POLYPUBLIE

Anomaly detection with switching Kalman filter and imitation learning

Zhanwen Xin, James Alexandre Goulet

Document en libre accès dans PolyPublie
Open Access document in PolyPublie

Date	2024
Type	Communication de conférence / Conference or Workshop Item Communication écrite/Paper
Version	Version officielle de l'éditeur / Published version
URL de PolyPublie **PolyPublie URL**	https://publications.polymtl.ca/62199/
Évaluation par les pairs **Peer reviewed**	Révisé par les pairs / Refereed
Conditions d'utilisation **Terms of use**	Creative Commons Attribution-Utilisation non commerciale-Pas d'oeuvre dérivée 4.0 International / Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND)
Référence **Citation**	Xin, Z., & Goulet, J. A. (juin 2024). Anomaly detection with switching Kalman filter and imitation learning [Communication écrite]. 12th International Conference on Bridge Maintenance, Safety and Management (IABMAS 2024), Copenhagen, Denmark. https://doi.org/10.1201/9781003483755-179

Publication officielle
Official publication

Nom de la conférence **Conference name**	12th International Conference on Bridge Maintenance, Safety and Management (IABMAS 2024)
Date et lieu **Date and location**	2024-06-24 - 2024-06-28, Copenhagen, Denmark
Maison d'édition **Publisher**	Taylor & Francis
URL officielle **Official URL**	https://doi.org/10.1201/9781003483755-179
Mention légale **Legal notice**	CC BY-NC-ND 4.0 license

Ce fichier a été téléchargé à partir de PolyPublie, le dépôt institutionnel de Polytechnique Montréal
This file has been downloaded from PolyPublie, the institutional repository of Polytechnique Montréal
https://publications.polymtl.ca

Bridge Maintenance, Safety, Management, Digitalization and Sustainability – Jensen, Frangopol & Schmidt (eds)
© 2024 The Author(s), ISBN 978-1-032-77040-6
Open Access: www.taylorfrancis.com, CC BY-NC-ND 4.0 license

Anomaly detection with switching Kalman filter and imitation learning

Zhanwen Xin & James-A. Goulet
Polytechnique Montreal, Montreal, Canada

ABSTRACT: The detection of anomalies in structural health monitoring (SHM) time-series is notoriously difficult, because one needs to separate reversible patterns caused by the environmental factors and loading, from the baseline irreversible degradation, whereas anomalies typically consist in long-term drifts that have an effect that is orders of magnitude smaller than the reversible patterns. The switching Kalman filter (SKF) has been shown to be an effective tool to address these challenges by quantifying the probability of a switch between different regimes corresponding to normal stationary conditions and abnormal non-stationary ones. Current applications of the SKF have relied on the probability of regime switch as a proxy for detecting anomalies. Despite its capacity to outperform threshold-based detection approaches, it remains prone to (1) missed alarms when the reversible and irreversible pattern separation fails, and (2) false alarms when the predictive capacity of the stationary model is too limited. Previous works have shown that these issues can be mitigated by replacing an alarm triggering policy based solely on the probability of regime switch, with imitation learning (IL) agents that build policies based on the probability of regime switch as well as the trend hidden state. This study further enhances the IL agents by including high-dimensional hidden states vectors as well as by leveraging multiple time steps from the data history. The IL agents that have learnt the underlying dependencies between the high-dimensional hidden state estimates and the labeled actions, which in our context, are whether to trigger an alarm or not, are tested on time series including synthetic and real ones collected on a bridge located in Canada. The results show that the new method provides anomaly detection agents that improve the detectability of anomalies, and are self-adaptive to new stationary conditions.

1 INTRODUCTION

Detecting anomalies in time series is an effective way to identify infrastructures' deterioration before they can be visually observed (Pimentel et al, 2014). Time series anomaly detection methods have the potential to be scaled to infrastructure networks because they do not require a physics-based model of the structure's behaviour (Hamida & Goulet, 2022). Schmidl et al. (2022) categorized the methodologies for time-series anomaly detection as distance, forecasting, and reconstruction-based methods. Among them, the distance and forecasting methods were shown to be effective (Sarmadi & Karamodin, 2020; Willm, 1967) in identifying outliers, yet prone to missed and false alarms in SHM anomaly detection, where the anomalies typically consist in long-term drifts of baselines overlapped by reversible patterns with much bigger magnitude.

