Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Machine Learning for Receding Horizon Observer Design: Application to Traffic Density Estimation ⋆

Didier Georges *

* *Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, 38000 Grenoble, France,*
(e-mail: didier.georges@grenoble-inp.fr).

Abstract: This paper is devoted to the application of a simple machine learning technique for the design of a receding horizon state observer. The proposed approach is based on a neural network trained to learn the inverse problem consisting in deriving the current system state from past measurements and inputs. The training data is obtained from simple integrations of the system dynamics to be observed. The approach is here applied to the problem of estimating the car density on a highway online. A comparison with the solution of an receding horizon observer based on an adjoint method and used as reference demonstrates the effectiveness of the proposed approach.

Keywords: Receding horizon observers, neural-network-based machine learning, traffic monitoring, adjoint method.

1. INTRODUCTION

Getting some accurate estimates of the state of a dynamical system apart from a limited number of sensor measurements appears to be of the greatest importance in many monitoring applications, especially for risk assessment. The early detection of hazards or vulnerabilities and the ability to predict spatial and temporal evolution of dynamical hazards still remain a big challenge for disaster risk assessment and mitigation (see Zio (2018)). A large number of applications are concerned: Weather hazards, environmental monitoring of landslides, earthquakes, flooding, wildfires, air pollution, water quality, or crops; critical infrastructures monitoring (see Alonso (2018)): power, gas, water, oil, traffic and transportation networks; large-scale structural health monitoring : Engineered structures of bridges, buildings and other related infrastructures submitted to various stresses (earthquakes, structural ageing, attacks); human or animal health monitoring: epidemics and pandemics.

Our goal in this paper is to propose an effective way to derive a state observer that will provide a state estimate from a limited number of measurements, as an alternative to the Receding Horizon Observer (RHO) approach, which has been studied by many authors (see for instance, Michalska (1995), Muske (1995), Rawlings (2006), Alamir (2007), Kuhl (2011), Rangegowda (2018)). In most of the cases and especially for medium to large scale systems with small time constants, solving the RHO problem (relying on the online computation of a nonlinear least-square optimization problem) appears to be uncompatible with real-time operations. In addition non convexity of the nonlinear least-square problem may lead to unsatisfactory solutions corresponding to local minima.

The approach proposed in this paper relies on simple integrations of the system dynamics and offline training of a neural network. For that reason, the online computational cost is limited to the evaluation of the approximate observer function trained offline. Furthermore, under the assumption of system observability, the existence of non accurate estimates due to local minima of nonlinear least-square RHO is mitigated by the neural network training process. Indeed, the data used for training is based on accurate sample pairs of initial states / outputs and the final accuracy only depends on the quality of the training process. This approach is different from Alessandrini (2011) which computes an approximate solution to a least-square RHO by using a neural network.

The organization of the paper is as follows: Section 2 provides some background on receding horizon observers. Section 3 is devoted to the design of receding horizon observer based on the training of a neural network. In section 4, the application of the proposed methodology is applied to the vehicle density estimation based on the Lighthill-Whitham conservation law governing macroscopic dynamics of traffic flows. The overall performance of the machine learning approach is compared to the one of an optimal receding horizon observer based on the calculus of variations. Finally, the paper ends with some conclusions and perspectives.

2. SOME BACKGROUND ON RECEDING HORIZON OBSERVERS (RHO)

We consider the following class of nonlinear systems:

⋆ This work is supported by the French National Research Agency in the framework of the Investissements d'Avenir program (ANR-15-IDEX-02). Corresponding author D. Georges.

$$\dot{x} = F(x,u),\ x \in \mathbb{R}^n,\ x(0) = x_0,\ u \in \mathbb{R}^m \qquad (1)$$
$$y = H(x),\ y \in \mathbb{R}^p,$$

where x denotes the state vector, u is the vector of known exogenous inputs, and y is the vector of measured outputs.

Observability can be defined as the *injectivity* of the operator initial state to output (the observability function $y(t) = \Phi(t, x(t_0))$), that can be also reformulated as *the preservation of initial state distinctness or the nonzero output sensitivity to initial state (see Besançon (2007) for more details*).

$$\Phi(t, x_1(t_0)) = \Phi(t, x_2(t_0)),\ \forall t \in [t_0, t_0 + T] \Rightarrow x_1(t_0) = x_2(t_0). \qquad (2)$$

The state estimation problem consists in using a sequence of output measurements over a time interval to retrieve the state vector x. The measurements are assumed to be obtained from physical sensors.

