Information-Preserving Continuous Occupancy Mapping with Variance-Weighted Submap Joining

Zhuhua Bai*, Yingyu Wang†, Liang Zhao†, and Shoudong Huang*

***Abstract*—Large-scale SLAM remains challenging due to accumulated trajectory drift and the increasing computational cost of maintaining global consistency. Submap joining alleviates these issues by constructing locally consistent submaps and subsequently fusing them into a global map. However, existing occupancy-based submap joining methods operate on discrete grids, resulting in non-smooth gradients during optimization and neglecting the uncertainty associated with occupancy estimates. We propose the first continuous probabilistic submap joining framework that jointly optimizes submap poses and a global occupancy field in the latent log-odds space. The framework employs an information-preserving sparse Bayesian formulation that compresses raw occupancy observations into sufficient-statistic log-odds tuples while retaining the posterior information of the original observations. This yields closed-form predictive mean and variance estimates for occupancy mapping, which directly enable a submap joining formulation with analytical Jacobians, leading to more accurate submap joining and yielding a closed-form optimal global map upon pose convergence. Experiments on both simulated and large-scale practical datasets demonstrate that the proposed method achieves higher pose accuracy and improved global consistency than state-of-the-art (SOTA) grid-based submap joining approaches, while producing more compact map representations and better-calibrated uncertainty estimates than the continuous occupancy mapping baseline.**

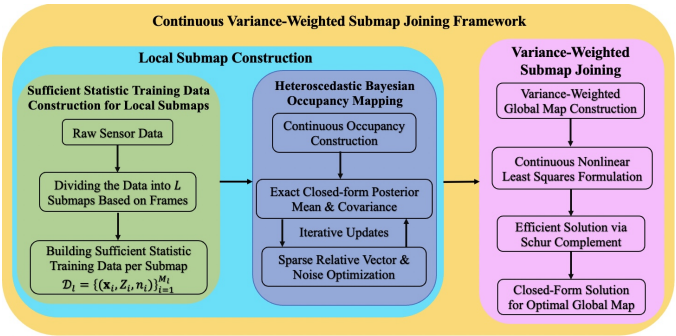

Fig. 1: Overview of the proposed framework. Each local submap compresses raw observations into sufficient-statistic log-odds tuples $\mathcal{D}_l = \{(\mathbf{x}_i, Z_i, n_i)\}_{i=1}^{M_l}$, where $\mathbf{x}_i$, Z_i, and n_i denote the spatial location, average log-odds value, and accumulated observation count, respectively. A heteroscedastic RVM regression then yields closed-form predictive mean and variance, which drive a variance-weighted joining formulation that jointly optimizes local submap frames and the global occupancy field via analytical Jacobians and a closed-form globally optimal map.

I. Introduction

Large-scale SLAM is a long-standing challenge in autonomous navigation: as the explored environment grows, accumulated trajectory drift and increasing computational cost make globally consistent estimation progressively difficult. Submap joining has emerged as a foundational strategy for addressing these challenges [4, 20, 16, 17]. State-of-the-art occupancy-based submap joining methods [16, 17] jointly optimize submap poses and a global occupancy map within a nonlinear least-squares (NLLS) formulation, but operate entirely on discrete occupancy grids and rely on bilinear interpolation during pose-map optimization. This introduces two structural limitations: (i) the interpolated gradient field is piecewise-defined and non-smooth at cell boundaries, which can adversely affect optimization convergence and accuracy; and (ii) the NLLS formulation fails to account for occupancy uncertainty, which compromises the reliability of residual weighting during optimization.

A natural question is whether a continuous probabilistic occupancy representation, which offers smooth gradients and predictive uncertainty by construction, could be combined with submap joining to address both limitations above. Existing continuous methods, however, are not designed for large-scale operation. Full Gaussian Process maps [7, 5] scale cubically and must retain all training data. Hilbert Maps [8, 10] and sparse Relevance Vector Machine (RVM) methods [14, 1] mitigate this through random Fourier features or relevance-vector sparsification, but their feature configurations are anchored to a single global frame and cannot absorb trajectory drift once the map has been built. More critically, virtually all existing continuous mapping methods take robot poses as given inputs and cannot refine them [8, 10, 1, 6]; long-trajectory odometry error therefore propagates directly into the map, producing structural inconsistency at loop closures that no local feature flexibility can correct. Continuous mapping at large scale thus depends on submap joining as much as grid-based mapping does, yet no submap joining formulation has been developed for continuous occupancy maps.

Among continuous representations, the RVM framework is uniquely suited for this role: relevance-vector sparsification yields a compact per-submap kernel set, and the kernel-based predictor admits analytical spatial derivatives required for pose Jacobians. Existing RVM-based methods [1, 15], however, formulate occupancy as binary classification with a non-conjugate

*Zhuhua Bai and Shoudong Huang are with the Robotics Institute, University of Technology Sydney, Australia (e-mail: Zhuhua.Bai@student.uts.edu.au; Shoudong.Huang@uts.edu.au).

†Yingyu Wang and Liang Zhao are with the School of Informatics, University of Edinburgh, UK(e-mail: Yingyu.Wang@uts.edu.au; liang.zhao@ed.ac.uk).

probit or logistic likelihood, forcing reliance on Laplace or variational approximation [12, 13]. The resulting predictive variance is unreliable, especially near occupancy boundaries, which is where submap joining critically requires trustworthy uncertainty to weight residuals during pose optimization.

