BARYSLAM: BARYCENTRIC CONSENSUS GEOMETRY FOR COLLABORATIVE 3D GAUSSIAN SLAM

Anonymous authors
Paper under double-blind review

ABSTRACT

Collaborative 3D Gaussian SLAM must reconcile independently reconstructed submaps whose overlap is often partial and whose Gaussian decompositions may differ across agents. Pairwise registration and sequential fusion can make the global estimate sensitive to uncertain intermediate geometry and fusion order. We present BarySLAM, which organizes cross-agent information around Barycentric Consensus Geometry, a partially supported latent geometric representation shared by agents observing the same scene structure. BarySLAM derives structural observations from local Gaussian fields, establishes partial cross-agent support through unbalanced optimal transport, and jointly refines submap poses and consensus geometry with held-out validation to suppress inconsistent cross-agent evidence. The validated consensus is then transferred to the original Gaussian primitives for global map construction, maintaining a common geometric basis from collaborative estimation to reconstruction. Experiments on synthetic and real multi-agent benchmarks show SOTA overall performance in trajectory estimation and Gaussian reconstruction, with improved robustness under limited overlap and corrupted cross-agent registration and reduced sensitivity to fusion order.

1 INTRODUCTION

Collaborative SLAM estimates a common multi-agent state from measurements acquired by independently operating robots, combining intra-agent observations such as odometry and loop closures with inter-agent measurements that couple trajectories across robot-specific reference frames (Schmuck & Chli, 2019; Tian et al., 2022; Lajoie & Beltrame, 2024). A standard probabilistic formulation separates these two sources of information as

$$\widehat{\mathcal{X}}=\arg\max_{\mathcal{X}} p(\mathcal{X})\prod_{r=1}^{R} p\big(\mathcal{Z}_r^{\mathrm{intra}}\,|\,\mathcal{X}_r\big)\prod_{(r,s)\in\mathcal{C}} p\big(\mathcal{Z}_{rs}^{\mathrm{inter}}\,|\,\mathcal{X}_r,\mathcal{X}_s\big), \tag{1}$$

where $\mathcal{X}_r$ denotes the state associated with agent r, $\mathcal{Z}_r^{\mathrm{intra}}$ collects its intra-agent measurements, and $\mathcal{Z}_{rs}^{\mathrm{inter}}$ contains observations coupling agents r and s. Collaborative 3D Gaussian SLAM extends this estimation setting to independently reconstructed and directly renderable Gaussian submaps, combining inter-agent association, global pose estimation, and Gaussian-map construction within a common multi-agent system (Yugay et al., 2025; Xu et al., 2025; Chen et al., 2026; Ding et al., 2026; Thomas et al., 2025).

Independently optimized Gaussian submaps do not preserve persistent primitive identities across agents. Densification, pruning, covariance adaptation, and appearance optimization may represent the same physical structure with different Gaussian decompositions, while cross-agent visibility determines which parts of these reconstructions are jointly observed (Xu et al., 2025; Cheng et al., 2025; Shorinwa et al., 2025). The geometric evidence available to collaborative estimation is therefore both heterogeneous and spatially incomplete.

Limited overlap further weakens direct cross-agent relations. An individual pair of submaps may share only a narrow scene region even when the complete agent collection contains sufficient complementary geometry. Registration uncertainty becomes more consequential when transformed geometry participates in an evolving collaborative state. Sequential map construction introduces

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

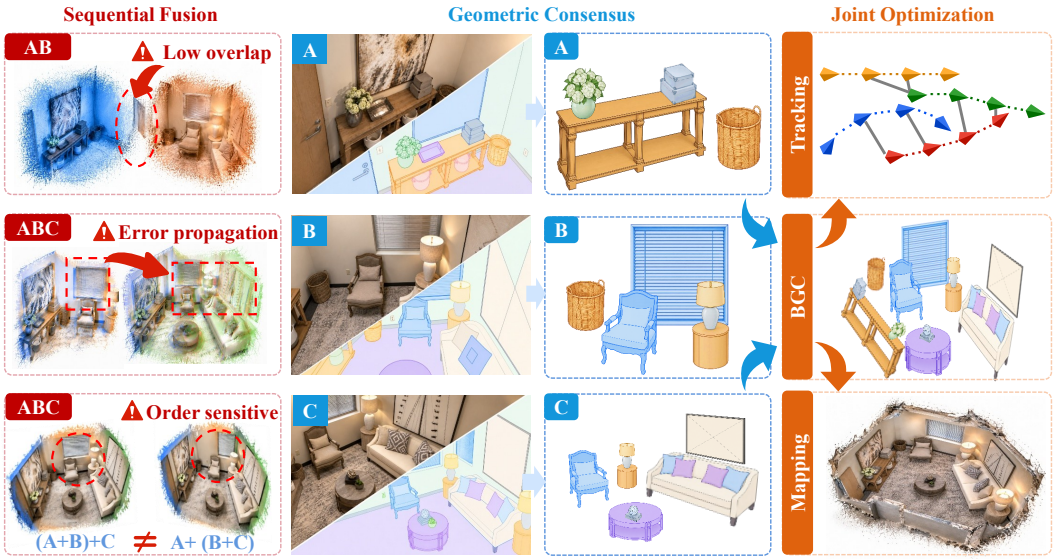


Figure 1: Comparison of cross-agent information organization in sequential fusion and BarySLAM. BarySLAM jointly estimates and preserves the Barycentric Consensus Geometry across independently reconstructed Gaussian submaps.

order-dependent dependencies because each incorporated submap changes the intermediate geometry available to later processing (Evangelidis et al., 2014; Gojcic et al., 2020; Jin et al., 2024; Yugay et al., 2025; Xu et al., 2025).

Across these settings, three recurring difficulties become coupled in collaborative Gaussian mapping. A relation restricted to two submaps can access only their common support, while a progressively updated reference already reflects earlier estimation and integration decisions.

Multiway registration offers a different organization in which several observations jointly constrain a shared latent representation (Evangelidis et al., 2014; Gojcic et al., 2020). Wasserstein barycenters provide a corresponding construction for distributions, with the Gaussian W_2 metric combining Euclidean displacement of the means and Bures–Wasserstein geometry of the covariances (Agueh & Carlier, 2011; Álvarez-Esteban et al., 2016; Malagò et al., 2018; Bhatia et al., 2019). This shared-state view motivates a geometric representation that can aggregate evidence from multiple independently reconstructed maps without assigning any single submap as a persistent reference.

Following this view, we introduce Barycentric Consensus Geometry (BCG) for collaborative 3D Gaussian SLAM. BCG therefore provides a common geometric reference without requiring persistent primitive correspondence across independently reconstructed Gaussian maps.

BarySLAM realizes this shared-state formulation through a unified estimation and reconstruction process. Experiments on synthetic and real multi-agent benchmarks show state-of-the-art overall performance in trajectory estimation and Gaussian reconstruction, together with improved robustness under limited overlap and corrupted cross-agent registration and reduced sensitivity to fusion order.

The main contributions are summarized as follows.

1. We introduce Barycentric Consensus Geometry, a partially supported latent state for multiway cross-agent geometric consensus.
2. We formulate joint estimation of submap poses and consensus geometry under partial overlap and uncertain cross-agent registration.
3. We transfer the validated consensus to Gaussian-map consolidation, linking collaborative estimation with global reconstruction while reducing fusion-order sensitivity.

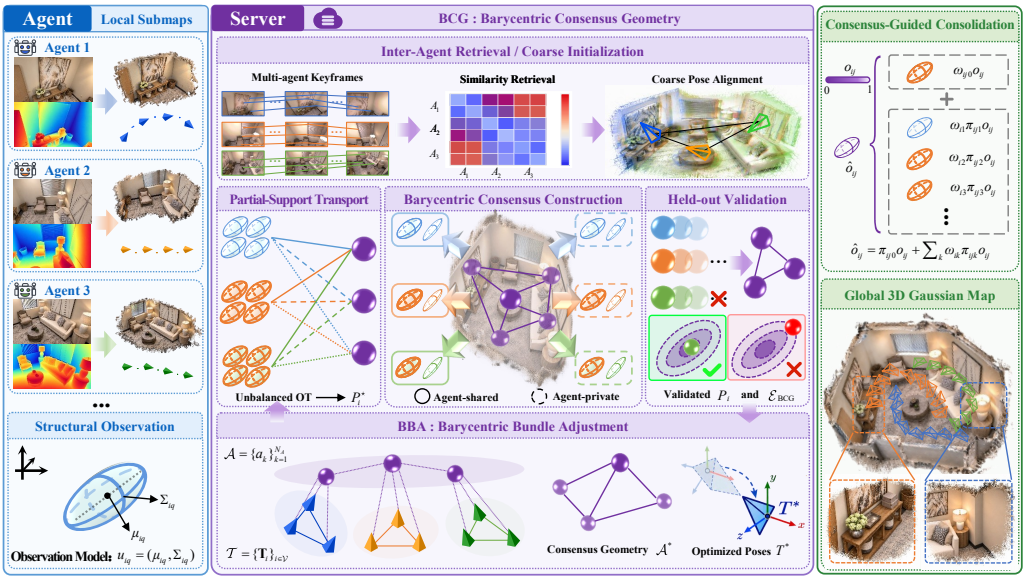


Figure 2: Overview of BarySLAM. Local Gaussian submaps are organized through Barycentric Consensus Geometry, where partial transport establishes cross-agent support and predictive validation removes inconsistent associations. Barycentric Bundle Adjustment jointly refines submap poses and consensus geometry, and the validated consensus subsequently guides global Gaussian-map consolidation.

2 METHOD

2.1 OVERVIEW

BarySLAM organizes cross-agent geometric estimation in collaborative Gaussian SLAM through Barycentric Consensus Geometry (BCG). Each agent provides its locally optimized Gaussian submaps together with the conventional odometry and intra-agent loop-closure constraints. Fig. 2 summarizes the resulting information flow.

A common collaborative SLAM formulation encodes cross-agent information through direct pose-to-pose or map-to-map relations. Such relations are effective when the participating maps share sufficient geometry, while incomplete overlap and uncertain registration make the geometric reference increasingly important. Multiway registration offers a different perspective by aligning several observations through a shared latent representation (Evangelidis et al., 2014; Gojcic et al., 2020). BarySLAM extends this principle to collaborative Gaussian SLAM through BCG, a spatially supported latent geometry placed between independently reconstructed submaps. Let $\mathcal{V}$ denote the uploaded submap nodes and $\mathcal{E}_{\text{intra}}$ the odometry and intra-agent loop-closure edges inherited from the host SLAM systems. BarySLAM connects a submap pose $\mathbf{T}_i \in \mathrm{SE}(3)$ to the latent consensus regions supported by that submap. The resulting collaborative graph is bipartite on its cross-agent side. Conventional pose nodes remain connected by $\mathcal{E}_{\text{intra}}$, while additional factors connect submap poses to shared latent geometry.

The latent nodes form a Barycentric Consensus Atlas (BCA), $\mathcal{A} = \{a_k\}_{k=1}^{N_A}$. Unlike a single global reference map, an atlas element is supported only by the agents that observe the corresponding scene structure. Geometry visible to only one agent remains outside the shared state. This spatially varying support allows cross-agent constraints to arise only where common geometric evidence is available. Independently optimized Gaussian maps first provide structural observations without requiring persistent splat identities. Partial transport then determines which observations support common latent geometry, and the resulting consensus factors are screened before entering the global estimator. The validated factors are optimized jointly with the original pose graph, after which the same support is transferred back to the appearance-bearing Gaussian primitives for global map construction.

2.2 BARYCENTRIC CONSENSUS GRAPH CONSTRUCTION

To support latent multiway registration across independently optimized Gaussian maps, BCG must handle unstable primitive identities and partial spatial overlap. Densification, pruning, and local map optimization prevent persistent splat correspondences across reconstructions, while different agents typically observe only partially shared geometry. We therefore represent each local Gaussian map through the continuous field induced by its primitives. At analysis site q in submap i, the localized field is summarized by a Gaussian structural observation $u_{iq} = (\boldsymbol{\mu}_{iq}, \boldsymbol{\Sigma}_{iq})$, where the mean encodes local position and the covariance captures anisotropic spatial structure. Under the current submap pose, each observation is transformed into the common frame and compared with a latent consensus element using

$$d^2_{\mathrm{BCG}}(\widetilde{u}_{iq}, a_k) = \frac{\left\|\widetilde{\boldsymbol{\mu}}_{iq} - \bar{\boldsymbol{\mu}}_k\right\|_2^2}{\sigma_\mu^2} + \frac{d^2_{\mathrm{BW}}\left(\widetilde{\boldsymbol{\Sigma}}_{iq}, \bar{\boldsymbol{\Sigma}}_k\right)}{\sigma_\Sigma^2}. \tag{2}$$

Here $\mathbf{T}_i = (\mathbf{R}_i, \mathbf{t}_i) \in \mathrm{SE}(3)$ transforms the observation as $\widetilde{\boldsymbol{\mu}}_{iq} = \mathbf{R}_i\boldsymbol{\mu}_{iq} + \mathbf{t}_i$ and $\widetilde{\boldsymbol{\Sigma}}_{iq} = \mathbf{R}_i\boldsymbol{\Sigma}_{iq}\mathbf{R}_i^\top$. The latent element $a_k = (\bar{\boldsymbol{\mu}}_k, \bar{\boldsymbol{\Sigma}}_k)$ belongs to the current BCA $\mathcal{A}$, while σ_μ and σ_Σ balance positional and anisotropic-shape disagreement. The covariance term follows the Bures–Wasserstein geometry induced by the Gaussian 2-Wasserstein metric (Agueh & Carlier, 2011; Malagò et al., 2018; Bhatia et al., 2019).

