

Decentralised Gradient-based Variational Inference for Multi-sensor Fusion and Tracking in Clutter

Qing Li[§], Runze Gan[§], Simon Godsill

Engineering Department

University of Cambridge

Cambridge, UK

{ql289, rg605, sjg30}@cam.ac.uk

Abstract—This paper investigates the task of tracking multiple objects in clutter under a distributed multi-sensor network with time-varying connectivity. Designed with the same objective as the centralised variational multi-object tracker, the proposed method achieves optimal decentralised fusion in performance with local processing and communication with only neighboring sensors. A key innovation is the decentralised construction of a locally maximised evidence lower bound, which greatly reduces the information required for communication. Our decentralised natural gradient descent variational multi-object tracker, enhanced with the gradient tracking strategy and natural gradients that adjusts the direction of traditional gradients to the steepest, shows rapid convergence. Our results verify that the proposed method is empirically equivalent to the centralised fusion in tracking accuracy, surpasses suboptimal fusion techniques with comparable costs, and achieves much lower communication overhead than the consensus-based variational multi-object tracker.

Index Terms—distributed sensor fusion, multiple object tracking, data association, variational inference, decentralised gradient descent, natural gradient, decentralised optimisation

I. INTRODUCTION

Multi-sensor fusion can greatly improve object tracking in challenging environments, whilst the ideal centralised fusion is often impractical due to communication constraints or dynamic networks e.g., disruptions and link failures. Several optimal distributed fusion implementations have been developed, including a decentralised Kalman filter [1] and track-to-track fusion techniques [2], [3]. However, they either require specific network topology [1] or cross-correlations information among sensors, thus being restricted to use in practice.

In light of this, two suboptimal fusion rules, generalised covariance intersection (GCI) [3] and arithmetic average (AA) [4] were proposed and integrated into existing trackers, e.g., probability hypothesis density (PHD) filter [5], and the multi-Bernoulli (MB) filter [6]. Typically, each sensor employs a local tracker to generate multi-object distributions, which are then fused using either GCI or AA fusion rules. In comparison, AA fusion is shown to have lower computational costs since GCI fusion normally has no closed form [7]. Additionally, AA is shown to be more robust and outperformed the GCI fusion in low detection probability environments [7]. To perform in a fully distributed manner, consensus algorithms [8], [9] have been introduced to GCI and AA fusion, allowing them

to function efficiently in larger sensor networks. Nonetheless, these approaches are suboptimal and are expected to lead to reduced sensor fusion and tracking performance.

In response, [10] devised an optimal distributed sensor fusion solution for tracking multiple objects in cluttered environments under non-homogeneous Poisson process (NHPP) measurement model [11]. Specifically, it is based on the variational multi-object tracker (VT) presented in [12], [13], which has demonstrated superior performance over leading tracking algorithms [14]–[17] in single-sensor fixed object number scenarios. This method enables sensors to operate independently using their local measurements while collaborating with neighboring sensors to assimilate global statistics through an average consensus algorithm [8]. It is verified to achieve tracking precision on par with centralised fusion. Despite its benefits, it may incur high communication costs, as sensors must undergo multiple runs of the consensus algorithm at each variational update iteration.

Therefore, this paper develops a more advanced decentralised gradient-based VT method that is grounded in optimal fusion, while allowing each sensor to work independently without awaiting consensus during variational inference iterations. A key concept is the locally maximised evidence lower bound (LM-ELBO) introduced in [18], [19] under different names, though it has seen limited discussion and application afterwards. For our tracking and sensor fusion tasks, we derive that this LM-ELBO can be decomposed to form a decentralised optimisation problem, leading to much more communication-efficient updates than those using the original ELBO. Note that our work is conceptually different from the existing method in [20], which directly decomposes the original global lower bound into a set of local lower bounds and, as pointed out by authors in [20], uses stochastic gradient methods for optimisation without theoretical performance analysis. Here we first establish the equivalence of optimising LM-ELBO and the original ELBO, by which we develop decentralised optimisation solutions for maximising LM-ELBO. Our framework is thus more theoretically robust and versatile, enabling the use of emerging techniques from the decentralised optimisation field alongside our proposed decentralised gradient-based methods. In particular, we propose an efficient decentralised natural gradient descent VT (DeNG-VT) algorithm that employs the natural gradient [21]

[§]Runze Gan and Qing Li contributed equally to this work.

and a gradient tracking strategy [22] to accelerate convergence. Simulation results demonstrate that our proposed DeNG-VT method is empirically equivalent to centralised fusion in tracking performance, requiring much less communication cost compared to [10] and maintaining comparable costs to suboptimal fusion techniques.

II. PROBLEM FORMULATION AND MODELLING

Assume that there are K objects in the surveillance area. At each time step n , their joint state is $X_n = [X_{n,1}^\top, X_{n,2}^\top, \dots, X_{n,K}^\top]^\top$, where each vector $X_{n,k}, k \in \{1, \dots, K\}$ denotes the kinematic state for the k -th object. Suppose that objects are observed by a sensor network consisting of N_s sensors, each capable of observing the entire tracking area. The time-varying sensor network at time t can be modelled as a graph $\mathcal{G}(t) = \{\mathcal{S}, \mathcal{E}(t)\}$ at any given continuous time t , where the sensor set is denoted by $\mathcal{S} = \{1, 2, \dots, N_s\}$, and $\mathcal{E}(t)$ is the set of edges with the existence of edge (i, j) meaning that the i -th sensor can communicate with the j -th sensor at time t . The set of neighbours of sensor i is denoted by $\mathcal{N}_i(t) = \{j \mid (i, j) \in \mathcal{E}(t)\}$. The degree $d_i(t)$ of the i -th sensor represents the number of its neighbouring sensors with which it can communicate, i.e., $d_i(t) = |\mathcal{N}_i(t)|$. In a sensor network, the measurements received from all sensors at time step n can be denoted by $Y_n = [Y_n^1, Y_n^2, \dots, Y_n^{N_s}]$. Each Y_n^s includes measurements acquired by the s -th sensor, and $Y_n^s = [Y_{n,1}^s, \dots, Y_{n,M_n^s}^s]$, where M_n^s is the total number of measurements received at the s -th sensor ($s = 1, \dots, N_s$). Subsequently, $M_n = [M_n^1, \dots, M_n^{N_s}]$ records the total number of measurements received from all sensors at time step n .

