

CSS-BA: Gate-Guided Column Space Search for Bundle Adjustment

Ayano Kaneda¹, Takafumi Taketomi², Shugo Yamaguchi¹, and
Shigeo Morishima¹

¹ Waseda University, Tokyo, Japan

² CyberAgent, Inc., Tokyo, Japan

Abstract. Bundle adjustment (BA) remains a critical refinement module for image-based 3D reconstruction and continues to improve geometric accuracy even in learning-based pipelines. However, in low-parallax and near-rotational regimes, classical Schur-based Levenberg–Marquardt (LM) often becomes ill-conditioned and yields unreliable pose and calibration estimates. We propose Gate-Guided CSS-BA, a solver-side modification of Schur-LM that preserves the classical BA objective and trust-region framework while constraining each update to a geometrically informed low-dimensional subspace. By integrating Column Space Search (CSS) with geometry-aware gating, the method stabilizes the Schur-LM update without altering the estimation problem. In contrast to keyframe or state-selection approaches, all camera and point parameters remain in the optimization problem; only the update direction is restricted. The method serves as a drop-in replacement for existing BA pipelines. Experiments on both generic and challenging weak-geometry scenarios show more stable optimization, improved relative pose accuracy, and competitive calibration behavior while maintaining reprojection quality.

1 Introduction

Despite rapid advances in learning-based and foundation-model approaches [5, 25, 26], bundle adjustment (BA) remains the dominant refinement stage for achieving high geometric accuracy. In practice, BA often determines the final pose and calibration quality. However, in low-parallax or near-rotational regimes, the Schur-reduced camera system becomes ill-conditioned. This is a BA-level issue: image observations constrain bearing directions but weakly separate translation, point depth, and focal calibration when the camera baseline is small. In a simplified two-view model, the disparity satisfies $d_{\text{disp}} \approx fB/Z$, where f is focal length, B is baseline, and Z is scene depth; since $\partial d_{\text{disp}}/\partial Z \approx -fB/Z^2$, depth sensitivity vanishes as $B \rightarrow 0$. This entanglement of translation, depth, and focal length creates weakly constrained directions in the linearized system that, after Schur point elimination, manifest as near-null camera-increment directions, regardless of the specific linear solver used. Classical Schur-based Levenberg–Marquardt (LM) may converge to solutions with small reprojection error while producing inaccurate relative poses and unstable calibration. This discrepancy

between reprojection accuracy and geometric reliability persists across both classical and learning-based pipelines.

Significant effort has been devoted to improving scalability and numerical efficiency of BA, including square root BA [4], matrix-free formulations [19], and inverse-expansion methods such as Power BA [27] and its variable-projection extensions [28]. While these approaches enhance conditioning or computational efficiency, they do not directly address the instability of the Schur-LM update under weak geometric conditioning.

We therefore take a complementary solver-side perspective. Rather than modifying the BA objective or its linearization, we retain the classical Schur-reduced camera system and trust-region framework, and instead restrict each LM step to a geometrically informed low-dimensional search subspace. In this way, we directly stabilize the update direction while preserving the full objective and variable set, aiming to improve relative pose quality without sacrificing reprojection accuracy.

We refer to the resulting method as **Gate-Guided CSS-BA**, a selective LM scheme operating entirely within the classical trust-region (TR) framework. At each iteration, a gated support of informative camera blocks is formed, and Column Space Search (CSS) constructs a compact LM basis from this support. Camera blocks are ranked using a Schur-consistent predicted reduction score. A low-dimensional search subspace is then constructed from the camera-block coordinate subspaces spanned by a small high-gain subset of the Schur camera system.

By concentrating updates on well-conditioned directions, the method improves trust-region acceptance and stability in weak geometry while remaining compatible with standard Schur-based BA infrastructures.

Our contributions are summarized as follows:

- **CSS instantiated on Schur-reduced BA.** We adapt Column Space Search to the Schur camera system and construct a compact search basis from gate-guided block statistics without modifying the objective, residuals, or Jacobians.
- **Direction-only restriction with full variables.** Top- k gated blocks are used only to localize the LM step direction; all camera variables remain in the optimization problem.
- **Stability gains in weak geometry.** The method improves relative pose accuracy in low-parallax regimes and can reduce focal-length calibration error under the hardest weak-geometry settings, while preserving reprojection quality and lowering the effective solve dimension.

2 Related Work

Bundle adjustment (BA) is the standard nonlinear least-squares refinement in multiview geometry [23]. Modern systems rely on Schur-complement reduction and sparse LM solvers, as implemented in widely used libraries such as SBA [14], g2o [12], Ceres [1], and COLMAP [20].

Prior work on BA has primarily focused on scalability, numerical conditioning, and alternative factorizations. In this section, we review efficiency- and stability-oriented approaches and position our solver-side update strategy within this landscape.

2.1 Efficiency-oriented methods

A major line of work improves the scalability of BA by reworking its linear algebra backend. Agarwal et al. [2] replace sparse Cholesky with preconditioned conjugate gradients on the Schur-reduced system, enabling large-scale Internet reconstructions. Subsequent efforts focus on parallelization and memory efficiency, including multicore acceleration [29], out-of-core optimization [15], and distributed solvers [18, 31].

More recent approaches explore approximate or matrix-free formulations. Zhou et al. [33] reduce computational cost via stochastic approximations, while Safari [19] proposes a matrix-free scheme for shared-intrinsic settings. In parallel, inverse-expansion methods such as Power BA [27] and its variable-projection extension [28] improve efficiency by approximating the Schur complement. Recent GPU-oriented work further accelerates the BA linearization and solve stages [30, 32].

These works primarily target scalability and conditioning through alternative linear solvers or factorizations.