Partial overlap requires the shared geometry to admit only supported local evidence. Unbalanced optimal transport allows part of the source evidence to remain unmatched by relaxing exact marginal conservation (Chizat et al., 2018). BarySLAM adopts this principle for graph construction and additionally imposes a source-capacity constraint so that local evidence cannot be reused beyond its available mass. For submap i, the partial coupling is obtained from

$$\mathbf{P}_i^\star = \underset{\mathbf{P}\geq 0,\ \mathbf{P}\mathbf{1}_{N_A}\leq \mathbf{m}_i}{\arg\min} \ \langle \mathbf{P}, \mathbf{C}_i\rangle + \varepsilon\ \mathrm{Ent}(\mathbf{P}) + \tau_s D_{\mathrm{KL}}(\mathbf{P}\mathbf{1}_{N_A}\|\mathbf{m}_i) + \tau_a D_{\mathrm{KL}}\left(\mathbf{P}^\top\mathbf{1}_{Q_i}\middle\|\mathbf{m}_A\right). \tag{3}$$

The coupling $\mathbf{P}_i \in \mathbb{R}_+^{Q_i\times N_A}$ transports soft structural mass from the Q_i observations of submap i to the current atlas, with $[\mathbf{C}_i]_{qk} = d^2_{\mathrm{BCG}}(\widetilde{u}_{iq}, a_k)$. The source mass $\mathbf{m}_i$ bounds the available evidence, while $\mathbf{m}_A$ serves as a soft atlas-occupancy reference. We use the negative entropy $\mathrm{Ent}(\mathbf{P}) = \sum_{q,k} P_{qk}(\log P_{qk} - 1)$ together with the generalized KL divergence for nonnegative masses. The KL terms relax the source and atlas marginals, whereas the constraint $\mathbf{P}\mathbf{1}_{N_A} \leq \mathbf{m}_i$ prevents repeated use of source evidence. Consequently, the residual $\mathbf{m}_i - \mathbf{P}_i^\star\mathbf{1}_{N_A} \geq 0$ represents local evidence excluded from the shared geometry. The resulting constrained entropic transport problem is solved using Bregman–Dykstra projections following regularized transport methods (Benamou et al., 2015); numerical details are provided in the Appendix.

Because entropic transport yields soft, generally dense couplings, nonzero transport mass does not directly define an active graph factor. A provisional incidence (i, k) is activated only when the transported support $m_{ik}^{\mathrm{tr}} = \sum_q P^\star_{i,qk}$ exceeds τ_{supp}, suppressing entropic tails while retaining meaningful shared evidence. Atlas elements supported by fewer than two distinct agents are excluded from the collaborative graph. For each active incidence, the transported observations are summarized by their first and second moments as a Gaussian structural observation $u_{ik} = (\boldsymbol{\mu}_{ik}, \boldsymbol{\Sigma}_{ik})$. When multiple submaps originate from the same agent, their support is pooled before normalization so that subdividing a local reconstruction does not artificially increase that agent's contribution.

Once the support of a latent region has been established, its shared geometry is estimated collectively from the participating observations. Wasserstein barycenters provide a canonical Fréchet mean for probability distributions (Agueh & Carlier, 2011; Álvarez-Esteban et al., 2016). BarySLAM therefore estimates the provisional consensus element as

$$\widetilde{a}_k = \underset{a\in\mathbb{R}^3\times\mathbb{S}^3_{++}}{\arg\min} \sum_{i\in\mathcal{S}_k} \omega_{ik}\, d^2_{\mathrm{BCG}}(\widetilde{u}_{ik}, a)\,. \tag{4}$$

The support set $\mathcal{S}_k$ contains the submaps associated with latent element k, while $\widetilde{u}_{ik}$ and ω_{ik} denote the corresponding common-frame observation and normalized agent-balanced weight. For Gaussian observations, Eq. (4) reduces to a weighted mean of centers and a Bures–Wasserstein barycenter of covariances. The resulting $\{\widetilde{a}_k\}$ form the provisional atlas $\widetilde{\mathcal{A}}$ for subsequent optimization.

2.3 CONSENSUS-VALIDATED JOINT OPTIMIZATION

The transport stage proposes submap–consensus incidences, and these provisional relations are screened before entering the global estimator. Robust multi-robot SLAM commonly evaluates uncertain inter-robot measurements before global optimization (Mangelson et al., 2018; Forsgren et al., 2024). BarySLAM carries the same consistency principle to the latent consensus graph. Directly testing a submap against the current consensus node would allow the tested agent to validate a factor using geometry that already contains its own evidence. BarySLAM avoids this self-support by rebuilding the corresponding consensus element after removing every submap of the tested agent and evaluating the provisional factor against this held-out geometry.

Conditioned on the current transport epoch, the predictive consistency score and retained support are

$$\begin{aligned} e_{ik} &= (\mathbf{r}_{ik}^{-r})^{\top} \boldsymbol{\Omega}_{ik}^{\mathrm{pred}} \mathbf{r}_{ik}^{-r}, \\ \widehat{\mathcal{S}}_k &= \{ i \in \mathcal{S}_k \,|\, e_{ik} \le \tau_{\mathrm{val}} \}. \end{aligned} \tag{5}$$

For a tested submap i belonging to agent r, the residual $\mathbf{r}_{ik}^{-r}$ compares its structural observation with the held-out consensus element a_k^{-r} reconstructed without any submap from that agent. The predictive information matrix $\boldsymbol{\Omega}_{ik}^{\mathrm{pred}}$ accounts for uncertainty in both the tested observation and the held-out consensus anchor. With two participating agents, the test reduces to reciprocal geometric consistency. With three or more agents, the same construction additionally exploits multiway redundancy in the held-out reference.

The surviving factors define $\widehat{\mathcal{E}}_{\mathrm{BCG}} = \{(i,k) \mid i \in \widehat{\mathcal{S}}_k\}$. Validation also masks the provisional transport plan. Transport entries associated with a rejected submap–consensus incidence are set to zero, yielding the validated coupling $\widehat{\mathbf{P}}_i$. If fewer than two distinct agents remain after validation, the corresponding latent node is removed. The agent-balanced weights are then recomputed and normalized over the validated support $\widehat{\mathcal{S}}_k$.

After validation, the collaborative graph consists of the original intra-agent pose factors and the retained submap–consensus factors. With this topology fixed for one continuous optimization epoch, the pose nodes and latent geometry are jointly updated through

$$(\mathcal{T}^{+}, \mathcal{A}^{+}) = \arg\min_{\mathcal{T},\mathcal{A}} \sum_{(i,j)\in\mathcal{E}_{\mathrm{intra}}} \left\| \mathbf{r}_{ij}^{\mathrm{pose}} \right\|_{\boldsymbol{\Omega}_{ij}}^{2} + \sum_{(i,k)\in\widehat{\mathcal{E}}_{\mathrm{BCG}}} \omega_{ik}\, \rho\Big(\left\| \mathbf{r}_{ik}^{\mathrm{BCG}} \right\|_2^2 \Big). \tag{6}$$

The pose block is $\mathcal{T} = \{\mathbf{T}_i\}_{i\in\mathcal{V}}$. For an intra-agent edge $(i,j) \in \mathcal{E}_{\mathrm{intra}}$, the host SLAM system provides the relative-pose measurement $\mathbf{Z}_{ij} \simeq \mathbf{T}_i^{-1}\mathbf{T}_j$, with residual $\mathbf{r}_{ij}^{\mathrm{pose}} = \mathrm{Log}(\mathbf{Z}_{ij}^{-1}\mathbf{T}_i^{-1}\mathbf{T}_j)$ and information matrix $\boldsymbol{\Omega}_{ij}$.

The residual $\mathbf{r}_{ik}^{\mathrm{BCG}}$ is the local tangent realization of the same center–Bures–Wasserstein discrepancy used to construct the consensus graph, satisfying $\|\mathbf{r}_{ik}^{\mathrm{BCG}}\|_2^2 = d_{\mathrm{BCG}}^2(\widetilde{u}_{ik}, a_k)$. The robust kernel ρ limits the influence of residual geometric inconsistency after factor screening, while ω_{ik} preserves the agent balancing established on the validated support. One submap pose is fixed in each connected component of the active collaborative graph to remove gauge freedom. Equation (6) has the characteristic bipartite structure of a bundle-adjustment-style estimator, with conventional pose nodes connected by the host SLAM factors and additional factors linking those poses to shared latent geometric nodes. We refer to this fixed-support continuous update as Barycentric Bundle Adjustment (BBA).

The discrete graph topology and continuous state are mutually dependent. Updating the submap poses changes the transport costs, while updating the latent geometry changes the reference against which each submap is associated and optimized. BarySLAM therefore alternates BCG graph updates and fixed-support BBA. Each BCG update comprises partial-support consensus construction followed by held-out incidence validation. After every BBA update, the BCG graph is reconstructed at the updated continuous state. Convergence is accepted when the continuous state has stabilized and the validated incidence set remains unchanged.

2.4 CONSENSUS-GUIDED GAUSSIAN CONSOLIDATION

The optimized consensus graph yields refined submap poses together with validated shared support, whereas rendering operates on the original appearance-bearing Gaussian primitives. Consensus-Guided Gaussian Consolidation (CGC) transfers the same validated support back through the local field participation from which the structural observations were constructed, maintaining the connection between collaborative estimation and global map construction. The resulting primitive-level responsibilities are

$$\begin{aligned} s_{ij} &= \boldsymbol{\lambda}_{ij}^{\top}\mathbf{m}_i, \\ \boldsymbol{\pi}_{ij} &= s_{ij}^{-1}\widehat{\mathbf{P}}_i^{\top}\boldsymbol{\lambda}_{ij}, \\ \pi_{ij0} &= 1 - \mathbf{1}_{N_A}^{\top}\boldsymbol{\pi}_{ij}. \end{aligned} \tag{7}$$

The vector $\boldsymbol{\lambda}_{ij} = [\lambda_{ijq}]_q$ contains the participation of primitive g_{ij} in the local structural observations, and s_{ij} is its accumulated structural support. For $s_{ij} > 0$, the validated transport plan $\widehat{\mathbf{P}}_i$ maps this participation back to the consensus elements through $\boldsymbol{\pi}_{ij} = [\pi_{ijk}]_k$. The residual responsibility π_{ij0} is the portion that remains private to the source submap. The source-capacity constraint guarantees $\mathbf{1}_{N_A}^{\top}\boldsymbol{\pi}_{ij} \le 1$ and hence $\pi_{ij0} \ge 0$. When $s_{ij} = 0$, we set $\boldsymbol{\pi}_{ij} = \mathbf{0}$ and $\pi_{ij0} = 1$, so that a primitive unsupported by the structural interface remains unaffected by consensus attenuation.

The BBA-refined pose transforms every retained primitive to the common frame through the standard rigid Gaussian push-forward. At the consolidation stage, BCG enters explicitly through the opacity retained from private and shared support,

$$\widehat{o}_{ij} = \pi_{ij0}o_{ij} + \sum_{k:\,(i,k)\in\widehat{\mathcal{E}}_{\mathrm{BCG}}} \omega_{ik}\pi_{ijk}o_{ij}. \tag{8}$$

Here o_{ij} is the renderer opacity. The private term preserves primitive-specific content, while the shared term is weighted by the validated agent-balanced consensus. The mean and covariance are transformed by the refined pose $\widehat{\mathbf{T}}_i$, with appearance parameters adjusted according to the renderer convention.

Equation (8) thus defines a responsibility-weighted consolidation rule that preserves the shared/private structure in the final Gaussian map. For a fixed optimized state, the consolidation is determined solely by the validated consensus and is invariant to a consistent relabeling of the input submaps.

