

Semi-Supervised Model-Free Bayesian State Estimation from Compressed Measurements

Anubhab Ghosh ^{*}, Yonina C. Eldar [†], Saikat Chatterjee ^{*}

^{*} Digital Futures Centre and School of Elect. Engg. & Comp. Sc., KTH Royal Institute of Technology, Sweden

[†] Faculty of Mathematics and Computer Science, The Weizmann Institute of Science, Israel

^{*} anubhabg@kth.se, [†] yonina.eldar@weizmann.ac.il, ^{*} sach@kth.se

Abstract—We consider data-driven Bayesian state estimation from compressed measurements (BSCM) of a model-free process. The dimension of the temporal measurement vector is lower than that of the temporal state vector to be estimated, leading to an under-determined inverse problem. The underlying dynamical model of the state’s evolution is unknown for a ‘model-free process.’ Hence, it is difficult to use traditional model-driven methods, for example, Kalman and particle filters. Instead, we consider data-driven methods. We experimentally show that two existing unsupervised learning-based data-driven methods fail to address the BSCM problem in a model-free process. The methods are – data-driven nonlinear state estimation (DANSE) and deep Markov model (DMM). While DANSE provides good predictive/forecasting performance to model the temporal measurement data as a time series, its unsupervised learning lacks suitable regularization for tackling the BSCM task. We then propose a semi-supervised learning approach and develop a semi-supervised learning-based DANSE method, referred to as SemiDANSE. In SemiDANSE, we use a large amount of unlabelled data along with a limited amount of labelled data, i.e., pairwise measurement-and-state data, which provides the desired regularization. Using benchmark chaotic dynamical systems, we empirically show that the data-driven SemiDANSE provides competitive state estimation performance for BSCM using a handful of different measurement systems, against a hybrid method called KalmanNet and two model-driven methods (extended Kalman filter and unscented Kalman filter) that know the dynamical models exactly.

Index Terms—Bayesian state estimation, semi-supervised learning, maximum-likelihood, recurrent neural network.

I. INTRODUCTION

Estimating the state of a nonlinear dynamical process from noisy measurements is an active area of research with many applications. Traditionally, a dynamical system is represented using a state space model (SSM). An SSM has two main parts – a dynamical model and a measurement system. The dynamical model of the process describes the state’s evolution over time. The measurement/observation system, illustrates the relation between the unobserved state and the observed, noisy measurement [1]–[3]. The Bayesian state estimation task is to find an estimate of the unknown state at a given time instant using available, noisy measurements [1, Chap. 1].

In this article, we assume a compressed, linear measurement system with additive measurement noise, with the measure-

ment system known a priori. The dimension of the measurement vector is smaller than that of the corresponding temporal state vector. The Bayesian state estimation from compressed measurements (BSCM) is an under-determined state estimation problem. Note that BSCM subsumes the special case of Bayesian state estimation from partial measurements where only some components of the state vector are observed.

Classical methods for addressing the BSCM problem are SSM-informed., namely they are *model-driven* and know the under-determined measurement system. Model-driven methods assume knowledge of the dynamical model of the underlying process. Typically, a dynamical model is considered Markovian. Prominent examples of classical methods include the celebrated Kalman filter (KF) and its extensions such as the extended KF (EKF), the unscented KF (UKF) and the sampling-based particle filter (PF) [4]–[8]. KF uses a linear SSM and EKF uses a linearized SSM. UKF and PF, on the other hand, can directly work with the nonlinear SSM. The success of these methods for tackling the BSCM problem is owing to the regularization provided by the dynamical model.

Our interest in this article is to address BSCM for a model-free process. Here, ‘model-free’ means that the dynamical model of the process is *unknown*. Therefore, the use of classical model-driven methods is challenging, motivating the need for *data-driven* solutions. Design of data-driven methods uses different learning approaches. Learning can be in a supervised, unsupervised or semi-supervised manner. Supervised learning uses a labelled training dataset that consists of pairwise state and noisy measurement trajectories. On the other hand, unsupervised learning considers an unlabelled training dataset comprising noisy measurement trajectories. A semi-supervised learning approach uses a mix of labelled and unlabelled training data – typically, a limited amount of labelled data along with a larger amount of unlabelled data.

Examples of data-driven methods for Bayesian state estimation are Gaussian process-based (GP-based) methods proposed in [9]–[11]. GP-based techniques seek to learn the underlying process dynamics in a data-driven manner, using either a supervised or an unsupervised framework with the help of particle Markov chain Monte Carlo (MCMC). A second approach involves learning the unknown process noise or measurement noise statistics using an end-to-end approach via recurrent neural networks (RNNs) [12]. Another strategy involves learning the inverse relation of the measurement system and performing inference using message passing and

This work has been submitted to the IEEE for future publication. Copyright may be transferred without notice, after which this version will no longer be available. The work is supported by a research grant from the Digital Futures Centre. Project title: “DLL: Data-limited learning of complex systems”. Website: <https://www.digitalfutures.kth.se/>

graph neural networks [13]. In [14], the Markovian constraint in the process state dynamics is relaxed, and long short-term memory networks (LSTMs) are used for learning the process state dynamics from linear measurements. A semi-supervised approach using differentiable particle filters (DPFs) to explicitly learn the Markovian state dynamics and the measurement system was proposed in [15]. A class of variational inference-based approaches referred to as dynamic variational autoencoders (DVAEs) that seek to model the underlying process state dynamics using an approximate posterior distribution was proposed in [16]–[19]. A relevant example of a DVAE applied to state estimation is the deep Markov model (DMM) proposed in [19]. Finally, an unsupervised learning-based data-driven nonlinear state estimation (DANSE) method for a model-free process was considered in [20], [21].

Hybrid methods combine model-driven approaches with data-driven techniques. One prominent hybrid method is KalmanNet [22], [23] and its varieties [24]–[26]. KalmanNet is usually partially SSM-informed (only knowledge of the process and measurement dynamics) and uses data-driven neural networks. In [27], a UKF is used together with RNNs to model the process state dynamics using a general nonlinear autoregressive moving-average (NARMA) model. We point the interested reader to [21, Section I-A.] for a more comprehensive review of Bayesian state estimation covering model-driven, data-driven and hybrid approaches.

Both data-driven and hybrid approaches require training data for addressing state estimation tasks. While collecting a sufficient amount of unlabelled training data is realistic (e.g., by observing a model-free process for a long period), collecting a sufficient amount of labelled training data is expensive. Hence, designing a fully supervised, data-driven method for a model-free scenario is challenging. While unsupervised, data-driven approaches seem feasible, using them for tackling the BSCM problem has limitations. For example, unsupervised approaches such as DMM [18], [19] and GP-based SSM [10], are known to be computationally expensive to train and infer due to ancestral sampling. In the same context, the DANSE method is computationally simple and has been shown to perform well when the Bayesian state estimation problem is not under-determined [21]. However, when the problem is indeed underdetermined, DANSE suffers in performance and this limitation was shown using a single experiment with the benchmark chaotic dynamical system – the Lorenz-63 system in [21, Section III-F].

The experimental evidence leads to the following *hypothesis*: an unsupervised learning-based data-driven method like DANSE fails to tackle the BSCM problem for a model-free process owing to a lack of suitable regularization. Our remark is that the regularization is lacking as DANSE neither has a dynamical model nor access to labelled data. While the evidence of failure is mentioned in [21], there was no comprehensive study to verify the hypothesis. In this article, we empirically verify this hypothesis for DANSE and another unsupervised learning-based method called DMM. We also address the limitation of unsupervised learning using a *semi-supervised* learning-based approach.

Typically, semi-supervised learning is utilized in scenarios

where only a limited amount of labelled data is insufficient to achieve satisfactory performance in machine learning tasks, e.g., classification, regression, clustering, etc. [28], [29]. While a large body of semi-supervised learning research predominantly considers object detection and image segmentation [30], [31], a comparatively limited number of approaches consider regression tasks [32]. *In this paper, we design a semi-supervised learning-based approach to address the BSCM problem for a model-free process.*

Our main contributions in this work are as follows. We develop a semi-supervised learning-based method called semi-supervised DANSE (SemiDANSE) to tackle the BSCM problem for a model-free process. SemiDANSE has similar modeling assumptions as DANSE and retains the same computational benefits at inference time. However, unlike the unsupervised learning-based DANSE, SemiDANSE requires a combination of a limited amount of labelled data along with a large amount of unlabeled data for offline training. Our results are empirical and we do not focus on providing any theoretical guarantees concerning the use of semi-supervised learning in this paper. We experimentally demonstrate the limitation of unsupervised learning-based methods to handle the BSCM task in the case of two existing approaches – DANSE and DMM using the benchmark Lorenz-63 system [33]. For the Lorenz-63 system, we experimentally show the competitive state estimation performance of SemiDANSE compared with the model-driven EKF, UKF, and the hybrid KalmanNet for a handful of different under-determined measurement systems. Finally, we also use three other challenging, chaotic dynamical systems – a Chen system [34], [35], a Rössler system [36] and a high-dimensional Lorenz-96 system [37], [38] to empirically show that while DANSE fails, SemiDANSE performs reasonably well for the BSCM task.

The outline of the paper is as follows: in Section II-A, we restate the BSCM problem mathematically. Next, in Section II, we present the proposed SemiDANSE method, illustrating the associated inference and learning problems. Then, we perform experiments in Section IV to show how SemiDANSE can tackle the BSCM problem for a model-free process and compare its performance with relevant model-driven, data-driven and hybrid methods. Finally, we provide our conclusions and scope of future work in Section V.

A. Notations

We use bold font, lowercase symbols to denote vectors and regular, lowercase font to denote scalars. For example, \mathbf{x} represents a vector while x_j represents the j 'th component of \mathbf{x} . A sequence of vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t$ is compactly denoted by $\mathbf{x}_{1:t}$, where t denotes a discrete time index and $x_{t,j}$ is the j 'th component of \mathbf{x}_t . Upper case symbols in bold font, like \mathbf{H} , represent matrices. The operator $(\cdot)^\top$ denotes the transpose. $\mathcal{N}(\cdot; \mathbf{m}, \mathbf{L})$ represents the probability density function of the Gaussian distribution with mean \mathbf{m} and covariance matrix \mathbf{L} . $\det(\cdot)$, $\mathbb{E}[\cdot]$ and $\text{tr}(\cdot)$ are the determinant, expectation, and the trace operator, respectively. The notation $\|\mathbf{x}\|_{\mathbf{C}}^2$ is the squared ℓ_2 norm of \mathbf{x} weighted by the matrix \mathbf{C} , i.e. $\|\mathbf{x}\|_{\mathbf{C}}^2 = \mathbf{x}^\top \mathbf{C} \mathbf{x}$. $\text{diag}(\mathbf{x})$ denotes a square, diagonal matrix

with \mathbf{x} in its main diagonal. $|\mathcal{D}|$ is the cardinality of the set \mathcal{D} . A set of N natural indices $\{1, 2, \dots, N\}$ is compactly denoted by $[N]$. $\exp(\mathbf{A})$ is the matrix exponential of \mathbf{A} . $\mathbf{I}_{m \times m}$ denotes the $m \times m$ identity matrix, $\mathbf{O}_{m \times m}$ denotes the matrix of zeros of size $m \times m$. We mention that our notations are similar to those in [21].

II. BSCM PROBLEM AND SEMIDANSE

In this section, we describe the proposed SemiDANSE method. First, we state the BSCM problem formulation mathematically. Then, we describe the inference problem for SemiDANSE. It is worth mentioning here that the semi-supervised learning mechanism (described later in Section III) is the primary distinguishing feature between SemiDANSE and DANSE since the inference mechanism for the two methods is the same. Therefore, some of the equations in the following section are similar to those in [21, Section II-B], which we restate for completeness.

A. BSCM Problem formulation

Consider a nonlinear dynamical process $\{\mathbf{x}_t\}$ generating a sequence of states $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t, \dots$. The m -dimensional state vector $\mathbf{x}_t \in \mathbb{R}^m$ is observed using an n -dimensional noisy linear measurement vector $\mathbf{y}_t \in \mathbb{R}^n$, s.t.

$$\mathbf{y}_t = \mathbf{H}\mathbf{x}_t + \mathbf{w}_t, \quad t = 1, 2, \dots, \quad (1)$$

where $\mathbf{w}_t \sim \mathcal{N}(\mathbf{w}_t; \mathbf{0}, \mathbf{C}_w)$ is Gaussian measurement noise with zero mean and covariance matrix $\mathbf{C}_w \in \mathbb{R}^{n \times n}$. Here $\mathbf{H} \in \mathbb{R}^{n \times m}$ denotes a fixed, measurement matrix. The measurement system in (1) is assumed to be known, which means \mathbf{H} and \mathbf{C}_w are known. The dynamical model of the process $\{\mathbf{x}_t\}$ is unknown. The BSCM problem of interest is to find a probabilistic estimate of the state \mathbf{x}_t given the sequence of noisy, linear measurements $\mathbf{y}_{1:t}$ till the present time instant t (i.e. causally), where $n < m$, leading to compressed measurements.

