POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Analytically Tractable Exponential Component for Monitoring the Alkali-Aggregate Reaction in Dams with State-Space Models

MICHEL WU
Département des génies civil, géologique et des mines

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
Génie civil

Août 2026

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Ce mémoire intitulé :

Analytically Tractable Exponential Component for Monitoring the Alkali-Aggregate Reaction in Dams with State-Space Models

présenté par **Michel WU**
en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
a été dûment accepté par le jury d'examen constitué de :

Rocio L. SEGURA, présidente
James-A. GOULET, membre et directeur de recherche
Dany LAUZON, membre

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my supervisor Prof. James-A. Goulet. It is through the knowledge he shared with me, the patience he showed in explaining concepts, and the guidance he offered that I was able to complete this research.

Major thanks to Hydro-Québec (HQ) for funding this research. I am also grateful to the people I met there, whose advice helped me to progress in my project.

I would like to thank my friends and colleagues at Polytechnique Montréal, Lucas Alric, Marina Decker, Miquel Florensa, Dai Vuong Van, David Wardan and Zhanwen Xin, for all the advice they gave me, both on this research and on life in general, for the laughter we shared, and for the kindness that kept me motivated right up to the end.

My thanks also go to my friends and roommates, Tiphaine Bloch and Timé Jaupitre, who came from France with me and with whom I shared these two years more closely than with anyone else. Thank you for standing with me from start to finish, through both the hard moments and the happy ones.

To all the friends I made here in Canada, thank you for making this experience unforgettable and so full of warmth.

Finally, the people I want to thank most of all are, without question, my family, for their unconditional support. To my parents who have always believed in me. To both of my sisters, who have taught me so much about life.

RÉSUMÉ

Les barrages en béton affectés par la réaction alcalis-granulats (AAR) subissent une lente expansion non linéaire, pouvant compromettre leur intégrité structurale sur plusieurs décennies. Ainsi, des capteurs sont installés sur les barrages afin d'obtenir des séries temporelles décrivant le comportement de ces derniers. Ces données sont interprétées en utilisant diverses méthodes. Ce mémoire s'intéresse particulièrement aux modèles espace-état (SSM) et au filtre de Kalman (KF), qui permettent de décomposer de façon probabiliste les signaux en composantes interprétables au fur et à mesure de leur acquisition. Toutefois, le KF a été développé pour fonctionner avec des systèmes linéaires et est ainsi inutilisable en tant que tel pour suivre des phénomènes non linéaires tels que l'AAR. Pour pallier cette limite, des extensions non linéaires du filtre de Kalman peuvent être employées telles que le filtre de Kalman étendu (EKF), sans parfum (UKF) et par cubature (CKF). Néanmoins, ces méthodes présentent des limites, telles que la divergence des estimations ou la perte de tractabilité analytique.

Motivé par les modèles cinétiques utilisant des fonctions exponentielles pour décrire l'AAR, ce mémoire développe une composante exponentielle purement analytique et l'intègre dans le formalisme des SSM et du KF. Pour cela, les propriétés de la loi log-normale et la *Gaussian multiplicative approximation* (GMA) sont utilisées pour obtenir la moyenne et la variance de la composante exponentielle, tandis qu'une stratégie de linéarisation analytique permet d'obtenir ses covariances croisées avec les autres composantes du SSM. Ensuite, les algorithmes du filtre de Kalman standard et du lissage sont adaptés pour intégrer ces formulations analytiques à chaque étape linéaire. Suite à cela, une vérification de la méthodologie est faite sur des données synthétiques où les résultats permettent de conclure que les algorithmes ont bien convergé vers les valeurs réelles des paramètres utilisés pour créer les cas synthétiques, à condition que le filtre soit initialisé avec $\pm 25\%$ de coefficient de variation par rapport aux valeurs réelles. Néanmoins, une vérification manuelle des résultats reste nécessaire car certains cas peuvent produire des valeurs aberrantes. Une validation sur des données réelles a démontré que la composante exponentielle est capable de modéliser le comportement à long terme attendu des barrages affectés par l'AAR. Une dernière étude est menée sur une série temporelle issue d'un barrage ayant subi une intervention. Cette étude a permis de montrer que la composante exponentielle peut aider à quantifier le changement dans la dynamique non linéaire et à prédire le comportement futur du barrage en injectant des moyennes et des variances a priori ciblées au moment de l'intervention.

ABSTRACT

Concrete dams affected by the alkali-aggregate reaction (AAR) exhibit slow, non-linear expansion that can compromise their structural integrity over several decades. To monitor this behaviour, dams are equipped with sensors that record their structural behaviour over time. Among the various methods available to interpret these measurements, this thesis focuses on state-space models (SSM) and the Kalman filter (KF), a probabilistic framework that decomposes measured signals into interpretable hidden states. However, because the KF assumes linear system dynamics, it cannot directly model the non-linear expansion caused by AAR. Existing non-linear extensions of the KF, such as the extended (EKF), unscented (UKF), and cubature (CKF) Kalman filters address this issue but introduce their own limitations, including estimation divergence and the loss of end-to-end analytical tractability.

Motivated by the exponential-based kinetic models used in the literature to describe AAR expansion, this thesis develops a purely analytical exponential component and integrates it within the SSM and KF formalism. The mean and variance for the hidden states of the exponential component are derived analytically using log-normal properties and a Gaussian multiplicative approximation (GMA). An analytical linearization strategy is then applied to relate the exponential component to the other state variables through their cross-covariances. The standard KF and smoother are adapted to incorporate these formulations after each linear step. Verification on synthetic data shows that the algorithms recover the true parameter values used to generate the synthetic cases, provided that the filter is initialized within a $\pm 25\%$ coefficient of variation of these values. Nevertheless, manual verification of the results remains necessary, as certain cases can converge to biased values. Validation on real dam-monitoring data shows that the exponential component can capture the long-term behaviour of AAR-affected dams. A case study conducted on a dam that had a structural intervention shows that, by injecting targeted means and variances to the hidden state vector at the time of the intervention, the exponential component can help quantify the resulting shift in the non-linear dynamics and predict the dam's future behaviour.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ACRONYMS

Symbols

a	Exponential scaling factor
$\mathbf{A}$	Transition matrix
$\mathbf{C}$	Observation matrix
$\text{cov}(\cdot)$	Covariance operator
$\mathbf{e}_i$	Vector of an orthonormal basis
$\mathbb{E}[\cdot]$	Expected value operator
$\exp(\cdot)$ or $e^{(\cdot)}$	Exponential operator
$f(\cdot)$	Probability density function
$f(\mathbf{x}_t \vert \mathbf{y}_{1:t-1})$	Prior PDF of hidden states at time t
$f(\mathbf{x}_t \vert \mathbf{y}_{1:t})$	Posterior PDF of hidden states at time t
$f(\mathbf{x}_t \vert \mathbf{y}_{1:\mathrm{T}})$	Smoothed PDF of hidden states at time t
$f(\mathbf{y}_t \vert \mathbf{x}_t)$	Likelihood
$f_i(\cdot)$	Time-dependent polynomial, rational, logarithmic, or exponential function
$\mathbf{f}(\cdot)$	Transition function
$\mathbf{G}$	Innovation covariance matrix
h	Hydrostatic pressure
$\mathbf{h}(\cdot)$	Observation function
$\mathbf{J}$	Kalman smoother gain matrix
$\mathbf{K}$	Kalman gain matrix
$\log(\cdot)$	Logarithmic operator
$\mathcal{N}(\cdot)$	Gaussian distribution
$\mathbf{Q}$	Process error covariance matrix
r	Amplitude parameter for the exponential in the automatic initialization algorithm
$\mathbf{r}$	Innovation vector
$\mathbf{R}$	Observation error covariance matrix
$\mathcal{R}$	Set of amplitude parameters for the exponential in the automatic initialization algorithm
$\sin(\cdot)$	Sine operator
t	Time stamp

t_{itv}	Time of intervention
v	Observation error
$\mathbf{v}$	Observation error vector
$\boldsymbol{V}$	Random vector representing observation error
$\mathbf{w}$	Process error vector
$\boldsymbol{W}$	Random vector representing process error
w_i	Weight of sigma or cubature points
$\mathbf{x}$	Hidden state vector
$x^{\mathtt{A}}$	Scaling factor hidden state
$\hat{x}^{\mathtt{A}}$	Estimated scaling factor hidden state
$x^{\mathtt{G}}$	Generic hidden state
$x^{\mathtt{L}}$	Local level hidden state
$x^{\mathtt{LA}}$	Local acceleration hidden state
$x^{\mathtt{LR}}$	Latent rate hidden state
$\hat{x}^{\mathtt{LR}}$	Estimated latent rate hidden state
$x^{\mathtt{LSTM}}$	Long short-term memory hidden state
$x^{\mathtt{LT}}$	Local trend hidden state
$x^{\mathtt{LTi}}$	Latent time hidden state
$x^{\mathtt{P}}$	Periodic hidden state
$x^{\mathtt{W2}}$	Square of the process error hidden state
$x^{\overline{\mathtt{W2}}}$	Process error's variance term hidden state
$x^{\mathtt{WN}}$	White noise hidden state
$x^{\mathtt{WN\_err}}$	White noise error hidden state
$\mathtt{X}$	Hidden state vector size
$\boldsymbol{\mathcal{X}}_i$	Sigma or cubature points
y	Observation
$\hat{y}$	Predicted observation
$\mathbf{y}$	Observation vector
$\hat{\mathbf{y}}$	Predicted observation vector
$\mathtt{Y}$	Observation vector size
$z^{\mathtt{E}}$	Shifted exponential hidden state
$z^{\mathtt{P}}$	Scaled exponential hidden state
α	Level parameter in the automatic initialization algorithm
$\boldsymbol{\beta}$	Research bounds in the automatic initialization algorithm
γ	Tolerance factor in the automatic initialization algorithm
Γ	Sensitivity factor

δ	Deformation
ϵ	Residual error term
ζ	Acceleration parameter in the automatic initialization algorithm
η	Trend parameter in the automatic initialization algorithm
θ	Temperature
$\boldsymbol{\theta}$	Parameter vector for the automatic initialization algorithm
$\boldsymbol{\theta}^*$	Optimal parameter vector for the automatic initialization algorithm
κ_i	Relative initial uncertainty coefficient for the automatic initialization algorithm
κ_{UKF}	Scaling parameter of the UKF
λ	Time constant parameter for the exponential in the automatic initialization algorithm
Λ	Set of time constant parameters for the exponential in the automatic initialization algorithm
μ	Expected value
$\boldsymbol{\mu}$	Expected value vector
$\mu^{\mathtt{A}}$	Scaling factor expected value
$\mu^{\mathtt{E}}$	Shifted exponential expected value
$\mu^{\mathtt{LR}}$	Latent rate expected value
$\mu^{\mathtt{LTi}}$	Latent time expected value
$\mu^{\mathtt{P}}$	Scaled exponential expected value
$\boldsymbol{\mu}_{t\|t-1}$	Prior mean vector for the hidden states at time t
$\boldsymbol{\mu}_{t\|t}$	Posterior mean vector for the hidden states at time t
$\boldsymbol{\mu}_{t\|\mathtt{T}}$	Smoothed mean vector for the hidden states at time t
ν	Amplitude parameter for the periodic component in the automatic initialization algorithm
$\xi(\cdot)$	Expansion function
ξ^{∞}	Ultimate expansion
ψ_i	Regression coefficient
σ	Standard deviation
$\sigma^{\mathtt{A}}$	Scaling factor standard deviation
$\sigma^{\mathtt{E}}$	Shifted exponential standard deviation
$\sigma^{\mathtt{LR}}$	Latent rate standard deviation
$\sigma^{\mathtt{LTi}}$	Latent time standard deviation
$\sigma^{\mathtt{P}}$	Scaled exponential standard deviation
σ_V	Observation error standard deviation

$\boldsymbol{\Sigma}$	Covariance matrix
$\boldsymbol{\Sigma}_{t\|t-1}$	Prior covariance matrix for the hidden states at time t
$\boldsymbol{\Sigma}_{t\|t}$	Posterior covariance matrix for the hidden states at time t
$\boldsymbol{\Sigma}_{t\|\mathrm{T}}$	Smoothed covariance matrix for the hidden states at time t
τ	Time constant
τ_c	Characteristic time
τ_l	Latency time
ϕ	Phase parameter for the periodic component in the automatic initialization algorithm

Acronyms

AAR	Alkali-Aggregate Reaction
ACR	Alkali Carbonate Reaction
AGVI	Approximate Gaussian Variance Inference
ASR	Alkali Silica Reaction
CKF	Cubature Kalman Filter
COV	Coefficient Of Variation
EKF	Extended Kalman Filter
GMA	Gaussian Multiplicative Approximation
HST	Hydrostatic-Seasonal-Time
KF	Kalman Filter
LL	Log-Likelihood
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MSE	Mean Squared Error
RTS	Rauch-Tung-Striebel
SHM	Structural Health Monitoring
SSM	State-Space Model
UKF	Unscented Kalman Filter

LIST OF APPENDICES

CHAPTER 1 INTRODUCTION

1.1 Motivation

Dams are essential infrastructure for the safety and development of societies [2], particularly in Canada, where hydroelectricity generated more than half of the electricity in 2021 [3]. In Quebec, these concrete structures were built over the last century using concrete formulations that often included high-alkali cements [4]. These conditions favour the development of the alkali-aggregate reaction (AAR), a pathology identified as early as the 1940s [5], which causes progressive and non-linear expansion of the material. This concrete expansion induces displacements that can compromise the long-term integrity of these structures [6].

To ensure the safety of these dams, their structural behaviour is monitored over time. To this end, these structures are equipped with sensors that record their displacements, cracks, etc. [7]. However, in the case of non-linear responses such as AAR, analyzing these measurements remains complex, as the overall observed displacement results from the superposition of several distinct physical phenomena, and its decomposition remains an open question. In this regard, the Hydrostatic-Seasonal-Time (HST) model, proposed by Willm and Beaujoint [8], popularized the decomposition of dam deformation. This model distinguishes between reversible effects, due to variations in hydrostatic loads or seasonal thermal fluctuations, and irreversible effects, which include phenomena such as AAR. Although this approach remains a reference for analyzing dam behaviour, its reliance on linear regression prevents real-time, online updating as new measurements are acquired.

To address this limitation, the use of state-space models (SSM) offers a recursive probabilistic framework suited for the online decomposition of time series [9, 10]. This framework decomposes time series into hidden states, i.e., interpretable components that, once combined, reconstruct the measured signal. The sequential estimation of these hidden states is classically performed using the Kalman filter (KF) [11]. However, while the KF is a powerful filtering tool for linear systems, it is no longer applicable in the presence of non-linearities. Given the non-linear nature of AAR-induced expansion, the standard KF cannot be used as-is for monitoring dams subjected to this pathology. To handle non-linear systems, extensions of the KF have been developed, such as the extended Kalman filter (EKF), the unscented Kalman filter (UKF), and the cubature Kalman filter (CKF). Nevertheless, these methods present limitations such as estimation divergence, or loss of end-to-end analytical tractability.

Motivated by the exponential-based kinetic models used to describe AAR expansion in the

literature [12], this thesis proposes the development of analytical formulations for an exponential component within the standard KF framework. This component enables the KF to model sensor data that exhibit exponential behaviour caused by AAR. By remaining within the standard KF formalism, this component can be combined with other SSM components. This methodology aims to provide a tool for the online monitoring of dams and other civil infrastructure subjected to exponential deterioration phenomena.

1.2 Research Objectives

The main objective of this research is to develop a non-linear exponential component and to integrate it within the formalism of SSM and the KF. To address this problem, this master's study focuses on the following objectives:

- The development of the exponential component and its analytical formulation, allowing its first two moments to be derived in closed form.
- The adaptation of the KF framework to integrate this new formulation.
- The verification of this methodology on synthetic data where the ground truth is known, to ensure that the adapted KF converges properly toward the solutions and to evaluate its robustness against different initialization configurations.
- The validation on real data by comparing the exponential component with classic SSM baseline components, such as local trend and local acceleration, as well as evaluating its ability to quantify dynamic changes following an intervention on the structure.

1.3 Thesis Outline

This thesis is organized as follows: Chapter 2 provides a literature review focusing on AAR to understand its mechanisms as well as existing approaches for modelling the reaction. It then reviews data-driven methods, notably HST-based models, as well as SSM relying on the KF, EKF, UKF, and CKF. Subsequently, Chapter 3 details the methodology for defining the proposed exponential component and its integration into the SSM and KF formalism. Chapter 4 presents the results of the method's verification on synthetic data, followed by its validation on real data. Finally, Chapter 5 draws the conclusion of the thesis, outlines the limitations of the proposed model, and suggests directions for future research.

CHAPTER 2 LITERATURE REVIEW

2.1 Introduction

This chapter first describes the Alkali Aggregate Reaction (AAR), a chemical pathology that causes concrete expansion in infrastructures. Next, it describes the Hydrostatic-Seasonal-Time (HST) based models conventionally used to analyze dam displacements in the context of structural health monitoring (SHM). Then it introduces state-space models (SSM), which provide a framework for the dynamic tracking of infrastructures' behaviour for SHM.

2.2 Alkali Aggregate Reaction

Concrete infrastructures are subjected to AAR [1], an expansive chemical reaction that occurs between the alkali of cement paste and minerals in aggregates and takes two forms : Alkali Carbonate Reaction (ACR) and Alkali Silica Reaction (ASR). ASR, the most common one, creates a gel that swells by absorbing moisture. This increases the pressure in concrete, leading to microcracking and thus loss of resistance and durability of the materials. In contrast, ACR is a less common reaction that follows a process called dedolomitization [1]. The expansion and its kinetics, both on the material and the structure scales, are dependent on multiple environmental and intrinsic factors, such as temperature, relative humidity, type of aggregate, composition of the concrete, and confinement conditions of a given structure [1,6]. For example, confinement-induced stress dictates the anisotropy of the swelling process [13]. As highlighted by Zahedi et al. [13], ASR-induced expansion does not occur uniformly as it redirects expansion to directions where structural stress is lower.

Multiple models have been developed to forecast the AAR reaction for laboratory specimens under static conditions. Among them, the model by Larive [1, 12, 14] is the most prevalent one

$$\xi(t,\theta) = \frac{1 - e^{-\frac{t}{\tau_c(\theta)}}}{1 + e^{-\frac{(t-\tau_l(\theta))}{\tau_c(\theta)}}} \cdot \xi^{\infty}. \tag{2.1}$$

Larive's model in Equation 2.1 follows a logistic function as illustrated in Figure 2.1. It characterizes AAR in four phases and is parameterized by three variables that are dependent on the temperature:

- **The latency time (τ_l)** governs phases 1 and 2. It represents the period during

which silica gels are formed and begin to saturate the concrete's internal porosity. The expansion remains negligible until the gel absorbs enough moisture to exert sufficient hydraulic pressure to overcome the tensile strength of the cement paste. At the end of the latency time, the function is at its inflection point, which represents the theoretical maximum of the expansion rate.

- **The characteristic time** (τ_c) governs phase 3. At this stage, the expansion rate decreases because the pressure has reached a level that leads to micro-cracking. Because the gel has new empty spaces to fill, it decreases the pressure and leads to thus less deformation.

- **The ultimate expansion** (ξ^∞) is the asymptotic limit of the reaction reached in phase 4. It is the stabilization of the phenomenon that occurs once the chemical reactants are fully consumed.

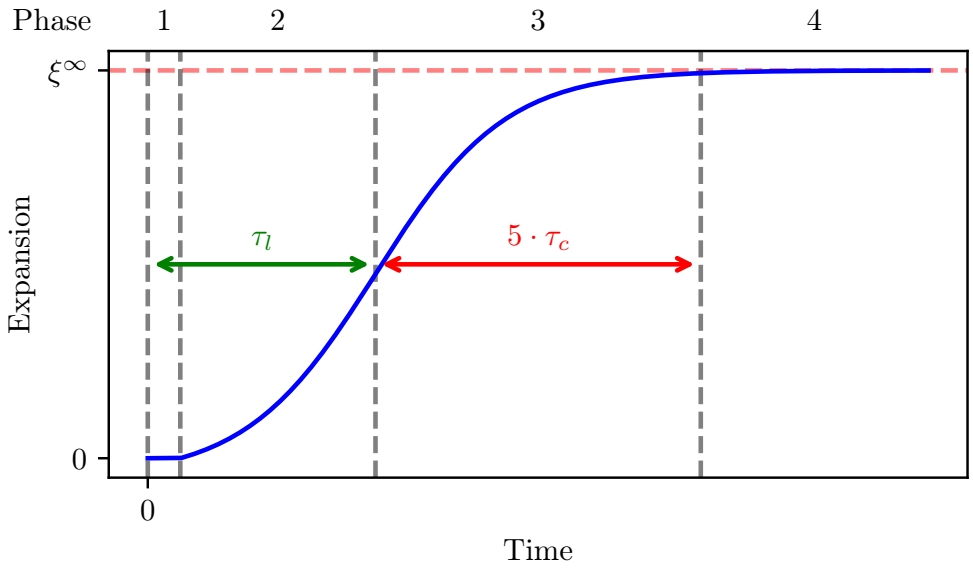

Figure 2.1 Evolution of AAR induced expansion over time according to Larive's model, illustrating the four reaction phases and displaying the three parameters of the model τ_l,τ_c and ξ^∞ (adapted from de Grazia et al. [1]).

De Grazia et al. [1] highlighted that Larive's equation is based on empirical mathematical parameters and thereby, it is only suited to predict AAR expansion and kinetics for conditions that are studied in laboratory conditions. Numerous researchers have since improved this model to account for more general conditions [14]. For instance, Ulm et al. [15] leveraged Larive's experimental data [12] and established a log-linear relationship between both τ_l and

τ_c with the inverse of the temperature, enabling the application of Larive's model across varying temperature profiles. De Grazia et al. [1] extended the study further; in addition to temperature, they investigated the influence of aggregate type, alkali content, and relative humidity, and proposed adding coefficients to the Larive model to account for these factors. Although these studies have enabled a generalization of Larive's model, it solely predicts AAR behaviour and induced expansion for concrete specimens at the laboratory scale, rather than at the structural scale.

One method to bridge this gap and model the displacement of concrete structures is to use finite element models [16, 17]. First, the environmental conditions of the infrastructure, such as concrete saturation degree or temperature, need to be modelled. Then, various constitutive laws need to be defined to translate the mechanical response of the concrete; in particular, specific swelling models, such as Larive's model, must be implemented to track the AAR. Using this methodology, Ben Ftima et al. [17] obtained satisfactory numerical results by validating their model against experimental tests from the literature. Subsequently, both Sellier et al. [16] and Ben Ftima et al. [17] demonstrated the validity of their models by generating structural displacement predictions that align with actual field monitoring measurements.

While finite element models provide a framework to extrapolate laboratory-scale laws to the structural scale, our work focuses on a different approach; given the extensive availability of historical monitoring data on dams, we are interested in data-driven SHM methods that do not rely on physics-based models. By directly exploiting these measurements, we aim to model the AAR phenomena through the observed structural behaviour. The following sections will detail some of the most prevalent methods used in the context of SHM for dams.

