



AFRL-RI-RS-TR-2013-174

BRAIN-BASED DEVICES FOR NEUROMORPHIC COMPUTER SYSTEMS

NEUROSCIENCES RESEARCH FOUNDATION, INC.
THE NEUROSCIENCES INSTITUTE

JULY 2013

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2013-174 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/ S /

CLARE D. THIEM
Work Unit Manager

/ S /

MARK H. LINDERMAN
Technical Advisor, Computing &
Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE**Form Approved
OMB No. 0704-0188**

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) JUL 2013			2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) SEP 2011 – JAN 2013	
4. TITLE AND SUBTITLE BRAIN-BASED DEVICES FOR NEUROMORPHIC COMPUTER SYSTEMS					5a. CONTRACT NUMBER FA8750-11-2-0255	
					5b. GRANT NUMBER N/A	
					5c. PROGRAM ELEMENT NUMBER 92189D	
6. AUTHOR(S) W. Einar Gall					5d. PROJECT NUMBER SETL	
					5e. TASK NUMBER NS	
					5f. WORK UNIT NUMBER IP	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Neurosciences Research Foundation, Incorporated The Neurosciences Institute 10640 John Jay Hopkins Drive San Diego, CA 92121					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RITB 525 Brooks Road Rome NY 13441-4505					10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
					11. SPONSOR/MONITOR'S REPORT NUMBER AFRL-RI-RS-TR-2013-174	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT The ultimate goal of this project was to construct mobile devices capable of navigation, surveillance, and report. Control of these devices rests with simulated neuronal networks based on knowledge of higher brain functions that comprise cognitive control systems. In pursuing this goal, we made specific attempts to apply already demonstrated software strategies to develop a versatile set of spiking models, apply them to a recognition task, and to demonstrate a working memory. In the course of this work a new analytical method for spiking data was devised. We also designed and built a new type of mobile device capable of both flight and ground navigation while carrying a payload of surveillance sensors or other devices setting the stage for more autonomous systems.						
15. SUBJECT TERMS neuronal networks, large-scale neural models, neuromorphic computing, UAV, UGV						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			CLARE D. THIEM	
U	U	U	SAR	115	19b. TELEPHONE NUMBER (Include area code) N/A	

TABLE OF CONTENTS

Section	Page
List of Figures.....	ii
1.0 SUMMARY	1
2.0 INTRODUCTION.....	1
3.0 METHODS, ASSUMPTIONS, AND PROCEDURES	2
3.1 Construction and Testing of Versatile Winner-Take-All Networks.....	2
3.2 Temporal Sequence Learning Demonstrated in a Brain-Based Device	2
3.3 Mental Imagery in a Brain-Based Device	3
3.3.1 Previous Research	3
3.3.2 Current Research	4
3.4 Spiking Neural Model Simulation of Working Memory	5
3.5 A Novel Method for Analysis of Time Series Data	6
3.6 A Novel Unmanned Air and Ground Vehicle.....	7
3.6.1 Overview.....	7
3.6.2 Engineering Considerations and Design Iterations.....	9
3.6.3 Simulator.....	12
4.0 RESULTS AND DISCUSSION	12
5.0 CONCLUSIONS	14
6.0 REFERENCES.....	14
APPENDIX - PUBLICATIONS.....	15
LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS	110

LIST OF FIGURES

Figure		Page
1	Human Response Time Is Proportional to the Degree of Rotation	4
2	QuadHopter™ UAGV	8
3	First Version of the QuadHopter™ Showing Wheels and Central Gimbel with Camera and Accessory Boom	9
4	Second Version of the QuadHopter™	10
5	QuadHopter™ Version Three	11
6	AeroSim RC Simulator Showing QuadHopter™	12
7	Evidence of Mental Imagery in a BBD	13

1.0 SUMMARY

The ultimate goal of this project was to construct mobile devices capable of navigation, surveillance, and report. Control of these devices rests with simulated neuronal networks based on knowledge of higher brain functions that comprise cognitive control systems. In pursuing this goal, we made specific attempts to apply already demonstrated software strategies to develop a versatile set of spiking models, apply them to a recognition task, and to demonstrate a working memory. In the course of this work a new analytical method for spiking data was devised. We also designed and built a new type of mobile device capable of both flight and ground navigation while carrying a payload of surveillance sensors or other devices setting the stage for more autonomous systems.

2.0 INTRODUCTION

We have previously demonstrated neuronal network models with the ability to navigate, recognize objects, and learn from experience. However, in order to obtain the fast responses necessary for control in real-world applications and to take advantage of important timing relationships within the networks, spiking models (where the activity of a neural unit is defined by a transient voltage "spike," rather than by the average activity rate used in many previous models) are necessary.

Thus, this project involved significant effort related to the development and testing of large-scale networks of spiking neuronal models. Obtaining optimal configurations for stability and scalability required the refining of physiological and anatomical parameters of networks with varying configurations and with different numbers of neural units. We found that achieving good results with this approach was more challenging than with the rate models. We met this challenge by using spiking applications of a winner-take-all (WTA) network design. Of the several variations in connectivity rules tested, we found that only a center annular surround type exhibited appropriate WTA behavior, with stable firing of a relatively small group of units in response to different inputs. This architecture was therefore used in the succeeding work.

We begin by presenting an analysis of various network geometries that give rise to WTA behavior and describe a type of anatomy and dynamics that can be used in various contexts to build functional networks. Learning to reach to a target is one example of an application of this approach. We then describe simulations of interconnected networks of spiking neurons that learn to generate patterns of activity in correct temporal order. Animal behavior often involves a temporally ordered sequence of actions learned from experience which can be demonstrated in a Brain-Based Device (BBD). This work was then extended to a demonstration of a "cognitive" task by using a test for mental imagery related to mental rotation. The integrated output of networks with specific functional properties is necessary for producing appropriate behavior in time and space. Such behavior requires the formation of a short-term or working

memory. Additionally, time was also spent on exploring a new method to analyze temporal patterns in time series data, such as spike trains.

Finally, effort was expended to develop the concept of the QuadHopter™, a new type of mobile device capable of both flight and ground operations. This concept of a mobile platform is applicable to missions where a single mode of transport, either flying or driving, is not sufficient.

3.0 METHODS, ASSUMPTIONS, AND PROCEDURES

3.1 Construction and Testing of Versatile Winner-Take-All Networks

This aspect of the effort focused on the simulations of large-scale networks of excitatory and inhibitory neurons that incorporate realistic spiking kinetics, connectivity, and synaptic plasticity. These networks can generate dynamically stable WTA behavior. In contrast to studies of networks composed of mean-firing-rate neurons in which center-surround connectivity is sufficient for WTA dynamics, we found that a singular type of microcircuit connectivity, center-annular-surround (CAS), gave rise to WTA behavior in large-scale spiking networks. We show that these networks can form smooth maps in response to patterned sensory input. In addition, we show that a humanoid Brain-Based-Device (BBD) under the control of a spiking WTA neural network can learn to reach to target positions in its visual field, thus demonstrating the acquisition of sensorimotor coordination. A full report on this work and relevant scholarly references can be found in Appendix A1.

3.2 Temporal Sequence Learning Demonstrated in a Brain-Based Device

Animal behavior often involves a temporally ordered sequence of actions learned from experience. Here we describe simulations of interconnected networks of spiking neurons that learn to generate patterns of activity in correct temporal order. The simulation consists of large-scale networks of thousands of excitatory and inhibitory neurons that exhibit short-term synaptic plasticity and spike-timing dependent synaptic plasticity (STDP). The neural architecture within each area is arranged to evoke WTA patterns of neural activity that persist for tens of milliseconds. In order to generate and switch between consecutive firing patterns in correct temporal order, a reentrant exchange of signals between these areas was necessary. To demonstrate the capacity of this arrangement, we used the simulation to train a BBD that responded to visual input by autonomously generating temporal sequences of motor actions.

Our previous models of WTA spiking networks were coupled together and trained to generate segmented and sequential neural activity. The neural system is composed of thousands of simulated biologically realistic excitatory and inhibitory spiking neurons. The single compartment neurons modeled in these simulations display voltage

dynamics similar to those seen in cortical neurons. Activity of the simulated neurons reflects the conductance of well-known ion channels. Synapses were subject both to short-term synaptic plasticity and to STDP, which modeled the long-term synaptic changes that allowed the system to learn temporal sequences. We found that networks composed of spiking neurons of this sort, when trained to respond to repeated sequences of sensory cues, generate temporally ordered patterns of neuronal activity consisting of brief steady states separated by sharp transitions that resemble those observed in functioning brains. We found that the model could be used to control specific motor sequences in a BBD. In this research we used the hominid BBD that we denote as APE-X. The population activity pattern in this modeled neuronal system has similarities to those observed in primate prefrontal cortex during multi-segmented limb movements. A full report on this work and related scholarly references can be found in Appendix A2.

3.3 Mental Imagery in a Brain-Based Device

We have chosen to model how mental functions involved in behavior can arise from brain mechanisms. Specifically, we have chosen to model a system capable of carrying out a mental rotation task. In this well-known behavioral experiment, a subject is shown a picture of an object in a certain orientation together with a picture of the same or a different object rotated into a position different from that of the first. The task is to declare whether or not the first object is the same as the second. It has been inferred that human subjects actually mentally rotate the image of one of the objects to make the comparison, since the observed time to make a decision correlates well with the degree of the object's rotation between the two images (Figure 1, [1]). Emulating this behavior requires the integration of multiple levels of neural models, including spiking networks, various mappings, and interactions with the environment. We believe that success in this approach can be applied to a variety of problems in cognitive neuroscience.

3.3.1 Previous Research

Previous research in mental imagery lead to initial results similar to those in human tests. Our results were preliminary in several ways: 1) they were generated entirely in a simulation and did not use the APE-X BBD, 2) the dopamine-dependent STDP "value" system was not included, and 3) the motor responses to "match" and "non-match" pairs of images were not learned. This latter point is important in order to rule out classical conditioning, rather than mental imagery, as the basis for the motor responses.

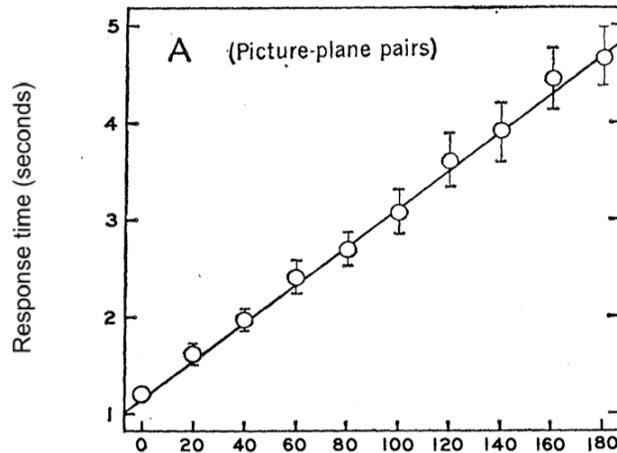


Figure 1. Human Response Time Is Proportional to the Degree of Rotation

3.3.2 Current Research

For the current effort we designed an experimental protocol to overcome the limitations identified above. First, APE-X was trained to make distinct movements when presented with pairs of objects that were either identical or not identical. Pairs where one object was a rotated version of the other were not used in this training. Second, the sequence generation network, described above, was trained to generate patterned activity when shown an object in a series of rotated positions. Finally, APE-X was tested with a pair consisting of an object and the same object rotated. It correctly reported a match. Thus the system was able to generalize from its training to new situations.

This series of experiments is described here in more detail. According to the delayed match-to-sample (DMS) protocol, APE-X was repeatedly presented with pairs of objects in succession: either two identical objects (a “match”) or two different objects (a “non-match”). A value system with dopamine-dependent STDP was used to modify synaptic strengths in the spiking networks so that APE-X successfully learned to make an appropriate “match” or “non-match” arm movement in response to the pair of objects. Sixteen stimuli (two different objects and their mirror image objects, each positioned at four orientations) were used. Pairs were presented with a one-second delay between each object. For training of match responses, the same object was presented twice. For training of non-match responses, each object was paired with its mirror object in the same orientation. It is important to note that two successive stimuli were never of the same object at different orientations.

After training, the value system was turned off so that synaptic modifications would not be made during testing. We then tested APE-X to verify that it would not give a match response unless the presented objects were exact matches. As expected, APE-X gave a “match” response in 39 of 40 trials in which the two stimuli were identical. It also correctly gave a “non-match” response in 276 of 280 trials where the stimuli were not identical, including 140 of the latter trials where the second stimulus was a rotated

version of the first. This result was expected, since the sequence generation network had not yet been trained.

To activate and train the sequence generation network, APE-X was presented with a series of stimuli in which the same object was rotated for 0, 90, 180, and 270 degrees about a central axis. The stimuli were the same set of sixteen used for the earlier training, but they were now presented to APE-X in an ordered sequence reflecting the rotation of the object.

APE-X was then tested, as before, using pairs of stimuli that were rotated versions of the same object. APE-X now reported that two stimuli matched. Thus the same stimuli that APE-X reported did not match prior to training the sequence generation network were now reported as matching. APE-X could not have been conditioned to respond as it did, since its response involved changes after conditioning had been completed. We conclude that it must be using its new ability to mentally rotate objects internally in order to produce this match response.

We repeated the entire training and testing procedure with five different “subjects,” each differing only in the initial conditions used to generate their detailed neuroanatomy. The overall performance for the five subjects was 92% correct for all test trials, whether match or non-match.

3.4 Spiking Neural Model Simulation of Working Memory

Retaining a fleeting perception for seconds or minutes after a stimulus disappears is critical for many forms of behavior, cognition, and learning. Working memory (WM) allows for the cognitive manipulation of stored information about stimuli, and such memories can be used in decision making. An interesting and well-known feature of WM is its ability to hold simultaneously only a limited number of different items.

In both human and animal studies, WM has often been investigated using a DMS paradigm. Typically an animal is shown a brief stimulus to be remembered for a few seconds or minutes. After a delay period, during which no stimulus is presented, the animal is shown a second stimulus that might or might not be identical to the first. A correct response, indicating whether the two stimuli match, results in the delivery of reward.

We made a large-scale spiking neural model with persistent activity that enables multi-item working memory. The network incorporates three distinct biological mechanisms for generating persistent activity. All three mechanisms operate simultaneously in real cortical circuits, and each mechanism has been shown to be independently capable of supporting persistent activity. These mechanisms are: (1) dense reentrant connectivity producing attractor dynamics, (2) short-term synaptic plasticity enabling robustness against brief drops of firing rate, and (3) relatively long-acting glutamatergic receptors maintaining excitation over durations longer than input

inter-spike intervals. Persistent activity in the network is characterized in relation to parameters controlling these mechanisms. DMS tasks also require a neural mechanism for detecting a match between persistent activity and activity evoked by the current stimulus. We proposed a matching mechanism based on the segregation of visual and memory-related inputs onto fast-responding glutamatergic and longer acting glutamatergic receptors of postsynaptic neurons. The ability of the network to perform visual DMS tasks was examined. Finally, we characterized the capacity of the network to store multiple items simultaneously as a function of network size. A full report on this work and relevant scholarly references can be found in Appendix A3.

3.5 A Novel Method for Analysis of Time Series Data

Simulations of biologically realistic neuronal networks or experimental studies all produce large amount of data, the majority of which can be characterized as showing neuronal spikes as a function of time, i.e., spike trains. Many techniques have been developed to identify correlations or patterns within such data. However, current techniques are not always successful at detecting and identifying patterns in lifelike scenarios: a large set of neurons, arbitrary delays between firings in the pattern, jitter in firing times, and firing failures. Determining patterns in the data is an important first step in analyzing the underlying phenomena and making a reasonable biological interpretation.

We have developed a neuronal pattern detection algorithm capable of detecting patterns in the data, where the only parameter required is the maximum duration of a pattern one is looking for. This method can be used on thousands of neurons to detect and identify patterns with arbitrary delays, and it is robust against jitter and firing failures. It is an extension of the method of Lopes-dos-Santos, et al. (2011).

For this purpose, a neuronal pattern P consists of a set of neurons ($N1, N2, \text{etc.}$) and associated delays ($D1, D2, \text{etc.}$). Activation of pattern P at time t is defined as the firing of neuron $N1$ at time $t+D1$, the firing of $N2$ at time $t+D2$, etc. A neuronal pattern can be activated at multiple times ($T1, T2, \text{etc.}$). If the firing pattern is activated enough times in the data set, then it is possible to statistically identify the set of neurons and corresponding delays to fully characterize P .

The algorithm is as follows: 1) Identify the “strongest” neuronal pattern P in the data, 2) Determine if P is statistically significant, 3) Remove all activations of P from the data, and 4) Repeat steps 1-2-3 until no statistically significant pattern is detected.

Two matrices are required to extract the “strongest” neuronal pattern: a correlation matrix C and a delay matrix D . Element $C(i,j)$ of matrix C corresponds to the maximum value of the Pearson correlation between the spike activity of neuron i against neuron j as the spike activity of neuron j is time-shifted from a delay of 0 to d ms. The amount of shift required to generate the maximum correlation in $C(i,j)$ is stored in $D(i,j)$. Given the results of Lopes-dos-Santos et al., we extract the strongest eigenvector from C , which represents the strongest neuronal pattern P .

To determine if P is statistically significant, we randomly time shuffle the spiking data and extract the corresponding correlation matrix C. The maximum eigenvalue E determines the significance of pattern P. If the eigenvalue associated with the eigenvector of P is greater than E, then P is statistically significant.

Once P is identified then all spiking data related to P are aligned using the associated delays. This alignment detects the activations of pattern P. The removal of these activations is then achieved by randomly shifting all spikes of P at activation times.

To test the method, we used a two-pattern case and a large synfire model. In the first test, we simulated a set of 10 neurons spiking over a 1,000 sec period. Two random patterns were generated; each was formed by 4 neurons spiking in order over 50 msec or less. Each pattern was activated 30 times over the simulation. To model a more lifelike biological scenario, each spike was randomly jittered within ± 2 ms, and each activation included only 3 of the 4 neurons (25% firing failure rate). Two hundred Monte Carlo simulations were carried out, and the 2 patterns were correctly detected and identified in 94.5% of the tests.

The second test case simulated a synfire chain within a simulation of 2,000 spiking neurons. The neurons were fully connected with generally weak synapses, while a small subset of neurons was connected with strong synapses to form the synfire chain. The synfire chain was built as 10 pools of 30 neurons each with a propagation delay of 3 msec. The chain was randomly activated 10 times with random current levels. The method detected and identified 99% of the neurons forming the synfire chain (298 out of the 300).

3.6 A Novel Unmanned Air and Ground Vehicle

A central and critical aim of this work has been to construct a mobile platform useful for surveillance both under manual control and potentially under various degrees of autonomous control by simulated neural systems. The platform described below is a proprietary design developed under this contract. Patents are pending.

3.6.1 Overview.

The QuadHopter™ (Figure 2) uniquely combines features of an unmanned aerial vehicle (UAV) and an unmanned ground vehicle (UGV) to form the first unmanned aerial-ground vehicle (UAGV). Thus it can carry out tasks to which neither type of traditional unmanned vehicle is suited. It has many potential applications including information gathering, surveillance, communications, payload delivery and retrieval, and search and rescue.

Based on a quad rotorcraft for flight stability and maneuverability, the QuadHopter™ incorporates four powered wheels for mobility on surfaces. In addition to

driving on the ground by using its wheels alone, the QuadHopter™ can travel along ceilings by taking advantage of the lift provided by the rotors to keep the wheels in contact with the ceiling.

Another key feature is the hook and central gear system that allows the QuadHopter™ to perch on structures such as building parapets and walls or to hang from overhead structures such as wires. Thus, the device can be used for surveillance operations, for example, while using minimal power. When the operation is complete, the QuadHopter™ simply flies away.

The QuadHopter™ design can be implemented at a range of scales depending on functional, payload, and power requirements. The current version described below is about 26 inches (66 cm) square and 11 inches (28 cm) high and weighs about 4.5 pounds (2 kg). This size is suitable for use inside buildings as well as outdoors. The frame is constructed of carbon fiber using novel puzzle-fit and tongue-and-groove methods. The radio-control, telemetry, and flight-stabilization systems are built from commercial off-the-shelf (COTS) components and can be implemented in several ways for different purposes. A variety of sensor or effector payloads can be envisioned. To aid in control system development and flight training, a computer simulation of the QuadHopter™ has been developed.



Figure 2. QuadHopter™ UAGV

3.6.2 Engineering Considerations and Design Iterations

The device has gone through three major design revisions in an evolutionary process. Brief descriptions of the three evolutionary stages of the design are presented below.

The initial design began with a standard quad X frame with added wheels to allow the device to drive on the ground like a four-wheel drive robot (Figure 4). Given that cameras on a flying robot are usually mounted on the bottom to see the ground below, and that cameras on standard ground robots are usually on the top or on a vertical arm to see as far ahead as possible, it seemed desirable to have a camera or other sensor that could pivot 360 degrees about the center of the robot. The central differential gimbel was designed to achieve this goal. A short boom attached to the gimbel would allow for mounting a camera or other sensor, and a second boom mounted at 180 degrees to the short boom would allow additional devices to be mounted (see Figure 4). Shortcomings of the initial design became apparent and were addressed in the next iteration.

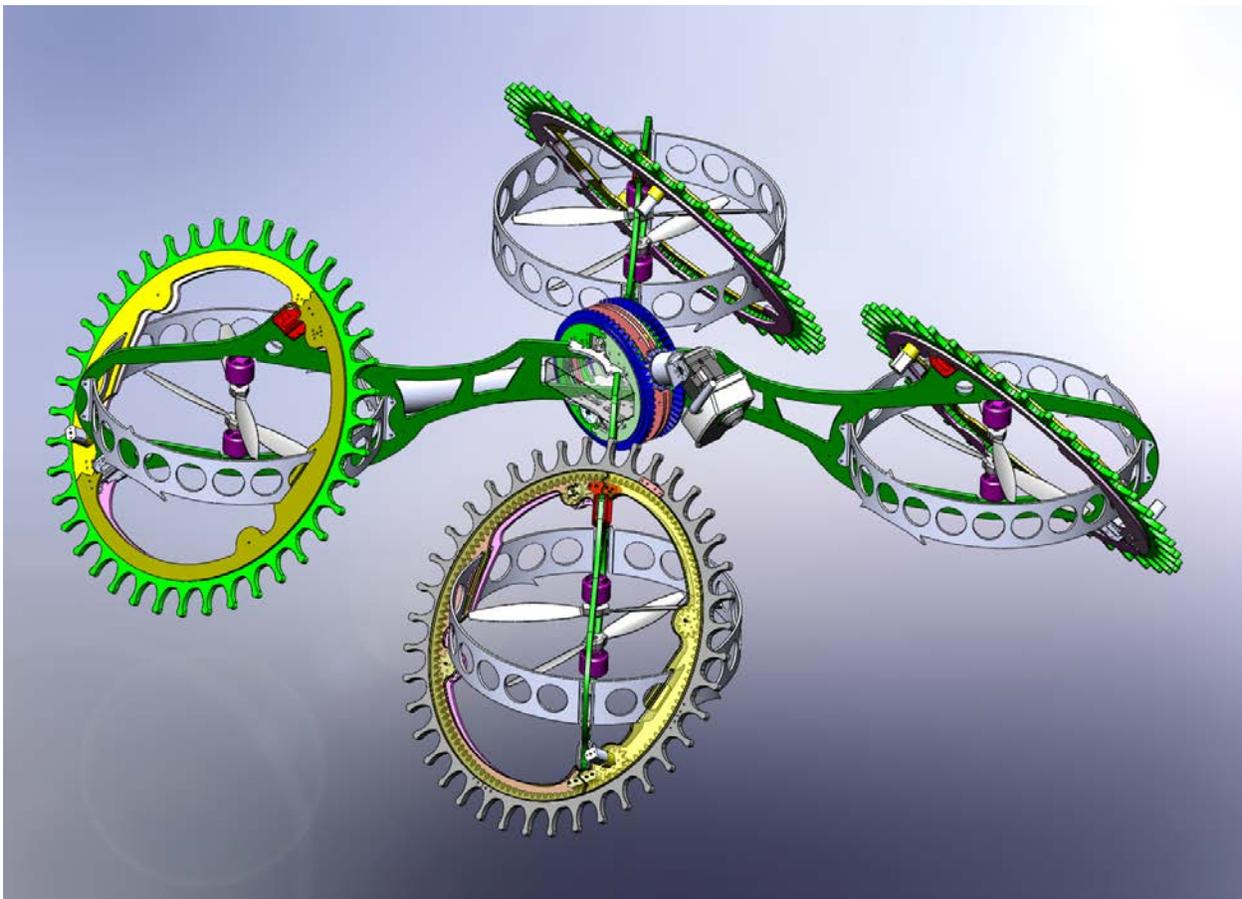


Figure 3. First Version of the QuadHopter™ Showing Wheels and Central Gimbel with Camera and Accessory Boom

The second revision of the design included many improvements in the overall concept (Figure 5). The frame was redesigned to make it smaller to better fit through standard doorways while retaining the 11-inch wheel diameter. The details of the wheel design were improved to make driving easier. An initial titanium frame was replaced by a carbon fiber frame. More powerful motors were used for the rotating center mechanism and for driving the wheels. The design included a provision for collective pitch for the props which was later removed to simplify construction.

The most significant addition was that of a hook, attached to the central gimbal as a second boom. Properly located, the hook allows the device to hang from wires or similar structures. This approach would save power or possibly allow the device to charge inductively from a power line. The dimensions of the hook and the gap between the wheels are such that the device can also perch on roof parapets, wall tops, or other such ledges. This capability would be useful to monitor activities in an area of interest from a variety of vantage points. While this second version was an improvement over the initial design, it was determined that there were areas where improvements that could be made in a third version.

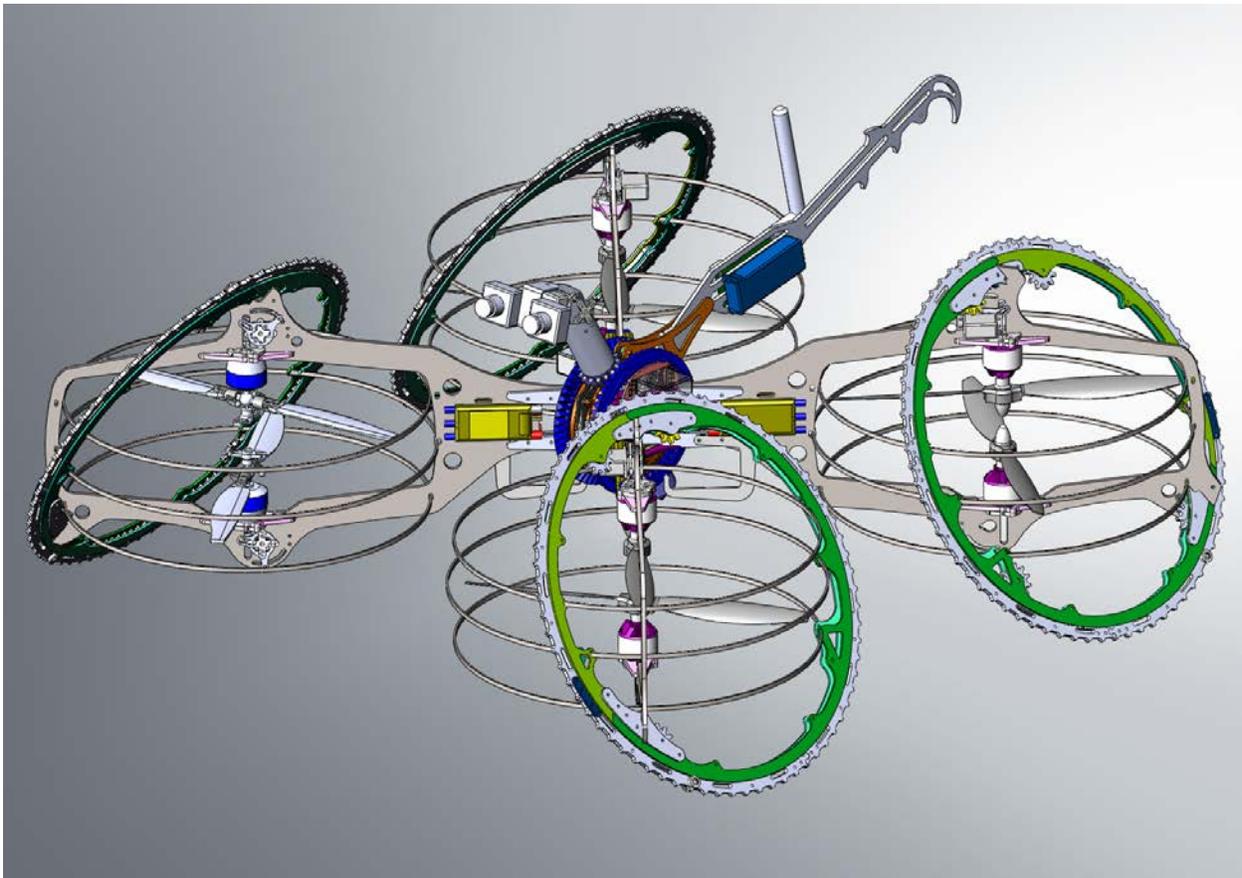


Figure 4. Second Version of the QuadHopter™

The third version (Figures 2 and 5) incorporates a significant revision to the way the wheels were driven. The ring gear used to drive each wheel in the previous version was inverted so that it could be driven from outside, and one center drive motor on each side was used. In addition to allowing for future modifications to the surface contact part of the wheel (e.g., adding tank treads), the addition of the drive gear made the carbon frame more rigid. Integration of the new Robotis AX-12 servos for the drive wheels and the center gimbel payload mechanism allowed for even more speed and torque.

Figure 5 also identifies some of the COTS components used for the control systems and sensors. We developed code for the Arduino Nano to drive the AX12 smart servos so that we could add more control options to the QuadHopter. Future plans include a better way to intuitively control the device and possibly to allow a single pilot to also operate a robot arm or other tool.