The following notation will be used in what follows to denote the forward solution of (1) at time τ, starting from state x at time $t - T$:

$$X(\tau, x, u(.)),\ \tau \geq t - T.$$

State Receding or Moving Horizon Observers (RHO) (see Michalska (1995), Muske (1995), Rawlings (2006), Alamir (2007), Kuhl (2011), Rangegowda (2018)) provide an estimate of x by minimizing the output prediction error in the least-square sense over a past receding horizon defined by horizon length T, at each time t:

$$\{\hat{x}(t-T)\} = arg \min_{z \in \mathcal{X}(t-T)} \{J(t, z, y(.), u(.)) = \int_{t-T}^{t} \|y(\tau) - H(X(\tau, z, u(.)))\|^2_{R^{-1}} d\tau + \|z\|^2_{M^{-1}}\}, \qquad (3)$$

where weighting matrix R can be interpreted in the Bayesian framework as the covariance matrix of a noise vector affecting the output measurement. R^{-1} can also be used to reflect the degree of trust in the measurements. R can also depend on time to introduce some forgetting factor with respect to past measurements. M can be viewed as a regularization matrix or the covariance matrix of uncertain z. $\mathcal{X}(t)$ is the set of admissible values of the state at time t (often the set is used to impose the state remains positive, for instance, in the case of physical densities). In what follows, $\mathcal{X}(t)$ is assumed to be the full state-space (no state constraints).

It is worth mentioning that (3) provides in fact an estimate of state x at time t, since it suffices to integrate (1) from estimate state $\hat{x}(t - T)$, knowing $u(.)$. In practice, the continuous time t will be replace by a sampled time $t_k = kT_s$ defined by some sampling period $T_s > 0$.

Rather than solving optimization problem (3), which can be computationally expensive for online applications, the idea here is to approximate the inverse observability function Ψ defined by

$$x(t) = \Psi(y(.), u(.)), \qquad (4)$$

where both $y(.)$ and $u(.)$ are defined on time interval $[t - T, t]$.

If this function is available, $x(t)$ can be obtained online for every receding horizon $[t - T, t]$.

2.1 Differences between the proposed approach and RHO design

The two approaches are fundamentally different:

Solving RHO problem (3) provides an estimate of any state $x(t)$ for a given set of $(u(.), y(.)))$ defined on any interval $[t - T, t]$ by computing a solution of the related nonlinear least-square problem.

The here-proposed approach consists in estimating the inverse observability function in order to provide all the estimates of any $x(t)$ belonging to a compact set of $\mathbb{R}^n$ corresponding to inputs/outputs $(u(.), y(.)))$ belonging to a compact set of $L_2([0, T], \mathbb{R}^{m+p})$.

3. RHO DESIGN BASED ON SUPERVISED LEARNING

In order to approximate the inverse function Ψ, three main ingredients are needed:

(1) The definition of a **one-hidden-layer neural network**, with N neurons in the hidden layer with $(m + p) \times N_T$ input neurons and n output neurons (where N_T is the number of time samples), used to approximate $\Psi(y(.), u(.))$:

$$\Psi(Z) \approx \sum_{i=1}^{N} w_i \sigma(\alpha_i^T Z + b_i) + c_i = \Phi_\theta(Z), \qquad (5)$$

where

$$Z = (y(t_1), y(t_2), ..., y(t_{N_T}), u(t_1), u(t_2), ..., u(t_{N_T}))$$

is the vector of $2N_T$ sampled inputs and outputs defined on a grid of time interval $[0, T]$. $\sigma(x) = \dfrac{1}{1 + e^{-x}}$ or $tanh(x)$ represents the neuron activation function, and $w_i \in \mathbb{R}^m$, $\alpha_i \in \mathbb{R}^n$, $b_i \in \mathbb{R}$, $c_i \in \mathbb{R}$, and $\theta = (w_i, \alpha_i, b_i, c_i)_{i=1,...,N}$.

It has been shown in Cybenko (1989) that such networks can be used as multi-dimensional approximants. In this paper, best results have been obtained with *tanh* activation function.