2.3 Hydrostatic-Seasonal-Time Based Models

Since its introduction by Willm and Beaujoint [8], the Hydrostatic-Seasonal-Time (HST) model has been widely used for analyzing and interpreting the behaviour of dams in the

context of SHM [18–20]. The expression of that model is

$$\delta = \psi_0 + \underbrace{\sum_{i=1}^{n} \psi_i^h h^i}_{\delta_1(h)} + \underbrace{\sum_{i=1}^{2} \left[\psi_i^c \cos\left(\frac{2i\pi d}{365}\right) + \psi_i^s \sin\left(\frac{2i\pi d}{365}\right) \right]}_{\delta_2(s=\frac{2\pi d}{365})} + \underbrace{\psi_1^t t + \psi_2^t (1 - \exp(-t)) + \psi_3^t \ln(t+1) + \frac{\psi_4^t t}{t+1}}_{\delta_3(t)} + \epsilon,$$

where δ is the total deformation, $\delta_1(h)$ and $\delta_2(s)$ are the reversible deformations due respectively to hydrostatic loads and temperature, and the terms $\psi_0, \psi_i^h, \psi_i^c, \psi_i^s$ and ψ_i^t are regression coefficients. The term $\delta_3(t)$ represents the irreversible deformation caused by phenomena such as AAR or creep, and ϵ is the residual error.

Note that $\delta_3(t)$ can take different forms as pointed out by Liu et al. [18]. Thus this component can be described by a linear combination of the form

$$\delta_3(t) = \sum_{i=1}^{m} \psi_i^t f_i(t),$$

where $f_i(t)$ represents a set of m time-dependent polynomial, rational, logarithmic and/or exponential functions.

The classical HST model presents limitations in its formulation. For example, its thermal component relies exclusively on a seasonal pattern with an annual periodicity, failing to capture actual temperature variations [18, 21]. To address this limitation, Penot et al. [21] proposed the Hydrostatic-Seasonal-Temperature-Time model, which incorporates daily air temperature to take into account temperature variations. Subsequently, Tatin et al. [22] extended this approach by accounting for the temperature of the water and the thermal gradient within the structure. Another alternative is the Hydrostatic-Temperature-Time model developed by Léger and Leclerc [23]. Instead of relying on a seasonal approximation to represent temperature-induced deformations, this model directly uses internal concrete temperatures measured in situ by embedded sensors.

The most fundamental limitation of HST-based models lies in the fact that they rely on a linear regression framework calibrated via the ordinary least squares method [18, 20]. During the training phase, this method identifies a fixed set of coefficients that minimizes the residual error ϵ. In order to make future predictions, the model relies on these coefficients. Because of the lack of a sequential updating mechanism, the parameters cannot be adapted to new scenarios unless a complete recalibration is performed whenever a new observation becomes

available.

2.4 State-Space Models

To overcome the fundamental limitation of the linear regression methods highlighted in Section 2.3, state-space models (SSM) provide a recursive probabilistic framework that enables continuous learning without requiring a complete reprocessing of past data when new observation becomes available [9, 10, 24]. Within this framework, the hidden states $\mathbf{x}_t = [x_1, x_2, \ldots, x_{\mathtt{X}}]_t^{\mathsf{T}}$ represent underlying components (such as a baseline level, trend, or periodic) that allow for the reconstruction of the system's observed responses $\mathbf{y}_t = [y_1, y_2, \ldots, y_{\mathtt{Y}}]_t^{\mathsf{T}}$. To achieve this, the SSM formalism is based on two equations

$$\mathbf{x}_t = \mathbf{f}(\mathbf{x}_{t-1}, \mathbf{w}_t), \tag{2.2}$$

$$\mathbf{y}_t = \mathbf{h}(\mathbf{x}_t, \mathbf{v}_t). \tag{2.3}$$

Equation 2.2 is the transition model. The function $\mathbf{f}(\cdot)$ describes the temporal evolution of the hidden states $\mathbf{x}$ from time $t-1$ to time t. This evolution relies on the Markov hypothesis which assumes that $\mathbf{x}_t$ depends only on $\mathbf{x}_{t-1}$ and not on earlier states. Moreover, because the transition model remains an approximation of the reality, the process errors $\mathbf{w}_t$ are introduced to account for this discrepancy.

Equation 2.3 is the observation model. The function $\mathbf{h}(\cdot)$ links the hidden states $\mathbf{x}_t$ to the observations of the system $\mathbf{y}_t$. The observation errors $\mathbf{v}_t$ account for the fact that sensors do not capture the exact true condition of a system.

Once the SSM framework is established, estimating the hidden states from the imperfect observations can be done using two operations: filtering and smoothing. Filtering refers to the estimation of the state $\mathbf{x}_t$ using only the observations available up to time t. In probabilistic terms, this corresponds to computing the posterior distribution $f(\mathbf{x}_t|\mathbf{y}_{1:t}) \equiv f(\mathbf{x}_t|\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_t)$, where the notation $1:t$ corresponds to the abbreviation of $\{1, 2, \ldots, t\}$. Smoothing retrospectively refines the estimation of $\mathbf{x}_t$ by using all the observations from the entire time series, which is done by computing $f(\mathbf{x}_t|\mathbf{y}_{1:\mathtt{T}})$, where $\mathtt{T}$ is the total number of available measurements. Recursive algorithms, such as the Kalman filter (KF) and the associated Rauch-Tung-Striebel (RTS) Kalman smoother, can be used to perform filtering and smoothing. These algorithms will be introduced in the following section.

2.4.1 Kalman Filter and Smoother

The KF [11] is based on two assumptions [25]:

- Linearity of the transition and observation models of the SSM, presented in equations 2.2 and 2.3.
- The process error $\boldsymbol{W}$ and the observation error $\boldsymbol{V}$ are mutually independent, zero-mean Gaussian that are strictly additive to the transition and observation equations.

Note that this second assumption guarantees that both the prior distribution $f(\mathbf{x}_t|\mathbf{y}_{1:t-1})$ and the likelihood $f(\mathbf{y}_t|\mathbf{x}_t)$ follow a Gaussian distribution. Consequently, by applying Bayes' theorem and the properties of conjugate priors, the resulting posterior distribution $f(\mathbf{x}_t|\mathbf{y}_{1:t})$ is also Gaussian.

Under these assumptions, the transition and observation models are written as

$$\begin{aligned}\mathbf{x}_t &= \mathbf{A}\mathbf{x}_{t-1} + \mathbf{w}_t, \quad \mathbf{w}_t : \boldsymbol{W} \sim \mathcal{N}(\mathbf{w}; \mathbf{0}, \mathbf{Q}),\\ \mathbf{y}_t &= \mathbf{C}\mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{v}_t : \boldsymbol{V} \sim \mathcal{N}(\mathbf{v}; \mathbf{0}, \mathbf{R}),\end{aligned}$$

where $\mathbf{A}$ is the transition matrix, $\mathbf{C}$ is the observation matrix, $\mathbf{Q}$ is the process error covariance matrix, and $\mathbf{R}$ is the observation error covariance matrix.

To compute the posterior distribution mean value $\boldsymbol{\mu}_{t|t} \equiv \mathbb{E}[\mathbf{X_t}|\mathbf{y}_{1:t}]$ and covariance $\boldsymbol{\Sigma}_{t|t} \equiv \text{cov}[\mathbf{X_t}|\mathbf{y}_{1:t}]$, the KF algorithm recursively performs two distinct operations:

- **The prediction step** computes the prior mean value $\boldsymbol{\mu}_{t|t-1}$ and covariance $\boldsymbol{\Sigma}_{t|t-1}$ from the previous posterior knowledge $\boldsymbol{\mu}_{t-1|t-1}$ and $\boldsymbol{\Sigma}_{t-1|t-1}$. This step is formulated as follows

$$\begin{aligned}\boldsymbol{\mu}_{t|t-1} &= \mathbf{A}\boldsymbol{\mu}_{t-1|t-1},\\ \boldsymbol{\Sigma}_{t|t-1} &= \mathbf{A}\boldsymbol{\Sigma}_{t-1|t-1}\mathbf{A}^\mathsf{T} + \mathbf{Q}.\end{aligned}$$

- **The update step** computes the posterior mean value $\boldsymbol{\mu}_{t|t}$ and covariance $\boldsymbol{\Sigma}_{t|t}$ by using

the observation $\mathbf{y}_t$. This step is formulated as follows

$$\begin{aligned}
\hat{\mathbf{y}}_t &= \mathbf{C}\boldsymbol{\mu}_{t|t-1}, \\
\mathbf{r}_t &= \mathbf{y}_t - \hat{\mathbf{y}}_t, \\
\mathbf{G}_t &= \mathbf{C}\boldsymbol{\Sigma}_{t|t-1}\mathbf{C}^\mathsf{T} + \mathbf{R}, \\
\mathbf{K}_t &= \boldsymbol{\Sigma}_{t|t-1}\mathbf{C}^\mathsf{T}\mathbf{G}_t^{-1}, \\
\boldsymbol{\mu}_{t|t} &= \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t r_t, \\
\boldsymbol{\Sigma}_{t|t} &= (\mathbf{I} - \mathbf{K}_t\mathbf{C})\boldsymbol{\Sigma}_{t|t-1}.
\end{aligned}$$

Once the entire sequence of $\mathtt{T}$ measurements has been processed by the KF, the Kalman smoother [26] can be employed to further refine the estimates from the KF filtering process. The Kalman smoother algorithm uses the filtered estimates $\boldsymbol{\mu}_{\mathtt{T}|\mathtt{T}}$ and $\boldsymbol{\Sigma}_{\mathtt{T}|\mathtt{T}}$ from the last time step $\mathtt{T}$ and propagates backward from $t = \mathtt{T}-1$ down to the initial time step. At each step, it corrects the filtered estimates $\boldsymbol{\mu}_{t|t}$ and $\boldsymbol{\Sigma}_{t|t}$ by incorporating the smoothed information from the subsequent time step $t+1$ using

$$\mathbf{J}_t = \boldsymbol{\Sigma}_{t|t}\mathbf{A}^\mathsf{T}\boldsymbol{\Sigma}_{t+1|t}^{-1}. \tag{2.4}$$

$$\boldsymbol{\mu}_{t|\mathtt{T}} = \boldsymbol{\mu}_{t|t} + \mathbf{J}_t(\boldsymbol{\mu}_{t+1|\mathtt{T}} - \boldsymbol{\mu}_{t+1|t}), \tag{2.5}$$

$$\boldsymbol{\Sigma}_{t|\mathtt{T}} = \boldsymbol{\Sigma}_{t|t} + \mathbf{J}_t(\boldsymbol{\Sigma}_{t+1|\mathtt{T}} - \boldsymbol{\Sigma}_{t+1|t})\mathbf{J}_t^\mathsf{T}. \tag{2.6}$$

As established in Sections 2.2 and 2.3, modelling the expansion induced by the AAR requires non-linear functions, specifically exponential components. Consequently, the standard Kalman filter and smoother cannot be applied directly to this problem, as both algorithms rely on the assumption that transition and observation models are linear. The following section reviews extensions to the KF which enable using non-linear transition and observation models.

2.4.2 Non-Linear Extensions of the Kalman Filter

This subsection reviews: the extended Kalman filter (EKF), the unscented Kalman filter (UKF), and the cubature Kalman filter (CKF).

Extended Kalman Filter

The EKF relies on a first-order Taylor series expansion to locally linearize the transition and observation models [24,27]. Specifically, the EKF directly applies the non-linear functions $\mathbf{f}(\cdot)$

and $\mathbf{h}(\cdot)$ to compute respectively the predicted mean $\boldsymbol{\mu}_{t|t-1}$ and the predicted observation $\hat{\mathbf{y}}_t$. Then, instead of using constant linear matrices $\mathbf{A}$ and $\mathbf{C}$, it evaluates the Jacobian matrices of $\mathbf{f}(\cdot)$ and $\mathbf{h}(\cdot)$ at each time step to propagate the predicted covariance and compute the Kalman gain for the update step and thus the posterior mean and covariance.

A limitation of the EKF is its reliance on local linearization for non-linear functions. When the system is highly non-linear, this approximation can lead to the divergence of the estimations made by the filter [28–30].

Sampling-Based: Unscented and Cubature Kalman Filters

To overcome the limitation of the EKF's local linearization, sampling-based approaches were developed. Instead of approximating the non-linear equations themselves, these filters approximate the resulting probability distribution after the non-linear operation. They generate a discrete set of points $\boldsymbol{\mathcal{X}}_i$ with associated weights w_i, which are directly propagated through a generic non-linear function $\mathbf{g}(\cdot)$ to reconstruct the transformed mean and covariance

$$\begin{aligned}\boldsymbol{\mu}_{\boldsymbol{Z}} &= \sum_i w_i \mathbf{g}(\boldsymbol{\mathcal{X}}_i),\\ \boldsymbol{\Sigma}_{\boldsymbol{ZZ}} &= \sum_i w_i (\mathbf{g}(\boldsymbol{\mathcal{X}}_i) - \boldsymbol{\mu}_{\boldsymbol{Z}})(\mathbf{g}(\boldsymbol{\mathcal{X}}_i) - \boldsymbol{\mu}_{\boldsymbol{Z}})^{\mathsf{T}}.\end{aligned}$$

The UKF relies on the unscented transform, which generates $2\mathtt{X}+1$ sigma points around the mean values of the hidden states $\mathbf{x} = [x_1, x_2, \ldots, x_{\mathtt{X}}]^{\mathsf{T}}$ [28]. The points and their associated weights are computed as follows

$$\begin{aligned}\boldsymbol{\mathcal{X}}_0 &= \boldsymbol{\mu}_{\boldsymbol{X}}, & w_0 &= \frac{\kappa_{\mathrm{UKF}}}{\mathtt{X}+\kappa_{\mathrm{UKF}}},\\ \boldsymbol{\mathcal{X}}_i &= \boldsymbol{\mu}_{\boldsymbol{X}} + (\sqrt{(\mathtt{X}+\kappa_{\mathrm{UKF}})\boldsymbol{\Sigma}_{\boldsymbol{XX}}})_i, & w_i &= \frac{1}{2(\mathtt{X}+\kappa_{\mathrm{UKF}})},\\ \boldsymbol{\mathcal{X}}_{\mathtt{X}+i} &= \boldsymbol{\mu}_{\boldsymbol{X}} - (\sqrt{(\mathtt{X}+\kappa_{\mathrm{UKF}})\boldsymbol{\Sigma}_{\boldsymbol{XX}}})_i, & w_{\mathtt{X}+i} &= \frac{1}{2(\mathtt{X}+\kappa_{\mathrm{UKF}})},\end{aligned}$$

where $i \in \{1, 2, \ldots, \mathtt{X}\}$, $\boldsymbol{\mu}_{\boldsymbol{X}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{XX}}$ are the mean and the covariance of $\boldsymbol{X}$, respectively. The term $\left(\sqrt{(\mathtt{X}+\kappa_{\mathrm{UKF}})\boldsymbol{\Sigma}_{\boldsymbol{XX}}}\right)_i$ denotes the i^{th} column of $\sqrt{(\mathtt{X}+\kappa_{\mathrm{UKF}})\boldsymbol{\Sigma}_{\boldsymbol{XX}}}$, and κ_{UKF} is a scaling parameter that controls the spread of the sigma points around $\boldsymbol{\mu}_{\boldsymbol{X}}$. For Gaussian distributions, its optimal value is chosen as $\kappa_{\mathrm{UKF}} = 3 - \mathtt{X}$.

Due to this formulation, the UKF encounters limitations as the dimension $\mathtt{X}$ of the hidden state vector increases. In particular, weights can be negative values that induce numerical errors and prevent the resulting covariance matrix from being guaranteed positive definite [30].

To address these, the scaled unscented Kalman filter was proposed, introducing additional scaling parameters [29, 31]. However, these parameters are heuristically determined [30].

The CKF overcomes the heuristic parameterization and the potential negative weight values of the UKF [30]. The CKF is derived from the spherical-radial cubature rule, an approach that shifts the problem into a spherical coordinate system. Through this transformation, a set of initial points is symmetrically generated with an orthonormal basis at a fixed distance of $\sqrt{\mathtt{X}}$ from the origin. These initial points are then mapped and scaled into the original state-space using the square-root of the covariance matrix. Consequently, the CKF produces a strictly symmetric set of $2\mathtt{X}$ cubature points while omitting the central mean point, unlike the UKF.

For a hidden state vector with mean $\boldsymbol{\mu_X}$ and covariance $\boldsymbol{\Sigma_{XX}}$, the cubature points $\boldsymbol{\mathcal{X}}_i$ are computed as follows [24]

$$\begin{aligned} \boldsymbol{\mathcal{X}}_i &= \boldsymbol{\mu_X} + \sqrt{\mathtt{X}\boldsymbol{\Sigma_{XX}}}\mathbf{e}_i, && \text{for i} \in \{1,2,...,\mathtt{X}\} \\ \boldsymbol{\mathcal{X}}_{i-\mathtt{X}} &= \boldsymbol{\mu_X} - \sqrt{\mathtt{X}\boldsymbol{\Sigma_{XX}}}\mathbf{e}_i, && \text{for i} \in \{\mathtt{X}+1,\mathtt{X}+2,...,2\mathtt{X}\} \end{aligned}$$

where $\mathbf{e}_i$ is the i^{th} vector of the orthonormal basis.

Regarding the weights associated with the cubature points, since the latter have been sampled along the orthogonal axes defining the hypersphere's space, each point contributes equally to the distribution. Thus, unlike the UKF which required calculating distinct and potentially negative weights, the CKF assigns a single and identical positive weight to each of its points, such that for $\forall i \in \{1, 2, \ldots, 2\mathtt{X}\}$ we have:

$$w_i = \frac{1}{2\mathtt{X}}.$$

By ensuring uniformly positive weights, the CKF avoids the numerical instabilities associated with the UKF. However, both of these methods are numerical approximations to evaluate the moments of a distribution. Consequently, they lose the inherent end-to-end analytical tractability of the original KF framework.

2.5 Conclusion

The existing methods reviewed in this chapter present distinct limitations for our study. Approaches that rely on finite element models to extrapolate laboratory-scale constitutive laws, such as Larive's one, are not data-driven. Although HST methods are data-driven, their reliance on linear regression prevents continuous online learning as the data comes in.

To address the need for continuous online learning, SSMs provide the foundation for our framework by enabling the decomposition of time series into interpretable hidden states, which is advantageous for structural health monitoring. However, the filtering methods reviewed in this chapter also present limitations. The standard KF is strictly limited to linear systems. Among non-linear extensions, the EKF can diverge in the presence of highly non-linear behaviours, while the UKF and CKF lose the end-to-end analytical tractability of the standard KF framework.

The objective of this master thesis is to establish an analytical framework capable of capturing exponential behaviours of dams, such as AAR-induced expansion, through continuous online learning from sensor data. To achieve this, the following chapter presents the main contribution of this research: a methodology based on closed-form solutions integrated within the KF framework. This approach allows the moments of the exponential formulation used to model the AAR to be computed in closed form (exactly in some cases, and through an analytically tractable approximation in others) thereby preserving the end-to-end analytical tractability of the standard KF framework.

CHAPTER 3 METHODOLOGY

3.1 Introduction

This section develops an analytically tractable framework for modelling exponential components within state-space models (SSM). First, we establish a generic mathematical foundation for the exponential term that will be denoted as $x_t^{\texttt{EXP}}$. Subsequently, we describe its integration into the SSM framework, specifically tailored to characterize the behaviour of alkali aggregate reaction (AAR).

Note that throughout this chapter, two properties that a derived moment, i.e. mean, variance, or covariance, may possess can be distinguished :

- Exactness. A moment is said to be exact when it is obtained in closed form directly from the true distributional properties of the underlying random variables, without introducing any additional approximation.
- Analytical tractability. A quantity is said to be analytically tractable when it can be expressed through a finite, closed-form algebraic expression

These two properties are not equivalent: a quantity can be analytically tractable without being exact. For each quantity derived in this chapter, we indicate explicitly whether it is exact or constitutes an approximation.

3.2 Modelling the Exponential Component

This section constructs a mathematical and probabilistic representation of the exponential phenomena observed in structural health monitoring (SHM) datasets. To achieve this, the development is divided into two main stages. We start by defining a structural formulation that enables modelling arbitrary shapes of decaying exponential. Afterwards, we bring this expression into a probabilistic framework, allowing for the estimation of its state variables.

3.2.1 Formulation of the Exponential Component

At the core of our modelling approach, we have the fundamental exponential decay term that serves as the basis for the exponential component,

$$x_{t,\texttt{base}}^{\texttt{EXP}} = \exp\left(-\frac{t}{\tau}\right).$$

This formulation allows representing a wide variety of trajectories by adjusting a single parameter, the *time constant* (τ). This parameter acts as the primary controller of the exponential's saturation rate, defining how rapidly the function evolves from its initial state toward its asymptotic limit.

A small τ leads to a rapid saturation, characterizing a phenomenon that reaches its steady state shortly after initiation. Conversely, as τ increases, the exponential decay slows down, allowing the model to capture long-term evolutions that persist over extended periods. As illustrated in Figure 3.1.a, by manipulating this ratio, we can tune the curvature of the component to match the specific kinetic signature of a time series.

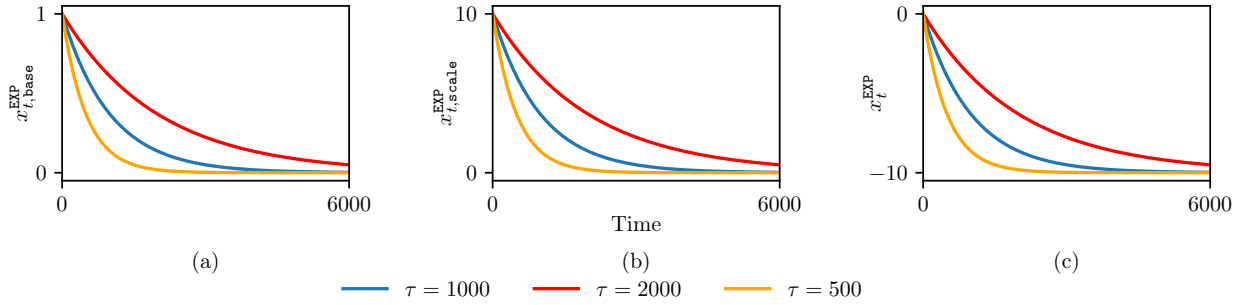

Figure 3.1 Effect of parameters on the exponential formulation. (a) Decay of the exponential term for different time constants τ. (b) Effect of the scaling factor a on the dynamics of the functions ($a = 10$). (c) Introduction of a unit shift applied to the scaled functions.

While the previous formulation captures the desired curvature, it possesses a unit amplitude that ranges from 1 to 0 as t evolves. To adapt this to more general cases, we must introduce a scaling factor a

$$x^{\texttt{EXP}}_{t,\texttt{scale}} = a \cdot \exp\left(-\frac{t}{\tau}\right).$$

However, as illustrated in Figure 3.1.b, applying this factor directly to $\exp(-\frac{t}{\tau})$ implies that the component starts at a non-zero value $x^{\texttt{EXP}}_0 = a$, creating a bias at $t = 0$.

