



Gated Recurrent Units Viewed Through the Lens of Continuous Time Dynamical Systems

Ian D. Jordan^{1,2}, Piotr Aleksander Sokół³ and Il Memming Park^{1,2,3*}

¹ Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY, United States, ² Institute of Advanced Computational Science, Stony Brook University, Stony Brook, NY, United States, ³ Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, NY, United States

Gated recurrent units (GRUs) are specialized memory elements for building recurrent neural networks. Despite their incredible success on various tasks, including extracting dynamics underlying neural data, little is understood about the specific dynamics representable in a GRU network. As a result, it is both difficult to know a priori how successful a GRU network will perform on a given task, and also their capacity to mimic the underlying behavior of their biological counterparts. Using a continuous time analysis, we gain intuition on the inner workings of GRU networks. We restrict our presentation to low dimensions, allowing for a comprehensive visualization. We found a surprisingly rich repertoire of dynamical features that includes stable limit cycles (nonlinear oscillations), multi-stable dynamics with various topologies, and homoclinic bifurcations. At the same time we were unable to train GRU networks to produce continuous attractors, which are hypothesized to exist in biological neural networks. We contextualize the usefulness of different kinds of observed dynamics and support our claims experimentally.

Keywords: recurrent neural network, dynamical systems, continuous time, bifurcations, time-series

OPEN ACCESS

Edited by:

Martin A. Giese,
University of Tübingen, Germany

Reviewed by:

Lee DeVille,
University of Illinois at
Urbana-Champaign, United States
Mario Negrello,
Erasmus Medical Center, Netherlands
J. Michael Herrmann,
University of Edinburgh,
United Kingdom

*Correspondence:

Il Memming Park
memming.park@stonybrook.edu

Received: 09 March 2021

Accepted: 28 June 2021

Published: 22 July 2021

Citation:

Jordan ID, Sokół PA and Park IM
(2021) Gated Recurrent Units Viewed
Through the Lens of Continuous Time
Dynamical Systems.
Front. Comput. Neurosci. 15:678158.
doi: 10.3389/fncom.2021.678158

1. INTRODUCTION

Recurrent neural networks (RNNs) can capture and utilize sequential structure in natural and artificial languages, speech, video, and various other forms of time series. The recurrent information flow within an RNN implies that the data seen in the past has influence on the current state of the RNN, forming a mechanism for having memory through (nonlinear) temporal traces that encode both *what* and *when*. Past works have used RNNs to study neural population dynamics (Costa et al., 2017), and have demonstrated qualitatively similar dynamics between biological neural networks and artificial networks trained under analogous conditions (Mante et al., 2013; Sussillo et al., 2015; Cueva et al., 2020). In turn, this brings into question the efficacy of using such networks as a means to study brain function. With this in mind, training standard vanilla RNNs to capture long-range dependences within a sequence is challenging due to the vanishing gradient problem (Hochreiter, 1991; Bengio et al., 1994). Several special RNN architectures have been proposed to mitigate this issue, notably the long short-term memory (LSTM) units (Hochreiter and Schmidhuber, 1997) which explicitly guard against unwanted corruption of the information stored in the hidden state until necessary. Recently, a simplification of the LSTM called the *gated recurrent unit* (GRU) (Cho et al., 2014) has become popular in the computational neuroscience and machine learning communities thanks to its performance in speech (Prabhavalkar et al., 2017), music (Choi et al., 2017), video (Dwivedi et al., 2018), and extracting nonlinear dynamics underlying neural data

(Pandarinath et al., 2018). However, certain mechanistic tasks, specifically unbounded counting, come easy to LSTM networks but not to GRU networks (Weiss et al., 2018).

Despite these empirical findings, we lack systematic understanding of the internal time evolution of GRU's memory structure and its capability to represent nonlinear temporal dynamics. Such an understanding will make clear what specific tasks (natural and artificial) can or cannot be performed (Bengio et al., 1994), how computation is implemented (Beer, 2006; Sussillo and Barak, 2012), and help to predict qualitative behavior (Beer, 1995; Zhao and Park, 2016). In addition, a great deal of the literature discusses the local dynamics (equilibrium points) of RNNs (Bengio et al., 1994; Sussillo and Barak, 2012), but a complete theory requires an understanding of the global properties as well (Beer, 1995). Furthermore, a deterministic understanding of a GRU network's topological structure will provide fundamental insight as to a trained network's generalization ability, and therefore help in understanding how to seed RNNs for specific tasks (Doya, 1993; Sokół et al., 2019).

In general, the hidden state dynamics of an RNN can be written as $\mathbf{h}_{t+1} = f(\mathbf{h}_t, \mathbf{x}_t)$ where \mathbf{x}_t is the current input in a sequence indexed by t , f is a nonlinear function, and \mathbf{h}_t represents the hidden memory state that carries all information responsible for future output. In the absence of input, \mathbf{h}_t evolves over time on its own:

$$\mathbf{h}_{t+1} = f(\mathbf{h}_t) \quad (1)$$

where $f(\cdot) := f(\cdot, \mathbf{0})$ for notational simplicity. In other words, we can consider the temporal evolution of memory stored within an RNN as a trajectory of an autonomous dynamical system defined by Equation (1), and use dynamical systems theory to further investigate and classify the temporal features obtainable in an RNN. In this paper, we intend on providing a deep intuition of the inner workings of the GRU through a continuous time analysis. While RNNs are traditionally implemented in discrete time, we show in the next section that this form of the GRU can be interpreted as a numerical approximation of an underlying system of ordinary differential equations. Historically, discrete time systems are often more challenging to analyze when compared with their continuous time counterparts, primarily due to their more *jumpy* nature, allowing for more complex dynamics in low-dimensions (Pasemann, 1997; Laurent and von Brecht, 2017). Due to the relatively continuous nature of many abstract and physical systems, it may be of great use to analyze the underlying continuous time system of a trained RNN directly in some contexts, while interpreting the added dynamical complexity from the discretization as anomalies from numerical analysis (LeVeque and Leveque, 1992; Thomas, 1995; He et al., 2016; Heath, 2018). Furthermore, the recent development of *Neural Ordinary Differential Equations* have catalyzed the computational neuroscience and machine learning communities to turn much of their attention to continuous-time implementations of neural networks (Chen et al., 2018; Morrill et al., 2021).

We discuss a vast array of observed local and global dynamical structures, and validate the theory by training GRUs to predict

time series with prescribed dynamics. As to not compromise the presentation, we restrict our analysis to low dimensions for easy visualization (Beer, 1995; Zhao and Park, 2016). However, given a trained GRU of any finite dimension, the findings here still apply, and can be applied with further analysis on a case by case basis (more information on this in the discussion). Furthermore, to ensure our work is accessible we will assume a pedagogical approach in our delivery. We recommend Meiss (Meiss, 2007) for more background on the subject.

2. UNDERLYING CONTINUOUS TIME SYSTEM OF GATED RECURRENT UNITS

The GRU uses two internal gating variables: the *update gate* \mathbf{z}_t which protects the d -dimensional hidden state $\mathbf{h}_t \in \mathbb{R}^d$ and the *reset gate* \mathbf{r}_t which allows overwriting of the hidden state and controls the interaction with the input $\mathbf{x}_t \in \mathbb{R}^p$.

$$\mathbf{z}_t = \sigma(\mathbf{W}_z \mathbf{x}_t + \mathbf{U}_z \mathbf{h}_{t-1} + \mathbf{b}_z) \quad (2)$$

$$\mathbf{r}_t = \sigma(\mathbf{W}_r \mathbf{x}_t + \mathbf{U}_r \mathbf{h}_{t-1} + \mathbf{b}_r) \quad (3)$$

$$\mathbf{h}_t = (1 - \mathbf{z}_t) \odot \tanh(\mathbf{W}_h \mathbf{x}_t + \mathbf{U}_h (\mathbf{r}_t \odot \mathbf{h}_{t-1}) + \mathbf{b}_h) + \mathbf{z}_t \odot \mathbf{h}_{t-1} \quad (4)$$

where $\mathbf{W}_z, \mathbf{W}_r, \mathbf{W}_h \in \mathbb{R}^{d \times p}$ and $\mathbf{U}_z, \mathbf{U}_r, \mathbf{U}_h \in \mathbb{R}^{d \times d}$ are the parameter matrices, $\mathbf{b}_z, \mathbf{b}_r, \mathbf{b}_h \in \mathbb{R}^d$ are bias vectors, \odot represents element-wise multiplication, and $\sigma(\mathbf{z}) = 1/(1 + e^{-\mathbf{z}})$ is the element-wise logistic sigmoid function. Note that the hidden state is asymptotically contained within $[-1, 1]^d$ due to the saturating nonlinearities, implying that if the state is initialized outside of this trapping region, it must eventually enter it in finite time and remain in it for all later time.


Note that the update gate \mathbf{z}_t controls how fast each dimension of the hidden state decays, providing an adaptive time constant for memory. Specifically, as $\lim_{z_t \rightarrow 1} \mathbf{h}_t = \mathbf{h}_{t-1}$, GRUs can implement perfect memory of the past and ignore \mathbf{x}_t . Hence, a d -dimensional GRU is capable of keeping a near constant memory through the update gate—near constant since $0 < [\mathbf{z}_t]_j < 1$, where $[\cdot]_j$ denotes j -th component of a vector. Moreover, the autoregressive weights (mainly \mathbf{U}_h and \mathbf{U}_r) can support time evolving memory (Laurent and von Brecht, 2017 considered this a hindrance and proposed removing all complex dynamical behavior in a simplified GRU).