We address this by reformulating sparse Bayesian occupancy mapping as a heteroscedastic Gaussian regression problem in the latent log-odds space, with raw observations compressed into sufficient-statistic log-odds tuples. We prove that this compression is lossless for the weight posterior, restoring Gaussian conjugacy and yielding exact closed-form posterior mean and covariance. The heteroscedastic noise precision is additionally updated jointly with the relevance set within a single marginal-likelihood maximization, rather than in a separate stage. Built on the resulting reliable predictive uncertainty, our variance-weighted submap joining formulation admits analytical Jacobians throughout the optimization and yields a closed-form optimal global map via variance-weighted fusion upon pose convergence. The overall pipeline is illustrated in Fig. 1.

The main contributions of this paper are as follows:

1) The first variance-weighted continuous submap joining framework that jointly optimizes local submap frames and a global occupancy field directly in a continuous probabilistic latent space, with analytical Jacobians and a closed-form optimal global map upon pose convergence.
2) An information-preserving heteroscedastic RVM occupancy formulation based on sufficient-statistic log-odds compression, theoretically proven to yield a weight posterior equivalent to that of the raw observations. This restores Gaussian conjugacy, removes the Laplace approximation of classification-based RVM, and supports joint update of the noise precision and the relevance set.
3) Extensive evaluation on 2D simulated and practical datasets demonstrates superior pose accuracy and global consistency to the SOTA grid-based baseline, and superior model compactness to the RVM-based baseline.

II. INFORMATION-PRESERVING CONTINUOUS OCCUPANCY MAPPING

To address limitations of existing RVM-based methods [1, 15], we reformulate occupancy estimation as a Gaussian-conjugate Bayesian regression in the latent log-odds space via sufficient-statistic compression, as presented next.

A. Bayesian Continuous Occupancy Mapping

1) Construction of the Training Data: We build a multi-resolution grid map that refines resolution by a factor $h > 1$ near obstacle boundaries and coarsens it (e.g., 0.5 m) in free space, balancing fine details and training sample count. For each laser ray, the endpoint is treated as an occupied (hit) observation, and all traversed cells as free observations.

For the i-th grid cell located at $\mathbf{x}_i$, the j-th observation of its occupancy value is represented in log-odds form as

$$z_{ij} = \ln \frac{p(\mathbf{x}_i \in \text{occ})}{1 - p(\mathbf{x}_i \in \text{occ})} \tag{1}$$

Following [3, 11], occupied and free observations are assigned fixed log-odds values, $z_{\text{hit}} = \ln(0.7/0.3)$ and $z_{\text{free}} = \ln(0.4/0.6)$, respectively. Given n_i observations of grid cell $\mathbf{x}_i$, the corresponding average log-odds value is computed as $Z_i = \frac{1}{n_i}\sum_{j=1}^{n_i} z_{ij}$. The resulting training dataset is constructed as $\mathcal{D} = \{(\mathbf{x}_i, Z_i, n_i)\}_{i=1}^{M}$, where M denotes the total number of observed grid cells.

2) Heteroscedastic Continuous Occupancy Construction: Each relevance vector defines a basis function via a kernel $k(\mathbf{x}, \mathbf{x}_m)$, mapping a grid cell $\mathbf{x}$ to a feature vector $\Phi_{\mathbf{x}} = [k(\mathbf{x}, \mathbf{x}_1), \ldots, k(\mathbf{x}, \mathbf{x}_M)]^\top \in \mathbb{R}^M$. We adopt the radial basis function kernel $k(\mathbf{x}, \mathbf{x}_m) = \eta \exp(-\gamma \|\mathbf{x} - \mathbf{x}_m\|^2)$ and model the latent occupancy field directly in the log-odds space as

$$F(\mathbf{x}) := \Phi_{\mathbf{x}}^\top \boldsymbol{w} + b \tag{2}$$

with weights $\boldsymbol{w} \in \mathbb{R}^M$ and bias $b \in \mathbb{R}$.

Observations falling into the same grid cell are assumed to share the latent value $F(\mathbf{x})$ and to be conditionally independent. A raw log-odds measurement is modeled as

$$z_{ij} = F(\mathbf{x}_i) + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \beta^{-1}) \tag{3}$$

with the noise precision β.

By the Gaussian closure property, the per-cell average Z_i satisfies

$$p(Z_i \mid \mathbf{x}_i, \boldsymbol{w}, \beta, n_i) = \mathcal{N}\left(Z_i \,\middle|\, \Phi_{\mathbf{x}_i}^\top \boldsymbol{w} + b,\ \frac{1}{\beta n_i}\right) \tag{4}$$

Therefore, the effective measurement precision accumulates linearly as βn_i. Stacking all M cells yields the heteroscedastic data likelihood

$$p(\mathbf{Z} \mid \mathbf{X}, \boldsymbol{w}, \beta, \mathbf{N}) = \mathcal{N}\left(\mathbf{Z} \,\middle|\, \Phi^\top \boldsymbol{w} + b\mathbf{1},\ \frac{1}{\beta \mathbf{N}}\right) \tag{5}$$

where $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_M]^\top \in \mathbb{R}^{M\times 2}$, $\mathbf{Z} = [Z_1, \ldots, Z_M]^\top \in \mathbb{R}^{M\times 1}$, $\mathbf{N} = \operatorname{diag}(n_1, \ldots, n_M) \in \mathbb{R}^{M\times M}$ is a diagonal information-weighting matrix containing the observation counts associated with each sample, and $\Phi \in \mathbb{R}^{M\times M}$ is the kernel matrix evaluated between all pairs of training inputs.

