

# AI-Aided Kalman Filters

Nir Shlezinger, *Senior Member, IEEE*, Guy Revach, *Senior Member, IEEE*, Anubhab Ghosh, *Student Member, IEEE*, Saikat Chatterjee, *Senior Member, IEEE*, Shuo Tang, *Student Member, IEEE*, Tales Imbiriba, *Member, IEEE*, Jindrich Dunik, *Senior Member, IEEE*, Ondrej Straka, *Member, IEEE*, Pau Closas, *Senior Member, IEEE*, and Yonina C. Eldar, *Fellow, IEEE*

**Abstract**— The Kalman filter (KF) and its variants are among the most celebrated algorithms in signal processing. These methods are used for state estimation of dynamic systems by relying on mathematical representations in the form of simple state-space (SS) models, which may be crude and inaccurate descriptions of the underlying dynamics. Emerging data-centric artificial intelligence (AI) techniques tackle these tasks using deep neural networks (DNNs), which are model-agnostic. Recent developments illustrate the possibility of fusing DNNs with classic Kalman-type filtering, obtaining systems that learn to track in partially known dynamics. This article provides a tutorial-style overview of design approaches for incorporating AI in aiding KF-type algorithms. We review both generic and dedicated DNN architectures suitable for state estimation, and provide a systematic presentation of techniques for fusing AI tools with KFs and for leveraging partial SS modeling and data, categorizing design approaches into *task-oriented* and *SS model-oriented*. The usefulness of each approach in preserving the individual strengths of model-based KFs and data-driven DNNs is investigated in a qualitative and quantitative study, whose code is publicly available, illustrating the gains of hybrid model-based/data-driven designs. We also discuss existing challenges and future research directions that arise from fusing AI and Kalman-type algorithms.

## I. INTRODUCTION

The Apollo program, which successfully landed the first humans on the moon, is still considered one of mankind’s greatest achievements. Among the technological innovations that played a role in the success of the Apollo program is a filtering method based on the extended version of an algorithm developed by Rudolf Kalman in the late 1950s [1], which was used to track and estimate the trajectory of the spaceship. This algorithm, known as the Kalman filter (KF), and its extension into the extended KF (EKF) [2], provide an accurate and reliable real-time trajectory estimation while still being simple to implement and applicable on the hardware-limited

navigation computer of the Apollo spaceship [3]. To date, the KF is among the most celebrated algorithms in signal processing and electrical engineering at large, with numerous applications including radar, biomedical systems, vehicular technology, navigation, wireless communications, etc.

The KF is a *model-based* method, namely, it leverages mathematical parametric representations of the environment. Specifically, the KF and its variants [4] rely on the ability to describe the underlying dynamics as a state-space (SS) model. Such model-based designs have several core advantages when the dynamics are faithfully captured using a known and tractable SS representation: (i) they can achieve *optimal performance*. For example, the KF achieves the minimal mean-squared error (MSE) for linear Gaussian SS models; (ii) model-based methods operate with *controllable and often reduced complexity*. In fact, the KF and its variants are implemented on devices such as sensors and mobile systems, where they operate in real-time using limited computational and power resources; (iii) decisions made by the KF and its variants are *interpretable*, in the sense that their internal features have concrete meaning and their reasoning can be explained from their observations and underlying SS model; (iv) they are inherently *adaptive*, as changes in the SS model are naturally incorporated into the operation; and (v) they reliably characterize *uncertainty* in their estimate, providing the error covariance alongside their estimates [5].

A key characteristic of KF-type algorithms is their reliance on knowledge of the underlying dynamics and specifically, on the ability to accurately capture these dynamics using a known and tractable SS representation. A common approach is to rely on simplifying assumptions (e.g., linear systems, Gaussian mutually and temporally independent process and measurement noises, etc.) that make models understandable and the associated algorithms computationally efficient, and then use data to estimate the unknown parameters. However, simple models frequently fail to represent some of the nuances and subtleties of dynamic systems and their associated signals. Practical applications are thus often required to operate with partially known SS models. Furthermore, the performance and reliability of KF-type algorithms degrade considerably when using a postulated SS model that deviates from the true nature of the underlying system.

The unprecedented success of machine learning (ML), and particularly deep learning, as the enabler technology for artificial intelligence (AI) in areas such as computer vision has initiated a general mindset geared towards data. It is now quite common practice to replace simple principled models with data-driven pipelines, employing ML architectures based

N. Shlezinger is with the School of ECE, Ben-Gurion University of the Negev, Beer-Sheva, Israel (e-mail: nirshl@bgu.ac.il). G. Revach is with the Institute for Signal and Information Processing, D-ITET, ETH Zürich, Switzerland (e-mail: grevach@ethz.ch). A. Ghosh and S. Chatterjee are with the School of EE and CS, KTH Royal Institute of Technology, Stockholm, Sweden (e-mail: {anubhabg; sach}@kth.se). J. Dunik and O. Straka are with the Faculty of Applied Science, University of West Bohemia, Czech Republic (email: {dunikj; straka30}@kky.zcu.cz). T. Imbiriba is with the Dept. of CS, University of Massachusetts Boston, Boston, MA, USA (email: tales.imbiriba@umb.edu). S. Tang and P. Closas are with the Dept. of ECE, Northeastern University, Boston, MA, USA (email: {tang.shu; closas}@northeastern.edu). Y. C. Eldar is with the Math and CS Faculty, Weizmann Institute of Science, Rehovot, Israel (e-mail: yonina.eldar@weizmann.ac.il).

This work has been partially supported by the Israel Innovation Authority, and by the National Science Foundation under Awards ECCS-1845833 and CCF-2326559.

on deep neural networks (DNNs) that are trained with massive volumes of labeled data. DNNs can be trained end-to-end without relying on analytical approximations, and therefore, they can operate in scenarios where analytical models are unknown or highly complex [6]. With the growing popularity of deep learning, recent years have witnessed the design of various DNNs for tasks associated with KF-type algorithms, including, e.g., the application of DNNs for state estimation [7], [8] and control [9].

AI systems can learn from data to track dynamic systems without relying on full knowledge of underlying SS models. However, replacing KF-type algorithms with deep architectures gives rise to several core challenges. For one, the *computational burden* of training and utilizing highly parametrized DNNs, as well as the fact that *massive data sets* are typically required for their training, often constitute major drawbacks in various signal-processing applications. This limitation is particularly relevant for hardware-constrained devices (such as mobile systems and sensors) whose ability to utilize highly parametrized DNNs is typically limited [10]. Furthermore, due to the complex and generic structure of DNNs, it is often *challenging to understand* how they obtain their predictions, or track the rationale leading to their decisions. Moreover, DNNs typically struggle to adapt to variations in the data distribution and are limited in their capability to provide estimation of uncertainty.

The challenges outlined above gave rise to a growing interest in using AI not to replace KF-type algorithms, but to empower them, as a form of hybrid model-based/data-driven design [11]. Various approaches have been proposed for combining KFs and deep learning including training DNNs to map data into features that obey a known SS model and can be tracked with a KF [12], [13]; employing deep models for identifying SS models to be used for KF-based tracking [14], [15] and converting the KF into an ML model that can be trained in a supervised [16]–[19] or unsupervised [20], [21] manner. Such AI-empowered KFs were already in various signal processing applications, ranging from brain-machine interface, acoustic echo cancellation, financial monitoring, beam tracking in wireless systems, and drone-based monitoring systems [22], [23]. These recent advances in combining AI and KFs, along with their implications on emerging technologies, motivate a systematic presentation of the different design approaches, as well as their associated signal processing challenges.

In this article, we provide a tutorial-style overview of the design approaches and potential benefits of combining deep learning tools with KF-type algorithms. We refer to this as ‘AI-aided Kalman Filters’. While highlighting the potential benefits AI-aided KFs, we also mention the main signal processing challenges that arise from such hybrid designs. For this goal, we begin by reviewing the fundamentals of KF that are relevant to understanding its fusion with AI. We briefly describe its core statistical representation – the SS model – and formulate the mathematical steps for filtering and smoothing. We then discuss the pros and cons associated with KF-type algorithms and pinpoint the main challenging aspects that motivate the usage of AI tools. We then shift

our focus to deep learning techniques that are suitable for processing time sequences, briefly reviewing both generic time-sequence architectures, such as recurrent neural networks (RNNs) and attention mechanisms, and proceeding to DNN architectures that are inspired by SS representations [24] and by KF processing flow [7]. We discuss the gains offered by such data-driven techniques, while also highlighting their limitations, in turn indicating the potential of combining AI techniques with classic SS model-based KF-type methods.

The bulk of the article is dedicated to presenting design approaches that combine deep learning techniques with KF-type processing based on a partial mathematical representation of the dynamics. We categorize existing design approaches, drawing inspiration from the ML paradigms of *discriminative* (task-oriented) learning and *generative* (statistical model-oriented) learning [25], [26]. Accordingly, the first part is dedicated to AI-augmented KFs via task-oriented learning, presenting in detail candidate approaches for converting KF-type filtering into an ML architecture via DNN-augmentation. These include designs that employ an external DNN, either sequentially [27] or in parallel [19], as well as DNNs augmented into a KF-type algorithm, e.g., for Kalman gain (KG) computation [16]. We then proceed to detail SS model-oriented AI-aided KFs designs. These can be viewed as a form of AI-aided system identification, i.e., techniques that utilize DNNs in the process of identifying SS models, which can then be used by model-based KF-type tracking [14], [15], [28]. We discuss several approaches with various structures connecting physics-based and data-based model components, such as incremental, hybrid, or fixed SS structures.

To highlight the gains of each approach for designing AI-aided KFs, and capture the interplay between the different algorithms and conventional model-based or contemporary data-driven methods, we provide a comparative study. We begin with a qualitative comparison, pinpointing the conceptual differences between the design approaches in terms such as the level of domain knowledge required, the type of data needed, and flexibility. We also provide a quantitative comparison, where we evaluate representative methods from the presented design approaches in a setting involving the tracking of the challenging Lorenz attractor. The article concludes with a discussion of open challenges and future research directions.

**Notations:** In the following, we use boldface lowercase for vectors, e.g.,  $\mathbf{x}$  and boldface uppercase letters, e.g.,  $\mathbf{X}$  for matrices. For a time sequence  $\mathbf{x}_t$  and time indices  $t_1 \leq t_2$ , we use the abbreviated form  $\mathbf{x}_{t_1:t_2}$  for the set of subsequent variables  $\{\mathbf{x}_t\}_{t=t_1}^{t_2}$ . We use  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  for the multivariate Gaussian distribution with mean  $\boldsymbol{\mu}$  and covariance  $\boldsymbol{\Sigma}$ , while  $p(\cdot)$  is a probability density function. The notation  $\|\mathbf{x}\|_{\mathbf{C}}^2$  denotes the squared  $\ell_2$  norm of  $\mathbf{x}$  weighted by the matrix  $\mathbf{C}$ , i.e.  $\|\mathbf{x}\|_{\mathbf{C}}^2 = \mathbf{x}^\top \mathbf{C} \mathbf{x}$ . The operations  $(\cdot)^\top$  and  $\|\cdot\|_2$  are used for transpose and  $\ell_2$  norm, respectively.

## II. FUNDAMENTALS OF KALMAN FILTERING

In this section, we review some basics of Kalman-type filtering. We commence with reviewing SS models, after which we recall the formulation of the KF and EKF. We conclude

this section by highlighting the challenges that motivate its augmentation with deep learning.

### A. State-Space Models

SS models are a class of mathematical models that describe the probabilistic behavior of dynamical systems. They constitute the fundamental framework for formulating a broad range of engineering problems in the areas of signal processing, control, and communications, and they are also widely used in environment studies, economics, and many more. The core of SS models lies in the representation of the dynamics of a system using a latent state variable that evolves over time while being related to the observations made by the system.

**Generic SS Models:** Focusing on discrete-time formulations of continuous-valued variables, SS models represent the interplay between the observations at time instant  $t$ , denoted  $\mathbf{y}_t$ , an input signal  $\mathbf{u}_t$ , and a latent state capturing the system dynamics  $\mathbf{x}_t$ . In general, SS models consist of: (i) a *state evolution* model of the form

$$\mathbf{x}_{t+1} = \tilde{f}_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{v}_t), \quad (1a)$$

representing how the state evolves in time; and (ii) an observation model of the form

$$\mathbf{y}_t = \tilde{h}_t(\mathbf{x}_t, \mathbf{w}_t), \quad (1b)$$

which relates the observations and the current system state. While the mappings  $\tilde{f}_t(\cdot)$  and  $\tilde{h}_t(\cdot)$  are deterministic, stochasticity is induced by the temporally independent noises  $\mathbf{v}_t$  and  $\mathbf{w}_t$ , and by the distribution of the initial state  $\mathbf{x}_0$ .

**Gaussian SS Models:** The most common special case of the above generic model is that of the linear Gaussian SS model, used by the celebrated KF. In this model, the state evolution takes the form

$$\mathbf{x}_{t+1} = \mathbf{F}_t \mathbf{x}_t + \mathbf{G}_t \mathbf{u}_t + \mathbf{v}_t, \quad (2a)$$

and the observation model is given by

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

where  $\mathbf{F}_t, \mathbf{G}_t, \mathbf{H}_t$  are matrices of appropriate dimensions, and  $\mathbf{v}_t$  and  $\mathbf{w}_t$  are temporally and mutually independent Gaussian noise signals, with zero-mean and covariances  $\mathbf{Q}_t$  and  $\mathbf{R}_t$ , respectively. Namely,  $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$  and  $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$ , while  $\mathbf{v}_t$  is independent  $\mathbf{w}_t$ , and both are independent of  $\mathbf{v}_\tau$  and  $\mathbf{w}_\tau$  for any  $\tau \neq t$ . The initial state  $\mathbf{x}_0$  is independent of the noises and also assumed Gaussian, with known mean  $\hat{\mathbf{x}}_0$  and covariance  $\hat{\Sigma}_0$ .

Similar models to (2) are frequently utilized for settings characterized by non-linear transformations. In the non-linear additive Gaussian SS model (termed henceforth as non-linear Gaussian), (1) takes the form

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{v}_t, \quad (3a)$$

$$\mathbf{y}_t = h_t(\mathbf{x}_t) + \mathbf{w}_t, \quad (3b)$$

where the noise signals  $\mathbf{v}_t$  and  $\mathbf{w}_t$  are as in (2), i.e., temporally independent with  $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$  and  $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$ .

**Tasks:** SS models are mostly associated with two main families of tasks. The first is *state estimation*, which deals

with the recovery of the state variable  $\mathbf{x}_t$  based on a set of observations  $\{\mathbf{y}_\tau\}$ , i.e., an open-loop system where the input  $\mathbf{u}_t$  is either absent or not controlled. Some of the most common state estimation tasks are [4]

- *Filtering*: estimate  $\mathbf{x}_t$  from  $\mathbf{y}_{1:t}$ .
- *Smoothing*: estimate  $\mathbf{x}_{1:T}$  from  $\mathbf{y}_{1:T}$  for some  $T > 0$ .

Additional related tasks are prediction, input recovery, and imputation. State estimation plays a key role in applications that involve tracking, localization, and denoising, ranging from target tracking in radar systems to monitoring biomedical signals.

The second family of SS model-based tasks are those that deal with *stochastic control* (closed-loop) policies. In this family, the SS framework is used to select how to set the input variable  $\mathbf{u}_t$  based on, e.g., past measurements  $\mathbf{y}_{1:t}$ . Such tasks are fundamental in robotics, vehicular systems, and aerospace engineering. As stochastic control policies often employ state estimation schemes followed by state regulators [29], we focus in this article on state estimation (i.e., the first family of tasks).

### B. Kalman Filtering

The representation of dynamic systems via SS models gives rise to some of the most celebrated algorithms in signal processing, particularly the family of Kalman-type filters. To describe these, we henceforth focus on *state estimation*. Therefore, for convenience, we omit the input signal  $\mathbf{u}_t$  from the following relations (as it can also be absorbed into the state evolution noise as a known bias term).

**KF:** The KF is the minimal MSE estimator for the filtering task in linear Gaussian SS models, i.e., estimating  $\mathbf{x}_t$  from  $\mathbf{y}_{1:t}$  when these are related via (2). In every time step  $t$ , the KF estimates  $\mathbf{x}_t$  using only the previous estimate  $\hat{\mathbf{x}}_{t-1}$  as a sufficient statistic and the new observation  $\mathbf{y}_t$ , thus its complexity does not grow in time.

