

# Liquid Time-Constant Networks

Ramin Hasani,<sup>1,3\*</sup> Mathias Lechner,<sup>2\*</sup> Alexander Amini,<sup>1</sup> Daniela Rus,<sup>1</sup> Radu Grosu<sup>3</sup>

<sup>1</sup> Massachusetts Institute of Technology (MIT)

<sup>2</sup> Institute of Science and Technology Austria (IST Austria)

<sup>3</sup> Technische Universität Wien (TU Wien)

rhasani@mit.edu, mathias.lechner@ist.ac.at, amini@mit.edu, rus@csail.mit.edu, radu.grosu@tuwien.ac.at

## Abstract

We introduce a new class of time-continuous recurrent neural network models. Instead of declaring a learning system’s dynamics by implicit nonlinearities, we construct networks of linear first-order dynamical systems modulated via nonlinear interlinked gates. The resulting models represent dynamical systems with varying (i.e., *liquid*) time-constants coupled to their hidden state, with outputs being computed by numerical differential equation solvers. These neural networks exhibit stable and bounded behavior, yield superior expressivity within the family of neural ordinary differential equations, and give rise to improved performance on time-series prediction tasks. To demonstrate these properties, we first take a theoretical approach to find bounds over their dynamics, and compute their expressive power by the *trajectory length* measure in a latent trajectory space. We then conduct a series of time-series prediction experiments to manifest the approximation capability of Liquid Time-Constant Networks (LTCs) compared to classical and modern RNNs.

## Introduction

Recurrent neural networks with continuous-time hidden states determined by ordinary differential equations (ODEs), are effective algorithms for modeling time series data that are ubiquitously used in medical, industrial and business settings. The state of a neural ODE,  $\mathbf{x}(t) \in \mathbb{R}^D$ , is defined by the solution of this equation (Chen et al. 2018):  $d\mathbf{x}(t)/dt = f(\mathbf{x}(t), \mathbf{I}(t), t, \theta)$ , with a neural network  $f$  parametrized by  $\theta$ . One can then compute the state using a numerical ODE solver, and train the network by performing reverse-mode automatic differentiation (Rumelhart, Hinton, and Williams 1986), either by gradient descent through the solver (Lechner et al. 2019), or by considering the solver as a black-box (Chen et al. 2018; Dupont, Doucet, and Teh 2019; Gholami, Keutzer, and Biros 2019) and apply the *adjoint method* (Pontryagin 2018). The open questions are: how expressive are neural ODEs in their current formalism, and can we improve their structure to enable better representation learning?

Rather than defining the derivatives of the hidden-state directly by a neural network  $f$ , one can determine a more stable continuous-time recurrent neural network (CT-RNN)

by the following equation (Funahashi and Nakamura 1993):  $\frac{d\mathbf{x}(t)}{dt} = -\frac{\mathbf{x}(t)}{\tau} + f(\mathbf{x}(t), \mathbf{I}(t), t, \theta)$ , in which the term  $-\frac{\mathbf{x}(t)}{\tau}$  assists the autonomous system to reach an equilibrium state with a time-constant  $\tau$ .  $\mathbf{x}(t)$  is the hidden state,  $\mathbf{I}(t)$  is the input,  $t$  represents time, and  $f$  is parametrized by  $\theta$ .

We propose an alternative formulation: let the hidden state flow of a network be declared by a system of linear ODEs of the form:  $d\mathbf{x}(t)/dt = -\mathbf{x}(t)/\tau + \mathbf{S}(t)$ , and let  $\mathbf{S}(t) \in \mathbb{R}^M$  represent the following nonlinearity determined by  $\mathbf{S}(t) = f(\mathbf{x}(t), \mathbf{I}(t), t, \theta)(A - \mathbf{x}(t))$ , with parameters  $\theta$  and  $A$ . Then, by plugging in  $\mathbf{S}$  into the hidden states equation, we get:

$$\frac{d\mathbf{x}(t)}{dt} = -\left[\frac{1}{\tau} + f(\mathbf{x}(t), \mathbf{I}(t), t, \theta)\right]\mathbf{x}(t) + f(\mathbf{x}(t), \mathbf{I}(t), t, \theta)A. \quad (1)$$

Eq. 1 manifests a novel time-continuous RNN instance with several features and benefits:

**Liquid Time-Constant.** A neural network  $f$  not only determines the derivative of the hidden state  $\mathbf{x}(t)$ , but also serves as an input-dependent varying time-constant ( $\tau_{sys} = \frac{\tau}{1 + \tau f(\mathbf{x}(t), \mathbf{I}(t), t, \theta)}$ ) for the learning system (Time constant is a parameter characterizing the speed and the coupling sensitivity of an ODE). This property enables single elements of the hidden state to identify specialized dynamical systems for input features arriving at each time-point. We refer to these models as *liquid time-constant* networks (LTCs). LTCs can be implemented by an arbitrary choice of ODE solvers. In Section 2, we introduce a practical fixed-step ODE solver that simultaneously enjoys the stability of the implicit Euler and the efficiency of the explicit Euler methods.

**Reverse-Mode Automatic Differentiation of LTCs.** LTCs realize differentiable computational graphs. Similar to neural ODEs, they can be trained by variform of gradient-based optimization algorithms. We settle to trade memory for numerical precision during a backward-pass by using a vanilla backpropagation through-time algorithm to optimize LTCs instead of an adjoint-based optimization method (Pontryagin 2018). In Section 3, we motivate this choice thoroughly.

**Bounded Dynamics - Stability.** In Section 4, we show that the state and the time-constant of LTCs are bounded to a finite range. This property assures the stability of the output dynamics and is desirable when inputs to the system relentlessly increase.

\* Authors with equal contributions

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

**Superior Expressivity.** In Section 5, we theoretically and quantitatively analyze the approximation capability of LTCs. We take a functional analysis approach to show the universality of LTCs. We then delve deeper into measuring their expressivity compared to other time-continuous models. We perform this by measuring the *trajectory length* of activations of networks in a latent trajectory representation. Trajectory length was introduced as a measure of expressivity of feed-forward deep neural networks (Raghu et al. 2017). We extend these criteria to the CT family.

**Time-Series Modeling.** In Section 6, we conduct a series of eleven time-series prediction experiments and compare the performance of modern RNNs to the time-continuous models. We observe improved performance on a majority of cases achieved by LTCs.

**Why This Specific Formulation?** There are two primary justifications for the choice of this particular representation: I) LTC model is loosely related to the computational models of neural dynamics in small species, put together with synaptic transmission mechanisms (Sarma et al. 2018; Gleeson et al. 2018; Hasani et al. 2020). The dynamics of non-spiking neurons’ potential,  $\mathbf{v}(t)$ , can be written as a system of linear ODEs of the form (Lapicque 1907; Koch and Segev 1998):  $d\mathbf{v}/dt = -g_l\mathbf{v}(t) + \mathbf{S}(t)$ , where  $\mathbf{S}$  is the sum of all synaptic inputs to the cell from presynaptic sources, and  $g_l$  is a leakage conductance.

