Effect of dilution in asymmetric recurrent neural networks

Viola Folli [a], Giorgio Gosti [a,*], Marco Leonetti [a,c], Giancarlo Ruocco [a,b]

[a] Center for Life Nanoscience, Istituto Italiano di Tecnologia, Viale Regina Elena 291, 00161, Rome, Italy
[b] Department of Physics, Sapienza University of Rome, Piazzale Aldo Moro 5, 00185, Rome, Italy
[c] CNR NANOTEC-Institute of Nanotechnology c/o Campus Ecotekne, University of Salento, Via Monteroni, 73100 Lecce, Italy

HIGHLIGHTS

- The dilution and symmetry of a recurrent neural network affect its limit behaviors.
- There are two optimal regions that optimize the number of limit behaviors.
- The first region is symmetric and fully connected as predicted by Hebbs' learning.
- The second region is asymmetric and diluted as found in the neocortex and hippocampus.

ARTICLE INFO

ABSTRACT

We study with numerical simulation the possible limit behaviors of synchronous discrete-time deterministic recurrent neural networks composed of N binary neurons as a function of a network's level of dilution and asymmetry. The network dilution measures the fraction of neuron couples that are connected, and the network asymmetry measures to what extent the underlying connectivity matrix is asymmetric. For each given neural network, we study the dynamical evolution of all the different initial conditions, thus characterizing the full dynamical landscape without imposing any learning rule. Because of the deterministic dynamics, each trajectory converges to an attractor, that can be either a fixed point or a limit cycle. These attractors form the set of all the possible limit behaviors of the neural network. For each network we then determine the convergence times, the limit cycles' length, the number of attractors, and the sizes of the attractors' basin. We show that there are two network structures that maximize the number of possible limit behaviors. The first optimal network structure is fully-connected and symmetric. On the contrary, the second optimal network structure is highly sparse and asymmetric. The latter optimal is similar to what observed in different biological neuronal circuits. These observations lead us to hypothesize that independently from any given learning model, an efficient and effective biologic network that stores a number of limit behaviors close to its maximum capacity tends to develop a connectivity structure similar to one of the optimal networks we found.

1. Introduction

Recurrent neural networks are able to store stimuli-response associations, and serve as a model of how live neural networks store and recall behaviors as responses to given stimuli. A discrete-time deterministic recurrent N binary-neuron neural network is completely characterized by its N^2 edges, and its instantaneous state is defined by a neuron activation vector σ, which is a binary vector of size N. In this paper, we consider a specific kind of recurrent neural network, which is initialized, analogously to a Hopfield network, by assigning to the network's neurons an initial pattern which is the network stimulus or input. The collection of all possible neuron activation vectors contains 2^N allowed vectors σ, these vectors can be partitioned in three categories: steady states, limit cycles, and transient states. Steady states are neuron activation states that do not change in time, and limit cycles are sequences of neuron activation vectors that repeat cyclically, with a period that we call cycle length. From now on, we will consider a steady state as a limit cycle of length 1. A network, given any initial activation vector, always evolves to a limit cycle, which for this reason we also refer to as attractor. In other words, a network associates a limit cycle to any initial neural activation state that is given as an input. For this reason, limit cycles can be considered as behaviors stored as responses to initial stimuli. In the case of length 1 cycles, limit behaviors are a single activation

* Corresponding author.
  E-mail addresses: viola.folli@iit.it (V. Folli), giorgio.gosti@iit.it (G. Gosti), marco.leonetti@roma1.infn.it (M. Leonetti), Giancarlo.Ruocco@roma1.infn.it (G. Ruocco).

state which in the case of Hopfield networks correspond to the recollection of a memory. In the case of cycles with a length greater than 1, stored limit behaviors are sequences of activation patterns which may correspond to a stored dynamical sequence, such as the performance of a complex motor task, or a dynamic sequence of static memories. In principle, a recurrent neural network stores a certain number of limit behaviors as vectors from a 2^N set in a data structure defined by N^2 parameters. Furthermore, these vectors can be recovered in responses to input stimuli. This clearly has intriguing analogies with content-addressable memory systems capable of indexing large strings of bits (Carpenter, 1989; Hopfield, 1987).

In the past, recurrent neural network, and specifically Hopfield neural networks have been used to model memory storage and recall, though more recently neurobiology models implemented recurrent neural networks to describe brain activity in different cognitive tasks. Mante, Sussillo, Shenoy, and Newsome (2013) use recurrent neural networks to model the integration of context information in the prefrontal cortex in discrimination tasks. Similarly, Carnevale, deLafuente, Romo, Barak, and Parga (2015) model with recursive neural networks the premotor cortex modulation of its response criteria in a detection task with temporal uncertainty. Furthermore recurrent neural networks are used to model phoneme acquisition (Kanda, Ogata, Takahashi, Komatani, & Okuno, 2009), and language acquisition (Heinrich & Wermter, 2018). Neuroscience proposes two fundamental conceptual frameworks that enable recurrent neural networks to store limit behaviors: the connectionist hypotheses and the innate hypotheses. Hebb (1949) proposed the connectionist hypothesis, which assumes that a neural network starts blank and forms new links or adjusts the existing ones each time it stores a new limit behaviors. In this framework limit behaviors are stable equilibria in the neural network dynamics. A criticism of Hebbs' networks is that as new limit behaviors are added the corresponding generated stable states start interfering with the stable states associated with older limit behaviors. This limits the maximum storage capacity C, which is defined as the maximum number of limit behaviors that can be stored. Notice that this definition is different from the usual definition used in associative memory networks, in which the storage capacity is defined as the number of uniformly distributed random vectors that can be stored in an associative memory (Hassoun, 1993; Hassoun & Watta, 1997). Amit, Gutfreund, and Sompolinsky (1985a) show that a Hebbian network has a storage capacity of $C = pN$ with $p \approx 0.14$. In contrast, innate network models assume that limit behaviors are stored using innate neural assemblies with a given connectivity. Among other innate memory models, Perin, Berger, and Markram (2011) propose that groups of pyramidal neurons in the rats' neocortex may be innate neuron assemblies that may only partially change their overall connectivity structure. Indeed, Perin et al. (2011) find that these assemblies have similar connectivity proprieties among different animals, and argue that these assemblies serve as building blocks for the formation of composite complex memories.