An RVM model imposes a Gaussian prior on each weight w_m with zero mean and precision α_m (i.e., variance $\frac{1}{\alpha_m}$).

$$p(\boldsymbol{w} \mid \boldsymbol{\alpha}) = \prod_{m=1}^{M} \mathcal{N}(w_m \mid 0, \alpha_m^{-1}) = \mathcal{N}(\boldsymbol{w} \mid \mathbf{0}, \mathbf{A}^{-1}) \tag{6}$$

where $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_M]^\top \in \mathbb{R}^{M\times 1}$ denotes the precision vector associated with the weight vector $\boldsymbol{w}$, and $\mathbf{A} = \operatorname{diag}(\alpha_1, \ldots, \alpha_M) \in \mathbb{R}^{M\times M}$ is the corresponding diagonal precision matrix.

With the heteroscedastic likelihood (5) and the Gaussian prior (6) both defined, we can now state the lossless-compression property that underpins the entire formulation.

Proposition 1. *The set of raw per-observation log-odds measurements $\{z_{ij}\}_{i=1,\, j=1}^{M,\, n_i}$ is compressed into the sufficient statistics $\mathcal{D} = \{(\mathbf{x}_i, Z_i, n_i)\}_{i=1}^{M}$. Under this model, for any*

hyperparameters $(\boldsymbol{\alpha}, \beta)$ the weight posteriors induced by the two representations coincide:

$$p\Big(\boldsymbol{w} \,\big|\, \{z_{ij}\}_{i=1,\,j=1}^{M,\,n_i}, \boldsymbol{\alpha}, \beta\Big) \;=\; p\big(\boldsymbol{w} \,\big|\, \mathcal{D},\, \boldsymbol{\alpha}, \beta\big) \tag{7}$$

Operationally, Proposition 1 guarantees that replacing the $\sum_i^M n_i$ raw measurements by the M tuples $(\mathbf{x}_i, Z_i, n_i)$ incurs no information loss for either posterior inference.

Due to Gaussian conjugacy, the posterior distribution remains Gaussian:

$$p(\boldsymbol{w} \mid \mathbf{Z}, \mathbf{X}, \boldsymbol{\alpha}, \beta, \mathbf{N}) = \mathcal{N}(\boldsymbol{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{8}$$

Therefore, the exact closed-form expressions for the posterior mean and covariance can be obtained directly.

$$\begin{aligned} \boldsymbol{\Sigma} = (\boldsymbol{\Sigma}^{-1})^{-1} = (\beta \Phi^\top \mathbf{N} \Phi + \mathbf{A})^{-1} \\ \boldsymbol{\mu} = \beta \boldsymbol{\Sigma} \Phi^\top \mathbf{N}(\mathbf{Z} - b\mathbf{1}) \end{aligned} \tag{9}$$

3) Joint Sequential Optimization of $(\boldsymbol{\alpha}, \beta)$: The hyperparameters $(\boldsymbol{\alpha}, \beta)$ are determined by type-II maximum likelihood, i.e., by maximizing the log marginal likelihood

$$\mathcal{L}(\boldsymbol{\alpha}, \beta) = -\tfrac{1}{2}\big[M \log(2\pi) + \log|\mathbf{C}| + \mathbf{r}^\top \mathbf{C}^{-1} \mathbf{r}\big] \tag{10}$$

with $\mathbf{C} = (\beta \mathbf{N})^{-1} + \Phi \mathbf{A}^{-1} \Phi^\top$ and $\mathbf{r} = \mathbf{Z} - b\mathbf{1}$.

Unlike prior RVM methods [1], which alternate between updating $\boldsymbol{\alpha}$ and β, we update both jointly at every iteration.

The sparsity update for each candidate $\mathbf{x}_m$ follows the standard Tipping–Faul criterion [15].

Updating of β. By taking the partial derivative of (10) with respect to β and setting $\partial \mathcal{L} / \partial \beta = 0$, we obtain β^{new} as

$$\beta^{\text{new}} = \frac{M - \sum_{m=1}^{N_{\text{RV}}} \gamma_m}{\sum_{i=1}^M n_i r_i^2} \tag{11}$$

where the relevant terms are defined as follows:

- N_{RV} is the number of currently active relevance vectors.
- $\gamma_m \equiv 1 - \alpha_m \Sigma_{mm}$ acts as a measure of well-determinedness for the corresponding parameter. Here, Σ_{mm} is the m-th diagonal element of the posterior covariance matrix $\boldsymbol{\Sigma}$, representing the posterior variance of the weight w_m.
- r_i is the i-th element of the residual error, calculated from $\mathbf{v} = \mathbf{Z} - b\mathbf{1} - \boldsymbol{\Phi}\boldsymbol{\mu}$.

B. Continuous Occupancy Prediction

After the model has been trained, for a query location $\mathbf{x}_*$, the predictive distribution is

$$\begin{aligned} p(Z_* \mid \mathbf{x}_*) = \mathcal{N}(Z_* \mid \mu_*, \sigma_*^2) \\ \mu_* = \Phi_{\mathbf{x}_*}^\top \boldsymbol{w} + b, \quad \sigma_*^2 = \Phi_{\mathbf{x}_*}^\top \boldsymbol{\Sigma}\, \Phi_{\mathbf{x}_*} \end{aligned} \tag{12}$$

The latent prediction for $\mathbf{x}_*$ can be converted to an occupancy probability via the probit approximation [9], which incorporates both predictive mean and variance:

$$\overline{P}_{\text{occ}}(\mathbf{x}_*) \approx \frac{1}{1 + \exp(-\kappa \mu_*)}, \quad \kappa = \frac{1}{\sqrt{1 + \pi \sigma_*^2 / 8}} \tag{13}$$

With thresholds $0 < e_{\text{free}} < e_{\text{occ}} < 1$, a query is classified as occupied if $\overline{P}_{\text{occ}} > e_{\text{occ}}$, free if $\overline{P}_{\text{occ}} < e_{\text{free}}$, and unknown otherwise.