All synaptic currents to the cell can be approximated in steady-state by the following nonlinearity (Koch and Segev 1998; Wicks, Roehrig, and Rankin 1996):  $\mathbf{S}(t) = f(\mathbf{v}(t), \mathbf{I}(t)), (A - \mathbf{v}(t))$ , where  $f(\cdot)$  is a sigmoidal nonlinearity depending on the state of all neurons,  $\mathbf{v}(t)$  which are presynaptic to the current cell, and external inputs to the cell,  $\mathbf{I}(t)$ . By plugging in these two equations, we obtain an equation similar to Eq. 1. LTCs are inspired by this foundation.

II) Eq. 1 might resemble that of the famous Dynamic Causal Models (DCMs) (Friston, Harrison, and Penny 2003) with a Bilinear dynamical system approximation (Penny, Ghahramani, and Friston 2005). DCMs are formulated by taking a second-order approximation (Bilinear) of the dynamical system  $d\mathbf{x}/dt = F(\mathbf{x}(t), \mathbf{I}(t), \theta)$ , that would result in the following format (Friston, Harrison, and Penny 2003):  $d\mathbf{x}/dt = (A + \mathbf{I}(t)B)\mathbf{x}(t) + C\mathbf{I}(t)$  with  $A = \frac{dF}{d\mathbf{x}}$ ,  $B = \frac{dF^2}{d\mathbf{x}(t)d\mathbf{I}(t)}$ ,  $C = \frac{dF}{d\mathbf{I}(t)}$ . DCM and bilinear dynamical systems have shown promise in learning to capture complex fMRI time-series signals. LTCs are introduced as variants of continuous-time (CT) models that show great expressivity, stability, and performance in modeling time series.

## LTCs Forward-Pass By A Fused ODE Solvers

Solving Eq. 1 analytically, is non-trivial due to the nonlinearity of the LTC semantics. The state of the system of ODEs, however, at any time point  $T$ , can be computed by a numerical ODE solver that simulates the system starting from a trajectory  $x(0)$ , to  $x(T)$ . An ODE solver breaks down the continuous simulation interval  $[0, T]$  to a temporal discretization,  $[t_0, t_1, \dots, t_n]$ . As a result, a solver’s step involves only the update of the neuronal states from  $t_i$  to  $t_{i+1}$ .

LTCs’ ODE realizes a system of stiff equations (Press

---

## Algorithm 1 LTC update by fused ODE Solver

---

**Parameters:**  $\theta = \{\tau^{(N \times 1)} = \text{time-constant}, \gamma^{(M \times N)} = \text{weights}, \gamma_r^{(N \times N)} = \text{recurrent weights}, \mu^{(N \times 1)} = \text{biases}\}$ ,  $A^{(N \times 1)} = \text{bias vector}$ ,  $L = \text{Number of unfolding steps}$ ,  $\Delta t = \text{step size}$ ,  $N = \text{Number of neurons}$ ,  
**Inputs:**  $M$ -dimensional Input  $\mathbf{I}(t)$  of length  $T$ ,  $\mathbf{x}(0)$   
**Output:** Next LTC neural state  $\mathbf{x}_{t+\Delta t}$   
**Function:** FusedStep( $\mathbf{x}(t), \mathbf{I}(t), \Delta t, \theta$ )  
 $\mathbf{x}(t + \Delta t)^{(N \times T)} = \frac{\mathbf{x}(t) + \Delta t f(\mathbf{x}(t), \mathbf{I}(t), t, \theta) \odot A}{1 + \Delta t (1/\tau + f(\mathbf{x}(t), \mathbf{I}(t), t, \theta))}$   
 $\triangleright f(\cdot)$ , and all divisions are applied element-wise.  
 $\triangleright \odot$  is the Hadamard product.  
**end Function**  
 $\mathbf{x}_{t+\Delta t} = \mathbf{x}(t)$   
**for**  $i = 1 \dots L$  **do**  
     $\mathbf{x}_{t+\Delta t} = \text{FusedStep}(\mathbf{x}(t), \mathbf{I}(t), \Delta t, \theta)$   
**end for**  
**return**  $\mathbf{x}_{t+\Delta t}$

---

et al. 2007). This type of ODE requires an exponential number of discretization steps when simulated with a Runge-Kutta (RK) based integrator. Consequently, ODE solvers based on RK, such as Dormand–Prince (default in torchdiffeq (Chen et al. 2018)), are not suitable for LTCs. Therefore, we design a new ODE solver that fuses the explicit and implicit Euler methods. Our discretization method results in greater stability, and numerically unrolls a given dynamical system of the form  $dx/dt = f(x)$  by:

$$x(t_{i+1}) = x(t_i) + \Delta t f(x(t_i), x(t_{i+1})). \quad (2)$$

In particular, we replace only the  $x(t_i)$  that occur linearly in  $f$  by  $x(t_{i+1})$ . As a result, Eq 2 can be solved for  $x(t_{i+1})$ , symbolically. Applying the Fused solver to the LTC representation, and solving it for  $\mathbf{x}(t + \Delta t)$ , we get:

$$\mathbf{x}(t + \Delta t) = \frac{\mathbf{x}(t) + \Delta t f(\mathbf{x}(t), \mathbf{I}(t), t, \theta) A}{1 + \Delta t (1/\tau + f(\mathbf{x}(t), \mathbf{I}(t), t, \theta))}. \quad (3)$$

Eq. 3 computes one update state for an LTC network. Correspondingly, Algorithm 1 shows how to implement an LTC network, given a parameter space  $\theta$ .  $f$  is assumed to have an arbitrary activation function (e.g. for a  $\tanh$  nonlinearity  $f = \tanh(\gamma_r \mathbf{x} + \gamma \mathbf{I} + \mu)$ ). The computational complexity of the algorithm for an input sequence of length  $T$  is  $O(L \times T)$ , where  $L$  is the number of discretization steps. Intuitively, a dense version of an LTC network with  $N$  neurons, and a dense version of a long short-term memory (LSTM) (Hochreiter and Schmidhuber 1997) network with  $N$  cells, would be of the same complexity.

## Training LTC Networks By BPTT

Neural ODEs were suggested to be trained by a constant memory cost for each layer in a neural network  $f$  by applying the adjoint sensitivity method to perform reverse-mode automatic differentiation (Chen et al. 2018). The adjoint method, however, comes with numerical errors when running in reverse mode. This phenomenon happens because

---

**Algorithm 2** Training LTC by BPTT

---

**Inputs:** Dataset of traces  $[I(t), y(t)]$  of length  $T$ , RNN-cell  $= f(I, x)$   
**Parameter:** Loss func  $L(\theta)$ , initial param  $\theta_0$ , learning rate  $\alpha$ , Output  $w = W_{out}$ , and bias  $= b_{out}$   
**for**  $i = 1 \dots$  number of training steps **do**  
   $(I_b, y_b) =$  Sample training batch,  $x := x_{t_0} \sim p(x_{t_0})$   
  **for**  $j = 1 \dots T$  **do**  
     $x = f(I(t), x)$ ,  $\hat{y}(t) = W_{out} \cdot x + b_{out}$ ,  $L_{total} = \sum_{j=1}^T L(y_j(t), \hat{y}_j(t))$ ,  $\nabla L(\theta) = \frac{\partial L_{total}}{\partial \theta}$   
     $\theta = \theta - \alpha \nabla L(\theta)$   
  **end for**  
**end for**  
**return**  $\theta$

---

the adjoint method forgets the forward-time computational trajectories, which was repeatedly denoted by the community (Gholami, Keutzer, and Biros 2019; Zhuang et al. 2020).