To avoid this initial offset which would adversely interact with the initial states of other components in a global SSM model, we must introduce a unit shift prior to scaling. This ensures that the exponential contribution is null at the origin, allowing the scale factor a to independently control the magnitude of the function without affecting the starting point.

The resulting formulation follows

$$x_t^{\texttt{EXP}} = a \cdot \left(\exp\left(-\frac{t}{\tau} \right) - 1 \right). \tag{3.1}$$

As illustrated in Figure 3.1.c, the final formulation successfully anchors the process at the origin $x_0^{\texttt{EXP}} = 0$. This property is essential for the integration of the exponential component into a SSM, as it ensures a decoupling between the exponential and the other component modelled in time series.

3.2.2 Probabilistic Treatment of the Exponential Component

Equation 3.1 provides a flexible way to model decaying exponential behaviour. However, its practical application to real-world time series presents fundamental challenges because the parameters a and τ are not known and thus cannot be treated as fixed constants. To account for measurement errors and the non-uniqueness of the possible parameter values, we must shift from a deterministic representation to a probabilistic framework. In this perspective, the parameters are redefined as latent random variables characterized by probability distributions rather than single values. This transformation allows the model to represent not only a point estimate for the exponential component but also the degree of uncertainty.

Furthermore, another major constraint associated with Equation 3.1 is its explicit dependence on the time variable t. In order to take full advantage of the dynamic modelling context of SSM, it is required to define the system's evolution based on its internal state rather than an external timestamp. To achieve this, we need to replace the deterministic time t and the constant τ.

To take into account everything we have set out previously, we need to define three latent state variables :

- **The scaling factor** ($x_t^{\texttt{A}}$). In many cases, the final magnitude toward which the exponential converges is unknown. By treating the scaling factor as a latent state, we allow the model to dynamically update its belief about the asymptotic value with every new observation. This variable follows the transition equation

$$x_{t+1}^{\texttt{A}} = x_t^{\texttt{A}}. \tag{3.2}$$

- **The latent time and the latent rate** ($x_t^{\texttt{LTi}}$,$x_t^{\texttt{LR}}$). Instead of using a fixed array for t and scaling it by $\frac{1}{\tau}$, we redefine the temporal progress itself as a local trend component [10]. The latent rate ($x_t^{\texttt{LR}}$) represents the step size, while the latent time

($x_t^{\texttt{LTi}}$) accumulates these steps. This allows the model to explore and learn the time scale of the phenomenon directly from the data, adapting the curvature of the exponential as the kinetic signature becomes clearer. Borrowing from the local trend formulation, these two variables follow the transition model

$$x_{t+1}^{\texttt{LTi}} = x_t^{\texttt{LTi}} + x_t^{\texttt{LR}}, \tag{3.3}$$

$$x_{t+1}^{\texttt{LR}} = x_t^{\texttt{LR}}. \tag{3.4}$$

By substituting these latent variables into Equation 3.1, the exponential component $x_t^{\texttt{EXP}}$ becomes

$$x_t^{\texttt{EXP}} = x_t^{\texttt{A}} \cdot \left(\exp(-x_t^{\texttt{LTi}}) - 1\right), \tag{3.5}$$

where we assume that the three state variables follow Gaussian distributions, defined as

$$\begin{aligned} X_t^{\texttt{A}} &\sim \mathcal{N}(\mu_{X_t^{\texttt{A}}}, \sigma^2_{X_t^{\texttt{A}}}), \\ X_t^{\texttt{LTi}} &\sim \mathcal{N}(\mu_{X_t^{\texttt{LTi}}}, \sigma^2_{X_t^{\texttt{LTi}}}), \\ X_t^{\texttt{LR}} &\sim \mathcal{N}(\mu_{X_t^{\texttt{LR}}}, \sigma^2_{X_t^{\texttt{LR}}}). \end{aligned}$$

Because our objective is to model the exponential component itself, we must propagate the moments of these state variables through the non-linear exponential function and the product between the exponential value and the scaling factor. This process is necessary to determine the resulting probability distribution of $X_t^{\texttt{EXP}}$ along with the joint moments with its latent state variables $\{X_t^{\texttt{A}}, X_t^{\texttt{LTi}}, X_t^{\texttt{LR}}\}$. Note that the transition dynamics of these latent states remain entirely linear: non-linearity is confined to the moment propagation operation, which accounts for the non-linear behaviour of the exponential component. It is this resulting value that is subsequently used in the observation model of the Kalman filter, which allows the standard filtering machinery to remain applicable.

3.3 Moment Propagation

To determine the resulting probability distribution of $X_t^{\texttt{EXP}}$, we must propagate the mean and variance of the latent states through the non-linear formulation. In the following developments, we define three auxiliary variables :

- $z_t^{\texttt{E}} = \exp(-x_t^{\texttt{LTi}}) - 1$ represents the shifted exponential term.

- $z_t^{\mathtt{P}} = x_t^{\mathtt{A}} \cdot z_t^{\mathtt{E}}$ represents the final scaled exponential component ($x_t^{\mathtt{EXP}}$).
- $x_t^{\mathtt{G}}$ represents a generic latent component (such as a linear trend or a seasonal effect) potentially correlated with the exponential states. This generic state variable is added in order to develop the equation required to couple the exponential component with any other existing ones.

3.3.1 Exponential Transformation of the Latent Time

Since $X_t^{\mathtt{LTi}}$ is assumed to follow a Gaussian distribution, $Z_t^{\mathtt{E}}$ follows a log-normal distribution. Using the properties of log-normal moments, the mean ($\mu_{Z_t^{\mathtt{E}}}$), variance ($\sigma^2_{Z_t^{\mathtt{E}}}$), and the covariance between the state and its transformation are derived exactly as follows

$$\mu_{Z_t^{\mathtt{E}}} = \exp\left(-\mu_{X_t^{\mathtt{LTi}}} + \frac{1}{2}\sigma^2_{X_t^{\mathtt{LTi}}}\right) - 1, \tag{3.6}$$

$$\sigma^2_{Z_t^{\mathtt{E}}} = \exp\left(-2\mu_{X_t^{\mathtt{LTi}}} + \sigma^2_{X_t^{\mathtt{LTi}}}\right) \cdot \left(\exp(\sigma^2_{X_t^{\mathtt{LTi}}}) - 1\right), \tag{3.7}$$

$$\mathrm{cov}(X_t^{\mathtt{LTi}}, Z_t^{\mathtt{E}}) = -\sigma^2_{X_t^{\mathtt{LTi}}} \cdot \exp\left(-\mu_{X_t^{\mathtt{LTi}}} + \frac{1}{2}\sigma^2_{X_t^{\mathtt{LTi}}}\right). \tag{3.8}$$

We use a linearization strategy to get a sensitivity factor Γ_t, which represents the ratio between the covariance of the latent time $X_t^{\mathtt{LTi}}$ with its exponential tranformation, and the variance of $X_t^{\mathtt{LTi}}$ itself,

$$\Gamma_t = \frac{\mathrm{cov}(Z_t^{\mathtt{E}}, X_t^{\mathtt{LTi}})}{\mathrm{var}(X_t^{\mathtt{LTi}})}.$$

This linearization provides an analytically tractable, though approximate, estimate of the covariance between $Z_t^{\mathtt{E}}$ and $X_t^{\mathtt{G}}$:

$$\mathrm{cov}(X_t^{\mathtt{G}}, Z_t^{\mathtt{E}}) = \Gamma_t \cdot \mathrm{cov}(X_t^{\mathtt{G}}, X_t^{\mathtt{LTi}}), \tag{3.9}$$

where $\mathrm{cov}(X_t^{\mathtt{G}}, X_t^{\mathtt{LTi}})$ is already known.

3.3.2 Product of the Scaling Factor and the Shifted Exponential Term

The final step involves the product of the scaling factor $x_t^{\mathtt{A}}$ and the auxiliary variable $z_t^{\mathtt{E}}$. To estimate the moments of their product $Z_t^{\mathtt{P}}$, we employ *a Gaussian Multiplicative Approximation* (GMA) [32]. Since $Z_t^{\mathtt{E}}$ follows a log-normal distribution, GMA provides analytically tractable approximations of the expected value, the variance, and the covariance between the

product and a generic component:

$$\mu_{Z_t^{\mathtt{P}}} = \mu_{X_t^{\mathtt{A}}}\mu_{Z_t^{\mathtt{E}}} + \mathrm{cov}(X_t^{\mathtt{A}}, Z_t^{\mathtt{E}}), \tag{3.10}$$

$$\sigma^2_{Z_t^{\mathtt{P}}} = \sigma^2_{X_t^{\mathtt{A}}}\sigma^2_{Z_t^{\mathtt{E}}} + \mathrm{cov}(X_t^{\mathtt{A}}, Z_t^{\mathtt{E}})^2 + 2\mathrm{cov}(X_t^{\mathtt{A}}, Z_t^{\mathtt{E}})\mu_{X_t^{\mathtt{A}}}\mu_{Z_t^{\mathtt{E}}} + \sigma^2_{X_t^{\mathtt{A}}}\mu^2_{Z_t^{\mathtt{E}}} + \sigma^2_{Z_t^{\mathtt{E}}}\mu^2_{X_t^{\mathtt{A}}}, \tag{3.11}$$

$$\mathrm{cov}(X_t^{\mathtt{G}}, Z_t^{\mathtt{P}}) = \mathrm{cov}(X_t^{\mathtt{A}}, X_t^{\mathtt{G}})\mu_{Z_t^{\mathtt{E}}} + \mathrm{cov}(Z_t^{\mathtt{E}}, X_t^{\mathtt{G}})\mu_{x_t^{\mathtt{A}}}. \tag{3.12}$$

3.4 Link between Exponential Component and Alkali-Aggregate Reaction

Before integrating the exponential component's analytical formulations into the SSM framework, it is essential to link the formulation of Equation 3.5 to Larive's model displayed in Equation 2.1. The proposed exponential formulation shares a resemblance with the exponential term found in the numerator of Larive's model, which governs the kinetics of the chemical reaction when the time t reaches large values. This similarity is not coincidental but rather reflects the practical reality of structural health monitoring (SHM) data for aging dams.

Larive's empirical model was designed for material level characterized under laboratory conditions, where the AAR expansion process can be observed in its entirety, from start to finish. In contrast, for the existing infrastructures analyzed in this work, the monitoring data typically begins decades after dams' commissioning. Consequently, the structures monitored have already reached an advanced stage of AAR expansion. Under this assumption, the available data only capture the temporal regime where the expansion kinetics are dominated by the characteristic time τ_c described in Section 2.2. Mathematically, the further the reaction progresses after the inflection point, the less influence the latency time τ_l has. Formally, as t increases beyond τ_l, the exponent $-\frac{t-\tau_l}{\tau_c}$ becomes increasingly negative, so that

$$\lim_{t\to\infty} e^{-\frac{t-\tau_l}{\tau_c}} = 0,$$

and thus the denominator of Equation 2.1 tends toward 1. Once the denominator has converged to 1, Equation 2.1 reduces to $\xi(t) \approx \xi^{\infty}(1 - e^{-\frac{t}{\tau_c}})$, an expression that no longer depends on τ_l. This behaviour justifies reducing the full Larive model to its numerator term.

Furthermore, even if the characteristic time τ_c can be identified from the available data, the latency time τ_l remains unknown. Due to the advanced stage of the reaction captured by the available data, there is no one-to-one relationship between the observed data trajectory and the latency time. Therefore, there are multiple values of τ_l that produce trajectories indistinguishable from one to another over the available data, rendering τ_l non-identifiable from field observations alone. This lack of identifiability is demonstrated in Figure 3.2 where

for the same τ_c value, several τ_l values can be used to fit the data.

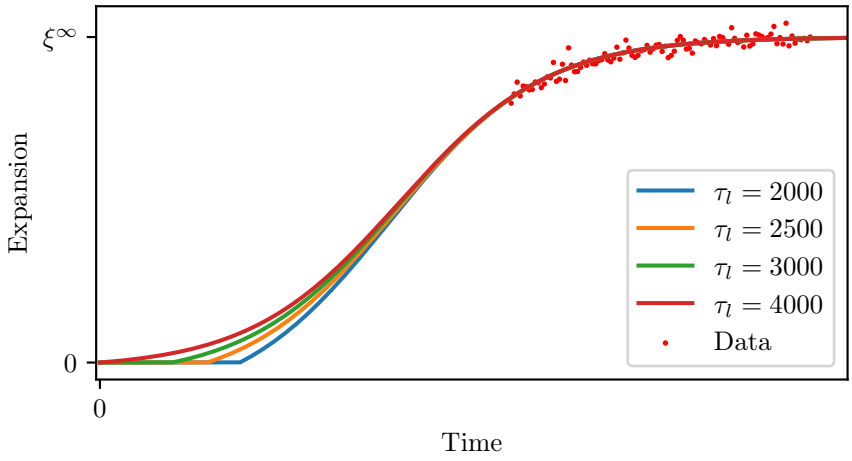

Figure 3.2 Illustration of the non-identifiability of the latency time τ_l in Larive's model. For a fixed characteristic time $\tau_c = 1000$ and ultimate expansion $\xi^\infty = 20$, distinct latency values generate curves that overlap during the advanced phase of the reaction demonstrating the impossibility of retrieving τ_l when early-age monitoring data is unavailable.

These observations clarify how the exponential component relates to Larive's model. Because field monitoring only captures the advanced phase of the AAR reaction, the exponential component retains the information that is identifiable in this regime, namely the characteristic time τ_c and the ultimate expansion ξ^∞. The latency time τ_l is thus left out of the formulation as a consequence of the identifiability limitations demonstrated above. The exponential component is therefore focused on describing AAR expansion using only the information available in field data.

3.5 Integration into the State-Space Model Framework

In Section 3.3, we established the analytical foundations for propagating the expected value, the variance and the covariance through both the exponential function and its product with a scaling factor. The main goal of this section is to exploit those results by integrating them into a SSM framework.

3.5.1 Adapted Kalman Filter Framework for the Analytical Exponential Component

Standard SSMs rely on linear transition equations where the hidden state vector $\mathbf{x}_t$ follows

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{w}_{t+1},\ \mathbf{w}_{t+1} : \boldsymbol{W} \sim \mathcal{N}(\mathbf{w}; \mathbf{0}, \mathbf{Q}). \tag{3.13}$$

In the context of Alkali-Aggregate Reaction (AAR), we assume that the process remains stable in the absence of external intervention. Such a stationary process is assumed to have a zero-valued process-error covariance matrix $\mathbf{Q}$. The linear transition model in Equation 3.13 is incompatible with the non-linear nature of the exponential component. To reconcile the recursive structure of the Kalman Filter (KF) with these non-linearities while maintaining analytical tractability, we propose a partitioned state-space formulation. Instead of the usual update, the state vector and the transition matrix are structured to separate the linear evolution of the latent state variables from the non-linear operations. This approach first propagates the linear dynamics of the latent space and subsequently projects these values using the analytical formulations developed in Section 3.3. We define the state vector $\mathbf{x}_t^{\mathtt{EXP}} = \begin{bmatrix} x_t^{\mathtt{LTi}} & x_t^{\mathtt{LR}} & x_t^{\mathtt{A}} & z_t^{\mathtt{E}} & z_t^{\mathtt{P}} \end{bmatrix}^\mathsf{T}$ for which the mean vector and the covariance matrix are

$$\boldsymbol{\mu}_t^{\mathtt{EXP}} = \begin{bmatrix} \mu_t^{\mathtt{LTi}} & \mu_t^{\mathtt{LR}} & \mu_t^{\mathtt{A}} & \mu_t^{\mathtt{E}} & \mu_t^{\mathtt{P}} \end{bmatrix}^\mathsf{T},$$

$$\boldsymbol{\Sigma}_t^{\mathtt{EXP}} = \begin{bmatrix} (\sigma_t^{\mathtt{LTi}})^2 & \mathrm{cov}(X_t^{\mathtt{LTi}}, X_t^{\mathtt{LR}}) & \mathrm{cov}(X_t^{\mathtt{LTi}}, X_t^{\mathtt{A}}) & \mathrm{cov}(X_t^{\mathtt{LTi}}, Z_t^{\mathtt{E}}) & \mathrm{cov}(X_t^{\mathtt{LTi}}, Z_t^{\mathtt{P}}) \\ \vdots & (\sigma_t^{\mathtt{LR}})^2 & \mathrm{cov}(X_t^{\mathtt{LR}}, X_t^{\mathtt{A}}) & \mathrm{cov}(X_t^{\mathtt{LR}}, Z_t^{\mathtt{E}}) & \mathrm{cov}(X_t^{\mathtt{LR}}, Z_t^{\mathtt{P}}) \\ \vdots & \ddots & (\sigma_t^{\mathtt{A}})^2 & \mathrm{cov}(X_t^{\mathtt{A}}, Z_t^{\mathtt{E}}) & \mathrm{cov}(X_t^{\mathtt{A}}, Z_t^{\mathtt{P}}) \\ \vdots & & \ddots & (\sigma_t^{\mathtt{E}})^2 & \mathrm{cov}(Z_t^{\mathtt{E}}, Z_t^{\mathtt{P}}) \\ \mathrm{sym.} & \dots & \dots & \ddots & (\sigma_t^{\mathtt{P}})^2 \end{bmatrix}.$$

Equation 3.13 holds if $\mathbf{x}_{t+1}$ can be expressed as a linear combination of the elements of $\mathbf{x}_t$. In our case, only the three latent variables $x_t^{\mathtt{A}}$, $x_t^{\mathtt{LTi}}$ and $x_t^{\mathtt{LR}}$ follow this linear dynamics, as established in Section 3.2.2 with Equations 3.2–3.4. In contrast, the auxiliary variables $z_t^{\mathtt{E}}$ and $z_t^{\mathtt{P}}$ cannot follow this transition model. As a result, the prediction step must be done in two steps. First, the linear transition model will be applied only for the latent variables with

the transition matrix $\mathbf{A}^{\texttt{LAT}}$ so that

$$\begin{bmatrix} x_{t+1}^{\texttt{LTi}} \\ x_{t+1}^{\texttt{LR}} \\ x_{t+1}^{\texttt{A}} \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{A}^{\texttt{LAT}}} \underbrace{\begin{bmatrix} x_t^{\texttt{LTi}} \\ x_t^{\texttt{LR}} \\ x_t^{\texttt{A}} \end{bmatrix}}_{\mathbf{x}_t^{\texttt{LAT}}}.$$

However, as we need to apply a transition matrix $\mathbf{A}$ that preserves the dimensions of our mean vector and covariance matrix, we extend $\mathbf{A}^{\texttt{LAT}}$ in such a way that it allows performing the transformations of the latent variables without propagating information to the auxiliary variables,

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}^{\texttt{LAT}} & \mathbf{0}_{3\times 2} \\ \mathbf{0}_{2\times 3} & \mathbf{0}_{2\times 2} \end{bmatrix}.$$

Following the linear prediction step of the Kalman filter,

$$\begin{aligned} \boldsymbol{\mu}_{t+1|t}^{\texttt{EXP}} &= \mathbf{A}\boldsymbol{\mu}_{t|t}^{\texttt{EXP}}, \\ \boldsymbol{\Sigma}_{t+1|t}^{\texttt{EXP}} &= \mathbf{A}\boldsymbol{\Sigma}_{t|t}^{\texttt{EXP}}\mathbf{A}^{\mathsf{T}}, \end{aligned}$$

we will have in our case the temporary predicted mean vector and covariance matrix

$$\boldsymbol{\mu}_{t+1|t}^{\texttt{EXP}} = \begin{bmatrix} \boldsymbol{\mu}_{t+1|t}^{\texttt{LAT}} \\ 0 \\ 0 \end{bmatrix}, \boldsymbol{\Sigma}_{t+1|t}^{\texttt{EXP}} = \begin{bmatrix} \mathbf{A}^{\texttt{LAT}}\boldsymbol{\Sigma}_{t|t}^{\texttt{LAT}}(\mathbf{A}^{\texttt{LAT}})^{\mathsf{T}} & \mathbf{0}_{3\times 2} \\ \mathbf{0}_{2\times 3} & \mathbf{0}_{2\times 2} \end{bmatrix},$$

where the auxiliary $z_t^{\texttt{E}}$ and $z_t^{\texttt{P}}$ are set to zero. To complete the prediction step, we leverage the analytical formulations developed in Section 3.3, Equations 3.6–3.12. Specifically, the prior predictive terms of $\boldsymbol{\mu}_{t+1|t}^{\texttt{EXP}}$ and $\boldsymbol{\Sigma}_{t+1|t}^{\texttt{EXP}}$ are used as inputs for these Equations. Once the prior predictive vector and the covariance matrix have been computed, we now compare the prediction with the observation. In a standard linear-Gaussian SSM framework, the observation model follows

$$y_t = \mathbf{C}\mathbf{x}_t + v_t,\ v_t : V \sim \mathcal{N}(v; 0, \sigma_V^2).$$

In our context, the latent variables governing the exponential behaviour are not directly observable; the only observable quantity is the scaled exponential $z_t^{\texttt{P}}$. Thus, we update the latent variables with the results of the comparison between the observation and the prediction.

We define the observation matrix as $\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \end{bmatrix}$, such that $\hat{y}_t = \mathbf{C}\mathbf{x}_t = z_t^{\mathtt{P}}$. The posterior state is then obtained by using the KF update step presented in Section 2.4. However, as the KF update is a linear operation, the auxiliary states $Z_{t+1|t+1}^{\mathtt{E}}$ and $Z_{t+1|t+1}^{\mathtt{P}}$ are corrected through linear combinations. Consequently, the posterior values in the mean vector and covariance matrix no longer satisfy the inherent non-linear definitions associated with the exponential function and the product operation.

To maintain consistency across the recursive iterations, the updated latent states are propagated again in Equations 3.6–3.12, ensuring that the auxiliary states reflect the non-linear definitions before proceeding to the next iteration. Figure 3.3 illustrates this recursive procedure.

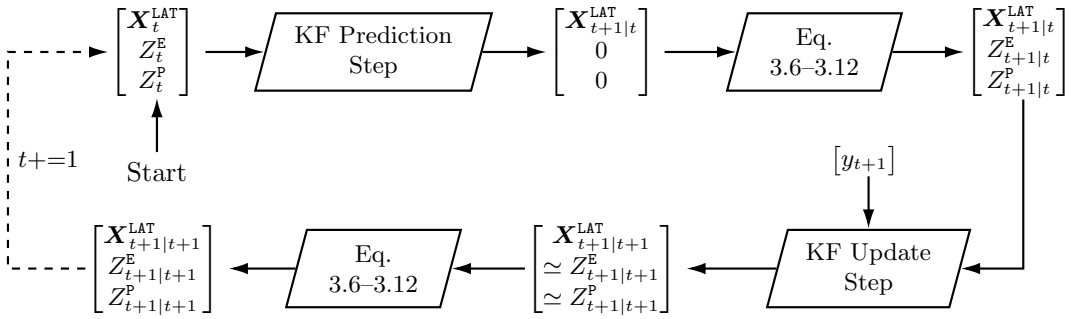

Figure 3.3 Flowchart for the estimation of the hidden states by using the Kalman Filter steps and the Equations 3.6–3.12 to take into account the non-linear definitions of the auxiliary hidden states.

