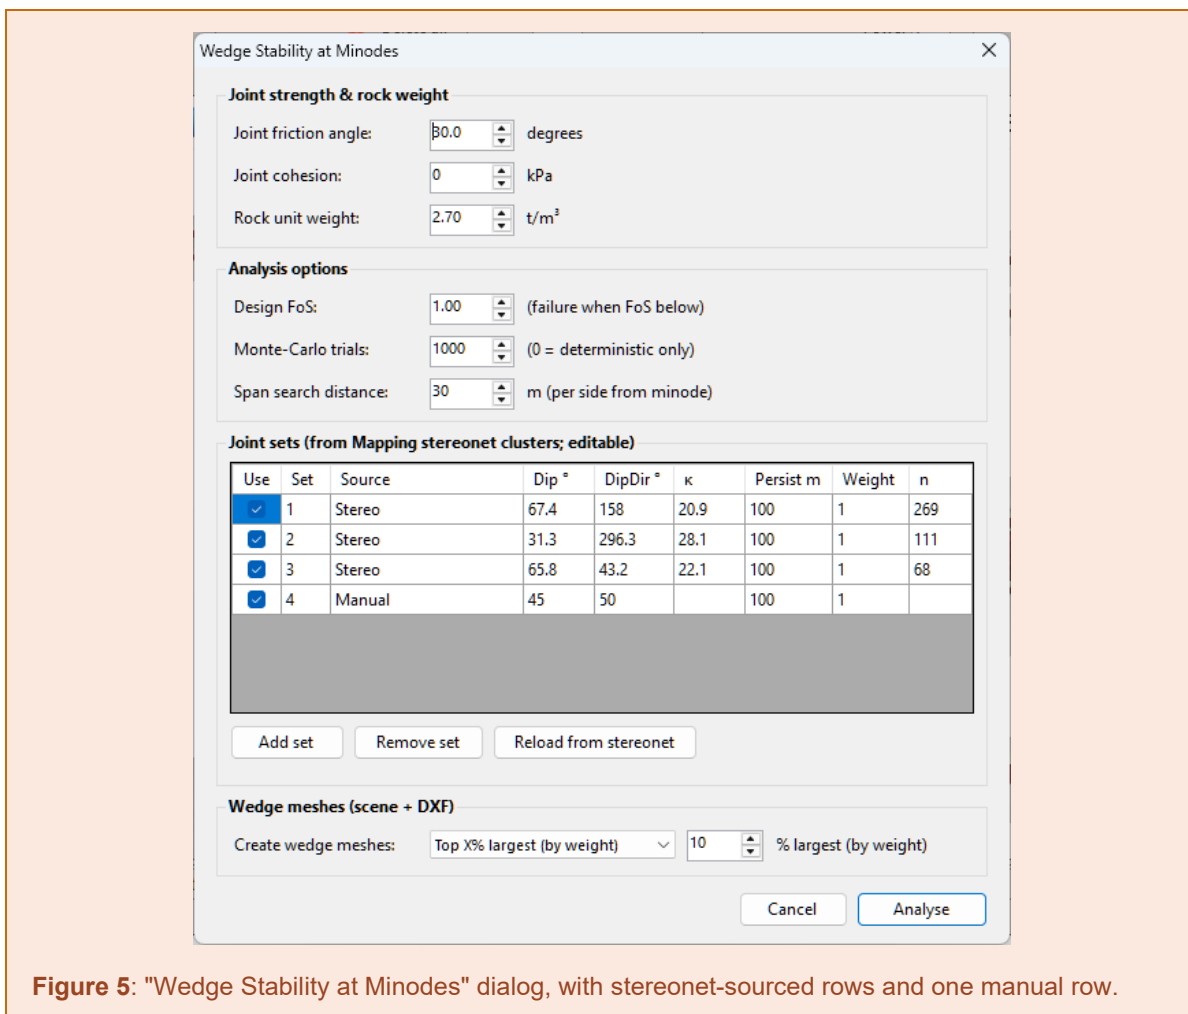


Figure 5: "Wedge Stability at Minodes" dialog, with stereonet-sourced rows and one manual row.

