Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A managed, uncertainty-aware pipeline from ground-penetrating radar to dimension-stone block yield in fractured quarries

Libish Murugesan

libish.1234@gmail.com

Alfaisal university

Research Article

Keywords: ground-penetrating radar, dimension stone, fracture detection, uncertainty propagation, guillotine packing, diamond wire saw, manufacturability, cutting-surface area, quarry yield, tolerance ladder

Posted Date: June 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10035624/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

A managed, uncertainty-aware pipeline from ground-penetrating radar to dimension-stone block yield in fractured quarries

Libish Murugean
Independent Research libish.1234@gmail.com

2026-06-09

Abstract

Dimension-stone quarries lose most of their extracted volume to fractures, and the fractures that matter most are deep, where ground-penetrating radar (GPR) is least certain. We present an end-to-end, fully managed (Rhino/Grasshopper) pipeline that maps GPR fracture picks to three-dimensional fracture surfaces, propagates a per-location *tolerance ladder* (reconstruction, interpolation and meshing uncertainty) to a confidence metric, adds a literature-calibrated *detection* rung (the probability a fracture is seen at all), and packs dimension blocks into the fracture-bounded intact zones. On a real Botticino-marble grid the position uncertainty is 0.065 m, 0.132 m and 0.278 m at 0.72 m, 1.96 m and 3.37 m depth. The detection model is calibrated to reproduce ≈74–79%, bracketing a measured ≈80% rate for granite open sub-horizontal fractures; detection and extraction were exercised on both marble and granite GPR. Requiring blocks to clear fractures by their measured, depth-dependent uncertainty collapses recoverable yield from 24% to 6.5% on this grid; scaling the clearance to the per-fracture σ, and computing its yield cost, is the contribution. An evolved mesh-bench packer raises geometric yield from 24% to 53.3% on this grid (matching a kerf-free idealisation while remaining kerf-correct). Under the on-site diamond *wire saw* (guillotine) constraint a three-stage packer reaches 49.8% yield at 100% saw-separability. A cost-versus-volume frontier shows bed-following (oblique) cuts recover +47–59% net over flat cuts on the ≈ 6° dipping beds (the *georeferencing prize*, which scales with dip and vanishes for flat-lying deposits). All quantitative yield results are single-site benchmarks pending multi-site and multi-lithology validation; the literature base and research gaps are reported in a PRISMA frame.

Keywords: ground-penetrating radar; dimension stone; fracture detection; uncertainty propagation; guillotine packing; diamond wire saw; manufacturability; cutting-surface area; quarry yield; tolerance ladder.

1 Introduction

Profitability of a dimension-stone quarry depends on the size and number of intact, marketable blocks that can be cut from the rock, which is controlled by the three-dimensional pattern of the discontinuity system (Mosch et al. 2010; Yavuz et al. 2005). GPR is, among surface geophysical methods, the one that resolves individual metre-scale fractures at quarry depth (Porsani et al. 2006; Grasmueck et al. 2005), but its fracture-location accuracy degrades with depth, and a fracture that is missed entirely (sealed, or steeply dipping relative to a surface scan) carries no measured uncertainty yet is the dominant yield risk (Molron et al. 2020; Dorn et al. 2012). Existing block-extraction optimisers place cuts *on* mapped discontinuities and maximise block volume (Yarahmadi et al. 2018; Ülker and Turanboy 2009; Kapageridis 2018); none scales a safety clearance to the *measured* per-fracture uncertainty, and the one ornamental-stone GPR recovery study that proposes any offset uses a generic constant "some centimetres" (Elkarmoty et al. 2017).

Contributions. (i) A managed tolerance ladder that propagates GPR reconstruction, geostatistical interpolation and meshing uncertainty to a per-location confidence; (ii) a literature-calibrated detection rung and detection-adjusted confidence, with per-stone base efficiencies; (iii) an uncertainty-safe block clearance equal to the measured fracture sigma, and its quantified yield cost; (iv) an evolved mesh-bench dimension-block packer that lifts geometric yield from a 24% fixed-grid baseline to a best-of-tested 53.3% (voxel-DLBF multi-size; the global-grid combined mode it supersedes reached 49.3%), topping two external packers head-to-head; (v) a diamond-wire-saw *manufacturable* variant: a three-stage guillotine packer (49.8% yield, 100% saw-separable) with the Jalalian I11 saw-cost metric and a tested multi-scale crack-aware recovery cascade; (vi) a verified PRISMA/SLM/ROSES literature base and an explicit research-gap accounting.

Novelty delta vs the closest precedents.

Precedent	What it does	What this work adds (measured)
Elkarmoty et al. (2017)	ψ-only pose, generic constant offset	full 3D pose + σ-scaled clearance; yield 24%→6.5%
Mosch et al. (2010)	natural-block size enumeration, no kerf	kerf-correct uniform-block mesh-bench yield, 53.3%
Jalalian et al. (2023)	I11 cut-area as a 2D Pareto axis	I11 on 3D wire-saw guillotine layouts, 49.8% at ϕ=100%
Yarahmadi et al. (2018)	conditional two-scale value optimiser	recursive multi-scale cascade, exact S=1 reduction

2 Literature review (ROSES synthesis)

We follow a three-tier review framework: PRISMA for the counted evidence statistics, SLM (systematic literature mapping) for per-algorithm math and code, and ROSES for the interdisciplinary synthesis. A rapid, verified-only search across four strands (GPR depth uncertainty; fracture-surface interpolation / discrete-fracture-network uncertainty; dimension-stone yield and fracture allowance; uncertainty propagation and decision) plus a per-lithology detection-calibration strand yielded 38 included sources, every value grounded in full text or a verified abstract with a Crossref DOI; extrapolations are flagged.

Evidence map (by discipline). Geophysics establishes that GPR fracture-location error grows with depth (velocity-to-depth term (Porsani et al. 2006; Xie et al. 2021), the $\lambda/4$ resolution floor (Annan 2009; Neal 2004), the detection dip cut-off and openness sensitivity (Molron et al. 2020; Dorn et al. 2012)). Geostatistics supplies the between-scan-line interpolation variance (kriging / Gaussian-process posterior std, zero at data and growing in the gaps (Cressie 1993; Rasmussen and Williams 2005; Lei et al. 2017)). Rock mechanics relates the intact block size to discontinuity spacing through the volumetric joint count and block volume (Palmstrom 2005); in open-pit mining, by contrast, block-model size is set by equipment and logistics (bench height, shovel width, shovel–truck efficiency) (Guo et al. 2024), whereas the dimension-stone problem here is set by per-fracture clearance and saw separability. Fabrication economics quantifies the low, fracture-driven recovery and the optimisation levers (Yarahmadi et al. 2018; Ülker and Turanboy 2009; Kapageridis 2018; Jalalian et al. 2023). Decision theory supplies the quadrature combination and coverage factor (the GUM (JCGM 2008)) and the value-of-information argument (Eidsvik et al. 2015). The block-packing tree algorithm (Kim 2025) sits on the fabricator side.

The gap the corpus reveals. No block-cutting optimiser imposes a geometric standoff scaled to

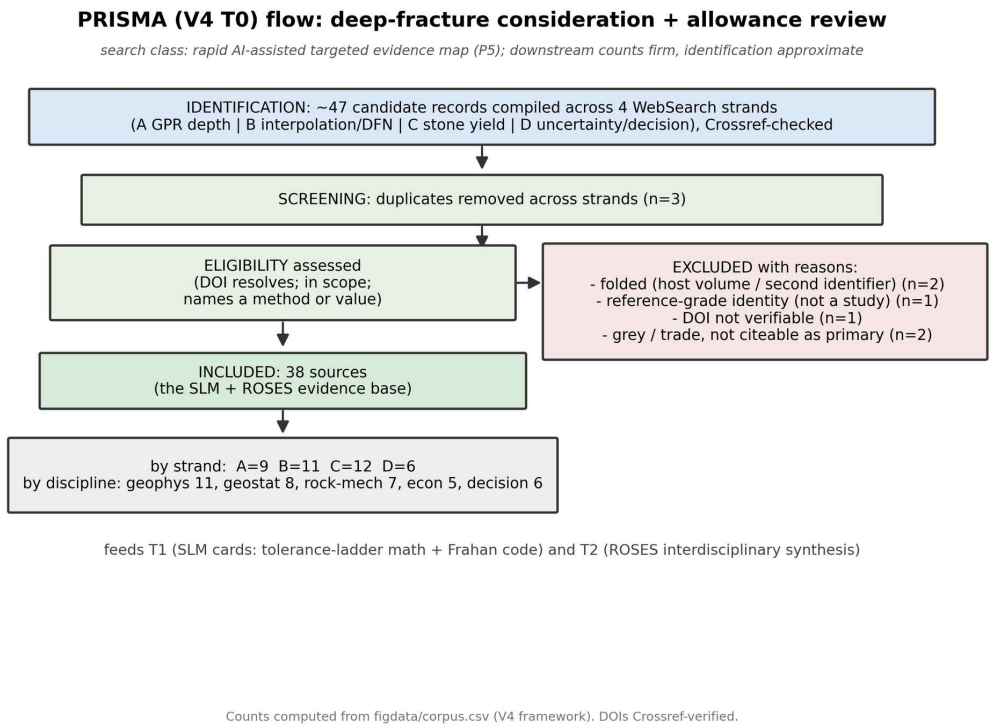


Figure 1: PRISMA identification-to-inclusion flow for the corpus (own work). Downstream counts are firm; identification is a rapid AI-assisted targeted map, not an exhaustive census.

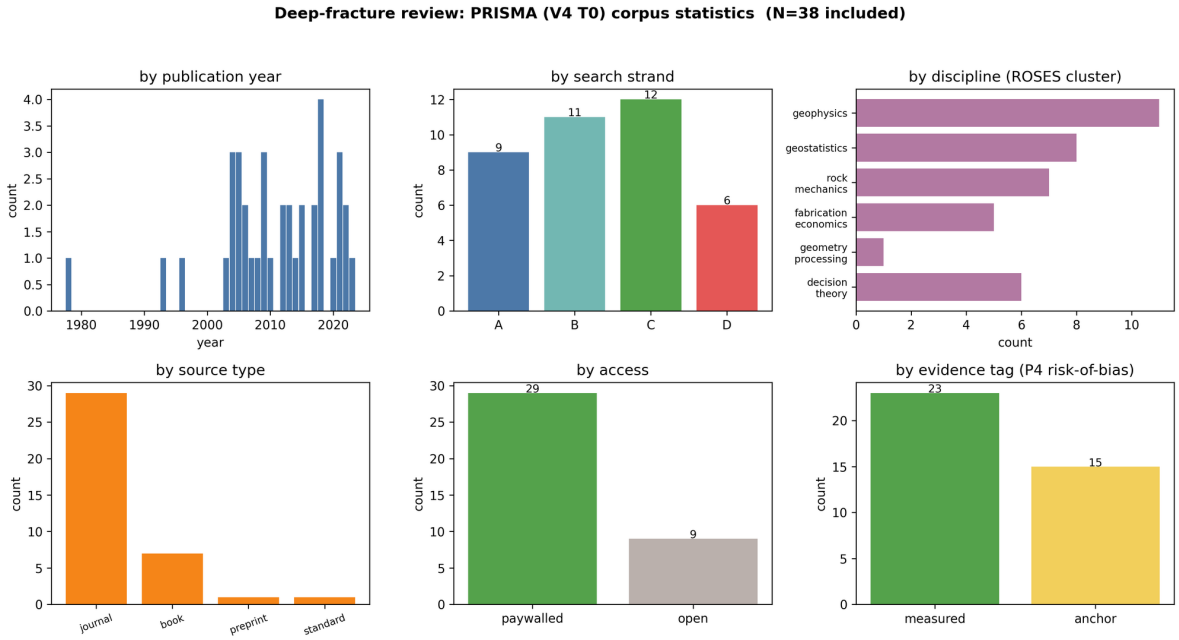


Figure 2: Corpus statistics (own work): by search strand, discipline, year, source type, access, and evidence tag (measured vs method-class anchor). The evidence base spans six disciplines and is majority measured (23 of 38).

the measured uncertainty of each fracture; the closest precedent (Elkarmoty et al. 2017) proposes a generic offset and was never executed for a yield number. This is the gap the present pipeline closes.