3.5.2 Adapted RTS Kalman Smoother for Analytical Exponential Component

In the case of the exponential component, the transition matrix $\mathbf{A}$ only allows propagating the latent variables, therefore, the terms in the cross-covariance matrix involving auxiliary variables are not computed. To overcome this structural limitation, these non-linear terms need to be explicitly computed. Different cases can be distinguished :

- **Latent-to-latent dependencies**. The terms involving only latent variables are correctly propagated by the initial linear product $\mathbf{\Sigma}_{t|t}\mathbf{A}^{\mathsf{T}}$ and remain unchanged.
- **Cross-covariance involving the shifted exponential** $z^{\mathtt{E}}$. Since latent-to-latent dependencies are accurately characterized — particularly those involving $X_{t|t}^{\mathtt{LTi}}$ and $X_{t+1|t}^{\mathtt{LTi}}$ — and given that the KF filtering process has already provided the quantities $\mathrm{cov}(X_{t|t}^{\mathtt{LTi}}, Z_{t|t}^{\mathtt{E}})$ and $\mathrm{cov}(X_{t+1|t}^{\mathtt{LTi}}, Z_{t+1|t}^{\mathtt{E}})$, Equation 3.9 is used to compute an analytically

tractable approximation for these specific cross-covariance terms. Note that for the term $\text{cov}(Z^{\mathtt{E}}_{t+1|t}, Z^{\mathtt{E}}_{t|t})$, the approximation from Equation 3.9 is applied twice to account for the non-linear projection at both time steps.

- **Cross-covariance involving scaled exponential** $z^{\mathtt{P}}$. Once the dependencies involving $X^{\mathtt{A}}_{t|t}$, $X^{\mathtt{A}}_{t+1|t}$, $Z^{\mathtt{E}}_{t|t}$, and $Z^{\mathtt{E}}_{t+1|t}$ are fully characterized, Equation 3.12 is applied to determine the cross-covariances involving the scaled exponential. Specifically, for the term $\text{cov}(Z^{\mathtt{P}}_{t+1|t}, Z^{\mathtt{P}}_{t|t})$, the application of GMA yields the following analytically tractable approximation

$$\begin{aligned}\text{cov}(Z^{\mathtt{P}}_{t+1|t}, Z^{\mathtt{P}}_{t|t}) =& \text{cov}(X^{\mathtt{A}}_{t+1|t}, X^{\mathtt{A}}_{t|t})\text{cov}(Z^{\mathtt{E}}_{t+1|t}, Z^{\mathtt{E}}_{t|t}) \\ &+ \text{cov}(X^{\mathtt{A}}_{t+1|t}, Z^{\mathtt{E}}_{t|t})\text{cov}(Z^{\mathtt{E}}_{t+1|t}, X^{\mathtt{A}}_{t|t}) + \text{cov}(X^{\mathtt{A}}_{t+1|t}, X^{\mathtt{A}}_{t|t})\mu_{Z^{\mathtt{E}}_{t+1|t}}\mu_{Z^{\mathtt{E}}_{t|t}} \\ &+ \text{cov}(X^{\mathtt{A}}_{t+1|t}, Z^{\mathtt{E}}_{t|t})\mu_{Z^{\mathtt{E}}_{t+1|t}}\mu_{X^{\mathtt{A}}_{t|t}} + \text{cov}(Z^{\mathtt{E}}_{t+1|t}, X^{\mathtt{A}}_{t|t})\mu_{X^{\mathtt{A}}_{t+1|t}}\mu_{Z^{\mathtt{E}}_{t|t}} \\ &+ \text{cov}(Z^{\mathtt{E}}_{t+1|t}, Z^{\mathtt{E}}_{t|t})\mu_{X^{\mathtt{A}}_{t+1|t}}\mu_{X^{\mathtt{A}}_{t|t}}.\end{aligned} \tag{3.14}$$

Once the smoother gain $\mathbf{J}_t$ is fully characterized, the smoothed state vector $\boldsymbol{\mu}_{t|\mathtt{T}}$ and covariance matrix $\boldsymbol{\Sigma}_{t|\mathtt{T}}$ are computed with Equations 2.5 and 2.6. As in the update step from the KF, the smoothed latent states are propagated in the analytical formulations from Equations 3.6–3.12, ensuring again that the auxiliary states reflect the non-linear definitions before proceeding to the next iteration of the smoother. Figure 3.4 illustrates the process needed to do one step of the adapted RTS Kalman Smoother.

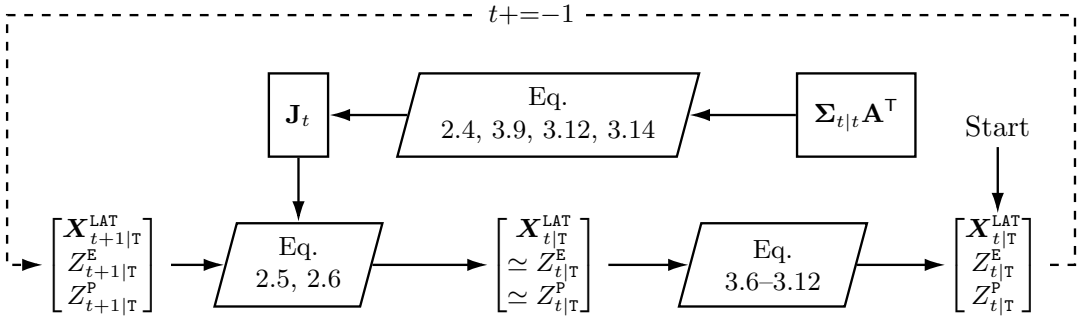

Figure 3.4 Flowchart for the estimation of the smoothed hidden states by using the adapted RTS Kalman smoother for the analytical exponential component. First, Equations 2.4, 3.9, 3.12 and 3.14 are applied to propagate the non-linear terms within the cross-covariance matrix. Then, Equations 2.5 and 2.6 are used to compute the smoothed values of the latent states. Finally, Equations 3.6–3.12 are applied to preserve the non-linear definitions of the auxiliary hidden states.

3.6 Conclusion

In this chapter, we developed a mathematical formulation to describe exponential behaviours, notably by establishing a link with AAR through the Larive model while accounting for the reality of available dam data. We also adapted the filtering and smoothing algorithms to address the non-linearity of the exponential. Through these developments, we established an analytically tractable exponential component and integrated it within a SSM framework. With these foundations now established, the following chapter will focus on the verification and validation of our approach by testing it on both synthetic and real data.

CHAPTER 4 VERIFICATION AND VALIDATION

4.1 Introduction

This chapter evaluates the performance and the robustness of the analytical framework developed in Chapter 3. The evaluation process is structured in three stages. First, the proposed exponential component is integrated into a Kalman filter (KF) and verified on synthetic datasets, where the true parameters and states are known. Second, the predictive capabilities of the model are validated on real monitoring datasets. Finally, the framework is applied to a complex scenario involving a corrective intervention on a dam affected by alkali aggregate reaction (AAR) in order to analyze how the model captures its exponential behaviour and adapts its hidden-states estimations after the intervention.

Note that throughout this chapter, displacement values are reported without an explicit physical unit. For the synthetic datasets, this is inconsequential, as the parameters are defined in a standardized, dimensionless space and could equally represent millimetres, centimetres, or any other unit of length without affecting the verification results. For the real datasets, the omission of physical units reflects a confidentiality requirement associated with the monitoring data; displacement values should be interpreted as being expressed in an arbitrary, but consistent, unit of length.

4.2 Verification on Synthetic Datasets

Before applying the proposed methodology to real structure health monitoring (SHM) datasets, we verify the model on synthetic ones. The latter provide an experimental setup where every parameter is known, allowing for a formal verification of the exponential component within the KF by directly comparing the estimations with the ground truth.

This verification process is structured in two steps. First, the analysis focuses on three synthetic scenarios of increasing complexity. The goal is to confirm the numerical stability of the filter and its ability to correctly retrieve exponential behaviours under different scenarios. For this stage, all components are initialized in a controlled manner representing realistic practical scenarios with a distance between the prior hyperparameters and the ground truth. Second, the study evaluates the robustness of the method across a broad variety of synthetic scenarios. To achieve this, the filter's prior hyperparameters are kept fixed, while the ground truth for hundred cases is generated by randomly drawing parameters from normal distributions describing the hidden-states' prior parameters. This approach allows to assess the

filter's ability to retrieve the true parameters across different synthetic cases.

4.2.1 Verification of the Exponential Component on Three Synthetic Datasets

To evaluate the exponential component, three distinct synthetic datasets are generated. Each dataset simulates 4 years of weekly data with zero-mean observation errors, $v_t : V \sim \mathcal{N}(v; 0, \sigma_V^2)$. These synthetic cases, whose mathematical configurations and true parameters are summarized in Table 4.1, are designed to incrementally replicate the complexity of the time series found in real dam monitoring data:

- Case (i) focuses on the ability of the exponential component to isolate the exponential behaviour from a noisy signal, which resembles a concrete specimen undergoing pure AAR expansion and where the only disturbance is the observation errors.
- Case (ii) adds a periodic component to simulate the real-world behaviour of dams, which are subject to seasonal deformations due to temperature variations. The objective is to verify that the exponential component is not biased by the presence of a cyclic pattern in the time series.
- Case (iii) introduces an initial offset of the signal. This represents a realistic scenario as monitoring time series do not necessarily begin at a zero displacement value. This case tests if the model is able to accurately identify the initial level of the dam without compromising the dynamics of the exponential component.

Table 4.1 Mathematical configurations and parameters for the synthetic datasets.

Case	**Mathematical Model (y_t)**	**True Parameters**
(i)	$a_1(e^{-bt} - 1) + v_t$	$a_1 = 10, b = 0.015, \sigma_V = 0.2$
(ii)	$a_1(e^{-bt} - 1) + a_2 \sin(\omega(t + t_0)) + v_t$	$a_2 = 1.75, \omega = \frac{2\pi}{52}, t_0 = 10$
(iii)	$d_0 + a_1(e^{-bt} - 1) + a_2 \sin(\omega(t + t_0)) + v_t$	$d_0 = 2$

To evaluate the convergence and predictive capabilities of the model, each dataset is split into a training and a test set, comprising respectively 80% and 20% of the total dataset. The training period allows the KF to compare its predictions with the incoming measurements to update its knowledge of the hidden states. Conversely, the test set relies solely on the final states estimated by the filter during the training set to forecast future behaviour. Furthermore, the performance of the KF heavily depends on the prior expected value $\boldsymbol{\mu}_0$ and the prior covariance matrix $\boldsymbol{\Sigma}_0$. To demonstrate if the filter is capable of recovering the true parameters in the synthetic cases, $\boldsymbol{\mu}_0$ is intentionally shifted away from the true values. Subsequently, the prior variance for each parameter is then chosen such that the true value of

each parameter falls within the ±1 prior standard deviation interval. Note that for the white noise component, its standard error is chosen equal to the standard deviation of the observation error v_t, and for the latent time $x^{\texttt{LTi}}$ an initial prior variance of 10^{-10} has been added for numerical stability. The initialization values are summarized in Table 4.2, and the resulting state estimations and forecasting performances for these scenarios are presented in Figure 4.1 and 4.2. Note that for components other than the exponential one, their formulations (hidden state vector, transition matrix, and observation matrix) are available in Appendix B.

Table 4.2 Prior state initialization parameters for the Kalman Filter components. The hidden states $x^{\texttt{LR}}$, $x^{\texttt{A}}$, $x^{\texttt{P}_1}$, and $x^{\texttt{L}}$ correspond to the parameters b, a_1, a_2, and d_0 defined in Table 4.1, respectively.

Component	**Hidden States (x)**	**Prior Mean ($\boldsymbol{\mu}_0$)**	**Prior Variance ($\boldsymbol{\Sigma}_0$)**
Exponential	$[x^{\texttt{LTi}}, x^{\texttt{LR}}, x^{\texttt{A}}, z^{\texttt{E}}, z^{\texttt{P}}]$	$[0, 0.005, 9.5, 0, 0]$	$\text{diag}([(10^{-5})^2, 0.012^2, 0.55^2, 0, 0])$
Periodic	$[x^{\texttt{P}_1}, x^{\texttt{P}_2}]$	$[1.6, 0]$	$\text{diag}([0.3^2, 0.3^2])$
Local Level	$[x^{\texttt{L}}]$	$[1.95]$	$[0.07^2]$

Figure 4.1 verifies the model's ability to recover the exponential behaviour within the datasets. Specifically, the estimation of the scaled exponential component, represented in cyan, successfully isolates the exponential signal from the observation errors and converges to the true values in the dashed lines in all 3 cases. The predictions of the signal, represented in magenta, also match the data during the test phase, which validates the filter's robustness against the superposition of periodic effects and initial offsets in the presence of an exponential trend. The right side of Figure 4.1 shows that each latent state of the exponential component successfully converges towards their true values. Figure 4.2 displays the smoothed version of the results.

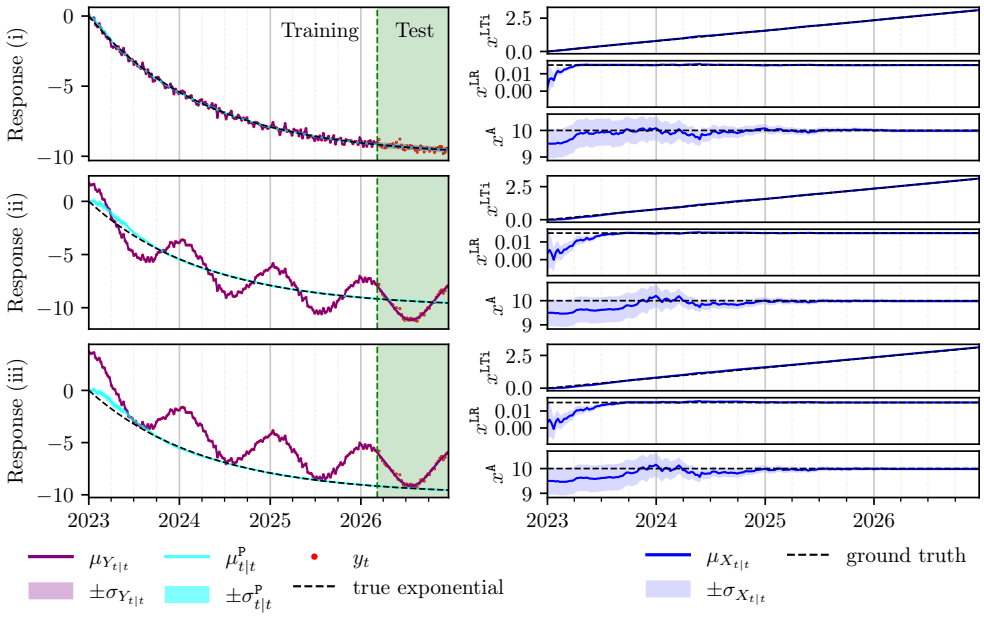

Figure 4.1 KF results on the synthetic data generated with the configurations described in Table 4.1. For each case, the left figure represents the training and testing phases of the KF, and the right figure shows the evolution of the predictions for the exponential latent states compared to the ground truth.

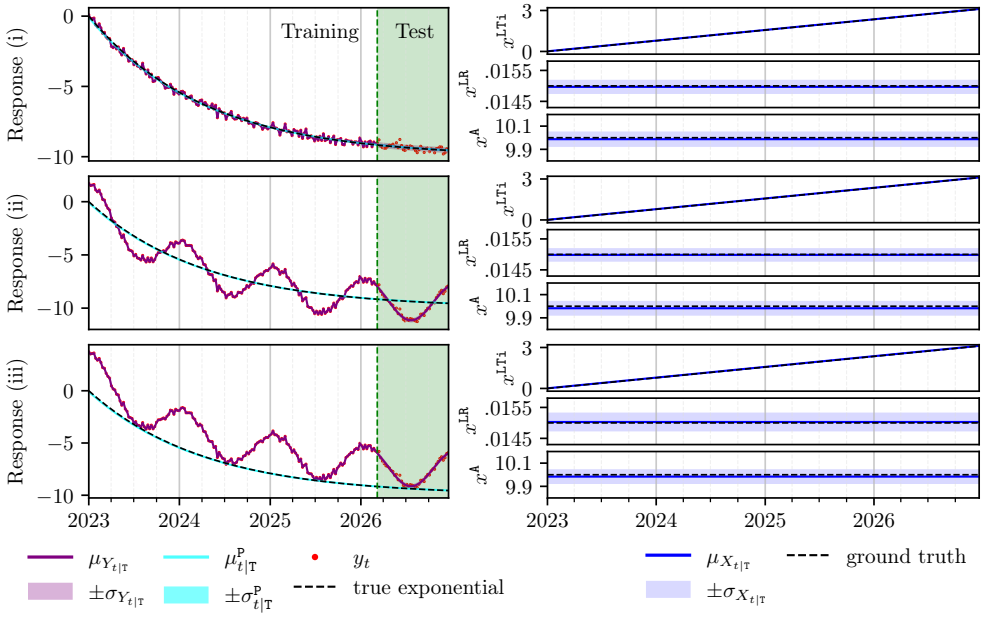

Figure 4.2 RTS smoother results on the synthetic data, showing the retrospective state estimations and reduced uncertainties.

4.2.2 Robustness Analysis With Respect to Prior Initialization

Section 4.2.1 demonstrated the ability of the exponential component to correctly retrieve the true values for each hidden state when initialized with a prior close to the ground truth. However, in real-world applications, the true values for the hidden state of the components are unknown. The objective of this section is to study the robustness of the exponential component when initialized with broad priors.

To achieve this, the analysis builds upon the case (iii) of Section 4.2.1, which includes an initial offset, a periodic component, an exponential behaviour and observation errors. The approach has been adapted so that we first define a prior, i.e., the initial state vector $\boldsymbol{\mu}_0$ and the covariance matrix $\boldsymbol{\Sigma}_0$, for the components. The priors remain constant for the local level and periodic components, while the variances associated with the exponential component specifically for $x^{\mathtt{LR}}$ and $x^{\mathtt{A}}$ are adjusted according to three levels of uncertainty. These levels of uncertainty correspond to coefficient of variation (COV) of 12.5%, 25%, and 50%, set at their respective prior expected values, representing low, medium and high uncertainty scenarios. For each uncertainty level, 100 random samples are drawn from the prior distribution, defined by the expected values and variances previously established, to get the true parameters used to generate the synthetic data following Table 4.1, on which we test our method. Note that for $x^{\mathtt{LR}}$ and $x^{\mathtt{A}}$, truncated normal distributions are used to ensure strictly positive values, as we are specifically interested in simulating the expansion of AAR, the phase t_0 of the periodic is chosen equal to 10 and each dataset simulates 4 years of weekly data. The complete set of hyperparameters used for this procedure is detailed in Table 4.3 and Figure 4.3 illustrates a selection of these synthetic time series.

Table 4.3 Hyperparameters of the normal distributions used for generating synthetic data. The notation follows the mathematical models defined in Table 4.1. The standard deviation for a_1 and b varies according to the three evaluated scenarios, for the other standard deviations, they are set at fixed values.

Description	**Variable**	**Expected Value** μ	**Standard Deviation** σ		
			$12.5\%\mu$	$25\%\mu$	$50\%\mu$
Exponential amplitude	a_1 ($x^{\mathtt{A}}$)	20	2.5	5	10
Exponential latent rate	b ($x^{\mathtt{LR}}$)	0.02	0.0025	0.005	0.01
Periodic amplitude	a_2	2.5		0.3	
Level offset	d_0	1.95		0.1	
Observation errors	v	0.0		0.2	

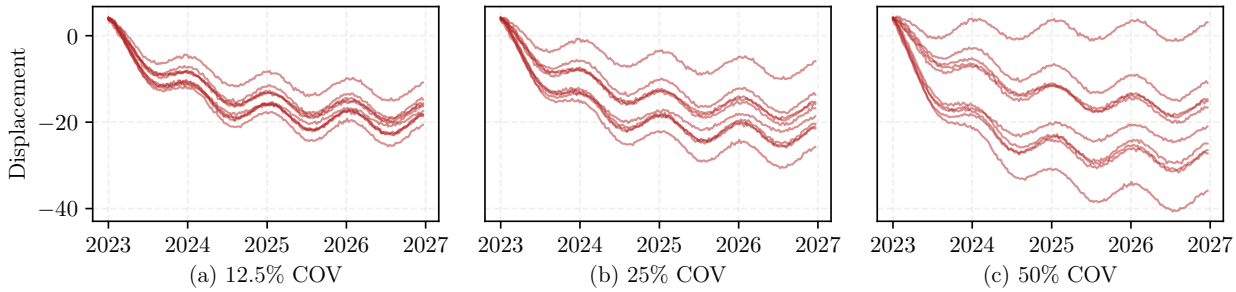

Figure 4.3 Examples of synthetic datasets generated with parameters defined in T able 4.3. The panels display 10 overlaid samples drawn with (a) 12.5%, (b) 25%, and (c) 50% COV on the exponential parameters.

To evaluate the robustness of the exponential component at retrieving true hidden state values, each of the 300 time series generated previously is divided into a training set and a testing set, comprising 80% and 20% of the corresponding dataset. The KF is applied to each case, and its components are initialized according to the expected values and standard deviations defined in Table 4.3. Figure 4.4 displays the scatter plots of $x^{\mathtt{LR}}$ and $x^{\mathtt{A}}$, obtained by the KF at the final time step, against the ground truth of the corresponding case across the three levels of uncertainty for the exponential parameters.