The KF updates its estimates of the first- and second-order statistical moments of the state, which at time  $t$  are denoted by  $\hat{\mathbf{x}}_{t|t}$  and  $\hat{\Sigma}_{t|t}$ , respectively. For  $t = 0$ , these moments are initialized to those of the initial state, namely  $\hat{\mathbf{x}}_0$  and  $\hat{\Sigma}_0$ , while for  $t > 0$  the estimates are obtained via a two-step procedure:

- 1) *Prediction*: The first step *predicts* the first- and second-order statistical moments of current *a priori* state and observation, based on the previous *a posteriori* estimate. Specifically, at time  $t$  the predicted moments of the state are computed via

$$\hat{\mathbf{x}}_{t|t-1} = \mathbf{F}_{t-1} \cdot \hat{\mathbf{x}}_{t-1|t-1}, \quad (4a)$$

$$\hat{\Sigma}_{t|t-1} = \mathbf{F}_{t-1} \cdot \hat{\Sigma}_{t-1|t-1} \cdot \mathbf{F}_{t-1}^\top + \mathbf{Q}_{t-1}. \quad (4b)$$

The predicted moments of the observations are computed as

$$\hat{\mathbf{y}}_{t|t-1} = \mathbf{H}_t \cdot \hat{\mathbf{x}}_{t|t-1}, \quad (5a)$$

$$\hat{\Sigma}_{t|t-1} = \mathbf{H}_t \cdot \hat{\Sigma}_{t|t-1} \cdot \mathbf{H}_t^\top + \mathbf{R}_t. \quad (5b)$$

- 2) *Update*: The predicted *a priori* moments are updated using the current observation  $\mathbf{y}_t$  into the *a posteriori* state

moments. The updated moments are computed as

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t \cdot \Delta \mathbf{y}_t, \quad (6a)$$

$$\boldsymbol{\Sigma}_{t|t} = \boldsymbol{\Sigma}_{t|t-1} - \mathbf{K}_t \cdot \mathbf{S}_{t|t-1} \cdot \mathbf{K}_t^\top. \quad (6b)$$

Here,  $\mathbf{K}_t$  is the KG, and it is given by

$$\mathbf{K}_t = \boldsymbol{\Sigma}_{t|t-1} \cdot \mathbf{H}_t^\top \cdot \mathbf{S}_{t|t-1}^{-1}. \quad (7)$$

The term  $\Delta \mathbf{y}_t \triangleq \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}$  is the defined as the *innovation*, representing the difference between the predicted observation and the observed value.

The output of the KF is the estimated state, i.e.,  $\hat{\mathbf{x}}_t \equiv \hat{\mathbf{x}}_{t|t}$ , and the error covariance is  $\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t}$ . Various proofs are provided in the literature for the MSE optimality of the KF [4]. A relatively compact way to prove this follows from the more general Bayes' rule and Chapman-Kolmogorov equation, recalled in the box entitled *From Chapman-Kolmogorov to the KF* on Page 5.

**Extension to Smoothing:** While the KF is formulated for the filtering task, it also naturally extends (while preserving its optimality) for smoothing tasks. A leading approach to realize a smoother is the Rauch-Tung-Striebel (RTS) algorithm [31], which employs two subsequent recursive passes to estimate the state, termed *forward* and *backward* passes. The forward pass is the standard KF, while refines the estimates for each time  $t$  using future observations at time instants  $\tau \in \{t+1, \dots, T\}$ .

The backward pass is similar in its structure to the update step in the KF. For each  $t \in \{T-1, \dots, 1\}$ , the forward belief is corrected with future estimates via

$$\hat{\mathbf{x}}_{t|T} = \hat{\mathbf{x}}_{t|t} + \overleftarrow{\mathbf{K}}_t \cdot (\hat{\mathbf{x}}_{t+1|T} - \hat{\mathbf{x}}_{t+1|t}), \quad (9a)$$

$$\boldsymbol{\Sigma}_{t|T} = \boldsymbol{\Sigma}_{t|t} - \overleftarrow{\mathbf{K}}_t \cdot (\boldsymbol{\Sigma}_{t+1|t} - \boldsymbol{\Sigma}_{t+1|T}) \cdot \overleftarrow{\mathbf{K}}_t^\top. \quad (9b)$$

Here,  $\overleftarrow{\mathbf{K}}_t$  is the backward KG, computed based on second-order statistical moments from the forward pass as

$$\overleftarrow{\mathbf{K}}_t = \boldsymbol{\Sigma}_{t|t} \cdot \mathbf{F}_t^\top \cdot \boldsymbol{\Sigma}_{t+1|t}^{-1}. \quad (10)$$

The output of the RTS is the estimated state for every time instant  $t \in \{1, \dots, T\}$ , with  $\hat{\mathbf{x}}_t \equiv \hat{\mathbf{x}}_{t|T}$ , as well as the error covariance  $\boldsymbol{\Sigma}_t \equiv \boldsymbol{\Sigma}_{t|T}$ .

**Extensions to Non-Linear SS Models:** The KF is MSE optimal for linear and Gaussian SS models. For non-linear settings, existing approaches vary between non-linear Gaussian SS models as in (3), and settings that are also non-Gaussian.

Filtering algorithms for state estimation in non-linear Gaussian SS models are typically designed to preserve the linear operation of the KF update step with respect to measurement. Specifically, the prediction of the first-order moments in (4a) and (5a), is replaced with

$$\hat{\mathbf{x}}_{t|t-1} = f_t(\hat{\mathbf{x}}_{t-1|t-1}), \quad (11a)$$

$$\hat{\mathbf{y}}_{t|t-1} = h_t(\hat{\mathbf{x}}_{t|t-1}). \quad (11b)$$

The key challenge in approximating the operation of the KF lies in the propagation of the second-order moments. Arguably the most common non-linear variant of the KF, known as the EKF, is based on local linearizations. Here, the matrices  $\mathbf{F}_t$

and  $\mathbf{H}_t$  in the KF formulations are respectively replaced with the Jacobian matrices of  $f_t(\cdot)$  and  $h_t(\cdot)$ , i.e.,

$$\hat{\mathbf{F}}_t = \nabla_{\mathbf{x}_{t-1}} f_{t-1}(\hat{\mathbf{x}}_{t-1|t-1}), \quad (12a)$$

$$\hat{\mathbf{H}}_t = \nabla_{\mathbf{x}_t} h_t(\hat{\mathbf{x}}_{t|t-1}). \quad (12b)$$

Alternatively, the propagation of second-order moments can be approximated using the unscented transform, resulting in the unscented KF (UKF), or using the cubature and Gauss-Hermite deterministic or stochastic quadrature rules.

The EKF and UKF are approximations of the KF designed for non-linear Gaussian SS models. When the SS is also non-Gaussian, state estimation algorithms typically aim at recursively updating the posterior distribution via the Chapman-Kolmogorov relation, without resorting to its Gaussian special case utilized by the KF (See box on page 5). Leading algorithms that operate in this manner include the family of particle filters, that are based on sequential sampling; Gaussian sum filters, based on Gaussian sum representation of all densities; and point-mass filters, numerically solving the Bayesian relation in a typically rectangular grid.

### C. Pros and Cons of Model-Based KF-Type Algorithms

The KF and its variants are a family of widely utilized and trusted state estimation algorithms [5]. The core of these model-based methods is the SS model, i.e., the mathematical representation of the system dynamics via closed-form equations, as those in (1). When the SS model faithfully captures the dynamics, these algorithms have several key desirable properties:

- P1 When the SS model is well described as being linear with Gaussian noise, KF-based algorithms can approach optimal state estimation performance, in the sense of minimizing the MSE.
- P2 Model-based methods are inherently adaptive to known variations in the SS model. For instance, the matrix  $\mathbf{H}_t$  can change with time index  $t$ , and one only needs to substitute the updated matrix in the corresponding equations.
- P3 Their operation is fully interpretable, in the sense that one can associate the internal features with concrete interpretation as they represent, e.g., statistical moments of prior and posterior predictions.
- P4 They provide reliable uncertainty measures via the error covariance in (6b) and (9b).
- P5 KF-type algorithms operate with relatively low complexity, that does not grow with time.

However, the reliance of KF-type algorithms on faithful mathematical modelling of the underlying dynamics, and their natural suitability with simplistic linear Gaussian models, also gives rise to several core challenges encountered in various applications. These challenges can be roughly categorized as follows:

- C1 The state evolution and observation models employed in SS models are often *approximations of the true dynamics*, whose fidelity can vary considerably between applications. For instance, while the temporal evolution of a state

## From Chapman-Kolmogorov to the KF

Consider a generic SS model as in (1) without an input signal  $\mathbf{u}_t$ . By Bayes' rule and the Markovian nature of (1), it holds that the conditional distribution of  $\mathbf{x}_t | \mathbf{y}_{1:t}$  satisfies

$$p(\mathbf{x}_t | \{\mathbf{y}_\tau\}_{\tau \leq t}) = \frac{p(\mathbf{y}_t | \mathbf{x}_t) p(\mathbf{x}_t | \{\mathbf{y}_\tau\}_{\tau \leq t-1})}{p(\mathbf{y}_t | \{\mathbf{y}_\tau\}_{\tau \leq t-1})}, \quad (8a)$$

where

$$p(\mathbf{y}_t | \{\mathbf{y}_\tau\}_{\tau \leq t-1}) = \int p(\mathbf{y}_t | \mathbf{x}_t) p(\mathbf{x}_t | \{\mathbf{y}_\tau\}_{\tau \leq t-1}) d\mathbf{x}_t, \quad (8b)$$

$$p(\mathbf{x}_t | \{\mathbf{y}_\tau\}_{\tau \leq t-1}) = \int p(\mathbf{x}_t | \mathbf{x}_{t-1}) p(\mathbf{x}_{t-1} | \{\mathbf{y}_\tau\}_{\tau \leq t-1}) d\mathbf{x}_{t-1}. \quad (8c)$$

Combining (8a)-(8b) with Equation (8c), referred to as the *Chapman-Kolmogorov* equation for SS models [30, Ch. 4.2], describes how the posterior distribution  $p(\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1})$  is recursively updated into  $p(\mathbf{x}_t | \mathbf{y}_{1:t})$ .

For the special case of a linear Gaussian SS model as in (1), all the considered variables are jointly Gaussian. Specifically, if  $\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1} \sim \mathcal{N}(\hat{\mathbf{x}}_{t-1|t-1}, \Sigma_{t-1|t-1})$ , then (8c) implies that  $\mathbf{x}_t | \mathbf{y}_{1:t-1} \sim \mathcal{N}(\hat{\mathbf{x}}_{t|t-1}, \Sigma_{t|t-1})$  (computed via (4)), which together with (8b) indicates that  $\mathbf{y}_t | \mathbf{y}_{1:t-1} \sim \mathcal{N}(\hat{\mathbf{y}}_{t|t-1}, \mathbf{S}_{t|t-1})$  (computed via (5)).

Combining these with (8a) reveals that  $\mathbf{x}_t | \mathbf{y}_{1:t} \sim \mathcal{N}(\hat{\mathbf{x}}_{t|t}, \Sigma_{t|t})$  with moments computed via (6). Accordingly,  $\hat{\mathbf{x}}_t$  computed by the KF is exactly the conditioned expectation of  $\mathbf{x}_t$  conditioned on  $\mathbf{y}_{1:t}$ , i.e., it is MSE optimal.

corresponding to the position and velocity of a vehicle can be represented as a linear transformation via mechanical relationships, e.g., a constant velocity model [32], such modeling of  $f_t(\cdot)$  is inherently a crude first-order approximation. Similarly, the relationship between the velocity of a vehicle and its sensed motor currents can be captured via an observation model of the form (3), the exact specification of the mapping  $h_t(\cdot)$  is likely to be elusive.

- C2 The purpose of the noise signals  $\mathbf{v}_t$  and  $\mathbf{w}_t$  is to (i) capture the inherent stochasticity in the state evolution and measurements, respectively; and (ii) model the discrepancy between the SS representation and the true system. Their actual distribution is thus often unknown, complex, and possibly intractable. Non-Gaussianity can have a notable effect on performance and reliability, especially since Kalman-type algorithms seek a linear filtering operation.
- C3 Even when the dynamics are faithfully characterized by a non-linear Gaussian SS model, Kalman-type algorithms are sub-optimal, with gaps from optimality largely depending on the nature and complexity of the underlying non-linearity.
- C4 Despite their relatively low complexity, non-linear variants of the KF, such as the EKF and UKF, induce some latency during filtering. This is due to the need to carry out, e.g., local linearization and matrix inversion, on each time instant (which the linear KF can do offline based on knowledge of the statistics).

These challenges motivate exploring data-driven approaches for tackling tasks associated with SS models, as detailed in the following section.

### III. COMBINING AI WITH KFS

Recent years have witnessed remarkable empirical success of deep learning, being the main enabler framework for AI, in various applications involving processing of time sequences. Data-driven DNNs were shown to be able to catch the subtleties of complex processes and replace the need to explicitly characterize the domain of interest. Therefore, an alternative strategy to implement state estimation while coping with C1-C4, namely, without requiring explicit and accurate knowledge of the SS representation, is to learn this task from data using deep learning.

#### A. Time Sequence Filtering with DNNs

A common strategy in ML is to utilize highly parameterized abstract models that are trained from data to find the parameterization that minimizes the empirical risk (with regularization introduced to prevent overfitting). Their mapping, denoted  $\mathcal{F}_\theta$ , is dictated by a set of parameters denoted  $\theta$ . In deep learning, the parametric model  $\mathcal{F}_\theta$  is a DNN, with  $\theta$  being the network parameters. Such highly-parametrized abstract models can effectively approximate any Borel measurable mapping, as follows from the universal approximation theorem [6, Ch. 6.4.1].

DNNs can learn various tasks from data without requiring any mathematical and statistical representation of the underlying dynamics and observations model. Accordingly, DNNs applied for tasks such as filtering and smoothing do not require the formulation of the dynamics as a SS representation. Despite this invariance, the resulting architecture of the DNN can still draw some inspiration from SS representation and KF-type processing. Consequently, we divide our presentation to *conventional DNNs*, which only account for the fact that the data being processed is a time sequence, and *SS/KF-inspired DNNs*, with both families being invariant of the underlying statistical modeling.

1) *Conventional DNNs*: Tasks associated with SS models, particularly filtering and smoothing, involve the processing of time sequences that exhibit temporal correlations. Accordingly, DNNs designed to process time sequences can conceptually be trained to carry out these tasks. These architectures, discussed next, involve some deviation from the basic DNN form, i.e., based on fully connected (FC) layers, which processes a fixed-size vector rather than a time sequence.

**RNNs**: A common DNN architecture tailored for time sequences is based on RNNs. Recurrent parametric models maintain an internal state vector, denoted  $s_t$ , representing the memory of the system [6, Ch. 10]. The parametric mapping is given by:

$$\hat{x}_t = \mathcal{F}_\theta(y_t, s_{t-1}), \quad (13)$$

where the internal state also evolves via a learned parametric mapping

$$s_t = \mathcal{G}_\theta(y_t, s_{t-1}). \quad (14)$$

The vanilla implementation of an RNN uses a hidden layer to map  $y_t$  and the latent  $s_{t-1}$  into an updated hidden variable  $s_t$ . Alternative RNN architectures, such as gated recurrent units (GRUs) and long short-term memories (LSTMs), employ several learned cells to update  $s_t$ . Then,  $s_t$  is used to generate the instantaneous output  $\hat{x}_t$  using another layer, as illustrated in Fig. 1(a).

**Attention**: A widely popular DNN architecture, which is the core of the transformer model, is based on attention mechanisms [33]. Such architectures were shown to be extremely successful in learning complex tasks in natural language processing, where the considered signals can be viewed as time sequences.

An attention mechanism jointly processes a query vector  $q$  and  $n_a$  key-value pairs  $\{k_i, \nu_i\}_{i=1}^{n_a}$ . In the common case of attention with scaled dot-product scoring, its output is  $\sum_{i=1}^{n_a} \text{softmax}(\text{const} \cdot q^T k_i) \nu_i$ . When applied for tasks such as filtering or smoothing, this mechanism is used as a form of *self-attention*. A self-attention head is an ML model that applies trained linear layers to map the input into the queries, keys, and values. A single head self-attention with parameters  $\theta = \{\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v\}$  applied for smoothing can be written as

$$\begin{aligned} \hat{x}_t &= \mathcal{F}_\theta(\{y_\tau\}_{\tau=1}^T) \\ &= \sum_{\tau=1}^T \text{softmax}\left(\left(\mathbf{W}^q y_t\right)^T \left(\mathbf{W}^k y_\tau\right)\right) \mathbf{W}^v y_\tau. \end{aligned} \quad (15)$$

This procedure is illustrated in Fig. 1(b).

