Proceedings of Machine Learning Research 120:1–9, 2020

Universal Simulation of Stable Dynamical Systems by Recurrent Neural Nets

Joshua Hanson JMH4@ILLINOIS.EDU
University of Illinois
Urbana, IL 61801

Maxim Raginsky MAXIM@ILLINOIS.EDU
University of Illinois
Urbana, IL 61801

Abstract

It is well known that continuous-time recurrent neural nets are universal approximators for continuous-time dynamical systems. However, existing results provide approximation guarantees only for finite-time trajectories. In this work, we show that infinite-time trajectories generated by dynamical systems that are stable in a certain sense can be reproduced arbitrarily accurately by recurrent neural nets. For a subclass of these stable systems, we provide quantitative estimates on the sufficient number of neurons needed to achieve a specified error tolerance.

Keywords: Dynamical systems, recurrent neural nets, continuous time, universal approximation, simulation, feedback, stability.

1. Introduction

We consider the problem of simulating the output trajectories of a given continuous-time dynamical system by a recurrent neural net. In practice, the parameters of the net can be learned from observed system behavior. In this work, we are primarily concerned with analyzing the expressivity of recurrent neural nets in the context of modeling stable systems. This is useful for a number of practical engineering applications, such as forecasting the behavior of a system with uncertain dynamics using simulated trajectories, replacing subsystems in a modular design with parametric models, or learning controllers for nonlinear plants where there may not be an obvious procedure to determine an optimal controller. Recurrent nets represent a suitable choice for model architecture because the underlying feedforward nets they are constructed from are universal function approximators (Pinkus, 1999), and the feedback connections naturally emulate the intrinsic recursive structure of dynamical systems.

Existing results have established that recurrent nets are capable of generating simulated trajectories that approximate true system trajectories within arbitrary error tolerance over a finite time interval (Sontag, 1992; Funahashi and Nakamura, 1993; Chow and Xiao-Dong Li, 2000; Xiao-Dong Li et al., 2005). All of these authors apply Grönwall's inequality to control the difference between the trajectories of the original system and its approximation, which incurs an exponential degradation of approximation accuracy over time. Without imposing any additional conditions, simulating a system over longer time intervals with the same error tolerance requires an exponentially more accurate model. This is problematic because the number of computation units in the network and number of training samples required to achieve the desired error tolerance will depend on the simulation time scale, which in many applications is unknown *a priori*.

However, for systems satisfying even modest stability conditions, Grönwall's inequality is overly conservative, and one can expect that a more delicate argument can show that approximation error is generally not compounded over time. One such notion of stability can be quantitatively described as saying that the initial condition has asymptotically neglible influence on the long-term behavior of the system trajectory. Many systems of practical interest naturally exhibit this sort of stability behavior. This property can be utilized to establish strict guarantees that the output of a simulating model will remain sufficiently close to the output of the true system for infinite time scales, rather than allowing performance degradation after a fixed time horizon.

In this paper, we consider the problem of universal simulation of dynamical systems that have the property of *uniform asymptotic incremental stability* (Pavlov et al., 2006). Our main result shows that, for any such system, one can find a recurrent neural net, such that the trajectories generated by this net are arbitrarily close to the trajectories generated by the system it approximates on an *infinite* time horizon. This is in contrast to many previous results which only guarantee approximation on a finite time horizon. For stable systems satisfying an additional regularity assumption, we derive quantitative bounds for the size of the recurrent net that achieves a desired error tolerance.

2. Dynamical systems and incremental stability

We are interested in controlled dynamical systems of the form

$$\begin{aligned} \dot{x} &= f(x,u) \qquad x(t) \in \mathbb{R}^n \quad u(t) \in \mathbb{R}^m \\ y &= h(x) \qquad y(t) \in \mathbb{R}^p \end{aligned} \tag{1}$$

where the state transition map $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ and the output map $h : \mathbb{R}^n \to \mathbb{R}^p$ are continuously differentiable. By augmenting the state space to include the equations $\dot{y} = \frac{\partial h}{\partial x} f(x,u)$ if necessary, we may assume without loss of generality that h is given by a linear map $x \mapsto Hx$ for some $H \in \mathbb{R}^{p\times n}$ (Sontag, 1992). We consider uniformly bounded inputs $u : \mathbb{R}_+ \to \mathbb{R}^m$ in the set