The use of multi-layer networks seems not to provide significant improvements in the here-proposed application according to the preliminary experiments. However when the number of input neurons (corresponding to the size of the sequences $u(t_i)$ and $y(t_i)$) is greater than 500, the use of Convolutional Neural Networks should be considered, which are now well known as being able to effectively process large input and output data sets in classification and regression applications (see Goodfellow (2016)).

(2) Some **low-discrepancy sequences** such as the ones proposed by Halton, Sobol, Faure (see Niederreiter (1988) for instance) to generate M learning samples of initial state $x_k(T)$, and inputs $U_k = (u_k(t_1), u_k(t_2), ..., u_k(t_{N_T}))$, for $k = 1, ..., M$.

Such sequences have been proposed to solve the problem of optimally choosing M samples x_i defined in a hypercube $C = [0,1]^n$ to "minimize holes" in the sense of the best possible approximation of integrals:

$$|\frac{1}{M}\sum_{i=1}^{M} f(x_i) - \int_C f(x)dx| \leq W(f)\frac{log(M)^n}{M} \quad (6)$$

$W(f)$ is the variation of f in the sense of Hardy & Krause.

Here the hypercube is defined as

$$\prod_i^N [a_i, b_i] \times \prod_i^N [e_i, f_i] \quad (7)$$

, where (a_i, b_i)'s is the bound of the components of state x_k, and (e_i, f_i) is the bound of the inputs in U_k.

(3) **The generation of output samples**

$$Y_k(T) = (y_k(t_1), y_k(t_2), ..., y_k(t_{N_T}))$$

from each sample pair $(x_k(T)$**,** $U_k(T))$**,** $k = 1, ..., M$ via a simple time integration of system (1) on time interval $[0, T]$.

A supervised learning technique is finally used to tune parameters θ of the neural network, as solution of the following nonlinear regression problem, using M samples $Z_k = (Y_k(T), U_k(T))$, $k = 1, ..., M$ generated above:

$$\min_\theta \frac{1}{2}\sum_{k=1}^{M} \|\Psi_\theta(Z_k(T)) - x_k(T)\|^2. \quad (8)$$

Many approaches can be used to solve this problem (stochastic gradient, quasi-Newton ...) (see Mohri (2012) for instance). The Levenberg-Marquardt algorithm (Kelley (1999)) implemented in MATLAB Deep Learning toolbox has been used. Due to non convexity, several learning trials may be needed to get an effective approximation of Ψ.

4. AN ILLUSTRATIVE CASE STUDY: OPTIMAL CAR TRAFFIC DENSITY ESTIMATION

Various traffic estimation problems have focused the interest of researchers in the past decades (see the extensive review proposed in Seo (2017)). In this context, an adjoint-based approach was proposed in Nguyen (2016) for traffic density estimation on a road section based on the well-known hyperbolic PDE for macroscopic flow dynamics proposed in Lighthill (2011):

$$\partial_t \rho(t,x) + \partial_x \phi(\rho(t,x)) = 0, \quad (9)$$

with flux $\phi(\rho) = \rho V(\rho)$, and where $\rho(t,x)$ is the density of cars on the highway, and $V(\rho)$ is the car speed defined by the fundamental diagram:

$$V(\rho) = V_{max}(1 - \frac{\rho}{\rho_{max}}), \quad (10)$$

where is V_{max} is the maximum speed, and ρ_{max} is the maximum car density. Even though this basic model is known to have some limitations which makes it impossible to reproduce complex phenomena, such as capacity drop, hysteresis effect, traffic instability, stop-and-go waves, etc, (see Seo (2017)), it is interesting to use it for demonstration purpose of the here-proposed approach.

For the purpose of the observer design, a finite-dimensional model is derived using the method of lines based on an finite-difference upwind scheme defined on a 1D grid:

$$\dot{\rho}_i(t) = \frac{1}{dx}(\rho_{i-1}(t)V(\rho_{i-1}(t)) - \rho_i(t)V(\rho_i(t))), \quad i = 1, ..., n, \quad (11)$$

where $\rho_i(t)$ denotes $\rho(t, x = idx)$, and dx is the spatial size of the grid discretizing domain $[0, L]$.

Here traffic flows $\phi(\rho(t,0)) = \rho(t,0)V(\rho(t,0))$, $\phi(\rho(t,L)) = \rho(t,L)V(\rho(t,L))$ at both extremities of the highway section of length L are supposed to be measured.

The upstream flow acts as a Dirichlet boundary condition of conservation law (9).