3 Methods

3.1 3D fracture-surface construction and meshing backends

The 3D fracture model is built from the raw GPR before any uncertainty is propagated. Each radargram is read from the field-instrument files with the open-source RGPR package (Huber and Hans 2018) and then processed by `RadargramProcessor` (dewow, background removal, time-power gain, automatic gain control, Stolt frequency–wavenumber (f–k) migration, then a Hilbert instantaneous-amplitude envelope), and the high-energy reflectors are picked by `FractureExtractor` (local maxima above the 0.985 energy quantile, kept only under a dip-aware lateral-continuity rule). Stolt migration remaps the data from angular frequency ω to vertical wavenumber k_z through the constant-velocity dispersion relation (exploding-reflector form)

$$k_z = \sqrt{\left(\tfrac{2\omega}{v}\right)^2 - k_x^2}, \qquad z = \tfrac{1}{2}\, v\, t, \tag{1}$$

with t the two-way travel time, v the (one-way) radar-wave velocity and k_x the horizontal wavenumber; the instantaneous energy $E(t) = |s(t) + i\,\mathcal{H}\{s\}(t)|^2$ (the squared envelope, $\mathcal{H}$ the Hilbert transform) makes the reflectors pick-able. The picked reflectors, interpreted as bedding-parallel discontinuities (partings) in the Botticino limestone, are clustered into beds, and each bed is fitted to a dipping plane $z = a + b_x x + b_y y$ by least squares, with dip $\delta = \arctan\sqrt{b_x^2 + b_y^2}$ and dip direction $\beta = \operatorname{atan2}(b_y, b_x)$ (up-dip gradient direction, math convention). Between the sparse survey lines the bed surface is densified by ordinary kriging, $\hat{z}(\mathbf{x}_0) = \sum_i \lambda_i\, z(\mathbf{x}_i)$, the weights λ_i solving the kriging system built from the fitted spatial covariance, and the continuous surfaces are triangulated to meshes.

All meshing and solid operations run through native computational-geometry libraries with a managed fallback. Surface reconstruction from picks and point clouds uses *geogram* (Lévy 2024), which bundles the Kazhdan screened-Poisson reconstruction (Kazhdan and Hoppe 2013, `GEO::PoissonReconstruction`), and geogram remeshing and hole-filling produce clean two-manifold surfaces. Block–fracture intersection and the digital-ravalement booleans use *CGAL* (The CGAL Project 2023; EPICK/Hybrid kernels, `frahan_cgal.dll`), with a managed binary-space-partition boolean as the fallback when the native kernel is unavailable. These backends keep every cut block a watertight closed solid, the precondition for the volume and surface-area economics of Section 4.7.

3.2 Tolerance ladder (position uncertainty)

Per surface location, three independent stages combine in quadrature (the GUM law of propagation of uncertainty (JCGM 2008)):

$$\sigma_{\text{total}} = \sqrt{\sigma_{\text{recon}}^2 + \sigma_{\text{interp}}^2 + \sigma_{\text{mesh}}^2}. \tag{2}$$

The reconstruction term carries a depth-growing velocity term, a vertical-resolution floor, and a time-zero term (Xie et al. 2021):

$$\sigma_{\text{recon}} = \sqrt{\left(d\,\frac{\sigma_v}{v}\right)^2 + \left(\frac{\lambda}{4}\right)^2 + \left(\frac{v\,\sigma_{t_0}}{2}\right)^2}, \quad \frac{\sigma_v}{v} = \frac{1}{2}\frac{|\delta\varepsilon_r|}{\varepsilon_r}, \quad \sigma_{t_0} = \frac{(t_E - t_A)/2}{\sqrt{3}}. \tag{3}$$

The two-term form (velocity + $\lambda/4$) is the shipped baseline; the time-zero component $v\sigma_{t_0}/2$ was added as a later refinement after Xie 2021, which shows time-zero dominates at shallow depth while $\lambda/4$ is a vertical-resolution *floor* (not a co-equal driver) and the velocity term leads at quarry depth. The $\lambda/4$ term is an instrumental vertical-resolution limit, not a statistically estimated variance; we enter it into the quadrature as a deliberately conservative one-sigma floor (it can only inflate, never deflate, σ_{recon}), following the standard quarter-wavelength GPR resolution bound (Grasmueck et al. 2005; Annan 2009). The grid3 results below use the two-term form (time-zero set to zero); the time-zero term is reported separately. This two-term form is appropriate for quarry-scale depth ($> 1\,\text{m}$, where the velocity term dominates – Porsani's ±8.5–9.5% at 25 m); the time-zero term would be required only at shallow depth ($< 0.5\,\text{m}$), outside the quarry-bench scope. σ_{interp} is the ordinary-kriging posterior std (managed Cholesky solver), $\sigma_{\text{mesh}} = (h^2/8)\kappa$ the triangulation chord error. The independence of the three stages (different physics) is assumed for the quadrature; an empirical correlation study is future work. The optimisation metric is the confidence within a cutting tolerance T:

$$C_{\text{pos}} = \operatorname{erf}\!\left(\frac{T}{\sigma_{\text{total}}\sqrt{2}}\right) = \mathbb{P}(|\text{deviation}| \le T). \tag{4}$$

3.3 Detection rung (whether a fracture is seen)

A detection rung was added to address the position ladder's implicit assumption that every fracture is observed: a missed fracture (sealed, or steeply dipping relative to a surface scan) has no σ but is the real risk. We add, calibrated to (Molron et al. 2020) (∼80% of open sub-horizontal granite fractures, *measured*) and (Dorn et al. 2012) (∼91% transmissive granite):

$$P_{\text{det}} = \eta_0\, p_\theta(\theta)\, p_{\text{open}}\, p_A(A), \qquad p_A(A) = \frac{A}{A + A_{\text{min}}}, \quad A_{\text{min}} = \tfrac{\lambda}{4}\tfrac{d}{2}, \tag{5}$$

with $p_\theta = 1$ for dip $\le 25°$ smoothly falling to 0.1 by 75°, $p_{\text{open}} = 1$ (open) or ≈0.1–0.2 (sealed), and a stone-specific base efficiency η_0 (Fig. 5). $A_{\text{min}} = (\lambda/4)(d/2)$ is a depth-aware Fresnel-fraction floor (the footprint widens with depth), calibrated to reproduce Molron's ∼80% for $1\,\text{m}^2$ to $10\,\text{m}^2$ fractures at 160 MHz to 750 MHz. The factor shapes p_θ and p_{open} are *engineering parameterisations* fixed to these measured anchors (Molron et al. 2020; Dorn et al. 2012), not independently derived functional forms; only η_0 and A are stress-tested (Section 4.3), and a curve-shape sensitivity is left to future work. The detection-adjusted (effective) confidence treats "detected *and* within tolerance" as approximately independent events, so the area-mean product $C_{\text{eff}} = C_{\text{pos}} \cdot \overline{P_{\text{det}}}$ is the joint coverage probability; this is the honest yield-safety metric.

3.4 Block-packing tree and the mesh-bench packer

The block-packing tree (Kim 2025) grows one tree per container; placing an element at a leaf splits the slab into three guillotine sub-slabs (one of six orderings), and the forest score is

$$S_k = \begin{cases} \sum p^E_{\text{packed}} & (E_k^{\text{packed}} \cap E) \subsetneq E \\ \sum p^E + \left(1 - \sigma(\sum p_k^{C*})\right) & (E_k^{\text{packed}} \cap E) = E \end{cases}, \quad \sigma(x) = \frac{x}{1 + |x|}, \tag{6}$$

selecting $\arg\max_k S_k$ (the published Eq. (9) prints arg min, an erratum: the prose and the cheap-container bonus both require arg max). Our mesh-bench problem differs: pack *uniform* marketable blocks into *irregular* fracture-bounded slabs. We combine the tree packer's rotation and best-of-N search with a heightmap-style irregular fit (a parity in/out test on the closed slab mesh), and evolve a combined multi-size fill.

Best-performing mesh-bench packer (derived; the new baseline). A head-to-head (Section 4.4) showed a per-block voxel placement beats a global grid, so the best-performing packer places each block independently on a voxel lattice of pitch h. Let M be the closed slab mesh and $L = (L_x, L_y, L_z)$ the block.

$$\text{occupancy: } O[i,j,k] = \mathbb{1}\big[\, c_{ijk} \in M \,\big], \quad c_{ijk} = \mathbf{min}_M + \big(i+\tfrac{1}{2},\, j+\tfrac{1}{2},\, k+\tfrac{1}{2}\big)\, h, \tag{7}$$

$$\text{cell span (size } f,\ \text{orientation } o \in R_6): \quad (w_x, w_y, w_z) = \big\lceil (f\, L_o + \text{kerf})/h \big\rceil, \tag{8}$$

placement (deepest-bottom-left-first, DLBF; scan k,j,i ascending): place at the first free cell where

$$\bigwedge\nolimits_{(p,q,r)\in \text{span}} O[p,q,r] = 1 \quad \wedge \quad \Phi(B,\, Cl), \qquad \text{then set those } O = 0, \tag{9}$$

$$\text{strict irregular fit: } \Phi(B, Cl) = \mathbb{1}\big[\, \text{centre}(B) \in M \ \wedge\ \forall\, 8 \text{ corners of } B^{+Cl} \in M \,\big], \tag{10}$$

$$\text{multi-size: } f \in \{1,\, 0.66,\, 0.5\} \text{ largest-first, sharing } O; \qquad Y = \frac{\sum_B \text{vol}(B)}{\sum_s \text{vol}(s)}. \tag{11}$$

R_6 is the six axis rotations (Kim's set); B is the placed block of size fL_o at the cell corner; the uncertainty-safe variant sets the clearance $Cl = \sigma_{\text{total}}$. Each evolution broadly subsumes the previous, so yield is empirically monotone non-decreasing across modes on this grid (not a strict nesting at every transition).

3.5 Wire-saw manufacturability: guillotine cutting, saw cost, and the recovery cascade

The voxel-DLBF packer maximises volume but places each block at an independent lattice cell, so its layout is not separable by full-span cuts. On site the only primary tool is a diamond *wire saw*: a taut bead wire sweeps a planar (gently ruled) surface in a single edge-to-edge traverse between two pilot holes or free faces, with kerf 7 mm to 11 mm (Konstanty 2021) and a mandatory horizontal-bed-cut-first sequence; it cannot make interior pockets, blind cuts or re-entrant corners (Konstanty 2021). These constraints are exactly the operations-research *guillotine* cut (Gilmore and Gomory 1965). A layout is wire-saw manufacturable iff it is guillotine-separable: every block is isolable by a recursive sequence of full-span axis-aligned planes, none crossing a block interior. We report the greedy guillotine-separable fraction

$$\phi = \frac{\#\{\text{blocks isolated by a full-span cut sequence}\}}{N} \in [0,1], \tag{12}$$

$\phi = 1$ for a manufacturable layout. In operations research the guillotine penalty is typically small: the asymptotic price of guillotinability for 2D bin packing is bounded between 4/3 and 1.692 (Khan and Pittu 2020) and the empirical loss is reported as marginal (Lodi et al. 1999). These are 2D bin-packing results and do not transfer as a guarantee to our 3D irregular-container yield; they are cited only as plausibility, and our measured 4-point gap (Section 4.6) is the direct evidence that a well-built guillotine packer sits just below the unconstrained optimum.

Staged guillotine packer. We pack each slab as a three-stage guillotine matching the wire saw: stage 1 horizontal bed cuts (layers), stage 2 vertical strips within each layer, stage 3 cross cuts within

each strip; each stage greedily takes the largest marketable size from a per-dimension catalog (the block stage trying smaller lengths against the irregular fit), over a best-of of the six orientations and the half-step phase offsets (Gilmore and Gomory 1965). Every cut is global per stage, so $\phi = 1$ by construction.