C. Variance-Weighted Continuous Submap Joining

The continuous occupancy field defined by (12) is differentiable in $\mathbf{x}$, so both the predictive mean and variance can be queried analytically at arbitrary locations. This enables direct continuous joint optimization of submap frames and the global occupancy field.

1) Variance-Weighted Global Map Construction: Assuming conditional independence between local submap predictions, global fusion is formulated as a minimum-variance Gaussian fusion problem. For N overlapping submaps, the global occupancy estimate at location $\mathbf{x}_f$ is constructed via variance-weighted fusion of the local predictive means $\mu_i(\mathbf{x}_f)$ and variances $\sigma_i^2(\mathbf{x}_f)$.

$$\begin{aligned} \mu_G(\mathbf{x}_f) &= \frac{\sum_{i=1}^N \mu_i(\mathbf{x}_f)/(\sigma_i^2(\mathbf{x}_f) + \epsilon)}{\sum_{i=1}^N 1/(\sigma_i^2(\mathbf{x}_f) + \epsilon)} \\ \sigma_G^2(\mathbf{x}_f) &= \frac{1}{\sum_{i=1}^N 1/(\sigma_i^2(\mathbf{x}_f) + \epsilon)} \end{aligned} \tag{14}$$

where ϵ is a small constant (e.g., 10^{-6}) to prevent division by zero or numerical instability.

2) NLLS Formulation: To formulate the submap joining problem as a NLLS problem, we first discretize the global map. Let the global map be represented by a grid of cells $\{\mathbf{M}(m_1), \ldots, \mathbf{M}(m_G)\}$, where $G = l_w \times l_h$ and m_j is the coordinate of the j-th cell. Denote the latent global log-odds value of j-th cell as $\mathbf{M}(m_j)$. Then the global map is fully described by the vector $\mathbf{M} = [\mathbf{M}(m_1), \ldots, \mathbf{M}(m_G)]^\top$.

The state vector of the NLLS problem consists of the poses of the local submap coordinate (except the first, which serves as the global reference) and the global log-odds map values:

$$\mathbf{V} = [(\mathbf{T}^r)^\top, \mathbf{M}^\top]^\top \tag{15}$$

where $\mathbf{T}^r = [(\mathbf{T}_2^r)^\top, \ldots, (\mathbf{T}_g^r)^\top]^\top$ contains the poses of all non-reference submaps. Each $\mathbf{T}_j^r = [\mathbf{t}_j^\top, \theta_j]^\top \in SE(2)$ represents the j-th pose of local submap coordinate , with $\mathbf{t}_j = [x_j, y_j]^\top$ and θ_j being the position and orientation, respectively; the corresponding rotation matrix is $\mathbf{R}_j$.

The objective function minimizes the variance-weighted residuals between the global map and each local submap:

$$f(\mathbf{V}) = \sum_{i=1}^N \sum_{j \in \{obs_i\}} \frac{1}{\sigma_i^2(\mathbf{p}_{im_j}) + \epsilon} \big(\mathbf{M}(m_j) - \mu_i(\mathbf{p}_{im_j})\big)^2 \tag{16}$$

where $\mathbf{p}_{im_j} = \mathbf{R}_i^\top (m_j - \mathbf{t}_i)$ is the projection of global cell m_j into the i-th submap coordinate frame, and $\mu_i(\cdot)$ and $\sigma_i^2(\cdot)$ are the predictive mean and variance of the i-th local RVM at the projected location.

3) Gauss-Newton Iteration: To solve the NLLS problem (16) using the Gauss-Newton method, we define the residual for global cell m_j observed by the i-th submap as:

$$r_{im_j} = \frac{\mathbf{M}(m_j) - \mu_i(\mathbf{p}_{im_j})}{\sqrt{\sigma_i^2(\mathbf{p}_{im_j}) + \epsilon}} \tag{17}$$