B. Inference mechanism of SemiDANSE

In this section, we describe the inference mechanism of the proposed SemiDANSE method. The inference procedure involves finding an estimate of \mathbf{x}_t using the available measurements $\mathbf{y}_{1:t}$. We formulate this by first parameterizing a prior distribution of \mathbf{x}_t based on a sequence of past measurements $\mathbf{y}_{1:t-1}$. For SemiDANSE, similar to DANSE, we parameterize this prior distribution $p(\mathbf{x}_t | \mathbf{y}_{1:t-1})$ as a multivariate Gaussian using an RNN. Specifically, the RNN recursively uses $\mathbf{y}_{1:t-1}$ as the input sequence and provides the parameters of the Gaussian prior – a mean vector and a covariance matrix. This is schematically shown in Fig. 1. The RNN (together with the feed-forward networks) has learnable parameters denoted by $\boldsymbol{\theta}$. The actual RNN that we used is a GRU [39]. A brief discussion about our GRU implementation, including its layers, is provided in Section IV-B and Appendix A.

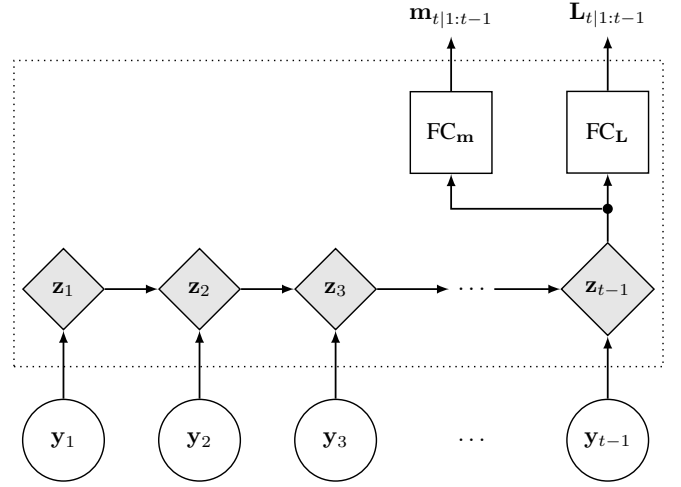


Fig. 1: Schematic of the parameterization of the Gaussian prior $p(\mathbf{x}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\theta})$ in (2) using an RNN (unfolded in time). The input to the RNN is the sequence $\mathbf{y}_{1:t-1}$ (unshaded, circle nodes). The internal states of the RNN are $\{\mathbf{z}_\tau\}$ (shaded, diamond nodes). At the last time instant the hidden state \mathbf{z}_{t-1} is mapped to $\mathbf{m}_{t|1:t-1}$ and $\mathbf{L}_{t|1:t-1}$ of the Gaussian prior using shallow, fully connected, feed-forward networks with suitable activation functions denoted by FC_m and FC_L respectively. Here, ‘FC’ stands for ‘Fully Connected.’ The mathematical details of the RNN and the feed-forward networks are described in Appendix A.

Mathematically, we have

$$\begin{aligned} p(\mathbf{x}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\theta}) &= \mathcal{N}(\mathbf{x}_t; \mathbf{m}_{t|1:t-1}(\boldsymbol{\theta}), \mathbf{L}_{t|1:t-1}(\boldsymbol{\theta})), \\ \text{s.t. } \{\mathbf{m}_{t|1:t-1}(\boldsymbol{\theta}), \mathbf{L}_{t|1:t-1}(\boldsymbol{\theta})\} &\triangleq \boldsymbol{\alpha}_{t|1:t-1}(\boldsymbol{\theta}), \\ \boldsymbol{\alpha}_{t|1:t-1}(\boldsymbol{\theta}) &= \text{RNN}(\mathbf{y}_{1:t-1}; \boldsymbol{\theta}), \end{aligned} \quad (2)$$

with the observation distribution as per (1) being $p(\mathbf{y}_t | \mathbf{x}_t) = \mathcal{N}(\mathbf{y}_t; \mathbf{H}\mathbf{x}_t, \mathbf{C}_w)$. Here, $\mathbf{m}_{t|1:t-1}(\boldsymbol{\theta}) \in \mathbb{R}^m$ and $\mathbf{L}_{t|1:t-1}(\boldsymbol{\theta}) \in \mathbb{R}^{m \times m}$ denote the mean and covariance matrix of the parameterized Gaussian prior respectively. We use $\mathbf{L}_{t|1:t-1}(\boldsymbol{\theta})$ as a diagonal covariance matrix for ease of practical implementation. In principle, we could have chosen any other parametric distribution for the prior $p(\mathbf{x}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\theta})$, but then we would possibly resort to a sampling-based approach using Monte-Carlo approximations for performing the necessary optimization of RNN parameters or subsequent calculation of posterior moments.

1) *Causal state estimation*: Using the prior in (2), the linear measurement setup (1) and the ‘completing the square’ approach [40, Chap. 2], we can compute the posterior $p(\mathbf{x}_t | \mathbf{y}_{1:t}; \boldsymbol{\theta})$ in closed-form Gaussian distribution.

$$\begin{aligned} p(\mathbf{x}_t | \mathbf{y}_{1:t}; \boldsymbol{\theta}) &= \mathcal{N}(\mathbf{x}_t; \mathbf{m}_{t|1:t}(\boldsymbol{\theta}), \mathbf{L}_{t|1:t}(\boldsymbol{\theta})), \\ \mathbf{m}_{t|1:t}(\boldsymbol{\theta}) &= \mathbf{m}_{t|1:t-1}(\boldsymbol{\theta}) + \mathbf{K}_{t|1:t-1} \boldsymbol{\varepsilon}_t, \\ \mathbf{L}_{t|1:t}(\boldsymbol{\theta}) &= \mathbf{L}_{t|1:t-1}(\boldsymbol{\theta}) - \mathbf{K}_{t|1:t-1} \mathbf{R}_{\varepsilon,t} \mathbf{K}_{t|1:t-1}^\top, \end{aligned} \quad (3)$$

where the second equation in (3) is obtained using the Woodbury matrix identity/matrix inversion lemma, with

$$\begin{aligned} \mathbf{K}_{t|1:t-1} &\triangleq \mathbf{L}_{t|1:t-1}(\boldsymbol{\theta}) \mathbf{H}^\top \mathbf{R}_{\varepsilon,t}^{-1}, \\ \mathbf{R}_{\varepsilon,t} &\triangleq \mathbf{H} \mathbf{L}_{t|1:t-1}(\boldsymbol{\theta}) \mathbf{H}^\top + \mathbf{C}_w, \\ \boldsymbol{\varepsilon}_t &\triangleq \mathbf{y}_t - \mathbf{H} \mathbf{m}_{t|1:t-1}(\boldsymbol{\theta}). \end{aligned} \quad (4)$$

The equations for $p(\mathbf{x}_t|\mathbf{y}_{1:t};\boldsymbol{\theta})$ are naturally the same as those for DANSE since the prior distribution in (2) is also parameterized as a multivariate Gaussian [21, Section II-B]. A possible point estimate of \mathbf{x}_t , denoted by $\hat{\mathbf{x}}_t(\boldsymbol{\theta})$, can be $\hat{\mathbf{x}}_t(\boldsymbol{\theta}) = \mathbf{m}_{t|1:t}(\boldsymbol{\theta})$ with uncertainty provided by $\mathbf{L}_{t|1:t}(\boldsymbol{\theta})$. In that case, we have minimum mean-square-error (MSE)

$$\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t(\boldsymbol{\theta})\|_2^2] = \mathbb{E}[\|\mathbf{x}_t - \mathbf{m}_{t|1:t}(\boldsymbol{\theta})\|_2^2] = \text{tr}(\mathbf{L}_{t|1:t}(\boldsymbol{\theta})). \quad (5)$$

2) *Comparison with inference in KF*: Similar to the KF updates, the linear, Gaussian measurement system in (1) allows for closed-form posterior updates as shown in (3) and (4) [1, Chap. 6]. However, we would like to point out two crucial differences in the inference steps of SemiDANSE compared to that of traditional KF-based variants. First, there is no explicit Gaussian propagation from the posterior moments at the current instant to the prior moments at the next time instant, as shown in Fig. 2. Second, from (2), we note that the prior moments at the t 'th instant are obtained from the RNN using $\mathbf{y}_{1:t-1}$ as opposed to a dynamical model-driven prediction step using previous posterior moments.

3) *Joint posterior estimation maintaining causality*: An additional task is to find the joint posterior $p(\mathbf{x}_{1:t}|\mathbf{y}_{1:t};\boldsymbol{\theta})$ because it is a crucial component for semi-supervised learning, shown later in Section III. We proceed as follows. First, noting that using the product rule, we have the following factorization:

$$p(\mathbf{x}_{1:t}|\mathbf{y}_{1:t};\boldsymbol{\theta}) = \prod_{\tau=1}^t p(\mathbf{x}_\tau|\mathbf{x}_{1:\tau-1}, \mathbf{y}_{1:t};\boldsymbol{\theta}). \quad (6)$$

In the absence of any simplifying assumptions, the prevalent approach for modelling $p(\mathbf{x}_{1:t}|\mathbf{y}_{1:t};\boldsymbol{\theta})$ using (6) involves ancestral sampling of past states $\mathbf{x}_{1:\tau-1}$ which is computationally expensive for moderately large T [16]. We avoid this computational burden by employing a conditional independence assumption – the current state \mathbf{x}_τ is assumed conditionally independent of the previous states $\mathbf{x}_{1:\tau-1}$ given the measurements $\mathbf{y}_{1:t}$. We point out that this is often used in variational inference-based approaches for obtaining simplified, mean-field approximations for variational distributions [40, Chap. 10], [19] and helps further simplify the right-hand side of (6). Then, we have

$$\prod_{\tau=1}^t p(\mathbf{x}_\tau|\mathbf{x}_{1:\tau-1}, \mathbf{y}_{1:t};\boldsymbol{\theta}) = \prod_{\tau=1}^t p(\mathbf{x}_\tau|\mathbf{y}_{1:t};\boldsymbol{\theta}). \quad (7)$$

Finally, since we are interested to find an estimate of \mathbf{x}_τ using only available measurements $\mathbf{y}_{1:\tau}$, for $\tau = 1, 2, \dots, t$ (as mentioned in Section II-A), we use a causality assumption. Under this assumption, we have

$$\prod_{\tau=1}^t p(\mathbf{x}_\tau|\mathbf{y}_{1:t};\boldsymbol{\theta}) = \prod_{\tau=1}^t p(\mathbf{x}_\tau|\mathbf{y}_{1:\tau};\boldsymbol{\theta}). \quad (8)$$

The equality mentioned in (8) holds under the causality assumption, which is a modeling assumption in our case. Combining (6), (7) and (8), we have

$$p(\mathbf{x}_{1:t}|\mathbf{y}_{1:t};\boldsymbol{\theta}) = \prod_{\tau=1}^t p(\mathbf{x}_\tau|\mathbf{y}_{1:\tau};\boldsymbol{\theta}) = \prod_{\tau=1}^t \mathcal{N}(\mathbf{x}_\tau; \mathbf{m}_{\tau|1:\tau}(\boldsymbol{\theta}), \mathbf{L}_{\tau|1:\tau}(\boldsymbol{\theta})), \quad (9)$$

where the last equality is obtained by using (3).

Using (9), we can conveniently incorporate RNNs in a semi-supervised learning framework and maintain causality. The use of the assumptions thus provides us with mathematical convenience and entails a trade-off with accuracy in case of modeling the joint posterior (9). In the following section, we describe the learning framework of SemiDANSE, which includes learning the unknown parameters $\boldsymbol{\theta}$ of the RNN.

III. SEMI-SUPERVISED LEARNING OF SEMIDANSE

In the following subsections, we formally describe the training dataset for the learning of SemiDANSE and our maximum-likelihood based optimization scheme. Finally, we provide intuitions regarding using semi-supervised and unsupervised learning for tackling the BSCM problem.

