14th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP14
Dublin, Ireland, July 9-13, 2023

Bayesian neural networks for large-scale infrastructure deterioration models

Said Ali Kamal Fakhri
M.Sc. Candidate, Dept. of Civil Engineering, Polytechnique Montreal, Montreal, Canada

Zachary Hamida
Postdoctoral Researcher, Dept. of Civil Engineering, Polytechnique Montreal, Montreal, Canada

James-A. Goulet
Professor, Dept. of Civil Engineering, Polytechnique Montreal, Montreal, Canada

ABSTRACT: State-space models (SSM) have been shown to be effective at modelling structural deterioration of transportation infrastructure based on visual inspections. The SSM approach was recently coupled with kernel regression (KR) to include structural attributes like age and location in the deterioration analysis to share information between similar structures. However, the existing SSM-KR method suffers from two major drawbacks: 1) it can only use a limited number of structural attributes and 2) it requires significant computational time and resources. This paper proposes a new method, titled SSM-TAGI, that uses a Bayesian neural network instead of KR for extracting information from structural attributes. The new SSM-TAGI approach is compared against SSM-KR using visual inspection data and structural attributes from a network of bridges in Canada. The new SSM-TAGI approach is shown to reduce the computational time by two orders of magnitude while maintaining comparable performance as measured by the test-set log-likelihood. SSM-TAGI also seamlessly incorporates additional structural attributes and does not require extensive preparation, making it better suited for modelling infrastructure deterioration based on visual inspections on a large scale.

1. INTRODUCTION

Deterioration of transportation infrastructure from ageing, usage, and environmental exposure is an issue faced by many industrialized countries (Boller et al., 2015). Monitoring and maintaining infrastructure is critical to prolonging its life, reducing economic costs, and ensuring public safety. A common approach used for monitoring structural condition is to conduct visual inspections (MTQ, 2014). Although visual inspections are widely adopted, they suffer from a few glaring issues. Namely, visual inspections are subjective, which results in inconsistent data over time. They are also infrequent, which results in very few data points per structure over a long period of time. Despite these limitations, numerous visual inspection-based deterioration models have been developed in the literature to help infrastructure owners interpret their data and make informed decisions. Examples include discrete Markov models, regression-based methods, and state-space models (SSM) (Soetjipto et al., 2017; Ying-Hua, 2010; Hamida and Goulet, 2020). The SSM approach is unique among these methods because it incorporates inspector uncertainty in the deterioration analysis. Although this allows quantifying the observation uncertainty, the visual inspection data is error-ridden and limited in quantity, which hinders the predictive capacity of SSM. This limitation was addressed by combining SSM with kernel regression (KR) to incorporate

structural attributes like age and location into the deterioration analysis, allowing information to be shared between structures with similar properties (Hamida and Goulet, 2021). However, the SSM-KR framework is sensitive to the selection and initialization of model parameters and can only incorporate a limited number of structural attributes while requiring significant computational time and resources. The high computational demand is particularly problematic when dealing with network-scale inspections of infrastructure, where datasets can be large.

This work proposes to replace kernel regression in the SSM-KR with a Bayesian neural network trained by tractable approximate Gaussian inference (TAGI), an analytical inference method developed by Goulet et al. (2021). Contrary to KR, TAGI is fast and can incorporate many structural attributes, removing the need for feature engineering. The proposed SSM-TAGI framework is compared against SSM-KR using visual inspection data and structural attributes from a network of bridges in the province of Quebec, Canada. The new approach is shown to be faster and scalable while maintaining comparable performance and requiring virtually no set-up.

2. METHODOLOGY

This section describes how the proposed deterioration model is formulated by coupling a state-space model with a Bayesian neural network.

2.1. Modelling deterioration using state-space models

A state-space model (SSM) is a probabilistic model that involves defining transition and observation equations tailored to the phenomena being modelled (Goulet, 2020). In the context of bridge deterioration, the transition equation describes the deterioration kinematics and the observation equation describes the visual inspection process (Hamida and Goulet, 2020),

$$\overbrace{\boldsymbol{x}_t = \boldsymbol{A}\boldsymbol{x}_{t-1} + \boldsymbol{w}_t}^{\text{transition model}},\ \overbrace{\boldsymbol{w}_t : \boldsymbol{W} \sim \mathcal{N}(\boldsymbol{w}; \boldsymbol{0}, \boldsymbol{Q})}^{\text{process errors}}, \quad (1)$$

$$\overbrace{y_t = \boldsymbol{C}\boldsymbol{x}_t + v_t}^{\text{observation model}},\ \overbrace{v_t : V \sim \mathcal{N}(v; \mu_V(I_i), \sigma_V^2(I_i))}^{\text{observation errors}}. \quad (2)$$

