

Variational Filtering with Copula Models for SLAM

John D. Martin^{*,1}, Kevin Doherty^{*,2}, Caralyn Cyr¹, Brendan Englot¹, John Leonard²

Abstract—The ability to infer map variables and estimate pose is crucial to the operation of autonomous mobile robots. In most cases the shared dependency between these variables is modeled through a multivariate Gaussian distribution, but there are many situations where that assumption is unrealistic. Our paper shows how it is possible to relax this assumption and perform simultaneous localization and mapping (SLAM) with a larger class of distributions, whose multivariate dependency is represented with a copula model. We integrate the distribution model with copulas into a Sequential Monte Carlo estimator and show how unknown model parameters can be learned through gradient-based optimization. We demonstrate our approach is effective in settings where Gaussian assumptions are clearly violated, such as environments with uncertain data association and nonlinear transition models.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) [1] is considered an essential capability for autonomous robots that operate in settings where precise maps and positioning are unavailable. SLAM enables robots to estimate their pose and landmark locations using a single inference procedure. Given generative models of how the latent pose evolves through time, and how observations relate to pose and landmarks, SLAM algorithms compute a posterior distribution and derive point estimates from its statistics. This paper studies a subclass of SLAM algorithms called *filters*, used to make incremental posterior updates whenever new data becomes available.

Conventional SLAM models impose several distributional and independence assumptions to keep inference tractable. One common assumption restricts poses and landmarks to be jointly Gaussian. Under a Gaussian model, multivariate dependency is represented as a linear function of covariances. Temporal dependency is often encoded with a Bayesian network or a factor graph [2], [3], [4] to sparsify the full joint Gaussian. In spite of these sometimes unrealistic modeling choices, there are many algorithms that can accurately and efficiently recover the hidden variables [5], [6], [7].

Previous work has also pointed out ways in which these assumptions can be unrealistic [8], [9], [10]. Gaussian distributions are appropriate for real-valued variables whose noise is symmetric, unimodal, and lightly-tailed. However, this unimodality is not always appropriate when measurements and landmarks have an uncertain correspondence [9]. This can produce multiple hypotheses, or modes in the distribution, arising from several possible associations. Gaussian assumptions are also unrealistic when there exist complex




Fig. 1: **Representing dependency for SLAM:** Graphical models are often employed to represent probabilistic dependence among latent variables and induce block structure in their (jointly Gaussian) information matrix (*top half*). To move beyond Gaussian assumptions, we express the full joint distribution in terms of two dependency models (i.e. copulas), $c(\tilde{s}_1, \tilde{l})$, $c(\tilde{s}_1, \tilde{s}_2)$, and a product of all the marginals, p_{s_1} , p_{s_2} , p_l . Copulas represent the low-level multivariate dependence for relationships expressed in the graph topology. Copulas are agnostic to the distribution of marginal variables, permitting construction of arbitrary joint relationships our method exploits for improved SLAM inference (*lower half*).

relationships between the latent variables that cannot be modeled with a linear function of covariances.

It stands to reason that by relaxing these assumptions, further improvements to accuracy could be realized. This is the main hypothesis our paper investigates; we introduce a family of factorized distributions whose dependency is modeled with a separate parametric function (Fig. 1) and argue that performing inference over this model class leads to improved accuracy and better representation of uncertainty.

The models we introduce exploit Sklar’s theorem [11], a foundational result from probability theory. Sklar’s theorem states that *any* joint distribution with a density, $p(X_1, \dots, X_N)$, can be expressed in terms of a copula, $c : [0, 1]^N \rightarrow [0, 1]$, and a product of the associated marginal distributions:

$$p(X_1, \dots, X_N) = c(\tilde{x}_1, \dots, \tilde{x}_N) \prod_{i=1}^N p_i(X_i). \quad (1)$$

^{*} Equal contribution, ¹ Stevens Institute of Technology, ² Massachusetts Institute of Technology, Correspondence to jmart13@stevens.edu




Fig. 2: **Copulas can give rise to many distributions:** We show three distributions parameterized by different Gaussian copulas. From left to right, the correlation is -0.9 , 0 , and 0.9 . The marginals are all the same: a two-component Gaussian mixture in x and a unimodal Gaussian in y .

Here $\tilde{x}_i = F_i(x_i)$ is a univariate variable, output from the marginal cumulative distribution function (CDF) $F_i(X_i)$. Equation 1 tells us that a joint probability can be decoupled into $N \in \mathbb{N}$ independent marginals $p_i(X_i)$ and a copula, $c(\cdot)$, to model the shared multivariate dependencies.

Our paper shows how copulas provide a powerful mechanism to fuse together a set of any marginal distributions, and how they permit SLAM algorithms to reason about probabilistic dependence and marginal distributions separately (Fig. 2). The main contributions of our paper are as follows.

A new class of models for SLAM: We introduce a rich class of distributional models based on a copula factorization of the SLAM posterior. Models express a hierarchy of the joint relationships between latent variables, and they can capture non-Gaussian distributions while still remaining fully factorized for efficient inference.

A variational filter for non-Gaussian SLAM: We introduce a new algorithm, called VC-SMC (Variational Copula Sequential Monte Carlo), that learns the parameters of a copula-factorized distribution model. VC-SMC is based on a particle filter. This sequentially generates data used to fit the variational parameters, making it possible to perform SLAM when the data has complex dependencies and the model parameters are unknown.

