MODEL-BASED UNKNOWN INPUT ESTIMATION VIA PARTIALLY OBSERVABLE MARKOV DECISION PROCESSES

Wei Liu[1,3], **Zhilu Lai**[4], **Charikleia D. Stoura**[2], **Kiran Bacsa**[2,3], **Eleni Chatzi**[2,3]
[1]Department of Industrial Systems Engineering and Management, National University of Singapore
[2]Department of Civil, Environmental and Geomatic Engineering, ETH Zürich
[3]Future Resilient Systems, Singapore-ETH Centre
[4]Internet of Things Thrust, Information Hub, HKUST(GZ)
weiliu@u.nus.edu, zhilulai@ust.hk, kiran.bacsa@sec.ethz.ch,
{charikleia.stoura, chatzi}@ibk.baug.ethz.ch

ABSTRACT

In the context of condition monitoring for structures and industrial assets, the estimation of unknown inputs, usually referring to acting loads, is of salient importance for guaranteeing safe and performant engineered systems. In this work, we propose a novel method for estimating unknown inputs from measured outputs, for the case of systems with a known or learned model of the underlying dynamics. The objective is to infer those system inputs that will reproduce the actual measured outputs; this can be reformulated as a Partially Observable Markov Decision Process (POMDP) problem and solved with well-established planning algorithms for POMDPs. The cross-entropy method (CEM) is adopted in this paper for solving the POMDP due to its efficiency and robustness. The proposed method is demonstrated using simulated dynamical systems for structures with known dynamics, as well as a real wind turbine with learned dynamics inferred through Neural Extended Kalman Filters (Neural EKF); a deep learning-based method for learning stochastic dynamics, previously proposed by the authors.

1 INTRODUCTION

The inference of the external inputs that are acting on dynamical systems, operating across domains (including engineering, robotics, economics, and biology), is essential for understanding the conditions under which systems operate. This is particularly true for the domains of Structural Health Monitoring (SHM) and Prognostics and Health Management (PHM). The assessment of performance or condition, such as fatigue accumulation and reliability, can be improved through the accurate estimation of acting loads (Leitner & Figuli, 2018). The direct measurement of these inputs is often challenging due to their distributed and continuous nature and the limited and noisy available observations (Vettori et al., 2023). Various methods have been proposed to estimate unknown inputs from measurable outputs, within an inverse problem setting. This task has been traditionally approached through input or input-state estimation methods, such as delayed observers (Sundaram & Hadjicostis, 2007), Bayesian filtering methods (Gillijns & De Moor, 2007; Azam et al., 2017; Tatsis et al., 2021; Maes et al., 2016). The majority of input-state estimation methods, tend to adopt linear filters, such as the Augmented Kalman Filter (Lourens et al., 2012) or the Dual Kalman Filter (Azam et al., 2015) schemes. In this case, an explicit definition of the system's (linear) state space model of the dynamics is required, which hinder integration with deep learning frameworks. While nonlinear filter settings, e.g. those relying on an Unscented or Particle filter (Dertimanis et al., 2019), relax this setting, they may compromise prediction accuracy for the sake of online (real-time) estimation.

In this paper, we explore the input estimation problem from a new perspective, which operates in an offline manner, but is flexible in admitting general (even neural network-based) representations of the system dynamics. The proposed scheme aims to furnish a high precision reconstruction of the system's input, by formulating it as a Partially Observable Markov Decision Process (POMDP), with the primary difference to a typical decision-making problem pertaining to the definition of reward/cost functions. This reformulation enables the use of state-of-the-art model-based reinforcement learning algorithms for policy search. The cross-entropy method is chosen for its efficiency and robustness. The proposed approach assumes the underlying dynamics model is either known or learned.

2 BACKGROUND

2.1 PARTIALLY OBSERVABLE MARKOV DECISION PROCESS

A partially observable Markov decision process (POMDP) is defined by the tuple $(\mathcal{Z}, \mathcal{X}, \mathcal{U}, f, g, r)$, where $\mathcal{Z}$ is the state space, $\mathcal{X}$ is the observation space, and $\mathcal{U}$ is the action (input) space, while f, g and r are the respective transition, observation and reward functions (Sutton et al., 1998). A Markovian transition model is required for describing the system dynamics. One possible representation is offered by the following stochastic equations:

$$\begin{aligned} \mathbf{z}_t &= f(\mathbf{z}_{t-1}, \mathbf{u}_{t-1}) + w_t, \\ \mathbf{x}_t &= g(\mathbf{z}_t) + v_t, \\ \mathbf{r}_t &= r(\mathbf{x}_t, \mathbf{u}_{t-1}). \end{aligned} \tag{1}$$