Algorithm 1 Collaborative Estimation and Map Consolidation

Input: Gaussian submaps $\{\mathcal{G}_i\}_{i\in\mathcal{V}}$, intra-agent factors $\mathcal{E}_{\mathrm{intra}}$, initial poses $\mathcal{T}^0$
Output: Refined poses $\widehat{\mathcal{T}}$, consolidated map $\mathcal{G}_{\mathrm{global}}$

$\mathcal{U} \leftarrow \{u_{iq} \mid i \in \mathcal{V},\ q = 1, \dots, Q_i\}$
$\mathcal{T}^{(0)} \leftarrow \mathcal{T}^0$; $\mathcal{A}^{(0)} \leftarrow \mathcal{A}_{\mathrm{init}}$
$(\widetilde{\mathcal{A}}^{(0)}, \{\widehat{\mathbf{P}}_i^{(0)}\}, \widehat{\mathcal{E}}_{\mathrm{BCG}}^{(0)}) \leftarrow \mathrm{BCG}(\mathcal{U}, \mathcal{T}^{(0)}, \mathcal{A}^{(0)})$
$n_\star \leftarrow N_{\max}$
for $n = 0, \dots, N_{\max} - 1$ **do**
  $(\mathcal{T}^{(n+1)}, \mathcal{A}^{(n+1)}) \leftarrow \mathrm{BBA}(\mathcal{T}^{(n)}, \widetilde{\mathcal{A}}^{(n)}, \mathcal{E}_{\mathrm{intra}}, \widehat{\mathcal{E}}_{\mathrm{BCG}}^{(n)})$
  $(\widetilde{\mathcal{A}}^{(n+1)}, \{\widehat{\mathbf{P}}_i^{(n+1)}\}, \widehat{\mathcal{E}}_{\mathrm{BCG}}^{(n+1)}) \leftarrow \mathrm{BCG}(\mathcal{U}, \mathcal{T}^{(n+1)}, \mathcal{A}^{(n+1)})$
  if $\Delta_{\mathcal{T}}^{(n)} < \epsilon_T$ and $\Delta_{\mathcal{A}}^{(n)} < \epsilon_A$ and $\widehat{\mathcal{E}}_{\mathrm{BCG}}^{(n+1)} = \widehat{\mathcal{E}}_{\mathrm{BCG}}^{(n)}$ **then**
    $n_\star \leftarrow n + 1$; **break**
  end if
end for
$\widehat{\mathcal{T}} \leftarrow \mathcal{T}^{(n_\star)}$
$s_{ij} \leftarrow \boldsymbol{\lambda}_{ij}^{\top}\mathbf{m}_i$
$\boldsymbol{\pi}_{ij} \leftarrow s_{ij}^{-1}(\widehat{\mathbf{P}}_i^{(n_\star)})^{\top}\boldsymbol{\lambda}_{ij}$ if $s_{ij} > 0$; otherwise $\boldsymbol{\pi}_{ij} \leftarrow \mathbf{0}$
$\pi_{ij0} \leftarrow 1 - \mathbf{1}_{N_A}^{\top}\boldsymbol{\pi}_{ij}$
$\mathcal{G}_{\mathrm{global}} \leftarrow \mathrm{CGC}(\{\mathcal{G}_i\}, \widehat{\mathcal{T}}, \{\boldsymbol{\pi}_{ij}, \pi_{ij0}\}, \widehat{\mathcal{E}}_{\mathrm{BCG}}^{(n_\star)})$
return $\widehat{\mathcal{T}}, \mathcal{G}_{\mathrm{global}}$

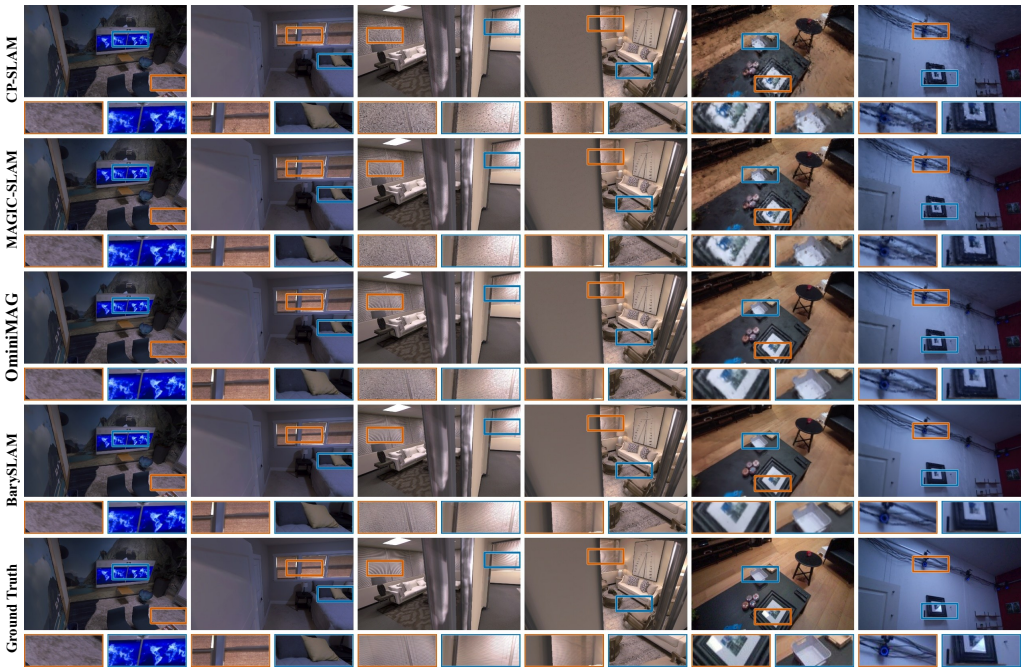


Figure 3: Qualitative reconstruction comparison on ReplicaMultiagent and AriaMultiagent. All methods use identical evaluation viewpoints and crop regions.

3 EXPERIMENTS

We evaluate BarySLAM on ReplicaMultiagent and AriaMultiagent. ReplicaMultiagent contains four collaborative RGB-D scenes derived from Replica (Straub et al., 2019), with ground-truth poses and metric depth for geometric evaluation. AriaMultiagent contains two real-world scenes from Aria Digital Twin (Pan et al., 2023), providing more challenging captured imagery and larger viewpoint variation.

Trajectory accuracy is measured by ATE RMSE in centimeters. All trajectories within a collaborative instance are evaluated in a common frame after a single $\mathrm{SE}(3)$ alignment. Global reconstruction is evaluated using PSNR, SSIM, LPIPS, and L_1 depth error where metric depth is available; training and novel views are reported separately on AriaMultiagent. All BarySLAM variants use identical local Gaussian submaps, intra-agent factors, initialization, optimization budgets, rendering protocols, and evaluation views.

Table 1: Collaborative trajectory estimation on ReplicaMultiagent and AriaMultiagent. ATE RMSE is reported in centimeters (↓). Each entry is averaged over all agents in the corresponding scene.

Method	ReplicaMultiagent					AriaMultiagent		
	Off-0	Apt-0	Apt-1	Apt-2	Avg.	R0	R1	Avg.
Swarm-SLAM	1.42	1.80	5.56	5.61	3.60	6.45	4.78	5.62
MNE-SLAM	0.31	0.48	1.02	0.59	0.60	1.81	0.95	1.38
MAGiC-SLAM	0.27	**0.16**	0.26	0.32	0.25	1.15	0.65	0.90
MultiSlam	0.58	0.47	0.65	0.36	0.52	1.05	0.85	0.95
MAGS-SLAM	0.29	0.46	0.46	0.31	0.38	0.86	**0.54**	0.70
BarySLAM	**0.21**	**0.16**	**0.25**	**0.17**	**0.20**	**0.61**	0.55	**0.58**

Table 2: Training-view reconstruction on ReplicaMultiagent.

Method	Metric	Office-0		Apartment-0		Apartment-1		Apartment-2		Avg.	
CP-SLAM	PSNR/SSIM	28.56	0.870	26.12	0.790	12.16	0.310	23.98	0.810	22.71	0.690
	Depth/LPIPS	02.74	0.290	19.93	0.410	66.77	0.970	02.47	0.390	22.98	0.510
MAGiC-SLAM	PSNR/SSIM	39.32	0.990	36.96	0.980	30.01	0.950	30.73	0.960	34.26	0.970
	Depth/LPIPS	00.41	0.050	00.64	0.090	03.16	0.180	00.99	0.170	01.30	0.120
GRAND-SLAM	PSNR/SSIM	43.12	0.990	44.15	0.990	38.65	**0.990**	39.46	0.990	41.35	0.990
	Depth/LPIPS	00.25	0.030	00.31	0.030	00.77	0.050	00.29	0.050	00.41	0.040
CoMA-SLAM	PSNR/SSIM	42.77	**1.000**	42.28	0.990	33.86	0.980	34.59	0.980	38.38	0.990
	Depth/LPIPS	00.22	**0.020**	00.27	0.030	00.93	0.080	00.48	0.080	00.48	0.050
OminiMAG	PSNR/SSIM	41.21	0.990	39.25	0.990	33.10	0.960	34.21	0.960	36.94	0.980
	Depth/LPIPS	00.32	0.040	00.55	0.070	01.82	0.160	00.99	0.170	00.92	0.110
BarySLAM	PSNR/SSIM	**43.48**	0.995	**44.62**	**0.996**	**39.12**	0.988	**40.03**	**0.991**	**41.81**	**0.993**
	Depth/LPIPS	**00.21**	0.025	**00.26**	**0.024**	**00.63**	**0.043**	**00.24**	**0.041**	**00.34**	**0.033**

3.1 Comparison with the State of the Art

3.1.1 Collaborative Trajectory Estimation

Table 1 shows that BarySLAM reduces the average ATE from the strongest baseline of 0.25 to 0.20 cm on ReplicaMultiagent (20%) and from 0.70 to 0.58 cm on AriaMultiagent (17%). The largest scene-level gains occur on Apartment-2 (0.17 vs. 0.31 cm, 45%) and Room-0 (0.61 vs. 0.86 cm, 29%). BarySLAM ties the best result on Apartment-0 and is only 0.01 cm behind MAGS-SLAM on Room-1, indicating that the average improvement is broad but not uniform across scenes.

3.1.2 Global Reconstruction on ReplicaMultiagent

On ReplicaMultiagent, BarySLAM attains averages of 41.81 dB PSNR, 0.993 SSIM, 0.033 LPIPS, and 0.34 cm Depth L_1 (Table 2). Relative to GRAND-SLAM, the strongest average baseline, this is a 0.46 dB PSNR gain with 17.5% lower LPIPS and 17.1% lower depth error. BarySLAM has the highest PSNR in every scene, while CoMA-SLAM and GRAND-SLAM retain slightly higher SSIM or lower LPIPS in isolated cases; the principal advantage is therefore the joint reduction of photometric, perceptual, and geometric error rather than dominance of every individual entry. Figure 3 provides the corresponding qualitative comparison under matched viewpoints and crops.

3.1.3 Training- and Novel-View Reconstruction on AriaMultiagent

On AriaMultiagent, BarySLAM improves the training-view average over OminiMAG from 25.08 to 26.00 dB PSNR and from 0.920 to 0.940 SSIM, while reducing depth error from 1.61 to 1.35 cm (16.1%) and LPIPS from 0.190 to 0.170 (10.5%). The gains persist on novel views: PSNR increases from 21.85 to 22.65 dB, depth error decreases from 3.50 to 2.87 cm (18.0%), and LPIPS improves from the best baseline value of 0.240 to 0.220 (8.3%). Thus, the reconstruction gain is retained under held-out viewpoints rather than being confined to training-view fitting.

3.2 Ablation Studies

Table 4 isolates BCG, BBA, and CGC on ReplicaMultiagent while keeping all other components fixed. The first two blocks quantify estimation changes; the CGC block fixes the optimized poses, so its identical ATE values isolate map consolidation.

BCG. Replacing Pairwise Reference with a Full-Support Barycenter reduces ATE from 0.62 to 0.41 cm (33.9%), increases PSNR by 0.66 dB, and lowers depth error from 0.86 to 0.73 cm. Allowing partial support further reduces ATE to 0.31 cm (24.4% beyond full support), LPIPS to 0.054, and depth error to 0.67 cm. The latent barycentric reference provides the larger trajectory gain, while partial support adds a consistent benefit by excluding unsupported evidence from the consensus.

Table 3: Training- and novel-view reconstruction on AriaMultiagent.

Method	Metric	Training Views						Novel Views					
		Room-0		Room-1		Avg.		Room-0		Room-1		Avg.	
MAGiC-SLAM	PSNR/SSIM	23.19	0.870	26.16	0.930	24.68	0.900	21.94	0.850	21.59	0.850	21.77	0.850
	Depth/LPIPS	02.03	0.260	01.43	0.170	01.73	0.220	03.02	0.270	05.32	**0.210**	04.17	0.240
OminiMAG	PSNR/SSIM	22.92	0.910	27.23	0.930	25.08	0.920	21.77	0.840	21.93	0.860	21.85	0.850
	Depth/LPIPS	01.90	0.200	01.32	0.170	01.61	0.190	02.85	0.320	04.14	**0.210**	03.50	0.270
BarySLAM	PSNR/SSIM	**24.70**	**0.930**	**27.30**	**0.950**	**26.00**	**0.940**	**22.70**	**0.870**	**22.60**	**0.870**	**22.65**	**0.870**
	Depth/LPIPS	**01.45**	**0.180**	**01.25**	**0.160**	**01.35**	**0.170**	**02.45**	**0.230**	**03.29**	**0.210**	**02.87**	**0.220**

Table 4: Component ablations on ReplicaMultiagent. We report ATE RMSE (cm) and training-view reconstruction metrics. CGC variants share fixed optimized poses.