Attention-based DNNs processing time sequences typically apply multiple mappings as in (15) in parallel, as a form of *multi-head attention*. As opposed to RNNs, attention mechanisms do not maintain an internal state vector that is sequentially updated. This makes attention-based DNNs more amenable to parallel training compared with RNNs. However, attention mechanisms as in (15) are invariant of the order of the processed signal samples, and are thus typically combined with additional pre-processing termed *positional embedding* that embed each sample while accounting for its position.

**CNNs**: Unlike RNNs and attention mechanisms,

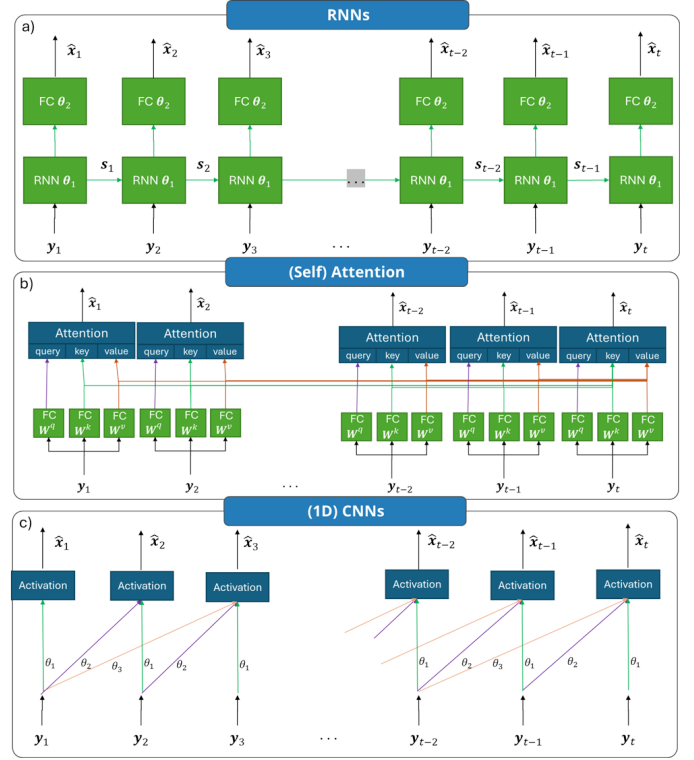


Fig. 1: Illustration of conventional DNN architectures for tasks related to filtering and smoothing, including (a) RNNs; (b) Attention; and (c) CNNs.

convolutional neural networks (CNNs) are DNN architectures that originate from image processing, and not time sequences. Specifically, CNNs are designed to learn the parameters of spatial kernels, aiming to exploit the locality and spatial stationarity of image data [6, Ch. 9]. Nonetheless, CNNs can also be applied for time sequence processing, and particularly for tasks such as filtering and smoothing. For once, the trainable kernel of CNNs can implement a learned finite impulse response filter (as a form of 1D CNN) and be applied to a time sequence, as illustrated in Fig. 1(c). Alternatively, one can apply CNNs to the time sequence by first converting the time sequence (or an observed window) into the form of an image via, e.g., short-time Fourier transform, and then use a CNN to process this representation.

2) *SS/KF-Inspired DNNs*: DNNs applied to process time sequences, are invariant of the underlying statistics governing the dynamics, and learn their operation purely from data. Nonetheless, one can still design DNNs for processing time sequences whose architecture is to some extent inspired by traditional model-based processing of such signals, particularly on SS representations and KF processing.

**SS-Inspired Architectures**: A form of DNN architecture that is inspired by SS representation models the filter (and not the dynamics as in (1)), as a *deterministic SS model*. These ML models, which we term SSM following [24], parameterize the mapping from  $y_t$  into the estimate  $\hat{x}_t$  using a latent state vector  $s_{t-1}$ . Particularly, the architecture operates using two

## Selective State-Space Models (SSMs)

An emerging SSM DNN, which is the core of the Mamba architecture and its variants [34], is the *selective SSM* [35]. Selective SSMs are ML models that process time sequences via (17). However, instead of using fixed linear mappings as in (17), here the matrices  $\mathbf{W}^y$ ,  $\mathbf{W}^s$ ,  $\mathbf{W}^x$  are obtained from learned mappings of the input  $\mathbf{y}_t$ , effectively implementing a time-varying and non-linear SSM.

Generally speaking, consider the discretization with sampling period  $\Delta$  of continuous-time deterministic SS representations. The resulting state evolution in (17a) is parameterized by

$$\mathbf{W}^y = \exp(\Delta \cdot \bar{\mathbf{W}}^y) \quad (16a)$$

$$\mathbf{W}^s = (\Delta \cdot \bar{\mathbf{W}}^y)^{-1} (\exp(\Delta \bar{\mathbf{W}}^y) - \mathbf{I}) \cdot \Delta \cdot \bar{\mathbf{W}}^s, \quad (16b)$$

where  $\bar{\mathbf{W}}^y$ ,  $\bar{\mathbf{W}}^s$  are the continuous-time state evolution matrices. Selective SSMs implement SSMs based on this representation, while using a dedicated layer with parameters  $\theta_1$  to map  $\mathbf{y}_t$  into  $\bar{\mathbf{W}}^s$ ,  $\Delta$ , and  $\mathbf{W}^x$ , where the former two are then substituted in (16) along with the learned  $\bar{\mathbf{W}}^y$ . The overall parameters are thus  $\theta = \{\theta_1, \bar{\mathbf{W}}^y\}$ . The resulting operation is illustrated in Fig. 2.




Fig. 2: Selective SSM illustration.

linear mappings, where first  $s_t$  is updated via

$$s_t = \mathbf{W}^y \mathbf{y}_t + \mathbf{W}^s s_{t-1}. \quad (17a)$$

Then, the predicted sequence is given by

$$\hat{x}_t = \mathbf{W}^x s_t. \quad (17b)$$

The parameters of the resulting ML architecture are  $\theta = \{\mathbf{W}^y, \mathbf{W}^s, \mathbf{W}^x\}$ .

The SSM operation in (17) can be viewed as a form of an RNN, with (17a) implementing (14), and (17b) implementing (13). While the restriction to linear time-invariant learned operators in (17) often results in limited effectiveness, an extension of SSMs, termed *Selective SSMs* (see box on page 7), was recently shown to yield efficient and high-performance architectures that are competitive to costly and complex transformers [34].

**KF-Inspired Architectures:** While SSMs parameterize a learned filter via a SS representation, one can also design DNNs whose architecture is inspired by filters suitable for tracking in SS models, e.g., KF-type algorithms. A representative example of such architectures is the recurrent Kalman network (RKN), proposed in [7].

The RKN is a DNN whose operation imitates the prediction-update steps of the KF and the EKF. However, while the model-based algorithms realize these computations based on the characterization of the dynamics as a SS model, the RKN parameterizes the predict and update stages as two interconnected DNNs. The predict DNN, whose parameters are  $\theta_1$ , replaces (4), and computes the a-priori state prediction via

$$\hat{x}_{t|t-1}, \Sigma_{t|t-1} = \mathcal{F}_{\theta_1}(\hat{x}_{t-1|t-1}, \Sigma_{t-1|t-1}). \quad (18a)$$

Similarly, the update DNN, parameterized by  $\theta_2$ , produces the posterior state estimate of (6) via

$$\hat{x}_{t|t}, \Sigma_{t|t} = \mathcal{G}_{\theta_2}(\hat{x}_{t|t-1}, \Sigma_{t|t-1}, \mathbf{y}_t). \quad (18b)$$

The resulting architecture can be combined with additional learned input and output processing, as proposed in [7].

3) *Pros and Cons of DNN-Based State Estimation:* The DNN architectures detailed so far can be trained to carry out state estimation tasks, i.e., map the observed time sequence into the latent state sequence. Specifically, provided data describing the task, e.g., labeled data set comprised of trajectories of observations and corresponding states, these DNNs learn their mapping from data without relying on any statistical modeling of the dynamics. This form of *discriminative learning*, i.e., leveraging data to learn to carry out a task end-to-end [26], excels where model-based methods struggle: it is not affected by inaccurate modeling (thus coping with C1-C2, and their abstractness allows DNNs to operate reliably in complex settings (C3). In terms of inference speed (C4), while DNNs involve lengthy and complex training, they often provide rapid inference (forward path), particularly when using relatively compact parameterization. This follows as trained DNNs are highly amenable to parallelization and acceleration of built-in hardware software accelerators, e.g., PyTorch.

Nonetheless, replacing KF-type algorithms with DNNs trained end-to-end gives rise to various shortcomings, particularly in losing some of the desired properties of model-based methods. For once, DNNs lack in adaptability (P2), as one cannot substitute time-varying parameters of SS models into their operation, and thus changes in the SS model may necessitate lengthy retraining. Moreover, DNNs are typically highly parameterized architectures viewed as black-boxes, that do not share the interpretability of model-based KF-type methods (P3), and are complex to train or even store on limited devices (P5). Furthermore, DNNs struggle in providing uncertainty (P4), for which there is typically no "ground truth" to learn from, and do not share the theoretical guarantees of KFs (P1).

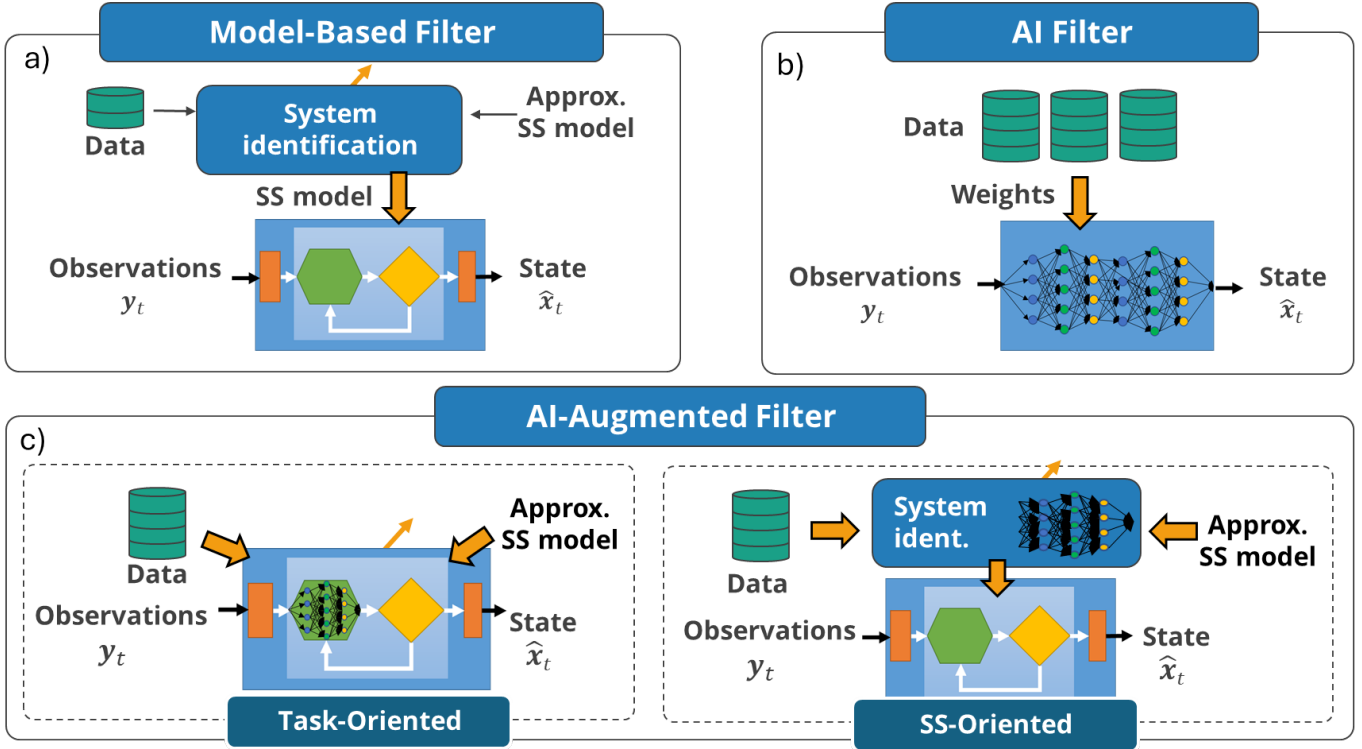


Fig. 3: Illustrative comparison between model-based Kalman-type filters (a); AI-based filters (b); and AI-augmented KF (c) divided into task-oriented and SS-oriented designs.

### B. AI-Augmented KFs

The DNN-based approaches discussed so far are highly data-driven, in the sense that they do not rely on any statistical characterization of the dynamics. Even SS or KF inspired architectures, such as the RKN whose operation generally follows the high-level stages of the KF, are ignorant of any SS modeling.

An alternative approach that aims to benefit from the best of both worlds is based on hybrid model-based/data-driven designs via model-based deep learning [11], [36]. This family of algorithms typically jointly leverage data along with some form of domain knowledge, i.e., partial knowledge of some components of the underlying SS model (which is often available to some degree as noted in C1). In the context of state estimation, such hybrid algorithms operate by *augmenting* KF-type algorithms with deep learning modules, rather than replacing them with DNNs.

As discussed above, the direct application of DNNs for state estimation reviewed so far is based on task-oriented end-to-end discriminative learning. In the same spirit, existing approaches for augmenting the operation of KFs with AI tools can be generally categorized following the ML paradigms of generative and discriminative learning [25], [26]:

- *SS-oriented* hybrid algorithms, that use data to learn the underlying statistical model as DNN-aided system identification (thus bearing similarity to generative learning).
- *Task-oriented* schemes, that directly learn to carry out the state estimation task (as a form of discriminative learning), while leveraging partial state knowledge and

principled KF stages as inductive bias.

This categorization, illustrated in Fig. 3, serves for our structured review of existing approaches in the subsequent sections.

## IV. AI-AUGMENTED KFs VIA TASK-ORIENTED LEARNING

The first family of AI-aided KFs converts Kalman-type state estimation into an ML architecture that is trainable end-to-end via DNN-augmentation. The key rationale is to leverage deep learning techniques to directly learn the state estimation as a form of discriminative learning. A common approach to designing such architectures is based on using an external DNN, operating alongside a classic state estimator, with recent architectures integrating deep learning modules into the internal processing of Kalman-type algorithms.

### A. External DNN Architectures

Utilizing external DNNs aims to enhance KF-type algorithms without altering their internal processing. This approach facilitates design, as one can separate the DNN components from the classic state estimator. Broadly speaking, the leading approaches to utilizing external DNNs employ them either sequentially, i.e., for pre-processing, or in parallel, e.g., as learned correction terms.

1) *Learned Pre-Processing:* A popular approach when dealing with complex measurement models, e.g., visual or multi-modal observations, builds on the ability of deep learning to extract meaningful features from complex data. Specifically, it uses a DNN pre-processor to map the observations into a latent space via a DNN [13], [27], [37].



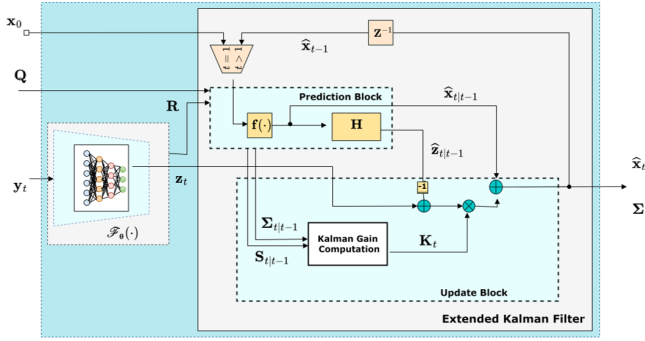


Fig. 4: EKF with DNN pre-processing illustration.

**Architecture:** Consider a SS model where the observation model  $h_t(\cdot)$  is complex and possibly intractable. In such cases, one can design a DNN with parameters  $\theta$  whose output is approximated as obeying a linear Gaussian observations model, i.e.,

$$z_t = \mathcal{F}_\theta(y_t) \approx \mathbf{H}x_t + w_t, \quad w_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \quad (19)$$

with the observation matrix  $\mathbf{H}$  being fixed a-priori. The latent  $z_t$ , combined with a known state evolution model, represent observations of a closed-form SS representation. Accordingly, the latent signal can be used as input to a model-based KF-type state estimator to track  $x_t$ , as illustrated in Fig. 4.

**Supervised Training:** In a supervised setting, one has access to a labeled data set comprised of multiple pairs of state and observation length  $T$  trajectories, e.g.,

$$\mathbb{D}_s = \left\{ \{x_t^{(i)}, y_t^{(i)}\}_{t=1}^T \right\}_{i=1}^N. \quad (20)$$

Given such labeled data,  $\theta$  can be trained by encouraging the DNN output to closely match  $\mathbf{H}x_t$ , i.e., by using a loss of the form

$$\mathcal{L}_{\mathbb{D}_s}(\theta) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left\| \mathcal{F}_\theta(y_t^{(i)}) - \mathbf{H}x_t^{(i)} \right\|_2. \quad (21)$$

When training via (21), the noise covariance  $\mathbf{R}$  can be estimated empirically from the validation data.