All settings are **sticky for the session** — the dialog reopens with your last values and table (including manual rows and edits). They are not saved across GEM4D restarts.

5.1 Joint strength & rock weight

Field	Default	Range	Notes
Joint friction angle	30°	5–89°	One global value for all sets.
Joint cohesion	0 kPa	0–1 000	0 = conservative, size-independent FoS (§2.5).
Rock unit weight	2.7 t/m ³	1–5	Converts volume to tonnes.

5.2 Analysis options

Field	Default	Range	Notes
Design FoS	1.0	0.1–10	The failure threshold for Monte-Carlo trials ("failure when FoS below"). Set it to your site acceptance criterion.
Monte-Carlo trials	1 000	0–100 000	0 = deterministic columns only (Pfail and the 95th percentiles stay blank). See §7.3 for choosing a count.
Span search distance	30 m	5–200	How far the span rays (and the back-presence probe) search from each Minode, per side . It also caps the reported span when a ray finds nothing. Use a value comfortably larger than your widest opening, but not so large that rays reach unrelated excavations.

5.3 Joint sets table

Prepopulated from the Mapping stereonet auto-clusters (largest clusters first, up to the 8-row cap), and fully editable:

Column	Meaning
Use	Tick = the set takes part in the analysis. At least 3 enabled sets are required.
Set	Display number (1, 2, 3 ...) — this is the digit used in the WdgSets result column. Rows renumber automatically when added/removed.
Source	Stereo (from a stereonet cluster), Stereo (girdle) (cluster flagged as girdle-shaped — its Fisher κ is unreliable, treat its scatter with suspicion), or Manual.
Dip ° / DipDir °	Mean plane orientation (dip 0–90, dip direction 0–360).
κ	Fisher concentration for the Monte-Carlo sampling. Blank = no scatter information — MC samples sit at the mean (effectively no orientation uncertainty for that set).
Persist m	Persistence: the set's maximum trace length on the back (metres). Blank = infinite (span-limited only). Prepopulated from the cluster's mean Mapping trace length (the Size column) when available.
Weight	Relative occurrence, greater than 0 and at most 1: the probability the set is present in each MC trial (§2.6). Blank = 1 = always present. The deterministic columns ignore weights — they always assume every enabled set is present (worst case).
n	Pole count of the source cluster (blank for manual rows). A set built from 8 poles deserves less confidence than one built from 300.

Buttons: **Add set** (a manual row, default 45°/000°), **Remove set** (or the Delete key), **Reload from stereonet** (replaces the Stereo rows with the current clusters, keeps Manual rows).

The table is capped at **8 sets** because every 3-set combination is solved at every location (8 sets = 56 combinations per Minode). A note under the table reports ignored smaller clusters and girdle warnings.

5.4 Wedge meshes (scene + DXF)

Option	Meaning
None (<i>default</i>)	No meshes — columns only.
FoS < design FoS	Build a mesh for every analysed Minode whose deterministic FoS is below the design FoS.
Top X% largest (by weight)	Build the X % heaviest governing wedges (X editable, default 10 %).
All analysed	One wedge per analysed Minode — readable on small runs, dense on big ones.

See §8 for what the meshes mean and their colour code.

6. Result columns reference

Nine numeric columns are written to the Marker table (created if absent, overwritten in place on re-runs, and CSV round-trip safe). Blank (NaN) means *not applicable* — the row was hidden, had no Dip/DipDir, found no mesh, or the statistic had no data (e.g. percentiles with zero failing trials).

Column	Units	Meaning
WdgSpan	m	The measured local back span the wedge was scaled to (§9).
WdgFoS	–	Deterministic Factor of Safety of the governing (lowest-FoS) wedge over all set combinations, at the mean set orientations. 0 = free fall; 99 = stable sentinel.
WdgMode	code	Failure mode of that governing wedge: 0 none · 1 fall · 2 plane slide · 3 edge slide · 4 stable (§2.4).
WdgSets	code	Which three sets form the governing wedge, encoded as their Set numbers in ascending order — e.g. 125 = sets 1, 2 and 5. (Kept numeric so colour-by-column and numeric filters work.)
WdgDepth	m	Apex depth (perpendicular height into the back) of the governing deterministic wedge.
WdgWeight	t	Weight of the governing deterministic wedge.
WdgPfail	%	Of the Monte-Carlo trials, the percentage whose worst wedge had FoS below the design FoS. Blank when trials = 0.
WdgDepth95	m	95th-percentile apex depth among the failing trials only . Blank when no trial failed.
WdgWeight95	t	95th-percentile weight among the failing trials only . Blank when no trial failed.