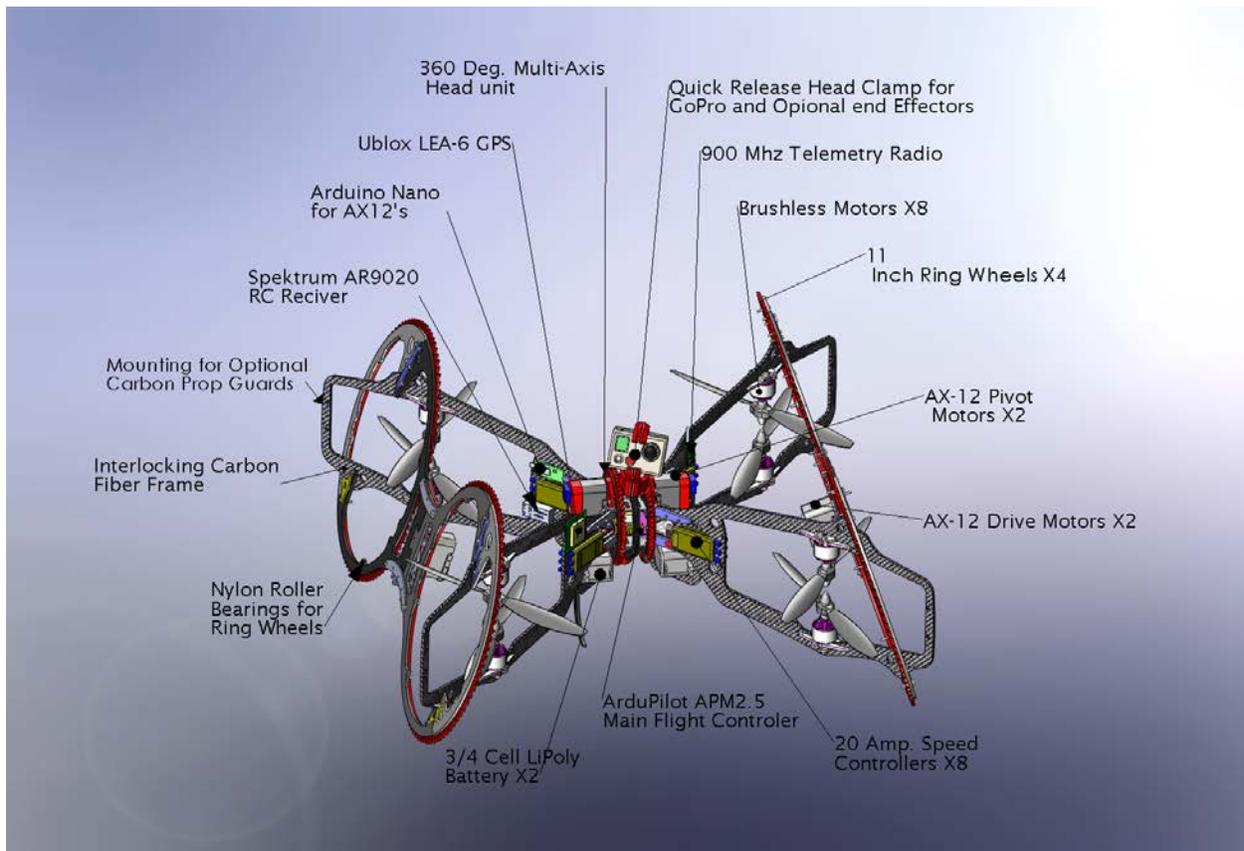


Figure 5. QuadHopter™ Version Three

3.6.3 Simulator

In order to facilitate testing and operator training, we have employed a commercial simulator called AeroSim RC. It interfaces with the ArduPilot Planner software from DIY Drones to get real data from the QuadHopter™. We converted the original SolidWorks design files to the Open Scene Graph format usable by AeroSimRC (Figure 6). This enabled “hardware in the loop” mode, where real QuadHopter can fly either in the simulated world of AeroSim RC or the real world. All the data that is coming from the real sensors will affect the simulated device. This can be very useful for debugging complicated code. More remains to be done with the simulation, but the center-pivoting camera is movable in three axes and can be stabilized by its own movement. Also, the wheels will rotate in contact with the ground. The simulator can also be useful for training people to fly the craft without risk.



Figure 6. AeroSim RC Simulator Showing QuadHopter™

4.0 RESULTS AND DISCUSSION

During our research we were able to implement and test WTA networks that exhibit the properties of a working memory. The research leads to an alternative

architecture being proposed for which WM capacity does not scale with respect to network size. This result is consistent with data showing that animal species with very different brains sizes may have similar visual WM capacity.

In the mental imagery research involving the APE-X BBD we were pleased to see results that were obtained. The timing results for match trials (Figure 7) indicated that the time to make a match response was proportional to the angular difference between the two presented objects, as was the case for the human subject data originally reported by Shepard and Metzler (Figure 1). The response time is the elapsed time between the offset of the second object image and the initiation of the BBD's "match" movement. Note that the data are from the newly completed series of experiments as described in this text. During this research, we overcame some of the limitations of our earlier simulation work by integrating the networks into APE-X, training APE-X to make appropriate motor responses, and testing APE-X in the mental rotation task. We found that APE-X successfully carried out the task, and we obtained behavioral data that matched experimental data from human subjects.

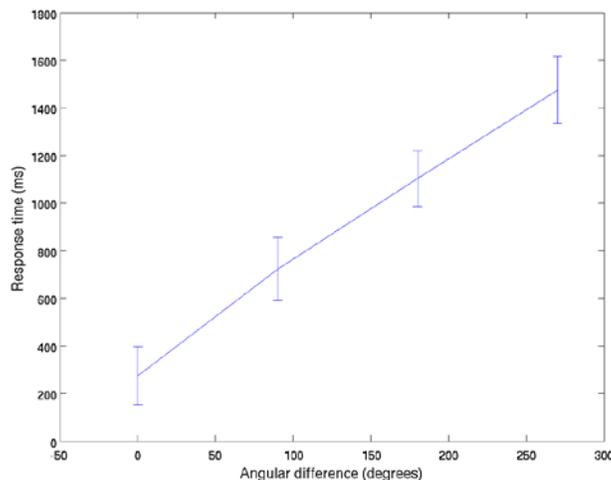


Figure 7. Evidence of Mental Imagery in a BBD

We believe the neuronal pattern detection algorithm can be extended to analyze temporal patterns in any type of time series such as patterns in weather data, video feeds, etc. In the case of video feeds, for example, this method could be adapted to identify objects moving with any type of repetitive patterns (people walking, cars driving, etc).

Additional detailed results and discussions can be found in in the manuscripts located in the Appendix.

5.0 CONCLUSIONS

This research involved the construction and testing of WTA networks which demonstrated the benefits of CAS connectivity over center-surround connectivity along with the acquisition of sensorimotor coordination. The results of the mental imagery portion of the research using APE-X suggest useful approaches to understanding the conscious generation of images that will be explored further. We also believe that the new analytical method for spiking data devised during this effort can be applied to any type of time series.

Furthermore, we believe that the QuadHopter™ concept presented above can be used in a variety of contexts, both for research and development, and for practical tasks. Further engineering development in terms of the best array of sensors, effectors, and scale will be needed to obtain optimum performance in a specific context.

This research demonstrates the ability to develop neuronal networks based on knowledge of higher brain functions. The results can be seen as taking a step closer to the vision of autonomous systems that many desire.

6.0 REFERENCES

1. Lopes-dos-Santos V, Conde-Ocazonez S, Nicolelis MAL, Ribeiro ST, Tort ABL (2011) "Neuronal Assembly Detection and Cell Membership Specification by Principal Component Analysis," *PLoS ONE* 6(6): e20996. doi:10.1371/journal.pone.0020996
2. Shepherd, R. and Metzler. J. (1971) "Mental rotation of three dimensional objects." *Science* 171(972):701-3.

(Note: See the publications in the Appendix for additional scholarly references.)

APPENDIX - PUBLICATIONS

A1. Chen, Y., McKinstry, J.L., and Edelman, G.M. (2013) "Versatile networks of simulated spiking neurons displaying winner-take-all behavior," *Frontiers in Computational Neuroscience*, doi: 10.3389/fncom.2013.00016.

Versatile Networks of Simulated Spiking Neurons Displaying Winner-Take-All Behavior

**Y. Chen, J.L. McKinstry, and G.M. Edelman
The Neurosciences Institute**

Abstract

We describe simulations of large-scale networks of excitatory and inhibitory spiking neurons that can generate dynamically stable winner-take-all (WTA) behavior. The network connectivity is a variant of center-surround architecture that we call center-annular-surround (CAS). In this architecture each neuron is excited by nearby neighbors and inhibited by more distant neighbors in an annular-surround region. The neural units of these networks simulate conductance-based spiking neurons that interact via mechanisms susceptible to both short-term synaptic plasticity and STDP. We show that such CAS networks display robust WTA behavior unlike the center-surround networks and other control architectures that we have studied. We find that a large-scale network of spiking neurons with separate populations of excitatory and inhibitory neurons can give rise to smooth maps of sensory input. In addition, we show that a humanoid Brain-Based-Device (BBD) under the control of a spiking WTA neural network can learn to reach to target positions in its visual field, thus demonstrating the acquisition of sensorimotor coordination.

1. Introduction

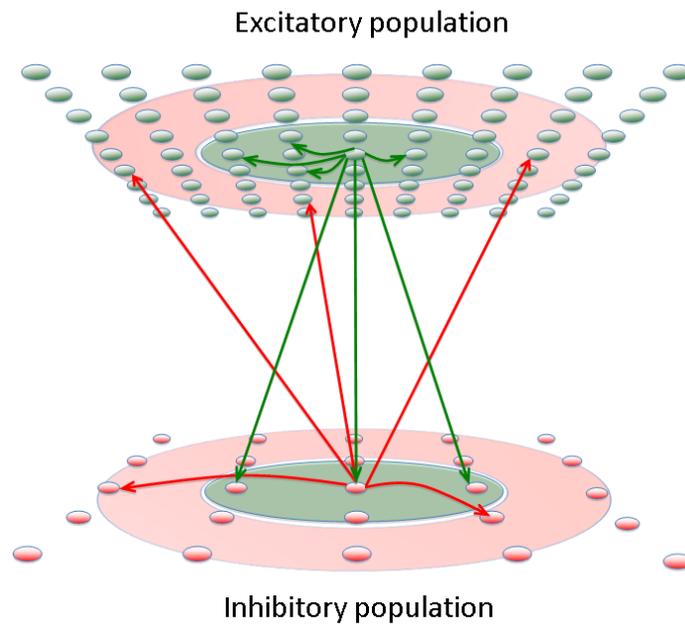
Analyses in computational neurobiology have successfully used mean-firing-rate neuronal models to simulate the spatiotemporal patterns of neural activity that arise in interconnected networks of excitatory and inhibitory neurons, such as those in the vertebrate cortex (von der Malsburg, 1973; Obermayer et al, 1990; Dayan and Abbot, 2001). Certain aspects of these systems may, however, require the modeling of the dynamic properties of large populations of individual neurons, each calculated with millisecond precision. Simulations of such systems are challenged with issues such as nonlinearity, instability, and resistance to scaling. Here we address these issues by simulating networks of spiking neurons that are capable of sensory map formation and sensorimotor interactions.

It has been proposed that local microcircuits of the cerebral cortex can function as Winner-Take-All (WTA) networks (Douglas and Martin, 2004). In such systems, an individual pattern of input can evoke network responses that suppress possible alternative responses. In addition, the population response to any sensory stimulus is sparse. This proposal is attractive for several reasons. On theoretical grounds, WTA networks have demonstrated utility in models of pattern recognition (von der Malsburg, 1973), map formation (Obermayer et al, 1990), selective attention (Itti et al, 1998), and working memory (Wilson and Cowan, 1973). The proposal is also supported by cortical anatomy. A characteristic structural feature of WTA networks is long range inhibition among cellular components coupled to short range excitation. Anatomical evidence exists

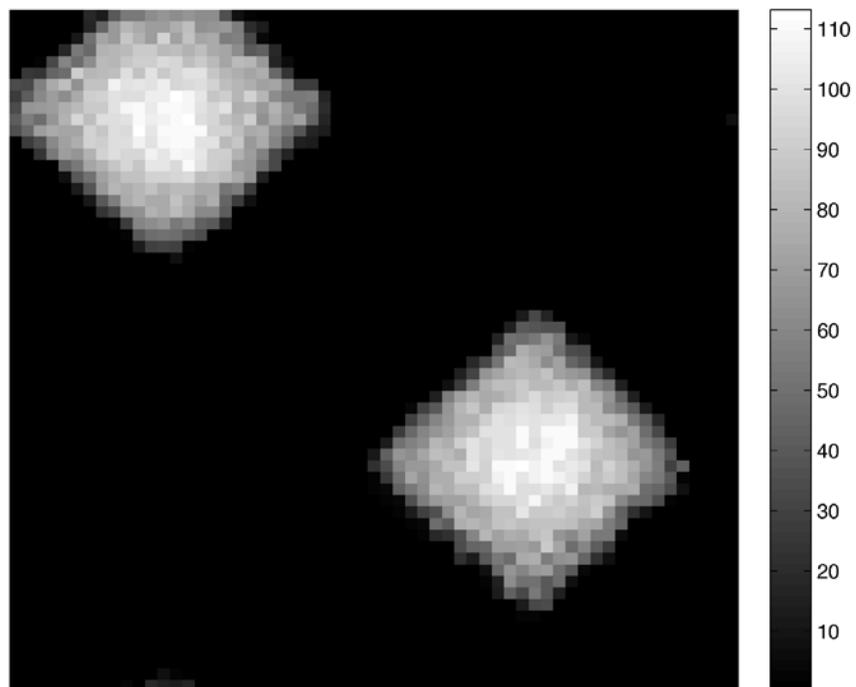
for such an architecture in animal nervous systems (Perin et al, 2011; Goldman-Rakic, 1995; Kisvárdy et al, 2000; Holmgren et al, 2003; Fino and Yuste, 2011). Indirect physiological evidence (Derdikman et al, 2003; Haider et al, 2010) has also been obtained for local excitation and surround inhibition in the cerebral cortex of mammals.

Rate-based WTA networks with center-surround architecture have been extensively explored (Dayan and Abbott, 2001). Although these networks have been shown to possess useful properties, they lack the temporal precision and biological realism of networks of spiking neurons. In some prior studies of spiking models capable of WTA behavior the neuronal network structure has been highly simplified. Networks are simulated as a one-dimensional chain or ring (Shriki et al, 2003; Laing and Chow, 2001). The inhibitory population may be reduced to one unit (Rutishauser et al, 2011; Oster et al, 2009), or the inhibitory population was removed altogether and modeled as direct inhibitory connections among excitatory neurons (Laing and Chow, 2001; Choe and Miikkulainen, 2004). One large-scale spiking model did produce smooth maps of orientation columns, but this model also combined excitatory and inhibitory neurons into a single population, and did not incorporate spike-timing dependent plasticity (STDP) (Choe and Miikkulainen, 2004). If the complex circuits of the cortex function as WTA networks, biologically realistic spiking models must exhibit robust WTA network dynamics that can explain behavior at the systems level.

In the present study we describe a general and robust computer simulation of the activity within neural networks containing thousands of excitatory and inhibitory spiking neurons in a variant of center-surround architecture that we call center-annular-surround (CAS). In this architecture each neuron is excited by nearby neighbors and inhibited by more distant neighbors in an annular-surround region (Figure 1A). The neural units of these networks simulate conductance-based spiking neurons that interact via mechanisms susceptible to both short-term synaptic plasticity and STDP. We show that such CAS networks display robust WTA behavior unlike the center-surround networks we have studied. We demonstrate for the first time that a large-scale network of spiking neurons with separate populations of excitatory and inhibitory neurons can give rise to smooth maps of sensory input (Obermayer et al, 1990). We also show that, a brain-based device (BBD) under the control of a system of such networks learns to reach to visual targets.



(A)



(B)

Figure 1

2. Materials and Methods

Spiking Neuronal Networks - Each modeled network (Figure 1A) is comprised of 3 interconnected populations of spiking neuronal units (Izhikevich, 2010) distributed over two-dimensional square grids. Each population is composed of units simulating one of three functional classes of spiking neurons: input (“thalamic”), excitatory, and inhibitory. The parameters of simulated neurons in each class are tuned so that the voltage waveform mimics its biological counterpart (Izhikevich, 2003). The synapses display STDP and short-term plasticity dynamics as previously described in detail (Izhikevich and Edelman, 2008). The neuron model equations, short-term synaptic plasticity equations, and STDP equations are presented below.

Neuronal Dynamics – Spiking dynamics of each neuron were simulated using the phenomenological model proposed by Izhikevich (2003). The model has only 2 equations and 4 dimensionless parameters that could be explicitly determined from neuronal resting potential, input resistance, rheobase current, and other measurable characteristics. We present the model in a dimensional form so that the membrane potential is in millivolts, the current is in picoamperes and the time is in milliseconds:

$$C\dot{v} = k(v - v_r)(v - v_t) - u - I_{syn} \quad (1)$$

$$\dot{u} = a\{b(v - v_r) - u\} \quad (2)$$

where C is the membrane capacitance, v is the membrane potential (in mV), v_r is the resting potential, v_t is the instantaneous threshold potential, u is the recovery variable (the difference of all inward and outward voltage-gated currents), I_{syn} is the synaptic current (in pA) defined below, a and b are parameters. When the membrane potential reaches the peak of the spike, i.e., $v > v_{peak}$, the model fires a spike, and all variables are reset according to $v \leftarrow c$ and $u \leftarrow u + d$, where c and d are parameters. Supplementary Table I lists each of the neuron model parameters used in all experiments. At the start of all simulations, v was set to -60 for all neurons, whereas u was set to a different random value for each neuron drawn uniformly from the range 0 to 100.

Short-Term Synaptic Plasticity – The strength of synapses varied as a function of the presynaptic neuron’s firing history. We assume that the synaptic conductance (strength) of each synapse can be scaled down (depression) or up (facilitation) on a short time scale (hundreds of milliseconds) by a scalar factor x . This scalar factor, different for each presynaptic cell, is modeled by the following one-dimensional equation

$$\dot{x} = (1 - x)/\tau_x, x \leftarrow px \text{ when presynaptic neuron fires.} \quad (3)$$

x tends to recover to the equilibrium value $x = 1$ with the time constant τ_x , and it is reset by each spike of the presynaptic cell to the new value px . Any value $p < 1$ decreases x and results in short-term synaptic depression, whereas $p > 1$ results in short-term synaptic facilitation. The parameters, τ_x and p , for each combination of presynaptic and postsynaptic neuron types were as follows: exc. \rightarrow exc.: 150, 0.8; exc. \rightarrow inh.: 150, 0.8;

inh.→exc.: 150, 0.8; inh.→inh.: 150, 0.8; thalamic→exc: 150, 0.7; thalamic→inh.: 200, 0.5.

Synaptic Kinetics – The total synaptic current to each neuron is simulated as

$$I_{syn} = g_{AMPA}(v - 0) + g_{NMDA} \frac{[(v + 80)/60]^2}{1 + [(v + 80)/60]^2}(v - 0) + g_{GABA_A}(v + 70) + g_{GABA_B}(v + 90) + g_{SH}(v + 90) \quad (4)$$

where v is the postsynaptic membrane potential, and the subscript indicates the receptor type. Each conductance g (here we omit the subscript for the sake of clarity) has first-order linear kinetics $g' = -g / \tau$ with $\tau = 5$ ms, 150 ms, 6 ms, 150 ms, and 5,000 ms for each of the simulated AMPA, NMDA, GABA_A, GABA_B, and SH receptors, respectively. The SH “receptors” were an *ad hoc* method for adding slow hyperpolarizing (SH) currents in order to bias cells to remain off for longer periods of time; this improved pattern separation, and was used only in the BBD experiments.

Each firing of an excitatory neuron increases g_{AMPA} by xc , where c is the synaptic conductance (synaptic weight) in nanoSiemens and x is the short-term depression/potential scaling factor as above; g_{NMDA} was increased by $nmda_gain xc$, where $nmda_gain$ is the ratio of NMDA to AMPA conductances and is found experimentally to be less than one (Myme et al, 2003). Similarly $gabab_gain$ and $gabash_gain$ are used to adjust the contribution of g_{GABA_B} and g_{SH} respectively, relative to g_{GABA_A} . The gain factor for g_{SH} was set to zero for all simulations except for the BBD experiments in which case the gain factor was set to 0.2 for the first 45 simulation seconds and was set to 0.0 for the remainder of the simulation.

STDP – The change in conductance (weight) of each synapse in the model is simulated according to spike-timing-dependent plasticity (STDP): the synapse is potentiated or depressed depending on the order of firing of the presynaptic and postsynaptic neurons (Bi and Poo, 1998). We use the following equations to update each plastic synapse, s , in the network:

$$\dot{c} = -c / \tau_c + \alpha STDP(t) \delta(t - t_{pre/post}) \quad (5)$$

$$\dot{s} = c \quad (6)$$

where $\delta(t)$ is the Dirac delta function that step-increases the variable c . Firings of pre- and postsynaptic neurons, occurring at times t_{pre}, t_{post} , respectively, change c by the amount $\alpha STDP(t)$ where α is the learning rate for the synapse, $t = t_{post} - t_{pre}$ is the interspike interval, and

$$STDP(t) = \begin{cases} A^+ \exp(-1/\tau^+)t, & t > 0 \\ A^- \exp(-1/\tau^-)|t|, & t \leq 0 \end{cases} \quad (7)$$

where $A^+ = 0.005$, $A^- = 0.001$, $\tau^+ = \tau^- = 20$ ms. The variable c decays to zero exponentially with the time constant $\tau_c = 1$ s, and s is updated once every 50 ms for computational efficiency. Note that for simplicity, each synapse was modeled with a

single weight, s ; therefore the STDP rule changed both AMPA and NMDA components of the synapse proportionally.

Synaptic scaling – Synaptic scaling was performed for each neuron in order to maintain the total of all synaptic strengths on a given connection pathway, S_{total} , at a constant value. This scaling was performed for every neuron every 50 ms during the simulation. In addition, each synapse was prevented from exceeding S_{max} or going below zero, regardless of learning rules and normalization.

Anatomy – The Input network is composed of 484 simulated ‘thalamic’ neurons that provide excitatory input to “cortical” excitatory and inhibitory neurons. “Thalamic” neurons project to both “cortical” populations with uniform random connectivity. Current levels to these “thalamic” cells were adjusted to evoke distinct patterns of activity in the input area with a maximum firing rate of approximately 100 Hz for either abstract patterns or video camera input.

The cortical network contained 3,481 excitatory cells and 900 inhibitory cells. All connections made from cortical excitatory neurons to other neurons followed local-type connectivity. In this connectivity, a two-dimensional Gaussian probability distribution, centered on each cell, determined the probability of forming an input synapse to surrounding neurons. This probability density function was scaled to generate, on average, a pre-specified number of excitatory synapses onto each cell (see Supplementary Material for details). The initial synaptic strength between connected neurons also varied as a Gaussian function of the distance between them. The total of all synaptic efficacies for each simulated neuron was scaled to sum to a constant value unique to each neuron type.

In contrast, inhibitory neurons in the system exhibited center-annular-surround (CAS) connectivity. For CAS connectivity, each neuron received synaptic input only from neurons located in a surrounding area specified by a minimum (r_{min}) and maximum (r_{max}) radial distance from the postsynaptic cell. The probability of forming a connection with a neuron in the annulus was a function of the distance separating the cells. The function used was a Gaussian with standard deviation σ , centered at $(r_{min}+r_{max})/2$. This probability distribution function was scaled to create a prespecified number of inhibitory synapses onto each neuron. The synaptic strengths for the surround-type connection were also initialized using the same function, with the same parameters. However, the synaptic strengths of this type were scaled to make their sum equal to a constant value under experimental control.

We found that this CAS connectivity arrangement confers WTA properties to the networks. Each distinctive pattern of neural activity in the “thalamic” network evoked enhanced neural activity in only a few localized patches in the “cortical” area due to competitive interactions between local neural populations (Figure 1B). Local patches of interconnected neurons that on average respond better than surrounding cells ‘win’ a dynamic competition and remain active. In contrast, neurons in the surround are suppressed by inhibition and do not fire. A detailed description of the network along with

all the parameter settings used in the experiments can be found in the Supplementary Material, and connectivity parameters can be found in Supplementary Tables II-IV.

Winner-Take-All measure – We use the following measure of population sparseness (S) to characterize WTA dynamics in the excitatory population:

$$S = \frac{1 - \left(\sum_{j=1}^N \frac{r_j}{N} \right)^2 / \left(\sum_{j=1}^N \frac{r_j^2}{N} \right)}{1 - \frac{1}{N}}$$

where r_j is the number of spikes emitted by neuron j during the measurement interval, (one second in this paper) and N was the number of neurons in the population (Willmore and Tolhurst, 2001).

Brain Based Device (BBD) - To demonstrate that a simulated network can control real-world behavior, we designed and constructed a humanoid brain-based device (BBD). The device is 50 cm high and uses a black and white wireless webcam for vision. Each arm of the BBD contains eight Dynamixel servomotors (Robotis, Irvine, CA, USA). In the specific experiments described here only the two shoulder joints function; all other joints remain stationary with the arm extended. Shoulder joint angles provided by the motors determine the posture of the arms. A miniature PC (VIA Technologies, Fremont, USA) mounted on the back of the BBD maintains wireless communication between the device and the neuronal networks simulated on a Mac-Pro ([Apple, Inc. Cupertino, CA](#)).

A simulated motor neural network constructed and incorporated into the BBD controlled its behavior. This network was similar to the sensory network, but was composed of only 1600 excitatory and 400 inhibitory spiking neurons. Different patterns of activity in the motor area neurons specified distinct equilibrium postures of the left arm. Since the camera of the BBD was aimed at the left robotic hand, each of these postures presented a distinct pattern of visual input to the system. The motor region received non-topographic connections from the output of the sensory network. By adjusting parameters of feed-forward connections to the motor area from the cortical area receiving camera input, this system came to associate the visual input evoked by different postures to the motoric output pattern that would generate and maintain those postures.

Position error calculation – We measure the position error of a given joint as follows. For every arm posture measured during testing, we find the closest posture found during the training period. We then measure the angular difference of the joint between these two postures. We report the median and the maximum joint position error across all joints, reaching trials, and subjects.

3. Results

3.1. Spiking Activity in a WTA Network.

We first characterized spiking activity in the network as a function of the parameters of network connectivity (Figure 1). All analyses were carried out under the assumption of CAS connectivity described above, and examined the effects of different patterns of relative synaptic strengths on the various pathways in the network.

In these analyses, each simulated network received identical random input to ‘thalamic’ cells and started with identical random neural states, but had different values of total excitatory-to-inhibitory and inhibitory-to-excitatory synaptic strengths. The total weight of inhibitory-to-inhibitory synapses was kept equal to 2.4 times the total weight of inhibitory-to-excitatory synapses to limit the parameter space. The strengths of excitatory-to-excitatory synapses were kept constant in all simulations. Connection strengths were not modulated by STDP but were subject to the short-term synaptic plasticity inherent in modeled neurons (Izhikevich and Edelman, 2008). Exact values of all parameters are given in supplementary Table II. All spikes that the networks emitted between 2 and 3 seconds after the onset of thalamic input were recorded, at which time most simulations had reached steady state.

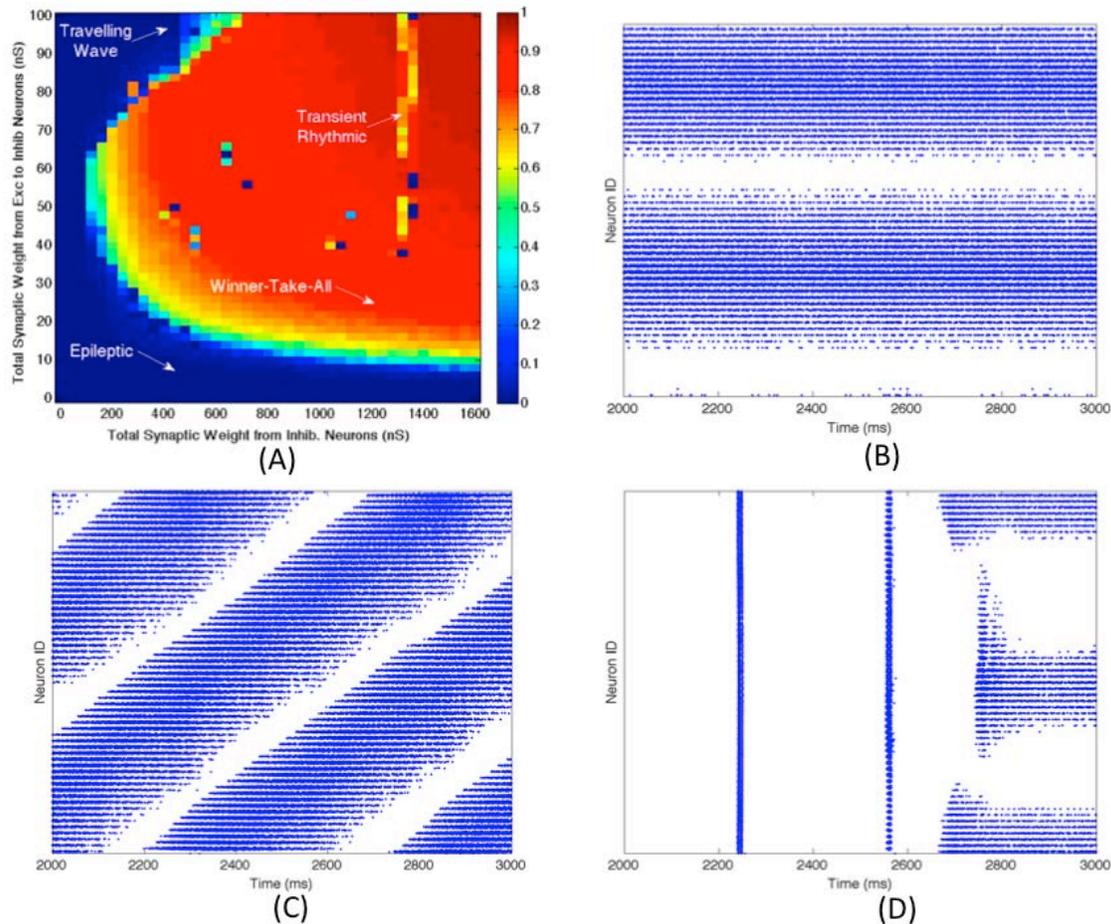


Figure 2

Figure 2A illustrates the dynamic behavior of these networks for 2,000 different combinations of excitatory-to-inhibitory and inhibitory-to-excitatory synaptic strengths.

The color of each pixel in Figure 2A is determined by a measure of the WTA behavior of the network dynamics in the same one second time period. Since WTA behavior entails sparse activity, we use a standard measure of population activity sparseness to characterize WTA behavior (see Materials and Methods). The measure will be close to one for networks in which only a small subset of neurons respond to the “thalamic” input with elevated firing rates. Parameters modeled in each raster plot in Figures 2B-D are indicated by a corresponding labeled arrow in Figure 2A.

When both excitatory and inhibitory connection weights were relatively high, local patches of excitatory neurons had a high maximal firing rate, as shown in the corresponding spike raster plot (Fig. 2B). However, only a localized subset, (25% of this neuronal population), maintained high firing rates; most neurons were silent. This outcome, in which a subset of neurons fires persistently at a high frequency and suppresses the activity of other neurons, defines a WTA network state. The majority of the parameter space explored corresponds to the WTA state as indicated by the predominance of warm colors in Fig 2A.

The spike raster plot in Fig. 2C shows activity within a network in a traveling wave state. The firing of both excitatory and inhibitory neurons moves as a localized “patch” through the network rather than remaining stationary in a WTA state. Fig. 2D shows a network that remained in an initial rhythmic, periodic state for a prolonged period after stimulus onset, but entered a WTA state towards the end of the third second of stimulus presentation. Single excitatory neurons maintained a state of high-frequency spiking activity only when connection strengths were within the WTA region delineated in Fig. 2A. Supplementary Figures 1-3 show close up plots from portions of Figures 2B-D.

For comparison, we also simulated the spiking behavior of networks of excitatory and inhibitory cells linked together in three different, non-CAS architectures. The three alternative network architectures analyzed were: (1) standard center surround architecture in which connectivity among all neurons was determined by a two-dimensional Gaussian probability distribution centered on each cell, inhibition having a larger σ than excitation; (2) an inverse connectivity in which the excitatory connections project to an annular surround and the inhibitory neurons connect locally, and (3) uniform random connectivity among all neuron types (excitatory-to-inhibitory, inhibitory-to-excitatory, excitatory-to-excitatory, and inhibitory-to-inhibitory). (See Supplementary Material for details of the parameters used.) In the same parameter space analyzed in Figure 2A, none of these connection types supported WTA behavior, characterized by stable patch activity. The maximum population sparseness measure for the three alternative network architectures listed above were 0.16, 0.54, and 0.21 respectively, whereas for the CAS network, the majority of the parameter space yielded population sparseness measures close to 1 (Fig. 2A). The most common firing patterns evoked in these neural networks were quasi-rhythmic firings of excitatory neurons in the 10 to 20 Hz range punctuated with short bursts of localized activity in inhibitory neurons. Among the different connectivities we analyzed, only the CAS motif gave rise to localized persistent activity that defines a WTA state.

3.2. Using CAS Architecture to Develop Maps of Orientation Selectivity

Smooth maps, in which nearby neurons have similar response properties, are ubiquitous in sensory and motor regions of the cerebral cortex (Obermayer et al, 1990; Kaschube et al, 2010). For example, in the primary visual area of many animals smooth retinotopic maps coexist with smooth maps of stimulus orientation. Computational neural models have successfully generated such smooth maps (Choe and Miikkulainen, 2004), but not, so far, with detailed networks of excitatory and inhibitory spiking neurons. It is therefore of interest to investigate whether such simulated networks of interconnected excitatory and inhibitory spiking neurons might produce such maps. We found that by slowly increasing inhibition in the model over time as experimentally observed (Ben-Ari et al, 2012), the CAS network described above develops smooth orientation maps when trained with oriented visual input.