To investigate the memory structure further, let us consider the dynamics of the hidden state in the absence of input, i.e., $\mathbf{x}_t = 0, \forall t$, which is of the form Equation (1). From a dynamical system's point of view, all inputs to the system can be understood as perturbations to the autonomous system, and therefore have no effect on the set of achievable dynamics. To utilize the rich descriptive language of continuous time dynamical systems theory, we recognize the autonomous GRU-RNN as a weighted forward Euler discretization to the following continuous time dynamical system:

$$\dot{\mathbf{z}}(t) = \sigma(\mathbf{U}_z \mathbf{h}(t) + \mathbf{b}_z) \quad (5)$$

$$\dot{\mathbf{r}}(t) = \sigma(\mathbf{U}_r \mathbf{h}(t) + \mathbf{b}_r) \quad (6)$$

$$\dot{\mathbf{h}} = (1 - \mathbf{z}(t)) \odot (\tanh(\mathbf{U}_h (\mathbf{r}(t) \odot \mathbf{h}(t)) + \mathbf{b}_h) - \mathbf{h}(t)) \quad (7)$$



where $\dot{\mathbf{h}} \equiv \frac{d\mathbf{h}(t)}{dt}$. Since both $\sigma(\cdot)$ and $\tanh(\cdot)$ are smooth, this continuous limit is justified and serves as a basis for further analysis, as all GRU networks are attempting to approximate this continuous limit. In the following, GRU will refer to the continuous time version (Equation 7). Note that the update gate $\mathbf{z}(t)$ again plays the role of a state-dependent time constant for memory decay. We note, however, that $\mathbf{z}(t)$ adjusts flow speed point-wise, resulting in non-constant nonlinear slowing of all trajectories, as $\mathbf{z}(t) \in (0, 1)$. Since $1 - \mathbf{z}(t) > 0$, and thus cannot change sign, it acts as a homeomorphism between (Equation 7) and the same system with this leading multiplicative term removed. Therefore, it does not change the topological structure of the dynamics (Kuznetsov, 1998), and we can safely ignore the effects of $\mathbf{z}(t)$ in the following theoretical analysis (sections 3, 4). In these sections we set $\mathbf{U}_z = 0$ and $\mathbf{b}_z = 0$. A derivation of the continuous time GRU can be found in section 1 of the **Supplementary Material**. Further detail on the effects of $\mathbf{z}(t)$ are discussed in the final section of this paper.

3. STABILITY ANALYSIS OF A ONE DIMENSIONAL GRU

For a 1D GRU¹ ($d = 1$), Equation (7) reduces to a one dimensional dynamical system where every variable is a scalar. The expressive power of a 1D GRU is quite limited, as only three stability structures (topologies) exist (see section 2 in the **Supplementary Material**): **(Figure 1A)** a single stable node, **(Figure 1B)** a stable node and a half-stable node, and **(Figure 1C)** two stable nodes separated by an unstable node (see **Figure 1**). The corresponding time evolution of the hidden state are (A) decay to a fixed value, (B) decay to a fixed value, but from one direction halt at an intermediate value until perturbed, or (C) decay to one of two fixed values (bistability). The bistability can be used to model a switch, such as in the context of simple decision making, where inputs can perturb the system back and forth between states.

The topology the GRU takes is determined by its parameters. If the GRU begins in a region of the parameter space corresponding to (A), we can smoothly vary the parameters to transverse (B) in the parameter space, and end up at (C). This is commonly known as a saddle-node bifurcation. Speaking generally, a bifurcation is

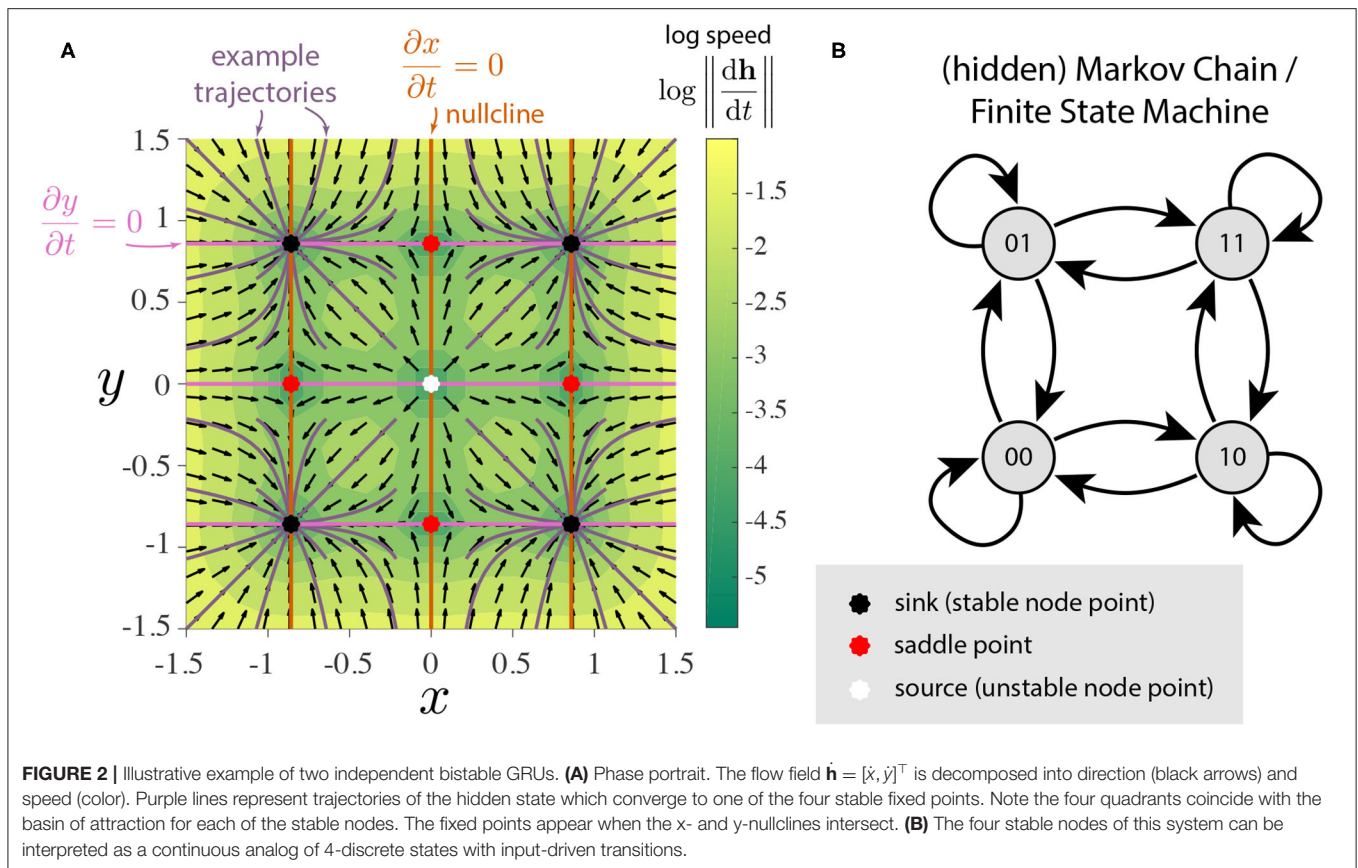
the change in topology of a dynamical system, resulting from a smooth change in parameters. The point in parameter space at which the bifurcation occurs is called the bifurcation point (e.g., **Figure 1B**), and we will refer to the fixed point that changes its stability at the bifurcation point as the *bifurcation fixed point* (e.g., the half-stable fixed point in **Figure 1B**). The codimension of a bifurcation is the number of parameters which must vary in order to remain on the bifurcation manifold. In the case of our example, a saddle-node bifurcation is codimension-1 (Kuznetsov, 1998). Right before transitioning to (B), from (A), the flow near where the half-stable node would appear can exhibit arbitrarily slow flow. We will refer to these as *slow points* (Sussillo and Barak, 2012). In this context, slow points allow for metastable states, where a trajectory will flow toward the slow point, remain there for a period of time, before moving to the stable fixed point.

4. ANALYSIS OF A TWO DIMENSIONAL GRU

We will see that the addition of a second GRU opens up a substantial variety of possible topological structures. For notational simplicity, we denote the two dimensions of \mathbf{h} as x and y . We visualize the flow fields defined by Equation (7) in 2-dimensions as *phase portraits* which reveal the topological structures of interest (Meiss, 2007). For starters, the phase portrait of two independent bistable GRUs can be visualized as **Figure 2A**. It clearly shows 4 stable states as expected, with a total of 9 fixed points. This could be thought of as a continuous-time continuous-space implementation of a finite state machine with 4 states (**Figure 2B**). The 3 types of observed fixed points—stable (sinks), unstable (sources), and saddle points—exhibit locally linear dynamics, however, the global geometry is nonlinear and their topological structures can vary depending on their arrangement.

We explored stability structures attainable by 2D GRUs. Due to the relatively large number of observed topologies, this section's main focus will be on demonstrating all observed local and global dynamical features obtainable by 2D GRUs. A catalog of all known topologies can be found in section 3 of the **Supplementary Material**, along with the parameters of every phase portrait depicted in this paper. We cannot say whether or not this catalog is exhaustive, but the sheer number of structures found is a testament to the expressive power of the GRU network, even in low dimensions.

¹The number/dimension of GRUs references to the dimension of the hidden state dynamics.



