# A nonlinear KKL framework for theoretical analysis and guarantees of neural network observers

V. Andrieu

Works in collaboration with
P. Bernard, L. Brivadis, V. Pachy, L. Praly

# Approaching the observation problem

## The Observation Problem

Estimate the state variables $x$ from the measured variables $y$.

- The algorithm solving this problem is called an **observer**.
- It uses **a posteriori information**: the real-time measurements $y(t)$.
- It also uses **a priori information**: a mathematical model of the system.

$$\dot{x} = f(x), \quad y = h(x), \quad x \in \mathbb{R}^n$$

## Observer Principle

# Dynamic Observer Approach

## Principle

- Measurement history is stored in an internal, finite-dim. state ($z$).
- The state estimate $\hat{x}$ is a static **function** of this internal state.

## Workflow

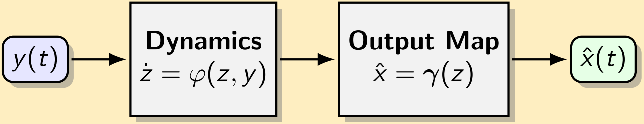


Key Question: How to design $\varphi$ and $\gamma$ for a good estimate?

Asking a computer science expert to solve the problem

The case of linear activation functions

The case of nonlinear activation functions

Why use nonlinear activation functions?

Conclusion

# A Popular Approach in Computer Science

### A Computer Science Observer Structure

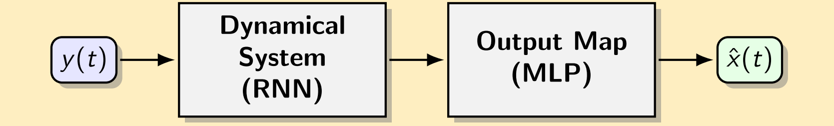


**RNN (Recurrent Neural Network)**

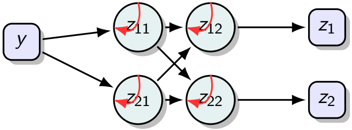


**MLP (Multilayer Perceptron)**

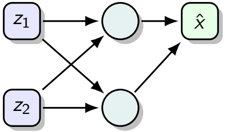

# A Universal Approach

- The structure (*RNN*, *MLP*) depends on two key elements:
  - Activation functions, denoted $\sigma$.
  - A set of **parameters** (weights, biases), denoted $\Omega$.

# A Universal Approach

- The structure (*RNN*, *MLP*) depends on two key elements:
  - Activation functions, denoted $\sigma$.
  - A set of **parameters** (weights, biases), denoted $\Omega$.
- A general continuous-time model for the RNN dynamics is given by:

$$\dot{z} = \mathbf{W_0}\sigma(\mathbf{W_1}z + \mathbf{W_2}y + \mathbf{b})$$

## The Computer Science Method: Supervised Learning

The parameters $\Omega$ are typically "learned" from data/model in two steps:

1. Define a **cost function** that quantifies the estimation error (e.g., $|\hat{x} - x|^2$).
2. Optimize the parameters $\Omega$ to minimize this cost, usually via gradient descent.

## The Control Theory Perspective

This data-driven approach often works, but it raises crucial questions: Can we give a formal **guarantee** of convergence? Is the observer **tunable**?

# Tunable Observers

## Definition: A Tunable Observer Structure

A structure is called **tunable** if for any desired precision ($\epsilon > 0$) and convergence time ($t_o > 0$)...

Given: A compact set of initial states $\mathcal{X}$, a time $t_o$, and a threshold $\epsilon$.

# Tunable Observers

## Definition: A Tunable Observer Structure

A structure is called **tunable** if for any desired precision ($\epsilon > 0$) and convergence time ($t_o > 0$)...

Given: A compact set of initial states $\mathcal{X}$, a time $t_o$, and a threshold $\epsilon$.

...we can prove the existence of parameters $\Omega$ that provide the following guarantee:

Guarantee: There **exist** parameters $\Omega$ such that for any initial condition in a compact set:

$$|\hat{x}(t) - x(t)| \leq \epsilon, \quad \forall t > t_o$$

## The Central Question of this Talk

For which classes of activation functions $\sigma$ can we formally prove this existence guarantee?

# A particular case: Linear activation

Taking a linear activation function $\sigma(v) = v$ in the *RNN* and choosing specific weights:

$$\dot{z}_i = W_0 \sigma(W_1 z_i + W_2 y + W_3) \quad \Rightarrow \quad \dot{z}_i = k\lambda_i z_i + y, \quad i = 1, \dots, m$$

$\Rightarrow$ We recognize the dynamics of a KKL observer.