An alternative training approach finds  $\theta$  based on its downstream state estimation task, leveraging the differentiable operation of KFs [38]. Here, by letting  $\hat{x}_t(\mathcal{F}_\theta(y_t))$  be the state estimate obtained by a KF/EKF with observations  $z_t = \mathcal{F}_\theta(y_t)$ , a candidate loss measure is

$$\mathcal{L}_{\mathbb{D}_s}(\theta) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left\| \hat{x}_t(\mathcal{F}_\theta(y_t^{(i)})) - x_t^{(i)} \right\|_2. \quad (22)$$

The covariance matrix  $\mathbf{R}$  can be learned during training by being incorporated into the trainable parameters of the architecture.

**Discussion:** Using an external pre-processing DNN that is separated from the Kalman-type state estimation algorithm is a design approach geared mostly towards handling complex observation models. As state estimation is carried out by a model-based algorithm, it preserves most of its favorable properties: The state estimation operation is interpretable (P3),

and uncertainty measures are provided (P4). The excessive complexity lies in the incorporation of the pre-processing DNN and thus depends on its parameterization, as well as on the dimensions of  $z_t$ . For instance, when  $z_t$  is of a much lower dimension compared to  $y_t$ , the complexity and inference latency savings of applying a KF to  $z_t$  compared to using  $y_t$  as observations may surpass the added complexity and latency of the pre-processing DNN.

The adaptivity of KFs (P2) is not necessarily preserved. Specifically, when the DNN is fully separate from the state estimator (e.g., when training via (21)), then the state evolution parameters only affect the model-based algorithm, and thus one can still operate with time-varying  $f_t(\cdot)$  (assuming its variations are known). This is not necessarily the case when training via (22), as the latent features are learned to be ones most supporting state estimation based on the evolution model used during training [17]. Moreover, variations in the observation model typically necessitate re-training of  $\theta$ .

Using an external pre-processing DNN, while being simple and straightforward to combine with model-based state estimation, relies on Gaussianity and known distribution of the state evolution. The resulting latent SS model is often non-Gaussian, which can impact the tracking accuracy in the latent space, thus not handling C2. Moreover, this approach does not cope with complexities in the state evolution, thus not being geared towards handling inaccuracies (C1) and dominant nonlinearities (C3) in  $f_t(\cdot)$ .

2) *Learned Correction Terms:* In scenarios where one has full characterization of the dynamic system as a SS model, such that KF-type algorithms are applicable, yet the characterization is not fully accurate, external DNNs can further enhance performance by providing correction terms to internal computations of the model-based method [19]. This is illustrated using a representative example, based on the augmented Kalman smoother proposed in [19].

**Architecture:** Consider a linear Gaussian SS model, and focus on the smoothing task, i.e., recovering a sequence of  $T$  state variables  $\mathbf{x}_{1:T} = \{x_t\}_{t=1}^T$  from the entire observed sequence  $\mathbf{y}_{1:T} = \{y_t\}_{t=1}^T$ . As discussed when presenting the RTS algorithm, various algorithms exist for such tasks, one of which involves smoothing by seeking to maximize the log-likelihood function via gradient ascent optimization, i.e., by iterating over

$$\mathbf{x}_{1:T}^{(q+1)} = \mathbf{x}_{1:T}^{(q)} + \gamma \nabla_{\mathbf{x}_{1:T}^{(q)}} \log p(\mathbf{y}_{1:T}, \mathbf{x}_{1:T}^{(q)}), \quad (23)$$

where  $\gamma > 0$  is a step-size, for  $q = 0, 1, 2, \dots$  denoting the iteration index. The Markovian nature of the SS model indicates that the gradients in (23) can be computed via message passing, such that for the  $t$ 'th index

$$\nabla_{\mathbf{x}_t^{(q)}} \log p(\mathbf{y}_{1:T}, \mathbf{x}_{1:T}^{(q)}) = \mu_{\mathbf{x}_{t-1}^{(q)} \rightarrow \mathbf{x}_t} + \mu_{\mathbf{x}_{t+1}^{(q)} \rightarrow \mathbf{x}_t} + \mu_{\mathbf{y}_t^{(q)} \rightarrow \mathbf{x}_t}, \quad (24)$$

where the summands, referred to as messages, are given by

$$\mu_{\mathbf{x}_{t-1} \rightarrow \mathbf{x}_t}^{(q)} = -\mathbf{Q}^{-1} \left( \mathbf{x}_t^{(q)} - \mathbf{F} \mathbf{x}_{t-1}^{(q)} \right), \quad (25a)$$

$$\mu_{\mathbf{x}_{t+1} \rightarrow \mathbf{x}_t}^{(q)} = \mathbf{F}^T \mathbf{Q}^{-1} \left( \mathbf{x}_{t+1}^{(q)} - \mathbf{F} \mathbf{x}_t^{(q)} \right), \quad (25b)$$

$$\mu_{\mathbf{y}_t \rightarrow \mathbf{x}_t}^{(q)} = \mathbf{H}^T \mathbf{R}^{-1} \left( \mathbf{y}_t - \mathbf{H} \mathbf{x}_t^{(q)} \right). \quad (25c)$$

The iterative procedure in (23), is repeated until convergence, and the resulting  $\mathbf{x}_{1:T}^{(q)}$  is used as the estimate. It is worth noting that in (25), the messages are obtained by assuming a time-invariant version of the SS model in (2) [19].

An external DNN augmentation suggested in [19] aims at facilitating the operation of such smoother under approximated SS characterization by learning to map the messages in (25) into a correction term  $\epsilon_{1:T}^{(q+1)}$ , replacing the update rule (23) with

$$\mathbf{x}_{1:T}^{(q+1)} = \mathbf{x}_{1:T}^{(q)} + \gamma \left( \nabla_{\mathbf{x}_{1:T}^{(q)}} \log p \left( \mathbf{y}_{1:T}, \mathbf{x}_{1:T}^{(q)} \right) + \epsilon_{1:T}^{(q+1)} \right). \quad (26)$$

Particularly, based on the representation of the smoothing operation as messages exchanged over a Markovian factor graph whose nodes are the state variables, [19] proposed a graph neural network (GNN)-RNN architecture with parameters  $\theta$ . This architecture maps the messages in (25) into the correction term  $\epsilon_{1:T}^{(q+1)}$ , namely,

$$\epsilon_{1:T}^{(q+1)} = \mathcal{F}_\theta \left( \mathbf{y}_{1:T}, \left\{ \mu_{\mathbf{x}_{t-1} \rightarrow \mathbf{x}_t}^{(q)}, \mu_{\mathbf{x}_{t+1} \rightarrow \mathbf{x}_t}^{(q)}, \mu_{\mathbf{y}_t \rightarrow \mathbf{x}_t}^{(q)} \right\}_{t=1}^T \right). \quad (27)$$

**Supervised Training:** The external DNN correction mechanism is trained end-to-end, such that the state predicted by the corrected algorithm best matches the true state. Particularly, since smoothing here is based on an iterative algorithm (gradient ascent over the likelihood objective), its training is carried out via a form of deep unfolding [36], fixing the number of iterations to some  $Q$ . Then, accounting for the fact that the iterative steps in (26) should provide gradually refined state estimates, the loss function accounts for the contribution of the intermediate iterations with a monotonically increasing contribution. Particularly, the loss function used for training  $\theta$  is

$$\mathcal{L}_{\mathbb{D}_s}(\theta) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \sum_{q=1}^Q \frac{q}{Q} \left\| \hat{\mathbf{x}}_t^{(q)}(\mathbf{y}_{1:T}; \theta) - \mathbf{x}_t^{(i)} \right\|_2, \quad (28)$$

where  $\hat{\mathbf{x}}_t^{(q)}(\mathbf{y}_{1:T}; \theta)$  is the smoothed estimate of  $\mathbf{x}_t^{(i)}$  produced by the  $q$ 'th iteration, i.e., via (26), with parameters  $\theta$  and  $\mathbf{y}_{1:T}$ .

**Discussion:** Augmenting a KF-type algorithm via a learned correction term is useful in settings where one can still apply the model-based algorithm quite reliably (up to some possible errors that external DNN corrects). For instance, the augmented algorithm detailed above assumes the dynamic system can be approximated as a linear Gaussian SS model, such that model-based smoothing can be applied based on this formulation, yet it is not MSE optimal (as is the case without discrepancy in the SS model). Consequently, this approach is suitable for tackling C1, yet it is less valid for handling complex dynamics (C3), and only adds excessive latency to the model-based algorithm, thus not handling C4.

## B. Integrated DNN Architectures

Unlike external architectures, which employ DNNs separately from Kalman-type algorithms, integrated architectures replace intermediate computations with DNNs. Doing so converts a Kalman-type state estimation algorithm into a trainable ML model which follows the operation of the classic state estimation method as an inductive bias. The key design rationale is to augment computations that depend on missing domain knowledge with dedicated DNNs. Accordingly, existing designs vary based on the absent domain knowledge or, alternatively, on the augmented computation.

1) *Learned Kalman Gain:* A key part of Kalman-type algorithms is the derivation of the KG  $\mathbf{K}_t$ . In particular, its computation via (7) encapsulates the need to propagate the second-order moments of the state and observations. This, in turn, induces the requirement to have full knowledge of the underlying stochasticity (C2); leads to some of the core challenges in dealing with non-linear SS models (C3); and results in excessive latency which is associated with propagating these moments (C4). Accordingly, a candidate approach for state estimation in partially known SS models, which is the basis for the KalmanNet algorithm [16] and its variants [17], [39], [40], augments the KG computation with a DNN.

**Architecture:** Consider a dynamic system represented as a SS model in which one has only a (possibly approximated) model of the time-invariant state evolution function  $f(\cdot)$  and the observation function  $h(\cdot)$ . An EKF with a learned KG employs a DNN with parameters  $\theta$  to compute the KG for each time instant  $t$ , denoted  $\mathbf{K}_t(\theta)$ . Using this learned KG, an EKF is applied, predicting only the first-order moments via (11), and estimating the state as

$$\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t(\theta)(\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}). \quad (29)$$

An illustration of the architecture is depicted in Fig. 5.

As the resulting architecture does not explicitly track the second-order moments, that are needed for the KG, and thus the learned computation must do so implicitly. Accordingly, the DNN has to

- 1) Process input features that are informative of the noise signals.
- 2) Possess some internal memory capabilities.

To meet the first requirement, the input features processed by the DNN typically include (i) differences in the observations, e.g.,  $\Delta \mathbf{y}_t = \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}$ ; and (ii) differences in the estimated state, such as  $\Delta \hat{\mathbf{x}}_{t-1} = \hat{\mathbf{x}}_{t-1} - \hat{\mathbf{x}}_{t-1|t-2}$ . In order to provide internal memory capabilities, architectures based on RNNs [16], [40], or even transformers [41] can be employed. For instance, a simple FC-RNN-FC architecture was proposed in [16], as well as a more involved interconnection of RNNs, while [40] used two RNNs – one for tracking  $\Sigma_{t|t-1}$  and one for tracking  $\mathbf{S}_{t|t-1}^{-1}$  – that are combined into the KG via (7).

While the above architecture is formulated for filtering, it can also be extended to smoothing tasks. Particularly, when smoothing via the RTS algorithm [31] employs a KF for a forward pass, followed by an additional backward pass (9) that uses the backward gain matrix  $\bar{\mathbf{K}}_t$  given in (10). Consequently, the DNN augmented methodology above extends to smoothing

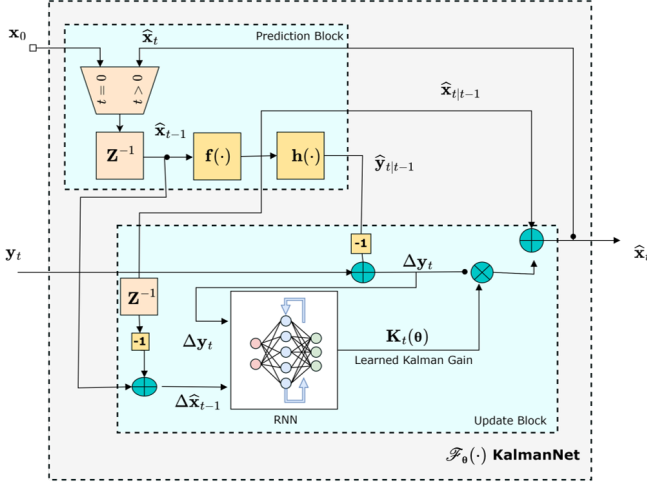


Fig. 5: EKF with learned KG illustration.

by introducing an additional KG DNN that learns to compute  $\hat{\mathbf{K}}_t$ . Moreover, the single forward-backward pass of the RTS algorithm is not necessarily MSE optimal in non-linear SS models, the learned RTS can be applied in multiple iterations, using the estimates produced at a given forward-backward pass be used as the input to the following pass [39]. Here, the fact that each pass is converted into a ML architecture allows to learn a different forward and backward gain DNNs for each pass, as a form of deep unfolding [36]

**Training:** KalmanNet and its variants use DNNs to compute the KG. While there is no ‘ground-truth’ KG when deviating from dynamics fully captured as linear Gaussian SS models, the overall augmented state estimation algorithm is trainable in a supervised manner. Particularly, given a labeled data set as in (20), a candidate loss measure is

$$\mathcal{L}_{\mathbb{D}_s}(\theta) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left\| \hat{\mathbf{x}}_t(\mathbf{y}_t^{(i)}; \theta) - \mathbf{x}_t^{(i)} \right\|_2, \quad (30)$$

where  $\hat{\mathbf{x}}_t(\mathbf{y}_t; \theta)$  is the state estimate obtained using the augmented EKF with parameters  $\theta$  and observations  $\mathbf{y}_t$ .

While KalmanNet is designed to be trained from labeled data via (30), in some settings it can also be trained without providing it with ground-truth state labels. This can be achieved via the following approaches:

- *Observation Prediction:* The fact that the DNN-augmented algorithm preserves the interpretable operation of Kalman-type state estimation can be leveraged for unsupervised learning, i.e., learning from a data set of the form

$$\mathbb{D}_u = \left\{ \left\{ \mathbf{y}_t^{(i)} \right\}_{t=1}^T \right\}_{i=1}^N. \quad (31)$$

Specifically, as Kalman-type algorithms with time-invariant SS models internally predict the next observation as  $\hat{\mathbf{y}}_{t+1|t} = h(f(\hat{\mathbf{x}}_t))$  in (5), a possible unsupervised loss is [21]

$$\mathcal{L}_{\mathbb{D}_u}(\theta) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=0}^{T-1} \left\| h(f(\hat{\mathbf{x}}_t(\mathbf{y}_t^{(i)}; \theta))) - \mathbf{y}_{t+1}^{(i)} \right\|_2. \quad (32)$$

- *Downstream Task:* Often, in practice, state estimation is carried out as part of an overall processing chain, with some downstream tasks. In various applications, one can evaluate an estimated state without requiring a ground truth value. When such evaluation is written (or approximated) as a function that is differentiable with respect to the estimated state, it can be used as a training loss. This approach was shown to enable unsupervised learning of KalmanNet when integrated in stochastic control systems, where it is combined with a linear quadratic regulator [42], as well as in financial pairs trading, where the state tracked are financial features used for trading policies [23].

**Discussion:** The design of DNN-aided Kalman-type algorithms by augmenting the KG computation is particularly suitable for tackling identified challenges in C1-C4. Unlike purely end-to-end DNNs designed for generic processing of time-sequences, augmenting the KG allows leveraging domain knowledge of the state evolution and observation models, as these are utilized in the prediction step. However, as the following processing of these predictions utilizes a DNN that is trained based on the accuracy of the overall algorithm, mismatches, and approximation errors are learned to be corrected, thus fully tackling C1. As the DNN augmentation bypasses the need to track second-order moments, the resulting algorithm does not require knowledge of the distribution of the noises. In fact, the resulting filter is not linear (as the KG depends on the observations), allowing to learn non-linear state estimators that are suitable for non-Gaussian SS models (C2) and non-linear dynamics (C3). Finally, as the computation of the KG and the propagation of the second-order moments induces most of the latency in Kalman-type algorithms for non-linear SS models, replacing these computations with a compact DNN often leads to more rapid inference (C4) [39].

