GLOBALLY OPTIMAL ROBUST ROTATION AVERAGING

Anonymous authors
Paper under double-blind review

ABSTRACT

Rotation averaging (recovering absolute orientations from pairwise relative rotations) underlies global Structure from Motion, multi-scan registration, and cryo-electron microscopy, and in all three the measurements are heavily contaminated by outliers. Robust solvers tolerate the contamination but certify nothing; certifiable relaxations certify a least-squares surrogate that outliers render meaningless. We certify the robust objective itself: maximum consensus, the NP-hard number of measurements explainable within a tolerance. Covering the rotation group $\mathrm{SO}(3)$ with cells of certified radius turns the continuous problem into two integer Markov random fields that bracket its optimum (cell-centre decoding from below, radius-inflated satisfiability from above), and a single branch and bound with an admissible node bound solves both at once. Because both ends are integers, subdivision does more than shrink the bracket: we prove it *closes* at a finite resolution set by the instance's own margin, for all but finitely many tolerances. The certificate is stated in the units that matter: the returned rotations explain C measurements, and no rotations explain more. On synthetic graphs, the bracket closes on every instance in about a minute, and the solver is the most accurate method at every outlier rate we test and the only one still standing at 90%. As an outlier-classification front end with a single parameter set, it improves the state of the art in Structure from Motion, multi-scan registration, and cryo-EM orientation recovery. The source code will be made public.

1 INTRODUCTION

Globally optimal solutions to robust geometric estimation problems are rare. Most formulations of interest, including maximum consensus (Fischler & Bolles, 1981) and its variants, are NP-hard and even APX-hard, admitting no polynomial-time approximation scheme unless P=NP (Chin et al., 2018). Practical approaches therefore rely on continuous surrogates – spectral relaxations (Singer, 2011), iteratively reweighted least squares (Chatterjee & Govindu, 2013), graduated non-convexity (Yang et al., 2020), semidefinite programming (Rosen et al., 2019; Eriksson et al., 2018; Dellaert et al., 2020) – whose tightness degrades when robustness matters most: at high outlier rates. We seek an estimate that comes with a *proof* of its quality, even under extreme contamination.

We target robust rotation averaging on $\mathrm{SO}(3)$, also known as rotation synchronization: recover absolute orientations $R_i \in \mathrm{SO}(3)$ from noisy pairwise measurements $\tilde{R}_{ij} \approx R_j R_i^\top$ on a graph, where a large fraction of the measurements may be gross outliers. The same problem arises in otherwise separate fields. (a) In Structure from Motion (SfM), rotation averaging initialises global pipelines, and its robustness bounds what the reconstruction can recover (Pan et al., 2024). (b) In multi-way point-cloud registration and pose-graph SLAM, pairwise scan alignments form a rotation graph whose gross failures (low-overlap pairs, repetitive structure) must be identified before averaging (Yew & Lee, 2021; Cadena et al., 2016). (c) In single-particle cryo-electron microscopy, common-line detections yield relative rotations with 30–80% outliers at realistic noise levels, and orientation recovery decides whether a molecule can be reconstructed (Singer, 2011).

Certificates for the coupled multi-rotation problem do exist, but for *relaxed* formulations: SE-Sync (Rosen et al., 2019), Shonan averaging (Dellaert et al., 2020), and strong duality (Eriksson et al., 2018) certify the global optimum of the least-squares (LS) objective whenever the convex relaxation is tight. Gross outliers destroy that tightness and make the LS optimum the wrong quantity to certify. *Robust* formulations have been certified only for a single unknown rotation, where branch and bound over one copy of $\mathrm{SO}(3)$ is tractable (Hartley & Kahl, 2009; Parra Bustos & Chin, 2015)

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

and truncated-LS relaxations stay manageable (Yang & Carlone, 2019) – never for the coupled problem over $\mathrm{SO}(3)^N$. The obstacle is dimension: the coupled problem lives on a $3N$-dimensional product manifold, where direct subdivision needs exponentially many cells. A second failure is quieter: even the robust certificates that exist are denominated in surrogate units (SDP suboptimality, a gap on a smoothed energy), numbers that answer no question a practitioner asks. Our question is which measurements are explained, and whether more could be.

Our route rests on a decoupling principle: *which measurements are inliers is a coarse, discrete decision that can be certified combinatorially; the estimation that remains once the consensus set is fixed is a benign LS problem.* We therefore certify maximum consensus (the number of measurements explained within a tolerance) in its own units. Cells with certified radii on a *single* copy of $\mathrm{SO}(3)$ turn the search into a pairwise Markov Random Field (MRF), two *integer* MRFs bracket the NP-hard count (cell-centre decoding from below, radius-inflated satisfiability from above), and one branch and bound solves both at once. Octree subdivision halves the inflation, and because both ends are integers, the bracket provably *closes* at a finite resolution for all but finitely many tolerances: the count is certified exactly. Under explicit margin and outlier-geometry conditions, the certified set is the true inlier set with high probability. We make two contributions:

1. A discretize-and-certify formulation of maximum-consensus rotation averaging: two integer MRF optima bracket the count at every resolution and provably close.
2. A certification stack: exact branch and bound with an admissible, potential-agnostic node bound; floored integer certificates from any admissible bound; closure and margin checks lifting the discrete conclusion to the continuum; safe integer pruning; and an exact-recovery guarantee.

With a single parameter set, the method improves or matches the state of the art in Structure from Motion, multi-scan registration, and cryo-EM orientation recovery, and a controlled synthetic study measures its accuracy and runtime across topology, outlier rate, and scale.

2 Related work

Rotation averaging and robust local solvers. Classical treatments formulate rotation averaging as least squares on $\mathrm{SO}(3)$ under various metrics (Hartley et al., 2013; Hartley & Zisserman, 2003), solved by spectral relaxation (Singer, 2011), Lie-algebraic Gauss–Newton, or iteratively reweighted least squares with a robust kernel (Chatterjee & Govindu, 2013; 2018). Robust variants down-weight suspect edges: Zhang et al. (2023) weight robust losses by edge uncertainty, graduated non-convexity anneals a surrogate towards a truncated loss (Yang et al., 2020), message-passing least squares (Shi & Lerman, 2020) and CEMP (Lerman & Shi, 2022) estimate per-edge corruption levels, and cycle-consistency filters (Zach et al., 2010; Shi et al., 2022; Li et al., 2024) reject edges violating loop constraints. Related switchable-constraint and consistency-check approaches appear in SLAM (Sünderhauf & Protzel, 2012; Agarwal et al., 2013; Mangelson et al., 2018; Lajoie et al., 2019). All are local: they report a solution, never a bound on its distance from the optimum, and beyond roughly half outliers they depend on initialisation.

Consensus maximisation and branch and bound. Maximum consensus is the historical formalisation of robust fitting (Fischler & Bolles, 1981). It is NP-hard and hard to approximate (Chin et al., 2018). Guaranteed-optimal consensus solvers exist for single-model fitting: mixed integer programming and tailored branch and bound (Li, 2009), tree search over outlier supports (Chin et al., 2017), and rotation-space search for problems with *one* unknown rotation (Hartley & Kahl, 2009; Bazin et al., 2012; Parra Bustos & Chin, 2015). All subdivide a single low-dimensional domain; over $\mathrm{SO}(3)^N$, direct subdivision needs exponentially many cells. We instead branch over per-node cell sets of a discrete MRF with an admissible combinatorial bound, and the certified quantity remains the consensus count itself.

Certifiable estimation. A parallel line certifies global optimality of *non-robust* objectives. SE-Sync (Rosen et al., 2019) and Shonan averaging (Dellaert et al., 2020) lift the problem to a higher-rank Riemannian staircase whose rank-deficient solutions certify the SDP optimum. Eriksson et al. (2018) derive strong duality with an explicit noise-dependent bound, rotation coordinate descent (Parra et al., 2021) solves the resulting relaxation at scale, and Bandeira et al. (2017; 2016) characterise when tightness holds, and Olsson et al. (2025) carry the certificate to anisotropic least squares. TEASER++ (Yang et al., 2021) and QUASAR (Yang & Carlone, 2019) extend certification to robust

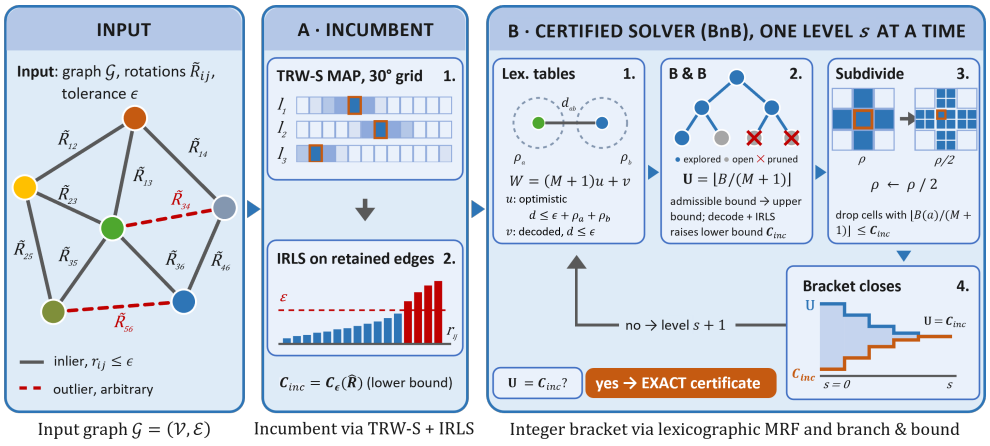

Figure 1: **Certified pipeline.** *Left:* input graph with relative rotations $\tilde{R}_{ij}$, tolerance ε, and outliers (dashed). *Middle:* MAP inference by TRW-S (Kolmogorov, 2006) on a 30° grid, followed by IRLS on the retained edges, gives the incumbent $\hat{\boldsymbol{R}}$, whose achieved count $\mathrm{C_{inc}} = C_\varepsilon(\hat{\boldsymbol{R}})$ is a valid lower bound however $\hat{\boldsymbol{R}}$ was obtained. *Right:* cells of certified radius ρ yield two integer potentials per edge, v_{ij} (decoded, $d_{ab} \le \varepsilon$) and u_{ij} (satisfiable, $d_{ab} \le \varepsilon + \rho_a + \rho_b$), folded into one lexicographic table $W = (M{+}1)u_{ij} + v_{ij}$. Branch and bound with an admissible node bound returns B, which floors to the integer cap $\mathrm{U} = \lfloor B/(M{+}1) \rfloor$; decoding and polishing its labelling can only raise $\mathrm{C_{inc}}$. Cells whose cap cannot beat $\mathrm{C_{inc}}$ are dropped, survivors are halved, and the loop repeats until the two integers meet – an exact certificate.

single-rotation and registration problems via truncated-least-squares SDPs, and two studies analyse orthogonal-group synchronization under adversarial corruption (Ling, 2023; Wang & Singer, 2013). These methods share two limitations: tightness holds only at low noise and corruption, and the certificate is an SDP gap that says nothing about which measurements are explained. Our bracket counts measurements and closes exactly.

Discrete inference for geometry. Crandall et al. (2013) pioneered treating SfM as MRF inference on a discretised pose space, solved by loopy belief propagation and refined continuously. Perry et al. (2018) analyse message passing for synchronization over compact groups. TRW-S (Kolmogorov, 2006) maximises a bound on the LP dual of MAP inference and certifies LP optimality at arc-consistent fixed points (Wainwright & Jordan, 2008). Graph cuts apply only to submodular energies (Kolmogorov & Zabih, 2004), which consensus potentials are not. In our pipeline, message passing initialises the incumbent, and any valid bound on the discrete problem, a message-passing dual included, floors into an integer certificate.