Saw cost (Jalalian I11). The cutting-surface area is the sum of the distinct planar cuts that free the blocks, each shared cut counted once,

$$A_{\text{cut}} = \tfrac{1}{2}\sum_{B}\sum_{\text{faces } f} a_f + \tfrac{1}{2}\sum_{B}\sum_{\text{waste faces } f} a_f, \qquad V_{\text{kerf}} = k\,A_{\text{cut}}, \tag{13}$$

and the Jalalian block-cutting-cost-per-block-value (BCSdbBV, our I11) is $\text{I11} = A_{\text{cut}}/\sum_B \text{value}(B)$ (Jalalian et al. 2023); the dominant wire-saw cost (consumable wire) scales with A_{cut} (Raza et al. 2024), so minimising I11 minimises cost. We report I11 with an equal class price, so it is the degenerate $A_{\text{cut}}/V_{\text{recovered}}$ form (cost per recovered cubic metre, units m^2/m^3); a value-weighted BCSdbBV needs per-class prices and is left to a market model. Three kerf scales coexist: the pure wire-bead kerf $\approx 10\,\text{mm}$ (Konstanty 2021), the packer's 30 mm (a conservative wire-plus-dressing value, the one used for the reported A_{cut} and the $V_{\text{kerf}} = k\,A_{\text{cut}} = 10.2\,\text{m}^3$; at the pure 10 mm bead it would be $\approx 3.4\,\text{m}^3$), and the $\approx 50\,\text{mm}$ total quarrying loss (kerf plus over-break and dressing) of BlockCutOpt (Elkarmoty et al. 2017). The invariant $V_{\text{kerf}} < V_{\text{slab}} - V_{\text{recovered}}$ (here 10.2 vs $47.5\,\text{m}^3$) is a sanity floor that the earlier double-counted accounting failed, not a tight bound.

BlockPacker (single-scale recovery solver). BlockPacker places a regular cutting grid of fixed-size candidate blocks in the tested rock volume and maximises the count N of *intact* blocks – those lying fully inside the volume and intersected by no fracture. The decision variables are the grid pose: the rotations ψ, θ, ϕ about Z, X, Y and the in-plane offsets d_x, d_y. The intact predicate is an AABB–BVH broad phase followed by a separating-axis test of each kerf-inflated candidate against the fracture triangles; the kerf-aware recovery is $R = N\,V_{\text{block}}/(V_{\text{tested}} - V_{\text{kerf}})$ (Elkarmoty et al. 2017, Eq. 7-1). The pose is searched by brute force over a discretised grid (e.g. $\psi \in [0, \pi]$ in 3° steps and $d_x, d_y \in [-1.5, 1.5]\,\text{m}$ in 0.5 m steps, $61 \times 7 \times 7 = 2989$ poses), returning the global optimum *over that grid* (the parallel search is bit-identical to the serial reference). BlockPacker is thus exact up to the pose-discretisation step (a finer grid can only weakly improve N), and it extends Elkarmoty's ψ-only BlockCutOpt (Elkarmoty et al. 2017) to full 3D rotation and the kerf-aware recovery. It is the $S = 1$ base of the cascade below.

Multi-scale recovery cascade. A block a fracture crosses is not discarded but logged and re-cut at a finer scale, recursively, around the fracture. With scales $s = 0 \ldots S{-}1$ (block, slab, tile) the value recovered from a region R is

$$W(R, s) = \sum_{b \in \text{kept}(R,s)} \text{RMV}_s(b) + \sum_{b \in \text{cracked}(R,s)} \begin{cases} W(\text{AABB}(b), s{+}1) & s{+}1 < S,\ \text{vol}(b) \geq V_{\min} \\ \text{residual} & \text{otherwise,} \end{cases} \tag{14}$$

with kept and cracked partitioning the winning grid by the fracture-intersection predicate. This couples the quarry-block and in-block scales (a standing open gap); it reduces to BlockPacker exactly at $S = 1$, generalises the conditional two-scale optimiser of Yarahmadi et al. (2018) and the usable-leftover threshold of cutting-stock (Cherri et al. 2009), and transfers the defect-aware cut-up logic of wood rough-milling. Oblique (dip-following) cuts, realizable by tilting the wire (Bennell et al. 2018), and the jointly-optimal cascade are named future work.

4 Results

4.1 Tolerance ladder on real data

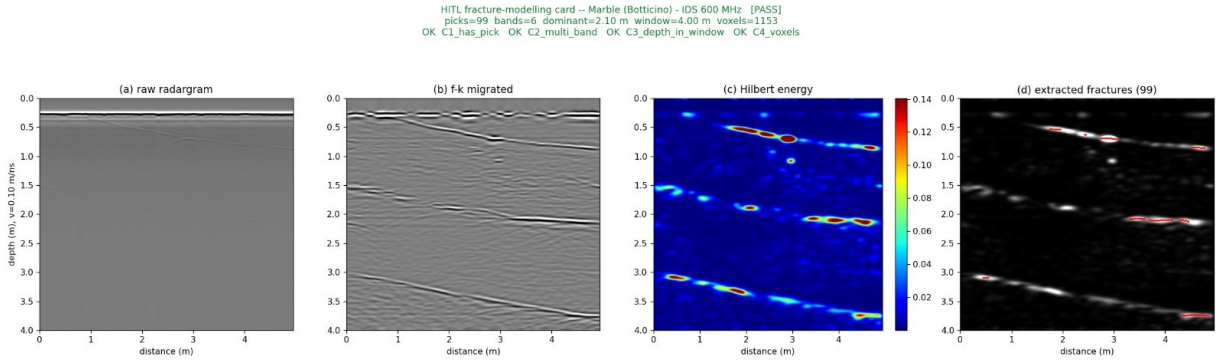


Figure 3: GPR fracture-extraction stage on the real Botticino marble line: (a) raw radargram, (b) f-k migrated, (c) Hilbert instantaneous energy, (d) 99 extracted fracture picks. These picks (the depth-converted reflectors) are the input to the tolerance ladder and the 3D surfaces below.

The extraction chain of Fig. 3 yields 99 depth-converted fracture picks; their per-location deviation from the true fracture is the tolerance ladder. On the Botticino grid3 stone (a compact Jurassic limestone marketed as marble; ε_r 9 – the Frahan marble preset, v 0.10 m/ns; the ± 1 used for $\sigma_v/v = \frac{1}{2}|\delta\varepsilon_r|/\varepsilon_r$ is a *local* relative-permittivity perturbation, not the full lithological spread, which across carbonate dimension stones runs ≈7.5–9.5 and brackets a measured marble value $\varepsilon_r \approx 7.7$ (Zanzi et al. 2023) at its lower end; 600 MHz; two-term σ_{recon}, time-zero off), the per-fracture total sigma and its components (m) are:

depth (m)	vel. $d\,\sigma_v/v$	$\lambda/4$ floor	σ_{recon}	$\sigma_{\text{interp+mesh}}$	σ_{total} (% depth)
0.72	0.040	0.042	0.058	0.030	0.065 (9.0%)
1.96	0.109	0.042	0.117	0.062	0.132 (6.7%)
3.37	0.187	0.042	0.192	0.201	0.278 (8.2%)

The $\lambda/4$ floor (0.042 m, at the marble survey's 600 MHz; the granite *detection* anchor of Section 3.3 uses 160 MHz and a correspondingly coarser $\lambda/4$) is the largest single term at 0.72 m (velocity 0.040 m); the velocity term dominates at 3.37 m; the interpolation term grows with depth (wider scan-line gaps). The %-of-depth ratios are 9.0/6.7/8.2 (the middle value is slightly lower), broadly in the range of Porsani's ±8.5–9.5% (but measured at 25 m, velocity-dominated (Porsani et al. 2006)) and Dorn's ≤4% (a maximum borehole-correlation error, not an RMS (Dorn et al. 2012)); the comparison is qualitative given the different depths and acquisition. Position confidence within ±2 cm is 16.8% (extraction) and 15.5% (3D surfaces); the detection-adjusted effective confidence is lower, ≈10% (=13.9% position × 74.4% detection completeness), which is the honest yield-safety number. Fig. 4 shows the sigma-coloured surfaces.

4.2 Detection calibration

The calibrated detection model recovers its anchor (a consistency check that the fit converged, not out-of-sample validation): for granite open sub-horizontal fractures at 9 m, $P_{\text{det}} = 73.8/77.7/78.6\%$

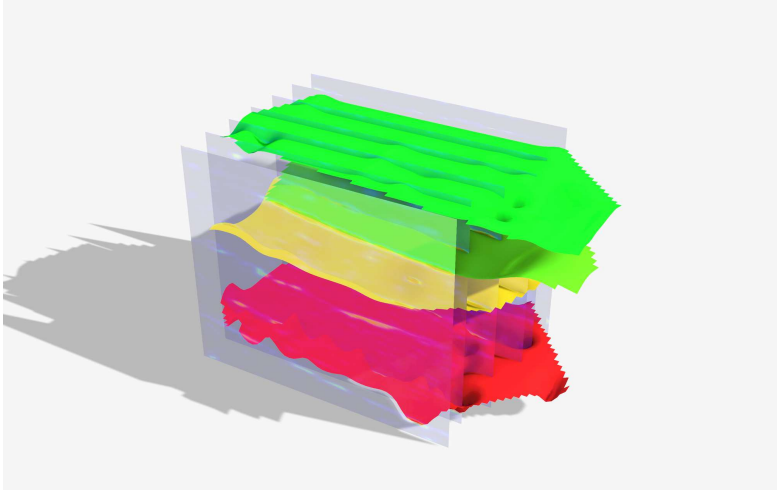


Figure 4: The four grid3 kriged fracture surfaces coloured by the total deviation σ – green (shallowest, within tolerance) grading to red (deepest, largest σ at 3.37 m) – over the translucent GPR radargram sections. Built and verified live in Rhino.

at 160/450/750 MHz (cf. Molron's 80% (Molron et al. 2020)); sealed 7.4%, sub-vertical 5.6%. Per-stone η_0 (Fig. 5) ranges from a measured 0.80 (granite) and an extrapolated 0.90 (low-loss limestone) down to 0.38 (tuff, a strong scatterer). The fracture-extraction chain itself was exercised on real GPR from more than one site and lithology: the three 600 MHz Botticino marble grids used throughout, and granite GPR from the Grimsel crystalline site (the source of the measured granite detection anchor). Detection and extraction therefore run on both marble and granite GPR; what remains single-site is the downstream dimension-stone block-yield economics (demonstrated on the one Botticino grid).

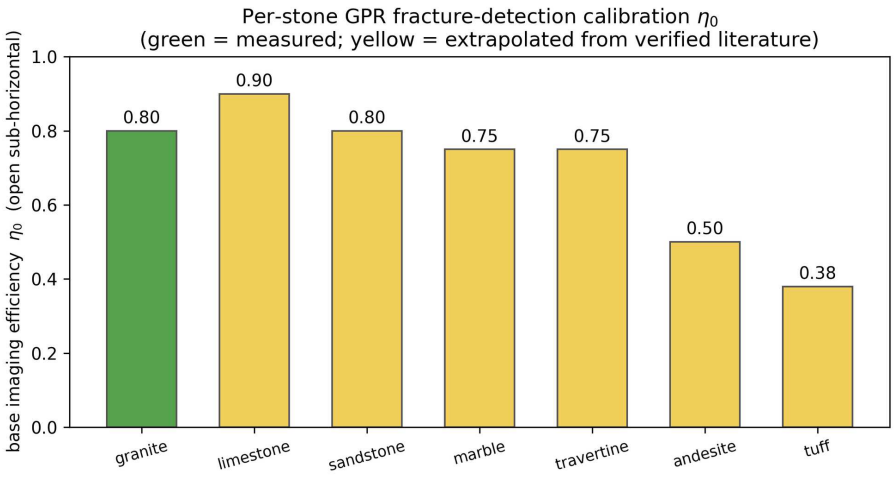


Figure 5: Per-stone GPR fracture-detection base efficiency η_0 (own work, from verified literature). Only granite is a measured rate; the rest are extrapolated and flagged.