On the contrary, direct backpropagation through time (BPTT) trades memory for accurate recovery of the forward-pass during the reverse mode integration (Zhuang et al. 2020). Thus, we set out to design a vanilla BPTT algorithm to maintain a highly accurate backward-pass integration through the solver. For this purpose, a given ODE solver’s output (a vector of neural states), can be recursively folded to build an RNN and then apply Algorithm 2 to train the system. Algorithm 2 uses a vanilla stochastic gradient descent (SGD). One can substitute this with a more performant variant of the SGD, such as Adam (Kingma and Ba 2014), which we use in our experiments.

**Complexity.** Table 1 summarizes the complexity of our vanilla BPTT algorithm compared to an adjoint method. We achieve a high degree of accuracy on both forward and backward integration trajectories, with similar computational complexity, at large memory costs.

### Bounds on $\tau$ and Neural State of LTCs

LTCs are represented by an ODE which varies its time-constant based on inputs. It is therefore important to see if LTCs stay stable for unbounded arriving inputs (Hasani et al. 2019; Lechner et al. 2020b). In this section, we prove that the time-constant and the state of LTC neurons are bounded to a finite range, as described in Theorems 1 and 2, respectively.

**Theorem 1.** *Let  $x_i$  denote the state of a neuron  $i$  within an LTC network identified by Eq. 1, and let neuron  $i$  receive  $M$  incoming connections. Then, the time-constant of the neuron,  $\tau_{sys_i}$ , is bounded to the following range:*

$$\tau_i / (1 + \tau_i W_i) \leq \tau_{sys_i} \leq \tau_i, \quad (4)$$

The proof is provided in Appendix. It is constructed based on bounded, monotonically increasing sigmoidal nonlinearity for neural network  $f$  and its replacement in the LTC network dynamics. A stable varying time-constant significantly enhances the expressivity of this form of time-continuous RNNs, as we discover more formally in Section 5.

	Vanilla BPTT	Adjoint
Time	$O(L \times T \times 2)$	$O((L_f + L_b) \times T)$
Memory	$O(L \times T)$	<b>O(1)</b>
Depth	$O(L)$	$O(L_b)$
FWD acc	High	High
BWD acc	<b>High</b>	Low

Table 1: Complexity of the vanilla BPTT compared to the adjoint method, for a single layer neural network  $f$ . Note:  $L$  = number of discretization steps,  $L_f$  =  $L$  during forward-pass.  $L_b$  =  $L$  during backward-pass.  $T$  = length of sequence, Depth = computational graph depth.

**Theorem 2.** *Let  $x_i$  denote the state of a neuron  $i$  within an LTC, identified by Eq. 1, and let neuron  $i$  receive  $M$  incoming connections. Then, the hidden state of any neuron  $i$ , on a finite interval  $Int \in [0, T]$ , is bounded as follows:*

$$\min(0, A_i^{min}) \leq x_i(t) \leq \max(0, A_i^{max}), \quad (5)$$

The proof is given in Appendix. It is constructed based on the sign of the LTC’s equation’s compartments, and an approximation of the ODE model by an explicit Euler discretization. Theorem 2 illustrates a desired property of LTCs, namely *state stability* which guarantees that the outputs of LTCs never explode even if their inputs grow to infinity. Next we discuss the expressive power of LTCs compared to the family of time-continuous models, such as CT-RNNs and neural ordinary differential equations (Chen et al. 2018; Rubanova, Chen, and Duvenaud 2019).

### On The Expressive Power of LTCs

Understanding the impact of a NN’s structural properties on their computable functions is known as the expressivity problem. The very early attempts on measuring expressivity of NNs include theoretical studies based on functional analysis. They show that NNs with three-layers can approximate any finite set of continuous mapping with any precision. This is known as the *universal approximation theorem* (Hornik, Stinchcombe, and White 1989; Funahashi 1989; Cybenko 1989). Universality was extended to standard RNNs (Funahashi 1989) and even continuous-time RNNs (Funahashi and Nakamura 1993). By careful considerations, we can also show that LTCs are also universal approximators.

**Theorem 3.** *Let  $x \in \mathbb{R}^n$ ,  $S \subset \mathbb{R}^n$  and  $\dot{x} = F(x)$  be an autonomous ODE with  $F : S \rightarrow \mathbb{R}^n$  a  $C^1$ -mapping*

Activations	Computational Depth		
	Neural ODE	CT-RNN	LTC
tanh	0.56 ± 0.016	4.13 ± 2.19	9.19 ± 2.92
sigmoid	0.56 ± 0.00	5.33 ± 3.76	7.00 ± 5.36
ReLU	1.29 ± 0.10	4.31 ± 2.05	56.9 ± 9.03
Hard-tanh	0.61 ± 0.02	4.05 ± 2.17	81.01 ± 10.05

Table 2: Computational depth of models. Note: # of tries = 100, input samples’  $\Delta t = 0.01$ ,  $T = 100$  sequence length. # of layers = 1, width = 100,  $\sigma_w^2 = 2$ ,  $\sigma_b^2 = 1$ .

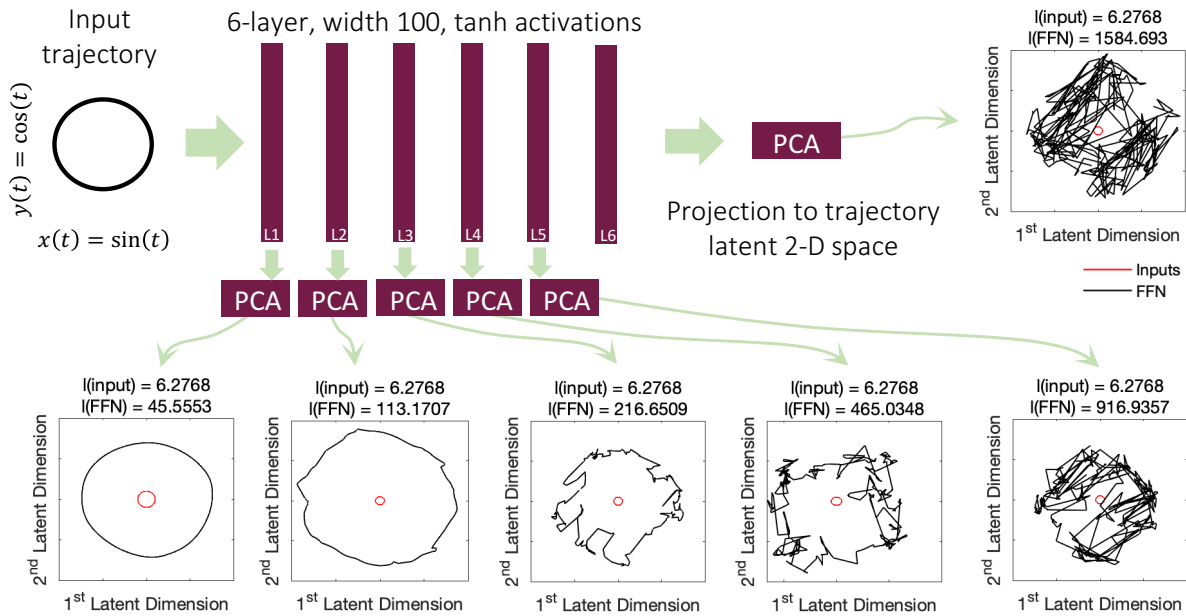


Figure 1: Trajectory’s latent space becomes more complex as the input passes through hidden layers.

on  $S$ . Let  $D$  denote a compact subset of  $S$  and assume that the simulation of the system is bounded in the interval  $I = [0, T]$ . Then, for a positive  $\epsilon$ , there exist an LTC network with  $N$  hidden units,  $n$  output units, and an output internal state  $\mathbf{u}(t)$ , described by Eq. 1, such that for any rollout  $\{\mathbf{x}(t)|t \in I\}$  of the system with initial value  $\mathbf{x}(0) \in D$ , and a proper network initialization,

$$\max_{t \in I} |\mathbf{x}(t) - \mathbf{u}(t)| < \epsilon \quad (6)$$

The proof defines an  $n$ -dimensional dynamical system and place it into a higher dimensional system. The second system is an LTC. The fundamental difference of the proof of LTC’s universality to that of CT-RNNs (Funahashi and Nakamura 1993) lies in the distinction of the semantics of both systems where the LTC network contains a nonlinear input-dependent term in its time-constant module which makes parts of the proof non-trivial.