The state of the structure at any time $t \in [1, \mathtt{T}]$ is represented by the vector $\boldsymbol{x}_t$, which contains the condition x_t, speed $\dot{x}_t$, and acceleration $\ddot{x}_t$. In Equation (1), $\boldsymbol{A}$ is the state transition matrix, $\boldsymbol{w}_t$ is the process error, and $\boldsymbol{Q}$ is the process error covariance matrix. In Equation (2), y_t represents the observation, $\boldsymbol{C}$ is the observation matrix, v_t is the observation error, and $\mu_V(I_i)$ and $\sigma_V^2(I_i)$ are the relative bias and variance that characterize the i^{th} inspector's observation error. The estimation of the deterioration state at each time step t is done by using the Kalman filter (KF) (Kalman, 1960), expressed concisely as,

$$(\boldsymbol{\mu}_{t|t}, \boldsymbol{\Sigma}_{t|t}, \mathscr{L}_t) = \text{Kalman filter}(\boldsymbol{\mu}_{t-1|t-1}, \boldsymbol{\Sigma}_{t-1|t-1}, y_t, \boldsymbol{A}, \boldsymbol{Q}, \boldsymbol{C}, \mu_V(I_i), \sigma_V^2(I_i)),$$

where $\boldsymbol{\mu}_{t|t}$ is the posterior expected value and $\boldsymbol{\Sigma}_{t|t}$ is the posterior covariance for the state at time t, given all the observations up to time t, and $\mathscr{L}_t$ is the log-likelihood for the observation y_t. After obtaining the estimates of the state at each time step t by the KF, the Kalman smoother (KS) (Rauch et al., 1965) is then applied to refine these estimates,

$$(\boldsymbol{\mu}_{t|\mathtt{T}}, \boldsymbol{\Sigma}_{t|\mathtt{T}}) = \text{Kalman smoother}(\boldsymbol{\mu}_{t+1|\mathtt{T}}, \boldsymbol{\Sigma}_{t+1|\mathtt{T}}, \boldsymbol{\mu}_{t|\mathtt{T}}, \boldsymbol{\Sigma}_{t|\mathtt{T}}, \boldsymbol{A}, \boldsymbol{Q}),$$

where $\boldsymbol{\mu}_{t|\mathtt{T}}$ and $\boldsymbol{\Sigma}_{t|\mathtt{T}}$ are the respective posterior mean and covariance of the smoothed state estimates at time t, given all the observations up to time $\mathtt{T}$. To correctly model deterioration, the condition x_t must always be declining. This constraint is enforced by restricting the deterioration rate $\dot{x}_t$ to be negative at each time step using a probability density function truncation method (Simon and Simon, 2010).

Despite the SSM framework effectively modelling deterioration, it relies only on the visual inspection data and is incapable of incorporating other information such as structural attributes. This limitation can be addressed by coupling a regression method with the SSM approach, as detailed in the next two sections.

2.2. Regression using tractable approximate Gaussian inference

Tractable approximate Gaussian inference (TAGI) is a framework for the analytical inference

of parameters associated with Bayesian neural networks (Goulet et al., 2021). TAGI assumes that the neural network parameters consisting of weights and biases are described by Gaussian random variables: $\boldsymbol{\theta}_{\texttt{TAGI}} \sim \mathcal{N}(\boldsymbol{\mu}_{\texttt{TAGI}}, \boldsymbol{\Sigma}_{\texttt{TAGI}})$. TAGI leverages this assumption, along with some approximations detailed in Goulet et al. (2021), to infer the neural network parameters $\boldsymbol{\theta}_{\texttt{TAGI}}$ in closed form. Moreover, TAGI is capable of predicting the heteroscedastic aleatory uncertainty — that is, uncertainty that varies across the input covariates and which is intrinsic to the process being modelled (Deka, 2022).

In this framework, TAGI is used to incorporate the structural attribute data in the deterioration analysis to share information between similar structures. Specifically, TAGI learns the relation between the structural attributes $\boldsymbol{z}$ and the structural elements' deterioration speeds $\dot{\boldsymbol{x}}_0$. The probabilistic nature of TAGI makes it possible to use the learnt relation to improve the initialization of the KF, and therefore, to improve the model predictions.