2.2 Stability-oriented methods

A complementary line of work focuses on improving numerical robustness of BA under ill-conditioning and weak initialization. Square Root BA replaces the standard Schur elimination with a QR-based null-space marginalization, improving conditioning and enabling accurate single-precision solves [4].

Variable projection (VarPro) methods analytically eliminate separable parameters to enlarge convergence basins. Hong et al. [6, 7] applied VarPro to BA to obtain initialization-free formulations, and Iglesias et al. [9] introduced expOSE, a projective factorization with exponential regularization. Weber et al. [27, 28] further extended this direction: PoBA approximates the inverse Schur complement via power-series expansion, while PoVar combines power expansion with VarPro for large-scale initialization-free optimization.

Other recent works explore robustness through probabilistic or calibrated formulations, such as ProBA [3] and calibrated BA without initialization [16]. As a representative stress case, spherical SfM [24] highlights the rank-deficiency arising under near-spherical motion, where classical LM can converge to geometrically unreliable solutions.

These approaches improve stability primarily by altering factorizations, reparameterizations, or eliminating subsets of variables. In contrast, our method retains the classical Schur-reduced system and objective, and instead stabilizes the solver by restricting each LM update to a geometrically informed search subspace.

2.3 Selective LM under Trust Region

Our approach is grounded in the trust-region (TR) interpretation of Levenberg–Marquardt (LM). Under standard assumptions, LM enjoys quadratic local convergence, and with rank deficiency one can still obtain fast local rates under suitable error-bound conditions [11]. The trust-region acceptance rule further ensures global descent even when the Jacobian is not locally full rank.

Motivated by this trust-region view, we adopt Column Space Search (CSS) [8] to construct a selective solver-side update policy within the classical LM framework. In its original form, CSS scores candidate coordinate directions using residual/Jacobian information, prioritizes data-consistent directions, and solves a reduced subproblem without modifying the objective.

Related lines of work restrict updates by selecting keyframes or states over time (e.g., incremental smoothing, double-window optimization, relative BA, and rolling-shutter BA) [10,13,21,22]. Unlike these approaches, which alter the active variable set or measurement support, our method keeps the full BA objective and variable space fixed, and restricts only the solver-side search subspace at each LM iteration.

In our formulation, a geometrically supported block set is identified per iteration, a compact search basis is constructed, and a damped LM subproblem is solved in this subspace before lifting the increment back to the full variable space. Because acceptance and damping updates follow the same trust-region rule as the baseline LM solver, accepted steps remain governed by the same objective-decrease test while reducing exposure to ill-conditioned directions under weak geometry.

3 Method: Gate-Guided CSS-BA

Figure 1 provides an overview of the proposed Gate-Guided CSS-BA pipeline. We design a solver-side Levenberg–Marquardt (LM) scheme tailored to the Schur-reduced bundle adjustment (BA) system. In classical Schur-LM, full-dimensional camera updates are computed at each iteration. Under weak geometric conditioning, such as low-parallax or near-rotational motion, the reduced camera system can become ill-conditioned, causing the LM step to move along poorly supported directions despite small reprojection error.

Rather than modifying the BA objective, altering its linearization, or eliminating variables, we retain the standard Schur-reduced formulation and trust-region mechanism. Instead, we restrict each LM step to a geometrically informed low-dimensional search subspace defined in the Schur-reduced camera increment coordinates.

Unlike keyframe or state-selection methods, the objective and variable set remain unchanged; only the search directions used to compute $\Delta \mathbf{c}$ are constrained. The resulting increment is lifted back to the full camera space, so all camera variables remain in the optimization problem, although each iteration restricts the admissible update directions.




Fig. 1: Comparison between standard BA and the proposed CSS-BA on the damped Schur camera system. Standard BA directly solves the full damped Schur system to obtain a full-dimensional camera update. In contrast, CSS-BA first selects a top- k camera-block support K_t , constructs a low-dimensional Lanczos/Ritz basis on the localized Schur operator, and solves the projected LM system in this subspace. The resulting step is lifted back to the full camera space, so all camera variables remain in the optimization problem, while each iteration restricts the admissible update directions. The selected blocks determine camera-block *coordinate subspaces* for the basis, not an active-variable mask.

3.1 Preliminaries: Schur-Reduced BA

We consider the standard bundle adjustment objective

$$E(\mathbf{x}) = \frac{1}{2} \sum_{i,j} \|\mathbf{r}_{ij}(\mathbf{c}_i, \mathbf{p}_j)\|_2^2, \quad \mathbf{x} = [\mathbf{c}; \mathbf{p}], \quad (1)$$

where $\mathbf{c}_i$ and $\mathbf{p}_j$ denote camera and point parameters. Linearization at each LM iteration yields the damped normal equations

$$(H + \lambda D_{\mathbf{x}}^\top D_{\mathbf{x}}) \Delta \mathbf{x} = -\mathbf{g}, \quad H = J^\top W J, \quad \mathbf{g} = J^\top W \mathbf{r}. \quad (2)$$

Here, $\mathbf{r}$ denotes reprojection errors, J the Jacobian of the reprojection errors, W a weighting matrix (e.g., from a robust loss), H the Gauss-Newton approximation to the Hessian, $\lambda > 0$ the LM damping parameter, and $D_{\mathbf{x}} := \text{blkdiag}(D_c, D_p)$ the diagonal LM scaling matrix for camera and point variables.