The universal approximation theorem broadly explores the expressive power of a neural network. The theorem however, does not yield a concrete measure on where the separation is between different neural network architectures. Therefore, a more rigorous measure of expressivity is demanded to compare models, specifically those networks specialized in spatiotemporal data processing, such as LTCs. The advances made on defining measures for the expressivity of static deep learning models (Pascanu, Montufar, and Bengio 2013; Montufar et al. 2014; Eldan and Shamir 2016; Poole et al. 2016; Raghu et al. 2017) could help measure the expressivity of time-continuous models, both theoretically and quantitatively, which we explore in the next section.

### Measuring Expressivity By Trajectory Length

A measure of expressivity has to take into account what degrees of complexity a learning system can compute, given

the network’s capacity (depth, width, type, and weights configuration). A unifying expressivity measure of static deep networks is the *trajectory length* introduced in (Raghu et al. 2017). In this context, one evaluates how a deep model transforms a given input trajectory (e.g., a circular 2-dimensional input) into a more complex pattern, progressively.

We can then perform principle component analysis (PCA) over the obtained network’s activations. Subsequently, we measure the length of the output trajectory in a 2-dimensional latent space, to uncover its relative complexity (see Fig. 1). The trajectory length is defined as the *arc length* of a given trajectory  $I(t)$ , (e.g. a circle in 2D space) (Raghu et al. 2017):  $l(I(t)) = \int_t \|dI(t)/dt\| dt$ . By establishing a lower-bound for the growth of the trajectory length, one can set a barrier between networks of shallow and deep architectures, regardless of any assumptions on the network’s weight configuration (Raghu et al. 2017), unlike many other measures of expressivity (Pascanu, Montufar, and Bengio 2013; Montufar et al. 2014; Serra, Tjandraatmadja, and Ramalingam 2017; Gabrié et al. 2018; Hanin and Rolnick 2018, 2019; Lee, Alvarez-Melis, and Jaakkola 2019). We set out to extend the trajectory-space analysis of static networks to time-continuous (TC) models, and to lower-bound the trajectory length to compare models’ expressivity. To this end, we designed instances of Neural ODEs, CT-RNNs and LTCs with shared  $f$ . The networks were initialized by weights  $\sim \mathcal{N}(0, \sigma_w^2/k)$ , and biases  $\sim \mathcal{N}(0, \sigma_b^2)$ . We then perform forward-pass simulations by using different types of ODE solvers, for arbitrary weight profiles, while exposing the networks to a circular input trajectory  $I(t) = \{I_1(t) = \sin(t), I_2(t) = \cos(t)\}$ , for  $t \in [0, 2\pi]$ . By looking at the first two principle components (with an average variance-explained of over 80%) of hidden layers’ activations, we observed consistently more complex trajectories for LTCs.

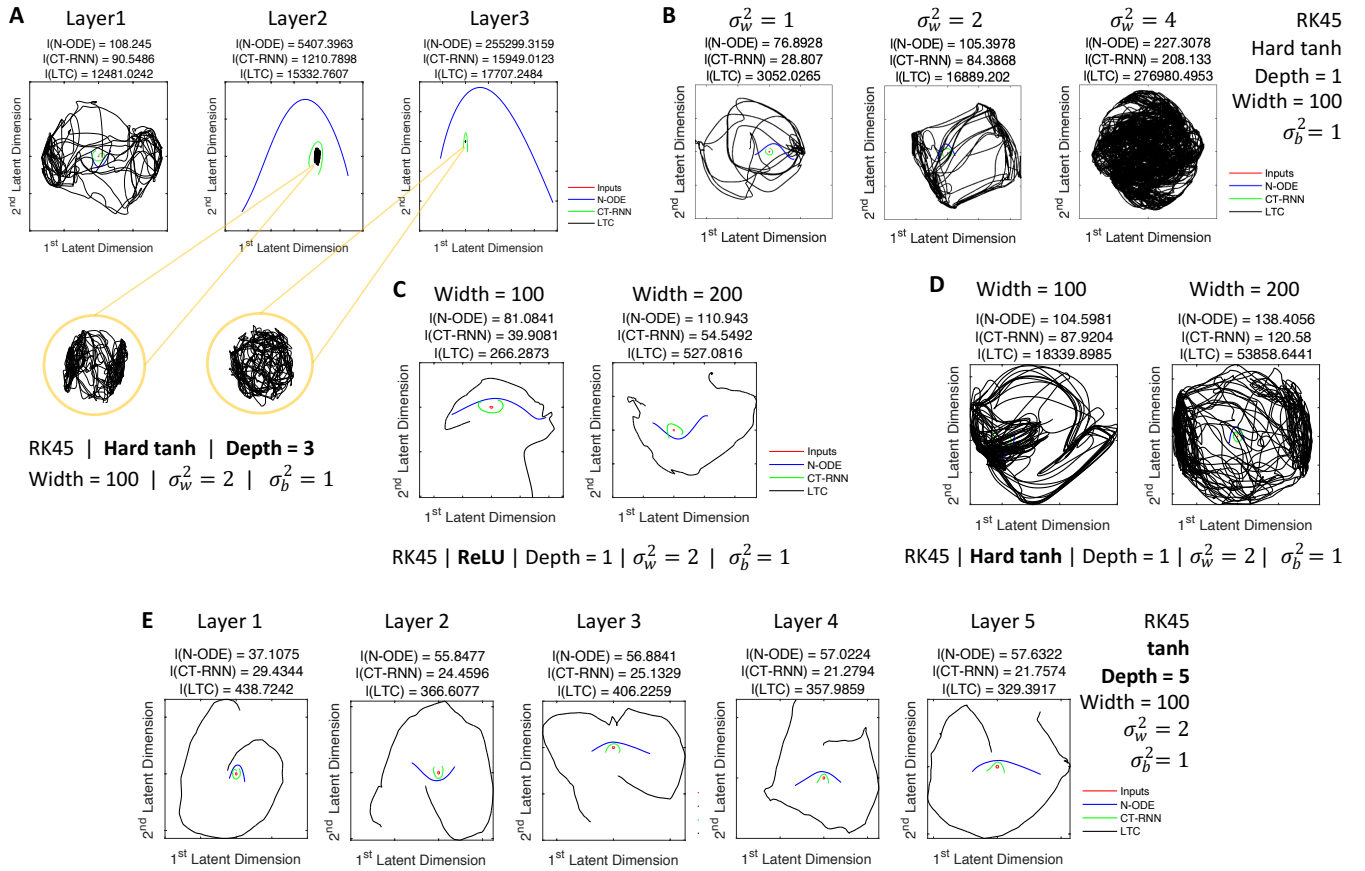


Figure 2: Trajectory length deformation A) in network layers with Hard-tanh activations, B) as a function of the weight distribution scaling factor, C) as a function of network width (ReLU), D) as a function of width (Hard-tanh), and E) in network layers with logistic-sigmoid activations.

