

Sparse Variable Projection in Robotic Perception: Exploiting Separable Structure for Efficient Nonlinear Optimization

Alan Papalia^{1*}, Nikolas Sanderson², Haoyu Han³,
 Heng Yang³, Hanumant Singh², Michael Everett²

Abstract—Robotic perception often requires solving large nonlinear least-squares (NLS) problems. While sparsity has been well-exploited to scale solvers, a complementary and underexploited structure is *separability* – where some variables (e.g., visual landmarks) appear linearly in the residuals and, for any estimate of the remaining variables (e.g., poses), have a closed-form solution. Variable projection (VarPro) methods are a family of techniques that exploit this structure by analytically eliminating the linear variables and presenting a reduced problem in the remaining variables that has favorable properties. However, VarPro has seen limited use in robotic perception; a major challenge arises from gauge symmetries (e.g., cost invariance to global shifts and rotations), which are common in perception and induce specific computational challenges in standard VarPro approaches. We present a VarPro scheme designed for problems with gauge symmetries that jointly exploits separability and sparsity. Our method can be applied as a one-time preprocessing step to construct a *matrix-free Schur complement operator*. This operator allows efficient evaluation of costs, gradients, and Hessian-vector products of the reduced problem and readily integrates with standard iterative NLS solvers. We provide precise conditions under which our method applies, and describe extensions when these conditions are only partially met. Across synthetic and real benchmarks in SLAM, SNL, and SfM, our approach achieves up to $2\times$ – $35\times$ faster runtimes than state-of-the-art methods while maintaining accuracy. We release an open-source C++ implementation and all datasets from our experiments.

I. INTRODUCTION

Robotic perception often requires solving large-scale optimization problems, typically posed as nonlinear least-squares (NLS) problems with up to millions of variables [1]–[4]. Such NLS formulations underpin key tasks such as simultaneous localization and mapping (SLAM) [5], structure from motion (SfM) [6], and sensor network localization (SNL) [7]. In such tasks, a solver’s performance directly impacts the reliability and scale of robotic deployments [2]–[4].

While important, it is difficult to efficiently solve these problems because: (i) they are large, often involving 10^5 – 10^6 variables [1]; (ii) applications typically require solving problems to high numerical accuracy [8]; and (iii) problems often exhibit ill-conditioning that reduces the convergence rate of standard methods [9]. To address these challenges, state-of-the-art systems [10]–[12] exploit problem sparsity to handle

Work was supported by the Northeastern University Institute for Experiential Robotics Postdoctoral Fellowship, Army Research Lab awards W911NF-24-2-006 and W911NF-24-2-0017, and Office of Naval Research grant N000142512322. ¹University of Michigan. Work done while at Northeastern University. apapalia@umich.edu. ²Northeastern University, USA. {sanderson.n, h.singh, m.everett}@northeastern.edu. ³Harvard University. {hyhan, hankyang}@harvard.edu.


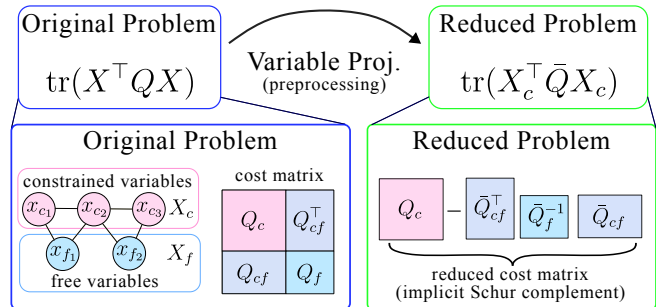


Fig. 1: **(Top) method overview:** Our approach exploits *separability* in optimization problem to perform variable projection and analytically eliminate a subset of variables (reducing problem size and improving conditioning), while preserving the efficiency of the original problem’s sparsity structure. Our approach can be applied as a one-time preprocessing step before passing the problem to a standard iterative solver. **(Bottom) runtime improvement:** Comparison of cost vs. time for our reduced problem (green) and the original problem (blue) on a real-world structure from motion dataset (BAL-1934). The variable projection step allows for substantial improvements in solver convergence and overall runtime.

scale [13], adopt second-order methods (e.g., Levenberg-Marquardt) to ensure solution accuracy [14], and employ numerically stable linear solvers to mitigate ill-conditioning [9]. Despite these techniques, scaling limitations remain, motivating approaches to further exploit problem structure.

One path to improved efficiency is to exploit *separable structure* [15], in which least-squares residuals depend linearly on a subset of unconstrained variables. Conditioning on the remaining variables allows these linear variables to be eliminated in closed form, a process known as variable projection (VarPro) [15]. This reduction both shrinks the problem size and improves convergence [16]. Yet in many perception problems, gauge symmetries in the cost (e.g., globally rotating the solution does not change the cost) introduce rank deficiencies in the cost matrix. This rank deficiency led prior approaches [17] to adopt workarounds that limited efficiency gains.

This paper. We show that a broad class of robotic perception problems (e.g., [18]–[21]) admits a VarPro scheme

that simultaneously exploits sparsity and separability. Our approach avoids challenges due to gauge symmetries and can be applied as a one-time step before optimization begins. We characterize this problem class by simple conditions on the cost and constraints that can be checked a priori:

- (i) there must be a set of unconstrained variables – these are the variables to be eliminated;
- (ii) the cost residuals must be linear functions; and
- (iii) the block of the cost residual Jacobian corresponding to the variables to be eliminated must have a specific graph-theoretic structure (see Section V-B).

Our approach produces a matrix-free Schur complement [22], i.e., an operator that can produce the reduced problem’s costs and gradients without forming the Schur complement (a large, dense matrix). This operator can be easily integrated into standard iterative solvers, preserving the scalability of modern systems. Finally, we discuss how the method extends when the ideal conditions are only partially satisfied (e.g., when some residuals are nonlinear).

Across a range of synthetic and real-world benchmarks in SLAM, SNL, and SfM, our method consistently outperforms state-of-the-art baselines, achieving runtime reductions of $2\times$ – $35\times$ and allowing for larger problems to be solved within memory limits. Our contributions are as follows:

- 1) a novel sparsity-preserving VarPro scheme to accelerate a broad class of perception problems;
- 2) a precise characterization of when this scheme applies, stated as simple conditions on costs and constraints; and
- 3) an open-source C++ implementation and all datasets used in the experiments, for reproducibility.¹

II. RELATED WORK

This review focuses on works which use variable projection (VarPro) in robotic perception. We categorize methods by how the reduced problem is constructed.

We do not discuss other techniques to improve solver efficiency such as general sparsity exploiting solvers [10]–[12], distributed computation [23]–[25], or initialization [26], [27].

