Enabling Reliable Estimation in Robotics Via Efficient Convex Relaxations

I. Motivation and Background

The near future promises an accelerated level of human-robot interaction, particularly in the field of mobile robots. In emergent collaborative environments (streets, factories, hospitals, etc.), it will be critical to develop autonomy that is not only efficient, but also *inherently safe and reliable*. Since optimization problems are ubiquitous at many layers of the robotics software stack, reliability means ensuring that our algorithms always provide *truly optimal* solutions. State-of-the-art robotics solutions often rely on local optimization techniques (e.g Gauss-Newton, etc.), which, although highly efficient, can converge to spurious local optima if not initialized well. Failures of this nature can cause troublesome bugs and inefficiencies in general, but can also lead to life-threatening situations once deployed.

Particularly in perception and state estimation, concerns with spurious local minima have led to a recent surge of interest in so-called *certifiable methods*, which *guarantee global optimality* by applying convex relaxations[1] to optimization problems. When successfully applied, certifiable methods either directly find *initialization-free, globally optimal solutions* to non-convex problems or *certify global optimality* of a solution found via local optimization. Since these approaches use convex optimization, they avoid the exponential complexity of other global methods like branch-and-bound.[2] While these methods have been applied to *a broad variety* of robotics problems (e.g., sensor calibration [2], [3], outlier robust perception [4], [5], pose-graph optimization [6], range-aided SLAM [7], trajectory planning [8], etc.), they still remain relatively underutilized by roboticists in practice.

My vision is to establish certifiable methods as mainstream solutions to key problems in state estimation for mobile robotics.

I posit that **two key challenges** have impeded broader use of certifiable methods:

Identifying Tight Relaxations: A convex relaxation is referred to as *tight* (or *exact*) when its optimal cost is equal to the optimal cost of the original, non-convex optimization problem. This is a necessary condition for the practical application of certifiable methods [9], but finding a tight relaxation can be quite a laborious trial-and-error process. Oftentimes, practitioners make restrictive assumptions (e.g., isotropic noise models [6], [10], simplified geometries [3], etc.) to yield tight relaxations, resulting in limited practical application. One promising approach is to *tighten* the relaxation by applying so-called *redundant constraints*. These constraints have been applied to certify several problems in vision and state-estimation [5], [10]–[13] as well as in the controls community [14]–[16]. However, identifying the tightening constraints is a bespoke process that largely depends on practitioner expertise. Alternatively, there are some *automated* methods to find tight relaxations, such as the Moment-SOS hierarchy, but they typically lead to formulations that are intractably large for robotics applications [4].

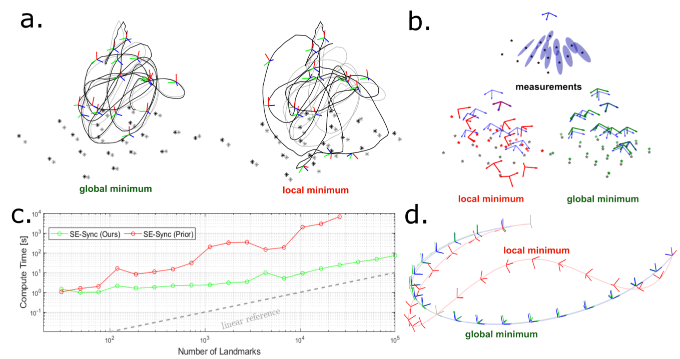

Fig. 1. Samples of my works in certifiable methods for SLAM and localization problems: **a.** Landmark-based SLAM [17]; **b.** Stereo-vision-based SLAM with anisotropic noise model [18]; **c.** Runtime complexity improvement in landmark-based SLAM [17] **d.** Continuous-time trajectory or continuum robot shape estimation [19].

Computational Efficiency: It is worth noting that even when a relaxation is tight, it may not be tractable to find or certify a solution *in real time*. This is because certifiable methods effectively lift the variables into a higher dimensional space, considerably increasing problem size. For some problems of sufficiently small dimension ($\leq$ 50 variables), the relaxation itself can often be solved directly without sacrificing realtime capability [10], [11], [20]. For large-scale problems like SLAM ($\geq$ 1000 variables), realtime performance can be obtained by efficiently *certifying global optimality* of a solution found via *fast, local optimization*. The *Riemannian Staircase* is one such approach and has been used to obtain realtime, globally optimal solutions to pose graph optimization (PGO) in different contexts [6], [7], [21]. Unfortunately, when redundant constraints are used to tighten relaxation, global optimality can only be obtained by directly solving relaxations using convex solvers. Though quite mature, these solvers are prohibitively slow for large scale problems [18].