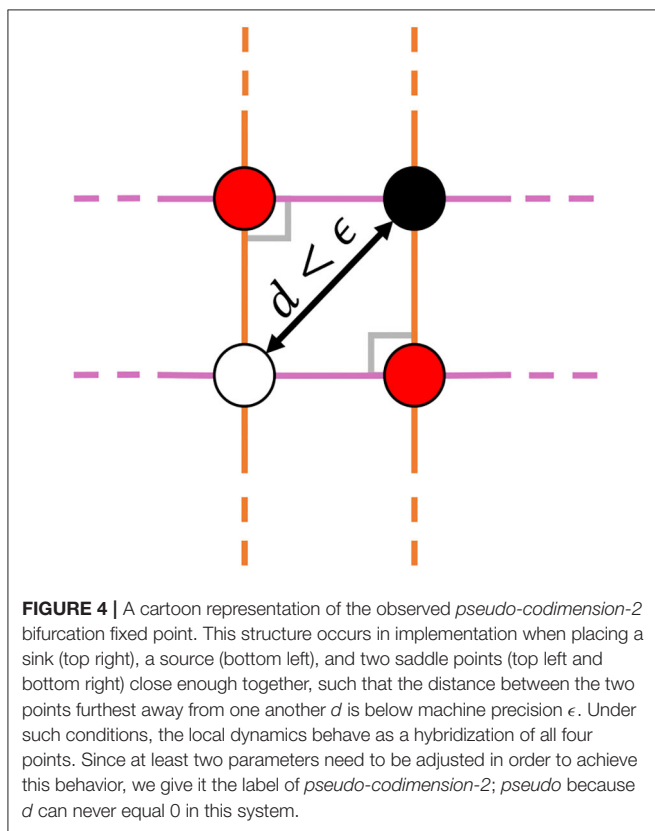
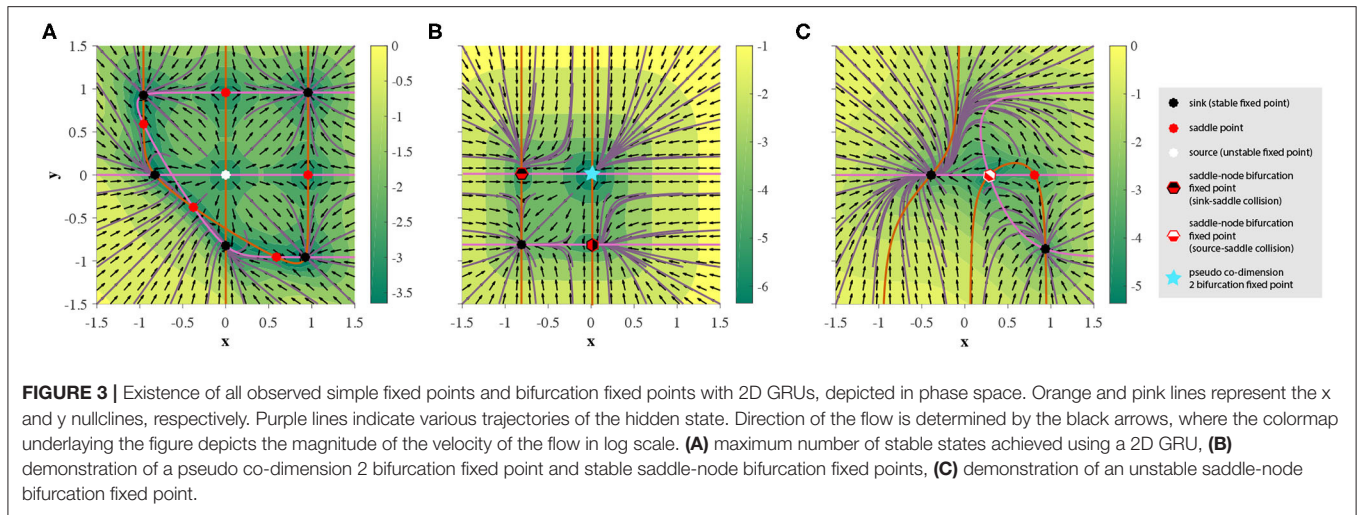
Before proceeding, let us take this time to describe all the local dynamical features observed. In addition to the previously mentioned three types of fixed points, 2D GRUs can exhibit a variety of bifurcation fixed points, resulting from regions of parameter space that separate all topologies restricted to simple fixed points (i.e. stable, unstable, and saddle points). Behaviorally speaking, these fixed points act as hybrids between the previous three, resulting in a much richer set of obtainable dynamics. In **Figure 3**, we show all observed types of fixed points². While no codimension-2 bifurcation fixed points were observed in the 2D GRU system, a sort of *pseudo-codimension-2* bifurcation fixed point was seen by placing a sink, source, and two saddle points sufficiently close together, such that, when implemented, all four points remain below machine precision, thereby acting as a single fixed point. **Figure 4** further demonstrates this concept, and **Figure 3B** depicts an example. We will discuss later that this sort of pseudo-bifurcation point allows the system to exhibit *homoclinic-like* behavior on a two dimensional compact set. In **Figure 3A**, we see 11 fixed points, the maximum number of fixed points observed in a 2D GRU system. A closer look at this system

²2D GRUs feature both codimension-1 and pseudo-codimension-2 bifurcation fixed points. In codimension-1, we have the saddle-node bifurcation fixed point, as expected from its existence in the 1D GRU case. These can be thought of as both the fusion of a stable fixed point and a saddle point, and the fusion of an unstable fixed point and a saddle point. We will refer to these fixed points as saddle-node bifurcation fixed points of the first kind and second kind, respectively.

reveals one interpretation as a continuous analog of 5-discrete states with input-driven transitions, similar to that depicted in **Figure 2**. This imposes a possible upper bound on the network's capacity to encode a finite set of states in this manner.

The addition of bifurcation fixed points opens the door to dynamically realize more sophisticated models. Take for example the four state system depicted in **Figure 3B**. If the hidden state is set to initialize in the first quadrant of phase space [i.e., $(0, \infty)^2$], the trajectory will flow toward the pseudo-codimension-2 bifurcation fixed point at the origin. Introducing noise through the input will stochastically cause the trajectory to approach the stable fixed point at $(-1, -1)$ either directly, or by first flowing into one of the two saddle-node bifurcation fixed points of the first kind. Models of this sort can be used in a variety of applications, such as perceptual decision making (Wong and Wang, 2006; Churchland and Cunningham, 2014).

We will begin our investigation into the non-local dynamics observed with 2D GRUs by showing the existence of an Andronov-Hopf bifurcation, where a stable fixed point bifurcates into an unstable fixed point surrounded by a limit cycle. A limit cycle is an attracting set with a well defined basin of attraction. However, unlike a stable fixed point, where trajectories initialized in the basin of attraction flow toward a single point, a limit cycle pulls trajectories into a stable periodic orbit. If the periodic orbit surrounds an unstable fixed point the attractor is *self-exciting*, otherwise it is a *hidden attractor* (Meiss, 2007). While hidden



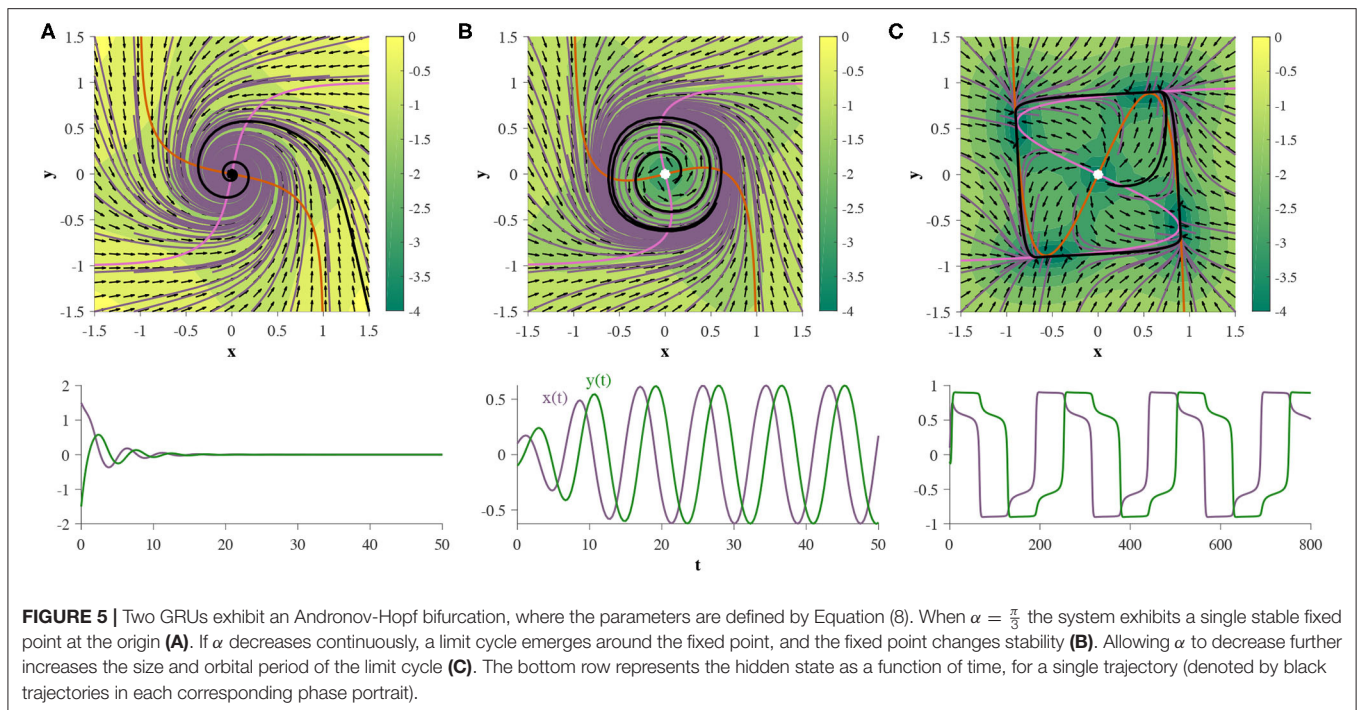
attractors have been observed in various 2D systems, they have not been found in the 2D GRU system, and we conjecture that they do not exist. If all parameters are set to zero except for the hidden state weights, which are parameterized as a rotation matrix with an associated gain, we can introduce rotation into the vector field as a function of gain and rotation angle. Properly tuning these parameters will give rise to a limit cycle; a result of the saturating nonlinearity impeding the rotating flow velocity

sufficiently distant from the origin, thereby pulling trajectories toward a closed orbit.

For $\alpha, \beta \in \mathbb{R}^+$ and $s \in \mathbb{R}$,

$$U_z, U_r, \mathbf{b}_z, \mathbf{b}_h = 0, U_h = \beta \begin{bmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{bmatrix}, \mathbf{b}_r = s \quad (8)$$

Let $\beta = 3$ and $s = 0$. If $\alpha = \frac{\pi}{3}$, the system has a single stable fixed point (stable spiral), as depicted in **Figure 5A**. If we continuously decrease α , the system undergoes an Andronov-Hopf bifurcation at approximately $\alpha = \frac{\pi}{3.8}$. As α continuously decreases, the orbital period increases, and as the nullclines can be made arbitrarily close together, the length of this orbital period can be set arbitrarily. **Figure 5B** shows an example of a relatively short orbital period, and **Figure 5C** depicts the behavior seen for slower orbits. If we continue allowing α to decrease, the system will undergo four simultaneous saddle-node bifurcations, and end up in a state topologically equivalent to that depicted in **Figure 2A**. **Figure 6A** depicts regions of the parameter space of Equation (7) parameterized by Equation (8), where the Andronov-Hopf bifurcation manifolds can be clearly seen. **Figure 6B** demonstrates one effect the reset gate can have on the frequency of the oscillations. If we alter the bias vector b_r , the expected oscillation period changes for regions of the $\alpha - \beta$ parameter space which exhibit a limit cycle. Computationally speaking, limit cycles are a common dynamical structure for modeling neuron bursting (Izhikevich, 2007), taking place in many foundational works including the Hodgkin-Huxley model (Hodgkin and Huxley, 1952) and the FitzHugh-Nagumo Model (FitzHugh, 1961). Such dynamics also arise in various population level dynamics in artificial tasks, such as sine wave generation (Sussillo and Barak, 2012). Furthermore, initializing the hidden state matrix U_h of an even dimensional continuous-time RNN (tanh or GRU) with 2×2 blocks along the diagonal and zeros everywhere else is theoretically shown to aid in learning



long-term dependencies, when all the blocks act as decoupled oscillators (Sokół et al., 2019).