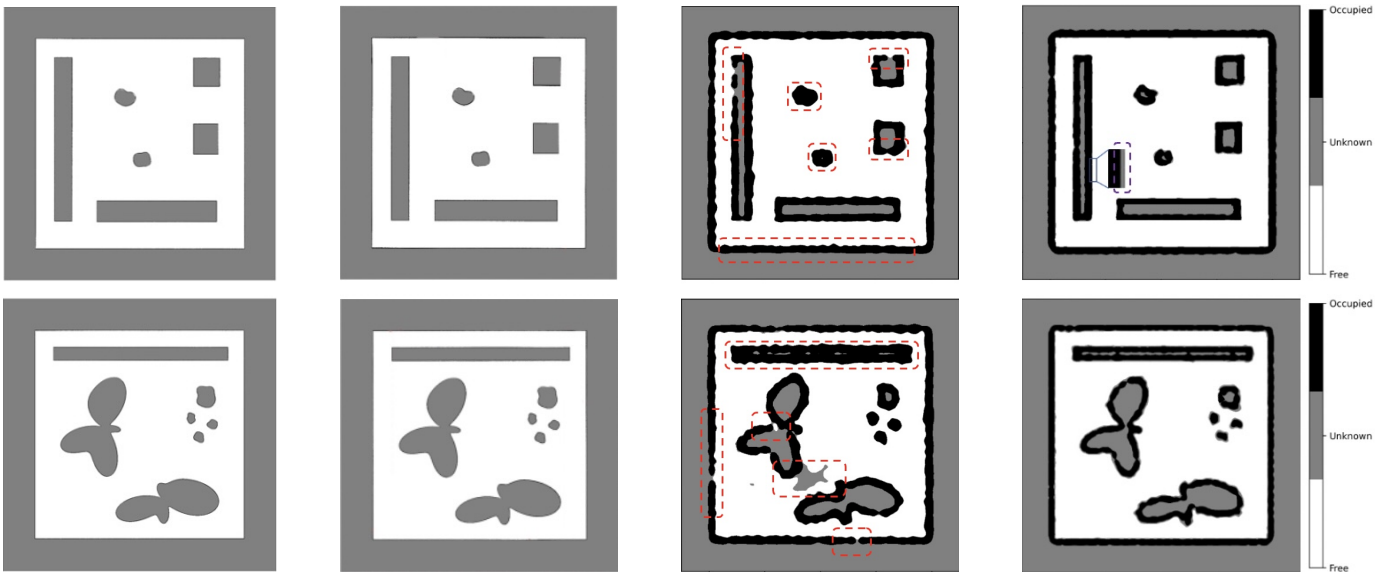

(a) Ground Truth (b) CBJS [16] (d) SBKM [1] (e) Ours

Fig. 2: Occupancy grid maps and continuous occupancy maps generated by submap joining. Initial poses come from Occupancy-SLAM. After optimization by CBSJ, SBKM, and our approach, the maps are presented for Simulation 1 (first row) and Simulation 2 (second row) on one dataset.

so that $f(\mathbf{V}) = \sum_{i=1}^{N} \sum_{j \in \{obs_i\}} r_{im_j}^2$.

The Gauss-Newton update $\Delta\mathbf{V} = [\Delta\mathbf{T}^\top, \Delta\mathbf{M}^\top]^\top$ is obtained by solving:

$$\mathbf{J}^\top \mathbf{J}\, \Delta\mathbf{V} = -\mathbf{J}^\top \mathbf{r} \tag{18}$$

where $\mathbf{r}$ is the stacked residual vector and $\mathbf{J} = [\mathbf{J}_T \;\; \mathbf{J}_M]$ is the Jacobian of $\mathbf{r}$ with respect to $\mathbf{T}^r$ and $\mathbf{M}$, respectively.

4) Efficient Solution via Schur Complement: The Eq. (18) has the block structure:

$$\begin{bmatrix} \mathbf{H}_{TT} & \mathbf{A}_{TM} \\ \mathbf{A}_{TM}^\top & \mathbf{H}_{MM} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{T} \\ \Delta\mathbf{M} \end{bmatrix} = - \begin{bmatrix} \mathbf{g}_T \\ \mathbf{g}_M \end{bmatrix} \tag{19}$$

Since $\mathbf{H}_{MM}$ is diagonal and can be inverted in $O(G)$ time. Applying the Schur complement [18] to eliminate $\Delta\mathbf{M}$:

$$\Delta\mathbf{T} = \left(\mathbf{H}_{TT} - \mathbf{A}_{TM}\mathbf{H}_{MM}^{-1}\mathbf{A}_{TM}^\top\right)^{-1}\left(-\mathbf{g}_T + \mathbf{A}_{TM}\mathbf{H}_{MM}^{-1}\mathbf{g}_M\right) \tag{20}$$

The global map update is recovered via:

$$\Delta\mathbf{M} = \mathbf{H}_{MM}^{-1}\left(-\mathbf{g}_M - \mathbf{A}_{TM}^\top \Delta\mathbf{T}\right) \tag{21}$$

After the submap poses $\hat{\mathbf{T}}^r$ converge, the residual objective with respect to the global map decouples across cells, yielding a closed-form expression for the optimal map. This is our second proposition.

Proposition 2. *Given converged submap poses $\hat{\mathbf{T}}^r$, the NLLS map objective* (16) *is strictly convex and separable in $\{\mathbf{M}(m_j)\}_{j=1}^{G}$, and its unique global minimizer at cell m_j is the variance-weighted fusion of the local submap predictions:*

$$\hat{\mathbf{M}}(m_j) \;=\; \frac{\sum_{i=1}^{N} \mu_i(\hat{\mathbf{p}}_{im_j}) / \left(\sigma_i^2(\hat{\mathbf{p}}_{im_j}) + \epsilon\right)}{\sum_{i=1}^{N} 1 / \left(\sigma_i^2(\hat{\mathbf{p}}_{im_j}) + \epsilon\right)} \tag{22}$$

This proposition implies that, given fixed submap poses, the globally optimal continuous occupancy map admits the closed-form solution in Eq. (22) via variance-weighted fusion.

III. Experiment

We evaluate our framework, which jointly optimizes local submaps and the global map within a continuous probabilistic occupancy latent space. Lacking direct continuous submap joining baselines, we compare against GBSJ [16], a state-of-the-art discrete grid method that outperforms Cartographer [2]. For continuous mapping, we benchmark against SBKM [1] due to its similar sparse Bayesian formulation. Other continuous mapping methods are based on substantially different representations, making direct comparison under the same submap joining framework nontrivial.