Under these measurement assumptions, state observability of traffic density (and flows) of finite-dimensional model (11) is ensured, provided that maximum density of cars is not reached (traffic congestion).

The performance of the NN observer is compared to the one of an optimal observer (which is a least-square RHO) defined as the solution of the following problem (see Nguyen (2016) for a closely-related formulation):

$$\begin{aligned}\min_{\rho(t-T,x),v(\tau,x)} \frac{1}{2}\int_{t-T}^{t} \|y(\tau) - y_m(\tau)\|^2 d\tau \\ + \frac{1}{2}\int_0^L \|v(\tau,x)\|^2 dx d\tau \\ + \frac{\epsilon}{2}\int_0^L \|\rho(t-T,x) - \rho_g(t-T,x)\|^2 dx\end{aligned} \quad (12)$$

subject to

$$\begin{aligned}\partial_t \rho(\tau,x) + \partial_x \phi(\rho(\tau,x)) = v(\tau,x), \\ \phi(\rho(\tau,0)) = u(\tau), \\ y(\tau) = \phi(\rho(\tau,L)),\end{aligned} \quad (13)$$

where $y_m(\tau)$ denotes the flux measured at the downstream end, $v(\tau,x)$ is a corrective source used to take input flow disturbances into account, and ρ_g is an intial state guess. ϵ is a regularization term. Then $\rho(t,x)$ is obtained by integrating (13) from estimate $\rho(t-T,x)$, using $v(t,x)$ on a discretized grid of both spatial domain $[0,L]$ and time interval $[t-T,t]$. By using the variational calculus and following a similar derivation as the one in Nguyen (2016) based on Lagrangian functional

$$\begin{aligned}L(\rho,v,\lambda) = \frac{1}{2}\int_{t-T}^{t}[\|y(\tau) - y_m(\tau)\|^2 + \int_0^L \|v(\tau,x)\|^2 dx]d\tau \\ + \frac{\epsilon}{2}\int_0^L \|\rho(t-T,x) - \rho_g(t-T,x)\|^2 dx \\ + \int_{t-T}^{t}\int_0^L \lambda(t,c)(\partial_t\rho(\tau,x) + \partial_x\phi(\rho(\tau,x)) - v(\tau,x))dx d\tau,\end{aligned} \quad (14)$$

the solution of (12)-(13) can be obtained from the following set of necessary conditions for optimality, provided that regularity properties are met (with no shock occurrence):

$$\partial_t \rho(\tau,x) + \partial_x \phi(\rho(\tau,x)) = v(\tau,x), \quad (15)$$
$$\phi(\rho(\tau,0)) = u(\tau), \quad (16)$$
$$\epsilon(\rho(t-T,x) - \rho_g(t-T,x)) - \lambda(t-T,x) = 0, \quad (17)$$
$$v(\tau,x) - \lambda(\tau,x) = 0, \quad (18)$$
$$\partial_t \lambda(\tau,x) + \phi'(\rho(\tau,x))\partial_x \lambda(\tau,x) = 0, \quad (19)$$
$$\lambda(\tau,L) = -(\phi(\rho(\tau,L)) - y_m(\tau))\phi'(\rho(\tau,L)), \quad (20)$$
$$\lambda(t,x) = 0, \quad (21)$$

where the left parts of (17) and (18) provide the gradient of the cost function w.r.t. initial state $\rho(t-T,x)$ and $v(\tau,x)$ respectively, by using a descent method (for instance a quasi-newton method) combined with the use of the method of lines for the computation of the two PDEs. Here $\lambda(t,x)$ refers to the adjoint of $\rho(t,x)$. The necessary conditions are obtained after performing two integrations by parts (in time and space respectively) and then computing the directional derivative of Lagrangian functional L in directions $\delta\rho(\tau,x)$, $\delta\rho(t-T,x)$, $\delta\rho(t,x)$, $\delta\rho(\tau,L)$, and $\delta v(\tau,x)$. Necessary conditions for optimality (15)-(21) are derived by cancelling the gradients associated to each direction.