Whether we assume a connectionist or an innate network scheme as our working framework, we implicitly assume that a recurrent neural network acts as a content-addressing memory which given an input pattern (stimulus) returns a limit behavior. This limit behavior can be a recovered memory or a more complex neural sequence of neural activations that may be integrated into a second neural network. To understand how well a recurrent neural network acts as a content-addressing memory, the memory storage and retrieval literature uses discrete-time recurrent neural networks with McCulloch–Pitts neurons (McCulloch & Pitts, 1943). Each discrete-time recurrent neural network, which in this literature is sometimes referred to as Hopfield neural network, is characterized by its connectivity matrix **J**, which schematically represents

the set of synapses and electrical junctions connecting couples of neurons. Deterministic discrete-time synchronous recursive neural networks are deterministic discrete dynamical systems. This implies three properties. First, each state in the neural network uniquely transits to another one. Second, the reverse is not true, different states can evolve to the same state. Third, each state belongs to a path that connects it to a stable activity pattern, *i.e.* a limit cycle. Given any initial neural state, or input, a discrete-time recurrent neural network dynamically falls into an attractor. In this framework, the attractor is the retrieved limit behavior (Amit, Gutfreund, & Sompolinsky, 1985b; Bastolla & Parisi, 1998; Folli, Leonetti, & Ruocco, 2017; Gutfreund, Reger, & Young, 1988; Hebb, 1949; McEliece, Posner, Rodemich, & Venkatesh, 1987; Sompolinsky, Crisanti, & Sommers, 1988; Wainrib & Touboul, 2013). Finally, it is important to consider that a recurrent network associates a limit behavior to each input from the set of all possible N-bit inputs, since the number of limit behaviors C is such that $C \ll 2^N$, it performs a many-to-few mapping. Recurrent neural network, and in particular the Hopfield model (Hopfield, 1987), show how information can be stored via attractor states. Indeed, there is some experimental support for discrete attractors in the rodents hippocampus cells' activity (Pfeiffer & Foster, 2015), and in monkey cells' activity during tasks (Fuster & Alexander, 1971; Miyashita, 1988).

To understand how well and how many limit behaviors a fully developed neural network can store, we explore how the structure properties of an arbitrary connectivity matrix **J** influences the attractor states of the network without imposing an a priori learning rules. Given a connectivity matrix **J**, to characterize the network structure, we define the network's asymmetry degree ϵ, and dilution degree ρ. The most understood properties on fixed discrete-time recurrent neural networks regard fully-connected Hopfield neural networks. Fully-connected recurrent Hopfield neural networks are networks with dilution degree $\rho = 0$, in which any couple of neurons is connected by two axons one in each direction. In contrast, we define diluted recurrent Hopfield neural networks as networks with $\rho > 0$, in which only a subset of all neurons couples are connected. The existing recurrent Hopfield neural network literature mostly discusses symmetric neural networks in which the weights of the two axons connecting neurons i and j in both directions are the same, and in only few cases researchers investigate asymmetric neural networks $\epsilon > 0$, in which the weights are no longer equal. Furthermore, most of the recurrent Hopfield neural network literature which studies the effect of asymmetry assumes binary neurons with activations state that can take values -1 and $+1$. Under these constraints, it is reported that symmetric fully-connected networks, $\epsilon = 0$, have several attractors, all of which are formed by cycles of length 1 and 2. As the network becomes less symmetric, $\epsilon > 0$, the attractors are composed of longer neural activation patterns. Increasing asymmetry in a fully connected neural network introduces severe drawbacks. Indeed, when the degree of asymmetry is increased above a certain threshold a neural network is subject to a transition from an ordered phase to a "chaotic" regime (Bastolla & Parisi, 1998; Gutfreund et al., 1988). In the chaotic regime, almost identical initial patterns can reach different attractors, and the network is characterized by a high sensitivity to initial conditions. Moreover, this chaotic regime causes exponentially longer recognition time, where the recognition time is the average number of discrete transitions required to reach the corresponding attractor from a generic point in its basin of attraction. Toyoizumi and Huang (2015) analyze asymmetric matrices $\epsilon = 0$ with neuron activation profile $\{-1, 1\}$, and show that under these conditions as the limit cycle length scales exponentially with N, the number of attractor scales linearly with N. It is important to point out that Bastolla and Parisi (1998), Gutfreund et al. (1988) and Toyoizumi and Huang (2015)