



HAL
open science

Towards Gain Tuning for Numerical KKL Observers

Mona Buisson-Fenet, Lukas Bahr, Valery Morgenthaler, Florent Di Meglio

► **To cite this version:**

Mona Buisson-Fenet, Lukas Bahr, Valery Morgenthaler, Florent Di Meglio. Towards Gain Tuning for Numerical KKL Observers. 22nd IFAC World Congress, Jul 2023, Yokohama, Japan. pp.4061 - 4067, 10.1016/j.ifacol.2023.10.1730 . hal-04410807

HAL Id: hal-04410807

<https://minesparis-psl.hal.science/hal-04410807>

Submitted on 22 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Gain Tuning for Numerical KKL Observers^{*}

Mona Buisson-Fenet^{*,**,*} Lukas Bahr^{**}
Valery Morgenthaler^{***} Florent Di Meglio^{*}

^{*} *Centre Automatique et Systèmes (CAS), Mines Paris – PSL, Paris, France (e-mail: mona.buisson@minesparis.psl.eu)*

^{**} *Institute for Data Science in Mechanical Engineering, RWTH Aachen University, Aachen, Germany*

^{***} *Ansys Research Team, Ansys France, Villeurbanne, France*

Abstract: This paper presents a first step towards tuning observers for general nonlinear systems. Relying on recent results around Kazantzis-Kravaris/Luenberger (KKL) observers, we propose an empirical criterion to guide the calibration of the observer, by trading off transient performance and sensitivity to measurement noise. We parametrize the gain matrix and evaluate this criterion over a family of observers for different parameter values. We then use neural networks to learn the mapping between the observer and the nonlinear system, and present a novel method to sample the state-space efficiently for nonlinear regression. We illustrate the merits of this approach in numerical simulations.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Nonlinear observers and filter design; Continuous time system estimation; Machine learning; Estimation and filtering; Observer design.

1. INTRODUCTION

In this paper, we propose a numerical method to calibrate Kazantzis-Kravaris-Luenberger (KKL) observers. The original design of Luenberger observers for linear systems can be found in Luenberger (1966). It consists in finding a linear mapping between the system dynamics and a linear filter of the measurement. Under appropriate observability assumptions and filter design, the Sylvester equation satisfied by the mapping has a unique injective solution. Its left-inverse, along with the filter, can be used to compute a convergent state estimate.

This design encompasses important degrees of freedom: the matrices defining the filter or, equivalently, the poles and zeros of the filter transfer function. To study their effect on state estimation performance, one must consider the effect of the mapping, which modifies the response, among others to measurement noise. For autonomous linear systems, the problem of tuning these degrees of freedom is essentially solved by the stationary Kalman Filter (Kalman and Bucy, 1961). Rather than directly assigning closed-loop eigenvalues, one can weigh the relative confidence in the measurement and the dynamic model and find the observer gains that are optimal for the metric defined by these weights.

The extension of these approaches to nonlinear systems is nontrivial. Indeed, there are few generic nonlinear observer designs; a review of these can be found in Bernard (2019); Bernard et al. (2022). Among the most commonly used are the High-Gain Observer (HGO) (Bornard and Hammouri, 1991; Khalil and Praly, 2014) and the Extended Kalman Filter (EKF) (Gelb, 1974). The EKF consists in linearizing

the observer dynamics around the current estimate to compute the optimal gain depending on chosen weights, akin to the linear case. There are, however, only local convergence guarantees (Krener, 2003). Conversely, the HGO relies on a change of variables to bring the system into canonical form, and high gains to “dominate” the Lipschitz constant of the nonlinearity. The stability guarantees come at the price of possibly poor transient performance (the so-called “peaking” phenomenon (Maggiore and Passino, 2003)) and high sensitivity to noise. While recent contributions aim at reducing these detrimental features thanks, e.g., to dynamic extensions (Astolfi et al., 2018), the question of gain tuning and performance criteria remains open. In particular, in Astolfi et al. (2018), the sensitivity to noise is examined a posteriori through numerous simulations.

In this paper, we develop a tuning methodology for Kazantzis-Kravaris/Luenberger (KKL) observers that does not rely on extensive tests, inspired by the Kalman filter or H_∞ control. The KKL design (Kazantzis and Kravaris, 1998; Andrieu and Praly, 2006) extends the results of Luenberger (1966) to nonlinear systems. It maps the system dynamics to a stable linear filter of the measured output, called the observer dynamics. The existence and injectivity of this mapping are guaranteed by mild observability conditions, which makes this design relatively generic. The contraction properties of the observer dynamics ensure convergence of the state estimates. The main challenge consists in computing said mapping and its left inverse, along with tuning the free parameters of the observer.

In Ramos et al. (2020), a method is proposed to approximate the mapping by performing nonlinear regression on datasets generated from trajectories of the system and the observer. Given fixed observer parameters, a neural

^{*} This work was supported by Ansys Inc.

network approximates the considered mapping, which is then used to compute state estimates from observer values.