Component	Variant	ATE↓	PSNR↑	SSIM↑	LPIPS↓	Depth↓
BCG	Pairwise Reference	0.62	38.92	0.971	0.068	0.86
	Full-Support Barycenter	0.41	39.58	0.978	0.059	0.73
	Partial-Support BCG	**0.31**	**39.90**	**0.980**	**0.054**	**0.67**
BBA	Fixed Atlas	0.31	39.90	0.980	0.054	0.67
	Alternating Pose–Barycenter	0.24	40.00	0.981	0.053	0.65
	Joint BBA	**0.20**	**40.08**	**0.981**	**0.052**	**0.63**
CGC	Rigid Union	**0.20**	40.08	0.981	0.052	0.63
	Hard CGC	**0.20**	40.96	0.987	0.041	0.47
	Soft CGC	**0.20**	**41.81**	**0.993**	**0.033**	**0.34**

BBA. Starting from the Fixed Atlas result of 0.31 cm, alternating pose–barycenter updates reduce ATE to 0.24 cm (22.6%), and Joint BBA further reduces it to 0.20 cm (16.7%). Over the same progression, PSNR changes by only 0.18 dB and depth error by 0.04 cm. BBA therefore contributes primarily to trajectory consistency, with a smaller downstream effect on rendering quality.

CGC. Because all CGC variants share the same 0.20 cm ATE, their differences arise solely from consolidation. Hard assignment improves over Rigid Union by 0.88 dB PSNR and reduces LPIPS and depth error by 21.2% and 25.4%, respectively. Soft CGC increases the total gain over Rigid Union to 1.73 dB, raises SSIM from 0.981 to 0.993, and reduces LPIPS from 0.052 to 0.033 (36.5%) and depth error from 0.63 to 0.34 cm (46.0%). Retaining continuous responsibilities is therefore materially better than either rigid concatenation or hard assignment.

4 Discussion and Conclusion

We presented BarySLAM, a collaborative 3D Gaussian SLAM framework that organizes independently reconstructed and partially overlapping Gaussian maps through shared Barycentric Consensus Geometry. BCG establishes spatially supported consensus while preserving agent-specific content, BBA jointly refines submap poses and consensus geometry, and CGC transfers the validated support to the final Gaussian map. Together, these components provide a unified geometric path from cross-agent association to global estimation and reconstruction. Experiments on ReplicaMultiagent and AriaMultiagent demonstrate state-of-the-art overall performance in trajectory estimation and Gaussian reconstruction, improved robustness to limited overlap and corrupted cross-agent registration, and reduced sensitivity to fusion order. The remaining limitation is geometric observability that extremely sparse overlap or persistent ambiguity can weaken cross-agent constraints.

AI USE STATEMENT

In this work, generative AI tools were used solely for language polishing and code optimization. No generative AI tools were involved in the writing of the paper, the formulation of claims, the design of the methodology, or the production of figures or data. We have reviewed all AI-assisted text and code, and we take responsibility for the final content of this work.

ETHICS STATEMENT

This work concerns collaborative SLAM algorithms and does not involve human subjects, personal data, or applications with direct societal harm. All datasets are drawn from existing open benchmarks and follow their original evaluation protocols. We do not foresee any violation of the Code of Ethics arising from this work.

REPRODUCIBILITY STATEMENT

The complete procedures and details needed to reproduce our results are described in the main text and the appendix: the algorithmic pseudocode, the closed-form field encoding, consensus construction, predictive validation, numerical safeguards, and per-agent and controlled experimental settings. All hyperparameters are selected on a disjoint development split and fixed before evaluation. We will release the source code upon request.

ACKNOWLEDGMENTS

We thank the developers and maintainers of the open-source systems and benchmarks that make collaborative SLAM research possible.

REFERENCES

Martial Agueh and Guillaume Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.

Pedro C Álvarez-Esteban, E Del Barrio, JA Cuesta-Albertos, and C Matrán. A fixed-point approach to barycenters in wasserstein space. *Journal of Mathematical Analysis and Applications*, 441(2): 744–762, 2016.

Roy Amoyal, Oren Freifeld, and Chaim Baskin. Cross-instance gaussian splatting registration via geometry-aware feature-guided alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4993–5002, June 2026.

Yiku Bai, Huy Tran, Steven B Damelin, and Soheil Kolouri. Partial transport for point-cloud registration. *Sampling Theory, Signal Processing, and Data Analysis*, 23(1):4, 2025.

Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015. doi: 10.1137/141000439.

Rajendra Bhatia, T. Jain, and Yongdo Lim. On the bures–wasserstein distance between positive definite matrices. *Expositiones Mathematicae*, 37(2):165–191, June 2019. ISSN 0723-0869. doi: 10.1016/j.exmath.2018.01.002. Publisher Copyright: © 2018 Elsevier GmbH.

Jérémie Bigot, Elsa Cazelles, and Nicolas Papadakis. Data-driven regularization of wasserstein barycenters with an application to multivariate density registration. *Information and Inference: A Journal of the IMA*, 8(4):719–755, 12 2019. ISSN 2049-8772. doi: 10.1093/imaiai/iaz023.

Jiahao Chang, Yinglin Xu, Yihao Li, Yuantao Chen, Wensen Feng, and Xiaoguang Han. Gaussreg: Fast 3d registration with gaussian splatting. In *European conference on computer vision*, pp. 407–423. Springer, 2024.

Lin Chen, Yongxin Su, Jvboxi Wang, Pengcheng Han, Zhenyu Xia, Shuhui Bu, Kun Li, Boni Hu, Shengqi Meng, and Guangming Wang. Coma-slam: Collaborative multi-agent gaussian slam with geometric consistency. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40 (4):2922–2929, Mar. 2026. doi: 10.1609/aaai.v40i4.37283. URL https://ojs.aaai.org/index.php/AAAI/article/view/37283.

Chong Cheng, Yu Hu, Sicheng Yu, Beizhen Zhao, Zijian Wang, and Hao Wang. Reggs: Unposed sparse views gaussian splatting with 3dgs registration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 8100–8109, October 2025.

Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Math. Comput.*, 87(314):2563–2609, 2018. doi: 10.1090/MCOM/3303. URL https://doi.org/10.1090/mcom/3303.

Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Neural Information Processing Systems*, pp. 2292–2300, 2013. URL https://mlanthology.org/neurips/2013/cuturi2013neurips-sinkhorn/.

Leqian Ding, Caibo Li, Yu Guo, and Fei Wang. Ominimag-slam : Unified online dual graph optimization for multi-agent gaussian slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Findings*, pp. 1431–1440, June 2026.

Georgios D Evangelidis, Dionyssos Kounades-Bastian, Radu Horaud, and Emmanouil Z Psarakis. A generative model for the joint registration of multiple point sets. In *European conference on computer vision*, pp. 109–122. Springer, 2014.

Brendon Forsgren, Michael Kaess, Ram Vasudevan, Timothy W McLain, and Joshua G Mangelson. Group-k consistent measurement set maximization via maximum clique over k-uniform hypergraphs for robust multi-robot map merging. *The International Journal of Robotics Research*, 43 (14):2245–2273, 2024.

Zan Gojcic, Caifa Zhou, Jan D. Wegner, Leonidas J. Guibas, and Tolga Birdal. Learning multiview 3d point cloud registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

Haobo Jiang, Jin Xie, Jian Yang, Liang Yu, and Jianmin Zheng. Fuser: Feed-forward multiview 3d registration transformer and se(3)$ˆn$ diffusion refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7393–7403, June 2026.

Shengze Jin, Iro Armeni, Marc Pollefeys, and Daniel Barath. Multiway point cloud mosaicking with diffusion and global optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20838–20849, June 2024.

Pierre-Yves Lajoie and Giovanni Beltrame. Swarm-slam: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems. *IEEE Robotics and Automation Letters*, 9(1):475–482, 2024. doi: 10.1109/LRA.2023.3333742.

Hugo Leblond, Gilles Simon, Renato Martins, Cédric Demonceaux, and M-O Berger. Gaussian splatting map registration with orthographic bird's-eye-view renderings. In *2026 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 6080–6089. IEEE, 2026.

Luigi Malagò, Luigi Montrucchio, and Giovanni Pistone. Wasserstein riemannian geometry of gaussian densities. *Information Geometry*, 1(2):137–179, 2018.

Joshua G. Mangelson, Derrick Dominic, Ryan M. Eustice, and Ram Vasudevan. Pairwise consistent measurement set maximization for robust multi-robot map merging. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2916–2923, Brisbane, Australia, May 2018.

Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Yuheng (Carl) Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 20133–20143, October 2023.

Gabriel Peyré and Marco Cuturi. Computational Optimal Transport. *Foundations and Trends in Machine Learning*, 11:355–607, 2019. doi: 10.1561/2200000073. URL https://mlanthology.org/ftml/2019/peyre2019ftml-computational/.

Patrik Schmuck and Margarita Chli. Ccm-slam: Robust and efficient centralized collaborative monocular simultaneous localization and mapping for robotic teams. *Journal of Field Robotics*, 36(4):763–781, 2019.

Olao Shorinwa, Jiankai Sun, Mac Schwager, and Anirudha Majumdar. Siren: Semantic, initialization-free registration of multi-robot gaussian splatting maps. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 3804–3827. PMLR, 27–30 Sep 2025. URL https://proceedings.mlr.press/v305/shorinwa25a.html.

Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.

Annika Thomas, Aneesa Sonawalla, Alex Rose, and Jonathan P. How. Grand-slam: Local optimization for globally consistent large-scale multi-agent gaussian slam. *IEEE Robotics and Automation Letters*, 10(12):13129–13136, 2025. doi: 10.1109/LRA.2025.3623026.

Yulun Tian, Yun Chang, Fernando Herrera Arias, Carlos Nieto-Granda, Jonathan P. How, and Luca Carlone. Kimera-multi: Robust, distributed, dense metric-semantic slam for multi-robot systems. *IEEE Transactions on Robotics*, 38(4):2022–2038, 2022. doi: 10.1109/TRO.2021.3137751.

Xiaohao Xu, Feng Xue, Shibo Zhao, Yike Pan, Sebastian Scherer, and Xiaonan Huang. Mac-ego3d: Multi-agent gaussian consensus for real-time collaborative ego-motion and photorealistic 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 854–863, June 2025.

H. Yang, P. Antonante, V. Tzoumas, and L. Carlone. Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):1127–1134, 2020.

Heng Yang, Jingnan Shi, and Luca Carlone. TEASER: fast and certifiable point cloud registration. *IEEE Trans. Robotics*, 37(2):314–333, 2021. doi: 10.1109/TRO.2020.3033695. URL https://doi.org/10.1109/TRO.2020.3033695.

Vladimir Yugay, Theo Gevers, and Martin R. Oswald. Magic-slam: Multi-agent gaussian globally consistent slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6741–6750, June 2025.

Liyuan Zhu, Yue Li, Erik Sandström, Shengyu Huang, Konrad Schindler, and Iro Armeni. Loopsplat: Loop closure by registering 3d gaussian splats. In *2025 International Conference on 3D Vision (3DV)*, pp. 156–167. IEEE, 2025.

A Related Work

A.1 Collaborative Gaussian SLAM

Collaborative Gaussian SLAM extends multi-agent pose estimation with dense Gaussian scene representations that can be optimized and rendered directly. MAGiC-SLAM (Yugay et al., 2025) combines inter-agent loop closure, relative submap registration, global pose-graph optimization, and Gaussian-map construction in a centralized pipeline. MAC-Ego3D (Xu et al., 2025) introduces Gaussian consensus across agents and couples collaborative ego-motion estimation with a global Gaussian representation. CoMA-SLAM (Chen et al., 2026) studies decentralized consistency in collaborative Gaussian mapping, while OminiMAG (Ding et al., 2026) incorporates server-side global

graph optimization for multi-agent estimation. GRAND-SLAM (Thomas et al., 2025) further addresses globally consistent collaborative Gaussian reconstruction. All these systems organize geometric information reconstructed independently by multiple agents.

Registration of independently reconstructed Gaussian representations has developed in parallel. GaussReg (Chang et al., 2024) performs registration directly on Gaussian splats, LoopSplat (Zhu et al., 2025) uses Gaussian registration for loop closure and trajectory correction, and SIREN (Shorinwa et al., 2025) combines semantic correspondence with geometric and photometric refinement for cross-robot Gaussian-map alignment. RegGS (Cheng et al., 2025) formulates Gaussian-mixture registration through an optimal-transport discrepancy based on the Mixture 2-Wasserstein distance. Related studies have also considered orthographic rendering for Gaussian alignment (Leblond et al., 2026) and geometry-aware registration of independently reconstructed Gaussian instances (Amoyal et al., 2026).