A. Training data

Consider two disjoint subsets of indices as $\mathcal{I}_s \subseteq [N]$ and $\mathcal{I}_u \subseteq [N]$, such that $\mathcal{I}_s \cup \mathcal{I}_u = [N]$ and $\mathcal{I}_s \cap \mathcal{I}_u = \emptyset$. Let us introduce two notations $N_s \triangleq |\mathcal{I}_s|$ and $N_u \triangleq |\mathcal{I}_u|$, such that $N_s + N_u = N$.

A T -length state sequence $\mathbf{x}_{1:T}$ has its corresponding measurement sequence $\mathbf{y}_{1:T}$. Assume there exists a genie that has access to a dataset

$$\mathcal{D} \triangleq \left\{ \left(\mathbf{x}_{1:T^{(i)}}^{(i)}, \mathbf{y}_{1:T^{(i)}}^{(i)} \right) \right\}_{i \in [N]}, \quad (10)$$

where $T^{(i)}$ is the length of the i 'th sequence. The dataset has N independent state-and-measurement pairwise data samples from the joint random process $p(\{\mathbf{x}_t\}, \{\mathbf{y}_t\})$. The genie makes a part of the full dataset \mathcal{D} accessible to a designer. When the genie provides the dataset

$$\mathcal{D}_u \triangleq \left\{ \left(\mathbf{y}_{1:T^{(i)}}^{(i)} \right) \right\}_{i \in \mathcal{I}_u} \subset \mathcal{D} \quad (11)$$

to the designer, \mathcal{D}_u consists of unlabelled (measurement-only) data for unsupervised learning. If $|\mathcal{I}_u| \triangleq N_u = N$, then the designer received the maximum amount of available unlabelled data. On the other hand, if the genie provides the labelled dataset

$$\mathcal{D}_s = \left\{ \left(\mathbf{x}_{1:T^{(j)}}^{(j)}, \mathbf{y}_{1:T^{(j)}}^{(j)} \right) \right\}_{j \in \mathcal{I}_s} \subseteq \mathcal{D}, \quad (12)$$

then, \mathcal{D}_s can be used for supervised learning. If $|\mathcal{I}_s| \triangleq N_s = N$, then the designer received the maximum amount of available labelled data. Finally, if we have access to a dataset $\mathcal{D}_{\text{semi}}$ comprised of two parts as

$$\mathcal{D}_{\text{semi}} = \{\mathcal{D}_s, \mathcal{D}_u\}, \quad (13)$$

such that $|\mathcal{I}_s| \triangleq N_s \ll N$ and $|\mathcal{I}_u| \triangleq N_u = N - N_s$, then \mathcal{D}_s is the part of the dataset that is labelled, and \mathcal{D}_u is the part of the dataset that is unlabelled. The dataset $\mathcal{D}_{\text{semi}}$ can be used for semi-supervised learning.

B. Maximum-likelihood based semi-supervised learning

We proceed to describe our maximum-likelihood-based semi-supervised learning approach. For semi-supervised learning, we start with the joint posterior

$$\begin{aligned} p(\mathbf{x}_{1:t}, \mathbf{y}_{1:t}) &\triangleq p(\mathbf{x}_{1:t}, \mathbf{y}_{1:t}; \boldsymbol{\theta}) \\ &= p(\mathbf{x}_{1:t} | \mathbf{y}_{1:t}; \boldsymbol{\theta}) p(\mathbf{y}_{1:t}; \boldsymbol{\theta}), \end{aligned} \quad (14)$$

where the joint distribution of the measurements is as follows:

$$p(\mathbf{y}_{1:t}) \triangleq p(\mathbf{y}_{1:t}; \boldsymbol{\theta}) = \prod_{\tau=1}^t p(\mathbf{y}_{\tau} | \mathbf{y}_{1:\tau-1}; \boldsymbol{\theta}). \quad (15)$$

The conditional marginal distribution $p(\mathbf{y}_{\tau} | \mathbf{y}_{1:\tau-1}; \boldsymbol{\theta})$ is expressed as in [21, Section II-B]

$$\begin{aligned} p(\mathbf{y}_t | \mathbf{y}_{1:t-1}) &= \int_{\mathbf{x}_t} p(\mathbf{y}_t, \mathbf{x}_t | \mathbf{y}_{1:t-1}) d\mathbf{x}_t \\ &= \int_{\mathbf{x}_t} p(\mathbf{y}_t | \mathbf{x}_t) p(\mathbf{x}_t | \mathbf{y}_{1:t-1}) d\mathbf{x}_t \\ &= \mathcal{N}(\mathbf{y}_t; \mathbf{H}\mathbf{m}_{t|1:t-1}(\boldsymbol{\theta}), \mathbf{C}_w + \mathbf{H}\mathbf{L}_{t|1:t-1}(\boldsymbol{\theta})\mathbf{H}^\top) \\ &\triangleq p(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\theta}). \end{aligned} \quad (16)$$

In the above, we used the prior form (2) and the Gaussian observation distribution as per (1). Here we introduced the notation $p(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\theta})$ to show the explicit dependency of $p(\mathbf{y}_t | \mathbf{y}_{1:t-1})$ on $\boldsymbol{\theta}$. Note that (16) is also a predictive distribution of \mathbf{y}_t given $\mathbf{y}_{1:t-1}$.

The maximum-likelihood based semi-supervised learning problem using (13) is

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg \max_{\boldsymbol{\theta}} \log p(\mathcal{D}_{\text{semi}}; \boldsymbol{\theta}) \\ &= \arg \max_{\boldsymbol{\theta}} \log p(\{\mathcal{D}_s, \mathcal{D}_u\}; \boldsymbol{\theta}) \\ &= \arg \max_{\boldsymbol{\theta}} \{\log p(\mathcal{D}_s; \boldsymbol{\theta}) + \log p(\mathcal{D}_u; \boldsymbol{\theta})\}. \end{aligned} \quad (17)$$

We further expand

$$\begin{aligned} \log p(\mathcal{D}_s; \boldsymbol{\theta}) &= \log \prod_{j \in \mathcal{I}_s} p(\mathbf{x}_{1:T^{(j)}}^{(j)}, \mathbf{y}_{1:T^{(j)}}^{(j)}; \boldsymbol{\theta}) \\ &= \log \prod_{j \in \mathcal{I}_s} p(\mathbf{x}_{1:T^{(j)}}^{(j)} | \mathbf{y}_{1:T^{(j)}}^{(j)}; \boldsymbol{\theta}) p(\mathbf{y}_{1:T^{(j)}}^{(j)}; \boldsymbol{\theta}), \end{aligned} \quad (18)$$

and

$$\log p(\mathcal{D}_u; \boldsymbol{\theta}) = \log \prod_{i \in \mathcal{I}_u} p(\mathbf{y}_{1:T^{(i)}}^{(i)}; \boldsymbol{\theta}). \quad (19)$$

Substituting (9), (16), (18), (19) in (17), and converting the logarithm of products to sums, we have the final expression of the semi-supervised learning optimization problem:

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg \max_{\boldsymbol{\theta}} \left\{ \sum_{j \in \mathcal{I}_s} \sum_{t=1}^{T^{(j)}} \log p(\mathbf{x}_t^{(j)} | \mathbf{y}_{1:t}^{(j)}; \boldsymbol{\theta}) \right. \\ &\quad \left. + \sum_{i \in \mathcal{I}_s \cup \mathcal{I}_u} \sum_{t=1}^{T^{(i)}} \log p(\mathbf{y}_t^{(i)} | \mathbf{y}_{1:t-1}^{(i)}; \boldsymbol{\theta}) \right\} \\ &= \arg \min_{\boldsymbol{\theta}} \left\{ \sum_{j \in \mathcal{I}_s} \mathcal{L}_s(\mathbf{x}_{1:T^{(j)}}^{(j)}, \mathbf{y}_{1:T^{(j)}}^{(j)}; \boldsymbol{\theta}) \right. \\ &\quad \left. + \sum_{i=1}^N \mathcal{L}_u(\mathbf{y}_{1:T^{(i)}}^{(i)}; \boldsymbol{\theta}) \right\}. \\ &= \arg \min_{\boldsymbol{\theta}} \mathcal{L}(\mathcal{D}_{\text{semi}}; \boldsymbol{\theta}). \end{aligned} \quad (20)$$

In the above expression, $\mathcal{L}_u(\cdot; \boldsymbol{\theta})$ denotes the loss due to unsupervised learning, and $\mathcal{L}_s(\cdot, \cdot; \boldsymbol{\theta})$ is the loss due to supervised learning. Together, we have a total loss

$$\begin{aligned} \mathcal{L}(\mathcal{D}_{\text{semi}}; \boldsymbol{\theta}) &\triangleq \sum_{j \in \mathcal{I}_s} \mathcal{L}_s(\mathbf{x}_{1:T^{(j)}}^{(j)}, \mathbf{y}_{1:T^{(j)}}^{(j)}; \boldsymbol{\theta}) \\ &\quad + \sum_{i=1}^N \mathcal{L}_u(\mathbf{y}_{1:T^{(i)}}^{(i)}; \boldsymbol{\theta}). \end{aligned} \quad (21)$$

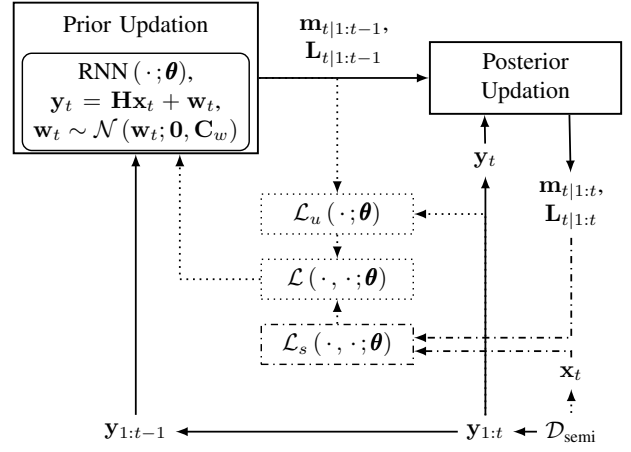


Fig. 2: Schematic of SemiDANSE at the t 'th time instant. The dotted lines represent information flow specifically during the learning phase, i.e. calculation of losses and gradients for RNN learning, $\mathcal{D}_{\text{semi}}$ represents the training dataset for SemiDANSE as defined in (13), \mathcal{L}_s and \mathcal{L}_u denote the supervised and the unsupervised loss respectively at the t 'th instant as defined in (22) and (23) respectively. The dash-dotted lines denote information flow specifically involving the limited \mathcal{D}_s in $\mathcal{D}_{\text{semi}}$ during the learning phase. Solid lines represent information flow involving other quantities and during inference.