The switching Kalman filter (SKF) is a good fit for the detection of baseline changes in structural health monitoring (SHM). Compared to the Hidden Markov Models (HMM) which often model hidden states as discrete variable (Coraça et al., 2023), the state space of the hidden variables in SKF is continuous. Previous studies have shown its effectiveness in baseline extraction (Goulet, 2017), estimation of periodic patterns (Nguyen & Goulet, 2018b), and detection of non-stationary regime quantified by probability values (Nguyen & Goulet, 2018a). However, the probability of the non-stationary

DOI: 10.1201/9781003483755-179

regime sometimes fails to identify regime switches, has delays, or triggers undesirable false alarms, especially when the triggering threshold is not well calibrated. This issue has been partially mitigated by an imitation learning (IL) decision-making framework, that triggers alarms by considering not only the probability of non-stationary regime, but also the baseline estimation of local trend (Khazaeli, 2022). The framework removes the necessity of setting a threshold on the probability of regime switch, and is capable to detect anomalies with lower magnitudes and longer durations compared to only using the probability of non-stationary regime. Yet, it is limited to low-dimensional features and time series with a single anomaly, which prevents it from being used in practice where multiple anomalies might happen during the life of an infrastructure.

The IL task in the decision-making framework typically employs neural networks, which are suitable for the problems having large dimensionalities. This allows us to further leverage the hidden states estimated by the SKF through the inclusion of higher dimensional features and history data in the input of neural networks instead of being limited to the two-dimensional features for the existing decision-making framework. On the other hand, the existing framework is not compatible with time series including multiple anomalies because the model is not able to adapt to new stationary regimes after a non-stationary one. In this paper, we proposed a reference regime, with which the models can be reset during a new stationary regime, and thus continuously detect anomalies.

The objectives of this paper are twofold: (1) it is to include high-dimensional features in the decision-making framework, and (2) to develop a detection method that can handle multiple subsequent anomalies. The structure of this paper is the following: Section 2 reviews the SKF and IL methods, Section 3 presents the methodology of this study, Section 4 presents experiments on a synthetic time series and a real one measured on a bridge, comparing the performance of the SKF and the new framework.

2 BACKGROUND

In this section, the principles behind the switching Kalman filter and imitation learning, as well as their limitations are introduced.

2.1 *Switching Kalman filter*

The Kalman filter is used for decomposing the signal of a stationary linear dynamic system by recursively estimating the moments of hidden states between pairs of consecutive time steps (Welch & Bishop, 2001). By denoting $\boldsymbol{X}$ as the vector including all the hidden states, $\boldsymbol{y}$ as the observed signal, $\mathbf{A}$, $\mathbf{F}$, $\mathbf{Q}$ and $\mathbf{R}$ as the model matrices, and $\mathscr{L}$ as the likelihood of observing $\boldsymbol{y}$ given the model prediction, the Kalman filter from time step $t-1$ to t is summarized by

$$\left(\mu^X_{t/t}, \Sigma^X_{t/t}, L_{t/t}\right) = \text{Filter}\left(\mu^X_{t-1/t-1}, \Sigma^X_{t-1/t-1}, y_t; A, F, Q, R\right) \tag{1}$$

where the subscript t/t denotes the updated estimate at time t using all the observations up to t. In Bayesian dynamic linear models (BDLM), the model matrices, $\mathbf{A}$, $\mathbf{F}$, $\mathbf{Q}$ and $\mathbf{R}$, can be built modularly using various generic components (Goulet, 2020) such as (a) baseline components that describe the local level (LL), trend(LT), and acceleration (LA) of time series; (b) periodic components that capture the harmonic periodic (PD) or non-harmonic kernel regression (KR) patterns, and (c) a residual component which represents the model errors that cannot fit in any other explainable components.