3 Maximum Consensus, Discretised and Certified

3.1 Problem: Which Measurements Are Inliers?

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a measurement graph with $N = |\mathcal{V}|$ nodes and $M = |\mathcal{E}|$ edges. Each node i carries an unknown orientation $R_i \in \mathrm{SO}(3)$; each edge (i, j) carries a measurement $\tilde{R}_{ij} \in \mathrm{SO}(3)$ of the relative rotation $R_j R_i^\top$. We use the bi-invariant angular metric $d(A, B) = \arccos\big((\mathrm{tr}(A^\top B) - 1)/2\big) \in [0, \pi]$, and write the residual of edge (i, j) at a configuration $\boldsymbol{R} = (R_1, \ldots, R_N)$ as

$$r_{ij}(\boldsymbol{R}) \;=\; d\big(\tilde{R}_{ij},\, R_j R_i^\top\big). \tag{1}$$

Inlier edges satisfy $r_{ij}(\boldsymbol{R}^*) \le \varepsilon$ at the true configuration $\boldsymbol{R}^*$ for a known tolerance ε. Outliers are arbitrary. Every residual is invariant to a global rotation $R_i \mapsto R_i Q$, so we fix the gauge by $R_1 = I$.

Our robust objective is *maximum consensus*, the historical formalisation of robust fitting (Fischler & Bolles, 1981), which seeks the configuration that explains the most measurements:

$$C_\varepsilon^* \;=\; \max_{\boldsymbol{R} \in \mathrm{SO}(3)^N} C_\varepsilon(\boldsymbol{R}), \qquad C_\tau(\boldsymbol{R}) \;=\; \sum_{(i,j)\in\mathcal{E}} \mathbf{1}\big[\, r_{ij}(\boldsymbol{R}) \le \tau \,\big]. \tag{2}$$

Problem (2) is NP-hard (Chin et al., 2018); certifying the consensus set is the load-bearing step, and the rotations follow by least squares on it.

3.2 Cells and the Integer Sandwich

Parameterisation. We parameterise rotations by angle-axis vectors $v \in \mathbb{R}^3$, $R(v) = \exp([v]_\times)$, and decompose the coordinate cube $[-\pi, \pi]^3$ (which covers the π-ball, hence all of $\mathrm{SO}(3)$) into axis-aligned cubes. A cube c with centre v_c and half-side h has representative $R_c = R(v_c)$ and *radius* $\rho_c = \sqrt{3}\,h$, certified by the following lemma. All proofs are in Appendix A.

Lemma 1 (Certified cell radius)**.** *Let $c \subset \mathbb{R}^3$ be an axis-aligned cube with centre v_c and half-side h, and let $R_c = R(v_c)$. Then $d\big(R(v), R_c\big) \le \sqrt{3}\,h$ for every $v \in c$. The bound holds uniformly, including for cubes intersecting the boundary $\|v\| = \pi$, and octree subdivision halves it.*

A *search region* assigns to each graph node i a finite set of cells $\mathcal{C}_i$. We write $\mathcal{U} = \prod_i \big(\bigcup_{c \in \mathcal{C}_i} c\big)$ for the corresponding subset of $\mathrm{SO}(3)^N$. Initially, $\mathcal{C}_i$ covers all of $\mathrm{SO}(3)$, so $\mathcal{U} = \mathrm{SO}(3)^N$.

Two integer potentials. Assigning every node one cell turns the search space into a finite pairwise Markov Random Field over the label sets $\mathcal{C}_i$. For each edge and each pair of cells (a, b), let $d_{ab} = d(\tilde{R}_{ij} R_a,\, R_b)$ be the residual the two cell *centres* produce, and attach two indicator potentials:

$$\underbrace{v_{ij}(a,b) \;=\; \mathbf{1}\big[\, d_{ab} \le \varepsilon \,\big]}_{\text{decoded consensus}}, \qquad \underbrace{u_{ij}(a,b) \;=\; \mathbf{1}\big[\, d_{ab} \le \varepsilon + \rho_a + \rho_b \,\big]}_{\text{optimistic satisfiability}}. \tag{3}$$

Here v_{ij} asks whether the edge is explained when both nodes sit at their cell centres; u_{ij} asks whether any placement in the two cells could explain it. Both define integer MRFs over the same labels, with optima V^* and U^*. Their opposite roles are one-line consequences of bi-invariance:

Lemma 2 (Displacement, and what each potential certifies)**.** *Fix an edge (i, j) and cells $a \in \mathcal{C}_i$, $b \in \mathcal{C}_j$ with radii ρ_a, ρ_b. For every $R_i \in a$, $R_j \in b$,*

$$\big|\, r_{ij}(\boldsymbol{R}) - d_{ab} \,\big| \;\le\; \rho_a + \rho_b. \tag{4}$$

Consequently: (a) if $r_{ij}(\boldsymbol{R}) \le \varepsilon$ then $u_{ij}(a, b) = 1$, so for every $\boldsymbol{R} \in \mathcal{U}$, choosing for each node a cell containing it, $C_\varepsilon(\boldsymbol{R}) \le \sum_e u_{ij}(\ell_i(\boldsymbol{R}), \ell_j(\boldsymbol{R})) \le U^*$. (b) The configuration that places each node at its cell centre has $r_{ij} = d_{ab}$ exactly, so a labelling's v-value is *the consensus count its decode achieves, and $V^* \le C^*_\varepsilon(\mathcal{U})$.*

The NP-hard count is thus squeezed between two computable integers:

Theorem 3 (Integer consensus bracket)**.** *Let $\mathcal{U}$ be any search region, $\hat{\boldsymbol{R}} \in \mathcal{U}$ any feasible configuration, and $\mathrm{U} \ge U^*$ any upper bound on the u-MRF optimum. Then*

$$\max\big\{V^*,\, C_\varepsilon(\hat{\boldsymbol{R}})\big\} \;\le\; C^*_\varepsilon(\mathcal{U}) \;\le\; U^* \;\le\; \mathrm{U}, \qquad \textit{and} \qquad U^* \;\le\; C^*_{\varepsilon + 2\rho_{\max}}(\mathcal{U}), \tag{5}$$

where $\rho_{\max}$ is the largest active cell radius. In particular, if $C_\varepsilon(\hat{\boldsymbol{R}}) = \mathrm{U}$, then $\hat{\boldsymbol{R}}$ attains the maximum consensus count over $\mathcal{U}$: its consensus set is a *maximum-consensus set, and no configuration anywhere in the region explains more measurements.*

Both ends are integers the solver computes, in the units that matter, and the last inequality caps the price of discretisation: an inflation of two cell radii that subdivision drives to zero.

One search for both bounds. The two MRFs share their label sets, so one search can serve both. We fold them into a single lexicographic table

$$W_{ij}(a,b) \;=\; (M{+}1)\, u_{ij}(a,b) + v_{ij}(a,b), \tag{6}$$

where $M = |\mathcal{E}|$ is the number of edges. Read W as a two-digit number in base $M{+}1$ with u as the high digit: a labelling's total v-score is at most M (one point per edge), so it can never carry into the u-digit. Maximising $\sum_e W$ once therefore maximises the u-count first and, among u-optimal labellings, the decoded count: one branch-and-bound run returns both the upper bound $U^* = \lfloor \max W/(M{+}1) \rfloor$ *and* the best available decode. The v-digit also breaks the ties that stall search on pure 0/1 potentials. Because subdivision halves every radius, the bracket (5) does more than tighten. It provably *closes*:

Theorem 4 (Finite-resolution exactness)**.** *For $\tau \geq 0$ let $g(\tau) = C^*_\tau(\mathcal{U})$. Then g is a non-decreasing, right-continuous, integer-valued step function with at most M breakpoints. If ε is not a breakpoint of g (equivalently, if g is constant on $[\varepsilon - \delta,\ \varepsilon + \delta]$ for some $\delta > 0$), then as soon as every active cell radius satisfies $\rho_{\max} \leq \min\{\delta, \varepsilon\}/2$,*

$$V^* \ = \ C^*_\varepsilon(\mathcal{U}) \ = \ U^*, \tag{7}$$

and the decode of any W-optimal labelling achieves the maximum count. Octree subdivision halves $\rho_{\max}$ per level, so the bracket closes exactly after at most $\lceil \log_2\big(2\rho_0 / \min\{\delta, \varepsilon\}\big) \rceil$ levels from root radius ρ_0. Moreover, under full octree refinement U^* is non-increasing from level to level. At a breakpoint, only the lower end is at risk: right-continuity still closes U^* onto $g(\varepsilon)$ at a finite resolution, while the guaranteed lower end may lag by up to the jump of g, a residual gap that the solver reports. The number of branch-and-bound nodes remains exponential in the worst case, consistent with the NP-hardness of (2) (Chin et al., 2018).

In words, the explainable count is a staircase in the tolerance (Figure 4). Unless ε lands exactly on a step edge, an event of probability zero under continuous noise, the staircase is flat around ε. Once cells are finer than half that flat stretch, optimism and achievement count the same edges. The margin also has an a posteriori face: if every residual of the incumbent clears the tolerance by more than $2(\rho_a + \rho_b)$ on its cell pair, no placement within those cells can flip any edge's classification.

3.3 Branch and bound with an admissible node bound

We search the lexicographic MRF, a maximisation over $\prod_i |\mathcal{C}_i|$ labellings, depth-first over partial assignments, with an upper bound that is cheap, tight, and *admissible*: it never underestimates the best completion of a partial assignment, so pruning against it never discards the optimum. Let X be the set of already-fixed nodes and F the free ones. Then

$$B_{\text{node}} = \sum_{\substack{(i,j)\in\mathcal{E}\\ i,j\in X}} W_{ij}(\ell_i, \ell_j) \ + \ \sum_{i\in F} \max_{a\in\mathcal{C}_i} \Big[\sum_{\substack{j\in N(i)\\ j\in X}} W_{ij}(a, \ell_j) \ + \ \tfrac{1}{2} \sum_{\substack{j\in N(i)\\ j\in F}} \max_{b\in\mathcal{C}_j} W_{ij}(a, b) \Big]. \tag{8}$$

The bound scores the decided part exactly and lets every undecided node claim the best it could collect: full credit from decided neighbours and half from undecided ones, whose other half is claimed at the far endpoint.

Theorem 5 (Admissible node bound)**.** *Let X be a set of fixed nodes with labels $\{\ell_i\}_{i\in X}$ and $F = \mathcal{V} \setminus X$. For every labelling $\boldsymbol{\ell}$, extending the partial assignment with $\ell_i \in \mathcal{C}_i$ for all $i \in F$, the table value satisfies $W(\boldsymbol{\ell}) \leq B_{\text{node}}$ of (8). The bound assumes nothing about the potentials or label sets.*

Because nothing about W enters, the bound applies at every level and to heterogeneous cell sets, at cost $\mathcal{O}(M\bar{F})$ per search node $\bar{F}$ the mean number of active cells per node. The integer structure of W pays off twice more. First, any bound converts to a count certificate by division:

Proposition 6 (Floored certificates without exhaustive search)**.** *Let B be any admissible upper bound on the lexicographic MRF (6) over $\mathcal{U}$ – the exact optimum, the root node bound (8) alone, or any bound in between. Then $C^*_\varepsilon(\mathcal{U}) \leq \lfloor B/(M{+}1) \rfloor$.*

Certification never requires solving the problem exactly: the root bound alone floors to a valid, if looser, cap; an exhausted search makes it exact. Second, unit spacing makes ties crisp: the search discards value-tied subtrees and returns *one* W-optimal labelling, all Theorem 4 needs.