2.3. Deterioration modelling using the hybrid SSM-TAGI framework

The succinct version of the proposed SSM-TAGI framework for estimating the structural deterioration $\tilde{\boldsymbol{x}}_t$ of a given structural element, once all the model parameters have been estimated, is presented in Figure 1. The estimation of the deterioration states starts with transforming the observations $\tilde{\boldsymbol{y}}_t$ into the unbounded domain $\boldsymbol{y}_t$ using the transformation function $o(.)$ specified in Hamida and Goulet (2020). Then, the prior for the initial state $\boldsymbol{x}_0$ is set, where TAGI is used to set the initial deterioration rate $\dot{x}_0$ using the structural attributes $\boldsymbol{z}$. The initial state $\boldsymbol{x}_0$ and the observations $\boldsymbol{y}_t$ are then passed into the KF to estimate the evolution of the deterioration states $\boldsymbol{x}_t$ over time. Finally, KS is applied to refine the KF estimates, and the smoothed estimates are back-transformed with $o(.)^{-1}$ to the original space $\tilde{\boldsymbol{x}}_t \in [l,u]$ for interpretation (Hamida and Goulet, 2020).

In the SSM-TAGI framework, the initial state $\boldsymbol{x}_0 : \boldsymbol{X}_0 \sim \mathcal{N}(\boldsymbol{x}_0; \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ is set as follows,

$$\begin{bmatrix} x_0 \\ \dot{x}_0 \\ \ddot{x}_0 \end{bmatrix} \sim \mathcal{N}\left(\overbrace{\begin{bmatrix} \lambda \\ \dot{\mu}(z) \\ 0 \end{bmatrix}}^{\boldsymbol{\mu}_0}, \overbrace{\begin{bmatrix} \alpha & 0 & 0 \\ 0 & \dot{\sigma}(z) & 0 \\ 0 & 0 & \ddot{\sigma}_0^* \end{bmatrix}^2}^{\boldsymbol{\Sigma}_0} \right), \quad (3)$$

where α and λ are the expected value and standard deviation of the condition, $\dot{\mu}(\boldsymbol{z})$ and $\dot{\sigma}(\boldsymbol{z})$ are the expected value and standard deviation of the element's deterioration rate produced by TAGI, and $\ddot{\sigma}_0^*$ is the optimized parameter for the standard deviation of acceleration. The α parameter is defined as,

$$\alpha = \max(\sigma_0^*, \sigma_{V,t=1}),$$

where σ_0^* is the optimized parameter for the standard deviation of condition and $\sigma_{V,t=1}$ is the uncertainty of the first observation. The maximum operation is used to set α to prevent having a prior with lower variance than the first observation. The λ in Equation 3 represents a fraction of the maximum observation performed on the element and is defined as,

$$\lambda = \left(1 - \frac{\mathbb{E}[\mu_V(1:\texttt{I})]}{u-l}\right) \times \max\{y_1, y_2, ..., y_{\texttt{T}}\},$$

where u and l are the respective upper and lower bounds for the structural element's condition and $\mathbb{E}[\mu_V(1:\texttt{I})]$ is the global inspector bias. The condition is initialized with λ to avoid having a prior estimate significantly lower than the observations at the following years, which could introduce numerical instability in the deterioration model predictions.

2.4. Model parameters

The full framework relies on the following parameters,

$$\boldsymbol{\theta} = \{\boldsymbol{\mu}_V(I_{1:\texttt{I}}), \boldsymbol{\sigma}_V(I_{1:\texttt{I}}), \sigma_W, n, \sigma_0, \ddot{\sigma}_0, p_1, p_2, \boldsymbol{\theta}_{\texttt{TAGI}}\}.$$

The parameters of the framework include the relative biases $\boldsymbol{\mu}_V(I_{1:\texttt{I}})$ and standard deviations $\boldsymbol{\sigma}_V(I_{1:\texttt{I}})$ of all the inspectors, the standard deviation of the process noise error σ_W, the space transformation parameter n, the TAGI parameters $\boldsymbol{\theta}_{\texttt{TAGI}}$, and the initial state standard deviation parameters

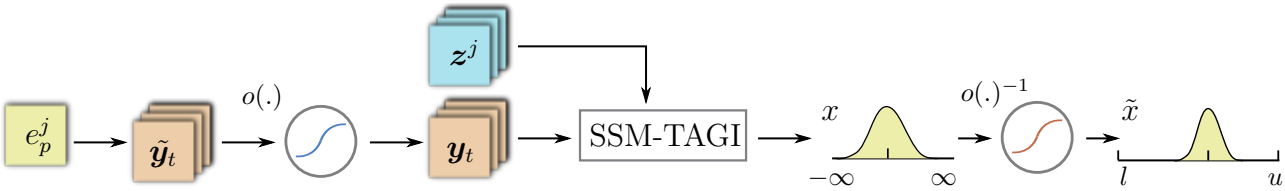

Figure 1: Full process of estimating the deterioration condition of the p^{th} element, part of the j^{th} bridge, using the SSM-TAGI framework.