The switching Kalman filter (SKF) is an extension of the Kalman filter that allows modeling non-stationary signals, made of multiple stationary regimes (Murphy, 1998). The Kalman filter transiting from regime i at time $t-1$ to regime j at t is represented by

$$\left(\mu^{i,j}_{t/t}, \Sigma^{i,j}_{t/t}, L^{i,j}_{t/t}\right) = \text{Filter}\left(\mu^i_{t-1/t-1}, \Sigma^i_{t-1/t-1}, y_t; A^j, F^j, Q^j, R^j\right) \tag{2}$$

Given the hidden states estimation in regime i at time $t-1$, it estimates the moments of the hidden states at time t by using model j, as well as the likelihood of the regime transition from i to j by comparing the observed signal $\boldsymbol{y}$ and the model prediction. By denoting $\pi^{l}_{t-1|t-1}$ as the probability of regime i at time $t-1$, and $Z^{i,j}$ as the pre-defined transition probability from $s_{t-1} = i$ to $s_t = j$, the probability of the regime j at time t is computed by

$$\pi^{j}_{t/t} = \sum_{i=1}^{S} M^{i,j}_{t-1,t/t}, \tag{3}$$

$$M^{i,j}_{t-1,t/t} = \frac{L^{i,j}_{t/t} \cdot z^{i,j} \cdot \pi^{i}_{t-1|t-1}}{\sum_{i}\sum_{j} L^{i,j}_{t/t} \cdot z^{i,j} \cdot \pi^{i}_{t-1|t-1}} \tag{4}$$

The hidden states estimated in different transitions can then be merged with $L^{i,j}_{t/t}$ and $\pi^{j}_{t/t}$ by using Gaussian mixture reduction (Runnalls, 2007). With the closed-form solution to estimate the probability of different regimes, SKF was shown to be effective in detecting anomalies in real time series by relying on the probability of the non-stationary regime as a criterion for triggering alarms (Nguyen & Goulet, 2018a). However, this probability may fail to capture the regime switches when time series are affected by unobserved external factors, or when it is subject to high level of noise. Figure 1 showcases a time series affected by a high level of noise, where the stationary regime is defined by a constant local trend. The non-stationary regime's probability, simply denoted as π, remains around zero at all times, yet the local trend hidden state (x^{LT}) switches to a new range after the regime change, indicated by the vertical dot-dashed lines. In such a case, an anomaly detection policy relying solely on the non-stationary regime's probability without considering other hidden states would be prone to missing alarms.

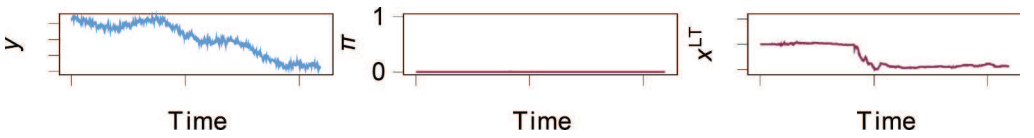


Figure 1. Limitation of the switching Kalman filter. The stationary regime in the time series y has a constant speed. The vertical dot-dashed lines indicate the time step when the speed changes. The non-stationary regime's probability π fails to capture the regime change. Yet, the local trend hidden state (x^{LT}) correctly estimates the speed change with a time delay.

2.2 *Imitation learning*

Imitation learning (IL), also known as behaviour cloning and learning from demonstration, maps a given state directly to an action without considering the long-term sequential effects (Zheng et al, 2022). The objective of an IL agent is to reproduce experts' demonstrations on how to act as a function of states. The expert's demonstration set consisting $\boldsymbol{D}$ instances is defined as $D = \{D_1, D_2, \ldots, D_D\}$, where each demonstration D_i is a state-action pair, i.e., $D_i = \left(s_i, a_i^{\mathrm{expert}}\right)$. For the anomaly detection in SHM, Khazaeli (2022) employed IL algorithms to determine whether a structure is in a normal or abnormal state given the monitoring data at time t. The state space at each time step t was defined by $s_t \in S = \left\{\mu^{LT}_{t|t}, \pi_{t|t}\right\}$, where $\mu^{LT}_{t|t}$ is the expected value of local trend hidden state, and $\pi_{t/t}$ is the probability of the non-stationary regime estimated by the switching Kalman filter. The actions are whether to trigger an alarm or not, $a_t^{expert} \in A = \left\{a^0 : \text{trigger an alarm}, a^1 : \text{not trigger an alarm}\right\}$, which are labeled from synthetic anomalies. His results showed that IL agents are capable to detect anomalies with lower magnitudes and longer durations compared to only using the probability of regime switch. However, the framework in this study is limited to low-dimensional features and deterministic setups. Time series with multiple anomalies, each defined by a regime switch, are not

compatible with the existing method because the local trend used in the state space cannot adapt to new stationary regimes. In the next section, we will introduce a new methodology, in order to overcome their limitations.