Regarding the second non-local dynamical feature, it can be shown that a 2D GRU can undergo a homoclinic bifurcation, where a periodic orbit (in this case a limit cycle) expands and collides with a saddle at the bifurcation point. At this bifurcation point the system exhibits a homoclinic orbit, where trajectories initialized on the orbit fall into the same fixed point in both forward and backward time. In order to demonstrate this behavior, let the parameters of the network be defined as follows:

For $\gamma \in \mathbb{R}$,

$$\mathbf{U}_z, \mathbf{U}_r, \mathbf{b}_z, \mathbf{b}_r = 0, \mathbf{U}_h = 3 \begin{bmatrix} \cos \frac{\pi}{20} & \sin \frac{\pi}{20} \\ -\sin \frac{\pi}{20} & \cos \frac{\pi}{20} \end{bmatrix}, \mathbf{b}_h = \begin{bmatrix} 0.32 \\ \gamma \end{bmatrix} \quad (9)$$

Under this parameterization the 2D GRU system exhibits a homoclinic orbit when $\gamma = 0.054085$ (Figure 7). In order to showcase this bifurcation as well as the previous Andronov-Hopf bifurcation sequentially in action we turn to Figure 8, where the parameters are defined by Equation (9) and γ is initialized at 0.051 in Figure 8A.

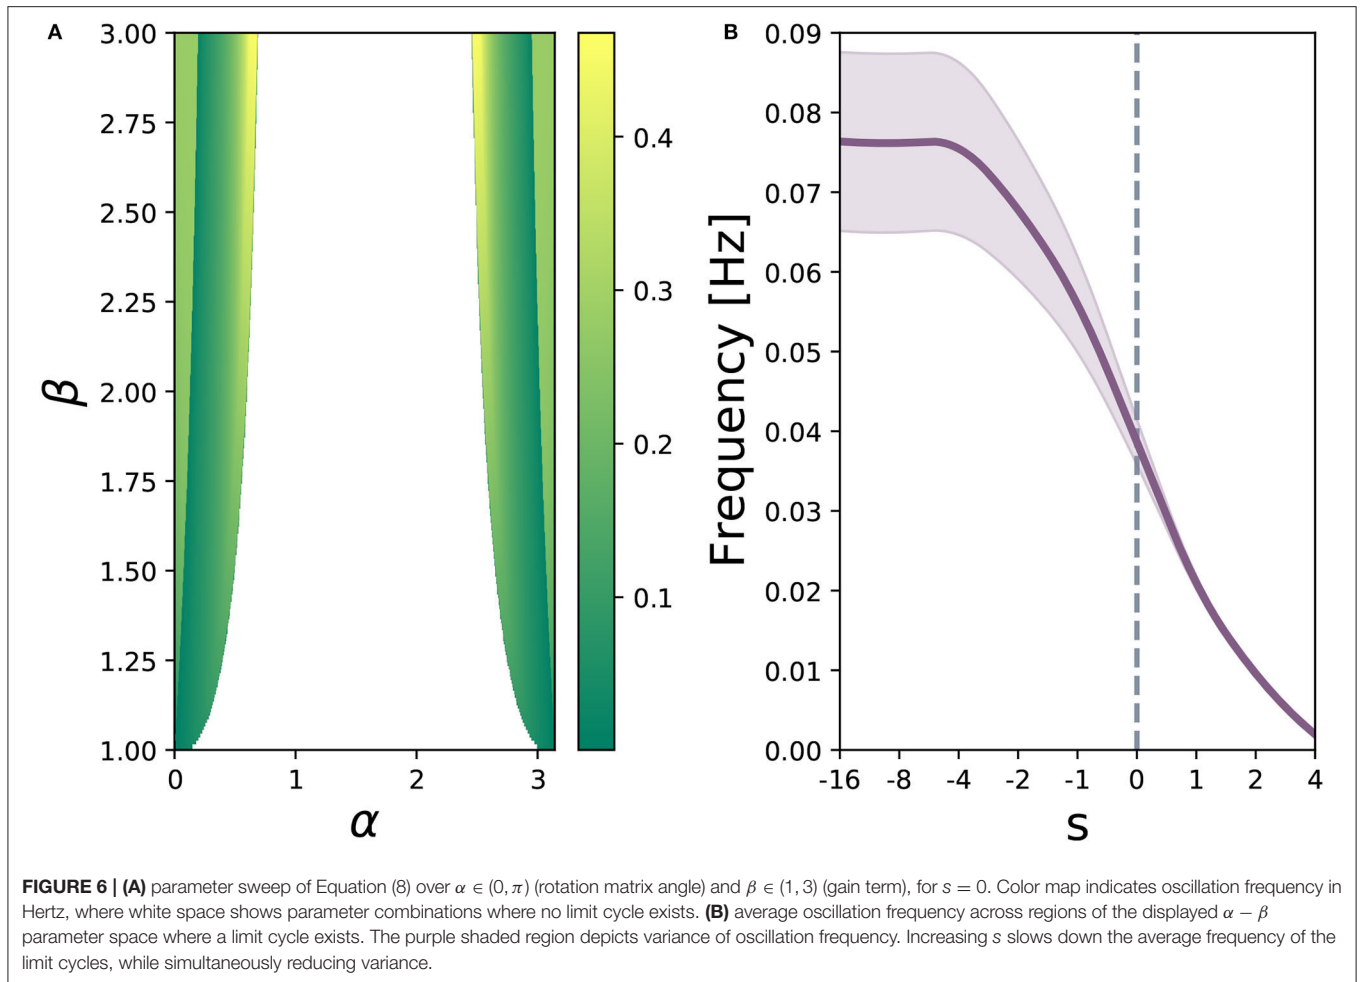
In addition to proper homoclinic orbits, we observe that 2D GRUs can exhibit one or two bounded planar regions of homoclinic-like orbits for a given set of parameters, as shown in Figures 9A,B, respectively. Any trajectory initialized in one of these regions will flow into the pseudo-codimension-2 bifurcation fixed point at the origin, regardless of which direction time flows in. Since the pseudo-codimension-2 bifurcation fixed point is technically a cluster of four fixed points, including one source and one sink, as demonstrated in Figure 4, there is actually no homoclinic loop. However, due to the close proximity of

these fixed points, trajectories repelled away from the source, but within the basin of attraction of the sink, will appear homoclinic due to the use of finite precision. This featured behavior enables the accurate depiction of various models, including neuron spiking (Izhikevich, 2007).

With finite-fixed point topologies and global structures out of the way, the next logical question to ask is *can 2D GRUs exhibit an infinite number of fixed points?* Such behavior is often desirable in models that require stationary attraction to non-point structures, such as line attractors and ring attractors. Computationally, movement along a line attractor may be interpreted as integration (Mante et al., 2013), and has been shown as a crucial population level mechanism in various tasks, including sentiment analysis (Maheswaranathan et al., 2019b) and decision making (Mante et al., 2013). In a similar light, movement around a ring attractor may computationally represent either modular integration or arithmetic. One known application of ring attractor dynamics in neuroscience is a representation of heading direction (Kim et al., 2017). While such behavior in the continuous GRU system has yet to be seen, an approximation of a line attractor can be made, as depicted in Figure 10. We will refer to this phenomenon as a *pseudo-line attractor*, where the nullclines remain sufficiently close on a small finite interval, thereby allowing for arbitrarily slow flow, by means of slow points.

5. EXPERIMENTS: TIME-SERIES PREDICTION

As a means to put our theory to practice, in this section we explore several examples of time series prediction of continuous



time planar dynamical systems using 2D GRUs. Results from the previous section indicate what dynamical features can be learned by this RNN, and suggest cases by which training will fail. All of the following computer experiments consist of an RNN, by which the hidden layer is made up of a 2D GRU, followed by a linear output layer. The network is trained to make a 29-step prediction from a given initial observation, and no further input through prediction. As such, to produce accurate predictions, the RNN must rely solely on the hidden layer dynamics.

We train the network to minimize the following multi-step loss function:

$$\mathcal{L}(\theta) = \frac{1}{T} \sum_{i=1}^{N_{\text{traj}}} \sum_{k=1}^T \|\hat{\mathbf{w}}_i(k; \mathbf{w}_i(0)) - \mathbf{w}_i(k)\|_2^2 \quad (10)$$

where θ are the parameters of the GRU and linear readout, $T = 29$ is the prediction horizon, $\mathbf{w}_i(t)$ is the i -th time series generated by the true system, and $\hat{\mathbf{w}}(k; \mathbf{w}_0)$ is the k -step prediction given \mathbf{w}_0 .

The hidden states are initialized at zero for each trajectory. The RNN is then trained for 4000 epochs, using ADAM (Kingma and Ba, 2014) in whole batch mode to minimize the loss function, i.e.,

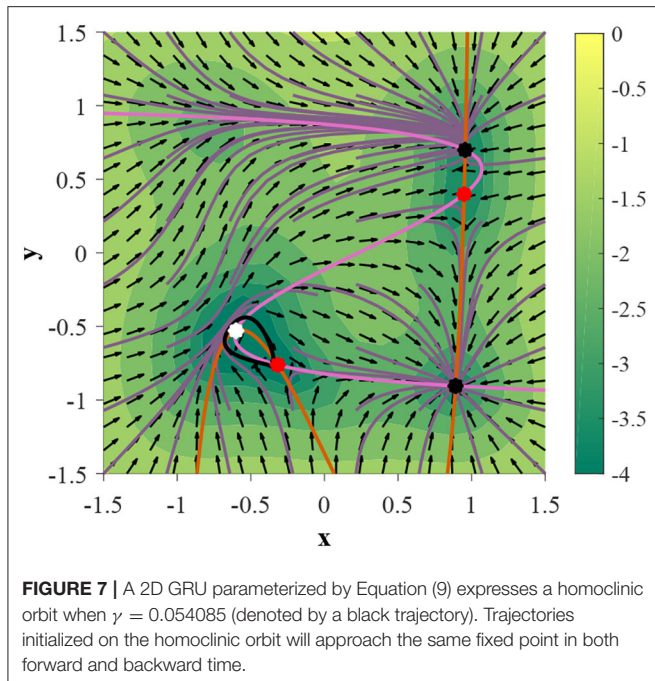
the mean square error between the predicted trajectory and the data. $N_{\text{traj}} = 667$ time series were used for training. **Figure 11** depicts the experimental results of the RNN's attempt at learning each dynamical system we describe below.