A. Dynamical Model

We assume that objects move in a 2D surveillance area with each $X_{n,k} = [x_{n,k}^1, \dot{x}_{n,k}^1, x_{n,k}^2, \dot{x}_{n,k}^2]^\top$, where $x_{n,k}^d$ and $\dot{x}_{n,k}^d$ ($d = 1, 2$) indicate the k -th object's position and velocity in the d -th dimension, respectively. We assume an independent linear Gaussian transition density for each object's states:

$$p(X_n|X_{n-1}) = \prod_{k=1}^K \mathcal{N}(X_{n,k}; F_{n,k}X_{n-1,k}, Q_{n,k}). \quad (1)$$

where $F_{n,k} = \text{diag}(F_{n,k}^1, F_{n,k}^2)$, $Q_{n,k} = \text{diag}(Q_{n,k}^1, Q_{n,k}^2)$.

B. NHPP Measurement Model and Association Prior

Here, we assume each sensor independently detects objects in accordance with the NHPP measurement model in [10], [11]. Notably, the NHPP model may vary for each sensor. Denote the set of Poisson rates for all sensors as $\Lambda = [\Lambda^1, \Lambda^2, \dots, \Lambda^{N_s}]$. For each sensor s , the Poisson rate vector is defined by $\Lambda^s = [\Lambda_0^s, \Lambda_1^s, \dots, \Lambda_K^s]$, where Λ_0^s is the clutter rate and Λ_k^s is the k -th object rate, $k = 1, \dots, K$.

Our independent measurement model assumption signifies that given X_n , the measurements of each sensor are conditionally independent, i.e., $p(Y_n|X_n) = \prod_{s=1}^{N_s} p(Y_n^s|X_n)$. We denote the associations of all measurements Y_n by $\theta_n = [\theta_n^1, \theta_n^2, \dots, \theta_n^{N_s}]$, with each $\theta_n^s = [\theta_{n,1}^s, \theta_{n,2}^s, \dots, \theta_{n,M_n^s}^s]$ ($s = 1, \dots, N_s$) representing the association vector for the s -th sensor's measurements. Each component $\theta_{n,j}^s$ ($j = 1, \dots, M_n^s$)

gives the origin of the measurement $Y_{n,j}^s$; $\theta_{n,j}^s = 0$ indicates that $Y_{n,j}^s$ is generated by clutter, and $\theta_{n,j}^s = k$ ($k = 1, \dots, K$) means that $Y_{n,j}^s$ is generated from the object k . The adopted conditionally independent NHPP model leads to the following properties according to [12]. First, $p(Y_n, \theta_n|X_n, M_n) = p(Y_n|\theta_n, X_n)p(\theta_n|M_n)$. Both joint association prior and joint likelihood are conditionally independent across sensors, i.e., $p(\theta_n|M_n) = \prod_{s=1}^{N_s} p(\theta_n^s|M_n^s)$, $p(Y_n|\theta_n, X_n) = \prod_{s=1}^{N_s} p(Y_n^s|\theta_n^s, X_n)$. Lastly, measurements are conditionally independent given associations and states

$$p(Y_n^s|\theta_n^s, X_n) = \prod_{j=1}^{M_n^s} \ell^s(Y_{n,j}^s|X_{n,\theta_{n,j}^s}^s), \quad (2)$$

where M_n^s is implicitly known from θ_n^s since $M_n^s = |\theta_n^s|$, and ℓ^s is the probability density function of a single measurement received in sensor s given its originator's state. Here we assume a linear and Gaussian model for object originated measurements and clutter measurements to be uniformly distributed in the observation area of volume V^s :

$$\ell^s(Y_{n,j}^s|X_{n,k}) = \begin{cases} \mathcal{N}(HX_{n,k}, R_k^s), & k \neq 0; \quad (\text{object}) \\ 1/V^s, & k = 0; \quad (\text{clutter}) \end{cases} \quad (3)$$

where H is the observation matrix, and R_k^s indicates the s -th sensor noise covariance. Moreover, the joint prior $p(\theta_n^s|M_n^s)$ can be factorised as the product of M_n^s independent association priors, i.e., $p(\theta_n^s|M_n^s) = \prod_{j=1}^{M_n^s} p(\theta_{n,j}^s)$, where $p(\theta_{n,j}^s)$ is a categorical distribution with $\theta_{n,j}^s \in \{0, \dots, K\}$

$$p(\theta_{n,j}^s) = \frac{1}{\sum_{k=0}^K \Lambda_k^s} \sum_{k=0}^K \Lambda_k^s \delta[\theta_{n,j}^s = k]. \quad (4)$$

III. VARIATIONAL FILTERING FOR MULTI-SENSOR FUSION

The objective is to sequentially estimate the posterior $p(X_n, \theta_n|Y_{1:n})$ at time step n given observations $Y_{1:n}$ from all sensors. Accordingly, the exact optimal filtering is recursively expressed as the following prediction and update steps:

$$p(X_n|Y_{1:n-1}) = \int p(X_n|X_{n-1})p(X_{n-1}|Y_{1:n-1})dX_{n-1},$$

$$p(X_n, \theta_n|Y_{1:n}) \propto p(Y_n|\theta_n, X_n)p(\theta_n|M_n)p(X_n|Y_{1:n-1}),$$

where $p(X_n|Y_{1:n})$ and $p(X_n, \theta_n|Y_{1:n})$ are the predictive prior and posterior, respectively. The parameters K, Λ , and $R_{1:K}^s$ in Section II are assumed to be known. Since this exact filtering recursion is intractable, here we use the variational filtering [12] to perform this task. Particularly, we replace $p(X_{n-1}|Y_{1:n-1})$ with the converged variational distribution $q_{n-1}^*(X_{n-1})$ obtained by variational inference [23] at time step $n-1$, and thus the predictive prior $\hat{p}_n(X_n)$ is written as

$$\hat{p}_n(X_n) = \int p(X_n|X_{n-1})q_{n-1}^*(X_{n-1})dX_{n-1}. \quad (5)$$

The target posterior of current approximate filtering step n is

$$\hat{p}_n(X_n, \theta_n|Y_n) \propto p(Y_n|\theta_n, X_n)p(\theta_n|M_n)\hat{p}_n(X_n). \quad (6)$$

A. Variational Inference and Evidence Lower Bound

Our objective is to find variational distribution $q_n^*(X_n, \theta_n)$ to approximate the target distribution $\hat{p}_n(X_n, \theta_n|Y_n)$ in (6). We first posit a mean-field assumption $q_n(X_n, \theta_n) = q_n(X_n)q_n(\theta_n)$. Then, we choose the variational distribution

$q_n^*(X_n, \theta_n)$ from the posited family that minimises the Kullback–Leibler (KL) divergence, which is equivalent to maximising the evidence lower bound (ELBO), defined as

$$\mathcal{F}(q_n) = \mathbb{E}_{q_n(X_n)q_n(\theta_n)} \log \frac{p(Y_n|\theta_n, X_n)p(\theta_n|M_n)\hat{p}_n(X_n)}{q_n(X_n)q_n(\theta_n)} \quad (7)$$

In the following, we discuss several possible ways to find the variational distribution q_n that maximises the objective ELBO $\mathcal{F}(q_n)$ using variational inference.

