Global Solvers for 3D Vision: Foundations, Frontiers, and a Call to the Robotics Community

Zhenjun Zhao
I3A
University of Zaragoza
Zaragoza, Spain
zhenjunz@unizar.es

Javier Civera
I3A
University of Zaragoza
Zaragoza, Spain
jcivera@unizar.es

***Abstract*—Global solvers for 3D vision, encompassing branch-and-bound, convex relaxation, and graduated non-convexity methods, provide something rare in modern robotics: certifiably optimal solutions to geometric estimation problems. Despite more than 60 years of theoretical development and demonstrated success across tasks ranging from the Wahba problem to bundle adjustment, these methods remain underutilized relative to their importance. This paper briefly presents the landscape of global solvers, identifies three fundamental open challenges, namely scalability, integration with deep learning, and standardized evaluation, and argues that the robotics optimization community is uniquely positioned to address them. As foundation models reshape 3D vision and safety-critical deployment demands grow, the need for certifiable geometric perception has never been more urgent. We invite the community to engage with this area. A continuously-updated literature summary and companion code tutorials are available at Awesome Global Solvers for 3D Vision.**

***Index Terms*—global optimization, certifiable algorithms, 3D vision, robotics, branch-and-bound, convex relaxation, graduated non-convexity**

I. Introduction

Robust and reliable perception is a cornerstone of autonomous robotics. Whether a surgical robot [1] must localize its instruments with millimeter precision or an autonomous vehicle [2] must estimate its pose in a GPS-denied environment, these systems operate through physical interaction with the 3D world, where perception errors may cause material damage or harm to humans. Such safety-critical settings demand not just accurate solutions, but *certifiably correct* ones. Yet, the optimization methods that underpin these tasks, collectively known as global solvers, remain underexplored relative to their importance.

Global solvers for 3D vision encompass a rich family of algorithms that seek guaranteed optimal solutions to geometric estimation problems such as point cloud registration [3]–[5], pose estimation [6], [7], structure-from-motion (SfM) [8], and simultaneous localization and mapping (SLAM) [9]–[13]. Unlike local methods [14], [15] that converge to whichever local minimum is closer to an initial guess, global solvers provide *optimality certificates*: a proof that the solution found is the best possible given the data. This property is not a luxury. It is a necessity for safety-critical robot deployment.

Despite over 60 years of development spanning branch-and-bound (BnB) [16], convex relaxation (CR) [17], [18], and graduated non-convexity (GNC) [19] paradigms, this field remains relatively niche. The recent explosion of interest in learning-based methods, while well justified by their impressive empirical success, has nevertheless drawn attention and resources away from principled optimization. As a result, a growing gap has emerged between what is theoretically achievable and what practitioners are deploying in practice.

This position paper is both a map and an invitation. Drawing on our recent survey of 250+ papers across ten 3D vision tasks [20], we present the landscape of global solvers, identify the most pressing open challenges, and argue that now is precisely the right moment for the robotics optimization community to engage deeply with this area. The convergence of foundation models, differentiable optimization, and safety-critical deployment demands has created a unique opportunity, one that global solvers are uniquely positioned to address.

II. A Landscape of Global Solvers for 3D Vision

Global solvers for 3D vision can be organized into three principal paradigms, each with distinct theoretical foundations, computational profiles, and practical trade-offs. Fig. 1 provides a structured overview of the field.

Branch-and-Bound (BnB) methods systematically partition the solution space and compute bounds on the optimal objective within each region, pruning regions that cannot contain the global optimum. They provide exact global optimality guarantees but suffer from worst-case exponential complexity. Recent advances in problem-specific bound tightening and adaptive branching have made BnB tractable for problems of moderate scale, including absolute pose estimation [21] and outlier-robust registration [22].

Convex Relaxation (CR) methods reformulate non-convex problems, typically involving rotation constraints or combinatorial outlier selection, as tractable convex programs, most commonly semidefinite programs (SDPs). When the relaxation is *tight*, the solution of the convex program recovers the global optimum of the original problem, along with a certificate of optimality. Landmark results have demonstrated tight relaxations for essential matrix estimation [23], point cloud registration [24], rotation averaging [25], and pose graph optimization [26], often achieving polynomial-time global optimality.