3.4 The certified solver

Figure 1 assembles the three stages; Algorithm 1 in Appendix B gives pseudocode.

Stage A: incumbent from message passing. TRW-S (Kolmogorov, 2006) performs MAP inference on a pairwise MRF over a 30° $\mathrm{SO}(3)$ grid with graded potentials, and IRLS on the retained edges refines it into a strong feasible configuration $\hat{\boldsymbol{R}}$. Its count $C_\varepsilon(\hat{\boldsymbol{R}})$ is the bracket's initial lower end, valid however $\hat{\boldsymbol{R}}$ was obtained. Any learned estimator or baseline could serve.

Stage B: search region. When N and the cell counts are small, we keep every root cell, and the certificate is *global*. Otherwise, we keep for each node the cells whose conditional reward, given the incumbent's neighbours (Remark 9), exceeds a fraction $\rho_{\text{ret}} = 0.3$ of the best, always including

the cell containing the incumbent. This heuristic restriction has a failsafe: if the bracket fails to close, the solver re-runs with the retention relaxed until, in the limit, every root cell is kept and the certificate is global. Our experiments never triggered the escalation.

Stage C: bound, decode, prune, subdivide. At each level, we build lexicographic tables (6), solve the MRF by branch and bound, and floor the returned bound to the integer U on $C^*_\varepsilon(\mathcal{U})$ (Proposition 6). We decode the labelling to cell centres, polish it with IRLS on the edges it marks satisfiable, and keep whichever configuration explains the most edges as the new incumbent, raising $\mathrm{C}_{\mathrm{inc}} = C_\varepsilon(\hat{\boldsymbol{R}})$. We discard cells that cannot contribute to a better solution:

Proposition 7 (Safe integer pruning)**.** *Let* $\mathrm{C}_{\mathrm{inc}}$ *be the consensus count of any feasible configuration and let* $B_i(a)$ *be an admissible upper bound on the lexicographic MRF restricted to* $\ell_i = a$*. If* $\lfloor B_i(a)/(M{+}1) \rfloor \le \mathrm{C}_{\mathrm{inc}}$*, then no configuration in the cell* a *of node* i *explains more than* $\mathrm{C}_{\mathrm{inc}}$ *edges. Removing* a *from* $\mathcal{C}_i$ *therefore preserves every configuration explaining at least* $\mathrm{C}_{\mathrm{inc}} + 1$ *edges, and the resulting bracket remains valid. Tied optima may be discarded: pruning preserves the certified* value*, not every argmax.*

A cell survives only by promising a full extra edge. We subdivide the survivors, halving their radius. The loop stops when the two integers meet: closure $C_\varepsilon(\hat{\boldsymbol{R}}) = \mathrm{U}$ is the certificate (Theorem 3), reached at a finite level for all but degenerate tolerances (Theorem 4).

The method is any-time. Both ends of the bracket are valid at every level: $\mathrm{C}_{\mathrm{inc}}$ is the count of a configuration in hand, and U floors from whatever admissible bound the search has reached, open subtrees included (Proposition 6). When compute is limited, the loop can stop after any level, or inside one, and still return the current best $\hat{\boldsymbol{R}}$ with a bracket that is only wider, never wrong (Appendix B.1, Figure 5).

3.5 When is the certified set the truth?

Everything so far certifies a *count*, not correctness: two edge sets can both have size C^*, and the solver may certify the wrong one. Whether the certified set is the true inlier set depends on the data: the inliers must have room to spare, and the outliers nowhere to hide.

Theorem 8 (Exact recovery of the inlier set, w.h.p.)**.** *Let a generator fix a configuration* $\boldsymbol{R}^*$ *and a partition of* $\mathcal{E}$ *into an inlier set* $\mathcal{I}^*$ *with* $C^* = |\mathcal{I}^*|$ *and* M_{out} *outlier edges. Fix a resolution: label sets* $\mathcal{C}_i$ *of cells of radius at most* ρ*. Assume: (i)* inner margin*: every inlier edge has* $\varepsilon - r_{ij}(\boldsymbol{R}^*) > 2\rho$*; (ii)* coverage*: for every node, some cell of* $\mathcal{C}_i$ *contains* R_i^**; (iii)* outer margin*: each outlier measurement avoids its* candidate-target neighbourhood $\mathcal{N}_{ij} = \bigcup_{a\in\mathcal{C}_i,\, b\in\mathcal{C}_j}\{R : d(R, R_b R_a^\top) \le \varepsilon + \rho_a + \rho_b\}$ *except with probability* $\Pr[\tilde{R}_{ij} \in \mathcal{N}_{ij}] \le \eta$*. Then, with probability at least* $1 - M_{\mathrm{out}}\,\eta$*, simultaneously: no configuration in* $\mathcal{U}$ *satisfies any outlier edge at* ε*, so* $\mathcal{I}^*$ *is the unique maximum-consensus set over* $\mathcal{U}$ *with* $C^*_\varepsilon(\mathcal{U}) = C^*$*; and every* W*-optimal labelling realises exactly* $\mathcal{I}^*$*.*

Why it holds. Snapping both endpoints of an edge to their cell centres moves its relative rotation by at most $\rho_a + \rho_b$, so with the inner margin (i) the truth's rounded labelling keeps every inlier. Hypothesis (iii) is the same inflation seen from the other side: $\mathcal{N}_{ij}$ is everything the active cells could predict for that edge, inflated by the slack that makes the upper bound sound, so no configuration or labelling can explain an outlier outside it. The maximum is then C^*, reached only by explaining all of $\mathcal{I}^*$ and nothing else, and the tie-break forces every optimal decode to do so; (ii) only asks that Stage B kept the truth's cells. The guarantee does not rest on the truth outscoring a runner-up: any other set of C^* edges contains an outlier, which nothing explains, so none exists.

What the probability means. Failure requires *some* outlier to land in *its* neighbourhood, so $M_{\mathrm{out}}\,\eta$ is a plain union bound that needs no independence between outliers. The event has volume, not combinatorial, character: by bi-invariance, $\mathcal{N}_{ij}$ lies in one ball of radius $\tau = \varepsilon + 2\rho + s_i + s_j$ about the incumbent's prediction for the edge, s_i being the spread of node i's active cells about the incumbent, and for Haar-uniform outliers that ball has mass $\eta \le (\tau - \sin\tau)/\pi \sim \tau^3/6\pi$. The bound is informative only for a compact region: at full retention, the spread is π and the theorem is silent, whereas pruning shrinks ρ and the spread together, and η tends to the mass of the ε-ball. Exact recovery is therefore a statement about the pruned region, whereas the bracket of Theorem 3 holds over any region at every resolution. Since the hypotheses are sufficient, not necessary, and concern the generator rather than anything checkable, we keep both layers: the bracket states *what* is certified on an instance, and this theorem explains *why* the certified set tends to be the truth.

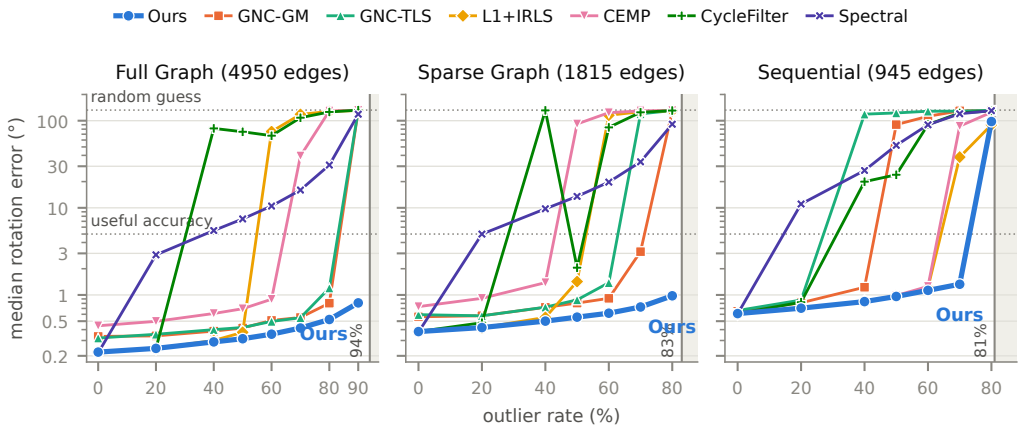

Figure 2: **Rotation synchronization on synthetic graphs.** Median relative-rotation error (log scale, ↓) against outlier rate over 20 random seeds, for three connectivity patterns; graphs have $N{=}100$ nodes, Gaussian inlier noise $\sigma{=}1°$, and outlier edges carrying uniformly random rotations. Every point is measured: methods that break down climb towards the random-guess error (132°, dotted), and the second reference line marks 5°. The vertical line and the shading to its right mark the identifiability limit (94% for full, 83% for sparse, 81% for sequential; Section 4.1). Our method (thick blue) is lowest everywhere, and the sparser the graph, the larger the margin.

4 Experiments

We test on synthetic graphs and in three domains built on rotation averaging: Structure from Motion, multi-scan point cloud registration, and cryo-electron microscopy. All experiments use a shared parameter set (30° grid, $\varepsilon = 5°$); nothing is tuned per scene.

4.1 Controlled evaluation on synthetic graphs

We generate graphs with $N{=}100$ nodes and three connectivity patterns: FULL (all pairs connected, 4,950 edges), SPARSE (20 random neighbors per node, 1,815 edges, mean degree 36), and SEQUENTIAL (sliding-window connectivity, 945 edges). Inlier edges carry Gaussian noise ($\sigma{=}1°$); outlier edges carry uniformly random rotations. We sweep outlier rates from 0 to 80% (to 90% on full graphs) over 20 random seeds, and report the median all-pairs relative-rotation error.

Figure 2 shows the results. Our method has the lowest error at every rate in every topology, and the informative comparison is *where* each baseline leaves the useful regime. On fully connected graphs, redundancy is abundant: the graduated-non-convexity solvers stay competitive to 80%, while CycleFilter (Zach et al., 2010), L1+IRLS (Chatterjee & Govindu, 2013), and CEMP (Lerman & Shi, 2022) collapse at 40%, 60%, and 70%, respectively, and spectral synchronization (Singer, 2011) degrades steadily; at 90% ours is the only method left standing (0.81° against 119–132°). As redundancy drops, the gap widens: on sparse graphs our error stays below one degree across the whole sweep (0.98° at 80%), whereas the best baseline is 4× worse at 70% and every baseline is at random-guess accuracy at 80%. On sequential graphs GNC-TLS fails by 40% and GNC-GM by 50%; L1+IRLS and CEMP track us to 60% and collapse at 70%, where we still return 1.33°. At 80%, every method, ours included, fails for an information-theoretic reason: a typical instance then contains a camera with fewer than two inlier edges, which nothing corroborates and no estimator can place. The vertical lines in Figure 2 mark this limit, measured on the same instances; our only failure, sequential at 80%, sits at its doorstep.

Method	t (s) ↓ $N{=}10^2$	$N{=}10^3$
Ours	0.38	**109**
GNC-GM	1.46	365
GNC-TLS	1.99	985
L1+IRLS	**0.07**	30†
CEMP	0.20†	2†
CycleFilter	0.24	1†
Spectral	0.01†	3†

Table 1: Mean runtime on sparse graphs at a 50% outlier rate, at $N{=}100$ (1,815 edges) and $N{=}1000$ (19,809 edges). †diverged at that size (median error above 5°); bold marks the fastest method that stays accurate.