$\{\sigma_0, \ddot{\sigma}_0, p_1, p_2\}$, with $\dot{\sigma}_0$ defined by the following linear relation,

$$\dot{\sigma}_0^2 = p_1^2 (u - \widetilde{\mu}_1) + p_2^2, \tag{4}$$

where $\widetilde{\mu}_1$ is the expected value of the condition at $t = 1$. The value of $\widetilde{\mu}_1$ is initially set to the first observation y_1; however, once the smoothed estimates are obtained, it is set as $\widetilde{\mu}_1 = \mu_{1|\mathtt{T}}$.

2.5. *Parameter estimation procedure*

The parameters of the proposed SSM-TAGI framework are estimated over multiple passes over the data. To avoid overfitting, the data is split into training, validation, and test sets with the constraint being that the structural elements of the same bridge cannot be simultaneously in the training set and the validation/test set. The validation set log-likelihood is used to determine the number of passes over the data and the test set log-likelihood is used to measure the generalization performance. The framework parameters are estimated using a mix of Bayesian updating and gradient-based optimization coupled with Maximum Likelihood Estimate (MLE). In MLE, the objective is to maximize the likelihood, which in this context corresponds to the joint probability of all the observations on the bridge network, given the model parameters $\boldsymbol{\theta}$ and that the observations are conditionally independent given the states $\boldsymbol{x}_t$. The natural logarithm of the likelihood is taken to ensure numerical stability; thus, the log-likelihood is given by,

$$\mathscr{L}(\boldsymbol{\theta}) = \sum_{j=1}^{\mathtt{B}} \sum_{p=1}^{\mathtt{E}_j} \sum_{t=1}^{\mathtt{T}_p} \ln f(y_{t,p}^j | y_{1:t-1,p}^j, \boldsymbol{\theta}),$$

where $\mathtt{B}$ denotes the total number of bridges in the network, $\mathtt{E}_j$ denotes the number of structural elements in the j^{th} bridge, and $\mathtt{T}_p$ denotes the number of visual inspections on the p^{th} element (Hamida and Goulet, 2021). This framework relies on the Newton-Raphson (NR) method (Hamida and Goulet, 2021) for gradient-optimization and approximate Gaussian variance inference (AGVI) (Deka, 2022) and TAGI for Bayesian updating.

The parameter estimation procedure is started with estimating $\{\sigma_W, \sigma_V, \sigma_0, \ddot{\sigma}_0, p_1, p_2\} \subset \boldsymbol{\theta}$ using NR. In this step, all the inspectors are assigned the same standard deviation $\boldsymbol{\sigma}_V(I_{1:\mathtt{I}}) = \sigma_V$ and their relative biases are fixed to zero $\boldsymbol{\mu}_V(I_{1:\mathtt{I}}) = 0$. Then, AGVI is used to estimate each inspector's parameters $\boldsymbol{\theta}_\mathtt{I} = \{\boldsymbol{\mu}_V(I_{1:\mathtt{I}}), \boldsymbol{\sigma}_V(I_{1:\mathtt{I}})\} \subset \boldsymbol{\theta}$ while the rest of the parameters remain fixed. The details of this step are outlined by Laurent (2022). Once the inspector parameters $\boldsymbol{\theta}_\mathtt{I}$ are estimated, they are fixed, and the SSM parameters $\boldsymbol{\theta}_\mathtt{S} = \{\sigma_W, \sigma_0, \ddot{\sigma}_0, p_1, p_2\} \subset \boldsymbol{\theta}$ are estimated using NR. Finally, $\boldsymbol{\theta}_\mathtt{I}$ and $\boldsymbol{\theta}_\mathtt{S}$ are fixed and the TAGI parameters $\boldsymbol{\theta}_{\mathtt{TAGI}}$ are estimated according to the procedure outlined in Section 2.6. The sequential estimation of $\boldsymbol{\theta}_\mathtt{I}$, $\boldsymbol{\theta}_\mathtt{S}$, and $\boldsymbol{\theta}_{\mathtt{TAGI}}$ is repeated until the improvement in the validation set log-likelihood is less than 0.1%. The transformation parameter n is found by repeating the estimation procedure for all the possible n values. The full estimation procedure is presented as a pseudocode in Algorithm 1.