In this paper, we build on the approach of Ramos et al. (2020) and propose a first step towards calibration of the observer. Our main contribution is a procedure to select the gain matrix using a tuning criterion that, in some sense, trades off the transient performance against the sensitivity to measurement noise. We start by setting this matrix based on a pre-defined filter, parametrized by its cut-off frequency. We then approximate the KKL mapping for different values of this parameter using neural networks, either independently or by learning the mapping as a function of the parameter. This approximation is enabled by appropriately sampling the state-space, improved upon Ramos et al. (2020). Computing the proposed criterion for all values of the parametrized gain matrix leads to an optimal calibration for the observer, in the sense of the proposed empirical criterion. Numerical simulations illustrate the approach.

The paper is organized as follows. In Sec. 2, we recall the main idea behind KKL observer design. In Sec. 3, we propose an empirical gain tuning criterion, then detail our numerical approach for state-space sampling, observer parametrization and nonlinear regression in Sec. 4. Finally, we illustrate the merits of the approach through numerical simulations in Sec. 5, before concluding in Sec. 6.

2. KKL OBSERVERS

Consider the autonomous nonlinear dynamical system

$$\begin{aligned} \dot{x} &= f(x) \\ y &= h(x) \end{aligned} \quad (1)$$

where $x \in \mathbb{R}^{d_x}$ is the state, $y \in \mathbb{R}^{d_y}$ is the measured output, f is a continuously differentiable function (C^1) and h is a continuous function. The goal of observer design is to compute an estimate of the state $x(t)$ from the knowledge of the past values of the output $y(s)$, $0 \leq s \leq t$. To ensure the feasibility of this task, KKL observers rely on the following two assumptions.

Assumption 1. There exists a compact set \mathcal{X} such that for any solution of interest x to (1), $x(t) \in \mathcal{X}$ for all $t \geq 0$.

Assumption 2. There exists an open bounded set \mathcal{O} containing \mathcal{X} such that (1) is *backward \mathcal{O} -distinguishable* on \mathcal{X} , namely for any trajectories x_a and x_b of (1) such that $(x_a(0), x_b(0)) \in \mathcal{X} \times \mathcal{X}$ and $x_a(0) \neq x_b(0)$, there exists $t < 0$ such that

$$h(x_a(t)) \neq h(x_b(t))$$

and $(x_a(\tau), x_b(\tau)) \in \mathcal{O} \times \mathcal{O}$ for all $\tau \in [t, 0]$. In other words, their respective outputs become different in backward finite time before leaving \mathcal{O} .

We now recall the following Theorem from Andrieu and Praly (2006) showing the existence of a KKL observer.

Theorem 1. (Andrieu and Praly (2006)). Suppose Assumptions 1 and 2 hold. Define $d_z = d_y(d_x + 1)$. Then, there exists $\ell > 0$ and a set S of zero measure in \mathbb{C}^{d_z} such that for any diagonalizable matrix $D \in \mathbb{C}^{d_z \times d_z}$ with eigenvalues $(\lambda_1, \dots, \lambda_{d_z})$ in $\mathbb{C}^{d_z} \setminus S$ with $\Re \lambda_i < -\ell$, and any $F \in \mathbb{C}^{d_z \times d_x}$ such that (D, F) is controllable, there exists a continuous injective mapping $\mathcal{T} : \mathbb{R}^{d_x} \rightarrow \mathbb{C}^{d_z}$ that satisfies the following equation on \mathcal{X}

$$\frac{\partial \mathcal{T}}{\partial x}(x) f(x) = D \mathcal{T}(x) + F h(x), \quad (2)$$

and its continuous pseudo-inverse $\mathcal{T}^* : \mathbb{C}^{d_z} \rightarrow \mathbb{R}^{d_x}$ such that the trajectories of (1) remaining in \mathcal{X} and any trajectory of

$$\dot{z} = Dz + Fy \quad (3)$$

satisfy

$$|z(t) - \mathcal{T}(x(t))| \leq M |z(0) - \mathcal{T}(x(0))| e^{-\lambda_{\min} t} \quad (4)$$

for some $M > 0$ and with

$$\lambda_{\min} = \min \{ |\Re \lambda_1|, \dots, |\Re \lambda_{d_z}| \}. \quad (5)$$

Due to the uniform continuity of \mathcal{T}^* , this yields:

$$\lim_{t \rightarrow +\infty} |\mathcal{T}^*(z(t)) - x(t)| = 0. \quad (6)$$

Note that according to this result, $z \in \mathbb{C}^{d_y(d_x+1)}$. Therefore, in order to represent this filter with real numbers only, we need $d_z = 2d_y(d_x + 1)$. However, in practice we assume that the $d_y(d_x + 1)$ complex eigenvalues needed for D are complex conjugates, such that we only need dimension $d_z = d_y(d_x + 1)$ to represent the real filter $z \in \mathbb{R}^{d_z}$.

Thus, implementing a KKL observer involves following the steps:

- (1) Choose matrices D and F
- (2) Compute the corresponding transformation \mathcal{T}^*
- (3) Simulate (3) from an arbitrary $z(0)$ and compute the estimate $\hat{x}(t) = \mathcal{T}^*(z(t))$.

In Ramos et al. (2020), a method to complete step 2 by performing nonlinear regression on trajectories of (1) and (3) is proposed. In the next section, we propose an approach to assist the user in completing step 1 by defining a performance criterion to optimize.