$$\mathcal{U} := \{u : \mathbb{R}_+ \to \mathbb{R}^m : \sup_{t\geq 0} |u(t)| \leq R\},$$

where $|\cdot|$ denotes the Euclidean norm. For an input $u \in \mathcal{U}$ and times $0 \leq s \leq t$, we denote the state $x(t)$ at time t that results from initial condition $x(s) = \xi$ at time s by $\varphi^u_{s,t}(\xi)$, referred to as the *flow* or *trajectory* generated by the system (1). We call a set $\mathcal{X} \subseteq \mathbb{R}^n$ *positively invariant* for inputs in $\mathcal{U}$ if, for all $\xi \in \mathcal{X}$, all $u \in \mathcal{U}$, and all $0 \leq s \leq t$, we have $\varphi^u_{s,t}(\xi) \in \mathcal{X}$. We will now restict our attention to systems with the following stability property:

Definition 1 *A dynamical system is* uniformly asympotically incrementally stable *for inputs in* $\mathcal{U}$ *on a positively invariant set* $\mathcal{X}$ *if there exists a function* $\beta : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ *of class* $\mathcal{KL}$[1] *such that*

$$|\varphi^u_{s,t}(\xi) - \varphi^u_{s,t}(\xi')| \leq \beta(|\xi - \xi'|, t-s) \tag{2}$$

holds for all $u \in \mathcal{U}$*, all* $\xi, \xi' \in \mathcal{X}$*, and all* $0 \leq s \leq t$*.*

1. A function $\beta : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ is of class $\mathcal{KL}$ if 1) for any t, the map $h \mapsto \beta(h,t)$ is continuous and strictly increasing and $\beta(0,t) = 0$ and 2) for any h, the map $t \mapsto \beta(h,t)$ is continuous and strictly decreasing and $\lim_{t\to\infty} \beta(h,t) = 0$.

This property quantitatively captures the idea that perturbations to the initial condition have asymptotically negligible influence on the long-term behavior of the system trajectory. One consequence for systems satisfying this definition is that imperfect system models may still be capable of generating outputs that uniformly approximate the outputs of the original system over infinite time intervals. Using this characterization, we can formulate the necessary assumptions of desired approximation and simulation results as regularity conditions on the function β. In constrast, for systems not satisfying this stability condition, a sharp bound on the approximation error degrades exponentially with time (Hirsch and Smale, 1974; Sontag, 1998).

3. Approximate simulation

In many practical applications of dynamical systems modeling, the main criterion for an effective model is that it approximately reproduces both the correct input/output relationships and the internal state dynamics. There are many different ways of expressing this criterion; in this work, we use the following formulation (Sontag, 1992):

Consider two systems Σ and $\tilde{\Sigma}$ described by the following dynamics

$$\Sigma : \quad \begin{aligned} \dot{x} &= f(x,u) \\ y &= Hx \end{aligned}$$

$$\tilde{\Sigma} : \quad \begin{aligned} \dot{\tilde{x}} &= \tilde{f}(\tilde{x},u) \\ \tilde{y} &= \tilde{H}\tilde{x} \end{aligned}$$

with inputs $u(t) \in \mathbb{R}^m$, outputs $y(t) \in \mathbb{R}^p$, and states $x(t) \in \mathbb{R}^n$ and $\tilde{x}(t) \in \mathbb{R}^{\tilde{n}}$. Suppose we are given a compact set $\mathcal{K} \subset \mathbb{R}^n$, a set $\mathcal{U}$ of admissible inputs, and a time interval $T \subseteq \mathbb{R}_+$. We say that $\tilde{\Sigma}$ *simulates* Σ *on sets* $\mathcal{K}$ *and* $\mathcal{U}$ *up to accuracy* ε *for times* $t \in T$ if there exist two continuous maps $\alpha : \mathbb{R}^{\tilde{n}} \to \mathbb{R}^n$ and $\gamma : \mathbb{R}^n \to \mathbb{R}^{\tilde{n}}$ such that, when Σ is initialized at $x(s) = \xi \in \mathcal{K}$, $\tilde{\Sigma}$ is initialized at $\tilde{x}(s) = \gamma(\xi)$, where $s := \inf T$, and any common input $u(\cdot) \in \mathcal{U}$ is supplied to both Σ and $\tilde{\Sigma}$, then