2.6. *Recursive estimation of TAGI parameters*

TAGI's parameters are initially estimated using the structural attributes $\boldsymbol{z}$ and the smoothed estimates for the deterioration speeds $(\dot{\boldsymbol{\mu}}_{0|\mathtt{T}}, \dot{\boldsymbol{\sigma}}_{0|\mathtt{T}})$ generated by the SSM approach. However, after obtaining the initial estimate of the TAGI parameters $\boldsymbol{\theta}_{\mathtt{TAGI}}$, the SSM approach is replaced by the SSM-TAGI framework to generate the smoothed initial speeds $(\dot{\boldsymbol{\mu}}_{\boldsymbol{z},0|\mathtt{T}}, \dot{\boldsymbol{\sigma}}_{\boldsymbol{z},0|\mathtt{T}})$. TAGI then continues

for $n \in \{1,2,3,4,5\}$ **do**
  optimize $\{\sigma_W, \sigma_V, \sigma_0, \ddot{\sigma}_0, p_1, p_2\}$ using **NR**
  while validation set log-likelihood improvement $\geq 0.1\%$ **do**
    optimize $\boldsymbol{\theta}_{\mathtt{I}} = \{\boldsymbol{\mu}_V(I_{1:I}), \boldsymbol{\sigma}_V(I_{1:I})\}$ using **AGVI**
    optimize $\boldsymbol{\theta}_{\mathtt{S}} = \{\sigma_W, \sigma_0, \ddot{\sigma}_0, p_1, p_2\}$ using **NR**
    optimize $\boldsymbol{\theta}_{\mathtt{TAGI}}$ using **recursive estimation** (Sec. 2.6)
  end
end

Algorithm 1: The pseudocode for estimation of the model parameters $\boldsymbol{\theta}$.

to learn $\boldsymbol{\theta}_{\mathtt{TAGI}}$ in a recursive manner until the validation set log-likelihood ceases to improve. This estimation procedure is illustrated in Figure 2.

To generate the initial set of smoothed deterioration speeds $(\dot{\boldsymbol{\mu}}_{0|\mathtt{T}}, \dot{\boldsymbol{\sigma}}_{0|\mathtt{T}})$, the initial state $\boldsymbol{x}_0$ is set as follows,

$$\begin{bmatrix} x_0 \\ \dot{x}_0 \\ \ddot{x}_0 \end{bmatrix} \sim \mathcal{N}\left(\overbrace{\begin{bmatrix} u \\ 0 \\ 0 \end{bmatrix}}^{\mu_0}, \overbrace{\begin{bmatrix} 5 & 0 & 0 \\ 0 & g(p_1,p_2) & 0 \\ 0 & 0 & \ddot{\sigma}_0 \end{bmatrix}^2}^{\Sigma_0} \right),$$

where $g(p_1, p_2)$ represents Equation (4) and u represents the upper bound for the condition (i.e. perfect condition). Initializing the condition with $\mu_0 = u$ and $\sigma_0^2 = 5^2$ is done to provide a broad prior for the initial condition. The estimation of the deterioration state over time $\boldsymbol{x}_t$ is done by propagating the initial state $\boldsymbol{x}_0$ using the KF, and then improving the KF estimates by applying the KS. The obtained smoothed estimates for the deterioration speeds $(\dot{\boldsymbol{\mu}}_{0|\mathtt{T}}, \dot{\boldsymbol{\sigma}}_{0|\mathtt{T}})$ and the corresponding structural attributes $\boldsymbol{z}$ are then passed into TAGI as training data. The details related to the TAGI training are outlined by Goulet et al. (2021) and Deka (2022). Once TAGI is trained, the SSM-TAGI framework is used to generate the smoothed estimates $(\dot{\boldsymbol{\mu}}_{\boldsymbol{z},0|\mathtt{T}}, \dot{\boldsymbol{\sigma}}_{\boldsymbol{z},0|\mathtt{T}})$ as detailed in Section 2.3, with one exception. Namely, during the training, the variance of the speed is initialized with Equation (4) instead of with the variance produced by TAGI. This is done to prevent TAGI potentially learning a poor first estimate of the speed variance, which would prevail in the framework given its recursive nature. The recursive procedure of generating the initial speed estimates $(\dot{\boldsymbol{\mu}}_{\boldsymbol{z},0|\mathtt{T}}, \dot{\boldsymbol{\sigma}}_{\boldsymbol{z},0|\mathtt{T}})$ with SSM-TAGI and then using these estimates to train TAGI is repeated until the validation set log-likelihood no longer improves.