In this paper, the car density estimation of a highway of length 100 km is considered. The domain is discretized into 10 spatial sections ($n = 10$). The observation horizon is $T = 1$ h. The sampling time is $dt = 0.0256$ h, which corresponds to $N_T = 40$ time samples. Maximum vehicle density $\rho_{max} = 300$ $vehicles/km$ and maximum speed $V_{max} = 150$ km/h. Boundary conditions $\phi(\rho(t,0)) = u(t)$ and horizon T are both chosen to avoid appearance of shocks (singularities in the solution). Regularization coefficient ϵ is chosen equal to $1e^{-7}$ without measurement noise and $1e^{-3}$ with measurement noise. Initial guess $\rho_g(t-T,x)$ is equal to zero. MATLAB function *fminunc* has been used to compute the solution of (12)-(13).

4.1 Comparison without measurement noise

3000 samples (Z_k, x_k) are generated using Sobol's sequences defined in the hypercube $[0,170]^n \times [0,1e^4]^{N_T}$.

The hidden layer of the neural network contained 10 neurons only. The setting and training of the network were performed by using MATLAB functions *feedforwardnet* and *train*, respectively. Fig. (1) provides an example of the traffic dynamics generated by a training sample (x_i, U_i).

The trained network is validated using a set of 100 new validation inputs Z_k generated randomly. For each Z_k, $k = 1, ..., 100$, problem (12) is solved. Then the relative root-square error (RRSE) defined by

$$RRSE = \frac{\|\hat{x}(t) - x(t)\|_2}{\|x(t)\|_2},$$

where $x \in \mathbb{R}^n$ will denote the vector of the $\rho_i(t)$'s and $\hat{x}$ is the estimate provided by each of the two methods, is computed for each of the two methods (neural network observer (NN) and optimal observer (OO)). Fig. (2) shows the residuals of the regression function using both training and validation data, which demonstrates the effectiveness of the approach. Fig. (3) and Fig. (4) provide some examples of car density estimation. Fig. (5) clearly shows that the NN observer performs very well (with RRSE less than 3%) even though the optimal observer performs the best as expected in average. It is worth mentioning that sometimes the optimal observer gives worse results (with RRSE greater than 20 % ; those results are not included in the comparison to be fair). This seems mainly due to non convexity that induces local minima.

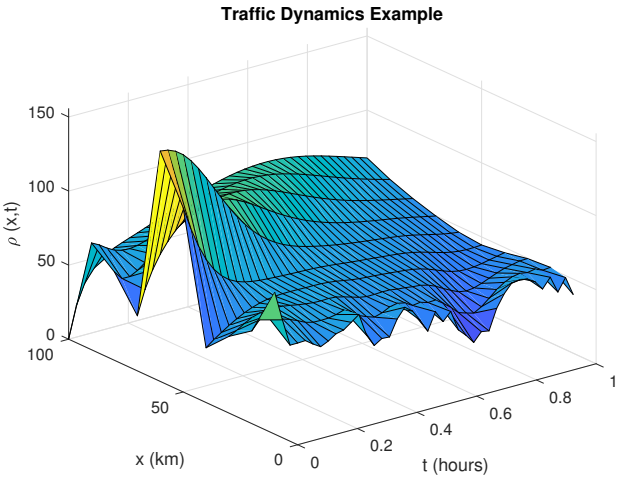

Fig. 1. Traffic dynamics from a training sample.

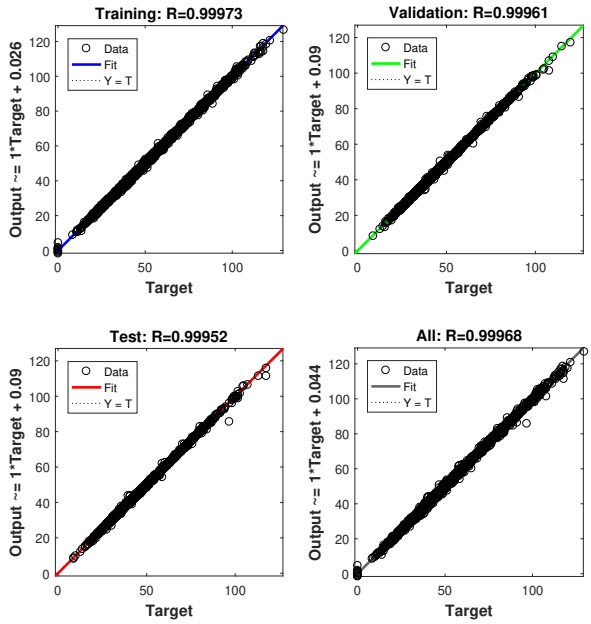

Fig. 2. Regression performance without Noise.