Let

$$H_{cc,\lambda} := H_{cc} + \lambda D_c^\top D_c, \quad H_{pp,\lambda} := H_{pp} + \lambda D_p^\top D_p.$$

Eliminating the point increments gives the damped Schur system

$$\tilde{H}_\lambda \Delta \mathbf{c} = -\tilde{\mathbf{g}}_\lambda, \quad \tilde{H}_\lambda := H_{cc,\lambda} - H_{cp} H_{pp,\lambda}^{-1} H_{pc}, \quad \tilde{\mathbf{g}}_\lambda := \mathbf{g}_c - H_{cp} H_{pp,\lambda}^{-1} \mathbf{g}_p. \quad (3)$$

The candidate update $\Delta \mathbf{x} = [\Delta \mathbf{c}; \Delta \mathbf{p}]$ is accepted or rejected using the standard trust-region ratio

$$\rho = \frac{E(\mathbf{x}) - E(\mathbf{x} + \Delta \mathbf{x})}{-\mathbf{g}^\top \Delta \mathbf{x} - \frac{1}{2} \Delta \mathbf{x}^\top H \Delta \mathbf{x}}. \quad (4)$$

We retain the baseline LM damping and trust-region update unchanged. Our modification targets only how $\Delta \mathbf{c}$ is computed from the damped Schur system in (3).

After Schur reduction, the linear solve is expressed over camera increments, and landmark information remains encoded through point elimination. This reduced system is naturally block-structured by camera blocks. As discussed in Sec. 1, under low-parallax or near-rotational motion this system becomes ill-conditioned, creating near-null camera-increment directions that cause the LM step to move along poorly supported directions. Intuitively, small baselines provide weak sensitivity to translation, depth, and focal-length changes; after point elimination, these weak modes appear as near-null directions in the reduced camera system.

3.2 Column-Space Search as LM Search-Space Design

Column Space Search (CSS) [8] is a general subspace strategy for nonlinear least squares. It scores candidate coordinate subspaces and computes an update in the span of a small, data-consistent subset.

In the Schur-reduced BA system, the unknown of the linear solve is the camera increment $\Delta \mathbf{c}$. Therefore, we instantiate CSS over camera-block coordinate subspaces of $\Delta \mathbf{c}$. At iteration t , CSS scores camera blocks and selects a top- k support K_t as described in Sec. 3.3. The selected support is then used to construct a basis V_t for a low-dimensional search subspace.

This support should not be interpreted as selecting columns of the Schur matrix or as solving a separate BA problem on selected cameras. Each selected camera block defines a coordinate subspace of the camera-increment vector. Equivalently, using the block projector P_{K_t} introduced in Sec. 3.3, the support localizes the symmetric Schur operator on both sides: $\tilde{H}_{\lambda, K_t} = P_{K_t} \tilde{H}_{\lambda} P_{K_t}$.

Let $\mathcal{U}_t \subseteq \mathbb{R}^{n_c}$ be the search subspace at iteration t , where n_c is the total dimension of the camera-parameter vector $\mathbf{c}$. Let $V_t \in \mathbb{R}^{n_c \times r_t}$ be a full column-rank basis of $\mathcal{U}_t$, constructed as described in Sec. 3.3. We restrict only the camera-update direction by parameterizing $\Delta \mathbf{c} = V_t \mathbf{y}_t$. Substituting this into the damped Schur system in Eq. (3) gives the projected LM system

$$(V_t^\top \tilde{H}_{\lambda} V_t) \mathbf{y}_t = -V_t^\top \tilde{\mathbf{g}}_{\lambda}, \quad \Delta \mathbf{c} = V_t \mathbf{y}_t. \quad (5)$$

After solving for $\Delta \mathbf{c}$, we recover $\Delta \mathbf{p}$ by back-substitution using $H_{pp, \lambda}$. Acceptance and damping updates follow the same trust-region rule in Eq. (4) as the baseline LM solver.

The BA objective, residuals, and variable set remain unchanged. The increment $\Delta \mathbf{c} = V_t \mathbf{y}_t$ is a full-sized camera vector, but its admissible direction is restricted to $\text{range}(V_t)$ at the current iteration. Thus, K_t is a basis-construction support, not an active-variable mask.

3.3 CSS Subspace Construction

To construct a stable search subspace, we first identify cameras that are geometrically reliable under weak-geometry conditions. These cameras provide direc-

tions that are better conditioned in the Schur-reduced system and are therefore suitable for constructing a stable LM step.

Let t denote the LM iteration index. We distinguish two camera index sets: $\mathcal{S}$ is the geometry-aware support, computed once before optimization, and $K_t \subseteq \mathcal{S}$ is the top- k basis seed selected by CSS scoring at iteration t . Here $k = \min(k_{\max}, |\mathcal{S}|)$, with $k_{\max} = 10$ in our experiments.

The set $\mathcal{S}$ collects cameras that pass geometric reliability tests.

At iteration t , CSS scores the camera blocks in $\mathcal{S}$ and selects K_t to construct the subspace basis; high-gain blocks are those with large local predicted decrease in the current damped Schur model.

Gate decides eligibility; CSS decides which eligible blocks are informative. The gate forms a reliable support $\mathcal{S}$ using view connectivity, parallax, and rotational consistency. CSS then ranks only blocks within $\mathcal{S}$ by their local Schur-LM predicted decrease and selects the top- k set K_t . Together, gating restricts the candidate pool for basis construction to well-constrained cameras, while CSS selects the most useful eligible blocks for the current LM update.

Importantly, K_t determines basis directions only; it is *not* an active-variable mask on camera parameters.

Geometry-aware support $\mathcal{S}$ We identify geometrically informative cameras using diagnostics commonly employed in SfM and SLAM keyframe selection, such as parallax magnitude, rotational consistency, and view connectivity. These quantities detect weakly constrained configurations that often yield ill-conditioned Schur systems.