$$|x(t) - \alpha(\tilde{x}(t))| < \varepsilon \qquad \text{and} \qquad |y(t) - \tilde{y}(t)| < \varepsilon$$

for all $t \in T$. In this paper, we consider the case when the simulating system $\tilde{\Sigma}$ is a (continuous-time) *recurrent neural net*, i.e., $\tilde{f}$ has the form

$$\tilde{f}(\tilde{x},u) = -\frac{1}{\tau}\tilde{x} + \sigma_{\tilde{n}}(A\tilde{x} + Bu),$$

where $\tau > 0$ is a positive constant, $A \in \mathbb{R}^{\tilde{n}\times\tilde{n}}$ and $B \in \mathbb{R}^{\tilde{n}\times m}$ are time-invariant matrices, and $\sigma_{\tilde{n}} : \mathbb{R}^{\tilde{n}} \to \mathbb{R}^{\tilde{n}}$ is a diagonal map of the form $\sigma_{\tilde{n}}(\tilde{x}) := [\sigma(\tilde{x}_1) \ \cdots \ \sigma(\tilde{x}_{\tilde{n}})]^{\mathsf{T}}$, where $\sigma : \mathbb{R} \to (0,1)$ is a continuous, strictly increasing function with $\lim_{h\to-\infty} \sigma(h) = 0$ and $\lim_{h\to\infty} \sigma(h) = 1$. Such functions are referred to as *sigmoidal* in the literature on neural nets (Barron, 1993).

4. Simulating stable systems with recurrent neural nets

Consider system (1) with an open positively invariant set $\mathcal{X} \subseteq \mathbb{R}^n$. We impose the following assumptions:

Assumption 2 *There exists a compact subset* $\mathcal{K} \subset \mathcal{X}$ *such that, for any initial condition* $\xi \in \mathcal{K}$, *there exists a compact subset* $\mathcal{X}_\xi \subset \mathcal{X}$, *such that* $\varphi^u_{s,t}(\xi) \in \mathcal{X}_\xi$ *for all* $u \in \mathcal{U}$ *and all* $t \geq s \geq 0$.

Assumption 3 *System (1) is uniformly asympotically incrementally stable on $\mathcal{X}$ for inputs in $\mathcal{U}$, and the function β in equation (2) satisfies the following conditions:*

1. *For any $t \geq 0$, the map $h \mapsto \beta(h,t)$ is differentiable from the right at $h = 0$.*
2. $\displaystyle\int_0^\infty \frac{\partial}{\partial h}\beta(h,t)\Big|_{h=0^+} \mathrm{d}t =: b < \infty.$

Assumption 3 is evidently satisfied by exponentially stable systems with $\beta(h,t) = che^{-\kappa t}$ for some $c, \kappa > 0$, but it also holds for systems with much longer transients, e.g., when $\beta(h,t) = \frac{ch}{(t+1)^{1+\kappa}}$.

Theorem 4 *Consider system (1) and suppose that Assumptions 2 and 3 are satisfied. Then, for any $\varepsilon > 0$, there exists a recurrent neural net of the form*

$$\dot{\tilde{x}} = -\frac{1}{\tau}\tilde{x} + \sigma_{\tilde{n}}(\tilde{A}\tilde{x} + \tilde{B}u)$$
$$\tilde{y} = \tilde{H}\tilde{x}$$

for some $\tau > 0$, $\tilde{A} \in \mathbb{R}^{\tilde{n}\times\tilde{n}}$, $\tilde{B} \in \mathbb{R}^{\tilde{n}\times m}$, and $\tilde{H} \in \mathbb{R}^{p\times\tilde{n}}$ that simulates system (1) on sets $\mathcal{K}$ and $\mathcal{U}$ up to accuracy ε for all $t \in \mathbb{R}_+$. Moreover, the mappings $\alpha : \mathbb{R}^{\tilde{n}} \to \mathbb{R}^n$ and $\gamma : \mathbb{R}^n \to \mathbb{R}^{\tilde{n}}$ that implement the approximate simulation are linear.

4.1. Technical lemmas

To prove the theorem, we will make use of the following lemmas.

Lemma 5 *Let $D\varphi^u_{s,t}(\xi)\cdot v$ denote the directional derivative of $\varphi^u_{s,t}(\xi)$ with respect to ξ in the direction of v. Suppose that Assumptions 2 and 3 are satisfied. Then for any $\xi \in \mathcal{X}$, the induced norm*

$$\|D\varphi^u_{s,t}(\xi)\| := \sup_{|v|=1} |D\varphi^u_{s,t}(\xi)\cdot v|$$

is integrable with respect to t on $[s,\infty)$.