5.1. Limit Cycle

To test if 2D GRUs can learn a limit cycle, we use a simple nonlinear oscillator called the FitzHugh-Nagumo Model (FitzHugh, 1961). The FitzHugh-Nagumo model is defined by: $\dot{x} = x - \frac{x^3}{3} - y + I_{\text{ext}}$, $\tau \dot{y} = x + a - by$, where in this experiment we will chose $\tau = 12.5$, $a = 0.7$, $b = 0.8$, and $I_{\text{ext}} = \mathcal{N}(0.7, 0.04)$. Under this choice of model parameters, the system will exhibit an unstable fixed point (unstable spiral) surrounded by a limit cycle (**Figure 11**). As shown in section 4, 2D GRUs are capable of representing this topology. The results of this experiment verify this claim (**Figure 11**), as 2D GRUs can capture topologically equivalent dynamics.

5.2. Line Attractor

As discussed in section 4, 2D GRUs can exhibit a pseudo-line attractor, by which the system mimics an analytic line attractor on a small finite domain. We will use the simplest representation



of a planar line attractor: $\dot{x} = -x$, $\dot{y} = 0$. This system will exhibit a line attractor along the y -axis, at $x = 0$ (**Figure 11**). Trajectories will flow directly perpendicular toward the attractor. While the hidden state dynamics of the trained network do not perfectly match that of an analytic line attractor, there exists a small subinterval near each of the fixed points acting as a pseudo-line attractor (**Figure 11**). As such, the added affine transformation (linear readout) can scale and reorient this subinterval on a finite domain. Since all attractors in a d -dimensional GRU are bound to $[-1, 1]^d$, no line attractor can extend infinitely in any given direction, which matches well with the GRUs inability to perform unbounded counting, as the continuous analog of such a task would require a trajectory to move along such an attractor.

5.3. Ring Attractor

For this experiment, a dynamical system representing a standard ring attractor of radius one is used: $\dot{x} = -(x^2 + y^2 - 1)x$; $\dot{y} = -(x^2 + y^2 - 1)y$. This system exhibits an attracting ring, centered around an unstable fixed point. We added Gaussian noise $\mathcal{N}(0, 0.1I)$ in the training data.

In our analysis we did not observe two GRUs exhibit this set of dynamics, and the results of this experiment, demonstrated in **Figure 3** experiments, suggest they cannot. Rather, the hidden state dynamics fall into an observed finite fixed point topology (see case xxix in section 3 of the **Supplementary Material**). In addition, we robustly see this over multiple initializations, and the quality of approximation improves as the dimensionality of GRU increases (**Figure 12**), suggesting that many GRUs are required to obtain a sufficient approximation of this set of dynamics for a practical task (Funahashi and Nakamura, 1993).

6. DISCUSSION

Through example and experiment we indicated classes of dynamics which are crucial in expressing various known neural computations and obtainable with the 2D GRU network. We demonstrated the system's inability to learn continuous attractors, seemingly in any finite dimension, a structure hypothesized to exist in various neural representations. While the GRU network was not originally made as a neuroscientific model, there has been considerable work done showing high qualitative similarity between the underlying dynamics of neural recordings and artificial RNNs on the population level (Mante et al., 2013; Sussillo et al., 2015). Furthermore, recent research has modified such artificial models to simulate various neurobiological phenomenon (Heeger and Mackey, 2019). One recent study demonstrated that trained RNNs of different architectures and nonlinearities express very similar fixed point topologies to one another when successfully trained on the same tasks (Maheswaranathan et al., 2019a), suggesting a possible connection in the dynamics of artificial networks and neural population dynamics. As such, an understanding of the obtainable dynamical features in a GRU network allow one to comment on the efficacy of using such an architecture as an analog of brain dynamics at the population level.

Although this manuscript simplified the problem by considering the 2D GRU, a lot of research has resulted in interpreting cortical dynamics as low dimensional continuous time dynamical systems (Harvey et al., 2012; Mante et al., 2013; Cueva et al., 2020; MacDowell and Buschman, 2020; Zhao and Park, 2020; Flesch et al., 2021). This is not to say that most standard neuroscience inspired tasks can be solved with such a low dimensional network. However, demonstrating that common dynamical features in neuroscience can arise in low dimensions can aid in one's ability to comment on attributes of large networks. These attributes include features such as sparsity of synaptic connections. For example, spiking models exhibiting sparse connectivity have been shown to perform comparably with fully connected RNNs (Bellec et al., 2018). Additionally, pruning (i.e., removing) substantial percentages of synaptic connections in a trained RNN is known to often result in little to no drop in the network's performance on the task it was trained on (Frankle and Carbin, 2019). This suggests two more examinable properties of large networks. The first is redundancy or multiple realizations of the dynamical mechanisms needed to enact a computation existing within the same network. For example, if only one limit cycle is sufficient to accurately perform a desired task, a trained network may exhibit multiple limit cycles, each qualitatively acting identically toward the overall computation. The second is the robustness of each topological structure to synaptic perturbation/pruning. For example, if we have some dynamical structure, say a limit cycle, how much can we move around in parameter space while still maintaining the existence of that structure?

In a related light, the GRU architecture has been used within more complex machine learning setups to interpret the real-time dynamics of neural recordings (Pandarinath et al., 2018; Willett