To maximise the objective ELBO $\mathcal{F}(q_n)$, in [10], we applied the centralised coordinate ascent variational inference (CAVI) method in a fully decentralised manner using an average consensus algorithm [8], which is in theory equivalent to the centralised version. See [10] for more details. Despite that, it requires a fully converged average consensus at each CAVI iteration. Therefore, this paper will develop a more communicationally efficient decentralised variational inference without waiting for consensus.

IV. LOCALLY MAXIMISED ELBO

This section introduces the locally maximised ELBO (LM-ELBO) as an alternative objective to the conventional ELBO in (7) to enhance optimisation efficiency. We will see that this is particularly beneficial in our decentralised sensor fusion.

A. Problem Setting and Background

Here we first set up the general problem of interest. Consider inferring two disjoint multivariate variables X, θ given Y , where the exact posterior $p(X, \theta|Y)$ is intractable but can be evaluated up to a constant, i.e., $p(X, \theta|Y) \propto f(X, \theta, Y)$ and the unnormalised posterior $f(X, \theta, Y)$ is computable. We seek to approximate $p(X, \theta|Y)$ with $q(X)q(\theta)$ through mean-field variational inference [23], where the objective is to find $q(X)$ and $q(\theta)$ that maximise the ELBO [23], defined as

$$\mathcal{F}(q(X), q(\theta)) := \mathbb{E}_{q(X)q(\theta)} \log \frac{f(X, \theta, Y)}{q(X)q(\theta)}, \quad (8)$$

If one variational distribution, e.g., $q(\theta)$, is allowed to take any form, then it has a unique global optimiser for \mathcal{F} while fixing $q(X)$ [23]:

$$q^*(\theta) \propto \exp(E_{q(X)} \log f(X, \theta, Y)), \quad (9)$$

$$\mathcal{F}(q(X), q^*(\theta)) = \max_{q(\theta)} \mathcal{F}(q(X), q(\theta)), \quad (10)$$

with the maximisation spanning all distribution forms of $q(\theta)$.

If we further assume the distribution forms of $q(X), q(\theta)$, and denote their respective governing parameters by vectors λ, ρ , then the ELBO in (8) can be reformulated as a fixed-form ELBO, denoted as $\mathcal{F}(\lambda, \rho)$:

$$\mathcal{F}(\lambda, \rho) := \mathbb{E}_{q(X;\lambda)q(\theta;\rho)} \log \frac{f(X, \theta, Y)}{q(X;\lambda)q(\theta;\rho)}. \quad (11)$$

This fixed-form ELBO enables more conventional optimisation techniques, such as gradient descent, particularly useful when the standard CAVI update like (9) is intractable.

However, optimising the fixed-form ELBO in (11) can become inefficient with high-dimensional parameters λ, ρ . This inefficiency motivates the exploration of alternative objectives with fewer parameters. The literature on variational inference has introduced two such objectives: the locally maximised

ELBO [19], [24] and the KL-corrected (KLC) bound (or marginalised variational bound) [18], [25], [26]. Despite their conceptual similarity and independent influence on subsequent research, to our knowledge, a discussion on their connections is absent. Our study finds that, compared to the original LM-ELBO in [19], the KLC bound in [18] is more convenient to implement and offers additional advantages that justify its use. Consequently, our definition of LM-ELBO is closely adhere the definition of the KLC bound in [18]. A detailed review of these two objectives will be presented in our full paper.

B. LM-ELBO and Its Properties

1) *Definition and Assumption:* Assume ρ is the parameter to eliminate from (11). Here we define LM-ELBO $\mathcal{L}(\lambda)$ by the fixed-form ELBO in (11), where $q(\theta; \rho)$ is replaced by the optimal distribution from the free-form update in (9), i.e.,

$$\mathcal{L}(\lambda) := \mathbb{E}_{q(X;\lambda)q^*(\theta)} \log \frac{f(X, \theta, Y)}{q(X;\lambda)q^*(\theta)}, \quad (12)$$

$$q^*(\theta) \propto \exp(E_{q(X;\lambda)} \log f(X, \theta, Y)). \quad (13)$$

Note that $q^*(\theta)$ is also a function of λ although not explicitly included in its argument. The LM-ELBO $\mathcal{L}(\lambda)$ in (12) is a well-defined function even if $q^*(\theta)$ cannot be evaluated in a closed-form. When f represents the exact joint density $p(X, \theta, Y)$, the expression of $\mathcal{L}(\lambda)$ in (12) can be shown to be identical to the original KL-corrected bound definition, i.e., the equation (4) of the [18], through mathematical manipulation.

Our LM-ELBO $\mathcal{L}(\lambda)$ will be used as an alternative objective to the conventional fixed-form ELBO $\mathcal{F}(\lambda, \rho)$ in (11). To establish their relationship, we will assume the distribution form of $q(\theta; \rho)$, which is employed in $\mathcal{F}(\lambda, \rho)$ in (11), encompasses the optimal distribution $q^*(\theta)$ in (13). We then denote $\rho^*(\lambda)$ as the parameter value (or one among several) that reproduces the $q^*(\theta)$ in (13) with λ held fixed, i.e., $q(\theta; \rho^*(\lambda)) = q^*(\theta)$.

2) *Properties:* Here we present properties of LM-ELBO; detailed analysis and proof are given in full paper. First, the definitions and assumptions in Section IV-B1 directly result in

$$\mathcal{L}(\lambda) = \mathcal{F}(\lambda, \rho = \rho^*(\lambda)) = \max_{\rho} \mathcal{F}(\lambda, \rho), \quad (14)$$

These properties, as will be shown in full paper, play a key role in offering simpler and more intuitive derivations of existing properties in [18], [19], and in establishing new properties that justify the use of LM-ELBO.

The second property offers a useful computational simplification, previously demonstrated in [18], [19]. Specifically,

$$\nabla_{\lambda} \mathcal{L}(\lambda) = \nabla_{\lambda} \mathcal{F}(\lambda, \rho)|_{\rho=\rho^*(\lambda)} \quad (15)$$

This property simplifies $\nabla_{\lambda} \mathcal{L}(\lambda)$ to the partial derivative of the fixed-form ELBO (11), where $q(\theta; \rho) = q^*(\theta)$ is treated as λ -independent during gradient evaluation.