We note that in the absence of $\mathcal{L}_s(\cdot, \cdot; \boldsymbol{\theta})$, the total loss is only due to unsupervised learning, which is indeed the case for DANSE [21]. We can further expand \mathcal{L}_s as

$$\begin{aligned} \mathcal{L}_s(\mathbf{x}_{1:T^{(j)}}^{(j)}, \mathbf{y}_{1:T^{(j)}}^{(j)}; \boldsymbol{\theta}) &\triangleq - \sum_{t=1}^{T^{(j)}} \log p(\mathbf{x}_t^{(j)} | \mathbf{y}_{1:t}^{(j)}; \boldsymbol{\theta}) \\ &= \sum_{t=1}^{T^{(j)}} \left\{ \frac{m}{2} \log 2\pi + \frac{1}{2} \log \det(\mathbf{L}_{t|1:t}^{(j)}(\boldsymbol{\theta})) \right. \\ &\quad \left. + \frac{1}{2} \|\mathbf{x}_t^{(j)} - \mathbf{m}_{t|1:t}^{(j)}(\boldsymbol{\theta})\|_{(\mathbf{L}_{t|1:t}^{(j)}(\boldsymbol{\theta}))^{-1}}^2 \right\}, \end{aligned} \quad (22)$$

where $\mathbf{m}_{t|1:t}^{(j)}(\boldsymbol{\theta})$ and $\mathbf{L}_{t|1:t}^{(j)}(\boldsymbol{\theta})$ are computed using (2), (3), (4) with input $\mathbf{y}_{1:t}^{(j)}$ where $j \in \mathcal{I}_s$. This means that the sequence of past measurements $\mathbf{y}_{1:t-1}^{(j)}$ is first passed to the RNN $(\cdot; \boldsymbol{\theta})$ to obtain the prior moments $\mathbf{m}_{t|1:t-1}^{(j)}(\boldsymbol{\theta})$, $\mathbf{L}_{t|1:t-1}^{(j)}(\boldsymbol{\theta})$. Subsequently the prior moments, together with the measurement system in (1) and the current measurement $\mathbf{y}_t^{(j)}$ are used to compute the posterior moments $\mathbf{m}_{t|1:t}^{(j)}(\boldsymbol{\theta})$ and $\mathbf{L}_{t|1:t}^{(j)}(\boldsymbol{\theta})$ according to (3), (4) where $j \in \mathcal{I}_s$. Similarly, we can expand \mathcal{L}_u as follows

$$\begin{aligned} \mathcal{L}_u(\mathbf{y}_{1:T^{(i)}}^{(i)}; \boldsymbol{\theta}) &\triangleq - \sum_{t=1}^{T^{(i)}} \log p(\mathbf{y}_t^{(i)} | \mathbf{y}_{1:t-1}^{(i)}; \boldsymbol{\theta}) \\ &= \sum_{t=1}^{T^{(i)}} \left\{ \frac{n}{2} \log 2\pi + \frac{1}{2} \log \det(\mathbf{C}_w + \mathbf{H}\mathbf{L}_{t|1:t-1}^{(i)}(\boldsymbol{\theta})\mathbf{H}^\top) \right. \\ &\quad \left. + \frac{1}{2} \|\mathbf{y}_t^{(i)} - \mathbf{H}\mathbf{m}_{t|1:t-1}^{(i)}(\boldsymbol{\theta})\|_{(\mathbf{C}_w + \mathbf{H}\mathbf{L}_{t|1:t-1}^{(i)}(\boldsymbol{\theta})\mathbf{H}^\top)^{-1}}^2 \right\}. \end{aligned} \quad (23)$$

The prior moments $\mathbf{m}_{t|1:t-1}^{(i)}(\boldsymbol{\theta})$, $\mathbf{L}_{t|1:t-1}^{(i)}(\boldsymbol{\theta})$ in (23) are obtained by passing the sequence of past measurements $\mathbf{y}_{1:t-1}^{(i)}$ to the RNN $(\cdot; \boldsymbol{\theta})$ for $i \in \mathcal{I}_u \cup \mathcal{I}_s$.

The overall optimization problem (20) is non-convex and we solve it using gradient descent. Finally, it is also worth noting that the two loss functions in (22) and (23) can be used with the same RNN $(\cdot; \theta)$. To see this, note that (23) requires $\mathbf{m}_{t|1:t-1}^{(i)}$, $\mathbf{L}_{t|1:t-1}^{(i)}$ for $i \in \mathcal{I}_u \cup \mathcal{I}_s$. Once the prior moments for $i \in \mathcal{I}_u \cup \mathcal{I}_s$ are computed, the specific prior moments and the current measurements for the indices $j \in \mathcal{I}_s$ can be extracted and passed on for additionally computing the posterior moments required in the supervised loss term (22). A schematic diagram is depicted in Fig. 2 for visualizing the flow of information during inference and learning.

Finally, we provide some comments on the loss terms in (17). From (20), note that the semi-supervised learning for SemiDANSE translates to the unsupervised learning of DANSE with the maximum amount of unlabelled data, when $\mathcal{D}_s = \emptyset$, that means $|\mathcal{I}_u| \triangleq N_u = N$. A small $N_s \triangleq |\mathcal{I}_s| \ll N$ with $\mathcal{D}_s \neq \emptyset$ is intended and interesting for semi-supervised learning. Also, while $\mathcal{L}_s(\cdot, \cdot; \theta)$ in (22) has (coincidentally) the same form as that of the loss term for deriving an empirical estimation performance limit in [21, Section II-E], we reiterate that the purpose of this loss term here is completely different. Lastly, the linear measurement system in enables us to compute the marginal likelihood and the posterior distribution analytically – leading to direct optimization instead of indirectly maximizing a lower bound of the log-likelihood, as is commonly done in the case of using an expectation lower bound (ELBO) [41].

C. Intuitions regarding unsupervised and semi-supervised learning-based methods

We provide intuitions behind the failure of unsupervised learning (DANSE) for addressing the BSCM problem and the success of semi-supervised learning (SemiDANSE) for the same. We reiterate that we do not focus on providing theoretical guarantees concerning supervised or semi-supervised learning, and our supporting results (as shown later in Section IV) are empirical. To the best of the authors’ knowledge, explainability and performance guarantees for methods involving deep architectures, such as LSTMs and GRUs, is non-trivial and an ongoing research problem in the theoretical machine learning community. The unsupervised learning problem in DANSE [21] is

$$\arg \min_{\theta} \sum_{i=1}^N \mathcal{L}_u \left(\mathbf{y}_{1:T(i)}^{(i)}; \theta \right), \quad (24)$$

where \mathcal{L}_u is shown in (23). A natural question is why the solution of the above learning problem does not work properly for the BSCM problem? We endeavor to provide an argument as follows: Let the total number of scalars in θ be N_{θ} . Assuming $T^{(i)} = T, \forall i \in [N]$, the total number of constraints in unsupervised learning is nNT . Typically $N_{\theta} < nNT$, and even $N_{\theta} \ll nNT$, hence the learning of θ may not be necessarily under-determined. However, learning an appropriate θ for state estimation remains an under-determined problem. We can learn θ by optimizing (24), and subsequently, RNN provides a set of parameters for the Gaussian prior that is consistent with the measurement/observation system. However, this does not imply the consistency of the learned parameters with




Fig. 3: Visualization of three chaotic dynamical systems: randomly chosen state trajectories from the test set of length $T_{\text{test}} = 2000$ for the Lorenz-63, Chen, and Rössler attractors. The state trajectories show variations of the three systems.

the underlying true states. Overall, this phenomenon can be explained conceptually as a limitation of unsupervised learning because we do not have access to any labelled data to provide a suitable regularization in (24).

For semi-supervised learning, we have a limited amount of labelled data available as shown in (13). This leads to the optimization problem in (20) which can be achieved from (24) with the additional regularization term $\sum_{j \in \mathcal{I}_s} \mathcal{L}_s \left(\mathbf{x}_{1:T(j)}^{(j)}; \mathbf{y}_{1:T(j)}^{(j)}; \theta \right)$. We believe that SemiDANSE could work by exploiting the available time-wise and dimension-wise correlation in the true state trajectories using the limited amount of labelled training data in (13) – something that is unavailable for a purely unsupervised learning-based method like DANSE. We hypothesize that this availability of a limited amount of labelled data enables the method to exploit the available time-wise and dimension-wise correlation in the true state trajectories from \mathcal{D}_s , and find a suitable set of learnable parameters. In the case of DANSE and DMM, the learning algorithm also finds a suitable set of parameters, but the methods are notably unsupervised, relying on only noisy measurement trajectories from \mathcal{D}_s . Note that a hybrid method like KalmanNet, knowing the underlying dynamical model, can also work if there are correlations in the true state trajectories. Finally, it is worthwhile noting that if there arises a scenario where two distinct state trajectories get mapped to *identical* measurement trajectories, then it would be very challenging for any method to address the BSCM problem satisfactorily. Without a known dynamical model to guide the state estimation algorithm, SemiDANSE would try to improve upon DANSE by exploiting the limited amount of labelled training data to learn a suitable θ^* as per (20). While we cannot provide any theoretical guarantees (such as an MMSE performance bound) for SemiDANSE owing to the use of neural networks, we note favourable empirical state estimation performances in the case of a handful of under-determined measurement systems, shown in the next section.

IV. EXPERIMENTS AND RESULTS

In this section, we demonstrate the empirical performance of SemiDANSE vis-à-vis other Bayesian state estimation methods. We perform experiments using complex nonlinear processes and generate data knowing their SSMs. While SemiDANSE does not know the dynamical models in the SSMs during learning or inference, the model-driven methods know

the SSMs (both the dynamical model and the measurement system).

The organization of this section is as follows: First, we describe three chaotic dynamical systems (stochastic in nature) that we simulate as nonlinear processes. Next, we discuss relevant information about the training and testing of SemiDANSE, performance measures and the choice of RNN. We then mention the competing Bayesian state estimation methods and finally illustrate experimental results related to the BSCM problem for the proposed SemiDANSE method.

A. Three chaotic dynamical systems

In this subsection, we briefly describe three chaotic dynamical systems (stochastic in nature) used for experiments as benchmark nonlinear systems. They are

- 1) The Lorenz-63 system [33],
- 2) The Chen system [35], and
- 3) The Rössler system [36].

In chaos theory literature, the above dynamical systems are often referred to as *strange attractors* [42]. Hence, we sometimes use the nomenclature ‘the Lorenz-63 attractor’ and ‘the Lorenz-63 system’ interchangeably. These three systems are 3-dimensional with highly complex state dynamics. We simulate each of the chaotic dynamical systems by discretizing their original continuous-time dynamics following [21]. We also incorporate i.i.d. process noise to ensure randomness in the simulated state trajectories. The detailed equations for their state dynamics (or dynamical models) and simulation procedures are described in Appendix B. All of them are Markovian processes. In all cases, the corresponding noisy measurement trajectories are generated using the linear measurement setup in (1). For visualization, we show randomly chosen test set trajectories of the three processes in Fig. 3. This figure helps to visualize how different the trajectories can be across the three chaotic dynamical systems and a qualitative justification for the use of them in state estimation experiments across different dynamical systems. Towards the end of the manuscript, we also experiment using a high-dimensional chaotic dynamical system – the Lorenz-96 system [37], [38], to explore the applicability of the proposed method.

B. Training, testing, performance measures and RNN

We have training datasets as described in Section III for DANSE, DMM, and SemiDANSE, and a testing dataset $\mathcal{D}_{\text{test}} = \left\{ \left(\mathbf{x}_{1:T_{\text{test}}}^{(j)}, \mathbf{y}_{1:T_{\text{test}}}^{(j)} \right) \right\}_{j=1}^{N_{\text{test}}}$. We mention that SemiDANSE is trained using $\mathcal{D}_{\text{semi}}$, thus involving a combination of labelled and unlabelled data. For training, we use a dataset where $N = 1000$ (total number of training trajectories) and having the same length $T = 100$ for each training trajectory. For testing, we use a dataset $\mathcal{D}_{\text{test}}$ where $N_{\text{test}} = 100$ (total number of testing trajectories) and having the same length $T_{\text{test}} = 2000$ for each testing trajectory. Denoting the estimated state as $\hat{\mathbf{x}}_t$ (for example using the posterior mean as a point estimate), we use the averaged normalized-mean-squared-error (NMSE) in decibels (dB) as the performance measure, defined

below as in [21, Section III]:

$$\text{NMSE} = \frac{1}{N_{\text{test}}} \sum_{j=1}^{N_{\text{test}}} 10 \log_{10} \frac{\sum_{t=1}^{T_{\text{test}}^{(j)}} \|\mathbf{x}_t^{(j)} - \hat{\mathbf{x}}_t^{(j)}\|_2^2}{\sum_{t=1}^{T_{\text{test}}^{(j)}} \|\mathbf{x}_t^{(j)}\|_2^2}. \quad (25)$$

For generating noisy measurements, we use i.i.d., white Gaussian noise as the measurement noise with $\mathbf{C}_w = \sigma_w^2 \mathbf{I}_n$. The process noise is also i.i.d., white Gaussian, with the process noise covariance \mathbf{C}_e in most cases being scaled diagonal, i.e. $\mathbf{C}_e = \sigma_e^2 \mathbf{I}_m$. Specific details for each kind of dynamical model (Lorenz, Chen and Rössler) are given in Appendix B. The signal-to-measurement noise ratio (SMNR) is calculated in dB, same as in [21, Section III] on $\mathcal{D}_{\text{test}}$, as

$$\begin{aligned} \text{SMNR} &= \frac{\sum_{j=1}^{N_{\text{test}}} 10 \log_{10} \left(\frac{\sum_{t=1}^{T_{\text{test}}^{(j)}} \mathbb{E} \{ \|\mathbf{H}\mathbf{x}_t^{(j)} - \mathbb{E} \{ \mathbf{H}\mathbf{x}_t^{(j)} \} \|_2^2 \}}{\text{tr}(\mathbf{C}_w)} \right)}{\sum_{j=1}^{N_{\text{test}}} 10 \log_{10} \left(\frac{\sum_{t=1}^{T_{\text{test}}^{(j)}} \mathbb{E} \{ \|\mathbf{H}\mathbf{x}_t^{(j)} - \mathbb{E} \{ \mathbf{H}\mathbf{x}_t^{(j)} \} \|_2^2 \}}{n\sigma_w^2} \right)} \\ &= \frac{N_{\text{test}}}{N_{\text{test}}}. \end{aligned} \quad (26)$$

For notational convenience, we also introduce a parameter

$$\kappa \triangleq \frac{N_s}{N}, \quad 0 \leq \kappa \leq 1. \quad (27)$$

The parameter κ provides a measure to quantify the relative amount of labelled data and unlabelled data in semi-supervised learning. As a concrete example, we use $N = 1000$ and $N_s = 20$. That means $\kappa = 0.02 = 2\%$, and SemiDANSE uses 2% labelled data compared to the total amount of labelled and unlabelled data. Unless the value of κ is mentioned specifically, we stick with $\kappa = 0.02 = 2\%$ for most of our experiments. Further, the semi-supervised learning in Section III is formulated in a way that when $\kappa = 0$ then the semi-supervised learning translates to unsupervised learning. That means, SemiDANSE translates to DANSE when $\kappa = 0$.