II. PRELIMINARIES

A. Simultaneous Localization And Mapping (SLAM)

We consider a standard observation process for navigation: at each step $t \in \mathbb{N}$ a mobile robot in state¹ $\mathbf{s}_t \in \mathcal{S}$ transitions to a state $\mathbf{s}_{t+1} \sim p(\cdot|\mathbf{s}_{1:t})$ and detects some set of $M_t \in \mathbb{N}$ landmarks², $\ell_{1:M_t}$, from the observation $\mathbf{z}_t \sim p(\cdot|\mathbf{s}_{1:t}, \ell_{1:M_t})$. The goal is to infer the robot's history of states and the landmark locations given the full sequence of observations. We aggregate the landmarks and robot state into the single latent variable $\mathbf{x}_t = (\mathbf{s}_t, \ell_{1:M_t})$. The joint distribution of this process is given by $p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}) =$

$$p(\mathbf{x}_1)p(\mathbf{z}_1|\mathbf{x}_1) \prod_{t=2}^T p(\mathbf{x}_t|\mathbf{x}_{1:t-1})p(\mathbf{z}_t|\mathbf{x}_{1:t}, \mathbf{z}_{1:t-1}). \quad (2)$$

¹In general states can be from any vector space. This work considers $\mathcal{S} = \text{SE}(2)$, for vehicles that operate in the two-dimensional plane.

²Each landmark ℓ is represented with a vector of spatial coordinates: two for planar navigation problems and three for 3D navigation.

This general model captures many common state-space systems, Hidden Markov Models (HMM), and non-Markov models, such as Gaussian Processes [12] and those represented by recurrent neural networks [13].

B. Sequential Monte Carlo

We consider Sequential Monte Carlo (SMC) methods to solve the SLAM problem (Fig. 3). SMC methods approximate the posterior with a weighted set of N sampled trajectories $p(\mathbf{x}_{1:T}|\mathbf{z}_{1:T}) \approx \frac{1}{N} \sum_{n=1}^N w_T^{[n]} \delta(\mathbf{x}_{1:T}^{[n]})$, which are drawn from a simpler proposal distribution $\{\mathbf{x}_{1:T}^{[n]}\}_{n=1}^N \sim q(\mathbf{x}_{1:T}|\mathbf{z}_{1:T})$ [14]. The joint distribution of latent variables and observations is given in (2). Although any distribution can be used for the proposal, it is beneficial to choose one with an autoregressive structure:

$$q(\mathbf{x}_{1:T}|\mathbf{z}_{1:T}) = q(\mathbf{x}_1|\mathbf{z}_1) \prod_{t=2}^T q(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{z}_{1:t}).$$

This makes it possible to decompose the complete proposal into T conditional distributions. The normalized importance weights of the n -th sample are defined as

$$w_t(\mathbf{x}_{1:t}^{[n]}) = \frac{p(\mathbf{x}_{1:t}^{[n]}|\mathbf{z}_{1:t})}{q(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]}, \mathbf{z}_{1:t})}.$$

They are computed with the following approximation:

$$w_t(\mathbf{x}_{1:t}^{[n]}) \approx \frac{\tilde{w}_t(\mathbf{x}_{1:T}^{[n]})}{\sum_{n=1}^N \tilde{w}_t(\mathbf{x}_{1:T}^{[n]})}, \quad (3)$$

$$\tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) = \tilde{w}_{t-1}(\mathbf{x}_{1:t-1}^{[n]}) \cdot \xi(\mathbf{x}_{1:t}^{[n]}, \mathbf{z}_{1:t}), \quad (4)$$

$$\xi(\mathbf{x}_{1:t}^{[n]}, \mathbf{z}_{1:t}) = \frac{p(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]})p(\mathbf{z}_t|\mathbf{x}_{1:t}^{[n]}, \mathbf{z}_{1:t-1})}{q(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]}, \mathbf{z}_{1:t})}. \quad (5)$$

The weighted set of particles is constructed sequentially. At time $t = 1$, we use standard importance sampling [14]. Beyond that, each step begins by resampling ancestor index variables $a_{t-1}^{[n]} \in \{1, \dots, N\}$ with probability proportional to the importance weights (3). Next, new values are proposed by drawing samples from $q(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[a_{t-1}^{[n]}]}, \mathbf{z}_{1:t})$. The new samples are appended to the end of the trajectory, $\mathbf{x}_{1:t}^{[n]} = (\mathbf{x}_{1:t-1}^{[a_{t-1}^{[n]}]}, \mathbf{x}_t^{[n]})$, and weights are updated with (3).

The accuracy of SMC methods critically relies on the proposal distribution. One common choice, known as the *Bootstrap Particle Filter (BPF)* [15], uses the transition model for the proposal:

$$q_{\text{BPF}}(\mathbf{x}_t|\mathbf{x}_{1:t-1}) = p(\mathbf{x}_t|\mathbf{x}_{1:t-1}). \quad (6)$$

The BPF has been employed in robotics settings [16], and is known to be inaccurate in high dimensions and when there are few particles available to sample. To remedy this issue, we propose a learning procedure that adjusts parameters of the proposal distribution, which is a copula-factorized model.