The “thalamic input” to the “cortical” cells were given a rough initial topographic bias (Choe and Miikkulainen, 2004) by limiting the maximum distance over which “thalamic” inputs traveled to synapse on “cortical” cells. This simulation allows a maximum radius of 0.65 mm in a simulated 2mm by 2mm cortical region. Training stimuli consisted of 4,000 images of computer-generated elongated Gaussian shapes distributed throughout the visual field at random locations and orientations as in Choe and Miikkulainen (2004). STDP was used to train a network of 60 by 60 excitatory and 30 by 30 inhibitory neurons for 40,000 simulated seconds. Each of the 4,000 stimuli was presented to the network 20 times, and each presentation lasted 500 msec.

To assure smoothness in the resulting maps, more abstract models of orientation map formation have generally made use of an annealing process (Obermayer et al, 1990). This annealing process takes the form of a slow decrease of the size of the subpopulation of neurons active during the presentation of a stimulus (Kohonen, 1984). We sought a biological mechanism to accomplish this slow decrease in the active population size. Recent experimental evidence suggests that early during development, GABAergic conductances are excitatory rather than inhibitory (Ben-Ari et al, 2012). We hypothesized that such a lack of inhibition would lead to a large fraction of the population becoming active, and that slowly increasing inhibition during map formation would cause a monotonic reduction in active neurons. This has the same effect as the more artificial annealing process implemented algorithmically in abstract models of map formation.

We approximated this mechanism in our simulations by slowly increasing the GABAergic conductance of synapses onto excitatory cells, from zero to a plateau value. This plateau is reached one fourth of the way through the simulation (See Supplementary Table III). This mechanism had the desired effect. Early during map development, nearly one-quarter of the neurons in the network responded to each stimulus. This number was reduced to a small fraction of the neurons when inhibition reached its maximal level, and the active population remained small for the remainder of the training period (data not shown).

Finally, we tested the proposed annealing mechanism in conjunction with the CAS architecture for the ability to develop smooth maps. As shown in the resulting map

(Figure 3), nearby neurons in the network tend to have similar orientation preferences, i.e. the map is smooth, a characteristic of the primary visual cortex of cat, ferret, tree shrew, and monkey (Obermayer et al, 1990; Kaschube et al, 2010). In addition, dark areas are found at the centers of so-called orientation pinwheels, around which cells responding to all of the different orientations are found. The fact that each color occurs multiple times in the map reflects the fact that groups of cells respond to all orientations at each location in the visual field. This simplified spiking model based on the visual cortex develops orientation columns qualitatively similar to those found in the animal species mentioned above.

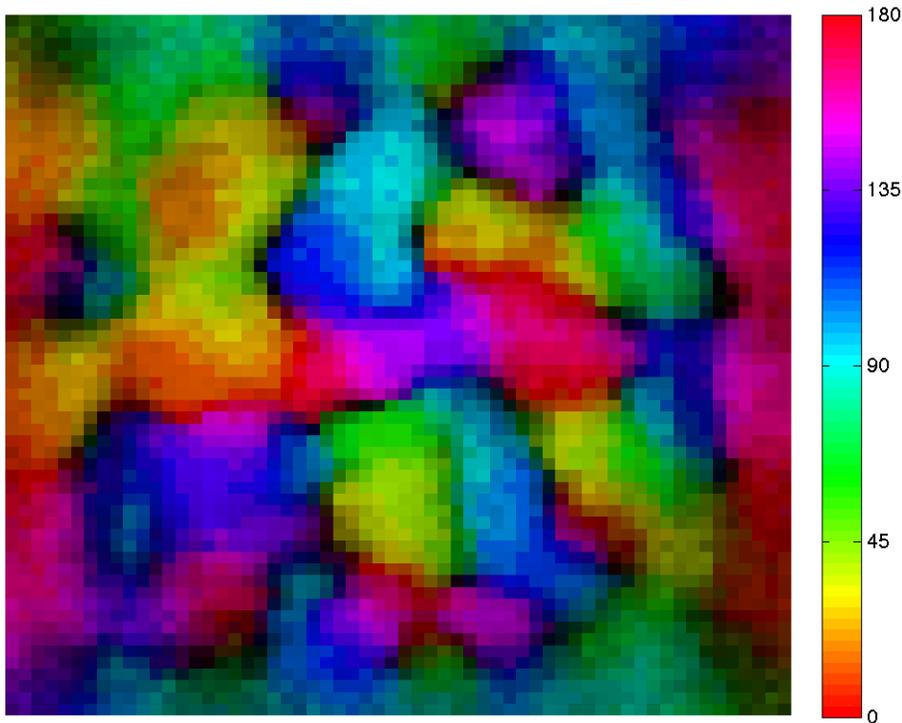


Figure 3.

3.3. Learning Hand-eye Coordination in a BBD Controlled by a Large Scale Spiking Network

The work of Davison and Frégnac (2006) demonstrated that STDP could be used to establish a mapping between two spiking networks with correlated spiking activity. We confirm that this finding holds in a real-world task in a large-scale model of approximately 7,000 spiking neurons, which was able to learn a mapping from visual targets to motor actions in a BBD.

To do so we coupled together two CAS networks to create a system that could learn the correlations between individual maps. After training, the output of a system of such networks controlled behavior in a real-world task: reaching to targets within the visual field of a BBD. To do this, we integrated a CAS-network motor map in a BBD. This motor map gave rise to autonomous arm movements, a form of “motor babbling”. With experience, this system came to correlate the location of the hand in its own visual field to the motor command needed to maintain the hand at that location, i.e. hand-eye-coordination.

The upper torso of the BBD maintained a seated posture that allowed a sufficient range of arm motion (see figure 4). The head unit containing a gray-scale video camera was aimed and held fixed during the experiment to allow the full range of motion of the left arm to fit into the camera’s field of view. A bright yellow object (5cm x 5 cm x 7 cm) attached to the end of the left arm allowed the visual system to detect the location of the end effector. The neural simulation controlled only the two shoulder joints of the BBD. Any given combination of the two joint angles yielded a unique arm posture and thus determined the location within the visual field of the bright object. The goal was to form a mapping between the visual input and joint angle commands that gave rise to that input.

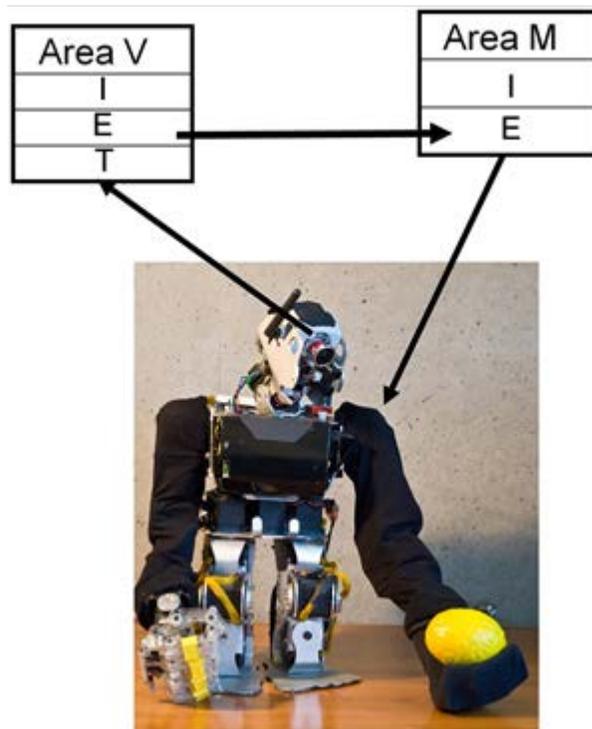


Figure 4.

The neural network controlling the behavior of the BBD consisted of the visual map area (V) and the motor area (M). Area V was a two-dimensional array of 3,481 excitatory and 900 inhibitory neurons. The network formed a topographic map of the visual input from the camera (see Supplementary Material for details of visual input processing). The

activity of each neuron in this array was roughly proportional to the brightness level of the corresponding pixel from the video input.

Area M, the motor area, contained 1600 excitatory and 400 inhibitory neurons. Each excitatory neuron was assigned a preferred set of angles for each of the two shoulder joints. Nearby neurons in this predetermined map responded to similar joint angles, but different patterns of activity among these cells could evoke all possible positions of the left arm. In order to translate from neuronal firings to joint angle in the left shoulder of the BBD, the output of these cells was pooled using population vector averaging (Georgopoulos et al, 1986). That is, for each joint, the preferred joint angles of all cells, weighted by the corresponding firing rate, were summed to determine an equilibrium posture. Joint angles were recalculated in this manner, and the angles of the shoulder joints were adjusted every 250 milliseconds.

To learn the mapping from visual input to motor output, area V was connected to area M with initially random one-way synaptic connections. In order to allow arbitrary mappings to form, the connections were all-to-all. STDP was calculated as described in Materials and Methods and was used to adjust the synaptic strengths during the learning process; short-term synaptic plasticity was used as described previously (Izhikevich and Edelman, 2008). In addition, the sum of the incoming synaptic strengths for each neuron was normalized to a constant value on this connection pathway. Supplementary Table IV gives all parameters used in this experiment.

In order to train the device to reach, a so-called motor-babbling reflex was incorporated in the BBD. During each movement trial of the training phase we directly stimulated one of nine different spots in the motor network by injecting current into excitatory neurons for 450 msec. This effectively drove the arm into a corresponding posture in open-loop fashion within approx. 100 msec., and the arm remained in a constant posture for nearly 400 msec. before the beginning of the next trial. A total of 15 repetitions, each generating nine postures, were used during this motor-babbling phase. During this time, STDP modulated the strength of connections between co-active neurons in the simulated visual and motor cortex, generating the visuomotor mapping.

After the training phase, direct motor cortex stimulation was turned off, and the target yellow object was detached from the BBD. With the arm of the BBD at its side, the target object was repeatedly placed by the experimenter in each of the nine spatial locations that it had occupied during training. This experiment was repeated five times; in each repetition, parameters and conditions were unchanged, except for initial synaptic strengths and connectivity that were controlled by a seed of the random number generator function from the standard C library (Kernighan and Ritchie, 1988). During the testing period, the arm moved in response to each new visual stimulus. Figure 5 shows the joint angles that correspond to the nine successive postures assumed by the BBD during training (blue) and testing (red) phases for all 5 experiments. The joint angles arrived at during testing cluster around those achieved in the training period, indicating an accurate mapping between visual and motor responses. To quantify the precision of equilibrium postures, a measure of the position error was recorded. We define position error at a given joint as the difference between the joint angles of the visually evoked

postures during testing and those recorded during the training period (See Materials and Methods). The median joint angle error, pooled across the two joints and across subjects, was 0.3 degrees; the maximum error was 13.6 degrees. Variability in manually positioning the stimulus in the visual field of the robot contributed to the variability in the motor error. A video clip showing the behavior of the system after being trained to reach to four positions is available in the online Supplementary Material.

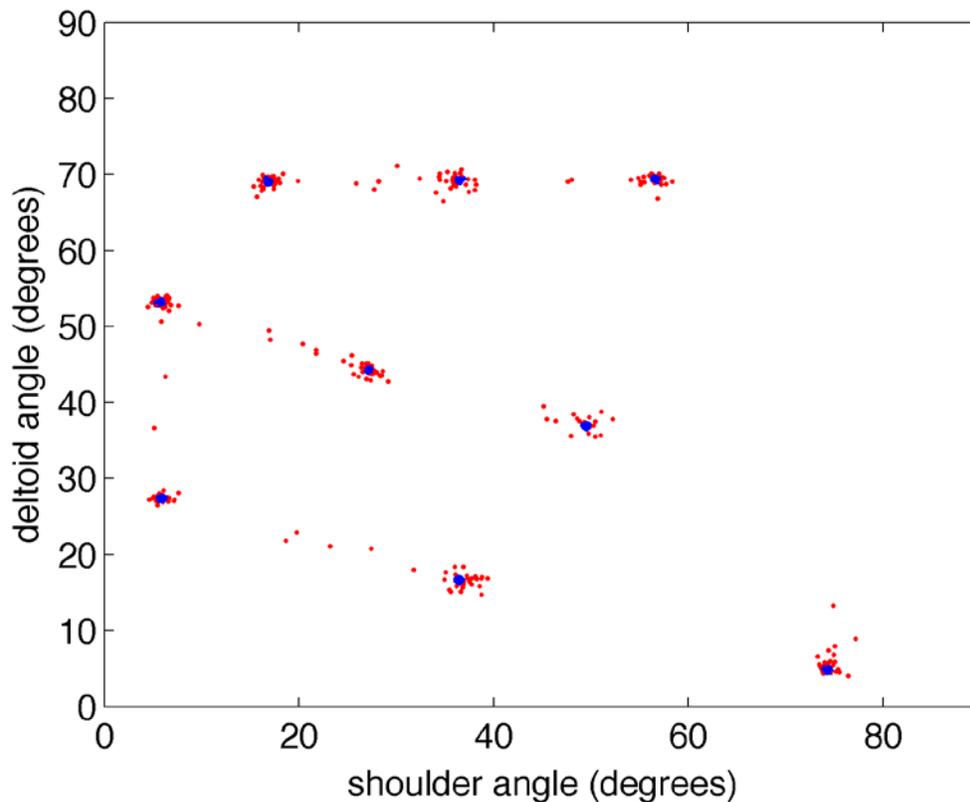


Figure 5.

4. Discussion

Our studies indicate that large-scale simulations of networks of excitatory and inhibitory spiking neurons incorporating center-annular-surround anatomy and synaptic plasticity can generate dynamically stable behavior. Such networks are versatile, as shown by their ability to form smooth maps, and they can serve as a basis for systems that learn sensorimotor coordination.

How did competitive interactions in a network of spiking neurons lead to a network that can categorize external inputs? Initial synaptic strengths were randomly distributed, so neurons were not tuned to specific stimuli. For any particular pattern of input, some local population of neurons will, by chance, be slightly more responsive than alternative groups, and active neuronal groups will suppress activity in surrounding neurons. The operation of STDP then acts to increase the synaptic drive from that input pattern of

activity. In addition, STDP and synaptic normalization force heterosynaptic reduction in the strength of synapses from uncorrelated input patterns.

The model networks described in this paper rely upon the presence of short-range excitation and long-range inhibition. This result is consistent with recent theoretical arguments that long-range inhibitory interactions are critical for cortical map formation (Kaschube et al, 2010). Among the three different connectivity topologies we analyzed, it was expected that the standard center-surround architecture would have also produced WTA network behavior (Dayan and Abbot, 2001). However, only the CAS inhibition motif gave rise to the generation of localized persistent activity that characterizes a WTA state. It is possible that connection architectures other than the ones we tried might produce WTA behavior. Although we did explore the parameter space for the standard center-surround model as we did for the CAS model in Figure 2a, it is also possible that even this connectivity might work under different parameter settings. It may prove informative to further explore analytically and empirically why the center-surround inhibition failed to produce WTA behavior in our simulations, and why the CAS architecture produced robust WTA behavior under these same conditions.

We have demonstrated the establishment of a mapping between two maps given spiking input from the real-world. The work of Davison and Frégnac (2006) demonstrated that STDP could be used to establish a mapping between two areas with correlated spiking activity. We confirm that this finding holds in a real-world task which, in our large-scale visuomotor model with approximately 7,000 spiking neurons, was able to learn a mapping from visual targets to motor actions in a BBD. Since STDP requires consistent firing of presynaptic before postsynaptic neurons to potentiate synaptic efficacies, one might not expect that STDP would strengthen synapses from the visual to the motor area, given that motor commands occur well before any visual feedback from the arm movement occurs. However at high firing rates STDP is purely facilitatory, so that all that was required to learn the mapping between visual and motor areas was a brief overlap between the time of bursts of spikes in the two areas. This was accomplished by maintaining the BBD in each posture long enough to assure that both motor area and visual area achieved equilibrium.

In the simulated network reported here at least one type of inhibitory neuron strongly inhibits an annulus in its surround while not inhibiting nearby neurons. This differs from computational models in which inhibitory connection profiles have a Gaussian distribution with the strongest inhibition occurring within the neighboring region (Laing and Chow, 2001; Dayan and Abbot, 2001). Such models are capable of WTA behavior because strong local excitation is greater than local inhibition, essentially removing that local inhibition. In our spiking model, however, we did not obtain WTA behavior with strong local inhibition. This may relate to the previous finding that spike synchronization can prevent competition in networks of spiking neurons (Lumer, 2000). Our simulations are in agreement with this finding (see for example, figure 2D). In addition, we have shown that WTA behavior can arise in large-scale spiking networks even in the presence of strong initial synchronization, if inhibitory neurons inhibit in an annular surrounding region rather than locally. We have found that WTA behavior still emerges in our CAS

network when we shrink the inner radius to zero indicating that some level of local inhibition may be tolerated (data not shown).

Our model with annular surround inhibition also appears to conflict with anatomical connections observed among certain inhibitory cells within the cortex. Reports of high connection probabilities between nearby basket and inhibitory neurons (Holmgren et al, 2003) and the finding that small basket cells tend to project little more than 100 microns from the cell body seem at odds with our model. However, local connections from small basket cells may perform a different role than do large basket cells that project up to one mm from their cell bodies, and that have been reported to mediate lateral inhibition in cortical networks (Crook et al, 1998). Regardless of the mechanisms, our simulations lead to the testable prediction that inhibition should be stronger in some annular region surrounding inhibitory neurons than it is within the local region from which it receives its excitatory inputs.

The behavior of our simulations demonstrates the versatility of networks of simulated spiking neurons endowed with CAS connectivity and activity-dependent synaptic plasticity. Further analyses of such simulations will undoubtedly prove to be a valuable tool leading to an understanding of brain function. They may also form a useful basis for more sophisticated brain-based devices, and for further theoretical studies of increasingly realistic brain networks.

Acknowledgements

We thank Joseph A. Gally for numerous helpful comments and suggestions on the manuscript. Our Engineer Donald Hudson built the custom BBD for us. Engineer Richard Martin helped in developing the computer interface to the BBD, and carrying out the BBD experiments. This work was supported in part by DARPA through ONR Grant N00014-08-1-0728 and by AFRL Cooperative Agreement FA8750-11-2-0255 to Neurosciences Research Foundation. For support of late developments we are grateful to the Mathers Charitable Foundation. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views, opinions, and/or findings contained in this article/presentation are those of the authors/presenters and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Air Force Research Laboratory, the Department of Defense, or the U. S. Government.

References

- Ben-Ari, Y., et al. (2012) Refuting the challenges of the developmental shift of polarity of GABA actions: GABA more exciting than ever! *Front Cell Neurosci.* 6:35.
- Bi, G.Q., Poo, M.M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464-10472.
- Binshtok, A.M., Fleidervish I.A., Sprengel, R., Gutnick, M.J. (2006). NMDA receptors in layer 4 spiny stellate cells of the mouse barrel cortex contain the NR2C subunit. *J Neurosci* 26(2):708-15.

- Choe, Y., Miikkulainen, R. (2004). Contour integration and segmentation in a self-organizing map of spiking neurons. *Biological Cybernetics*, 90:75-88.
- Crook, J.M., Kisvárdy, Z.F., Eysel, U.T. (1998). Evidence for a contribution of lateral inhibition to orientation tuning and direction selectivity in cat visual cortex: reversible inactivation of functionally characterized sites combined with neuroanatomical tracing techniques. *Eur. J. Neurosci.*10(6):2056-75.
- Davison, A.P., Frégnac, Y. (2006). Learning cross-modal spatial transformations through spike timing-dependent plasticity. *J. Neurosci.* 26(21):5604-15.
- Dayan, P., and Abbot, L.F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, Massachusetts, pp. 255-256.
- Derdikman, D., Hildesheim, R., Ahissar, E., Arieli, A., and Grinvald, A. (2003). Imaging spatiotemporal dynamics of surround inhibition in the barrels somatosensory cortex. *J. Neurosci.* 23(8):3100-5.
- Douglas, R.J., and Martin, K.A.C. (2004). Neural circuits of the neocortex. *Ann. Rev. of Neurosci.*, 27:419-51.
- Fino, E., Yuste, R. (2011). Dense inhibitory connectivity in neocortex. *Neuron*. 69(6):1188-203.
- Georgopoulos, A.P., Schwartz, A.B., Kettner, and R.E. (1986). Neuronal population coding of movement direction. *Science*. 233(4771):1416-9.
- Goldman-Rakic, P.S. (1995). Cellular basis of working memory. *Neuron*. 14(3):477-85.
- Haider, B., Krause, M.R., Duque, A., Yu, Y., Touryan, J., Mazer, J.A., and McCormick, D.A. (2010). Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron*. 65(1):107-21.
- Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J. Physiol.* 551(Pt 1):139-153.
- Itti, L., Koch, C., Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 20(11), 1254-1259.
- Izhikevich, E.M. (2003) Simple Model of Spiking Neurons. *IEEE Transactions on Neural Networks* 14:1569-1572.
- Izhikevich, E.M., and Edelman, G.M. (2008) Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. U. S. A.* 105(9):3593-3598.
- Izhikevich, E.M. (2010). Hybrid spiking models. *Philos Transact A Math Phys Eng Sci.* 368(1930):5061-70.
- Kaschube, M., Schnabel, M., Löwel, S., Coppola, D.M., White, L.E., and Wolf, F. (2010). Universality in the evolution of orientation columns in the visual cortex. *Science* 330(6007):1113-6.
- Kernighan, B.W., Ritchie, D.M. (1988). *The C programming Language*, Prentice-Hall, Upper Saddle River, NJ.
- Kisvárdy, Z.F., Crook, J.M., Buzás, P., and Eysel, U.T. (2000). Combined physiological-anatomical approaches to study lateral inhibition. *J. Neurosci. Methods.* 103(1):91-106.
- Kohonen, T. (1984). *Self-Organization and Associative Memory*. Springer-Verlag, Berlin.
- Laing, C.R., Chow, C.C. (2001). Stationary bumps in networks of spiking neurons. *Neural Comput.* 13(7):1473-94.

- Lumer, E.D. (2000). Effects of spike timing on winner-take-all competition in model cortical circuits. *Neural Comput.* 12(1):181-94.
- Malsburg, Ch. Von der (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14, 85-100.
- Myme C.I., Sugino K., Turrigiano G.G., Nelson S.B. (2003). The NMDA-to-AMPA ratio at synapses onto layer 2/3 pyramidal neurons is conserved across prefrontal and visual cortices. *J Neurophysiol.* 90(2):771-9.
- Obermayer. K., Ritter, H., Schulten, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proc. Natl. Acad. Sci. U. S. A.* 87(21):8345-9.
- Oster, M., Douglas, R., Liu, S. (2009). Computation with Spikes in a Winner-Take-All Network. *Neural Comput.* 21, 2437-2465.
- Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proc. Natl. Acad. Sci. U.S.A.* 108, 5419–5424.
- Rutishauser, U., Douglas, R.J., Slotine, J. (2011). Collective stability of networks of winner-take-all circuits. *Neural Computation.* 23:735-773.
- Shriki, O., Hansel, D., and Sompolinsky, H. (2003). Rate models for conductance-based cortical neuronal networks. *Neural Computation* 15:1809-1841.
- Willmore B., Tolhurst D.J. (2001). Characterizing the sparseness of neural codes. *Network* 12(3):255-70.
- Wilson, H.R., Cowan, J.D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13, 55-80.

Figure 1. The Center-Annular-Surround (CAS) spiking network architecture leads to Winner-Take-All (WTA) dynamics. (A) The CAS network architecture consists of interconnected spiking neurons, excitatory (green ovals) and inhibitory (red ovals). Each population is arranged in a two-dimensional grid. Connections from representative cells are illustrated. Axons from excitatory neurons (green arrows) project to neurons within green areas. Axons from inhibitory neurons (red arrows) project to neurons in the transparent red annular areas. The sensory input projecting nontopographically to both the excitatory and inhibitory “cortical” populations is not shown. (B) The CAS connectivity leads to WTA dynamics: small areas of high activity are surrounded by large regions with little activity. The firing rates of excitatory neurons in the network are shown as pixels with brightness proportional to firing-rate indicated by the scale bar to the right (in Hz). The number and size of the winning regions are functions of a variety of network parameters.

Figure 2. WTA dynamics can occur in large regions of the parameter space of CAS networks. (A) A measure of winner-take-all behavior in a network is plotted as a function of synaptic weights coupling the excitatory and inhibitory neural populations. The measure we use is the highest firing rate of any neuron in the network, subject to a sparseness constraint that at least half of the neurons in the network are firing at less than 2 Hz; otherwise the measure is defined to be zero. The total synaptic conductance in nano-siemens (nS) (Izhikevich and Edelman, 2008) in each individual inhibitory neuron from excitatory neurons is on the y-axis, and total inhibitory conductance received by each neuron, excitatory or inhibitory, is on the x-axis. The orange and red areas indicate regions of the parameter space in which the network exhibits WTA behavior. The lower left region of the parameter space, labeled “Epileptic”, defines networks exhibiting epileptic dynamics in which all neurons fire indiscriminately to the stimulus. (B) to (D) are raster plots which show all spikes (blue dots) during the third second of the simulation for each excitatory neuron in the network. (B) All spikes of a network in a WTA state at parameters labeled “Winner-Take-All in Figure 2A. During this state some excitatory neurons (horizontal band of blue dots) fire persistently in response to a constant stimulus while others are silent. (C) At certain values of parameters, labeled “Traveling Wave” in Figure 2A, region of the parameter space, the network exhibits moving patches of activity instead of the stable patches shown in (B); this results in diagonal bands in the one-dimensional raster plot. (D) Occasionally the network requires more than two simulated seconds for a winner to emerge. These “Transient Rhythmic” states result in all neurons firing synchronously and rhythmically for some time before a winning group emerges. See Supplementary Fig. 1-3 for close up plots from Figures 2B-D. A raster plot corresponding to the epileptic activity state is not shown.

Figure 3. A simulated neural network develops a smooth orientation map similar to those of cat and primate visual cortex. The map shows the preferred orientation of individual excitatory neurons arrayed in a 60 X 60 neuron grid. Pixel colors relate location of each neuron to its preferred orientation as indicated by the color bar at right of the map. Adjacent neurons in the network tend to have similar orientation preferences. Brightness varies with orientation selectivity (dark = low selectivity, bright = high selectivity); dark areas are found at the centers of so-called orientation pinwheels.

Figure 4. STDP plus synaptic scaling forms a mapping between visual and motor maps. CAS networks were used in a humanoid BBD to demonstrate that such a system could learn sensorimotor coordination. CAS networks consist of populations of excitatory (E) and inhibitory (I) neurons synaptically coupled as described in the text. Visual input from the video camera provided patterned input to “thalamic” (T) neurons of the visual area (V), while the output of excitatory neurons in the motor area (M) were used to control the two shoulder joints of the left arm. After repeatedly stimulating the motor area in one of nine different locations, and thus moving the arm to one of nine different postures, a mapping formed from the visual area responses to the location of the hand to the motor area output that drove the hand to those locations.

Figure 5. The BBD reaches accurately towards visual targets after training. During the testing period, the arm consistently moved in response to the visual stimulus. To demonstrate the accuracy of the movements, the joint angles of the commanded movements made during training (blue) and testing (red) are plotted in two-dimensional joint angle space every 200 ms for all 5 subjects. Note that the joint angles achieved during testing cluster around those achieved in the training period showing the accuracy of the visually guided, learned movements.

Supplementary Material.

Neuron parameters.

Table I. Neuron parameters.

Neuron type	Area	C	k	V_r	V_t	V_{peak}	a	b	c	d
Excitatory	V	80	3	-60	-50	50	0.01	5	-60	10
Inhibitory	V	20	1	-55	-40	25	0.15	8	-55	200
Excitatory	Motor	100	0.7	-60	-50	0	0.03	-2	-60	100
Inhibitory	Motor	20	1	-55	-40	25	0.15	8	-55	200
Thalamic	Input	200	1.6	-60	-50	40	0.01	15	-60	10

Anatomy.

The connectivity between model neurons fell into two classes: either local-type or surround-type. For local-type connectivity, a two-dimensional Gaussian probability distribution, centered on each postsynaptic cell, determines the probability of forming a synapse between each potential presynaptic neuron within a specified maximum distance, r_{max}

$$f(d) = ae^{-\frac{(d-\mu)^2}{2\sigma^2}}$$

where a is a scale factor set to generate, on average, a target number of synapses on each postsynaptic cell, d is the distance between the presynaptic neuron and the postsynaptic neuron, μ is 0, and σ is the standard deviation. In a similar manner, a two-dimensional Gaussian function was also used to specify the synaptic strength between connected neurons as a function of the distance between them in the network. The total of all synaptic efficacies was scaled to sum to a constant parameter with units in nanoSiemens (nS). Thus both connection probability and strength were maximal between nearest neighbors, and fell off as a function of distance, controlled by the same parameter, the standard deviation of a Gaussian.

For surround-type connectivity, a postsynaptic neuron receives synaptic connections from neurons located in a surrounding annular region specified by a minimum (r_{\min}) and maximum (r_{\max}) radial distance from the postsynaptic cell. (This is equivalent to saying that each presynaptic neuron sends projections to postsynaptic neurons in an annular region). The probability of forming a connection with a neuron in the annulus is determined as a function of distance from the postsynaptic cell. The function used is a Gaussian with standard deviation σ , centered at $\mu = (r_{\min} + r_{\max})/2$. Thus a postsynaptic neuron connects with no neurons in the center of the annulus, has minimal connection probability at the minimum radius, increasing to the maximum probability half-way between the inner and outer radius, and falling off once again with increasing distance up to the outer radius, beyond which the connection probability is forced to zero. This probability distribution function is scaled to create a target number of synapses for each postsynaptic neuron. The synaptic strengths for the surround-type connection are also initialized using the same function, with the same parameters. However, the sum of all synaptic strengths of this type was scaled to make the total equal to a constant value under experimenter control.

In order to avoid boundary conditions in the network, the network was treated as a torus. Thus connections from neurons that would go outside of the network instead “wrap around” to connect with neurons on the opposite edge.

Table II shows the parameters defining the anatomy and synaptic parameters of the CAS network used in the parameter space analysis in Results section 1. The table defines two types of information for every neural area: the neuron composition, and the synaptic connectivity for each neuron type. The first four columns of the table list, for each separate neural population in the simulation, the type of neuron, the area in which it is located, the number of neurons in the population, and the total number of synapses per neuron.