## KKL Paradigm:

*If the system is observable, by picking m sufficiently large, there exists a map* $\mathbf{T}^{\mathrm{inv}} : \mathbb{R}^m \mapsto \mathbb{R}^n$ *such that* $\hat{x}(t) = \mathbf{T}^{\mathrm{inv}}(z(t))$ *gives an asymptotic observer!*

- Local version: Shoshitaishvili (1990), Kazantzis-Kravaris (1998)
- Global version: Kreisselmeier-Engel (2003), Andrieu-Praly (2006), Brivadis-Andrieu-Bernard-Serres (2022)
- Time-varying version: Bernard-Andrieu (2019)
- Discrete-time version: Tran-Bernard (2024)

# KKL Observers: Step 1

Given $m$ linear filters:

$$\dot{z}_i = k\lambda_i z_i + y \ , \ \ k > 0, \lambda_i < 0 \quad i = 1, \dots, m$$

The filter's state $z(t)$ converges to a function of the system's state $x(t)$ (for bounded trajectories).

**Theorem (VA-Praly, 2006)**

There exists a $C^0$ map $\mathbf{T}_k : \mathbb{R}^n \mapsto \mathbb{R}^m$ such that for a constant $C$:

$$|z(t) - \mathbf{T}_k(x(t))| \leq Ce^{-k \min_i |\lambda_i| t} |z_0 - \mathbf{T}_k(x_0)|$$

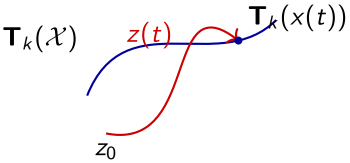


$\Rightarrow$ If $\mathbf{T}_k$ is invertible, we can recover $x$ from $z$!

# KKL observers: Invertibility of $\mathbf{T}_k$

Step 2: Ensure $\mathbf{T}_k$ is invertible by choosing $k$ large enough.

**Assumption: Differential observability on $\mathcal{X}$**

There exists an integer $m \geq 1$ such that the map $\mathbf{H}_m : \mathbb{R}^n \to \mathbb{R}^m$ defined by:

$$\mathbf{H}_m : x \mapsto \begin{pmatrix} h(x) & L_f h(x) & \dots & L_f^{m-1} h(x) \end{pmatrix}^\top$$

is Lipschitz injective on $\mathcal{X}$.

**Theorem (Andrieu-Praly, 2006; Andrieu, 2014)**

Let $\mathcal{X} \subset \mathbb{R}^n$ be a compact invariant set. Under the observability assumption, there exists $k^* > 0$ such that for all $k \geq k^*$, the map $\mathbf{T}_k$ is $C^1$ and Lipschitz injective.

If $\mathbf{T}_k$ is injective, there exists an inverse map $\mathbf{T}^{\mathrm{inv}}$ such that $\mathbf{T}^{\mathrm{inv}}(\mathbf{T}_k(x)) = x$.

# KKL observers: The final result

An (asymptotic) observer is given by:

$$\hat{x}(t) = \mathbf{T}^{\mathrm{inv}}(z(t)), \quad \dot{z}_i = k\lambda_i z_i + y, \quad i = 1, \ldots, m$$

## Theorem (Andrieu, 2014)

Let $\mathcal{X} \subset \mathbb{R}^n$ be a compact invariant set. There exists $k^* > 0$ such that for all $k \geq k^*$, there exists a $C^1$ mapping $\mathbf{T}^{\mathrm{inv}} : \mathbb{R}^m \mapsto \mathbb{R}^n$ and a constant $C$ such that

$$|\mathbf{T}^{\mathrm{inv}}(z(t)) - x(t)| \leq Ce^{-k \min_i |\lambda_i| t}(|z_0 - \mathbf{T}_k(x_0)|), \quad \forall (z_0, x_0) \in \mathbb{R}^m \times \mathcal{X}.$$

$\Rightarrow$ For each $(\epsilon, t_o)$, there exists $k^*$ such that for all $k \geq k^*$:

$$|\mathbf{T}^{\mathrm{inv}}(z(t)) - x(t)| \leq \epsilon, \quad \forall t > t_o, \quad \forall (x_0, z_0) \in \mathcal{X} \times \mathcal{Z}_0$$

$\Rightarrow$ We have a **tunable** asymptotic observer.

Question: How do we compute $\mathbf{T}^{\mathrm{inv}}$?

# MLP as an Approximator of $\mathbf{T}^{\text{inv}}$

The KKL observer provides a theoretical map $\mathbf{T}^{\text{inv}}$, but it is generally impossible to compute analytically.

However, since $\mathbf{T}^{\text{inv}}$ is a smooth function ($C^1$), we can approximate it!

## Universal Approximation Theorem (Cybenko, 1989)

An MLP can approximate any continuous function to any desired precision $\epsilon$ on a compact set.