Table 1 reports runtimes at the plotted size and at $N{=}1000$. On the small graphs, our solver runs within the same order of magnitude as the baselines, well under a second per instance. On the

Dataset	Method	AUC@5°	@10°	t (min)
ETH3D	GLOMAP	80.6	85.8	**0.4**
	+ Ours	**85.0**	**90.9**	0.5
Haiper	GLOMAP	67.0	75.5	**0.1**
	+ Ours	**75.3**	**83.9**	**0.1**
Heritage	GLOMAP	75.2	78.9	1.5
	+ Ours	**78.4**	**82.3**	**1.1**
PhotoT.	GLOMAP	77.0	82.2	10.3
	+ Ours	**77.6**	**82.7**	**9.7**
IMC24	GLOMAP	21.3	24.3	23.1
	+ Ours	**22.6**	**28.4**	**13.3**

Table 2: **Structure-from-motion pose accuracy.** Our method integrated into GLOMAP (Pan et al., 2024). AUC@T (%) is the normalised area under the cumulative recall curve up to threshold T; runtime t in minutes. Bold: best per benchmark.

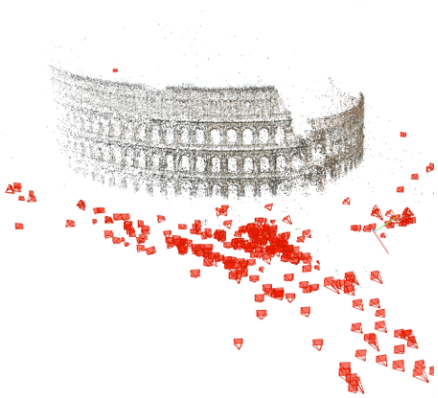

Figure 3: **Example SfM reconstruction.** Camera poses and sparse point cloud for the Colosseum scene from the PhotoTourism dataset, reconstructed by GLOMAP with our rotation-averaging front end.

large graphs, the only methods faster than ours are those that collapse there (daggered); among the methods that still return a usable estimate, ours is the fastest, 3–9× ahead of the GNC pair.

4.2 Structure from Motion

In global Structure from Motion, matched features (Lowe, 1999; DeTone et al., 2018) and essential-matrix fitting (Hartley & Zisserman, 2003) produce a *view graph* of relative rotations from which rotation averaging recovers each camera's orientation. Symmetric structures and repetitive textures produce relative rotations that are internally consistent but wrong, and under local optimisation these mutually reinforcing outliers fold the reconstruction onto itself.

We integrate our solver into the GLOMAP pipeline (Pan et al., 2024) as an outlier-rejection filter preceding its standard L1+IRLS (Chatterjee & Govindu, 2018) refinement. The baseline and ours use identical features and matching, so only the rotation averaging step differs. Table 2 reports the area under the cumulative recall curve (AUC) of the pose error at 5° and 10° on five benchmarks. On ETH3D (Schöps et al., 2017) (13 scenes, laser-scanned ground truth) we gain $+4.4$ AUC at 5° and $+5.1$ at 10° (Figure 3). The improvement is largest on heavily corrupted view graphs: $+8.3$ at 5° on IMC23-HAIPER and $+4.1$ at 10° on IMC24. On PhotoTourism (15 large tourist collections) all thresholds improve, with near-identical results on well-conditioned scenes. Total reconstruction time, rotation averaging included, is comparable or lower with our front end, since GLOMAP converges faster on the cleaned edge set.

4.3 Multi-scan Point Cloud Registration

Assembling overlapping 3D scans into one model requires estimating each pairwise rigid transformation and enforcing global consistency. Many pairwise registrations fail silently (limited overlap, repetitive geometry) and produce outlier edges that only global reasoning can tell from correct ones. We obtain pairwise alignments from PREDATOR features (Huang et al., 2021) with LO-RANSAC (Chum et al., 2003), pass the pose graph to each rotation averaging method, and report recall at 10° and 15° (fraction of pairwise rotations below the threshold).

On 3DMatch (Yew & Lee, 2021) (8 indoor scenes) our method achieves the highest recall at both thresholds, 79.4% at 10° and 87.9% at 15°, ahead of L1+IRLS (73.7/79.4) and GNC-GM (68.2/82.5); see Table 3. On the harder 3DLoMatch (same scenes, 10–30% overlap), the advantage widens to nearly 10 and 11 points: global outlier rejection matters most where pairwise registration fails silently. On KITTI, our method again scores highest.

Dataset	Metric	L1+IRLS	Spectral	CycleFlt	GNC-TLS	GNC-GM	Ours
3DMatch	Recall@10°	73.7	69.4	58.6	64.7	68.2	**79.4**
	Recall@15°	79.4	77.7	66.2	80.3	82.5	**87.9**
3DLoMatch	Recall@10°	39.7	33.1	26.1	14.3	25.5	**49.3**
	Recall@15°	49.5	48.4	34.5	22.2	38.0	**60.7**
KITTI	Recall@10°	58.4	34.8	57.3	38.8	52.2	**60.1**
	Recall@15°	64.8	41.4	63.5	42.3	58.1	**66.3**

Table 3: **Multi-scan point cloud registration.** Recall (%, ↑) of pairwise rotations at 10° and 15°, averaged across scenes; pose graphs from PREDATOR + LO-RANSAC. Bold: best per row.

SNR (out.)	Spectral	CycleFlt	CEMP	L1+IRLS	GNC-TLS	GNC-GM	**Ours**
1.0 (20%)	**32.0**	**32.0**	**32.0**	**32.0**	**32.0**	**32.0**	**32.0**
1/16 (36%)	28.6	10.4	12.6	30.2	25.4	28.8	**30.2**
1/32 (56%)	20.8	7.6	6.6	24.6	18.8	24.0	**24.8**
1/64 (77%)	11.0	6.2	6.0	10.0	9.0	11.2	**16.6**

Table 4: **Cryo-EM reconstruction from common-line pairwise rotations.** Highest resolved Fourier shell (FSC>0.143, out of 32; higher is better), averaged over 5 seeds; $N{=}500$ projections of the 80S ribosome (EMD-2660) at $65{\times}65$ pixels. Bold: best per row.

4.4 Cryo-electron microscopy orientation determination

Single-particle cryo-EM images many copies of a molecule, each a 2D projection at an unknown orientation. Recovering these orientations is the central problem of *ab initio* reconstruction, and it is a rotation synchronization problem (Shkolnisky & Singer, 2012): common lines constrain the relative rotation of two projections, and triplet voting resolves the remaining ambiguity. At realistic signal-to-noise ratios, 30–80% of the resulting estimates are outliers (Wang et al., 2013). We compute pairwise rotations with ASPIRE's triplet voting (Shkolnisky & Singer, 2012) on common lines perturbed by Gaussian noise of standard deviation $\min(20^\circ, 5/\sqrt{\mathrm{SNR}})$ and partially replaced at random, following Wang et al. (2013) (from $\sim$20% outliers at SNR$=1$ to $\sim$77% at $1/64$). We reconstruct from $N{=}500$ clean projections of the 80S ribosome (EMD-2660) and report the highest Fourier shell with FSC>0.143 against ground truth, out of 32.

At 20% outliers every method achieves full resolution (Table 4). At 36% CycleFilter and CEMP collapse to 10–13 shells while we maintain 30; at 56% both fail entirely ($\leq$8 shells) while our method and L1+IRLS reach 25. In the critical regime, SNR$=1/64$ (77% outliers, characteristic of small proteins and heterogeneous complexes), every baseline loses most high-frequency information (best competitor 11 shells) while ours resolves 17, a $1.5\times$ improvement; ours is the only solver that retains more than half the available detail.

5 Conclusion

We revisited robust rotation averaging through a decoupling principle: which measurements are inliers is a discrete decision that can be certified combinatorially; the estimation that remains is benign. A cell decomposition of a single copy of $\mathrm{SO}(3)$ brackets the NP-hard consensus count between two integer MRF optima, solved by one admissible branch and bound and driven to *exact* closure by subdivision at a resolution the instance's own margin determines. The certificate is stated in the units that matter: the returned rotations explain a stated number of measurements, and no rotations explain more. It costs nothing in accuracy: the same solver has the lowest error at every outlier rate and topology we tested, is the only method standing at 90%, runs 3–9$\times$ faster at 1000 nodes than the two baselines that do not collapse, and, with one parameter set, improves Structure from Motion, multi-scan registration, and cryo-EM orientation recovery. Certification costs time: seconds to minutes against milliseconds for local solvers. In return, the any-time loop turns "the solver converged" into a closed integer bracket that a pipeline can act on. The construction applies to consensus problems on any compact group with certified covering radii. The code will be released.

AI USE STATEMENT

We used large language models as tools throughout this work, under our direction and with full responsibility for the content: (i) to draft and polish the text of the paper; (ii) to implement the benchmarking and evaluation scripts (data generation, baseline wrappers, parameter sweeps); (iii) to optimise the solver implementation at a low level, without changing the algorithm; and (iv) to verify and polish the proofs, checking each step of the arguments in Appendix A and tightening their presentation. The formulation, the theoretical results, the algorithm design, and the experimental conclusions are our own, and we thoroughly reviewed all AI-produced text, code, and proof edits.

ETHICS STATEMENT

This work concerns certified optimisation for rotation averaging, a geometric estimation primitive. All experiments use synthetic data or public benchmarks (ETH3D, Image Matching Challenge, PhotoTourism, 3DMatch, KITTI, and the public EMD-2660 cryo-EM density map) under their published licences; no human subjects, personally identifiable information, or sensitive attributes are involved. Improved and certifiable 3D reconstruction has broadly beneficial uses (mapping, robotics, structural biology, cultural-heritage digitisation); we are not aware of ethical risks specific to this method beyond those generic to 3D reconstruction.

REPRODUCIBILITY STATEMENT

The method is fully specified in the paper: the objective and cell construction in Section 3, the complete procedure in Algorithm 1, the admissible bound in (8), and proofs of all formal claims in Appendix A. The synthetic data-generation protocol (graph topologies, noise and outlier models, and metrics) is given in Section 4.1 and Appendix B, and every default hyperparameter ships with the released code. Certificates are per-instance and checkable from the solver's outputs alone: the reported lower end is the consensus count of the returned configuration (recomputable from the returned rotations in one pass over the measurements), and the reported upper end is an integer that dominates it by construction. Source code (a C++ core with Python bindings, and the evaluation scripts) will be released upon publication.

REFERENCES

P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard. Robust map optimization using dynamic covariance scaling. In *Proc. ICRA*, pp. 62–69, 2013.

A. S. Bandeira, N. Boumal, and V. Voroninski. On the low-rank approach for semidefinite programs arising in synchronization and community detection. In *Proc. COLT*, pp. 361–382, 2016.

A. S. Bandeira, N. Boumal, and A. Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *Math. Program.*, 163:145–167, 2017.

Jean-Charles Bazin, Yongduek Seo, and Marc Pollefeys. Globally optimal consensus set maximization through rotation search. In *Proc. ACCV*, pp. 539–551, 2012.

C. Cadena et al. Past, present, and future of simultaneous localization and mapping: toward the robust-perception age. *IEEE Trans. Robot.*, 32:1309–1332, 2016.

A. Chatterjee and V. M. Govindu. Efficient and robust large-scale rotation averaging. In *Proc. ICCV*, pp. 521–528, 2013.

Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):958–972, 2018.

T.-J. Chin, Z. Cai, and F. Neumann. Robust fitting in computer vision: easy or hard? In *Proc. ECCV*, pp. 715–730, 2018.

Tat-Jun Chin, Pulak Purkait, Anders Eriksson, and David Suter. Efficient globally optimal consensus maximisation with tree search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):758–772, 2017.