Fig. 3: **Sequential Monte Carlo sampling:** These methods construct a weighted set of particles whose sequences form trajectories. Circle size is proportional to the weight w_t . Arrows denote ancestor links, and the black path denotes the chosen trajectory. Our procedure selects a trajectory based on its final weight size.

C. Copula models

We consider proposal distribution models whose multivariate dependency is modeled with a copula. Different copula models give rise to different types of multivariate relationships. One common model we use is the Gaussian copula, which is given in terms of the zero-mean multivariate CDF $\Phi_\Sigma(\cdot)$ and its inverse $\Phi_\Sigma^{-1}(\cdot)$:

$$C_\Sigma(\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N) = \Phi_\Sigma(\Phi^{-1}(\tilde{\mathbf{x}}_1), \dots, \Phi^{-1}(\tilde{\mathbf{x}}_N)). \quad (7)$$

Here Σ is the covariance matrix, and $\tilde{\mathbf{x}}_i = \Phi_\Sigma(\mathbf{x}_i)$. Notice the transformed variables $\tilde{\mathbf{x}}_n$ are uniformly distributed, since CDFs map their inputs to the unit interval. Figure 2 shows how the Gaussian copula can give rise to many complex distributions which are not Gaussian themselves. Other copula models that focus on two-dimensional relationships are the Student-t, Clayton, Gumbel, Frank, and Joe copulas [17]. Copula models can be applied in a hierarchy, and their parameters can be learned with variational inference [18].

III. COPULA-FACTORIZED DISTRIBUTIONS FOR SLAM

In this section we introduce a class of distributions for SLAM, aimed to model the joint relationship among latent variables with a hierarchy of copula models. We apply Sklar's theorem (1) to factor the full SLAM posterior into several copulas and a product of univariate marginals. The copula hierarchy we present is one of several plausible dependency structures that could be useful for estimation. We propose four time-invariant relationships intended to promote parameter sharing. Other structures could be chosen to trade off computational constraints for model accuracy. Throughout the section we denote univariate variables with superscript parentheses, $\mathbf{x}^{(i)}$, and the copula parameters relating a and b are $\theta(a, b)$. For instance, the parameters of a Gaussian copula are the elements of the correlation matrix.

a) States and landmarks: The first relationship we consider is between the states and landmarks. We apply the decomposition of (1) by factoring these two variables and introducing a copula function $c_{\theta(s, \ell)}$ to model their joint relationship:

$$p(\mathbf{s}_t, \ell_{1:M_t} | \mathbf{z}_{1:t}) = c_{\theta(s, \ell)}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) p(\mathbf{s}_t | \mathbf{z}_{1:t}) p(\ell_{1:M_t} | \mathbf{z}_{1:t}).$$

In scenarios where perturbations in either s or ℓ cause variations to the other, dependency should be represented in the model. In such cases, $c_{\theta(s, \ell)}$ stands as a means to model that relationship.

b) Landmark dependency: Relationships commonly exist among landmarks when their positions are influenced by the same phenomena. For instance, two landmarks may be a part of the same dynamic structure. When they do, their joint distribution can be modeled as

$$p(\ell_{1:M_t} | \mathbf{z}_{1:t}) = c_{\theta(\ell, \ell)}(\tilde{\ell}_{1:M_t}) \prod_{j=1}^{M_t} p(\ell_j | \mathbf{z}_{1:t}).$$

This assumes homogeneity among the landmarks' relationships, so that dependence can be modeled with a common copula $c_{\theta(\ell, \ell)}$. Without homogeneity, there could be $\binom{M_t}{2}$ copulas in the worst case.

c) Component dependency: Another common source of dependency exists between the scalar components of the landmarks $c_{\theta(\ell)}$ and states $c_{\theta(s)}$. These relationships represent more familiar dependency encoded in the transition and observation models. Each model factorizes as follows

$$p(\ell_j | \mathbf{z}_{1:t}) = c_{\theta(\ell)}(\tilde{\ell}_t^{(1:d_\ell)}) \prod_{i=1}^{d_\ell} p(\ell_t^{(i)} | \mathbf{z}_{1:t}),$$

$$p(\mathbf{s}_t | \mathbf{z}_{1:t}) = c_{\theta(s)}(\tilde{\mathbf{s}}_t^{(1:d_s)}) \prod_{i=1}^{d_s} p(\mathbf{s}_t^{(i)} | \mathbf{z}_{1:t}).$$

d) Full factorization: The full copula hierarchy is based on four primary dependencies: states and landmarks $c_{\theta(s, \ell)}$, landmarks and landmarks $c_{\theta(\ell, \ell)}$, and component dependency $c_{\theta(s)}$ and $c_{\theta(\ell)}$. The complete structure is given by the product

$$c_{\theta}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) = c_{\theta}^{(1)}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) \cdot c_{\theta}^{(2)}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}), \quad (8)$$

$$c_{\theta}^{(1)}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) = c_{\theta(s, \ell)}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) c_{\theta(\ell, \ell)}(\tilde{\ell}_{1:M_t}),$$

$$c_{\theta}^{(2)}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) = c_{\theta(s)}(\tilde{\mathbf{s}}_t^{(1:d_s)}) \prod_{j=1}^{M_t} c_{\theta(\ell)}(\tilde{\ell}_j^{(1:d_\ell)}).$$