The latent states, $\mathbf{z}_t$, evolve according to the transition function, f, for an imposed instantaneous input and a Gaussian random noise disturbance $w_t \sim \mathcal{N}(0, \mathbf{Q})$. Subsequently, we observe a noisy or partially observed quantity according to the observation function, g, as well contaminated with a Gaussian random noise $v_t \sim \mathcal{N}(0, \mathbf{R})$, reflecting measurement and modeling imprecision; finally, a reward, $\mathbf{r}_t$, is received based on the reward function r, the generated observation, and the imposed input. The ultimate goal is to search for a policy $\mathbf{u}_{1:T}$ that maximizes the sum of rewards (or minimizes costs).

At a high level, all standard reinforcement learning algorithms follow the same loop: the agent interacts with the POMDP by using a trail policy, which may or may not match the true policy, by observing the current observation $\mathbf{x}_t$, selecting an action $\mathbf{u}_t$, then observing the resulting next observation $\mathbf{x}_{t+1}$ and a reward value $\mathbf{r}_{t+1} = r(\mathbf{x}_{t+1}, \mathbf{u}_t)$. This procedure is repeated for multiple iterations, and the agent uses the observed tuple $(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}, \mathbf{r}_{t+1})$ to update its policy. In this paper, we propose to apply a similar framework for solving the input estimation (reconstruction) problem for dynamical systems, given an observed batch of data. In this case, the policies to be optimized are the system inputs.

2.2 THE UNKNOWN INPUT ESTIMATION PROBLEM RECAST AS A POMDP

To formulate the unknown input estimation problem as a POMDP, the reward function is defined as

$$r(\hat{\mathbf{u}}_t) = \|\hat{\mathbf{x}}_t - \mathbf{x}_t\|. \tag{2}$$

The objective is to find a candidate input $\hat{\mathbf{u}}_t$ that can minimize the difference between the observation that is generated from this candidate input and the true (measured) observation. The key idea is to define the reward function as the difference between the generated observation from the candidate input and the true observation, and then solve the POMDP based on a priori known or an inferred dynamics model.

Reformulating the input estimation problem as a POMDP offers two distinct advantages: 1) When the dynamics model is readily available, it is straightforward to solve the reformulated POMDP with different well-established model predictive control (MPC) algorithms, offering ample solution schemes. 2) When the dynamics model is not available, it is possible to apply model-based reinforcement learning (MBRL) methods, thus coupling this scheme with deep learning-based dynamics models in order to accomplish simultaneous input estimation and dynamics model learning. For the latter scenario, reinforcement learning is preferred over optimal control methods due to its ability to handle uncertainty in data and modeling errors arising from fitted deep learning models. In this work, we mainly investigate the use of the proposed framework for cases where the dynamics is known a priori or inferred (learned) by means of deep learning methods. The task of simultaneous estimation and learning is left for future work.

Algorithm 1 Unknown Input Estimation with cross-entropy method (CEM)

Input: H Input estimation horizon $\quad I$ Optimization iterations
K Candidates per iteration $\quad n$ Number of top candidates to be selected
Initialize the belief distribution over inputs $\mathbf{u}_{t:t+H} \sim N(\mathbf{0}, \mathbf{I})$
for iteration $i = 1, 2, ..., I$ **do**
  Draw a set of candidate solutions $\mathbf{u}^{(k)}_{t:t+H}$, $k = 1, ..., K$, from the current belief $N(\mu_i, \text{diag}(\sigma_i))$
  Evaluate the rewards r_k for $k = 1, ..., K$
  Rank the K rewards r_k and note their indices k an elite set $\mathcal{N} = \{k \in \{1, ..., K\} : r_k \text{ is one of the minimal } n \text{ rewards}\}$
  Update $\mu_{i+1} = \frac{1}{n} \sum_{k \in \mathcal{N}} \mathbf{u}^{(k)}_{t:t+H}$ and $\sigma_{i+1} = \frac{1}{n-1} \sum_{k \in \mathcal{N}} |\mathbf{u}^{(k)}_{t:t+H} - \mu_{i+1}|$
end for
return the first input $\mathbf{u}_t$

---

3 Simulation and Experimental Results

3.1 Road profile

A degrading road surface (pavement) condition reduces the driving comfort, induces disruptions compromising traffic safety, while further leading in substantial financial costs. Although high-accuracy road profilers equipped with lasers, inertia sensors, and cameras have been developed, their use is not practical for frequently evaluating the road network, owing to their high initial and operation costs. Lower cost alternatives often suffer from poorer precision, unless boosted with an appropriate processing or data analysis technique. Here, we suggest use of our suggested POMDP scheme, for high accuracy estimation of road roughness, in the form of input estimation.

