What Makes a Brain Smart? Reservoir Computing as an Approach for General Intelligence

Janelle Szary, Bryan Kerster, and Christopher Kello

University of California, Merced, Cognitive and Information Sciences Program,
5200 North Lake Road, Merced, CA 95343, USA
{jszary,bkerster,ckello}@ucmerced.edu

Abstract. Recurrent connectivity, balanced between excitation and inhibition, is a general principle of cortical connectivity. We propose that balanced recurrence can be achieved by tuning networks near their *critical branching* (CB) points when spike propagation is formalized as a branching process. We consider critical branching networks as foundations for artificial general intelligence when they are analyzed as *reservoir computing* models. Our reservoir models are based on principles of metastability and criticality that were developed in statistical mechanics in order to account for long-range correlations in activities exhibited by many types of complex systems. We discuss reservoir models and their computational properties, and we demonstrate their versatility by reviewing a number of applications.

Keywords: Reservoir computing; metastability; critical branching; neural networks

1 Introduction

Different brain areas are characterized by different neural circuitry, and some believe that different circuitry means different computations [1]. Others, however, have focused on similarities in circuitry across cortical areas [2], leading to the concept of a 'canonical cortical microcircuit' that embodies common features of neural computations [3–5]. These features are expressed at the level of thresholded spike signals (i.e. action potentials) sent through the synaptic connections of neural networks. Connections are characterized by recurrent pathways varying in spatial and temporal scales [6], and recurrent spiking activity is ongoing. Here we simulate recurrent activity as a basis for memory and computation, on the time scales of individual spikes over networks ranging widely in size. We examine memory and computational capacity and its link to evidence for so-called "avalanches" of neural activity in real neural tissue. Our model exhibits basic earmarks of general intelligence in that recurrent dynamics can support a diverse range of perceptual and cognitive functions and applications.

2 Reservoir Computing

The *reservoir computing* framework, developed by Maass [7], Jaeger [8], and colleagues, involves a network of generic units with random connections in order to produce the cortex's recurrent looping structure [9]. Spikes propagate across synapses and (sometimes) branch into further spikes, and may continue branching repeatedly and (sometimes) recurrently through the network. Future spike patterns are nonlinear functions of past spike patterns, which means that spikes carry and transform information about past inputs. With large numbers of units, this transform can be viewed as a projection into a high-dimensional space, akin to a support vector machine. If the projection sufficiently separates and organizes inputs, then a linear readout function can be used to make classifications that would otherwise not be linearly separable [7]. Fig. 1 shows an example of the general framework. The memory-less readout is the only task-dependent part of the system, whereas the reservoir can be completely task-general. In fact, multiple readout functions can simultaneously perform different computations on the same internal state. This allows the network to be flexible with respect to the functions it supports.

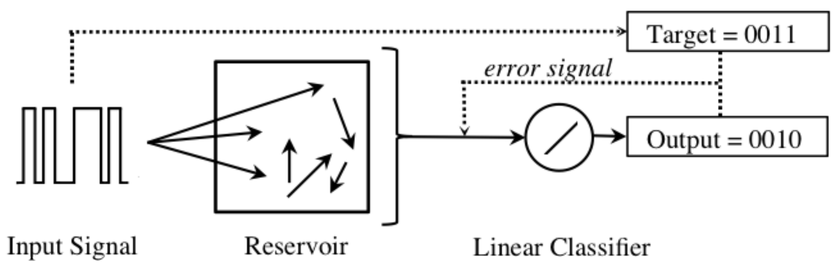

Fig. 1. A reservoir transforms a time-varying input signal onto an internal state, and a readout function (*linear classifier*) is trained to map states onto target outputs.

3 Critical Branching and Metastability

If a recurrent spiking network is overly excitable, spikes may multiply to the point of maxing out the neuron firing rates. If a network is overly inhibited, spikes will not branch and propagate. Thus, the network needs to strike a balance between these extremes at which neurons lose their information coding capabilities. In statistical mechanics, striking such a balance may poise a system near a critical state between two phases, in this case between convergent and divergent spike dynamics. When the critical state is achieved in a branching process such as a spiking neural network, so-called 'neural avalanches' are predicted to occur in the network's spontaneous, intrinsic spiking activities [10]. The distribution of avalanche sizes is predicted to follow a $-3/2$ power law [11] $P(n) \sim n^{\alpha}$, where n is the size of the avalanche (i.e. the number of units involved), $P(n)$ is the probability of observing a size-n avalanche, and the exponent α gives the slope of the relationship between $P(n)$ and n (in log-log coordinates) [12].