The remaining columns define the connectivity for each type of neuron in the area. Multiple rows are necessary to define the connectivity for each postsynaptic type; one row is needed for each presynaptic neuron type forming synapses on the postsynaptic neurons. Pre-area and pre-type specify the presynaptic area and type of the neuronal

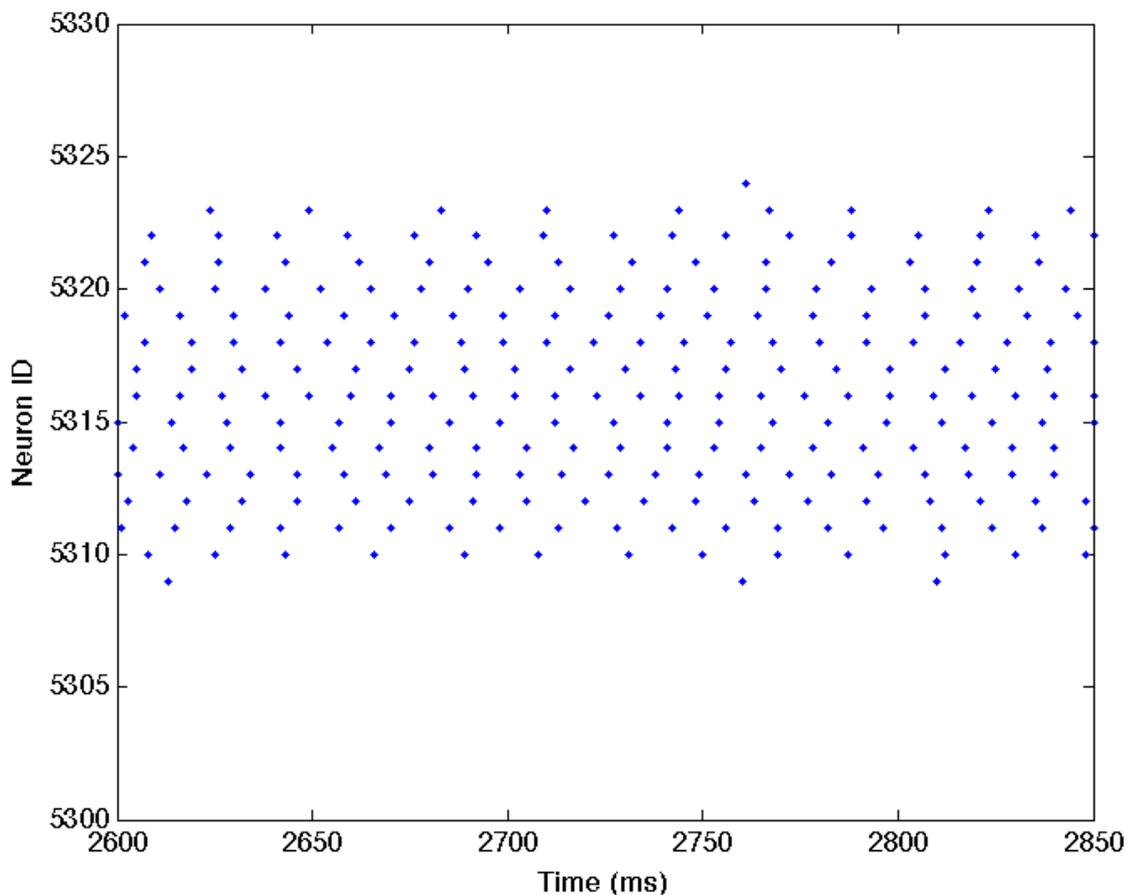
group projecting to the postsynaptic group. The next column specifies the percentage of the postsynaptic cell's synapses allocated to this pathway. The remaining columns provide all of the parameters used to specify details of the synaptic pathways as described in the paragraphs above.

Table III shows the parameters defining the anatomy and synaptic parameters of the orientation selective map experiment in the Results section 2; the format is the same as that for Table II.

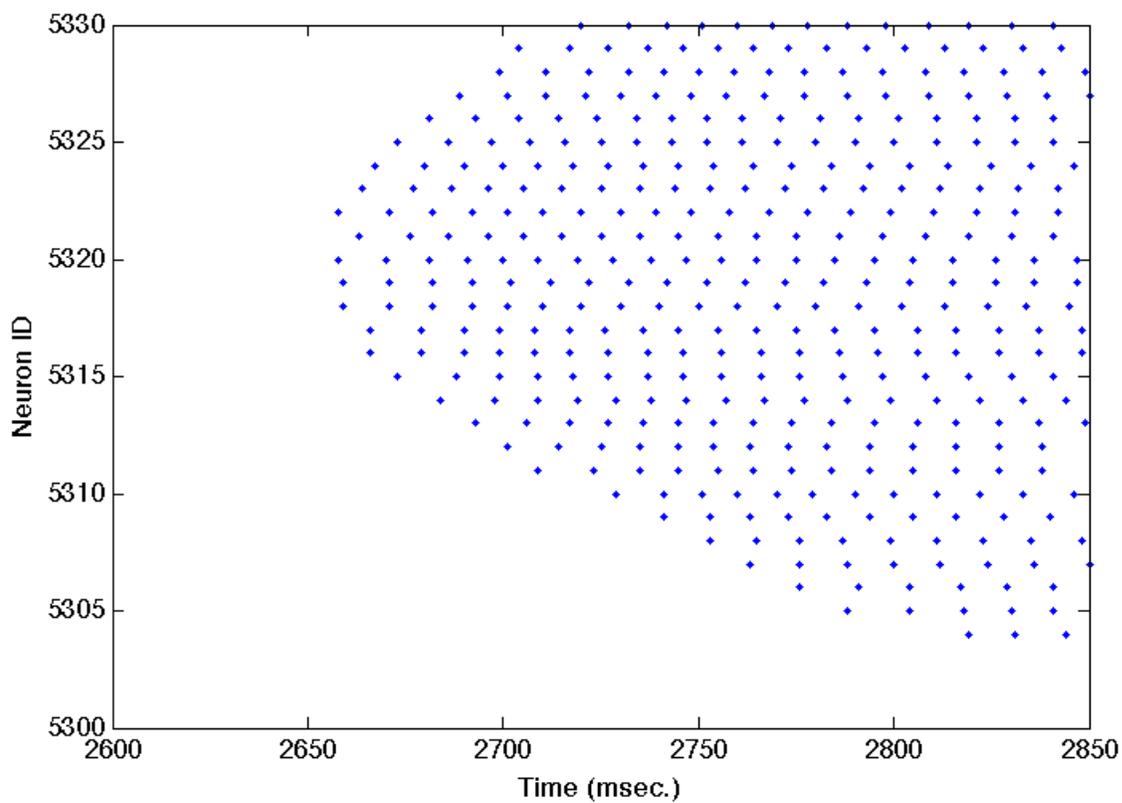
Table IV shows the parameters defining the anatomy and synaptic parameters of the visuomotor coordination network used to control the arm of The BBD in the Results section 3; the format is the same as that for Table II.

Close up of a portion of Figures 2B-D.

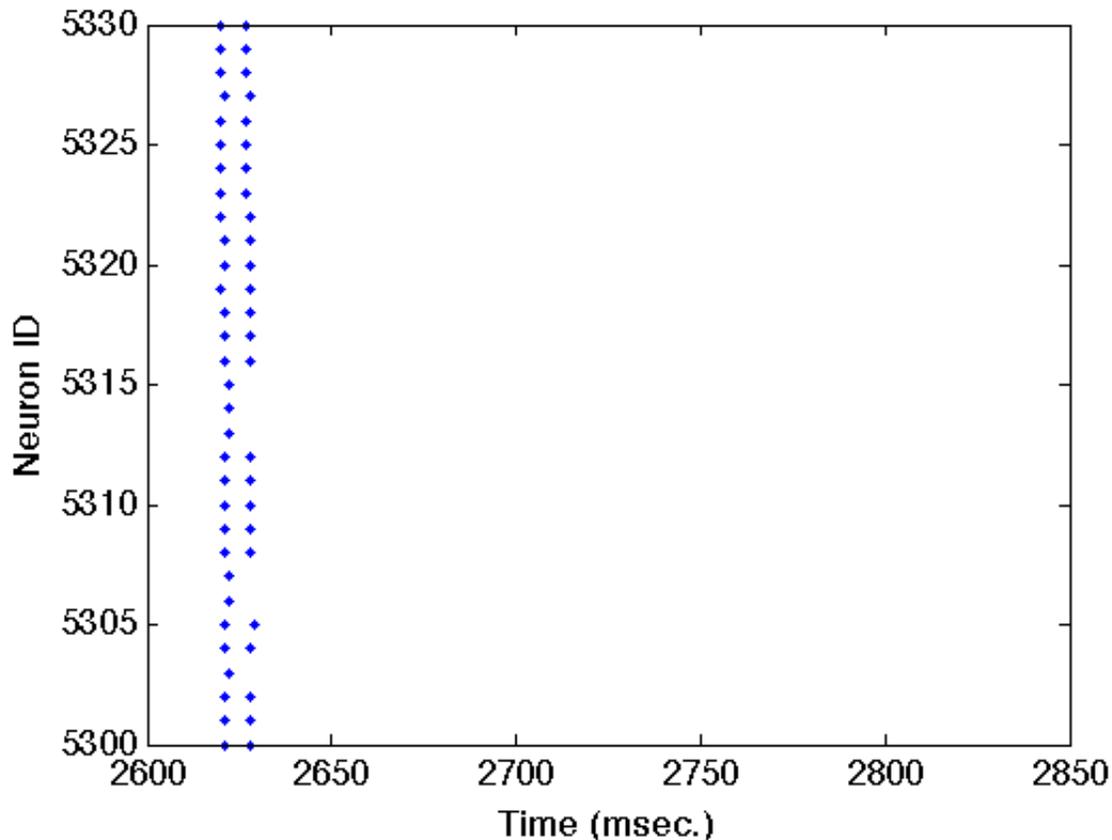
Supplementary Figures 1-3 show close ups of figures 2B-D.



Supplementary Figure 1. Close up of Figure 2B showing individual spike trains for a small subset of excitatory neurons. See figure 2B caption for a complete description.



Supplementary Figure 2. Close up of Figure 2C showing individual spike trains for a small subset of excitatory neurons. See figure 2C caption for a complete description.



Supplementary Figure 3. Close up of Figure 2D showing individual spike trains for a small subset of excitatory neurons. See figure 2D caption for a complete description.

Control experiment parameters.

The following parameters were used in control experiments to demonstrate that the CAS architecture made an improvement in WTA behavior in our simulations.

Standard Center-Surround architecture:

In the classical center-surround topology, both excitatory and inhibitory neurons have local connection type but with different standard deviations. This connection architecture has been reported to produce WTA dynamics (see main text for references). In this control experiment, excitatory neurons connected with inhibitory neurons in a Gaussian distribution with maximal distance (r_{max}) of 0.33mm and a standard deviation of 0.16mm. Inhibitory neurons connect to excitatory neurons and themselves in a wider Gaussian distribution with maximal distance (r_{max}) at 1.44 mm and larger standard deviations ($\sigma=0.8$ mm) than excitatory connections.

Excitatory surround, inhibitory center architecture:

An additional control experiment was conducted with a connection architecture is a reverse version of our CAS topology. That is, excitatory connections are annular surround type specified by $r_{\min}=0.1\text{mm}$, $r_{\max}=1\text{mm}$ and a Gaussian distribution centered at $(r_{\min}+r_{\max})/2$ with a standard deviation of 0.3333 (sigma). Inhibitory connection, on the other hand, are local Gaussian type with $r_{\max}=0.333$ and standard deviation of 0.16.

Uniform random excitatory and inhibitory architecture. In this control experiment, excitatory and inhibitory neurons have an equal probability of connecting to any other excitatory or inhibitory neuron. This is implemented as a local connection in which the maximal connection distance is set to cover the entire area ($r_{\max}=1.44\text{ mm.}$, $r_{\min}=0\text{ mm.}$) and standard deviation of the Gaussian distribution used to generate the connection probability is large enough to approximate a uniform distribution (sigma = 10 mm.).

Visual input to the BBD.

Video was recorded with an Axis 207MW wifi camera. Black and white images with a resolution of 320x240 were transmitted at 30fps. The central portion of the video frames were used as input to a two-dimensional grid of on-center Retinal Ganglion Cells (RGC). The grid size was 21x21 neurons with a center area size of 3x3 and the surround area of 6x6 neurons. Each RGC receives a current that is computed following the algorithm of Wohrer and Kornprobst (2009). These currents were constantly injected at each integration step until the next video frame was received. RGCs were modeled with the Izhikevich model (Izhikevich and Edelman, 2008) with the following parameters: $C=100$, $V_r=-70\text{mV}$, $V_t=-50\text{mV}$, $k=1$, $a=0.005$, $b=0$, $c=-75\text{mV}$, $d=250$, and $V_{\text{peak}}=10\text{mV}$.

Supplementary Video

The online supplementary material includes a video showing the behavior of the BBD during testing, after it has been trained to reach to 4 visual locations.

References

Izhikevich EM and Edelman GM (2008). Proc Natl Acad Sci U S A. 105(9):3593-3598.
Wohrer A, Kornprobst P (2009). Virtual Retina: a biological retina model and simulator, with contrast gain control. *Journal of Computational Neuroscience* 26(2), 219–249.

Table II. Anatomical and synaptic parameters for the CAS network used in the parameter space analysis in results section 1.

Postsynaptic neuron type	Postsynaptic area	Number of neurons	Average synapses per neuron	Presynaptic area	Presynaptic neuron type	Percentage of total synapses	r_{min} (mm)	r_{max} (mm)	σ (mm)	ϵ	S_{total} (nS)	S_{max} (nS)	Initial α	Final α	Learning start time (ms)	Learning end time (ms)	nmda_gain	gabab_gain
Exc.	V	3481	3520	V	Exc.	12.5	-	0.1	0.05	0	22	10	0	0	0	0	0.5	0.1
				V	Inh.	25	0.1	1	0.8	0	0 to 1600	20	0	0	0	0	0.5	0.1
				Input	Thalamic	62.5	-	1.44	2.5	0	900	50	0	0	0	0	0.5	0.1
Inh.	V	900	2000	V	Exc.	20	-	0.33	0.16	0	0 to 100	5	0	0	0	0	0.5	0.1
				V	Inh.	40	0.1	1	0.3333	0	0 to 240	15	0	0	0	0	0.5	0.1
				Input	Thalamic	40	-	4.0	10	0	10	10	0	0	0	0	0.5	0.1
Thalamic	Input	441	0	-	-	0	-	0	0	0	0	0	0	0	0	0	0.5	0.1

Table III. Anatomy and synaptic parameters for the orientation selective map experiment.

Postsynaptic neuron type	Postsynaptic area	Number of neurons	Average synapses per neuron	Presynaptic area	Presynaptic neuron type	Percentage of total synapses	r_{min} (mm)	r_{max} (mm)	σ (mm)	ϵ	S_{total} (nS)	S_{max} (nS)	Initial α	Final α	Learning start time (ms)	Learning end time (ms)	mnda_gain	gabab_gain	
Exc.	V	3600	2200	V	Exc.	20	-	0.035	0.07	0	10	2	0	0	0	0	0.5	0.1	
				V	Inh.	40	0.035	1	0.3333	0	400	10	0	0	0	0	0	0.5	0.1
				Input	Thalamic	40	-	0.65	0.25	0.5	20	10	0.01	0.01	0	10000000	0.5	0.1	
Inh.	V	900	2000	V	Exc.	20	-	0.33	0.16	0	20	10	0	0	0	0	0.5	0.1	
				V	Inh.	40	0.07	1	0.3333	0	180	10	0	0	0	0	0	0.5	0.1
				Input	Thalamic	40	-	1.44	10	0.5	30	2	0	0	0	0	0	0.5	0.1
Thalamic	Input	400	0	-	-	0	-	0	0	0	0	0	0	0	0	0.5	0.1		

Table IV. Anatomy and synaptic parameters for the visuomotor coordination network used to control the arm of the BBD.

Postsynaptic neuron type	Postsynaptic area	Number of neurons	Average synapses per neuron	Presynaptic area	Presynaptic neuron type	Percentage of total synapses	r_{min} (mm)	r_{max} (mm)	σ (mm)	ϵ	S_{total} (nS)	S_{max} (nS)	Initial α	Final α	Learning start time (ms)	Learning end time (ms)	nmda_gain	gabab_gain	
Exc.	V	3600	2200	V	Exc.	20	-	0.035	0.07	0	10	2	0	0	0	0	0.5	0.1	
				V	Inh.	40	0.035	1	0.3333	0	400	10	0	0	0	0	0	0.5	0.1
				Input	Thalamic	40	-	0.65	0.25	0.5	20	10	0.01	0.01	0	10000000	0.5	0.1	
Inh.	V	900	2000	V	Exc.	20	-	0.33	0.16	0	20	10	0	0	0	0	0.5	0.1	
				V	Inh.	40	0.07	1	0.3333	0	180	10	0	0	0	0	0	0.5	0.1
				Input	Thalamic	40	-	1.44	10	0.5	30	2	0	0	0	0	0	0.5	0.1
Thalamic	Input	400	0	-	-	0	-	0	0	0	0	0	0	0	0	0.5	0.1		

A2. McKinstry, J.L. and Edelman, G.M. (2013) “Temporal Sequence Learning in Winner-Take-All Networks of Spiking Neurons Demonstrated in a Brain-Based Device,” *Frontiers in Neurorobotics*, accepted for publication.

Temporal Sequence Learning in Winner-Take-All Networks of Spiking Neurons Demonstrated in a Brain-Based Device

ABSTRACT

Animal behavior often involves a temporally ordered sequence of actions learned from experience. Here we describe simulations of interconnected networks of spiking neurons that learn to generate patterns of activity in correct temporal order. The simulation consists of large-scale networks of thousands of excitatory and inhibitory neurons that exhibit short-term synaptic plasticity and spike-timing dependent synaptic plasticity. The neural architecture within each area is arranged to evoke winner-take-all (WTA) patterns of neural activity that persist for tens of milliseconds. In order to generate and switch between consecutive firing patterns in correct temporal order, a reentrant exchange of signals between these areas was necessary. To demonstrate the capacity of this arrangement, we used the simulation to train a brain-based device responding to visual input by autonomously generating temporal sequences of motor actions.

INTRODUCTION

A growing body of neurophysiological evidence suggests that patterns of activity in vertebrate brains observed during movement are commonly composed of temporal sequences of periods with steady-state firing rates lasting several hundred milliseconds separated by sharp transitions (Tanji, 2001; Averbek et al., 2002; Nakajima et al., 2009). This pattern of activity is also observed during sensory perception in gustatory cortex (Jones et al., 2007), and the operation of working memory (Seidemann et al., 1996). Although network models composed of mean-firing-rate neurons have been used to model sequential neural activity (Rhodes et al., 2004; Salinas, 2009; Verduzco-Flores et al., 2012), biological networks are composed of spiking neurons. Therefore understanding spiking networks with this capability requires further exploration (Liu and Buonomano, 2009; Chersi et al., 2011). Given open questions regarding the stability and robustness of networks which learn to generate sequences (Verduzco-Flores et al., 2012), testing such networks in Brain-Based-Devices (BBD) is warranted (Edelman, 2007; McKinstry et al., 2008).

In this paper we describe how our previous models of Winner-Take-All (WTA) spiking networks (Chen et al., 2013) can be coupled together and trained to generate segmented and sequential neural activity (See Rutishauser and Douglas, 2009 for a mean-firing rate WTA network that generates sequences). The neural system is composed of thousands of simulated biologically realistic excitatory and inhibitory spiking neurons. The single compartment neurons modeled in these simulations display voltage dynamics similar to those seen in cortical neurons. Activity of the simulated neurons reflects the conductance of ion channels in the model including: AMPA, NMDA, GABA_A and GABA_B (Izhikevich and Edelman, 2008). Model synapses were subject to short-term synaptic plasticity (Zucker,

1989). Spike-timing dependent plasticity (STDP) modeled long-term synaptic changes that allowed the system to learn temporal sequences.

We found that networks composed of spiking neurons of this sort, when trained to respond to repeated sequences of sensory cues, generate temporally ordered patterns of neuronal activity consisting of brief steady states separated by sharp transitions that resemble those observed in functioning brains. We found that the present model could be used to control specific motor sequences in a brain-based device. The population activity pattern in this modeled neuronal system has similarities to those observed in primate prefrontal cortex during multi-segmented limb movements (Averbeck et al., 2002).

MATERIALS AND METHODS

Spiking Neuronal Networks - Each modeled network (figure 1A) is comprised of up to three interconnected populations of spiking neuronal units (Izhikevich, 2010) distributed over two-dimensional square grids. Each population is composed of units simulating one of three functional classes of spiking neurons: input, excitatory, and inhibitory. The parameters of simulated neurons in each class are tuned so that the voltage waveform mimics its biological counterpart. The synapses display STDP and short-term plasticity dynamics as previously described in detail (Izhikevich and Edelman, 2008). The neuron model equations, short-term synaptic plasticity equations, and STDP equations are presented after a description of the network connectivity.

Neuronal Network Architecture – Each of the three major structural and functional components of the modeled nervous system (figure 1B) consisted of a network of spiking neuronal units (Izhikevich) distributed over a two-dimensional (2mm by 2mm) grid. The networks function as analogs of a thalamic nucleus (Input area), and two interconnected cortical areas, (Area A and Area B).

The Input network contained 484 simulated neurons providing topographic excitatory input to Area A. Current levels to cells of the Input area were adjusted by trial and error to assure that the network responded to abstract patterns or video camera input by generating distinct response patterns of neuronal activity with a maximum firing rate of approximately 100 Hz. The Area A and Area B networks were each made up of 1600 excitatory cells as well as 400 inhibitory cells having fast-spiking behavior.

Areas A and B had similar connectivity. Each was composed of a Center-Annular-Surround (CAS) network, a variant of center-surround networks, that we have found (Chen et al., 2013) to effectively generate WTA dynamics (Dayan and Abbott, 2001) in large-scale networks of spiking neurons. Any distinctive pattern of neural activity in the input area evoked enhanced neural activity within a few localized patches in both areas A and B. This CAS network architecture is illustrated in figure 1A. Connectivity between the model neurons fell into two

classes: either local-type or surround-type. Local-type connections are between nearby neighbors, whereas surround-type connections come from neighbors in a surrounding annular region. Excitatory cells receive both local-type projections from excitatory cells and surround-type inhibitory projections (figure 1A). Inhibitory cells also received local-type projections from the excitatory cells and surround-type input from other inhibitory cells. The CAS connectivity confers WTA properties to both areas A and B. A complete description of all connectivity parameters is provided in the supplementary material.

To create a network capable of storing and generating sequences of neural activity, we added reentrant connections between Areas A and B in the following way. In Area A (figure 1B), both excitatory and inhibitory cells also receive simulated feed-forward input that was approximately all-to-all. Area B neurons, on the other hand, do not receive connections from the input. Instead, they receive non-plastic, local-type input that is topographic from Area A. Area B excitatory neurons project back to Area A with plastic and widespread surround-type connectivity. Synaptic changes resulting from STDP at these connections form a link between temporally adjacent patterns of neural activity within the sequence. These excitatory reentrant connections from Area B to Area A are widespread and cover most of the region since each activity pattern in Area B has two bumps, similar to the activity pattern shown in Supplementary Figure 5. This widespread connectivity enables the network to learn to associate arbitrary temporally adjacent patterns. This was useful for the BBD experiment, since the patterns that emerged within Area A during the initial training phase were not under experimenter control.

Neuronal Dynamics – Spiking dynamics of each neuron were simulated using the phenomenological model proposed by Izhikevich (2003). The model has only 2 equations and 4 dimensionless parameters that could be explicitly found from neuronal resting potential, input resistance, rheobase current, and other measurable characteristics. We present the model in a dimensional form so that the membrane potential is in millivolts, the current is in picoamperes and the time is in milliseconds:

$$C\dot{v} = k(v - v_r)(v - v_t) - u - I_{syn} \quad (1)$$

$$\dot{u} = a\{b(v - v_r) - u\} \quad (2)$$

where C is the membrane capacitance (in picofarads (pF)), v is the membrane potential (in mV), v_r is the resting potential (in mV), v_t is the instantaneous threshold potential (in mV), u is the recovery variable (the difference of all inward and outward voltage-gated currents in pA), I_{syn} is the synaptic current (in pA) defined below, a and b are unitless parameters. When the membrane potential reaches the peak of the spike, i.e., $v > v_{peak}$, the model is said to fire a spike, and all variables are reset according to $v = v_{peak} - c$ and $u = u_{peak} - d$, where c (mV) and d (pA) are parameters. Supplementary Table I lists each of the neuron model parameters used in all experiments. At the start of all simulations, v was set to -

60 for all neurons, whereas u was set to a different random value for each neuron, drawn uniformly from the range 0 to 100.

Short-Term Synaptic Plasticity – The strength of synapses varied as a function of the presynaptic neuron's recent firing history independent of long-term synaptic changes as found in biological synapses (Zucker, 1989). We assume that the synaptic conductance (strength) of each synapse can be scaled down (depression) or up (facilitation) on a short time scale (hundreds of milliseconds) by a scalar factor x . This scalar factor, different for each presynaptic cell, is modeled by the following one-dimensional equation

$$\dot{x} = (1 - x) / \tau_x, \quad x \leftarrow px \text{ when presynaptic neuron fires. (3)}$$

x tends to recover to the equilibrium value $x = 1$ with the time constant τ_x (in ms), and it is reset by each spike of the presynaptic cell to the new value px . Any value $p < 1$ decreases x and results in short-term synaptic depression, whereas $p > 1$ results in short-term synaptic facilitation. The parameters, τ_x , in ms, and scale factor p , for each combination of presynaptic and postsynaptic neuron type were as follows: exc. \rightarrow exc.: 150, 0.8; exc. \rightarrow inh.: 150, 0.8; inh. \rightarrow exc.: 150, 0.8; inh. \rightarrow inh.: 150, 0.8; thalamic \rightarrow exc: 150, 0.7; thalamic \rightarrow inh.: 200, 0.5.

Synaptic Kinetics – The total synaptic current to each neuron is simulated as

$$I_{syn} = g_{AMPA}(v - 0) + g_{NMDA} \frac{[(v + 80)/60]^2}{1 + [(v + 80)/60]^2}(v - 0) + g_{NMDA_{VI}} \frac{[(v + 100)/60]^2}{1 + [(v + 100)/60]^2}(v - 0) + g_{GABA_A}(v + 70) + g_{GABA_B}(v + 90) \quad (4)$$

where v is the postsynaptic membrane potential, and the subscript indicates the receptor type. Each millisecond, each synaptic conductance is updated according to equation (5).

$$g_r(t) = \begin{cases} g_r(t-1) - g_r(t-1)/\tau_r + gain_r x s(t-1) & \text{when the presynaptic neuron fires} \\ g_r(t-1) - g_r(t-1)/\tau_r & \text{otherwise} \end{cases} \quad (5)$$

where subscript r indicates the receptor type, $\tau_r = 5, 150, 6,$ and 150 ms for the simulated AMPA, NMDA, GABA_A and GABA_B receptors, respectively. The voltage-independent NMDA channel (NMDA_{VI}) is based loosely on the type of channel found between excitatory cells in layer 4 of visual cortex (Binshtok, 2006); we used $\tau_r = 150$ ms for this simulated receptor as well. $s(t)$ is the synaptic weight at time t . x is the short-term depression/potential scaling factor as above; $gain_{NMDA}$ is the ratio of NMDA to AMPA conductance and is found experimentally to be less than one (Myme et al., 2003). Similarly, $gain_{GABAB}$ is the ratio of GABAB to GABAA receptors. The values of $gain_{AMPA}$ and $gain_{GABAA}$ were

always one. The values of $gain_{NMDA}$ and $gain_{GABAB}$ used in the simulations are shown in Tables II and III.

STDP – The long-term change in conductance (weight) of each synapse in the model is simulated according to spike-timing-dependent plasticity (STDP): the synapse is potentiated or depressed depending on the order of firing of the presynaptic and post-synaptic neurons (Bi and Poo, 1998). We use the following equations to update the state of each plastic synapse, $s(t)$, in the network every millisecond:

$$y(t) = y(t-1) - y(t-1)/\tau_c + \alpha STDP(t) \delta(t - t_{pre/post}) \quad (6)$$

$$s(t) = \begin{cases} s(t-1) + y(t) & \text{if } \text{mod}(t, 50) = 0 \\ s(t-1) & \text{otherwise} \end{cases} \quad (7)$$

where $\delta(t)$ is the Dirac delta function that step-increases the variable y . Firings of pre- and postsynaptic neurons, occurring at times t_{pre}, t_{post} , respectively, change y by the amount $\alpha STDP(t)$ where α is the learning rate for the synapse, $t = t_{post} - t_{pre}$ is the interspike interval, and

$$STDP(t) = \begin{cases} A^+ \exp(-1/\tau^+)^t, & t > 0 \\ A^- \exp(-1/\tau^-)^{|t|}, & t \leq 0 \end{cases}. \quad (8)$$

where $A^+ = 0.005$, $A^- = 0.001$, $\tau^+ = \tau^- = 20$ ms. The variable c decays to zero exponentially with the time constant $\tau_c = 1000$ ms. Each synapse is updated only once every 50 ms for computational efficiency. Note that for simplicity, each synapse was modeled with a single weight, s ; therefore the STDP rule changed both AMPA and NMDA components of the synapse proportionally. In addition, each synapse was prevented from exceeding s_{max} or going below zero, regardless of learning rules and normalization (see synaptic scaling). Values of s_{max} for each connection pathway are provided in Supplementary Tables II and III.

Synaptic scaling – Synaptic strengths at time t , $s_i^j(t)$, were scaled for each synapse i , in order to maintain the total of all synaptic strengths on a given connection pathway to neuron j , s_{total} , at a constant value:

$$s_i^j(t) = s_i^j(t-1) \left(\frac{s_{total}}{\sum_{k=1}^{n_j} s_k^j(t-1)} \right)$$

where n_j is the number of synapses on the connection pathway to neuron j . This scaling was performed for every neuron each time the synapses were updated with equation (7). Values of s_{total} for each connection pathway are provided in Supplementary Tables II and III.

Data analysis – To evaluate how accurately the network regenerated individual activity patterns within a sequence, we calculated the similarity between the

network response to each individual segment (pattern), and the population response during sequence training and recall. To measure similarity between two neural activity patterns in a given population at two different times, t1 and t2, the following steps were performed. The mean firing rate of each neuron in the population at time t1 was calculated within some small window, yielding a number for each neuron; this list of numbers formed a vector, **f1**. The same was done at time t2, yielding vector **f2**. A match score was computed between the two population vectors by taking the normalized dot-product as follows:

$$match = \frac{\sum_{i=1}^n f1_i \cdot f2_i}{\|f1\| \|f2\|}$$

where n is the number of neurons in the population, and $\|x\|$ computes the length of the vector \bar{x} . This match score provides a measure of similarity where one is a perfect match, and zero is a complete mismatch.

The mean firing rate of each Area A excitatory neuron in response to each input stimulus in the sequence was recorded during the first epoch of sequence training during which there was no overlap in the input patterns presented or in the corresponding network responses to those patterns. Subsequently, during sequence training and free recall phases, these templates were used to quantify how closely an observed pattern of neural activity resembled each individual segment of a sequence. To do this, the mean firing rate vector of Area A neurons was computed every 50 msec of sequential behavior. A match score was then calculated between each of the sub-pattern templates and the template of each 50-msec population firing rate segment. This method can detect whether ongoing spiking activity reflects multiple sub-patterns of a sequence at the same time. It makes no assumptions about the time-course of sequence generation.

Brain Based Device – To investigate a simulated nervous system in a real-world device, we designed and constructed a humanoid BBD (figure 2). This device is approximately 20 inches high and uses a black and white wireless webcam for vision. Each arm contains eight Dynamixel motors (Robotis, Irvine, CA, USA). In the experiments described here only the two shoulder motors function; all other joints remain stationary with the arm extended. Shoulder joint angles provided by the motors determine the posture of the arms. A miniature PC (VIA Technologies, Fremont, USA) mounted on the back of the BBD maintained wireless communication between the device and the spiking neuronal networks simulated on a Mac-Pro (Apple, Inc. Cupertino, CA). The robot operated approximately three times slower than real-time during experiments.

To test the sequence generation network in the BBD, a motor area in addition to area A and B was added to enable the system to generate motor sequences (see Supplementary Figure 1). This network was the same size as the excitatory-inhibitory networks in Areas A and B, with 1,600 excitatory, and 400 inhibitory spiking neurons, and had similar parameters as well. Different patterns of

spiking of motor area neurons specified distinct equilibrium postures of the left arm using population vector coding as described in the Supplementary Material. Since the video camera was aimed at the robotic hand and remained fixed during the experiments, each of these postures, in turn, evoked a distinct pattern of visual input to the system. The motor region received non-topographic connections from the output of Area B in the sequence generation network. These connections were also subject to STDP and homeostatic plasticity, which allowed arbitrary sensorimotor transformations to develop during training. A more detailed description of the network along with the parameter settings used in the experiments can be found in the Supplementary Material.