For camera pose $(R_i, \mathbf{t}_i)$ with center

$$\mathbf{C}_i := -R_i^\top \mathbf{t}_i, \quad (6)$$

we define valid two-view edges

$$(i, j) \in \mathcal{E} \iff |\mathcal{P}_{ij}| \geq \tau_{\text{shared}}, \quad (7)$$

where $\mathcal{P}_{ij}$ denotes the set of shared tracks observed in both cameras i and j , with neighborhood $\mathcal{N}(i) = \{j \mid (i, j) \in \mathcal{E}\}$.

For each shared track $q \in \mathcal{P}_{ij}$ with 3D point $\mathbf{X}_q$, define rays

$$\mathbf{u}_{iq} = \frac{\mathbf{X}_q - \mathbf{C}_i}{\|\mathbf{X}_q - \mathbf{C}_i\|}, \quad \mathbf{u}_{jq} = \frac{\mathbf{X}_q - \mathbf{C}_j}{\|\mathbf{X}_q - \mathbf{C}_j\|}, \quad (8)$$

and pairwise median parallax

$$\text{Par}_{ij} = \text{median}_{q \in \mathcal{P}_{ij}} \arccos(\text{clamp}(\mathbf{u}_{iq}^\top \mathbf{u}_{jq}, -1, 1)). \quad (9)$$

We define the parallax-filtered neighborhood

$$\mathcal{N}_E(i) = \{j \in \mathcal{N}(i) \mid \text{Par}_{ij} \geq \tau_{\text{par}}^{\text{edge}}\}. \quad (10)$$

The camera-level parallax score is the median over parallax-filtered neighboring views,

$$\text{Par}_i = \text{median}_{j \in \mathcal{N}_E(i)} \text{Par}_{ij}. \quad (11)$$

If $\mathcal{N}_E(i) = \emptyset$, we set $\text{Par}_i = 0$. For rotational consistency, let $\bar{R}_i$ denote the mean rotation of cameras in $\mathcal{N}(i)$. We compute the Rotation Agreement (RA) score as the geodesic distance between R_i and $\bar{R}_i$:

$$\text{RA}_i = \arccos\left(\text{clamp}\left(\frac{\text{tr}(R_i^\top \bar{R}_i) - 1}{2}, -1, 1\right)\right). \quad (12)$$

Here $\text{RA}_i = 0$ means that R_i is aligned with the average orientation of its neighbors; larger values indicate lower local rotational consistency. We use RA as a consistency filter, while parallax provides the main baseline-related geometric test.

The geometry-aware support is

$$\mathcal{S} = \{i \mid |\mathcal{N}_E(i)| \geq \tau_{\text{nbr}}, \quad \text{RA}_i \leq \tau_{\text{ra}}, \quad \text{Par}_i \geq \tau_{\text{par}}^{\text{cam}}\}. \quad (13)$$

Although similar geometric criteria are often used for keyframe selection in SfM and SLAM, here they serve only to define a stable support for subspace construction. No variables or measurements are removed from the optimization.

CSS Camera-Block Scoring for K_t At iteration t , blocks in $\mathcal{S}$ are scored on the damped Schur system to select the top- k set K_t . We partition the n_c -dimensional camera vector $\mathbf{c} = [\mathbf{c}_1; \dots; \mathbf{c}_B]$ into B camera blocks, where block i has dimension d_i and $n_c = \sum_{i=1}^B d_i$. Let $P_i \in \mathbb{R}^{n_c \times d_i}$ embed block i into the full vector.

Let $\tilde{\mathbf{g}}_{\lambda,i} := P_i^\top \tilde{\mathbf{g}}_\lambda$ and $\tilde{H}_{\lambda,ii} := P_i^\top \tilde{H}_\lambda P_i$ denote the corresponding blocks of the damped Schur system.

For each camera $i \in \mathcal{S}$, CSS computes

$$s_i = \frac{1}{2} \tilde{\mathbf{g}}_{\lambda,i}^\top \tilde{H}_{\lambda,ii}^{-1} \tilde{\mathbf{g}}_{\lambda,i}, \quad (14)$$

which is the local predicted decrease under single-block activation for the current damped Schur system.

We choose $K_t \subseteq \mathcal{S}$ as the top- k indices by descending score.

Basis Construction from K_t Given K_t , define the block projector

$$P_{K_t} = \sum_{i \in K_t} P_i P_i^\top, \quad (15)$$

and the localized Schur operator

$$\tilde{H}_{\lambda,K_t} = P_{K_t} \tilde{H}_\lambda P_{K_t}. \quad (16)$$

Thus, top- k selection localizes the operator rather than defining a coordinate-only active-variable set.

Algorithm 1 Gate-Guided CSS-BA

Require: Current state $(\mathbf{c}, \mathbf{p})$; precomputed support $\mathcal{S}$; options $(k_{\max}, m)$; LM damping λ ; trust-region parameters