For readers seeking broader background on VarPro, see Golub and Pereyra [15] for algorithms and applications and Ruhe and Wedin [16] for convergence analysis.

Schur complement. The most common VarPro approach in practice is the *Schur complement trick*: eliminate the linear variables via block elimination of the normal equations, leaving a reduced problem in the remaining variables [8], [22]. Schur methods come in two flavors. Explicit/direct Schur methods fully form the reduced matrix and optimize over the (dense) reduced problem. This can be effective in SfM instances when the reduced problem is relatively small [20], [28]. Implicit/matrix-free methods do not form the reduced matrix; they apply it as an operator via a series of matrix operations. This can better exploit problem sparsity and has been used in SfM [1], [29] and SLAM [18]. For SLAM, [17] gives a sparsity-preserving scheme that does

not use the Schur complement but is iteration-equivalent to the Schur-based approach of [30].

Projection/QR. Instead of modifying normal equations, projection methods *remove the effect of the linear variables* by projecting the residuals onto a subspace that implies the optimal assignment of the linear variables (obtained by thin QR factorization). From this the residual only depends on the remaining variables and optimization proceeds on that reduced problem. This has been used in appearance modeling [31] and computer vision applications of matrix factorization [32]–[34]. While numerically stable, the projection step can be costly and it is more difficult to exploit sparsity.

Placement of our work. Our work is situated among a large body of literature on matrix-free Schur complement methods [15], [16]. Our contribution is showing that (a) a broad class of perception problems (PGO, RA-SLAM, SNL, SfM) admits a sparse, matrix-free Schur complement operator despite the gauge symmetries that limited prior approaches, and (b) this operator can be constructed efficiently by removing a single row and column from specific problem matrices, followed by a sparse Cholesky factorization of the result. To our knowledge, only [17], [18] consider problems with gauge symmetries and try to preserve sparsity. Our work generalizes [18] to a broader problem class via a framework that clarifies when the approach applies and how it may be applied when conditions are only partially satisfied. Our work is more limited in scope than [17], which considers general separable residuals; we consider a more specific class of residuals, that admits additional structure that we exploit to improve computational efficiency.

III. PROBLEM FORMULATION

Many robotic perception problems estimate variables (e.g., robot poses, landmark locations) from noisy measurements (e.g., odometry, visual observations, range measurements) by minimizing a sum of squared residuals, where each residual measures the discrepancy between a predicted and actual measurement. We formalize this as a nonlinear least squares (NLS) problem, where the quantities to estimate are represented as matrices.

Problem 1 (Nonlinear Least Squares (NLS) Estimation).

$$\min_{X=[X_c; X_f]} \sum_{i=1}^m \|r_i(X)\|_{\Omega_i}^2 \quad (1)$$

$$\text{subject to } X_c \in \mathcal{X}_c \subseteq \mathbb{R}^{n_c \times d}$$

$$X_f \in \mathbb{R}^{n_f \times d}$$

$$X \in \mathbb{R}^{n \times d} \triangleq \begin{bmatrix} X_c \\ X_f \end{bmatrix} \in \begin{bmatrix} \mathcal{X}_c \subseteq \mathbb{R}^{n_c \times d} \\ \mathbb{R}^{n_f \times d} \end{bmatrix}, \quad (2)$$

where X_c are variables constrained to the domain \mathcal{X}_c , X_f are unconstrained variables, $r_i : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{k_i \times d}$ is the residual function associated with the i -th measurement, $\Omega_i \in \mathcal{S}_{++}^{k_i}$ is the positive definite concentration matrix indicating the precision of the i -th measurement, and $\|r_i(X)\|_{\Omega_i}^2 = \text{tr}(r_i(X)^\top \Omega_i r_i(X))$ is a weighted squared Frobenius norm.

¹Code: github.com/UMich-RobotExploration/variable-projection

TABLE I: Example linear residuals in state-of-the-art formulations. In these problems the variables are: R_i (rotation matrices), t_i (translation vectors), u_{ij} (unit bearing vectors), and RS_{ij} (scaled rotation matrices). The quantities \tilde{R}_{ij} , \tilde{t}_{ij} , \tilde{u}_{ij} are noisy measurements of the same objects and \tilde{d}_{ij} are noisy distance measurements. All of these residuals are homogeneous linear functions (i.e., $r_{ij}(X) = A_{ij}X$).

Measurement Type	Residual $r_{ij}(X)$
Relative rotation [18]	$R_j - R_i \tilde{R}_{ij}$
Relative translation [18]	$t_j - t_i - R_i \tilde{t}_{ij}$
Scale-free relative translation [20]	$t_j - t_i - RS_i \tilde{t}_{ij}$
Range [19], [21]	$t_j - t_i - u_{ij} \tilde{d}_{ij}$

This matrix-valued formulation generalizes the standard vector-valued NLS formulation, enabling us to address problems involving constrained variables that are difficult to express in vector form (e.g., rotations satisfying $R^\top R = I$ [18], [19]). We focus on a subclass of Problem 1 where the residuals are linear functions of the variables, enabling us to exploit separable structure via a one-time preprocessing step.

The case of linear residuals. When residuals $r_i(X)$ are linear, the NLS problem (Problem 1) reduces to a *quadratic* cost. We focus on *homogeneous* linear residuals, i.e., $r_i(X) = A_i X$ for some matrix $A_i \in \mathbb{R}^{k_i \times n}$, as (i) it simplifies notation, (ii) many common perception residuals are homogeneous (see Table I), and (iii) inhomogeneous residuals can be made homogeneous through variable augmentation [35].

By vertically stacking the Jacobians of the residuals A_i and block-diagonally arranging the concentration matrices Ω_i , we can rewrite the cost as

$$\sum_i \|r_i(X)\|_{\Omega_i}^2 = \|AX\|_{\Omega}^2 = \text{tr}(X^\top A^\top \Omega AX), \quad (3)$$

$$A \triangleq [A_1^\top \cdots A_m^\top]^\top, \quad (4)$$

$$\Omega \triangleq \text{blkdiag}(\Omega_1, \dots, \Omega_m), \quad (5)$$

where the $\text{blkdiag}(\cdot)$ operator constructs a block-diagonal matrix from its arguments.

We can reformulate the NLS problem (Problem 1) as follows, where $Q \triangleq A^\top \Omega A$:

Problem 2 (Constrained Quadratic Cost Problem).

$$\begin{aligned} \min_{X \in \mathbb{R}^{n \times d}} \quad & \text{tr}(X^\top QX) \\ \text{subject to} \quad & X_c \in \mathcal{X}_c, \end{aligned} \quad (6)$$

Importantly, Problem 2 possesses a globally quadratic cost (i.e., is quadratic without linearization). This enables VarPro as a one-time preprocessing step, as shown next.