3 METHODOLOGY

This section presents a decision-making framework that leverages the high-dimensional hidden states from SKF, and a reset mechanism for the local trend in order to allow decision-making agents to adapt to different trend-stationary regimes.

3.1 *State-action pairs as the input and output of decision-making agents*

Decision-making agents for SHM must be able to handle high-dimensional states as the input for the reasons that (1) multiple hidden states other than the local trend and the probability of the non-stationary regime in SKF might contain useful patterns that contribute to the anomaly detection as shown in Figure 1; (2) patterns are typically time-dependant as anomaly develops with time, so the hidden state estimates not only at the current time step, but also at previous ones should be considered in the decision making. By using neural networks as the supervised imitation learning model, the state space can be expanded to higher dimensions simply by increasing the number of input nodes.

The action output of triggering an alarm or not is labeled according to the non-stationary (**ns**) and stationary (**s**) regimes in time series. Figure 2 depicts the labeling of the two regimes partitioned by different range of the synthetic anomaly's local trend, which gradually shifts from 0 to a negative value $-m$ following a Gaussian cumulative distributed function with random mean and variance scaled by $-m$. The anomaly duration is defined by the 95% confidence interval bounded by $-0.025m$ and $-0.975m$. The non-stationary regime begins with the anomaly, and ends when the SKF detects it and returns again to the new stationary regime. The pairs of the high dimensional input state and the output action defined above form the demonstration set D to train the IL decision-making agents.

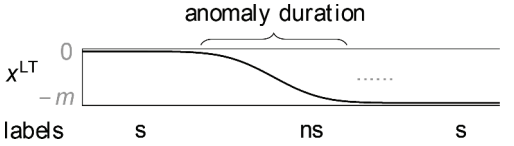


Figure 2. Labeling of non-stationary (ns) and stationary (s) regimes, which corresponds to the action of triggering an alarm or not respectively, according to the anomaly local trend(x^{LT}).

3.2 *Reset mechanism for the local trend*

As presented in Figure 1, the local trend hidden state X^{LT} is a key feature in our imitation learning decision making framework. The departure of X^{LT} from the trend baseline of the stationary regime contributes to the anomaly detection policy. However, when the time series is in a new stationary regime, where the X^{LT} baseline changes to a new level, the model trained on the previous baseline will no longer be valid and will fail to continuously detect anomalies unless it is retrained on the new regime. This causes issues in practice because SHM time series typically spans from years to decades, during which multiple anomalies may happen. To avoid retraining models in each stationary regime and to have a continuous anomaly detecting system, it is more effective to reset X^{LT} to a fixed baseline for different stationary regimes.

In order to develop such a reset mechanism and adapt X^{LT} to new trend-stationary regimes, we propose two new hidden states namely the local-trend reference $\boldsymbol{X}^{LTr}$ and the local trend in agent's vision X^{LTa}, where $X^{LTa} = X^{LT} - X^{LTr}$. The agent makes decisions based on the difference between the actual local trend estimated by the SKF, and the reference value, which is estimated by the Kalman smoother (Murphy, 1999). The stages to determine X^{LTr} are