3. A GAIN TUNING CRITERION

Consider the dynamical system (1) and associated observer dynamics (3). Denote x, z their solutions starting respectively at $x(0)$ and $\mathcal{T}(x(0))$. Assume now that the measurement y is corrupted by an unknown noise vector $\epsilon \in \mathbb{R}^{d_y}$, so that $y(t) = h(x(t)) + \epsilon(t)$. Denote \hat{z} the corresponding solution of (3) starting at an arbitrary initial condition z_0 , and $\tilde{z} = \hat{z} - z$ the estimation error due to both the initial error and the measurement noise. In general, we aim to choose D such that the overall error on the estimated state \hat{x} is minimized, where $\hat{x} = \mathcal{T}^*(\hat{z})$, similarly to Henwood (2014). The following result provides a criterion for tuning D , which we then apply to the approximated transformation.

Proposition 1. Suppose Assumptions 1 and 2 are verified, such that Theorem 1 holds. Further, assume that \mathcal{T}^* is Lipschitz continuous of constant L . Then, we have

$$|\hat{x} - x|_{L^2} \leq L \left(|G_\epsilon|_\infty |\epsilon|_{L^2} + |G_z|_{H^2} |\tilde{z}(0)| \right) \quad (7)$$

where $|\cdot|$ is the Euclidean norm, $|\cdot|_{L^2}$ is the L^2 norm, and the H^2 respectively H_∞ norms are defined as

$$|G|_{H^2}^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} |G(j\omega)|^2 d\omega, \quad |G|_\infty = \sup_{\omega} |G(j\omega)| \quad (8)$$

with $G_\epsilon(s) = (sI_{d_z} - D)^{-1}F$ the transfer function from ϵ to \tilde{z} , and $G_z(s) = (sI_{d_z} - D)^{-1}$ from $\tilde{z}(0)$ to \tilde{z} .

Proof. By Lipschitz continuity of \mathcal{T}^* , we have

$$\begin{aligned} |\hat{x} - x|_{L^2}^2 &= \int_0^\infty |\mathcal{T}^*(\hat{z}(t)) - \mathcal{T}^*(z(t))|^2 dt \\ &\leq L^2 \|\hat{z}\|_{L^2}^2. \end{aligned} \quad (9)$$

The Laplace transform applied to the dynamics of \hat{z} yields

$$\begin{aligned} \hat{z}(s) &= (sI_{d_z} - D)^{-1} F \underline{\epsilon}(s) + (sI_{d_z} - D)^{-1} \hat{z}(0) \\ &= G_\epsilon(s) \underline{\epsilon}(s) + G_z(s) \hat{z}(0), \end{aligned} \quad (10)$$

where we denote the Laplace transform of a signal $f(t)$ by $\underline{f}(s)$. Applying standard results on signal norms for linear systems (Toivonen, 2010) yields

$$\|\hat{z}\|_{L^2} = \|\hat{z}\|_{L^2} \leq |G_\epsilon|_\infty \|\epsilon\|_{L^2} + |G_z|_{H^2} \|\hat{z}(0)\|. \quad (11)$$

Replacing (11) in (9) concludes the proof.

Proposition 1 exhibits a standard trade-off in linear system theory, between sensitivity to noise through the term in $\|\epsilon\|_{L^2}$ and convergence speed through the term in $\|\hat{z}(0)\|$. In this paper, we propose a heuristic that guides the choice of D such that the error on the estimate \hat{x} is minimized.

Remark 1. Proposition 1 relies on the assumption that \mathcal{T}^* is Lipschitz continuous. This is not true in general; however, we approximate \mathcal{T}^* with the neural network model \mathcal{T}_θ^* , which is Lipschitz if its activation function is and if its weights are bounded (Scaman and Virmaux, 2018). Its Lipschitz constant can be approximated empirically, for example by computing its maximum over a regular grid of n samples z_j . However, the maximum value is subject to outliers and tends to vary strongly between models.

In the light of this remark, we propose to monitor the following empirical criterion

$$\begin{aligned} \alpha(D) &:= |J| (|G_\epsilon|_\infty + |G_z|_{H^2}) \\ J &:= \left(\left| \frac{\partial \mathcal{T}^*}{\partial z}(z_j) \right| \right)_{j \in \{1, \dots, n\}} \end{aligned} \quad (12)$$

where we consider the l_2 -norm of J rather than its infinity norm. This is an approximate bound for $\|\hat{x} - x\|_{L^2}$. This heuristic trades off the transient through $|G_z|_{H^2}$, and the performance and noise sensitivity through $|G_\epsilon|_\infty$ and $|J|$. In our experiments, we consider a family of matrices D indexed by a scalar parameter ω_c . We then compute α for different $D(\omega_c)$ and pick the value of ω_c that minimizes it.

Remark 2. The bound (9) is conservative, and the choice of the L^2 norm is somewhat arbitrary. In practice, one could consider a variety of criteria by weighting different norms of $\frac{\partial \mathcal{T}^*}{\partial z}$, G_ϵ and G_z . For example, in the linear case where \mathcal{T} , \mathcal{T}^* are matrices, we have

$$\begin{aligned} \hat{x}(s) - \underline{x}(s) &= (sI_{d_x} - \mathcal{T}^* D \mathcal{T})^{-1} \mathcal{T}^* F \underline{\epsilon}(s) \\ &\quad + (sI_{d_x} - \mathcal{T}^* D \mathcal{T})^{-1} (\hat{x}(0) - x(0)). \end{aligned} \quad (13)$$

Another criterion could be the H_∞ norm of an analogy of this transfer function (13) for the nonlinear case using the empirical gradients. Note also that there are more advanced methods to estimate the Lipschitz constant of \mathcal{T}_θ^* (Scaman and Virmaux, 2018); we focus on the simpler criterion (12), which is enough to exhibit some of the trade-offs faced when tuning D .