II. Prior Work

Initial Works: My initial projects involved extending the catalogue of existing efficient certifiable solutions. Together with my collaborators, I showed that *range-only localization* (e.g., for UWB-based drone flight) could be certified efficiently

[1] A *convex relaxation* extends or modifies the feasible set of a non-convex optimization problem to yield a convex problem, which may have the same optimum.

[2] Convex problems can be solved in polynomial time [1] whereas exactly solving a non-convex problem is NP-hard.

via a *recursive LDL-decomposition scheme* [22]. In [17], I introduced landmark estimation into a highly performant method for PGO [6], yielding an efficient certifiable method for *landmark-based SLAM* (see Figure 1-a). Using an efficient marginalization of variables via a Schur-complement trick and a sparsity-exploiting Cholesky decomposition, I achieved a 10x-100x speed improvement compared to prior methods (see 1-c for runtime comparison). Crucially, in both projects, I was able to establish **linear** runtime complexity with respect to main problem variables.

Identifying Tight Relaxations: The extension these initial works required led to non-tight formulations that required redundant constraints. Realizing that a *principled approach* to finding these constraints did not exist in the literature, I developed a method to find *all possible redundant constraints* for a broad class of problems [23]. This approach allowed practitioners to *exactly* assess if and how a given relaxation could be tightened using standard linear algebra tools. My specific contributions included theoretical *completeness* proofs and the insight that a *permuted QR decomposition* could be used to find constraints that were both interpretable and scalable.

Armed with this new tool, I was able to find tight convex relaxations for SLAM and localization problems with more practical sensor models (e.g., anisotropic noise and stereo-vision models as shown in Figure 1) and establish new connections between tightness and the posterior uncertainty of pose estimates [18]. This tool also allowed my collaborators and I to apply certifiable methods to *Lie-group-based, continuous-time state estimation* (see Figure 1-d) [19]. In these problems, the use of redundant constraints was crucial to obtaining tightness, but rendered them too slow for realtime operation. This led me to explore two potential recourses: improving speed by **exploiting sparsity** and using these methods offline for **reliable learning**.

Exploiting Problem Sparsity: Exploitation of problem structure and sparsity have been crucial to the development of performant solvers for SLAM over the past decades [24]–[26]. Though less well-known in robotics, *chordal sparsity*[3] of convex relaxations facilitate a decomposition that can yield great improvements to efficiency for certifiable methods [27]. In [28], I showed that the batch-based localization problem I explored in [18] exhibit a *chain-like*, chordal structure. Despite the use of tightening redundant constraints, I showed that by taking advantage of *chordal decomposition techniques*, the runtime complexity in this problem could be kept **linear** in the number of landmarks, making it competitive with state-of-the-art local solvers.

Reliable Optimization in Learning: When certifiable methods are not appropriate for realtime robotic applications, they can still be quite useful in an offline setting. In particular, my recent work embeds certifiable methods in learning networks as a differentiable optimization layer [29]. In this paper, I showed that *local* differentiable optimization layers (e.g, Theseus [30]) are prone to spurious local minima that corrupt the training process. On the other hand, our certifiable approach used *implicit differentiation* to efficiently provide gradient information during backpropagation and *guaranteed correctness* of this information. I also used this method to substantially improve the accuracy of an existing deep-learned visual localization pipeline for field robotics.

[3]For the purposes of this note, chordal sparsity can be thought of as a *tree-like graph structure* connecting interacting groups of variables

III. Future Work

I have shown that certifiable methods improve the reliability of state estimation algorithms and can be adapted to even complex sensing modalities by using redundant constraints. In some cases, I have shown that these methods can remain performant *at scale*. However, these methods still lag behind other modern state-estimation approaches in terms of exploiting problem structure.