- 1: **while** not converged **do**
- 2: Linearize residuals; apply LM damping and Schur $\Rightarrow$ damped reduced $(\tilde{H}_\lambda, \tilde{\mathbf{g}}_\lambda)$
- 3: $\mathbf{rhs} \leftarrow -\tilde{\mathbf{g}}_\lambda$ $\triangleright$ CSS scoring and selection inside the geometry-aware support
- 4: **for all** $i \in \mathcal{S}$ **do**
- 5: $s_i \leftarrow \frac{1}{2} \tilde{\mathbf{g}}_{\lambda,i}^\top \tilde{H}_{\lambda,ii}^{-1} \tilde{\mathbf{g}}_{\lambda,i}$
- 6: **end for**
- 7: $k \leftarrow \min(k_{\max}, |\mathcal{S}|)$ $\triangleright$ fixed top- k budget
- 8: $K_t \leftarrow$ top- k indices in $\mathcal{S}$ by s_i
- 9: $P_{K_t} \leftarrow \sum_{i \in K_t} P_i P_i^\top$ $\triangleright$ block projector, $P_i \in \mathbb{R}^{n_c \times d_i}$
- 10: $\tilde{H}_{\lambda,K_t} \leftarrow P_{K_t} \tilde{H}_\lambda P_{K_t}$ $\triangleright$ localized operator
- 11: Lanczos($\tilde{H}_{\lambda,K_t}, m$) $\Rightarrow (Q_m, T_m)$ $\triangleright$ Low-dimensional basis via Lanczos/Ritz
- 12: $C_k \leftarrow$ top- k eigenvectors of T_m $\triangleright$ Q_m : Krylov basis; T_m : tridiagonal matrix
- 13: $V_K \leftarrow Q_m C_k$ $\triangleright$ k Ritz vectors
- 14: $z_g \leftarrow (I - P_{K_t})(-\tilde{\mathbf{g}}_\lambda)$
- 15: **if** $\|z_g\|_2 > \epsilon_g$ **then**
- 16: orthogonalize and normalize z_g against V_K to obtain v_g
- 17: $V_t \leftarrow [V_K, v_g]$ $\triangleright$ $k + 1$ basis vectors
- 18: **else**
- 19: $V_t \leftarrow V_K$ $\triangleright$ k basis vectors
- 20: **end if** $\triangleright$ Projected LM solve and lifting (full-sized update)
- 21: Solve $(V_t^\top \tilde{H}_\lambda V_t) y = V_t^\top \mathbf{rhs}$
- 22: $\Delta \mathbf{c} \leftarrow V_t y$ $\triangleright$ full camera increment; not restricted to K_t or $\mathcal{S}$
- 23: Recover $\Delta \mathbf{p}$ by back-substitution using $H_{pp,\lambda}$
- 24: Evaluate the trial state obtained by applying $(\Delta \mathbf{c}, \Delta \mathbf{p})$ on the manifold
- 25: *Additional prediction safeguards and hybrid fallback are described in the supplementary material.*
- 26: Compute ρ by Eq. (4)
- 27: **if** $\rho > \eta_{\text{acc}}$ **then**
- 28: commit the trial update and decrease damping
- 29: **else**
- 30: discard the trial update and increase damping
- 31: **end if**
- 32: **end while**

Lanczos/Ritz basis. The localized operator $\tilde{H}_{\lambda,K_t}$ concentrates the Schur curvature on the selected camera blocks. Lanczos is initialized from the negative reduced gradient restricted to the selected support.

Running m -step Lanczos on $\tilde{H}_{\lambda,K_t}$ generates an orthonormal Krylov basis $Q_m \in \mathbb{R}^{n_c \times m}$ and a tridiagonal matrix $T_m \in \mathbb{R}^{m \times m}$. We take the top- k eigenvectors C_k of T_m and form

$$V_K = Q_m C_k \in \mathbb{R}^{n_c \times k}. \quad (17)$$

These vectors capture the dominant curvature directions within the selected support and provide a compact basis for the LM step before adding the optional complement direction. This converts the selected camera-block support into smooth joint update directions, rather than optimizing the selected cameras independently.

With the default setting, we also consider one complement direction outside the selected support:

$$z_g = (I - P_{K_t})(-\tilde{\mathbf{g}}_\lambda). \quad (18)$$

If $\|z_g\|_2 > \epsilon_g$, we orthogonalize and normalize it against V_K to obtain v_g and set

$$V_t = [V_K, v_g], \quad (19)$$

which has $k + 1$ columns. Otherwise, we set $V_t = V_K$, which has k columns. The complete procedure is summarized in Algorithm 1.

This basis V_t is plugged into (5) to solve the projected LM step and lift a full-sized increment.

Implementation details. Our implementation reuses the standard BA components: linearization, LM damping, Schur reduction, and manifold updates. CSS-BA only changes the camera-step subspace. In preprocessing, we construct the geometry-aware support set $\mathcal{S}$ from the diagnostics RA_i , $|\mathcal{N}_E(i)|$, and Par_i .

During LM iterations, each step scores blocks in $\mathcal{S}$ by (14), selects the top- k set K_t , constructs the Ritz basis, solves the projected Schur-LM system, and lifts the full camera increment.

The method uses runtime hyperparameters $(k_{\max}, m)$ and gate thresholds $(\tau_{\text{shared}}, \tau_{\text{nbr}}, \tau_{\text{ra}}, \tau_{\text{par}}^{\text{edge}}, \tau_{\text{par}}^{\text{cam}})$. The concrete values used in our experiments are reported in Sec. 4; implementation-level details of prediction safeguards and hybrid fallback are provided in the supplementary material.

4 Experiments

We evaluate Gate-Guided CSS-BA on two complementary regimes designed to test stability and generality. First, we consider PhoneSweep [24], a low-parallax dataset with near-spherical camera motion, to assess robustness under geometric degeneracy and calibration ambiguity. Second, we evaluate on Bundle Adjustment in the Large (BAL) [2], which contains large-scale Internet-photo reconstructions with more diverse motion, to assess solver behavior beyond near-rotational regimes.

Our implementation is built in C++ on top of the publicly available PoBA codebase [27], ensuring that Schur reduction, linearization, LM damping, and manifold updates are identical across solvers. Unless otherwise stated, all methods share the same trust-region strategy and stopping criteria to isolate differences in search-space design.

Experiments were conducted on Ubuntu 22.04.3 LTS (Linux 5.4.0-190-generic) with dual-socket AMD EPYC 7742 CPUs (256 hardware threads), 1TB RAM, and a single NVIDIA A100-SXM4 GPU (40GB, CUDA 12.2).