As public practical 2D LiDAR datasets lack ground-truth poses, we first evaluate our method in two simulated environments (quantitatively and qualitatively), then validate on two public practical datasets. Table I summarizes all datasets used.

In all experiments, local submaps are built using our continuous occupancy mapping framework, with internal poses optimized by Occupancy-SLAM [17]. Depending on dataset scale, the number of local submaps ranges from 3 to 10.

TABLE I: The parameters of all datasets.

Dataset	No. Scans	Duration (s)	Map Size (m)
Simulation 1	3640	117	50×50
Simulation 2	2680	83	50×50
Museum b0 1G	5522	152	100×100
Museum b0 EG	22650	615	200×160

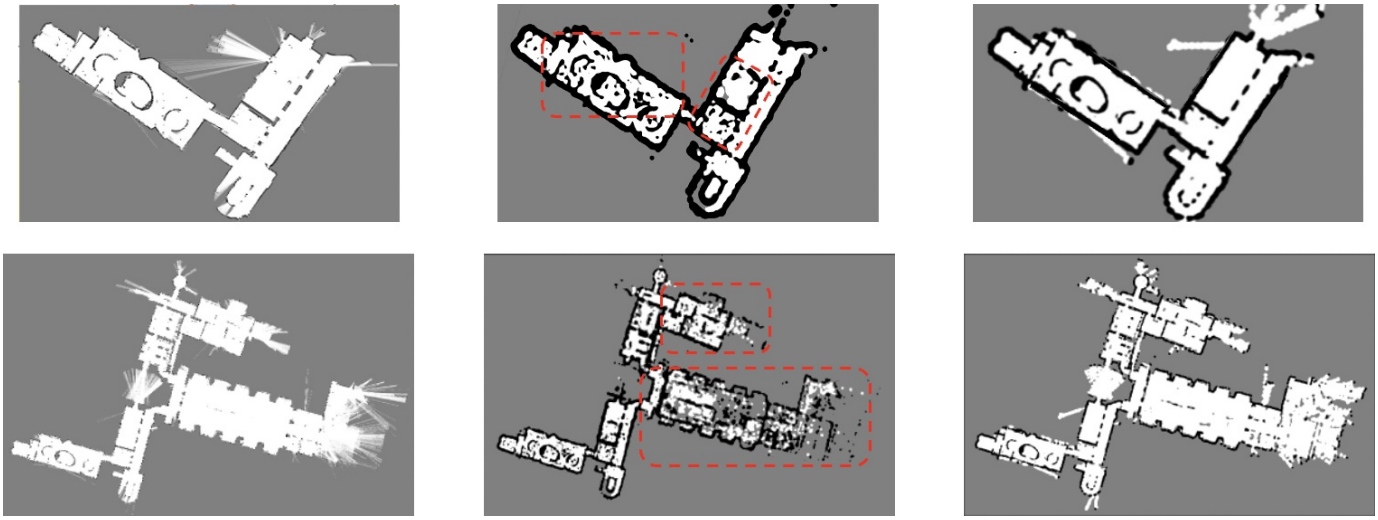

(a) CBSJ (b) SBKM (c) Ours

Fig. 3: Experimental results on the practical datasets: Deutsches Museum b0 1G (first row) and b0 EG (second row). Our method achieves comparable performance to CBSJ and outperforms SBKM in both detail and overall quality.

1) Simulation Experiments: Following [16], we use two simulation environments (varying nonlinearity, non-convex obstacles, long corridors). Zero-mean Gaussian noises are added to ground-truth scan beams and odometry inputs, generating five random noise datasets per environment.

Table II presents the quantitative pose error comparison for GBSJ, SBKM, and our method over five simulation runs. Translation error (in meters) and rotation error (in radians) are evaluated using Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). Our method consistently outperforms all discrete grid-based algorithms across every metric. This clearly validates that submap joining on a continuous map is advantageous. Furthermore, our method also achieves substantially better results than SBKM, demonstrating that our mapping formulation is more effective for submap joining.

TABLE II: Pose estimation errors.

Dataset	Metric	GBSJ	SBKM	Our
Simulation 1	MAE (Trans/m)	**0.0453**	0.2837	**0.0164**
	MAE (Rot/rad)	**0.0012**	0.0068	**0.0009**
	RMSE (Trans/m)	**0.0609**	0.3823	**0.0279**
	RMSE (Rot/rad)	**0.0015**	0.0174	**0.0011**
Simulation 2	MAE (Trans/m)	**0.0151**	0.2532	**0.0051**
	MAE (Rot/rad)	**0.0011**	0.0106	**0.0003**
	RMSE (Trans/m)	**0.0205**	0.2506	**0.0105**
	RMSE (Rot/rad)	**0.0018**	0.0023	**0.0013**

Red and **blue** represent the best and second-best results, respectively.

Fig. 2 compares occupancy grid maps and continuous occupancy maps from submap joining in Simulations 1 and 2. Unlike traditional grid methods, continuous maps represent object boundaries more reasonably with smooth transitions between free and occupied space. Our method also outperforms SBKM, especially in the red-dotted regions. As shown in the zoomed-in view (Fig. 2(e)), our method classifies transitional areas as unknown (occupancy probability ≈ 0.5) rather than forcing binary decisions, which is more reasonable.