Fig. 2 gives a glimpse of our empirical observations. All networks are implemented by the Dormand-Prince explicit Runge-Kutta(4,5) solver (Dormand and Prince 1980) with a variable step size. We had the following **observations**:

**I)** Exponential growth of the trajectory length of Neural ODEs and CT-RNNs with Hard-tanh and ReLU activations (Fig. 2A) and unchanged shape of their latent space regardless of their weight profile.

**II)** LTCs show a slower growth-rate of the trajectory length when designed by Hard-tanh and ReLU, with the compromise of realizing great levels of complexity (Fig. 2A, 2C and 2D).

**III)** Apart from multi-layer time-continuous models built by Hard-tanh and ReLU activations, in all cases, we observed a longer and a more complex latent space behavior for the LTC networks (Fig. 2B to 2D).

**IV)** Unlike static deep networks (Fig. 1), we witnessed that the trajectory length does not grow by depth in multi-layer continuous-time networks realized by tanh and sigmoid (Fig. 2E).

**V)** conclusively, we observed that the trajectory length in TC models varies by a model's activations, weight and bias distributions variance, width and depth. We presented this

more systematically in Fig. 3.

**VI)** Trajectory length grows linearly with a network's width (Fig. 3B - Notice the logarithmic growth of the curves in the log-scale Y-axis).

**VII)** The growth is considerably faster as the variance grows (Fig. 3C).

**VIII)** Trajectory length is reluctant to the choice of ODE solver (Fig. 3A).

**IX)** Activation functions diversify the complex patterns explored by the TC system, where ReLU and Hard-tanh networks demonstrate higher degrees of complexity for LTCs. A key reason is the presence of recurrent links between each layer's cells.

**Definition of Computational Depth ( $L$ ).** For one hidden layer of  $f$  in a CT network,  $L$  is the average number of integration steps by the solver for each incoming input sample. Note that for an  $f$  with  $n$  layers we define the total depth as  $n \times L$ . These observations allow us to formulate lower bounds on the trajectory length growth of CT networks.

**Theorem 4.** Trajectory Length growth Bounds for Neural ODEs and CT-RNNs. Let  $dx/dt = f_{n,k}(x(t), I(t), \theta)$  with  $\theta = \{W, b\}$ , represent a Neural ODE and  $\frac{dx(t)}{dt} =$

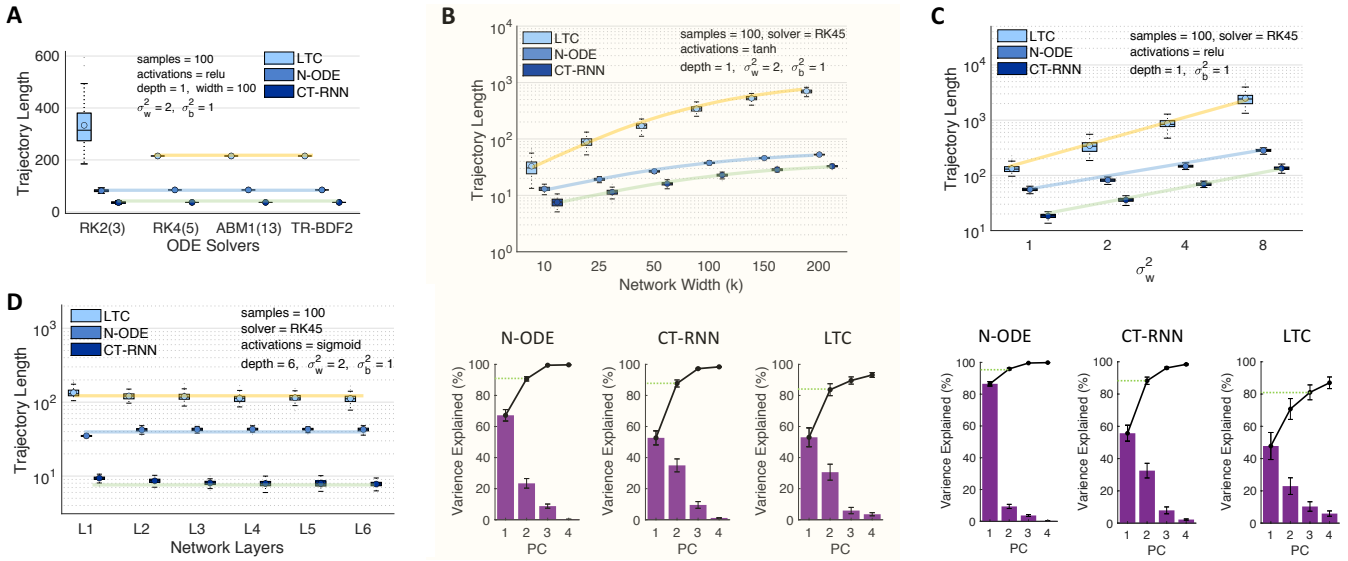


Figure 3: Dependencies of the trajectory length measure. A) trajectory length vs different solvers (variable-step solvers). RK2(3): Bogacki-Shampine Runge-Kutta (2,3) (Bogacki and Shampine 1989). RK4(5): Dormand-Prince explicit RK (4,5) (Dormand and Prince 1980). ABM1(13): Adams-Bashforth-Moulton (Shampine 1975). TR-BDF2: implicit RK solver with 1st stage trapezoidal rule and a 2nd stage backward differentiation (Hosea and Shampine 1996). B) Top: trajectory length vs network width. Bottom: Variance-explained of principle components (purple bars) and their cumulative values (solid black line). C) Trajectory length vs weights-distribution variance. D) trajectory length vs layers. (More results in the supplements)

$-\frac{\mathbf{x}(t)}{\tau} + f_{n,k}(\mathbf{x}(t), \mathbf{I}(t), \theta)$  with  $\theta = \{W, b, \tau\}$  a CT-RNN.  $f$  is randomly weighted with Hard-tanh activations. Let  $\mathbf{I}(t)$  be a 2D input trajectory, with its progressive points (i.e.  $\mathbf{I}(t + \delta t)$ ) having a perpendicular component to  $\mathbf{I}(t)$  for all  $\delta t$ , with  $L =$  number of solver-steps. Then, by defining the projection of the first two principle components' scores of the hidden states over each other, as the 2D latent trajectory space of a layer  $d$ ,  $z^{(d)}(\mathbf{I}(t)) = z^{(d)}(t)$ , for Neural ODE and CT-RNNs respectively, we have:

$$\mathbb{E} \left[ l(z^{(d)}(t)) \right] \geq O \left( \frac{\sigma_w \sqrt{k}}{\sqrt{\sigma_w^2 + \sigma_b^2 + k \sqrt{\sigma_w^2 + \sigma_b^2}}} \right)^{d \times L} l(\mathbf{I}(t)), \quad (7)$$

$$\mathbb{E} \left[ l(z^{(d)}(t)) \right] \geq O \left( \frac{(\sigma_w - \sigma_b) \sqrt{k}}{\sqrt{\sigma_w^2 + \sigma_b^2 + k \sqrt{\sigma_w^2 + \sigma_b^2}}} \right)^{d \times L} l(\mathbf{I}(t)). \quad (8)$$

The proof is provided in Appendix. It follows similar steps as (Raghu et al. 2017) on the trajectory length bounds established for deep networks with piecewise linear activations, with careful considerations due to the continuous-time setup. The proof is constructed such that we formulate a recurrence between the norm of the hidden state gradient in layer  $d + 1$ ,  $\|dz/dt^{(d+1)}\|$ , in principle components domain, and the expectation of the norm of the right-hand-side of the differential equations of neural ODEs and CT-RNNs. We then roll back the recurrence to reach the inputs.