4.2 Comparison with measurement noise

The neural network is now trained with noisy sample pairs (Y_k, U_k) (a centered Gaussian noise with standard deviation 100 is added to each component of the training samples).

It appears that the NN observer still performs remarkably well in the presence of the noise as seen in both Fig. (6) and Fig. (9), although the optimal observer provides more accurate results in average (except again when local minima

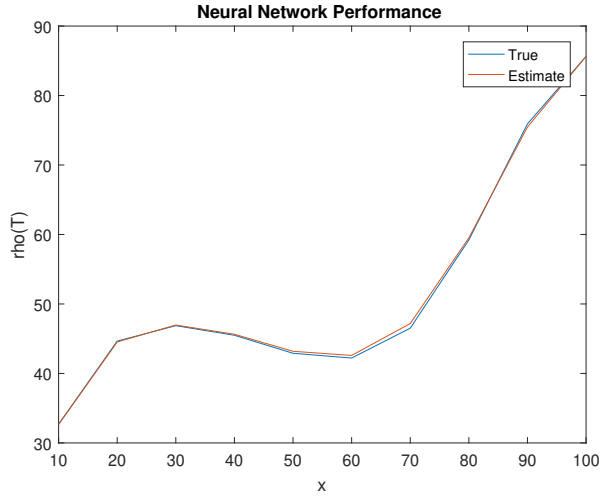

Fig. 3. Example with NN observer without noise.

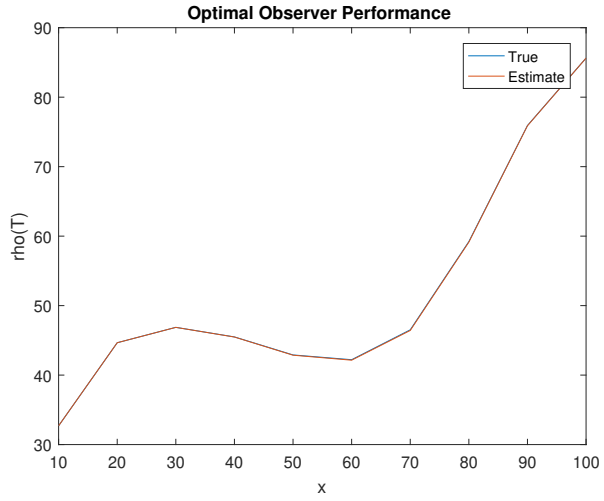

Fig. 4. Example with optimal observer without noise.

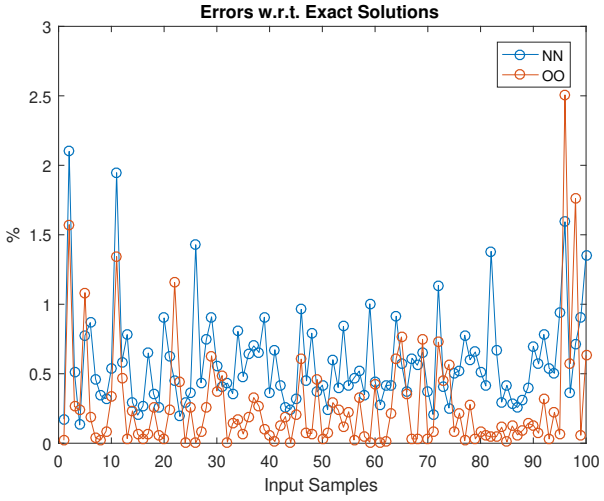

Fig. 5. RRSE comparison with 100 validation samples without noise.

are obtained). Fig. (7) and Fig. (8) show a comparison of the car density estimates obtained with the two methods in the presence of noise in the same conditions.