G. S. Chirikjian and A. B. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, 2001.

Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized RANSAC. In *Proc. DAGM*, pp. 236–243. Springer, 2003.

D. Crandall, A. Owens, N. Snavely, and D. Huttenlocher. SfM with MRFs: discrete-continuous optimization for large-scale structure from motion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35 (12):2841–2853, 2013.

F. Dellaert, D. Rosen, J. Wu, R. Mahony, and L. Carlone. Shonan rotation averaging: global optimality by surfing $\mathrm{SO(p)}^n$. In *Proc. ECCV*, pp. 292–308, 2020.

Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proc. CVPR Workshops*, pp. 224–236, 2018.

A. Eriksson, C. Olsson, F. Kahl, and T.-J. Chin. Rotation averaging and strong duality. In *Proc. CVPR*, pp. 127–135, 2018.

Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24 (6):381–395, 1981.

R. Hartley, J. Trumpf, Y. Dai, and H. Li. Rotation averaging. *Int. J. Comput. Vis.*, 103:267–305, 2013.

Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

Richard I. Hartley and Fredrik Kahl. Global optimization through rotation space search. *International Journal of Computer Vision*, 82(1):64–79, 2009.

W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.*, 58(301):13–30, 1963.

Shengyu Huang, Zan Gojcic, Mikhail Usvyatsov, Andreas Wieser, and Konrad Schindler. PREDATOR: Registration of 3d point clouds with low overlap. In *Proc. CVPR*, pp. 4267–4276, 2021.

V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28:1568–1583, 2006.

V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Trans. Pattern Anal. Mach. Intell.*, 26:147–159, 2004.

P.-Y. Lajoie, S. Hu, G. Beltrame, and L. Carlone. Modeling perceptual aliasing in SLAM via discrete–continuous graphical models. *IEEE Robot. Autom. Lett.*, 4:1232–1239, 2019.

G. Lerman and Y. Shi. Robust group synchronization via cycle-edge message passing. *Found. Comput. Math.*, 22:1665–1741, 2022.

Hongdong Li. Consensus set maximization with guaranteed global optimality for robust geometry estimation. In *Proc. ICCV*, pp. 1074–1080, 2009.

S. Li, Y. Shi, and G. Lerman. Efficient detection of long consistent cycles and its application to distributed synchronization. In *Proc. CVPR*, pp. 5260–5269, 2024.

S. Ling. Solving orthogonal group synchronization via convex and low-rank optimization: tightness and landscape analysis. *Math. Program.*, 200:589–628, 2023.

David G Lowe. Object recognition from local scale-invariant features. In *Proc. ICCV*, volume 2, pp. 1150–1157. Ieee, 1999.

J. G. Mangelson, D. Dominic, R. M. Eustice, and R. Vasudevan. Pairwise consistent measurement set maximization for robust multi-robot map merging. In *Proc. ICRA*, pp. 2916–2923, 2018.

C. Olsson, Y. Lochman, J. Malmport, and C. Zach. Certifiably optimal anisotropic rotation averaging. In *Proc. ICCV*, pp. 14856–14865, 2025.

L. Pan, D. Barath, M. Pollefeys, and J. L. Schönberger. Global structure-from-motion revisited. In *Proc. ECCV*, 2024.

Á. Parra, S.-F. Ch'ng, T.-J. Chin, A. Eriksson, and I. Reid. Rotation coordinate descent for fast globally optimal rotation averaging. In *Proc. CVPR*, pp. 4298–4307, 2021.

Alvaro Parra Bustos and Tat-Jun Chin. Guaranteed outlier removal for rotation search. In *Proc. ICCV*, pp. 2165–2173, 2015.

A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra. Message-passing algorithms for synchronization problems over compact groups. *Commun. Pure Appl. Math.*, 71:2275–2322, 2018.

D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard. SE-Sync: a certifiably correct algorithm for synchronization over the special Euclidean group. *Int. J. Robot. Res.*, 38:95–125, 2019.

T. Schöps et al. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proc. CVPR*, pp. 3260–3269, 2017.

Y. Shi and G. Lerman. Message passing least squares framework and its application to rotation synchronization. In *Proc. ICML*, pp. 8796–8806, 2020.

Y. Shi, C. Wyeth, and G. Lerman. Robust group synchronization via quadratic programming. In *Proc. ICML*, pp. 20095–20105, 2022.

Y. Shkolnisky and A. Singer. Viewing direction estimation in cryo-EM using synchronization. *SIAM J. Imaging Sci.*, 5:1088–1110, 2012.

A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30:20–36, 2011.

N. Sünderhauf and P. Protzel. Switchable constraints for robust pose graph SLAM. In *Proc. IROS*, pp. 1879–1884, 2012.

M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.*, 1:1–305, 2008.

L. Wang and A. Singer. Exact and stable recovery of rotations for robust synchronization. *Inf. Inference*, 2:145–193, 2013.

L. Wang, A. Singer, and Z. Wen. Orientation determination of cryo-EM images using least unsquared deviations. *SIAM J. Imaging Sci.*, 6(4):2450–2483, 2013.

H. Yang, P. Antonante, V. Tzoumas, and L. Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. *IEEE Robot. Autom. Lett.*, 5: 1127–1134, 2020.

H. Yang, J. Shi, and L. Carlone. TEASER: fast and certifiable point cloud registration. *IEEE Trans. Robot.*, 37:314–333, 2021.

Heng Yang and Luca Carlone. A quaternion-based certifiably optimal solution to the wahba problem with outliers. In *Proc. ICCV*, pp. 1665–1674, 2019.

Z. J. Yew and G. H. Lee. Learning iterative robust transformation synchronization. In *Proc. 3DV*, pp. 1206–1215, 2021.

C. Zach, M. Klopschitz, and M. Pollefeys. Disambiguating visual relations using loop constraints. In *Proc. CVPR*, pp. 1426–1433, 2010.

G. Zhang, V. Larsson, and D. Barath. Revisiting rotation averaging: uncertainties and robust losses. In *Proc. CVPR*, pp. 17215–17224, 2023.

A PROOFS

How to read these proofs. Two facts carry everything: a cell is a ball of known radius around its representative (Lemma 1), and moving a configuration to cell centres changes every residual by at most the two cell radii (Lemma 2). Every later proof is bookkeeping on top of these: the bracket chains the displacement bound with definitions, exactness asks when the displacement is smaller than the room the instance leaves around the tolerance, the search bounds are counting arguments, and recovery asks when no outlier can survive the displacement. Each proof opens with its idea in words and names its steps by what they establish.

Throughout, $d(A,B) = \arccos\big((\mathrm{tr}(A^\top B)-1)/2\big)$ is the angular metric on $\mathrm{SO}(3)$. It is bi-invariant, $d(QA,QB) = d(AQ,BQ) = d(A,B)$ for every $Q \in \mathrm{SO}(3)$, and symmetric under transposition, $d(A^\top,B^\top) = d(A,B)$; these two properties are all we use of it.

A.1 PROOF OF LEMMA 1 (CERTIFIED CELL RADIUS)

Idea. The angle-axis map $v \mapsto R(v)$ never stretches distances: vectors $\|u-v\|$ apart give rotations at most $\|u-v\|$ apart in angle ($\mathrm{SO}(3)$ has non-negative curvature, so the exponential map is distance non-increasing). The farthest point of a cube from its centre is a corner, at distance $\sqrt{3}\,h$, which is therefore a valid radius on all of $\mathbb{R}^3$; the boundary of the π-ball needs no correction.

Step 1 (reduce to a Lipschitz bound). We prove the global statement

$$d\big(R(u),R(v)\big) \;\le\; \|u-v\| \qquad \text{for all } u,v \in \mathbb{R}^3. \tag{9}$$

Given it, every v in a cube with centre v_c and half-side h has $\|v-v_c\| \le \sqrt{3}\,h$, hence $d(R(v),R_c) \le \sqrt{3}\,h$.

Step 2 (a rotation angle is twice a distance on the 3-sphere). Represent rotations by unit quaternions $q(w) = \big(\cos(\|w\|/2),\, \sin(\|w\|/2)\, w/\|w\|\big) \in S^3$, with $q(0) = (1,0,0,0)$. The covering $S^3 \to \mathrm{SO}(3)$ is two-to-one (q and $-q$ are the same rotation), and

$$\begin{aligned} d\big(R(u),R(v)\big) &= 2\arccos\big|\langle q(u),q(v)\rangle\big| = 2\min\big\{d_{S^3}(q(u),q(v)),\, d_{S^3}(q(u),-q(v))\big\} \\ &\le 2\, d_{S^3}\big(q(u),q(v)\big), \end{aligned} \tag{10}$$

where d_{S^3} is the great-circle metric. It therefore suffices to show that $q:\mathbb{R}^3 \to S^3$ is $\frac{1}{2}$-Lipschitz.

Step 3 (the quaternion map is $\frac{1}{2}$-Lipschitz). Fix $w \neq 0$ with $r = \|w\|$, $n = w/r$, and split a perturbation into a radial and a tangential part, $\delta = a\,n + b\,m$ with $m \perp n$, $\|m\| = 1$. Differentiating,

$$\mathrm{d}q(w)[\delta] = \Big(-\tfrac{a}{2}\sin\tfrac{r}{2},\;\; \tfrac{a}{2}\cos\tfrac{r}{2}\,n + \tfrac{\sin(r/2)}{r}\,b\,m\Big), \tag{11}$$

so that

$$\big\|\mathrm{d}q(w)[\delta]\big\|^2 = \frac{a^2}{4} + b^2\frac{\sin^2(r/2)}{r^2} \;\le\; \frac{a^2+b^2}{4}, \tag{12}$$

using $\sin x \le x$; the case $w = 0$ follows by continuity. Hence the image of the segment from u to v has length at most $\|u-v\|/2$ on S^3, so $d_{S^3}(q(u),q(v)) \le \|u-v\|/2$ and Step 2 gives (9).

Step 4 (halving, boundary, coverage). Octree subdivision halves h, hence the radius. Nothing above restricts $\|v\| \le \pi$, so cubes straddling the boundary of the π-ball need no special treatment. Coverage: the cubes tile $[-\pi,\pi]^3$, every rotation has an angle-axis representative with $\|v\| \le \pi$, and a cube whose closest point to the origin is farther than π contains no such representative, so discarding it loses no rotation. □

The lemma replaces the ad hoc inflation factors of earlier discretised formulations (typically a $1.2\times$ margin on $\sqrt{3}h$, added to compensate for boundary effects that, by Step 4, need no compensation).

A.2 PROOF OF LEMMA 2 (DISPLACEMENT AND THE TWO POTENTIALS)

Idea. Move the two endpoints to their cell centres one at a time: by bi-invariance each move shifts the relative rotation by at most that cell's radius, and the measurement does not move, so the residual

changes by at most $\rho_a + \rho_b$. Hence an explained edge cannot be missed by the inflated table u, and the v-value of a labelling is the count of a genuine configuration, the cell centres themselves.

Step 1 (two moves). Fix an edge (i, j), cells $a \in \mathcal{C}_i$, $b \in \mathcal{C}_j$ with representatives R_a, R_b and radii ρ_a, ρ_b, and any $R_i \in a$, $R_j \in b$. By right-invariance, $d_{ab} = d(\tilde{R}_{ij} R_a, R_b) = d(\tilde{R}_{ij}, R_b R_a^\top)$. Two triangle inequalities and bi-invariance give

$$\left|d(\tilde{R}_{ij}, R_j R_i^\top) - d(\tilde{R}_{ij}, R_b R_a^\top)\right| \;\le\; d\big(R_j R_i^\top, R_b R_a^\top\big) \;\le\; d(R_j, R_b) + d(R_i, R_a) \;\le\; \rho_a + \rho_b, \tag{13}$$

where the middle step is the two moves: $d(R_j R_i^\top, R_b R_a^\top) \le d(R_j R_i^\top, R_b R_i^\top) + d(R_b R_i^\top, R_b R_a^\top) = d(R_j, R_b) + d(R_i, R_a)$. This is (4).