4.3 Uncertainty-safe yield

Requiring every block to clear the fractures by their measured sigma drops the recoverable grid3 yield from 24% (geometric) to 6.5% (uncertainty-safe): the deep, high-sigma zones vanish. The yield-vs-clearance relation is monotone; for the combined packer (Section 4.4) the geometric 49.3% falls to 34.6%, 28.8%, 24.0% and 16.4% at a uniform clearance of 0.05, 0.10, 0.15 and 0.20 m (so the 6.5% is not an artefact of one clearance choice, and the evolved packer at a clearance $\approx$ mean σ recovers about the original baseline's *geometric* yield). This is the business case for the ladder: because the velocity term dominates σ_{recon} at depth (Section 4.1 table), a CMP/WARR velocity calibration that halves σ_v/v shrinks the deep σ and buys back safe yield (a value-of-information decision (Eidsvik et al. 2015)); a quantitative sounding-vs-yield VoI analysis is named future work.

4.4 Mesh-bench packing: a comparative study

We ran a head-to-head of seven packers on the same four watertight grid3 slabs (93.6 m^3 intact, block $0.9 \times 0.7 \times 0.4$ m, kerf 0.03 m), all with the same strict 8-corner irregular fit (Table 1, Fig. 6). Two are external: the Kim forest (`TreePackForest`, the Kim 2025 guillotine port) and a Mosch-style voxel greedy (a faithful re-implementation, since the 3D-BlockExpert software was unavailable). The naive fixed grid (24.0%) is a deliberately simple reference. The Kim forest scores only 19.7%: its guillotine slabs fragment against the wavy fracture boundary, so most placements straddle it and are rejected by the irregular fit, confirming Kim is a cuboid-into-cuboid method rather than a mesh-bench one. The Mosch-style greedy (52.2%) beat our global-grid combined packer (49.3%) by placing each block at its own lattice cell; per the standing directive we *adopted* that per-block placement, added our kerf-correct multi-size fill, and the resulting best-performing packer (Eq. 11, voxel-DLBF + multi) reaches **393 blocks / 53.3% geometric yield**, topping every method including the (kerf-free, hence optimistic) Mosch-style greedy, with no runtime errors. The packer runs are reproduced headless by a committed script (the recovered volumes and counts in `scripts/paper/results/marquee_packbench.csv`); a leave-one-slab-out hold-out confirms the 53.3% is not carried by one slab – the four fracture-bounded slabs individually yield 50.5/55.7/53.3/51.9% (spread 5.1 points). These are VOLUME yields; the 0.66× and 0.5× fill cuts trade at lower per-m^3 prices. Weighting the recovered blocks by the catalogue price (primary 2200, 0.66× 1650, 0.5× 1100 USD/m^3) gives a *value yield* of 48.8% (blended 2015 USD/m^3, 0.916 of the all-primary value) against the 53.3% volume yield – the fill cuts cost about 4.5 value points (the block-shape-factor / class weighting of Yarahmadi et al. (2018)). Under the uncertainty-safe clearance the packer yields 16–35% depending on the standoff (Section 4.3, monotone in Cl).

4.5 Wire-saw-manufacturable packing and the yield-manufacturability trade

The voxel-DLBF+multi packer wins on geometric yield (53.3%) but is *not* wire-saw cuttable: on the four grid3 slabs its layout is 0% guillotine-separable (Fig. 8), so much of that volume cannot be freed by full-span saw cuts. Re-posing the problem under the diamond-wire-saw constraint, the three-stage guillotine packer reaches **49.8% yield (**46.6 m^3 **of** 93.6 m^3**) at** $\phi = 100\%$ (every block separable), only **3.5 yield points** below the unconstrained optimum and at essentially the recovered volume of the global-grid combined packer while being fully cuttable (the combined-grid is only 63.6% separable, voxel-DLBF 0%). The two coincide on volume but the guillotine reaches it with far fewer, larger blocks (221 vs 363), which is the manufacturability feature, not a coincidence. Its sawing cost is $A_{\text{cut}} = 341\,\text{m}^2$ (equal-price I11 $= A_{\text{cut}}/V_{\text{recovered}} = 341/46.6 = 7.3\,\text{m}^2/\text{m}^3$); the

Table 1: Comparative packer study on the grid3 slabs (geometric yield; same irregular fit).

Packer	Yield	Blocks	kerf-ok	irreg. fit	rotation	multi-size	key limitation
Kim forest (TreePackForest)	19.7%	–	yes	post-filter	6	no	guillotine fragments on wavy boundary
Frahan grid (naive)	24.0%	89	yes	yes	no	no	one orientation, fixed lattice
Frahan best-of	35.3%	131	yes	yes	6	no	global grid (whole-row phase only)
Frahan voxel-DLBF	39.3%	146	yes	yes	6	no	single block size
Frahan combined-grid	49.3%	363	yes	yes	6	yes	global grid phase only
Mosch-style greedy	52.2%	–	*no*	yes	6	no	no kerf (optimistic); single size
Frahan voxel-DLBF+multi	**53.3%**	**393**	yes	yes	6	yes	voxel-pitch waste; single-site so far

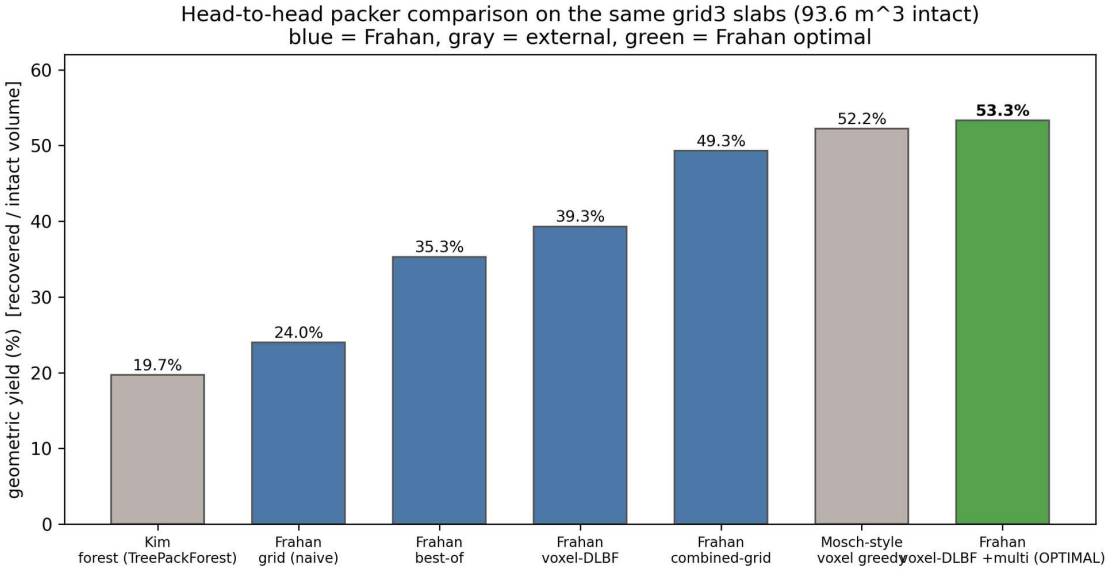
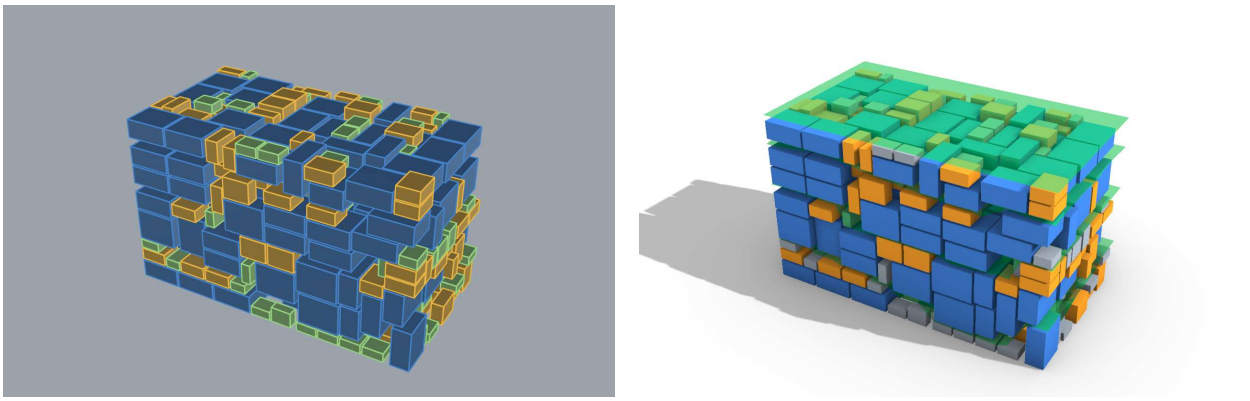


Figure 6: Head-to-head packer comparison on the same grid3 slabs. Frahan's best-of-tested voxel-DLBF+multi (green) leads on this grid; blue = other Frahan modes, gray = external packers. Own benchmark, all verified live.

kerf volume is $V_{\text{kerf}} = k\,A_{\text{cut}} = 0.03 \times 341 = 10.2\,\text{m}^3$ at the packer's 30 mm kerf (Section 3.5). Table 2 is the comparative manufacturability study; Figs. 9a and 9b show the staged packing and a section with the GPR fracture surfaces progressing across the block stack.

Multi-scale recovery cascade (closing the loop on cracked blocks). The cascade (Section 3.5) is implemented over BlockPacker, our block-recovery solver (which extends the BlockCutOpt formulation of Elkarmoty et al. (2017)), and verified by a headless test battery (six tests, all passing): it reduces bit-identically to BlockPacker at a single scale; a finer scale strictly recovers volume the coarse scale rejected (on a *synthetic* granite-style DFN benchmark of random fracture facets generated in the manner of standard stochastic DFN models (Berrone et al. 2016; Davy et al. 2013) and amenable to the same conditioning and uncertainty-quantification treatment (Dorn et al. 2013), not the real granite dataset, a slab scale added 97 pieces, lifting recovered



(a) the packing alone (benches hidden) **(b)** the same packing against the GPR fracture surfaces

Figure 7: The best-of-tested voxel-DLBF+multi packing, 393 blocks coloured by size: primary 0.9×0.7×0.4 m (blue) plus the 0.66× (orange) and 0.5× (green/grey) marketable fill cuts. (a) the packing recovers the wavy-boundary gaps; (b) the translucent green GPR fracture surfaces are what the packing must fight against, bounding each intact zone.

Table 2: Manufacturability comparative study on the grid3 slabs: geometric yield vs diamond-wire-saw (guillotine) separability ϕ and saw cost. ϕ is the greedy guillotine-separable fraction (1.0 = fully wire-saw cuttable). The staged guillotine matches the combined-grid yield while being the only high-yield packer that is fully cuttable.

Packer	Yield	Blocks	wire-saw sep. ϕ	saw cost
Frahan grid (naive, one size)	24.0%	89	100%	low yield, fully cuttable
Frahan combined-grid (multi-size)	49.3%	363	63.6%	partly cuttable
Frahan voxel-DLBF+multi (max yield)	53.3%	393	0%	greedy ϕ=0; not separable
Frahan staged guillotine	**49.8%**	**221**	**100%**	A_{cut} 341 m^2, I11 7.3

volume by ≈24% on that geometry); recovery is monotone in scale depth; the marketable-volume threshold gates recursion; the run is deterministic; and a fracture-free bench triggers no recursion. We report only the measured per-tier recovery; no universal recovery-uplift figure is claimed (it is geometry-specific). A reproducibility asymmetry is worth stating: the cascade ships committed C# with six registered headless tests, whereas the staged-guillotine numbers (ϕ, A_{cut}, 221 blocks) are a live-Rhino result under the visual-validation truth criterion, reproducible from the committed `guillotine_headtohead.cs` but not yet a headless regression; and ϕ is a greedy separability lower bound, so the 0%/63.6%/100% figures are conservative.