All columns behave like normal Marker data: colour the Markers by them, filter on them, add calculated columns from them, and export them to CSV.

	Eastng	Northng	0_Elevation	1_Dip	2_DpDr	3_ERF	4_CrRad	5_EstHR	6_VertZ	7_numPnts	WdgSpan	WdgFoS	WdgMode	WdgSets	WdgDepth	WdgWeight	WdgPfail	WdgDepth9	WdgWeight9
1	60200.76	9371.18	292.48	-2.22	176.61	2.94	2.67	2.7	5.27	6	10.35	0	1	123	0.88	1.99	100	1.28	2.79
2	60620.9	9195.95	293.96	6.78	188.98	3.41	4.39	3.12	6.47	6	10.06	0	1	123	0.76	1.42	100	1.25	2.53
3	60205.33	9324.47	294.45	-0.75	253.51	3.38	2.8	3.07	6.51	8	12.98	0	1	123	0.91	2.16	99.9	1.28	2.64
4	60187.63	9320.09	294.02	0.59	278.3	3.56	4.73	3.26	5.28	6	11.44	0	1	123	0.89	2.12	100	1.25	2.55
5	60173.11	9321.03	294.52	0.35	279.76	3.39	4.28	3.08	5.81	6	10.5	0	1	123	0.9	2.14	100	1.28	2.63
6	60204.78	9104.59	294.65	2.43	287.58	3.38	2.82	3.24	5.85	6	9.23	0	1	123	0.86	2.06	100	1.27	2.58
7	60198.49	9106.58	294.37	-2.62	273.71	3.43	3.46	3.24	5.63	7	9.7	0	1	123	0.93	2.12	99.9	1.32	2.59
8	60220.33	9303.8	294.62	1.51	347.65	3.24	3.2	3.08	5.29	6	10.56	0	1	123	0.89	2.07	100	1.27	2.6
9	60176.71	9109.65	295.19	-0.19	283.24	3.34	3.57	3.27	5.61	8	7.4	0	1	123	0.91	2.17	100	1.26	2.7
10	60217.57	9104.65	295.41	3.42	262.86	3.65	4.92	3.32	6.7	6	11.28	0	1	123	0.81	1.83	100	1.25	2.65
11	60239.73	9158.99	295.19	-0.1	1.13	4.19	6.02	3.78	5.85	8	12.23	0	1	123	0.9	2.15	100	1.33	2.71
12	60205.92	9171.98	296.52	13.48	319.72	0.77	0.3	0.79	3.54	8	7.14	0	1	123	0.6	1.12	100	1.11	2.23
13	60185.7	9247.76	295.52	-1.83	23.51	2.89	1.91	2.69	4.93	6	10.31	0	1	123	0.86	1.92	100	1.22	2.62
14	60239.59	9206.38	296.76	0	359.95	2.74	2.66	2.99	5.95	8	5.67	0	1	123	0.9	2.16	100	1.29	2.57
15	60210.59	9216.64	295.76	0	0	0.79	0.79	0.88	72	0	2.95	0	1	123	0.9	2.11	100	1.21	2.48
16	60239.69	9249.7	296.17	0.1	358.56	3.98	4.37	3.62	5.53	8	10.22	0	1	123	0.9	2.17	100	1.26	2.65
17	60172.55	9240.53	296.05	-1.15	0.05	3.08	3.25	3.36	5.28	6	7.15	0	1	123	0.88	2.03	100	1.27	2.55
18	60178.2	9081.47	295.99	-8.97	201.05	3.49	4.04	3.17	5.45	6	10.58	0	1	123	0.8	1.42	100	1.27	2.53
19	60174.23	9169.34	297.11	7.66	354.23	3.25	4.18	2.96	7.36	6	9.88	0	1	123	0.77	1.48	100	1.25	2.56
20	60189.59	9232.7	296.49	3.39	350.5	2.95	3.78	2.72	6.44	5	9.74	0	1	123	0.85	1.89	100	1.26	2.55
21	60239.43	9103.46	297	-1.36	265.47	3.22	3.86	2.94	6.62	6	9.87	0	1	123	0.92	2.15	100	1.3	2.59
22	60250.61	9371.81	298.65	3.41	82.1	3.51	4.71	3.21	6.92	6	11.5	0	1	123	0.92	2.04	100	1.28	2.75
23	60274.03	9101.65	298.44	3.45	268.23	3.35	3.75	3.15	5.87	6	10.36	0	1	123	0.82	1.86	100	1.25	2.53
24	60281.15	9097.68	298.95	2.94	304.3	3.31	3.33	3.11	5.7	6	10.99	0	1	123	0.87	2.11	100	1.24	2.57
25	60268.09	9372.2	299.9	-7.11	69.88	3.49	4.47	3.16	6.37	6	11.76	0	1	123	0.68	1.3	100	1.15	2.37
26	60307.21	9380.87	300.78	3.27	173.63	3.82	4.98	3.58	5.96	6	11.98	0	1	123	0.86	1.85	100	1.34	2.64