illustrated in Figure 3, along with time-series observations (y) and the probability of regime switch (π). At the initialization stage of the analysis, $\boldsymbol{X}^{\mathrm{LTr}}$ is defined as the smoothed local trend at time step t_0 using all the future observations until the end of the initialization stage corresponding to the time t_1 when the standard deviation of the smoothed estimation $\sigma^{LT}_{t_0|t_1}$ is smaller than a pre-defined threshold $\tilde{\sigma}^{\mathrm{LT}}$. After the initialization, $\boldsymbol{X}^{\mathrm{LTr}}$ remains constant in the stationary stage. The non-stationary stage begins when π is bigger than a threshold $\tilde{\pi}$ corresponding to the time t_2, when the model realizes that a new $\boldsymbol{X}^{\mathrm{LTr}}$ is needed, but cannot be estimated yet, because the local trend estimate is not stable enough to be used as a reference for the new regime. Hidden states are considered as stable at the time step t_4, $\tilde{n}_w$ steps away from t_3, when $\pi < \tilde{\pi}$. $\tilde{n}_w$ is a pre-defined number of waiting steps. The Kalman smoother is employed to estimate the local trend at t_4 using all the future observations until time t_5, where $\sigma^{\mathrm{LT}}_{t_4|t_5} < \tilde{\sigma}^{\mathrm{LT}}$. Eventually, $\boldsymbol{X}^{\mathrm{LTr}}$ stays constant for the remaining of the new stationary stage.

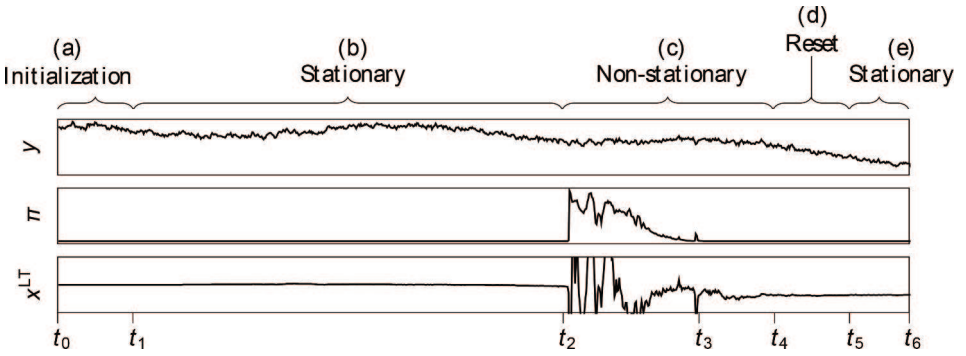


Figure 3. Reset mechanism for x^{LT} in one time series including five stages.

Figure 4 presents an example estimating the $\boldsymbol{X}^{\mathrm{LTr}}$ and $\boldsymbol{X}^{\mathrm{LTa}}$ in a time series including three regimes with reset hyperparameters $\tilde{\sigma}^{\mathrm{LT}} = 0.01$, $\tilde{\pi} = 0.1$ and $\tilde{n}_{\mathrm{w}} = 50$. The local-trend reference for the first regime converges quickly to the true value thanks to the accurate prior knowledge. As the time series enter the second regime, the x^{LT} estimate fluctuates, and the reset process for x^{LTr} and x^{LTa} is triggered. When x^{LT} becomes stable again, x^{LTr} is estimated using the Kalman smoother, and converges gradually to the true value. It is reset again when the time series is in the third regime, after which x^{LTr} is estimated and converges again. The reset local trend $\boldsymbol{X}^{\mathrm{LTa}}$ has a baseline around zero in different regimes, as required by the continuous IL agents.

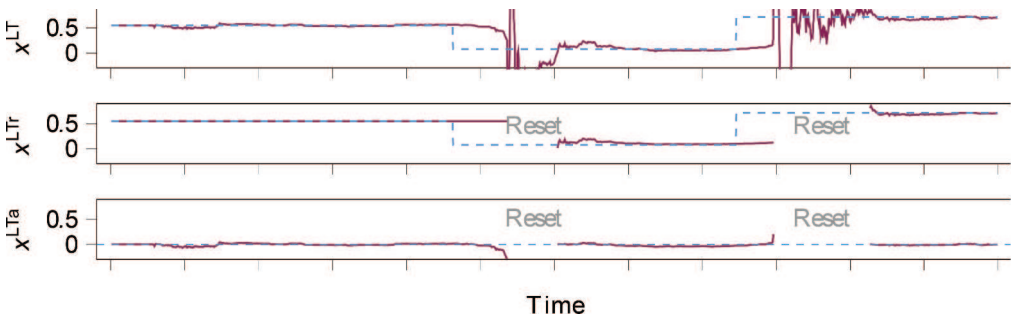


Figure 4. Estimation of local-trend reference x^{LTr} and local trend in agent's vision X^{LTa} in one time-series having three regimes with different local trends.