Proof From definitions,

$$\begin{aligned}
\|D\varphi^u_{s,t}(\xi)\| &= \sup_{|v|=1} |D\varphi^u_{s,t}(\xi)\cdot v| \\
&= \sup_{|v|=1} \lim_{h\downarrow 0} \frac{1}{|hv|}|\varphi^u_{s,t}(\xi + hv) - \varphi^u_{s,t}(\xi)| \\
&\leq \sup_{|v|=1} \lim_{h\downarrow 0} \frac{1}{|hv|}\beta(|hv|, t-s) \\
&= \lim_{h\downarrow 0} \frac{1}{h}\beta(h, t-s) \\
&= \frac{\partial}{\partial h}\beta(h,t-s)\Big|_{h=0^+}
\end{aligned}$$

and, by Assumption 3, $\frac{\partial}{\partial h}\beta(h,t-s)|_{h=0^+}$ is integrable with respect to t on $[s,\infty)$. ■

Lemma 6 *Consider two dynamical systems $\dot{x} = f(x,u)$ and $\dot{\hat{x}} = \hat{f}(\hat{x},u)$ with $x(t), \hat{x}(t) \in \mathbb{R}^n$, which generate flows $\varphi^u_{s,t}(\xi)$ and $\hat{\varphi}^u_{s,t}(\xi)$, respectively. Then the following inequality holds for all $t \geq s \geq 0$:*

$$|\varphi^u_{s,t}(\xi) - \hat{\varphi}^u_{s,t}(\xi)| \leq \int_s^t \|D\varphi^u_{r,t}(\hat{\varphi}^u_{s,r}(\xi))\| \cdot |f(\hat{\varphi}^u_{s,r}(\xi), u(r)) - \hat{f}(\hat{\varphi}^u_{s,r}(\xi), u(r))| \, \mathrm{d}r. \tag{3}$$

The proof can be found in van Handel (2007), Chapter 3, Proposition 3.1.3.

Lemma 7 *Consider the C^1 map $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^m$ from system (1) satisfying Assumption 2. Then, for any $\varepsilon > 0$, we can construct:*

- *compact sets $\mathcal{X}_1 \subset \mathcal{X}_2 \subset \mathcal{X}$;*
- *a C^∞ bump function $\rho : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ that satisfies $\rho|_{\mathcal{Z}_1} \equiv 1$ and $\rho|_{(\mathcal{Z}_2)^c} \equiv 0$ for $\mathcal{Z}_i := \mathcal{X}_i \times B^m_R(0)$, $i \in \{1,2\}$, where $B^m_R(0) := \{v \in \mathbb{R}^m : |v| \leq R\}$;*
- *a C^1 map $\hat{f}$ that vanishes outside $\mathcal{Z}_2$, such that, for $(x,u) \in \mathcal{Z}_1$,*

  $$\hat{f}(x,u) = -\frac{1}{\tau}x + T\sigma_\ell(Ax + Bu + \mu) + \nu$$

  for some $\tau > 0$, $T \in \mathbb{R}^{n\times\ell}$, $A \in \mathbb{R}^{\ell\times n}$, $B \in \mathbb{R}^{\ell\times m}$, $\mu \in \mathbb{R}^\ell$, and $\nu \in \mathbb{R}^n$, and

  $$\sup_{(x,u)\in\mathbb{R}^n\times\mathbb{R}^m} |\rho(x,u)f(x,u) - \hat{f}(x,u)| \leq \varepsilon.$$

The proof is omitted due to space limitations; see the full version (Hanson and Raginsky, 2020).