IV. SEPARABLE STRUCTURE & VARIABLE PROJECTION

Separable structure in an optimization problem refers to the situation where, if certain variables are held fixed, the remaining variables can be efficiently solved in closed-form. This structure partitions the variables into two sets: those that are more difficult to optimize over (e.g., due to constraints) and those that are easily optimized.

Variable projection (VarPro) methods [15] exploit separable structure by iteratively optimizing over the ‘difficult’ variables while implicitly considering the ‘easy’ variables at their optimal values conditioned on the current iterate of the ‘difficult’ variables. VarPro effectively reduces the dimensionality of the problem to just the ‘difficult’ variables, and has been both theoretically and empirically shown to improve convergence rates of optimization [16].

The problems we consider exhibit separable structure that is particularly well suited for VarPro. Because the ‘easy’ variables are unconstrained and the cost is quadratic, a closed-form elimination of the unconstrained variables can be performed once and holds at every iteration. We now show how unconstrained variables in a quadratic cost induce separable structure and how the Schur complement can eliminate these variables.

With the variable ordering $X = [X_c^\top \mid X_f^\top]^\top$, we can partition the matrix Q as

$$Q = \begin{bmatrix} Q_c & Q_{cf} \\ Q_{cf}^\top & Q_f \end{bmatrix}. \quad (7)$$

For a fixed X_c , we set the gradient of the cost with respect to the unconstrained variables X_f to zero and solve for the optimal value X_f^* as a function of X_c :

$$\frac{\partial}{\partial X_f} \text{tr}(X^\top QX) = 2Q_f X_f + 2Q_{cf}^\top X_c = 0, \quad (8)$$

$$X_f^* = -Q_f^\dagger Q_{cf}^\top X_c, \quad (9)$$

where Q_f^\dagger is the Moore-Penrose pseudoinverse of Q_f .²

Plugging this optimal value X_f^* into the cost, we obtain a cost that depends only on the constrained variables X_c :

$$\text{tr}(X_c^\top Q_{sc} X_c) = \text{tr}\left(X_c^\top \left(Q_c - Q_{cf} Q_f^\dagger Q_{cf}^\top\right) X_c\right), \quad (10)$$

where $Q_{sc} \triangleq Q_c - Q_{cf} Q_f^\dagger Q_{cf}^\top$ is commonly called the *Schur complement* [22] of Q_f in Q .

This leads to the following reduced problem, which depends only on the constrained variables X_c ,

Problem 3 (*Reduced Constrained Quadratic Cost Problem*).

$$\min_{X_c \in \mathcal{X}_c} \text{tr}(X_c^\top Q_{sc} X_c) \quad (11)$$

Computational challenges in the reduced cost. The steps above yield a reduced problem (Problem 3) in the constrained variables X_c . However, two obstacles appear: (i) forming the Schur complement Q_{sc} requires a pseudoinverse, which is costly and numerically fragile; and (ii) Q_{sc} is typically dense even when the original matrix Q is sparse, making storage and operations expensive. Either issue can erase the

²This exact solution of X_f^* requires that $Q_f X_f = -Q_{cf}^\top X_c$ is consistent (solvable for all X_c), which holds iff $\text{range}(Q_{cf}^\top) \subseteq \text{range}(Q_f)$. From Eq. (12), $Q_f = A_f^\top \Omega A_f$ and $Q_{cf}^\top = A_f^\top \Omega A_c$. Since $\Omega \succ 0$, it follows that $\text{range}(Q_f) = \text{range}(A_f^\top)$ and that $\text{range}(Q_{cf}^\top) \subseteq \text{range}(A_f^\top)$, thus implying $\text{range}(Q_{cf}^\top) \subseteq \text{range}(Q_f)$. Additionally, since Q_f may be rank-deficient and the system may have multiple solutions, the pseudoinverse in Eq. (9) returns the minimum-norm solution X_f^* .

benefits of eliminating the unconstrained variables. We avoid these challenges by leveraging iterative methods that do not require explicitly forming Q_{Sc} .

Proposed approach: implicit (matrix-free) Schur via iterative methods. Rather than explicitly forming the dense Schur complement Q_{Sc} , we solve the reduced problem (Problem 3) using iterative methods [36], which only require computing matrix-vector products $Q_{Sc} X_c$. Our key contribution is an efficient, matrix-free routine for computing these products without ever forming Q_{Sc} explicitly or computing a pseudoinverse. Specifically, we derive an exactly equivalent reformulation of Q_{Sc} that can be applied via sparse matrix operations and triangular solves with Cholesky factors of the original problem data.

V. MATRIX-FREE SCHUR COMPLEMENT PRODUCTS

This section identifies the specific term that creates computational challenges in computing Schur complement products $Q_{Sc} X_c$, namely a pseudoinverse that creates a large, dense matrix. We then show how least-squares structure admits a reformulation that replaces the pseudoinverse with a positive definite matrix inverse, which can be efficiently represented via Cholesky factorizations. This reformulation allows us to compute products with Q_{Sc} via a series of sparse matrix products and triangular solves.

We revisit the quadratic cost matrix $Q = A^\top \Omega A$. By ordering the variables as $X = [X_c^\top \mid X_f^\top]^\top$, the Jacobian matrix becomes $A = [A_c \mid A_f]$, where A_c and A_f are the stacked Jacobians of the residuals with respect to the constrained and unconstrained variables, respectively. This then reframes Q and the Schur complement Q_{Sc} as:

$$Q = \begin{bmatrix} Q_c & A_c^\top \Omega A_f \\ A_f^\top \Omega A_c & A_f^\top \Omega A_f \end{bmatrix}, \quad (12)$$

$$Q_{Sc} = Q_c - A_c^\top \Omega A_f (A_f^\top \Omega A_f)^\dagger A_f^\top \Omega A_c. \quad (13)$$

Importantly, Q and the constituent matrices Q_c, A_c, A_f, Ω naturally inherit the sparsity of the graphical structure of the underlying problem. We can write products with the Schur complement $Q_{Sc} X_c$ as:

$$Q_{Sc} X_c = \underbrace{Q_c}_{\text{sparse}} X_c - \underbrace{\left(\underbrace{A_c^\top \Omega A_f}_{\text{sparse}} \underbrace{(A_f^\top \Omega A_f)^\dagger}_{\text{dense}} \underbrace{A_f^\top \Omega A_c}_{\text{sparse}} \right)}_{Q_2} X_c, \quad (14)$$

$$Q_2 \triangleq A_c^\top \Omega A_f (A_f^\top \Omega A_f)^\dagger A_f^\top \Omega A_c, \quad (15)$$

where we have highlighted the sparsity of each term assuming the original problem is sparse.