Graduated Non-Convexity (GNC) methods solve a sequence of progressively non-convex surrogate problems, starting from a convex approximation and gradually sharpening toward the original objective. GNC achieves dramatically better scalability than BnB or CR but sacrifices formal global guarantees, instead offering empirical robustness under high outlier ratios.

These three paradigms are complementary rather than competing. BnB provides exact guarantees at small scale; CR provides polynomial-time certificates for structured problems; GNC provides scalable approximate solutions that can be certified post-hoc. Together, they form a unified toolkit for certifiable geometric perception across the full spectrum of problem scales and structures, covering ten fundamental 3D vision tasks as shown in Fig. 1.

III. Open Challenges

Despite substantial progress, three fundamental challenges limit the practical adoption of global solvers in real-world robotics systems.

A. Scalability

The computational cost of BnB and CR methods grows prohibitively with problem size, making real-time or large-scale deployment challenging. For BnB, tighter problem-specific bounds and learned branching strategies offer promising acceleration paths. For CR, recent advances, including block-coordinate descent [27], the Riemannian staircase via low-rank Burer-Monteiro factorization [28]–[30], chordal relaxation and sparsity exploitation [31], GPU-parallel solvers [32], and incremental formulations [33], [34], have reduced costs significantly, but these gains do not yet extend cleanly to extremely large-scale settings where problem size grows with the number of observations. GNC scales well but lacks guarantees.

A key open problem is designing *hybrid pipelines*: fast GNC-style solvers paired with lightweight certifiers that balance speed and certifiability at deployment scale. Such pipelines could enable certifiable perception for online applications including autonomous driving and aerial navigation, where both real-time performance and provable correctness are required.

B. Integration with Deep Learning

Foundation models such as DUSt3R [35] and VGGT [36] can produce high-quality geometric predictions from images without explicit optimization, representing a paradigm shift in 3D vision. However, these learned predictions may violate multi-view geometric constraints and carry no optimality certificates. This is precisely where global solvers become indispensable: not as replacements for learning, but as *verifiers and refiners* of learned outputs.

This integration is bidirectional. Convex relaxations can serve as differentiable layers within learning pipelines [37], certifying global optimality while enabling end-to-end training. Conversely, learned priors can warm-start or prune the search space of global solvers, dramatically accelerating convergence [38]. Both directions remain largely underexplored and represent one of the most exciting frontiers in certifiable geometric perception.

C. Standardized Evaluation

The field currently lacks shared benchmarks and evaluation protocols enabling fair comparison across methods, problem scales, and noise regimes. Most evaluations rely on small synthetic experiments, making it difficult to assess real-world readiness or track community-wide progress.

Beyond solution accuracy, certifiable methods require new metrics: the fraction of instances solved to certified global optimality, relaxation tightness gaps, runtime-accuracy tradeoff curves across problem scales, and robustness under adversarial outlier configurations. Establishing such infrastructure, common datasets, reference implementations, and evaluation scripts, would substantially lower the barrier to entry and accelerate community-wide adoption.

IV. Why Now?

Three converging trends make this a particularly opportune moment for global solvers in robotics.

Safety-critical robotics is scaling up. Autonomous driving, surgical robotics, and aerial navigation are moving from research labs to real-world deployment, where perception failures carry serious consequences. The demand for provably correct estimation, not just empirically accurate estimation, has never been higher. Certifiability is transitioning from a theoretical nicety to an operational requirement.

Foundation models have exposed a certification gap. The success of large-scale geometric predictors like DUSt3R and VGGT demonstrates that data-driven methods can produce impressive geometric outputs. But impressive is not certifiable. As these models are integrated into safety-critical pipelines, the need to verify and certify their outputs will become acute. Global solvers are the natural tool for this role, providing the formal guarantees that learned models cannot.

The optimization community has the tools. Recent years have seen remarkable advances in SDP solvers, differentiable optimization layers, and scalable relaxation techniques. The theoretical machinery is mature. What is needed is the community's sustained attention to connecting it with practical robotics problems at realistic scales.