3) *Validating LM-ELBO with Optimality Properties:* We introduce a novel optimality alignment property, validating our LM-ELBO \mathcal{L} as an equally reasonable alternative objective:

If λ^ is a global maximum, a local maximum, or a stationary point of $\mathcal{L}(\lambda)$, then $[\lambda^*, \rho^*(\lambda^*)]$ is, respectively, a global maximum, a local maximum, or a stationary point of $\mathcal{F}(\lambda, \rho)$.*

This demonstrates that any optimum found by optimising $\mathcal{L}(\lambda)$ is inherently an optimum within the original ELBO $\mathcal{F}(\lambda, \rho)$, thereby verifying the optimisation of our LM-ELBO.

V. DECENTRALISED NATURAL GRADIENT DESCENT VARIATIONAL INFERENCE FOR UPDATE STEP

A. The Rule of Decentralised Gradient Descent

First, we present a brief introduction of the fundamental DGD strategy [27], [28] for the decentralised optimisation problem: N_s sensors cooperatively minimise $f(x) = \sum_{s=1}^{N_s} f_s(x)$, where $x \in \mathbb{R}^p$ and each f_s is only known to sensor s . The DGD algorithm employs consensus ideas for estimating the gradient of the global objective function $\nabla f(x)$. Specifically, the update rule for each sensor s at iteration i is:

$$x^s(i+1) = \sum_{j=1}^{N_s} w_{sj} x^j(i) + \alpha \nabla_{x^s} f_s(x^s(i)), \quad (16)$$

where α is the stepsize, w_{sj} is nonzero only if s and j are neighbours or $s = j$ and the matrix $W = [w_{sj}] \in \mathbb{R}^{N_s \times N_s}$ is symmetric and doubly stochastic [28]. Each sensor s updates its local variable x^s by combining the average of its neighbours' with a local gradient $\alpha \nabla f_s(x^s)$. This DGD method has guaranteed convergence for both convex and non-convex functions; after convergence, all sensor reach the same solution, and the solution is the stationary point of $f(x)$ under diminishing step sizes [29].

The motivation for developing the LM-ELBO in Section IV-A now becomes evident. For the considered tracking tasks, directly applying the DGD update (16) to the fixed-form ELBO in (11) would force sensors to share extensive high-dimensional data association information, since the optimisation variable would include parameters of $q_n(\theta_n)$. We will see that the LM-ELBO can be optimised in a decentralised manner, enabling the exchange of only object state information, thus greatly reducing communication overhead.

B. Decentralisation of LM-ELBO for Multi-sensor Fusion

Define λ_n and η_n as parameters of the variational distribution $q_n(X_n; \lambda_n)$ and the prior $\hat{p}_n(X_n; \eta_n)$, respectively. Using the definition in (12), the LM-ELBO's specific form for our task follows from the ELBO in (7) and models in Section II:

$$\mathcal{L}(\lambda_n) = \sum_{s=1}^{N_s} \mathbb{E}_{q_n(X_n; \lambda_n) q_n^*(\theta_n)} \log p(Y_n^s | \theta_n^s, X_n) + \mathbb{E}_{q_n(X_n; \lambda_n)} \log \frac{\hat{p}_n(X_n; \eta_n)}{q_n(X_n; \lambda_n)} + \mathbb{E}_{q_n^*(\theta_n)} \log \frac{p(\theta_n | M_n)}{q_n^*(\theta_n)} \quad (17)$$

Assume $q_n(X_{n,k}; \lambda_{n,k}) = \mathcal{N}(X_{n,k}; \mu_{n|n}^k, \Sigma_{n|n}^k)$, $k = 1, \dots, K$. The optimal distribution $q_n^*(\theta_n)$ in (17) is computed by CAVI update in (13), where $q_n^*(\theta_n) \propto \prod_{s=1}^{N_s} \prod_{j=1}^{M_n^s} q_n^*(\theta_{n,j}^s)$, and each $\theta_{n,j}^s$ is updated independently and in parallel as follows

$$q_n^*(\theta_{n,j}^s) \propto \frac{\Lambda_0^s}{V^s} \delta[\theta_{n,j}^s = 0] + \sum_{k=1}^K \Lambda_k^s l_k^s \delta[\theta_{n,j}^s = k], \quad (18)$$

$$l_k^s = \mathcal{N}(Y_{n,j}^s; H \mu_{n|n}^k, R_k^s) \exp(-0.5 \text{Tr}((R_k^s)^{-1} H \Sigma_{n|n}^k H^\top)),$$

Note that $q_n^*(\theta_{n,j}^s)$ is also a function of λ_n .

To formulate a consensus optimisation task compatible with many existing decentralised algorithms, we decentralise $\mathcal{L}(\lambda_n)$ in (17) into a sum of local LM-ELBOs $\mathcal{L}_s(\lambda_n)$ at s -th sensor

$$\mathcal{L}(\lambda_n) = \sum_{s=1}^{N_s} \mathcal{L}_s(\lambda_n) \quad (19)$$

$$\mathcal{L}_s(\lambda_n) = \mathbb{E}_{q_n(X_n; \lambda_n) q_n^*(\theta_n^s)} \log p(Y_n^s | \theta_n^s, X_n) + \mathbb{E}_{q_n^*(\theta_n^s)} \log \frac{p(\theta_n^s | M_n^s)}{q_n^*(\theta_n^s)} + \frac{1}{N_s} \mathbb{E}_{q_n(X_n; \lambda_n)} \log \frac{\hat{p}_n(X_n; \eta_n)}{q_n(X_n; \lambda_n)} \quad (20)$$

This design transforms it into a decentralised optimisation problem, where each local $\mathcal{L}_s(\lambda_n)$ depends only on local data Y_n^s , and thus computations with $\mathcal{L}_s(\lambda_n)$ (e.g., gradients) can be performed fully locally. Thus, it enables the application of numerous established decentralised optimisation algorithms from the growing field to collectively optimise $\mathcal{L}_s(\lambda_n)$.