Given the dependency structure (8) and a fully-factorized distribution of univariate marginals, our new model class is defined as the set of all distributions $p(\mathbf{s}_t, \ell_{1:M_t} | \mathbf{z}_{1:t}) =$

$$c_{\theta}(\tilde{\mathbf{s}}_t, \tilde{\ell}_{1:M_t}) \prod_{i=1}^{d_s} p(\mathbf{s}_t^{(i)} | \mathbf{z}_{1:t}) \prod_{j=1}^{M_t} \left[\prod_{k=1}^{d_\ell} p(\ell_j^{(k)} | \mathbf{z}_{1:t}) \right]. \quad (9)$$

IV. VARIATIONAL COPULA SMC

In this section we introduce a variational SMC method that solves the SLAM problem using proposal distributions of the form (9). Each proposal $q(\mathbf{x}_t | \mathbf{x}_{1:t-1}; \boldsymbol{\lambda})$ replaces $p(\cdot | \mathbf{z}_{1:t})$ and is indexed by the variational parameters $\boldsymbol{\lambda} = \{\boldsymbol{\theta}, \boldsymbol{\eta}\}$, representing the copula parameters $\boldsymbol{\theta}$ and the marginal parameters $\boldsymbol{\eta}$. In what follows, we describe how to sample from the proposals and how to learn their variational parameters through gradient-based optimization.

A. A variational objective for fitting proposal distributions

Our method fits proposal parameters to a set of observations using the Variational SMC [19] framework. The objective is to minimize the Kullback Leibler (KL) divergence between the expected SMC approximation $\mathbf{E}[q(\mathbf{x}_{1:T}; \boldsymbol{\lambda})]$ and the target distribution $p(\mathbf{x}_{1:T}|\mathbf{z}_{1:T})$. This objective is intractable, but it can be upper-bounded by

$$\text{KL}(\mathbf{E}[q(\mathbf{x}_{1:T}; \boldsymbol{\lambda})] \parallel p(\mathbf{x}_{1:T}|\mathbf{z}_{1:T})) \leq -\mathbf{E} \left[\log \frac{\hat{Z}_T}{Z_T} \right].$$

The expectation here is taken with respect to the joint distribution of all variables generated by the SMC sampling procedure: $\phi(\mathbf{x}_{1:T}^{[1:N]}, a_{1:T}^{[1:N]}; \boldsymbol{\lambda}) =$

$$\prod_{n=1}^N q(\mathbf{x}_1^{[n]}; \boldsymbol{\lambda}) \cdot \prod_{t=2}^T \prod_{n=1}^N \rho_{t-1}^{[n]} q(\mathbf{x}_t | \mathbf{x}_{1:t-1}^{[n]}; \boldsymbol{\lambda}) \cdot \rho_T^{[n]},$$

where $\rho_t^{[n]} = \tilde{w}_t^{a_t^{[n]}} / \sum_j \tilde{w}_t^{a_t^{[j]}}$. We describe this sampling procedure in Alg. 2. The log-normalization constant is

$$\log \hat{Z}_T = \sum_{t=1}^T \log \left(\frac{1}{N} \sum_{n=1}^N \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) \right). \quad (10)$$

The dependence on the data and parameters enters the upper bound implicitly through the weights, proposed samples, and resampling procedure. Because Z_T does not depend on the variational parameters, it can be treated a constant in the optimization. Therefore, minimizing the upper bound (10) is equivalent to maximizing

$$\mathbf{E}[\log \hat{Z}_T] = \sum_{t=1}^T \mathbf{E}_{\phi_t} \left[\log \left(\frac{1}{N} \sum_{n=1}^N \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) \right) \right]. \quad (11)$$

In contrast to previous adaptive SMC methods that minimize the KL of the target directly with the proposal [13], this approach optimizes the fit of the final SMC distribution to the true target.

B. Computing gradients

We assume the proposals can be reparameterized such that sampling $\mathbf{x}_t \sim q(\mathbf{x}_t|\mathbf{x}_{1:t-1}; \boldsymbol{\lambda})$ is equivalent to evaluating the deterministic function $\mathbf{x}_t = f(\varepsilon_t, \mathbf{x}_{1:t-1}; \boldsymbol{\lambda})$ with a sample $\varepsilon_t \sim q(\varepsilon; \boldsymbol{\lambda})$ [20]. We then update the variational parameters using stochastic gradient steps, where the gradient is

$$\begin{aligned} \mathbf{g}_{\boldsymbol{\lambda}} &= \nabla_{\boldsymbol{\lambda}} \mathbf{E}[\log \hat{Z}_T], \\ &= \underbrace{\mathbf{E}_{q(\varepsilon)}[\log \hat{Z}_T \nabla_{\boldsymbol{\lambda}} \log q(\varepsilon; \boldsymbol{\lambda})]}_{\mathbf{g}_{\text{score}}} + \underbrace{\mathbf{E}_{q(\varepsilon)}[\nabla_{\boldsymbol{\lambda}} \log \hat{Z}_T]}_{\mathbf{g}_{\text{rep}}}. \end{aligned}$$