Beggs and Plenz showed that neural spiking activity followed an avalanche-like power law, with an estimated $-3/2$ exponent, both *in vitro* and in stochastic models of branching processes [12]. This finding is consistent with the hypothesis that branching networks are tuned to their critical states, or CB points, in order to achieve balance between convergent and divergent spike dynamics [12]. As Beggs and Plenz also argued, this balance may be adaptive because CB maximizes the transmission of information across networks under certain conditions [13]. Branching processes can be measured using the "branching ratio", which is the ratio of descendent spikes to ancestor spikes: $R = N_{desc.}/N_{anc.}$. When $R = 1$, each ancestor spike causes an average of one descendent spike, and the network is said to be at its CB point [12].

3.1 A Critical Branching Reservoir Model

To test whether CB can be usefully self-tuned in a spiking neural network, we developed a reservoir computing model that uses a self-tuning algorithm to maintain spike dynamics near their CB point. The model is based on Kello and Mayberry [11] (but see also [14, 15]), and is composed of 900 leaky integrate-and-fire (LIF) reservoir units and 100 input units. To create recurrent loops, each input and reservoir unit was randomly connected to each (other) reservoir unit with probability 0.5 (excluding self-connections)[1]. After initializing the network, input units were forced to spike randomly with some probability (e.g. 0.5) in order to spur network activity, and synapses were activated and de-activated probabilistically so that each neuron locally approaches $R = 1$. The basic idea of the algorithm is to count the number of postsynaptic spikes for each presynaptic spike on a given presynaptic neuron. Synaptic connections are activated with some probability when the neuron's postsynaptic spike count ($N_{desc.}$) is low, and de-activated with some probability when high. In particular, when $N_{desc.} < 1$, each synapse is activated with probability:

$$\eta f(s_i)|N_{desc.,i} - 1|/U. \tag{1}$$

where η is a global tuning rate parameter (fixed at 0.1), and U is the number of synapses available for activation. $f(s_i) = 1 - e^{-\lambda_i(t-t')}$ for excitatory neurons, and $f(s_i) = e^{-\lambda_i(t-t')}$ for inhibitory neurons. If $N_{desc.} > 1$ then each synapse is de-activated with the same probability as in Eq. 1, except U is the number of synapses available for de-activation, and the assignment of $f(s_i)$ is switched for excitatory versus inhibitory neurons. Over the tuning phase, this algorithm results in a network with global $R = 1$. As input spikes decrease in number, spike activity becomes more burst-like and produces neural avalanche[2]

[1] Synapses can be excitatory or inhibitory, with randomized weight values.

[2] Avalanche sizes are measured as the number of spikes occurring during a period of unusually high (threshold-exceeding) activity.

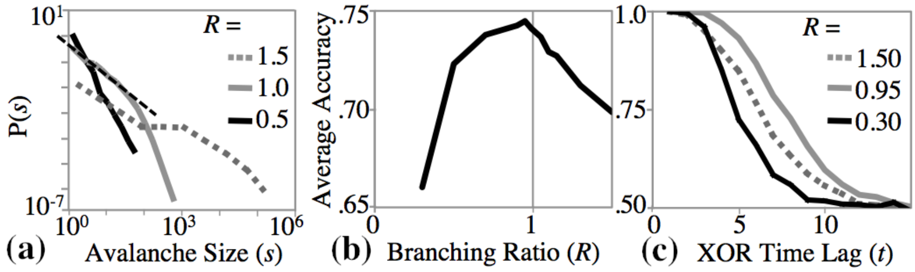

Fig. 2. **(a)** Avalanche histogram showing power laws for networks with $R = 1.5$ ($\alpha \sim -1$), $R = 1.0$ ($\alpha \sim -3/2$), and $R = 0.5$ ($\alpha \sim -3$). The dashed line shows $\alpha = -3/2$. **(b)** XOR accuracy is maximum in networks with $R \sim 1$. **(c)** XOR accuracy for past inputs (with varying time lag) is maximum in networks with $R \sim 1$.

behavior following a power law with $\alpha \sim -3/2$. This indicates that the CB algorithm produces dynamics similar to an *in vitro* cortical circuit [16]. As shown in Fig. 2(a), when the network is tuned to supercritical ($R = 1.5$) or subcritical ($R = 0.5$) levels, this result disappears, as the α no longer matches that found in neural recordings. The computational performance of the model, in terms of representational and memory capacity, was tested at various levels of R using the nonlinearly separable XOR logic task as a diagnostic[3]. Representational capacity refers to the network's ability to represent complex, nonlinearly separable patterns such that a linear classifier can extract relevant information, while memory capacity refers to the ability to maintain these representations over time. As shown in Fig. 2(b-c), computational performance is maximal near $R \sim 1$, indicating that the performance of reservoir models is enhanced when spiking dynamics approach their CB point [17][4].