C. Decentralised Natural Gradient Descent for $\mathcal{L}(\lambda_n)$

Natural gradients scales the gradient with the inverse of its Fisher Information Matrix (FIM), $G(\lambda_n) = \mathbb{E}_{\lambda_n}[(\nabla_{\lambda_n} \ln q_n(X_n | \lambda_n)) (\nabla_{\lambda_n} \ln q_n(X_n | \lambda_n))^\top]$, providing a direction of steepest ascent that is more aligned with the underlying statistical manifold, which has been demonstrated to enhance convergence over traditional gradients [19], [21]. Hence, we propose a decentralised natural gradient descent (DNGD) scheme where local sensors collaboratively solve the optimisation task of maximising the LM-ELBO $\mathcal{L}(\lambda_n) = \sum_{s=1}^{N_s} \mathcal{L}_s(\lambda_n)$ in (20). Subsequently, we replace the standard gradient of DGD method in Section V-A with the natural gradient. Thus, the update equation at each iteration i at each sensor s for jointly optimising the LM-ELBO $\mathcal{L}(\lambda_n)$ is

$$\lambda_n^s(i+1) = \sum_{j=1}^{N_s} w_{sj}(i) \lambda_n^j(i) + \alpha \hat{\nabla}_{\lambda_n^s} \mathcal{L}_s(\lambda_n^s(i)) \quad (21)$$

$$\hat{\nabla}_{\lambda_n^s} \mathcal{L}_s(\lambda_n^s(i)) = G(\lambda_n^s)^{-1} \nabla_{\lambda_n^s} \mathcal{L}_s(\lambda_n^s(i)) \quad (22)$$

where the weight $w_{sj}(i)$ is chosen as the Metropolis weight in [9]. Note that $w_{sj}(i)$ depends on the connectivity of our time-varying sensor network $\mathcal{G}(t)$ at different time t .

1) *Computing the local natural gradient $\hat{\nabla}_{\lambda_n^s} \mathcal{L}_s(\lambda_n^s)$:* For the convenience of computing natural gradients, we rewrite prior $\hat{p}_n(X_n; \eta_n^s) = \prod_{k=1}^K \hat{p}_n(X_{n,k}; \eta_{n,k}^s)$ and the variational distribution $q_n(X_n; \lambda_n^s) = \prod_{k=1}^K q_n(X_{n,k}; \lambda_{n,k}^s)$ at s -th sensor into canonical exponential family distributions [30]

$$q_n(X_{n,k}; \lambda_{n,k}^s) = (2\pi)^{-\frac{d}{2}} \exp(\lambda_{n,k}^s T(X_{n,k}) - A(\lambda_{n,k}^s)) \quad (23)$$

$$\hat{p}_n(X_{n,k}; \eta_{n,k}^s) = (2\pi)^{-\frac{d}{2}} \exp(\eta_{n,k}^s T(X_{n,k}) - A(\eta_{n,k}^s)) \quad (24)$$

where $A(\cdot)$ is the log partition function, $\eta_{n,k}^s$ and $\lambda_{n,k}^s$ are natural parameters and their mappings to Gaussian forms are

$$\lambda_{n,k}^s = \begin{bmatrix} \lambda_{n,k}^{s,1} \\ \lambda_{n,k}^{s,2} \end{bmatrix} = \begin{bmatrix} (\Sigma_{n|n}^{k,s})^{-1} \mu_{n|n}^{k,s} \\ -\frac{1}{2} (\Sigma_{n|n}^{k,s})^{-1} \end{bmatrix} \quad (25)$$

$$\eta_{n,k}^s = \begin{bmatrix} \eta_{n,k}^{s,1} \\ \eta_{n,k}^{s,2} \end{bmatrix} = \begin{bmatrix} (\Sigma_{n|n-1}^{k*,s})^{-1} \mu_{n|n-1}^{k*,s} \\ -\frac{1}{2} (\Sigma_{n|n-1}^{k*,s})^{-1} \end{bmatrix} \quad (26)$$

where prior $\hat{p}_n(X_{n,k}; \eta_{n,k}^s) = \mathcal{N}(X_{n,k}; \mu_{n|n-1}^{k*,s}, \Sigma_{n|n-1}^{k*,s})$ is computed using (5) and the converged variational distribution $q_{n-1}^*(X_{n-1,k}) = \mathcal{N}(X_{n-1,k}; \mu_{n-1|n-1}^{k*,s}, \Sigma_{n-1|n-1}^{k*,s})$:

$$\mu_{n|n-1}^{k*,s} = F_{n,k} \mu_{n-1|n-1}^{k*,s}, \quad (27)$$

$$\Sigma_{n|n-1}^{k*,s} = F_{n,k} \Sigma_{n-1|n-1}^{k*,s} F_{n,k}^\top + Q_{n,k}. \quad (28)$$

Note that the definition of $\hat{p}_n(X_n; \eta_n^s)$ in (24) assumes that consensus is reached at all time steps; thus, all sensors have

Algorithm 1: DeNG-VT at time step n for each sensor s

Input: $q_{n-1}^*(X_{n-1}; \lambda_{n-1}^{s,*}), Y_n^s$, maximum iteration I_{max} .
Output: $q_n^*(X_n; \lambda_n^{s,*}) = \prod_{k=1}^K q_{n,k}^*(X_n, k; \lambda_{n,k}^{s,*})$
for $k = 1, 2, \dots, K$ **do**
 Prediction step: $\hat{p}_n(X_n, k) = \mathcal{N}(X_n, k; \mu_{n|n-1}^{k*,s}, \Sigma_{n|n-1}^{k*,s})$ using (27)
Initialisation: $q_n(X_n, k; \lambda_{n,k}^s), k = 1, 2, \dots, K$:
 $\lambda_{n,k}^{s,1}(0) = (\Sigma_{n|n-1}^{k*,s})^{-1} \mu_{n|n-1}^{k*,s}, \lambda_{n,k}^{s,2}(0) = -\frac{1}{2}(\Sigma_{n|n-1}^{k*,s})^{-1}$
 $g_{n,k}^{s,1}(0) = \hat{\nabla}_{\lambda_{n,k}^s} \mathcal{L}_s(\lambda_{n,k}^s(0)), g_{n,k}^{s,2}(0) = \hat{\nabla}_{\lambda_{n,k}^s} \mathcal{L}_s(\lambda_{n,k}^s(0))$
for $i = 0, 1, \dots, I_{max}$ **do**
 For $j = 1, \dots, M_n$, compute $q_n^{s,*}(\theta_{n,j})$ (defined below (30)).
 for $k = 1, 2, \dots, K$ **do**
 Update $\lambda_{n,k}^{s,1}(i+1), \lambda_{n,k}^{s,2}(i+1)$ according to (31)
 Update $\hat{g}_{n,k}^{s,1}(i+1), \hat{g}_{n,k}^{s,2}(i+1)$ according to (32).
 $q_{n,k}^*(X_n, k; \lambda_{n,k}^{s,*}) = \mathcal{N}(X_n, k; \mu_{n|n}^{k*,s}, \Sigma_{n|n}^{k*,s})$
 where $\mu_{n|n}^{k*,s} = -\frac{1}{2}(\lambda_{n,k}^{s,2})^{-1} \lambda_{n,k}^{s,1}, \Sigma_{n|n}^{k*,s} = -\frac{1}{2}(\lambda_{n,k}^{s,2})^{-1}$

the same prior $\hat{p}_n(X_n; \eta_n)$ as stated in (20) since all $\{\eta_n^s\}_{s=1}^{N_s}$ are equal. For limited iteration cases please see Section V-C3.