Here, $\mathbf{g}_{\text{score}}$ is the score function gradient and $\mathbf{g}_{\text{rep}}$ is the standard reparameterization gradient [21]. To keep the optimization tractable, we approximate $\mathbf{g}_{\boldsymbol{\lambda}} \approx \mathbf{g}_{\text{rep}}$. Similar to others [19], we found this to work well in practice. The

reparameterization gradient is derived as

$$\begin{aligned} \mathbf{g}_{\text{rep}} &= \sum_{t=1}^T \mathbf{E}_{\phi_t} \left[\nabla_{\boldsymbol{\lambda}} \log \left(\frac{1}{N} \sum_{n=1}^N \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) \right) \right], \\ &= \sum_{t=1}^T \mathbf{E}_{\phi_t} \left[\sum_{n=1}^N w_t(\mathbf{x}_{1:t}^{[n]}) \nabla_{\boldsymbol{\lambda}} \log \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) \right], \\ &\approx \hat{\mathbf{g}}_{\text{VSMC}}^{(k)} = \sum_{t=1}^T \sum_{n=1}^N w_t(\mathbf{x}_{1:t}^{[n]}) \nabla_{\boldsymbol{\lambda}} \log \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}). \quad (12) \end{aligned}$$

A derivation (12) is provided in the extended version of this paper [22]. Given an iterate $\boldsymbol{\lambda}_{k-1}$ we estimate the gradient by running SMC with our proposals $q(\mathbf{x}_t|\mathbf{x}_{1:t-1}; \boldsymbol{\lambda})$ and evaluate $\hat{\mathbf{g}}_{\text{VSMC}}^{(k)}$. The iterate is then updated with $\boldsymbol{\lambda}_k \leftarrow \boldsymbol{\lambda}_{k-1} + \alpha_k \hat{\mathbf{g}}_{\text{VSMC}}^{(k)}$, where the step sizes are scheduled to satisfy the Robbins Monroe conditions: $\sum_k \alpha_k = \infty$, $\sum_k \alpha_k^2 < \infty$.

The gradient $\nabla_{\boldsymbol{\lambda}} \log \tilde{w}_t(\mathbf{x}_{1:t}^{[n]})$ will change depending on the variational parameter and whether we consider the marginal parameters or the copula parameters. We have

$$\begin{aligned} \nabla_{\boldsymbol{\lambda}} \log \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) &= \frac{\nabla_{\boldsymbol{\lambda}} \tilde{w}_t(\mathbf{x}_{1:t}^{[n]})}{\tilde{w}_t(\mathbf{x}_{1:t}^{[n]})}, \\ \nabla_{\boldsymbol{\lambda}} \tilde{w}_t(\mathbf{x}_{1:t}^{[n]}) &= \frac{\gamma(\mathbf{x}_{1:t}^{[n]})}{[q(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]}; \boldsymbol{\lambda})]^2} \nabla_{\boldsymbol{\lambda}} q(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]}; \boldsymbol{\lambda}). \end{aligned}$$

Under the reparameterization and the dependency model, the optimization naturally splits into two problems:

$$\begin{aligned} \nabla_{\mathbf{x}} \log q(\mathbf{x}_t|\mathbf{x}_{1:t-1}; \boldsymbol{\lambda}) &= \nabla_{\mathbf{x}} \log c_{\boldsymbol{\theta}}(\tilde{\mathbf{x}}_{1:d}; \boldsymbol{\lambda}) \\ &\quad + \sum_{i=1}^d \nabla_{\mathbf{x}} \log q(\mathbf{x}_{i,t}|\mathbf{x}_{i,1:t-1}; \boldsymbol{\lambda}). \end{aligned}$$

When optimizing the copula parameters $\boldsymbol{\theta}$, the second term disappears, and both terms contribute to the gradient by optimizing the marginals. This leads to the following expectation-maximization procedure, which we call VC-SMC (Alg. 1).

VC-SMC holds the marginal parameters fixed while it optimizes the copula parameters. It then alternates: holding the copula parameters fixed to fit the marginals. The process repeats until the ELBO is maximized at convergence. After obtaining a good fit, we can sample from our proposals using Alg. 2 to solve the SLAM problem. In practice, we compute the gradients with automatic differentiation, reducing the model requirements and permitting the algorithm to be implemented as a library.

Algorithm 1 VC-SMC

```

1: input: Observations  $\mathbf{z}_{1:T}$ , Model  $p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T})$ , Proposals  $q(\mathbf{x}_t^{(i)}|\mathbf{x}_{1:t-1}; \lambda)$ , Copulas  $c_{\theta(\cdot, \cdot)}$ 
2: while  $\mathcal{L}(\theta, \eta)$  has not converged do
3:   # Learn the dependency structure
4:   while not converged do
5:      $\theta_{k+1} \leftarrow \theta_k + \alpha_k \hat{\mathbf{g}}_{\text{VSMC}, \theta}$ 
6:   # Learn the marginal structure
7:   while not converged do
8:      $\eta_{k+1} \leftarrow \eta_k + \alpha_k \hat{\mathbf{g}}_{\text{VSMC}, \eta}$ 
9:   # Sample from the VC-SMC posterior (Alg. 2)
10: output:  $\mathbf{x}_{1:T} \sim q(\mathbf{x}_{1:T}|\lambda)$ 