Figure 6 — Capture: the Marker FlexGrid scrolled to show the nine Wdg\ columns after a run.*

7. Interpreting the results

7.1 The deterministic triplet: WdgFoS, WdgMode, WdgDepth, WdgWeight

These describe **one specific wedge**: the worst (lowest-FoS) combination of sets, solved at the cluster *mean* orientations, with every enabled set assumed present. Read them together:

- **WdgFoS vs your design FoS** is the primary screen. Colour by WdgFoS with a legend band at your design value.
- **WdgMode** tells you *why*. Mode 1 (fall) zones have zero frictional reserve — support is the only thing holding such wedges, and bolt design is governed by dead weight. Mode 2/3 zones retain frictional resistance, and FoS responds to the friction angle you chose — worth a sensitivity check.
- **WdgDepth and WdgWeight** size the hazard and the support. A 0.4 m, 2 t wedge and a 4 m, 300 t wedge can have the same FoS but completely different support implications (bolt length must pass the apex depth with anchorage to spare; plate/strand capacity must carry the weight).
- **WdgSets** tells you *which* sets to blame. If one set number appears in nearly every governing triplet across an area, that set's orientation and persistence deserve the most mapping attention — and the most scepticism (check its n and k in the dialog).

Remember the deliberate conservatisms: the wedge is assumed to exist at every point, at its maximum size, fully formed, and dry. Deterministic FoS < design over large areas is therefore an instruction to *look closer*, not proof of imminent failure.

7.2 WdgPfail — probability of failure

WdgPfail answers a different question from WdgFoS. The deterministic FoS uses mean orientations; Pfail asks: *given how scattered the mapped orientations actually are (κ), and how often each set occurs (Weight), in what fraction of plausible realities does the worst wedge here fail the design FoS?*

Practical readings:

- **High FoS + low Pfail** — comfortably fine.
- **Low FoS + high Pfail** — the mean-orientation wedge fails and most perturbed versions do too. A genuine, robust problem area.
- **FoS above design but Pfail noticeably > 0** (say 5–20 %) — the *mean* geometry is fine but the scatter reaches failing configurations. Common with diffuse sets (low κ). The deterministic columns alone would under-call these locations.

- **FoS below design but moderate Pfail** (say 30–60 %) — the mean is unfavourable, but many perturbations escape. Often a marginal geometry sitting near a mode boundary.

Because the deterministic solve always assumes all sets are present, while trials with sub-unity Weights sometimes drop sets, **Pfail can be low even where the deterministic FoS is poor** when the offending set is rare. That is exactly the risk-weighting the Weight column exists for.

There is no universal acceptance threshold for Pfail — it is a *screening discriminator*, not a calibrated annual probability. Within one run with consistent inputs, ranking locations by Pfail is meaningful; comparing absolute Pfail values against another mine's criteria is not.