Note that to reduced the complexity of the problem, we only bounded the orthogonal components of the hidden state

image  $\|dz/dt_{\perp}^{(d+1)}\|$ , and therefore we have the assumption on input  $\mathbf{I}(t)$ , in the Theorem's statement (Raghu et al. 2017). Next, we find a lower-bound for the LTC networks.

**Theorem 5.** Growth Rate of LTC's Trajectory Length. Let Eq. 1 determine an LTC with  $\theta = \{W, b, \tau, A\}$ . With the same conditions on  $f$  and  $\mathbf{I}(t)$ , as in Theorem 4, we have:

$$\mathbb{E} \left[ l(z^{(d)}(t)) \right] \geq O \left( \left( \frac{\sigma_w \sqrt{k}}{\sqrt{\sigma_w^2 + \sigma_b^2 + k \sqrt{\sigma_w^2 + \sigma_b^2}}} \right)^{d \times L} \times \left( \sigma_w + \frac{\|z^{(d)}\|}{\min(\delta t, L)} \right) \right) l(\mathbf{I}(t)). \quad (9)$$

The proof is provided in Appendix. A rough outline: we construct the recurrence between the norm of the hidden state gradients and the components of the right-hand-side of LTC separately which progressively build up the bound.

## Discussion of The Theoretical Bounds

**I)** As expected, the bound for the Neural ODEs is very similar to that of an  $n$  layer static deep network with the exception of the exponential dependencies to the number of solver-steps,  $L$ . **II)** The bound for CT-RNNs suggests their shorter trajectory length compared to neural ODEs, according to the base of the exponent. This results consistently matches our experiments presented in Figs. 2 and 3. **III)** Fig. 2B and Fig. 3C show a faster-than-linear growth for LTC's trajectory length as a function of weight distribution variance. This is confirmed by LTC's lower bound shown in Eq.

Dataset	Metric	LSTM	CT-RNN	Neural ODE	CT-GRU	LTC (ours)
Gesture	(accuracy)	64.57% $\pm$ 0.59	59.01% $\pm$ 1.22	46.97% $\pm$ 3.03	68.31% $\pm$ 1.78	<b>69.55% <math>\pm</math> 1.13</b>
Occupancy	(accuracy)	93.18% $\pm$ 1.66	94.54% $\pm$ 0.54	90.15% $\pm$ 1.71	91.44% $\pm$ 1.67	<b>94.63% <math>\pm</math> 0.17</b>
Activity recognition	(accuracy)	95.85% $\pm$ 0.29	95.73% $\pm$ 0.47	<b>97.26%</b> $\pm$ 0.10	96.16% $\pm$ 0.39	95.67% $\pm$ 0.575
Sequential MNIST	(accuracy)	<b>98.41%</b> $\pm$ 0.12	96.73% $\pm$ 0.19	97.61% $\pm$ 0.14	98.27% $\pm$ 0.14	97.57% $\pm$ 0.18
Traffic	(squared error)	0.169 $\pm$ 0.004	0.224 $\pm$ 0.008	1.512 $\pm$ 0.179	0.389 $\pm$ 0.076	<b>0.099</b> $\pm$ 0.0095
Power	(squared-error)	0.628 $\pm$ 0.003	0.742 $\pm$ 0.005	1.254 $\pm$ 0.149	<b>0.586</b> $\pm$ 0.003	0.642 $\pm$ 0.021
Ozone	(F1-score)	0.284 $\pm$ 0.025	0.236 $\pm$ 0.011	0.168 $\pm$ 0.006	0.260 $\pm$ 0.024	<b>0.302</b> $\pm$ 0.0155

Table 3: Time series prediction. Mean and standard deviation, n=5

9. **IV)** LTC’s lower bound also depicts the linear growth of the trajectory length with the width,  $k$ , which validates the results presented in 3B. **V)** Given the computational depth of the models  $L$  in Table 2 for Hard-tanh activations, the computed lower bound for neural ODEs, CT-RNNs and LTCs justify a longer trajectory length of LTC networks in the experiments of Section 5. Next, we assess the expressive power of LTCs in a set of real-life time-series prediction tasks.

### Experimental Evaluation

In this section, we evaluate the performance of the LTCs compared to the state-of-the-art RNN models in a series of time-series benchmarks.

#### Time Series Predictions

We evaluated the performance of LTCs realized by the proposed Fused ODE solver against the state-of-the-art discretized RNNs, LSTMs (Hochreiter and Schmidhuber 1997), CT-RNNs (ODE-RNNs) (Funahashi and Nakamura 1993; Rubanova, Chen, and Duvenaud 2019), continuous-time gated recurrent units (CT-GRUs) (Mozer, Kazakov, and Lindsey 2017), and Neural ODEs constructed by a 4<sup>th</sup> order Runge-Kutta solver as suggested in (Chen et al. 2018), in a series of diverse real-life supervised learning tasks. The results are summarized in Table 3. The experimental setup are provided in Appendix. We observed between 5% to 70% performance improvement achieved by the LTCs compared to other RNN models in four out of seven experiments and comparable performance in the other three (see Table 3).

#### Person Activity Dataset

We use the “Human Activity” dataset described in (Rubanova, Chen, and Duvenaud 2019) in two distinct frameworks. The dataset consists of 6554 sequences of activity of humans (e.g. lying, walking, sitting), with a period of 211 ms. we designed two experimental frameworks to evaluate models’ performance. In the *1st Setting*, the baselines are the models described before, and the input representations are unchanged (details in Appendix). LTCs outperform all models and in particular CT-RNNs and neural ODEs with a large margin (Table 4. Note that the CT-RNN architecture is equivalent to the ODE-RNN described in (Rubanova, Chen, and Duvenaud 2019), with the difference of having a state damping factor  $\tau$ .