4.1 Datasets

PhoneSweep [24] contains thirteen real-world sequences with near-spherical camera motion (close to pure rotation, extremely low translational parallax), a regime that induces rank deficiency and calibration ambiguity in BA and provides a challenging stress test for solver stability. **BAL** [2] provides large-scale Internet-photo reconstructions with diverse camera trajectories (ladybug, trafilgar, dubrovnik families), and serves as a complementary generality check under well-conditioned motion.

4.2 Evaluation Metrics

For PhoneSweep, we report relative rotation and translation accuracies, $\text{RRA}@ \tau$ and $\text{RTA}@ \tau$, defined as the percentage of camera pairs whose relative rotation or translation angular error is below the threshold $\tau \in \{5^\circ, 15^\circ, 30^\circ\}$; higher is better. We also report $\text{AUC}@30$ (higher is better) and Absolute Focal Error $\text{AFE} = |\hat{f} - f_{\text{gt}}|/f_{\text{gt}}$ (lower is better). Formal definitions follow Ventura et al. [24] and are recalled in the supplementary material.

For BAL, we report per-instance objective-improvement statistics $r_i = f_i^{\text{final}}/f_i^0$ (smaller is better) and $\Delta_i = \log_{10}(f_i^0/f_i^{\text{final}})$ (larger is better), aggregated as family-wise means.

4.3 Experimental Setup

Common solver settings. We compare three solvers at the bundle adjustment stage: a standard Levenberg–Marquardt baseline (**Normal-LM**), Power Bundle Adjustment (**PoBA**) [27], and our **CSS-BA**. Normal-LM is the classical Schur-reduced LM solver that solves the full damped system (3) at each iteration without any subspace restriction. All solvers operate on the same Schur-reduced camera system.

We use identical trust-region parameters across solvers (initial radius 40.0, acceptance threshold $\eta = 0.1$) and stop when the relative decrease in reprojection cost falls below 10^{-2} or after at most 150 iterations.

Unless otherwise noted, we fix $\tau_{\text{nbr}} = 2$ and $\tau_{\text{par}}^{\text{cam}} = 2.0$ across datasets and adapt only the rotation gate to the motion regime.

For PoBA, we use 20 power iterations per linear solve. For CSS-BA, we use $\text{top-}k = 10$ with 32 Lanczos iterations. At each iteration, support cameras are selected from the geometry-aware support $\mathcal{S}$ using the CSS block-gain score.

PhoneSweep protocol. We evaluate on the iPhone13Mini and Nexus5X subsets. Each sequence is initialized using a general SfM pipeline with GLOMAP [17]. Due to the near-rotational, low-parallax motion pattern, we use a stricter rotation gate ($\tau_{\text{ra}} = 8.0^\circ$).

The original PhoneSweep benchmark assumes shared intrinsics across views. To isolate solver behavior, our primary comparison among Normal-LM, PoBA, and CSS-BA is conducted without shared intrinsics. We report SphericalSfM+BA only as a reference baseline, not as a directly matched configuration.

Table 1: Results on the PhoneSweep dataset. SphericalSfM with the shared-intrinsic assumption is used as a reference baseline. $\uparrow$ indicates higher is better and $\downarrow$ lower is better. Compared with Normal-LM and PoBA, CSS-BA substantially improves relative-pose accuracy on both devices; focal calibration is competitive but dataset-dependent.

Dataset	Method	RRA $\uparrow$			RTA $\uparrow$			AUC@30 $\uparrow$	AFE $\downarrow$
		@5	@15	@30	@5	@15	@30		
iPhone13Mini	Normal-LM	17.40	53.49	81.38	6.24	28.73	51.26	13.34	115.56
	PoBA [27]	41.58	73.98	87.62	17.22	52.99	62.13	42.62	169.43
	CSS-BA	85.79	86.55	87.82	77.24	86.32	89.68	79.97	158.01
	Baseline	100.00	100.00	100.00	86.65	98.77	99.77	91.45	0.25
Nexus5X	Normal-LM	8.06	44.32	76.29	5.10	21.38	43.32	12.06	265.35
	PoBA [27]	14.15	66.90	88.10	16.14	52.90	75.04	33.01	121.18
	CSS-BA	86.32	91.09	95.98	81.12	94.79	98.41	83.78	0.89
	Baseline	100.00	100.00	100.00	83.48	96.56	98.83	90.43	0.97

Table 2: BAL family-wise improvement summary. For each problem i , we compute $r_i = f_i^{\text{final}}/f_i^0$ (smaller is better) and $\Delta_i = \log_{10}(f_i^0/f_i^{\text{final}})$ (larger is better), then aggregate within each family (mean over problems).

Family	#Probs.	Normal-LM		PoBA		CSS-BA	
		mean $r \downarrow$	mean $\Delta \uparrow$	mean $r \downarrow$	mean $\Delta \uparrow$	mean $r \downarrow$	mean $\Delta \uparrow$
ladybug	20	0.01556	1.8713	0.016339	1.860373	0.0172	1.8089
trafalgar	14	0.00484	2.3542	0.004707	2.338294	0.0046	2.3400
dubrovnik	16	0.00484	2.3254	0.004907	2.319353	0.0053	2.2919

BAL protocol. We use the official BAL problem files and preserve the original camera parameterization, 3D points, and observations. All solvers are initialized from the same provided state and optimized with identical trust-region parameters and stopping criteria. Since BAL contains more diverse camera motion than PhoneSweep, we adopt a looser rotation gate ($\tau_{\text{ra}} = 20^\circ$) to avoid over-pruning informative cameras.

5 Results

5.1 PhoneSweep Results

Table 1 compares Normal-LM, PoBA, and CSS-BA on PhoneSweep. SphericalSfM+BA with shared intrinsics is shown as a reference.