Lemma 8 *The state-space dynamics*

$$\dot{\hat{x}} = -\frac{1}{\tau}\hat{x} + T\sigma_\ell(A\hat{x} + Bu + \mu) + \nu \tag{4}$$

can be simulated with zero loss in accuracy by a system in the form of a recurrent net

$$\dot{\tilde{x}} = -\frac{1}{\tau}\tilde{x} + \sigma_{\tilde{n}}(\tilde{A}\tilde{x} + \tilde{B}u) \tag{5}$$

for some $\tilde{n} \in \mathbb{N}$, $\tilde{A} \in \mathbb{R}^{\tilde{n}\times\tilde{n}}$, $\tilde{B} \in \mathbb{R}^{\tilde{n}\times m}$. That is, there exist matrices $F \in \mathbb{R}^{n\times\tilde{n}}$ and $G \in \mathbb{R}^{\tilde{n}\times n}$, such that $\hat{x}(t) = F\tilde{x}(t)$ for all $t \geq 0$, with initial conditions $\hat{x}(0) = \xi$ and $\tilde{x}(0) = G\xi$.

Proof Following Sontag (1992), we will construct the recurrent net (5) and the matrices F and G in several steps.

Step 1 - Eliminating T: We may assume without loss of generality that the matrix T takes the form $[T_1^\top\ 0]^\top$ with T_1 having full row rank. Then the system (4) can be written as

$$\begin{aligned} \dot{x}_1 &= -\frac{1}{\tau}x_1 + T_1\sigma_\ell(A_1x_1 + A_2x_2 + Bu + \mu) + \nu_1, \qquad & x_1(0) = \xi_1 \\ \dot{x}_2 &= -\frac{1}{\tau}x_2 + \nu_2, & x_2(0) = \xi_2 \end{aligned} \tag{6}$$

where $\xi = [\xi_1^\top\ \xi_2^\top]^\top$ and $\nu = [\nu_1^\top\ \nu_2^\top]^\top$. Since T_1 is surjective, there exist vectors $\tilde{\nu}_1, \tilde{\xi}_1 \in \mathbb{R}^\ell$ such that $T_1\tilde{\nu}_1 = \nu_1$ and $T_1\tilde{\xi}_1 = \xi_1$. Consider the following transformed system:

$$\begin{aligned} \dot{z}_1 &= -\frac{1}{\tau}z_1 + \sigma_\ell(A_1T_1z_1 + A_2x_2 + Bu + \mu) + \tilde{\nu}_1, \qquad & z_1(0) = \tilde{\xi}_1 \\ \dot{x}_2 &= -\frac{1}{\tau}x_2 + \nu_2, & x_2(0) = \xi_2 \end{aligned} \tag{7}$$

The trajectory $(x_1(t), x_2(t))$ of system (6) can be recovered from the trajectory $(z_1(t), x_2(t))$ of system (7) via the transformation $x_1(t) := T_1z_1(t)$. Let $\kappa := \sigma(0)$; then the equation for the dynamics of x_2 may be rewritten as

$$\dot{x}_2 = -\frac{1}{\tau}x_2 + \sigma_{n-r}(0x + 0u) + (\nu_2 - [\kappa\ \cdots\ \kappa]^\top),$$

where $r := \operatorname{rank}(T_1)$. This permits us to combine the two equations in (7) into

$$\dot{\bar{x}} = -\frac{1}{\tau}\bar{x} + \sigma_{\bar{n}}(\bar{A}\bar{x} + \bar{B}u + \bar{\mu}) + \bar{\nu}$$

for suitable matrices $\bar{A} \in \mathbb{R}^{\bar{n}\times\bar{n}}$, $\bar{B} \in \mathbb{R}^{\bar{n}\times m}$, vectors $\bar{\mu}, \bar{\nu} \in \mathbb{R}^{\bar{n}}$, and the initial condition $\bar{x}(0) = \bar{\xi}$, where $\bar{n} := \ell + n - r$ and $\bar{\xi} := [\tilde{\xi}_1^\top\ \xi_2^\top]^\top$.

Step 2 - Eliminating $\bar{\nu}$: Define $\underline{x} := \bar{x} - \tau\bar{\nu}$ and $\theta := \tau\bar{A}\bar{\nu} + \bar{\mu}$. It follows that

$$\dot{\underline{x}} = -\frac{1}{\tau}\underline{x} + \sigma_{\bar{n}}(\bar{A}\underline{x} + \bar{B}u + \theta)$$

with $\underline{x}(0) = \underline{\xi} := \bar{\xi} - \tau\bar{\nu}$, and the trajectory $\bar{x}(t)$ from Step 1 is recovered via $\bar{x}(t) = \underline{x}(t) + \tau\bar{\nu}$.