A.2 Multiway Registration and Barycentric Geometry

Multiway registration replaces a collection of isolated pairwise alignments with estimation over several observations in a common reference structure. JR-MPC (Evangelidis et al., 2014) models multiple point sets as rigidly transformed observations of a latent Gaussian mixture and jointly estimates the mixture parameters and rigid transformations. Gojcic et al. (Gojcic et al., 2020) combine pairwise correspondences with globally consistent multiview refinement, while multiway mosaicking (Jin et al., 2024) and FUSER (Jiang et al., 2026) estimate globally consistent poses from partially overlapping scan collections.

Wasserstein geometry provides a related construction when geometric observations are represented as probability distributions. Wasserstein barycenters define Fréchet means under the W_2 metric (Agueh & Carlier, 2011), with fixed-point formulations available for Gaussian and location–scatter families (Álvarez-Esteban et al., 2016). For Gaussian measures, the covariance component of the W_2 distance corresponds to the Bures–Wasserstein geometry on symmetric positive-definite matrices (Malagò et al., 2018; Bhatia et al., 2019). Regularized Wasserstein barycenters have also been used as shared objects for aligning multiple distributions (Bigot et al., 2019). For Gaussian observations, this geometry accounts jointly for spatial displacement and anisotropic covariance structure, which is directly relevant to consensus over independently reconstructed Gaussian representations.

A.3 Partial Correspondence and Robust Estimation

Partial correspondence arises when only part of one geometric observation has valid support in another. Entropic optimal transport provides a soft-coupling formulation for correspondence estimation (Cuturi, 2013; Peyré & Cuturi, 2019), while unbalanced optimal transport relaxes exact marginal preservation and allows unmatched mass to remain outside the correspondence (Chizat et al., 2018). This relaxation has been applied to point-cloud registration under incomplete visibility and missing correspondences (Bai et al., 2025).

Robust estimation addresses a different source of uncertainty after candidate relations have been established. TEASER (Yang et al., 2021) provides certifiable registration under high outlier rates, while graduated non-convexity (Yang et al., 2020) offers a general strategy for robust geometric estimation. In collaborative SLAM, Pairwise Consistent Measurement Set Maximization (Mangelson et al., 2018) filters mutually inconsistent inter-robot measurements before global optimization, and group-k consistency (Forsgren et al., 2024) evaluates consistency over groups of measurements. This motivates the shared-state formulation developed below.

B Field Geometry and Structural Encoding

B.1 Closed-Form Localized Gaussian Moments

For submap i, the geometry and opacity of its Gaussian primitives induce the finite measure

$$\begin{aligned} d\eta_i(\boldsymbol{x}) &= \sum_{j=1}^{M_i} \alpha_{ij}^{\mathrm{fld}} \mathcal{N}\big(\boldsymbol{x}; \boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij}\big)\, d\boldsymbol{x}, \\ \alpha_{ij}^{\mathrm{fld}} &= (2\pi)^{3/2} o_{ij} \det(\boldsymbol{\Sigma}_{ij})^{1/2}. \end{aligned} \tag{9}$$

This measure is used only as a geometric analysis interface and is distinct from the renderer's alpha-compositing model.

Around analysis location $\boldsymbol{c}_{iq}$, we use the isotropic Gaussian window $\kappa_{iq}^h(\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}; \boldsymbol{c}_{iq}, h^2\mathbf{I}_3)$. The product between primitive j and the analysis window is

$$\begin{aligned} &\mathcal{N}\big(\boldsymbol{x}; \boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij}\big)\, \mathcal{N}\big(\boldsymbol{x}; \boldsymbol{c}_{iq}, h^2\mathbf{I}_3\big) \\ &\quad = \mathcal{N}\big(\boldsymbol{c}_{iq}; \boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij} + h^2\mathbf{I}_3\big)\, \mathcal{N}\Big(\boldsymbol{x}; \boldsymbol{\mu}_{ijq}^h, \boldsymbol{\Sigma}_{ijq}^h\Big), \end{aligned} \tag{10}$$

where

$$\begin{aligned} \boldsymbol{\Sigma}_{ijq}^h &= \big(\boldsymbol{\Sigma}_{ij}^{-1} + h^{-2}\mathbf{I}_3\big)^{-1}, \\ \boldsymbol{\mu}_{ijq}^h &= \boldsymbol{\Sigma}_{ijq}^h \big(\boldsymbol{\Sigma}_{ij}^{-1}\boldsymbol{\mu}_{ij} + h^{-2}\boldsymbol{c}_{iq}\big). \end{aligned} \tag{11}$$

The unnormalized localized response of primitive j is

$$\xi_{ijq} = \alpha_{ij}^{\mathrm{fld}} \mathcal{N}\big(\boldsymbol{c}_{iq}; \boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij} + h^2\mathbf{I}_3\big). \tag{12}$$

The source mass and normalized primitive participation are

$$m_{iq} = \sum_j \xi_{ijq}, \qquad \lambda_{ijq} = \frac{\xi_{ijq}}{m_{iq}}, \qquad \sum_j \lambda_{ijq} = 1. \tag{13}$$

For retained sites with $m_{iq} > 0$, the structural observation $u_{iq} = (\boldsymbol{\mu}_{iq}, \boldsymbol{\Sigma}_{iq})$ used in Sec. 2.2 is

$$\begin{aligned} \boldsymbol{\mu}_{iq} &= \sum_j \lambda_{ijq} \boldsymbol{\mu}_{ijq}^h, \\ \boldsymbol{\Sigma}_{iq} &= \sum_j \lambda_{ijq} \Big[\boldsymbol{\Sigma}_{ijq}^h + \Delta\boldsymbol{\mu}_{ijq}^h (\Delta\boldsymbol{\mu}_{ijq}^h)^\top\Big], \end{aligned} \tag{14}$$

where $\Delta\boldsymbol{\mu}_{ijq}^h = \boldsymbol{\mu}_{ijq}^h - \boldsymbol{\mu}_{iq}$. The construction requires neither voxelization nor sampling of the induced field. The same coefficients λ_{ijq} are retained for the responsibility transfer used by CGC. Since each localized component covariance is positive definite and at least one component has positive weight, the resulting covariance remains in $\mathbb{S}_{++}^3$.

B.2 Rigid-Motion Equivariance

Proposition B.1 (Rigid-motion equivariance)**.** *Let $g = (\mathbf{R}, \mathbf{t}) \in \mathrm{SE}(3)$ transform a Gaussian map by $\boldsymbol{\mu}_j \mapsto \mathbf{R}\boldsymbol{\mu}_j + \mathbf{t}$ and $\boldsymbol{\Sigma}_j \mapsto \mathbf{R}\boldsymbol{\Sigma}_j\mathbf{R}^\top$, with the analysis location transformed by the same rigid motion. The localized structural observation satisfies*

$$(\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q) \mapsto \big(\mathbf{R}\boldsymbol{\mu}_q + \mathbf{t}, \mathbf{R}\boldsymbol{\Sigma}_q\mathbf{R}^\top\big). \tag{15}$$

Proof. The rigid change of variables preserves the Gaussian Mahalanobis exponent and covariance determinant, so the induced measure is transformed by push-forward. Because the analysis window is isotropic, its value is unchanged under the corresponding rigid change of coordinates. The localized mean therefore transforms affinely, and centering removes translation from the second moment, yielding Eq. (15). □

C BCG Construction and Partial Support

C.1 Consensus Atlas Initialization

Reindexing may compact the storage indices after pruning, while a persistent latent-node identifier is retained across outer epochs for support and convergence comparisons. Atlas initialization uses the coarse common-frame pose block $\mathcal{T}^0$ supplied by the upstream cross-agent stage. Each structural observation is transformed as

$$\widetilde{\boldsymbol{\mu}}_{iq}^0 = \mathbf{R}_i^0 \boldsymbol{\mu}_{iq} + \mathbf{t}_i^0, \qquad \widetilde{\boldsymbol{\Sigma}}_{iq}^0 = \mathbf{R}_i^0 \boldsymbol{\Sigma}_{iq} (\mathbf{R}_i^0)^\top. \tag{16}$$

The transformed centers are pooled and thinned deterministically at the analysis scale h, without selecting a reference agent. Each retained seed initializes one atlas element, and its covariance is initialized from the BW mean of covariances assigned to the same seed neighborhood. Elements that fail the cross-agent support criterion after the first BCG update are removed. Active atlas elements retain their latent-node identity through an outer epoch and are reindexed after pruning. Accordingly, N_A denotes the cardinality of the active atlas in the current epoch.

The auxiliary atlas-side mass $\mathbf{m}_A^{(0)}$ is initialized from the same deterministic seeding. It is used only by the transport solver and is not part of the geometric BCA state.

C.2 Source-Capacitated Unbalanced Transport

The atlas-side mass $\mathbf{m}_A$ is fixed within one transport solve. Let

$$\mathbf{K}_i = \exp(-\mathbf{C}_i/\varepsilon) \tag{17}$$

elementwise. For nonnegative vectors or matrices, the generalized KL divergence is

$$D_{\mathrm{KL}}(\mathbf{a}\|\mathbf{b}) = \sum_j \left[a_j \log \frac{a_j}{b_j} - a_j + b_j \right], \tag{18}$$

with its standard continuous extension at zero.

Up to terms independent of $\mathbf{P}$, Eq. (3) is equivalent to

$$\begin{aligned} \mathbf{P}_i^\star = \underset{\mathbf{P}\geq 0,\ \mathbf{P}\mathbf{1}_{N_A} \leq \mathbf{m}_i}{\arg\min} \quad & \varepsilon D_{\mathrm{KL}}(\mathbf{P}\|\mathbf{K}_i) + \tau_s D_{\mathrm{KL}}(\mathbf{P}\mathbf{1}_{N_A}\|\mathbf{m}_i) \\ & + \tau_a D_{\mathrm{KL}}\big(\mathbf{P}^\top \mathbf{1}_{Q_i} \big\| \mathbf{m}_A\big)\,. \end{aligned} \tag{19}$$

We solve this convex subproblem using generalized Sinkhorn proximal steps for the relaxed marginals and an interleaved KL projection for the source-capacity constraint within a Bregman–Dykstra scheme (Benamou et al., 2015; Chizat et al., 2018). For a positive provisional matrix $\mathbf{Q}$, the capacity projection is row-wise,

$$\big[\mathrm{Proj}_{\mathrm{cap}}^{\mathrm{KL}}(\mathbf{Q})\big]_{qk} = Q_{qk} \min\left(1, \frac{m_{iq}}{\sum_{\ell=1}^{N_A} Q_{q\ell}}\right), \tag{20}$$

with the factor set to one for a zero row. Because this projection is part of the iterative constrained solve rather than a post-hoc modification, the converged coupling satisfies the source-capacity constraint in Eq. (3). One complete scaling/projection sweep costs $O(Q_i N_A)$.

C.3 Support Selection and Within-Submap Reduction

The transported support from submap i to atlas element k is

$$m_{ik}^{\mathrm{tr}} = \sum_{q=1}^{Q_i} P_{i,qk}^\star. \tag{21}$$

The provisional support and participating-agent sets are

$$\begin{aligned} \mathcal{S}_k &= \big\{ i \in \mathcal{V} \;\big|\; m_{ik}^{\mathrm{tr}} > \tau_{\mathrm{supp}} \big\}, \\ \mathcal{R}_k &= \{ \mathrm{ag}(i) \mid i \in \mathcal{S}_k \}, \qquad |\mathcal{R}_k| \geq 2. \end{aligned} \tag{22}$$

Here $\mathrm{ag}(i)$ denotes the agent that generated submap i, and $\mathcal{V}_r = \{i \in \mathcal{V} \mid \mathrm{ag}(i) = r\}$ denotes the submaps belonging to agent r. The threshold suppresses the positive numerical tails produced by entropic regularization. An atlas element with fewer than two distinct supporting agents is removed from the collaborative graph.

For $i \in \mathcal{S}_k$, define

$$\alpha_{iqk} = \frac{P^{\star}_{i,qk}}{m^{\mathrm{tr}}_{ik}}. \tag{23}$$

The corresponding transport-reduced structural observation is

$$\begin{aligned} \boldsymbol{\mu}_{ik} &= \sum_q \alpha_{iqk} \boldsymbol{\mu}_{iq}, \\ \boldsymbol{\Sigma}_{ik} &= \sum_q \alpha_{iqk} \left[\boldsymbol{\Sigma}_{iq} + (\boldsymbol{\mu}_{iq} - \boldsymbol{\mu}_{ik})(\boldsymbol{\mu}_{iq} - \boldsymbol{\mu}_{ik})^{\top} \right]. \end{aligned} \tag{24}$$

The positive between-component term preserves the complete second moment of the transported local structural mixture.