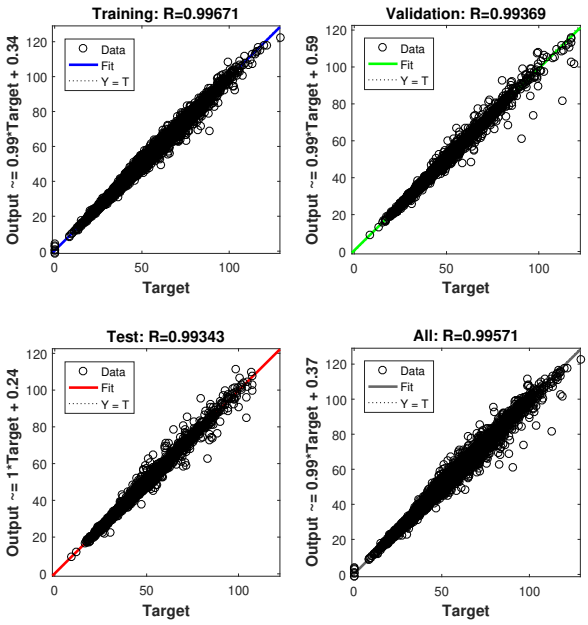

Fig. 6. Regression performance with noise.

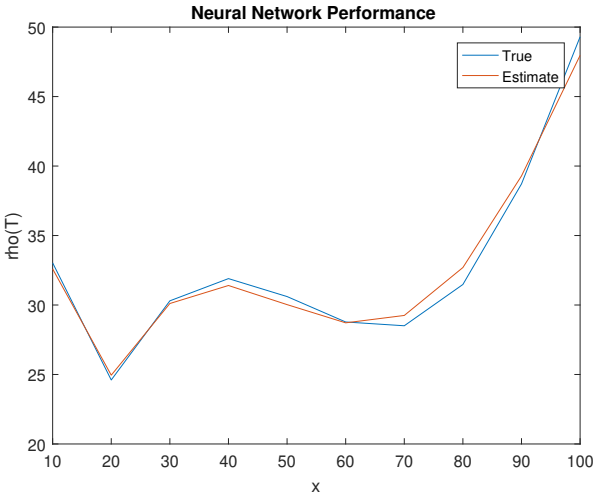

Fig. 7. Example with NN observer with noise.

5. CONCLUSIONS AND PERSPECTIVES

A simple approach for designing a Receding Horizon Observer has been proposed in this paper. A neural network has been used which makes it possible to approximate the inverse of the observability function. The training samples have been obtained from simple integrations of the system to be observed. The estimation of the traffic density over an highway section based on Lighthill-Whitham model, has been investigated to demonstrate the effectiveness of the proposed approach at least for medium-scale systems compared to an optimal observer design.

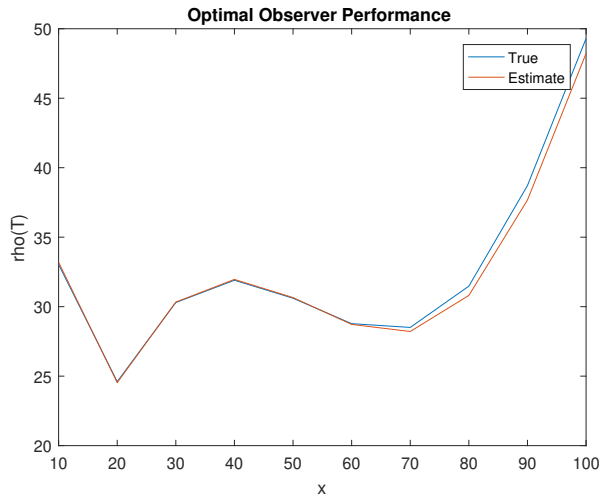

Fig. 8. Example with optimal observer with noise.

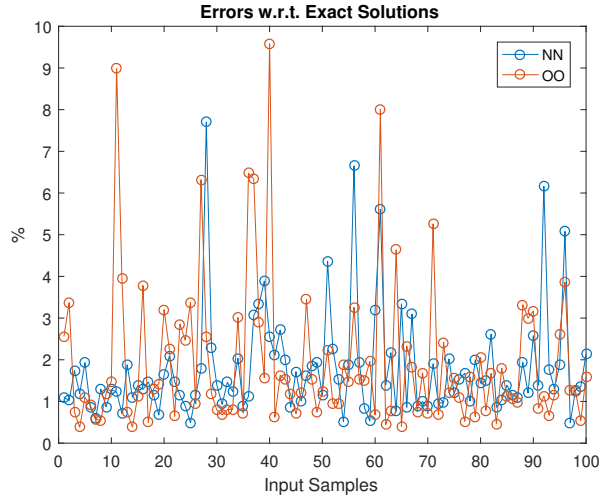

Fig. 9. RRSE comparison with 100 with noise.