The fact that the resulting state estimator is converted into a trainable discriminative ML architecture makes it natural to be combined with DNN-based pre-processing (as in Fig. 4), e.g., for processing images or high dimensional data. Particularly, the learned state estimator can be trained jointly with the pre-processing stage, thus having it learn latent features that are most useful for processing with the learned filter, without requiring one to approximate the distribution of the latent observation model noise [17].

Unlike state estimation based on end-to-end DNNs, which also tackles C1-C4 (as discussed in the previous section), the usage of Kalman-type algorithms as an inductive bias allows to also preserve some of the desirable properties of model-based state estimators. In particular, the interpretable operation is preserved, in the sense that the internal features are associated with concrete meaning (P3), which can be exploited for, e.g., unsupervised learning. Moreover, while the error covariance is explicitly tracked, in some cases it can actually be recovered from the learned KG (see box entitled ‘*Uncertainty Extraction from Learned KG*’ on Page 12), thus providing P4 to some extent. In addition, the fact that only an internal computation is learned allows utilizing relatively compact DNNs, striking a balance between the excessive complexity of end-to-end DNNs and the relatively low one of model-based methods (P5).

## Uncertainty Extraction from Learned KG

Kalman-type algorithms use both the prior covariance  $\Sigma_{t|t-1}$  as well as the KG to compute the posterior error covariance  $\Sigma_t$ , which represents the estimation uncertainty. Particularly, by (6b) and (7), the EKF computes the error covariance as

$$\Sigma_t = (\mathbf{I} - \mathbf{K}_t \hat{\mathbf{H}}_t) \Sigma_{t|t-1}. \quad (33a)$$

Consequently, despite the fact that AI-aided filters with learned KG do not explicitly track the second-order moments, in some cases, these can still be recovered via (33a), as shown in [43]. In particular, the architecture provides the learned KG as  $\mathbf{K}_t(\theta)$ . Accordingly, when the matrix  $\tilde{\mathbf{H}}_t = (\hat{\mathbf{H}}_t^\top \hat{\mathbf{H}}_t)^{-1}$  exists, then, combining (5b) and (7), the prior covariance can be recovered via

$$\Sigma_{t|t-1} = (\mathbf{I} - \mathbf{K}_t(\theta) \tilde{\mathbf{H}}_t)^{-1} \mathbf{K}_t(\theta) \mathbf{R} \hat{\mathbf{H}}_t \tilde{\mathbf{H}}_t. \quad (33b)$$

This indicates that when one accurately computes  $\hat{\mathbf{H}}_t$  (via (12)) and has knowledge of the observations noise covariance  $\mathbf{R}$ , then the learned KG can be used to recover the error covariance via (33a) and (33b), as illustrated in Fig. 10. The extracted covariance can be encouraged to match the empirical estimation error in training by, e.g., adding an additional loss term that penalizes uncertainty extraction, or by imposing a Gaussian prior on the error and minimizing its likelihood, see [43].

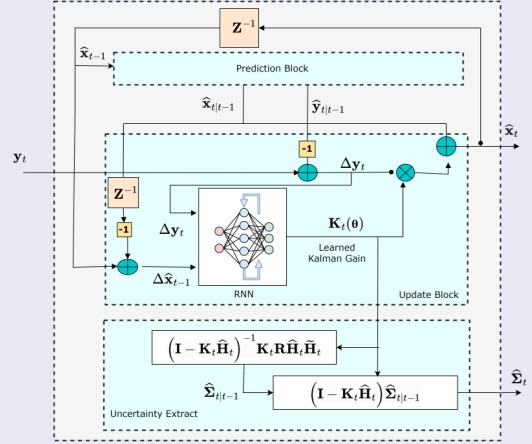


Fig. 6: EKF with learned KG and covariance extraction illustration.

The core limitations of designing AI-aided state estimator by augmenting the KG computations are associated with its adaptivity. The design is particularly geared towards time-invariant state evolution and observation functions. As the KG computation is coupled with these models, any deviation, even known ones, is expected to necessitate re-training. This can be alleviated under some forms of variations using hypernetworks [44], i.e., having an additional DNN that updates the KG DNN, while inducing some complexity increase. Moreover, the architecture is designed for supervised learning. The ability to train in an unsupervised manner via (32) relies on full knowledge of  $f(\cdot)$  and  $h(\cdot)$ , thus not being applicable in the same range of settings as its supervised counterpart. The same also holds for uncertainty extraction via (33), which requires some additional domain knowledge compared to that needed for merely tracking the state via (29).

The introduced learned KG algorithms calculate the complete gain by a DNN. However, if a Kalman-type filter, such as the unscented Kalman filter or divided-difference filter, depends on a *user-defined* scaling parameter, a DNN can be used for the parameter prediction (instead of the complete gain calculation) [45], [46]. As a consequence, such DNN-reasoned scaling parameter affects all elements of the KG and is able, up to a certain extent, to compensate for model discrepancies or linearisation errors. This DNN-augmented filter inherently provides the estimate error covariance matrix, but under the assumption of the known SS model (3).

2) *Learned State Estimation*: Another class of integrated DNN architectures includes methods that learn the task of state estimation in a data-driven manner using the known ob-

servation model without any knowledge of the state evolution model. The state evolution model in (1a) may not be linear or require  $\mathbf{v}_t$  to be Gaussian noise. A candidate approach in this category is the data-driven nonlinear state estimation (DANSE) method that provides a closed-form posterior of the underlying state using linear state measurements under Gaussian noise [20]. The noisy measurements follow the linear observation model similar to (2b), namely,

$$\mathbf{y}_t = \mathbf{H}_t \mathbf{x}_t + \mathbf{w}_t, \quad (34)$$

where the measurement noise is  $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$  with known covariance  $\mathbf{R}$ . The matrix  $\mathbf{H}_t$  is assumed to be full column rank and known  $\forall t$ . Note that,  $p(\mathbf{y}_t | \mathbf{x}_t) = \mathcal{N}(\mathbf{H}_t \mathbf{x}_t, \mathbf{R})$ .

**Architecture:** The core of the DANSE method relies on the sequential modeling capability of RNNs [6, Ch. 10]. DANSE consists of an RNN with parameters  $\theta$  that parameterizes a Gaussian prior on  $\mathbf{x}_t$  given  $\mathbf{y}_{1:t-1}$  at a given time  $t$ . Concretely, the prior distribution  $p(\mathbf{x}_t | \mathbf{y}_{1:t-1}; \theta)$  is

$$p(\mathbf{x}_t | \mathbf{y}_{1:t-1}; \theta) = \mathcal{N}(\hat{\mathbf{x}}_{t|t-1}(\theta), \Sigma_{t|t-1}(\theta))$$

$$\text{s.t. } \{\hat{\mathbf{x}}_{t|t-1}(\theta), \Sigma_{t|t-1}(\theta)\} = \mathcal{F}_\theta(\mathbf{y}_{t-1}), \quad (35)$$

where  $\mathcal{F}_\theta(\mathbf{y}_{t-1})$  refers to the RNN that recursively processes the sequence of past observations  $\mathbf{y}_{1:t-1}$  as described in (13), (14). Also,  $\hat{\mathbf{x}}_{t|t-1}(\theta)$  and  $\Sigma_{t|t-1}(\theta)$  denote the mean vector and the covariance matrix respectively of the RNN-based Gaussian prior at time  $t$ . The covariance matrix  $\Sigma_{t|t-1}(\theta)$  can be designed to be full or diagonal. More accurately, the hidden state of the RNN in (35) at time  $t$  is non-linearly transformed using feed-forward networks with appropriate

## DANSE Architecture

The DANSE method utilizes an RNN for recursively processing the input sequence  $\mathbf{y}_{1:t-1}$  [20]. An example of this sequential processing at time  $t$  is shown in Fig. 7. One can use prominent RNNs for sequential processing, such as GRUs or LSTMs. In DANSE, GRU was used due to simplicity and good expressive power compared to vanilla RNNs. Specifically, for modeling the mean vector  $\hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta})$  and diagonal covariance matrix  $\boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta})$  of the parameterized Gaussian prior in (35), the hidden state of the GRU was nonlinearly transformed using feed-forward networks with suitable activation functions. In Fig. 7, the ‘Accumulator’ block is shown for ease of understanding; in practice, the block is embedded in the recursive operation of the RNN.

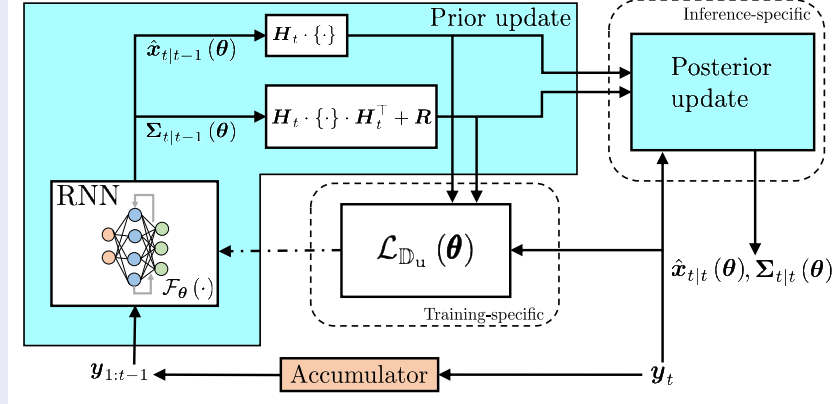


Fig. 7: A detailed schematic diagram of the DANSE method at time  $t$ , highlighting both training-specific and inference-specific blocks [20]. The dash-dotted line represents the gradient flow to  $\mathcal{F}_\theta$  during training, solid lines indicate information flow during training/inference.

activation functions to obtain  $\hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta})$  and  $\boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta})$  [20]. In [20], DANSE uses GRU for implementation purposes owing to its simplicity and modeling capability.

Using the fact that the measurement system in (34) is linear and also Gaussian, we have  $p(\mathbf{y}_t|\mathbf{x}_t) = \mathcal{N}(\mathbf{H}_t\mathbf{x}_t, \mathbf{R})$ . This allows the use of the Chapman-Kolmogorov equation (8a) to obtain  $p(\mathbf{x}_t|\mathbf{y}_{1:t}; \boldsymbol{\theta})$  in closed-form [20]. One can show that the posterior distribution  $p(\mathbf{x}_t|\mathbf{y}_{1:t}; \boldsymbol{\theta})$  is also Gaussian with

$$\begin{aligned} p(\mathbf{x}_t|\mathbf{y}_{1:t}; \boldsymbol{\theta}) &= \mathcal{N}(\hat{\mathbf{x}}_{t|t}(\boldsymbol{\theta}), \boldsymbol{\Sigma}_{t|t}(\boldsymbol{\theta})), \\ \hat{\mathbf{x}}_{t|t}(\boldsymbol{\theta}) &= \hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta}) + \mathbf{K}'_t(\boldsymbol{\theta}) \Delta \mathbf{y}_t(\boldsymbol{\theta}), \\ \boldsymbol{\Sigma}_{t|t}(\boldsymbol{\theta}) &= \boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta}) \\ &\quad - \mathbf{K}'_t(\boldsymbol{\theta}) \mathbf{S}_{t|t-1}(\boldsymbol{\theta}) (\mathbf{K}'_t(\boldsymbol{\theta}))^\top. \end{aligned} \quad (36)$$

Notably, (36) resembles the Kalman update equations in (6), (7) where

$$\begin{aligned} \mathbf{K}'_t(\boldsymbol{\theta}) &\triangleq \boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta}) \mathbf{H}_t^\top \mathbf{S}_{t|t-1}^{-1}(\boldsymbol{\theta}), \\ \mathbf{S}_{t|t-1}(\boldsymbol{\theta}) &\triangleq \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta}) \mathbf{H}_t^\top + \mathbf{R}, \\ \Delta \mathbf{y}_t(\boldsymbol{\theta}) &\triangleq \mathbf{y}_t - \mathbf{H}_t \hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta}). \end{aligned} \quad (37)$$

Note that  $\mathbf{K}'_t(\boldsymbol{\theta})$  is conceptually similar as the traditional Kalman gain term in (7), but computed in a data-driven manner, and that the computation approach is different from the one illustrated in Fig. 5 due to (35).

In the backdrop of KF, a crucial aspect of DANSE is that there is no Gaussian propagation of the posterior in (36) to the prior because DANSE does not use any state evolution model like (1a). The parameters  $\hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta})$ ,  $\boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta})$  of the prior are obtained directly from the RNN, which does not use any explicit state evolution model as in (1a), or require any knowledge regarding the same, e.g. first-order Markovian, Gaussian process noise, etc. A simplified schematic of DANSE

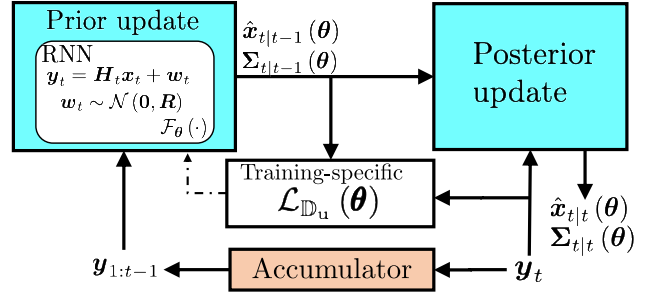


Fig. 8: Simplified schematic of DANSE [20].

is shown in Fig. 8 with a detailed schematic and details regarding architectural choices present in the box ‘DANSE Architecture’ on Page 13.

**Training:** The parameters  $\boldsymbol{\theta}$  in DANSE are learned in an unsupervised manner using a training dataset consisting of only noisy measurement trajectories  $\mathbb{D}_u = \{\mathbf{y}_{1:T}^{(i)}\}_{i=1}^N$ , where  $N$  is the number of training samples, where every  $i$ 'th sample is assumed to have the same trajectory length  $T$ . The learning mechanism is based on maximizing the likelihood of the dataset  $\mathbb{D}$ , where one computes the joint likelihood of a full sequence  $\mathbf{y}_{1:T}$ . In order to achieve this, one first calculates the conditional marginal distribution  $p(\mathbf{y}_t|\mathbf{y}_{1:t-1}; \boldsymbol{\theta})$  using the Chapman-Kolmogorov equation (8b) as follows

$$\begin{aligned} p(\mathbf{y}_t|\mathbf{y}_{1:t-1}; \boldsymbol{\theta}) &= \int p(\mathbf{y}_t|\mathbf{x}_t)p(\mathbf{x}_t|\mathbf{y}_{1:t-1}; \boldsymbol{\theta}) d\mathbf{x}_t \\ &= \mathcal{N}(\mathbf{H}_t \hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta}), \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1}(\boldsymbol{\theta}) \mathbf{H}_t^\top + \mathbf{R}) \\ &= \mathcal{N}(\mathbf{H}_t \hat{\mathbf{x}}_{t|t-1}(\boldsymbol{\theta}), \mathbf{S}_{t|t-1}(\boldsymbol{\theta})), \end{aligned} \quad (38)$$

where we use the notation for  $\mathbf{S}_{t|t-1}(\boldsymbol{\theta})$  in (37). Then for the

sequence  $\mathbf{y}_{1:T}$ , one can calculate the joint likelihood as

$$p(\mathbf{y}_{1:T}; \boldsymbol{\theta}) = \prod_{t=1}^T p(\mathbf{y}_t | \mathbf{y}_{1:t-1}; \boldsymbol{\theta}). \quad (39)$$

Thus, using  $\mathbb{D}_u$ , (38) and (39), the optimization problem can be formulated as maximization of the logarithm of the joint log-likelihood of  $\mathbb{D}_u$  as follows

$$\max_{\boldsymbol{\theta}} \log \prod_{i=1}^N \prod_{t=1}^T p(\mathbf{y}_t^{(i)} | \mathbf{y}_{1:t-1}^{(i)}; \boldsymbol{\theta}) = \min_{\boldsymbol{\theta}} \mathcal{L}_{\mathbb{D}_u}(\boldsymbol{\theta}), \quad (40)$$

where the loss term  $\mathcal{L}_{\mathbb{D}_u}(\boldsymbol{\theta})$  is obtained using (38) as follows

$$\begin{aligned} \mathcal{L}_{\mathbb{D}_u}(\boldsymbol{\theta}) = & \sum_{i=1}^N \sum_{t=1}^T \left\{ \frac{n}{2} \log 2\pi + \frac{1}{2} \log \det \left( \mathbf{S}_{t|t-1}^{(i)}(\boldsymbol{\theta}) \right) \right. \\ & \left. + \frac{1}{2} \left\| \mathbf{y}_t - \mathbf{H}_t \hat{\mathbf{x}}_{t|t-1}^{(i)}(\boldsymbol{\theta}) \right\|_{\left( \mathbf{S}_{t|t-1}^{(i)}(\boldsymbol{\theta}) \right)^{-1}} \right\}. \quad (41) \end{aligned}$$

In the above  $\{\hat{\mathbf{x}}_{t|t-1}^{(i)}(\boldsymbol{\theta}), \boldsymbol{\Sigma}_{t|t-1}^{(i)}(\boldsymbol{\theta})\}$  denote the parameters of the Gaussian prior (35) obtained using the  $i$ 'th sample  $\mathbf{y}_{1:t-1}^{(i)}$  as input at time  $t$ . The parameters  $\boldsymbol{\theta}$  are thus learned by minimizing  $\mathcal{L}_{\mathbb{D}_u}(\boldsymbol{\theta})$  with respect to  $\boldsymbol{\theta}$  in (40). Practically, this is achieved by using mini-batch stochastic gradient descent as explained in [20, Sec. II-C].