```

Algorithm 2 VC-SMC Posterior Sampling

```

1: input: Targets  $p(\mathbf{x}_{1:t}, \mathbf{z}_{1:t})$ , Proposals  $q(\mathbf{x}_t^{(i)}|\mathbf{x}_{1:t-1})$ , Copulas  $c_{\theta(\cdot, \cdot)}$ , number of particles  $N$ 
2: # Initialize trajectory particles
3: for  $n = 1, \dots, N$  do
4:    $\mathbf{x}_1^{[n]} \sim q(\mathbf{x}_1; \lambda)$ 
5:    $\tilde{w}_1^{[n]} \leftarrow \frac{p(\mathbf{x}_1^{[n]})p(\mathbf{z}_1|\mathbf{x}_1^{[n]})}{q(\mathbf{x}_1^{[n]}; \lambda)}$ 
6: for  $t = 1, \dots, T$  do
7:   for  $n = 1, \dots, N$  do
8:      $a_{t-1}^{[n]} \sim \text{Categorical}(N, \rho_{t-1}^{[j]})$ 
9:      $\mathbf{x}_t^{[n]} \sim q(\mathbf{x}_t|\mathbf{x}_{1:t-1}^{(a_{t-1}^{[n]})})$ 
10:     $\tilde{w}_t^{[n]} \leftarrow \frac{p(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]})p(\mathbf{z}_t|\mathbf{x}_{1:t}^{[n]}, \mathbf{z}_{1:t-1})}{q(\mathbf{x}_t^{[n]}|\mathbf{x}_{1:t-1}^{[n]}, \mathbf{z}_{1:t})}$ 
11:    $a_T \sim \text{Categorical}(N, \rho_T^{[j]})$ 
12: output:  $\mathbf{x}_{1:T}^{a_T}$ 

```

Algorithm 3 VC-SMC Proposal Sampling

```

1: input:  $\varepsilon, \lambda = \{\theta, \eta\}$ ,
2: # Sample from the reparameterized function
3:  $\mathbf{x}_t^{(i)} = f(\varepsilon, \mathbf{x}_{1:t-1}; \eta) \forall i = 1, \dots, d$ 
4: output:  $\mathbf{x}_t = \mathbf{x}_t^{(1:d)}$ 

```

V. RELATED WORK

A. Variational Inference

Our variational algorithm (VC-SMC) is an extension of variational sequential Monte Carlo [19] to support the class of copula-factorized proposal distributions. Others have also considered hierarchical copula models [18], [23] in a variational setting to preserve multivariate dependence between latent variables. Normalizing flows are another related set of inference methods, where the parameters of a nonlinear function mapping samples from a standard normal distribution to the posterior of interest are learned [24]. More related methods fall in the area of *optimal transport* [25], [26], where again, parameters of a transformation between a ‘simple’ distribution and the posterior of interest are obtained. The application of these methods to non-Gaussian




Fig. 4: **VC-SMC can handle uncertain data association:** We show the KL divergence between the true posterior and pose distributions, and the root mean-squared error between ground truth landmark values and the posterior mean estimates of VC-SMC and the BPF. Data was gathered over 50 independent trials using 100 particles.

SLAM problems has yet to be realized and could serve as an interesting area for future work.

B. Non-Gaussian SLAM

The broad focus of non-Gaussian SLAM research has largely been on nonparametric methods [5], [10], [16], [27], [28]. Similar to our algorithm, these methods make no assumptions about the state or landmark distributions. However, their model classes are not able to separate the joint dependency from the marginal distributions in a way that permits decoupled modeling and inference. For example, the sum of Gaussians method [29] is a filtering-based approach that uses mixtures of coupled dependency and marginal behavior. The proposed copula decomposition enables the development of inference algorithms that flexibly incorporate a variety of assumptions about the marginals or dependency structure of the joint distribution, bridging the gap between fully parametric methods and nonparametric methods.

VI. EXPERIMENTS

In this section, we validate the main claims of our paper using experimental data gathered in simulation. We test if VC-SMC can improve inference accuracy in SLAM settings where uncertainty is non-Gaussian, and nonlinearities induce complex relationships between latent variables. Each experiment uses the Bootstrap Particle Filter (BPF) as its baseline (6). In each experiment, VC-SMC is trained using the Adam optimizer [30], with a fixed learning rate of $\alpha = 10^{-2}$. More details about each experiment are provided in the extended version of this paper [22].

A. Uncertain Data Association

This experiment tests the hypothesis that VC-SMC can perform accurate inference in the presence of uncertain data association, where there is multi-modal, non-Gaussian uncertainty in the robot pose. The problem is set within the classic 3Doors environment [31]. A robot moves along a single axis with constant velocity and observes the range to one of three doors serving as the navigation landmarks. Their absolute locations are $\ell_1 = 0$, $\ell_2 = 2$, $\ell_3 = 6$.

This setting is challenging, because the robot does not know from which door its measurement is associated, and the number of possibilities compound at each time step. With Gaussian uncertainty in the transitions and observations, the




Fig. 5: **Training progress in the 3Doors problem:** The variational loss (11) converges from randomly-initialized parameters at a reasonable rate: in approximately 400 out of the 1000 total training iterations.

exact posterior is a Gaussian mixture of three modes at $t = 1$, nine modes at $t = 2$, and twenty-seven modes at $t = 3$. To infer the latent variables and represent their uncertainty accurately, the robot must consider more than one hypothesis about its position and landmark locations.