4.6 Cost-wise vs volume-wise block yield and the georeferencing prize

Re-posing the packer with a per-block net-value objective (sale price minus the 200 USD/m^2 diamond-saw cost) and a volume credit W (USD/m^3) that sweeps from pure profit ($W = 0$) through balanced ($W = 500$ USD/m^3) to pure throughput ($W \to \infty$) turns the same real Botticino grid into a cost-vs-volume frontier. The catalogue prices (1100–2200 USD/m^3 by block size) and the 200 USD/m^2 saw cost are illustrative engineering assumptions: quarry-block USD/m^3 prices are not separately published (only finished-product retail), and the cost *structure* follows the relative shares of Raza et al. (2024). The W-sweep, and a ±30% swing in the price band, leave the objective ranking

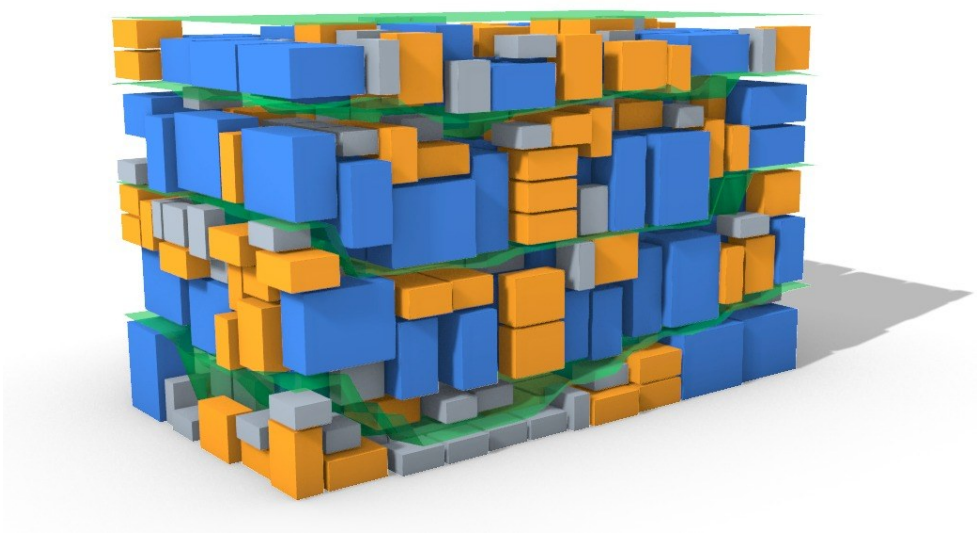
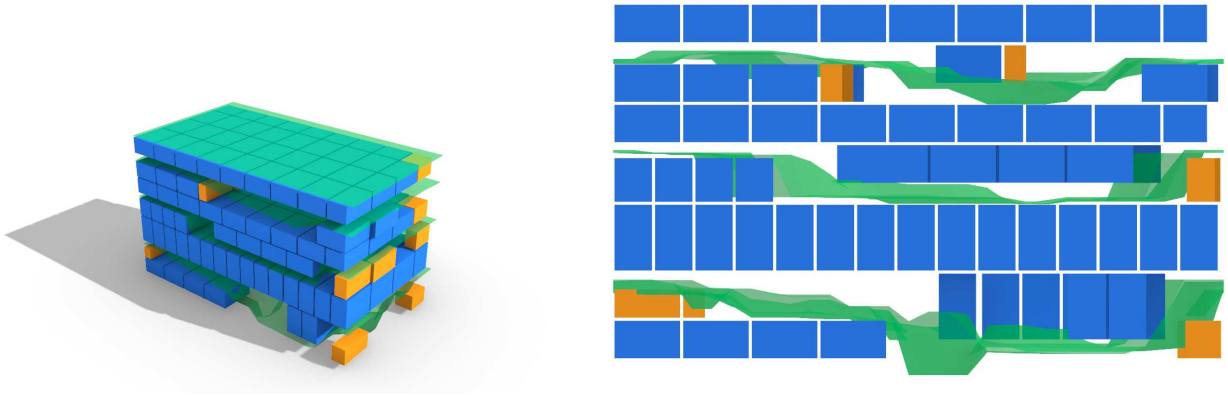


Figure 8: Max-yield voxel-DLBF+multi packing (53.3%): blocks coloured by size (blue primary, orange 0.66×, grey 0.5×) with the GPR fracture surfaces (green) weaving through. The blocks are staggered at independent lattice cells with no shared cut lines, so the layout is 0% wire-saw separable.



(a) staged three-stage guillotine packing **(b)** front elevation with the GPR fracture surfaces

Figure 9: The staged three-stage guillotine packing (49.8% yield, fully wire-saw cuttable). (a) blocks sit in clean horizontal beds with shared full-span cut lines (the diamond-wire bed-cut-then-vertical sequence), the GPR fracture surfaces (green) separating the intact slabs; (b) front elevation, the fracture surfaces progressing across the stack, every block separable by an edge-to-edge wire cut.

and the sign of the georeferencing prize unchanged, so the conclusion does not hinge on the exact price vector. The two LA01000x profiles give 280 fracture picks clustered into three dipping beds (0.72 m/6.1°, 2.10 m/0.9°, 3.70 m/6.1°; plane-fit RMS 5–15 mm), and the bed spacing (0.72, 1.38, 1.59, 0.30 m) caps the block height. Two guillotine modes are compared: *flat* (orthogonal, full-span axis-aligned cuts at the dip-safe envelope, fabricable on any gangsaw today) and *oblique* (full-span cuts tilted to follow each dipping bed, recovering the wedge the flat cut wastes). Table 3 and Figure 10 are the result.

Two findings. First, the cost-vs-volume trade is real and asymmetric: within the oblique mode, moving from max-cost to balanced buys 4.2 m^3 for only −\$1,478 (−5.1% net), whereas pushing to

Table 3: Cost/volume/balanced block yield on the real 3-bed Botticino grid (bench 49.9 m^3), for flat (today-fabricable) and oblique (bed-following) guillotine cutting. W is the volume credit.

Mode	Objective	Blocks	Volume (m^3)	Net (illustrative USD)
Flat (orthogonal)	Max cost	9	20.3	17,454
	Balanced	14	24.9	15,327
	Max volume	20	27.1	12,317
Oblique (bed-following)	Max cost	13	32.2	**28,741**
	Balanced	15	36.3	27,263
	Max volume	20	38.9	25,010

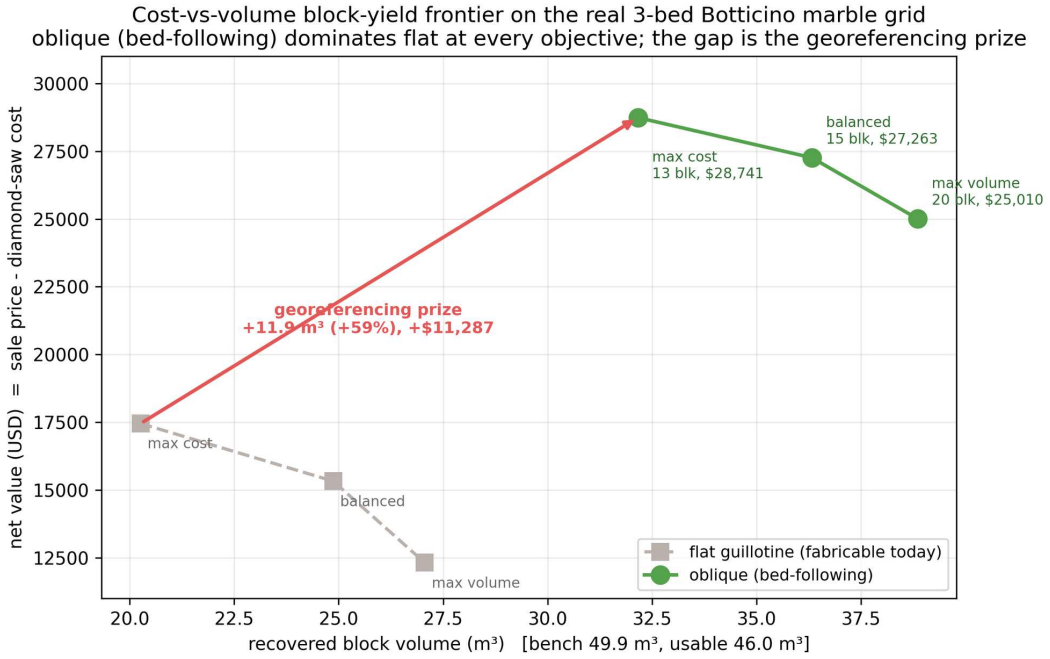
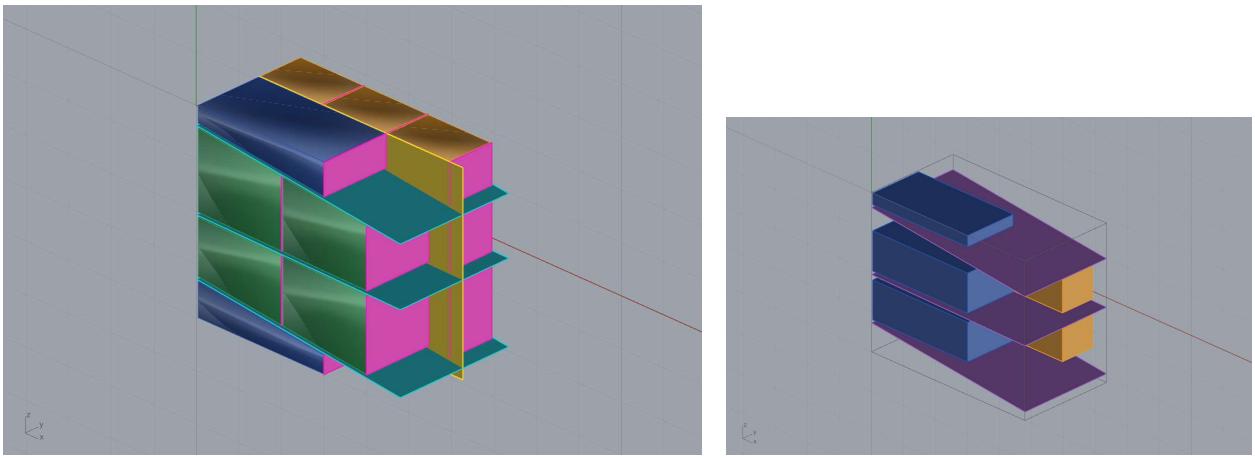


Figure 10: Cost-vs-volume yield frontier on the real 3-bed Botticino marble grid (same data as Table 3). The oblique (bed-following) frontier dominates the flat (today-fabricable) one at every objective; the red gap at the max-cost setting is the georeferencing prize (11.9 m^3, +59%, +$11,287). Generated at 300 DPI from `scripts/paper/results/cost_volume.csv` (results-run 2026-06-08).

max-volume swaps every premium 3.0×1.5 block for two 2.0×1.5 blocks, trading dollar density for 6.7 m^3 at −13% net. Balanced is the recommended operating point. Second, and larger, is the *georeferencing prize*: at the max-cost setting the oblique (bed-following) plan recovers 11.9 m^3 **(+59%) and** +$11,287 over flat cuts on this one bench, entirely because the beds dip $\approx 6°$; the prize scales with dip and collapses toward zero for flat-lying beds. This is the nominal figure; recomputed independently from the baked block geometry on the conservative basis (true tilted-hexahedron volume rather than the nominal footprint-times-spacing volume), the prize is still 9.4 m^3 (+47%) and +$8,910, so its direction and order of magnitude do not depend on the volume convention. Capturing that prize requires cutting along a sloped bed, which is only markable if the extracted block is scanned and georeferenced to the GPR-fitted bed planes so the tilted saw line can be transferred to the real stone. The prize quantifies the value of that scanning-and-marking last mile (Section 6, gap ix); flat guillotine is the fall-back that needs no such infrastructure. Localised

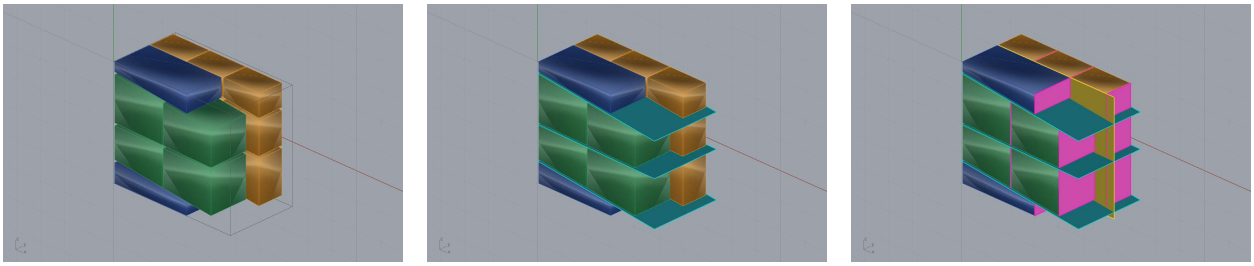
adjustment of each mark remains with the stonemason; the system supplies the georeferenced plan. Full method and per-bed statistics accompany the reproducibility deposit (Data availability).