Across both devices, CSS-BA substantially improves relative-pose accuracy over Normal-LM and PoBA. On iPhone13Mini, AUC@30 rises to 79.97 from 13.34 (Normal-LM) and 42.62 (PoBA). On Nexus5X, the gain is larger: 83.78 vs. 12.06 and 33.01. On Nexus5X, CSS-BA also reduces AFE from 265.35/121.18 to 0.89, indicating that the baselines largely fail to recover focal calibration under this extreme degeneracy.

Figure 2 shows that Normal-LM and PoBA often produce collapsed or fragmented trajectories, while CSS-BA recovers the characteristic near-spherical cir-




Fig. 2: Qualitative comparison of reconstructed camera trajectories on two PhoneSweep sequences: `Nexus5X/engineering_west` (top) and `iPhone13Mini/scott_shand` (bottom). Ground-truth camera layouts exhibit the characteristic near-spherical circular trajectory of PhoneSweep. Normal-LM and PoBA often fail to preserve this geometry, producing collapsed, fragmented, or strongly distorted camera configurations. In contrast, CSS-BA recovers camera trajectories that remain close to the ground-truth circular layout, indicating substantially improved robustness under low-parallax, near-rotational motion.

cular layout. CSS-BA operates without shared intrinsics or explicit motion constraints, yet approaches the strongly regularized SphericalSfM+BA baseline.

5.2 BAL Results

Table 2 shows that CSS-BA obtains objective reductions comparable to Normal-LM and PoBA on BAL across all three families. CSS-BA is not consistently the best solver on these less-degenerate problems, but the results suggest that the subspace restriction does not introduce severe degradation in standard large-scale BA settings. The Venice BAL family is reported in the supplementary material.

5.3 Ablation

Gate ablation. Table 3 ablates the geometric gate on PhoneSweep (GLOMAP init.). The gate is a stability filter: it helps most on the harder Nexus5X setting, while iPhone13Mini variants are close due to near-saturation. The AUC/AFE ordering is not strictly monotone with gate complexity because the update direction is determined by CSS ranking and the Ritz basis, not by the gate alone. **Top- k sensitivity.** Table 4 reports AUC@30 and AFE for $k \in \{4, 6, 8, 10, 12, 16\}$ on three PhoneSweep sequences under GLOMAP initialization. Metrics vary

Table 3: Gate ablation on PhoneSweep (GLOMAP initialization). The gate acts as a stability filter; it helps most on the harder Nexus5X setting.

Gate	iPhone13Mini		Nexus5X	
	AUC@30↑	AFE↓	AUC@30↑	AFE↓
Full gate	79.97	158.01	83.78	0.89
No gate	79.98	157.82	79.56	1.06
RA only	79.91	157.30	79.65	0.88
Par only	79.94	158.25	79.61	0.94

Table 4: Top- k sensitivity on three PhoneSweep sequences.

k	N5X/eng.		i13/scott		i13/plaza	
	AUC@30↑	AFE↓	AUC@30↑	AFE↓	AUC@30↑	AFE↓
4	97.747	0.503	95.365	0.601	94.327	0.692
6	97.746	0.128	95.394	0.299	93.872	0.586
8	97.762	0.430	95.365	0.601	93.426	0.576
10	97.753	0.447	95.363	0.601	94.249	0.654
12	97.754	0.447	95.365	0.601	93.294	0.580
16	97.757	0.467	95.391	0.379	93.367	0.578

little across the range, supporting a fixed default $k = 10$ without per-scene tuning.

Runtime analysis. CSS-BA introduces per-iteration overhead from block scoring, Lanczos basis construction, and the projected solve. Table 5 reports representative wall-clock timings on two PhoneSweep sequences under the same 150-iteration cap. Aggregating these timed cases gives Normal-LM / PoBA / CSS-BA ≈ 6495 / 1930 / 13142s, so CSS-BA costs about $2.0\times$ Normal-LM and $6.8\times$ PoBA on this weak-geometry timing subset. Together with the accuracy gains reported above, this shows that CSS-BA trades throughput for improved pose stability in the targeted ill-conditioned regime. Calibration gains are strongest on Nexus5X and remain dataset-dependent.

On BAL (Table 2), the results show comparable objective improvement; the added subspace construction overhead is therefore mainly justified when weak geometry makes accuracy the primary concern.

6 Conclusion and Limitations

We introduced Gate-Guided CSS-BA, a solver-side Column Space Search variant of Schur-LM tailored to bundle adjustment under weak geometric conditioning. Operating on the Schur-reduced camera system, the method combines geometry-aware support selection with a subspace-restricted update while preserving the original objective, variable set, and trust-region mechanism.

Table 5: Representative wall-clock runtime on two PhoneSweep sequences under the 150-iteration cap.

Case	Normal-LM	PoBA	CSS-BA	vs. Normal	vs. PoBA
Seq. A	2306 s	683 s	4478 s	$1.94\times$	$6.56\times$
Seq. B	4189 s	1247 s	8664 s	$2.07\times$	$6.95\times$
Total	6495 s	1930 s	13142 s	$2.02\times$	$6.81\times$

On near-spherical, low-parallax PhoneSweep sequences, CSS-BA improves relative pose accuracy over standard LM and PoBA and, in the Nexus5X case, substantially reduces focal-length error. It approaches the strongly regularized SphericalSfM+BA reference in relative-pose and trajectory-layout metrics without shared intrinsics or explicit motion priors.

Because CSS-BA only changes the LM update subspace, it is compatible with standard robust estimators and outlier handling. Its main benefit is in ill-conditioned regimes where classical Schur-LM is prone to poorly supported updates; in well-conditioned cases, gains are more modest and subspace construction adds overhead. Adaptive gating and subspace sizing remain future work.