We estimated the ground truth mean and variance parameters of the true posterior using a standard Bayes filter, and we compared the KL divergence between the pose distribution and estimates obtained with each test method. The root mean squared error was used for the unimodal landmark distributions. Distributions were parameterized with 100 particles. VC-SMC uses a three-component Gaussian mixture over the pose marginal and univariate Gaussians for the landmarks. These are linked through a Gaussian copula. The variational parameters of both methods were randomly initialized, and we trained them offline with 1000 gradient steps.

Figure 4 shows the distribution of root mean squared error, which was gathered from 100 independent trials. We can see that VC-SMC predictions have a significantly smaller error than the BPF method, which assumes a unimodal Gaussian proposal distribution. A shortcoming of the Gaussian proposal in the 3Doors environment is that it provides inadequate coverage of the true posterior support – assigning very little probability mass to regions of the latent space that potentially have high probability. In contrast, the VC-SMC proposal provides more modeling flexibility, which allows us to choose expressive models that provide better posterior coverage. Adjusting the mixture means is also straightforward through our variational algorithm.

In Fig. 6 we show qualitative results of the belief distribution at each step. We observe that even in this relatively simple problem BPF fails to accurately represent the full belief over robot poses and quickly loses modes of the posterior. This is due to the fact that the bootstrap proposal poorly covers the target distribution. In contrast, by optimizing the marginal and dependency parameters of non-Gaussian proposals VC-SMC is able to represent more modes in the posterior resulting from the data association uncertainty. In Figure 5 we show the variational loss (11) as it decreases throughout training. Based on the results of this section, we conclude that VC-SMC can support non-Gaussian inference in settings involving uncertain data association.

B. Nonlinear Planar Navigation

In this experiment we test the hypothesis that copula-factored distribution models can lead to improved inference accuracy in the presence of non-linearly distorted uncertainty. A robot wishes to localize its three-dimensional state over several time steps using measurements from the map origin, which are corrupted by additive Gaussian noise. States evolve according to the nonlinear transition model:

$$x_{t+1} = x_t + v \cos \theta_t, \quad y_{t+1} = y_t + v \sin \theta_t, \quad \theta_{t+1} = \theta_t.$$

The robot moves with constant speed v and heading θ . This model has been used to simulate marine vehicles localizing with an ultra-short baseline acoustic system [32].

In Fig. 7 we compare the total root mean-square error between the true reference trajectory and the estimates from both VC-SMC and BPF. We find that VC-SMC is able to achieve a lower error, since it is able to adjust the variational parameters at each time step, to fit the observed data. The plots represent data gathered over ten trials.

VII. DISCUSSION

This section concludes our paper with a brief discussion of the results, computational complexity, and possible avenues for future work.

A. What have we learned?

Starting from the observation that many SLAM problems are unrealistically approximated with Gaussian uncertainty, we asked if it was possible to improve predictive accuracy with more expressive models. To answer this, we introduced a new class of copula-factorized distributions that can represent any distribution with the added flexibility of separating multivariate dependency and marginal uncertainty into independent models. Indeed, these distributions were found to better represent domain uncertainty and lead to more accurate predictions compared to standard approaches when the model parameters were fit to maximize a lower bound to the log likelihood.

B. Computational Complexity

The per-iteration complexity of Algorithm 1 is dictated by the total number of parameters $|\lambda|$. With N particles evolving over T time steps, a parameter update requires $\mathcal{O}(TN)$ gradient computations with respect to each component in λ , as we show in the extended version of this paper [22]. Each call of Algorithm 1 computes $\mathcal{O}(TN|\lambda|)$ partial derivatives. Since the number of parameters $|\lambda|$ per time-step is roughly constant in typical applications, and the number of particles N is fixed a priori, we expect that the per-iteration computational complexity of Algorithm 1 is linear in time, in terms of the number of gradient computations.

C. Future work

While the VC-SMC procedure permits inference with arbitrary copula models and non-Gaussian univariate marginals, the modeling freedom of the copula factorization lends itself to many new avenues for future work. Specifically,




Fig. 6: **Time evolution of the approximate target distribution in the 3Doors environment:** We show histograms of the three landmarks ℓ_i and robot pose at times $t = 1, 2, 3$. The BPF fails to capture the multi-modal nature of the state distribution, and under-represents the posterior support. Each histogram is comprised of 500 particles.




Fig. 7: **Comparison of MAP estimates for Nonlinear Planar Navigation:** We compare the root mean-squared error of the predictions from BPF and VC-SMC. The variational method is able to achieve a superior fit since it can adjust the marginal and copula parameters.




Fig. 8: **Training progress in the Nonlinear Planar Navigation problem:** The variational loss (11) converges from randomly-initialized parameters at a reasonable rate, in approximately 80 out of the 1000 total training iterations.

the factorization permits separate assumptions to be made about the joint dependency parameters than are made about the univariate marginal distributions. This is in contrast to assumptions of full joint Gaussianity. In this work, a particle-based inference method was chosen to preserve as much modeling flexibility as possible, but this came at the expense of efficiency. By developing algorithms tailored to a specific subset of copula models (for example, by considering specifically Gaussian copula models) or selection of univariate marginals, more efficient procedures for inference may be developed. Nonparanormal belief propagation [33] is one example of such an existing method. Furthermore, leveraging data structures like the Bayes tree [4] is of interest for improving the scalability of copula-based approaches.