In the next section, we present a method to improve the regression process by carefully generating the dataset and propose a possible parameterization of D , before illustrating the merits of the criterion in Sec. 5.




Fig. 1. Schematics of the learned KKL observer. First, we solve the ordinary differential equation (3) for the measurement y generated from the original system (1). Then, the estimate \hat{x} is computed as $\mathcal{T}_\theta^*(z)$, where \mathcal{T}_θ^* approximates \mathcal{T}^* .

4. NUMERICAL METHODS

As in Ramos et al. (2020), we approximate the transformation \mathcal{T}^* by a neural network¹ of weights θ . The resulting observer is illustrated in Fig. 1: we feed the measurement y into the observer dynamics (3), then apply the neural network model \mathcal{T}_θ^* . To train \mathcal{T}_θ^* , i.e. perform nonlinear regression, a dataset of N pairs (x_i, z_i) , $i \in \{1, \dots, N\}$ needs to be generated from trajectories of (1), (3). The construction of this dataset poses an important challenge, as the observer state z converges towards $\mathcal{T}(x)$ only after a transient period whose length depends on D . This transient is not suitable for gathering data to learn the transformation, since we do not have $x \simeq \mathcal{T}^*(z)$ during the transient. However, for autonomous nonlinear systems, the trajectories tend to converge towards the ω -limit sets (Rouche et al., 1977) of the dynamics, so that the points (x_i, z_i) after the transient tend to be close to these ω -limit sets, leading to an uninformative dataset. We solve this problem in Sec. 4.1.

Further, in order to calibrate the observer using the gain tuning criterion (12), a parametrization of the gain matrix D by a scalar ω_c is needed. Then, one can either learn a model \mathcal{T}_θ^* for each value of ω_c independently, or learn the transformation as a function of ω_c . This yields a harder regression problem, but avoids needing to learn a new transformation each time the pair (D, F) is changed. This is discussed in Sec. 4.2.

4.1 Backward-forward sampling

The choice of (x_i, z_i) pairs is critical to numerically approximate \mathcal{T}^* . In Ramos et al. (2020), inspired by Marconi and Praly (2008), the authors propose to first generate an arbitrary grid of initial conditions $(x(0), z(0))$ using standard statistical methods such as Latin Hypercube Sampling (LHS). Then, relying on the observer's stability, meaning that it forgets its arbitrary initial condition $z(0)$ after some time, the dynamics $x(t)$ and $z(t)$ are simulated forward in time for t_c , where t_c is chosen large enough such that $z(t_c)$ is "close" to its steady-state. Finally, the beginning of the solutions $(x(t), z(t))$ for $t < t_c$ is removed from the dataset.

¹ Note that \mathcal{T} can also be approximated using the same methodology, but is not necessary for state estimation.

Unfortunately, this approach lets the dynamics dictate the position of the (x_i, z_i) pairs: for large values of t_c , they are bound to be located close to the ω -limit sets of the system (Rouche et al., 1977). However, it is desirable to have training samples all over the state-space, especially in regions where the function \mathcal{T}^* are less smooth and therefore more difficult to approximate.

We propose the following methodology to generate an arbitrary dataset of (x_i, z_i) pairs.

- (1) Choose N initial conditions $x_i(0) \in \mathcal{X}$, $i \in \{1, \dots, N\}$ using a uniform grid, LHS sampling, or any other method.
- (2) Simulate the system $\dot{x} = f(x)$ from $x_i(0)$ backward in time for t_c seconds.
- (3) If the system diverges in backward finite time, then f should be saturated smoothly outside of \mathcal{X} as suggested in Andrieu and Praly (2006); Bernard and Andrieu (2019). An example of saturation is provided in Sec. 5.2.
- (4) Simulate both systems $\dot{x} = f(x)$ and $\dot{z} = Dz + Fy$ with $y = h(x)$, starting from $x_i(-t_c)$ obtained previously and $z_i(-t_c) = z_0$, where z_0 is an arbitrary initial condition, for t_c seconds forward in time.
- (5) Set the training dataset to $(x_i, z_i) = (x_i(0), z_i(0))$ as obtained from backward-forward simulation.

With this approach, the user can set the training points x_i a priori and obtain the corresponding z_i without the system dynamics modifying the desired state-space grid.