First, chordal sparsity techniques used in [28] have limited applicability in large-scale problems with structured loops. Indeed, when applied to the SLAM problems with *loop closures* in [18], I found that these techniques did not lead to a runtime speedup. Inspired by recent distributed algorithms for SLAM problems [31]–[33], I am currently investigating a **more general approach to exploiting chordal sparsity** that still retains global optimality guarantees.

Second, online state estimation approaches are designed to efficiently update their estimates as new data becomes available [26], [34], but, to my knowledge, there has been no certifiable method that is **explicitly incremental**. Since chordal sparsity techniques break relaxations into smaller subproblems, they could be coupled with warm-starting solvers (e.g., SCS [35]) to enable *incremental* certification. Equally promising is the fact that these subproblems can actually be **processed in parallel**, making them amenable to acceleration via modern GPU architectures [8].

In the short term, I plan to apply these ideas in the context of the extended SLAM approaches discussed in [18]. As a representative case study, the lessons learned in this project will be applicable to many large-scale robotics problems. Advances in these areas will firmly establish certifiable methods as mainstream tools for reliable and efficient state estimation.

Looking more towards long term goals, I plan to **further explore the combination of versatile deep-learning frameworks with the reliability of model-based optimization layers** explored in [29]. For example, large language models (LLM) and their variants have sparked a massive shift in state estimation towards *metric-semantic understanding*, particularly for SLAM and localization [36]–[38]. By replacing the backends of these methods with *certifiable and differentiable* optimization layers, we can achieve *versatile and reliable autonomy* that is capable of *learning and adapting online*.

In summary, certifiable methods enhance the reliability of robotics, a property that often has lower priority than speed and scalability. By unlocking the efficiency of these methods and making them more accessible, we can also unlock autonomy stacks that are safe, adaptive and performant.

References

[1] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK ; New York: Cambridge University Press, 2004.

[2] E. Wise, M. Giamou, S. Khoubyarian, A. Grover, and J. Kelly, "Certifiably Optimal Monocular Hand-Eye Calibration," in *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, Karlsruhe, Germany: IEEE, Sep. 2020, pp. 271–278.

[3] M. Giamou, F. Marić, D. M. Rosen, *et al.*, "Convex Iteration for Distance-Geometric Inverse Kinematics," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1952–1959, Apr. 2022.

[4] H. Yang and L. Carlone, "Certifiably Optimal Outlier-Robust Geometric Perception: Semidefinite Relaxations and Scalable Global Optimization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 2816–2834, Mar. 2023.

[5] H. Yang, J. Shi, and L. Carlone, "TEASER: Fast and Certifiable Point Cloud Registration," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 314–333, Apr. 2021.

[6] 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," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 95–125, Mar. 2019.

[7] A. Papalia, A. Fishberg, B. W. O'Neill, J. P. How, D. M. Rosen, and J. J. Leonard, "Certifiably Correct Range-Aided SLAM," *IEEE Transactions on Robotics*, vol. 40, pp. 4265–4283, 2024. (visited on 10/03/2024).

[8] S. Kang, X. Xu, J. Sarva, L. Liang, and H. Yang, "Fast and Certifiable Trajectory Optimization," *arXiv:2406.05846*, Jun. 2024. arXiv: 2406.05846 [cs, math].

[9] L. Carlone and F. Dellaert, "Duality-based verification techniques for 2D SLAM," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA: IEEE, May 2015, pp. 4589–4596.

[10] J. Briales and J. Gonzalez-Jimenez, "Convex Global 3D Registration with Lagrangian Duality," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI: IEEE, Jul. 2017, pp. 5612–5621.

[11] J. Zhao, W. Xu, and L. Kneip, "A Certifiably Globally Optimal Solution to Generalized Essential Matrix Estimation," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, USA: IEEE, Jun. 2020, pp. 12 031–12 040.

[12] J. Briales, L. Kneip, and J. Gonzalez-Jimenez, "A Certifiably Globally Optimal Solution to the Non-minimal Relative Pose Problem," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, UT: IEEE, Jun. 2018, pp. 145–154.

[13] M. Garcia-Salguero, J. Briales, and J. Gonzalez-Jimenez, "A Tighter Relaxation for the Relative Pose Problem Between Cameras," *J Math Imaging Vis*, Apr. 2022.