## Consequence on the Total Error

The total estimation error can be split into two parts:

$$|\hat{x}(t) - x(t)| \leq \underbrace{|\gamma(z(t)) - \mathbf{T}^{\text{inv}}(z(t))|}_{\text{Approximation Error } (\leq \epsilon)} + \underbrace{|\mathbf{T}^{\text{inv}}(z(t)) - x(t)|}_{\text{Convergence Error}}$$

We control the first term by augmenting the MLP, and the second by tuning the observer gain $k$.

# Conclusion for the Linear Case: A Tunable Structure

By combining the KKL linear filters with an MLP as a universal approximator, we obtain a complete and practical observer structure.

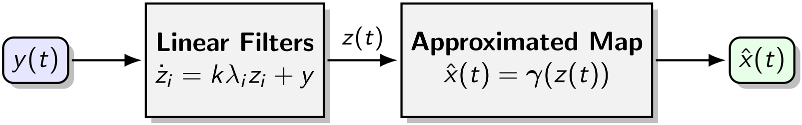


## Main Conclusion

The combined **Linear Filter + MLP** architecture is a tunable observer structure. It possesses theoretical convergence guarantees while being practically implementable.

# Starting Point: RNNs as Contracting Systems

Recent work has provided a crucial bridge between Recurrent Neural Networks and control theory.

## Key Result (e.g., Galimberti et al. 2023)

Under certain conditions on the activation functions ($\sigma$) and weight matrices ($\Omega$), it is possible to guarantee that a continuous-time RNN behaves as a **contracting system**.

**RNN with specific $\sigma, \Omega$** $\Longrightarrow$ **A Guaranteed Contracting System**

## Our Approach

We model the RNN part of our observer as a general contracting system, making the contraction rate tunable with a high gain parameter $k$.

# Our Working Hypothesis: The Model

We formalize the observer dynamics as:

$$\dot{z} = k \cdot g(z, y)$$

with $k \gg 1$ (high gain) and a base dynamics $g(z, y)$.

## Key Assumptions on $g(z, y)$

1. Contraction Property: Ensures stability and convergence to a unique solution.
$$\frac{\partial g}{\partial z} + \left(\frac{\partial g}{\partial z}\right)^\top \leq -2I_m$$
2. Sufficient Smoothness: The partial derivatives of $g$ are assumed to be bounded.

Allows analysis via Contraction Theory.

# Nonlinear Case, Step 1: Convergence

## Principle (from Contraction Theory)

- For any bounded input $y(t)$, our contracting filter has a **unique, exponentially attractive, bounded solution**: $\boldsymbol{\theta}_*(t)$.
- *Ref: Pavlov et al., 2004, Praly 2025 for its regularity*

## Our Definition of the Map $\mathbf{T}_k$

We define our map by identifying it with this unique solution:

$$\mathbf{T}_k(x(t)) := \boldsymbol{\theta}_*(t)$$

## Formal Result (Andrieu, Bernard, Brivadis, Praly, 2025)

This construction yields the exponential convergence guarantee:

$$|z(t) - \mathbf{T}_k(x(t))| \leq Ce^{-\alpha k \min_i |\lambda_i| t}|z_0 - \mathbf{T}_k(x_0)|.$$

# Step 2, Part A: The Filter Rank Condition

In addition to system observability, we need a structural condition on the filter itself to guarantee injectivity.

## Assumption 2: Filter Rank Condition

The filter's base dynamics $g(z, y)$ must have a sufficiently "rich" structure.

- Let $\varphi_0(y)$ be the unique solution to $g(\varphi_0(y), y) = 0$.
- We construct a matrix $C(y)$ from the Jacobians of $g$ evaluated at this point:

$$C(y) = \left( \dots \left[ \frac{\partial g}{\partial z} \right]^{-i} \frac{\partial g}{\partial y} \dots \right)_{i=1..m-1}$$

- **Condition:** This matrix $C(y)$ must be **left-invertible**.

## Practical Implication for RNNs

This condition, while technical, is not restrictive. It can be **generically satisfied** by an appropriate choice of the RNN's weight matrices ($\Omega$).

# Step 2, Part B: The Injectivity Result

With both assumptions (System Observability & Filter Rank) now in place, we obtain the main injectivity theorem.

## Theorem (Andrieu, Bernard, Brivadis, Praly, 2025)

There exists $k^* > 0$ such that for all $k \geq k^*$:

- ✓ The map $\mathbf{T}_k$ becomes $C^1$ and **Lipschitz injective**.
- ✓ This guarantees the existence of a stable inverse map $\mathbf{T}^{\mathrm{inv}}$.