Several useful properties are applied when calculating the natural gradients. First, for $q_n(X_n; \lambda_n^s)$ defined in (24), $G(\lambda_n^s)$ equals to the second derivative of the log normaliser, that is, $G(\lambda_n^s) = \nabla_{\lambda_n^s}^2 A(\lambda_n^s)$. We can also derive a similar property as in (15) to bypass the laborious gradient calculations related to λ_n^s -dependent $q_n^*(\theta_n^s)$.

Subsequently, we calculate the local natural gradient of local LM-ELBO with respect to natural parameters $\lambda_{n,k}^{s,1}$ and $\lambda_{n,k}^{s,2}$:

$$\hat{\nabla}_{\lambda_{n,k}^{s,1}} \mathcal{L}_s(\lambda_{n,k}^s) = H^\top (R_k^s)^{-1} \sum_{j=1}^{M_n^s} Y_{n,j}^s q_n^{s,*}(\theta_{n,j}^s = k) \quad (29)$$

$$+ \frac{1}{N_s} \left[\eta_{n,k}^{s,1} - \lambda_{n,k}^{s,1} + (\Sigma_{n|n-1}^{k*,s})^{-1} \mu_{n|n-1}^{k*,s} - (\Sigma_{n|n}^{k*,s})^{-1} \mu_{n|n}^{k*,s} \right]$$

$$\hat{\nabla}_{\lambda_{n,k}^{s,2}} \mathcal{L}_s(\lambda_{n,k}^s) = -\frac{1}{2} H^\top (R_k^s)^{-1} H \sum_{j=1}^{M_n^s} q_n^{s,*}(\theta_{n,j}^s = k)$$

$$+ \frac{1}{N_s} (\eta_{n,k}^{s,2} - \lambda_{n,k}^{s,2}) + \frac{1}{2N_s} [(\Sigma_{n|n}^{k*,s})^{-1} - (\Sigma_{n|n-1}^{k*,s})^{-1}] \quad (30)$$

where $q_n^{s,*}(\theta_{n,j}^s)$ is the distribution in (18) with local parameter λ_n^s . It has the form of (18) with $\mu_{n|n}^k$ and $\Sigma_{n|n}^k$ replaced by $\mu_{n|n}^{k*,s}$ and $\Sigma_{n|n}^{k*,s}$ at sensor s . Here we can see that natural gradients are easy to calculate, as they bypass the need for calculating the inverse of the FIMs, which are effectively cancelled out. Full derivations will be given in full paper.

2) *Gradient tracking strategy for speeding up convergence:* Here we adopt a gradient tracking strategy in [22] to speed up the convergence of DNGD algorithm since it is shown to have a notably rapid convergence speed. It is particularly advantageous as it ensures convergence with a constant step size [29], simplifying the tuning process for practical applications.

The gradient tracking strategy relies on tracking differences of gradients. To use it in our settings, for each sensor s and each iteration i , we will update both the variational parameters $\lambda_{n,k}^{s,1}(i)$ and $\lambda_{n,k}^{s,2}(i)$, and an additional gradient estimate $\hat{g}_{n,k}^{s,1}(i)$ and $\hat{g}_{n,k}^{s,2}(i)$. In this setting, the update equations for each variational parameter $\lambda_{n,k}^{s,m}(i)$ and gradient estimate $\hat{g}_{n,k}^{s,m}(i)$, $m = 1, 2$ are as follows




Fig. 1: Simulation scenarios; the left figure is a sensor network at a single time step; red circles are sensors and grey lines indicate connectivity. The right figure is ground-truth tracks (blue lines) and their initial positions (blue dots); grey dots are measurements received at one time step at one sensor.

$$\lambda_{n,k}^{s,m}(i+1) = \sum_{j=1}^{N_s} w_{sj}(i) \lambda_{n,k}^{j,m}(i) + \alpha \hat{g}_{n,k}^{s,m}(i), m = 1, 2 \quad (31)$$

$$\hat{g}_{n,k}^{s,m}(i+1) = \sum_{j=1}^{N_s} w_{sj}(i) \hat{g}_{n,k}^{j,m}(i) + \hat{\nabla}_{\lambda_{n,k}^s} \mathcal{L}_s(\lambda_{n,k}^s(i+1))$$

$$- \hat{\nabla}_{\lambda_{n,k}^s} \mathcal{L}_s(\lambda_{n,k}^s(i)), m = 1, 2. \quad (32)$$

Finally, the full procedure for DeNG-VT with the gradient tracking technique can be seen in Algorithm 1.

3) *Robust decentralised tracking: explainable performance in limited iterations:* Ideally, achieving consensus in prior steps ensures identical prior $\hat{p}_n(X_n; \eta_n^s)$ across different sensors at time step n . When using limited DNGD iterations for efficiency, sensors may hold unique priors $\hat{p}_n(X_n; \eta_n^s)$. Even so, our approach remains interpretable since it still maximises LM-ELBO in (20) with $\hat{p}_n(X_n; \eta_n)$ being replaced by an effective prior, which is the geometric average (GA) fusion of individual sensors' priors, $\hat{p}_{eff}(X_n) \propto \prod_{s=1}^{N_s} \hat{p}_n(X_n; \eta_n^s)^{1/N_s}$. In DeNG-VT, this GA fusion occurs automatically without extra processing steps. This contrasts with traditional GA fusion approaches, e.g., [3], which necessitate separate consensus algorithms to implement a fully distributed GA fusion rule.

VI. RESULTS

This section studies the communication efficiency and tracking performance of the proposed DeNG-VT method under a time-varying multi-sensor network. We present a detailed comparison with several methods, including individual VT with no sensor fusion (I-VT), (optimal) centralised VT that receives all measurements from all sensors in [10] (C-VT), distributed consensus-based VT proposed in [10] (DeC-VT). To verify the advantage of the proposed DeNG-VT algorithm, we also included the suboptimal distributed VT with the AA fusion strategy in [31] (DeAA-VT), where each sensor infers a multi-object posterior distribution using the variational tracker in [12] based on local measurements, which is then fused with other posteriors from sensors with a distributed average consensus algorithm under the AA fusion rule.