Step 2 (consequence (a): u never misses an explained edge). If $r_{ij}(\boldsymbol{R}) \le \varepsilon$ then $d_{ab} \le r_{ij}(\boldsymbol{R}) + \rho_a + \rho_b \le \varepsilon + \rho_a + \rho_b$, i.e. $u_{ij}(a, b) = 1$. For $\boldsymbol{R} \in \mathcal{U}$, coverage lets us choose for each node a cell of $\mathcal{C}_i$ containing it; summing the implication over the edges that $\boldsymbol{R}$ explains gives $C_\varepsilon(\boldsymbol{R}) \le \sum_e u_{ij}(\boldsymbol{\ell}(\boldsymbol{R})) \le U^*$.

Step 3 (consequence (b): the v-value is the count of a real configuration). Placing each node at its cell centre realises $r_{ij} = d_{ab}$ for every edge (the centre *is* a point of the cell), so the decoded configuration's consensus count equals $\sum_e v_{ij}(\boldsymbol{\ell})$. Maximising over labellings, V^* is the best count achievable by any cell-centre configuration, hence a valid and attained lower bound on $C^*_\varepsilon(\mathcal{U})$. □

A.3 Proof of Theorem 3 (integer consensus bracket)

Idea. Every inequality in (5) is a definition or Lemma 2: the lower ends are counts of actual configurations, so they cannot exceed the maximum; for the upper end, the cells a maximiser occupies earn a unit of u on every edge it explains, so the u-optimum is at least the maximum. The one delicate point, that a maximiser exists, needs the closed threshold $r \le \varepsilon$.

Step 1 (lower ends). $\hat{\boldsymbol{R}} \in \mathcal{U}$ gives $C_\varepsilon(\hat{\boldsymbol{R}}) \le C^*_\varepsilon(\mathcal{U})$ by definition, and $V^* \le C^*_\varepsilon(\mathcal{U})$ by Lemma 2(b), since the maximising labelling decodes to a configuration in $\mathcal{U}$ of count V^*.

Step 2 (the maximum is attained). Each indicator $\mathbf{1}[r_{ij} \le \varepsilon]$ can only jump *downwards* as the configuration moves continuously (the set $\{r_{ij} \le \varepsilon\}$ is closed), so C_ε is upper semicontinuous; and $\mathcal{U}$ is a finite union of products of closed cells, hence compact. An upper semicontinuous function on a compact set attains its maximum. (With a strict threshold $r < \varepsilon$ the maximum could be approached without ever being reached.) Let $\boldsymbol{R}^\circ$ attain $C^*_\varepsilon(\mathcal{U})$.

Step 3 (upper end). By Lemma 2(a) applied to $\boldsymbol{R}^\circ$, $C^*_\varepsilon(\mathcal{U}) = C_\varepsilon(\boldsymbol{R}^\circ) \le \sum_e u_{ij}(\boldsymbol{\ell}(\boldsymbol{R}^\circ)) \le U^* \le \mathrm{U}$ for any admissible U.

Step 4 (refinement bound). Let $\boldsymbol{\ell}^u$ be a u-optimal labelling and decode it to centres, giving $\boldsymbol{R}^c \in \mathcal{U}$. Every edge with $u_{ij} = 1$ has $r_{ij}(\boldsymbol{R}^c) = d_{ab} \le \varepsilon + \rho_a + \rho_b \le \varepsilon + 2\rho_{\max}$, so $U^* \le C_{\varepsilon + 2\rho_{\max}}(\boldsymbol{R}^c) \le C^*_{\varepsilon + 2\rho_{\max}}(\mathcal{U})$: the optimistic count is at most what maximum consensus itself reaches at a tolerance inflated by two radii.

Step 5 (closure). If $C_\varepsilon(\hat{\boldsymbol{R}}) = \mathrm{U}$, the chain $C_\varepsilon(\hat{\boldsymbol{R}}) \le C^*_\varepsilon(\mathcal{U}) \le \mathrm{U}$ collapses to equalities, so $\hat{\boldsymbol{R}}$ attains the maximum count over the region. □

A.4 Proof of Theorem 4 (finite-resolution exactness)

Idea. Figure 4 is the whole argument: the best achievable count is an integer staircase in the tolerance, flat around ε unless ε is one of its finitely many jumps. The upper end lives at the inflated tolerance $\varepsilon + 2\rho$, the lower end (after snapping to centres) at the deflated one $\varepsilon - 2\rho$; once both lie on the flat window, both ends are the same integer.

Step 1 (the count is a staircase in the tolerance). For each fixed τ the maximum defining $g(\tau) = C^*_\tau(\mathcal{U})$ is attained (Step 2 of the previous proof). g is non-decreasing in τ (the closed constraint set only grows), integer valued in $\{0, \dots, M\}$, and right-continuous: for $\tau_n \downarrow \tau$ pick maximisers $\boldsymbol{R}_n$; since there are finitely many edge subsets, pass to a subsequence along which the satisfied set $S = \{e : r_e(\boldsymbol{R}_n) \le \tau_n\}$ is constant and, by compactness, $\boldsymbol{R}_n \to \bar{\boldsymbol{R}}$; every $e \in S$ has $r_e(\boldsymbol{R}_n) \le \tau_n$, hence $r_e(\bar{\boldsymbol{R}}) \le \tau$ in the limit, so $\lim_n g(\tau_n) = |S| \le C_\tau(\bar{\boldsymbol{R}}) \le g(\tau)$, while monotonicity gives $g(\tau_n) \ge g(\tau)$. A non-decreasing integer step function bounded by M has at most M jump points. If

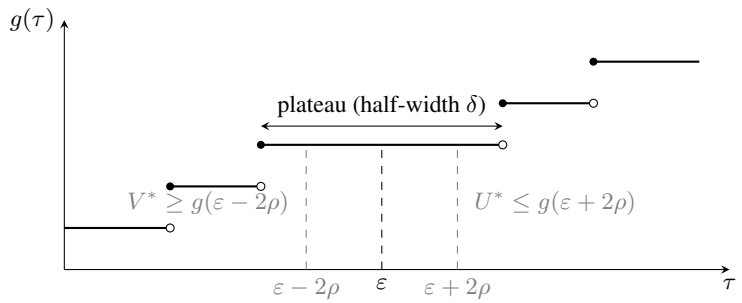

Figure 4: **The staircase behind Theorem 4.** The best achievable count $g(\tau) = C^*_\tau(\mathcal{U})$ against the tolerance is an integer step function, right-continuous (filled dots), with at most M jumps. Unless ε is a jump, g is flat on a window of half-width δ around it. The upper end of the bracket is at most g at the inflated tolerance $\varepsilon + 2\rho$; snapping to cell centres costs at most 2ρ of tolerance, so the lower end is at least g at the deflated tolerance $\varepsilon - 2\rho$. Once $2\rho \leq \delta$, both shifted tolerances lie on the plateau and the bracket closes.

ε is not one of them, g is continuous at ε, and a right-continuous step function continuous at a point is constant on a two-sided window $[\varepsilon - \delta, \varepsilon + \delta]$ for some $\delta > 0$.

Step 2 (inflating the tolerance lands on the plateau, so the upper end closes). Suppose $\rho_{\max} \leq \min\{\delta, \varepsilon\}/2$. By the refinement bound of Theorem 3, $g(\varepsilon) \leq U^* \leq g(\varepsilon + 2\rho_{\max}) = g(\varepsilon)$, the last equality because $2\rho_{\max} \leq \delta$ keeps $\varepsilon + 2\rho_{\max}$ on the plateau. So $U^* = g(\varepsilon) = C^*_\varepsilon(\mathcal{U})$.

Step 3 (deflating the tolerance and snapping to centres closes the lower end). Since $2\rho_{\max} \leq \varepsilon$, the deflated tolerance $\varepsilon - 2\rho_{\max} \geq 0$ is admissible, and since $2\rho_{\max} \leq \delta$ it lies on the plateau: $g(\varepsilon - 2\rho_{\max}) = g(\varepsilon)$. Let $\boldsymbol{R}^\circ$ attain it, label each node by a cell of $\mathcal{C}_i$ containing R°_i (coverage), and call the labelling $\boldsymbol{\ell}^\circ$. Every edge with $r_e(\boldsymbol{R}^\circ) \leq \varepsilon - 2\rho_{\max}$ has centre residual $d_{ab} \leq r_e(\boldsymbol{R}^\circ) + \rho_a + \rho_b \leq \varepsilon$ by (4), hence $v_{ij}(\boldsymbol{\ell}^\circ) = 1$: snapping costs at most $2\rho_{\max}$ of tolerance, which the deflation has already paid for. So $V^* \geq \sum_e v_{ij}(\boldsymbol{\ell}^\circ) \geq C_{\varepsilon - 2\rho_{\max}}(\boldsymbol{R}^\circ) = g(\varepsilon)$, and $V^* \leq g(\varepsilon)$ always (Lemma 2(b)). Together with Step 2, $V^* = C^*_\varepsilon(\mathcal{U}) = U^*$, which is (7).

Step 4 (the decode achieves it). A W-optimal labelling is u-optimal and, among u-optimal labellings, v-optimal (Section 3.2: $\sum_e v \leq M < M{+}1$, so the v-digit never carries). By Step 3 the labelling $\boldsymbol{\ell}^\circ$ has v-value exactly $g(\varepsilon)$ (it is at least $g(\varepsilon)$ and at most $V^* = g(\varepsilon)$), and since $v_{ij} \leq u_{ij}$ pointwise, its u-value is at least $g(\varepsilon) = U^*$: $\boldsymbol{\ell}^\circ$ is u-optimal. Any W-optimal labelling therefore has u-value U^* and, by the tie-break, v-value at least that of $\boldsymbol{\ell}^\circ$, i.e. at least $g(\varepsilon)$; as no v-value exceeds $V^* = g(\varepsilon)$, it equals $g(\varepsilon)$, and its decode achieves that count by Lemma 2(b).

Step 5 (how many levels, and monotonicity under full subdivision). Lemma 1 halves $\rho_{\max}$ per octree level, so $\rho_0 2^{-s} \leq \min\{\delta, \varepsilon\}/2$ holds from level $s = \lceil \log_2\big(2\rho_0 / \min\{\delta, \varepsilon\}\big) \rceil$ on. For monotonicity: an octree child's centre differs from its parent's by $(\pm h/2, \pm h/2, \pm h/2)$ in coordinates, so by (9) $d(R_{a'}, R_a) \leq \sqrt{3}\, h/2 = \rho_a/2$ for a child a' of a. If $u_{a'b'} = 1$ at the child level then, by the two moves of Lemma 2, $d_{ab} \leq d_{a'b'} + \rho_a/2 + \rho_b/2 \leq \varepsilon + (\rho_{a'} + \rho_a/2) + (\rho_{b'} + \rho_b/2) = \varepsilon + \rho_a + \rho_b$, using $\rho_{a'} = \rho_a/2$; hence $u_{ab} = 1$ at the parent level. Every child-level labelling therefore maps to a parent-level labelling with pointwise-dominating u-value, so U^* cannot increase under full subdivision. (Under heuristic region restriction the level regions are not nested and no such claim is made; the solver reports the last level's bound.)