**Discussion:** The DANSE approach can be used to address the challenges associated with model-based KF-type algorithms explained in C1-C4. DANSE can be used to partly address C1 as it does not require knowledge of the state evolution model (1a) or the involved process noise in (1a). However, it requires knowledge of the linear observation model shown in (34), where  $\mathbf{H}_t$  should be full-rank and noise covariance  $\mathbf{R}$  should be known *a priori*. The parameterization of the Gaussian prior in (35) ensures that DANSE can capture long-term temporal dependency in the state model so that one is not constrained to state estimation for Markovian SS models. This is also mentioned by the authors in [20] where they introduce DANSE in a ‘Model-free process’ setting. In [20], the authors define a ‘Model-free process’ as a process where the knowledge of the state evolution model is absent. This also partly tackles C2 since there is no requirement on the distribution of  $\mathbf{v}_t$  in (1a) to be Gaussian. At the same time, it is required that  $\mathbf{w}_t$  is Gaussian with a fixed covariance  $\mathbf{R}$  as shown in (34).

The use of recurrent neural networks and the parameterization of the Gaussian prior also ensures that one can learn a nonlinear state estimation method for nonlinear SS models similar to that in the case of KalmanNet. This in turn helps to tackle the challenge described in C3. The training of DANSE, as explained in the previous paragraph, is unsupervised and offline as it requires access to a dataset  $\mathbb{D}_u$  consisting of noisy measurements. Once the training is completed, the inference step in DANSE is causal and separated from the offline training step. This ensures that one has a rapid inference at test time. This helps address C4, as DANSE doesn't require any linearization of the state evolution model as in EKF or sampling sigma points or particles in UKF or PF, respectively.

The DANSE method is also partly interpretable as the posterior update of DANSE as shown in (36) is tractable and bears similarity to the KF posterior update in (6). This is ensured by using the linear observation model in (34) and the choice of the Gaussian prior for  $\mathbf{x}_t$  in (35). This ensures tractable posterior updates, resonating with the feature P3 regarding the interpretability of conventional model-based approaches using first and second-order moments. Furthermore, the advantage P4 regarding providing uncertainty estimates is also present in DANSE as we have both prior and posterior covariance estimates in (35), (36). As mentioned earlier, the key difference compared to model-based approaches is the absence of propagating the posterior moments to the prior at the next time point.

Lastly, it is worth noting that DANSE cannot immediately adapt to underlying changes in the SS model and would require re-training. This also applies to changes in the observation model as it requires complete knowledge of the same. Hence DANSE does not have the advantage P2, which is inherent in model-based approaches and AI-aided approaches that utilize the knowledge of the state evolution model or additional hypernetworks as mentioned earlier in section IV-B1.

## V. AI-AUGMENTED KFS VIA SS-ORIENTED LEARNING

The second family of AI-aided KFs uses data to learn, refine, or augment the underlying statistical model using deep learning tools. Then, the data-augmented SS can be directly used in Kalman-type state estimation. Instead of learning the state estimation or the specific parameters in the Kalman filter, e.g., KG, the key rationale of SS-oriented learning is to exploit the data to obtain a more accurate model. At the same time, SS-oriented learning preserves explainability of the state and brings associated benefits of the statistical estimation<sup>1</sup> such as the inherent calculation of the covariance matrix of the estimate error [15].

**Architecture:** The SS-oriented AI-augmented KFs focus on data-augmented modeling of the state equation, while the measurement (or observation) model is assumed to be known. The motivation behind this formulation is twofold: (i) the fact that the sensors can often be well modeled on the basis of first principles but a state dynamics model is typically approximate and widely depends on a user decision (for example, object kinematic can be modeled by nearly constant velocity/acceleration model or a Singer model) [15]; (ii) the need to estimate state dynamics that carry physical meaning. The concept of data-augmented modeling thus resides in the definition of four models of state dynamics;

- *True model (TM)*, or a data generator, is a complex and context-dependent model that cannot be expressed in a finite-dimensional form. The TM is thus more of a theoretical construct.
- *Physics-based model (PBM)* is defined by (3a) and can be understood as the ‘‘best achievable’’ model of the given complexity found from the first principles.

<sup>1</sup>Such augmented model can directly be used also for design a fault-detection or control algorithm.

- *Data-driven model (DDM)* is solely identified from data. The use of ML in system identification has been studied for several decades [47], [48] and the DDM can be written as

$$\mathbf{x}_{t+1} = g(\mathbf{x}_t; \boldsymbol{\theta}^{\text{DNN}}) + \mathbf{v}_t^{\text{DDM}}, \quad (42)$$

where  $\mathbf{v}_t^{\text{DDM}}$  is a state noise,  $g(\cdot)$  is a vector-valued function, possibly a DNN, which is parametrized by the vector  $\boldsymbol{\theta}^{\text{DNN}}$  designed to minimize the discrepancy between outputs of (42) and the TM.

- The current trend is to consider models that are not purely trained from data, as in DDMs, but that make use of the available information that PBMs offer. These hybrid models are discussed in more detail in the remainder of the section.

The approaches to data-augmented modeling of the state equation can be classified in terms of the SS components that the DNN characterizes or is embedded in. In the general SS structures (see Fig. 9a), the whole model of state dynamics is identified from data and replaced by an DNN<sup>2</sup> (42). In [49], FC DNNs were used to represent the function  $f_t(\cdot)$ , and in [50], the authors use predictive partial autoencoders to find the state and its mapping to the measurement. In [51], FC DNNs trained by the expectation-maximization algorithm were used to represent both functions while [52] utilized LSTM DNNs for this purpose. A comparison of the performance of several general SS structures was presented in [14], considering LSTM, GRU, and variational autoencoders (VAEs). The authors assumed Gaussian likelihood and combined the RNN and VAE into a deep state space model. Stochastic RNNs were introduced for system identification in [53] to account for stochastic properties of the SS model. A similar approach is followed physics-informed neural network (PINN) [54] modeling the state equation by a DNN which training is constrained by the physics describing the PBM (see Fig. 9b).

Another approach to model state equation is to consider a fixed structure such as a linear parameter-varying model (2a) and use the DNNs to represent the matrices [55] (see Fig. 9c). Such an approach has also been adopted in [38], where the covariance matrices  $\mathbf{Q}$  and  $\mathbf{R}$  were provided by a DNN.

A similar idea was elaborated in [56] where an incremental SS structure was proposed to separate  $f_t(\cdot)$  into linear and nonlinear parts and use a DNN to represent the nonlinear part (see Fig. 9d). In [57], a hybrid model (see Fig. 9e) was proposed, which combines the PBM for some state elements and the DNN for the rest. It leverages the fact that the dynamic of some states is precisely known (e.g., the position derivative is velocity).

The data-augmented modeling represented by the *data-augmented physics-based model (APBM)* combines the physics- and data-driven components into a single model, which reads

$$\mathbf{x}_{t+1} = g(f_t(\mathbf{x}_t), \mathbf{x}_t, \mathbf{d}_t; \boldsymbol{\theta}_t^{\text{APBM}}) + \mathbf{v}_t^{\text{APBM}}, \quad (43)$$

<sup>2</sup>A structure similar to the general one can be considered for the continuous-time models, where the DNN represents the function in the differential state equation, and then the Euler discretization is used to obtain a discrete-time model used for the estimation.

where  $\mathbf{v}_t^{\text{APBM}}$  is a state noise,  $\mathbf{d}_t$  represents additional data<sup>3</sup>, the function  $g(\cdot)$  is aware of the PBM part, and the vector  $\boldsymbol{\theta}_t^{\text{APBM}}$  is designed to minimize the discrepancy between outputs of (43) and the TM. The APBM compensates the PBM structure and parameter mismodelling using information extracted from available data. Consequently, the APBM preserves the physical meaning of the state components and exploits actual system behavior dependencies, which were ignored in the PBM design. An example of this modeling versatility was shown in [58], where APBMs were used to filter a high-order Markov process without the need for order selection. One important characteristic of the APBM formulation (43) is the flexibility to cope with the non-stationarity of the model over time or space ( $\boldsymbol{\theta}_{t-\tau}^{\text{APBM}} \neq \boldsymbol{\theta}_t^{\text{APBM}}$ ), which requires adaptive estimation strategies [59].

The general APBM structure (43) can be simplified into an additive form with explicitly controlled contribution of the data-based component (see box entitled ‘*APBM with Controlled Additive Structure*’ on Page 17). Bounding the contribution of the data-driven component of the model is an essential feature of the APBM, which prevents the data-based component from overruling the PBM component contribution and, thus, preserving APBM explainability.

The APBM state dynamics model (43) together with the observation model (any of (1b), (2b), (3b)) can directly be used for joint state  $\mathbf{x}_t$  and parameter  $\boldsymbol{\theta}_t^{\text{APBM}}$  estimation by a regular state estimator such as the EKF or the UKF [30].

**Training APBMs:** Training APBMs can be performed under different paradigms depending on data availability, architecture, and assumptions regarding the stationarity of the system’s dynamics. Although the learning strategy can be supervised when  $\mathbb{D}_s$  is available, we will focus on the more common problem in the Bayesian filtering literature, where only noisy observations,  $\mathbb{D}_u$ , are accessible, setting up an unsupervised learning scenario. In this context, parameter estimation can be achieved by (i) obtaining and maximizing the marginal posterior  $p(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ , often using the energy function ( $\varphi_T(\boldsymbol{\theta}) = -\log p(\mathbf{y}_{1:T}|\boldsymbol{\theta}) - \log p(\boldsymbol{\theta})$ ) [30]; (ii) the joint posterior  $p(\mathbf{x}_t, \boldsymbol{\theta}|\mathbf{y}_{1:T})$  [59]; or (iii) obtaining a point estimate  $\hat{\boldsymbol{\theta}}$  through deterministic optimization strategies often aiming at maximizing the variational lower bound of the log-likelihood  $p(\mathbf{y}_{1:T}; \boldsymbol{\theta})$  [60].

In [15], [59], the authors opted for a state-augmentation approach aiming at obtaining the joint posterior distribution  $p(\mathbf{x}_t, \boldsymbol{\theta}_t|\mathbf{y}_{1:T})$  through Bayesian filtering recursion. For such, system states are augmented with the APBM parameters:

$$\begin{bmatrix} \mathbf{x}_{t+1} \\ \boldsymbol{\theta}_{t+1}^{\text{APBM}} \end{bmatrix} = \begin{bmatrix} g(f_t(\mathbf{x}_t), \mathbf{x}_t, \mathbf{d}_t; \boldsymbol{\theta}_t^{\text{APBM}}) \\ \boldsymbol{\theta}_t^{\text{APBM}} \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^{\mathbf{x}} \\ \mathbf{v}_t^{\boldsymbol{\theta}} \end{bmatrix} \quad (44)$$

where a near-constant state transition process is introduced for  $\boldsymbol{\theta}^{\text{APBM}}$ , allowing one to cast the APBM learning as a filtering problem.  $\mathbf{v}_t^{\boldsymbol{\theta}}$  is a “small” noise and is introduced to avoid numerical issues [30]. Note that such noise can also allow  $\boldsymbol{\theta}_t^{\text{APBM}}$  to drift over time and eventually evolve if the model is time-varying.

<sup>3</sup>These denote data related to the system available to the user but not used in the PBM as additional inputs to avoid overly complex models (e.g., ionospheric models, weather forecasts).

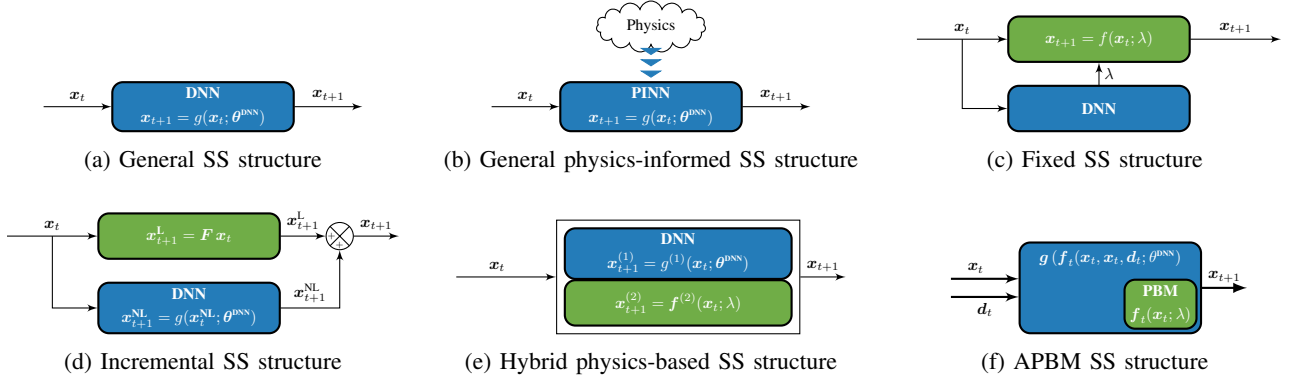


Fig. 9: Structures for SS-oriented learning.

Another important component regarding the APBM learning process is the need for a control mechanism to prevent the DNN component from overpowering the PBM. The prevention can be achieved by imposing appropriate constraints over  $\theta^{\text{APBM}}$ . A Bayesian filtering based approach is to introduce pseudo-observation equations into the observation model [61]. Let  $\bar{\theta}$  be a vector such that  $g(f_t(x_t), x_t, d_t; \theta_t^{\text{APBM}} = \bar{\theta}) = f_t(x_t)$ , then the observation model in (3b) can be augmented as:

$$\begin{bmatrix} \mathbf{y}_t \\ \bar{\theta} \end{bmatrix} = \begin{bmatrix} h_t(x_t) \\ \theta_t^{\text{APBM}} \end{bmatrix} + \begin{bmatrix} \mathbf{w}_t \\ \mathbf{w}_t^\theta \end{bmatrix} \quad (45)$$

where the bottom equation acts as a constraint, forcing  $\theta_t^{\text{APBM}}$  to be in the vicinity of  $\bar{\theta}$ ; and the distribution of  $\mathbf{w}_t^\theta$  governs the trade-off between regularization and data fit. In the case where the noise terms in the SS model defined in Equations (44) and (45) are Gaussian,  $\text{cov}\{\omega_t\} = \mathbf{R}$  and  $\text{cov}\{\omega_t^\theta\} = \eta \mathbf{I}$ , the Bayesian filtering solution can be cast as a sequence of minimization problems for every time step  $t$ :

$$\begin{aligned} (\hat{x}_t, \hat{\theta}_t^{\text{APBM}}) &= \arg \min_{(x_t, \theta_t)} \|\mathbf{y}_t - h_t(x_t)\|_{\mathbf{R}}^2 + \eta \|\bar{\theta} - \theta_t\|^2 \\ &+ \|\mathbf{x}_t - g(f_t(x_{t-1}), x_{t-1}, d_{t-1}; \hat{\theta}_{t-1}^{\text{APBM}})\|_{[\hat{\mathbf{P}}_{t|t-1}^x]^{-1}}^2 \\ &+ \|\theta_t - \hat{\theta}_{t-1}^{\text{APBM}}\|_{[\hat{\mathbf{P}}_{t|t-1}^\theta]^{-1}}^2 \end{aligned} \quad (46)$$

where  $\eta \geq 0$  controls the strictness of the regularization added by the pseudo-observation equation in (45).