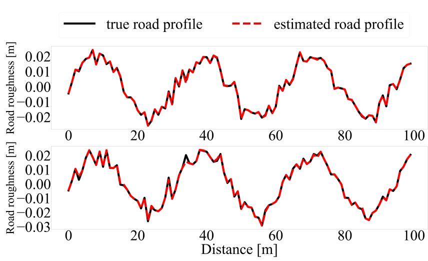

Figure 1: Road profile estimation.

In accounting for the interacting vehicle/road-surface system, a half-car model is used. The system matrices are detailed in Appendix A.3. We assume an accelerometer is mounted on the car body, allowing to track acceleration. By checking the invertibility condition stated in Appendix A.2, it is verified that a mere measurement of the acceleration of the car body is enough for identifying the input (i.e. the road profile). The road is simulated as a sinusoidal function, which represents large wavelength variation in terms of an uphill and downhill profile, with addition of random noise, for representing shorter-wavelength variation, corresponding to the local roughness of the road. The estimation results are shown in Fig. 1, with the average RMSE of 4.14×10^{-4} and the average R^2 of 0.9991. It can be observed that both large-scale patterns and the local roughness can be accurately identified, using only the acceleration of the car body, demonstrating the applicability and performance of the proposed method.

3.2 Wind turbine

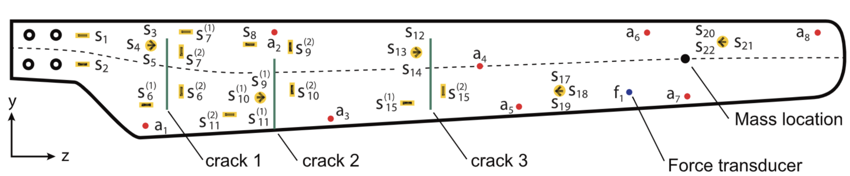

Figure 2: Position of sensors on the experimentally tested wind turbine blade. A total of eight accelerometers, a_i, are mounted on the blade, marked in red. Strain information is also collected, s_{ij}, but remains unused here. The figure is reused from (Ou et al., 2021).

We further demonstrate the value of the proposed input estimation framework in the context of SHM applications for wind energy infrastructure. We validate use of the proposed scheme for vibration monitoring of operational wind turbines. The data used in this paper were obtained and illustrated in (Ou et al., 2021) by experimentally testing a small-scale wind turbine blade. The sensor placement adopted during the experiments is shown in Fig. 2.

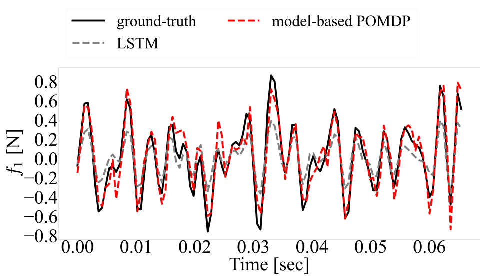

Figure 3: Input estimation of the wind turbine blade.

Neural Extended Kalman Filters (Neural EKF), adopted in (Liu et al., 2022) for learning structural dynamics (detailed in Appendix A.4), is a deep learning framework to capture the dynamics of complex systems. Here, we use Neural EKF as a means to demonstrating applicability of the herein suggested POMDP approach to input estimation, even for cases where a model is not known a priori (as was the case in the previous examples). Thus, the Neural EKF serves for recovering an underlying (latent) dynamics model under availability of data. With this learned model, we can use the presented POMDP framework to conduct model-based input estimation. In this example, we infer a dynamics model using Neural EKF based on the acceleration measurements collected from accelerometers a_1 to a_8, which serve as system outputs, while the data collected from the force transducer f_1 serve as information on the system input. For the training purpose, the system inputs are also required to be measured for learning an input-output dynamics model, and simultaneous model learning and input estimation will be considered for future work. Then, based on the learned dynamics model, we utilize the proposed POMDP approach to conduct input estimation on further test datasets, this time assuming that the input on f_1 is unmeasured (unknown) and to be inferred. The results are shown in Fig. 3. To compare the performance of our proposed method with other baseline models, we also trained a long short-term memory (LSTM) network. It is observed that our proposed method yielded significantly more accurate input estimation, with an RMSE of 0.129, while the RMSE for LSTM is 0.187. Our method is model-based, in contrast to LSTM which is model-free, but it does not require any prior knowledge of the model and relies solely on available data. We first estimate a deep learning-based dynamics model from the available data, and then use a POMDP to perform model-based input estimation. The results are accurate even for use of a learned dynamics model, which intrinsically contains modeling errors and approximations. Since in a real-world scenario, such modeling errors are ubiquitous, it is important to establish flexible inference schemes, which account for uncertainties. This is a trait of the proposed POMDP approach, where, further, the accumulated errors during the input estimation process do not grow to be unbounded.