We quantitatively compare occupancy maps using Area Under the ROC Curve (AUC) and precision, following [16] for ground-truth labels and unknown cells. Table III shows that our method achieves comparable performance to GBSJ and outperforms SBKM. An ablation study on Eq. (13) compares our model with a mean-only variant ($\kappa = 1$) that ignores predictive variance. Incorporating variance consistently improves both Precision and AUC, demonstrating that uncertainty information suppresses overconfident predictions in ambiguous regions, leading to more reliable occupancy classification and map accuracy.

TABLE III: Accuracy of the occupancy grid maps (reslution 10cm).

Dataset	Method	AUC	Precision
Simulation 1	GBSJ	**0.9731**	**0.9806**
	SBKM	0.9341	0.9467
	Ours w/o variance	0.9223	0.9297
	Ours	**0.9839**	**0.9812**
Simulation 2	GBSJ	**0.9846**	**0.9941**
	SBKM	0.9149	0.9419
	Ours w/o variance	0.9227	0.9399
	Ours	**0.9802**	**0.9915**

2) Comparisons Using Practical Datasets: We evaluate our method against CBSJ and SBKM on two large-scale practical datasets from [21]: Deutsches Museum b0 1G [2] and Museum b0 EG [19] (map sizes: 100×100 to 200×160 m²).

Fig. 3 shows that our method achieves accuracy comparable to CBSJ on two large-scale datasets while significantly outperforming SBKM, especially in the red-dot regions. SBKM performs better on small simulations than on large practical

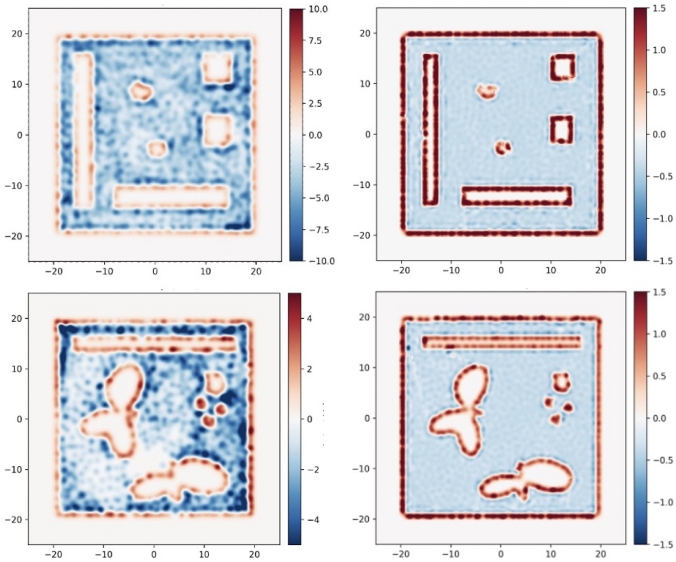

(a) SBKM (b) Ours

Fig. 4: Visualization Comparison of Our Method and SBKM for Continuous Mapping across two datasets: Simulation 1 (first row) and Simulation 2 (second row).

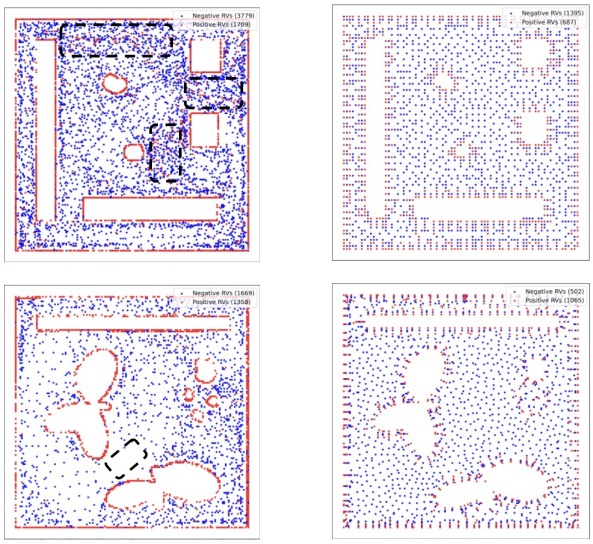

(a) SBKM (b) Ours

Fig. 5: Comparison of optimized relevance vector (RV) distributions between SBKM and the proposed method. Positive and negative RVs are shown in red and blue, respectively. The proposed method produces a more compact and spatially organized RV distribution with clearer separation between occupied and free-space structures.

ones (e.g., b0 EG) because it cannot handle the increasing observation noise from longer trajectories, which degrades its performance.

3) Efficiency and Convergence: On two simulation datasets, we compare the proposed method against SBKM, with a focus on map quality, model compactness, submap iteration count, and computational efficiency.

Map Quality. As shown in Fig. 4, our method produces smoother occupancy maps with better boundary continuity and more consistent free-to-occupied transitions, while SBKM exhibits local fluctuations and irregular boundaries due to approximate posterior inference under non-conjugate likelihood.

Fig. 5 compares the optimized RV distributions. Positive RVs correspond to occupied structures and negative RVs to free space. Compared with SBKM, our method gives a more spatially organized sparse representation: positive RVs are concentrated near occupancy boundaries, and negative RVs are more uniformly distributed over free space, especially in the black-dotted region. This structured RV distribution improves map smoothness and boundary consistency.

Model Compactness. Table IV reports average construction time and number of RVs per submap. Our method is $2.1\times$ faster than SBKM while using nearly $1/3$ the RVs per submap. The speed gain comes from our Gaussian-conjugate formulation, which avoids iterative Laplace approximation. The fewer RVs indicate a more compact sparse representation with lower redundancy and memory.

Convergence Speed. Table IV also compares convergence during submap joining. Our method converges with far fewer iterations and achieves lower final errors (Table II). This is due to our smoother continuous occupancy representation, which provides more stable gradient information.