We found the RNN architecture that includes appropriate feed-forward networks by cross-validation. For SemiDANSE, we experimented with both LSTM and GRU and found the suitable RNN architecture to be a GRU [39], [43]. The GRU has 1 hidden layer with 30 hidden nodes. The output of the RNN is passed to shallow feed-forward networks FC_m, FC_L as described in Fig. 1 for calculating the means and the covariances. The feed-forward networks have 2 hidden layers, where the first hidden layer has 30 hidden nodes and is shared between FC_m and FC_L . The output layer was designed to have 32 hidden nodes and distinct for FC_m and FC_L . The mathematical details of the feed-forward networks are described in (32) in Appendix A. We implemented SemiDANSE in Python and PyTorch [44] and trained the architecture using a single NVIDIA Tesla P100 GPU card¹ The training algorithm uses a mini-batch gradient descent with a batch size of 64. The optimizer chosen was Adam [45] with an adaptive learning rate set at a starting value 5×10^{-4} and decreased by 10% every $1/6^{\text{th}}$ of the maximum number of training epochs. The maximum number of training epochs was set at 2000, and an

¹The code will be made available upon request.

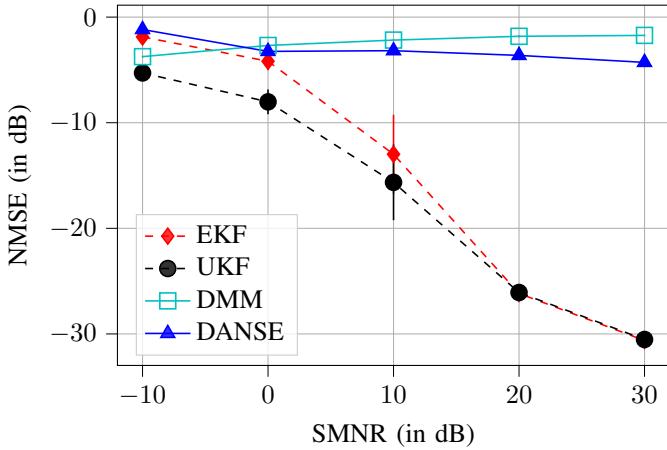


Fig. 4: The average NMSE (in dB) on $\mathcal{D}_{\text{test}}$ versus SMNR (in dB) performances to illustrate failure of unsupervised learning-based DANSE and DMM for the BSCM problem setup described in Section IV-D, with σ_e^2 corresponding to -10 dB. The model-driven EKF and UKF perform well.

early stopping criterion based on the MSE of the validation set was used to avoid overfitting, similar to DANSE.

C. Competing methods

In this subsection, we compare SemiDANSE vis-à-vis a few other Bayesian state estimation methods like [21] on the BSCM task. The methods that we compare are:

- 1) model-driven EKF and UKF,
- 2) data-driven causal DMM and DANSE,
- 3) hybrid method KalmanNet.

All the above methods (including SemiDANSE) have knowledge of the linear measurement system given by (1). Additionally, the model-driven EKF and UKF have full knowledge of the dynamical model of the underlying process and are used for benchmarking the performance of data-driven methods. They are implemented using Python, PyTorch and FilterPy [46]. The DMM and DANSE methods are implemented in PyTorch and trained by unsupervised learning that uses \mathcal{D}_u . DMM assumes that the underlying process is Markovian, but it does not know the exact dynamical model. A brief description of the learning problem for the DMM that we implemented in this article is given in Appendix C. DANSE does not have any knowledge of the underlying process. The KalmanNet is trained by unsupervised learning that uses \mathcal{D}_u and it also has the full knowledge of the process dynamics [22]. The training and simulation strategies for these competing methods follow the same as described in [21, Section III].

D. Failure of unsupervised DANSE and DMM for BSCM

We begin our experiments with the Lorenz-63 system, where we first show the failure of two unsupervised learning-based methods - DANSE and DMM for the BSCM task. For the experiment, we use 2×3 -dimensional \mathbf{H} matrix where each element is drawn from i.i.d. Gaussian source $\mathcal{N}(0, 1)$ and then




Fig. 5: The average NMSE (in dB) on $\mathcal{D}_{\text{test}}$ versus SMNR (in dB) performances to illustrate the success of SemiDANSE for the BSCM problem setup described in Section IV-D, with σ_e^2 corresponding to -10 dB. SemiDANSE ($\kappa = 0.02$) is compared with the model-driven EKF and UKF, the DMM [19], the hybrid KalmanNet [22] and the unsupervised learning-based DANSE.

fixed. Therefore we have $n = 2$ and $m = 3$, and the \mathbf{H} matrix we use for the BSCM problem is

$$\mathbf{H} = \begin{bmatrix} 0.37992 & 0.34099 & 1.04317 \\ 0.98070 & -0.70477 & 2.17908 \end{bmatrix}. \quad (28)$$

In Fig. 4, we show the NMSE versus SMNR performances of DANSE and DMM, and only include the performances of EKF and UKF for comparison. We note that both DANSE and DMM were trained on different SMNRs as shown in Fig. 4 and tested on the same SMNRs. It shows that the model-driven EKF and UKF perform well, but the unsupervised learning-based DANSE and DMM - both of them - fail to perform satisfactorily. Note that the Lorenz-63 system that we simulated is a Markovian process and DMM uses the knowledge that the underlying dynamical model is Markovian. Despite that, DMM is unable to address the BSCM problem.

Here, for completeness, we mention that if the \mathbf{H} matrix would be a full column-rank matrix, for example, a 3×3 -dimensional full-rank matrix, then DANSE and DMM could perform well, see [21, Section III-D].

E. Success of SemiDANSE for BSCM

For the same setup of the Section IV-D, we now show the performance of SemiDANSE and compare it with EKF, UKF and KalmanNet in Fig. 5. As in Section IV-D, SemiDANSE was also trained and tested on different SMNRs using $\kappa = 0.02$ to keep the comparison consistent. We observe that SemiDANSE is competitive and overcomes the limitation of DANSE for the BSCM problem. We remark that using 2% labelled data brings forth the desired regularization in SemiDANSE, leading to an improved state estimation performance. Additionally, we empirically provide an ablation study for assessing the individual contributions of the semi-supervised

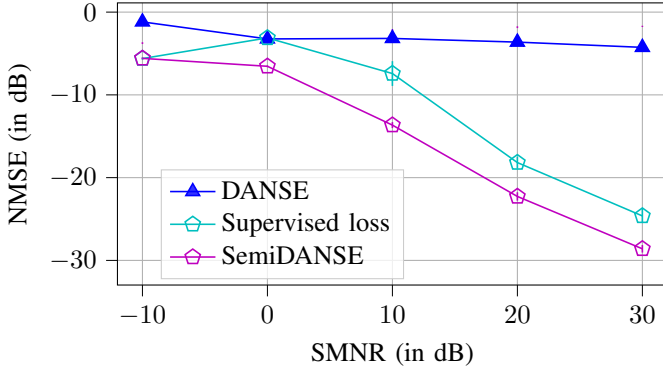


Fig. 6: The average NMSE (in dB) on $\mathcal{D}_{\text{test}}$ versus SMNR (in dB) performances to illustrate the success of SemiDANSE ($\kappa = 0.02$) for the BSCM problem setup for the BSCM problem setup described in Section IV-D, with σ_e^2 corresponding to -10 dB. SemiDANSE ($\kappa = 0.02$) is compared with the unsupervised learning-based DANSE and the supervised baseline using only \mathcal{L}_s in (22).

loss, namely the supervised and the unsupervised loss terms. The results are shown in Fig. 6. SemiDANSE is empirically found to perform better than a supervised baseline that is optimized using only \mathcal{L}_s as shown in (22).

F. Partial measurement system

In the previous two subsections, we had a dense \mathbf{H} matrix providing the 2-dimensional measurement vector \mathbf{y}_t that has influences from all the three components of the state vector \mathbf{x}_t . Instead of a dense \mathbf{H} , we can have scenarios where only some components of the state vectors are measured (observed). We refer to these scenarios as partial measurement scenarios.

Let us assume that a partial measurement matrix is

$$\mathbf{H} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (29)$$

That means we have measurements from the second and third components of the 3-dimensional state vector. Fig. 7 shows a 3-dimensional trajectory of the Lorenz-63 process, its 2-dimensional measurement at SMNR = 10 dB, the estimated trajectory using DANSE, and the same using SemiDANSE. It is clear that DANSE fails for the BSCM problem and SemiDANSE performs well.

In Fig. 8, we show the time trajectory of the first component $x_{t,1}$ of \mathbf{x}_t , and estimates of UKF, DANSE and SemiDANSE. DANSE fails to track $x_{t,1}$, and we observe that the model-driven UKF and the data-driven SemiDANSE perform well.

Note that the 2-dimensional measurement \mathbf{y}_t has observations of $x_{t,2}$ and $x_{t,3}$ under noise. From our experiments, we have seen that while DANSE fails to estimate $x_{t,1}$, it is able to estimate $x_{t,2}$ and $x_{t,3}$. We do not show the plots of the estimates of $x_{t,2}$ and $x_{t,3}$ for brevity. On the other hand, SemiDANSE could track all three components of \mathbf{x}_t .

Before proceeding to show further results, we seek to analyse the performance of SemiDANSE in case of the partial measurement scenario for different κ (defined in (27)). To

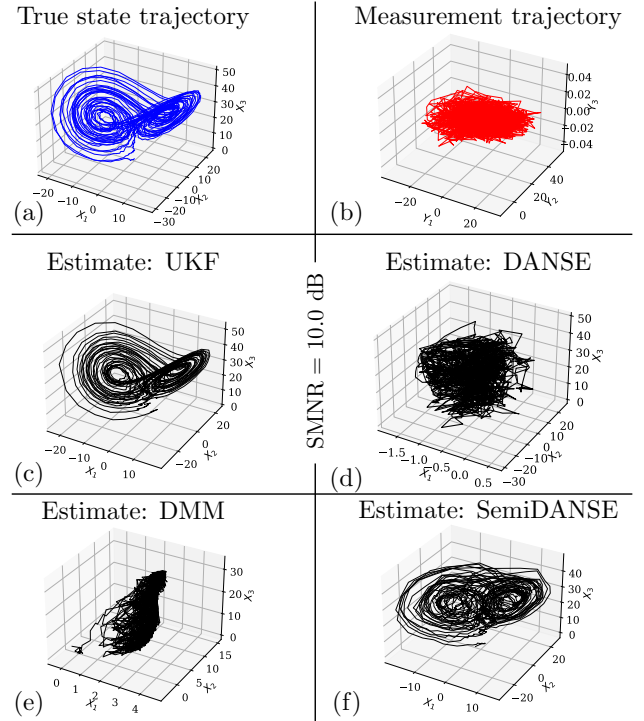


Fig. 7: Demonstrating failure of DANSE and success of SemiDANSE for BSCM problem with the partial measurement of Lorenz-63 process. We use the BSCM problem setup described in Section IV-F, where \mathbf{H} is shown in (29), SMNR = 10 dB and σ_e^2 corresponding to -10 dB. (a) A 3-dimensional true state trajectory instance of Lorenz-63 system from $\mathcal{D}_{\text{test}}$. (b) The corresponding 2-dimensional (noisy) measurement trajectory (shown in three-dimensions for visualization). (c) The estimate of UKF. (d) The estimate of DMM. (e) The estimate of DANSE (posterior mean). (f) The estimate of SemiDANSE (posterior mean) with $\kappa = 0.02$.

illustrate the behaviour of SemiDANSE for different values of κ , we plot the NMSE vs. SMNR for three different values of $\kappa = 0.002, 0.02$ and 0.05 in Fig. 9 for the partial measurement scenario described in section IV-F. We observe that there is a noticeable gain in the estimation performance moving from $\kappa = 0.002$ to $\kappa = 0.02$, and even minor improvement from $\kappa = 0.02$ to $\kappa = 0.05$ in certain SMNR values, indicating the benefit of using labelled data in a semi-supervised setup.