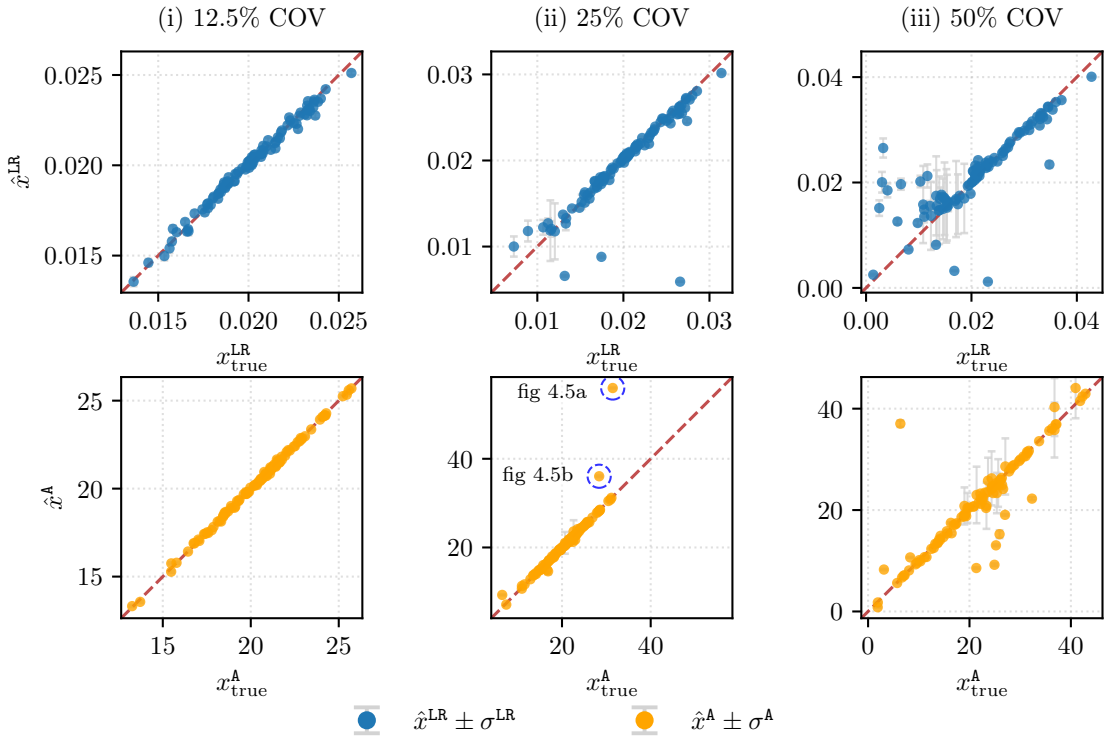

Figure 4.4 Robustness evaluation of the KF against initial uncertainties. The scatter plots compare the ground truth values on the horizontal axis to the final filter estimations on the vertical axis for the parameters x^{LR} on the top row and x^{A} on the bottom row. The grey error bars, where visible, indicate ± 1 standard deviation of the estimation. The columns illustrate the performance degradation when priors are initialized with a standard deviation of (i) 12.5%, (ii) 25%, and (iii) 50% of the expected value, respectively. The red dashed line represents the identity line, the closer a point is to this line, the more accurate the estimation.

For the low uncertainty scenario, the estimated values for x^{A} and x^{LR} by the KF closely match their ground truth, as displayed by their alignment along the diagonal line in Figure 4.4(i). This demonstrates that our prior values effectively constrain the search space, allowing the KF to successfully converge toward the ground truth. In the medium uncertainty scenario, it can be observed in Figure 4.4(ii) that the majority of cases still converge toward the ground truth, although biased estimations begin to appear, especially when x^{LR} is small. However, in these failing cases, the KF's failure is visually detectable by observing its predictions. To illustrate this, Figure 4.5 shows the prediction versus the observations for two cases, labelled (a) and (b), which correspond directly to the two outliers identified in Figure 4.4(ii). Indeed, during the testing phase, the predictions diverge from the observation points. Thus, when the exponential component is initialized with this degree of uncertainty, it is crucial to perform sanity checks on the filter's predictions to assess the quality of the estimated parameters. Observing Figure 4.4(iii), the transition to the high uncertainty scenario leads

to a significant number of cases where the filter fails to recover the ground truth especially when $x^{\mathtt{LR}}$ is small.

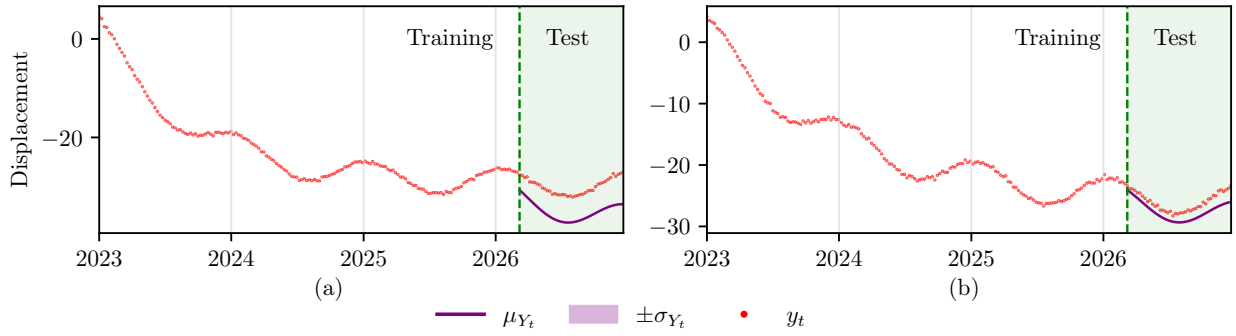

Figure 4.5 Examples of visually detectable convergence failures for the 25% initial uncertainty scenarios presented in Figure 4.4(ii). The purple line represents the mean prediction, surrounded by a shaded area corresponding to ±1 standard deviation, while the red dots indicate the observed data.

In practical applications, we can expect a user to have enough information to initialize the filter within a range of ±25% of the true value. However, in scenarios of complete prior ignorance where one might need to assume a high degree of uncertainty, such as 50% in our case, the filter's convergence can be compromised, making this specific initialization less effective for obtaining reliable estimates. To improve performance and overcome this limitation, an automatic initialization algorithm can be employed. Such an automatic initialization procedure is presented in Appendix C, and the specific hyperparameters used for this test case are detailed in Table D.1 from Appendix D. This algorithm uses the training data to automatically estimate the priors $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$ for the hidden states, providing the filter with an optimized starting point. Figure 4.6 presents the final KF estimations obtained after applying the auto-initialization on the high uncertainty scenario.

In contrast to the results previously observed in Figure 4.4(iii), the majority of the estimations for $x^{\mathtt{LR}}$ and $x^{\mathtt{A}}$ are now gathered closer to the diagonal line. This improvement confirms that the auto-initialization has found better priors, which have led the hidden states to converge toward the correct values. However, despite this overall trend, Figure 4.6 reveals the presence of two outliers for the amplitude parameter $x^{\mathtt{A}}$. These cases differ from the instabilities encountered in Figure 4.4(ii) and Figure 4.4(iii), although the estimated parameters deviate from the ground truth, the model's prediction fit the observed data during the testing phase, making the error visually undetectable during a standard sanity check. Figure 4.7 presents the predictions for these two cases. These two outliers may not be representative of the method's general performance, as their specific parameter combinations create severe identifiability

issues. In the first case shown in Figure 4.7a, the amplitude of the exponential appears to be too small to be distinguished from the other components. In the second case displayed in Figure 4.7b, the low kinetic rate results in an exponential curvature that remains almost linear over the 4-year observation window. Therefore, for cases that display marginal exponential behaviour, the user will need to manually initialize the hidden states in order to remain within a ±25% range of the correct values.

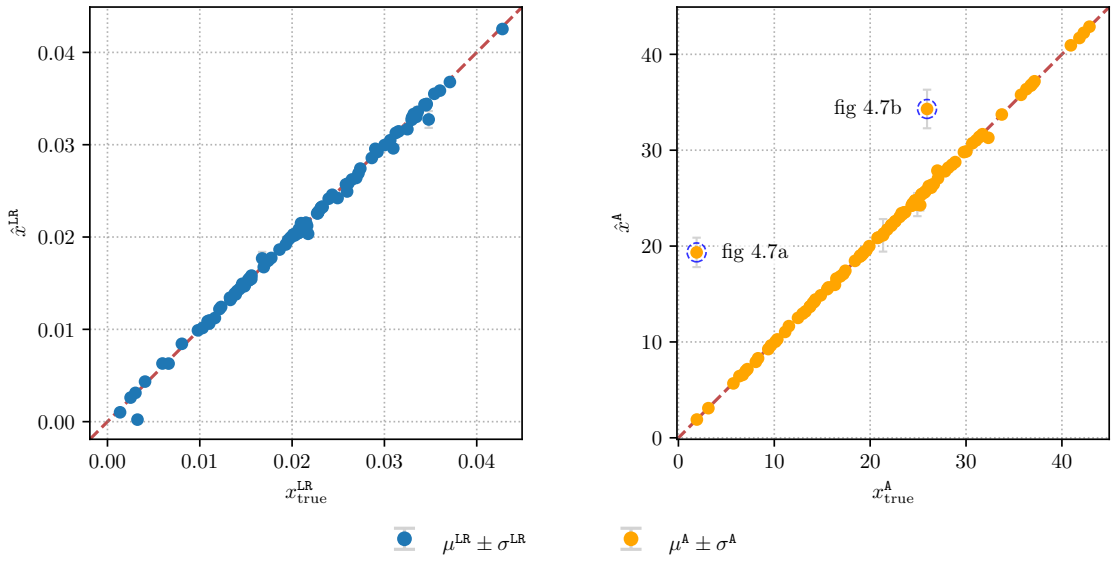

Figure 4.6 Impact of the auto-initialization algorithm on the 50% uncertainty scenario. As in Figure 4.4, the scatter plots compare the ground truth values on the horizontal axis to the final filter estimations on the vertical axis. The grey error bars, where visible, are ±1 standard deviation of the estimation. The red dashed line represents the identity line.

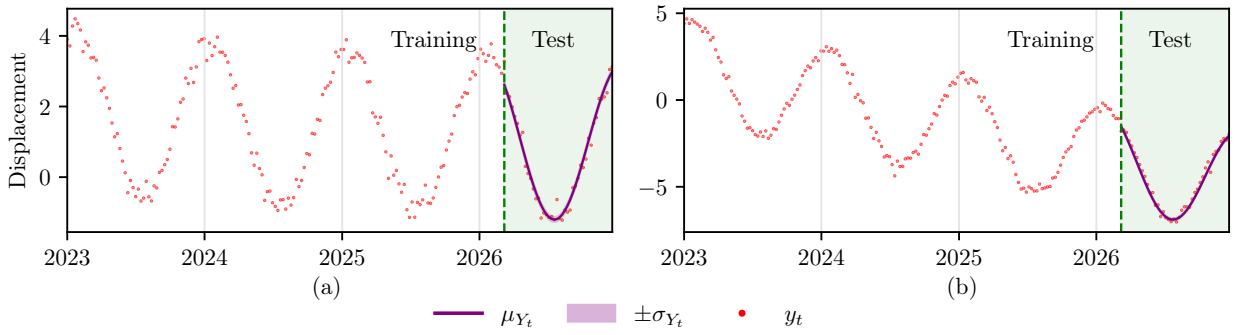

Figure 4.7 Prediction plots for the outliers of the parameter x^{A} displayed in Figure 4.6. The purple line represents the mean prediction, surrounded by a shaded area corresponding to ±1 standard deviation, while the red dots indicate the observed data.

Note that the codes to reproduce all the synthetic cases are available at `https://github.com/MichelWU99/canari-MW-memoire2026` (branch `research/memoire-2026`)

4.3 Validation of the Exponential Component on Real Data

To validate the proposed exponential component, we use monitoring datasets representing the displacements of concrete dams located in the province of Quebec, Canada. Four datasets displaying a non-linear behaviour consistent with an exponential evolution over time were selected through a preliminary visual screening of several hundred candidate monitoring datasets. The remaining datasets are excluded for one of three reasons: some display only a linear trend, some contain too few observations to be usable, and others do not have any discernible pattern.

In the first phase of this validation, the exponential component is evaluated against standard baseline models, i.e., the trend and acceleration components. Subsequently, we study a dataset involving a structural intervention that significantly changes the dynamic of the time series. The objective of this final case is to determine whether the exponential component is capable of quantifying the effect of an intervention on a dam subjected to AAR.

4.3.1 Comparative Analysis on Real Monitoring Datasets

The objective of this comparative analysis is to demonstrate that existing baseline components such as local trend and local acceleration are inadequate for capturing and predicting the long-term AAR behaviour. Consequently, evaluating these baseline components against the proposed exponential one justifies the necessity of the latter.

The comparative analysis is conducted on three datasets collected by sensors measuring the monthly displacement of concrete dams. The data span several decades: from January 1995 to February 2023 for the first dataset, from September 1991 to March 2023 for the second, and from September 1993 to April 2023 for the third. These captured displacement profiles exhibit seasonal patterns and a non-linear behaviour.

A preliminary preprocessing phase is done to transform the raw data into regularized series. First, the raw displacement data contain outliers, identified through visual inspection: these are isolated points whose displacement value deviates sharply from the neighbouring displacement data. Then, since the sensor readings for these three datasets are originally acquired at irregular dates within each month, with some months containing multiple readings and others none, they are temporally regularized onto a strictly regular monthly grid: readings within the same calendar month are aggregated by averaging, and months without any

reading are left as missing observations. This regularization is necessary because the exact acquisition day varies across readings (e.g., a reading on January 31 and one on February 1 would nominally belong to different months despite being only one day apart), which would otherwise result in an irregular time step inconsistent with the constant time-step transition model of the Kalman filter. The missing observations are handled directly within the Kalman filter recursion. For time steps without an available observation, only the prediction step is performed without a corresponding update step. For all the datasets, the last four years data are considered as a test set shown by the green regions in Figure 4.8

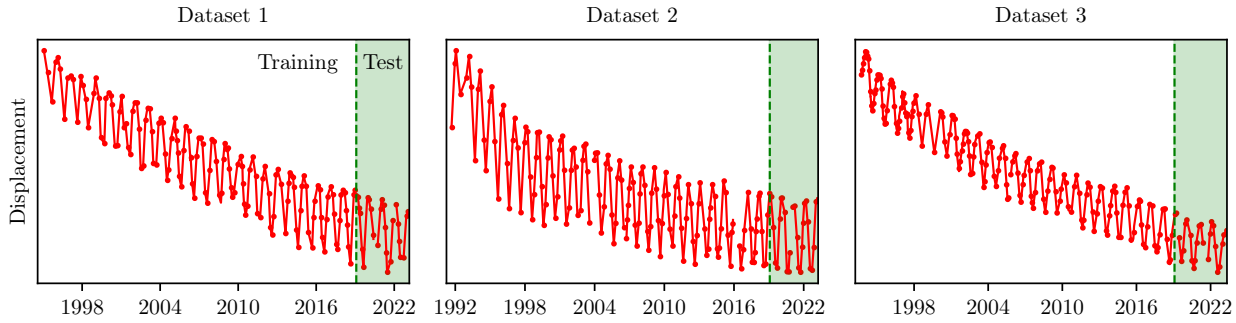

Figure 4.8 Preprocessed cumulative displacements datasets used for the comparative analysis.

With the datasets prepared, three models are built to conduct the comparative analysis. The components used for each architecture are summarized in Table 4.4. As the local trend, local acceleration, periodic, and white noise components are standard SSM formulations documented in the literature [10], and not a contribution of this thesis, their formulations (hidden state vector, transition matrix, and observation matrix) are provided in Appendix B.

Table 4.4 Summary of the components and hidden states used for each state-space model architecture. The checkmarks indicate whether the component is used in the model.

Component	Hidden States	Model Architecture		
		Exponential	Trend	Acceleration
Exponential	$[x^{\mathtt{LTi}}, x^{\mathtt{LR}}, x^{\mathtt{A}}, z^{\mathtt{E}}, z^{\mathtt{P}}]$	✓		
Local Level	$[x^{\mathtt{L}}]$	✓		
Local Trend	$[x^{\mathtt{L}}, x^{\mathtt{LT}}]$		✓	
Local Acceleration	$[x^{\mathtt{L}}, x^{\mathtt{LT}}, x^{\mathtt{LA}}]$			✓
Periodic	$[x^{\mathtt{P1}}, x^{\mathtt{P2}}]$	✓	✓	✓
White noise	$[x^{\mathtt{WN}}, x^{\mathtt{WN\_err}}, x^{\mathtt{W2}}, x^{\overline{\mathtt{W2}}}]$	✓	✓	✓

A notable difference from the verification phase lies in the treatment of the white noise component. In the synthetic scenarios, the exact magnitude of the errors' standard deviation is known. However, for the real monitoring data, this parameter value is uncertain. Therefore, rather than imposing an arbitrary standard deviation, it is treated as an unknown state variable that is estimated through the KF using the AGVI method [33] already implemented in the library canari [34].

The automatic initialization algorithm detailed in Appendix C is applied to the training set of each dataset in order to obtain the prior hidden state values used to initialize the KF. The hyperparameters for configuring the initialization algorithm are summarized in Table D.2 from Appendix D, whereas the prior values for the three datasets and the differents model architectures are presented in Table E.1, E.2 and E.3 respectively from Appendix E. Following this initialization step, we evaluate the performance of the different models by computing the MSE, MAE, and LL on the test sets. The results are summarized in Table 4.5, and the resulting Kalman filter predictions are illustrated in Figure 4.9.

Table 4.5 Test performance metrics for the exponential, trend, and acceleration models. For each model and dataset, the mean squared error (MSE), mean absolute error (MAE), and log-likelihood (LL) are evaluated. The best result for each metric is highlighted in bold. This corresponds to the lowest value for MSE and MAE, and the highest value for LL.

Dataset	Exponential Model			Trend Model			Acceleration Model		
	MSE ↓	MAE ↓	LL ↑	MSE ↓	MAE ↓	LL ↑	MSE ↓	MAE ↓	LL ↑
#1	**0.02**	0.11	10.42	0.14	0.35	-60.94	0.02	**0.11**	**11.22**
#2	**0.15**	**0.29**	**-12.77**	10.09	3.14	-219.43	0.29	0.46	-20.85
#3	0.25	0.38	**-18.18**	5.03	2.18	-132.42	**0.22**	**0.36**	-18.56

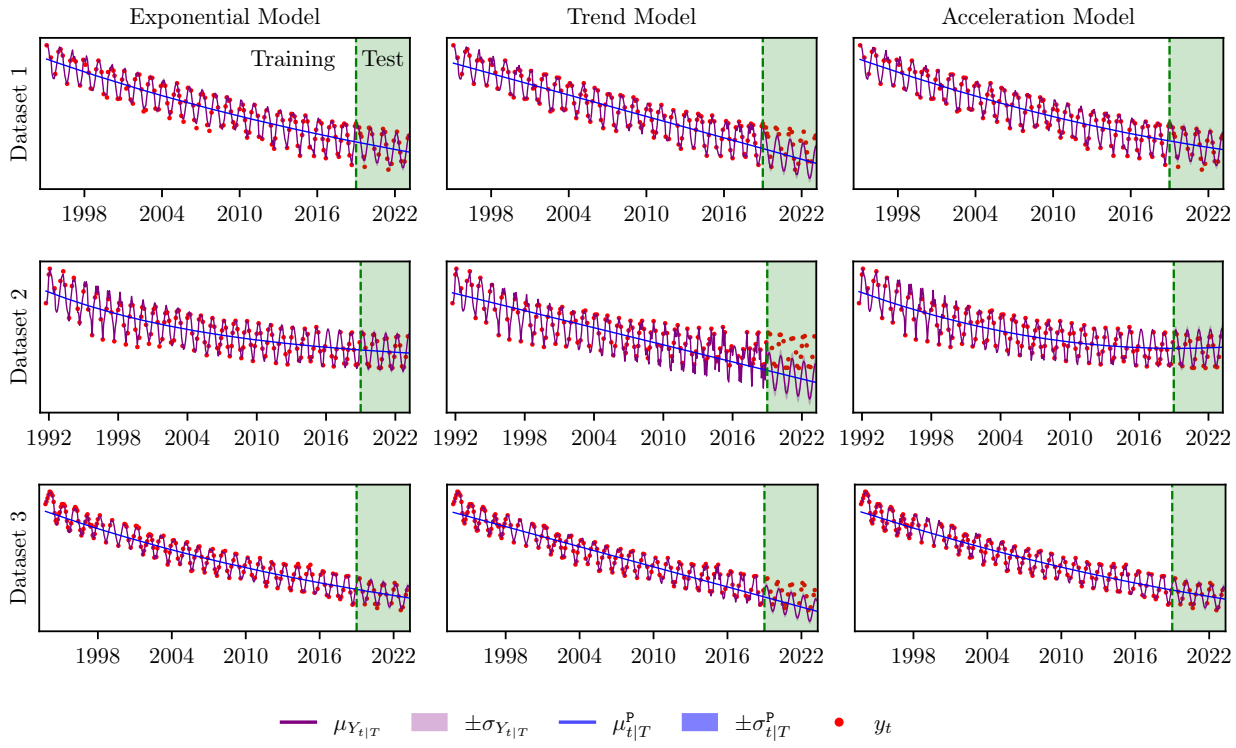

Figure 4.9 Comparative analysis of the exponential, trend, and acceleration models across three datasets. The columns correspond to the respective state-space models, while the rows represent the dataset used.

The results in Table 4.5 reveal that the predictive performance of the trend model is significantly worse across all datasets compared with the exponential and acceleration models. These latter two models consistently yield the best metrics. Furthermore, the performance gaps between them remain marginal, making it difficult to choose between them solely based on their predictive performance over the test sets. The fundamental advantage of using the exponential model becomes more apparent when we forecast further into the future, as displayed in Figure 4.10 for the second dataset. When the predictions are extrapolated beyond the test window, the acceleration model reverses its trajectory. This contradicts the physical reality of AAR expansion where we expect that the concrete structure will keep expanding. Conversely, the long-term predictions of the exponential model align with the expected behaviour, the expansion is continuous with a gradually decreasing rate.

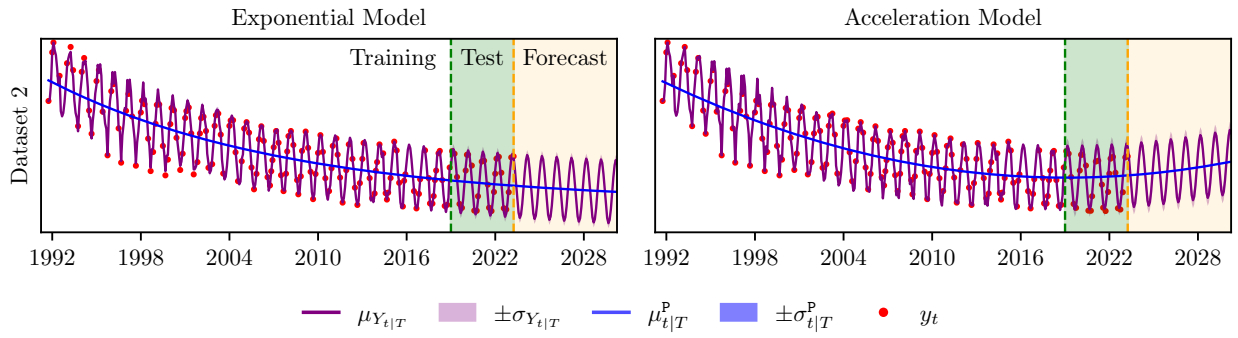

Figure 4.10 Long-term forecasting comparison between the exponential and the acceleration models applied to dataset 2. Compared to Figure 4.9, an orange shaded area has been added to illustrate pure forecast generated by the KF. No observational data is present in this region, as its purpose is solely to visualize the extrapolated long-term forecast of each model.

4.3.2 Adaptability of the Exponential Component to Intervention

In this section, a new monitoring dataset, referred to as dataset 4, is investigated to evaluate the potential use of the exponential component for quantifying the change in exponential decay after an intervention. The data represents the displacement of a dam from December 2002 to August 2014. Unlike the datasets analyzed in the previous section, this time series exhibits a regime change caused by an intervention on the dam, occurring at a known timestamp, denoted as t_{itv}. The time series is temporally regularized by resampling the measurements into weekly intervals. Following the same testing protocol, the dataset, displayed in Figure 4.11, is partitioned to use the last two years of data as test set and the remaining data as a training set.