D. Conclusion

In this work we sought to move beyond the Gaussian paradigm, to improve filtering accuracy in settings where states and landmarks have complex probabilistic dependence. We were able to exploit a powerful result from probability theory that factorizes joint relationships into marginal distributions and dependency models, called copulas. This copula factorization was used in a variational SMC filter (VC-SMC), where samples from the SLAM posterior were approximated, and parameters were fit through gradient based optimization.

Our experimental results support the hypothesis that VC-SMC indeed improves inference accuracy when there is uncertain data association and when there was nonlinear propagation of uncertainty. We believe our results highlight the importance of modeling multivariate dependency for SLAM, as well as exemplify the benefits that copulas can offer the SLAM paradigm as a whole.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their helpful feedback. Additionally we would like to thank Qiangqiang Huang for his comments on an early draft. This research was supported in part by the National Science Foundation, grant number IIS-1652064.

REFERENCES

- [1] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," *Autonomous robot vehicles*, pp. 167–193, 1990.
- [2] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Autonomous Robots*, vol. 4, no. 4, pp. 333–349, 1997.
- [3] S. Thrun and M. Montemerlo, "The GraphSLAM algorithm with applications to large-scale mapping of urban structures," *The International Journal on Robotics Research (IJRR)*, vol. 25, no. 5/6, pp. 403–430, 2005.
- [4] M. Kaess, V. Ila, R. Roberts, and F. Dellaert, "The Bayes tree: An algorithmic foundation for probabilistic robot mapping," *Algorithmic Foundations of Robotics IX*, pp. 157–173, 2010.
- [5] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with Rao-Blackwellized particle filters," *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.
- [6] R. M. Eustice, H. Singh, and J. Leonard, "Exactly sparse delayed-state filters," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2417–2424, 2005.
- [7] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental smoothing and mapping," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008.
- [8] D. M. Rosen, M. Kaess, and J. Leonard, "Robust incremental online inference over sparse factor graphs: Beyond the Gaussian case," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1025–1032, 2013.
- [9] K. Doherty, D. Fourie, and J. Leonard, "Multimodal semantic SLAM with probabilistic data association," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2419–2425, 2019.
- [10] D. Fourie, J. Leonard, and M. Kaess, "A nonparametric belief solution to the Bayes tree," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2189–2196, 2016.
- [11] A. Sklar, "Fonctions de répartition à n dimensions et leurs marges," *Publications de l'Institut de Statistique de l'Université de Paris*, 1959.
- [12] R. Frigola, Y. Chen, and C. E. Rasmussen, "Variational Gaussian process state-space models," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3680–3688, 2014.
- [13] S. S. Gu, Z. Ghahramani, and R. E. Turner, "Neural adaptive sequential Monte Carlo," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 2629–2637, 2015.
- [14] A. Doucet, N. De Freitas, and N. Gordon, "An introduction to sequential monte carlo methods," in *Sequential Monte Carlo methods in practice*. Springer, 2001, pp. 3–14.
- [15] N. J. Gordon, D. J. Salmond, and A. F. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings F - Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, 1993.
- [16] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," *American Association for Artificial Intelligence (AAAI) Conference on Artificial Intelligence*, 2002.
- [17] R. B. Nelson, *An Introduction to Copulas*, 2nd ed. Springer, 2006.
- [18] D. Tran, D. Blei, and E. M. Airoldi, "Copula variational inference," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015, pp. 3564–3572.
- [19] C. Naesseth, S. Linderman, R. Ranganath, and D. Blei, "Variational sequential Monte Carlo," *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- [20] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *International Conference on Learning Representations (ICLR)*, 2013.
- [21] F. Ruiz, M. Titsias, and D. Blei, "The generalized reparameterization gradient," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2016, pp. 460–468.
- [22] J. D. Martin, K. Doherty, C. Cyr, B. Englot, and J. Leonard, "Variational filtering with copula models for SLAM," *arXiv preprint*, 2020.
- [23] S. Han, X. Liao, D. B. Dunson, and L. Carin, "Variational Gaussian copula inference," *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [24] D. J. Rezende and S. Mohamed, "Variational inference with normalizing flows," *arXiv preprint arXiv:1505.05770*, 2015.
- [25] T. A. El Moselhy and Y. M. Marzouk, "Bayesian inference with optimal maps," *Journal of Computational Physics*, vol. 231, no. 23, pp. 7815–7850, 2012.
- [26] M. Pulido and P. J. van Leeuwen, "Sequential Monte Carlo with kernel embedded mappings: the mapping particle filter," *Journal of Computational Physics*, vol. 396, pp. 400–415, 2019.
- [27] D. Fourie, "Multi-modal and inertial sensor solutions for navigation-type factor graphs," Ph.D. dissertation, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution, 2017.
- [28] A. I. Eliazar and R. Parr, "DP-SLAM 2.0," *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, pp. 1314–1320, 2004.
- [29] H. Durrant-Whyte, S. Majumder, S. Thrun, M. de Battista, and S. Scheding, "A Bayesian algorithm for simultaneous localisation and map building," in *Robotics Research*. Springer, 2003, pp. 49–60.
- [30] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [31] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [32] J. D. Martin and B. Englot, "Extending model-based policy gradients for robots in heteroscedastic environments," *Conference on Robot Learning (CoRL)*, pp. 438–447, 2017.
- [33] G. Elidan and C. Cario, "Nonparanormal belief propagation (NPNBP)," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 899–907, 2012.