Several submaps from one agent may support the same atlas element. Before validation, equal total influence is assigned to each participating agent and distributed among that agent's submaps according to transported support,

$$\omega^{\mathrm{pre}}_{ik} = \frac{1}{|\mathcal{R}_k|} \frac{m^{\mathrm{tr}}_{ik}}{\sum_{t \in \mathcal{S}_k :\, \mathrm{ag}(t) = \mathrm{ag}(i)} m^{\mathrm{tr}}_{tk}}. \tag{25}$$

Consequently,

$$\sum_{i \in \mathcal{S}_k} \omega^{\mathrm{pre}}_{ik} = 1, \qquad \sum_{i \in \mathcal{S}_{k,r}} \omega^{\mathrm{pre}}_{ik} = \frac{1}{|\mathcal{R}_k|}, \tag{26}$$

where $\mathcal{S}_{k,r} = \{i \in \mathcal{S}_k : \mathrm{ag}(i) = r\}$. The main text suppresses the superscript when the stage is unambiguous.

C.4 Bures–Wasserstein Barycenter Updates

For $\{\mathbf{C}_j\}_{j=1}^N \subset \mathbb{S}^3_{++}$ and normalized nonnegative weights a_j, the BW Fréchet mean is

$$\mathbf{S}^{\star} = \arg\min_{\mathbf{S} \in \mathbb{S}^3_{++}} \sum_{j=1}^N a_j d^2_{\mathrm{BW}}(\mathbf{C}_j, \mathbf{S}). \tag{27}$$

Starting from $\mathbf{S}^{(0)} \succ 0$, a standard fixed-point iteration is

$$\mathbf{C}^{(\ell)} = \sum_{j=1}^N a_j \left[(\mathbf{S}^{(\ell)})^{1/2} \mathbf{C}_j (\mathbf{S}^{(\ell)})^{1/2} \right]^{1/2}, \tag{28}$$

followed by

$$\mathbf{S}^{(\ell+1)} = (\mathbf{S}^{(\ell)})^{-1/2} (\mathbf{C}^{(\ell)})^2 (\mathbf{S}^{(\ell)})^{-1/2}. \tag{29}$$

These equations give the standard fixed-point iteration for Gaussian Wasserstein barycenters (Agueh & Carlier, 2011; Álvarez-Esteban et al., 2016).

Writing the provisional element in Eq. (4) as $\widetilde{a}_k = (\widetilde{\boldsymbol{\mu}}_k, \widetilde{\boldsymbol{\Sigma}}_k)$, its two components are

$$\begin{aligned} \widetilde{\boldsymbol{\mu}}_k &= \sum_{i \in \mathcal{S}_k} \omega^{\mathrm{pre}}_{ik} \widetilde{\boldsymbol{\mu}}_{ik}, \\ \widetilde{\boldsymbol{\Sigma}}_k &= \arg\min_{\boldsymbol{\Sigma} \in \mathbb{S}^3_{++}} \sum_{i \in \mathcal{S}_k} \omega^{\mathrm{pre}}_{ik} d^2_{\mathrm{BW}}\left(\widetilde{\boldsymbol{\Sigma}}_{ik}, \boldsymbol{\Sigma} \right). \end{aligned} \tag{30}$$

The fixed positive scales σ_μ and σ_Σ in Eq. (2) multiply the two separable component objectives by constants and do not alter their minimizers.

Collecting the provisional elements gives

$$\widetilde{\mathcal{A}} = \{\widetilde{a}_k\}_{k=1}^{N_A}, \tag{31}$$

which initializes the latent geometry of the subsequent fixed-support BBA update.

D JOINT ESTIMATION AND PREDICTIVE VALIDATION

D.1 BURES–WASSERSTEIN TANGENT GEOMETRY

For $\mathbf{S} \in \mathbb{S}^3_{++}$ and a symmetric matrix $\mathbf{X}$, let $\mathcal{L}_{\mathbf{S}}(\mathbf{X})$ denote the unique symmetric solution of

$$\mathcal{L}_{\mathbf{S}}(\mathbf{X})\mathbf{S} + \mathbf{S}\mathcal{L}_{\mathbf{S}}(\mathbf{X}) = \mathbf{X}. \tag{32}$$

The BW tangent inner product is

$$\langle \mathbf{X}, \mathbf{Y} \rangle^{\mathrm{BW}}_{\mathbf{S}} = \frac{1}{2} \operatorname{tr}[\mathcal{L}_{\mathbf{S}}(\mathbf{X})\mathbf{Y}]\,. \tag{33}$$

For $\mathbf{C} \in \mathbb{S}^3_{++}$, define

$$\mathbf{A}_{\mathbf{S}\to\mathbf{C}} = \mathbf{S}^{-1/2}\left(\mathbf{S}^{1/2}\mathbf{C}\mathbf{S}^{1/2}\right)^{1/2}\mathbf{S}^{-1/2} \tag{34}$$

and

$$\mathrm{Log}^{\mathrm{BW}}_{\mathbf{S}}(\mathbf{C}) = (\mathbf{A}_{\mathbf{S}\to\mathbf{C}} - \mathbf{I}_3)\mathbf{S} + \mathbf{S}(\mathbf{A}_{\mathbf{S}\to\mathbf{C}} - \mathbf{I}_3). \tag{35}$$

This convention satisfies

$$\left\|\mathrm{Log}^{\mathrm{BW}}_{\mathbf{S}}(\mathbf{C})\right\|^2_{\mathbf{S},\mathrm{BW}} = d^2_{\mathrm{BW}}(\mathbf{S}, \mathbf{C}). \tag{36}$$

For a symmetric matrix, we use the Frobenius-isometric vectorization

$$\operatorname{svec}(\mathbf{X}) = \left[\, X_{11}\ X_{22}\ X_{33}\ \sqrt{2}X_{12}\ \sqrt{2}X_{13}\ \sqrt{2}X_{23}\,\right]^\top. \tag{37}$$

Let $\mathbf{G}_{\mathbf{S}}$ satisfy

$$\operatorname{svec}(\mathbf{X})^\top \mathbf{G}_{\mathbf{S}} \operatorname{svec}(\mathbf{Y}) = \langle \mathbf{X}, \mathbf{Y} \rangle^{\mathrm{BW}}_{\mathbf{S}}. \tag{38}$$

If $\mathbf{S} = \mathbf{U}\operatorname{diag}(s_1, s_2, s_3)\mathbf{U}^\top$, then in the eigenbasis of $\mathbf{S}$ the metric matrix is diagonal with entries

$$\operatorname{diag}\left(\frac{1}{4s_1}, \frac{1}{4s_2}, \frac{1}{4s_3}, \frac{1}{2(s_1+s_2)}, \frac{1}{2(s_1+s_3)}, \frac{1}{2(s_2+s_3)}\right). \tag{39}$$

For an active incidence (i,k), choose $\mathbf{L}_k$ such that $\mathbf{L}_k^\top \mathbf{L}_k = \mathbf{G}_{\bar{\mathbf{\Sigma}}_k}$. The 9-dimensional BCG residual is

$$\begin{aligned} \mathbf{r}^{\mathrm{BCG}}_{ik} &= \begin{bmatrix} \sigma_\mu^{-1}(\widetilde{\boldsymbol{\mu}}_{ik} - \bar{\boldsymbol{\mu}}_k) \\ \sigma_\Sigma^{-1}\mathbf{L}_k \operatorname{svec}\left[\mathrm{Log}^{\mathrm{BW}}_{\bar{\mathbf{\Sigma}}_k}(\widetilde{\mathbf{\Sigma}}_{ik})\right] \end{bmatrix}, \\ \|\mathbf{r}^{\mathrm{BCG}}_{ik}\|_2^2 &= d^2_{\mathrm{BCG}}(\widetilde{u}_{ik}, a_k). \end{aligned} \tag{40}$$

Thus association, barycentric consensus, predictive validation, and BBA share the same center–Bures–Wasserstein residual geometry.

D.2 BBA LINEARIZATION

We use a left perturbation of submap pose $\mathbf{T}_i$,

$$\mathbf{T}_i(\delta\boldsymbol{\xi}_i) = \exp(\delta\boldsymbol{\xi}_i^\wedge)\mathbf{T}_i, \qquad \delta\boldsymbol{\xi}_i = \begin{bmatrix} \delta\boldsymbol{v}_i \\ \delta\boldsymbol{\omega}_i \end{bmatrix}. \tag{41}$$

For the transformed reduced observation $(\widetilde{\boldsymbol{\mu}}_{ik}, \widetilde{\mathbf{\Sigma}}_{ik})$, the first-order pose variations are

$$\delta\widetilde{\boldsymbol{\mu}}_{ik} = \delta\boldsymbol{v}_i - [\widetilde{\boldsymbol{\mu}}_{ik}]_\times \delta\boldsymbol{\omega}_i \tag{42}$$

and

$$\delta\widetilde{\mathbf{\Sigma}}_{ik} = [\delta\boldsymbol{\omega}_i]_\times \widetilde{\mathbf{\Sigma}}_{ik} - \widetilde{\mathbf{\Sigma}}_{ik}[\delta\boldsymbol{\omega}_i]_\times. \tag{43}$$

Derivatives of matrix square roots are evaluated through the Sylvester equation

$$\mathbf{A}^{1/2}\mathbf{X} + \mathbf{X}\mathbf{A}^{1/2} = \mathbf{E}, \qquad \mathbf{X} = D(\mathbf{A}^{1/2})[\mathbf{E}]. \tag{44}$$

Composing Eq. (44) with Eqs. (34)–(35) gives the covariance-residual Jacobians. The atlas covariance is updated in a local SPD chart; this parameterization affects only the numerical retraction and not the BW residual itself.

After robust reweighting, a fixed-support linearization gives

$$(\mathbf{H} + \lambda_{\mathrm{GN}}\mathbf{I})\,\delta\boldsymbol{\theta} = -\mathbf{g}, \tag{45}$$

where $\mathbf{H}$ and $\mathbf{g}$ are the robustified Gauss–Newton normal matrix and gradient vector, $\lambda_{\mathrm{GN}} \geq 0$ is the damping coefficient, and $\delta\boldsymbol{\theta}$ collects the active submap-pose and atlas increments. The intra-agent edges $\mathcal{E}_{\mathrm{intra}}$ and validated BCG incidences $\widehat{\mathcal{E}}_{\mathrm{BCG}}$ determine the sparse normal-equation structure. Transport plans and validated support are held fixed during this continuous solve. One pose is fixed in each connected component of the active factor graph to remove gauge freedom.

D.3 Conditional Held-Out-Agent Predictive Information

Fix an active incidence (i, k) with $r = \mathrm{ag}(i)$ and define

$$\mathcal{S}_k^{-r} = \mathcal{S}_k \setminus \mathcal{V}_r. \tag{46}$$

The held-out agent-balanced weights are recomputed and normalized over $\mathcal{S}_k^{-r}$; denote them by ω_{tk}^{-r}. The held-out anchor is an auxiliary validation state. It is therefore reconstructed from the remaining support available for the test and is not itself required to satisfy the multi-agent activation criterion used for persistent BCG nodes.

The held-out anchor a_k^{-r} is reconstructed using the same barycentric objective as Eq. (4). Let $\delta\mathbf{a}_k \in \mathbb{R}^9$ be a local chart of the held-out anchor and let $\mathbf{H}_{t,a}$ be the Jacobian of the whitened BCG residual with respect to this chart. Its local information is approximated by

$$\mathbf{I}_{a,k}^{-r} = \sum_{t \in \mathcal{S}_k^{-r}} \omega_{tk}^{-r}\mathbf{H}_{t,a}^{\top}\mathbf{H}_{t,a}. \tag{47}$$

On the locally observable subspace,

$$\boldsymbol{\Sigma}_{a,k}^{-r} \approx (\mathbf{I}_{a,k}^{-r})^{\dagger}. \tag{48}$$

Let $\boldsymbol{\Sigma}_{ik}^{\mathrm{res}}$ denote the residual covariance assigned to the tested structural relation and $\mathbf{H}_a$ its derivative with respect to the held-out anchor. First-order uncertainty propagation gives

$$\begin{aligned} \boldsymbol{\Sigma}_{ik}^{\mathrm{pred}} &= \boldsymbol{\Sigma}_{ik}^{\mathrm{res}} + \mathbf{H}_a\boldsymbol{\Sigma}_{a,k}^{-r}\mathbf{H}_a^{\top}, \\ \boldsymbol{\Omega}_{ik}^{\mathrm{pred}} &= (\boldsymbol{\Sigma}_{ik}^{\mathrm{pred}})^{\dagger}. \end{aligned} \tag{49}$$

The held-out residual $\mathbf{r}_{ik}^{-r}$ is obtained from Eq. (40) by replacing a_k with a_k^{-r}. The resulting covariance and score are conditional on the current pose, candidate, and transport state. The resulting quantity is used only for conditional geometric consistency testing within the current outer epoch.

D.4 Validation Mask and Agent Rebalancing

The consistency gate in Eq. (5) defines the retained support $\widehat{\mathcal{S}}_k$. The provisional transport plan is masked consistently with this support,

$$\widehat{P}_{i,qk} = \begin{cases} P^{\star}_{i,qk}, & i \in \widehat{\mathcal{S}}_k, \\ 0, & i \notin \widehat{\mathcal{S}}_k. \end{cases} \tag{50}$$

If fewer than two distinct agents remain in $\widehat{\mathcal{S}}_k$, the atlas element is removed and the corresponding column is discarded.