4.2 Parametrization of D

In order to evaluate the proposed gain tuning criterion, we parametrize D by a scalar ω_c . Several parametrizations can be considered, for example choosing D as a given diagonal matrix multiplied by a factor. In this paper, we propose to use a d_z -order Bessel filter with cut-off frequency $2\pi\omega_c$, while $F = \mathbf{1}_{d_z \times d_y}$ is fixed to guarantee the controllability of (D, F) . We choose D by setting its eigenvalues to be the filter's poles. For any set of poles (p_1, \dots, p_n) where p poles are real and m poles are complex conjugates such that $n = p + 2m$, we choose D as the following block-diagonal matrix:

$$D = \begin{pmatrix} D_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & D_{p+m} \end{pmatrix}, D_i = \begin{cases} p_i & \text{if } p_i \text{ is real} \\ \begin{pmatrix} \Re p_i & \Im p_i \\ -\Im p_i & \Re p_i \end{pmatrix} & \text{otherwise} \end{cases} \quad (14)$$

The choice of parametrization influences the performance of the obtained models; analyzing these different possibilities further could be an interesting topic for future work.

We can then compute the gain tuning criterion (12) for different values of ω_c , by learning a model \mathcal{T}_θ^* for each value of interest. However, this requires training several neural networks independently for each value of D , which can be tedious. Also, if the observer needs to be fine-tuned, a new model will be required. Instead, it is also possible to treat ω_c as an extra input to the network, so that the transformation to approximate is $\mathcal{T}_\theta^*(z, w_c)$. This yields a harder regression problem, so that training will require more data and a careful design, but also a single model for all values of D . The user can then choose an acceptable value of D for the use case at hand




Fig. 2. Proposed gain tuning criterion (12) for the reverse Duffing oscillator, divided by $n = 10,000$ points used to compute $\frac{\partial \mathcal{T}_\theta^*}{\partial z}(z_j)$. The infinity and H_2 norms are high for low values of ω_c , while the gradient of the approximate transformation is high for high values. Choosing $\omega_c = 0.15$ appears to be optimal with respect to this metric.

and directly use the previous model. Alternatively, they can train again for this specific value of D to obtain a more accurate approximation for this particular choice. This approach can be advantageous for low-dimensional problems or when the observer will be needed in different experimental conditions without re-training.

In the next section, we illustrate the relevance of criterion (12) for choosing D in numerical simulations, using the proposed sampling scheme and parametrization of D .

5. RESULTS

We now evaluate the proposed approach on simulations of two nonlinear systems². We demonstrate that D can be tuned a posteriori by optimizing a metric such as (12), and show that it is a relevant criterion for choosing D so as to limit the noise sensitivity of the state estimate. We learn the observer as a function of ω_c for the first system, independently for different values of ω_c for the second system. Note that the model can eventually be trained again after selecting ω_c to reach higher accuracy.³

5.1 Reverse Duffing oscillator

The reverse Duffing oscillator

$$\begin{cases} \dot{x}_1 = x_2^3 \\ \dot{x}_2 = -x_1 \end{cases} \quad y = x_1 \quad (15)$$

is a nonlinear system whose solutions evolve on invariant compact sets. We choose a set of hundred values of ω_{c_i} in $[0.03, 1]$. Then, LHS is used to select $N = 5,000$ samples $x_i \in [-1, 1]^2$ for each value of ω_c . The corresponding z_i samples are computed using backward-forward sampling

² Note that for our empirical criterion (12) to be meaningful, the variables should be normalized (Skogestad and Postlethwaite, 2005). In these academic examples, the variables can be considered scaled.

³ Code for reproducing the results is available at https://github.com/Centre-automatique-et-systemes/learn_observe_KKL.git.