7.3 WdgDepth95 / WdgWeight95 — how big, when it fails

These are the 95th-percentile apex depth and weight computed **only over the failing trials**. They pair *severity* with *likelihood*:

- WdgPfail says how often things go wrong; WdgDepth95/WdgWeight95 say **how bad the bad cases are** — a near-worst-case size among the wedges that actually fail the design FoS, ignoring the harmless trials.
- They can legitimately **exceed** the deterministic WdgDepth/WdgWeight: a perturbed orientation can produce a steeper, deeper wedge than the mean geometry. Use Depth95 (not WdgDepth) when checking that bolt lengths pass the credible failing wedges.
- They are **blank when nothing failed** — no failing trials means no failing-size statistic, which is itself good news.
- Interpret them alongside Pfail: a 50 t Weight95 with Pfail = 0.4 % is a rare tail case; the same Weight95 with Pfail = 45 % is a design driver.
- Statistical health warning: the percentile is computed from $Pfail \times trials$ samples. With 1 000 trials and Pfail = 1 %, that is 10 samples — the "95th percentile" of 10 numbers is just their second largest, and it will jump between runs. If you intend to *design from* Depth95/Weight95 at low-Pfail locations, raise the trial count (10 000+) until the values stabilise.

7.4 A suggested screening workflow

1. Run with defaults (cohesion 0), your site design FoS, and 1 000 trials.
2. Colour by WdgFoS → identify the broad problem areas and the governing modes/sets.
3. Colour by WdgPfail → find the scatter-driven locations the deterministic map missed, and de-prioritise deterministic-only hits caused by rare sets.
4. Filter to the worst cells and read WdgDepth/WdgDepth95 against your standard bolt length, WdgWeight/WdgWeight95 against support capacity.
5. Re-run the handful of critical locations in a dedicated single-location tool (e.g. UNWEDGE) with site-specific strength data, water, and support, before changing any design. Expect the *mode and FoS* to agree; expect *sizes* to differ where the other tool scales wedges to the full excavation perimeter rather than the measured local span.


8. Wedge meshes

When a mesh mode is selected (§5.4), the governing deterministic wedge of each selected Minode is built into the scene, anchored with its face triangle on the measured back (at the point where the vertical-ish probe from the Minode hit the mesh). All wedges of one failure mode are consolidated into **one mesh actor per mode**:

Actor	Mode	Colour
WedgesFall	1 — fall	Red
WedgesPlaneSlide	2 — plane slide	Orange
WedgesEdgeSlide	3 — edge slide	Yellow
WedgesStable	4 — stable	Green

They are first-class Create meshes: you are prompted to save them as DXF (one save dialog per run; cancelling skips mesh creation but keeps the columns), and they appear in the object list with normal visibility/colour controls — so they can go straight into design plots or be reloaded later.

Minodes whose back probe found no mesh hit cannot be anchored and are skipped (the summary dialog counts them).



9. How the local span is measured

The wedge size depends on the local opening size, which GEM4D measures per Minode with a horizontal **azimuth fan** of ray pairs:

- 12 pairs of opposite horizontal rays are cast every 15°, out to the span search distance per side.
- A pair only counts if **both** rays hit a mesh — a miss means that direction runs into open excavation (e.g. down a crossing drive), which would not bound a wedge.
- The span is the **smallest** fully bounded pair total. Geometrically this approximates the diameter of the largest circle that fits in the local opening: a straight drive returns its width (even when the Minode sits off-centre), and an intersection returns its diagonal span rather than a runaway measurement down the crosscut.
- If no pair is fully bounded, a fallback uses the perpendicular-to-drive pair with misses capped at the search distance.

- A separate probe along the local back normal confirms a back exists at all (and anchors the wedge mesh, §8). Rows with no mesh contact at all are skipped with a count in the summary.

Implications: spans (and therefore wedge sizes) are largest at intersections — which matches reality and is where the red wedges usually appear. Check `WdgSpan` itself when a result looks odd; an unexpectedly large span usually means the rays escaped through a hole or a missing mesh in that area.