[14] A. Papachristodoulou and S. Prajna, "A tutorial on sum of squares techniques for systems analysis," in *Proceedings of the 2005, American Control Conference, 2005.*, Portland, OR, USA: IEEE, 2005, pp. 2686–2700.

[15] S. Kang, Y. Chen, H. Yang, and M. Pavone, "Verification and Synthesis of Robust Control Barrier Functions: Multilevel Polynomial Optimization and Semidefinite Relaxation," *arXiv:2303.10081*, Mar. 2023. arXiv: 2303.10081 [cs, eess, math].

[16] P. A. Parrilo and S. Lall, "Semidefinite Programming Relaxations and Algebraic Optimization in Control," *European Journal of Control*, vol. 9, no. 2-3, pp. 307–321, Jan. 2003.

[17] Anonymous, "Title removed for anonymous review," *IEEE Robot. Autom. Lett.*, Mar. 2023.

[18] Anonymous, "Title removed for anonymous review," *IEEE Transactions on Robotics*, 2024.

[19] Anonymous, "Title removed for anonymous review," *The International Journal of Robotics Research*, Sep. 2024.

[20] C. Olsson and A. Eriksson, "Solving quadratically constrained geometrical problems using lagrangian duality," in *2008 19th International Conference on Pattern Recognition*, Dec. 2008, pp. 1–5.

[21] J. Briales and J. Gonzalez-Jimenez, "Cartan-Sync: Fast and Global SE(d)-Synchronization," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2127–2134, Oct. 2017.

[22] Anonymous, "Title removed for anonymous review," *IEEE Robotics and Automation Letters*, Feb. 2023.

[23] Anonymous, "Title removed for anonymous review," *IEEE Transactions on Robotics*, 2024.

[24] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *IEEE Robotics & Automation Magazine*, vol. 13, no. 3, pp. 108–117, Sep. 2006.

[25] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part I," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, Jun. 2006.

[26] F. Dellaert and M. Kaess, "Factor Graphs for Robot Perception," *FNT in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.

[27] Y. Zheng, G. Fantuzzi, and A. Papachristodoulou, "Chordal and factor-width decompositions for scalable semidefinite and polynomial optimization," *Annual Reviews in Control*, vol. 52, pp. 243–279, 2021.

[28] Anonymous, "Title removed for anonymous review," *arXiv (To Be Published in WAFR)*, Jun. 2024.

[29] Anonymous, "Title removed for anonymous review," *arXiv (Under Review by T-RO)*, May 2024.

[30] L. Pineda, T. Fan, M. Monge, *et al.*, "Theseus: A library for differentiable nonlinear optimization," *Advances*

in Neural Information Processing Systems, vol. 35, pp. 3801–3818, 2022.

[31] D. McGann and M. Kaess, "iMESA: Incremental Distributed Optimization for Collaborative Simultaneous Localization and Mapping," in *Robotics: Science and Systems XX*, Robotics: Science and Systems Foundation, Jul. 2024.

[32] D. McGann, J. G. Rogers, and M. Kaess, "Robust Incremental Smoothing and Mapping (riSAM)," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, London, United Kingdom: IEEE, May 2023, pp. 4157–4163.

[33] Y. Tian, K. Khosoussi, D. M. Rosen, and J. P. How, "Distributed Certifiably Correct Pose-Graph Optimization," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 2137–2156, Dec. 2021.

[34] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, Feb. 2012.

[35] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, "Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding," *J Optim Theory Appl*, vol. 169, no. 3, pp. 1042–1068, Jun. 2016.

[36] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: An Open-Source Library for Real-Time Metric-Semantic Localization and Mapping," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, Paris, France: IEEE, May 2020, pp. 1689–1696.

[37] Y. Tian, Y. Chang, F. Herrera Arias, C. Nieto-Granda, J. P. How, and L. Carlone, "Kimera-Multi: Robust, Distributed, Dense Metric-Semantic SLAM for Multi-Robot Systems," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2022–2038, Aug. 2022.

[38] M. B. Peterson, Y. X. Jia, Y. Tian, A. Thomas, and J. P. How, "ROMAN: Open-Set Object Map Alignment for Robust View-Invariant Global Localization," *arXiv:2410.08262*, Oct. 2024. arXiv: 2410.08262.