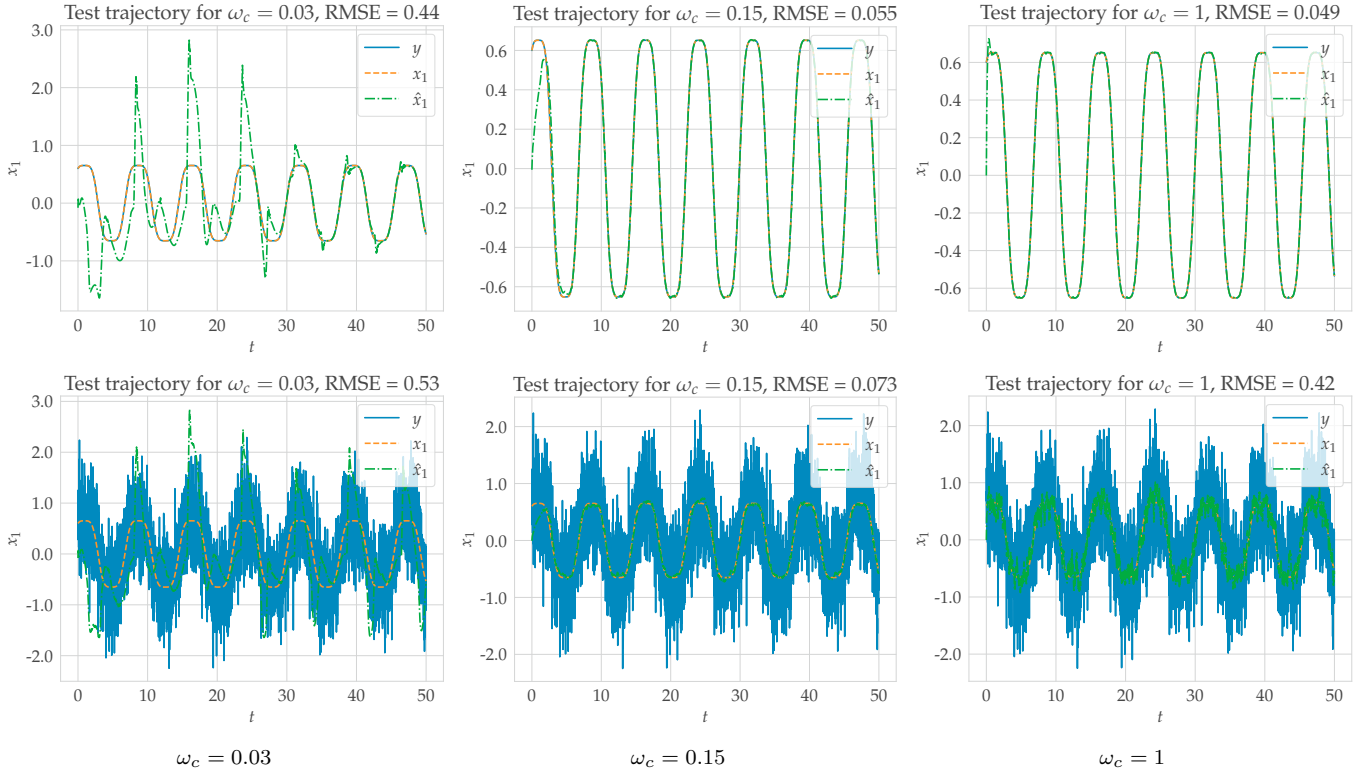


Fig. 3. Estimated trajectories of the reverse Duffing oscillator for $x(0) = (0.6, 0.6)$, without measurement noise at the top, with noise $\mathcal{N}(0, 0.5)$ at the bottom. For each setting, we compute the Root Mean Squared Error (RMSE) over the whole trajectory. For low ω_c (left), we observe long transients. For high ω_c (right), the estimate is sensitive to high-frequency noise. For $\omega_c = 0.15$ (middle), it is accurate and relatively robust to measurement noise.

as described in Sec. 4.1. The training data is normalized to ease the optimization process. For each given ω_c , D is computed following (14), while $F = (1 \ 1 \ 1)^T$. The time t_c after which we consider that the observer has converged is set to $\frac{10}{\lambda_{\min}(D)}$, where $\lambda_{\min}(D)$ is the minimum absolute value of the real part of the eigenvalues of D , such that it is different for each value of ω_c . The neural networks are multi-layer perceptrons with five hidden layers of size 50 and *SiLU* activation, which is Lipschitz continuous and shows good performance. We train \mathcal{T}_θ^* by minimizing

$$L(\theta) = \frac{1}{2} \sum_{x_i, z_i, \omega_{c_i}} |x_i - \mathcal{T}_\theta^*(z_i, \omega_{c_i})|^2. \quad (16)$$

We approximate $\mathcal{T}^*(z, \omega_c)$ over the training data as a function of ω_c , then compute the criterion (12) for each value of ω_c over a uniform grid of $n = 10,000$ test points z_j , also obtained with backward-forward sampling. The empirical criterion is shown in Fig. 2.

The choice of ω_c greatly influences the performance of the learned observer, as seen in Fig. 3. In our simulations, lower values of ω_c lead to a long convergence time and large overshoot, which corresponds to high values of $|G_z|_{H^2}$. However, low ω_c also yields a high signal to noise ratio in z , such that the observer is relatively robust to measurement noise. This is illustrated in the left column of Fig. 3. On the other hand, high values of ω_c lead to a high gradient of \mathcal{T}_θ^* : the approximate transformation is not smooth and therefore very sensitive to changes in z , hence to measurement noise. The signal to noise ratio in z is

also low due to the fast eigenvalues of D . This is depicted at the bottom right of Fig. 3. In the central column of Fig. 3, we select $\omega_c = 0.15$ the optimal value according to criterion (12). This setting yields an acceptable trade-off between these different aspects: both overshoot and noise sensitivity remain limited. Hence, the proposed gain turning criterion leads to satisfying performance for this use case.

5.2 Quanser Qube

We then consider simulations of a rotational inverted pendulum: the Qube Servo 2 by Quanser (2022), illustrated in Fig. 4. Its state of dimension four consists of two angles (θ_1, θ_2) and two angular velocities $(\dot{\theta}_1, \dot{\theta}_2)$; we measure $y = \theta_1$. Its trajectories diverge in finite backward time. Hence, as suggested in Andrieu and Praly (2006); Bernard and Andrieu (2019), we consider the modified system

$$\begin{aligned} \dot{x} &= f(x)g(x), \\ g(x) &= \begin{cases} 1 & \text{if } |x| \leq r \\ 0 & \text{if } |x| \geq r + d \\ p(|x| - r) & \text{otherwise} \end{cases} \quad (17) \end{aligned}$$

where f is the dynamics model of the Qube and g is a saturation function. We set $r = 50$ and $d = 100$. The function $p(\cdot)$ is a polynomial of order three chosen such that g be C^1 . This modified system has the same trajectories as the original system inside \mathcal{X} but does not blow up in backward time from any initial condition in \mathcal{X} .




Fig. 4. Qube Servo 2 by Quanser (2022).