TABLE IV: Computational Efficiency and Convergence Comparison (Intel Core i9-14900HX).

Dataset	Method	Time/Sub	RVs/Sub	Iterations
Simulation 1	SBKM	45 min	5488	12
	Our	21 min	2082	4
Simulation 2	SBKM	35 min	4299	16
	Our	17 min	1576	6

Assessment of Robustness to Initial Guess. In Simulation 1, zero-mean uniformly distributed noise with varying bounds was added to the ground truth poses of the non-reference local submap coordinates to evaluate convergence under initialization errors. Three noise levels were tested: Level 1 (translation: $[-1, 1]$ m, rotation: $[-0.15, 0.15]$ rad), Level 2 ($[-2, 2]$ m, $[-0.3, 0.3]$ rad), and Level 3 ($[-3, 3]$ m, $[-0.45, 0.45]$ rad). Our method achieved a 100% convergence rate under Levels 1 and 2, and maintained a high convergence rate of 70% even under the most challenging Level 3 noise, demonstrating strong robustness to highly noisy initial guesses.

IV. Conclusion

We propose the first continuous probabilistic submap joining framework. Meanwhile, we improve upon existing RVM-based mapping approaches to make them more suitable for our joint optimization framework. Experiments demonstrate the effectiveness of our method. Future work includes adaptive submap partitioning via information gain and extension to 3D.

Please refer to the full version of this paper for more details (https://github.com/xiaobaitu123321/paper/tree/master/files).

References

[1] Thai Duong, Michael Yip, and Nikolay Atanasov. Autonomous navigation in unknown environments with sparse bayesian kernel-based occupancy mapping. *IEEE Transactions on Robotics*, 38(6):3694–3712, 2022.

[2] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar slam. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 1271–1278. IEEE, 2016.

[3] Armin Hornung, Kai M Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. Octomap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous robots*, 34(3):189–206, 2013.

[4] Shoudong Huang, Zhan Wang, and Gamini Dissanayake. Sparse local submap joining filter for building large-scale maps. *IEEE Transactions on Robotics*, 24(5):1121–1130, 2008.

[5] Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Dissanayake. Warped gaussian processes occupancy mapping with uncertain inputs. *IEEE Robotics and Automation Letters*, 2(2):680–687, 2017.

[6] Peter Zhi Xuan Li, Sertac Karaman, and Vivienne Sze. Gmmap: Memory-efficient continuous occupancy map using gaussian mixture model. *IEEE Transactions on Robotics*, 40:1339–1355, 2024.

[7] Simon T O'Callaghan and Fabio T Ramos. Gaussian process occupancy maps. *The International Journal of Robotics Research*, 31(1):42–62, 2012.

[8] Fabio Ramos and Lionel Ott. Hilbert maps: Scalable continuous occupancy mapping with stochastic gradient descent. *The International Journal of Robotics Research*, 35(14):1717–1730, 2016.

[9] Matthias Seeger. Gaussian processes for machine learning. *International journal of neural systems*, 14(02):69–106, 2004.

[10] Ransalu Senanayake and Fabio Ramos. Bayesian hilbert maps for dynamic continuous occupancy mapping. In *Conference on Robot Learning*, pages 458–471. PMLR, 2017.

[11] Sebastian Thrun. Probabilistic robotics. *Communications of the ACM*, 45(3):52–57, 2002.

[12] Luke Tierney and Joseph B Kadane. Accurate approximations for posterior moments and marginal densities. *Journal of the american statistical association*, 81(393): 82–86, 1986.

[13] Luke Tierney, Robert E Kass, and Joseph B Kadane. Fully exponential laplace approximations to expectations and variances of nonpositive functions. *Journal of the american statistical association*, 84(407):710–716, 1989.

[14] Michael Tipping. The relevance vector machine. *Advances in neural information processing systems*, 12, 1999.

[15] Michael E Tipping and Anita C Faul. Fast marginal likelihood maximisation for sparse bayesian models. In *International workshop on artificial intelligence and statistics*, pages 276–283. PMLR, 2003.

[16] Yingyu Wang, Liang Zhao, and Shoudong Huang. Grid-based submap joining: An efficient algorithm for simultaneously optimizing global occupancy map and local submap frames. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10121–10128. IEEE, 2024.

[17] Yingyu Wang, Liang Zhao, and Shoudong Huang. Occupancy-slam: An efficient and robust algorithm for simultaneously optimizing robot poses and occupancy map. *IEEE Transactions on Robotics*, 41:4057–4077, 2025.

[18] Fuzhen Zhang, editor. *The Schur Complement and Its Applications*, volume 4 of *Numerical Methods and Algorithms*. Springer, New York, 2005. ISBN 978-0-387-24271-2. doi: 10.1007/b105056.

[19] Jiaheng Zhao, Liang Zhao, Shoudong Huang, and Yue Wang. 2d laser slam with general features represented by implicit functions. *IEEE Robotics and Automation Letters*, 5(3):4329–4336, 2020.

[20] Liang Zhao, Shoudong Huang, and Gamini Dissanayake. Linear slam: A linear solution to the feature-based and pose graph slam based on submap joining. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 24–30. IEEE, 2013.

[21] Liang Zhao, Yingyu Wang, and Shoudong Huang. Occupancy-slam: Simultaneously optimizing robot poses and continuous occupancy map. In *Robotics: Science and Systems (RSS)*, pages 1–13. MIT Press Journals, 2022.