3.3 *SKF-IL framework for anomaly detection*

The proposed anomaly detection framework for coupling SKF and IL is designed for online operation on cases presenting small changes in the irreversible baseline, which are developing over long periods, such as corrosion, crack-opening, foundation settlements, etc. The IL agents

are trained using synthetic time series where the true labels are available as defined in Section 3.1. The pretrained agents then operate on real time series where labels are unavailable. The input hidden states of the IL agents in the training and testing process are estimated by the SKF presented in Section 2.1 and the reset mechanism in Section 3.2. Figure 5 shows a flow chart of the above-mentioned process.

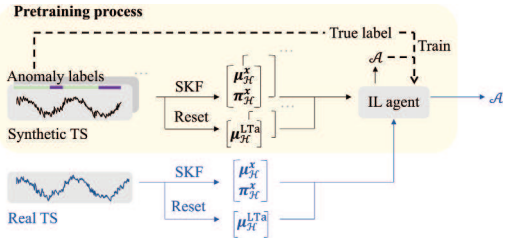


Figure 5. The training and testing process of the SKF-IL anomaly detection framework. μ_H^x represents the mean values of hidden states over time history H, π_H^x is the probability of the non-stationary regime.

4 EXPERIMENTS

An IL decision-making agent modeled by a neural network with two fully connected hidden layers of 64 hidden units, is trained with 900,000 demonstrations, $D_t = \left(s_t, a_t^{\text{expert}}\right)$, pulled out from 1,000 synthetic time series. The state space consists of the non-stationary regime's probability (π), the expected value of the local trend in the agent's vision (μ^{LTa}), as well as the expected value of the residual component (μ^{res}), at previous time steps with a logarithmic spacing from t to $t-128$, i.e., $S_t = \left\{\pi_{t-128,t-64,\ldots t-1,t}, \mu^{\text{LTa}}_{t-128,t-64,\ldots t-1,t}, \mu^{\text{res.}}_{t-128,t-64,\ldots t-1,t}\right\}$. The activation function is the rectified linear unit, and the objective function is defined by the cross-entropy loss. The pre-trained agent is tested on unseen synthetic time series in Section 4.1, and on real time series measured on a bridge in Section 4.2.

4.1 *Synthetic data*

This section verifies the anomaly detectability of the IL agent and the effectiveness of the reset mechanism for handling multiple anomalies. A synthetic time series (y) spanning 20 years is

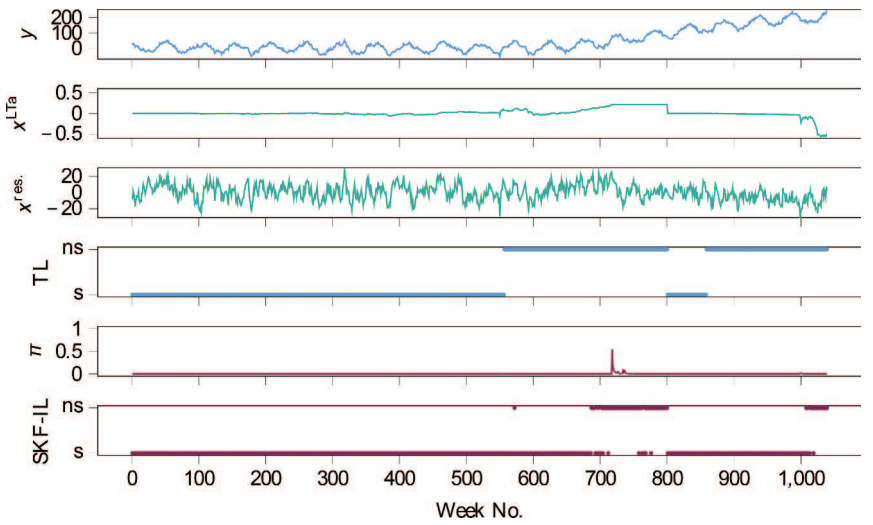


Figure 6. Testing the IL decision-making agent on a synthetic time series with two anomalies.