To evaluate the tracking performance, we use the generalised optimal sub pattern assignment (GOSPA) [32], where the order $p = 1$, $\alpha = 2$, and the cut-off distance $c = 50$. Concurrently, GOSPA metric returns localisation errors for




Fig. 2: GOSPA over iteration number at a single time step ($T=10$); black dotted line represents the performance of the optimal centralised solution C-VT, and colored lines correspond to the performance of 20 local sensor nodes of DeNG-VT




Fig. 3: Mean GOSPA over 50 time steps; for all methods, lines are means of GOSPA averaged over all sensors and shaded areas indicate ± 1 standard deviation.

well-tracked objects, the missed object errors and false object errors. Note that here we have a fixed number of objects in the scene; thus, the missed and false object errors denote the track loss rather than the disappearance or appearance of objects. We define a MGOSPA metric, which is the mean GOSPA averaged over all sensors and all time steps. To show the communication cost, we define the communication iteration (CI), that is, the total iteration number that sensors pass messages to its neighbours at a time step, averaged over total time steps and Monte Carlo runs. Specifically, for DeNG-VT, CI equals to the total DNGD iterations at each time step n ; For DeC-VT in [10], CI equals to the total variational update iterations at each time step n multiplies the consensus algorithm iterations at each variational update iteration. For the compared suboptimal DeAA-VT, CI equals to total iterations of consensus algorithm performed at one time step.

A. Simulation and Parameter Settings

We simulate a sensor network consisting of 20 nodes shown as in Figure 1, and their connectivity is randomly generated at each time step. All sensors observe the same surveillance area where 50 objects move under the constant velocity model with $F_{n,k}^d = \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix}$, $Q_{n,k}^d = 25 \begin{bmatrix} \tau^3/3 & \tau^2/2 \\ \tau^2/2 & \tau \end{bmatrix}$ ($d = 1, 2$), and ground-truth tracks are shown in Figure 1. The total time steps are 50, and the time interval between observations is $\tau = 1$. We follows the NHPP measurement model in Section II-B with $R_k^s = 100I$, where I is a 2-D identity matrix. For each sensor, the object Poisson rates are 1 and the clutter rate is 500. An example measurement data of one sensor (sensor node 1) at one time step is shown in Figure 1.

TABLE I: Performance of compared methods over 50 runs

method	MGOSPA	location	missed	false	CI
C-VT	144.6 ± 1	144.6 ± 1	0 ± 0	0 ± 0	—
I-VT	1502 ± 19	370 ± 6	566 ± 12	566 ± 12	—
DeAA-VT ₁	443.8 ± 16	423.9 ± 6	9.9 ± 8	9.9 ± 8	20
DeAA-VT ₂	443.4 ± 16	423.0 ± 6	10.2 ± 8	10.2 ± 8	100
DeC-VT	144.6 ± 1	144.6 ± 1	0 ± 0	0 ± 0	1000
DeNG-VT ₁	155.0 ± 1	155.0 ± 1	0 ± 0	0 ± 0	20
DeNG-VT ₂	145.4 ± 1	145.4 ± 1	0 ± 0	0 ± 0	50
DeNG-VT ₃	144.6 ± 1	144.6 ± 1	0 ± 0	0 ± 0	100

B. Result Analysis

1) Tracking and fusion performance at one example run:

Figure 2 shows the convergence speed of the proposed DeNG-VT method, where we plot the GOSPA values of all local sensors over the DNGD iterations at a single time step in a Monte Carlo run. It is observed that the GOSPA values for all nodes exhibit a pronounced decrease within the initial 10 iterations, indicating a rapid convergence towards accurate estimations. Additionally, despite the different convergence speeds among sensors, they all approach the performance of the centralised C-VT solution after around 50 iterations, which empirically demonstrates its equivalent performance to C-VT.

Figure 3 illustrates mean GOSPA with its one standard deviation over 50 time steps for each compared methods. Here we set the DNGD iteration for DeNG-VT to 50, according to the convergence performance in Figure 2. The variational iteration number for DeC-VT, DeAA-VT, C-VT, I-VT is 20. The average consensus iteration for DeC-VT and DeAA-VT are 50 and 100, respectively. We can see that for all compared methods except I-VT, the mean GOSPA has zero standard deviation at each time step, which means that 20 sensor nodes converge to the same values and thus demonstrate the trackers can converge to a local optimum. It can be seen that the discrepancy of mean GOSPA between the suboptimal DeAA-VT and towards-optimal solutions (DeNG-VT and DeC-VT) are large, verifying that DeNG-VT and DeC-VT having a much better tracking accuracy. Most importantly, Figure 3 empirically demonstrates the equivalence of centralised C-VT and two decentralised solutions DeNG-VT and DeC-VT at every time step with regards to tracking performance.

2) *Tracking and fusion performance over all 50 runs:* We verify the robustness of the proposed method by testing it over 50 Monte Carlo runs with different measurement sets under the settings and ground-truth tracks in Section VI-A. Table I shows the performance of the compared methods in both tracking accuracy and communication efficiency. The average consensus iterations for DeC-VT and two configurations of DeAA-VTs are 50, 20, 100. The variational inference iterations for DeC-VT, DeAA-VT, C-VT, and I-VT are 20. The DNGD iterations for three configurations of DeNG-VTs are 20, 50, and 100. We record in Table I the mean and one standard deviation of MGOSPA and its submetric (location error, missed object and false object error), averaged over 50 runs.

We can see that C-VT, DeC-VT, and all versions of DeNG-VTs show very accurate tracking, with MGOSPA scores around 144.6 with minimal deviation. In contrast, the tracking

accuracy of I-VT and 2 versions of suboptimal DeAA-VTs is much lower. It is clear to see that I-VT MGOSPA errors mainly come from the track loss, and the MGOSPA errors of DeAA-VTs are mainly due to location inaccuracy. The estimation results also confirm the equivalence of the proposed DeNG-VT with the centralised C-VT solution when it converges.

With regards to communication costs, we can see from CI values the great improvement of the proposed DeNG-VT compared with the DeC-VT, under the same optimal tracking accuracy. Compared to the suboptimal DeAA-VT method, we can see that our method still greatly outperforms DeAA-VT in tracking accuracy even using the same communication iteration number, which showcases its advantages in both tracking accuracy and communication efficiency.