The robotics optimization community gathered at this workshop is uniquely positioned to drive this agenda forward. The problems are hard, the stakes are high, and the tools are ready.

V. Conclusion

Global solvers for 3D vision offer something rare in modern robotics: the ability to *prove* that a solution is optimal. After 60 years of development across branch-and-bound, convex relaxation, and graduated non-convexity paradigms, the field stands at an inflection point. Scalability barriers are falling,

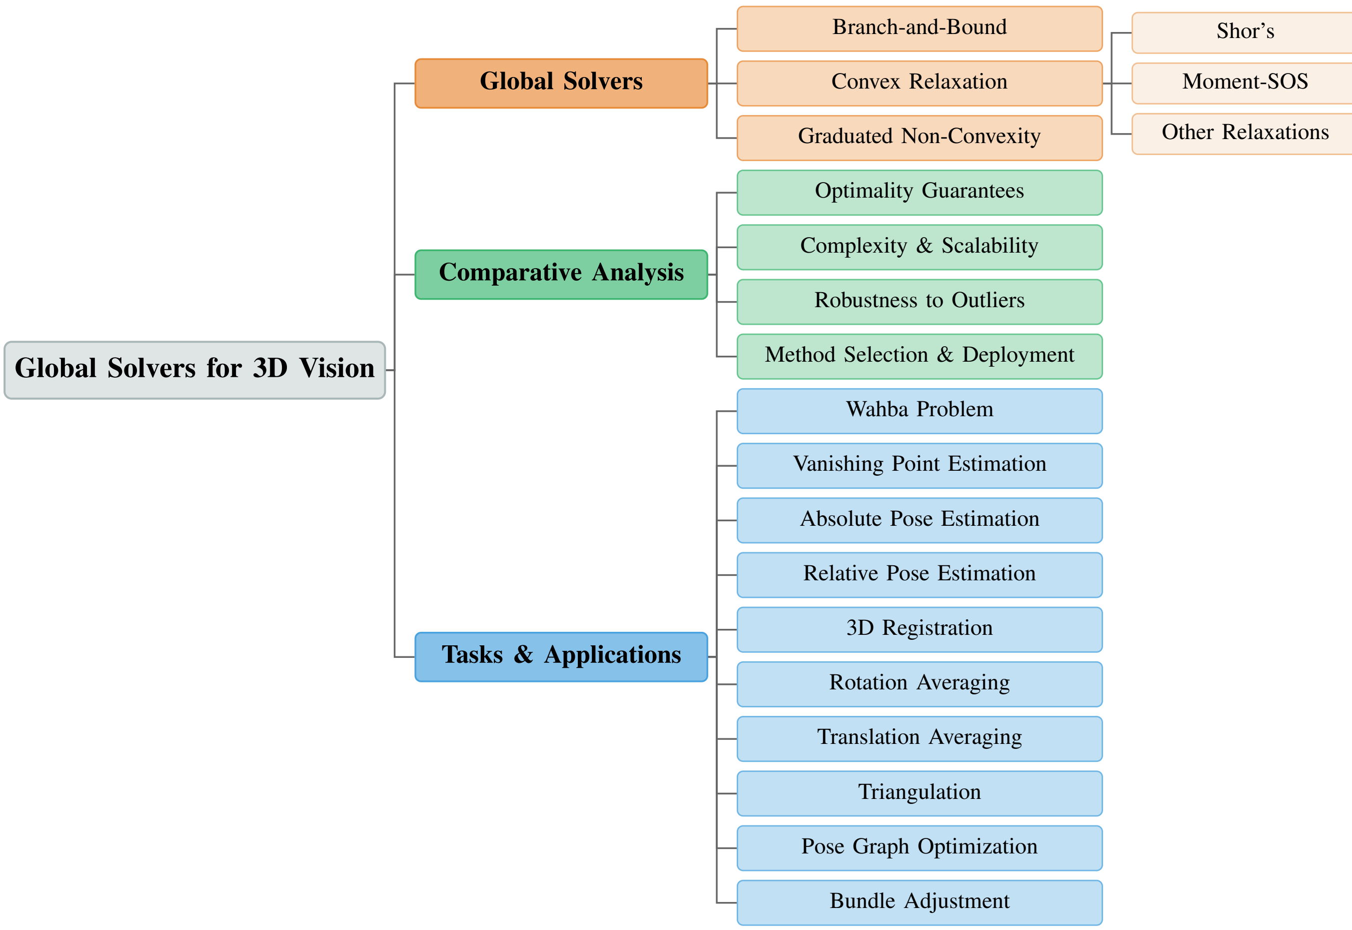

Fig. 1. Taxonomy of global solvers for 3D vision [20]. The field is organized around three components: **Global Solvers**, encompassing Branch-and-Bound, Convex Relaxation (Shor's relaxation, Moment-SOS relaxation, and other relaxations), and Graduated Non-Convexity; **Comparative Analysis**, examining optimality guarantees, complexity and scalability, robustness to outliers, and deployment considerations; and **Tasks & Applications**, covering ten fundamental problems in 3D vision.

integration with deep learning is becoming tractable, and safety-critical deployment is creating genuine demand for certification.

We invite the robotics optimization community to engage with this area and bring expertise in scalable algorithms, differentiable optimization, and real-world deployment to bear on the open challenges outlined here. The next generation of certifiable perception systems for autonomous driving, surgical robotics, and aerial navigation will require exactly this combination of skills. The opportunity is now.

Acknowledgment

The authors thank the broader global solvers community whose work is surveyed in [20].

References

[1] I. O. Sunmola, Z. Zhao, S. Schmidgall, Y. Wang, P. M. Scheikl, V. Pham, and A. Krieger, "Surgical gaussian surfels: Highly accurate real-time surgical scene rendering using gaussian surfels," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2026, pp. 4515–4524.

[2] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2012, pp. 3354–3361.

[3] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-d point sets," *IEEE Trans. Pattern Anal. Mach. Intell.*, no. 5, pp. 698–700, 1987.