4 Conclusion

In this work, we investigate input estimation for dynamical systems from a new perspective, by reformulating this as a Partially Observable Markov Decision Process (POMDP). The ground-truth system inputs are shown to be well-approximated by iteratively selecting those candidate inputs, whose corresponding outputs can best approximate the actual measurements, and then updating the belief distribution of the inputs. We show the applicability of the proposed methodology in theory and real-world applications, and adopt a straight forward algorithm, the cross-entropy method, to solve the reformulated POMDP. Different model-based reinforcement learning frameworks and dynamics modeling methods can be integrated into the proposed methodology. The results of this study demonstrate the potential of this new approach in improving the safety and performance of engineered systems in the field of structural health monitoring. This method can be leveraged to identify the loads acting on the structure during its operation, which can be used to ensure safety, but also for improving the design of structural systems. This work aims to set the idea of such a use case for POMDPs in place. The influence of various reinforcement learning methods and a thorough comparison against further input estimation frameworks are left for further work.

Acknowledgments

The research was conducted at the Future Resilient Systems at the Singapore-ETH Centre, which was established collaboratively between ETH Zurich and the National Research Foundation Singapore. This research is supported by the National Research Foundation Singapore (NRF) under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. The project is also supported by the Stavros Niarchos Foundation through the ETH Zurich Foundation and the ETH Zurich Postdoctoral Fellowship scheme.

References

Saeed Eftekhar Azam, Eleni Chatzi, and Costas Papadimitriou. A dual kalman filter approach for state estimation via output-only acceleration measurements. *Mechanical systems and signal processing*, 60:866–886, 2015.

Saeed Eftekhar Azam, Eleni Chatzi, Costas Papadimitriou, and Andrew Smyth. Experimental validation of the kalman-type filters for online and real-time state and input estimation. *Journal of vibration and control*, 23(15):2494–2519, 2017.

Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. *arXiv preprint arXiv:1506.02216*, 2015.

Vasilis K Dertimanis, EN Chatzi, S Eftekhar Azam, and Costas Papadimitriou. Input-state-parameter estimation of structural systems from limited output information. *Mechanical Systems and Signal Processing*, 126:711–746, 2019.

Marco Fraccaro, Simon Kamronn, Ulrich Paquet, and Ole Winther. A disentangled recognition and nonlinear dynamics model for unsupervised learning. *arXiv preprint arXiv:1710.05741*, 2017.

Steven Gillijns and Bart De Moor. Unbiased minimum-variance input and state estimation for linear discrete-time systems with direct feedthrough. *Automatica*, 43(5):934–937, 2007.

Laurent Girin, Simon Leglaive, Xiaoyu Bie, Julien Diard, Thomas Hueber, and Xavier Alameda-Pineda. Dynamical variational autoencoders: A comprehensive review. *arXiv preprint arXiv:2008.12595*, 2020.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017. URL `https://openreview.net/forum?id=Sy2fzU9gl`.

Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. In *International Conference on Learning Representations*, 2017. URL `https://openreview.net/forum?id=HyTqHL5xg`.

Rahul Krishnan, Uri Shalit, and David Sontag. Structured inference networks for nonlinear state space models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

Bohuš Leitner and Lucia Figuli. Fatigue life prediction of mechanical structures under stochastic loading. In *MATEC Web of Conferences*, volume 157, pp. 02024. EDP Sciences, 2018.

Wei Liu, Zhilu Lai, Kiran Bacsa, and Eleni Chatzi. Neural extended kalman filters for learning and predicting dynamics of structural systems. *arXiv preprint arXiv:2210.04165*, 2022.

E Lourens, Edwin Reynders, Guido De Roeck, Geert Degrande, and Geert Lombaert. An augmented kalman filter for force identification in structural dynamics. *Mechanical systems and signal processing*, 27:446–460, 2012.

Kristof Maes, AW Smyth, Guido De Roeck, and Geert Lombaert. Joint input-state estimation in structural dynamics. *Mechanical Systems and Signal Processing*, 70:445–466, 2016.