The oblique guillotine cut plan for the balanced packing is 27 full-span passes: 3 oblique bed-parallel passes (tilted to the dipping beds, cutting along the natural parting) plus 1 vertical and 23 rip cuts that separate the blocks within each layer (Fig. 11a). Only the 3 tilted passes differ from the flat plan, and only they require the georeferenced marking; the flat fall-back (Fig. 11b) uses horizontal passes at the dip-safe envelope and abandons the wedges. The staged cutting sequence is shown in Fig. 12.



(a) oblique (bed-following) cut plan **(b)** flat (orthogonal) fall-back

Figure 11: Guillotine cut plans over the balanced packing on the real 3-bed Botticino grid. (a) the oblique plan: cyan = the 3 bed-parallel passes tilted to the dipping beds, yellow = the vertical beam-split, magenta = the 23 vertical rip cuts; the tilt of the cyan passes is the georeferenced-marking requirement. (b) the flat fall-back: axis-aligned blocks at the dip-safe envelope, the triangular wedges between the flat block tops and the dipping beds (violet) wasted. Flat is fabricable today but recovers 59% less than oblique (Table 3).



(a) placed bed-bounded blocks **(b)** the 3 oblique bed-parallel passes **(c)** all 27 full-span passes

Figure 12: The staged oblique guillotine cutting sequence for the balanced packing: (a) the placed bed-bounded blocks; (b) the 3 oblique bed-parallel passes tilted to the dipping beds (the only passes needing georeferenced marking); (c) all 27 full-span passes once the single vertical split and the 23 vertical rip cuts are added. Each stage is one set of full-span wire-saw passes.

4.7 Sensitivity analyses

Three reviewer-requested sensitivities, all analytic on the grid3 marble components (Fig. 13; time-zero is held off throughout, as justified in Section 3.1 – it is sub-dominant beyond 0.5 m and would

only widen the shallowest bin): *(1) Detection* η_0 (the largest extrapolation for non-granite stone): the effective confidence is near-linear in η_0, spanning 8.3% ($\eta_0 = 0.60$) to 12.4% (0.90) about the 10.3% at the marble extrapolate 0.75, so the headline ≈10% is bounded within ±2 points across the plausible range. *(2) Quadrature independence:* relaxing the zero-correlation assumption to $\rho = 0.1/0.3/0.5$ between σ_{recon} and σ_{interp} grows σ_{total} by up to 22% and lowers the mean within-2 cm confidence only from 13.9% to 11.7% – robust to the assumption. *(3) Cutting tolerance* T*:* the within-T confidence rises 7.0 / 13.9 / 33.1 / 56.8% at $T = 1/2/5/10$ cm; the 2 cm headline is a conservative ornamental-stone cutting tolerance, and the methodology (clearance = σ) is the contribution, not the single T. *(4) Detection area* A*:* the detection probability is governed by $p_A(A) = A/(A + A_{\text{min}})$ with $A_{\text{min}} = (\lambda/4)(d/2) = 0.84\,\text{m}^2$ at the 9 m granite anchor (160 MHz). The measured anchor corresponds to a large continuous fracture ($A \approx 10\,\text{m}^2$, $p_A = 0.92$); reducing the fracture area to $A = 1/2/5\,\text{m}^2$ drops P_{det} to ≈43/56/68% (from 74% at $\eta_0 = 0.80$), a ≈30-point swing, so area is the dominant single lever on detection. Picks with $A < A_{\text{min}}$ are flagged sub-resolution and excluded; the per-pick A is estimated from the along-line continuity of the energy maximum. *Kriging diagnostics:* the ordinary kriging uses an exponential variogram fitted to the bed-pick cloud (practical range ≈ 3.75 m, the empirical correlation length; the nugget is folded into the LOOCV-calibrated floor below). A leave-one-out cross-validation of the σ_{interp} over the 180 bed-assigned picks across the four grid3 surfaces gives an LOOCV held-out RMSE of 0.026 m (mean |error| 0.022 m), which exceeds the near-pick posterior std (0.011 m) as expected (kriging variance is zero at data and grows in the gaps). Since the raw posterior std therefore under-estimates the between-line error, we adopt the LOOCV RMSE as the calibrated σ_{interp} floor (a ≈2.4× inflation over the raw std). In practice this floor is *non-binding* at the three representative depths: the grid3 $\sigma_{\text{interp+mesh}}$ values used above (0.030/0.062/0.201 m) already exceed the 0.026 m floor, so it changes none of the headline σ_{total} or confidence numbers and binds only very near picks (where the raw posterior std would otherwise dip toward zero). The 0.011 m near-pick under-estimate therefore never contaminated the reported three-depth summary; the fitted variogram practical range is ≈ 3.75 m. *The 8–12%* η_0 *band itself rests on the extrapolated marble* η_0*; quarry-face validation, not the sweep, is what would confirm it.*

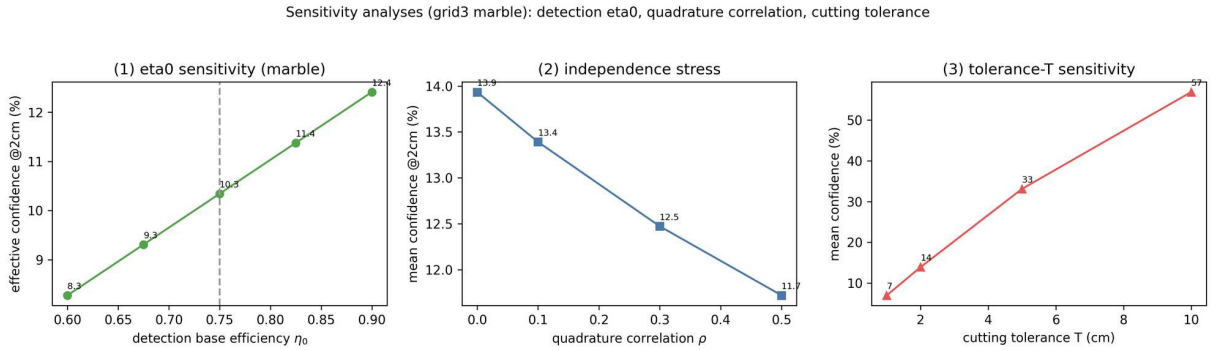


Figure 13: Sensitivity analyses (own work, grid3 marble): (1) effective confidence vs detection efficiency η_0; (2) mean confidence vs the quadrature correlation ρ (independence stress); (3) confidence vs cutting tolerance T. The headline numbers are bounded and robust.

5 Research gaps

Closed by this work. (i) No prior block-cutting optimiser scaled the cut clearance to the *measured* per-fracture uncertainty and computed the yield cost; we do ($24\% \to 6.5\%$). (ii) The detection rung that the position ladder omitted is now a calibrated model with a detection-adjusted confidence. (iii) The mesh-bench packer is evolved past a fixed grid to a best-of-tested voxel-DLBF multi-size algorithm ($24\% \to 53.3\%$) that, in a head-to-head on the same slabs, beats the external Kim forest (19.7%) and a Mosch-style greedy (52.2%) – closing the "compare to state-of-the-art packers" gap with a comparative metrics table (Table 1). (iv) The manufacturability gap is closed for a real diamond-wire-saw quarry (no robot): a three-stage guillotine packer reaches 49.8% yield at 100% wire-saw separability (3.5 points below the uncuttable max-yield optimum), with the Jalalian I11 saw-cost metric and a tested multi-scale crack-aware recovery cascade that re-cuts cracked blocks at finer scales (Table 2, Section 4.6). (v) The block-cutting objective is extended from pure geometric yield to a cost-vs-volume net-value frontier (price minus saw cost, with a swept volume credit) on the same real Botticino grid, and the flat-vs-oblique (dip-following) guillotine yield gap is quantified for the first time at +59%/+$11,287 per bench (Table 3, Section 4.7).

Limitations (prominent). Every quantitative result in this paper is a benchmark on a *single* Botticino-marble quarry grid with one dimension-block size; generalisation to other quarries, lithologies (granite, limestone) and geometries requires further data, and the per-stone η_0 for the six non-granite stones is the largest single uncertainty (extrapolated, not measured). The benchmarks demonstrate the *method*, not a production-validated yield curve.

The real GPR data are Botticino marble, used because it is the only openly licensed quarry-scale GPR grid with mapped fractures we could obtain; the motivating application is South Indian (Tamil Nadu) granite and charnockite dimension stone (Suresh 2020; Anbazhagan 2011), for which open GPR is absent. This is not a bait-and-switch: the method is lithology-agnostic and parametric in radar-wave velocity and relative permittivity (marble $\varepsilon_r \approx 9$, $v \approx 0.10\,\mathrm{m/ns}$; granite $\varepsilon_r \approx 5$–6, $v \approx 0.12\,\mathrm{m/ns}$), and the detection model is in fact calibrated on *measured* granite rather than marble. The granite and Tamil-Nadu cases are synthetic discrete-fracture networks built from field joint orientations (Anbazhagan 2011) with spacing in the heavily-fractured sub-metre tail of the granite range of Sousa (2007) (0.98 m to 2.8 m), consistent with the joint- and crack-controlled yield of Krishnagiri dimension-stone granite (Suresh 2020); they are illustrative, not validated field results. Transfer to real granite quarries requires a GPR field campaign, the priority next step.