## Final Consequence

The existence of $\mathbf{T}^{\mathrm{inv}}$ allows us to define the observer and prove its exponential convergence:

$$|\hat{x}(t) - x(t)| \leq C e^{-\alpha k \min_i |\lambda_i| t} |z_0 - \mathbf{T}_k(x_0)|$$

# Conclusion for the Nonlinear Case

- Like the linear case, the map $\mathbf{T}^{\mathrm{inv}}$ exists and is smooth (Lipschitz).
- Problem: Not analytically computable.
- Solution: Approximate it with a Multilayer Perceptron (MLP).

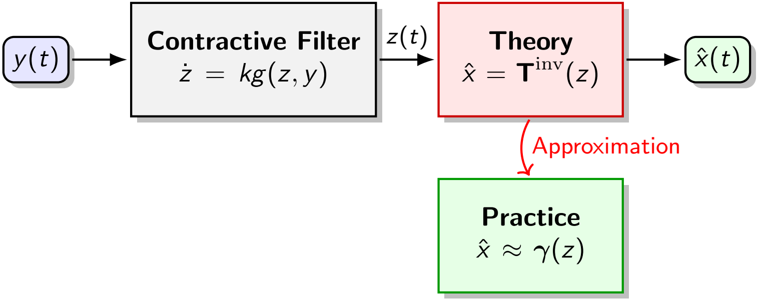


## Main Result

- ✓ The **Nonlinear Filter + MLP** architecture is a **tunable observer**.
- ✓ Backed by formal guarantees (convergence & injectivity).
- ✓ Practically implementable.

# Linear vs. Nonlinear Activation for the RNN

For both linear and monotonic nonlinear activation functions, we get a tunable observer structure.

*So, is it better to use linear or nonlinear functions?*

Consider a linear KKL observer: $\dot{z}_i = k\lambda_i z_i + y$. There is a well-known trade-off:

- If $k$ **is large**:
  - Convergence rate is high.
  - Sensitivity to measurement noise is high.
- If $k$ **is small**:
  - Convergence rate is slow.
  - Robustness to measurement noise is better.

Question: How can we combine the advantages of both?

# A nonlinear gain scheduling approach

We want an observer that is:

- **Fast** during the transient phase (when the error $z - y$ is large).
- **Slow/robust** at steady state (when the error $z - y$ is small).

A possible nonlinear structure for the filter that achieves this is:

$$\dot{z} = \lambda \left( \underbrace{a_{\text{fast}}(z-y)}_{\text{High-gain term}} + \underbrace{(a_{\text{slow}} - a_{\text{fast}}) \tanh(z-y)}_{\text{Saturation for small errors}} \right)$$

This defines a monotonic function $\sigma(z, y)$!

$$\sigma(z,y) = a_{\text{fast}}(z-y) + (a_{\text{slow}} - a_{\text{fast}}) \tanh(z-y)$$

$\Rightarrow$ Our theoretical results apply, and we can learn the mapping $\mathbf{T}^{\mathrm{inv}}$.

# Simulation example: Duffing oscillator

Consider a nonlinear Duffing oscillator:

$$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = -0.2x_1 - x_1^3 \end{cases}, \quad y = x_1.$$

We compare three activation functions:

1. **Nonlinear**: $\dot{z} = \lambda(a_{\text{fast}}(z-y) + (a_{\text{slow}} - a_{\text{fast}})\tanh(z-y))$
2. **Fast Linear**: $\dot{z} = \lambda a_{\text{fast}}(z-y)$
3. **Slow Linear**: $\dot{z} = \lambda a_{\text{slow}}(z-y)$

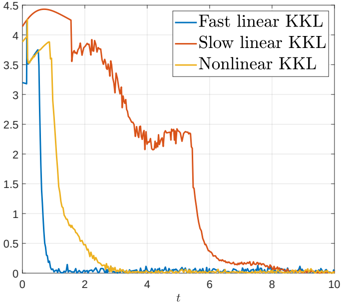


(a) Scenario 1: Convergence without noise. The nonlinear observer is as fast

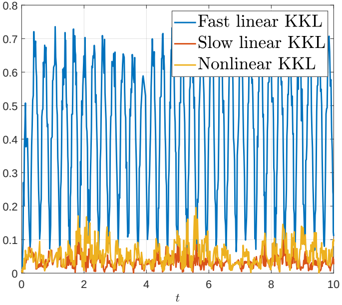


(b) Scenario 2: Estimation with measurement noise. The nonlinear

# In Conclusion

- It is possible to show that a continuous-time model of an observer based on RNNs and MLPs results in a **tunable observer structure**.
- The proof relies on a nonlinear extension of the KKL observer theory, leveraging properties of **contracting systems**.
- The use of specific nonlinear activation functions is not just a theoretical generalization; it can be practically motivated to combine desirable **behaviors** like fast convergence and noise robustness.
- Open question: What about rigorous guarantees for discrete-time **versions** of these observers?