generated with BDLM as displayed in Figure 6. Two anomalies that shift the local trend are randomly introduced in the 557th and 859th week as indicated by the dot-dashed lines. The pre-processed hidden states x^{LTa}, x^{res}, the true labels (TL) of anomalies, the regime-switching probability π, and the decisions made by the SKF-IL agent are plotted along with y. The SKF relying only on π is capable to identify the first anomaly in the 718th week, and misses the second one with a threshold of 0.1 as in the reset mechanism. By contrast, the SKF-IL agent triggers one alarm much earlier in the 572nd week and again continues from the 687th week during the first anomaly. In addition, it detects the second anomaly missed by π in the 1008th week.

4.2 *Bridge data*

In this section, the performance of the pre-trained agent is tested on a time series measuring the elongation on a bridge. The elongation measurements are dependent on the temperature, which can be modeled by using the observation matrix $\mathbf{F}$ in the Kalman filter. The data is continuous, and the raw time series, covering the period between 2019/08/16 and 2021/11/16 is aggregated during a preprocessing step to a unified time step interval of two data points per week. It is assumed to have a zero local trend, given that no structural anomalies were observed during this period. A synthetic anomaly that shifts the zero local trend to 0.023 mm/year is introduced on 2020/11/12, as indicated by the dot-dashed lines in Figure 7, whereas the changes in the dashed local level baseline is difficult to be visually identified in the y-plot.

The hidden states of x^{LTa}, x^{res}, the true anomaly labels (TL), the probability of the non-stationary regime π estimated using the SKF, the anomalies detected by the forecasting method Prophet (Taylor & Letham, 2018), as well as the anomalies identified by our SKF-IL agent, are plotted following the elongation measurements y. The estimation of x^{LTa} remains around zero, but has fluctuations after the anomaly is introduced. The values of π gradually drifts from 0 to 1 in July 2021. Following the threshold that defines a non-stationary regime in the reset mechanism, $\bar{\pi} = 0.1$, the SKF detects the anomaly on 2021/07/16 with a delay of more than eight months, which will be even longer if the threshold is set higher. The Prophet algorithm fails to detect the anomalies, and provides two false alarms that are triggered before the anomaly is introduced with an interval width threshold of 99%. Without the need of

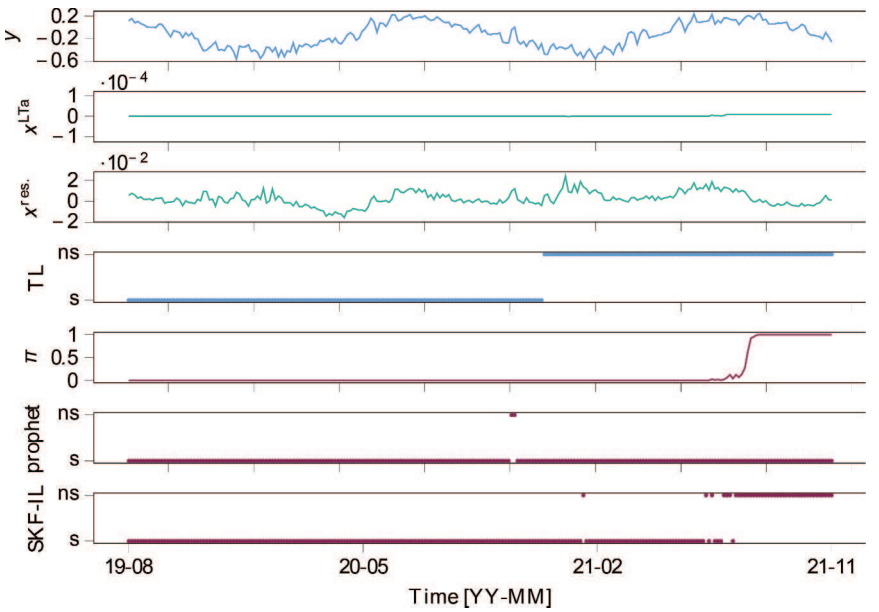


Figure 7. Testing the IL decision-making agent on a real time series measured on a bridge with one anomaly that shift the zero local trend to 0.023 mm/year.

defining a triggering threshold for the non-stationary regime's probability, the pre-trained agent detects the anomaly firstly on 2021/01/01 with a delay of less than two months, and continuously triggers more alarms after 2021/07/05. Results obtained from the real data validate our experiment on the synthetic one, where the hidden state estimates' history contributes to a decision-making policy that is capable to detect anomalies earlier than the one relying only on the non-stationary regime's probability.