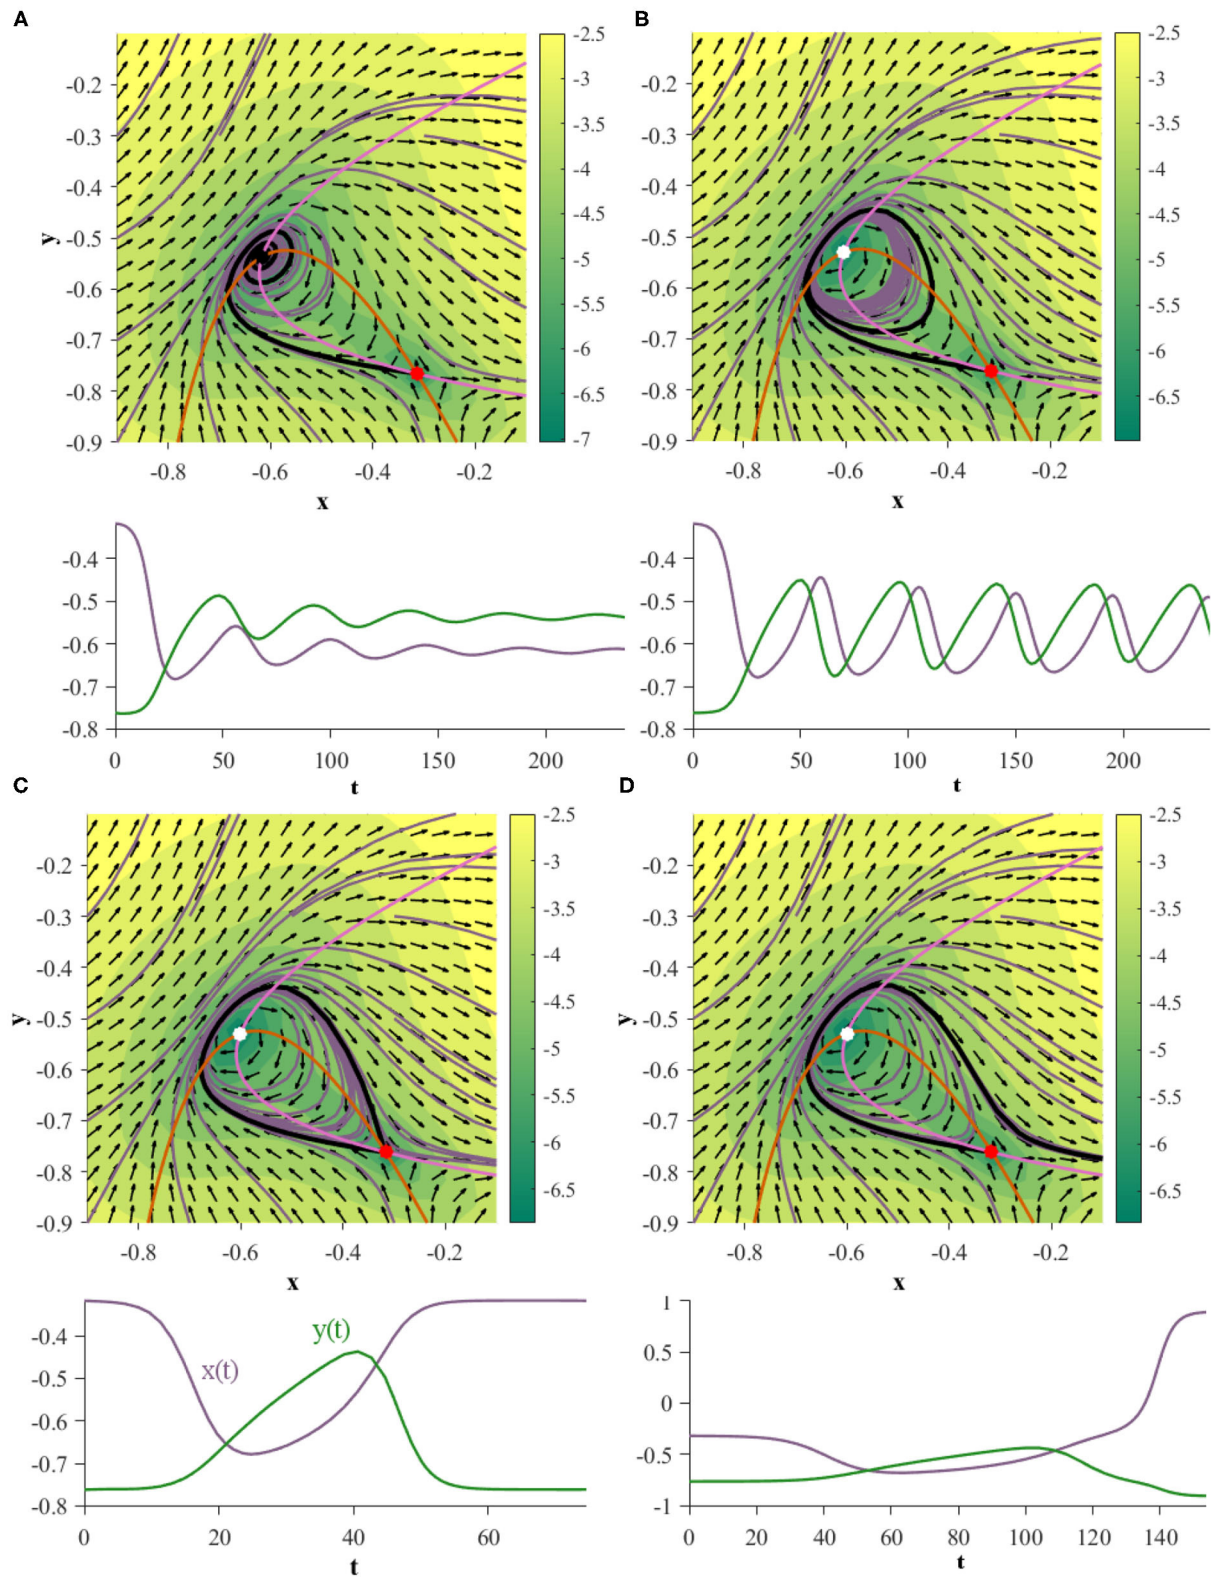
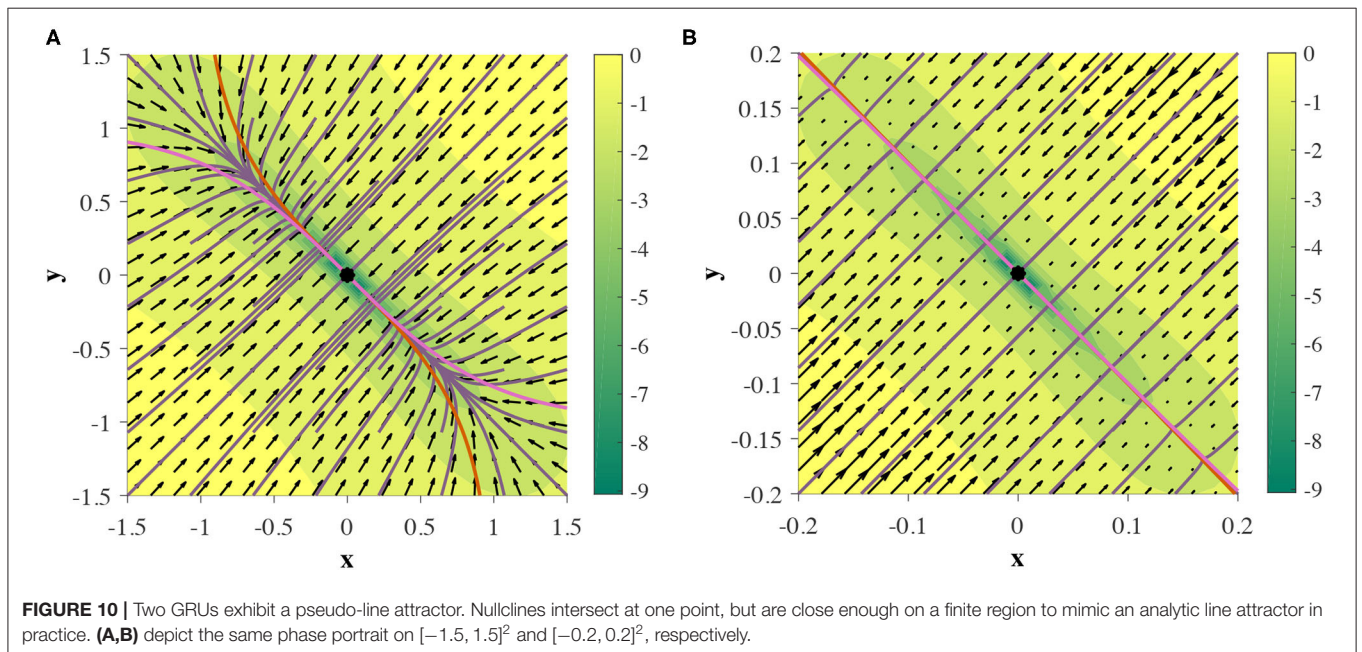
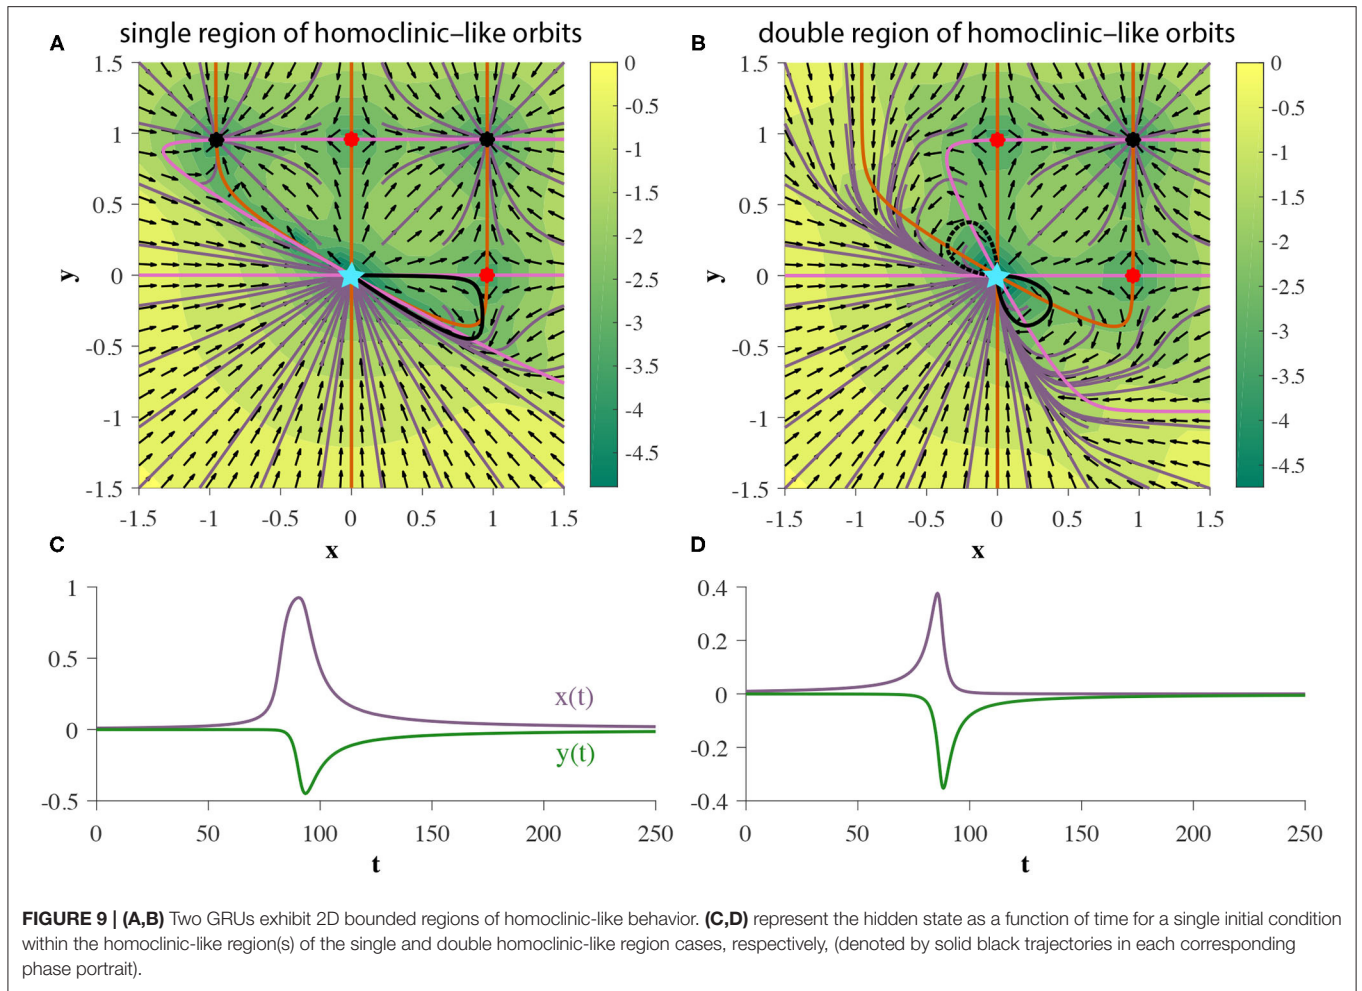
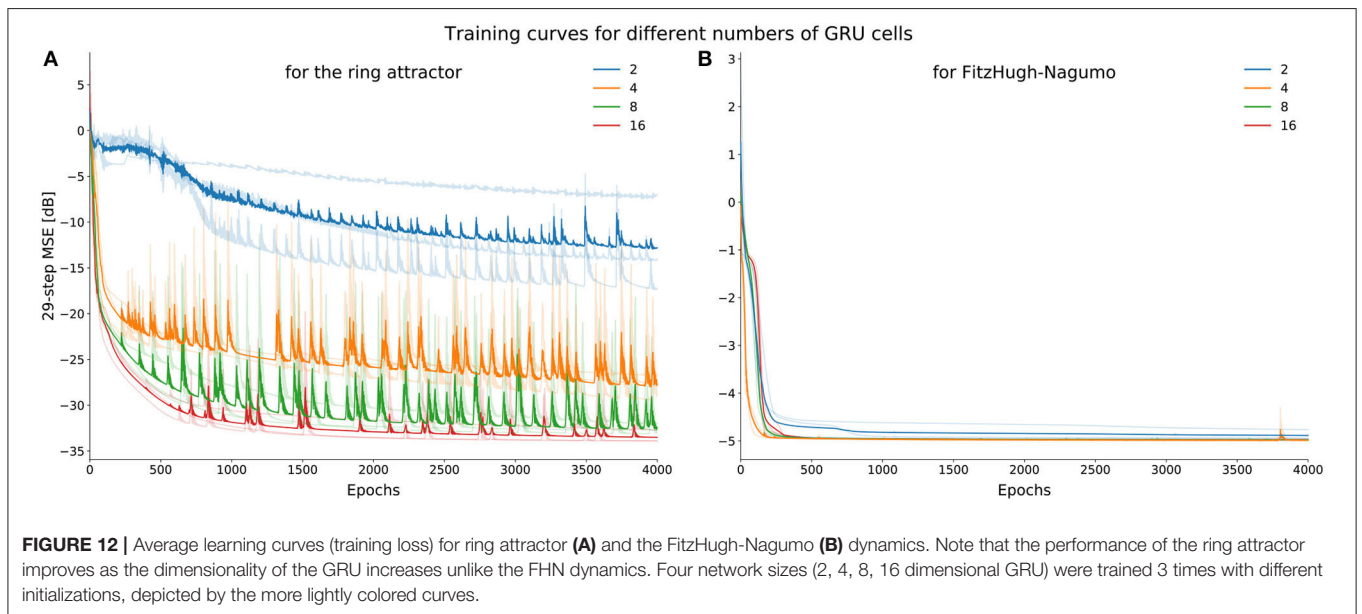
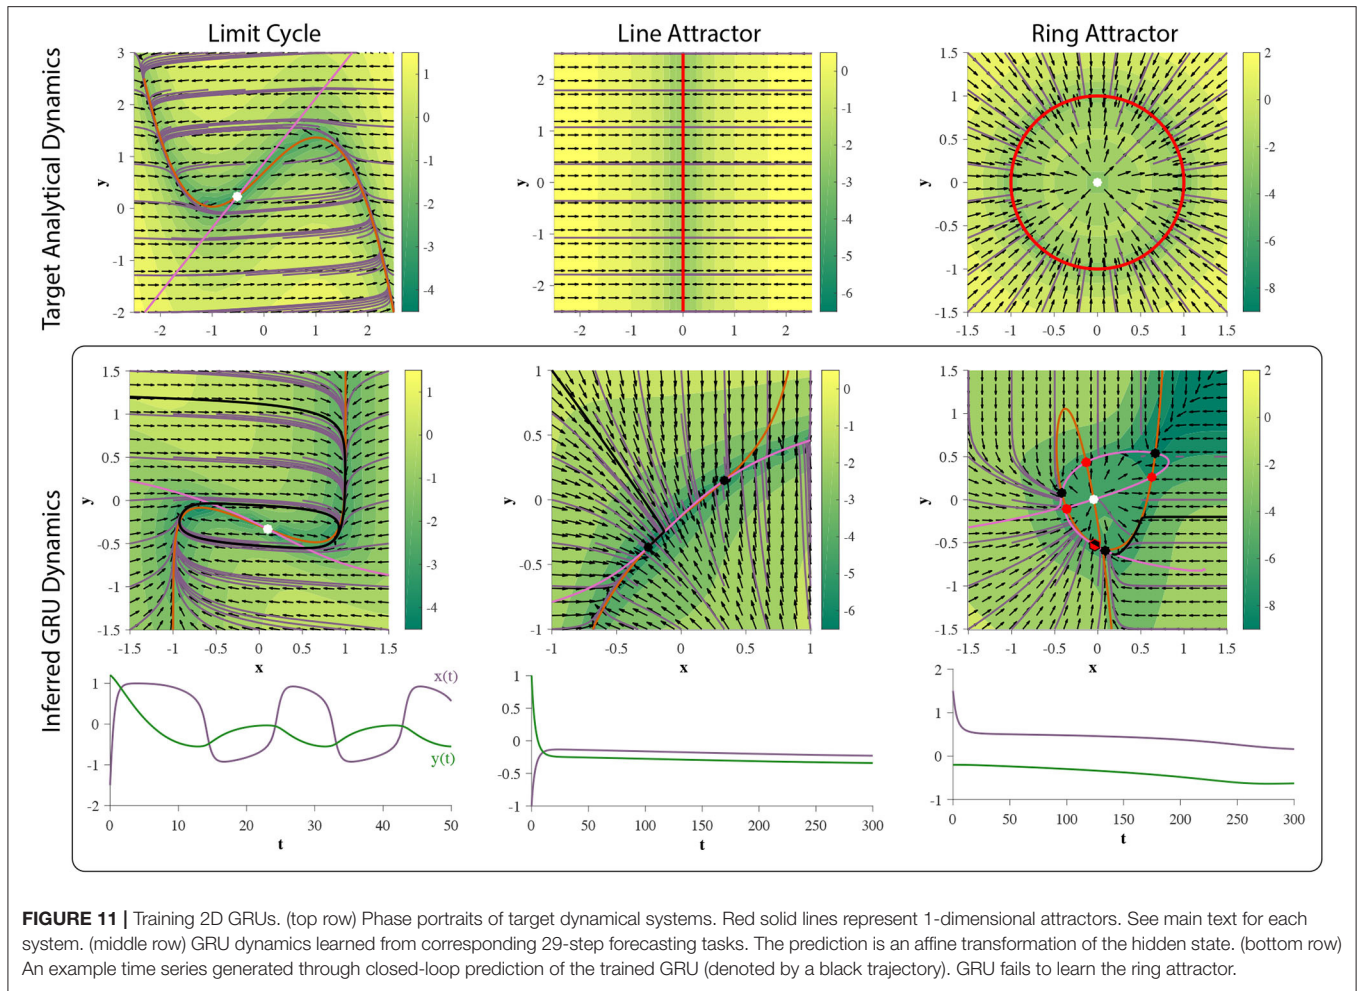