Open (named for future work). (i) Only granite has a *measured* GPR detection rate; the other six lithologies' η_0 are extrapolated and need quarry-face / borehole ground truth. (ii) The fabrication-economics evidence cluster is the thinnest (5 of 38) and warrants an exhaustive registered review. (iii) Detection uncertainty (missed clay / sub-vertical fractures) is modelled probabilistically but not field-validated outside granite. (iv) A quantitative value-of-information analysis (one CMP sounding vs recovered yield) is set up but not run. (v) The statistical sensitivities flagged in review are now ADDRESSED (Section 4.5): kriging LOOCV (RMSE 0.026 m, range $\approx$ 3.75 m, adopted as the inflated σ_{interp} floor since the raw posterior std under-estimates the between-line error); the η_0 $\pm$20% sweep (effective confidence 8–12%); the quadrature-correlation stress (ρ 0–0.5: mean confidence 13.9%$\to$11.7%); and the tolerance-T curve. What remains is a hold-out / multi-quarry packer validation to remove best-of-N search bias. (vi) *Cutting-plane cost is now modelled* (Sections 3.5, 4.6): guillotine-cuttability, the diamond-wire-saw saw-path / kerf cost, and the Jalalian I11 objective, with a tested crack-aware recovery cascade. The *value* of oblique (dip-following) cuts is now quantified (Section 4.7: +59%/+$11,287 per bench over flat cuts on the real dipping-bed grid); what remains is *executing* them physically – realizable by tilting the wire but needing

a wire-sag flatness model and a kerf-aware tilted-plane derivation (Bennell et al. 2018), plus the block-scanning and georeferenced-marking last mile (gap ix) to transfer the sloped saw line to the real stone – and the jointly-optimal (rather than greedy-per-scale) cascade via staged column generation. (vii) *LiDAR mesh-bench integration.* The fracture-bounded container geometry is currently a constructed slab set; coupling the GPR fracture surfaces with a LiDAR / photogrammetry scan of the bench (the scan-ingest → mesh-bench chain) would supply the true bounding volume to pack into directly from reality, to be explored. (viii) *Georeferencing and mining-plan support.* Tying the georeferenced (UTM) fracture model and yield into a quarry reserve / mining-plan model would turn the per-bench yield into a site-scale extraction-planning tool; folding equipment and logistics efficiency into the block sizing, as done for open-pit mining (Guo et al. 2024), is future work. (ix) *Customisation to small-scale, stone-mason quarrying (the principal application).* Our main target is the small fractured quarry worked largely by hand rather than by a multi-axis robot, where the cut plan still follows a mason's thumb-rules: a dimension block must exceed 1 m on every axis and be at least 3 m in length (about a $3\,\mathrm{m}^3$ minimum, e.g. $3 \times 1 \times 1$ up to $3 \times 3 \times 1$ m and variable), a slab is at least 3 m long by 1 m wide at 1, 1.5 or 2 inch thickness, and monuments take varied bespoke sizes, so the packer and cascade remain to be constrained to these catalogues and to the manual extraction sequence (strip the overburden, clean the outcrop, expose two faces of the rock, mark the bore-hole and cut lines, extract a large primary block, then gangsaw it into the required sizes), repeated bench by bench. A natural and high-value extension is to drive those marking decisions directly from the GPR and laser / photogrammetry model so that crack-prone zones are deliberately avoided for yield, economics and planning, and to deepen the survey to the second working layer (on the order of 20 m). Realising this in the field also means co-designing the exact scan positions with the on-site ground team and stone masons, so the survey data lands precisely where the extraction plan needs it. (x) *Aesthetic figure, grain and pattern are not modelled.* The yield here is purely geometric: it maximises intact-block recovery against the fracture network and ignores the visual quality that sets dimension-stone value. Commercial granite is graded by colour uniformity, grain direction, vein and wavy patterns, and characteristic figure – the speckled "galaxy" stones (Black, Golden and Mercury Galaxy), the flower-and-dot textures, and banded flow patterns – and a quarry mason orients, selects and trims each block at the face to capture or align these features, discarding blocks with patches, moles or cross-cutting intrusions even when they are mechanically sound (Suresh 2020). Mapping grain and figure (from photogrammetry or hyperspectral imaging) and matching them to the cut plan, so the layout honours both fracture avoidance and aesthetic value, is left as future work; this is the natural complement to the geometric yield reported here.

Resource efficiency. Raising recovery from the 24% fixed-grid baseline to the 53.3% mesh-bench (or the wire-saw-manufacturable 49.8%) roughly doubles the marketable stone won per cubic metre of extracted rock, so the waste rock and overburden handled per saleable block falls by about half; the $11.9\,\mathrm{m}^3$ georeferencing prize per bench is stone that would otherwise become spoil. For an energy- and transport-intensive industry this is a direct cut in the embodied impact per finished cubic metre, and the uncertainty-safe clearance protects that gain from re-cuts on blocks that later prove fractured – a quantified contribution to lower-waste, responsible dimension-stone sourcing.

6 Conclusion

The pipeline turns a GPR survey into an honest, uncertainty-aware dimension-block yield: it states not only where the fractures are but how certain that is and whether any were missed, and it packs the safe intact volume efficiently. Every stage is managed and was verified live in Rhino.

The headline numbers (position sigma 7–9% of depth; detection ≈80% reproduced (calibrated); uncertainty-safe yield 24% → 6.5%; packer 24% → 53.3% best-of-tested) are real benchmarks on a real marble grid, and the literature base and its gaps are reported transparently.

Author contributions

Libish Murugean: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Visualization, Writing – original draft, Writing – review & editing.

Competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

The Botticino marble GPR data are openly available under CC BY 4.0 (Bonduà et al. 2024, *Data* 9(3):42, `10.3390/data9030042`; Mendeley `10.17632/w26n6nftxs.3`). The ETH1100 drystone meshes (Zenodo `10.5281/zenodo.10038881`, CC BY 4.0) and the Loviisa fracture-trace map (Chudasama 2022, Zenodo) are open. The results-run code, the per-stone results CSV, and the figure scripts are archived at Zenodo (`10.5281/zenodo.20608279`) as a reviewer-testable reproducibility package: the BlockPacker batch runner, the `marquee_packbench.csx` packer-comparison script, the input `blocks_grid3.3dm`, the results CSVs, and the figure scripts, with documentation mapping each packer-yield and cost/volume number and figure to the command that produces it; the live-Rhino guillotine ϕ, A_{cut} and block count are flagged separately (Section 4.6) and are not yet headless-reproducible.

All underpinning measurements are openly published, third-party datasets acquired at operational quarry and field sites and accessed independently from public repositories; no proprietary, restricted, or institutionally collected data were used. The uncertainty-aware pipeline is built entirely from open-source software—the RGPR package (R) for reading the raw GPR records, the *geogram* and *CGAL* computational-geometry libraries for meshing and Boolean operations, the Frahan StonePack plugin for packing and block-cutting (running in Rhinoceros/Grasshopper, Robert McNeel & Associates 2023), and Python (NumPy, SciPy, Matplotlib) for the statistics and figures—and runs on commodity hardware, requiring no commercial or licensed engineering, GIS, GPR-processing, or CAD/CAM software.

The Frahan StonePack Rhino/Grasshopper plugin that implements these GPR-to-block-yield workflows will be released as open source upon publication, so the full pipeline is reproducible by practitioners.

Use of generative AI and AI-assisted technologies

During the preparation of this work, the author used a generative AI assistant, Claude by Anthropic, as a supporting tool for software development.

All datasets used in this study are real, openly published measurements. The governing equations were developed by the author, and the analyses, numerical results, and conclusions were independently reviewed and verified by the author.

The author carefully reviewed all AI-assisted outputs and takes full responsibility for the content of the publication. No AI system is listed as an author.

Acknowledgements

The author thanks the providers of the open datasets used in this work: the Botticino marble GPR survey (Bonduà et al., *Data* 2024, CC BY 4.0), the ETH Zurich dry-stone masonry mesh dataset, and the Loviisa granite fracture-trace map (Chudasama et al.).

References

Anbazhagan, Guru, Biswal (2011) Remote sensing in delineating deep fractured aquifer zones. In: Geoinformatics in applied geomorphology. CRC Press, Boca Raton, ch 12. ISBN 9781439830598

Annan AP (2009) Electromagnetic principles of ground penetrating radar. In: Jol HM (ed) Ground penetrating radar: theory and applications. Elsevier, Amsterdam, pp 3–40

Bennell JA, Cabo M, Martínez-Sykora A (2018) A beam search approach to solve the convex irregular bin packing problem with guillotine cuts. European Journal of Operational Research 270:89–102. https://doi.org/10.1016/j.ejor.2018.03.029

Berrone S, Pieraccini S, Scialò S (2016) Towards effective flow simulations in realistic discrete fracture networks. Journal of Computational Physics 310:181–201. https://doi.org/10.1016/j.jcp.2016.01.009

Bonduà S, Monteiro Klen A, Pilone M, Asimopolos L, Asimopolos NS (2024) A set of ground penetrating radar measures from quarries. Data 9(3):42. https://doi.org/10.3390/data9030042

The CGAL Project (2023) CGAL user and reference manual. CGAL Editorial Board. https://www.cgal.org

Cherri AC, Arenales MN, Yanasse HH (2009) The one-dimensional cutting stock problem with usable leftover – A heuristic approach. European Journal of Operational Research 196:897–908. https://doi.org/10.1016/j.ejor.2008.04.039

Chudasama B (2022) Loviisa rapakivi-granite fracture and lineament dataset, southern Finland. Zenodo (open dataset)

Cressie NAC (1993) Statistics for Spatial Data. Wiley. https://doi.org/10.1002/9781119115151

Davy P, Le Goc R, Darcel C (2013) A model of fracture nucleation, growth and arrest, and consequences for fracture density and scaling. JGR Solid Earth 118:1393–1407. https://doi.org/10.1002/jgrb.50120

Dorn C, Linde N, Doetsch J, Le Borgne T, Bour O (2012) Fracture imaging within a granitic rock aquifer using multiple-offset single-hole and cross-hole GPR reflection data. Journal of Applied Geophysics 78:123–132. https://doi.org/10.1016/j.jappgeo.2011.01.010

Dorn C, Linde N, Borgne TL, Bour O, de Dreuzy JR (2013) Conditioning of stochastic 3-D fracture networks to hydrological and geophysical data. Advances in Water Resources 62:79–89. https://doi.org/10.1016/j.advwatres.2013.10.005

Eidsvik J, Mukerji T, Bhattacharjya D (2015) Value of Information in the Earth Sciences. Cambridge University Press. https://doi.org/10.1017/CBO9781139628785

Elkarmoty M, Colla C, Gabrielli E, Kasmaeeyazdi S, Tinti F, Bonduà S, Bruno R (2017) Mapping and modelling fractures using ground penetrating radar for ornamental stone assessment and recovery optimization: two case studies. Min-Geol-Pet Eng Bull 32:63–76. https://doi.org/10.17794/rgn.2017.4.7

Gilmore PC, Gomory RE (1965) Multistage Cutting Stock Problems of Two and More Dimensions. Operations Research 13:94–120. https://doi.org/10.1287/opre.13.1.94

Grasmueck M, Weger R, Horstmeyer H (2005) Full-resolution 3D GPR imaging. Geophysics 70:K12–K19. https://doi.org/10.1190/1.185

Guo W, Liu G, Li J, Chai S, Guo S (2024) Research on the method of determining the block size for an open-pit mine integrating mining parameters and shovel-truck's operation efficiency. Sci Rep 14. https://doi.org/10.1038/s41598-024-52815-9

Huber E, Hans G (2018) RGPR—an open-source package to process and visualize GPR data. In: 2018 17th International Conference on Ground Penetrating Radar (GPR), IEEE, pp 1–4. https://doi.org/10.1109/ICGPR.2018.8441658

Jalalian MH, Bagherpour R, Khoshouei M (2023) Environmentally sustainable mining in quarries to reduce waste production and loss of resources using the developed optimization algorithm. Sci Rep 13. https://doi.org/10.1038/s41598-023-49633-w

JCGM (2008) JCGM 100:2008 evaluation of measurement data—guide to the expression of uncertainty in measurement (GUM). Joint Committee for Guides in Metrology, BIPM

Kapageridis I (2018) Resource and Reserve estimation for a marble quarry using quality indicators. J. South. Afr. Inst. Min. Metall. 118:39–45. https://doi.org/10.17159/2411-9717/2018/v118n1a5

Kazhdan M, Hoppe H (2013) Screened Poisson surface reconstruction. ACM Trans Graph 32(3):29:1–29:13. https://doi.org/10.1145/248

Khan A, Pittu E (2020) On guillotine separability of squares and rectangles. In: Approximation, randomization, and combinatorial optimization (APPROX/RANDOM 2020), LIPIcs vol 176. Schloss Dagstuhl, pp 47:1–47:22. https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2020.47

Kim T (2025) Packing and Cutting Stone Blocks Based on the Nonlinear Programming of Tree Cases. Computation 13:211. https://doi.org/10.3390/computation13090211