In the *2nd Setting*, we carefully set up the experiment to match the modifications made by (Rubanova, Chen, and Duvenaud 2019) (See supplements), to obtain a fair compari-

Algorithm	Accuracy
LSTM	83.59% $\pm$ 0.40
CT-RNN	81.54% $\pm$ 0.33
Latent ODE	76.48% $\pm$ 0.56
CT-GRU	85.27% $\pm$ 0.39
LTC (ours)	<b>85.48% <math>\pm</math> 0.40</b>

Table 4: Person activity, 1st setting - n=5

Algorithm	Accuracy
RNN $\Delta_t$ *	0.797 $\pm$ 0.003
RNN-Decay*	0.800 $\pm$ 0.010
RNN GRU-D*	0.806 $\pm$ 0.007
RNN-VAE*	0.343 $\pm$ 0.040
Latent ODE (D enc.)*	0.835 $\pm$ 0.010
ODE-RNN *	0.829 $\pm$ 0.016
Latent ODE(C enc.)*	0.846 $\pm$ 0.013
LTC (ours)	<b>0.882 <math>\pm</math> 0.005</b>

Table 5: Person activity, 2nd setting, n=5 Note: Accuracy for algorithms indicated by \*, are taken directly from (Rubanova, Chen, and Duvenaud 2019) with: RNN  $\Delta_t$  = classic RNN + input delays, D-enc. = RNN encoder C-enc = ODE encoder. RNN-Decay = RNN with exponential decay on hidden states (Mozer, Kazakov, and Lindsey 2017). GRU-D = gated recurrent unit + exponential decay + input imputation (Che et al. 2018).

son between LTCs and a more diverse set of RNN variants discussed in (Rubanova, Chen, and Duvenaud 2019). LTCs show superior performance with a high margin compared to other models. The results are summarized in Table 5).

#### Half-Cheetah Kinematic Modeling

We intended to evaluate how well continuous-time models can capture physical dynamics. To perform this, we collected 25 rollouts of a pre-trained controller for the HalfCheetah-v2 gym environment (Brockman et al. 2016), generated by the MuJoCo physics engine (Todorov, Erez, and Tassa 2012). The task is then to fit the observation space time-series in an autoregressive fashion (Fig. 4). To increase the difficulty, we overwrite 5% of the actions by random actions. The test results are presented in Table 6, and root for the superiority of the performance of LTCs compared to




Figure 4: Half-cheetah physics simulation

other models.

## Related Works

**Time-Continuous Models.** TC networks have become unprecedentedly popular. This is due to the manifestation of several benefits such as adaptive computations, better continuous time-series modeling, memory, and parameter efficiency (Chen et al. 2018). A large number of alternative approaches have tried to improve and stabilize the adjoint method (Gholami, Keutzer, and Biros 2019), use neural ODEs in specific contexts (Rubanova, Chen, and Duvenaud 2019; Lechner et al. 2019) and to characterize them better (Dupont, Doucet, and Teh 2019; Durkan et al. 2019; Jia and Benson 2019; Hanshu et al. 2020; Holl, Koltun, and Thuerey 2020; Quaglino et al. 2020). In this work, we investigated the expressive power of neural ODEs and proposed a new ODE model to improve their expressivity and performance.

**Measures of Expressivity.** Many works have tried to address why deeper networks and particular architectures perform well, and where is the boundary between the approximation capability of shallow and deep networks? In this context, (Montufar et al. 2014) and (Pascanu, Montufar, and Bengio 2013) proposed counting the linear regions of NNs as a measure of expressivity, (Eldan and Shamir 2016) showed that there exists a class of radial functions that smaller networks fail to produce, and (Poole et al. 2016) studied the exponential expressivity of NNs by transient chaos.

These methods are compelling; however, they are bound to particular weight configurations of a given network in order to lower-bound expressivity (Serra, Tjandraatmadja, and Ramalingam 2017; Gabri el et al. 2018; Hanin and Rolnick 2018, 2019; Lee, Alvarez-Melis, and Jaakkola 2019). (Raghu et al. 2017) introduced an interrelated concept which quantifies expressivity by trajectory length. We extended their analysis to CT networks and provided lower-bound for the growth of the trajectory length, proclaiming the superior approximation capabilities of LTCs.

## Conclusions, Scope and Limitations

We introduced liquid time-constant networks. We showed that they could be implemented by arbitrary variable and fixed step ODE solvers, and be trained by backpropagation through time. We demonstrated their bounded and stable dynamics, superior expressivity, and superseding performance in supervised learning time-series prediction tasks, compared to standard and modern deep learning models.

Algorithm	MSE
LSTM	$2.500 \pm 0.140$
CT-RNN	$2.838 \pm 0.112$
Neural ODE	$3.805 \pm 0.313$
CT-GRU	$3.014 \pm 0.134$
LTC (ours)	<b><math>2.308 \pm 0.015</math></b>

Table 6: Sequence modeling. Half-Cheetah dynamics  $n=5$

## Long-term Dependencies

Similar to many variants of time-continuous models, LTCs express the vanishing gradient phenomenon (Pascanu, Mikolov, and Bengio 2013; Lechner and Hasani 2020), when trained by gradient descent. Although the model shows promise on a variety of time-series prediction tasks, they would not be the obvious choice for learning long-term dependencies in their current format.

## Choice of ODE Solver

Performance of time-continuous models is heavily tied to their numerical implementation approach (Hasani 2020). While LTCs perform well with advanced variable-step solvers and the Fused fixed-step solver introduced here, their performance is majorly influenced when off-the-shelf explicit Euler methods are used.

## Time and Memory

Neural ODEs are remarkably fast compared to more sophisticated models such as LTCs. Nonetheless, they lack expressivity. Our proposed model, in their current format, significantly enhances the expressive power of TC models at the expense of elevated time and memory complexity which must be investigated in the future.

## Causality

Models described by ODE semantics inherently possess causal structures (Sch olkopf 2019), especially models that are equipped with recurrent mechanisms to map past experiences to next-step predictions. Studying causality of performant recurrent models such as LTCs would be an exciting future research direction as their semantics resemble *DCMs* (Friston, Harrison, and Penny 2003) with a *bilinear dynamical system* approximation (Penny, Ghahramani, and Friston 2005). Accordingly, a natural application domain would be the control of robots in continuous-time observation and action spaces where causal structures such as LTCs can help improve reasoning (Lechner et al. 2020a).

## Acknowledgments

R.H. and D.R. are partially supported by Boeing. R.H. and R.G. were partially supported by the Horizon-2020 ECSEL Project grant No. 783163 (iDev40). M.L. was supported in part by the Austrian Science Fund (FWF) under grant Z211-N23 (Wittgenstein Award). A.A. is supported by the National Science Foundation (NSF) Graduate Research Fellowship Program. This research work is partially drawn from the PhD dissertation of R.H.