G. On predictive performance / forecasting

In this subsection, we show that DANSE provides a good predictive / forecasting performance for tracking the measurement sequence \mathbf{y}_t despite failing to estimate the underlying state sequence \mathbf{x}_t , which means failing to address the BSCM for the partial measurement system in Section IV-F. In Fig. 10, we plot the time-wise and coordinate-wise values of the measurement vector \mathbf{y}_t and the corresponding one-step ahead predicted measurement $\hat{\mathbf{y}}_t$ given $\mathbf{y}_{1:t-1}$ using DANSE and SemiDANSE at 10 dB SMNR for the Lorenz-63 process. Here $\hat{\mathbf{y}}_t$ is the predictive mean of $p(\mathbf{y}_t|\mathbf{y}_{1:t-1})$ (see (16) for SemiDANSE.). A notable fact is that DANSE provides

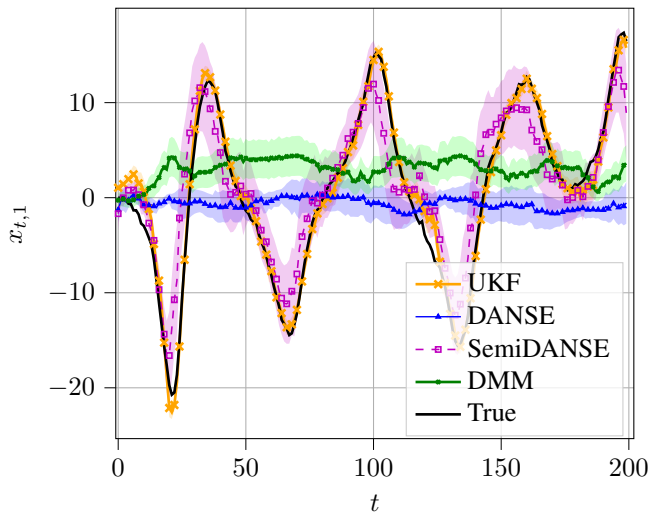


Fig. 8: Time-wise plot of a (true) state trajectory $x_{t,1}$ from $\mathcal{D}_{\text{test}}$ and its estimates using UKF, DANSE, DMM and SemiDANSE ($\kappa = 0.02$) at SMNR = 10 dB and σ_e^2 corresponding to -10 dB for the Lorenz-63 process. The results are for the BSCM problem setup described in Section IV-F. The solid/dashed line represents the corresponding posterior mean estimate. Note that DMM and DANSE fail to track the true state, whereas UKF and SemiDANSE can track it. A short portion of the test trajectory is shown for ease of visualization.






Fig. 9: NMSE (in dB) vs. SMNR (in dB) for SemiDANSE on $\mathcal{D}_{\text{test}}$ for the partial measurement case in IV-F different values of $\kappa = N_s/N$, with σ_e^2 corresponding to -10 dB.

reasonably good predictive performance compared to the proposed SemiDANSE, which shows that the DANSE is able to capture the underlying dynamics of the measurements, but the structure of \mathbf{H} in (29) renders it difficult to track the unobserved state vector coordinate $x_{t,1}$ as shown in Fig. 8.

This result for DANSE concludes a fact: unsupervised learning may provide a good predictive / forecasting performance on measurement data as a consequence of efficient time-series modelling, but that may not translate to a good state estimation performance. Estimating the hidden state from measurements proves to be a harder task than modeling measurement data.



(a) Plot of $y_{t,1}$ versus t



(b) Plot of $y_{t,2}$ versus t

Fig. 10: Time-wise plots of the coordinates of \mathbf{y}_t from $\mathcal{D}_{\text{test}}$ and the one-step ahead predicted estimates using DANSE and SemiDANSE ($\kappa = 0.02$) at SMNR = 10 dB and σ_e^2 corresponding to -10 dB. The results are related to the partial measurement system described in subsections IV-F, IV-G and for the Lorenz-63 process. The shaded region in each plot shows the $\pm 1\sigma$ point uncertainty of the corresponding predicted estimated of $\{\mathbf{y}_t\}$.

H. Partial measurement scenario - an extreme BSCM case

In the previous subsection, we had measurements from two components of the three-dimensional state vector of the Lorenz-63 process. An extreme BSCM problem case is: we have only measurements from one state component. For example, we have the following measurement matrix

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}. \quad (30)$$

That means we have the measurement of the first component of the 3-dimensional state vector. In this extreme case, Fig. 11 shows a 3-dimensional trajectory of the Lorenz-63 process, and the 3-dimensional state estimates from DANSE, UKF and SemiDANSE at SMNR = 10 dB. We observe again that data-driven SemiDANSE and model-driven UKF perform well, but unsupervised learning-based DANSE fails. In this regard, we remark that the success of SemiDANSE is because SemiDANSE can exploit the time-wise correlation in the state trajectories in $\mathcal{D}_{\text{semi}}$ and the dimension-wise correlation in each state vector, e.g., in the case of Lorenz-63 system, the second and third components of \mathbf{x}_t depend on the first component.

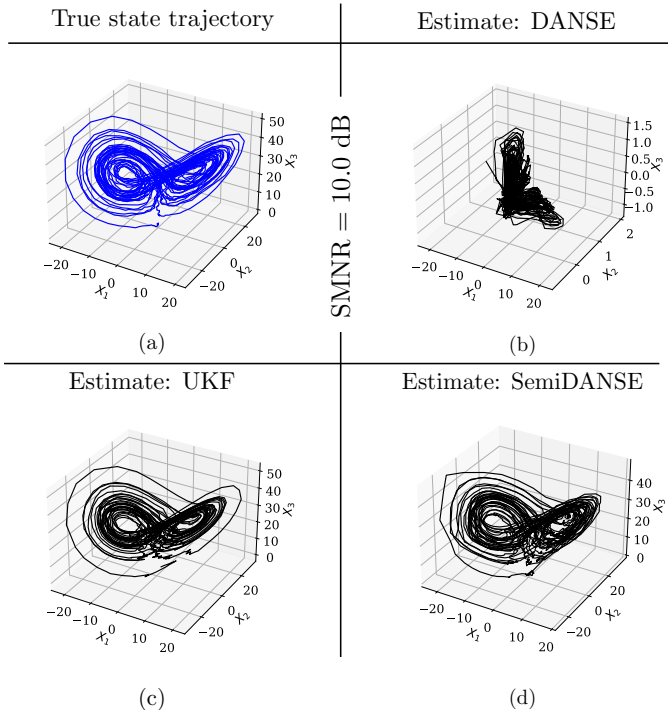


Fig. 11: Demonstrating failure of DANSE and success of SemiDANSE for the extreme BSCM problem with one-dimensional partial measurement of 3-dimensional Lorenz-63 process. We use the BSCM problem setup described in Section IV-H, where the 1×3 measurement matrix \mathbf{H} is shown in (30), $\text{SMNR} = 10$ dB and σ_e^2 corresponding to -10 dB. A 3-dimensional true state trajectory from $\mathcal{D}_{\text{test}}$ is shown in (a). The estimates for DANSE, UKF and SemiDANSE ($\kappa = 0.02$) are all posterior mean estimates shown in (b), (c) and (d).

I. Demonstration for more chaotic dynamical systems

We also show the performance of SemiDANSE on two other 3-dimensional chaotic dynamical systems - the Chen attractor and the Rössler attractor - for partial measurement setups. We use the 2×3 -dimensional \mathbf{H} matrix shown in (29) and $\text{SMNR} = 10$ dB. From Fig. 12 and Fig. 13, we observe that while DANSE fails, SemiDANSE succeeds.

Additionally, we performed experiments on a 20-dimensional, stochastic Lorenz-96 system [37], [38]. The system details can be found in [21, Appendix B], and we followed a similar experimental setup as described in [21, Section H]. We investigated the quantitative performance of DANSE and SemiDANSE for the measurement system in (1), with different sizes of the measurement matrix, starting with $\mathbf{H} = \begin{bmatrix} \mathbf{I}_{15 \times 15} & \mathbf{O}_{15 \times 5} \\ \mathbf{O}_{5 \times 15} & \mathbf{O}_{5 \times 5} \end{bmatrix}$ (corresponds to a partial measurement scenario with $n = 15, m = 20$) by varying n , keeping m fixed at 20. We used σ_e^2 corresponding to -10 dB and the SMNR corresponding to 10 dB. Thereafter, we reduced the measurement dimension n to 10, 5, 2, and in each case, evaluated the performance for both DANSE and SemiDANSE. DANSE was trained using \mathcal{D}_u with $N = 250, T = 1000$, SemiDANSE was trained using $\mathcal{D}_{\text{semi}}$ with $\kappa = 0.08$ with the same N, T as in the case of DANSE. Both the methods were tested on the same

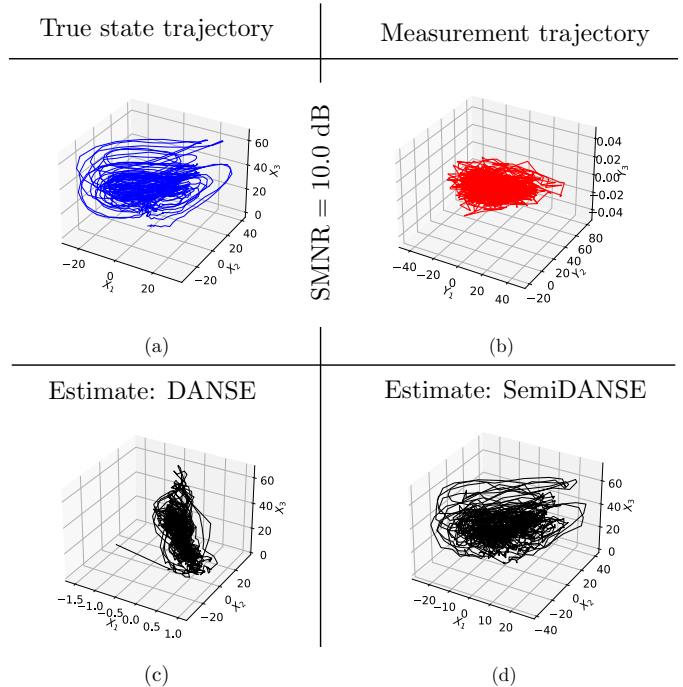


Fig. 12: Additional example of the failure of DANSE and success of SemiDANSE for BSCM problem with partial measurement of 3-dimensional Chen process. We use the 2×3 measurement matrix \mathbf{H} shown in (29). A 3-dimensional true state trajectory from $\mathcal{D}_{\text{test}}$ with σ_e^2 corresponding to -10 dB is shown in (a). The corresponding measurement trajectory at $\text{SMNR} = 10$ dB is shown in (b). The estimates for DANSE and SemiDANSE ($\kappa = 0.02$) are posterior mean estimates shown in (c) and (d) respectively.

$\mathcal{D}_{\text{test}}$ with $N_{\text{test}} = 100, T_{\text{test}} = 2000$. The natural expectation is that the difficulty of the BSCM problem increases with the decreasing value of n . The results are shown in Fig. 14. We observe that SemiDANSE achieves a reasonable performance gain over DANSE, which becomes noticeable as n decreases, thus demonstrating the applicability of SemiDANSE to high-dimensional state estimation problems.