The validated transported support is

$$\widehat{m}_{ik}^{\mathrm{tr}} = \sum_q \widehat{P}_{i,qk}. \tag{51}$$

Let

$$\widehat{\mathcal{R}}_k = \left\{ \mathrm{ag}(i) \;\middle|\; i \in \widehat{\mathcal{S}}_k \right\}. \tag{52}$$

The agent-balanced weight is recomputed on the retained support,

$$\widehat{\omega}_{ik} = \frac{1}{|\widehat{\mathcal{R}}_k|} \frac{\widehat{m}_{ik}^{\mathrm{tr}}}{\displaystyle\sum_{t \in \widehat{\mathcal{S}}_k:\, \mathrm{ag}(t) = \mathrm{ag}(i)} \widehat{m}_{tk}^{\mathrm{tr}}}. \tag{53}$$

Hence

$$\sum_{i \in \widehat{\mathcal{S}}_k} \widehat{\omega}_{ik} = 1. \tag{54}$$

After validation, the main text overwrites $\omega_{ik} \leftarrow \widehat{\omega}_{ik}$, so all subsequent BBA and CGC expressions use weights normalized over the validated support.

D.5 Validation Calibration

The score in Eq. (5) is used as a predictive geometric consistency gate. The threshold is selected on a disjoint development split from a calibration set $\mathcal{C}_{\mathrm{val}}$ of geometrically compatible incidences,

$$\tau_{\mathrm{val}} = Q_{1-\alpha_{\mathrm{val}}}(\{e_{ik} : (i,k) \in \mathcal{C}_{\mathrm{val}}\}), \tag{55}$$

where $Q_{1-\alpha}$ denotes the empirical quantile. Under a locally Gaussian residual model with correctly calibrated covariance, the corresponding Mahalanobis score has the usual chi-square interpretation on the observable residual subspace. This relation is used only as a diagnostic reference. No post-selection, distribution-free, or independent-validation significance guarantee is claimed.

D.6 Atlas-Side Auxiliary Mass

In Alg. 1, $\mathbf{m}_A$ is carried internally by the BCG update. It is an auxiliary transport-solver state rather than a geometric BCA variable and is therefore omitted from the external BCG interface.

The atlas-side mass is auxiliary to the transport solver and is not part of the BCA geometric state. After validation, it is updated from the retained transported support,

$$m_{A,k}^{(n+1)} = \frac{1}{|\widehat{\mathcal{S}}_k^{(n)}|} \sum_{i \in \widehat{\mathcal{S}}_k^{(n)}} \widehat{m}_{ik}^{\mathrm{tr},(n)}. \tag{56}$$

Elements with insufficient validated cross-agent support are removed before the next BCG update. Updating $\mathbf{m}_A$ only between outer epochs prevents the transport objective from depending implicitly on the coupling currently being solved and prevents rejected incidences from influencing the next atlas-occupancy reference.

E Consensus-Guided Gaussian Consolidation and System Details

E.1 Primitive-to-Observation Participation

The primitive participation used in Eq. (7) is the localized response from Eq. (13),

$$\lambda_{ijq} = \frac{\alpha_{ij}^{\mathrm{fld}} \mathcal{N}\big(\boldsymbol{c}_{iq}; \boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij} + h^2 \mathbf{I}_3\big)}{\sum_t \alpha_{it}^{\mathrm{fld}} \mathcal{N}\big(\boldsymbol{c}_{iq}; \boldsymbol{\mu}_{it}, \boldsymbol{\Sigma}_{it} + h^2 \mathbf{I}_3\big)}. \tag{57}$$

It relates a primitive only to structural observations derived from the same submap. Writing $\boldsymbol{\lambda}_{ij} = [\lambda_{ijq}]_q$ makes its composition with the validated transport plan explicit.

E.2 Responsibility and Opacity Bounds

For $s_{ij} = \boldsymbol{\lambda}_{ij}^\top \mathbf{m}_i > 0$, Eq. (7) gives

$$\mathbf{1}_{N_A}^\top \boldsymbol{\pi}_{ij} = \frac{\boldsymbol{\lambda}_{ij}^\top \widehat{\mathbf{P}}_i \mathbf{1}_{N_A}}{s_{ij}}. \tag{58}$$

The source-capacity constraint and validation mask imply

$$\mathbf{0} \le \widehat{\mathbf{P}}_i \mathbf{1}_{N_A} \le \mathbf{m}_i. \tag{59}$$

Hence

$$0 \le \mathbf{1}_{N_A}^\top \boldsymbol{\pi}_{ij} \le 1, \qquad 0 \le \pi_{ij0} \le 1. \tag{60}$$

When $s_{ij} = 0$, the convention $\boldsymbol{\pi}_{ij} = \mathbf{0}$ and $\pi_{ij0} = 1$ satisfies the same bounds.

For the effective renderer opacity $o_{ij} \in [0, 1]$, and because $0 \le \omega_{ik} \le 1$ and all responsibilities are nonnegative, Eq. (8) satisfies

$$\begin{aligned}\widehat{o}_{ij} &= o_{ij}\left[\pi_{ij0} + \sum_{k:\,(i,k)\in\widehat{\mathcal{E}}_{\mathrm{BCG}}} \omega_{ik}\pi_{ijk}\right] \\ &\le o_{ij}\left[\pi_{ij0} + \mathbf{1}_{N_A}^{\top}\boldsymbol{\pi}_{ij}\right] = o_{ij}. \end{aligned} \tag{61}$$

Therefore,

$$0 \le \widehat{o}_{ij} \le o_{ij}. \tag{62}$$

CGC can retain or attenuate the effective opacity contribution of a primitive but cannot amplify it beyond the original value. This bound does not imply exact radiance or alpha-compositing conservation.

E.3 Rigid Push-Forward and Appearance Transformation

Let $\gamma_{ij} = \mathcal{N}(\boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij})$. The optimized submap pose induces the standard rigid Gaussian push-forward

$$(\widehat{\mathbf{T}}_i)_{\#}\gamma_{ij} = \mathcal{N}\Big(\widehat{\mathbf{R}}_i\boldsymbol{\mu}_{ij} + \widehat{\mathbf{t}}_i, \widehat{\mathbf{R}}_i\boldsymbol{\Sigma}_{ij}\widehat{\mathbf{R}}_i^{\top}\Big). \tag{63}$$

View-independent appearance is carried forward unchanged. If the local renderer uses directional appearance such as spherical harmonics, the optimized rotation requires the renderer-consistent basis transformation

$$\widehat{\mathbf{c}}_{ij} = \mathcal{R}_{\mathrm{app}}(\widehat{\mathbf{R}}_i)\mathbf{c}_{ij}. \tag{64}$$

For view-independent RGB, $\mathcal{R}_{\mathrm{app}}(\mathbf{R}) = \mathbf{I}$. For spherical harmonics, $\mathcal{R}_{\mathrm{app}}$ is the block-diagonal real-SH rotation consistent with the renderer's direction convention, coefficient ordering, and basis normalization. CGC does not average appearance coefficients across agents.

E.4 Global Gaussian Map Construction

Let $\widehat{g}_{ij}$ denote the primitive obtained by applying Eq. (63), Eq. (8), and the corresponding appearance transformation to g_{ij}. The global renderable map is the indexed multiset union

$$\mathcal{G}_{\mathrm{global}} = \biguplus_{i\in\mathcal{V}} \{\widehat{g}_{ij}\}_{j=1}^{M_i}. \tag{65}$$

The multiset construction preserves the original primitive records and does not require a one-to-one primitive correspondence across independently optimized Gaussian maps. Shared content is regulated through the validated consensus responsibilities rather than through explicit primitive merging.

E.5 Integration with a Gaussian SLAM System

BarySLAM receives the locally reconstructed Gaussian submaps $\{\mathcal{G}_i\}$, the intra-agent factor graph $\mathcal{E}_{\mathrm{intra}}$, and the coarse common-frame initialization $\mathcal{T}^0$. These quantities provide the geometric and pose information used to construct the BCG state and initialize collaborative estimation.

The resulting optimized submap poses $\widehat{\mathcal{T}}$ and consolidated global Gaussian map $\mathcal{G}_{\mathrm{global}}$ are returned as summarized in Alg. 1. The formulation does not prescribe a particular implementation for local Gaussian reconstruction or inter-agent retrieval.

E.6 Computational Complexity

For Q_i structural observations in submap i and N_A active atlas elements, one constrained transport scaling/projection sweep costs

$$O(Q_i N_A). \tag{66}$$

The transport problems are independent across submaps. All BW covariance operations involve only 3×3 SPD matrices. With support fixed, each active BCG incidence contributes a 9-dimensional

local geometric residual, and the normal equations preserve the sparsity of the bipartite submap–consensus graph.

For atlas element k, held-out validation requires at most $|\mathcal{R}_k|$ distinct anchor reconstructions because all submaps from the same tested agent share the same held-out anchor. Primitive responsibility transfer is linear in the retained primitive-to-observation participation records. The final CGC step preserves the primitive count of the contributing submaps unless a separate renderer-side pruning policy is applied.

E.7 Numerical Safeguards

Before matrix square roots, eigendecompositions, or Sylvester solves, covariance matrices are symmetrized and a fixed eigenvalue floor is applied only to preserve positive definiteness in finite precision. Constrained UOT and BW barycenter iterations use fixed convergence tolerances and maximum iteration counts.

During one BBA solve, validated support sets and transport plans remain fixed. A single submap pose is anchored in each connected component of the active factor graph to remove gauge freedom. The discrete BCG graph is updated only between continuous BBA epochs.

The analysis bandwidth, structural scales, UOT parameters, support threshold τ_{supp}, validation calibration level, robust-loss parameters, solver tolerances, and iteration limits are selected on the development split and fixed before evaluation.

F Experimental Protocols and Reproducibility

This appendix records the procedures used in Sec. 3, separating benchmark-defined quantities from choices specific to BarySLAM.

F.1 Benchmark Protocols and Numerical Provenance

A published baseline is included only when its originating evaluation block is identified. ReplicaMultiagent follows the GRAND-SLAM protocol (Thomas et al., 2025), while AriaMultiagent follows OminiMAG (Ding et al., 2026).

BarySLAM uses the scene split and metric implementation associated with each source-defined block. If the required view list or evaluation script cannot be reproduced unambiguously, the BarySLAM entry is left unreported. A value from a different paper is never substituted solely because it uses the same dataset name.

Published numbers support descriptive benchmark comparison only. Statistical tests are restricted to methods rerun under the same local implementation and evaluation protocol.

F.2 Trajectory and Rendering Evaluation

For each collaborative scene, the source-defined per-trajectory ATE RMSE values are retained together with the reported scene statistic. When the source table provides the scene statistic directly, we retain it rather than recomputing the value from rounded trajectory entries.

For AriaMultiagent and the controlled experiments, all trajectories belonging to one collaborative instance are kept in a common frame and evaluated after a single $\mathrm{SE}(3)$ alignment of the complete set. Independent trajectory-wise alignments are excluded because they would remove cross-agent inconsistency from the reported error.

Rendering metrics follow the view category specified by the source protocol, with training and novel views treated as separate populations. PSNR and SSIM use the RGB renderings, LPIPS uses the same image pairs, and Depth L_1 is computed in metric depth before conversion to centimeters for reporting.

Sequence names, frame ranges, and evaluation-view lists are stored in the rerun configuration and frozen before the final BarySLAM results are generated.

F.3 Controlled Cross-Agent Overlap

The low-overlap experiment requires an overlap measure independent of the retrieval descriptor. We compute it directly from ground-truth geometry.

Let $\mathcal{V}_s$ denote the set of ground-truth surface voxels observed by submap s. The set is obtained by back-projecting valid ground-truth depth samples into the global frame and quantizing them using a fixed voxel width r_{v}. The directed normalized overlap between submaps s and t is symmetrized as

$$O_{st} = \frac{|\mathcal{V}_s \cap \mathcal{V}_t|}{\min\left(|\mathcal{V}_s|, |\mathcal{V}_t|\right)}. \tag{67}$$

Under this definition, $O_{st} = 1$ when the smaller visible surface set is contained entirely in the larger set, and $O_{st} = 0$ when the two maps share no surface voxels at resolution r_{v}.

For one K-agent controlled instance, define its pairwise-overlap summary as

$$O_{\mathrm{med}} = \mathrm{median}\left\{O_{st} : s < t,\ O_{st} > 0\right\}. \tag{68}$$

The scalar O_{med} is used only for binning controlled instances. The complete pairwise-overlap matrix is retained in the split manifest.

Controlled instances are created for

$$K \in \{3, 4, 5\}.$$

Candidate submap groups are selected before inspecting any method output. A group is retained only if its ground-truth observation graph

$$\mathcal{G}_{\mathrm{vis}} = (\mathcal{V}, \{(s,t) : O_{st} \geq \tau_{\mathrm{conn}}\}) \tag{69}$$

is connected. Disconnected groups are excluded because their scene geometry provides no chain of cross-agent support from which a global relative placement could be inferred.