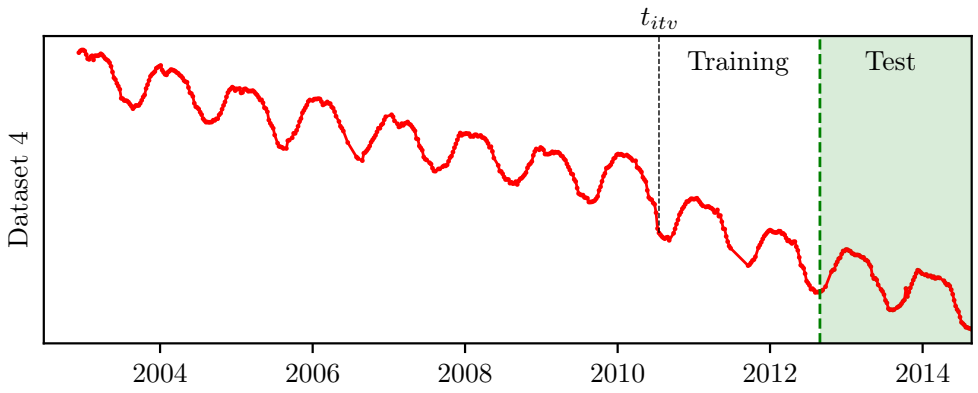

Figure 4.11 Illustration of the regime change in dataset 4. Compared to the previous section, the temporal window has been extended by four years. The black dashed line indicates the time of the intervention, and the red curve represents the observed data.

The objective of this case study is to evaluate the ability of the exponential component to help quantify the change in the exponential decay before and after the intervention. It is hypothesized that the structural intervention acted as a catalyst for the AAR, resulting in a stronger exponential decay. Consequently, the latent states $x^{\mathtt{A}}$ and $x^{\mathtt{LR}}$ should be adjusted to higher values during the post-intervention phase. To test this hypothesis, we propose the following model containing

- Exponential component $[x^{\mathtt{LTi}}, x^{\mathtt{LR}}, x^{\mathtt{A}}, z^{\mathtt{E}}, z^{\mathtt{P}}]$,
- White noise component $[x^{\mathtt{WN}}, x^{\mathtt{WN\_err}}, x^{\mathtt{W2}}, x^{\overline{\mathtt{W2}}}]$,
- Long short-term memory (LSTM) component $[x^{\mathtt{LSTM}}]$,
- Local level 1 component $[x^{\mathtt{L1}}]$,
- Local level 2 component $[x^{\mathtt{L2}}]$.

Note that for components other than the exponential one, their formulations (hidden state vector, transition matrix, and observation matrix) are available in Appendix B.

The use of an LSTM component instead of a standard periodic component allows the model to capture complex seasonal patterns [35]. The role of the local level 1 component $x^{\mathtt{L1}}$ is to track the initial level of the displacement of the structure and to capture the jump caused by the intervention. Therefore, at t_{itv}, a deterministic shift of -1.4 is added to the mean of $X^{\mathtt{L1}}_{t_{itv}|t_{itv}}$. At the same time step, variance injections of $+1^2$ and $+0.0003^2$ are applied to $X^{\mathtt{A}}_{t_{itv}|t_{itv}}$ and $X^{\mathtt{LR}}_{t_{itv}|t_{itv}}$, respectively. Based on our initial hypothesis, after the intervention the latent parameters $x^{\mathtt{A}}$ and $x^{\mathtt{LR}}$ should have higher values. Therefore, when the KF and RTS smoothing algorithms are applied, the model identifies two different regimes,

one pre-intervention and one post-intervention, with different parameter values. Because the state variables adapt to these regimes, their smoothed values will have a sudden jump at t_{itv}. This effect is particularly critical for $x^{\mathtt{A}}$. Because the intervention occurs after a given period of time, the latent time $x^{\mathtt{LTi}}$ is greater than zero. Consequently the shifted exponential term, defined $z^{\mathtt{E}} = \exp(-x^{\mathtt{LTi}}) - 1$ has a non-zero value. Since the final scaled exponential component $z^{\mathtt{P}}$ is the direct product of the amplitude and the shifted exponential, i.e., $z^{\mathtt{P}} = z^{\mathtt{E}} \cdot x^{\mathtt{A}}$, multiplying by $x^{\mathtt{A}}$ propagates its sudden value change to $z^{\mathtt{P}}$, as illustrated by a synthetic case in Figure 4.12. Note that $x^{\mathtt{LR}}$ also undergoes a regime change for the different plots, but it is not displayed since this modification only modifies the subsequent curvature of the exponential and does not induce any jump in the component's value.

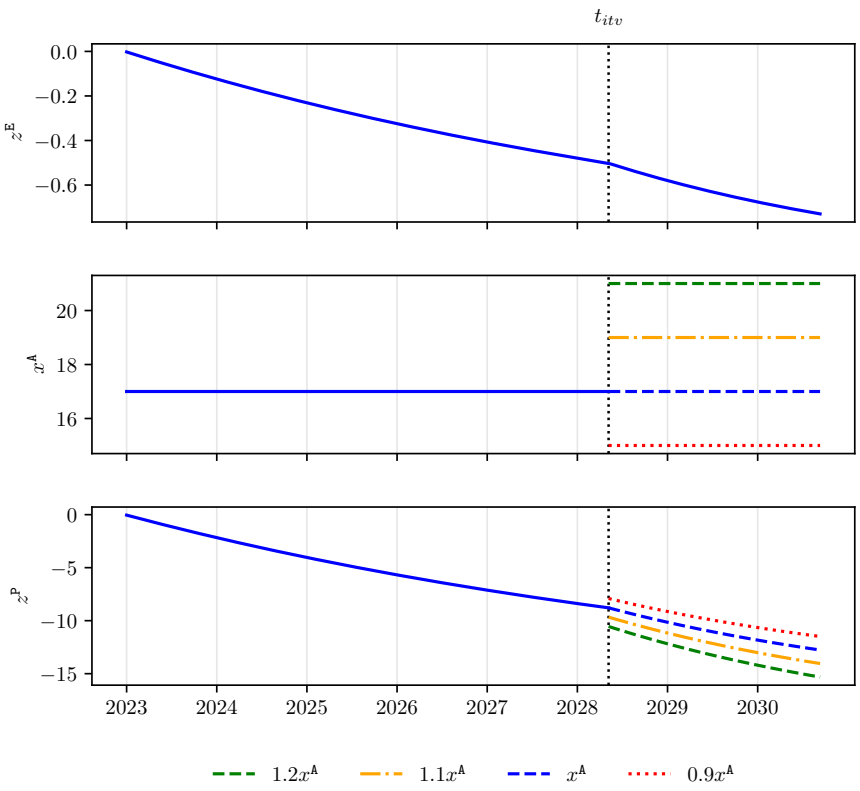

Figure 4.12 Synthetic case of how a sudden value change in $x^{\mathtt{A}}$ propagates to $z^{\mathtt{P}}$ at t_{itv}. The figure displays four scenarios for $x^{\mathtt{A}}$ in the middle subplot that will be applied to calculate $z^{\mathtt{P}} = z^{\mathtt{E}} \cdot x^{\mathtt{A}}$. Since $z^{\mathtt{E}}$ has a non-zero value at t_{itv}, the product directly inherits the discontinuity, resulting in a proportional jump in the last subplot. Note that while $x^{\mathtt{LR}}$ simultaneously undergoes a regime change, reaching nearly twice its value before intervention, this only alters the subsequent slope of the exponential and does not contribute to the instantaneous jump.

If the jump in $z^{\mathtt{P}}$ is left uncorrected, the model will attempt to compensate for this effect

in order to match the observed data. Consequently, $x^{\texttt{L1}}$ would be forced to shift to absorb this artificial jump induced by the exponential. This would bias its value and hence its interpretation. To prevent this interference, the level component $x^{\texttt{L2}}$ is activated at the intervention time step, by a variance injection of $+0.7^2$ to $X^{\texttt{L2}}_{t_{itv}|t_{itv}}$, in order to compensate the jump induced by the exponential component, ensuring that $x^{\texttt{L1}}$ remains accurate. All these explicit adjustments are summarized in Table 4.6.

Table 4.6 Summary of the explicit adjustments applied to the hidden states at the intervention timestamp (t_{itv}).

Hidden state	Parameter Adjusted	Value
$X^{\texttt{A}}_{t_{itv}\|t_{itv}}$	Variance	$+1^2$
$X^{\texttt{LR}}_{t_{itv}\|t_{itv}}$	Variance	$+0.0003^2$
$X^{\texttt{L1}}_{t_{itv}\|t_{itv}}$	Mean	-1.4
$X^{\texttt{L2}}_{t_{itv}\|t_{itv}}$	Variance	$+0.7^2$

To implement this framework, an initialization process is first conducted on the pre-intervention training set, i.e., data from December 2002 to t_{itv}. The Prophet Python library [36] extracts a detrended version of the training data to train the LSTM component, while the automatic initialization algorithm presented in Appendix C is used to establish the prior estimations for the remaining components. The hyperparameters for configuring the automatic initialization algorithm are displayed in Table D.3 from Appendix D, just note that for the local level initialization, the algorithm initializes it to $x^{\texttt{L1}}$, while $x^{\texttt{L2}}$ mean value and variance remain at zero-value. The prior values used to initialize the Kalman filter for this case are presented in Table E.4 from Appendix E.

Following this setup, the KF is applied to the pre-intervention data. Subsequently, at t_{itv} the deterministic shift and variance injections are performed. Thus, we apply the KF on the remaining training data, forecast on the test data and use the RTS smoothing on all data. Figure 4.13 displays the hidden states found and Figure 4.14 the predictions made by the KF.

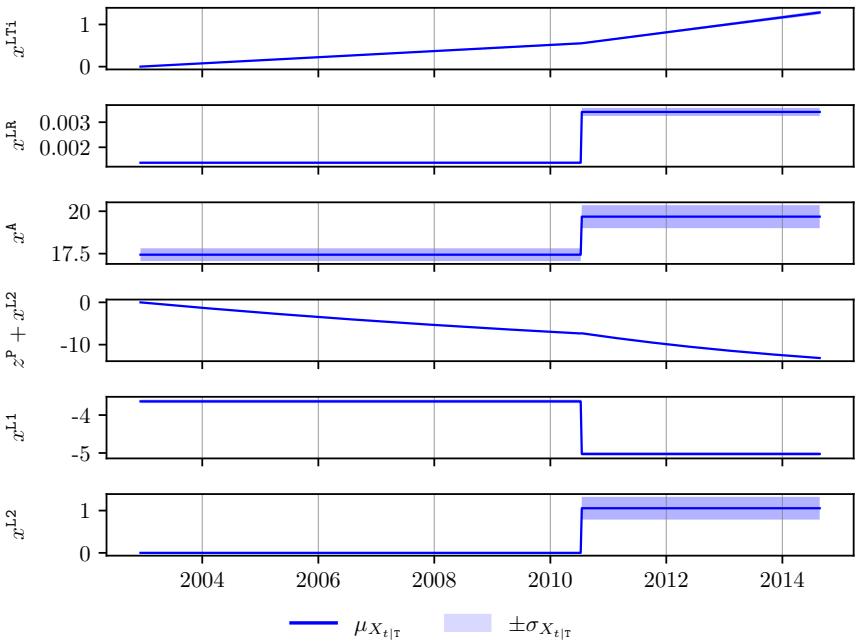

Figure 4.13 RTS smoother evolution of the hidden states for the intervention case. The fourth subplot illustrates the compensated exponential component $z^{\mathrm{P}} + x^{\mathrm{L2}}$ used to absorb the artificial jump.

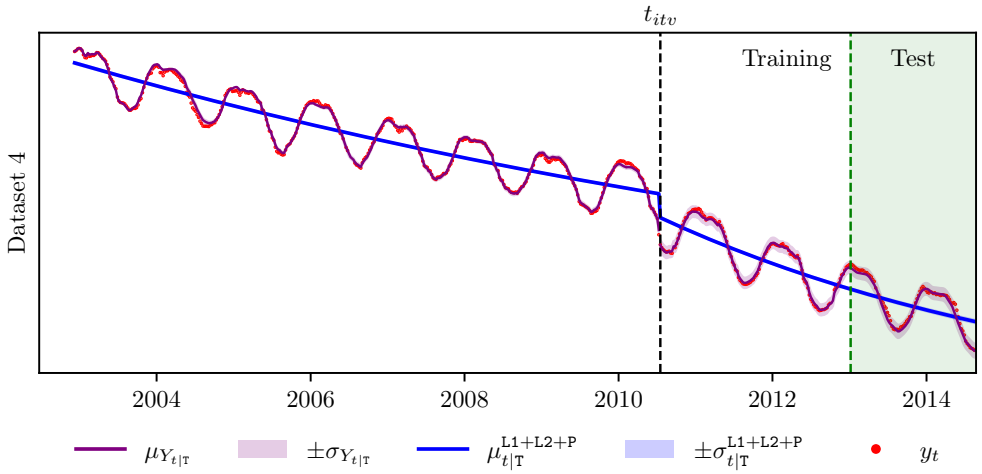

Figure 4.14 Prediction plot for the intervention case. The blue plot represents the trend baseline of the dam $z^{\mathrm{P}} + x^{\mathrm{L1}} + x^{\mathrm{L2}}$, highlighting the effect of the intervention on the dam's exponential behaviour.

The analysis of the hidden states in Figure 4.13 confirms our initial hypothesis regarding the post-intervention evolution of x^{A} and x^{LR}. Furthermore, the first level component x^{L1}

has modelled the physical jump in the data, while the secondary level x^{L2} has absorbed the artificial jump since $z^{\mathsf{P}} + x^{\mathsf{L2}}$ remains continuous. Leveraging these dynamically readjusted parameters, the filter generates predictions that closely match the observations in the test set, as illustrated in Figure 4.14.

However, it is essential to interpret these results with nuance. While the variance injection strategy has been effective for this specific dataset, this demonstration relies on a single case study. Consequently, this approach does not yet allow for generalized conclusions or the establishment of a universal adjustment rule for all infrastructures subjected to similar interventions. Nevertheless, this case study provides evidence that the proposed exponential component can support quantifying the effect of intervention on dams subject to AAR.

4.4 Conclusion

In this chapter, the analytical framework developed in Chapter 3 is verified and validated using synthetic and real data.

The study conducted on synthetic data shows that our methodology is capable of recovering the latent parameters of an exponential component if the latter is initialized within a $\pm 25\%$ coefficient of variation from the true value. Outside this range, a refinement of the initialization hyperparameters is required. This can be done manually or by employing the auto-initialization algorithm presented in Appendix C. However, although this algorithm improves the convergence rate, some cases of non-convergence persist, notably when the exponential behaviour is marginal, making human verification indispensable.

The study on real data demonstrates that the proposed component performs better on data presenting an exponential behaviour than usual baseline components, such as the local trend or local acceleration. However, these results demonstrate the component's effectiveness specifically in cases where an exponential behaviour is already identified through visual inspection. There is no method yet that automatically determines whether, for a new monitoring dataset, applying the exponential component is required. Finally, the analysis also demonstrates that the model is capable of quantifying the modifications in the dynamics of the exponential behaviour following an intervention on the structure.

CHAPTER 5 CONCLUSION

5.1 Thesis Conclusion

In response to the objectives listed in Chapter 1, this thesis has proposed an exponential component whose first two moments are obtained in closed form, exactly through the properties of the log-normal distribution and approximately through the Gaussian multiplicative approximation (GMA). Its cross-covariance terms with other SSM components are obtained in closed form, using an analytical linearization strategy that is also approximate. This component is integrated into the KF and smoother through a partitioned state-space formulation, where closed-form equations are applied after each linear step of both algorithms.

The verification on synthetic data confirmed that the methodology can recover the latent parameters of the exponential component, provided that the initialization remains within a $\pm 25\%$ coefficient of variation of the true values. In practice, users are expected to have sufficient knowledge to satisfy this condition, or can conduct a study to refine their prior knowledge. Nevertheless, manual verifications remain necessary to ensure the reliability of the results, especially when modelling weak exponential phenomena with short time series. The validation on real dam surveillance data, selected for their clearly identifiable exponential trend, demonstrated that the exponential component captures the expected long-term behaviour caused by AAR more effectively than standard baseline components. Finally, a case study demonstrated that, by injecting targeted prior means and variances to the hidden state vector at the time of an intervention, the exponential component can support quantifying the resulting change in the non-linear dynamics and predict the structure's future behaviour.

5.2 Limitations

This section points out the limitations of the proposed methodology.

5.2.1 Scope and Applicability Conditions

First, in the context of this study, the scaling factor $x^{\mathtt{A}}$ is constrained to positive values, as the AAR-induced displacement datasets used in Chapter 4 exhibit a decaying trend. The formulation of the exponential component is not restricted to this case: choosing a negative $x^{\mathtt{A}}$ instead would produce a growing trend toward an upper asymptote. Regardless of the sign

chosen, however, the second derivative of the exponential component, $\frac{d^2 x_t^{\mathtt{EXP}}}{dt^2} = x^{\mathtt{A}}(x^{\mathtt{LR}})^2 e^{-x_t^{\mathtt{LTi}}}$, has the same sign of $x^{\mathtt{A}}$ at all times, since $(x^{\mathtt{LR}})^2$ and $e^{-x_t^{\mathtt{LTi}}}$ are always positive. The component can therefore never exhibit the change in curvature, from an accelerating to a decelerating phase, that characterizes the sigmoidal shape of Larive's model.

Second, and as a direct consequence, the available monitoring data must begin after the inflection point of the reaction has been reached, as discussed in Section 3.4. Note that this condition itself relies on the assumption that AAR-induced expansion at the structural scale follows kinetics analogous to the logistic curve observed at the material scale as described in Larive's model. Since this model was derived from laboratory specimens rather than full structures, this extrapolation has not been independently established and should be treated as a working hypothesis rather than a verified property of structural AAR expansion.

Third, the AAR reaction is assumed to be stationary in the absence of external intervention, which leads to a zero-valued process-error covariance matrix $\mathbf{Q}$. As a consequence, no formulation has been developed to account for potentially non-stationary exponential behaviours, where the parameters of the exponential component would change over time. This restricts the applicability of the current formulation to stationary phenomena.

5.2.2 Lack of Automated Initialization

One of the main limitations of the proposed methodology is the lack of a fully automated robust initialization strategy. Even under realistic initialization conditions, corresponding to a $\pm 25\%$ coefficient of variation of the true values, the filter can converge to wrong values that can only be detected through manual verifications. This limitation is amplified when the exponential behaviour is marginal (i.e., the scaling factor is small or the latent rate is low): the filter may converge toward an incorrect solution whose predictions remain visually consistent with the observations, making the error undetectable.

5.2.3 Absence of an Empirical Comparison with Usual Non-Linear Kalman Filter Extensions

The exclusion of the EKF, UKF, and CKF as candidate frameworks for this thesis is justified on theoretical grounds. The EKF relies on a first-order linearization that can diverge under strong non-linearities. The UKF and CKF, despite being more robust, rely on numerical sampling that breaks the end-to-end analytical tractability that the proposed exponential component is specifically designed to preserve, consistent with the modelling approach of the canari library [34] within which this work is situated. However, these arguments are not

directly validated through an empirical benchmark comparison to measure how EKF-, UKF-, or CKF-based implementations of the exponential component would perform against the proposed methodology, in terms of estimation accuracy, convergence behaviour, or computational cost. Establishing such a benchmark would make it possible to situate the performance of our work relative to these usual non-linear Kalman filter extensions.

5.2.4 Calibration of Errors

The predictive uncertainty presented in this thesis relies on error terms that are estimated directly by the filter, without a dedicated calibration study to verify that the resulting predictive intervals are reliable on new, unseen data. In Section 4.3, the predictive intervals for the real datasets are sometimes narrow and do not always contain the observed test data. This suggests that the estimated uncertainty may not always be reliable, and that a proper calibration study, evaluating and correcting the predictive intervals so that they contain the expected proportion of observations, would be needed to improve their practical accuracy.

5.3 Future Research

This section suggests future research directions to improve the proposed methodology.

5.3.1 Process-Error Covariance for the Exponential Component

To extend the applicability of the exponential component $\mathbf{x}_t^{\texttt{EXP}} = \begin{bmatrix} x_t^{\texttt{LTi}} & x_t^{\texttt{LR}} & x_t^{\texttt{A}} & z_t^{\texttt{E}} & z_t^{\texttt{P}} \end{bmatrix}^\top$ to non-stationary phenomena, it can be noted that its latent variables follow linear dynamics already defined in the SSM literature [10]: $x^{\texttt{LTi}}$ and $x^{\texttt{LR}}$ follow local-trend dynamics, and $x^{\texttt{A}}$ follows local level dynamics. Their respective process-error covariance formulations can be directly adopted to construct a block-structured $\mathbf{Q}^{\texttt{exp}}$ matrix for the exponential component:

$$\mathbf{Q}^{\texttt{exp}} = \begin{bmatrix} \frac{(\sigma_W^{\texttt{LR}})^2}{4} & \frac{(\sigma_W^{\texttt{LR}})^2}{2} & 0 & 0 & 0 \\ \frac{(\sigma_W^{\texttt{LR}})^2}{2} & (\sigma_W^{\texttt{LR}})^2 & 0 & 0 & 0 \\ 0 & 0 & (\sigma_W^{\texttt{A}})^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

where $\sigma_W^{\texttt{LR}}$ and $\sigma_W^{\texttt{A}}$ are the standard deviations of the process error associated with the latent rate $x^{\texttt{LR}}$ and the scaling factor $x^{\texttt{A}}$, respectively.

Since the process-error covariance matrix is applied during the standard prediction step

of the KF, the process noise will be propagated into $z^{\mathtt{E}}$ and $z^{\mathtt{P}}$ through the application of the analytical formulations described in Chapter 3. This constitutes an extension of the framework developed in this master's thesis to account for non-stationary exponential behaviours, although its formal verification remains to be done.

Beyond extending the applicability of the component itself, developing this non-stationary formulation would also align the exponential component with the core purpose of the canari library within which this work is situated, namely the online detection of changepoints in infrastructure monitoring data.

5.3.2 Improved Update Step via Logarithmic Transformation

A potential improvement to the proposed framework concerns the update of the latent variables $x^{\mathtt{LTi}}$ and $x^{\mathtt{LR}}$. In the current formulation, during the standard KF update step, these variables are updated via a linear correction based on their covariance with the observed scaled exponential state, yielding $x^{\mathtt{LTi}}_{t|t}$ and $x^{\mathtt{LR}}_{t|t}$. However, this correction only partially accounts for the non-linear relationship $z^{\mathtt{E}}_{t|t} = \exp(-x^{\mathtt{LTi}}_{t|t}) - 1$ between $x^{\mathtt{LTi}}_{t|t}$ and $z^{\mathtt{E}}_{t|t}$, introducing an approximation.