Yaowen Ou, Konstantinos E Tatsis, Vasilis K Dertimanis, Minas D Spiridonakos, and Eleni N Chatzi. Vibration-based monitoring of a small-scale wind turbine blade under varying climate conditions. part i: An experimental benchmark. *Structural Control and Health Monitoring*, 28 (6):e2660, 2021.

Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. *Advances in neural information processing systems*, 31:7785–7794, 2018.

Michael Sain and James Massey. Invertibility of linear time-invariant dynamical systems. *IEEE Transactions on automatic control*, 14(2):141–149, 1969.

Shreyas Sundaram and Christoforos N Hadjicostis. Delayed observers for linear systems with unknown inputs. *IEEE Transactions on Automatic Control*, 52(2):334–339, 2007.

Richard S Sutton, Andrew G Barto, et al. Introduction to reinforcement learning. 1998.

KE Tatsis, Vasilis K Dertimanis, Costas Papadimitriou, E Lourens, and EN Chatzi. A general substructure-based framework for input-state estimation using limited output measurements. *Mechanical Systems and Signal Processing*, 150:107223, 2021.

Silvia Vettori, Emilio Di Lorenzo, Bart Peeters, MM Luczak, and Eleni Chatzi. An adaptive-noise augmented kalman filter approach for input-state estimation in structural dynamics. *Mechanical Systems and Signal Processing*, 184:109654, 2023.

A Appendix

A.1 Additional results on a simulated structural system

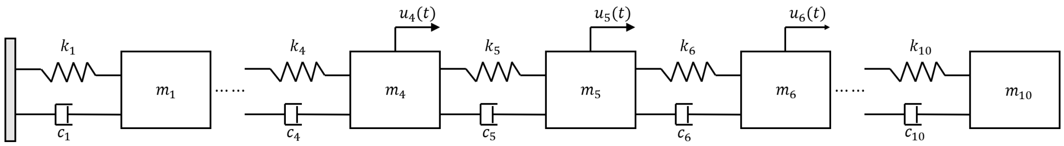

Figure 4: Illustration of the simulated 10-DOF structural system.

In this section, we implement the proposed framework on a simulated 10 Degree of Freedom (10-DOF) structural system, which is shown in Fig. 4. The structural system is governed by the following differential equations:

$$\mathbf{M}\ddot{\mathbf{q}}(t) + \mathbf{C}\dot{\mathbf{q}}(t) + \mathbf{K}\mathbf{q}(t) = \mathbf{S_u}\mathbf{u}(t), \tag{3}$$

where the displacement vector $\mathbf{q} = [q_1, ..., q_{10}]^T$; the mass matrix $\mathbf{M} = \text{diag}(m_1, ..., m_{10})$, and $m_1 = ... = m_{10} = 1$; the damping matrix $\mathbf{C} = \text{diag}(c_1, ..., c_{10})$, and $c_1 = ... = c_{10} = 1$; and the stiffness matrix

$$\mathbf{K} = \begin{bmatrix} k_1 + k_2 & -k_2 & 0 & \dots & 0 & 0 \\ -k_2 & k_2 + k_3 & -k_3 & \dots & 0 & 0 \\ 0 & -k_3 & k_3 + k_4 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & k_9 + k_{10} & -k_{10} \\ 0 & 0 & 0 & \dots & -k_{10} & k_{10} \end{bmatrix},$$

where $k_1 = ... = k_{10} = 1$. In order to embed this within the POMDP framework, the aforementioned differential equation can be brought in state-space form, by introducing the state vector $\mathbf{z} = [\mathbf{q}^T, \dot{\mathbf{q}}^T]^T$. Consequently, Eq. 3 can be rewritten in a first order ODE form:

$$\dot{\mathbf{z}}(t) = \mathbf{A}_c\mathbf{z}(t) + \mathbf{B}_c\mathbf{u}(t), \tag{4}$$

where the system matrices are

$$\mathbf{A}_c = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}\mathbf{C} \end{bmatrix}, \quad \mathbf{B}_c = \begin{bmatrix} \mathbf{0} \\ \mathbf{M}^{-1}\mathbf{S_u} \end{bmatrix}. \tag{5}$$