Konstanty JS (2021) The mechanics of sawing granite with diamond wire. Int J Adv Manuf Technol 116:2591–2597. https://doi.org/10.1007/s00170-021-07577-3

Lei Q, Latham JP, Tsang CF (2017) The use of discrete fracture networks for modelling coupled geomechanical and hydrological behaviour of fractured rocks. Computers and Geotechnics 85:151–176. https://doi.org/10.1016/j.compgeo.2016.12.024

Lévy B (2024) Geogram: a programming library of geometric algorithms. https://github.com/BrunoLevy/geogram

Lodi A, Martello S, Vigo D (1999) Heuristic and Metaheuristic Approaches for a Class of Two-Dimensional Bin Packing Problems. INFORMS Journal on Computing 11:345–357. https://doi.org/10.1287/ijoc.11.4.345

Molron J, Linde N, Baron L, Selroos JO, Darcel C, Davy P (2020) Which fractures are imaged with Ground Penetrating Radar? Results from an experiment in the Äspö Hardrock Laboratory, Sweden. Engineering Geology 273:105674. https://doi.org/10.1016/j.enggeo.2020.105674

Mosch S, Nikolayew D, Ewiak O, Siegesmund S (2010) Optimized extraction of dimension stone blocks. Environ Earth Sci 63:1911–1924. https://doi.org/10.1007/s12665-010-0825-7

Neal A (2004) Ground-penetrating radar and its use in sedimentology: principles, problems and progress. Earth-Science Reviews 66:261–330. https://doi.org/10.1016/j.earscirev.2004.01.004

Palmstrom A (2005) Measurements of and correlations between block size and rock quality designation (RQD). Tunnelling and Underground Space Technology 20:362–377. https://doi.org/10.1016/j.tust.2005.01.005

Porsani JL, Sauck WA, Júnior AOS (2006) GPR for mapping fractures and as a guide for the extraction of ornamental granite from a quarry: A case study from southern Brazil. Journal of Applied Geophysics 58:177–187. https://doi.org/10.1016/j.jappgeo.2005.05.010

Rasmussen CE, Williams CKI (2005) Gaussian Processes for Machine Learning. The MIT Press. https://doi.org/10.7551/mitpress/3206

Raza MA, Raza S, Khan MU, Emad MZ, Jalil K, Saki SA (2024) Cost modelling for dimension stone quarry operations. J. S. Afr. Inst. Min. Metall. 123:521–525. https://doi.org/10.17159/2411-9717/1578/2023

Robert McNeel & Associates (2023) Rhinoceros 3D, version 8 [computer software]. Seattle, WA. https://www.rhino3d.com

Sousa LMO (2007) Granite fracture index to check suitability of granite outcrops for quarrying. Engineering Geology 92:146–159. https://doi.org/10.1016/j.enggeo.2007.04.001

Suresh, Uma Maheswaran, Tamilarasan, Ranjith Kumar, Anbazhagan (2020) Quality assessment and grading of dimension stone in Krishnagiri District, Tamil Nadu, India. J Sci Tech 5(2):76

Ülker E, Turanboy A (2009) Maximum volume cuboids for arbitrarily shaped in-situ rock blocks as determined by discontinuity analysis—A genetic algorithm approach. Computers & Geosciences 35:1470–1480. https://doi.org/10.1016/j.cageo.2008.08

Xie F, Lai WWL, Dérobert X (2021) GPR-based depth measurement of buried objects based on constrained least-square (CLS) fitting method of reflections. Measurement 168:108330. https://doi.org/10.1016/j.measurement.2020.108330

Yarahmadi R, Bagherpour R, Taherian SG, Sousa LMO (2018) Discontinuity modelling and rock block geometry identification to optimize production in dimension stone quarries. Engineering Geology 232:22–33. https://doi.org/10.1016/j.enggeo.2017.

Yavuz AB, Turk N, Koca MY (2005) Geological parameters affecting the marble production in the quarries along the southern flank of the Menderes Massif, in SW Turkey. Engineering Geology 80:214–241. https://doi.org/10.1016/j.enggeo.2005.05.003

Zanzi L, Izadi-Yazdanabadi M, Karimi-Nasab S, Arosio D, Hojat A (2023) Time-lapse GPR measurements to monitor resin injection. Sensors 23(20):8490. https://doi.org/10.3390/s23208490

Appendix A: Nomenclature and abbreviations

Abbreviations

AABB	axis-aligned bounding box
AGC	automatic gain control
BVH	bounding-volume hierarchy
BlockCutOpt	block-cutting optimisation of Elkarmoty et al. 2017/2020 (the cited prior method)
BlockPacker	our block-recovery solver, which extends BlockCutOpt
CC-BY	Creative Commons Attribution 4.0 (the open licence of the GPR data)
CGAL	Computational Geometry Algorithms Library
CRA	coupled rigid-block analysis
DFN	discrete fracture network
DLBF	deepest/bottom-left-fill (packing heuristic; voxel-DLBF variant)
DPI	dots per inch
f–k	frequency–wavenumber domain (Stolt migration)
GPR	ground-penetrating radar
NFP-BLF	no-fit-polygon bottom-left-fill (2D nesting)
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
RBE	rigid-block equilibrium
RMS	root-mean-square (error)
ROSES	RepOrting standards for Systematic Evidence Syntheses
SAT	separating-axis theorem
SLM	systematic literature map (algorithm + code tier)
TLS	terrestrial laser scanning
USGS	United States Geological Survey

Symbols

ε_r	relative permittivity of the rock mass
v	GPR wave velocity (m/ns)
$\lambda,\ \lambda/4$	radar wavelength; quarter-wavelength vertical resolution limit
$\psi,\ \theta,\ \phi$	block-grid rotation about the Z, X, Y axes
$d_x,\ d_y$	block-grid translation offsets (m)
kerf	material width lost to the saw cut (m)
$\sigma_{\text{recon}},\ \sigma_{\text{interp}},\ \sigma_{\text{total}}$	reconstruction, interpolation, and total position uncertainty (m)
η_0	base GPR fracture-imaging efficiency (open sub-horizontal)
ρ	quadrature correlation in the uncertainty sum
T	cutting tolerance (m)
W	volume credit in the cost objective (USD/m^3)
N_{ni}	number of non-intersected (intact) blocks
R	recovery (%) $= N_{\text{ni}}\, V_{\text{block}}/(V_{\text{tested}} - V_{\text{kerf}})$
φ	wire-saw (guillotine) separability fraction

Appendix B: BlockPacker recovery results-run

BlockPacker, our block-recovery solver (which extends the BlockCutOpt formulation (Elkarmoty et al. 2017)), was run headless over a batch of fracture benches by a Rhino-free console (it consumes the solver read-only; inputs are PLY fracture meshes and synthetic discrete-fracture networks). The pose search used the component defaults: $\psi \in [0, 180]^\circ$ in 3° steps and $d_x, d_y \in [-1.5, 1.5]$ m in 0.5 m steps (a $61 \times 7 \times 7 = 2989$ pose grid), kerf 0.05 m. Every run is deterministic and completes in well under a second.

The validation anchor is the real Loviisa rapakivi-granite fracture-trace map (Chudasama 2022), clipped to a $12.5 \times 13.2 \times 3.0$ m window and extruded to a 704-triangle fracture mesh. At the A4 block size ($0.7 \times 0.6 \times 0.4$ m) the run recovers $N = 158$ intact blocks (5.43% kerf-aware recovery). A prior validation log recorded $N = 163$ for the same configuration; the $\sim 3\%$ difference is consistent with single-precision round-trip of the re-exported fracture mesh, since the pose grid is bit-identical (2989 evaluations). Recovery falls monotonically as the block grows (Figure 14), as expected on a fixed fracture spacing.

Across benches (Figure 15), the three synthetic DFN seeds bracket the real bench, and a literature-grounded closely jointed Tamil-Nadu granite bench yields zero intact $2.0 \times 1.0 \times 1.0$ m blocks, the expected floor for sub-metre jointing. That bench uses field joint orientations from Anbazhagan et al. (Anbazhagan et al. 2011) (Hosur–Rayakottai charnockite/granite) with sub-metre spacing in the heavily-fractured tail of the Sousa granite range (0.98–2.8 m), consistent with the joint- and crack-controlled yield reported for Krishnagiri dimension-stone granite (Suresh et al. 2020). All nine runs completed with no solver failures.

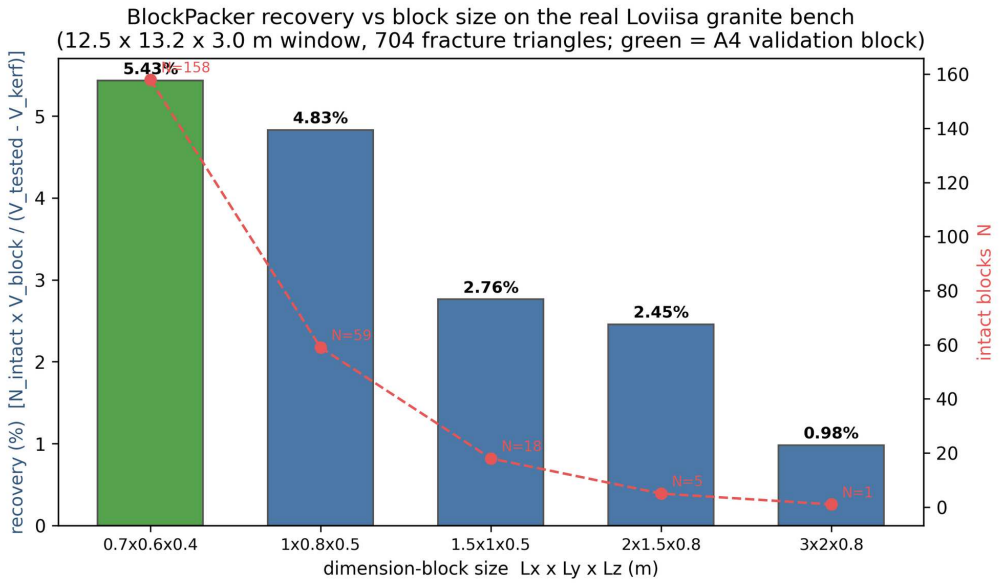


Figure 14: BlockPacker recovery versus dimension-block size on the real Loviisa granite bench (the catalogue sweep; green = the A4 validation block). Bars: kerf-aware recovery. Dashed line: intact-block count N.

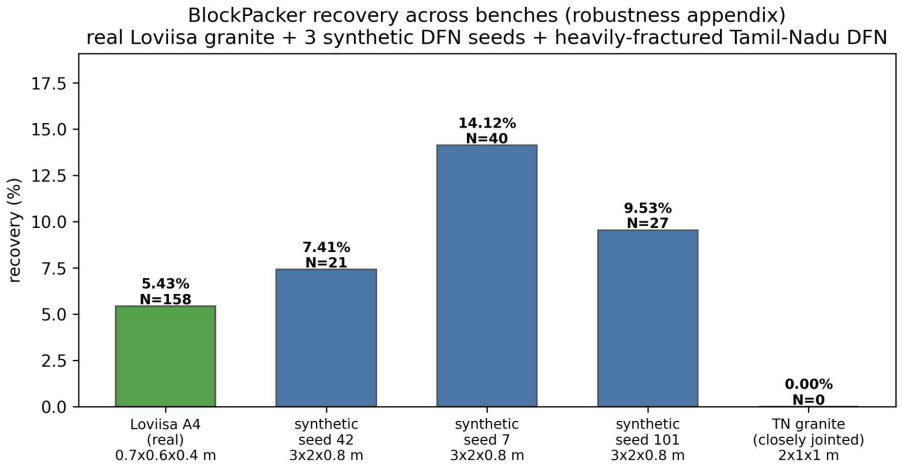


Figure 15: BlockPacker recovery across benches (robustness appendix): the real Loviisa A4 anchor, three synthetic DFN seeds, and a heavily fractured Tamil-Nadu DFN.