Key challenge: dense pseudoinverse in Schur complement products. Eq. (14) emphasizes that the pseudoinverse $(A_f^\top \Omega A_f)^\dagger$ creates a dense matrix that is the computational bottleneck in computing products with the Schur complement. If $(A_f^\top \Omega A_f)$ were full rank and positive definite, as in many VarPro applications [15], the pseudoinverse becomes an inverse. Since $(A_f^\top \Omega A_f)$ inherits the sparsity of the residuals, this inverse could be efficiently computed via sparse

Cholesky factorization and applied via sparse triangular solves. However, in many robotic perception problems the matrix is rank-deficient due to inherent symmetries.

A. Exact Reformulation of Q_2 via CR Decomposition

Our reformulation relies on the CR decomposition of a matrix [37], [38], which factorizes a matrix by its column and row spaces and can be computed e.g., via rank-revealing QR decomposition [9]. We first factorize A_f as

$$A_f = CR, \quad (16)$$

where the columns of C form a basis for the column space of A_f and the rows of R form a basis for the row space of A_f . The matrix C has full column rank and the matrix R has full row rank. Using this decomposition, we can rewrite the inner portion of Q_2 as

$$A_f (A_f^\top \Omega A_f)^\dagger A_f^\top = (CR)(R^\top C^\top \Omega CR)^\dagger (CR)^\top. \quad (17)$$

Let $M \triangleq C^\top \Omega C$. Since $\Omega \succ 0$ and C has full column rank, M is positive definite and therefore invertible. Because R has full row rank, the following identity holds

$$(R^\top MR)^\dagger = R^\dagger M^{-1} (R^\dagger)^\top. \quad (18)$$

Substituting Eq. (18) into Eq. (17) and recognizing that $RR^\dagger = I$ (since R has full row rank) yields

$$A_f (A_f^\top \Omega A_f)^\dagger A_f^\top = CR (R^\dagger M^{-1} (R^\dagger)^\top) R^\top C^\top \quad (19)$$

$$= C (RR^\dagger) M^{-1} (RR^\dagger)^\top C^\top \quad (20)$$

$$= CM^{-1} C^\top \quad (21)$$

$$= CL^{-\top} L^{-1} C^\top, \quad (22)$$

where the last step replaces the inverse via Cholesky factorization $M = LL^\top$.

Plugging this reformulation into Q_2 in Eq. (14) obtains

$$Q_{Sc} X_c = Q_c X_c - (A_c^\top \Omega C) L^{-\top} L^{-1} (C^\top \Omega A_c) X_c \quad (23)$$

$$= Q_c X_c - (BL^{-\top} L^{-1} B^\top) X_c,$$

where $B \triangleq A_c^\top \Omega C$. This can be performed as a series of matrix products and forward- and back-substitution with the Cholesky factors L and L^\top as described in Algorithm 1.

When is this reformulation valid? The reformulation in Eq. (23) is exact and only depends on the least-squares structure (which induces the block structuring of Q in Eq. (7)). This could be applied at each iteration of an iterative solver that linearizes the problem, though this incurs the cost of recomputing C and the Cholesky factorization. Because our problems have linear residuals, we perform this reformulation once as preprocessing.

When is the Schur product in Eq. (23) efficient? Efficiency is driven by the sparsity of C , which determines the density of both $B = A_c^\top \Omega C$ and the Cholesky factors $L^{-\top} L^{-1}$. As C is a linearly independent basis for the column space of A_f , there are many ways to compute it with sparsity in mind [39]. These approaches typically require either non-trivial computation (e.g., factorization) or *a priori* knowledge of the rank of A_f (e.g., to determine stopping in greedy algorithms).

Algorithm 1 Matrix-Free Schur Complement Products

Input: Matrices Q_c , A_c , A_f , Ω , and X_c .

Output: Product $Q_{Sc}X_c$
Precomputation (once, given Q_c , A_c , A_f , Ω):

- 1: $C \leftarrow \text{CR}(A_f)$. \triangleright CR decomposition (Section V)
- 2: $L \leftarrow \text{Cholesky}(C^\top \Omega C)$. \triangleright sparse factorization
- 3: $B \leftarrow A_c^\top \Omega C$.

Online computation (per new X_c):

- 1: $Y \leftarrow B^\top X_c$ \triangleright sparse matrix product
 - 2: $Z \leftarrow L^{-1}(L^{-\top} Y)$ \triangleright sparse triangular solves
 - 3: **return** $Q_{Sc}X_c = Q_cX_c - BZ$.
-

B. Leveraging Graph Structure for Efficient Schur Products

We now establish specific, yet common, graph-theoretic conditions on the Jacobian A_f under which the implicit Schur complement product in Eq. (23) retains the sparsity of the original problem. Furthermore, these conditions allow for closed-form computation of C and $(C^\top \Omega C)$ by simply removing rows and columns of the original matrices A_f and $(A_f^\top \Omega A_f)$, making the reformulation in Eq. (23) efficient.

Consider a graph G where each unconstrained variable is a node and any two variables that appear together in a residual are connected by an edge. Many robotic perception problems yield a connected graph over the unconstrained variables.

Furthermore, for residuals based on relative differences between pairs of unconstrained variables (i.e., $r_i(X) = g(X_c) + (X_{f_i} - X_{f_j})$ for some function $g(X_c)$), the Jacobian A_f can be interpreted as the *incidence matrix* of this graph G . I.e., each row of A_f corresponds to an edge in G and each column to a variable (node). For a row corresponding to an edge between nodes i and j , the k -th entry is

$$A_{f,(i,j)}(k) = \begin{cases} 1, & k = i \\ -1, & k = j \\ 0, & \text{otherwise} \end{cases}. \quad (24)$$


We leverage one key property of incidence matrices for connected graphs [40]: removing any one column of the incidence matrix yields a linearly independent basis for the column space of the matrix.³

Furthermore, the matrix $A_f^\top \Omega A_f$ is the *weighted graph Laplacian* of the graph, which is positive semidefinite with a single zero eigenvalue [40]. Closely related to the property above, the product $C^\top \Omega C$ when C is a reduced incidence matrix is called a *reduced graph Laplacian* and is positive definite [40]. The reduced graph Laplacian is formed by removing the same row and column from the graph Laplacian that was removed from A_f to form C .