Step 6 (at a jump point, and why the search is exponential). If ε is a breakpoint with jump $J = g(\varepsilon) - g(\varepsilon^-)$, right-continuity still gives a $\delta' > 0$ with g constant on $[\varepsilon, \varepsilon + \delta']$, so Step 2 applies one-sidedly: $U^* = g(\varepsilon)$ once $\rho_{\max} \leq \delta'/2$: the upper end always closes at a finite resolution. The deflation of Step 3, however, now steps down off the plateau: once $\varepsilon - 2\rho_{\max}$ lies on the constant stretch to the left of the jump, it only guarantees $V^* \geq g(\varepsilon - 2\rho_{\max}) = g(\varepsilon^-) = g(\varepsilon) - J$, so the guaranteed lower end may lag by up to J however fine the resolution, and the solver reports the residual gap. Finally, a polynomial-time algorithm producing closed brackets on arbitrary instances would solve (2), which is impossible unless P=NP (Chin et al., 2018); accordingly the branch-and-bound search is exponential in the worst case. □

Where the hypotheses enter. "ε is not a breakpoint" supplies the two-sided plateau (Steps 2–3); $2\rho_{\max} \le \delta$ keeps both shifted tolerances on it; $2\rho_{\max} \le \varepsilon$ only keeps the deflated tolerance non-negative. Nested subdivision is used solely for the monotonicity of Step 5.

A.5 Proof of Theorem 5 (admissible node bound)

Idea. Charge every edge to its endpoints: fixed–fixed edges at their known value, fixed–free edges to the free endpoint at its best label, free–free edges half to each endpoint at the best label of the far end. Each maximum dominates the true value, so the charges cover the edge.

Proof. Partition $\mathcal{E} = \mathcal{E}_{XX} \cup \mathcal{E}_{XF} \cup \mathcal{E}_{FF}$ by whether zero, one, or both endpoints are free, and decompose $W(\boldsymbol{\ell}) = W_{XX} + W_{XF}(\boldsymbol{\ell}) + W_{FF}(\boldsymbol{\ell})$ accordingly. The first sum of (8) equals W_{XX} exactly. For each free node i, replacing the maximum over $a \in \mathcal{C}_i$ by the evaluation at the actual label ℓ_i can only decrease the bracket, so

$$\sum_{i\in F} \max_a[\cdots] \;\ge\; \sum_{i\in F}\Big[\sum_{j\in N(i)\cap X} W_{ij}(\ell_i,\ell_j) + \tfrac{1}{2}\sum_{j\in N(i)\cap F} \max_{b\in\mathcal{C}_j} W_{ij}(\ell_i,b)\Big]. \tag{14}$$

The first inner sum, over $i \in F$, counts each mixed edge once (at its unique free endpoint), giving $W_{XF}(\boldsymbol{\ell})$. In the second inner sum, $\max_{b\in\mathcal{C}_j} W_{ij}(\ell_i,b) \ge W_{ij}(\ell_i,\ell_j)$ since $\ell_j \in \mathcal{C}_j$; each free–free edge appears once from each endpoint, so the halved sum is at least $W_{FF}(\boldsymbol{\ell})$. Adding the three contributions proves $B_{\text{node}} \ge W(\boldsymbol{\ell})$; no step used a property of W or of the label sets. □

A.6 Proof of Proposition 6 (floored certificates)

Idea. W is a two-digit number in base $M{+}1$: the high digit is the u-count, the low digit the v-count. Because the low digit is at most M, it never carries, so dividing by $M{+}1$ and rounding down reads off the high digit exactly, both from the optimum and from any number above it.

Proof. For any labelling, $W(\boldsymbol{\ell}) = (M{+}1)\sum_e u_{ij} + \sum_e v_{ij}$ with $0 \le \sum_e v_{ij} \le M$, so $\sum_e u_{ij} \le W(\boldsymbol{\ell})/(M{+}1) < \sum_e u_{ij} + 1$, i.e. $\lfloor W(\boldsymbol{\ell})/(M{+}1)\rfloor = \sum_e u_{ij}$. Taking the maximum, $\lfloor \max_{\boldsymbol{\ell}} W/(M{+}1)\rfloor = U^*$: the lexicographic maximum sits in $[(M{+}1)U^*,\ (M{+}1)U^* + M]$. Any admissible $B \ge \max_{\boldsymbol{\ell}} W$ then satisfies $\lfloor B/(M{+}1)\rfloor \ge U^* \ge C^*_\varepsilon(\mathcal{U})$ by monotonicity of the floor and Theorem 3. (In floating point we floor $B + 10^{-6}$; the exact values are half-integers, because of the $\frac{1}{2}$ on free–free edges, spaced $\ge \frac{1}{2}$ apart, so this can loosen the certificate by at most one count and never tightens it.) □

A.7 Proof of Proposition 7 (safe integer pruning) and the $O(1)$ per-label bound

Idea. If the most optimistic completion of "node i sits in cell a" cannot beat the incumbent, no configuration through that cell can either, so deleting it loses at most a tie; hence the guarantee concerns the value, not every maximiser. The per-label bound is the root bound with one node's contribution swapped from its best label to a: one subtraction and one addition per cell.

Step 1 (pruning is safe). Let $\boldsymbol{R}$ be any configuration in the region with R_i in cell a, and let $\boldsymbol{\ell}(\boldsymbol{R})$ be covering cells with $\ell_i = a$. By Lemma 2(a) and Proposition 6 applied to the region restricted to $\ell_i = a$, $C_\varepsilon(\boldsymbol{R}) \le \lfloor B_i(a)/(M{+}1)\rfloor \le \mathrm{C}_{\text{inc}}$. Hence every configuration using cell a at node i explains at most C_{inc} edges, and removing a cannot remove any configuration explaining more; in particular, if C^*_ε over the region exceeds C_{inc}, it is attained in the pruned region, and since the incumbent's own cell is never pruned (its bound floors to at least C_{inc}, which the rule alone might prune, so Algorithm 1 keeps it explicitly), the pruned region always contains a configuration of count C_{inc}, so all subsequent brackets remain valid. Configurations that *tie* the incumbent may be discarded, which is why the guarantee concerns the certified value and not the set of argmaxes.

Step 2 (the per-label bound). Instantiate Theorem 5 at the root (only the gauge node fixed) with $\mathcal{C}_i$ replaced by $\{a\}$, but keep the *neighbours'* half-max terms ranging over the full $\mathcal{C}_i$; enlarging the sets inside maxima can only increase the bound, so the result is still admissible for the restricted problem. Writing $T_i(a)$ for node i's bracket evaluated at label a and $B_{\text{root}} = W_{XX} + \sum_j \max_b T_j(b)$ for the root bound, the restricted bound is

$$B_i(a) \;=\; B_{\text{root}} - \max_b T_i(b) + T_i(a), \tag{15}$$

computable for all (i, a) in one pass over the tables after precomputing row and column maxima. □

A.8 Proof of Theorem 8 (exact recovery)

Idea. Two margins. The inner margin gives every true inlier 2ρ of room below the tolerance, so snapping the truth to cell centres loses no inlier. The outer margin keeps every outlier outside everything the grid could predict for its edge, inflated by the same 2ρ, so, by the displacement bound read backwards, nothing in the region can explain an outlier. The truth then explains all C^* inliers, nothing explains anything else, and the lexicographic tie-break forces every optimal labelling to decode to exactly the true inlier set. The only randomness is whether the outer margin holds, and that is a union bound.

The probability bound. Let $\mathcal{E}_{\text{good}}$ be the event that no outlier measurement lies in its candidate-target neighbourhood $\mathcal{N}_{ij}$. By hypothesis (iii) and a union bound over the M_{out} outlier edges, $\Pr[\mathcal{E}_{\text{good}}] \geq 1 - M_{\text{out}}\,\eta$; no independence between outliers is used. The neighbourhood is the union over the active pairs $\mathcal{C}_i \times \mathcal{C}_j$, exactly the pairs Step 1 ranges over, so nothing larger is needed; a union over the full grid would be valid but vacuous, since its balls cover $\mathrm{SO}(3)$. For Haar-uniform outliers, integrating the angle density $(1 - \cos t)/\pi$ over $[0, \tau]$ gives the per-ball mass $p_{\text{vol}}(\tau) = (\tau - \sin\tau)/\pi$ (Chirikjian & Kyatkin, 2001). Fix reference rotations $\hat{R}_i$ (the incumbent) and let $s_i = \max_{a \in \mathcal{C}_i} d(R_a, \hat{R}_i)$ be the spread of node i's active cells. For $a \in \mathcal{C}_i$, $b \in \mathcal{C}_j$, bi-invariance gives $d(R_b R_a^\top, \hat{R}_j \hat{R}_i^\top) \leq d(R_b, \hat{R}_j) + d(R_a, \hat{R}_i) \leq s_i + s_j$, so every ball of $\mathcal{N}_{ij}$, of radius $\varepsilon + \rho_a + \rho_b \leq \varepsilon + 2\rho$, lies in the ball of radius $\tau = \varepsilon + 2\rho + s_i + s_j$ about $\hat{R}_j \hat{R}_i^\top$, whence $\eta \leq p_{\text{vol}}(\tau)$. All steps below hold deterministically on $\mathcal{E}_{\text{good}}$. The *realised inlier set* of a labelling $\boldsymbol{\ell}$ is $\{e : r_e(\boldsymbol{R}^c(\boldsymbol{\ell})) \leq \varepsilon\}$, where $\boldsymbol{R}^c(\boldsymbol{\ell})$ places each node at its cell representative.

Step 1 (nothing can explain an outlier). Let $\boldsymbol{R} \in \mathcal{U}$ satisfy an outlier edge $e = (i, j)$ at ε, and let (a, b) be cells of $\mathcal{C}_i \times \mathcal{C}_j$ containing (R_i, R_j). By the displacement bound (4), $d(\tilde{R}_{ij}, R_b R_a^\top) = d_{ab} \leq r_e(\boldsymbol{R}) + \rho_a + \rho_b \leq \varepsilon + \rho_a + \rho_b$, so $\tilde{R}_{ij} \in \mathcal{N}_{ij}$, which $\mathcal{E}_{\text{good}}$ excludes. Hence no configuration in $\mathcal{U}$ satisfies any outlier at ε, and likewise no labelling has $u_{ij} = 1$ (a fortiori $v_{ij} = 1$) on any outlier edge, since $u_{ij} = 1$ means exactly $d_{ab} \leq \varepsilon + \rho_a + \rho_b$.

Step 2 (the truth is the unique maximiser). By hypothesis (ii), $\boldsymbol{R}^* \in \mathcal{U}$, and by the inner margin every inlier edge has $r_e(\boldsymbol{R}^*) < \varepsilon$, so $C_\varepsilon(\boldsymbol{R}^*) = C^*$. By Step 1, any $\boldsymbol{R} \in \mathcal{U}$ satisfies only inlier edges at ε, so $C_\varepsilon(\boldsymbol{R}) \leq C^*$: $C^*_\varepsilon(\mathcal{U}) = C^*$, and any maximiser's ε-inlier set is a subset of $\mathcal{I}^*$ of size C^*, i.e. equals $\mathcal{I}^*$: the maximum-consensus set over $\mathcal{U}$ is unique and true.