Eq. 4 can be further discretized in time, using for example a zero order hold (zoh) scheme, which leads into the following formulation:

$$\mathbf{z}_{t+1} = \mathbf{A}\mathbf{z}_t + \mathbf{B}\mathbf{u}_t \tag{6}$$

As for the observation equation, we consider the most general case, where a combination of the displacement, velocities and accelerations of this system can be measured, and thus the measurement vector can be formulated as the following form by using Eq. 3:

$$\mathbf{x}_t = \mathbf{C}\mathbf{z}_t + \mathbf{D}\mathbf{u}_t, \tag{7}$$

where

$$\mathbf{C} = \begin{bmatrix} \mathbf{S_d} & \mathbf{0} \\ \mathbf{0} & \mathbf{S_v} \\ \mathbf{S_a}\mathbf{M}^{-1}\mathbf{K} & \mathbf{S_a}\mathbf{M}^{-1}\mathbf{C} \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{S_a}\mathbf{M}^{-1}\mathbf{S_u} \end{bmatrix}. \tag{8}$$

Here, $\mathbf{S}_d$, $\mathbf{S}_v$ and $\mathbf{S}_a$ are the selection matrices that determine which DOFs in terms of available displacements, velocities and accelerations are measured.

We conduct a comprehensive investigation on the proposed method in terms of various traits: 1) input type (random, sinusoidal, and sine-swept inputs) and input localization; and 2) measurement availability.

To illustrate the workings of the proposed method, we evaluate the method for different input types, including random noise (sampled from $\mathcal{N}(0, 1)$), harmonic inputs (generated as $\sin(t)$) and swept sine inputs (generated by $\sin(t^2)$). We first set the fourth, fifth and sixth DOFs to be loaded with the same type of inputs, for the three different types respectively, while all other DOFs remain unloaded, as shown in Fig. 4, so as to investigate its capability of input localization. Then we also test scenarios where all three different types of inputs are applied to the fourth, fifth and sixth DOFs respectively and simultaneously, while other DOFs of the inputs are set to be zero. For a comparison across various input types, accelerations of all DOFs are assumed as available measurements. We evaluate the root mean square error (RMSE) for the simulation cases. The results are presented in the left side of Table 1. As an instance, the input estimation results for mixed-type inputs are plotted in Fig. 5. The results reveal that the proposed method can effectively tackle diverse input types, and can further estimate whether external inputs are applied on remaining DOFs, thus delivering input localization information.

Table 1: Performance for different input types

Input type	**RMSE**	**Measurement availability**	**RMSE**
Random noise	.0219	Displacement	.0105
Sine	.0068	Velocity	.0317
Sinesweep	.0064	Acceleration	.0553
Mixed	.0116	Mixed	.0341

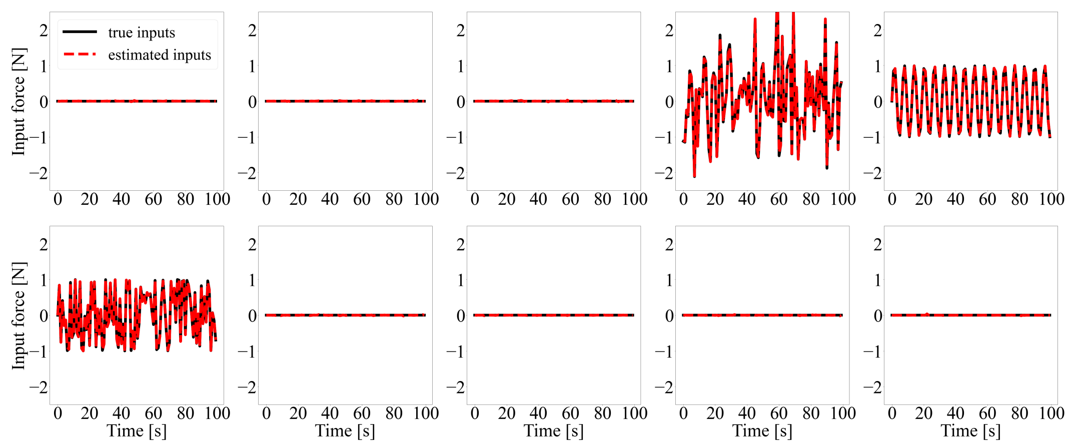

Figure 5: Input estimation results for scenario with mixed input types.

Additionally, we wish to investigate the efficacy of the proposed approach under availability of diversified measurements, i.e., different quantities from the possible set of displacement, velocity and acceleration outputs. In accounting for such a mixed measurement case, it is assumed that the displacements of the first 5 DOFs and the accelerations of the last 5 DOFs are measured. The results are presented in the right side of Table 1. It is observed that the proposed method achieved high accuracy for all the considered scenarios, which demonstrates the effectiveness of the proposed method under different availability of measurements.