As both the reduced incidence matrix and reduced graph Laplacian are formed by removing rows and columns of the original matrices A_f and $(A_f^\top \Omega A_f)$, they are easily

³The rank-deficiency of A_f is because the graphs inherently capture relative information. This can be seen from the fact that the all-ones vector is the null space of A_f ; adding a constant offset to all variables does not affect relative differences.

Schur Complement Products

Original Matrix (Q)

Explicit (Dense)


Implicit (Sparse)


Fig. 2: **Matrix-Free Schur Complement Products:** Here we demonstrate the difference between explicitly forming the Schur complement Q_{Sc} as a dense matrix versus performing a series of sparse operations to implicitly compute the product $Q_{Sc}X_c$ without forming the matrix. The top row shows the sparsity pattern of the original matrix Q . The bottom two rows show the sparsity patterns of the dense (explicit) and sparse (implicit) Schur complement approaches with the data from the ‘Garage’ dataset [27].

constructed and inherit the sparsity of the original matrices. As a result, the reduced incidence matrix is a natural choice for C in Eq. (23). With this choice, the only non-negligible computation required to form the operator for the Schur complement product (as in Algorithm 1) is a (sparse) Cholesky factorization of the reduced graph Laplacian.

In Fig. 2, we visualize the sparsity patterns of the matrices involved in the Schur complement product $Q_{Sc}X_c$ for a specific problem instance. In contrast to forming the dense Schur complement Q_{Sc} , the reformulation in Eq. (23) retains the sparsity of the original problem.

If A_f is not incidence-like (i.e., X_f does not appear in residuals solely through pairwise differences of its entries), C can instead be computed using basis-construction techniques [39]. This may incur additional computational cost and yield a denser C , but the reformulation in Eq. (23) remains exact.

VI. EXPERIMENTS

We analyze how our methodology impacts computational efficiency for four common state estimation problems: pose-graph optimization (PGO), range-aided SLAM (RA-SLAM), structure from motion (SfM), and sensor network localization (SNL). We consider both *iteration count* and *wall-clock time* to convergence.

All experiments use residuals from Table I, where our methodology fully exploits the problem structure: cost residuals are homogeneous linear functions (Section III) and the Jacobian block A_f corresponding to unconstrained variables is a directed incidence matrix (Section V-B).

A. Implementation and Baselines

Our approach. We implement our approach within a Riemannian trust-region (RTR) framework [41] using preconditioned truncated conjugate gradients (pTCG) [42, Ch. 6.5] for trust-region subproblems. We use the matrix-free Schur complement product algorithm (Section V-B) to efficiently compute matrix products, with Riemannian gradients and Hessian-vector products derived via standard conversions from their Euclidean counterparts [42]. We precondition

TABLE II: **Runtime and Iterations Results for all Experiments:** The median runtime (in seconds) and the median number of iterations each solver required for the proposed method (Ours) and the Original, Original + VarPro, and GTSAM approaches across all datasets, evaluated from random initializations over 5 trials. The runtime and iteration improvement factors are defined as $\frac{\text{Baseline}}{\text{Ours}}$. Values > 1 indicate our method required less time/iterations than the baseline and < 1 indicates our method required more (e.g., 2 means our method runs in half the time or with half the iterations). Each row’s fastest runtime is bolded. For runtime improvement factors, values > 1 are green and values < 1 are red. We use (—) to indicate a method failed to converge to the global minimum and (X) to indicate a method ran out of memory.

Dataset	Runtime (s) ↓				Solver Iterations ↓				Improvement Factor ↑		
	Ours	Original	Orig. + VP	GTSAM	Ours	Original	Orig. + VP	GTSAM	Runtime (Orig./Orig + VP/GTSAM)	Iterations (Orig./Orig + VP/GTSAM)	
PGO	Intel	0.19	0.83	0.5	0.63	21	69	40	12	4.37/2.63/3.32	3.29/1.9/0.57
	Garage	12.48	—	—	7.51	219	—	—	63	—/—/0.6	—/—/0.29
	Grid3D	3.6	10.22	4.26	—	14	22	13	—	2.84/1.18/—	1.57/0.93/—
	MIT	0.09	0.25	0.16	0.28	20	39	31	11	2.78/1.78/3.11	1.95/1.55/0.55
	M3500	0.72	11.48	8.06	2.85	31	183	120	22	15.94/11.19/3.96	5.9/3.87/0.71
	City10000	4.09	26.09	10.09	18.43	31	88	54	29	6.38/2.47/4.51	2.84/1.74/0.94
	Torus	0.89	1.09	1.13	37.24	12	16	13	36	1.22/1.27/41.84	1.33/1.08/3
	Sphere	1.06	2.06	1.83	14.68	20	30	26	39	1.94/1.73/13.85	1.5/1.3/1.95
RA-SLAM	TIERS	10.49	48.2	20.96	—	67	133	74	—	4.59/2/—	1.99/1.1/—
	Single Drone	1.03	5.53	3.62	4.91	32	130	79	57	5.37/3.51/4.77	4.06/2.47/1.78
	Plaza2	1.45	5.53	2.11	4.73	37	85	34	57	3.81/1.46/3.26	2.3/0.92/1.54
	Plaza1	12.76	22.07	16.87	10.91	62	140	74	63	1.73/1.32/0.86	2.26/1.19/1.02
	MR.CLAM2	6.17	—	13.57	11.59	28	—	34	18	—/2.2/1.88	—/1.21/0.64
	MR.CLAM4	4.53	33.09	11.77	8.02	26	178	30	15	7.3/2.6/1.77	6.85/1.15/0.58
	MR.CLAM6	3.07	17.03	5.21	5.89	31	219	34	25	5.55/1.7/1.92	7.06/1.1/0.81
	MR.CLAM7	3.18	23.42	7.65	10.32	32	224	32	36	7.36/2.41/3.25	7/—/1.12
SNL	Intel	0.13	—	0.78	—	32	—	121	—	—/6/—	—/3.78/—
	Garage	3.94	—	15.29	—	245	—	249	—	—/3.88/—	—/1.02/—
	Grid3D	10.56	—	16.94	—	63	—	83	—	—/1.6/—	—/1.32/—
	MIT	0.01	—	0.08	0.21	13	—	60	12	—/8/21	—/4.62/0.92
	M3500	0.69	—	—	—	61	—	—	—	—/—/—	—/—/—
	City10000	28.46	—	—	—	208	—	—	—	—/—/—	—/—/—
	Torus	1.99	—	20.83	—	40	—	137	—	—/10.47/—	—/3.42/—
	Sphere	2.65	—	—	—	70	—	—	—	—/—/—	—/—/—
SfM	BAL-93	0.39	4.5	1.45	10.09	17	213	63	52	11.54/3.72/25.87	12.53/3.71/3.06
	BAL-392	9.6	—	93.32	—	21	—	149	—	—/9.72/—	—/7.1/—
	BAL-1934	43.61	—	—	X	23	—	—	X	—/—/—	—/—/—
	IMC Gate	18.74	—	—	X	33	—	—	X	—/—/—	—/—/—
	IMC Temple	12.59	91.13	43.21	X	20	86	48	X	7.24/3.43/—	4.3/2.4/—
	IMC Rome	55.41	—	—	X	41	—	—	X	—/—/—	—/—/—
	Rep. Office0-100	0.33	2.07	0.92	8.39	11	40	20	10	6.27/2.79/25.42	3.64/1.82/0.91
	Rep. Office1-100	0.17	2.32	0.71	3.14	11	114	30	9	13.65/4.18/18.47	10.36/2.73/0.82
	Rep. Room0-100	0.54	3.12	1.4	12	13	38	19	9	5.78/2.59/22.22	2.92/1.46/0.69
	Rep. Room1-100	0.46	4.6	1.63	13.63	12	62	37	13	10/3.54/29.63	5.17/3.08/1.08
	Mip-NeRF Garden	1.48	9.53	3.98	19.76	11	48	25	12	6.44/2.69/13.35	4.36/2.27/1.09
	Mip-NeRF Room	17.82	—	54.79	—	20	—	48	—	—/3.07/—	—/2.4/—
	Mip-NeRF Kitchen	13.03	—	83.61	—	14	—	47	—	—/6.42/—	—/3.36/—
	TUM Room	15.39	—	89.11	X	23	—	60	X	—/5.79/—	—/2.61/—
	TUM Desk	8.12	—	33.32	X	14	—	51	X	—/4.1/—	—/3.64/—
	TUM Comp-R	5.38	83.09	20.09	X	16	249	47	X	15.44/3.73/—	15.56/2.94/—
	TUM Comp-T	10.59	—	56.92	X	17	—	62	X	—/5.37/—	—/3.65/—