FIGURE 8 | Two GRUs exhibit an Andronov-Hopf bifurcation followed by a homoclinic bifurcation under the same parameterization. The plots directly under each phase portrait depict the time evolution of the black trajectory for the corresponding system. **(A)** ($\gamma = 0.051$): the system exhibits a stable fixed point. **(B)** ($\gamma = 0.0535$): the system has undergone an Andronov-Hopf bifurcation and exhibits a stable limit cycle. **(C)** ($\gamma = 0.054085$): the limit cycle collides with the saddle point, creating a homoclinic orbit. **(D)** ($\gamma = 0.0542$): the system has undergone a homoclinic bifurcation exhibits neither a homoclinic orbit nor a limit cycle.





et al., 2021). These tools allow researchers to better understand and study the differences between neural responses, trial to trial. Knowledge of the inner workings and expressive power of GRU networks can only further our understanding of the limitations and optimization of such setups by the same line of reasoning previously stated, thereby helping to advance this class of technologies, aiding the field of neuroscience as a whole.

The most compared RNN architecture to the GRU is LSTM, as GRU was designed as both a model and computational simplification of this preexisting design in discrete time implementation. LSTM, for a significant period of time, was arguably the most popular discrete time RNN architecture, outperforming other models of the time on many benchmark tasks. However, there is one caveat when comparing the continuous time implementations of LSTM and GRU. A one dimensional LSTM (i.e., a single LSTM unit) is a two dimensional dynamical system, as information is stored in both the system's hidden state and cell state (Hochreiter and Schmidhuber, 1997). With the choice of analysis we use to dissect the GRU in this paper, LSTM is a vastly different class of system. We would expect to see a different and more limited array of dynamics for an LSTM unit when compared with the 2D GRU. However, we wouldn't consider this a fair comparison.

One attribute of the GRU architecture we chose to disregard in this manuscript was the influence of the update gate $\mathbf{z}(t)$. As stated in section 2, every element of this gate is bound to $(0, 1)^d$. Since Equation (7) only has one term containing the update gate, $[1 - \mathbf{z}(t)]$, which can be factored out, the fixed point topology does not depend on $\mathbf{z}(t)$, as this term is always strictly positive. The role this gate plays is to adjust the point-wise speed of flow, and therefore can bring rise to slow manifolds. Because each element of $\mathbf{z}(t)$ can become arbitrarily close to the value of one, regions of phase space associated with an element of the update-gate sufficiently close to one will experience seemingly no motion in the directions associated with those elements. For example, in the 2D GRU system, if the first element of $\mathbf{z}(t)$ is sufficiently close to one, the trajectory will maintain a near fixed value in x . These slow points are not actual fixed points. Therefore, in the autonomous system, trajectories traversing them will eventually overcome this stoppage given sufficient time. However, this may add one complicating factor for analyzing implemented continuous time GRUs in practice. The use of finite precision allows for the flow speed to dip below machine precision, essentially creating *pseudo-attractors* in these regions. The areas of phase space containing these points will qualitatively behave as attracting sets, but not by traditional dynamical systems terms, making them more difficult to analyze. If needed, we recommend looking at $\mathbf{z}(t)$ separately, because this term acts independently from the remaining terms in the continuous time system. Therefore, any slow points found can be superimposed with the traditional fixed points in phase space. In order to avoid the effects of finite precision all together, the system can be realized through a hardware implementation (Jordan and Park, 2020). However, proper care needs to be given in order to mitigate analog imperfections.

Unlike the update gate, we demonstrated that the reset gate $\mathbf{r}(t)$ affects the network's fixed point topology, allowing for

more complicated classes of dynamics, including homoclinic-like orbits. These effects are best described through the shape of the nullclines. We will keep things qualitative here as to help build intuition. In 2D, if every element of the reset gate weight matrix \mathbf{U}_r and bias \mathbf{b}_r is zero, nullclines can form two shapes. First is a *sigmoid-like* shape (Figures 5A, 10, 11; inferred limit cycle and line attractor), allowing them to intersect a line (or hyperplane in higher dimensions) orthogonal to their associated dimension a single time. The second is an *s-like* shape (Figures 5B,C, 7, 11; limit cycle), allowing them to intersect a line orthogonal to their associated dimension up to three times. The peak and trough of the *s-like* shape can be stretched infinitely as well (Figure 2A). In this case, two of the three resultant seemingly disconnected nullclines associated with a given dimension can be placed arbitrarily close together (Figure 3B). Varying $\mathbf{r}(t)$ allows the geometry of the nullclines to take on several additional shapes. The first of these additional structures is a *pitchfork-like* shape (Figures 3A,C, 9). By disconnecting two of the prongs from the pitchfork we get our second structure, simultaneously exhibiting a sigmoid-like shape and a *U-like* shape (Figure 3C). Bending the ends of the "U" at infinity down into \mathbb{R}^2 connects them, forming our third structure, an *O-like* shape (Figure 3 periments; inferred ring attractor–orange nullcline). This *O-like* shape can then also intersect the additional segment of the nullcline, creating one continuous curve (Figure 3 periments; inferred ring attractor–pink nullcline). One consequence of the reset-gate is the additional capacity to encode information in the form of stable fixed points. If we neglect $\mathbf{r}(t)$, we can obtain up to four sinks (Figure 2A), as we are limited to the intersections of the nullclines; two sets of three parallel lines. Incorporating $\mathbf{r}(t)$ increases the number of fixed points obtainable (Figure 3A). Refer to section 3 of the **Supplementary Material** to see how these nullcline structures lead to a vast array of different fixed point topologies.