RESULTS

Simulated neural activity during temporally segmented behavior – Before describing the BBD experiment, we illustrate the capability of the sequence network to learn to generate sequences of simulated responses to sensory inputs. The system of reentrantly coupled CAS networks can learn to reproduce a temporal sequence of eight consecutive input patterns. These individual patterns were simulated by means of current injections into Area A excitatory cells (see Supplementary Figure 2 for resulting network activity patterns). Multiple presentations of a given temporal sequence to the network constituted a training regimen. A given sequence consisted of an ordered series of eight distinct, randomly generated input patterns. Each pattern was presented for one second. After thirty-two seconds of training, the input was discontinued. At this point, the network continued to regenerate the eight patterns in correct order. Figure 3A shows a raster plot of the spiking of all excitatory and inhibitory neurons in the simulation as the system autonomously cycled through the trained sequence for 700 ms. The pattern of activity of the neuronal population consisted of a series of stable microstates – periods in which each neuron fires at a steady rate – each lasting ~100 msec, flanked by briefer, more complex transition states.

A brief account of the mechanisms by which the network develops sequence generation ability will aid in understanding what follows. One way to form a network that recognizes and generates temporal sequences of input patterns is to establish serial connections between distinct neuronal groups. If each neuronal group responds to a different input pattern – due to WTA dynamics – and a sequence of unique patterns is presented, then the neuronal groups will be activated successively. Given sufficient temporal overlap between the activity in successively responding neuronal groups, Hebbian mechanisms will act to strengthen the connections between them. These connections favor activation of the next neuronal group in the sequence in the absence of the external input, allowing for internal pattern generation of an arbitrary temporal sequence learned through experience. Separating the network into two populations, Area A reflecting the current pattern, and Area B reflecting the prior pattern, allows simultaneous activity in temporally adjacent neuronal groups, one in each WTA area, facilitating synaptic change via Hebbian learning.

Figure 3A illustrates the mechanisms underlying the microstate transitions between neuronal group activations within a temporal sequence after training. At the time labeled T_1 in figure 3A, the activity in area A that reflected pattern 4 (blue dots) ceased. Active neurons in area B no longer received input and ceased to fire at time T_2 when voltage *independent* NMDA currents, which characterize this network, decayed. Due to the loss of lateral inhibition, neurons in Area B giving rise to pattern 5 (green) began to fire at time T_3 in response to input from area A. Once these Area B cells for pattern 5 were activated, they triggered the firing of cells in Area A that correspond to pattern 6 (red) at time T_4 . At time T_5 , cells in Area A corresponding to pattern 5 no longer received input from Area B and ceased to fire. The network continued to advance through a series of microstates in this fashion until all patterns were generated.

Figure 3B reflects an analysis of spiking data from this simulated network, acquired over a longer period, 24 seconds. Each row in the figure plots the match score (in 50 ms time bins) to one of the eight training patterns. White is a perfect match, while black indicates a complete mismatch. The last training repeat is from $t=24$ to $t=32$. Subsequently, external stimulation was removed and STDP was discontinued in order to test whether training was successful. Nevertheless network activity continued autonomously. After presenting any one pattern in the sequence, the network repeatedly cycled through the patterns until another input stimulus was presented. Because the network had been presented with repeated transitions from pattern 8 to pattern 1 during training, the network cycled through all eight patterns repeatedly until it was interrupted. In order to test that these results were reproducible, the simulation was performed five times in total using different pseudorandom number seeds from the standard C library (Kernighan and Ritchie, 1988) to distribute the initial synaptic connectivities and strengths in the networks; in every case the system recalled eight patterns in the correct order.

Although the system of networks repeatedly regenerated the learned sequence autonomously, it nonetheless remained responsive to novel external input. To demonstrate this, we interrupted the autonomous activity every eight seconds by presenting the input corresponding to a different member of the set of learned patterns. For example, as shown in figure 3B at $t=37$ sec., pattern 6 was presented out of order for one second to reset network activity. Subsequently the sequence continued in the trained order. Thus, after being presented with a repeated series of input patterns, this system of networks correctly anticipated the next pattern in a temporal sequence. Figure 4 shows plots of the average match score of each pattern in a sequence during the one second presentation of the previous pattern. During the second presentation of the sequence from $t=9$ to $t=16$, the match score is zero (blue solid line), but during the fourth training trial from $t=25$ to $t=32$ the match score to the anticipated pattern increases after 250 ms (red dashed line), indicating that the system has formed an association between temporally adjacent patterns in the sequence. Similar results were obtained in all five simulations with different initial conditions.

Motor control of a BBD - To demonstrate the use of this system of simulated neuronal networks to regulate real-world behavior in real-time, the spiking output of the sequence generation network was used to control three-dimensional movements of a BBD. The task for the device was to learn to move its hand autonomously in a pre-specified order through four different locations in its visual field. Visual input from the BBD's camera was used to drive a retina model that projected topographically to the Input area of the sequence generation network. This allowed the BBD to learn a sequence of visual stimuli (See Supplementary Material for details of the retina model). The BBD was placed in a seated position with its camera looking towards its left side. A bright object was placed in its hand to provide salient stimulation at the location of the BBD's left hand in the visual field. Figure 5 shows examples of the raw video input from the camera with the arm in each of the four postures used in the experiment. To generate the desired sequence of segmented arm movements, the pattern of spiking excitatory cells of Area B was used as input to the simulated motor network. To establish the hand-eye coordination that this task requires, two angles of the right shoulder of the robot were successively manipulated to position its hand for 1 second in each of the four locations in the visual field. This was accomplished by injecting appropriate current into groups of neurons within the simulated motor area. This was repeated a total of 5 times. During this first training stage from $t=1$ to $t=20$ seconds, this system learned to discriminate between the four different spatial visual patterns. STDP was activated on the pathway from the Input area to Area A. The CAS network operating in Area A developed sparse activity patterns discriminating these four Input area patterns (Chen et al., 2013) (See Supplementary Figures 3-5 for example of activity patterns from one simulation). Non-plastic topographic connections from Area A to Area B essentially create a copy of Area A's pattern of activity in Area B.

A second training stage was used to allow the system to learn hand-eye coordination. The stimulation patterns from stage one were repeated from $t=21$ to $t=40$ seconds while STDP was activated on the pathway from Area B to the motor area. During this stage, this system came to associate the visual responses in Area B evoked by different postures with the pattern of motoric output that generated and maintained these postures (figure 5). After this training stage, hand-eye coordination was established, but sequence learning had not yet been achieved.

A final training stage was used to train the visual sequence network to generate the sequence of visual patterns corresponding to a sequence of arm movements. During this stage, from $t=41$ to $t=60$ seconds, STDP was activated on the pathway from Area B to area A for the first time. The system was trained by moving the arm of the device once again five times through the sequence of four postures, pausing one second at each posture. Subsequently, after the camera input to the system was discontinued, the BBD continued to autonomously generate motor commands that evoked movements similar to those used during the training phase. Each segment of the autonomous gestures lasted ~400 ms. The experiment was performed five times incorporating different initial network

parameters. Each time it reproduced the correct continuous sequence of movements. Figure 6 shows a trace of the movements made by the hand of the BBD during the five experiments, plotted in Cartesian coordinates both during the last training stage (green), and for 20 seconds after training during autonomous motor sequence generation (red). Positions were calculated from joint angles recorded every 200 ms during the simulation. One of the five subjects showed some error and consistently “cut” the upper corner, generating a different shape than the other four subjects. The self generated arm trajectories approximate the training trajectories.

We verified that the system remained responsive to external visual stimulation while it continued to generate the trained sequence autonomously. For each of the five subjects, visual stimulation was resumed at $t=100$ seconds. While the BBD continued to cycle through the trained sequence, the bright object used for training was moved sequentially by the experimenter to each of the locations in space it occupied during training and was held in place from 3 to 10 seconds. In 19 of 20 trials (five subjects tested at four locations) the BBD moved its hand to the location of the object and held it there until the experimenter removed the object from the visual field, at which point the BBD resumed the learned sequence from its present location. In one trial, the BBD moved its hand to the object, but then resumed the sequence prior to stimulus removal.

DISCUSSION

The robust recognition and regeneration of motor sequences known to occur in animals is accomplished by networks of spiking neurons. Here we show that this basic capability can be simulated using large-scale networks of spiking neurons. The computational model employed here can be further elaborated to explore sequence recognition and generation in networks consisting of groups of reentrantly connected neurons. The simulations demonstrate how networks composed of thousands of densely interconnected spiking neurons can respond adaptively to patterned sensory input by generating autonomous, temporally-ordered sequences of neural activity. We found that the operation of STDP to shape the distribution of synaptic strengths within and among WTA networks can give rise to network responses able to control complex behavior in a robotic device.

The system described here builds upon our previous work with large-scale spiking networks (Chen et al., 2013). The prior work explored the use of WTA networks for visual pattern categorization and feed-forward mappings between a sensory and motor map. The prior system was not capable of learning sequences. The present work demonstrates that coupling two WTA spiking networks together with specific reentrant connections leads to the ability to regenerate sequences after experience. The prior system was entirely sensory driven, while the present work allows internally generated network activity in the absence of sensory input (figure 3), yet remains responsive to external input. Further, the

rapid microstate transitions observed in this network are consistent with cortical microstate transitions.

Several theoretical models of behavioral sequence generation have been reported in the literature. Rhodes et al. (2004) proposed a mean-firing-rate model, N-STREAMS, which reproduces the physiological results of Averbeck et al., (2002). Salinas studied a mean-firing-rate simulation that incorporates rank-order-selective (ROS) neurons into a network and showed that the model could learn sequential motor actions given such neural responses. The activity of the ROS neuronal units was built-into the model, and did not emerge through learning. Verduzco-Flores et al. (2012) created a small mean-firing-rate network with 200 neurons that could learn multiple sequences with shared subsequences. Their model required temporally adjacent input patterns to partially overlap in time. Finally, Chersi et al. (2011) investigated a spiking network model that generates chains of temporal sequences of neural activity similar to those in our model and comparable to neurophysiological responses found in the intra-parietal lobe in primates. They used four separate pools of 500 neurons each to represent one of 4 different actions. Sparse connections between the 4 pools were subject to STDP. They showed that repeated activation of the neurons in the 4 pools in a given temporal order via simulated current injection eventually lead to correct recall of the remaining sequence after injecting only the first pattern. Our model does not require the use of discrete pools of neurons; rather such pools emerge automatically within each network through a WTA competition in the CAS architecture.

It is interesting to consider whether there is a benefit to using spiking neurons instead of rate based neurons in the brain-based device. The sequence generation network may have worked just as well with a model incorporating mean-firing rate neurons. Nevertheless it is important to demonstrate that spiking networks can generate such behavior, because animal nervous systems incorporate spiking neurons. This work demonstrates that spiking networks incorporating STDP can be reliably trained to generate sequences in the real-world.

By using simulated neuronal networks to control the behavior of a BBD we found that a real world device can be trained to generate autonomous, multi-segmented behavior. After training the system by presenting the target pattern of video input in one-second time steps, the BBD regenerates this sequential input pattern, but at a faster rate. The BBD was able to recreate movements composed of four consecutive steps in the correct order. Although the device can remember and reproduce multiple sequences of behavior, each posture within any sequence must be unique. Otherwise the subsequent posture would be ambiguous. Learning more complex behaviors will require incorporation of longer temporal contexts than those provided by the immediately preceding pattern.

The spiking activity corresponding to consecutive equilibrium postures in the behavioral sequence overlap in time, similar to activity reported in primates (Averbeck et al., 2002). This can be seen in figure 3B from second 24 to 32. For

example, shortly after the network responds with a high match score to input pattern 3 at $t=26$, network activity begins also to match pattern 4. The match scores (shades of gray) to each subsequent pattern begin to increase well before the pattern is presented to the network. In primates, this overlap in neural responses reflects current and future gestures made by the animal as it draws shapes “in the air”. Averbeck et al. (2002) also reported that the neural activity pattern corresponding to the current gesture was more strongly represented than the activity pattern reflecting the upcoming movement. This behavior is seen in our network. The match score for the current pattern, pattern 3, in figure 3 at $t=26$ is higher than the match score for the upcoming pattern 4.

Over time, spiking activity in the model network transits through a series of microstates, each characterized by a stable unique pattern of steady-state firing rates (figure 3A). Similar behavior has been observed in mammalian cortex. For example, neurons in the gustatory cortex in rodents (Jones et al., 2007), and in the prefrontal cortex of primates (Seidemann et al., 1996) progress through sequences of states, identifiable in examinations of simultaneously recorded neuronal ensembles. As in the simulation reported here, these states lasted for hundreds of milliseconds, with rapid transitions on the order of 50 ms. Our model suggests that such microstate transitions may be explained as reentrant interactions (Edelman, 1978) between multiple WTA networks.

Acknowledgements

We thank Joseph A. Gally for numerous helpful comments and suggestions on the manuscript. We are grateful to Yanqing Chen for numerous helpful discussions and development of WTA networks. Donald Hutson built the custom BBD for us. Richard Martin developed the computer interface to the BBD, and assisted with the BBD experiments. This work was supported in part by DARPA through ONR Grant N00014-08-1-0728 and by AFRL Cooperative Agreement FA8750-11-2-0255 to Neurosciences Research Foundation. For support of late developments we are grateful to the Mathers Charitable Foundation. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views, opinions, and/or findings contained in this article/presentation are those of the authors/presenters and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Air Force Research Laboratory, the Department of Defense, or the U. S. Government.

References

- Averbeck, B.B., Chafee, M.V., Crowe, D.A., and Georgopoulos, A.P. (2002). Parallel processing of serial movements in prefrontal cortex Proc Natl Acad Sci U S A. 99: 13172-13177.
- Bi GQ and Poo MM (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J. Neurosci. 18, 10464-10472.

- Binshtok AM, Fleidervish IA, Sprengel R, Gutnick MJ. (2006). NMDA receptors in layer 4 spiny stellate cells of the mouse barrel cortex contain the NR2C subunit. *J Neurosci* 26(2):708-15.
- Chen Y, McKinstry JL, Edelman GM (2013). Versatile networks of simulated spiking neurons displaying winner-take-all behavior. *Front. Comput. Neurosci.* 7:16. doi: 10.3389/fncom.2013.00016
- Chersi F, Ferrari PF, Fogassi L. (2011). Neuronal chains for Actions in the Parietal Lobe: A Computational Model. *PLoS ONE*, 6(11):e27652.
- Dayan P and Abbott LF (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, Massachusetts, pp. 255-256.
- Edelman GM (1978). Group Selection and Phasic Reentrant signaling: A Theory of Higher Brain Function in *The Mindful Brain*, MIT Press.
- Edelman GM (2007). Learning in and from brain-based devices. *Science* 16:318:1103-5.
- Izhikevich EM (2003) Simple Model of Spiking Neurons. *IEEE Transactions on Neural Networks*. 14:1569-1572.
- Izhikevich EM, Edelman GM (2008). Large-scale model of mammalian thalamocortical systems. *Proc Natl Acad Sci U S A*. 105(9):3593-3598, 2008.
- Jones LM, Fontanini A, Sadacca BF, Miller P, and Katz DB (2007). Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A*. 104(47):18772-18777.
- Kernighan BW, Ritchie DM (1988). *The C programming Language*, Prentice-Hall, Upper Saddle River, NJ.
- Liu JK, Buonomano DV (2009). Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *J Neurosci*. 29(42):13172-81.
- McKinstry JL, Seth AK, Edelman GM, Krichmar JL (2008). Embodied models of delayed neural responses: Spatiotemporal categorization and predictive motor control in brain-based devices. *Neural Networks* 21:553-61.
- Myme CI, Sugino K, Turrigiano GG, Nelson SB (2003). The NMDA-to-AMPA ratio at synapses onto layer 2/3 pyramidal neurons is conserved across prefrontal and visual cortices. *J Neurophysiol.* 90(2):771-9.
- Nakajima T, Hosaka R, Mushiake H, Tanji J. (2009). Covert representation of second-next movement in the pre-supplementary motor area of monkeys. *J Neurophysiol.* 101(4):1883-9.
- Rhodes BJ, Bullock D, Verwey WB, Averbeck BB, Page MPA (2004). Learning and production of movement sequences: Behavioral, neurophysiological, and modeling perspectives. *Human Movement Science* 23:699–746.
- Rutishauser U and Douglas R (2009). State-dependent computation using coupled recurrent networks. *Neural Computation*, 21(2).
- Salinas E (2009). Rank-Order-Selective Neurons Form a Temporal Basis Set for the Generation of Motor Sequences. *Journal of Neuroscience*. 29(14):4369-4380
- Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996). Simultaneously Recorded Single Units in the Frontal Cortex Go through

Sequences of Discrete and Stable States in Monkeys Performing a Delayed Localization Task. *Journal of Neuroscience* 76(2):752-768.

Tanji J (2001). Sequential organization of multiple movements: Involvement of Cortical Motor Areas. *Annu. Rev. Neurosci.* 24:631–51.

Verduzco-Flores SO, Bodner M, Ermentrout B (2012). A model for complex sequence learning and reproduction in neural populations. *J Comput Neurosci.* 32(3):403-23.

Zucker RS (1989). Short-term synaptic plasticity. *Annu Rev Neurosci.* 12:13-31.

Figure 1. Sequence generation network architecture. (A) Center-Annular-Surround (CAS) network architecture that produces WTA dynamics. The CAS network architecture consists of interconnected spiking neurons, excitatory (green ovals) and inhibitory (red ovals). Each population is arranged in a two-dimensional grid. Connections from representative cells are illustrated. Axons from excitatory neurons (green arrows) project to neurons within green areas. Axons from inhibitory neurons (red arrows) project to neurons in the transparent red annular areas. The CAS connectivity leads to WTA dynamics, in which small regions of high activity are surrounded by large regions with little activity. (B) The sequence generation network is comprised of two reentrantly interconnected Center-Annular-Surround (CAS) spiking networks, Areas A and B. Arrows indicate directions of neural connectivity, while the circle and the donut shape indicate the inter-network connectivity (projection field) from single points in the projecting network. The input area projects non-topographically to Area A. Area A projects topographically to Area B, as indicated by the small oval in Area B. In turn, Area B projects topographically and widely back to area A, but not to the same spot from which it received input, as indicated by the donut-shaped ring in Area A. Avoiding projections to the corresponding spot helped prevent the network from locking into a single activity pattern due to self-amplification. Rather it allowed the network to switch smoothly between patterns in a sequence.

Figure 2. Custom humanoid robot, or brain-based device, used for behavioral tests of the sequence generation network. The BBD has a grayscale camera which monitors the location of the bright object in its hand in order to learn “hand-eye” coordination of its left hand. During experiments the left arm was moved repeatedly in a sequence of four different postures. See Materials and Methods for a detailed description of the device.

Figure 3. A large-scale network of approximately 4,000 spiking neurons autonomously transitions between states reflecting a learned sequence.

(A) Spike rastergram of all neurons in Networks A and B showing the population activity recorded over 700 ms. as the network spontaneously generated the learned sequence. Each spike is shown as a colored dot and each neuron is assigned a color to indicate the pattern to which it responds maximally. Networks A and B transitioned spontaneously between stable-states corresponding to three learned input patterns in the sequence as indicated by the three colors blue, green, and red associated with patterns four, five, and six, respectively. (The few magenta dots are associated with neurons responding best to another pattern, but which are also activated by pattern four (blue)). The four neural populations in the network are labeled on the right of the diagram. Inh. = Inhibitory, Exc. = Excitatory. The labels T_1 through T_5 mark transition times referred to in the text.

(B) After training, the coupled networks spontaneously generate a sequence of 8 patterns in the correct order, and that sequence can be interrupted or shifted by presenting an external stimulus. Each row of the figure indicates, by brightness, the match score over time for one of the eight patterns that make up the sequence. The match score for pattern X for example indicates how closely the neural population activity pattern in the Area A excitatory neurons matched the activity in the same population recorded when pattern X was presented to the network for the first time during training. White is a perfect match, while black indicates a complete mismatch. The network was trained from $t=0$ -32 seconds by stimulating the 8 patterns in order. The times of stimulus presentation are indicated by black bars under the figure. The internally generated sequence is interruptible. When presented with one pattern in the sequence, pattern 6, for one second at $t=37$ s and $t=45$ s, the network activity immediately reflected the stimulus, and when the stimulus was removed the network generated the sequence from pattern 6 onward.

Figure 4. After training, the simulated neural system anticipates upcoming patterns in the sequence. The average match strength of individual patterns in the sequence is plotted over time during the presentation of the *prior* pattern. Data was obtained during the second presentation of the sequence from $t=9$ s to $t=16$ s (blue solid line), and during the fourth presentation of the sequence from $t=25$ s to $t=32$ s (red dashed line). During the fourth presentation, 250 msec after each new pattern was presented, the activity of neurons in Area A began to match that of the next pattern in the sequence. Error bars show the standard error of the mean match response.

Figure 5. The four arm postures of the BBD as viewed through its video camera. The BBD was trained to move its hand consecutively in numerical order to the 4 spots outlined in red. The images show the all four arm postures of the BBD. During training, a bright object placed in the hand made the hand positions salient against the dark background (The lighting was increased when these images were taken to provide sufficient contrast to see the arm). During a test for spontaneous recall of the learned sequence, the bright object was removed to eliminate visual input.

Figure 6. The BBD self generates a learned sequence of arm motions using the sequence generation network model. The BBD was trained to move its arm through four different postures such that its hand traced out a quadrilateral in space. The figure shows the superposition of hand positions recorded during 20 seconds of training (green line) and during 20 seconds of self-generated movements after training (red line) for all 5 subjects. Lines are drawn between temporally adjacent data points recorded at 50 ms intervals. One of the five subjects showed some error and consistently “cut” the upper corner,

generating a different shape than the other four subjects. The self generated arm trajectories approximate the training trajectories.

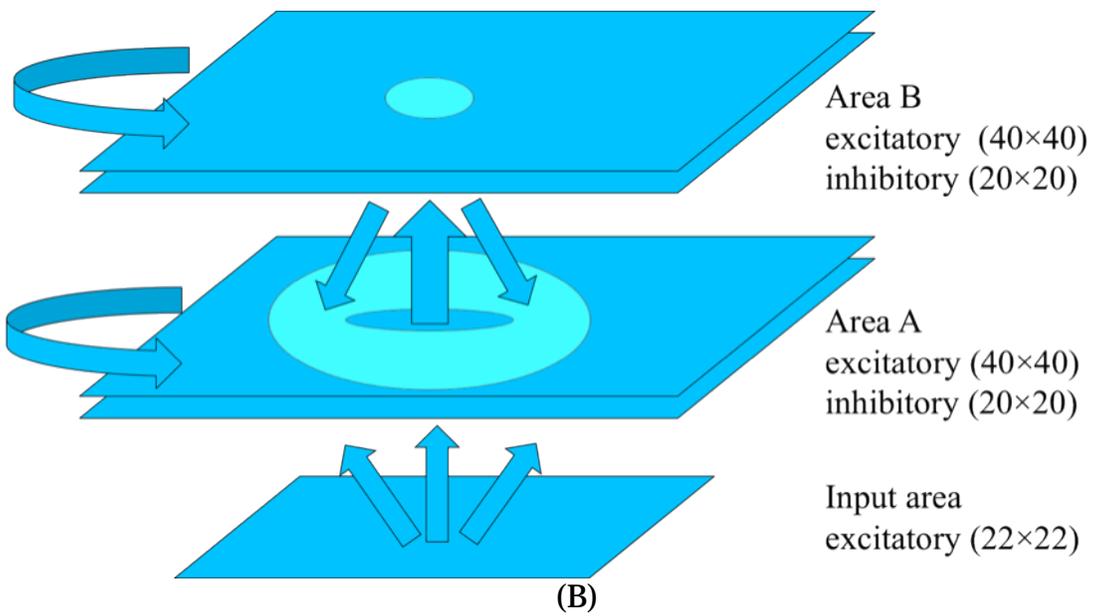
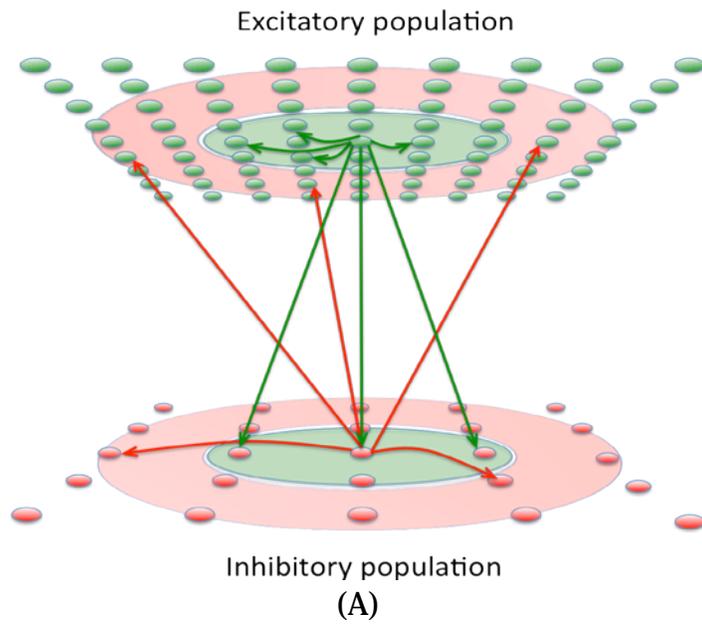
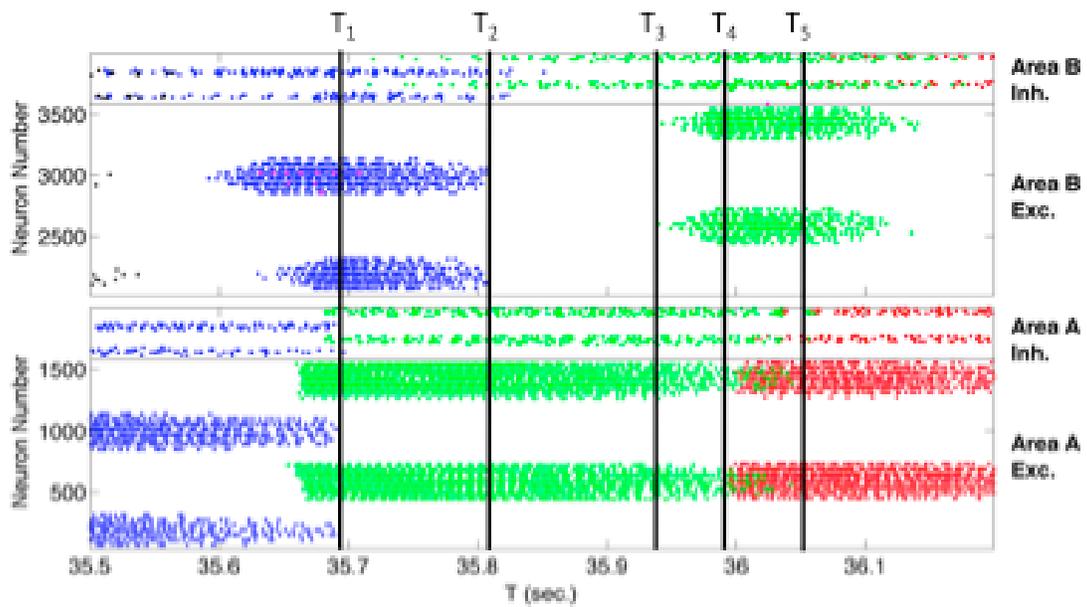


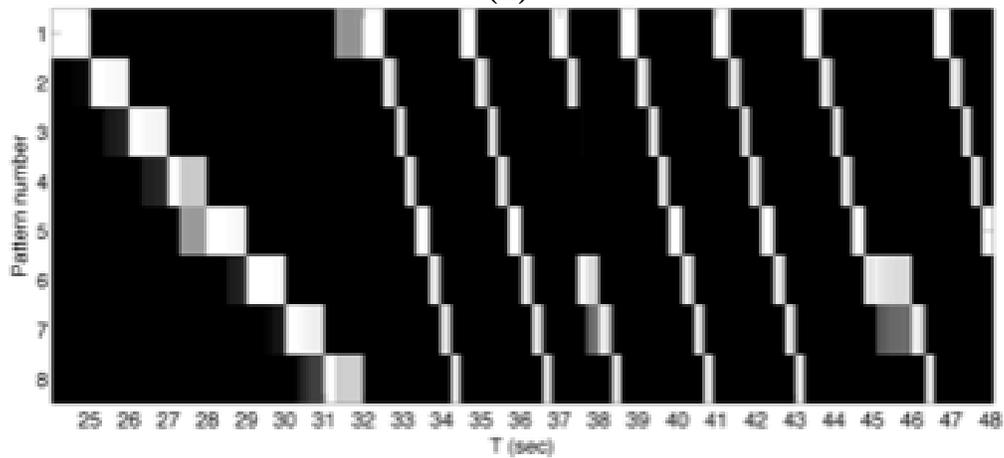
figure 1



figure 2



(A)



(B)

figure 3

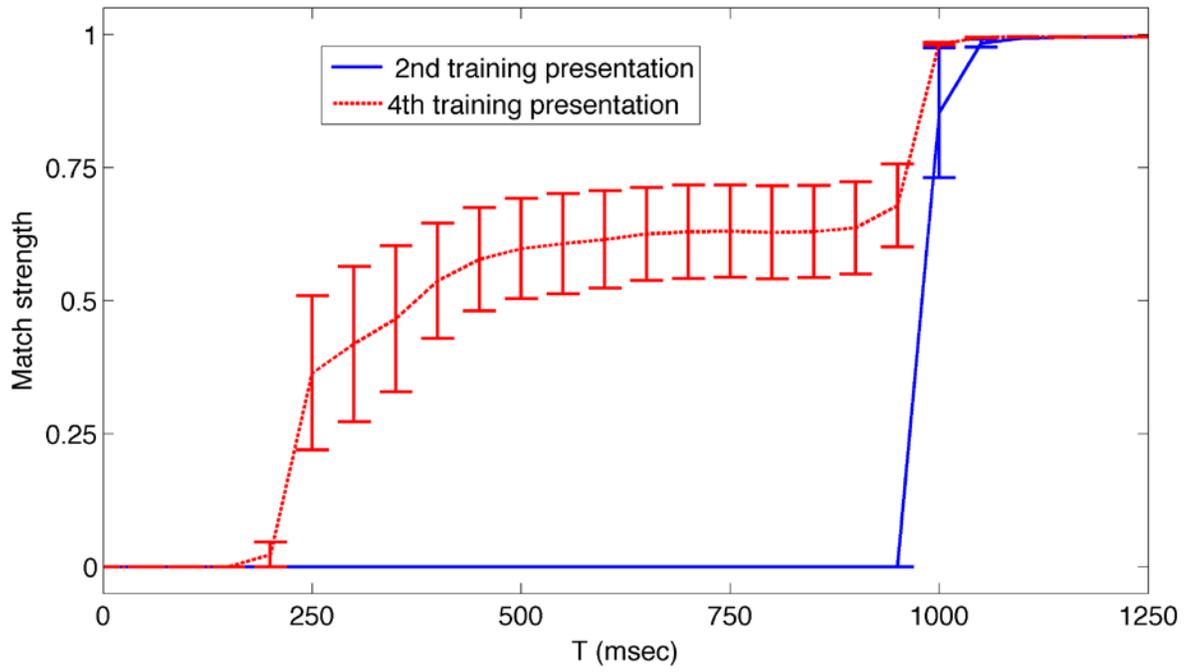


figure 4

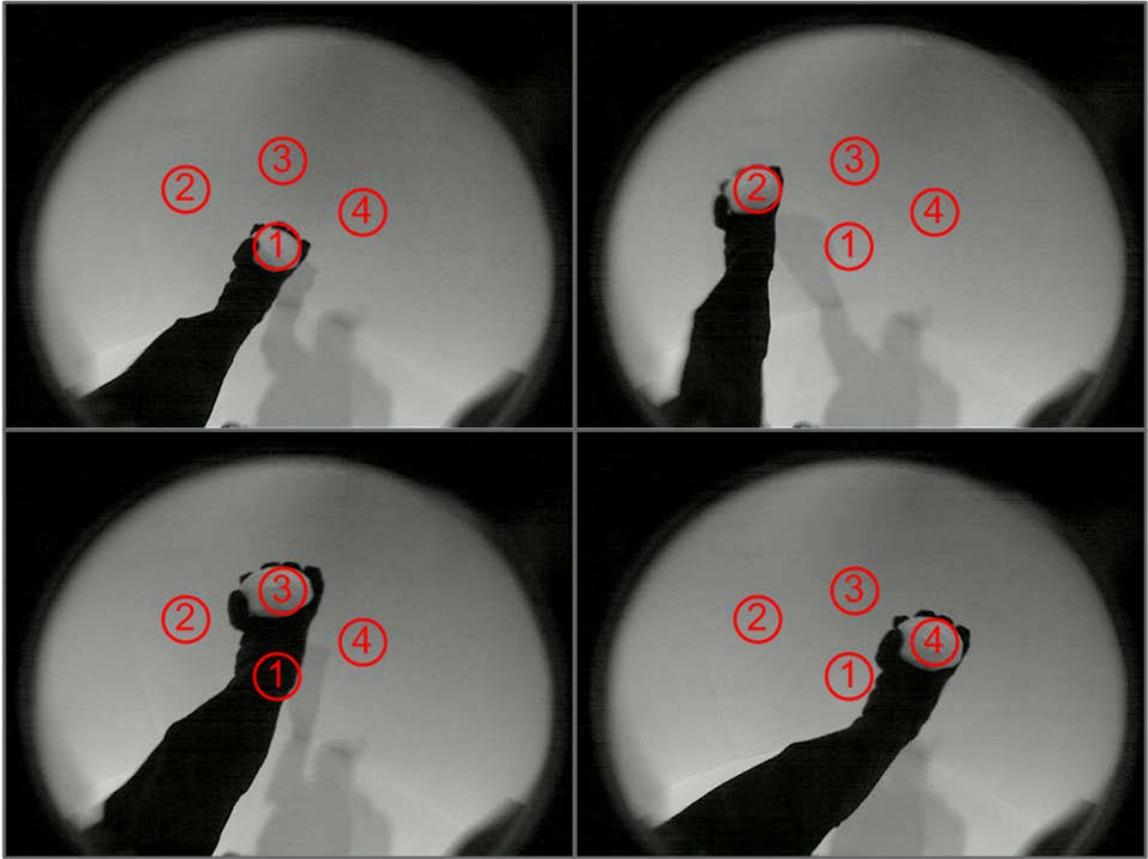


figure 5

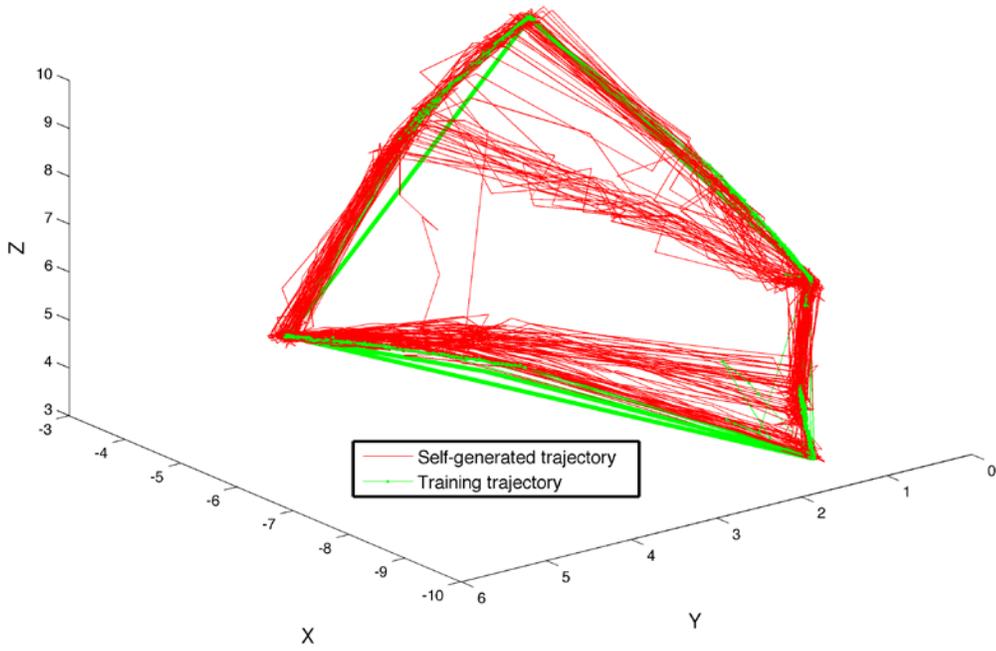


figure 6

Supplementary Material.