*Offline vs Online Training of APBMs:* One interesting aspect of leveraging the Bayesian filtering approach for joint parameter and state estimation lies in its equivalence to a second-order Newton method [62], leading to fast convergence and making it suitable for both online and offline training. The training methodology seamlessly allows for either offline, online, or both (offline-online) training of the system dynamics. In the offline scenario, the filtering with the augmented model can be performed over the multiple sequences (and multiple epochs) in  $\mathbb{D}_u$  if system states,  $\mathbf{x}_0$ , are properly initialized for every data sequence in  $\mathbb{D}_u$  [63]. Once training is completed, a standard Bayesian filtering approach can be used to update only the system states,  $\mathbf{x}_t$ , over time while

keeping the APBM parameters fixed. In the online setting, the training approach becomes a conventional filtering problem that continuously adapts both states and model parameters over time, starting from some initial condition  $\hat{x}_0, \hat{\theta}_0^{\text{APBM}}$ . Again, the fast convergence of the Bayesian filter allows for quick reaction of the data-driven component over time, correcting the PBM to adapt to new conditions and improving the state estimation performance. This feature is extremely important when dealing with non-stationary dynamics that continuously change over time. Both strategies can be combined, where the offline training solution is used as the initial condition for the online procedure, which, in turn, keeps updating both states and parameters during the test.

**Discussion:** The SS-oriented AI-augmented KF design relies on augmenting the physics-based component, namely the “deterministic” part of the state equation with a data component. As seen in the box entitled ‘APBM with Controlled Additive Structure’, a solution might lead to an additive APBM state equation (47). The data component is designed to extract a *time-correlated* component of the discrepancy between the pure PBM (3a) and the TM, that is usually embraced by the overbounded state noise  $\mathbf{v}_t$ . As a consequence, the APBM state noise  $\mathbf{v}_t^{\text{APBM}}$  has different statistical properties from the PBM  $\mathbf{v}_t$  and for optimal performance of the KF, the noise properties have to be identified. In [64], the noise properties identification of the APBM was discussed and illustrated using correlation and maximum likelihood methods. Utilizing the identified state noise covariance matrix in the KF led to significant improvement of the estimate consistency.

Following the terminology used in system identification [65], the PBM augmentation with DNN belongs into *block-oriented “slate-gray” models*. The idea of block-oriented models is “to build up structures from simple building blocks” [65], which allows physical insight and data-oriented flexible complement. Note that besides the DNN, the PBM can be augmented with the nonlinear autoregressive moving-average model as in [66], where the data component is used for drift compensation. The concept of the APBM can also be found in the area of deterministic models with a completely measured state. In this area, the “deterministic” version of the APBM



## APBM with Controlled Additive Structure

The general APBM in (43) can be particularized as an additive augmentation of the PBM [15]

$$\mathbf{x}_{t+1} = \phi_0 f_t(\mathbf{x}_t) + \phi_1 g(\mathbf{x}_t; \boldsymbol{\theta}^{\text{DNN}}) + \mathbf{v}_t^{\text{APBM}}, \quad (47)$$

where the APBM parameters  $\boldsymbol{\theta}^{\text{APBM}} = [\phi_0, \phi_1, \boldsymbol{\theta}^{\text{DNN}}]$  may be estimated offline or online, with supervised or unsupervised training strategies. To ensure, that the data-based component does not overrule the PBM component, we can use *constrained* state estimation algorithms where the parameters of the DNN are *encouraged* to be close to a nominal value  $\bar{\boldsymbol{\theta}}$  that keeps the DNN contribution limited. State estimation can be interpreted as regularized optimization problems, such as

$$(\hat{\mathbf{x}}_{1:t}, \boldsymbol{\theta}) = \arg \min_{(\mathbf{x}_{1:t}, \boldsymbol{\theta})} \mathcal{L}_{\mathbb{D}_u}(\mathbf{y}_{1:t}, \hat{\mathbf{y}}_{1:t}) + \eta \mathcal{R}(\mathbf{x}_{1:t}, \boldsymbol{\theta})$$

where  $\mathcal{L}_{\mathbb{D}_u}$  and  $\mathcal{R}$  are cost and regularization functionals, and  $\eta \in \mathbb{R}_+$  is a regularization parameter governing the trade-off between model fit and regularization. In the context of Kalman filtering, the minimized cost function is the MSE of the state given observations and dynamics (47) model and the control of the DNN can be effectively controlled through the regularization term, for which several implementation options are possible [15].

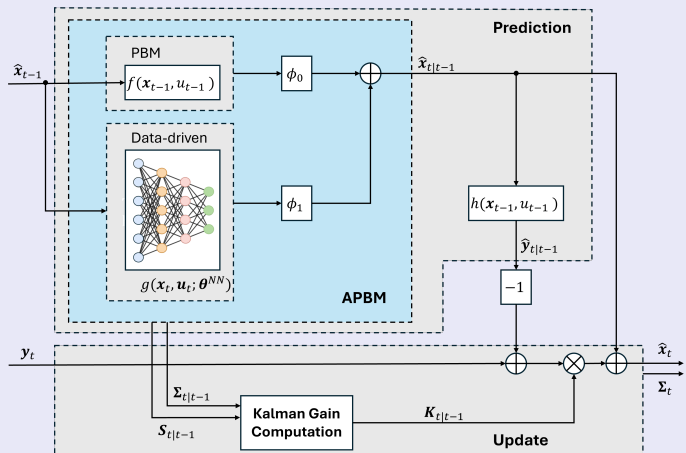


Fig. 10: Illustrative APBM with additive structure.

was successfully used for characterization of the unmodelled acceleration caused, e.g., by the quadrotor drag [67].

## VI. COMPARATIVE STUDY

The previous sections presented a set of design approaches and concrete algorithms fusing Kalman-type filtering and AI techniques. While all these methods tackle the same task of state estimation in dynamic systems, they notably vary in their strengths, requirements, and implications. To highlight the interplay between the approaches mentioned above and provide an understanding of how practitioners should prioritize one approach over the others, we next provide a comparative study. We divide this study into two parts: we first present a *qualitative comparison*, which pinpoints the conceptual differences between the approaches in light of the identified desired properties **P1-P5** and challenges **C1-C4**. We then detail a *quantitative study* for a representative scenario of tracking in challenging dynamics to capture the regimes in which each approach is expected to be preferable.

### A. Qualitative Comparison

Here, we discuss the relationship and individual gains of the different approaches over each other in terms of key conceptual figures of merit that are not quantifiable in the same sense as estimation performance on a given test-bed is. Based on the identified desired properties **P1-P5** and challenges **C1-C4** of conventional Kalman-type algorithms, we focus on qualitative comparison in terms of domain knowledge, interpretability, uncertainty extraction, adaptability, target family of SS models, and learning framework. The comparison detailed below is summarized in Table I.

1) *Domain Knowledge*: The presented methodologies substantially vary in the level of knowledge and characterization of the underlying SS model required, corresponding to challenges **C1** and **C2**. The extreme cases are those of fully model-based filters (such as the standard KF and its variants) that require full and accurate knowledge of the SS model, and those of end-to-end DNNs, that are full model-agnostic (yet do not incorporate available characterization when such is provided).

Among the methodologies representing AI-augmented KFs, the usage of an external DNN to classic Kalman-type tracking typically requires the same level of domain knowledge needed to apply its classic counterpart. When the DNN is applied in parallel as a learned correction term, then the complete SS model should be known, while applying a learned pre-processing module can compensate for unknown and intractable observation modeling. Settings in which one has full knowledge of the observations model but does not know the state evolution model are the focus of methods based on learned state estimation and the SS-oriented DDM and PINN. Knowledge (though possibly an approximated one, as in **C1**) of the functions  $f(\cdot)$  and  $h(\cdot)$  is required by techniques that learn the KG, while some modeling of these functions is needed by SS-oriented approaches based on parameter learning and APBM.

2) *Interpretability*: In interpretability of a state estimation algorithm, we refer to the ability to explain the operation of each computation and associate its internal features with concrete meaning, as in **P3**. This high level of interpretability is naturally provided by purely model-based KF-type algorithms, as well as by SS-oriented AI techniques that learn the parameters of a pre-determined parametric SS model [55], similarly to conventional system identification. On the contrary, algorithms employing end-to-end DNNs for state estimation

Approach		SS Knowledge	Interpretability	Uncertainty	Adaptability	SS Model	Learning
End-to-end DNN	Generic	Fully model-agnostic	Black-box	Only outputs state	Not adaptive	Can be non-linear and non-Gaussian	Supervised learning from large data sets
	KF-inspired (e.g., RKN [7])	Fully model-agnostic	Connection of black-box DNNs	Estimates both state and uncertainty	Not adaptive	Can be non-linear and non-Gaussian	Supervised learning from large data sets
External DNN	Learned pre-process	Requires state evolution	Input transformation is black-box, while tracking based on the latent representation is fully interpretable	Tracks both state and uncertainty	Adapts to known variations only in the state evolution	State evolution should be simple (preferably linear) and Gaussian	Preferably supervised for end-to-end learning, though DNN can also be trained as a feature extractor, possibly unsupervised
	Learned correction (e.g., [19])	Requires (estimated) SS representation	Highly interpretable	Tracks both state and uncertainty	Designed for a specific SS model	Should be Gaussian with simple non-linearities	Supervised learning from moderate data sets
Integrated DNN	Learned KG (e.g., KalmanNet [16])	Requires (estimates) of $h(\cdot)$ and $f(\cdot)$	Processing follows the same pipeline as standard EKF, with KG computation being black-box	Can be extracted in some scenarios [43]	Requires re-training [21] or hypernetworks [44] for adaptivity	Can be non-linear and non-Gaussian	Preferably supervised for end-to-end learning, can also be trained unsupervised with additional domain knowledge [21]
	Learned State Estimation (e.g., DANSE [20])	Requires SS observation model and does not require SS state evolution model	Posterior update is interpretable and resembles that of SS-oriented KF, while state prediction is black-box	Tracks both state and uncertainty	Requires re-training for adaptivity	State evolution model can be non-linear, observation model should be linear, Gaussian	Unsupervised learning
SS-Oriented DNN	DDM (e.g., fully data-driven model or PINN [68])	Fully model-agnostic (DNN) or some knowledge of the model required (PINN). The observation model is assumed to be known.	The state evolution is black-box, while the filter operation is interpretable	Tracks both state and uncertainty	Not adaptive	Can be non-linear and non-Gaussian	Supervised learning
	Parameter learning (e.g., [55])	Requires models of $h(\cdot)$ and $f(\cdot)$	Fully interpretable	Tracks both state and uncertainty	Not adaptive	Can be non-linear and non-Gaussian	Unsupervised learning
	APBM (e.g., [15])	Requires models of $h(\cdot)$ and $f(\cdot)$	Highly interpretable	Tracks both state and uncertainty	Fully adaptive	Can be non-linear and non-Gaussian	Unsupervised learning
Model-Based	Kalman-type filters (e.g., EKF)	Requires (accurate) SS representation	Fully interpretable	Tracks both state and uncertainty	Fully adaptive	Should be Gaussian with mild non-linearities	Fully model-based

TABLE I: Qualitative comparison between the considered approaches.

are essentially black-box methods.

The level of interpretability varies when considering algorithms that augment classic algorithms with DNNs, such that some of the computations are based on principled statistical models, while some are based on black-box data-driven pipelines. For instance, the internal features of task-oriented designs with integrated DNNs are exactly those of standard Kalman-type algorithms, while some of the internal computations – such as the computation of the KG in KalmanNet or the propagation of the prior state moments in DANSE – are carried out by black-box DNNs. A higher level of interpretability is provided when the DNN is applied in parallel to a fully model-based and interpretable algorithm for learned correction, as in [19].

3) *Uncertainty*: The ability to provide faithful *uncertainty* measures (as in P4) arises in classic Kalman-type algorithms from their inherent tracking of the error (posterior) covariance matrix. As the posterior update is an integral aspect of the *update* step of conventional Kalman filtering, any algorithm that fully implements this computation provides uncertainty measures along with its state estimate. This property is exhibited by fully model-based Kalman-type filters, as well as by some AI-aided KFs. Specifically, both SS-oriented designs,

which implement conventional filtering on top of a learned state evolution model, as well as task-oriented architecture that augments the prior state prediction (as in [69]) or utilize DNNs external to model-based filters provide uncertainty in the update computation.

Among AI-aided algorithms that do not fully preserve the posterior covariance propagation of classic KFs, the ability to provide uncertainty varies between the approaches. Generic end-to-end DNNs that track only the state do not offer such measures, while KF-inspired architectures are designed to output error covariances via their update DNNs. AI-aided KFs with learned KG (such as KalmanNet) are specifically designed to bypass the need to propagate second-order moments such as the posterior covariance. Still, as noted in the box entitled ‘*Uncertainty Extraction from Learned KG*’ on Page 12, in some settings, one can still extract uncertainty measures from their internal KG features.

4) *Adaptability*: As noted in P2, classic KF-type algorithms are adaptable to temporal variations in the statistical model governing the underlying dynamic system. Assuming that one can identify and characterize the variations, the updated SS model parameters are simply substituted into the filter equations. End-to-end DNNs, in which there is no explicit

dependence on the SS model parameters, and these are implicitly embedded in the learned weights, adapting the SS models that deviate from those observed in training necessitates re-training, i.e., they are not adaptive.

Task-oriented AI-aided KFs are trained in a manner that is entangled with the operation of the augmented Kalman-type algorithm. Accordingly, while algorithms with integrated and external DNNs typically use the SS model parameters in their overall processing, their DNN modules are fixed, and are suitable for the SS models observed in training. Accordingly, adaptations require re-training, where some level of variations can still be coped with using hypernetworks [44]. An exception is the usage of external DNNs for feature extraction, which is typically designed to map complex observations into a simplified observation model and thus are not affected by variations in the state evolution model. Among SS-oriented approaches, adaptivity is provided by the design of APBM, while alternatives based on, e.g., DDM or PINN architectures need retraining when the state evolution statistics change.

5) *SS Model Type*: A key distinct property among the presented methodologies lies in the family of *SS models* for which each algorithm is suitable. As noted in C3, fully model-based KF-type algorithms are most suitable for Gaussian SS models with simple and preferably not highly non-linear state evolution and observation models. On the opposite edge of the spectrum are end-to-end model-agnostic DNNs, that can be applied in complex and intractable dynamic systems, provided with sufficient data corresponding to that task.

Among task-oriented AI-augmented KFs, architectures employing external DNNs as learned correction terms are suitable for similar SS models as classic KFs, being geared towards settings where the latter is applicable. External DNNs that pre-process observations facilitate coping with complex observation models, while integrated DNNs for learned state estimation overcome complex state evolution models, with each approach requiring the remainder of the SS model to be simple (preferably linear), fully known, and Gaussian. Augmenting KFs with learned KG notably enhances the family of applicable SS models, not requiring Gaussianity and modeling of any of the noise terms, and learning to cope with non-linear and possibly approximated state evolution and observation functions. The same holds for the reviewed SS-oriented methods that can all cope with non-linear and non-Gaussian dynamics.

6) *Learning Framework*: The learning framework is specific to algorithms that incorporate AI, and encapsulates the amount of data and the reliance on its labeling. Accordingly, fully model-based algorithms that rely on given mathematical modeling of the SS representation do not involve learning in their formulation, while black-box DNNs typically require learning from large volumes of labeled data sets.

Designs combining partial domain knowledge with deep learning techniques typically result in a dominant inductive bias that allows training with limited data, compared to black-box end-to-end DNNs. Methods based SS-oriented DDM/PINN, external DNN, and learned integrated KG DNN, typically require this data to be labeled, though for learned pre-processing and for KalmanNet one can also train unsuper-

vised in some settings. Methods such as DANSE, SS-oriented parameter learning, and APBM are specifically designed to be trained based solely on observations, i.e., unsupervised.

## B. Quantitative Comparison

To illustrate the performance of the considered state estimation algorithms, we provide a dedicated experimental study<sup>4</sup>. We consider the task of tracking the nonlinear movement of a free particle in three-dimensional space ( $m = 3$ ) from noisy position observations. The Lorenz Attractor, a chaotic solution to the Lorenz system of ordinary differential equations, defines the continuous-time state evolution of the particle's trajectory. To enhance the challenge of this tracking case, we introduce additional uncertainty due to a sampling-time mismatch in the observation process. While the underlying ground truth synthetic trajectory is generated using a high-resolution time interval ( $\Delta\tau = 10^{-5}$ ), the tracking filter can access only noisy observations decimated at a rate of  $\frac{1}{2000}$ , resulting in a decimated process with  $\Delta t = 0.02$ .

To ensure a fair evaluation, we restrict any populated evolution model of the tracking filter, if it exists, to discrete time with  $\Delta t \geq 0.02$ . Further details about the Lorenz Attractor evolution model can be found in [39]. The averaged MSE values and their standard deviation values for filtering 10 sequences with a length of  $T = 3000$  time steps are reported in Table II. There, we compare the MSE achieved by using the noisy observations as state estimate (*Noise*), that of filtering via the model-based EKF and particle filter (PF); the DNN-integrated KalmanNet and DANSE; and two forms of the SS-oriented APBM, employing offline and online learning.