VII. CONCLUSION

This paper presents a distributed implementation of tracking multiple objects in cluttered environments under a time-varying multi-sensor network. This approach matches the centralised fusion tracking performance, surpasses the traditional suboptimal distributed VT strategies in tracking accuracy, and demonstrates significant reductions in communication costs when compared to existing average consensus VT methods. It is noted that our framework can cope with emerging decentralised optimisation algorithms, where user can select the most updated version to refine the inference performance. Future work will extend it to accommodate unknown object numbers and multimodal sensors with varying coverage.

Acknowledgements: This research is sponsored by the US Army Research Laboratory and the UK MOD University Defence Research Collaboration (UDRC) in Signal Processing under the SIGNeTS project. It is accomplished under Cooperative Agreement Number W911NF-20-2-0225. The views and conclusions in this document are of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, the MOD, the U.S. Government or the U.K. Government. The U.S. Government and U.K. Government are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

- [1] B. Rao, H. F. Durrant-Whyte, and J. Sheen, "A fully decentralized multi-sensor system for tracking and surveillance," *The International Journal of Robotics Research*, vol. 12, no. 1, pp. 20–44, 1993.
- [2] C.-Y. Chong, "Distributed multitarget multisensor tracking," *Multitarget-multisensor tracking: Advanced applications*, pp. 247–296, 1990.
- [3] R. P. Mahler, "Optimal/robust distributed data fusion: a unified approach," in *Signal Processing, Sensor Fusion, and Target Recognition IX*, vol. 4052. SPIE, 2000, pp. 128–138.
- [4] A. K. Gostar, R. Hoseinnezhad, and A. Bab-Hadiashar, "Cauchy-Schwarz divergence-based distributed fusion with poisson random finite sets," in *2017 International Conference on Control, Automation and Information Sciences (ICCAIS)*. IEEE, 2017, pp. 112–116.
- [5] M. Ueney, D. E. Clark, and S. J. Julier, "Distributed fusion of PHD filters via exponential mixture densities," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 521–531, 2013.
- [6] B. Wang, W. Yi, R. Hoseinnezhad, S. Li, L. Kong, and X. Yang, "Distributed fusion with multi-Bernoulli filter based on generalized covariance intersection," *IEEE Transactions on Signal Processing*, vol. 65, no. 1, pp. 242–255, 2016.
- [7] T. Li, X. Wang, Y. Liang, and Q. Pan, "On arithmetic average fusion and its application for distributed multi-Bernoulli multitarget tracking," *IEEE Transactions on Signal Processing*, vol. 68, pp. 2883–2896, 2020.
- [8] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on automatic control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [9] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Fourth International Symposium on Information Processing in Sensor Networks*. IEEE, 2005, pp. 63–70.
- [10] Q. Li, R. Gan, and S. Godsill, "Consensus-based distributed variational multi-object tracker in multi-sensor network," in *2023 Sensor Signal Processing for Defence Conference (SSPD)*. IEEE, 2023, pp. 1–5.
- [11] K. Gilholm, S. Godsill, S. Maskell, and D. Salmond, "Poisson models for extended target and group tracking," in *Signal and Data Processing of Small Targets 2005*, vol. 5913, 2005.
- [12] R. Gan, Q. Li, and S. Godsill, "A variational Bayes association-based multi-object tracker under the non-homogeneous Poisson measurement process," in *25th International Conference on Information Fusion (FUSION)*. IEEE, 2022, pp. 1–8.
- [13] R. Gan, Q. Li, and S. J. Godsill, "Variational tracking and redetection for closely-spaced objects in heavy clutter," *IEEE Transactions on Aerospace and Electronic Systems*, 2024.
- [14] F. Meyer and M. Z. Win, "Scalable data association for extended object tracking," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 6, pp. 491–507, 2020.
- [15] Q. Li, J. Liang, and S. Godsill, "Scalable data association and multi-target tracking under a Poisson mixture measurement process," in *IEEE ICASSP*, 2022, pp. 5503–5507.
- [16] K. Granström, M. Fatemi, and L. Svensson, "Poisson multi-Bernoulli mixture conjugate prior for multiple extended target filtering," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 1, 2019.
- [17] Q. Li, R. Gan, J. Liang, and S. J. Godsill, "An adaptive and scalable multi-object tracker based on the non-homogeneous Poisson process," *IEEE Transactions on Signal Processing*, 2023.
- [18] J. Hensman, M. Rattray, and N. Lawrence, "Fast variational inference in the conjugate exponential family," *Advances in neural information processing systems*, vol. 25, 2012.
- [19] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *Journal of Machine Learning Research*, 2013.
- [20] J. Hua and C. Li, "Distributed variational bayesian algorithms over sensor networks," *IEEE Transactions on Signal Processing*, vol. 64, no. 3, pp. 783–798, 2015.
- [21] S.-I. Amari, "Natural gradient works efficiently in learning," *Neural computation*, vol. 10, no. 2, pp. 251–276, 1998.
- [22] A. Nedic, A. Olshevsky, and W. Shi, "Achieving geometric convergence for distributed optimization over time-varying graphs," *SIAM Journal on Optimization*, vol. 27, no. 4, pp. 2597–2633, 2017.
- [23] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *Journal of the American statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.
- [24] D. Durante and T. Rigon, "Conditionally conjugate mean-field variational Bayes for logistic models," *Statistical Science*, vol. 34, no. 3, pp. 472–485, 2019.
- [25] N. J. King and N. D. Lawrence, "Fast variational inference for Gaussian process models through KL-correction," in *17th European Conference on Machine Learning*. Springer, 2006, pp. 270–281.
- [26] M. Lázaro-Gredilla, S. Van Vaerenbergh, and N. D. Lawrence, "Overlapping mixtures of Gaussian processes for the data association problem," *Pattern recognition*, vol. 45, no. 4, pp. 1386–1395, 2012.
- [27] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [28] J. Zeng and W. Yin, "On nonconvex decentralized gradient descent," *IEEE Transactions on signal processing*, vol. 66, no. 11, 2018.
- [29] T.-H. Chang, M. Hong, H.-T. Wai, X. Zhang, and S. Lu, "Distributed learning in the nonconvex world: From batch data to streaming and beyond," *IEEE Signal Processing Magazine*, 2020.
- [30] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [31] T. Li and F. Hlawatsch, "A distributed particle-PHD filter using arithmetic-average fusion of gaussian mixture parameters," *Information Fusion*, vol. 73, pp. 111–124, 2021.
- [32] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *2017 20th International Conference on Information Fusion*. IEEE, 2017, pp. 1–8.