10. Assumptions and limitations

Keep these in mind whenever results feed a decision:

1. **Back (roof) wedges only.** The analysed face is the excavation back, following the local tunnel grade. Sidewall wedges are not analysed (a known candidate for a future version).
2. **Tetrahedral wedges only.** Each wedge is formed by exactly three joints plus the face. Truncated/composite blocks bounded by 4+ joints are not modelled; the tetrahedron is the bounding (largest, most conservative) case.
3. **Gravity-only, dry.** No water pressure, no in-situ/induced stress clamping (which usually *helps* shallow wedges and can *hurt* under low-stress backs), no blast or seismic loading.
4. **Ubiquitous joints.** Wedges are assumed possible everywhere (§2.2) — occurrence is overstated by design. The Weight column is the only occurrence control, and it is statistical, not spatial.
5. **One global strength.** A single friction/cohesion pair applies to all sets.
6. **Mean-pole sets.** Each set enters as a single mean plane (plus Fisher scatter in MC). Local orientation anomalies that mapping hasn't captured are invisible to the analysis.
7. **Persistence caps the face traces** of the wedge, not the apex-edge lengths — the UNWEDGE joint-length-scaling convention.
8. **Span is measured horizontally** at the Minode elevation; very irregular profiles (e.g. tall stope brows) may need a sanity check of `WdgSpan`.
9. **Screening tool.** Results rank and localise risk; they do not replace site standards, detailed wedge analysis with support, or geotechnical judgement.

11. Messages and troubleshooting

Message / symptom	Cause and fix
"Load Minode Markers first..."	The Marker table is empty. Create Minodes via Create 2 → Minodes from development meshes (with size calculations enabled).
"The Marker table has no Dip / DipDir columns..."	The loaded Markers lack the tunnel-axis orientation columns. Re-create the Minodes with the size calculations enabled, or load a Marker CSV that has Dip and DipDir columns.
"Wedge analysis needs at least 3 enabled joint sets..."	Tick more rows, add manual sets, or <i>Reload from stereonet</i> in the dialog.
"No visible mesh geometry to measure the tunnel spans against..."	Make the tunnel meshes visible (in clipping mode, the visible clipped fragments are used).
"No visible Marker rows to analyse."	Your Marker filter hides everything. Relax the filter.
"Cannot write the wedge columns: an existing non-numeric column named '...' is in the way."	A text column already uses one of the Wdg* names (typically from a hand-edited CSV). Rename or delete it and re-run.
Summary reports many "Skipped — no mesh within X m" rows	Those Minodes found no mesh within the search distance in any direction — often stale Minodes from a deleted mesh, or a too-small search distance.
wdgPfail differs slightly between identical runs	Expected Monte-Carlo sampling noise (each run draws a fresh seed). Increase the trial count for stabler values.
Wedges look absurdly large at one spot	Check wdgSpan there — span rays probably escaped through a gap/hole in the mesh or into an unmeshed void. Repair or hide the offending geometry.
A girdle-flagged set (Stereo (girdle)) dominates results	Girdle clusters have unreliable Fisher statistics; consider replacing it with manual sets that represent the structure properly.

12. References

- Hoek, E. & Brown, E.T. (1980). *Underground Excavations in Rock*. IMM, London. (Wedge limit-equilibrium procedure.)
- Goodman, R.E. (1989). *Introduction to Rock Mechanics*, 2nd ed. Wiley. (Block kinematics and limit equilibrium.)
- Goodman, R.E. & Shi, G. (1985). *Block Theory and its Application to Rock Engineering*. Prentice-Hall. (Removable-block theory underpinning the mode classification.)
- Fisher, R.A. (1953). Dispersion on a sphere. *Proc. Royal Society A*, 217, 295–305. (Orientation statistics used in the Monte-Carlo sampling.)
- Rocscience UNWEDGE documentation — the joint-length (persistence) scaling and span-scaling conventions GEM4D follows.

Revision history

Date	GEM4D	Change
2026-06-12	1.9.0	First issue — covers v2 of the feature (editable joint-set table, persistence capping, occurrence weights, azimuth-fan span, mode-coloured wedge meshes).