pTCG with the *regularized Cholesky preconditioner* [19, Sec. VI.A] of Q : $P = L^{-T}L^{-1} = (Q + \mu I)^{-1}$, where μ is chosen so the condition number of P is below 10^6 . The preconditioner is stored as a sparse Cholesky factor $LL^T = (Q + \mu I)$ computed once at initialization. Since the reduced variable X_c does not match the preconditioner dimensions, we bottom-pad with zeros ($X = [X_c; 0]$). With respect to X_c , this is equivalent to preconditioning with the Schur complement of $(Q + \mu I)$, making it a principled preconditioner.

Baselines. We benchmark against three baselines: (i) *Original*, which directly optimizes the full problem without variable elimination; (ii) *Original + VarPro* [17], which optimizes the same problem as *Original* but uses variable projection to update unconstrained variables in closed form at each iteration; and (iii) the Levenberg-Marquardt implementation in GTSAM [12], a state-of-the-art solver using direct factorization-based linear solvers.

The first two baselines (Original and Original + VarPro) are implemented as alternative options within the same RTR

framework as our approach. Our GTSAM implementations uses custom manifolds and residuals from [43] to ensure identical problem formulations.

Hardware and software. Experiments were conducted on a laptop running Ubuntu 20.04 with an Intel i7-12700H CPU (20 threads) and 40 GB of RAM. All implementations are available as open-source C++ libraries along with data and scripts to reproduce the experiments.

B. Experiments

Dataset construction. The pose graph optimization (PGO) datasets are from [27]. The range-aided SLAM (RA-SLAM) datasets are from [19]. Sensor network localization (SNL) datasets were generated synthetically from the PGO datasets by converting all poses to points and all measurements to range measurements with identical noise levels. Structure from motion (SfM) datasets were generated according to [20, Sec. IV]. All dataset names match original sources.

Problem formulations. We use problem formulations empirically found to possess benign optimization landscapes

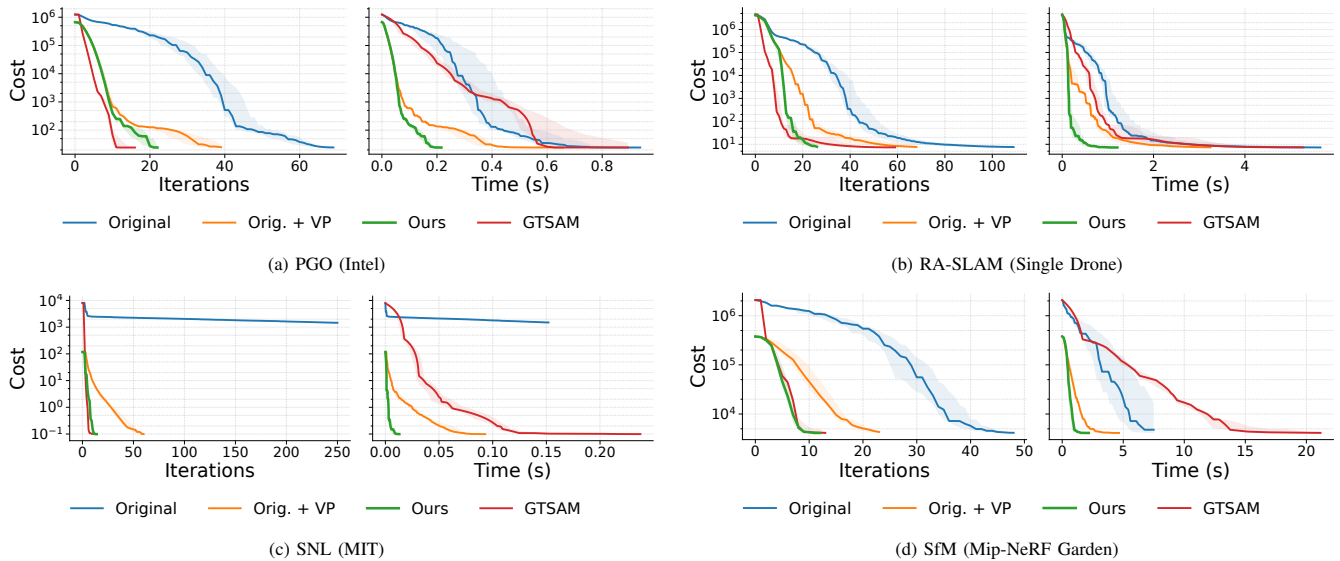


Fig. 3: **Convergence behavior on select problems.** Each panel pair shows (left) cost vs. iterations and (right) cost vs. time for our method and the baseline methodologies on representative datasets chosen from (a) pose-graph optimization (Intel), (b) RA-SLAM (Single Drone), (c) SNL (MIT), and (d) SfM (Mip-NeRF Garden). Notably, the most iteration-efficient method is not always the most time-efficient, as some methods have higher per-iteration costs.