In the comparative performance reported in Table II, the model-based filters, i.e., the EKF and PF manage to improve upon using noisy observations, but suffer from an error floor due to their sampling mismatch. All considered AI-aided filters manage to improve upon this error floor. Specifically, KalmanNet, which is trained in a supervised manner, achieves the best performance, improving by approximately 5 dB in MSE compared to the model-based algorithms. Among the AI-aided filters that are trained from unlabeled data, DANSE achieves the best performance. APBM, which is designed to boost adaptivity, achieves MSE within a small gap of DANSE when trained offline. These quantitative results complement the conceptual comparison provided in Table I, in revealing the interplay between the reviewed methods for combining deep learning with classic KF-type filtering.

## VII. FUTURE RESEARCH DIRECTIONS

The tutorial-style presentation of designs, algorithms, and experiments of AI-aided KFs indicate potential gains of proper fusion of model-based tracking algorithms and data-driven deep learning techniques. These, in turn, give rise to several core research directions that can be explored to further unveil the potential, strengths, and prospective use cases of such designs. Exploring these directions is expected to further advance

<sup>4</sup>The source code for all simulated algorithms and the hyperparameters used can be found online at [https://github.com/ShlezingerLab/AI\\_Aided\\_KFs](https://github.com/ShlezingerLab/AI_Aided_KFs)

TABLE II: MSE [dB] - Lorenz Attractor with Sampling Mismatch.

Noise	EKF	PF	KalmanNet	DANSE	APBM (offline)	APBM (online)
-0.024	-6.316	-5.333	-11.106	-10.115	-9.457	-7.452
$\pm 0.049$	$\pm 0.135$	$\pm 0.136$	$\pm 0.224$	$\pm 0.163$	$\pm 0.231$	$\pm 0.548$

the development and understanding of algorithmic tools to jointly leverage classic SS model-based KF-type algorithms with emerging deep learning techniques, and preserving the individual strengths of each approach that has a bearing on the potential of alleviating some of the core challenges shared by a broad range technologies tracking dynamic systems and state estimation. We thus conclude this article with a discussion of some of the open challenges that can serve as future research directions.

- 1) *Time-varying SS models and adaptation*: A practical state estimation method might need to cater for a scenario where the observation model (1b) is mismatched between training and testing stages, and/or the observation or state evolution functions are time-varying without a clear pattern of time-varying nature. Typically AI-aided KFs are tied to a fixed observation model, typically not time-varying, and remains same between the the training stage and inference stage. Therefore, the future research question is how to design AI-aided state estimation methods that can adapt to time-varying SS models, mainly for time-varying observation models.
- 2) *Non-Markovian SS models*: Throughout this article and being consistent with usual practice, the prevalent state evolution model is Markovian in nature (see (1a)). Naturally, it is a quest to design methods that can exploit short and long-term memories in state evolution, which means non-Markovian state evolution.
- 3) *Non-Gaussian SS models*: Most of the directions discussed in the article focused on Gaussian noises in SS models. Extensions to non-Gaussian noises encountered in real problems will require more complex algorithms with high computational complexity, which is challenging, especially when learning the large-scale DNN parameters.
- 4) *Distributed AI-aided KFs*: While there exists considerable research on distributed KFs, such as the work of [70] for sensor networks, there is little attention currently to design distributed AI-aided KFs. Design of them for federated learning and edge computing can be a new research direction.
- 5) *Robust training*: Outliers appearing in the SS models, mainly in the observation model (1b), may substantially affect the training process. The approaches should address this issue by providing robust training.

#### REFERENCES

- [1] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [2] S. F. Schmidt, "The Kalman filter-its recognition and development for aerospace applications," *Journal of Guidance and Control*, vol. 4, no. 1, pp. 4–7, 1981.
- [3] M. S. Grewal and A. P. Andrews, "Applications of Kalman filtering in aerospace 1960 to the present [historical perspectives]," *IEEE Control Syst. Mag.*, vol. 30, no. 3, pp. 69–78, 2010.
- [4] J. Durbin and S. J. Koopman, *Time Series Analysis by State Space Methods*. Oxford University Press, 05 2012, vol. 38.
- [5] S. Gannot and A. Yeredor, "The Kalman filter," *Springer Handbook of Speech Processing*, pp. 135–160, 2008.
- [6] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- [7] P. Becker, H. Pandya, G. Gebhardt, C. Zhao, C. J. Taylor, and G. Neumann, "Recurrent Kalman networks: Factorized inference in high-dimensional deep feature spaces," in *International Conference on Machine Learning*, 2019, pp. 544–552.
- [8] R. Krishnan, U. Shalit, and D. Sontag, "Structured inference networks for nonlinear state space models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [9] S. Kuutti, R. Bowden, Y. Jin, P. Barber, and S. Fallah, "A survey of deep learning applications to autonomous vehicle control," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 2, pp. 712–733, 2020.
- [10] J. Chen and X. Ran, "Deep learning with edge computing: A review," *Proc. IEEE*, vol. 107, no. 8, pp. 1655–1674, 2019.
- [11] N. Shlezinger, J. Whang, Y. C. Eldar, and A. G. Dimakis, "Model-based deep learning," *Proc. IEEE*, vol. 111, no. 5, pp. 465–499, 2023.
- [12] A. Klushyn, R. Kurlle, M. Soelch, B. Cseke, and P. van der Smagt, "Latent matters: Learning deep state-space models," *Advances in Neural Information Processing Systems*, vol. 34, pp. 10234–10245, 2021.
- [13] H. Coskun, F. Achilles, R. DiPietro, N. Navab, and F. Tombari, "Long short-term memory Kalman filters: Recurrent neural estimators for pose regularization," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5524–5532.
- [14] D. Gedon, N. Wahlström, T. B. Schön, and L. Ljung, "Deep state space models for nonlinear system identification," *IFAC-PapersOnLine*, vol. 54, no. 7, pp. 481–486, 2021.
- [15] T. Imbiriba, O. Straka, J. Dunik, and P. Closas, "Augmented physics-based machine learning for navigation and tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 60, no. 3, pp. 2692–2704, 2024.
- [16] G. Revach, N. Shlezinger, X. Ni, A. L. Escoriza, R. J. van Sloun, and Y. C. Eldar, "KalmanNet: Neural network aided Kalman filtering for partially known dynamics," *IEEE Trans. Signal Process.*, vol. 70, pp. 1532–1547, 2022.
- [17] I. Buchnik, G. Revach, D. Steger, R. J. van Sloun, T. Rottengberg, and N. Shlezinger, "Latent-KalmanNet: Learned Kalman filtering for tracking from high-dimensional signals," *IEEE Trans. Signal Process.*, vol. 72, pp. 352–367, 2023.
- [18] K. Pratik, R. A. Amjad, A. Behboodi, J. B. Soriaga, and M. Welling, "Neural augmentation of Kalman filter with hypernetwork for channel tracking," in *IEEE Global Communications Conference (GLOBECOM)*, 2021.
- [19] V. G. Satorras, Z. Akata, and M. Welling, "Combining generative and discriminative models for hybrid inference," in *Advances in Neural Information Processing Systems*, 2019, pp. 13802–13812.
- [20] A. Ghosh, A. Honoré, and S. Chatterjee, "DANSE: Data-driven nonlinear state estimation of model-free process in unsupervised learning setup," *IEEE Trans. Signal Process.*, vol. 72, pp. 1824–1838, 2024.
- [21] G. Revach, N. Shlezinger, T. Locher, X. Ni, R. J. G. van Sloun, and Y. C. Eldar, "Unsupervised learned Kalman filtering," in *European Signal Processing Conference (EUSIPCO)*, 2022, pp. 1571–1575.
- [22] S. N. Aspragkathos, G. C. Karras, and K. J. Kyriakopoulos, "Event-triggered image moments predictive control for tracking evolving features using UAVs," *IEEE Robot. Autom. Lett.*, vol. 9, no. 2, pp. 1019–1026, 2024.
- [23] A. Milstein, G. Revach, H. Deng, H. Morgenstern, and N. Shlezinger, "Neural augmented Kalman filtering with Bollinger bands for pairs trading," *IEEE Trans. Signal Process.*, vol. 72, pp. 1974–1988, 2024.
- [24] A. Gu, I. Johnson, K. Goel, K. Saab, T. Dao, A. Rudra, and C. Ré, "Combining recurrent, convolutional, and continuous-time models with linear state space layers," *Advances in neural information processing systems*, vol. 34, pp. 572–585, 2021.
- [25] T. Jebara, *Machine learning: discriminative and generative*. Springer Science & Business Media, 2012, vol. 755.

- [26] N. Shlezinger and T. Rauttenberg, “Discriminative and generative learning for linear estimation of random signals [lecture notes],” *IEEE Signal Process. Mag.*, vol. 40, no. 6, pp. 75–82, 2023.
- [27] L. Zhou, Z. Luo, T. Shen, J. Zhang, M. Zhen, Y. Yao, T. Fang, and L. Quan, “KFNet: Learning temporal camera relocation using Kalman filtering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4919–4928.
- [28] S. Jouaber, S. Bonnabel, S. Velasco-Forero, and M. Pilte, “NNAKF: A neural network adapted Kalman filter for target tracking,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 4075–4079.
- [29] A. Bensoussan, *Stochastic control of partially observable systems*. Cambridge University Press, 1992.
- [30] S. Särkkä and L. Svensson, *Bayesian filtering and smoothing*. Cambridge University Press, 2023, vol. 17.
- [31] H. E. Rauch, F. Tung, and C. T. Striebel, “Maximum likelihood estimates of linear dynamic systems,” *AIAA Journal*, vol. 3, no. 8, pp. 1445–1450, 1965.
- [32] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.
- [33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [34] A. Gu and T. Dao, “Mamba: Linear-time sequence modeling with selective state spaces,” *arXiv preprint arXiv:2312.00752*, 2023.
- [35] A. Gu, K. Goel, and C. Ré, “Efficiently modeling long sequences with structured state spaces,” *International Conference on Learning Representations (ICLR)*, 2022.
- [36] N. Shlezinger, Y. C. Eldar, and S. P. Boyd, “Model-based deep learning: On the intersection of deep learning and optimization,” *IEEE Access*, vol. 10, pp. 115 384–115 398, 2022.
- [37] A. H. Li, P. Wu, and M. Kennedy, “Replay overshooting: Learning stochastic latent dynamics with the extended Kalman filter,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [38] L. Xu and R. Niu, “EKFNet: Learning system noise statistics from measurement data,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 4560–4564.
- [39] G. Revach, X. Ni, N. Shlezinger, R. J. van Sloun, and Y. C. Eldar, “RTSNet: Learning to Smooth in Partially Known State-Space Models,” *IEEE Trans. Signal Process.*, vol. 71, pp. 4441–4456, 2023.
- [40] 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 Trans. Veh. Technol.*, vol. 72, no. 9, pp. 12 326–12 331, 2023.
- [41] J. Wang, X. Geng, and J. Xu, “Nonlinear Kalman filtering based on self-attention mechanism and lattice trajectory piecewise linear approximation,” *arXiv preprint arXiv:2404.03915*, 2024.
- [42] A. N. Putri, C. Machbub, D. Mahayana, and E. Hidayat, “Data driven linear quadratic Gaussian control design,” *IEEE Access*, vol. 11, pp. 24 227–24 237, 2023.
- [43] Y. Dahan, G. Revach, J. Dunik, and N. Shlezinger, “Uncertainty quantification in deep learning based Kalman filters,” *arXiv preprint arXiv:2309.03058*, 2023.
- [44] X. Ni, G. Revach, and N. Shlezinger, “Adaptive KalmanNet: Data-driven Kalman filter with fast adaptation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 5970–5974.
- [45] A. Sakai and Y. Kuroda, “Discriminative parameter training of unscented Kalman filter,” *IFAC Proceedings Volumes*, vol. 43, no. 18, pp. 677–682, 2010, 5th IFAC Symposium on Mechatronic Systems.
- [46] Z. Fan, D. Shen, Y. Bao, K. Pham, E. Blasch, and G. Chen, “RNN-UKF: Enhancing hyperparameter auto-tuning in unscented kalman filters through recurrent neural networks,” in *27th International Conference on Information Fusion*, Venice, Italy, 2024.
- [47] A. Chiuso and G. Pillonetto, “System identification: A machine learning perspective,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, no. 1, pp. 281–304, 2019.
- [48] L. Ljung, C. Andersson, K. Tiels, and T. B. Schön, “Deep learning and system identification,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 1175–1181, 2020, 21st IFAC World Congress.
- [49] J. Suykens, B. Moor, and J. Vandewalle, “Nonlinear system identification using neural state space models, applicable to robust control design,” *International Journal of Control*, vol. 62, no. 1, p. 129–152, 1995.
- [50] D. Masti and A. Bemporad, “Learning nonlinear state–space models using autoencoders,” *Automatica*, vol. 129, pp. 0005–1098, 2021.
- [51] A. Gorji and M. Menhaj, “Identification of nonlinear state space models using an MLP network trained by the EM algorithm,” in *2008 IEEE International Joint Conference on Neural Networks*, 2008, p. 53–60.
- [52] Y. Wang, “A new concept using LSTM neural networks for dynamic system identification,” in *2017 American Control Conference (ACC)*, 2017, pp. 5324–5329.
- [53] M. Fraccaro, S. Sønderby, U. Paquet, and O. Winther, “Sequential neural models with stochastic layers,” in *Advances in neural information processing systems*, 2016, p. 29.
- [54] F. Arnold and R. King, “State–space modeling for control based on physics-informed neural networks,” *Engineering Applications of Artificial Intelligence*, vol. 101, p. 104195, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0952197621000427>
- [55] Y. Bao, J. M. Velni, A. Basina, and M. Shahbakhti, “Identification of state-space linear parameter-varying models using artificial neural networks,” *IFAC-PapersOnLine, Volume*, vol. 53, pp. 5286–5291, 2405–8963, 2020.
- [56] M. Schoukens, “Improved initialization of state-space artificial neural networks,” *arXiv preprint arXiv:2103.14516*, 2021.
- [57] M. Forgione and D. Piga, “Model structures and fitting criteria for system identification with neural networks,” in *IEEE International Conference Application of Information and Communication Technologies*, 2020.
- [58] S. Tang, T. Imbiriba, J. Duník, O. Straka, and P. Closas, “Augmented Physics-Based Models for High-Order Markov Filtering,” *Sensors*, vol. 24, no. 18, 2024. [Online]. Available: <https://www.mdpi.com/1424-8220/24/18/6132>
- [59] T. Imbiriba, A. Demirkaya, J. Duník, O. Straka, D. Erdoğan, and P. Closas, “Hybrid neural network augmented physics-based models for nonlinear filtering,” in *International Conference on Information Fusion (FUSION)*. IEEE, 2022.
- [60] R. Krishnan, U. Shalit, and D. Sontag, “Deep Kalman filters,” *arXiv preprint arXiv:1511.05121*, 2015.
- [61] D. Simon, “Kalman filtering with state constraints: a survey of linear and nonlinear algorithms,” *IET Control Theory & Applications*, vol. 4, no. 8, pp. 1303–1318, 2010.
- [62] J. Humpherys, P. Redd, and J. West, “A fresh look at the kalman filter,” *SIAM review*, vol. 54, no. 4, pp. 801–823, 2012.
- [63] A. Bemporad, “Recurrent neural network training with convex loss and regularization functions by extended kalman filtering,” *IEEE Trans. Autom. Control*, vol. 68, no. 9, pp. 5661–5668, 2023.
- [64] J. Dunik, O. Straka, O. Kost, S. Tang, T. Imbiriba, and P. Closas, “Noise identification for data-augmented physics-based state-space models,” in *Accepted for IEEE Workshop on Signal Processing Systems (SiPS 2024)*, Cambridge, MA, USA, 2024.
- [65] L. Ljung, “Perspectives on system identification,” *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 7172–7184, 2008, 17th IFAC World Congress.
- [66] D. Li, J. Zhou, and Y. Liu, “Recurrent-neural-network-based unscented Kalman filter for estimating and compensating the random drift of MEMS gyroscopes in real time,” *Mechanical Systems and Signal Processing*, vol. 147, p. 107057, 2021.
- [67] G. Torrente, E. Kaufmann, P. Föhn, and D. Scaramuzza, “Data-driven MPC for quadrotors,” *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 3769–3776, 2021.
- [68] M. Raissia, P. Perdikaris, and G. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [69] A. Ghosh, J. Hong, D. Yin, and K. Ramchandran, “Robust federated learning in a heterogeneous environment,” *arXiv preprint arXiv:1906.06629*, 2019.
- [70] R. Olfati-Saber, “Distributed kalman filtering for sensor networks,” in *2007 46th IEEE Conference on Decision and Control*, 2007, pp. 5492–5498.