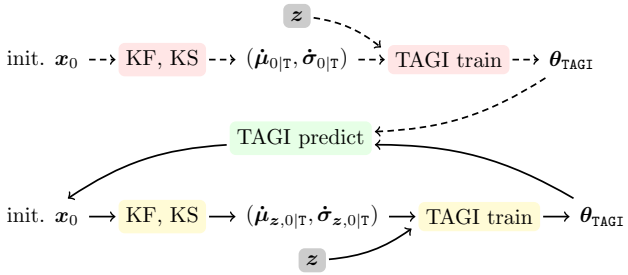

Figure 2: The estimation of TAGI parameters. The recursive estimation, represented by the solid arrows, is terminated once the validation set log-likelihood starts to decline.

3. CASE STUDY

This section provides a brief overview of the data used in this study and illustrates the results of SSM-TAGI versus SSM-KR.

3.1. Data description

The visual inspection and structural attribute data used for this study is obtained from the network of bridges in the province of Quebec, Canada. The data consists of $\mathtt{B} \approx 10,000$ bridges, each characterized with 11 structural attributes, which are listed in the left column of Table 1. The methods were compared using only the beam elements of the bridges. The beam data consists of $\mathtt{B} = 48,824$ elements and $\mathtt{I} = 295$ inspectors. The inspectors report the condition of an element by visually examining it and assigning fractions of its area to four damage categories: A: *Nothing to little*, B: *Medium*, C: *Important*, and D: *Very Important* (MTQ, 2014). For simplicity, the four categories are aggregated into a single metric $\tilde{y}$ as specified by Hamida and Goulet (2020), which results in $\tilde{y} = 100$ corresponding to the perfect condition and $\tilde{y} = 25$ corresponding to the worst condition.

3.2. Results

SSM-TAGI is initially compared against SSM-KR using 4 covariates, which include 3 structural attributes and the condition at $t = 1$, which is estimated as the average of the first three inspections (Hamida and Goulet, 2021). The attributes are the

structure's material type, latitude, and age at the time of the first inspection. Since TAGI is inherently scalable, SSM-TAGI is also tested with all the available attributes to determine the effect of additional attributes on the performance.

The hyperparameter selection and training of the KR is done as specified in Hamida and Goulet (2021). As for TAGI, the employed architecture consists of a single hidden layer with 128 hidden units and ReLU activation function. The TAGI weights and biases are initialized using the He initialization method (He et al., 2015). A batch size of one is used and the input covariates for TAGI are processed as follows, the material type is encoded using one-hot encoding since there is no natural order between different material categories while the rest of the covariates and observations are standardized to have zero mean and unit variance. The training data used for TAGI is further split into the training and validation sets following an 85/15 split, where the validation set is used to signal the end of training once its log-likelihood is no longer improving.

The performance of SSM-KR versus SSM-TAGI on a single beam element is illustrated in Figure 3. The plot includes the predictions by SSM-KR and SSM-TAGI with 4 covariates, as well as SSM-TAGI with 12 covariates. The beam element examined in Figure 3 is taken from bridge $\mathscr{B}_4$ and has an estimated condition of approximately 96 at the time of the first inspection. The structural attributes of this element with their corresponding values are listed in Table 1. As evident from the plot in Figure 3, the predictions of all three models are closely aligned, indicating similar performance on this element; the estimates for the deterioration condition also demonstrate good adaption to the inspection data. The models are also compared on an unseen observation that was never used during the training, shown in red in Figure 3. Such an observation serves as a testing point since the true state of the element is not known. In this example, the unseen observation appears to align best with the prediction of the SSM-TAGI model with 4 covariates. Although examining the difference between an unseen observation $\tilde{y}_{t|\mathrm{T}}$ and the model prediction $\tilde{\mu}_{t|\mathrm{T}}$ on a single element does not convey much information about the model's performance, examining this difference on a larger scale can convey if the model is biased to overestimation or underestimation. To that end, 10899 of unseen observations corresponding to forecast durations ranging from 1 year to 4 years were used to produce the scatter plots for all three models in Figure 4. Given that majority of the scatter points are approximately equally spread on either side of the diagonals, none of the examined models appear to be biased.