Neuron parameters

Table I shows each of the neuron model parameters used in all experiments.

Table I. Neuronal parameters. a.u.: arbitrary units.

Neuron type	Area	C (pF)	k (a.u.)	V_r (mV)	V_t (mV)	V_{peak} (mV)	a (a.u.)	b (a.u.)	c (mV)	d (pA)
Exc.	A	80	3	-60	-50	50	0.01	5	-60	10
Inh.	A	20	1	-55	-40	25	0.15	8	-55	200
Exc.	Motor	100	0.7	-60	-50	0	0.03	-2	-60	100
Inh.	Motor	20	1	-55	-40	25	0.15	8	-55	200
Exc.	B	80	3	-60	-50	50	0.01	5	-60	10
Inh.	B	20	1	-55	-40	25	0.15	8	-55	200
Thalamic	Input	200	1.6	-60	-50	40	0.01	15	-60	10

Anatomy.

The connectivity between model neurons fell into two classes: either local-type or surround-type. For local-type connectivity, a two-dimensional Gaussian probability distribution, centered on each post-synaptic cell, determined the probability of forming a synapse between each potential pre-synaptic neuron within a specified maximum distance, r_{max}

$$f(d) = ae^{-\frac{(d-\mu)^2}{2\sigma^2}} \quad (8)$$

where a is a scale factor set to generate, on average, a target number of synapses on each post-synaptic cell, d is the distance between the pre-synaptic neuron and the post-synaptic neuron, μ is 0, and σ is the standard deviation. In a similar manner, a two-dimensional Gaussian function was also used to specify the synaptic strength between connected neurons as a function of the distance between them in the network. The total of all synaptic efficacies was scaled to sum to a constant parameter with units in nanoSiemens (nS). Both connection probability and strength were maximal between nearest neighbors, and fell off as a function of distance, controlled by the same parameter, the standard deviation of a Gaussian.

For surround-type connectivity, a post-synaptic neuron receives synaptic connections from neurons located in a surrounding annular region specified by a

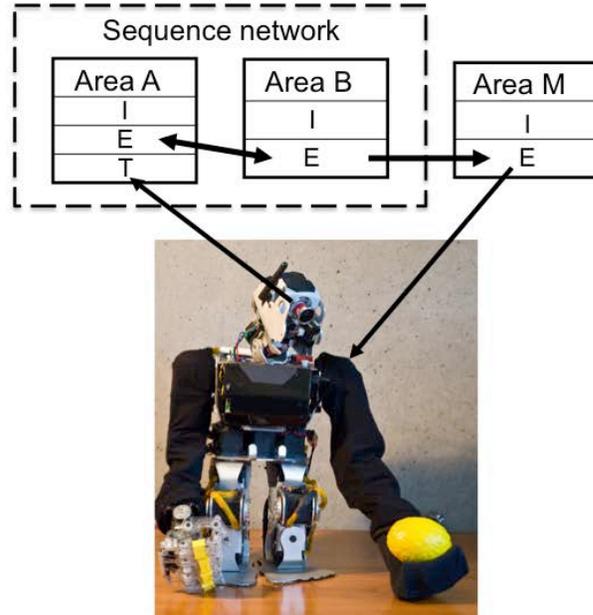
minimum (r_{\min}) and maximum (r_{\max}) radial distance from the post-synaptic cell. (This is equivalent to saying that each pre-synaptic neuron sends projections to post-synaptic neurons in an annular region). The probability of forming a connection with a neuron in the annulus is determined as a function of distance from the post-synaptic cell. The function used is a Gaussian with standard deviation σ , centered at $\mu=(r_{\min}+r_{\max})/2$. Thus a post-synaptic neuron does not connect with a neuron within the central ring, has minimal connection probability at the inner ring of the annulus, maximal probability half-way between the inner and outer ring of the annulus, and no connections beyond the outer ring. This probability density function is scaled to create a prespecified number of synapses for each post-synaptic neuron. The synaptic strengths for the surround-type connection are also initialized using the same function, with the same parameters. However, the sum of all synaptic strengths of this type was scaled to make the total equal to a constant value under experimenter control.

In order to avoid boundary conditions in the network, the network was treated as a torus. Thus connections from neurons that would go outside of the network instead “wrap around” to connect with neurons on the opposite edge.

Table II shows the parameters defining the anatomy and synaptic parameters of the sequence generation network for the network with simulated sequence inputs. The table defines two types of information for every neural area: the neuron composition, and the synaptic connectivity for each neuron type. The first four columns of the table list, for each separate neural population in the simulation, the type of neuron, the area in which they are located, the number of neurons in the population, and the total number of synapses per neuron.

The remaining columns define the connectivity for each type of neuron in the area. Multiple rows are necessary to define the connectivity for each post-synaptic type; one row is needed for each presynaptic neuron type forming synapses on the post-synaptic neurons. Pre-area and pre-type specify the presynaptic area and type of a neuronal group projecting to the post-synaptic group. The next column specifies the percentage of the post-synaptic cell's synapses allocated to this pathway. The remaining columns provide all of the parameters used to specify details of the synaptic pathways as described in the paragraphs above.

Table III show the parameters defining the anatomy and synaptic parameters of the motor sequence network used to control the arm of the BBD; the format is the same as Table II.



Supplementary Figure 1. The neural architecture for the BBD experiments.

Neural architecture for the BBD experiment.

Supplementary Figure 1 illustrates the neural architecture used for the BBD experiments. The sequence network consists of two CAS networks, Area A and B, each consisting of populations of excitatory (E) and inhibitory (I) neurons synaptically coupled as described in the paper. Visual input from the video camera provided patterned input to “thalamic” (T) neurons of Area A, while the output of excitatory neurons in the motor area (M) were used to control the two shoulder joints of the left arm. After repeatedly stimulating the motor area in one of four different locations in the same order, and thus moving the arm to one of four different postures, a mapping formed from the visual area responses to the location of the hand to the motor area output that drove the hand to those locations. The sequence of four visual patterns was learned by the sequence network, Areas A and B. As this network recalled the visual sequence after training, associated area M activity caused the arm to move to the posture associated with the next visual pattern in the sequence, as shown in Figure 6.

Mapping motor network activity to arm postures.

In order to translate from neuronal firings to joint angle in the left shoulder of the BBD, the output of the motor network cells was pooled using a form of population vector averaging described next.

Each excitatory neuron of the motor network, Area M, was assigned a preferred set of angles for each of the two shoulder joints. Nearby neurons in this predetermined map evoked similar joint angles, but different patterns of activity among these cells could evoke all possible positions of the left arm. Neuron i at grid position u, v in the network prefers shoulder angles $(w_i^1, w_i^2) = (u/40, v/40)$.

For each joint, j , the preferred joint angles of all cells, weighted by the corresponding firing rate, were added together to determine an equilibrium posture as follows:

$$angle_j = \begin{cases} \frac{\sum_{i=1}^n w_i^j r_i}{\sum_{i=1}^n r_i}, & \text{when } \geq 10 \text{ neurons have } r_i > 15 \text{ hz} \\ \alpha_j, & \text{otherwise} \end{cases}$$

where w_i^j is the angle for joint j preferred by cell i , and r_i is the estimated firing rate of the cell estimated with a time constant of 100 ms. Joint angles were recalculated in this manner, and shoulder joints were adjusted every 250 milliseconds.

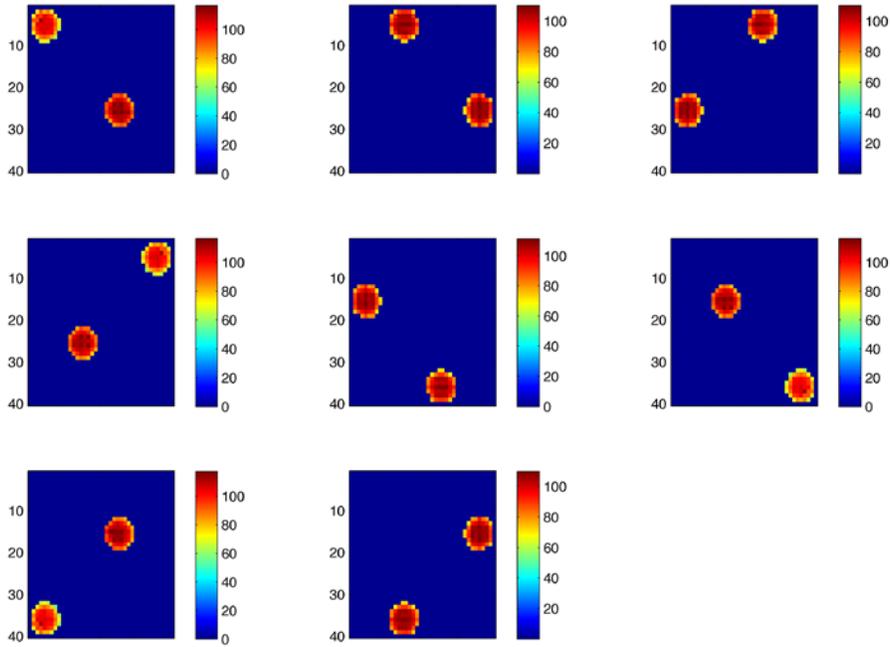
Visual input.

Video was recorded with an Axis 207MW wifi camera. Grayscale images with a resolution of 320x240 were transmitted at 30fps. The pixels from the central portion of the video frames were used to mimic retinal photoreceptors. These pixels provided input to a two-dimensional grid of on-center Retinal Ganglion Cells (RGC). The grid size was 21x21 neurons with a center area receptive field size of 3x3 and the surround area of 6x6 pixels. Each RGC receives a current that is computed following the algorithm of Wohrer and Kornprobst (2009). These currents were constantly injected at each integration step until the next video frame was received. RGCs were modeled with the Izhikevich model (Izhikevich and Edelman, 2008) with the following parameters: $C=100$, $V_r=-70\text{mV}$, $V_t=-50\text{mV}$, $k=1$, $a=0.005$, $b=0$, $c=-75\text{mV}$, $d=250$, and $V_{\text{peak}}=10\text{mV}$.

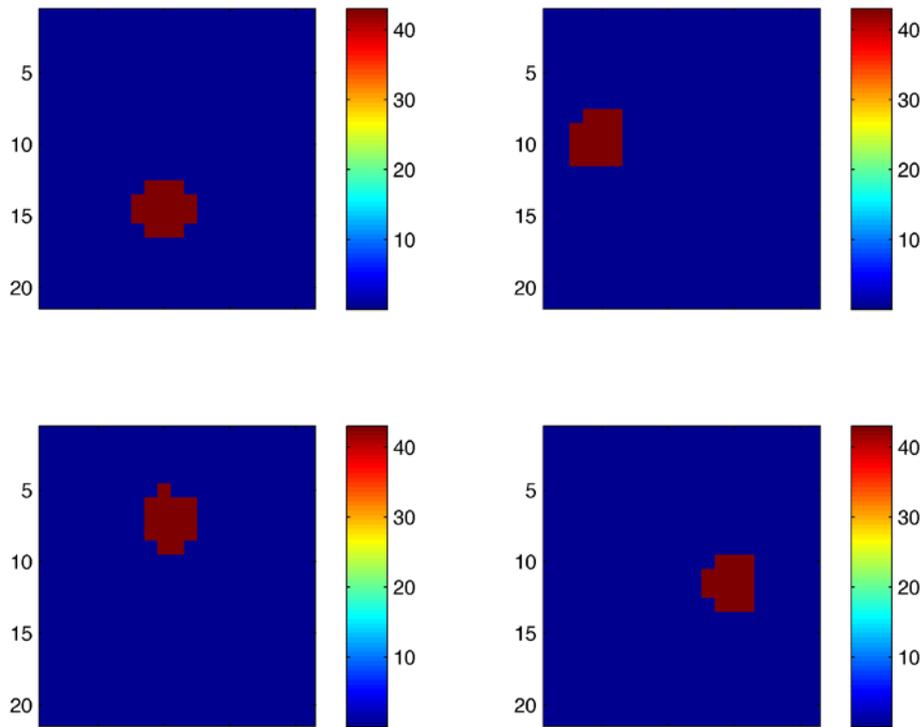
Training stimulation patterns used in the initial sequence experiment.

The first experiment uses eight artificial stimulation patterns to train the sequence network. We assume that each population pattern corresponds to a unique network response to an external sensory input. Supplementary Figure 2 plots the 8 individual population firing rate patterns resulting from current stimulation into a subset of the Area A excitatory neurons. Within each diagram, the color of each pixel indicates the firing rate of one neuron in the population, as

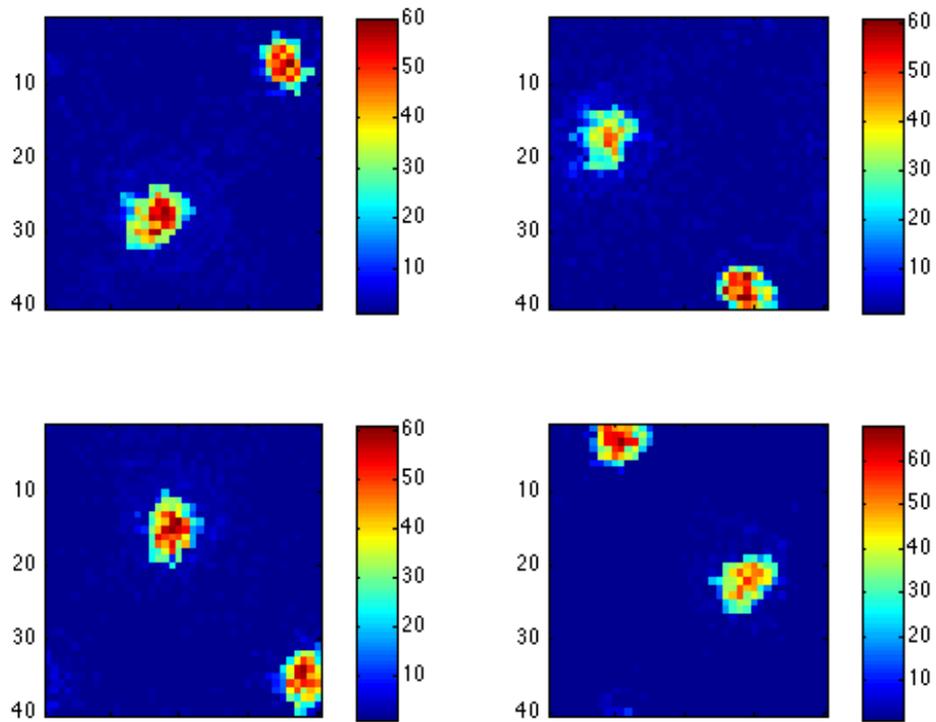
indicated by the color scale to the right of each diagram. Each simulated neuron with a non-zero firing rate was stimulated with a current injection of 1,000 pA for one second. These eight patterns were present one at a time and repeated in the network in the sequence shown in the figure from right to left, top to bottom.



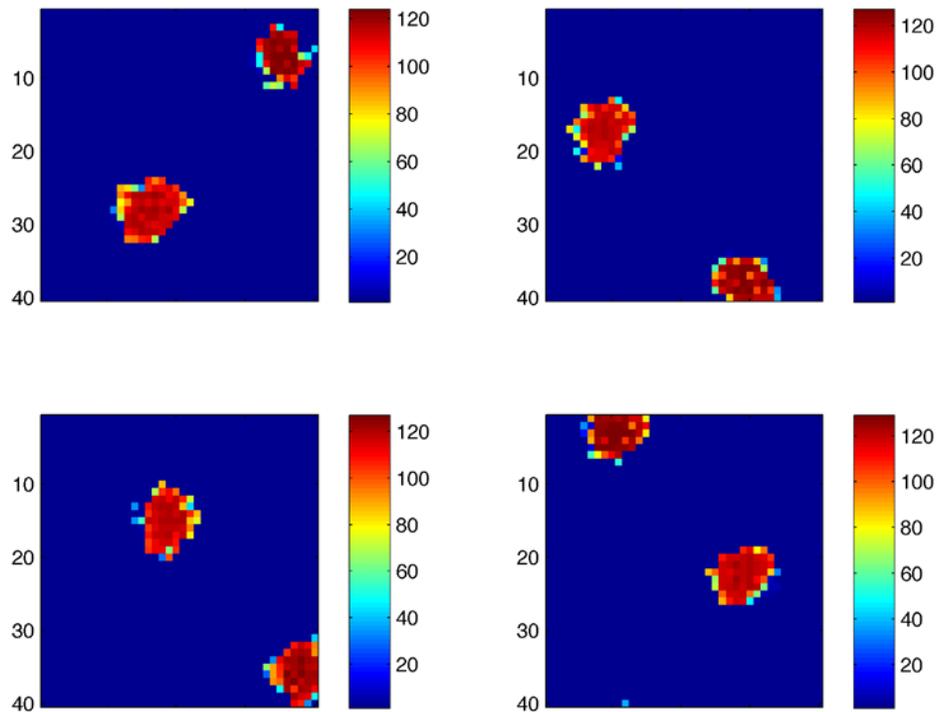
Supplementary Figure 2. The simulated population activity patterns used to train the sequence network. See supplementary text for details.



Supplementary Figure 3. The four population activity patterns in the Input area used in the first training stage in the BBD experiment from $t=1$ to $t=20$ simulation seconds (the network runs slower than real-time, and we measure time in network simulation cycles with a time step of 1ms). Plots are similar to supplementary figure 2. These four firing rate patterns reflect the visual patterns from the grayscale camera sensing the hand of APE-X in four different locations during 1 experiment. The color scale to the right of the four map indicates the mean firing rate associated with each color.



Supplementary Figure 4. The four population activity patterns in area A in response to the four stimulus patterns shown in Supplementary Figure 3 the first time they are presented, from $t=1$ to $t=4$ s for one BBD subject. Color scale to the right indicates the mean firing rate associated with each color.



Supplementary Figure 5. Similar to Supplementary Figure 4 but showing population responses during the fifth presentation of each pattern, from $t=16$ to $t=19$ s. The increased firing rates are due to STDP acting on connections from the Input area to excitatory cells in Area A, and local connections between excitatory cells within Area A.

References

- Izhikevich EM and Edelman GM (2008). Large-scale model of mammalian thalamocortical systems. *Proc Natl Acad Sci U S A.* 105(9):3593-3598.
- Wohrer A, Kornprobst P (2009). Virtual Retina: a biological retina model and simulator, with contrast gain control. *Journal of Computational Neuroscience* 26(2), 219–249.

Table II. Anatomical and synaptic parameters for the sequence generation network used in experiments with simulated input.

Post-synaptic neuron type	Post-synaptic area	Number of neurons	Average synapses per neuron	Pre-synaptic area	Pre-synaptic neuron type	Percentage of total synapses	r_{min} (mm)	r_{max} (mm)	σ (mm)	ϵ	S_{total} (nS)	S_{max} (nS)	Initial α	Final α	Learning start time (ms)	Learning end time (ms)	$gain_{wada}$	$gain_{wadaVI}$	$gain_{cavab}$	
Exc.	A	1600	3080	A	Exc.	14	-	0.1	0.05	0	22	10	0	0	0	0	0.5	0	0.1	
				A	Inh.	29	0.1	1	0.8	0	1200	20	0	0	0	0	0	0.5	0	0.1
				B	Exc.	29	0.36	1	0.3	0	60	30	0.9	0.9	0	32000	0	0.5	0.1	
				Input	Thalamic	29	-	1.44	2.5	0	200	50	0.9	0.9	0	32000	0.5	0	0.1	
Exc.	B	1600	2200	B	Exc.	40	-	0.15	0.07	0	30	5	0.9	0.9	0	32000	0.5	0	0.1	
				B	Inh.	40	0.15	1	0.3333	0	900	10	0	0	0	0	0.5	0	0.1	
				A	Exc.	20	-	0.15	0.07	0	22	5	0	0	0	0	0.5	0	0.1	
Inh.	A	400	2000	A	Exc.	20	-	0.33	0.16	0	25	5	0	0	0	0	0.5	0	0.1	
				A	Inh.	40	0.15	1	0.3333	0	180	3	0	0	0	0	0.5	0	0.1	
				Input	Thalamic	40	-	1.44	10	0	15	10	0	0	0	0	0.5	0	0.1	
Inh.	B	400	2000	B	Exc.	20	-	0.33	0.16	0	15	2	0	0	0	0	0.5	0	0.1	
				B	Inh.	40	0.15	1	0.3333	0	180	3	0	0	0	0	0.5	0	0.1	
				A	Exc.	40	-	0.15	100	0	15	50	0	0	0	0	0.5	0	0.1	
Thalamic	Input	484	0	-	-	0	-	0	0	0	0	0	0	0	0	0.5	0	0.1		

Table III. Anatomical and synaptic parameters for the sequence generation network used in experiments with the BBD.

Post-synaptic neuron type	Post-synaptic area	Number of neurons	Average synapses per neuron	Pre-synaptic area	Pre-synaptic neuron type	Percentage of total synapses	r_{min} (mm)	r_{max} (mm)	σ (mm)	ϵ	S_{total} (nS)	S_{max} (nS)	Initial α	Final α	Learning start time (ms)	Learning end time (ms)	$g_{dHPVmax}$	$g_{dHPVAVT}$	$g_{dHPVdAB}$	
Exc.	A	1600	3080	A	Exc.	14	-	0.1	0.05	0	22	10	0	0	0	0	0.5	0	0.1	
				A	Inh.	29	0.1	1	0.8	0	1200	20	0	0	0	0	0	0.5	0	0.1
				B	Exc.	29	0.36	1	0.3	0	60	30	0.9	0.9	40000	60000	0	0.5	0	0.1
				Input	Thalamic	29	-	1.44	2.5	0	200	50	0.9	0.9	0	20000	0.5	0	0	0.1
Exc.	B	1600	2200	B	Exc.	40	-	0.15	0.07	0	30	5	0.9	0.9	0	20000	0.5	0	0.1	
				B	Inh.	40	0.15	1	0.3333	0	900	10	0	0	0	0	0.5	0	0.1	
				A	Exc.	20	-	0.15	0.07	0	22	5	0	0	0	0	0.5	0	0.1	
				A	Exc.	20	-	0.33	0.16	0	25	5	0	0	0	0	0.5	0	0.1	
Inh.	A	400	2000	A	Exc.	20	-	0.33	0.16	0	25	5	0	0	0	0	0.5	0	0.1	
				A	Inh.	40	0.15	1	0.3333	0	180	3	0	0	0	0	0.5	0	0.1	
				Input	Thalamic	40	-	1.44	10	0	15	10	0	0	0	0	0.5	0	0.1	
				B	Exc.	20	-	0.33	0.16	0	15	2	0	0	0	0	0.5	0	0.1	
Inh.	B	400	2000	B	Inh.	40	0.15	1	0.3333	0	180	3	0	0	0	0	0.5	0	0.1	
				A	Exc.	40	-	0.15	100	0	15	50	0	0	0	0	0.5	0	0.1	
				Input	Thalamic	40	-	1.44	10	0	15	10	0	0	0	0	0.5	0	0.1	
				Thalamic	Input	484	0	-	-	0	-	0	0	0	0	0	0	0.5	0	0.1
Exc.	Motor	1600	1800	A	Exc.	33	-	0.33	0.16	0	5	5	0.9	0.1	20000	40000	0.5	0	0.1	
				A	Inh.	45	0.33	1.44	10	0	1200	10	0	0	0	0	0.5	0	0.1	
				B	Exc.	22	-	1.41	4.5	0.5	20	5	0.9	0.1	20000	40000	0.5	0	0.1	
				A	Inh.	40	0.33	1.44	10	0	800	20	0	0	0	0	0.5	0	0.1	
Inh.	Motor	400	1000	A	Exc.	60	-	0.33	0.16	0	10	2	0	0	0	0	0.5	0	0.1	
				A	Inh.	40	0.33	1.44	10	0	800	20	0	0	0	0	0.5	0	0.1	
				B	Exc.	22	-	1.41	4.5	0.5	20	5	0.9	0.1	20000	40000	0.5	0	0.1	
				A	Inh.	40	0.33	1.44	10	0	800	20	0	0	0	0	0.5	0	0.1	

A3. Fleischer, J.G., Gally, J.A., and Edelman, G.M. (2013) "A Spiking Neural Network Simulation of Working Memory," *Frontiers in Computational Neuroscience*, submitted for publication.

A spiking neural network simulation of working memory

Jason G. Fleischer^{1,2}, Joseph A. Gally¹, and Gerald M. Edelman¹

Abstract

Working memory (WM) is associated with persistent neural activity in prefrontal and parietal cortices and is often assessed in animals through delayed match-to-sample (DMS) tasks. In this paper we simulate a large-scale spiking neural network that can elicit persistent WM-related activity and carry out visual DMS tasks. We examine the ability of the network to remember stimuli accurately over a period of seconds through persistent activity. The network uses three known and distinct biological mechanisms that support persistent activity: dense reentrant excitatory connections, NMDA receptor activation, and short term synaptic plasticity. We assess the effects of manipulating these persistent activity mechanisms on the model's WM functionality. We describe a neural mechanism for detecting a match between persistent activity and activity evoked by the current stimulus, and we apply this mechanism to DMS tasks. We investigate how the network's capacity for storing multiple stimuli scales with the number of neurons, and discuss the implications of these results in relation to our network and others in the literature.

Introduction

Retaining a fleeting perception for seconds or minutes after a stimulus disappears is critical for many forms of behavior, cognition, and learning. Working memory (WM) allows for the cognitive manipulation of stored information about stimuli, and such memories can be used in decision making (Baddeley, 2012). While there is evidence for some separation between the neural substrates of WM and decision making (Bechara et al., 1998), it has been suggested that both WM and decision making involve active maintenance of information by prefrontal cortex (Miller and Cohen, 2001). Such active maintenance of information is associated with persistent neural activity (Fuster and

¹ The Neurosciences Institute, 800 Silverado Street #302, La Jolla, California 92037 USA.
Email addresses: {last name}@nsi.edu

² Corresponding author

Alexander, 1971; Curtis and Lee, 2010). An interesting feature of WM is its ability to simultaneously hold a limited number of different stimuli (Cowan, 2001).

In both human and animal studies, WM has often been investigated using a delayed match-to-sample (DMS) paradigm. Typically an animal is shown a brief stimulus to be remembered for a few seconds or minutes. After a delay period, during which no stimulus is presented, the animal is shown a second stimulus that might or might not be identical to the first. A correct response, indicating whether the two stimuli match, results in the delivery of reward.

In primates persistent firing discharges in prefrontal cortex correlate with the remembered stimuli during visual DMS tasks (Fuster and Alexander, 1971). Lesions in this area are known to disrupt WM performance (Passingham, 1975; Mishkin and Pribram, 1956), and the observation of persistent activity in prefrontal cortex suggests a plausible mechanism by which such memories can be maintained in the absence of the stimulus. Moreover, persistent discharge related to a stimulus during the delay period is not restricted to prefrontal cortex, but has also been observed in parietal cortex and sensory areas (Miyashita and Chang, 1988; Ferrera et al., 1994; Bisley et al., 2004; Romo and Salinas, 2003). One significant difference between memory responses in different areas is that prefrontal and parietal cells maintain stimulus-selective activity in the presence of distractors, while higher order sensory areas such as inferotemporal cortex (Miller et al., 1996) and S2 (Romo and Salinas, 2003) do not.