To address this approximation, an alternative approach would consist of recovering the updated moments of $x^{\mathtt{LTi}}_{t|t}$ by applying the inverse logarithmic transformation directly to the updated $z^{\mathtt{E}}_{t|t}$, i.e., $x^{\mathtt{LTi}}_{t|t} = -\log(z^{\mathtt{E}}_{t|t} + 1)$, rather than relying on the linear KF update step. Note that although $z^{\mathtt{E}}_{t|t}$ is itself approximated by the linear update step, this approach is expected to yield a more accurate estimate of $x^{\mathtt{LTi}}_{t|t}$ by preserving the exact non-linear relationship.

Furthermore, since $x^{\mathtt{LTi}}_{t|t-1} = x^{\mathtt{LTi}}_{t-1|t-1} + x^{\mathtt{LR}}_{t-1|t-1}$, any difference between the predicted $x^{\mathtt{LTi}}_{t|t-1}$ and the updated $x^{\mathtt{LTi}}_{t|t}$ obtained by the inverse transformation carries information about the correction to apply to $x^{\mathtt{LR}}$. However, the formal translation of this difference into an exact update for $x^{\mathtt{LR}}_{t|t}$ remains to be determined.

REFERENCES

[1] M. De Grazia, N. Goshayeshi, R. Gorga, L. Sanchez, A. Santos, and D. Souza, "Comprehensive semi-empirical approach to describe alkali aggregate reaction (AAR) induced expansion in the laboratory," *Journal of Building Engineering*, vol. 40, p. 102298, 2021.

[2] C. Tortajada, "Dams: an essential component of development," *Journal of Hydrologic Engineering*, vol. 20, no. 1, p. A4014005, 2015.

[3] Régie de l'énergie du Canada. (2026) Profil énergétique du canada. [Online]. Available: https://www.cer-rec.gc.ca/fr/donnees-analyse/marches-energetiques/profils-energetiques-provinces-territoires/canada.html

[4] M.-A. Bérubé, B. Durand, D. Vézina, and B. Fournier, "Alkali-aggregate reactivity in Québec (Canada)," *Canadian Journal of Civil Engineering*, vol. 27, no. 2, pp. 226–245, 2000.

[5] T. E. Stanton, "Expansion of Concrete through Reaction between Cement and Aggregate," *Transactions of the American Society of Civil Engineers*, vol. 107, no. 1, pp. 54–84, 1942.

[6] P. Léger, P. Côté, and R. Tinawi, "Finite element analysis of concrete swelling due to alkali-aggregate reactions in dams," *Computers & Structures*, vol. 60, no. 4, pp. 601–611, 1996.

[7] M. Scaioni, M. Marsella, M. Crosetto, V. Tornatore, and J. Wang, "Geodetic and Remote-Sensing Sensors for Dam Deformation Monitoring," *Sensors*, vol. 18, no. 11, p. 3682, Nov. 2018.

[8] G. Willm and N. Beaujoint, "Les méthodes de surveillance des barrages au service de la production hydraulique d'électricité de France, problèmes anciens et solutions nouvelles," in *9th International Congress on Large Dams, ICOLD*, Istanbul, Turkey, 1967, pp. 529–550.

[9] J. Durbin and S. J. Koopman, *Time series analysis by state space methods*. Oxford University Press, 2012.

[10] J.-A. Goulet, *Probabilistic Machine Learning for Civil Engineers*. MIT Press, 2020.

[11] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[12] C. Larive, "Apports combinés de l'expérimentation et de la modélisation à la compréhension de l'alcali-réaction et de ses effets mécaniques," Theses, École Nationale des Ponts et Chaussées, 1997.

[13] A. Zahedi, C. Trottier, L. F. M. Sanchez, and M. Noël, "Condition assessment of alkali-silica reaction affected concrete under various confinement conditions incorporating fine and coarse reactive aggregates," *Cement and Concrete Research*, vol. 153, p. 106694, 2022.

[14] T. N. Nguyen, L. F. M. Sanchez, J. Li, B. Fournier, and V. Sirivivatnanon, "Correlating alkali-silica reaction (ASR) induced expansion from short-term laboratory testings to long-term field performance: A semi-empirical model," *Cement and Concrete Composites*, vol. 134, p. 104817, 2022.

[15] F.-J. Ulm, O. Coussy, L. Kefei, and C. Larive, "Thermo-chemo-mechanics of ASR expansion in concrete structures," *Journal of Engineering Mechanics*, vol. 126, no. 3, pp. 233–242, 2000.

[16] A. Sellier, E. Bourdarot, S. Multon, M. Cyr, and E. Grimal, "Combination of structural monitoring and laboratory tests for assessment of alkali-aggregate reaction swelling: application to gate structure dam," *ACI Materials Journal*, vol. 106, no. 3, 2009.

[17] M. Ben Ftima, P. Bernard, A. Chéruel, and M. Argouges, "Concrete hydraulic structures affected by alkali–aggregate reaction: numerical and experimental investigations," *Journal of Performance of Constructed Facilities*, vol. 37, no. 6, p. 04023059, 2023.

[18] X. Liu, Z. Li, L. Sun, E. Y. Khailah, J. Wang, and W. Lu, "A critical review of statistical model of dam monitoring data," *Journal of Building Engineering*, vol. 80, p. 108106, 2023.

[19] L. Chouinard and V. Roy, "Performance of statistical models for dam monitoring data," in *Joint international conference on computing and decision making in civil and building engineering,*, Montréal, Canada, 2006, pp. 14–16.

[20] A. Simon, E. Buchoud, and J. Guaus, "Prediction of the monitoring measurements: the use of HSTT-models (Hydrostatic-Seasonal-Temperature-Time) to predict dam behaviour during exceptional meteorological events," *IOP Conference Series: Earth and Environmental Science*, vol. 1575, no. 1, p. 012024, 2025.

[21] I. Penot, B. Daumas, and J.-P. Fabre, "Monitoring behaviour," *International water power & dam construction*, vol. 57, no. 12, pp. 24–27, 2005.

[22] M. Tatin, M. Briffaut, F. Dufour, A. Simon, and J.-P. Fabre, "Thermal displacements of concrete dams: Accounting for water temperature in statistical models," *Engineering Structures*, vol. 91, pp. 26–39, 2015.

[23] P. Léger and M. Leclerc, "Hydrostatic, temperature, time-displacement model for concrete dams," *Journal of Engineering Mechanics*, vol. 133, no. 3, pp. 267–277, 2007.

[24] S. Särkkä and L. Svensson, *Bayesian filtering and smoothing*. Cambridge university press, 2023.

[25] H. H. Afshari, S. A. Gadsden, and S. Habibi, "Gaussian filters for parameter and state estimation: A general review of theory and recent trends," *Signal Processing*, vol. 135, pp. 218–238, 2017.

[26] 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.

[27] M. I. Ribeiro, "Kalman and extended Kalman filters: concept, derivation and properties," *Institute for Systems and Robotics*, vol. 43, no. 46, pp. 3736–3741, 2004.

[28] S. Julier, J. Uhlmann, and H. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 477–482, 2000.

[29] E. A. Wan and R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Proceedings of the IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000, pp. 153–158.

[30] I. Arasaratnam and S. Haykin, "Cubature Kalman filters," *IEEE Transactions on automatic control*, vol. 54, no. 6, pp. 1254–1269, 2009.

[31] S. Julier, "The scaled unscented transformation," in *Proceedings of the 2002 American Control Conference*, vol. 6, 2002, pp. 4555–4559.

[32] B. Deka, L. H. Nguyen, S. Amiri, and J.-A. Goulet, "The Gaussian multiplicative approximation for state-space models," *Structural Control & Health Monitoring*, vol. 29, no. 3, p. e2904, 2022.

[33] B. Deka and J.-A. Goulet, "Approximate Gaussian variance inference for state-space models," *International Journal of Adaptive Control and Signal Processing*, vol. 37, no. 11, pp. 2934–2962, 2023.

[34] V.-D. Vuong, Z. Xin, D. Wardan, M. Wu, L.-H. Nguyen, and J.-A. Goulet, "Canari: an open-source library for probabilistic anomaly detection in time series," https://github.com/Bayes-Works/canari, 2025.

[35] V.-D. Vuong, L.-H. Nguyen, and J.-A. Goulet, "Coupling LSTM neural networks and state-space models through analytically tractable inference," *International Journal of Forecasting*, vol. 41, no. 1, pp. 128–140, 2025.

[36] S. J. Taylor and B. Letham, "Forecasting at scale," *The American Statistician*, vol. 72, no. 1, pp. 37–45, 2018.

APPENDIX A MOMENTS OF THE SHIFTED EXPONENTIAL TERM

Consider a random variable X following a normal distribution, $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$. We define the transformation $Z = \exp(-X)$. Since $\ln(Z) = -X$ and $-X \sim \mathcal{N}(-\mu_X, \sigma_X^2)$, the variable Z follows a log-normal distribution with underlying normal parameters $\mu_{\ln Z} = -\mu_X$ and $\sigma_{\ln Z}^2 = \sigma_X^2$. Because the Kalman filter framework works under the Gaussian assumption, we approximate Z by a normal distribution $Z \sim \mathcal{N}(\mu_Z, \sigma_Z^2)$ via moment matching, where μ_Z and σ_Z^2 are obtained from the true moments of the log-normal distribution.

To find the mean μ_Z and variance σ_Z^2, we use the standard relations for a log-normal distribution where the underlying normal parameters are

$$\begin{aligned} -\mu_X &= \ln(\mu_Z) - \frac{\sigma_X^2}{2}, \\ \sigma_X^2 &= \ln\left(1 + \frac{\sigma_Z^2}{\mu_Z^2}\right). \end{aligned}$$

By isolating μ_Z and σ_Z^2, we obtain the system

$$\begin{aligned} \exp\left(-\mu_X + \frac{\sigma_X^2}{2}\right) &= \mu_Z, \\ \left(\exp(\sigma_X^2) - 1\right)\mu_Z^2 &= \sigma_Z^2. \end{aligned}$$

Substituting μ_Z into the second equation yields the explicit form for the variance

$$\begin{aligned} \mu_Z &= \exp\left(-\mu_X + \frac{\sigma_X^2}{2}\right), \\ \sigma_Z^2 &= \exp\left(-2\mu_X + \sigma_X^2\right) \cdot \left(\exp(\sigma_X^2) - 1\right). \end{aligned}$$

By definition, the covariance is given by

$$\operatorname{cov}(X, \exp(-X)) = \mathbb{E}[Xe^{-X}] - \mathbb{E}[X]\mathbb{E}[e^{-X}].$$

To evaluate $\mathbb{E}[Xe^{-X}]$, we use the moment-generating function $M_X(t) = \mathbb{E}[e^{tX}]$ and its deriva-

tive $M'_X(t) = \mathbb{E}[Xe^{tX}]$. For a normal distribution

$$M_X(t) = \exp\left(t\mu_X + \frac{\sigma_X^2 t^2}{2}\right),$$
$$M'_X(t) = (\mu_X + \sigma_X^2 t)\exp\left(t\mu_X + \frac{\sigma_X^2 t^2}{2}\right).$$

Setting $t = -1$, we find

$$\mathbb{E}[Xe^{-X}] = M'_X(-1) = (\mu_X - \sigma_X^2)\exp\left(-\mu_X + \frac{\sigma_X^2}{2}\right).$$

Substituting this into the covariance formula

$$\begin{aligned}\operatorname{cov}(X, e^{-X}) &= (\mu_X - \sigma_X^2)\mu_Z - \mu_X\mu_Z, \\ &= -\sigma_X^2 \exp\left(-\mu_X + \frac{\sigma_X^2}{2}\right).\end{aligned}$$

Finally, for $W = Z - 1$, its moments follow directly from those of Z derived above. Following the same Gaussian approximation via moment matching, W is approximated by a normal distribution $W \sim \mathcal{N}(\mu_W, \sigma_W^2)$, and we have:

$$\begin{aligned}\mu_W &= \exp\left(-\mu_X + \frac{\sigma_X^2}{2}\right) - 1, \\ \sigma_W^2 &= \exp\left(-2\mu_X + \sigma_X^2\right)\left[\exp(\sigma_X^2) - 1\right], \\ \operatorname{cov}(X, W) &= -\sigma_X^2 \exp\left(-\mu_X + \frac{\sigma_X^2}{2}\right).\end{aligned}$$

APPENDIX B FORMULATION FOR OTHER STATE-SPACE MODEL COMPONENTS

This appendix details the formulations of the state vectors and matrices for each state-space model (SSM) component, other than the exponential one, used in this thesis. For each component, we define the hidden state vector $\mathbf{x}$, the transition matrix $\mathbf{A}$, and the observation matrix $\mathbf{C}$. These formulations are derived from [10, 33, 35]

Local Level

$$\begin{aligned}\mathbf{x}^{\mathtt{L}} &= \begin{bmatrix} x^{\mathtt{L}} \end{bmatrix}, \\ \mathbf{A}^{\mathtt{L}} &= \begin{bmatrix} 1 \end{bmatrix}, \\ \mathbf{C}^{\mathtt{L}} &= \begin{bmatrix} 1 \end{bmatrix}.\end{aligned}$$

Local Trend

$$\begin{aligned}\mathbf{x}^{\mathtt{LT}} &= \begin{bmatrix} x^{\mathtt{L}} & x^{\mathtt{LT}} \end{bmatrix}^{\mathsf{T}}, \\ \mathbf{A}^{\mathtt{LT}} &= \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix}, \\ \mathbf{C}^{\mathtt{LT}} &= \begin{bmatrix} 1 & 0 \end{bmatrix}.\end{aligned}$$

Local Acceleration

$$\begin{aligned}\mathbf{x}^{\mathtt{LA}} &= \begin{bmatrix} x^{\mathtt{L}} & x^{\mathtt{LT}} & x^{\mathtt{LA}} \end{bmatrix}^{\mathsf{T}}, \\ \mathbf{A}^{\mathtt{LA}} &= \begin{bmatrix} 1 & \Delta t & \frac{\Delta t^2}{2} \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{bmatrix}, \\ \mathbf{C}^{\mathtt{LA}} &= \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}.\end{aligned}$$

Periodic

$$\begin{aligned}
\mathbf{x}^{\mathtt{P}} &= \begin{bmatrix} x^{\mathtt{P1}} & x^{\mathtt{P2}} \end{bmatrix}^{\mathsf{T}}, \\
\mathbf{A}^{\mathtt{P}} &= \begin{bmatrix} \cos\omega & \sin\omega \\ -\sin\omega & \cos\omega \end{bmatrix}, \\
\mathbf{C}^{\mathtt{P}} &= \begin{bmatrix} 1 & 0 \end{bmatrix},
\end{aligned}$$

where $\omega = \frac{2\pi}{p}$ and p is the period.

White Noise

If the errors' standard deviation is known:

$$\begin{aligned}
\mathbf{x}_t^{\mathtt{WN}} &= \begin{bmatrix} x^{\mathtt{WN}} \end{bmatrix}, \\
\mathbf{A}^{\mathtt{WN}} &= \begin{bmatrix} 0 \end{bmatrix}, \\
\mathbf{C}^{\mathtt{WN}} &= \begin{bmatrix} 1 \end{bmatrix}.
\end{aligned}$$

If the errors' standard deviation is unknown:

$$\begin{aligned}
\mathbf{x}^{\mathtt{WN}} &= \begin{bmatrix} x^{\mathtt{WN}} & x^{\mathtt{WN_err}} & x^{\mathtt{W2}} & x^{\overline{\mathtt{W2}}} \end{bmatrix}^{\mathsf{T}}, \\
\mathbf{A}^{\mathtt{WN}} &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \\
\mathbf{C}^{\mathtt{WN}} &= \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix}.
\end{aligned}$$

LSTM

$$\begin{aligned}
\mathbf{x}^{\mathtt{LSTM}} &= \begin{bmatrix} x^{\mathtt{LSTM}} \end{bmatrix}, \\
\mathbf{A}^{\mathtt{LSTM}} &= \begin{bmatrix} 0 \end{bmatrix}, \\
\mathbf{C}^{\mathtt{LSTM}} &= \begin{bmatrix} 1 \end{bmatrix}.
\end{aligned}$$

APPENDIX C AUTOMATIC INITIALIZATION ALGORITHM

The automatic initialization algorithm aims to find a robust first estimate for the initial state vector $\boldsymbol{\mu}_0$ and the initial covariance matrix $\boldsymbol{\Sigma}_0$ of the Kalman filter. To achieve this, it relies on an optimization process aimed at minimizing an objective function using only the provided training dataset. This initialization covers the following components: local level, local trend, local acceleration, periodic, exponential and white noise.

C.1 Data Splitting and Preparation

The first step of the algorithm is to subdivide the initial train dataset into two distinct sets: a training set and a validation set, controlled by a user-defined ratio. The training set is used to minimize the objective function (see Section C.2), while the validation set is used to select the optimal starting point for the exponential component. Note that if no exponential component is included in the model, this splitting step is not required and the full dataset is used directly for optimization. (see Section C.4)

C.2 Definition of the Objective Function

We define the objective function that allows obtaining the optimal hyperparameters $\boldsymbol{\theta}^*$ to setup $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$. The objective is to fit a parametric function to the provided dataset.

The formulation of $\hat{y}_t(\boldsymbol{\theta})$, the predicted signal at time t given the parameter vector $\boldsymbol{\theta}$, depends on the components constituting the Kalman filter model used. The complete expression of $\hat{y}_t(\boldsymbol{\theta})$ if all components are called (without redundancy of the states of local level, local trend, and local acceleration) would be:

$$\hat{y}_t(\boldsymbol{\theta}) = \alpha + \eta t + \frac{1}{2}\zeta t^2 + \nu \sin\left(\frac{2\pi}{p}(t+\phi)\right) + r(e^{-\lambda t} - 1),$$

where p is the period and the parameters to be optimized are α the level, η the trend, ζ the acceleration, ν the amplitude of the periodic signal, ϕ the phase, r the amplitude of the exponential, and λ the time constant of the exponential. Note that if an LSTM component is used in the Kalman filter model to represent cyclic patterns, the periodic term is replaced by s_t, an estimate of the seasonal effects obtained through decomposition via Prophet [36] (see Section C.3).

The cost function chosen to evaluate this fit is the Mean Squared Error (MSE), expressed as follows:

$$f(\boldsymbol{\theta}) = \frac{1}{\mathtt{M}} \sum_{t=1}^{\mathtt{M}} \left(y_t - \hat{y}_t(\boldsymbol{\theta})\right)^2,$$

where $\mathtt{M}$ is the size of the dataset on which the optimization is done.

This function is minimized using the L-BFGS-B algorithm, which requires starting points and search bounds associated with these points. The following section defines these quantities.

C.3 Warm-up Estimations

The mathematical optimization of a complex function requires a good starting point to prevent the algorithm from converging to an aberrant solution due to a local minimum. The first step thus consists of defining a macro-statistical parameter vector $\tilde{\boldsymbol{\theta}}$ and search bounds $\boldsymbol{\beta}$ using a heuristic approach. To do this, the algorithm analyzes the raw dataset provided to extract component-specific heuristics.

Before estimating the initial parameters for each component, the algorithm checks for the presence of a local level, local trend, or LSTM component. If any of these are detected, the Prophet library [36] is used to decompose the signal y_t into an additive model comprising a baseline trend g_t, a seasonality s_t, and a residual noise ϵ_t such that $y_t = g_t + s_t + \epsilon_t$.

Local Level

The initial level estimate $\tilde{\alpha}$ is extracted directly from the first point of the baseline trend g_t

$$\tilde{\alpha} = g_1.$$

The bounds are defined by a tolerance factor γ^α such that:

$$\boldsymbol{\beta}^\alpha = \left[\tilde{\alpha} - \gamma^\alpha|\tilde{\alpha}|,\ \tilde{\alpha} + \gamma^\alpha|\tilde{\alpha}|\right].$$

Local Trend

The local trend represents the rate of change of a signal. The algorithm exploits the trend $\mathbf{g}_t$ from Prophet to calculate the mean of the first-order finite differences to obtain the initial

trend estimate:

$$\tilde{\eta} = \frac{1}{\mathtt{N}} \sum_{t=1}^{\mathtt{N}} (g_{t+1} - g_t),$$

where $\mathtt{N}$ is the size of the training set.

The bounds are defined by the factor γ^{η} such that:

$$\boldsymbol{\beta}^{\eta} = \left[\tilde{\eta} - \gamma^{\eta}|\tilde{\eta}|,\ \tilde{\eta} + \gamma^{\eta}|\tilde{\eta}|\right].$$

Local Acceleration

The local acceleration describes the curvature of the signal. Since g_t from Prophet is a piecewise approximation, the algorithm favors an estimation using the raw data curvature by computing the mean of the second-order finite differences to obtain the initial acceleration estimate:

$$\tilde{\zeta} = \frac{1}{\mathtt{N}} \sum_{t=1}^{\mathtt{N}} ((y_{t+1} - y_t) - (y_t - y_{t-1})),$$

where $\mathtt{N}$ is the size of the training set. The bounds are defined by the factor γ^{ζ} such that:

$$\boldsymbol{\beta}^{\zeta} = \left[\tilde{\zeta} - \gamma^{\zeta}|\tilde{\zeta}|,\ \tilde{\zeta} + \gamma^{\zeta}|\tilde{\zeta}|\right].$$

Periodic

The signal is first filtered by a moving average centered on a window of size equal to the period p, yielding m_t, an approximation of the baseline and noise. By subtracting this approximation from the raw signal y_t, we isolate the periodic signal $h_t = y_t - m_t$.

We iterate over the signal h_t using a window of size of the period p moving with a global index $j \in [1, \mathtt{N} - p]$, where $\mathtt{N}$ is the size of the training set. For each window starting at j:

- The maximum $h_{\max,j}$ and the minimum $h_{\min,j}$ are stored.
- For each internal sample of the window, identified by a local index $k \in [0, p-1]$ (such that $t = j + k$), we calculate the phase $\phi_{j,k}$:

$$\phi_{j,k} \equiv \left[\frac{1}{\omega} \cdot \operatorname{atan2}\left(h_k \sin(\omega), h_{k+1} - h_k \cos(\omega)\right) - (j+k)\right] \pmod{p},$$

  where $\omega = \frac{2\pi}{p}$.