Step 3 (every W-optimal labelling decodes to the truth). Let $\hat{\boldsymbol{\ell}}$ label each node with a cell containing R_i^* (hypothesis (ii)). By the inner margin and the displacement bound, every inlier edge has $d_{ab} \leq r_e(\boldsymbol{R}^*) + \rho_a + \rho_b < \varepsilon$, so $v_{ij} = u_{ij} = 1$: $\sum_e u_{ij}(\hat{\boldsymbol{\ell}}) = \sum_e v_{ij}(\hat{\boldsymbol{\ell}}) = C^*$ (outlier terms vanish by Step 1). By Step 1, *every* labelling has $\sum_e u_{ij} \leq C^*$ (only inlier edges can contribute), so $\hat{\boldsymbol{\ell}}$ is u-optimal, and any W-optimal labelling $\boldsymbol{\ell}^*$ has $\sum_e u_{ij}(\boldsymbol{\ell}^*) = C^*$ and, by the lexicographic tie-break, $\sum_e v_{ij}(\boldsymbol{\ell}^*) \geq \sum_e v_{ij}(\hat{\boldsymbol{\ell}}) = C^*$; combined with $\sum_e v_{ij} \leq \sum_e u_{ij} = C^*$, equality holds. A v-value of C^* carried entirely by inlier edges means the decode of $\boldsymbol{\ell}^*$ satisfies every edge of $\mathcal{I}^*$ at ε; by Step 1 it satisfies no outlier. Its realised set is therefore exactly $\mathcal{I}^*$. Upon closure, the certified consensus set is the ε-inlier set of a configuration in $\mathcal{U}$ achieving $C^*_\varepsilon(\mathcal{U})$, which by Step 2 equals $\mathcal{I}^*$. □

Where the hypotheses enter. The inner margin (i) lets the snapped truth keep every inlier in Step 3; coverage (ii) puts the truth in the region (Step 2) and provides the cells to snap to (Step 3); the outer margin (iii) is Step 1 in its entirety, and the only hypothesis with a probability attached.

Remark 9 (Retention safety). Stage B's reward-based retention is heuristic, and the solver reports it as such. The conditional reward of a root cell c at node i is $r_i(c) = \sum_j w(d(\tilde{R}_{ij} R_c, \hat{R}_j))$, the sum over the neighbours j of i (and symmetrically for edges into i), with $\hat{R}_j$ the incumbent, $w(d) = \exp(-d^2/2s^2)$, and $s = \max\{\varepsilon, \rho_0\}$; a cell is kept if $r_i(c) \geq \rho_{\text{ret}} \max_c r_i(c)$. Were the incumbent fixed independently of node i's edges, a Hoeffding argument (Hoeffding, 1963) would bound the probability that the cell containing R_i^* falls below the threshold when a constant fraction of the neighbours are correctly estimated; the Stage A incumbent depends on every edge, so the

Algorithm 1 Certified maximum-consensus rotation averaging.

Input: graph $\mathcal{G}$, measurements $\{\tilde{R}_{ij}\}$, tolerance ε
$\hat{\boldsymbol{R}} \leftarrow$ TRW-S on the graded grid MRF, then IRLS on retained edges (Stage A)
$\mathrm{C}_{\mathrm{inc}} \leftarrow C_\varepsilon(\hat{\boldsymbol{R}})$ (achieved count; valid global lower bound)
$\mathcal{C}_i \leftarrow$ root cells (all, or retained by reward; always the cell containing $\hat{R}_i$) (Stage B)
for level $s = 0, 1, 2, \ldots$ **do**
  build lexicographic tables $W_{ij} = (M{+}1)u_{ij} + v_{ij}$ over $\mathcal{C}_i \times \mathcal{C}_j$ (Eq. 6)
  $(\boldsymbol{\ell}, B) \leftarrow$ BranchAndBound(W) with admissible node bound (Thm. 5)
  $\mathrm{U} \leftarrow \lfloor B/(M{+}1) \rfloor$ (integer cap on $C^*_\varepsilon(\mathcal{U})$; Prop. 6)
  $\boldsymbol{R}_{\mathrm{dec}} \leftarrow$ cell centres of $\boldsymbol{\ell}$; $\boldsymbol{R}_{\mathrm{pol}} \leftarrow$ IRLS polish of $\boldsymbol{R}_{\mathrm{dec}}$ on its satisfiable edges
  $\mathrm{C}_{\mathrm{inc}} \leftarrow \max\{\mathrm{C}_{\mathrm{inc}}, C_\varepsilon(\boldsymbol{R}_{\mathrm{dec}}), C_\varepsilon(\boldsymbol{R}_{\mathrm{pol}})\}$; update $\hat{\boldsymbol{R}}$ accordingly
  if $\mathrm{U} = \mathrm{C}_{\mathrm{inc}}$ **then break** (exact certificate)
  $B_i(c) \leftarrow$ per-label root bounds (15) for all i, $c \in \mathcal{C}_i$ (Eq. 15)
  $\mathcal{C}_i \leftarrow$ subdivide $\{c : \lfloor B_i(c)/(M{+}1) \rfloor \geq \mathrm{C}_{\mathrm{inc}} + 1\}$, keeping the cell containing $\hat{R}_i$ (Prop. 7)
end for
Return: $\hat{\boldsymbol{R}}$, its consensus set, certified count $\mathrm{C}_{\mathrm{inc}} = \mathrm{U}$, scope flag (global / region)

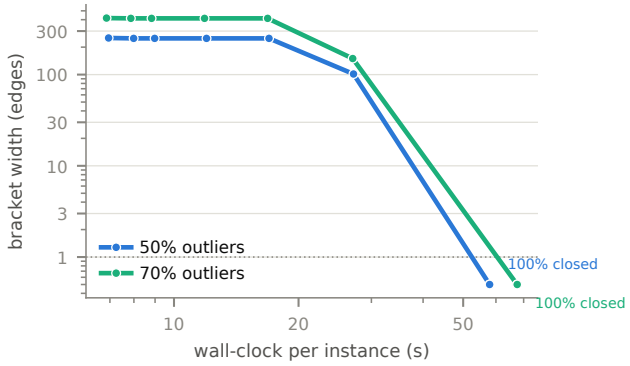

Figure 5: **The bracket closes with compute.** Median integer bracket width $\mathrm{U} - \mathrm{C}_{\mathrm{inc}}$ over 20 seeds when a wall-clock limit stops the solver, plotted against the time actually used, on the sparse $N{=}100$ graphs of Figure 2 at 50% and 70% outliers. Each drop is one subdivision level; runs end at closure, so the two largest limits coincide in measured time and every instance closes.

terms are not independent, and the remark is motivation only. The cell containing the incumbent itself is always retained by construction. We do not rely on this remark for any reported certificate: region-conditional certificates are labelled as such throughout.

B The certified solver in practice

Algorithm 1 assembles the stages of Section 3.4.

Stage A in detail. The initialiser is a pairwise MRF on the 30° grid of SO(3) (voxel radius $r \approx 31°$): node i takes a voxel a with representative R_a, and edge (i, j) scores the pair (a, b) by the graded potential $\phi_{ij}(a, b) = \exp(-d^2/2\sigma^2)\,\mathbf{1}[d \leq \tau_s]$ with $d = d(\tilde{R}_{ij}R_a, R_b)$, $\tau_s = \varepsilon + 2r \approx 67°$, and $\sigma = \tau_s/5$. TRW-S runs 50 iterations; the decoded labelling retains the edges whose residual is at most $0.8\,\tau_s$, an L1 solve on them initialises IRLS with kernel width ε, and least-residual edges bridge disconnected inlier components. Theorem 3 holds whatever this stage returns.

B.1 Certificate against compute

We run the certified solver on the sparse $N{=}100$ graphs of Section 4.1 (the same instances, seeds, and noise as Figure 2) at 50% and 70% outliers, with a wall-clock limit per instance from zero (Stage A plus the root bound alone) to 120 s, and record what the solver holds when the limit stops it (Figure 5). Two quantities behave in opposite ways.

The estimate is immediate. Stage A delivers the incumbent in about 7 s, and no later stage changes its median accuracy: the median relative-rotation error is $0.55°$ at 50% and $0.73°$ at 70% outliers at every limit, against $0.81°$ and $3.14°$ for the best baseline (GNC-GM). This is the decoupling principle at work: the estimate follows from the inlier set, which is already in hand.

The certificate is what compute buys. The bracket starts wide: after level 0, where the inflation $\varepsilon + 2\rho$ is $57°$ and nearly every cell pair is optimistically satisfiable, it spans 252 edges at 50% and 416 at 70%. It falls with every subdivision: to 102 and 150 edges after the first, to a handful after the fourth, and to zero on every instance, within a median of 58 and 68 s and at most 99 s. Most instances close after five or six levels; a few need up to ten, and these are the instances in which some edge's residual lies within a fraction of a degree of the tolerance, so that the inflation must shrink below that margin before the satisfiability table can decide it. Theorem 4 ties this finite closing resolution to the instance's own margin.

Cost. Closing the bracket costs about a minute here; Table 5 breaks the time into stages (a separate run of the same instances). (Table 1 times the Stage A estimate alone, on the GPU; the certified loop runs Stage A on the CPU, which is the 7 s above.) Branch and bound dominates at the finest levels, where the incumbent is already optimal and the search is spent proving it. Two to three orders of magnitude for a per-instance proof is the trade, and it is paid only when the proof is wanted: every intermediate bracket is valid, and a limit stops the loop with the current best configuration in hand.

B.2 When the certified count is not the truth

A closed bracket certifies a count; Theorem 8 states when that count identifies the truth. Two kinds of instance separate them, both observable from the solver's output.

An outlier that is accidentally consistent. On 3 of the 20 instances at 50% and 2 at 70%, the closed bracket certifies a configuration that explains *one edge more* than the true inlier set at $\varepsilon = 5°$ and lies 2.2–2.9° from it. Maximum consensus itself prefers this answer: an outlier fell inside the inflated candidate set of Theorem 8, hypothesis (iii) failed, and the certificate is correct about a count that does not identify the truth. With ground truth in hand, the tell is a certified count above the truth's; without it, the case is indistinguishable from a correct certificate, as it should be.

A region that excludes the truth. On 9 of the 20 instances at 50% and 12 at 70%, the certified count over the retained region is *below* the truth's own count, so the truth's cells were not all retained (hypothesis (ii) failed) and the closed bracket speaks about the region only. On most of these the region's optimum is still within $0.1°$ of the truth's accuracy; on one of them at 50% and five at 70%, it is the best configuration of a nearby basin, 2–3° away. The solver's region flag marks all of them, and the retention escalation of Section 3.4 is the remedy. Medians are untouched in both cases; the upper quartile at 70% rises to $2.6°$.

The information-theoretic floor. The extreme case is a tiny graph: on full graphs with $N = 10$ at 70% outliers (45 edges, 9 unknowns), the bracket closes while the median error is $90°$. The inlier subgraph has average degree ≈ 3 and is often disconnected, so the data admit configurations explaining as many measurements as the truth does; the solver proves that nothing explains more edges than the returned configuration, and the failure is that the best explanation is not unique. At $N = 20$ the same setting recovers to $0.6°$. The framework makes this distinction observable: it certifies the count, and correctness is a separate question.

Table 5: **Per-stage runtime** (medians over seeds). Top: full graphs at 50% outliers, by N. Bottom: the sparse $N{=}100$ graphs of Figure 5, by outlier rate. "tables", "BnB", and "polish" sum over levels.

graph	Stage A	tables	BnB	polish	levels	total
full, $N{=}10$	0.2s	0.1s	2.1s	0.0s	3	2s
full, $N{=}20$	0.7s	0.1s	30.1s	0.0s	4	31s
full, $N{=}50$	4.3s	0.0s	50.3s	0.0s	6	55s
full, $N{=}100$	17.4s	0.4s	73.4s	0.3s	8	92s
sparse, $N{=}100$, 50%	7.8s	0.3s	51.0s	0.6s	6	65s
sparse, $N{=}100$, 70%	11.2s	0.3s	61.5s	0.4s	6	72s