Step 3 - Eliminating θ: Since $\sigma : \mathbb{R} \to (0,1)$ is bounded, positive, and continuous, the fixed-point equation $z = \tau\sigma(z)$ has at least one nonzero solution ζ, by Brouwer's fixed-point theorem. Consider the following dynamics for $\tilde{x}(t) \in \mathbb{R}^{\bar{n}+1}$:

$$\begin{aligned} \dot{\tilde{x}}_{1:\bar{n}} &= -\frac{1}{\tau}\tilde{x}_{1:\bar{n}} + \sigma_{\bar{n}}(\bar{A}\tilde{x}_{1:\bar{n}} + \frac{1}{\zeta}\theta\tilde{x}_{\bar{n}+1} + \bar{B}u), \qquad & \tilde{x}_{1:\bar{n}}(0) = \underline{\xi} \\ \dot{\tilde{x}}_{\bar{n}+1} &= -\frac{1}{\tau}\tilde{x}_{\bar{n}+1} + \sigma(\tilde{x}_{\bar{n}+1}), & \tilde{x}_{\bar{n}+1}(0) = \zeta \end{aligned}$$

where evidently $\underline{x}(t) = \tilde{x}_{1:\tilde{n}}(t)$ and $\tilde{x}_{\bar{n}+1}(t) \equiv \zeta$ for all t. With $\tilde{n} := \bar{n} + 1$, this system can be represented in the desired form $\dot{\tilde{x}} = -\frac{1}{\tau}\tilde{x} + \sigma_{\tilde{n}}(\tilde{A}\tilde{x} + \tilde{B}u)$ by choosing

$$\tilde{A} := \begin{bmatrix} \bar{A} & \frac{1}{\zeta}\theta \\ 0 & 1 \end{bmatrix}, \quad \tilde{B} := \begin{bmatrix} \bar{B} \\ 0 \end{bmatrix}.$$

Altogether, we have showed that the trajectory of the system (4) with $x(0) = \xi$ can be reproduced with zero loss in accuracy by a recurrent net (5) by expanding the dimension of the state space from n to $n + \ell - r$ and adding one more neuron, for a total of $\tilde{n} = \ell + n - r + 1$ neurons. The matrices F and G can be constructed by retracing the above steps backwards from $\tilde{x}$ to $\hat{x}$ and then forwards from ξ to $\tilde{\xi} := [\underline{\xi}^\mathsf{T}\ \zeta]^\mathsf{T}$. (The affine map $\bar{x}(t) = \underline{x}(t) + \tau\bar{\nu}$ can be implemented as a linear map $\bar{x}(t) = \tilde{x}_{1:\bar{n}}(t) + \frac{\tau}{\zeta}\bar{\nu}\tilde{x}_{\bar{n}+1}$, since $\tilde{x}_{\bar{n}+1}(t) \equiv \zeta$ for all t.) ■

4.2. Proof of Theorem 4

Fix some $\eta > 0$ to be chosen later. Consider system (1) with an open positively invariant set $\mathcal{X} \subseteq \mathbb{R}^n$ satisfying Assumptions 2 and 3. By Lemma 7 (with $\varepsilon \leftarrow \frac{\eta}{b}$), there exist a map $\hat{f} : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ and compact sets $\mathcal{Z}_1 \subset \mathcal{Z}_2 \subset \mathcal{X} \times \mathbb{R}^m$, such that $\hat{f}|_{\mathcal{Z}_1}(x, u) = -\frac{1}{\tau}x + T\sigma_\ell(Ax + Bu + \mu) + \nu$ and $\|\rho f - \hat{f}\|_\infty \le \frac{\eta}{b}$, where ρ is a C^∞ bump function satisfying $\rho|_{\mathcal{Z}_1} = 1$ and $\rho_{(\mathcal{Z}_2)^c} = 0$.