Table 1: Structural attributes of the third beam element taken from bridge $\mathscr{B}_4$.

structural attribute	value
material	concrete
age	49 years
latitude	48.5374
longitude	-78.1311
total length	18.3m
slab length	12.2m
total width	10.8m
surface area	131 m^2
number of lanes	2
percentage of trucks	10%
annual average daily traffic	2800 cars/day

To get a better sense of the global performance, all three models were evaluated on an independent test set of bridges containing $\mathrm{E}_{\text{test}} = 1004$ beam elements that were not used during the training. The results are summarized in Table 2. Although, there is not a major difference between the models in terms of log-likelihood values, there is a significant difference in terms of computational time, as SSM-TAGI outperformed SSM-KR by two orders of magnitude. Moreover, the introduction of additional structural attributes did not have a significant impact on the performance, as SSM-TAGI with 12 covariates produced virtually the same log-likelihood as SSM-TAGI with 4 covariates. Intuitively, increasing the number of covariates should improve the performance, as this results in the model having access to more information. However, it is possible that in this case all the information was already encapsulated by the few struc-

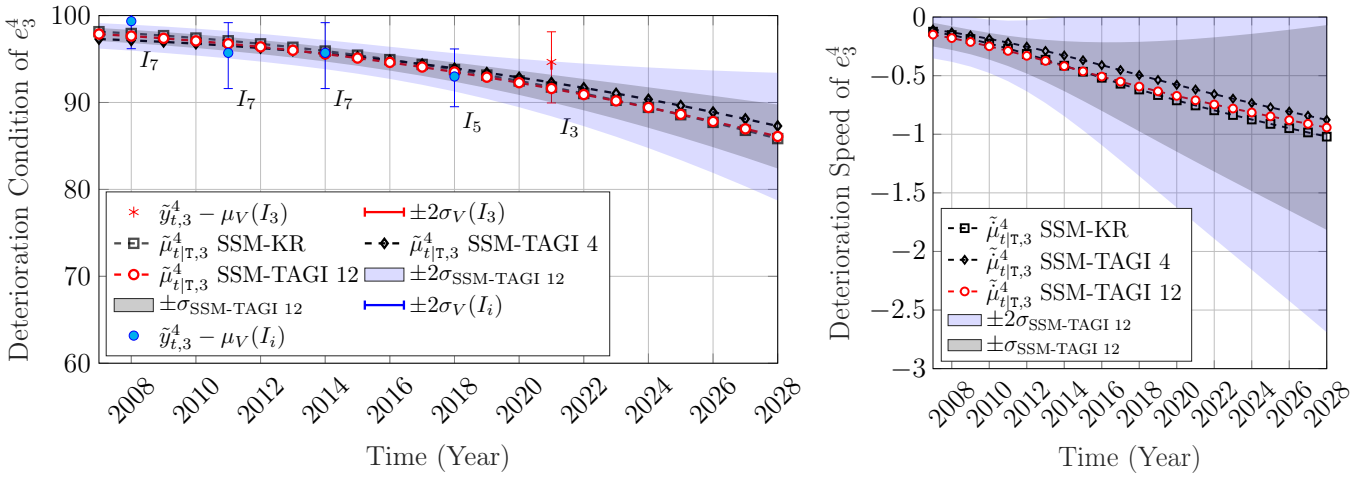

Figure 3: Deterioration condition and speed estimates for a beam element. The circle marker represents the SSM-TAGI estimates using all the available covariates. The square marker represents SSM-KR estimates and the diamond marker represents SSM-TAGI estimates, both of which rely on four covariates only: material, age, latitude, and structural condition at the time of the first inspection.

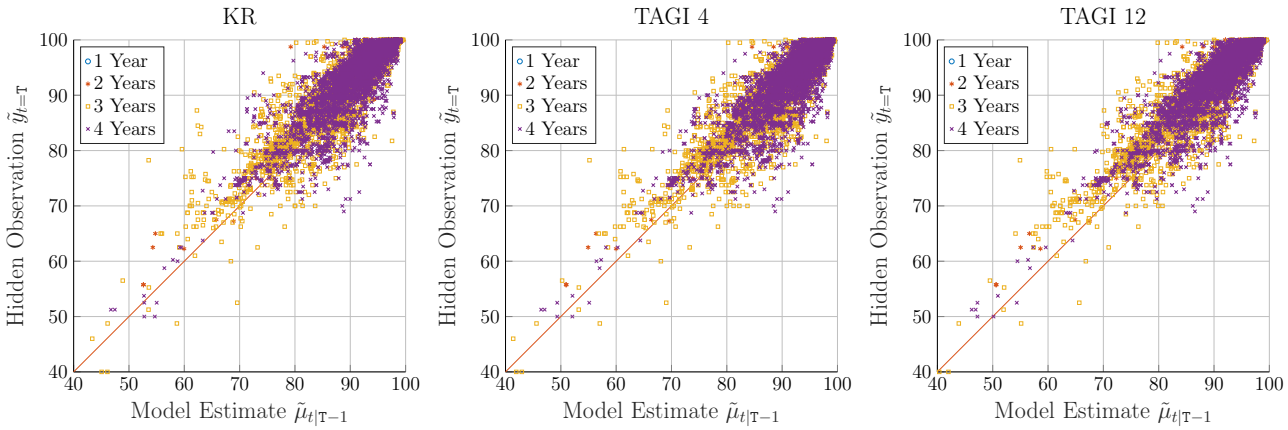

Figure 4: Model forecast $\tilde{\mu}_{t|\mathrm{T}}$ vs. unseen observations $\tilde{y}_{t|\mathrm{T}}$ with different symbol sizes representing different forecast durations.