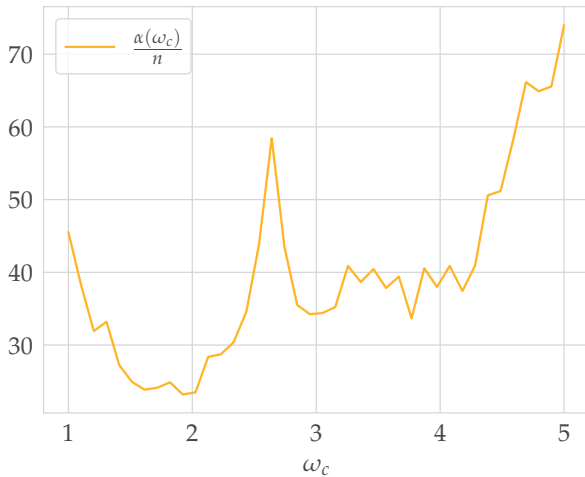


Fig. 5. Proposed gain tuning criterion (12) for the Qube, divided by $n = 10,000$. Choosing $\omega_c = 1.9$ appears to be optimal with respect to this metric.

Due to the curse of dimensionality, a large amount of data is necessary to learn \mathcal{T}^* with $d_x = 4$. In order to limit the computations, we generate data along realistic trajectories for the autonomous pendulum: we select 500 samples in a hypercube around the upward equilibrium position, use backward-forward sampling to obtain the corresponding z values, then run a joint simulation of both the x and z trajectories for 8s, sampled with time steps of 0.04s. This leads to $N = 10^6$ points for each of 41 values of $\omega_c \in [1, 5]$, for which we learn one model \mathcal{T}_θ^* independently.

We then compute the empirical criterion (12) for each value of ω_c independently on a grid of $n = 50,000$ points and obtain Fig. 5. The minimum is reached at $\omega_c = 1.9$, which again seems to be a good compromise between long transients and sensitivity to measurement noise. This is illustrated in Fig. 6: high values of ω_c lead to sensitivity to high frequency measurement noise, low values to long transients whenever the estimate is off, and $\omega_c = 1.9$ to an acceptable trade-off.

These results constitute a first step towards gain tuning for nonlinear observers. They can be considered as a proof of concept, showing that it is possible to tune the gains of KKL observers by parametrizing the gain matrix with a scalar ω_c then optimizing this scalar w.r.t. certain metrics. Note that many such metrics could be considered depending on the use case at hand. We propose the gain tuning criterion (12), which displays relevant aspects of the trade-off faced when choosing D as illustrated by our results, but other quantities could also be helpful.




Fig. 6. Estimation RMSE $|\hat{x} - x|$ for a simulated test trajectory of the Qube starting at $x(0) = (0.1, 0.1, 0, 0)$, with Gaussian noise $\mathcal{N}(0, 0.025)$ on the measurement. We observe a high sensitivity to noise for high values of ω_c (RMSE = 2.8 for $\omega_c = 5$). In contrast, the sensitivity to noise is lower for low values of ω_c , but we observe long transients, which lead to RMSE = 3.7 for $\omega_c = 1$. The value $\omega_c = 1.9$ compromises between performance and sensitivity to noise (RMSE = 1).

5.3 Discussion on the numerical results

Independently of the chosen parametrization of D or the strategy for generating the training data, approximating \mathcal{T}^* with neural networks bears the risk of overfitting. This risk is limited by monitoring the training loss (16) on a validation set and stopping the training early if this validation loss starts increasing, along with other standard techniques to restrict overfitting in supervised learning. For the Quanser Qube in Sec. 5.2, due to the difference between the simulation model used for generating the training data and the hardware, the performance of the learned observer decreases when used on experimental rather than simulated test trajectories. This illustrates that overfitting \mathcal{T}_θ^* to the particular system at hand, i.e. the robustness of the numerical KKL observer to model error, remains an issue. Recent works propose performance improvements of such observers (Niazi et al., 2022; Miao and Gatsis, 2022), which could help alleviate this. Fine-tuning the learned observer based on measurements could also help adapt it to the physical system at hand, for instance by retraining it inside an output predictor as in Janny et al. (2021).

6. CONCLUSION AND PERSPECTIVES

In this paper, we tackle the problem of gain tuning for KKL observers of autonomous nonlinear systems. We propose to numerically approximate the observer from simulation data, as introduced in Ramos et al. (2020), with an improved backward-forward sampling scheme. We parametrize the observer dynamics matrix D with a scalar ω_c , derive an empirical criterion for tuning it, and demonstrate on two numerical examples that it encompasses some relevant aspects of its influence on the performance.

We propose either to learn an observer for each value of ω_c of interest, or to directly learn a family of models that also takes this parameter as an input.

Similarly to Peralez and Nadri (2021); Lusch et al. (2018), it is also possible to learn a model of \mathcal{T} and \mathcal{T}^* jointly using an autoencoder structure, such that the latent variable z verifies (3). The cost function is then made up of a reconstruction loss and a loss on the PDE (2) verified by \mathcal{T} , such that an invertible solution to (2) is approximated on a grid of samples of x . This approach enables the user to optimize D jointly with the models \mathcal{T}_θ , \mathcal{T}_θ^* and to add terms to the cost functions to penalize other aspects, such as the criterion (12). However, it is also harder to train than the supervised approach. Further research aims at improving the accuracy of learning-based KKL observers, such as Niazi et al. (2022).