Moreover, by Lemma 7, $\mathcal{Z}_i = \mathcal{X}_i \times B_R^m(0)$. Let $\hat{\varphi}^u_{s,t}(\xi)$ denote the flow generated by the system $\dot{\hat{x}} = \hat{f}(\hat{x}, u)$. Since $\hat{f}$ is a C^1 map that vanishes outside $\mathcal{Z}_2$, we clearly have $\hat{\varphi}^u_{s,t}(\xi) \in \mathcal{X}_2 \subset \mathcal{X}$ for all $\xi \in \mathcal{K}$, all $u \in \mathcal{U}$, and all $t \ge s \ge 0$, since if the trajectory reaches the boundary $\partial\mathcal{X}_2$, it must stop and remain there permanently because $\hat{f}|_{\partial\mathcal{X}_2 \times B_R^m(0)} \equiv 0$. Furthermore, since $\varphi^u_{s,t}(\xi) \in \mathcal{X}_\mathcal{K} = \mathrm{cl}(\cup_{\xi\in\mathcal{K}}\mathcal{X}_\xi)$ for all $t \ge s \ge 0$ by Assumption 2, the flow generated by the system $\dot{x} = (\rho f)(x, u)$ is identically equal to $\varphi^u_{s,t}(\xi)$ because $\rho f|_{\mathcal{X}_\mathcal{K} \times B_R^m(0)} \equiv f|_{\mathcal{X}_\mathcal{K} \times B_R^m(0)}$. Therefore by applying Lemmas 5 and 6, we have

$$
\begin{aligned}
|\varphi^u_{s,t}(\xi) - \hat{\varphi}^u_{s,t}(\xi)| &\le \int_s^t \|D\varphi^u_{r,t}(\hat{\varphi}^u_{s,r}(\xi))\| \cdot |(\rho f)(\hat{\varphi}^u_{s,r}(\xi), u(r)) - \hat{f}(\hat{\varphi}^u_{s,r}(\xi), u(r))|\,\mathrm{d}r \\
&\le \int_s^t \frac{\partial}{\partial h}\beta(h, r-s)\Big|_{h=0^+} \sup_{(x,u)\in\mathcal{Z}_2} |(\rho f)(x, u) - \hat{f}(x, u)|\,\mathrm{d}r \\
&\le b \cdot \frac{\eta}{b} = \eta.
\end{aligned}
$$

By Lemma 8, the system

$$
\dot{\hat{x}} = -\frac{1}{\tau}\hat{x} + T\sigma_\ell(A\hat{x} + Bu + \mu) + \nu
$$

can be simulated with zero loss in accuracy by a system in the form of a recurrent net

$$
\dot{\tilde{x}} = -\frac{1}{\tau}\tilde{x} + \sigma_{\tilde{n}}(\tilde{A}\tilde{x} + \tilde{B}u)
$$

For the above recurrent net, let $\tilde{y}(t) := \tilde{H}\tilde{x}(t)$ with $\tilde{H} := HF$, where $H \in \mathbb{R}^{p\times n}$ is the linear output map of the original system (1) and $F \in \mathbb{R}^{n\times\tilde{n}}$ is the linear map given by Lemma 8. Then $H\hat{x}(t) = HF\tilde{x}(t) = \tilde{H}\tilde{x}(t)$ for all $t \ge 0$, and consequently

$$
\begin{aligned}
|y(t) - \tilde{y}(t)| &= |Hx(t) - \tilde{H}\tilde{x}(t)| \\
&= |Hx(t) - H\hat{x}(t)| \\
&\le \|H\||x(t) - \hat{x}(t)| \\
&\le \|H\|\eta.
\end{aligned}
$$

Choosing $\eta < \min(\varepsilon, \frac{\varepsilon}{\|H\|})$ gives $|x(t) - F\tilde{x}(t)| < \varepsilon$ and $|y(t) - \tilde{y}(t)| < \varepsilon$ for all $t \ge 0$, with $x(0) = \xi$ and $\tilde{x}(0) = G\xi$, which completes the proof.

5. Quantitative approximation bounds for Barron-class systems

Utilizing quantitative approximation bounds developed for feedforward nets, we can develop similar results for recurrent nets. For these bounds to hold, it is necessary for the vector field $f(x, u)$ of the original system (1) to satisfy certain regularity conditions (Barron, 1993):

Definition 9 *We say that a continuous function $f : \mathbb{R}^d \to \mathbb{R}$ belongs to the* Barron class *if*

$$C_f := \int_{\mathbb{R}^d} |\omega||\tilde{f}(\omega)| \, \mathrm{d}\omega < \infty,$$

where $\tilde{f} : \mathbb{R}^d \to \mathbb{R}$ is the Fourier transform of f.