4 Reservoir Computing Applications

The CB reservoir model described here has successfully been applied to visual object and motion classification tasks [18][5]. While biologically-inspired concepts such as metastability and CB might offer additional insights in achieving brain-like intelligence from reservoirs, the larger body of work has already established reservoirs as useful computational tools. They have been applied to a diverse range of tasks, from synthetic to real-world environments, and have proven to be uniquely well suited for processing data sets that are complex and time-varying. In engineering, reservoirs have been used for noise modeling to equalize wireless

[3] To express XOR, a single bit representing 0 or 1 was input to the network on each time step. Classification was performed on temporally adjacent bits into the past.

[4] Simple parameter adjustments such as increasing the number of reservoir units can increase overall performance, but here we adjusted parameters to produce results below ceiling in order to more clearly demonstrate the effects of CB.

[5] For a video demonstrating real-time object classification, see [30].

communication channels [19], online monitoring of a multi-machine power system [20], and rapid, online detection epileptic seizure onset from EEG recordings [21]. At the 2008 World Conference on Computational Intelligence, a simple reservoir model developed over only a couple of days was competitive in a target-detection competition using data from the Ford Motor Company [22].

Reservoir models have also been used to perform cognitive and perceptual functions. Our visual perception work was inspired by Maass and colleagues' work with object and motion classification [23] and Burgsteiner and colleagues' work with real-time object tracking and prediction in the RoboCup competition [24]. Reservoirs were also used for motor control in RoboCup [22], and for the control of a simulated robotic arm [25] and an artificial hand [26]. Robotics applications have included autonomous agent navigation, localization, and event detection [31, 32]. In linguistics, reservoirs have been used to classify spoken words and digits [22] and to generate grammatical structure [27], written-word sequences [28], and even musical sequences [29].

5 Conclusion

The flexibility of the reservoir computing framework is demonstrated by its successful application to a diverse range of functions. Many of these applications have involved environmentally realistic data sets, and required the mapping of complex, time-varying input signals onto stable outputs. Biological cortex, of course, must be good at this mapping, and in general it must be able to achieve a metastable balance whereby computations are somewhat stable (and in this sense robust to noise), yet can quickly adapt to reflect important changes in the input environment [33]. By focusing not only on the structure of neural circuitry but on the dynamic and metastable activity patterns produced by them, we have shown that information can be 'stored' in patterns of ongoing activity which act as the substrate for memory and computation. This capacity for generic computation in neural networks is captured by the reservoir framework, and in this way allows reservoirs to be a potential computational equivalent of the 'canonical cortical microcircuit' and a useful approach for investigating artificial general intelligence.

Acknowledgments. This work was supported by NSF, the Keck Foundation, and the DARPA SyNAPSE program (prime contract to IBM). The authors would like to thank reviewers for their helpful comments.

References

1. Cosmides, L., Tooby, J.: Origins of Domain Specificity: The Evolution of Functional Organization. In Hirschfeld, L.A., Gelman, S.A. (eds.) Mapping the Mind. New York: Cambridge University Press (1994)
2. Ebdon, M.: Is the Cerebral Neocortex a Uniform Cognitive Architecture? Mind and Language 8(3), 368–403 (1993)