[4] S. Yan, P. Shi, Z. Zhao, K. Wang, K. Cao, J. Wu, and J. Li, "Turboreg: Turboclique for robust and efficient point cloud registration," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2025, pp. 26 371–26 381.

[5] S. Yan, Y. Wang, K. Zhao, P. Shi, Z. Zhao, Y. Zhang, and J. Li, "Hemora: Unsupervised heuristic consensus sampling for robust point cloud registration," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2025, pp. 1363–1373.

[6] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Ep n p: An accurate o (n) solution to the p n p problem," *Int. J. Comput. Vis.*, vol. 81, no. 2, pp. 155–166, 2009.

[7] R. I. Hartley, "In defense of the eight-point algorithm," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 6, pp. 580–593, 1997.

[8] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 4104–4113.

[9] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam," *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, 2021.

[10] X. Meng, P. Hou, Z. Zhao, J. Civera, D. Cremers, H. Wang, and H. Li, "Dream-slam: Dreaming the unseen for active slam in dynamic environments," *arXiv preprint arXiv:2602.21967*, 2026.

[11] M. Li, D. Li, S. Hu, K. Wang, Z. Zhao, and H. Wang, "Slam-x: Generalizable dynamic removal for nerf and gaussian splatting slam," in *Proc. ACM Int. Conf. Multimedia*, 2025, pp. 1132–1140.

[12] H. Li, J. Zhao, J.-C. Bazin, P. Kim, K. Joo, Z. Zhao, and Y.-H. Liu, "Hong kong world: Leveraging structural regularity for line-based slam," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 11, pp. 13 035–13 053, 2023.

[13] Y. Gu, S. Yan, Z. Zhao, Y. Kou, J. Luo, P. Shi, and J. Li, "Ulf-loc: Unbiased landmark feature for robust visual localization with 3d gaussian splatting," *arXiv preprint arXiv:2605.04730*, 2026.

[14] Å. Björck, *Numerical methods for least squares problems*. SIAM, 2024.

[15] J. J. Moré, "The levenberg-marquardt algorithm: implementation and theory," in *Numerical analysis: proceedings of the biennial Conference held at Dundee, June 28–July 1, 1977*, 2006, pp. 105–116.