Future works will be devoted to the extension of the approach to large-scale systems (for instance physical phenomena governed by 2D PDEs such as wildfires (Georges (2019)) or pollution spreading (Georges (2013))) using Deep Learning approaches such as Convolutional Neural Networks.

REFERENCES

M. Alamir (2007), Nonlinear Moving Horizon Observers: Theory & Real-Time Implementation. In Nonlinear Observers and Applications, Gildas Besançon (Ed). Lecture Notes on Communication and Information Science. Springer-Verlag-Series 2007.

A. Alessandrini, M. Baglietto, G. Battistelli, and M. Gaggero (2011), Moving-Horizon State Estimation for Nonlinear Systems Using Neural Networks, in IEEE Trans. on Neural Networks, Vol. 22, No. 5, pp. 768-780, May 2011.

L. Alonso, J. Barbaran, J. Chena, M. Diaz, L. Llopis, and B. Rubio (2018), Middleware and communication technologies for structural health monitoring of critical infrastructures: A survey, Computer Standards & Interfaces Volume 56, February 2018, Pages 83-100.

G. Besançon (Ed.) (2007), An Overview on Observer Tools for Nonlinear Systems. In Nonlinear Observers and Applications, Lecture Notes on Communication and Information Science. Springer-Verlag-Series 2007.

G. Cybenko (1989). Approximation by superpositions of a sigmoidal function. Math. Control Signals Systems,2 (1989), 303–314.

D. Georges (2013). Optimal Location of Mobile Sensors for Environmental Monitoring, 2013 - 12th biannual European Control Conference (ECC 2013), Switzerland (2013) [hal-00834642 - version 1].

D. Georges (2019). A Variational Calculus Approach to Wildfire Monitoring Using a Low-Discrepancy Sequence-Based Deployment of Sensors. 58th IEEE Conference on Decision and Control, December 11-13, 2019, Nice, France.

I. Goodfellow, Y. Bengio, and A. Courville (2016). Deep Learning, MIT Press, http://www.deeplearningbook.org.

M. Mohri, A. Rostamizadeh, and A. Talawalkar (2012). Foundations of Machine Learning, MIT Press.

H. Niederreiter (1988). Low-Discrepancy and Low-Dispersion Sequences, Journal of Number Theory 30: 51–70.

C.T. Kelley (1999). Iterative Methods for Optimization, SIAM Frontiers in Applied Mathematics, no 18, 1999, ISBN 0-89871-433-8.

P. Khul, M. Diehl, T. Kraus, J. P. Schloder, H. G. Bock (2011). A Real-Time Algorithm for Moving Horizon State and Parameter Estimation, Computers & Chemical Engineering Volume 35, Issue 1, 10 January 2011, Pages 71-83.

M. J. Lighthill, and G. B. Whitham (1955). On Kinematic Waves. II. A Theory of Traffic Flow on Long Crowded Roads. In Proceedings of the royal society of london. series a. mathematical and physical sciences: vol. 229 (pp. 317–345).

H. Michalska, and D. Q. Mayne (1995). Moving Horizon Observers and Observer-Based Control, IEEE Transactions on Automatic Control, VOL. 40, No. 6, June 1995.

K. Muske, and J.B. Rawlings, Nonlinear Receding Horizon State Estimation. In: Methods of Model-Based Control (R. Berber, Ed.). NATO-ASI Series. Kluwer Press. Dordrecht. Netherlands, 1995.

V.T. Nguyen, D. Georges, and G. Besançon (2016). State and Parameter Estimation in 1-d hyperbolic PDEs Based on an Adjoint Method. Automatica, 67, 185–191.

P. H. Rangegowda, J. Valluru, S. C. Patwardhan, and S. Mukhopadhya (2018). Simultaneous State and Parameter Estimation using Receding-horizon Nonlinear Kalman Filter, IFAC PapersOnLine 51-18 (2018) 411–416.

J. B. Rawlings, and B. R. Bakshib (2006). Particle Filtering and Moving Horizon Estimation, Computers & Chemical Engineering Volume 30, Issues 10–12, 12 September 2006, Pages 1529-1541

T. Seo, A. M. Bayen, T. Kusakabe, and Y. Asakura (2017). Traffic State Estimation on Highway: A comprehensive survey. Annual Reviews in Control 43 (2017) 128–151.

E. Zio (2018), The Future of Risk Assessment, Reliability Engineering & System Safety Volume 177, September 2018, Pages 176-190.