5 CONCLUSION

This paper introduces a decision-making framework that detects the regime switches in time series, which is commonly encountered when anomalies develop in infrastructures. The existing switching Kalman filter identifies the regime switch by defining a threshold for the non-stationary regime's probability, which is prone to missed or delayed alarms. The imitation learning (IL) agent presented in this paper leverages hidden states in addition to the non-stationary regime's probability, as well as the history data in the decision-making process. Experiments on a synthetic time series and a real one measuring the elongation of a bridge have shown that the proposed agent can detect the regime switches earlier, reduce missed alarms, and self-adapt to new stationary regime, which allows it to continuously detect anomalies. Although much work remains to be done to probabilistically evaluate the agents' performance in big datasets, its reliability showcased in this paper makes it a promising anomaly detection framework for structural health monitoring.

REFERENCES

Coraça, E. M. & Ferreira, J. V. & Nóbrega, E. G.O. 2023. An unsupervised structural health monitoring framework based on Variational Autoencoders and Hidden Markov Models. *Reliability Engineering & System Safety* 231:109025.

Goulet, J.-A. 2017. Bayesian Dynamic Linear Models for Structural Health Monitoring. *Structural Control and Health Monitoring* 24: 1545–2263.

Goulet, J.-A. 2020. *Probabilistic Machine Learning for Civil Engineers*. Cambridge: MIT Press.

Hamida, Z. & Goulet, J.-A. 2022. A stochastic model for estimating the network-scale deterioration and effect of interventions on bridges. *Structural Control and Health Monitoring* 29: e2916.

Khazaeli, S. 2022. *Damage Detection for Structural Health Monitoring Using Reinforcement and Imitation Learning*. Montreal: Polytechnique Montréal.

Murphy, K.P. 1998. *Switching Kalman filters*. Vancouver: The University of British Columbia.

Murphy, K.P. 1999. *Filtering, smoothing and the junction tree algorithm*. Berkeley: University of California, Berkeley.

Nguyen, L.H. & Goulet, J.-A. 2018a. Anomaly Detection with the Switching Kalman Filter for Structural Health Monitoring. *Structural Control and Health Monitoring* 25: e2136.

Nguyen, L.H. & Goulet, J.-A. 2018b. Structural Health Monitoring with dependence on non-harmonic periodic hidden covariates. *Engineering Structures* 166: 187–194.

Pimentel, M.A.F. & Clifton, D.A. & Clifton, L. & Tarassenko, L. 2014. A review of novelty detection. *Signal Processing* 99: 215–249.

Runnalls, A. 2007. Kullback-Leibler approach to Gaussian mixture reduction. *IEEE Transactions on Aerospace and Electronic Systems* 43: 989–999.

Sarmadi, H. & Karamodin A. 2020. A novel anomaly detection method based on adaptive Mahalanobis-squared distance and one-class kNN rule for structural health monitoring un- der environmental effects. *Mechanical Systems and Signal Processing* 40: 106495.

Schmidl, S. & Wenig, P. & Papenbrock T. 2022. Anomaly detection in time series: A comprehensive evaluation. *Proceedings of the VLDB Endowment*.

Taylor, S. J. & Letham, B. 2018. Forecasting at scale. *The American Statistician* 72: 37–45.

Willm, N. B. G. 1967. Les méthodes de surveillance des barrages au service de la production hydraulique d'electricité de france-problèmes ancients et solutions nouvelles. *Proceedings of 9th ICOLD Congres*.

Welch, G. & Bishop, G. 2001. *An introduction to the Kalman filter*. Chapel Hill: University of North Carolina.

Zheng, B. & Verma, S. & Zhou, J. & Tsang, I. & Chen, F. 2022. Imitation learning: Progress, taxonomies and challenges. *IEEE Transactions on Neural Networks and Learning Systems*.