A.2 INVERTIBILITY CONDITIONS

If the unknown input $\mathbf{u}$ can be uniquely identified from the output $\mathbf{x}$, the system is called invertible. The conditions for invertibility have already been well studied in the literature (Sain & Massey, 1969). Here, we state a sufficient and necessary condition for invertibility of linear dynamical systems with different kinds of measurements.

Theorem 1. *The system is invertible if and only if the matrix* $\mathbf{N}$ *has full column rank, where*

$$\mathbf{N} = \begin{bmatrix} \mathbf{D} & 0 & \dots & 0 \\ \mathbf{CB} & \mathbf{D} & \dots & 0 \\ \mathbf{CAB} & \mathbf{CB} & \dots & 0 \\ \vdots & \vdots & & \vdots \\ \mathbf{CA}^{n-1}\mathbf{B} & \mathbf{CA}^{n-2}\mathbf{B} & \dots & \mathbf{D} \\ \mathbf{CA}^{n}\mathbf{B} & \mathbf{CA}^{n-1}\mathbf{B} & \dots & \mathbf{CB} \\ \vdots & \vdots & & \vdots \\ \mathbf{CA}^{2n-1}\mathbf{B} & \mathbf{CA}^{2n-2}\mathbf{B} & \dots & \mathbf{CA}^{n-1}\mathbf{B} \end{bmatrix} \tag{9}$$

and n *is the dimension of the state space.*

If the system is not invertible, there would be infinitely many candidate inputs that can generate the same output, which makes the search for the true system inputs impossible, since the problem is intrinsically ill-conditioned.

A.3 Half-car model

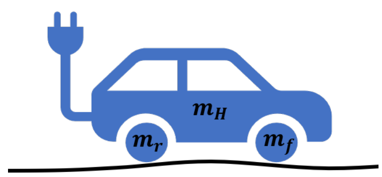

Figure 6: Road condition evaluation via input estimation algorithms. m_f, m_r, and m_H represent the masses of the front tire, rear tire, and the car body, respectively.

A half-car model is here used to simulate the vehicle dynamics. The state vector is defined as $\mathbf{z} = [\mathbf{z}_H, \theta, \mathbf{z}_f, \mathbf{z}_r]$, where $\mathbf{z}_H$, $\mathbf{z}_f$ and $\mathbf{z}_r$ are the displacements of the car body, front tire and rear tire, respectively, corresponding to the half car structure shown in Fig. 6; and θ is the pitching angle of the car body. The corresponding system matrices serve as the dynamics model describing the vehicle, which allows to conduct model-based input estimation for the road profile. The system matrices for the half-car model are as follows:

$$\begin{aligned} \mathbf{M} &= \begin{bmatrix} m_H & 0 & 0 & 0 \\ 0 & I_y & 0 & 0 \\ 0 & 0 & m_f & 0 \\ 0 & 0 & 0 & m_r \end{bmatrix}, \\ \mathbf{C} &= \begin{bmatrix} c_f + c_r & L_r c_r - L_f c_f & -c_f & -c_r \\ L_r c_r - L_f c_f & L_f^2 c_f + L_r^2 c_r & L_f c_f & -L_r c_r \\ -c_f & L_f c_f & c_f & 0 \\ -c_r & -L_r c_r & 0 & c_r \end{bmatrix}, \\ \mathbf{K} &= \begin{bmatrix} k_f + k_r & L_r k_r - L_f k_f & -k_f & -k_r \\ L_r k_r - L_f k_f & L_f^2 k_f + L_r^2 k_r & L_f k_f & -L_r k_r \\ -k_f & L_f k_f & k_f + k_{tf} & 0 \\ -k_r & -L_r k_r & 0 & k_r + k_{tr} \end{bmatrix}, \end{aligned} \tag{10}$$

where $m_H = 2200, I_y = 1100, m_f = 106, m_r = 152, c_f = c_r = 2500, k_f = 2 \times 10^4, k_r = 2.6 \times 10^4, k_{tf} = k_{tr} = 4 \times 10^5$.

A.4 Neural Extended Kalman Filters

Neural Extended Kalman Filters (Neural EKF), adopted in Liu et al. (2022) for learning structural dynamics, is a deep learning framework to capture the dynamics of complex systems. The framework can be described in the form of a nonlinear state space model:

$$\mathbf{z}_t = f_{\theta_\mathbf{t}}(\mathbf{z}_{t-1}, \mathbf{u}_{t-1}) + w_t, \ (transition) \tag{11}$$

$$\mathbf{x}_t = g_{\theta_\mathbf{o}}(\mathbf{z}_t) + v_t, \qquad (observation) \tag{12}$$

where $f_{\theta_\mathbf{t}}$ and $g_{\theta_\mathbf{o}}$ are learnable functions (i.e., not defined a priori) governing the transition and observation models, both parameterized by neural networks with parameters $\theta = \theta_\mathbf{t} \bigcup \theta_\mathbf{o}$. The process, w_t, and observation, v_t, noise sources are assumed to follow Gaussian distributions, with respective covariances set as learnable parameters during the training process.