Proposition 10 *Let a continuous function $f : \mathbb{R}^d \to \mathbb{R}$ be given, with $C_f < \infty$. Then for every $r > 0$ and every $N \in \mathbb{N}$, there exists a feedforward neural net $g : \mathbb{R}^d \to \mathbb{R}$ of the form*

$$g(z) = \sum_{k=1}^{N} c_k \sigma(a_k \cdot z + b_k) + c_0,$$

such that

$$\sup_{z \in B_r^d(0)} |f(z) - g(z)| \le \frac{2rC_f}{\sqrt{N}}.$$

The proof can be found in Barron (1993) or in Yukich et al. (1995). Note that the constant C_f depends implicitly on the input-space dimension d. If each coordinate of the state transistion map $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ from system (1) belongs to the Barron class, then we can bound the number of computation units (neurons) in a recurrent net that simulates system (1).

Proposition 11 *The number of computation units sufficient to guarantee the result of Theorem 4 with accuracy ε is*

$$\tilde{n} \ge n + 1 + \frac{16(C_f b \|H\| \Delta)^2 n}{\varepsilon^2}$$

where C_f and b are defined earlier, and $\Delta := \sup_{x \in \mathcal{X}_{\mathcal{K}}} |x| + R + \frac{\varepsilon}{2\|H\|}$.

Remark 12 The constant C_f may implicitly depend on the total dimension $n + m$.

Proof The desired underlying feedforward net is constructed in the proof of Lemma 7, such that

$$\sup_{(x,u) \in \mathcal{Z}_1} |(\rho f)(x, u) - g(x, u)| \le \frac{\eta}{2b},$$

where $\mathcal{Z}_1$ is a compact subset of $\mathbb{R}^n \times \mathbb{R}^m$ contained in the ball of radius Δ. On the other hand, Proposition 10 gives

$$\sup_{(x,u) \in \mathcal{Z}_1} |(\rho f)(x, u) - g(x, u)| \le \frac{2C_f \Delta \sqrt{n}}{\sqrt{\ell}},$$

where ℓ is the number of neurons in g. To achieve the desired inequality, it suffices to take $\ell \ge \frac{16C_f^2 b^2 \Delta^2 n}{\eta^2}$. From the proof of Theorem 4, we set $\eta < \frac{\varepsilon}{\|H\|}$, and from Lemma 8 we know that $\tilde{n} \ge n + \ell + 1$ neurons suffice. ■

Acknowledgments

This work was supported in part by the National Science Foundation under the Center for Advanced Electronics through Machine Learning (CAEML) I/UCRC award no. CNS-16-24811.

References

Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993. ISSN 1557-9654. doi: 10.1109/18.256500.

T. W. S. Chow and Xiao-Dong Li. Modeling of continuous time dynamical systems with input by recurrent neural networks. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 47(4):575–578, April 2000. doi: 10.1109/81.841860.

Ken-Ichi Funahashi and Yuichi Nakamura. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, 6(6):801 – 806, 1993. ISSN 0893-6080. doi: https://doi.org/10.1016/S0893-6080(05)80125-X.

Joshua Hanson and Maxim Raginsky. Universal simulation of stable dynamical systems by recurrent neural nets, May 2020. URL https://uofi.box.com/s/3zs9q2jb7xpbs8qflyco6er3muvwfagi.

Morris W. Hirsch and Stephen Smale. Nonautonomous equations and differentiability of flows. In *Differential equations, dynamical systems, and linear algebra*, chapter 15, pages 296–303. Academic Press, 1974.

Alexey Pavlov, Nathan van de Wouw, and Henk Nijmeijer. *Uniform Output Regulation of Nonlinear Systems: A Convergent Dynamics Approach*. Springer, 2006.

Allan Pinkus. Approximation theory of the MLP model in neural networks. *Acta Numerica*, pages 143–195, 1999.

Eduardo D. Sontag. Neural Nets as Systems Models and Controllers. In *Workshop on Adaptive and Learning Systems*, 1992.

Eduardo D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer, 2nd edition, 1998.

Ramon van Handel. *Filtering, Stability, and Robustness*. PhD thesis, California Institute of Technology, 2007.

Xiao-Dong Li, J. K. L. Ho, and T. W. S. Chow. Approximation of dynamical time-variant systems by continuous-time recurrent neural networks. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 52(10):656–660, Oct 2005. doi: 10.1109/TCSII.2005.852006.

Joseph E. Yukich, Maxwell B. Stinchcombe, and Halbert White. Sup-norm approximation bounds for networks through probabilistic methods. *IEEE Transactions on Information Theory*, 41(4): 1021–1027, July 1995.