V. CONCLUSIONS

We propose SemiDANSE - a semi-supervised data-driven state estimation method for tackling the BSCM problem in a model-free scenario. We empirically demonstrate the limitation of unsupervised learning for addressing the BSCM task in the case of two notable methods - DANSE and DMM. While DANSE provides satisfactory forecasting performance, the state estimation performance for the BSCM problem is unsatisfactory. Then, we empirically show the success of SemiDANSE in tackling the BSCM problem using a handful of different measurement systems and benchmark chaotic dynamical systems. We show this by employing only a limited amount of labelled data - for instance, 2% of the total number of training data samples together with unlabelled data for training SemiDANSE. Compared to DANSE, SemiDANSE shows noticeable improvement in learning, owing to the desired

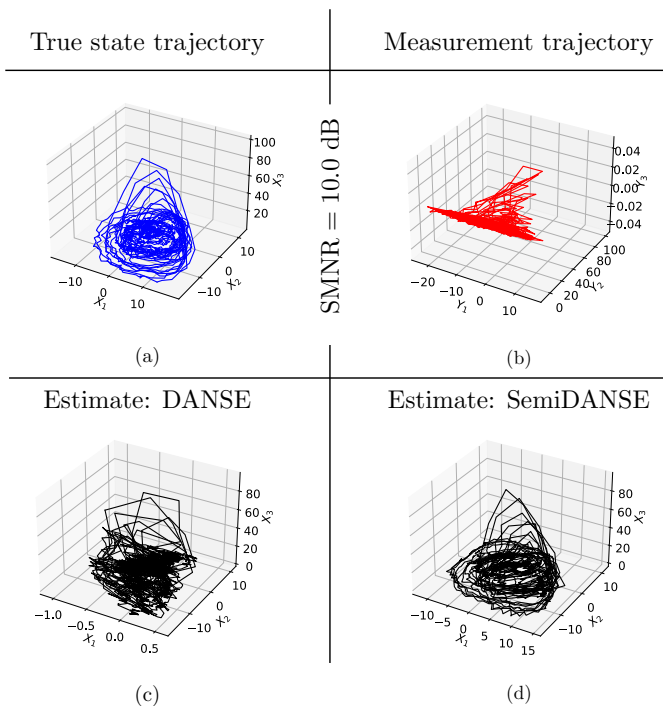


Fig. 13: Additional example of the failure of DANSE and success of SemiDANSE for BSCM problem with partial measurement of 3-dimensional Rössler process. We use the 2×3 measurement matrix \mathbf{H} shown in (29). A 3-dimensional true state trajectory from $\mathcal{D}_{\text{test}}$ with σ_e^2 corresponding to -15 dB is shown in (a). The corresponding measurement trajectory at SMNR = 10 dB is shown in (b). The estimates for DANSE and SemiDANSE ($\kappa = 0.02$) are posterior mean estimates shown in (c) and (d), respectively.

regularization brought about by the limited amount of labelled data. The performance of SemiDANSE also improves with the increased availability of supervision. Future work includes further exploration of the effect of semi-supervised learning for the general case of nonlinear, non-Gaussian measurements.

VI. ACKNOWLEDGEMENT

The authors would like to thank Dr. Sara Saeidian, KTH, for helpful comments on the structure and flow of the manuscript.

APPENDIX

A. Parameterization of the Gaussian prior using RNN

In this subsection, we describe the parameterization of the Gaussian prior using RNN in (2). The schematic representation is shown in Fig. 1. From Fig. 1, we note that the prior parameters $\mathbf{m}_{t|1:t-1}$ and $\mathbf{L}_{t|1:t-1}$ are calculated using an RNN together with feed-forward networks using $\mathbf{y}_{1:t-1}$ as input.

Let the internal state of the RNN at time τ be $\mathbf{z}_\tau \in \mathbb{R}^p$. For a causal system, we use unidirectional RNNs, where \mathbf{z}_τ generally depends on the *previous* internal state $\mathbf{z}_{\tau-1}$ and the current input \mathbf{y}_τ as

$$\mathbf{z}_\tau = \varphi(\mathbf{W}_{zz}\mathbf{z}_{\tau-1} + \mathbf{b}_{zz} + \mathbf{W}_{zy}\mathbf{y}_\tau + \mathbf{b}_{zy}), \quad (31)$$




Fig. 14: The average NMSE (in dB) on $\mathcal{D}_{\text{test}}$ versus n performances to illustrate the success of SemiDANSE ($\kappa = 0.04, 0.08$) for the BSCM problem setup for the 20-dimensional Lorenz-96 system (described in Section IV-D), with σ_e^2 corresponding to -10 dB and SMNR corresponding to 10 dB. n is varied from 15, 10, 5, 2 with $m = 20$ kept fixed. The comparison is against the unsupervised learning-based DANSE.

where $\mathbf{W}_{zz} \in \mathbb{R}^{p \times p}$, $\mathbf{b}_{zz} \in \mathbb{R}^{p \times 1}$, $\mathbf{W}_{zy} \in \mathbb{R}^{p \times m}$, $\mathbf{b}_{zy} \in \mathbb{R}^{p \times 1}$ denote the learnable weights and biases for the connections between the hidden states and the inputs and $\varphi(\cdot)$ denotes a nonlinear function [47, Chap. 10]. For the simplest RNN architectures such as a vanilla RNN, $\varphi(\cdot)$ can be as simple as an element-wise activation like a hyperbolic tangent function ($\tanh(\cdot)$) or can be much more complicated for sophisticated architectures such as GRUs and LSTMs involving gating functions [39], [43]. In our case, $\varphi(\cdot)$ is the nonlinear function for a GRU [39]. The parameters of the RNN \mathbf{W}_{zz} , \mathbf{b}_{zz} , \mathbf{W}_{zy} , \mathbf{b}_{zy} are shared across time and learning happens in accordance with backpropagation through time (BPTT) [48]. The internal state at the $(t-1)$ 'th time instant \mathbf{z}_{t-1} captures the information of the sequence $\mathbf{y}_{1:t-1}$. We then map \mathbf{z}_{t-1} to the mean $\mathbf{m}_{t|1:t-1}$ and the covariance $\mathbf{L}_{t|1:t-1}$ using two shallow feed-forward networks FC_m, FC_L through an intermediate representation \mathbf{z}'_{t-1} as follows

$$\begin{aligned} \mathbf{z}'_{t-1} &= \text{ReLU}(\mathbf{W}_{z'z}\mathbf{z}_{t-1} + \mathbf{b}_{z'z}), \\ \mathbf{m}_{t|1:t-1} &= \text{FC}_m(\mathbf{z}_{t-1}) \\ &= \mathbf{W}_{mz'}\mathbf{z}'_{t-1} + \mathbf{b}_{mz'}, \\ \mathbf{L}_{t|1:t-1} &= \text{diag}(\text{FC}_L(\mathbf{z}_{t-1})) \\ &= \text{diag}\left(\text{softplus}\left(\mathbf{W}_{Lz'}\mathbf{z}'_{t-1} + \mathbf{b}_{Lz'}\right)\right), \end{aligned} \quad (32)$$

where $\mathbf{W}_{mz'} \in \mathbb{R}^{m \times p'}$, $\mathbf{b}_{mz'} \in \mathbb{R}^{m \times 1}$, $\mathbf{W}_{Lz'} \in \mathbb{R}^{m \times p'}$, $\mathbf{b}_{Lz'} \in \mathbb{R}^{m \times 1}$, $\mathbf{W}_{z'z} \in \mathbb{R}^{p' \times p}$, $\mathbf{b}_{z'z} \in \mathbb{R}^{p' \times 1}$ are learnable weights and biases, $\text{ReLU}(x) = \max(0, x)$ is the element-wise, rectified linear unit function and $\text{softplus}(x) = \log_e(1 + \exp(x))$ is the element-wise, smooth approximation to $\text{ReLU}(x)$ [44]. As explained in section II-B, $\mathbf{L}_{t|1:t-1}$ is modelled as a diagonal covariance matrix as per (32). Thus, as per the notation in (2), the full set of learnable parameters is $\boldsymbol{\theta} = \{\mathbf{W}_{zz}, \mathbf{b}_{zz}, \mathbf{W}_{zy}, \mathbf{b}_{zy}, \mathbf{W}_{mz'}, \mathbf{b}_{mz'}, \mathbf{W}_{Lz'}, \mathbf{b}_{Lz'}, \mathbf{W}_{z'z}, \mathbf{b}_{z'z}\}$.

The specific RNN architectures and the feed-forward networks are chosen by grid-search. and described in Section IV-B.

B. Three chaotic dynamical systems

In this subsection, we briefly describe the mathematical models of the three chaotic dynamical systems used for experiments in section IV-A. The first two chaotic dynamical systems - the Lorenz-63 attractor [33] and the Chen attractor [35] were developed independently but have been shown to possess a generalized Lorenz canonical form in [34]. We used the discretized form of the Lorenz-63 attractor described in [21], as follows:

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{f}_t^{(L)}(\mathbf{x}_t) + \mathbf{e}_t \in \mathbb{R}^3, \\ &= \mathbf{F}_t^{(L)}(\mathbf{x}_t)\mathbf{x}_t + \mathbf{e}_t \in \mathbb{R}^3, \end{aligned} \quad (33)$$

$$\text{s.t. } \mathbf{F}_t^{(L)}(\mathbf{x}_t) = \exp \left(\begin{bmatrix} -10 & 10 & 0 \\ 28 & -1 & -x_{t,1} \\ 0 & x_{t,1} & -\frac{8}{3} \end{bmatrix} \Delta \right),$$

where the process noise $\mathbf{e}_t \sim \mathcal{N}(\mathbf{e}_t; \mathbf{0}, \mathbf{C}_e)$ with $\mathbf{C}_e = \sigma_e^2 \mathbf{I}_3$, the step-size $\Delta = 0.02$ seconds. σ_e^2 was set corresponding to -10 dB. The subscript (L) refers to the Lorenz attractor. In our simulations, we use a finite-Taylor series approximation of 5'th order for $\mathbf{F}_t^{(L)}(\mathbf{x}_t)$.

The Chen attractor has a mathematically similar form to the Lorenz-63 attractor except that the constants in the dynamical model are different. Similar to (33), it has the following mathematical form

$$\mathbf{x}_{t+1} = \mathbf{F}_t^{(C)}(\mathbf{x}_t)\mathbf{x}_t + \mathbf{e}_t^{(C)} \in \mathbb{R}^3, \quad (34)$$

$$\text{s.t. } \mathbf{F}_t^{(C)}(\mathbf{x}_t) = \exp \left(\begin{bmatrix} -35 & 35 & 0 \\ -7 & 28 & -x_{t,1} \\ 0 & x_{t,1} & -3 \end{bmatrix} \Delta' \right),$$

where $\mathbf{e}_t^{(C)} \sim \mathcal{N}(\mathbf{e}_t^{(C)}; \mathbf{0}, \mathbf{C}_e^{(C)})$ with $\mathbf{C}_e^{(C)} = \sigma_e^2 \mathbf{I}_3$, the step-size $\Delta' = 0.002$ seconds. σ_e^2 was set corresponding to -10 dB. In our simulations, we use a finite-Taylor series approximation of 5'th order for $\mathbf{F}_t^{(C)}(\mathbf{x}_t)$.

The Rössler attractor [36] appears simpler than the above two in terms of non-linear relationships among the state variables. Similar to (33), (34), the mathematical form is

$$\mathbf{x}_{t+1} = \mathbf{F}_t^{(R)}(\mathbf{x}_t)\mathbf{x}_t + \mathbf{e}_t^{(R)} \in \mathbb{R}^3, \quad (35)$$

$$\text{s.t. } \mathbf{F}_t^{(R)}(\mathbf{x}_t) = \exp \left(\begin{bmatrix} 0 & -1 & -1 \\ 1 & 0.2 & 0 \\ 0 & 0 & \frac{0.2}{x_{t,3}} + (x_{t,1} - 5.7) \end{bmatrix} \Delta'' \right),$$

where $\mathbf{e}_t^{(R)} \sim \mathcal{N}(\mathbf{e}_t^{(R)}; \mathbf{0}, \mathbf{C}_e^{(R)})$ with $\mathbf{C}_e^{(R)} = \begin{bmatrix} \sigma_e^2 \mathbf{I}_2 & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{1 \times 2} & \epsilon \end{bmatrix}$, $\epsilon \in \mathbb{R}_+$ is constant, the step-size $\Delta'' = 0.008$ seconds. In our simulations, we use a finite-Taylor series approximation of 5'th order for $\mathbf{F}_t^{(R)}(\mathbf{x}_t)$. For all the three systems, the measurement system is the linear measurement system in (1), with the measurement noise is $\mathbf{w}_t \sim \mathcal{N}(\mathbf{w}_t; \mathbf{0}, \mathbf{C}_w)$ with $\mathbf{C}_w = \sigma_w^2 \mathbf{I}_3$. σ_e^2 was set corresponding to -15 dB.