This paper presents a large-scale spiking neural model with persistent activity that enables multi-item working memory. The network incorporates three distinct biological mechanisms for generating persistent activity. All three mechanisms operate simultaneously in real cortical circuits, and each mechanism has been shown to be independently capable of supporting persistent activity. These mechanisms are (1) dense reentrant connectivity producing attractor dynamics (Amit and Brunel, 1997), (2) short-term synaptic plasticity enabling robustness against brief drops of firing rate (Mongillo et al., 2008), and (3) NMDA receptors maintaining excitation over durations longer than input inter-spike intervals (Wang, 1999). Persistent activity in the network is characterized in relation to parameters controlling these mechanisms. DMS tasks also require a neural mechanism for detecting a match between persistent activity and activity evoked by the current stimulus. We propose a matching mechanism based upon a segregation of visual and memory-related inputs onto AMPA and NMDA receptors of postsynaptic neurons, and we examine the ability of the network to perform visual DMS tasks. Finally, we characterize the capacity of the network to store multiple items simultaneously as a function of network size. An alternative architecture is proposed for which WM capacity does not scale with respect to network size, a result consistent with data showing that animal species with very different brains sizes may have similar visual WM capacity.

Material and Methods

Full details of the neural simulations can be found in the Supplementary Material.

Network. Neurons are grouped into areas that are distributed over a two-dimensional grid with toroidal topology. Neural areas have similar local circuitry and dynamics. Within each area, the cells are interconnected in an arrangement that we call a Center-Annular-Surround (CAS) architecture, which we have found to effectively generate winner-take-all or attractor dynamics in large-scale networks of spiking neurons (Chen et al., 2013). Such dynamics result from the CAS circuitry (see Figure 1A): neurons receive connections from nearby excitatory cells, whereas connections from inhibitory cells come from more distant neighbors in the surrounding annular region. When stimulated, the CAS circuitry gives rise to sparse firing organized into a few localized patches of neural activity. Patches form when some active neurons “win” by suppressing firing in their neighbors via the annular inhibition. The location of these patches is determined by the combination of random initial conditions and experience-dependent synaptic plasticity. CAS networks are similar in some ways to attractor neural networks arranged in two dimensions, e.g. (York and van Rossum, 2009). However, the network presented here has more biologically detailed synaptic and neural dynamics.

The network has four neural areas as shown in Figure 1B: visual input area (VIA), visual features area (VFA), working memory area (WMA), and match detection area (MA). VIA neurons are organized retinotopically, have firing rates proportional to pixel luminance in an on-center off-surround fashion similar to visual inputs arising from the retina (Wohrer and Kornprobst, 2009) and send random, non-retinotopic projections to VFA. VFA responses are selective to features of the visual stimulus, such as occurs throughout the visual system (Albright et al., 1984; Desimone et al., 1984). VFA sends topographic connections to both WMA and MA. WMA has persistent activity related to those visual features, such as is observed in parietal and prefrontal cortex (Fuster and Alexander, 1971; Miyashita and Chang, 1988), and sends topographic connections to MA. Afferents to MA from VFA arrive on AMPA receptors, while those from WMA arrive on NMDA receptors. These projections are balanced such that MA neurons respond only when a match occurs between a currently perceived stimulus and persistent activity in WMA. The responses of MA reflecting decision making in match-to-sample tasks can be observed in a wide range of cortical and subcortical areas (Curtis and Lee, 2010). Details of the connectivity within and between areas can be found in the Supplementary Material. In vertebrate brains, networks that correspond to these simulated areas may be distributed among a wide range of cortical and non-cortical areas, and network topology may be in the form of functional connectivity rather than the spatial location of cell bodies.

Neural model. We use a spiking neuron model (Izhikevich, 2007) that enables efficient computation for large-scale networks while accurately simulating single neuron membrane voltage and spike dynamics. Neural areas consist of 80% excitatory and 20% inhibitory neurons, and each cell type has parameters tuned to match its biological counterpart.

Neurons have conductance-based inputs from simulated AMPA, NMDA, GABA_A, and GABA_B receptors. Presynaptic spikes release a quantity of neurotransmitter onto postsynaptic receptors in proportion to the strength of the particular synapse. Excitatory (inhibitory) neurotransmitters affect AMPA and NMDA (GABA_A and GABA_B) receptors differentially, as determined by gains for each receptor type at that synapse. In the absence of presynaptic spikes, receptor activation dies away with a fixed time constant for each receptor type. NMDA current is gated by the voltage level inside the neuron. Synaptic strengths are modified by a short term plasticity rule.

Simulation experiments. We performed simulations using three different sets of visual inputs (Figure 2). Each set tested different aspects of network function:

- 1) Digits: 320x240 pixel JPEG images of the digits 0 through 9. This data set had the largest number of exemplars and tested discrimination among many visual stimuli. There were many sections of overlap in bright pixels between digits, c.f. the top of digits 2 and 3, and the right side of 3 and 8.
- 2) Natural images: Five 320x240 pixel JPEG images of natural scenes selected at random from Google Images and normalized for equal average luminosity. These images have much more complex spatial patterns of luminosity than the other data sets and therefore tested discrimination among complex stimuli.
- 3) J/mirror-J: 240x320 pixel video from an Axis 207MW wireless camera pointed at a set of wood blocks arranged in either a J shape or its mirror image, presented at different orientations. This data set has the most overlap in terms of pixels between categories: all images share half of their bright pixels with two other images, and a quarter of their pixels with three more images. In addition, the noise inherent in video capture added variability to induced spike trains.

During a simulation, connectivity among neurons was generated by randomly drawing from two-dimensional probability distributions. The random number generator was seeded with different values in each simulation, causing similar overall patterns of connectivity but different wiring in detail.

Synapses projecting to and within VFA are modified by a spike timing dependent plasticity learning rule (Song et al., 2000) during a training period at the beginning of each simulation. During this period all images in a data set were presented sequentially for 1 second each, and the sequence was repeated until 100 seconds had elapsed. During training a slow hyperpolarizing current in each VFA neuron helped ensure that the patches of neural activity generated by each stimulus were topographically well-separated from those of other stimuli.

After VFA training, the network was tested on a DMS task using the same image set presented during training. During a DMS trial, an image from the data set was presented for one second. After a one second delay period when no stimulus was presented, a second stimulus from the set was selected and presented for one second. A two second response period followed, during which no visual stimulus was presented. Any spikes in MA during this period indicate a “matching” response for that trial; otherwise a “non-matching” response was recorded. This process was iterated over

every permutation of images in the first and second presentation periods that are possible in the data set. Errors were recorded as either false positives (the two stimuli did not match, but a matching response was recorded) or false negatives (the two stimuli matched, but a non-matching response was recorded).

Analyses. We tested the persistence of patterns in WMA by locating the centroids of neural activity patches generated by stimuli. A correctly persisting WMA activity pattern had at least one activity patch whose centroid did not shift during the delay period by more than three times the distance between neighboring neurons. This criterion was empirically determined as the maximum shift that caused no overlap between any pair of stimulus patterns.

To measure WM capacity we simulated a set of 18 unique VFA neural activity patterns and tested whether those patterns were correctly stored as persistent activity in WMA. The 18 patterns, each consisting of two activity patches, were the maximum number of patterns that could exist in WMA without any overlap. Patterns were presented in random order to the network for one second each. After all patterns were presented, and an additional two second delay period had elapsed, we checked if WMA had correctly stored the patterns presented in VFA. A VFA pattern was correctly stored if we found a persistent activity patch in WMA whose location corresponded with that pattern (i.e., had not shifted as described above).

Results

Persistent activity generates working memory

When a trained network was presented with a visual stimulus, a pattern of neural activity was generated in VFA and WMA that reflected the identity of the stimulus (Figure 3A). After the stimulus was removed, the pattern of activity persisted in WMA but disappeared from VFA (Figure 3B).

Figure 4 shows that persistent activity in WMA remained unchanged during the presentation of a second stimulus, while activity in VFA changed to reflect the new stimulus. After presentation of a stimulus to be remembered, WMA holds the relevant pattern of activity for the rest of the trial, both during and after the presentation of a second stimulus. The firing rate of the persistent WMA activity slows briefly at the initial presentation of the second stimulus, but it returns to a higher firing rate even before stimulus offset. These results are consistent with observations of the difference between stimulus selective activity in prefrontal cortex and inferotemporal cortex (Miller et al., 1996). Although WMA eventually displays some persistent activity related to the second stimulus, this activity is lower in firing rate, involves fewer neurons, and appears after a longer latency than the activity from the first stimulus.

We examined the contribution of each of the three distinct mechanisms sustaining WMA activity: strong reentrant connectivity, NMDA receptor activation, and short term plasticity. To do so, we systematically varied the three simulation parameters controlling the strength of each mechanism, and tested how parameter changes affected persistent

activity. The parameters examined for excitatory neurons in WMA were long-term synaptic strength s , magnitude of short-term plasticity per-spike p , and the proportion of NMDA to AMPA receptors, $nmda_gain / ampa_gain$, while maintaining a constant amount of synaptic efficacy, $nmda_gain + ampa_gain$ (see Supplementary Material for details on how these parameters affect network dynamics). Each parameter was varied independently over 11 values resulting in 1331 simulations. During a simulation, an activity patch was stimulated at a random location in VFA for one second, and we recorded whether that patch persisted at the same location (using the centroid test described in Methods) after three additional seconds. Ten different patches were stimulated in each simulation.

Figure 5 shows the percentage of patterns that persisted from these simulations. On the “low synaptic strength” side of the parameter space (low long-term synaptic strength, low ratio of NMDA receptors to AMPA receptors, highly depressing short term plasticity), persistent activity does not initiate at all or slowly dissipates before the 3 second delay period is over. On the “high synaptic strength” side of the parameter space (high long term synaptic strength, high ratio of NMDA receptors to AMPA receptors, highly facilitating short term plasticity), the initial pattern of activity dissipates into various forms of traveling waves or epileptic, synchronous whole-network activity.

Supplementary Video 1 demonstrates the qualitatively different ways in which persistent activity fails, and Supplementary Figure 1 shows where in the parameter space the different types of failures occur.

Testing working memory in a visual delayed match-to-sample task

We tested the WM functionality of the network through a DMS task. In addition to WM, DMS tasks require the ability to distinguish stimuli from one another, to compare two stimuli, and to signal whether or not the stimuli match.

We investigated whether VFA activity was sufficient to distinguish among visual stimuli in the digits and natural image datasets. After a network was trained for a particular dataset, the patterns of VFA activity were unique for each stimulus in that dataset. These patterns were reliable: successive presentations of the same input rarely (3% out of 1500 total trials: 10 repetitions for each unique stimulus, 15 unique stimuli, repeated for 10 random seeds of network initial conditions) resulted in a VFA activity pattern that had shifted far enough to be confused with one generated by a different stimulus (see Methods for details). Therefore we conclude that VFA activity is sufficient to reliably distinguish stimuli from one another.

We constructed a network (MA) capable of signaling matches between currently perceived stimuli and persistent activity held in WMA while ignoring distractor stimuli. The segregation of MA afferents onto different receptors is the basis for this network’s ability to detect matches (see Discussion for relevant experimental support). Inputs from WMA arrive on NMDA receptors, while input from the VFA arrives on AMPA receptors. The first stimulus of a DMS task generates persistent activity in WMA, which accumulates large NMDA receptor activation on target neurons in MA during the delay

period. Because of the NMDA voltage gate this WMA input alone will not generate any firing in MA. Likewise, the visual input alone is insufficiently powerful to fire MA neurons. However, once AMPA receptor input depolarizes a match detection cell with activated NMDA receptors, that cell will begin to fire. In this way, MA will be activated when a VFA activity patch location overlaps that of a persistent activity patch in WMA. Any activity in MA is taken to indicate a match between the first and second stimuli in a trial (see Methods). Error rates for all three datasets were low when this network was applied to DMS tasks, as shown in Table 1.

Figure 6 shows examples of spike rasters during a typical DMS trial in which stimuli match. It also shows mean firing rate traces accumulated over all matching trials of the digits dataset. The figure shows that around 100 milliseconds after stimulus onset, VFA responded with a unique pattern of activity for each stimulus in the training set. By 500 milliseconds, WMA activity mirrored VFA activity. This WMA activity pattern persisted in the absence of the stimulus during the delay period and for the rest of the trial. MA activity began around 900 - 1200 milliseconds after the presentation of the second matching stimulus and persisted during the rest of the trial.

Capacity of the working memory network

A characteristic feature of working memory is its ability to store multiple items at the same time (Cowan, 2001). The mean number of items that can be stored simultaneously in WM is often referred to as WM capacity. It is clear from Figure 4B that our network is capable of simultaneously storing multiple patterns of persistent activity. Here we present simulations detailing how network properties relate to WM capacity.

We investigated how WM capacity (see Methods) was affected by the parameters of long-term synaptic strength, short-term plasticity, and the ratio of NMDA/AMPA receptors. The results of these simulations can be seen in Figure 7. Parameter choices that enabled the persistence of a high percentage of single patterns (see Figure 5) also had better capacity for multiple patterns (see Figure 7). In both Figure 5 and Figure 7, the parameter choices that enable persistent activity while also allowing long-term synaptic weight to vary widely are synaptic depression and NMDA/AMPA receptor ratios close to one.

We also examined the scaling of working memory capacity with the number of neurons in the working memory area. WMA networks with different numbers of neurons were tested; to ensure that firing dynamics remained constant, each had the same ratio of excitatory to inhibitory neurons and connectivity probability distribution functions. Simulations showed that capacity had a roughly log-linear relationship with the number of neurons in WMA, as shown by the black line in Figure 8. When the WMA network had sufficient numbers of neurons, mean capacity reached the size of the data set.

A scaling of capacity with neuron number would seem to be at odds with data showing that visual WM capacity is similar in species with very different brain sizes (Gibson et al., 2011; Buschman et al., 2011; Cowan, 2001). We were curious if a network could be constructed where capacity did not increase with the number of neurons. One solution

we found is shown in Figure 9: WMA was split into k subpopulations of m neurons each, where km equals the number of neurons in the single-population WMA network discussed so far in this paper. Each subpopulation was organized topographically, and neurons that shared the same position in different subpopulations would receive similar inputs. In this alternate model there was strong tightly-topographic mutual inhibition between subpopulations, but within a subpopulation the connectivity remained the same as in the previous model. Details of the multiple subpopulations can be found in the Supplementary Material.

The colored lines in Figure 8 show that in this alternate model WM capacity is not determined by network size. As the number of neurons in the whole network increases, capacity asymptotically approaches an upper limit. The number of neurons in a subpopulation, not the network, determines the limit. The limit arises as a result of inhibition between subpopulations: when an activity pattern persists in one subpopulation, inhibition suppresses activity in that neighborhood in the other $k-1$ subpopulations. Thus the k populations of m neurons respond more similarly to a single population of m neurons than to a single population of km neurons.

Although the subpopulation architecture is only one of many possible solutions for preventing linear scaling of capacity with network size, it is potentially illuminating to see what such connectivity would look like, were it to be observed in animals. This hypothetical account predicts that species with more neurons that function in working memory would have more subpopulations of neurons, yet those subpopulations would be of similar size across species. It would also predict more inhibitory connections overall in those species with more subpopulations.

Discussion

We are interested in testing hypotheses about the mechanisms behind the phenomena of working memory. Because real nervous systems are composed of large, degenerate (Edelman and Gally, 2001) networks operating via selectional principles (Edelman, 1987), we have built a large-scale neural network to investigate potential mechanisms of working memory.

This paper describes a spiking neural network that exhibits stimulus selective persistent activity after the stimulus is removed. In our neural model, stimulus selective persistent activity is the *sine qua non* of working memory. Although it is possible to conceive of mechanisms for maintaining a working memory trace without persistent activity during the delay period (Sugase-Miyamoto et al., 2008), such a mechanism would not influence other neural circuits during the delay period. Indeed, working memory is intimately tied to decision making (Baddeley, 2012), and it has been proposed that persistent activity is involved in decision making as well as the maintenance of the working memory trace (Curtis and Lee, 2010). Our network demonstrates a mechanism consistent with these observations, in which persistent activity both maintains the

memory trace and prepares match-detection neurons to respond in a visual delay match-to-sample task.

Spiking activity in simulated networks can become persistent through several mechanisms. Amit and Brunel (1997) demonstrated that strong connectivity in the local circuit can produce attractor dynamics. Mongillo et al. (2008) showed that short term synaptic plasticity enables persistent activity to be robust against brief drops of firing rate. Wang (1999) demonstrated that NMDA receptors allow a neuron to maintain excitatory current over time periods longer than the inter-spike intervals of its inputs. The network presented here incorporates all three mechanisms, and persistent activity was produced over a broad range of parameter values for implementing these mechanisms.

In real brains, synaptic strengths are highly variable, both among different circuits and in the same circuit over the course of time, and yet working memory must be maintained in spite of these variations. In our simulations, the combination of depressing short term plasticity and a near unitary ratio of NMDA to AMPA enables synaptic strength between excitatory neurons to vary over the largest possible range without affecting the persistence of neural activity. Interestingly, our current understanding of connections between pyramidal neurons in cortex includes depressing short term plasticity (Markram et al., 1998) and near unitary NMDA/AMPA ratio (Myme et al., 2003). Our results suggest that the known physiology of cortical short term plasticity and NMDA to AMPA receptor ratios might be critical for maintaining persistent activity despite the variability of synaptic strengths over time and between circuits.

While the view of working memory as intimately tied to decision making and behavior is prevalent, very few spiking models of working memory address decision making. One of these is the work of Compte et al. (2000), in which behavior was selected based upon an algorithmic read-out of persistent WM activity to perform an oculomotor delayed response task. Engel and Wang (2011) presented a variation of the same WM model that produced behavior in DMS tasks by integrating a neural mechanism based on competition between go and no-go circuits. The present paper describes a simple and effective neural mechanism for detecting a match between current perceptions and remembered stimuli: segregation of sensory and WM related inputs onto, respectively, AMPA and NMDA receptors. This segregation is consistent with the experimental evidence for 'silent' synapses where post-synaptic sites have only NMDA receptors (Isaac, 2003), and with the observation of systematic variations in AMPA and NMDA receptor prevalence at various points along the dendritic tree (Monaghan and Cotman, 1985; Nusser, 2000).

The network presented in this paper is able to sustain multiple patterns of neural activity at the same time, allowing us to explore working memory capacity limits in simulation. Previous spiking neural network models of working memory (Mongillo et al., 2008; Edin et al., 2009; Dempere-Marco et al., 2012) have also maintained persistent activity for multiple simultaneous stimuli. Edin et al. (2009) addressed the issue of capacity analytically for their model and show that it arises from lateral inhibition: as the number

of excitatory neurons involved in persistent activity increases, the amount of lateral inhibition also rises. Eventually elevated inhibition prevents persistent activity.

In our single-population working memory network, capacity scaled with the number of neurons in the working memory network. This is problematic in a model of working memory since animal species with very different number of neurons have very similar visual working memory capacities (Gibson et al., 2011; Buschman et al., 2011; Cowan, 2001). Although Edin et al. (2009) found that capacity for their network scaled with the fraction of neurons activated by a stimulus, under different assumptions their network would also display a similar scaling of capacity with network size. If one considers the number of neurons representing a stimulus to be fixed, while the number of neurons in the network increased, then the fraction of neurons representing a stimulus would effectively decrease and capacity would increase by their analysis. To our knowledge there is no experimental evidence establishing whether there is, across species, either a constant number or a constant proportion of neurons activated by a stimulus.

It is our belief that such a scaling will arise in any network where capacity is determined solely through lateral inhibition arising from persistent activity. Models where different processes limit capacity would not be subject to this effect, e.g. (Lisman and Idiart, 1995) as well as the multiple subpopulations network presented in this paper. Additionally, a different scaling of sensory coding capacity may arise if networks with lateral inhibition are employed in familiarity memory (Cortes et al., 2010).

While advances in neurophysiology and neuroanatomy are critical to understanding neural mechanisms, it remains technically difficult to observe the simultaneous spiking activity of more than ~100 neurons in a behaving animal. It is even more difficult to observe or infer the connectivity between neurons that are recorded. We therefore expect that further development of biologically-plausible, large-scale neural models such as this one will aid us in interpreting biological data and suggest new directions for investigating working memory.

Acknowledgements

We thank Alexandar E. Kozarev and Donald B. Hutson for their technical assistance. This work was supported in part by DARPA through ONR Grant N00014-08-1-0728 and by AFRL Cooperative Agreement FA8750-11-2-0255 to the Neurosciences Research Foundation. For support of late developments we are grateful to the Mathers Charitable Foundation. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Air Force Research Laboratory, the Department of Defense, or the U.S. Government.

Supplementary Material

Parameters in the spiking neuron equations were selected to model cortical pyramidal cells, basket cells, and thalamocortical neurons (see Supplementary Table 1). These simulated cells are arranged into one input (thalamocortical) area and three neural areas with similar CAS circuitry (Chen et al., 2013), consisting of 80% pyramidal and 20% basket cell types. Supplementary Table 2 describes the connections within and between neural areas.

Neuronal Dynamics. Spiking dynamics of each neuron were simulated using a computationally efficient phenomenological model (Izhikevich, 2007). The model has only 2 equations and 9 parameters that could be explicitly found from neuronal resting potential, input resistance, rheobase current, and other measurable characteristics. We present the model in a dimensional form: membrane potential is in millivolts, the current is in picoamperes and the time step is in milliseconds:

$$C\dot{v} = k(v - v_r)(v - v_t) - u - I_{syn}$$
$$\dot{u} = a\{b(v - v_r) - u\}$$

where C is the membrane capacitance, v is the membrane potential, v_r is the resting potential, v_t is the instantaneous threshold potential, u is the recovery variable (the difference of all inward and outward voltage-gated currents), I_{syn} is the synaptic current defined below, a and b are dimensionless parameters. When the membrane potential reaches the peak of the spike, i.e., $v > v_{peak}$, the neuron is said to fire a spike, and all variables are reset according to $v \leftarrow c$ and $u \leftarrow u + d$, where c and d are dimensionless parameters.

Supplementary Table 1 lists the parameters used to model different cell types. Cortical excitatory and inhibitory types populated neural areas VFA, WMA, and MA. Thalamocortical input neurons made up area VIA.

Short-Term Synaptic Plasticity. The strength of synapses varied as a function of the presynaptic neuron's firing history. We assume that the synaptic conductance (strength) of each synapse scales down (depression) or up (facilitation) on a short time scale (hundreds of milliseconds) by a multiplicative factor x . This factor, different for each presynaptic cell, is modeled by the following one-dimensional equation

$$\dot{x} = (1 - x)/\tau_x, \quad x \leftarrow px \text{ when a presynaptic neuron fires.}$$

x tends to recover to the equilibrium value $x = 1$ with the time constant τ_x , and it is reset by each spike of the presynaptic cell to the new value px . Any value $p < 1$ decreases x and results in short-term synaptic depression, whereas $p > 1$ results in short-term

synaptic facilitation. This equation was introduced in (Izhikevich and Edelman, 2008), and can be seen as a simplification of more detailed models, e.g. (Markram et al., 1998). In our model a given synapse can display either depression or facilitation but not both.

The parameters $\tau_x=150$ and $p=0.8$ were used on all excitatory and inhibitory connections within VFA and WMA. Within MA there is no short-term plasticity on connections originating from excitatory neurons and $\tau_x=150$ and $p=0.8$ on connections originating from inhibitory neurons. The connections from VIA to VFA had $\tau_x=150$ and $p=0.7$. The connections from VFA to MA and from WMA to MA had $\tau_x=150$ and $p=0.8$.

Synaptic Kinetics. The total synaptic current to each neuron is simulated as

$$I_{syn} = \left(g_{AMPA} + g_{NMDA} \frac{((v+80)/60)^2}{1+((v+80)/60)^2} \right) (v-0) + g_{GABA_A} (v+70) + (g_{GABA_B} + g_{SH})(v+90)$$

where v is the postsynaptic membrane potential, and the subscript indicates the receptor type. Each conductance g has first-order linear kinetics $g' = -g/\tau$ with $\tau = 5$ ms, 150 ms, 6 ms, 150 ms, and 15,000 ms for each of the simulated AMPA, NMDA, GABA_A, GABA_B, and SH receptors, respectively. The SH “receptors” were an *ad hoc* method for adding slow hyperpolarizing currents in order to bias cells that had already responded to one stimulus to remain off for a period of time, and thus improve pattern separation. This SH mechanism only operated in VFA for the first 100 seconds of a simulation when STDP organized neural responses to stimuli. In all other neurons, and in VFA neurons for $t > 100$ seconds, $g_{SH} = 0$.

Each firing of a presynaptic excitatory neuron increases g_{AMPA} by $ampa_gain \times s$ on the postsynaptic cell, where s is the long-term synaptic weight in nanoSiemens and x is the short-term depression/potential scaling factor as above; g_{NMDA} increases by $nmda_gain \times s$, where $nmda_gain / ampa_gain$ is the ratio of NMDA to AMPA receptors. The same computation is performed for inhibitory conductances, g_{GABA_A} , g_{GABA_B} , and g_{SH} , with their respective gains. The gain factor for g_{SH} was set to 0.3 for the first 100 simulation seconds and was set to zero for the remainder of the simulation.

Long Term Plasticity. Synaptic strengths of synapses between VIA and VFA excitatory neurons were modified using a spike timing dependent plasticity rule (Bi and Poo, 1998) for the first 100 seconds of each simulation. During this time the network was being sequentially shown each stimulus in the dataset, and each synaptic strength s was updated:

$$\dot{X} = -c/\tau_c + \alpha STDP(t) \delta(t - t_{pre/post})$$

$$X = c$$

where $\delta(t)$ is the Dirac delta function that step-increases the variable c . Firings of pre- and postsynaptic neurons, occurring at times t_{pre}, t_{post} , respectively, change c by the

amount $\alpha STDP(t)$ where α is the learning rate for the synapse, $t = t_{post} - t_{pre}$ is the interspike interval, and

$$STDP(t) = \begin{cases} A^+ \exp(-1/\tau^+)^t, & t > 0 \\ A^- \exp(-1/\tau^-)^{|t|}, & t \leq 0 \end{cases}. \quad (7)$$

where $A^+ = 0.005$, $A^- = 0.001$, $\tau^+ = \tau^- = 20$ ms. The variable c decays to zero exponentially with the time constant $\tau_c = 1$ s, and s is updated once every 50 ms for computational efficiency.

Synaptic scaling. Synaptic scaling is a biologically-observed mechanism (Turrigiano et al., 1998), which allows a neuron to maintain the same amount of total synaptic input while redistributing input strength among synapses during plasticity. Synaptic scaling was performed for each neuron to maintain the total of all synaptic strengths arriving on a given neuron, syn_{total} , at a constant value. This scaling was performed for every neuron every 50 ms during the simulation. Each synapse was prevented from exceeding syn_{max} or going below zero, regardless of learning rules and normalization.

Network connections. The local excitation and longer range inhibition structure described below is characteristic of the CAS architecture (Chen et al., 2013). Anatomical (Goldman-Rakic, 1995; Kisvárdy et al., 2000; Holmgren et al., 2003; Fino and Yuste, 2011) as well as functional evidence (Kaschube et al., 2010; Haider et al., 2010; Derdikman et al., 2003) exists for such an architecture in the cortex. However, some evidence points to inhibition acting in a strictly local fashion, e.g. (Hirsch and Gilbert, 1991), and models have addressed the effect of changing the reach of inhibition on neural dynamics and stimulus coding (Compte et al., 2003; Cortes et al., 2012). Unpublished experiments indicate that surround inhibition is necessary for our network to generate stable patches of activity that support WM.

The connectivity between model neurons fell into two classes: either local-type or surround-type. For local-type connectivity, a two-dimensional Gaussian probability distribution, centered on each postsynaptic cell, determines the probability of forming a synapse between each potential presynaptic neuron within a specified maximum distance, r_{max}

$$f(d) = ae^{-\frac{(d-\mu)^2}{2\sigma^2}}$$

where a is a scale factor set to generate, on average, a target number of synapses on each postsynaptic cell, d is the distance between the presynaptic neuron and the postsynaptic neuron, μ is 0, and σ is the standard deviation. In a similar manner, a two-dimensional Gaussian function was used to specify the synaptic strength between connected neurons as a function of the distance between them in the network. The total of all synaptic efficacies was scaled in order to sum to a constant parameter with units in nanoSiemens. Thus both connection probability and strength were maximal between nearest neighbors, and fell off as a function of distance, controlled by the same parameter, the standard deviation of a Gaussian.

For surround-type connectivity, a postsynaptic neuron receives synaptic connections from neurons located in a surrounding annular region specified by a minimum (r_{min}) and maximum (r_{max}) radial distance from the postsynaptic cell. The probability of forming a connection with a neuron in the annulus is determined as a function of distance from the postsynaptic cell. The function used is a Gaussian with standard deviation $\hat{\sigma}$, centered at $\hat{\mu}=(r_{min}+r_{max})/2$. This probability distribution function is scaled to create a target number of synapses for each postsynaptic neuron. The synaptic strengths for the surround-type connection are initialized using the same function. The sum of all synaptic strengths was scaled to make the sum equal to a constant value under experimenter control, syn_{total} .

In order to avoid boundary conditions in the network, the network was treated as a torus. Thus connections from neurons that would otherwise go outside of the network instead “wrap around” to connect with neurons on the opposite edge. Euclidean distance between neurons (on the torus) determined axonal conduction delays. Inside a neural area, the maximum conduction delay occurred between neurons that were the maximum possible distance apart, d_{max} . The delay on a given synapse between two neurons in the same area distance d apart was $5\lceil d/d_{max} \rceil$ milliseconds. Between neural areas, the delay was $5\lceil d/d_{max} \rceil + delay_{ave}$, where d was calculated as if the postsynaptic neuron’s location was in the presynaptic neural area. The value $delay_{ave}$ was 5ms from VIA to VFA, 10ms from VFA to WM, 15ms from VFA to MA, and 20ms from WM to MA.

Supplementary Table 2 lists the parameters defining the anatomy and synaptic parameters of the network as used in the DMS tasks. The connectivity definition for a postsynaptic neuron type spans multiple rows, one row for each possible presynaptic neuron type. The first four columns of the table list the type of neuron, the area in which it is located, the number of neurons in the population, and the total number of synapses per neuron. The next columns specify presynaptic areas and cell types, along with the percentage of the postsynaptic cell’s synapses allocated to this pathway. The minimum radius (r_{min}), maximum radius (r_{max}), and sigma columns ($\hat{\sigma}$) specify the connectivity parameters in units of Euclidean distance, such that 1 indicates the distance between two neighboring neurons in the presynaptic neural group. Percentage noise indicates how much random uniform variability there is in the initial synaptic weights. syn_{total} is the value used in the synaptic scaling operation for the maximal total conductance per postsynaptic neuron. syn_{max} is the maximum value of any single synapse. The gain columns indicate how synaptic strength is distributed to post-synaptic receptors as described previously.

Similar networks with different sized WMA (500 – 42,000 neurons) were used in some simulations. In those cases VFA and MA remained the same size, and topographic projections between neural areas were made based on relative position within the area’s 2D grid of neurons. Local connectivity within neural areas remained the same.

Alternative WMA architecture. Some simulations used a form of the working memory network that was divided up into k subpopulations of m neurons each, such that km

equals the number of neurons in WMA. Each subpopulation had the same ratio of the number of excitatory neurons to inhibitory neurons as the previously described single-population WMA network. Each subpopulation had local connections within the subpopulation as previously described for the connections inside the single-population WMA network. Inputs from VFA remain as previously described as well, but go to all k subpopulations. Each subpopulation i , inhibits all other subpopulations $j \neq i$ with tightly topographic projection from WMA _{i} inhibitory to WMA _{j} excitatory neurons using the parameters $r_{min} = 0$, $r_{max} = 2.1$, $\hat{r} = 1.4$, percent noise=0, $syn_{total} = 150$, $syn_{max} = 5$, GABAa gain=1, and GABAb gain=0. Each of these $k-1$ inhibitory projections adds on average 14 synapses per WMA excitatory neuron.

Visual input. Video was recorded with an Axis 207MW wifi camera using 320x240 pixel, 8 bit greyscale images at 30fps. The central portion of the video frames were used to drive the 21x21 retinotopic grid of input neurons in the VIA. Each VIA neuron received current injection proportional to the grayscale intensity of pixels in its local receptive field, and the centers of adjacent neurons' receptive fields mapped onto adjacent image pixels. The current injection function of each neuron was based on a model of on-center/off-surround retinal ganglion cells (Wohrer and Kornprobst, 2009), with a 3x3 on-center area and a 6x6 off-surround area. These currents were constantly

injected at each numerical integration step until the next video frame was received.