## References

- Bogacki, P.; and Shampine, L. F. 1989. A 3 (2) pair of Runge-Kutta formulas. *Applied Mathematics Letters* 2(4): 321–325.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.
- Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; and Liu, Y. 2018. Recurrent neural networks for multivariate time series with missing values. *Scientific reports* 8(1): 1–12.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, 6571–6583.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* 2(4): 303–314.
- Dormand, J. R.; and Prince, P. J. 1980. A family of embedded Runge-Kutta formulae. *Journal of computational and applied mathematics* 6(1): 19–26.
- Dupont, E.; Doucet, A.; and Teh, Y. W. 2019. Augmented neural odes. In *Advances in Neural Information Processing Systems*, 3134–3144.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019. Neural spline flows. In *Advances in Neural Information Processing Systems*, 7509–7520.
- Eldan, R.; and Shamir, O. 2016. The power of depth for feedforward neural networks. In *Conference on learning theory*, 907–940.
- Friston, K. J.; Harrison, L.; and Penny, W. 2003. Dynamic causal modelling. *Neuroimage* 19(4): 1273–1302.
- Funahashi, K.-I. 1989. On the approximate realization of continuous mappings by neural networks. *Neural networks* 2(3): 183–192.
- Funahashi, K.-i.; and Nakamura, Y. 1993. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural networks* 6(6): 801–806.
- Gabrié, M.; Manoel, A.; Luneau, C.; Macris, N.; Krzakala, F.; Zdeborová, L.; et al. 2018. Entropy and mutual information in models of deep neural networks. In *Advances in Neural Information Processing Systems*, 1821–1831.
- Gholami, A.; Keutzer, K.; and Biros, G. 2019. Anode: Unconditionally accurate memory-efficient gradients for neural odes. *arXiv preprint arXiv:1902.10298*.
- Gleeson, P.; Lung, D.; Grosu, R.; Hasani, R.; and Larson, S. D. 2018. c302: a multiscale framework for modelling the nervous system of *Caenorhabditis elegans*. *Phil. Trans. R. Soc. B* 373(1758): 20170379.
- Hanin, B.; and Rolnick, D. 2018. How to start training: The effect of initialization and architecture. In *Advances in Neural Information Processing Systems*, 571–581.
- Hanin, B.; and Rolnick, D. 2019. Complexity of linear regions in deep networks. *arXiv preprint arXiv:1901.09021*.
- Hanshu, Y.; Jiawei, D.; Vincent, T.; and Jiashi, F. 2020. On Robustness of Neural Ordinary Differential Equations. In *International Conference on Learning Representations*.
- Hasani, R. 2020. *Interpretable Recurrent Neural Networks in Continuous-time Control Environments*. PhD dissertation, Technische Universität Wien.
- Hasani, R.; Amini, A.; Lechner, M.; Naser, F.; Grosu, R.; and Rus, D. 2019. Response characterization for auditing cell dynamics in long short-term memory networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Hasani, R.; Lechner, M.; Amini, A.; Rus, D.; and Grosu, R. 2020. The natural lottery ticket winner: Reinforcement learning with ordinary neural circuits. In *Proceedings of the 2020 International Conference on Machine Learning. JMLR. org*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8): 1735–1780.
- Holl, P.; Koltun, V.; and Thuerey, N. 2020. Learning to Control PDEs with Differentiable Physics. *arXiv preprint arXiv:2001.07457*.
- Hornik, K.; Stinchcombe, M.; and White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2(5): 359–366.
- Hosea, M.; and Shampine, L. 1996. Analysis and implementation of TR-BDF2. *Applied Numerical Mathematics* 20(1-2): 21–37.
- Jia, J.; and Benson, A. R. 2019. Neural jump stochastic differential equations. In *Advances in Neural Information Processing Systems*, 9843–9854.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Koch, C.; and Segev, K. 1998. *Methods in Neuronal Modeling - From Ions to Networks*. MIT press, second edition.
- Lapicque, L. 1907. Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *Journal de Physiologie et de Pathologie Generale* 9: 620–635.
- Lechner, M.; and Hasani, R. 2020. Learning Long-Term Dependencies in Irregularly-Sampled Time Series. *arXiv preprint arXiv:2006.04418*.
- Lechner, M.; Hasani, R.; Amini, A.; Henzinger, T. A.; Rus, D.; and Grosu, R. 2020a. Neural circuit policies enabling auditable autonomy. *Nature Machine Intelligence* 2(10): 642–652.
- Lechner, M.; Hasani, R.; Rus, D.; and Grosu, R. 2020b. Gershgorin Loss Stabilizes the Recurrent Neural Network Compartment of an End-to-end Robot Learning Scheme. In *2020 International Conference on Robotics and Automation (ICRA)*. IEEE.
- Lechner, M.; Hasani, R.; Zimmer, M.; Henzinger, T. A.; and Grosu, R. 2019. Designing worm-inspired neural networks for interpretable robotic control. In *2019 International Conference on Robotics and Automation (ICRA)*, 87–94. IEEE.

- Lee, G.-H.; Alvarez-Melis, D.; and Jaakkola, T. S. 2019. Towards robust, locally linear deep networks. *arXiv preprint arXiv:1907.03207* .
- Montufar, G. F.; Pascanu, R.; Cho, K.; and Bengio, Y. 2014. On the number of linear regions of deep neural networks. In *Advances in neural information processing systems*, 2924–2932.
- Mozer, M. C.; Kazakov, D.; and Lindsey, R. V. 2017. Discrete Event, Continuous Time RNNs. *arXiv preprint arXiv:1710.04110* .
- Pascanu, R.; Mikolov, T.; and Bengio, Y. 2013. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, 1310–1318.
- Pascanu, R.; Montufar, G.; and Bengio, Y. 2013. On the number of response regions of deep feed forward networks with piece-wise linear activations. *arXiv preprint arXiv:1312.6098* .
- Penny, W.; Ghahramani, Z.; and Friston, K. 2005. Bilinear dynamical systems. *Philosophical Transactions of the Royal Society B: Biological Sciences* 360(1457): 983–993.
- Pontryagin, L. S. 2018. *Mathematical theory of optimal processes*. Routledge.
- Poole, B.; Lahiri, S.; Raghu, M.; Sohl-Dickstein, J.; and Ganguli, S. 2016. Exponential expressivity in deep neural networks through transient chaos. In *Advances in neural information processing systems*, 3360–3368.
- Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; and Flannery, B. P. 2007. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. New York, NY, USA: Cambridge University Press, 3 edition.
- Quaglino, A.; Gallieri, M.; Masci, J.; and Koutník, J. 2020. SNODE: Spectral Discretization of Neural ODEs for System Identification. In *International Conference on Learning Representations*.
- Raghu, M.; Poole, B.; Kleinberg, J.; Ganguli, S.; and Dickstein, J. S. 2017. On the expressive power of deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2847–2854. JMLR.org.
- Rubanova, Y.; Chen, R. T.; and Duvenaud, D. 2019. Latent odes for irregularly-sampled time series. *arXiv preprint arXiv:1907.03907* .
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning representations by back-propagating errors. *nature* 323(6088): 533–536.
- Sarma, G. P.; Lee, C. W.; Portegys, T.; Ghayoomie, V.; Jacobs, T.; Alicea, B.; Cantarelli, M.; Currie, M.; Gerkin, R. C.; Gingell, S.; et al. 2018. OpenWorm: overview and recent advances in integrative biological simulation of *Caenorhabditis elegans*. *Phil. Trans. R. Soc. B* 373(1758): 20170382.
- Schölkopf, B. 2019. Causality for Machine Learning. *arXiv preprint arXiv:1911.10500* .
- Serra, T.; Tjandraatmadja, C.; and Ramalingam, S. 2017. Bounding and counting linear regions of deep neural networks. *arXiv preprint arXiv:1711.02114* .
- Shampine, L. F. 1975. Computer solution of ordinary differential equations. *The Initial Value Problem* .
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033. IEEE.
- Wicks, S. R.; Roehrig, C. J.; and Rankin, C. H. 1996. A dynamic network simulation of the nematode tap withdrawal circuit: predictions concerning synaptic function using behavioral criteria. *Journal of Neuroscience* 16(12): 4017–4031.
- Zhuang, J.; Dvornek, N.; Li, X.; Tatikonda, S.; Papademetris, X.; and Duncan, J. 2020. Adaptive Checkpoint Adjoint Method for Gradient Estimation in Neural ODE. In *Proceedings of the 37th International Conference on Machine Learning*. PMLR 119.