[44], [45], meaning local optimization methods can reliably reach global minima from random initializations. Since all solvers obtain the global minimum, our analysis avoids complications from local minima.

Specifically, we use formulations from [18] (PGO), [19] (RA-SLAM), [21] (SNL), and [20] (SfM). These formulations are closely related: RA-SLAM generalizes both PGO and SNL. Our SfM formulation differs from [20] by omitting scale estimation in the constraints; the SfM formulation matches PGO but follows SfM’s bipartite sparsity pattern.

Experimental setup. We generate 5 random initializations per experiment and run each method on all 5 trials. Solvers run until convergence or 600 seconds elapsed. We compute optimal cost by solving a tight convex relaxation and consider a solver converged if it reaches within 1% of this minimum. Times and iteration counts in Table II are medians over 5 trials. Solver failures due to non-convergence (not reaching the global optimum within 600 seconds) or memory exhaustion are indicated by (–) and (X), respectively. In all instances where a solver failed in a trial, it was found to converge in none of the 5 trials.

C. Results

Runtime and iteration efficiency. Our approach achieved the fastest runtimes on 39/41 datasets (Table II), gaining efficiency through both improved per-iteration cost and reduced iteration counts (Fig. 3). On two datasets (Garage in PGO and Plaza 1 in RA-SLAM), GTSAM was faster by factors of 1.66 and 1.16, respectively. This is likely because these relatively small datasets had structure that was well-suited to GTSAM’s direct linear solvers.

Runtime improvements are particularly notable on SfM datasets, which typically have far more unconstrained variables (3D points) than constrained (orientations). Here, our method was often an order of magnitude faster than GTSAM and at least $2\times$ faster than *Original* and *Original + VarPro*.

Notably, despite being faster, our method often requires more iterations than GTSAM. This likely stems from GTSAM’s direct linear solvers, which are robust to conditioning and exactly solve subproblems (up to floating point precision), enabling more accurate steps per iteration. However, these direct solvers are often more computationally expensive, making GTSAM slower overall despite fewer iterations.

Memory efficiency. The SfM experiments show our method is more memory efficient than GTSAM, which often exhausted available memory. This is because our reduced problem preserves the sparsity of the original, and we solve it with iterative linear solvers (pTCG) rather than direct factorizations (as in GTSAM), which often require storing large dense matrices.

VII. CONCLUSION

This paper introduced a novel method to accelerate the solution of large-scale, separable nonlinear least-squares problems prevalent in robotic perception. Our approach leverages a sparsity-preserving VarPro scheme that can be applied as an efficient, one-time preprocessing step. The scheme represents the reduced problem with an efficient matrix-free Schur complement operator that readily integrates with standard iterative optimization frameworks.

Results, takeaways, and contributions. Experiments on diverse problems demonstrated the practical benefits of our approach. We achieved significant reductions in solver runtime compared to existing methods, including a prior VarPro technique [17]. Our method proved more scalable and memory-efficient than the widely deployed GTSAM solver, most notably on large-scale SfM instances which GTSAM was unable to solve due to memory constraints. This suggests that our method is well-suited for real-world robotic deployments.

The core of our method’s success lies in its ability to exploit both problem separability and sparsity, overcoming

the challenges posed by gauge symmetries that have limited previous approaches (see Section V-B for details). We also provided a clear characterization of the class of problems for which our method is applicable, offering a practical guide for its adoption on new problems.

To support further research and application, our work is released as an open-source C++ library, including all datasets and baselines from our experiments.

Limitations. Our method assumes residuals are separable and demonstrates specific advantages when the residuals are linear and the unconstrained Jacobian A_f admits an incidence-like structure (described in Section V-B). While these conditions are surprisingly common, there are many cases where they do not hold, with the most common violation being nonlinear residuals (e.g., reprojection errors in bundle adjustment or robust loss functions). In such cases, the general VarPro scheme can still be applied, but many of the efficiency benefits we demonstrate will not be realized. Studies of the practical performance of VarPro in these cases, and how much of the efficiency benefits can be retained, may be a fruitful direction for future work.