It is worth noting that, while most of the dynamical VAE setups (Girin et al., 2020; Krishnan et al., 2017; Rangapuram et al., 2018; Fraccaro et al., 2017; Chung et al., 2015; Karl et al., 2017; Higgins et al., 2017), which extend VAEs to a dynamical version by considering the temporal evolution of the latent variables, or deep state space models, use an additional neural network as the inference model. However, here the inference model of the Neural EKF follows the format of an Extended Kalman Filter, which is a well-established Bayesian filter approximation for nonlinear systems of known nonlinear functions. The Neural EKF in essence extends the EKF framework to a learnable observer representation. This alleviates the need for deriving a separate inference network, q_ϕ, parameterized by a parameter vector ϕ that is independent of parameters within f_{θ_t} and g_{θ_o}. Since the objective ELBO largely depends on the goodness of reconstruction and inference, a separate inference network can weaken the training of the transition and observation models.

Within this framework, the model of Eq. 12 can be learned by maximizing an evidence lower bound (ELBO) of the data log-likelihood

$$\log p(\mathbf{x}) \geq \mathcal{L}(\theta, \phi; \mathbf{x}) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})}[\log p_{\theta_\mathbf{o}}(\mathbf{x}|\mathbf{z})] - \mathrm{KL}\big(q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})||p_{\theta_\mathbf{t}}(\mathbf{z},\mathbf{u})\big), \tag{13}$$

Since the posterior distributions $q_\theta(\mathbf{z}_t|\mathbf{x},\mathbf{u})$ can be computed in closed form by EKF, the above ELBO can be computed in a surrogate way as:

$$\begin{aligned}
\mathcal{L}(\theta; \mathbf{x}) = &-\frac{1}{2}\sum_{t=1}^{T}\Big[\log|\mathbf{C}_{t|T}\mathbf{\Sigma}_{t|T}\mathbf{C}_{t|T}^T + \mathbf{R}| \\
&+ (\mathbf{x}_t - g_{\theta_\mathbf{o}}(\boldsymbol{\mu}_{t|T}))^T(\mathbf{C}_{t|T}\mathbf{\Sigma}_{t|T}\mathbf{C}_{t|T}^T)^{-1}(\mathbf{x}_t - g_{\theta_\mathbf{o}}(\boldsymbol{\mu}_{t|T})) + d_x\log(2\pi) \\
&+ \log\frac{|\mathbf{A}_{t-1|T}\mathbf{\Sigma}_{t-1|T}\mathbf{A}_{t-1|T}^T + \mathbf{Q}|}{|\mathbf{\Sigma}_{t|T}|} - d_z + \mathrm{Tr}((\mathbf{A}_{t-1|T}\mathbf{\Sigma}_{t-1|T}\mathbf{A}_{t-1|T}^T + \mathbf{Q})^{-1}\mathbf{\Sigma}_{t|T}) \\
&+ (f_{\theta_\mathbf{t}}(\boldsymbol{\mu}_{t-1|T}) - \boldsymbol{\mu}_{t|T})^T(\mathbf{A}_{t-1|T}\Sigma_{t-1|T}\mathbf{A}_{t-1|T}^T + \mathbf{Q})^{-1}(f_{\theta_\mathbf{t}}(\boldsymbol{\mu}_{t-1|T}) - \boldsymbol{\mu}_{t|T})\Big],
\end{aligned} \tag{14}$$

where $\mathbf{A}_{\cdot|\cdot} = \frac{\partial f(\boldsymbol{\mu}_{\cdot|\cdot}, \mathbf{u}_{t-1})}{\partial \boldsymbol{\mu}_{\cdot|\cdot}}$ is the Jacobian of f at $\boldsymbol{\mu}_{\cdot|\cdot}$, and $\mathbf{C}_{\cdot|\cdot} = \frac{\partial g(\boldsymbol{\mu}_{\cdot|\cdot})}{\partial \boldsymbol{\mu}_{\cdot|\cdot}}$ is the Jacobian of g at $\boldsymbol{\mu}_{\cdot|\cdot}$.