The truncated, reduced process noise variance is used in the simulation of the Rössler attractor in (35) to avoid numerical problems owing to incorporating noise in the dynamics. Since all the systems do not have the same step-size, in order to get a trajectory of length T with a similar degree of captured dynamics as the Lorenz-63 attractor, the other ones viz. Chen and Rössler are simulated for longer T , and then decimated in time by an appropriate factor (depending on the step-size) to get a T -length trajectory.

C. DMM - Learning problem

The deep Markov model (DMM) method proposed in [19] is well-suited to the modeling of complex time-series signals. A detailed description of the different inference modes and training architectures of the DMM can be found in [19], [16, Chap. 5]. Here, we provide a brief description of the DMM that we used for comparison in section IV. The DMM method relies on a choice of an *approximate* posterior distribution $q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})$ parameterized using a deep recurrent neural network with parameters $\boldsymbol{\phi}$. The factorization of this approximate posterior defines the inference mode used. E.g. in this work, we employed the structured left-information (ST-L) mode, as it gives rise to a causal filtering scenario [19]. The approximate posterior would then factorize as

$$q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi}) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}\mathbf{y}_{1:t}; \boldsymbol{\phi}). \quad (36)$$

In addition to this, the DMM method assumes that the process state dynamics is Markovian and tries to learn a model $p(\mathbf{x}_t|\mathbf{x}_{t-1}; \boldsymbol{\psi})$ with learnable parameters $\boldsymbol{\psi}$. This is achieved by optimizing a variational lower bound (VLB) using (1), (36) and the Markovian assumption as follows:

$$\begin{aligned} & \log p(\mathbf{y}_{1:T}; \boldsymbol{\psi}) \\ &= \log \left(\int \cdots \int p(\mathbf{y}_{1:T}, \mathbf{x}_{1:T}; \boldsymbol{\psi}) d\mathbf{x}_1 \cdots d\mathbf{x}_T \right) \\ &= \log \left(\mathbb{E}_{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \left[\frac{p(\mathbf{y}_{1:T}, \mathbf{x}_{1:T}; \boldsymbol{\psi})}{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \right] \right) \\ &\stackrel{(a4)}{\geq} \mathbb{E}_{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \left[\log \left(\frac{p(\mathbf{y}_{1:T}|\mathbf{x}_{1:T}; \boldsymbol{\psi}) p(\mathbf{x}_{1:T}; \boldsymbol{\psi})}{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \right) \right] \\ &= \mathbb{E}_{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \left[\log p(\mathbf{y}_{1:T}|\mathbf{x}_{1:T}; \boldsymbol{\psi}) \right] \\ &\quad + \mathbb{E}_{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \left[\log \left(\frac{p(\mathbf{x}_{1:T}; \boldsymbol{\psi})}{q(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}; \boldsymbol{\phi})} \right) \right] \\ &\stackrel{(b4)}{=} \sum_{t=1}^T \mathbb{E}_{q(\mathbf{x}_t|\mathbf{y}_{1:t}; \boldsymbol{\phi})} [\log \mathcal{N}(\mathbf{y}_t; \mathbf{H}\mathbf{x}_t; \mathbf{C}_w)] \\ &\quad - \mathbb{E}_{q(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}; \boldsymbol{\phi})} [\text{D}_{\text{KL}}(q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{y}_{1:t}; \boldsymbol{\phi}) || p(\mathbf{x}_t|\mathbf{x}_{t-1}; \boldsymbol{\psi}))] \\ &\triangleq \mathcal{L}(\mathbf{y}_{1:T}; \boldsymbol{\psi}, \boldsymbol{\phi}), \end{aligned}$$

where the inequality in (a4) is obtained by Jensen's inequality and the detailed derivation of the equality in (b4) is detailed in [16, Chap. 5]. In practice, we evaluate the above expectations using Monte-Carlo approximations and sequential sampling with the reparameterization trick [49], making it quite computationally expensive. For a training dataset $\mathcal{D} = \{\mathbf{y}_{1:T}^{(i)}\}_{i=1}^N$, the optimization problem then becomes

$$\boldsymbol{\psi}^*, \boldsymbol{\phi}^* = \arg \min_{\boldsymbol{\psi}, \boldsymbol{\phi}} \sum_{i=1}^N -\mathcal{L}(\mathbf{y}_{1:T}^{(i)}; \boldsymbol{\psi}, \boldsymbol{\phi}). \quad (37)$$

The inference follows by sampling sequentially from the learned, approximate posterior $q(\mathbf{x}_t|\mathbf{x}_{t-1}\mathbf{y}_{1:t}; \boldsymbol{\phi}^*)$ using the reparameterization trick.

REFERENCES

- [1] S. Särkkä and L. Svensson, *Bayesian filtering and smoothing*, vol. 17, Cambridge university press, 2023.
- [2] D. Simon, *Optimal state estimation: Kalman, H_∞ , and nonlinear approaches*, John Wiley & Sons, 2006.
- [3] T. D. Barfoot, *State estimation for robotics*, Cambridge University Press, 2024.
- [4] R.E. Kalman, “New results in linear filtering and prediction theory,” *J. Basic Eng.*, vol. 83, pp. 95–108, 1961.
- [5] R.E. Kalman, “A new approach to linear filtering and prediction problems,” *Trans. ASME, D*, vol. 82, pp. 35–44, 1960.
- [6] M. Gruber, “An approach to target tracking,” Tech. Rep., MIT Lexington Lincoln Lab, 1967.
- [7] S.J. Julier and J.K. Uhlmann, “Unscented filtering and nonlinear estimation,” *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.
- [8] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, “A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking,” *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [9] J. Ko and D. Fox, “GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models,” *Autonomous Robots*, vol. 27, pp. 75–90, 2009.
- [10] R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen, “Bayesian inference and learning in gaussian process state-space models with particle mcmc,” *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [11] A. Svensson, A. Solin, S. Särkkä, and T. B. Schön, “Computationally efficient bayesian learning of gaussian process state space models,” in *Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2016, pp. 213–221.
- [12] L. Xu and R. Niu, “EKFNet: Learning system noise statistics from measurement data,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 4560–4564.
- [13] V. Garcia Satorras, Z. Akata, and M. Welling, “Combining generative and discriminative models for hybrid inference,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [14] S. Jung, I. Schlangen, and A. Charlish, “A mnemonic Kalman filter for non-linear systems with extensive temporal dependencies,” *IEEE Signal Processing Letters*, vol. 27, pp. 1005–1009, 2020.
- [15] H. Wen, X. Chen, G. Papagiannis, C. Hu, and Y. Li, “End-to-end semi-supervised learning for differentiable particle filters,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 5825–5831.
- [16] L. Girin, S. Leglaive, X. Bie, J. Diard, T. Hueber, and X. Alameda-Pineda, “Dynamical variational autoencoders: A comprehensive review,” *Foundations and Trends in Machine Learning*, vol. 15, no. 1-2, pp. 1–175, 2021.
- [17] M. Fraccaro, S. Kamronn, U. Paquet, and O. Winther, “A disentangled recognition and nonlinear dynamics model for unsupervised learning,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [18] R.G. Krishnan, U. Shalit, and D. Sontag, “Deep Kalman filters,” *arXiv preprint arXiv:1511.05121*, 2015.
- [19] R. Krishnan, U. Shalit, and D. Sontag, “Structured inference networks for nonlinear state space models,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017, vol. 31.
- [20] A. Ghosh, A. Honoré, and S. Chatterjee, “DANSE: Data-Driven Non-Linear State Estimation of Model-Free Process in Unsupervised Bayesian Setup,” in *2023 31st European Signal Processing Conference (EUSIPCO)*, 2023, pp. 870–874.
- [21] A. Ghosh, A. Honoré, and S. Chatterjee, “DANSE: Data-Driven Non-Linear State Estimation of Model-Free Process in Unsupervised Learning Setup,” *IEEE Transactions on Signal Processing*, vol. 72, pp. 1824–1838, 2024.
- [22] G. Revach, N. Shlezinger, T. Locher, X. Ni, R.J.G. van Sloun, and Y.C. Eldar, “Unsupervised learned Kalman filtering,” in *2022 30th European Signal Processing Conference (EUSIPCO)*. IEEE, 2022, pp. 1571–1575.
- [23] G. Revach, N. Shlezinger, X. Ni, A. L. Escoriza, R.J.G. Van Sloun, and Y.C. Eldar, “KalmanNet: Neural network aided Kalman filtering for partially known dynamics,” *IEEE Transactions on Signal Processing*, vol. 70, pp. 1532–1547, 2022.
- [24] X. Ni, G. Revach, and N. Shlezinger, “Adaptive Kalmannet: Data-Driven Kalman Filter with Fast Adaptation,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 5970–5974.
- [25] G. Choi, J. Park, N. Shlezinger, Y. C. Eldar, and N. Lee, “Split-KalmanNet: A Robust Model-Based Deep Learning Approach for State Estimation,” *IEEE Transactions on Vehicular Technology*, vol. 72, no. 9, pp. 12326–12331, Sept. 2023.
- [26] S. Chen, Y. Zheng, D. Lin, P. Cai, Y. Xiao, and S. Wang, “MAML-KalmanNet: A Neural Network-Assisted Kalman Filter Based on Model-Agnostic Meta-Learning,” *IEEE Transactions on Signal Processing*, 2025.
- [27] T. Li, Y. Song, and H. Fan, “From target tracking to targeting track: A data-driven yet analytical approach to joint target detection and tracking,” *Signal Processing*, vol. 205, pp. 108883, 2023.
- [28] O. Chapelle, B. Schölkopf, and A. Zien, *Semi-Supervised Learning*, The MIT Press, 1st edition, 2010.
- [29] X. Zhu and A. B. Goldberg, “Introduction to semi-supervised learning,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2009.
- [30] X. Yang, Z. Song, I. King, and Z. Xu, “A survey on deep semi-supervised learning,” *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [31] J. E. Van Engelen and H. H. Hoos, “A survey on semi-supervised learning,” *Machine learning*, vol. 109, no. 2, pp. 373–440, 2020.
- [32] G. Kostopoulos, S. Karlos, S. Kotsiantis, and O. Ragos, “Semi-supervised regression: A recent review,” *Journal of Intelligent & Fuzzy Systems*, vol. 35, no. 2, pp. 1483–1500, 2018.
- [33] E.N. Lorenz, “Deterministic nonperiodic flow,” *Journal of atmospheric sciences*, vol. 20, no. 2, pp. 130–141, 1963.
- [34] S. Čelikovský and G. Chen, “On the generalized lorenz canonical form,” *Chaos, Solitons & Fractals*, vol. 26, no. 5, pp. 1271–1276, 2005.
- [35] G. Chen and T. Ueta, “Yet another chaotic attractor,” *International Journal of Bifurcation and chaos*, vol. 9, no. 07, pp. 1465–1466, 1999.
- [36] O.E. Rössler, “An equation for continuous chaos,” *Physics Letters A*, vol. 57, no. 5, pp. 397–398, 1976.
- [37] Edward N Lorenz, “Predictability: A problem partly solved,” in *Proc. Seminar on predictability*. Reading, 1996, vol. 1, pp. 1–18.
- [38] Edward N Lorenz and Kerry A Emanuel, “Optimal sites for supplementary weather observations: Simulation with a small model,” *Journal of the Atmospheric Sciences*, vol. 55, no. 3, pp. 399–414, 1998.
- [39] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder–decoder approaches,” in *Proceedings of 8th Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8)*, 2014, pp. 103–111.
- [40] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*, vol. 4, Springer, 2006.
- [41] J. Gordon and J. M. Hernández-Lobato, “Bayesian semisupervised learning with deep generative models,” *arXiv preprint arXiv:1706.09751*, 2017.
- [42] J. C. Sprott, *Chaos and time-series analysis*, Oxford university press, 2003.
- [43] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [44] A. Paszke et al., “PyTorch: An imperative style, high-performance deep learning library,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [45] D.P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations (ICLR)*, 2015.
- [46] R. Labbe, “FilterPy - Kalman and Bayesian filters in Python,” URL: <https://filterpy.readthedocs.io/en/latest/>, 2018.
- [47] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*, MIT press, 2016.
- [48] P. J. Werbos, “Backpropagation through time: what it does and how to do it,” *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [49] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, (ICLR)*, Yoshua Bengio and Yann LeCun, Eds., 2014.