REFERENCES

- [1] S. Agarwal, N. Snavely, S. M. Seitz, and R. Szeliski, "Bundle adjustment in the large," en, in *Computer Vision – ECCV 2010*, vol. 6312, Series Title: Lecture Notes in Computer Science, Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 29–42, ISBN: 978-3-642-15551-2 978-3-642-15552-9.
- [2] K. Ebadi, L. Berreiter, H. Biggie, et al., "Present and future of SLAM in extreme environments: The DARPA SUBT challenge," *IEEE Transactions on Robotics*, vol. 40, pp. 936–959, 2023.
- [3] L. Kunze, N. Hawes, T. Duckett, M. Hanheide, and T. Krajnı́k, "Artificial intelligence for long-term robot autonomy: A survey," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4023–4030, 2018.
- [4] M. Tranzatto, T. Miki, M. Dharmadhikari, et al., "Cerberus in the DARPA subterranean challenge," *Science Robotics*, vol. 7, no. 66, eabp9742, 2022.
- [5] C. Cadena, L. Carlone, H. Carrillo, et al., "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on robotics*, vol. 32, no. 6, pp. 1309–1332, 2017.
- [6] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.
- [7] G. Mao, B. Fidan, and B. D. Anderson, "Wireless sensor network localization techniques," *Computer networks*, vol. 51, no. 10, pp. 2529–2553, 2007.
- [8] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment—a modern synthesis," en, in *Vision Algorithms: Theory and Practice*, G. Goos, J. Hartmanis, J. Van Leeuwen, B. Triggs, A. Zisserman, and R. Szeliski, Eds., vol. 1883, Series Title: Lecture Notes in Computer Science, Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 298–372, ISBN: 978-3-540-67973-8 978-3-540-44480-0.
- [9] G. H. Golub and C. F. Van Loan, *Matrix computations*. JHU press, 2013.
- [10] S. Agarwal, K. Mierle, and T. C. S. Team, *Ceres Solver*, version 2.2, Oct. 2023.
- [11] R. Kummerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "G²o: A general framework for graph optimization," en, in *2011 IEEE International Conference on Robotics and Automation*, Shanghai, China: IEEE, May 2011, pp. 3607–3613, ISBN: 978-1-61284-386-5.
- [12] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," *Georgia Institute of Technology, Tech. Rep.*, vol. 2, no. 4, 2012.
- [13] F. Dellaert, M. Kaess, et al., "Factor graphs for robot perception," *Foundations and Trends® in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [14] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 2006.
- [15] G. Golub and V. Pereyra, "Separable nonlinear least squares: The variable projection method and its applications," en, *Inverse Problems*, vol. 19, no. 2, R1–R26, Apr. 2003, ISSN: 0266-5611, 1361-6420.
- [16] A. Ruhe and P. Å. Wedin, "Algorithms for separable nonlinear least squares problems," *SIAM review*, vol. 22, no. 3, pp. 318–337, 1980.
- [17] K. Khosoussi, S. Huang, and G. Dissanayake, "A sparse separable SLAM back-end," en, *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1536–1549, Dec. 2016, ISSN: 1552-3098, 1941-0468.
- [18] 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," en, *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 95–125, Mar. 2019, ISSN: 0278-3649, 1741-3176.
- [19] A. Papalia, A. Fishberg, B. W. O'Neill, J. P. How, D. M. Rosen, and J. J. Leonard, "Certifiably correct range-aided SLAM," en, *IEEE Transactions on Robotics*, vol. 40, pp. 4265–4283, 2024, ISSN: 1552-3098, 1941-0468.
- [20] H. Han and H. Yang, "Building Rome with convex optimization," *Robotics: Science and Systems (RSS)*, Jul. 2025, arXiv:2502.04640 [cs].
- [21] T. Halsted and M. Schwager, "The Riemannian elevator for certifiable distance-based localization," *Preprint*, 2022, Accessed: Jan. 20, 2023. [Online] Available https://msl1.stanford.edu/papers/halsted_riemannian_2022.pdf.
- [22] F. Zhang, *The Schur complement and its applications*. Springer Science & Business Media, 2006, vol. 4.
- [23] Y. Tian, A. Koppel, A. S. Bedi, and J. P. How, "Asynchronous and parallel distributed pose graph optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5819–5826, 2020.
- [24] D. McGann, K. Lassak, and M. Kaess, "Asynchronous distributed smoothing and mapping via on-manifold consensus ADMM," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 4577–4583.
- [25] T. Fan and T. D. Murphey, "Majorization minimization methods for distributed pose graph optimization," *IEEE Transactions on Robotics*, vol. 40, pp. 22–42, 2024.
- [26] A. Papalia, J. Morales, K. J. Doherty, D. M. Rosen, and J. J. Leonard, "SCORE: A second-order conic initialization for range-aided SLAM," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 10637–10644.
- [27] L. Carlone, R. Tron, K. Daniilidis, and F. Dellaert, "Initialization techniques for 3D SLAM: A survey on rotation estimation and its use in pose graph optimization," en, in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA: IEEE, May 2015, pp. 4597–4604, ISBN: 978-1-4799-6923-4.
- [28] O. J. Woodford and E. Rosten, "Large scale photometric bundle adjustment," en, in *2020 British Machine Vision Conference*, arXiv:2008.11762 [cs], arXiv, Sep. 2020.
- [29] S. Weber, J. H. Hong, and D. Cremers, *Power variable projection for initialization-free large-scale bundle adjustment*, en, arXiv:2405.05079 [cs], Aug. 2024.
- [30] R. H. Barham and W. Drane, "An algorithm for least squares estimation of nonlinear parameters when some of the parameters are linear," *Technometrics*, vol. 14, no. 3, pp. 757–766, 1972.
- [31] I. Matthews and S. Baker, "Active appearance models revisited," en, *International Journal of Computer Vision*, vol. 60, no. 2, pp. 135–164, Nov. 2004, ISSN: 0920-5691, 1573-1405.
- [32] J. H. Hong, C. Zach, and A. Fitzgibbon, "Revisiting the variable projection method for separable nonlinear least squares problems," en, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI: IEEE, Jul. 2017, pp. 5939–5947, ISBN: 978-1-5386-0457-1.
- [33] T. Okatani, T. Yoshida, and K. Deguchi, "Efficient algorithm for low-rank matrix factorization with missing components and performance comparison of latest algorithms," en, in *2011 International Conference on Computer Vision*, Barcelona, Spain: IEEE, Nov. 2011, pp. 842–849, ISBN: 978-1-4577-1102-2 978-1-4577-1101-5 978-1-4577-1100-8.
- [34] J. H. Hong, C. Zach, A. Fitzgibbon, and R. Cipolla, "Projective bundle adjustment from arbitrary initialization using the variable projection method," en, in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., vol. 9905, Series Title: Lecture Notes in Computer Science, Cham: Springer International Publishing, 2016, pp. 477–493, ISBN: 978-3-319-46447-3 978-3-319-46448-0.
- [35] Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 20–34, 2010.
- [36] Y. Saad, *Iterative methods for sparse linear systems*. SIAM, 2003.
- [37] G. Strang and D. Drucker, "Three matrix factorizations from the steps of elimination," en, *Analysis and Applications*, vol. 20, no. 06, pp. 1147–1157, Nov. 2022, ISSN: 0219-5305, 1793-6861.
- [38] G. Strang, "Elimination and factorization," en, *Mathematics Magazine*, vol. 97, no. 5, pp. 484–487, Oct. 2024, ISSN: 0025-570X, 1930-0980.
- [39] T. F. Coleman and A. Pothén, "The null space problem II. algorithms," *SIAM Journal on Algebraic Discrete Methods*, vol. 8, no. 4, pp. 544–563, 1987.
- [40] F. R. Chung, *Spectral graph theory*. American Mathematical Soc., 1997, vol. 92.
- [41] P.-A. Absil, C. G. Baker, and K. A. Gallivan, "Trust-region methods on Riemannian manifolds," *Foundations of Computational Mathematics*, vol. 7, no. 3, pp. 303–330, 2007.
- [42] N. Boumal, *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- [43] Z. Xu, N. R. Sanderson, H. J. Zhang, and D. M. Rosen, *Certifiable estimation with factor graphs*, 2026. arXiv: 2603.01267 [cs.RO].
- [44] A. D. McRae and N. Boumal, "Benign landscapes of low-dimensional relaxations for orthogonal synchronization on general graphs," *SIAM Journal on Optimization*, vol. 34, no. 2, pp. 1427–1454, 2024.
- [45] C. Criscitiello, A. D. McRae, Q. Rebjock, and N. Boumal, "Sensor network localization has a benign landscape after low-dimensional relaxation," *arXiv preprint arXiv:2507.15662*, 2025.