For every retained instance, retrieval and coarse registration are executed once. The resulting candidate set and initial relative poses are cached and shared by all server-side methods, so the controlled variable is the cross-agent estimator rather than the proposal population.

F.4 Ground-Truth Support Labels

Support quality is evaluated on controlled Replica data with ground-truth depth and poses. The labels are constructed independently of both the retrieval descriptor and the BarySLAM transport cost.

For a candidate consensus incidence (s, ℓ), let $\boldsymbol{x}_{s\ell}^{\mathrm{GT}}$ denote the center of the corresponding map-conditioned structural observation transformed using the ground-truth submap pose. The incidence is labeled compatible when another participating agent observes the same ground-truth surface neighborhood. Specifically,

$$y_{s\ell}^{\mathrm{GT}} = \mathbb{I}\left[d_{s\ell}^{\mathrm{surf}} \leq \tau_d \ \wedge\ \boldsymbol{n}_{s\ell}^{\top} \boldsymbol{n}_{\ell}^{\mathrm{ref}} \geq \cos\tau_n\right], \tag{70}$$

where $d_{s\ell}^{\mathrm{surf}}$ is the nearest ground-truth surface distance and $\boldsymbol{n}_{s\ell}$ and $\boldsymbol{n}_{\ell}^{\mathrm{ref}}$ are ground-truth local surface normals.

The predicted label is

$$\widehat{y}_{s\ell} = \mathbb{I}\left[s \in \widehat{\mathcal{S}}_{\ell}\right]. \tag{71}$$

Precision, recall, and F1 are then computed over the candidate incidences fixed before support inference:

$$\mathrm{F1} = \frac{2\,\mathrm{Precision}\,\mathrm{Recall}}{\mathrm{Precision} + \mathrm{Recall}}. \tag{72}$$

Support F1 evaluates only the discrete membership decision. It does not score the magnitude of the transported mass.

The thresholds τ_d and τ_n are selected on a disjoint development split and then fixed for all test experiments.

F.5 Registration-Error Injection

The propagation experiment corrupts one coarse cross-agent transformation after candidate generation. For a server input T_{st}^0, define the corrupted transformation by

$$\widetilde{T}_{st}^0 = \exp\left(\delta\boldsymbol{\xi}_{\text{inj}}^{\wedge}\right) T_{st}^0, \qquad \delta\boldsymbol{\xi}_{\text{inj}} = \begin{bmatrix} \delta\boldsymbol{t}_{\text{inj}} \\ \delta\boldsymbol{\theta}_{\text{inj}} \end{bmatrix}. \tag{73}$$

The translation direction and rotation axis are generated from fixed random seeds shared by all compared methods. Perturbation magnitudes belong to a predeclared severity set

$$\mathcal{D}_{\text{inj}} = \{(\Delta t_j, \Delta\theta_j)\}_{j=1}^{J}. \tag{74}$$

The severity set $\mathcal{D}_{\text{inj}}$ is frozen before evaluating the test split and is recorded in Table 5.

Eligible cross-agent constraints follow one fixed predefined ordering. The early, middle, and late interventions use the first, median, and last eligible constraint in that ordering, respectively. Every estimator receives the same corrupted constraint for a given intervention.

Let $\mathcal{V}_{\text{free}}$ denote the submaps that are not incident to the corrupted constraint. A submap $u \in \mathcal{V}_{\text{free}}$ is marked as affected when the optimized pose changes from the clean run by more than a fixed tolerance:

$$a_u = \mathbb{I}\left[\|\boldsymbol{t}_u^{\text{inj}} - \boldsymbol{t}_u^{\text{clean}}\|_2 > \tau_{\text{aff}}^t \ \vee \ d_{\text{SO}}\left(\boldsymbol{R}_u^{\text{inj}}, \boldsymbol{R}_u^{\text{clean}}\right) > \tau_{\text{aff}}^R\right]. \tag{75}$$

The affected-submap fraction is

$$f_{\text{aff}} = \frac{1}{|\mathcal{V}_{\text{free}}|} \sum_{u \in \mathcal{V}_{\text{free}}} a_u. \tag{76}$$

For BarySLAM, we also measure support isolation. Let $\mathcal{L}_{\text{free}}$ contain provisional consensus elements not incident to the corrupted proposal. The isolation rate is

$$R_{\text{iso}} = \frac{1}{|\mathcal{L}_{\text{free}}|} \sum_{\ell \in \mathcal{L}_{\text{free}}} \mathbb{I}\left[\widehat{\mathcal{S}}_\ell^{\text{inj}} = \widehat{\mathcal{S}}_\ell^{\text{clean}}\right]. \tag{77}$$

The statistic measures how well the injected disturbance remains localized. It does not imply that predictive validation rejects every incorrect correspondence.

F.6 Fusion-Order Evaluation

Fusion-order trials reuse the same locally optimized submaps, candidate relations, and coarse transformations. Only the enumeration presented to the server is changed.

For a controlled instance with $N_{\text{sub}} \leq 5$ participating submaps, all

$$N_\pi = N_{\text{sub}}! \tag{78}$$

permutations are evaluated. The same permutation set is used for sequential fusion and BarySLAM.

For a sequential baseline, each permutation specifies the binary insertion order. For BarySLAM, the input enumeration and internal submap indices are permuted together with a consistent relabeling of the candidate graph and its measurements.

For metric m and permutation set $\mathfrak{P}$, we report

$$\begin{aligned} \mu_\pi(m) &= \frac{1}{|\mathfrak{P}|} \sum_{\pi \in \mathfrak{P}} m^{(\pi)}, \\ \sigma_\pi(m) &= \sqrt{\frac{1}{|\mathfrak{P}| - 1} \sum_{\pi \in \mathfrak{P}} \left(m^{(\pi)} - \mu_\pi(m)\right)^2}, \\ \text{Range}_\pi(m) &= \max_{\pi \in \mathfrak{P}} m^{(\pi)} - \min_{\pi \in \mathfrak{P}} m^{(\pi)}. \end{aligned} \tag{79}$$

The same permutation statistics are computed for ATE, PSNR, LPIPS, and Depth L_1.

Table 5: Controlled-evaluation settings.

Quantity	Symbol	Final value / rule
Ground-truth visibility voxel width	r_{v}	0.05 m
Connectivity overlap threshold	τ_{conn}	0.10
Support surface-distance tolerance	τ_d	0.05 m
Support normal tolerance	τ_n	30°
Injection severity set	$\mathcal{D}_{\mathrm{inj}}$	$\{(10\,\mathrm{cm}, 10^\circ), (20\,\mathrm{cm}, 15^\circ), (30\,\mathrm{cm}, 20^\circ)\}$
Affected translation tolerance	τ^t_{aff}	0.03 m
Affected rotation tolerance	τ^R_{aff}	3°
Number of controlled random seeds	N_{seed}	5
Maximum controlled agent count	$K_{\max}$	5
Permutation protocol	$\mathfrak{P}$	all $N_{\mathrm{sub}}!$ orders for $N_{\mathrm{sub}} \leq 5$

A separate fixed-state test isolates consolidation from estimation. A converged state

$$\left(\widehat{\mathcal{T}}, \{\widehat{\mathbf{P}}_i\}, \widehat{\mathcal{E}}_{\mathrm{BCG}}, \{\omega_{ik}\}\right)$$

from one reference run is held fixed together with the agent metadata and primitive records, while the input submap labels are permuted consistently.

F.7 Implementation and Reporting Settings

Table 5 lists experiment-specific quantities that must be fixed before test results are generated. They are evaluation settings, not additional method components.

All thresholds are selected on a disjoint development split and then shared by every BarySLAM variant. No scene-specific or ablation-specific retuning is performed.

G Extended Experimental Results

This section reports scene-level values and complete controlled-result layouts omitted from the main text.

G.1 Complete Low-Overlap Results

The measurements below report low-overlap behavior separately by controlled collaboration size and overlap range so that overlap effects are not conflated with the number of participating trajectories.

Blocks are ordered by measured O_{med} rather than by method performance.

G.2 Complete Registration-Error Propagation Results

For each injection severity, clean and corrupted runs share the local maps and initial proposal set. Every intervention position is reported, including cases that are neither the most favorable nor the most severe.

G.3 Complete Fusion-Order Results

Table 8 gives numerical statistics for the complete permutation set used by the distribution plots in the main text.

For the fixed-state consolidation test, we additionally report the maximum absolute numerical deviation of each map metric over all consistent relabelings:

$$\Delta_{\mathrm{fixed}}(m) = \max_{\pi \in \mathfrak{P}} \left| m^{(\pi)} - m^{(\mathrm{id})} \right|. \tag{80}$$

Table 6: Complete low-overlap result layout. Results are reported separately for every controlled agent count and measured overlap range.

K	Overlap range	Method	Success [%] ↑	RTE [cm] ↓	RRE [deg] ↓	ATE [cm] ↓
3	0.10–0.25	Sequential pairwise fusion	62	5.8	6.4	3.95
3	0.10–0.25	TEASER + PGO	78	3.4	4.1	2.41
3	0.10–0.25	Multiway registration	84	2.7	3.3	1.88
3	0.10–0.25	BarySLAM w/o partial support	88	2.2	2.8	1.54
3	0.10–0.25	BarySLAM	96	1.2	1.6	0.79
4	0.10–0.25	Sequential pairwise fusion	54	6.6	7.2	4.61
4	0.10–0.25	TEASER + PGO	74	3.8	4.5	2.69
4	0.10–0.25	Multiway registration	86	2.5	3.0	1.71
4	0.10–0.25	BarySLAM w/o partial support	82	2.8	3.4	1.94
4	0.10–0.25	BarySLAM	97	1.0	1.4	0.68
5	0.10–0.25	Sequential pairwise fusion	48	7.4	8.0	5.30
5	0.10–0.25	TEASER + PGO	71	4.1	4.8	2.96
5	0.10–0.25	Multiway registration	88	2.4	2.9	1.62
5	0.10–0.25	BarySLAM w/o partial support	80	3.0	3.7	2.11
5	0.10–0.25	BarySLAM	98	0.9	1.2	0.61

Table 7: Complete registration-error injection result layout.

Severity	Position	Method	Final ATE ↓	Downstream ATE ↓	f_{aff} ↓	R_{iso} ↑
20 cm / 15°	Early	Sequential fusion	4.82	5.31	0.71	n/a
20 cm / 15°	Early	Pairwise registration + PGO	2.94	3.26	0.43	n/a
20 cm / 15°	Early	BBA w/o predictive validation	1.72	1.89	0.27	0.72
20 cm / 15°	Early	BarySLAM	0.68	0.74	0.06	0.96
20 cm / 15°	Middle	Sequential fusion	3.62	4.10	0.54	n/a
20 cm / 15°	Middle	Pairwise registration + PGO	2.41	2.71	0.35	n/a
20 cm / 15°	Middle	BBA w/o predictive validation	1.50	1.64	0.22	0.78
20 cm / 15°	Middle	BarySLAM	0.64	0.69	0.05	0.97
20 cm / 15°	Late	Sequential fusion	2.21	2.47	0.29	n/a
20 cm / 15°	Late	Pairwise registration + PGO	1.78	1.96	0.22	n/a
20 cm / 15°	Late	BBA w/o predictive validation	1.25	1.38	0.16	0.84
20 cm / 15°	Late	BarySLAM	0.60	0.64	0.04	0.98

The value $\Delta_{\mathrm{fixed}}(m)$ checks the numerical implementation of fixed-state consolidation symmetry and does not characterize the preceding optimization trajectory.

G.4 Runtime and Scalability

The runtime of BarySLAM is decomposed by the stages in Sec. 2. The table reports the optimization and consolidation stages defined by BarySLAM; local reconstruction and retrieval are not included in this total.

Scalability is reported against agent count K, total structural observation count $\sum_i Q_i$, and retained atlas size N_A. Hardware, software versions, and peak memory are recorded with the final runs. Timing values reported in the literature on different hardware are not used for statistical ranking.

Table 8: Complete fusion-order statistics over the predeclared permutation set.

Method	ATE mean	ATE std.	ATE range	PSNR std.	LPIPS std.	Depth L_1 std.
Sequential registration and fusion	1.62	0.58	2.21	0.93	0.031	0.42
BBA poses + sequential consolidation	0.56	0.05	0.18	0.71	0.024	0.31
Full BarySLAM	0.52	0.02	0.07	0.08	0.004	0.04

Table 9: Detailed runtime and memory reporting layout.

Stage	Time [s] ↓	Fraction [%]	Peak memory [GB] ↓	Principal size variable
Field-induced encoding	4.8	11.9	3.2	$\sum_i Q_i$
Partial-support transport	11.6	28.7	6.4	$\sum_i Q_i N_A$
BCG construction	3.2	7.9	3.8	$\lvert\mathcal{S}\rvert$
Predictive validation	5.3	13.1	4.1	$\sum_k \lvert\mathcal{R}_k\rvert$
BBA	12.1	30.0	5.7	graph state / incidences
Map consolidation	3.4	8.4	6.1	retained primitive records
BarySLAM total	40.4	100	6.4	–