Table 2: The performance and training time of SSM-KR on beam elements against SSM-TAGI with 4 covariates and SSM-TAGI with 12 covariates. The computational times are estimated using a system equipped with CPU Intel® Xeon® CPU E5-2687W v4, 128GB memory and NVIDIA Quadro RTX 5000.

model	number of regression covariates	test-set log-likelihood	total training time (hours)
SSM-KR	4	-14180	189.3
SSM-TAGI	4	**-13942**	4.11
SSM-TAGI	12	-13945	**2.87**

tural attributes, and that the additional structural attributes did not carry much weight in the regression analysis. The conclusion that using more structural attributes does not improve the predictive capacity cannot be drawn from these results since they were obtained by examining only the beam elements. To conclusively determine the effects of adding more structural attributes in the regression analysis, the

analysis should also be performed on other elements such as slabs, pavement, and so on.

4. Conclusion

The results of this study showed the newly developed SSM-TAGI framework to outperform SSM-KR, as it demonstrated comparable predictive capacity while reducing the computational time by two orders of magnitude. Moreover, SSM-TAGI was shown to be a more scalable method as it seamlessly included additional structural attributes in the deterioration analysis without affecting the computational time. Overall, this new method is a step towards efficient large-scale visual inspection-based infrastructure deterioration models. It is worth noting that this study only examined these methods on the beam elements of bridges. Future work should compare these models on other element types such as slabs, where the additional structural attributes could prove to be more influential covariates in the regression analysis. Additionally, using TAGI to predict the full initial state based on the structural attributes, instead of just the initial speed has the potential to improve the predictive capacity of the deterioration model.

Acknowledgements

This project is funded by the Ministry of Transportation Quebec (MTQ). The first author was supported by a research grant from the Institute for Data Valorization (IVADO).

5. References

Boller, C., Starke, P., Dobmann, G., Kuo, C.-M., and Kuo, C.-H. (2015). "Approaching the assessment of ageing bridge infrastructure." *Smart Structures and Systems*, 15(3), 593 – 608.

Deka, B. (2022). "Analytical bayesian parameter inference for probabilistic with engineering applications." Ph.D. thesis, Polytechnique Montréal, Montreal, Canada.

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

Goulet, J.-A., Nguyen, L., and Amiri, S. (2021). "Tractable approximate gaussian inference for bayesian neural networks." *Journal of Machine Learning Research*, 22(251), 1–23.

Hamida, Z. and Goulet, J.-A. (2020). "Modeling infrastructure degradation from visual inspections using network-scale state-space models." *Structural Control and Health Monitoring*, 27(9), e2582 e2582 stc.2582.

Hamida, Z. and Goulet, J.-A. (2021). "Network-scale deterioration modelling of bridges based on visual inspections and structural attributes." *Structural Safety*, 88, 102024.

He, K., Zhang, X., Ren, S., and Sun, J. (2015). "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification." *2015 IEEE International Conference on Computer Vision (ICCV)*, 1026–1034.

Kalman, R. E. (1960). "A new approach to linear filtering and prediction problems." *Journal of basic Engineering*, 82(1), 35–45

Laurent, B. (2022). "Analytical inference for visual inspection uncertainty in the context transportation infrastructures." M.S. thesis, Polytechnique Montréal, Montreal, Canada.

MTQ (2014). *Manuel d'Inspection des Structures*. Ministère des Transports, de la Mobilité Durable et de l'Électrification des Transports.

Rauch, H. E., Striebel, C., and Tung, F. (1965). "Maximum likelihood estimates of linear dynamic systems." *AIAA journal*, 3(8), 1445–1450

Simon, D. and Simon, D. L. (2010). "Constrained kalman filtering via density function truncation for turbofan engine health estimation." *International Journal of Systems Science*, 41(2), 159–171

Soetjipto, J. W., Adi, T. J. W., and Anwar, N. (2017). *Bridge Deterioration Prediction Model Based On Hybrid Markov-System Dynamic*, Vol. 138. EDP Sciences.

Ying-Hua, H. (2010). "Artificial neural network model of bridge deterioration." *Journal of Performance of Constructed Facilities*, 24(6), 597–602.