Many other questions remain open. As often in machine learning, it is unclear how to sample the state-space to generate the dataset optimally. Iterative active learning procedures can be envisioned, for example by learning the observer, then resampling in the parts of the state-space with the highest error, and learning again until the desired accuracy is achieved everywhere. Selecting the state-space grid a priori to achieve a given accuracy on the transformations could also be considered, as investigated in Marconi and Praly (2008). Extending KKL observers to nonautonomous systems is investigated in Bernard and Andrieu (2019); adapting the learning-based methodology to such systems is also a topic for future research.

ACKNOWLEDGEMENTS

The authors would like to thank Sebastian Giedyk at the Institute for Data Science in Mechanical Engineering (RWTH Aachen University) for his help with the experimental data. Thanks also to Pauline Bernard, Philippe Martin and Laurent Praly for all the fruitful discussions around this paper.

REFERENCES

- Andrieu, V. and Praly, L. (2006). On the existence of a Kazantzis-Kravaris/Luenberger observer. *SIAM Journal on Control and Optimization*, 45(2), 422–456.
- Astolfi, D., Marconi, L., Praly, L., and Teel, A.R. (2018). Low-power peaking-free high-gain observers. *Automatica*, 98, 169–179.
- Bernard, P. (2019). *Observer Design for Nonlinear Systems*. Springer International Publishing.
- Bernard, P. and Andrieu, V. (2019). Luenberger Observers for Nonautonomous Nonlinear Systems. *IEEE Transactions on Automatic Control*, 64(1), 270–281.
- Bernard, P., Andrieu, V., and Astolfi, D. (2022). Observer Design for Continuous-Time Dynamical Systems. *Annual Reviews in Control*, 53, 224–248.
- Bornard, G. and Hammouri, H. (1991). A high gain observer for a class of uniformly observable systems. In *Proceedings of the 30th IEEE Conference on Decision and Control*, 1494–1496.
- Gelb, A. (1974). *Applied optimal estimation*. MIT press.
- Henwood, N. (2014). *Estimation en ligne de paramètres de machines électriques pour véhicule en vue d'un suivi de la température de ses composants*. Ph.D. thesis, Mines ParisTech.
- Janny, S., Andrieu, V., Nadri, M., and Wolf, C. (2021). Deep KKL: Data-driven Output Prediction for Non-Linear Systems. In *Proceedings of the IEEE Conference on Decision and Control*.
- Kalman, R.E. and Bucy, R.S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83, 95–108.
- Kazantzis, N. and Kravaris, C. (1998). Nonlinear observer design using Lyapunov's auxiliary theorem. *Systems & Control Letters*, 34(5), 241–247.
- Khalil, H.K. and Praly, L. (2014). High-gain observers in nonlinear feedback control. *International Journal of Robust and Nonlinear Control*, 24(6), 993–1015.
- Krener, A.J. (2003). The convergence of the extended Kalman filter. In *Directions in mathematical systems theory and optimization*, 173–182. Springer.
- Luenberger, D. (1966). Observers for multivariable systems. *IEEE Transactions on Automatic Control*, 11(2), 190–197.
- Lusch, B., Kutz, J.N., and Brunton, S.L. (2018). Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1).
- Maggiore, M. and Passino, K.M. (2003). A separation principle for a class of non-UCO systems. *IEEE Transactions on Automatic Control*, 48(7), 1122–1133.
- Marconi, L. and Praly, L. (2008). Uniform practical nonlinear output regulation. *IEEE Transactions on Automatic Control*, 53(5), 1184–1202.
- Miao, K. and Gatsis, K. (2022). Learning Robust State Observers using Neural ODEs (longer version). *Preprint arXiv:2212.00866*.
- Niazi, M.U.B., Cao, J., Sun, X., Das, A., and Johansson, K.H. (2022). Learning-based Design of Luenberger Observers for Autonomous Nonlinear Systems. *Preprint arXiv:2210.01476*.
- Peralez, J. and Nadri, M. (2021). Deep Learning-based Luenberger observer design for discrete-time nonlinear systems. In *Proceedings of the IEEE Conference on Decision and Control*, 4370–4375. IEEE.
- Quanser (2022). Quanser courseware and resources. URL <https://www.quanser.com/products/qube-servo-2/>.
- Ramos, L.D.C., Meglio, F.D., Morgenthaler, V., da Silva, L.F.F., and Bernard, P. (2020). Numerical design of Luenberger observers for nonlinear systems. In *Proceedings of the 59th IEEE Conference on Decision and Control*, 5435–5442.
- Rouche, N., Habets, P., and Laloy, M. (1977). *Stability theory by Liapunov's direct method*. Springer.
- Scaman, K. and Virmaux, A. (2018). Lipschitz regularity of deep neural networks: Analysis and efficient estimation. *Advances in Neural Information Processing Systems* 32, 3835–3844.
- Skogestad, S. and Postlethwaite, I. (2005). *Multivariable Feedback Control: Analysis and Design*. John Wiley & Sons.
- Toivonen, H. (2010). Signal and system norms. *Lecture Notes for the Course "Advanced Control Methods"*. URL <http://users.abo.fi/htoivone/courses/robust/rob2.pdf>.