Several interesting extensions to this work immediately come to mind. For one, the extension to a 3D continuous time GRU network opens up the door for the possibility of more complex dynamical features. Three spatial dimensions are the minimum required to experience chaotic dynamics in nonlinear systems (Meiss, 2007), and due to the vast size of the GRU parameter space, even in low dimensions, such behavior is probable. Similarly, additional types of bifurcations may be present, including bifurcations of limit cycles, allowing for more complex oscillatory behavior (Kuznetsov, 1998). Furthermore, higher dimensional GRUs may bring rise to complex center manifolds, requiring center manifold reduction to better analyze and interpret the phase space dynamics (Carr, 1981). While we considered the underlying GRU topology separate from training, considering how the attractor structure influences learning can bring insight into successfully implementing RNN models (Sokół et al., 2019). As of yet, this topic of research is mostly uncharted. We believe such findings, along with the work presented in this manuscript, will unlock new avenues of research on the trainability of recurrent neural networks and help to further understand their mathematical parallels with biological neural networks.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/**Supplementary Material**, further inquiries can be directed to the corresponding author/s.

AUTHOR CONTRIBUTIONS

IJ performed the analysis. IJ, PS, and IP wrote the manuscript. PS performed the numerical experiments. IP conceived the idea, advised, and edited the manuscript. All authors have read and approved the final manuscript.

FUNDING

This work was supported by NIH EB-026946, and NSF IIS-1845836. IJ was supported partially by the Institute of Advanced

Computational Science Jr. Researcher Fellowship, at Stony Brook University.

ACKNOWLEDGMENTS

We thank Josue Nassar, Brian O'Donnell, David Sussillo, Aminur Rahman, Denis Blackmore, Braden Brinkman, Yuan Zhao, and D.S for helpful feedback and conversations regarding the analysis and writing of this manuscript.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fncom.2021.678158/full#supplementary-material>

REFERENCES

- Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adapt. Behav.* 3, 469–509. doi: 10.1177/105971239500300405
- Beer, R. D. (2006). Parameter space structure of continuous-time recurrent neural networks. *Neural Comput.* 18, 3009–3051. doi: 10.1162/neco.2006.18.12.3009
- Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., and Maass, W. (2018). Long short-term memory and learning-to-learn in networks of spiking neurons. *arXiv:1803.09574 [cs, q-bio]*. *arXiv: 1803.09574*.
- Bengio, Y., Simard, P., and Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Netw.* 5, 157–166. doi: 10.1109/72.279181
- Carr, J. (1981). *Applications of Centre Manifold Theory, 1982nd Edn*. New York, NY; Heidelberg, Berlin: Springer.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). “Neural ordinary differential equations,” in *Advances in Neural Information Processing Systems, Vol. 31*, eds S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Montreal, QC: Curran Associates, Inc.).
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., et al. (2014). learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv:1406.1078 [cs, stat]*. *arXiv: 1406.1078*. doi: 10.3115/v1/D14-1179
- Choi, K., Fazekas, G., Sandler, M., and Cho, K. (2017). “Convolutional recurrent neural networks for music classification,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (New Orleans, LA: IEEE), 2392–2396.
- Churchland, M. M., and Cunningham, J. P. (2014). A dynamical basis set for generating reaches. *Cold Spring Harb. Symp. Quant. Biol.* 79, 67–80. doi: 10.1101/sqb.2014.79.024703
- Costa, R., Assael, I. A., Shillingford, B., de Freitas, N., and Vogels, T. (2017). “Cortical microcircuits as gated-recurrent neural networks,” in *Advances in Neural Information Processing Systems 30*, eds I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. and Garnett (Long Beach, CA: Curran Associates, Inc.), 272–283.
- Cueva, C. J., Saez, A., Marcos, E., Genovesio, A., Jazayeri, M., Romo, R., et al. (2020). Low-dimensional dynamics for working memory and time encoding. *Proc. Natl. Acad. Sci. U.S.A.* 117, 23021–23032. doi: 10.1073/pnas.1915984117
- Doya, K. (1993). Bifurcations of recurrent neural networks in gradient descent learning. *IEEE Trans. Neural Netw.* 1, 75–80.
- Dwivedi, D., Sermanet, P., and Tompson, J. (2018). “Temporal reasoning in videos using convolutional gated recurrent units,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (Salt Lake City, UT).
- FitzHugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophys. J.* 1, 445–466. doi: 10.1016/S0006-3495(61)86902-6
- Flesch, T., Juechems, K., Dumbalska, T., Saxe, A., and Summerfield, C. (2021). *Rich and Lazy Learning of Task Representations in Brains and Neural Networks. bioRxiv, 2021.04.23.441128*. Cold Spring Harbor Laboratory Section: New Results.
- Frankle, J., and Carbin, M. (2019). “The lottery ticket hypothesis: finding sparse, trainable neural networks,” in *International Conference on Learning Representations* (New Orleans, LA). Available online at: <https://openreview.net/forum?id=rjl-b3RcF7>
- Funahashi, K. -I., and Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Netw.* 6, 801–806. doi: 10.1016/S0893-6080(05)80125-X
- Harvey, C. D., Coen, P., and Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62–68. doi: 10.1038/nature10918
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Las Vegas, NV: IEEE), 770–778.
- Heath, M. T. (2018). “Scientific computing: an introductory survey, revised second edition,” in *SIAM-Society for Industrial and Applied Mathematics, Philadelphia, 2nd Edn* (New York, NY).
- Heeger, D. J., and Mackey, W. E. (2019). Oscillatory recurrent gated neural integrator circuits (ORGaNICs), a unifying theoretical framework for neural dynamics. *Proc. Natl. Acad. Sci.* 116, 22783–22794. doi: 10.1073/pnas.1911633116
- Hochreiter, S. (1991). *Untersuchungen zu Dynamischen Neuronalen Netzen* (Ph.D. thesis), TU Munich. Advisor J. Schmidhuber.
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780. doi: 10.1162/neco.1997.9.8.1735
- Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117, 500–544. doi: 10.1113/jphysiol.1952.sp004764
- Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience*. Cambridge, MA: MIT Press.
- Jordan, I. D., and Park, I. M. (2020). Birhythmic analog circuit maze: A nonlinear neurostimulation testbed. *Entropy* 22:537. doi: 10.3390/e22050537
- Kim, S. S., Rouault, H., Druckmann, S., and Jayaraman, V. (2017). Ring attractor dynamics in the *Drosophila* central brain. *Science* 356, 849–853. doi: 10.1126/science.aal4835
- Kingma, D. P., and Ba, J. (2014). Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*. *arXiv: 1412.6980*.
- Kuznetsov, Y. A. (1998). *Elements of Applied Bifurcation Theory 2nd Edn*. Berlin; Heidelberg: Springer-Verlag.
- Laurent, T., and von Brecht, J. (2017). “A recurrent neural network without chaos,” in *5th International Conference on Learning Representations, ICLR 2017* (Toulon).

- LeVeque, R. J., and LeVeque, R. (1992). *Numerical Methods for Conservation Laws, 2nd Edn.* Basel; Boston, MA: Birkhäuser.
- MacDowell, C. J., and Buschman, T. J. (2020). Low-dimensional spatiotemporal dynamics underlie cortex-wide neural activity. *Curr. Biol.* 30, 2665.e8–2680.e8. doi: 10.1016/j.cub.2020.04.090
- Maheswaranathan, N., Williams, A., Golub, M., Ganguli, S., and Sussillo, D. (2019a). “Universality and individuality in neural dynamics across large populations of recurrent networks,” in *Advances in Neural Information Processing Systems, Vol. 32*, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox and R. Garnett (Vancouver, BC: Curran Associates, Inc.).
- Maheswaranathan, N., Williams, A., Golub, M. D., Ganguli, S., and Sussillo, D. (2019b). Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics. *arXiv:1906.10720 [cs, stat]*. *arXiv: 1906.10720*.
- Mante, V., Sussillo, D., Shenoy, K., and Newsome, W. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. doi: 10.1038/nature12742
- Meiss, J. (2007). *Differential Dynamical Systems. Mathematical Modeling and Computation. Society for Industrial and Applied Mathematics.* Boulder, CO: University of Colorado.
- Morrill, J., Salvi, C., Kidger, P., Foster, J., and Lyons, T. (2021). Neural rough differential equations for long time series. *arXiv:2009.08295 [cs, math, stat]*. *arXiv: 2009.08295*.
- Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* 15, 805–815. doi: 10.1038/s41592-018-0109-9
- Pasemann, F. (1997). A simple chaotic neuron. *Phys. D Nonlinear Phenomena* 104, 205–211. doi: 10.1016/S0167-2789(96)00239-4
- Prabhavalkar, R., Rao, K., Sainath, T. N., Li, B., Johnson, L., and Jaitly, N. (2017). “A comparison of sequence-to-sequence models for speech recognition,” in *Interspeech 2017 (ISCA)* (Stockholm), 939–943.
- Sokół, P. A., Jordan, I., Kadile, E., and Park, I. M. (2019). “Adjoint dynamics of stable limit cycle neural networks,” in *2019 53rd Asilomar Conference on Signals, Systems, and Computers* (Pacific Grove, CA), 884–887.
- Sussillo, D., and Barak, O. (2012). Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Comput.* 25, 626–649. doi: 10.1162/NECO_a_00409
- Sussillo, D., Churchland, M., Kaufman, M. T., and Shenoy, K. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* 18:1025–1033. doi: 10.1038/nn.4042
- Thomas, J. W. (1995). *Numerical Partial Differential Equations: Finite Difference Methods, 1st Edn.* New York, NY: Springer.
- Weiss, G., Goldberg, Y., and Yahav, E. (2018). On the practical computational power of finite precision RNNs for language recognition. *arXiv:1805.04908 [cs, stat]*. *arXiv: 1805.04908*. doi: 10.18653/v1/P18-2117
- Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., and Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature* 593, 249–254. doi: 10.1038/s41586-021-03506-2
- Wong, K.-F., and Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328. doi: 10.1523/JNEUROSCI.3733-05.2006
- Zhao, Y., and Park, I. M. (2016). “Interpretable nonlinear dynamic modeling of neural trajectories,” in *Advances in Neural Information Processing Systems (NIPS)* (Barcelona).
- Zhao, Y., and Park, I. M. (2020). Variational online learning of neural dynamics. *Front. Comput. Neurosci.* 14:71. doi: 10.3389/fncom.2020.00071

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Jordan, Sokół and Park. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.