References

1. Agarwal, S., Mierle, K., Team, T.C.S.: Ceres solver. <http://ceres-solver.org> (2022), accessed: 2026-03-01
2. Agarwal, S., Snavely, N., Seitz, S.M., Szeliski, R.: Bundle adjustment in the large. In: Eur. Conf. Comput. Vis. (2010)
3. Chui, J., Cremers, D.: Proba: Probabilistic bundle adjustment with the bhat-tacharyya coefficient (2025), arXiv preprint arXiv:2505.20858
4. Demmel, N., Sommer, C., Cremers, D., Usenko, V.: Square root bundle adjustment for large-scale reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. (2021)
5. Duisterhof, B.P., Zust, L., Weinzaepfel, P., Leroy, V., Cabon, Y., Revaud, J.: MAST3R-SfM: A fully-integrated solution for unconstrained structure-from-motion. In: Proceedings of International Conference on 3D Vision (3DV). pp. 1–10 (2025)
6. Hong, J.H., Zach, C., Fitzgibbon, A., Cipolla, R.: Projective bundle adjustment from arbitrary initialization using the variable projection method. In: Eur. Conf. Comput. Vis. Lecture Notes in Computer Science, vol. 9905, pp. 477–493 (2016)
7. Hong, J.H., Zach, C., Fitzgibbon, A., Cipolla, R.: pOSE: Pseudo object space error for initialization-free bundle adjustment. In: IEEE Conf. Comput. Vis. Pattern Recog. (2018)
8. Hyde, D.A.B., Bao, M., Fedkiw, R.: On obtaining sparse semantic solutions for inverse problems, control, and neural network training. *Journal of Computational Physics* **443**, 110498 (2021)
9. Iglesias, J.P., Nilsson, A., Olsson, C.: expOSE: Accurate initialization-free projective factorization using exponential regularization. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)

10. Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J.J., Dellaert, F.: iSAM2: Incremental smoothing and mapping using the bayes tree. *The International Journal of Robotics Research* **31**(2), 216–235 (2012)
11. Kanzow, C., Yamashita, N., Fukushima, M.: Levenberg–Marquardt methods with strong local convergence properties for solving nonlinear equations with convex constraints. *Journal of Computational and Applied Mathematics* **172**, 375–397 (2004)
12. Kümmerle, R., Grisetti, G., Strasdat, H., Konolige, K., Burgard, W.: g2o: A general framework for graph optimization. In: *ICRA* (2011)
13. Liao, B., Qu, D., Xue, Y., Zhang, H., Lao, Y.: Revisiting rolling shutter bundle adjustment: Toward accurate and fast solution. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
14. Lourakis, M.I.A., Argyros, A.A.: SBA: A software package for generic sparse bundle adjustment. *ACM Transactions on Mathematical Software* **36**(1) (2009)
15. Ni, K., Steedly, D., Dellaert, F.: Out-of-core bundle adjustment for large-scale 3D reconstruction. In: *Int. Conf. Comput. Vis.* (2007)
16. Olsson, C., Nilsson, A.: Towards initialization-free calibrated bundle adjustment (2025), arXiv preprint arXiv:2506.23808
17. Pan, L., Barath, D., Pollefeys, M., Schönberger, J.L.: Global structure-from-motion revisited. In: *Eur. Conf. Comput. Vis.* (2024)
18. Ren, J., Liang, W., Yan, R., Mai, L., Liu, S., Liu, X.: MegBA: A GPU-based distributed library for large-scale bundle adjustment. In: *Eur. Conf. Comput. Vis.* (2022)
19. Safari, D.: Matrix-free shared intrinsics bundle adjustment. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2025)
20. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2016)
21. Sibley, G.: Relative bundle adjustment. Tech. Rep. Technical Report 2307/09, Oxford University Computing Laboratory (2009)
22. Strasdat, H., Davison, A.J., Montiel, J.M.M., Konolige, K.: Double window optimisation for constant time visual SLAM. In: *Int. Conf. Comput. Vis.* (2011)
23. Triggs, B., McLauchlan, P.F., Hartley, R.I., Fitzgibbon, A.W.: Bundle adjustment—a modern synthesis. In: *Vision Algorithms: Theory and Practice*, pp. 298–372. Springer (2000)
24. Ventura, J., Larsson, V., Kahl, F.: Uncalibrated structure from motion on a sphere. In: *Int. Conf. Comput. Vis.* (2025)
25. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: VGGT: Visual geometry grounded transformer. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2025)
26. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: DUS3R: Geometric 3D vision made easy. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2024)
27. Weber, S., Demmel, N., Chan, T.C., Cremers, D.: Power bundle adjustment for large-scale 3D reconstruction. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
28. Weber, S., Hong, J.H., Cremers, D.: Power variable projection for initialization-free large-scale bundle adjustment. In: *Eur. Conf. Comput. Vis.* (2024)
29. Wu, C., Agarwal, S., Curless, B., Seitz, S.M.: Multicore bundle adjustment. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2011)
30. Zhan, Z., Xu, H., Fang, Z., Wei, X., Hu, Y., Wang, C.: Bundle adjustment in the eager mode (2024), arXiv preprint arXiv:2409.12190
31. Zhang, R., Zhu, S., Fang, T., Quan, L.: Distributed very large scale bundle adjustment by global camera consensus. In: *Int. Conf. Comput. Vis.* (2017)

32. Zhong, J., Zhan, Z., Gao, Q., Chen, Z., Lou, H., Mao, J., Neumann, U., Wang, C., Wang, Y.: InstantSfM: Towards GPU-native SfM for the deep learning era (2025), arXiv preprint arXiv:2510.13310
33. Zhou, L., Luo, Z., Zhen, M., Shen, T., Li, S., Huang, Z., Fang, T., Quan, L.: Stochastic bundle adjustment for efficient and scalable 3D reconstruction. In: Eur. Conf. Comput. Vis. (2020)