[16] E. L. Lawler and D. E. Wood, "Branch-and-bound methods: A survey," *Operations research*, vol. 14, no. 4, pp. 699–719, 1966.

[17] N. Z. Shor, "Quadratic optimization problems," *Soviet Journal of Computer and Systems Sciences*, vol. 25, no. 6, pp. 1–11, 1987.

[18] J. B. Lasserre, "Global optimization with polynomials and the problem of moments," *SIAM J. Optim.*, vol. 11, no. 3, pp. 796–817, 2001.

[19] M. J. Black and A. Rangarajan, "On the unification of line processes, outlier rejection, and robust statistics with applications in early vision," *Int. J. Comput. Vis.*, vol. 19, no. 1, pp. 57–91, 1996.

[20] Z. Zhao, H. Yang, B. Liao, Y. Zeng, S. Yan, Y. Gu, P. Liu, Y. Zhou, H. Li, and J. Civera, "Advances in global solvers for 3d vision," *arXiv preprint arXiv:2602.14662*, 2026.

[21] C. Long, Q. Hu, C. Jiang, D. Li, and Z. Ouyang, "Bnb-based robust pnp pose estimation method for outliers," *IEEE Robot. Autom. Lett.*, 2025.

[22] M. Xu, J. kai Wang, K. Wang, and Z. Chen, "Hierarchical and validated branch-and-bound method for global point cloud registration," *IEEE Trans. Ind. Inform.*, vol. 21, pp. 940–949, 2025.

[23] J. Zhao, "An efficient solution to non-minimal case essential matrix estimation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 4, pp. 1777–1792, 2020.

[24] H. Yang, J. Shi, and L. Carlone, "Teaser: Fast and certifiable point cloud registration," *IEEE Trans. Robot.*, vol. 37, no. 2, pp. 314–333, 2020.

[25] F. Dellaert, D. M. Rosen, J. Wu, R. Mahony, and L. Carlone, "Shonan rotation averaging: Global optimality by surfing so (p) n," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 292–308.

[26] 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.*, vol. 38, no. 2-3, pp. 95–125, 2019.

[27] Y. Tian, K. Khosoussi, and J. P. How, "Block-coordinate minimization for large sdps with block-diagonal constraints," *arXiv preprint arXiv:1903.00597*, 2019.

[28] N. Boumal, "A riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints," *arXiv preprint arXiv:1506.00575*, 2015.

[29] S. Burer and R. D. Monteiro, "A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization," *Math. Program.*, vol. 95, no. 2, pp. 329–357, 2003.

[30] N. Boumal, V. Voroninski, and A. Bandeira, "The non-convex burer-monteiro approach works on smooth semidefinite programs," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016.

[31] F. Dümbgen, C. Holmes, and T. D. Barfoot, "Exploiting chordal sparsity for fast global optimality with application to localization," *arXiv preprint arXiv:2406.02365*, 2024.

[32] H. Han and H. Yang, "Building rome with convex optimization," in *Proc. Robot.: Sci. Syst.*, 2025.

[33] B. Liao, Z. Zhao, L. Chen, H. Li, D. Cremers, and P. Liu, "Globalpointer: Large-scale plane adjustment with bi-convex relaxation," in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 360–376.

[34] B. Liao, Z. Zhao, H. Li, Y. Zhou, Y. Zeng, H. Li, and P. Liu, "Convex relaxation for robust vanishing point estimation in manhattan world," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2025, pp. 15 823–15 832.

[35] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud, "Dust3r: Geometric 3d vision made easy," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2024, pp. 20 697–20 709.

[36] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, "Vggt: Visual geometry grounded transformer," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2025, pp. 5294–5306.

[37] C. Holmes, F. Dümbgen, and T. D. Barfoot, "Sdprlayers: Certifiable backpropagation through polynomial optimization problems in robotics," *IEEE Trans. Robot.*, 2025.

[38] J. Mai, B. Liao, Z. Zhao, Y. Zeng, H. Li, J. Civera, T. Wu, Y. Zhou, and P. Liu, "Neural predictor-corrector: Solving homotopy problems with reinforcement learning," *arXiv preprint arXiv:2602.03086*, 2026.