  Then, we obtain the mean phase ϕ_j of this window by performing a circular mean of

all phases:

$$\phi_j \equiv \left(\frac{1}{\omega} \cdot \operatorname{atan2} \left(\frac{1}{p} \sum_{k=0}^{p-1} \sin\left(\omega \phi_k\right), \frac{1}{p} \sum_{k=0}^{p-1} \cos\left(\omega \phi_k\right) \right) \right) \pmod{p}$$

To obtain the final initial estimate for the entire signal, we compute the amplitude $\tilde{\nu}$ as follows

$$\tilde{\nu} = \frac{1}{2(\mathtt{N}-p)} \sum_{j=1}^{\mathtt{N}-p} \left(h_{\max,j} - h_{\min,j} \right),$$

and for the phase $\tilde{\phi}$:

$$\tilde{\phi} \equiv \left(\frac{1}{\omega} \cdot \operatorname{atan2} \left(\frac{1}{(\mathtt{N}-p)} \sum_{j=1}^{\mathtt{N}-p} \sin(\omega \phi_j), \frac{1}{(\mathtt{N}-p)} \sum_{j=1}^{\mathtt{N}-p} \cos(\omega \phi_j) \right) \right) \pmod{p}.$$

The bounds are defined by

$$\boldsymbol{\beta}^{\nu} = \left[\tilde{\nu} - 0.5,\, \tilde{\nu} + 0.5 \right],$$
$$\boldsymbol{\beta}^{\phi} = \left[-\frac{p}{2}, p + \frac{p}{2} \right].$$

Exponential

For the exponential function, we do not have access to an analytical formulation to find a starting point; therefore, we use a grid search defined by:

- $\mathcal{R} = \{r_1, \ldots, r_{\mathtt{N}_r}\}$: a set of $\mathtt{N}_r$ candidate amplitude points uniformly spaced between r_1 and $r_{\mathtt{N}_r}$, with an associated search bound $\boldsymbol{\beta}^{r} = [r_{\min}, r_{\max}]$.
- $\Lambda = \{\lambda_1, \ldots, \lambda_{\mathtt{N}_\lambda}\}$: a set of $\mathtt{N}_\lambda$ candidate time constants uniformly or logarithmically spaced between λ_1 and $\lambda_{\mathtt{N}_\lambda}$, with an associated search bound $\boldsymbol{\beta}^{\lambda} = [\lambda_{\min}, \lambda_{\max}]$.

C.4 Searching Optimal Hyperparameters

The search for optimal parameters $\boldsymbol{\theta}^*$ depends on the components included in the model. The algorithm distinguishes two distinct strategies.

With an Exponential Component

When the model includes an exponential component, the algorithm must first determine the optimal starting point from the grids defined for this component in the previous section. To

do this, for every possible combination of elements in $\mathcal{R}$ and Λ with the other previously found heuristics, we minimize the cost function on the training set and calculate the validation set MSE using the obtained parameters. The pair yielding the best MSE on the validation set is thus selected as the starting point for the exponential and integrated into the initial search point vector $\tilde{\boldsymbol{\theta}}$.

Finally, we minimize the cost function again with $\tilde{\boldsymbol{\theta}}$ and $\boldsymbol{\beta}$ over the entire dataset (training set + validation set) to obtain $\boldsymbol{\theta}^*$.

Without an Exponential Component

In the absence of an exponential component, we simply minimize the cost function with $\tilde{\boldsymbol{\theta}}$ and $\boldsymbol{\beta}$ over the entire dataset to obtain $\boldsymbol{\theta}^*$.

C.5 Definition of Prior Mean and Variance

We first define the initial state means and variances based on the optimized parameters $\boldsymbol{\theta}^*$

$$\begin{aligned}\mu'_{0_i} &= \theta_i^*,\\ \Sigma'_{0_{i,i}} &= (\kappa_i \theta_i^*)^2,\end{aligned}$$

where κ_i is a relative initial uncertainty coefficient and i is the index of the initiated hidden state.

Note that if the periodic component is used, we must perform the following operations using the optimal parameters ν^* and ϕ^*

$$\begin{aligned}\mu_0^{'\mathtt{P1}} &= \nu^* \sin(\omega\phi^*), & \sigma_0^{'\mathtt{P1}} &= \kappa_{\mathtt{P1}}\nu^*,\\ \mu_0^{'\mathtt{P2}} &= \nu^* \cos(\omega\phi^*), & \sigma_0^{'\mathtt{P2}} &= \kappa_{\mathtt{P2}}\nu^*.\end{aligned}$$

However, our optimized values $\boldsymbol{\theta}^*$ correspond to an estimate of the hidden states for the first observed data point. The structure of the Kalman filter requires $\boldsymbol{\mu}_0$ to be an a priori state, meaning an estimate allowing the filter to make its first prediction such that $\boldsymbol{\mu}_{1|0} = \mathbf{A}\boldsymbol{\mu}_0$. Thus, if we directly use our optimized values as $\boldsymbol{\mu}_0$, we would be one time step ahead. To

correct this, we apply a backward time projection step:

$$\begin{aligned}\boldsymbol{\mu}_0 &= \mathbf{A}^{-1}\boldsymbol{\mu}_0^{'},\\ \boldsymbol{\Sigma}_0 &= \mathbf{A}^{-1}\boldsymbol{\Sigma}_0^{'}(\mathbf{A}^{-1})^T.\end{aligned}$$

Note that to ensure numerical stability, the initial standard deviation $\sigma^{\texttt{LTi}}$ must be defined manually for the exponential component.

Finally, the optimization process provides a global residual mean square error metric, denoted MSE_{opt}, we exploit this information to calibrate the *white noise* component by separating the cases:

- If white noise is used as in the synthetic data case in Section 4.2, then we have:

$$\mathbf{Q}^{\texttt{WN}} = [\text{MSE}_{\text{opt}}]$$

- If white noise is used as in the real data case in Section 4.3, then we define:

$$\mu_0^{\texttt{WN}} = 0, \quad \sigma_0^{\texttt{WN}} = \sqrt{\text{MSE}_{\text{opt}}},$$

$$\mu_0^{\texttt{WN\_err}} = 0, \quad \sigma_0^{\texttt{WN\_err}} = \sqrt{\text{MSE}_{\text{opt}}},$$

$$\mu_0^{\texttt{W2}} = 0, \quad \sigma_0^{\texttt{W2}} = 10^{-3},$$

$$\mu_0^{\overline{\texttt{W2}}} = \text{MSE}_{\text{opt}}, \quad \sigma_0^{\overline{\texttt{W2}}} = \frac{\text{MSE}_{\text{opt}}}{2}.$$

C.6 Pseudo Algorithm of the Automatic Initialization

Require: Dataset y_t, model components, ratio_training
Split y_t into y^{train} and y^{val} using ratio_training (C.1)
if model contains local level, local trend or LSTM **then**
  Decompose y_t via Prophet: $y_t = g_t + s_t + \epsilon_t$
end if
Compute heuristic parameters $\tilde{\boldsymbol{\theta}}$ and bounds $\boldsymbol{\beta}$ for each component (C.3)
if model contains exp **then**
  Build grids $\mathcal{R}$ and Λ with associated bounds $\boldsymbol{\beta}^r = [r_{\min}, r_{\max}]$ and $\boldsymbol{\beta}^\lambda = [\lambda_{\min}, \lambda_{\max}]$
  for each $(r, \lambda) \in \mathcal{R} \times \Lambda$ **do**
    Minimize $f(\boldsymbol{\theta})$ on y^{train} with $\tilde{\boldsymbol{\theta}}$ and $\boldsymbol{\beta}$
    Evaluate MSE on y^{val}
  end for
  Select $(\tilde{r}, \tilde{\lambda})$ yielding lowest validation MSE, update $\tilde{\boldsymbol{\theta}} \leftarrow [\tilde{\boldsymbol{\theta}}, \tilde{r}, \tilde{\lambda}]$ and $\boldsymbol{\beta} \leftarrow [\boldsymbol{\beta}, \boldsymbol{\beta}^r, \boldsymbol{\beta}^\lambda]$
end if
Minimize $f(\boldsymbol{\theta})$ with $\tilde{\boldsymbol{\theta}}$ and $\boldsymbol{\beta}$ over full dataset $\rightarrow \boldsymbol{\theta}^*$ (C.2 and C.4)
Compute $\boldsymbol{\mu}_0'$ and $\boldsymbol{\Sigma}_0'$ from $\boldsymbol{\theta}^*$ and $\boldsymbol{\kappa}$ (C.5)
Apply backward projection: $\boldsymbol{\mu}_0 = \mathbf{A}^{-1}\boldsymbol{\mu}_0', \quad \boldsymbol{\Sigma}_0 = \mathbf{A}^{-1}\boldsymbol{\Sigma}_0'(\mathbf{A}^{-1})^T$
if model contains white noise **then**
  Calibrate white noise states using MSE_{opt} (C.5)
end if
Return $\boldsymbol{\mu}_0$, $\boldsymbol{\Sigma}_0$

APPENDIX D HYPERPARAMETERS FOR AUTOMATIC INITIALIZATION ALGORITHM

This appendix provides a comprehensive summary of the hyperparameter values to setup for the automatic initialization algorithm.

Table D.1 Summary of automatic initialization algorithm hyperparameters to reproduce the Figure 4.6 in Section 4.2.2

Component	**Hyperparameter**	**Value**
	ratio training	0.8
Local Level	γ^{α}	0.05
	κ_{α}	0.15
Periodic	$\kappa_{\texttt{P1}}, \kappa_{\texttt{P2}}$	0.075
Exponential	$\texttt{N}_r$	15
	r_1	5
	$r_{\texttt{N}_r}$	35
	$[r_{\min}, r_{\max}]$	$[1, 50]$
	$\texttt{N}_{\lambda}$	15 (uniformly spaced)
	λ_1	0.05
	$\lambda_{\texttt{N}_{\lambda}}$	0.035
	$[\lambda_{\min}, \lambda_{\max}]$	$[-0.1, 0.1]$
	$\sigma_r = \kappa_r \theta_r^*$	$\sqrt{5}$
	$\sigma_{\lambda} = \kappa_{\lambda} \theta_{\lambda}^*$	0.001
	$\sigma^{\texttt{LTi}}$	0.001

Note for this case, for the exponential component the values of $\sigma_r = \kappa_r \theta_r^*$ and $\sigma_{\lambda} = \kappa_{\lambda} \theta_{\lambda}^*$ are set instead of just setting κ_r and κ_{λ}

Table D.2 Summary of automatic initialization algorithm hyperparameters for Section 4.3.1

Component	**Hyperparameter**	**Value**
	ratio training	0.6
Local Level	γ^{α}	0.10
	κ_{α}	0.15
Local Trend	γ^{η}	2.00
	κ_{η}	0.45
Local Acceleration	γ^{ζ}	2.00
	κ_{ζ}	0.45
Periodic	$\kappa_{\texttt{P1}}, \kappa_{\texttt{P2}}$	0.075
Exponential	$\texttt{N}_r$	10
	r_1	1
	$r_{\texttt{N}_r}$	35
	$[r_{\min}, r_{\max}]$	$[1, 50]$
	$\texttt{N}_{\lambda}$	6 (logarithmically spaced)
	λ_1	10^{-6}
	$\lambda_{\texttt{N}_{\lambda}}$	10^{-1}
	$[\lambda_{\min}, \lambda_{\max}]$	$[10^{-6}, 1]$
	κ_r	0.40
	κ_{λ}	0.45
	$\sigma^{\texttt{LTi}}$	10^{-10}

Table D.3 Summary of automatic initialization algorithm hyperparameters for Section 4.3.2

Component	**Hyperparameter**	**Value**
	ratio training	0.5
Local Level	γ^{α}	0.10
	κ_{α}	0.15
Local Trend	γ^{η}	2.00
	κ_{η}	0.45
Local Acceleration	γ^{ζ}	2.00
	κ_{ζ}	0.45
Periodic	$\kappa_{\texttt{P1}}, \kappa_{\texttt{P2}}$	0.075
Exponential	$\texttt{N}_r$	10
	r_1	1
	$r_{\texttt{N}_r}$	35
	$[r_{\min}, r_{\max}]$	$[1, 50]$
	$\texttt{N}_{\lambda}$	6 (logarithmically spaced)
	λ_1	10^{-6}
	$\lambda_{\texttt{N}_{\lambda}}$	10^{-1}
	$[\lambda_{\min}, \lambda_{\max}]$	$[10^{-6}, 1]$
	κ_r	0.40
	κ_{λ}	0.45
	$\sigma^{\texttt{LTi}}$	10^{-10}

APPENDIX E PRIOR VALUES FOR KALMAN FILTER INITIALIZATION

This appendix presents the prior state estimates (μ_0) and initial variances (σ_0^2) used to initialize the Kalman filter.

Table E.1 Prior mean (μ_0) and variance (σ_0^2) values used to initialize the Exponential model in Section 4.3.1

Component	Hidden State	Dataset 1		Dataset 2		Dataset 3	
		μ_0	σ_0^2	μ_0	σ_0^2	μ_0	σ_0^2
Exponential	$x^{\texttt{LTi}}$	$-1.44 \cdot 10^{-3}$	10^{-10}	$-9.25 \cdot 10^{-3}$	10^{-10}	$-4.15 \cdot 10^{-3}$	10^{-10}
	$x^{\texttt{LR}}$	$1.44 \cdot 10^{-3}$	$4.19 \cdot 10^{-7}$	$9.25 \cdot 10^{-3}$	$1.73 \cdot 10^{-5}$	$4.15 \cdot 10^{-3}$	$3.49 \cdot 10^{-6}$
	$x^{\texttt{A}}$	8.56	11.71	8.11	10.52	27.44	$1.20e^2$
	$z^{\texttt{E}}$	0	0	0	0	0	0
	$z^{\texttt{P}}$	0	0	0	0	0	0
Local Level	$x^{\texttt{L}}$	-2.90	0.19	1.4	$7.84 \cdot 10^{-4}$	-1.75	$1.23 \cdot 10^{-3}$
Periodic	$x^{\texttt{P1}}$	0.49	$2.33 \cdot 10^{-3}$	-2.33	$3.07 \cdot 10^{-2}$	-2.81	$4.43 \cdot 10^{-2}$
	$x^{\texttt{P2}}$	0.46	$2.33 \cdot 10^{-3}$	-0.16	$3.07 \cdot 10^{-2}$	$-4.40 \cdot 10^{-2}$	$4.43 \cdot 10^{-2}$
White Noise	$x^{\texttt{WN}}$	0	$2.57 \cdot 10^{-2}$	0	0.20	0	0.60
	$x^{\texttt{WN\_err}}$	0	$2.57 \cdot 10^{-2}$	0	0.20	0	0.60
	$x^{\texttt{W2}}$	0	10^{-6}	0	10^{-6}	0	10^{-6}
	$x^{\overline{\texttt{W2}}}$	$2.57 \cdot 10^{-2}$	$1.66 \cdot 10^{-4}$	0.20	$1.00 \cdot 10^{-2}$	0.60	$8.86 \cdot 10^{-2}$

Note that μ_0 and σ_0^2 of the local level for datasets 2 and 3 have been set up manually.

Table E.2 Prior mean (μ_0) and variance (σ_0^2) values used to initialize the Local Trend model in Section 4.3.1

Component	Hidden State	Dataset 1		Dataset 2		Dataset 3	
		μ_0	σ_0^2	μ_0	σ_0^2	μ_0	σ_0^2
Local Trend	$x^{\texttt{L}}$	-2.99	0.20	1.4	$7.84 \cdot 10^{-4}$	-1.75	$1.23 \cdot 10^{-3}$
	$x^{\texttt{LT}}$	-0.01	$2.04 \cdot 10^{-5}$	$2.90 \cdot 10^{-2}$	$1.70 \cdot 10^{-4}$	$-7.45 \cdot 10^{-2}$	$1.12 \cdot 10^{-3}$
Periodic	$x^{\texttt{P1}}$	0.47	$2.36 \cdot 10^{-3}$	-2.09	$2.47 \cdot 10^{-2}$	-2.76	$4.27 \cdot 10^{-2}$
	$x^{\texttt{P2}}$	0.45	$2.36 \cdot 10^{-3}$	-0.14	$2.47 \cdot 10^{-2}$	$-8.75 \cdot 10^{-3}$	$4.27 \cdot 10^{-2}$
White Noise	$x^{\texttt{WN}}$	0	0.03	0	1.51	0	2.80
	$x^{\texttt{WN\_err}}$	0	0.03	0	1.51	0	2.80
	$x^{\texttt{W2}}$	0	10^{-6}	0	10^{-6}	0	10^{-6}
	$x^{\overline{\texttt{W2}}}$	0.03	$2.27 \cdot 10^{-04}$	1.51	0.57	2.80	1.96

Note that μ_0 and σ_0^2 of the local level for datasets 2 and 3 have been set up manually.

Table E.3 Prior mean (μ_0) and variance (σ_0^2) values used to initialize the Local Acceleration model in Section 4.3.1

Component	Hidden State	Dataset 1		Dataset 2		Dataset 3	
		μ_0	σ_0^2	μ_0	σ_0^2	μ_0	σ_0^2
Local Accel.	$x^{\texttt{L}}$	-2.88	0.19	1.4	$7.84 \cdot 10^{-4}$	-1.75	$1.23 \cdot 10^{-3}$
	$x^{\texttt{LT}}$	$-1.27 \cdot 10^{-2}$	$3.25 \cdot 10^{-5}$	$-5.39 \cdot 10^{-2}$	$5.85 \cdot 10^{-4}$	-0.11	$2.32 \cdot 10^{-3}$
	$x^{\texttt{LA}}$	$1.87 \cdot 10^{-5}$	$7.06 \cdot 10^{-11}$	$1.93 \cdot 10^{-4}$	$7.53 \cdot 10^{-9}$	$2.88 \cdot 10^{-4}$	$1.68 \cdot 10^{-8}$
Periodic	$x^{\texttt{P1}}$	0.41	$2.19 \cdot 10^{-3}$	-2.25	$2.86 \cdot 10^{-2}$	-2.79	$4.39 \cdot 10^{-2}$
	$x^{\texttt{P2}}$	0.47	$2.19 \cdot 10^{-3}$	-0.16	$2.86 \cdot 10^{-2}$	$-7.00 \cdot 10^{-2}$	$4.39 \cdot 10^{-2}$
White Noise	$x^{\texttt{WN}}$	0	$2.67 \cdot 10^{-2}$	0	0.32	0	0.70
	$x^{\texttt{WN\_err}}$	0	$2.67 \cdot 10^{-2}$	0	0.32	0	0.70
	$x^{\texttt{W2}}$	0	10^{-6}	0	10^{-6}	0	10^{-6}
	$x^{\overline{\texttt{W2}}}$	$2.67 \cdot 10^{-2}$	$1.79 \cdot 10^{-04}$	0.32	$2.49 \cdot 10^{-2}$	0.70	0.12

Note that μ_0 and σ_0^2 of the local level for datasets 2 and 3 have been set up manually.

Table E.4 Prior mean (μ_0) and variance (σ_0^2) values used to initialize the Exponential model in Section 4.3.2

Component	Hidden State	Dataset 4	
		μ_0	σ_0^2
Exponential	$x^{\texttt{LTi}}$	$-8.93 \cdot 10^{-4}$	10^{-10}
	$x^{\texttt{LR}}$	$8.93 \cdot 10^{-4}$	$1.62 \cdot 10^{-7}$
	$x^{\texttt{A}}$	23.67	89.62
	$z^{\texttt{E}}$	0	0
	$z^{\texttt{P}}$	0	0
Local Level 1	$x^{\texttt{L1}}$	-3.97	0.35
Local Level 2	$x^{\texttt{L2}}$	0	0
LSTM	$x^{\texttt{LSTM}}$	0	0
White Noise	$x^{\texttt{WN}}$	0	$2.39 \cdot 10^{-2}$
	$x^{\texttt{WN\_err}}$	0	$2.39 \cdot 10^{-2}$
	$x^{\texttt{W2}}$	0	10^{-6}
	$x^{\overline{\texttt{W2}}}$	$2.39 \cdot 10^{-2}$	$1.43 \cdot 10^{-4}$

APPENDIX F CODE AVAILABILITY

The exponential component developed in this thesis was implemented as a modification of Canari, an open-source library for online change-point detection in univariate time series [34]. The modified version, together with the scripts used to produce the verification results of Chapter 4, is hosted at:

`https://github.com/MichelWU99/canari-MW-memoire2026`

All modifications specific to this thesis are kept on the branch `research/memoire-2026`.

APPENDIX G RAW REAL DATASETS

Figure G.1 presents the raw, unprocessed measurements for the four real monitoring datasets used in Chapter 4, prior to outlier removal or temporal regularization. As discussed in Section 4.1, vertical axis values are omitted for confidentiality; the relative pattern of variation, including irregular acquisition dates and any isolated outliers, is preserved as originally recorded.

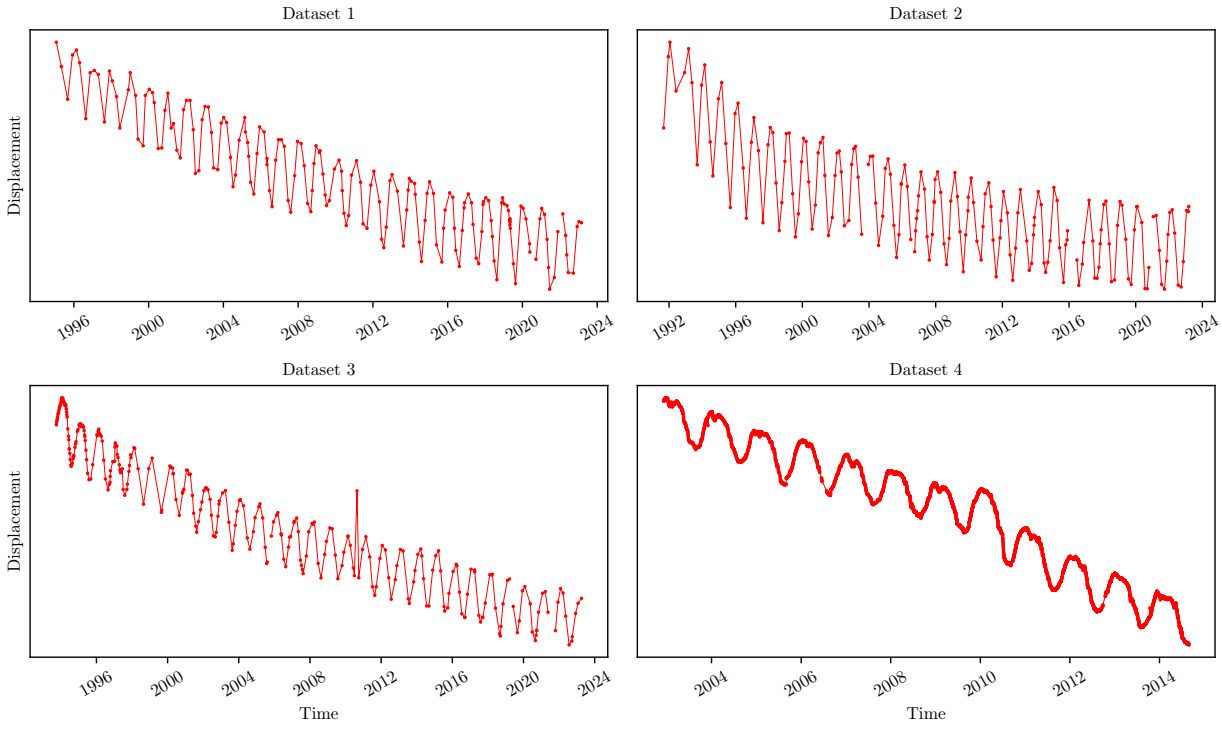

Figure G.1 Raw displacement measurements for the four real datasets studied in Chapter 4, prior to outlier removal and temporal regularization. Vertical axis values are omitted for confidentiality.