3. Douglas, R.J., Martin, K.A.C.: Neuronal Circuits Of The Neocortex. Annual Review of Neuroscience, 27, 419–451 (2004)
4. Helmstaedter, M., de Kock, C.P.J., Feldmeyer, D., Bruno, R.M., Sakmann, B.: Reconstruction of an Average Cortical Column in Silico. Brain Research Reviews 55(2), 193–203 (2007)
5. Johansson, C., Lansner, A.: Towards Cortex Sized Artificial Neural Systems. Neural Networks 20(1), 48–61 (2007)
6. Buzsáki, G.: Rhythms of the Brain. Oxford University Press, Oxford (2006)
7. Maass, W., Natschläger, T., Markram, H.: Real-time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. Neural Computation 14(11), 2531–2560 (2002)
8. Jaeger, H.: The "Echo State" Approach To Analysing And Training Recurrent Neural Networks. Technical Report GMD Report 148, German National Research Center for Information Technology (2001)
9. Verstraeten, D., Schrauwen, B., D'Haene, M., Stroobandt, D.: A Unifying Comparison of Reservoir Computing Methods. Neural Networks 20, 391–403 (2007)
10. Bak, P., Tang, C., Wiesenfeld, K.: Self-organized Criticality: An explanation of the $1/f$ noise. Physics Review Letters 59, 381–384 (1987)
11. Kello, C.T., Mayberry, M.R.: Critical Branching Neural Computation. In: IEEE World Congress on Computational Intelligence, 1475–1481 (2010)
12. Beggs, J.M., Plenz, D.: Neuronal Avalanches in Neocortical Circuits. Journal of Neuroscience 23(35), 11167–11177 (2003)
13. Beggs, J.M.: The Criticality Hypothesis: How Local Cortical Networks Might Optimize Information Processing. Philosophical Transactions of the Royal Society A 366(1864) 329–343 (2008)
14. Kello, C., Kerster, B.: Power Laws, Memory Capacity, and Self-Tuned Critical Branching in an LIF Model with Binary Synapses. Computational and Systems Neuroscience 2011, Salt Lake City, UT (2011)
15. Kello, C.T., Kerster, B., Johnson, E.: Critical Branching Neural Computation, Neural Avalanches, and $1/f$ Scaling. In: Proceedings of the 33rd Annual Conference of the Cognitive Science Society (in press)
16. Beggs, J.M., Plenz, D.: Neuronal Avalanches Are Diverse And Precise Activity Patterns That Are Stable For Many Hours In Cortical Slice Cultures. Journal of Neuroscience 24(22), 5216–5229 (2004)
17. Bertschinger, N., Natschläger, T.: Real-Time Computation at the Edge of Chaos in Recurrent Neural Networks. Neural Computation 16(7), 1413–1436 (2004)
18. Szary, J.K., Kello, C.T.: Visual Motion Perception using Critical Branching Neural Computation. In: Proceedings of the 33rd Annual Conference of the Cognitive Science Society (in press)
19. Jaeger, H., Haas, H: Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Telecommunication. Science 308, 78–80 (2007)
20. Venayagamoorthy, G.: Online Design of an Echo State Network Based Wide Area Monitor for a Multimachine Power System. Neural Networks 20(3), 404–413 (2007)
21. Buteneers, P., Schrauwen, B., Verstraeten, D., Stroobandt, D.: Real-Time Epileptic Seizure Detection on Intra-Cranial Rat Data Using Reservoir Computing. In: Köppen, M., Kasabov, N., Coghill, G. (eds.) Advances in Neuro-Information Processing. LNCS, vol. 5507, pp. 56–63. Springer, Heidelberg (2009)
22. Verstraeten, D.: Reservoir Computing: Computation With Dynamical Systems, Electronics and Information Systems. Thesis: Gent, Ghent University, (2009)

23. Maass, W., Natschläger, T., Markram, H.: A New Approach Towards Vision Suggested By Biologically Realistic Neural Microcircuit Models. In: Bülthoff, H.H., Lee, S.-W., Poggio, T., Wallraven, C. (eds.) Biologically Motivated Computer Vision. LNCS, vol. 2525, pp. 1–6. Springer, Heidelberg (2002)
24. Burgsteiner, H., Kröll, M., Leopold, A., Steinbauer, G.: Movement Prediction from Real-World Images using a Liquid State Machine. Applied Intelligence 26(2), 99–109 (2007)
25. Joshi, P., Maass, W.: Movement Generation with Circuits of Spiking Neurons. Neural Computation 17(8), 1715–1738 (2005)
26. Rao, Y.N., Kim S.-P., Sanchez, J.C., Erdogmus, D., Principe, J.C., Carmena, J.M., Lebedev, M.A., Nicolelis, M.A.: Learning Mappings in Brain Machine Interfaces with Echo State Networks. In: 2005 IEEE International Conference on Acoustics, Speech, and Signal Processing, pp. 233–236, (2005)
27. Tong, M.H., Bickett, A.D., Christiansen, E.M., Cottrell, G.W.: Learning Grammatical Structure with Echo State Networks. Neural Networks 20(3), 424–432 (2007)
28. Jaeger, H.: Short Term Memory in Echo State Networks. GMD Report 152, German National Research Institute for Computer Science, (2001)
29. Eck, D.: Generating Music Sequences with an Echo State Network. In: Neural Information Processing Systems 2006 Workshop on Echo State Networks and Liquid State Machines (2006)
30. LSM Lock and Keyring Demo, `http://www.youtube.com/watch?v=QDe1Tu8kLb0`
31. Antonelo, E.A., Schrauwen, B., Stroobandt, D.: Event Detection and Localization for Small Mobile Robots Using Reservoir Computing. Neural Networks 21(6), 862871 (2008)
32. Antonelo, E., Schrauwen, B., Stroobandt, D.: Mobile Robot Control in the Road Sign Problem using Reservoir Computing Networks. In: 2008 IEEE International Conference on Robotics and Automation, pp. 911916, (2008)
33. Kelso, J.A.S.: Dynamic Patterns: The Self-Organization of the Brain and Behavior. MIT Press, Cambridge, MA (1995)