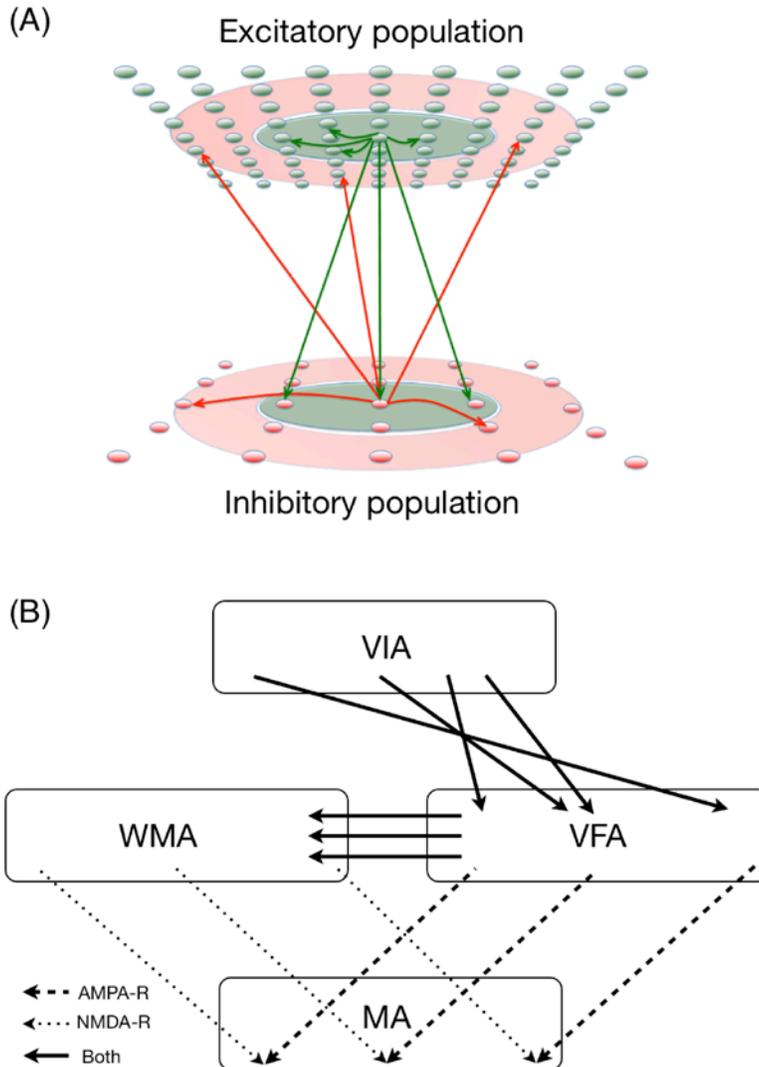


Figure 1

Figure 1: Architecture of the neural model. (A) Local network connectivity. Each neural area consists of excitatory neurons (green dots) and inhibitory neurons (red dots) arranged on a 2D grid. For each kind of connection a representative presynaptic neuron is shown surrounded by an overlay indicating the postsynaptic neurons to which it may connect. The color of the overlay indicates the kind of presynaptic neuron. The actual connections made between neurons are dictated by random draws from probability density functions. (B) Inter-area connectivity. The visual input area (VIA) simulates retinotopic spiking activity with firing rates proportional to the local brightness of the image in an on-center off-surround fashion. The visual features area (VFA) receives random (non-retinotopic) connections from VIA, and sends a topographic projection to a working memory area (WMA). WMA and VFA both send topographic projections to the match detection area (MA), but to different receptor types (AMPA or NMDA).

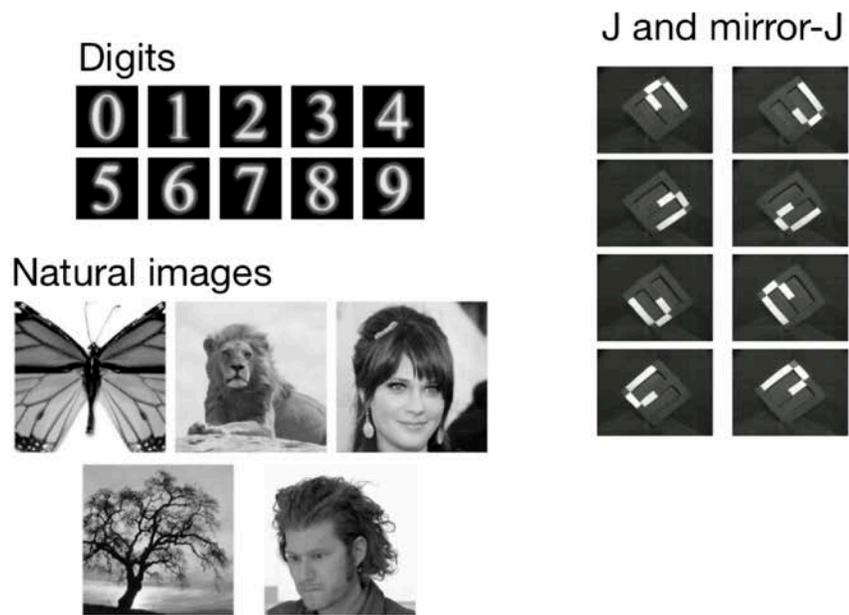


Figure 2

Figure 2. The digits, natural images, and J/mirror-J stimuli sets that were used in the simulations.

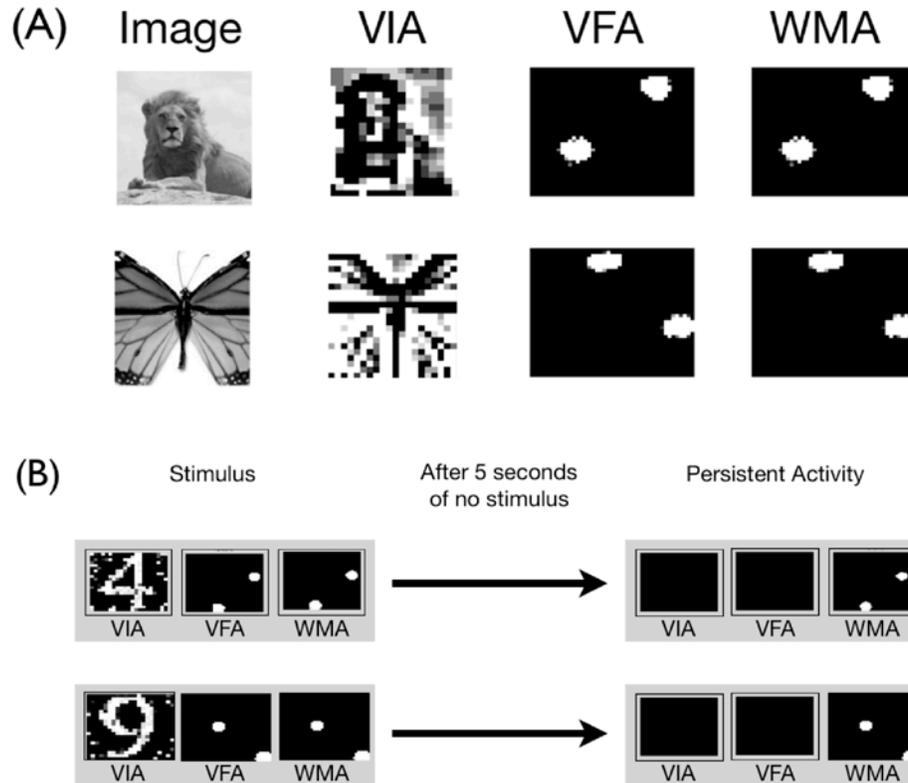


Figure 3

Figure 3. Examples of neural activity in response to stimuli. Each neural activity plot shows, at a particular time, pixels arranged to represent the firing rate of neurons on a 2D grid. Pixels with brighter luminance indicate neurons with higher firing rate. (A) Neural activity after one second of stimulus presentation. Visual input neurons (VIA) respond in an on-center off-surround way to visual inputs. In visual features (VFA) and working memory (WMA) areas, each stimulus generates a unique pattern of activity. Stimulus patterns consist of two patches or attractors formed by the CAS architecture of the network. (B) As before a stimulus in VIA generates a unique pattern of activity in VFA and WMA. After stimulus offset the pattern persists in WMA, whereas VIA and VFA become silent.

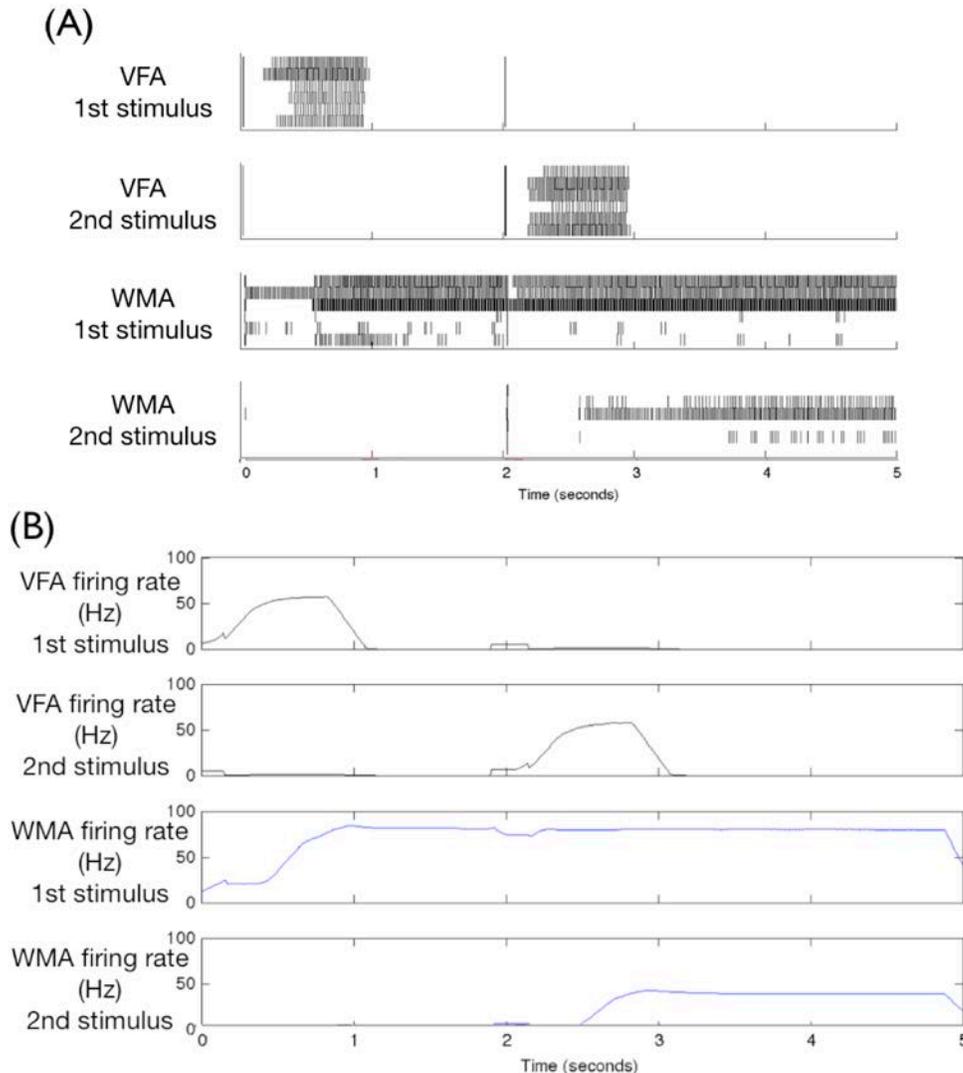


Figure 4

Figure 4. Examples of rasters and mean-firing rate traces of persistent working memory activity in the presence of a second, distracting stimulus. (A) Spike rasters of neurons that, after training on the digits data set, responded selectively to the first stimulus presented in the trial (top rastergram and 3rd from top) and second stimulus presented in the trial (2nd from top and 4th from top). For each rastergram we randomly selected six neurons that responded to the first time the stimulus was presented in a sequence of trials. This raster shows those same neurons on a subsequent trial, where the first stimulus is presented from 0-1 seconds, and the second stimulus from 2-3 seconds. WMA neurons responding to the first stimulus have persistent activity even when the second stimulus is presented. WMA neurons responding to the second stimulus do not respond as strongly, and those weaker responses occur later after stimulus onset. (B) Mean firing rate traces over all neurons responding to the digits data set. For each possible non-matching stimulus pair, the activity of all neurons that respond to the relevant stimulus were averaged together. Firing rates (in Hz) are calculated over a 200 ms window centered on the given time slice.

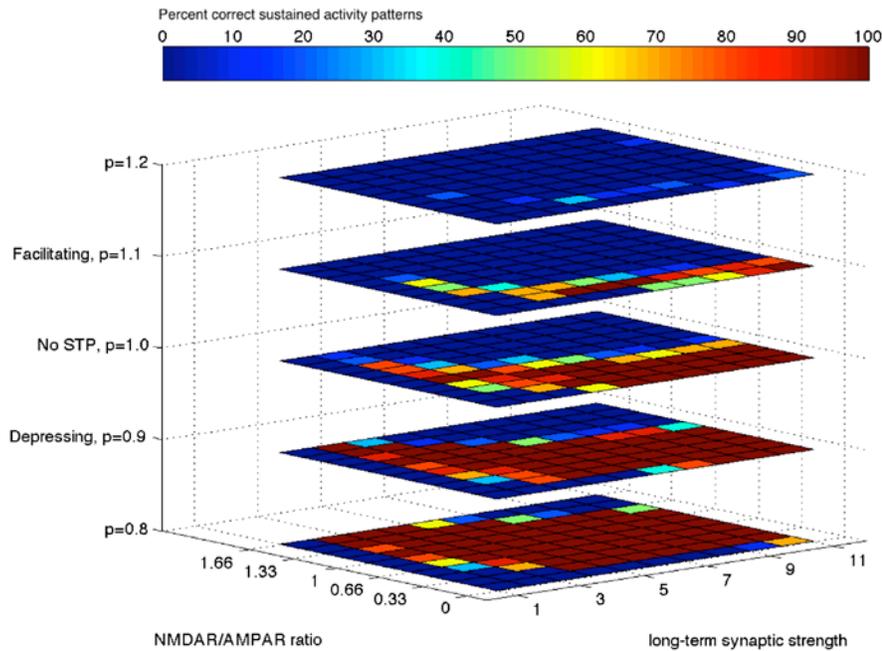


Figure 5

Figure 5. The reliability of persistent activity is affected by network parameters. The parameters varied were long-term synaptic strength, s , of local reentry between excitatory neurons, the ratio of simulated NMDA to AMPA receptors on excitatory neurons $nmda_gain/ampa_gain$, and short term plasticity (STP) parameter p for connections originating from excitatory neurons. When $p < 1$ presynaptic spikes produce synaptic depression, when $p = 1$ there is no STP, and when $p > 1$ presynaptic spikes produce synaptic facilitation. Each 2D slice shows the variation of synaptic strength (x-axis) vs NMDA/AMPA ratio (y-axis) for a single value p (z-axis). The colors of graph squares indicate the percentage of WM patterns that were still firing correctly 3 seconds after the stimuli initiating them were gone. Synaptic strength is able to vary over the widest range without affecting persistent activity when excitatory-excitatory synapses are depressing and have physiologically observed values of NMDA/AMPA ratio.

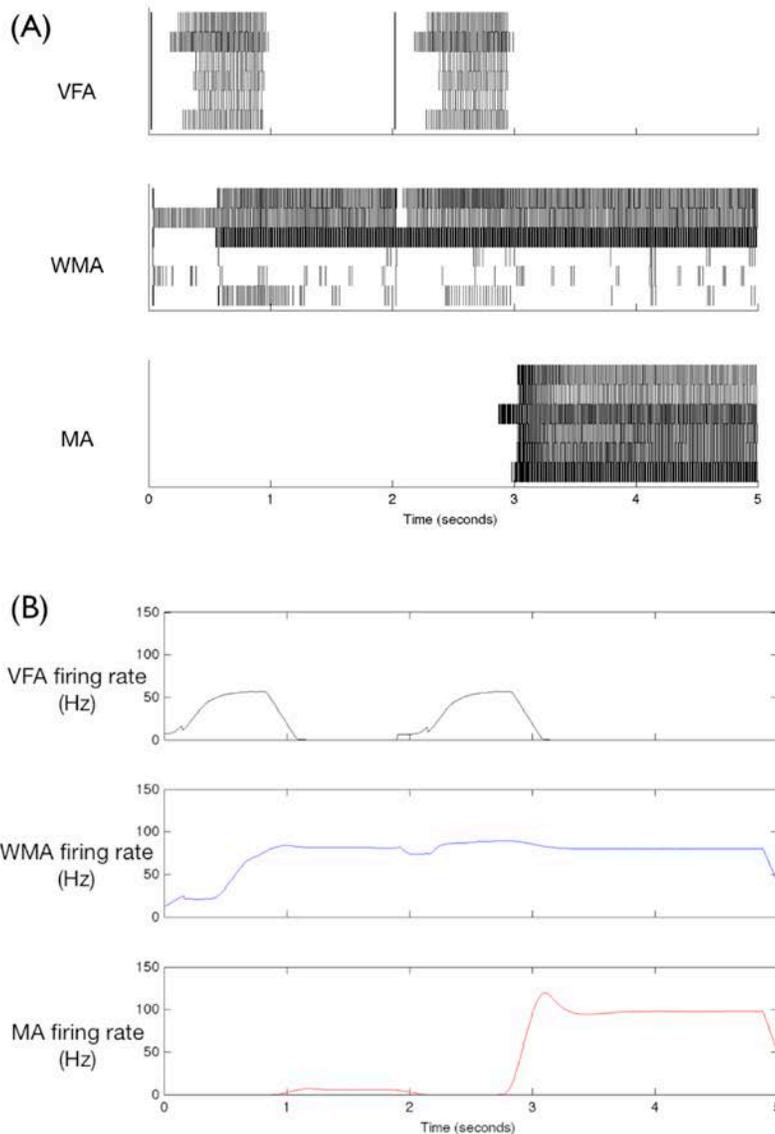


Figure 6

Figure 6. Time course of neural activity during matching DMS trials with the digits data set. (A) Spike rasters of randomly selected neurons responsive to the stimulus “4” from a correct match trial. In each neural region we randomly selected six neurons that responded to the very first presentation of the stimulus during training. This raster shows those same neurons on a subsequent trial where the same stimulus is presented in both the first (0-1 seconds) and second (2-3 seconds) presentation periods. WMA neurons responding to this stimulus are active during the delay period (1-2 seconds) and throughout the matching period (3-5 seconds). MA neurons begin to respond only towards the end of the second stimulus presentation. (B) Mean firing rate over all neurons responding to match trials in the digits data set. For each of the 10 match trials (0-0, 1-1, etc.), the responses of all neurons that respond to the given stimulus were averaged together. Firing rates are calculated over a 200 ms window centered on the given time slice with zero padding before and after the duration of the trial.

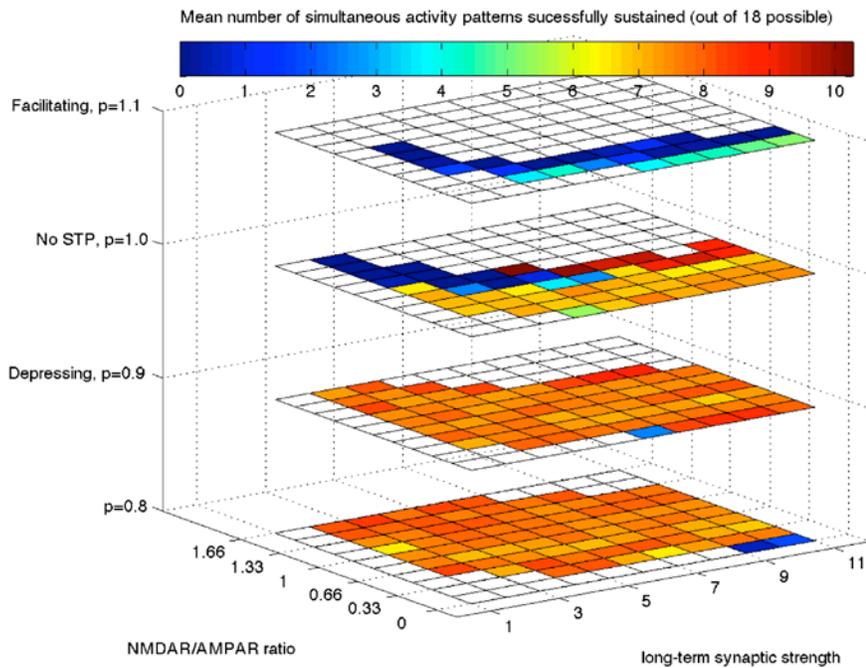


Figure 7

Figure 7. WM capacity is affected by network parameters. The parameters varied were baseline synaptic strength of local reentry between excitatory neurons, the NMDA/AMPA receptor ratio on excitatory neurons, and short term plasticity (STP) parameter p for connections originating from excitatory neurons. When $p < 1$ presynaptic spikes produce synaptic depression, when $p = 1$ there is no STP, and when $p > 1$ presynaptic spikes produce synaptic facilitation. Each 2D slice shows the variation of synaptic strength (x-axis) vs NMDA/AMPA ratio (y-axis) for a single value p (z-axis). The colors of graph squares indicate the percentage of WM patterns that persisted after being presented sequentially with 18 patterns for one second each, followed by a two second period of no stimulus. Blank squares represent parameter combinations that were not run because there were no successful single patterns stored for that parameter combination in the simulations of Figure 5. For each parameter combination, 10 networks were created from different random seeds, and each network was presented with the 18 patterns 10 times in a different random order each time.

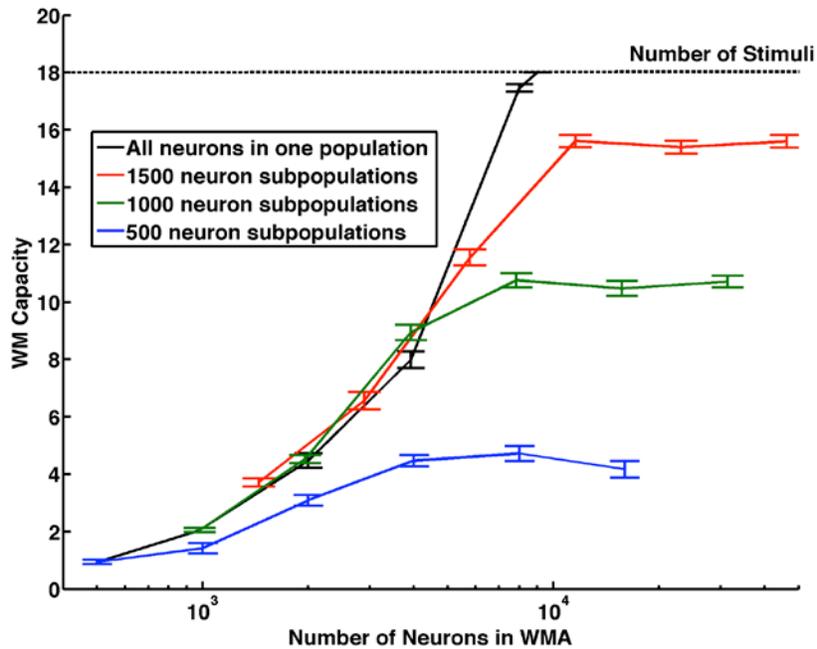


Figure 8

Figure 8. The relationship between WM capacity, number of neurons in WMA, and the size of subpopulations. The capacity of simultaneous WM storage increases with the number of neurons in WMA when all neurons are in a single population. The single population WMA reaches perfect capacity (all 18 stimuli) with sufficient network size. However, if the same number of neurons are arranged into multiple subpopulations that mutually inhibit each other in a topographic fashion, then capacity rises to some asymptotic limit, where the limit is proportional to the number of neurons in a subpopulation. X-axis values reflect the total number of neurons (excitatory and inhibitory) in WMA. Y-axis values reflect mean capacity over 100 trials (10 seeds, 10 random permutations of order each seed) with 95% confidence interval error bars.

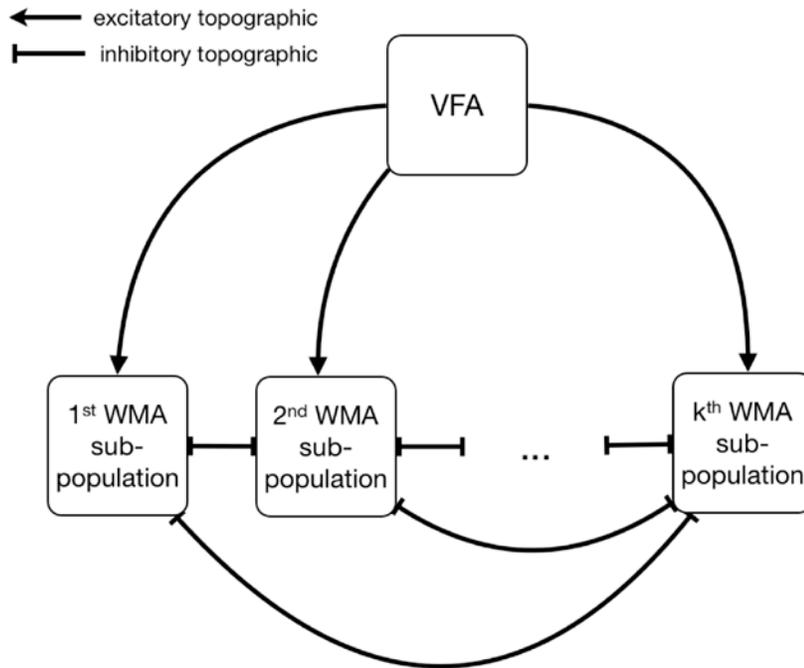


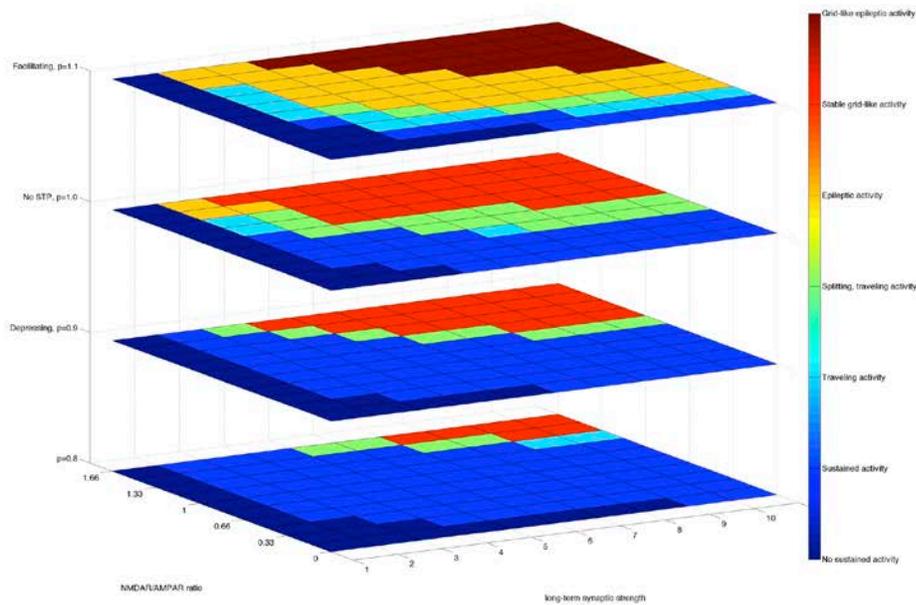
Figure 9

Figure 9. Alternate architecture where WMA is broken into k subpopulations, each of which has m neurons. Each subpopulation inhibits all others with tightly topographic projections. WMA neurons receive inputs from VFA as dictated by topographic location, regardless of subpopulation. Inside each subpopulation local connectivity is the same as previously described for the entire single population WMA network.

Data Set	Correct	False Positives	False Negatives
J/mirror-J 10 seeds x 8 stimulus #1 x 8 stimulus #2	93.8% (600/640)	5.5% (31/560)	11.3% (9/80)
Digits 10 seeds x 10 stimulus #1 x 10 stimulus #2	99.1% (931/1000)	7.6% (68/900)	1.0% (1/100)
Natural Images 10 seeds x 5 stimulus #1 x 3 presentations / stimulus x 5 stimulus #2 x 3 presentations / stimulus	94.6% (2240/2250)	4.9% (89/1800)	7.3% (33/450)

Table 1

Table 1. Matching and error rates for the three visual data sets in a DMS task. For each result, the table records the percentage (raw number out of maximum possible in parentheses) of each outcome possibility: correct match or non-match between trial stimuli, false positive match, or false negative non-match. Each data set was run for 10 different neural networks, with each network initialized by a different seed for the random number generator.



Supplementary Figure 1

Supplementary Figure 1. Neural activity dynamics when WM related activity fails to persist. After an activity patch was stimulated for one second, the network was given no stimulation for a further five seconds. In the fifth second the dynamics of neural activity were evaluated by the experimenter. Activity either (1) did not persist, (2) successfully persisted and remained at the topographic location where it was stimulated, (3) persisted and traveled across the network topology, (4) traveled and split into multiple patches that continued to travel, (5) split into multiple traveling patches that recursively split again until the entire network was filled with epileptic activity, (6) split into multiple traveling patches that became stable unmoving patches in a grid-like arrangement, or (7) split into multiple traveling patches that became stable in location but displayed epileptic activity. All of these behaviors are illustrated in Supplementary Video 1.

Neuron Type	C	k	Vr	Vt	Vpeak	a	b	c	d
Excitatory	80	3	-60	-50	50	0.01	5	-60	10
Inhibitory	20	1	-55	-40	25	0.15	8	-55	200
Input	200	1.6	-60	-50	40	0.01	15	-60	10

Supplementary Table 1

Supplementary Table 1. Parameter values for neural dynamics.

Post Neuron Area	Post Neuron Type	# Neurons	Average # Synapses per Post Neuron	Pre Neuron Area	Pre Neuron Type	Percentage of Total Synapses	Minimum Radius (neurons in pre area)	Maximum Radius (neurons in pre area)	Sigma (neurons in pre area)	Percentage Noise	Syn Total (nS)	Syn Max (nS)	AMPA Gain	NMDA Gain	GABAA Gain	GABAB Gain	
VFA	Exc.	1849	1054	VFA	Exc.	1.2	0.0	2.2	1.1	0	22	10	1	0.5	0	0	
				VFA	Inh.	57	1.5	15.0	12.0	0	1200	20	0	0	1	0.1	
	Inh.	900	779	V/A	Input	41.8	0.0	15.1	105.0	50	200	50	1	0.5	0	0	
				VFA	Exc.	20.4	0.0	7.1	3.4	0	100	20	1	0.5	0	0	
WMA	Exc.	3136	134	VFA	Inh.	76.4	1.5	15.0	12.0	0	180	30	0	0	1	0.1	
				V/A	Input	3.2	0.0	42.0	105.0	50	10	10	1	0.5	0	0	
	Inh.	784	300	WMA	Exc.	44	0.0	4.2	2.0	0	75	45	1	0.5	0	0	0
				WMA	Inh.	31.4	2.1	4.2	4.7	0	600	30	0	0	1	0.1	
				VFA	Exc.	24.6	0.0	3.2	1.5	0	7	7	1	0.5	0	0	
				WMA	Exc.	85	0.0	9.2	4.5	0	80	2	1	0.5	0	0	
MA	Exc.	3136	672	WMA	Inh.	15	2.1	4.2	4.7	0	120	3	0	0	1	0.1	
				MA	Exc.	3.6	0.0	2.8	0.8	0	22	10	1	0.5	0	0	
	Inh.	784	1235	MA	Inh.	90.6	1.4	14.0	4.7	0	900	10	0	0	1	0.1	
				VFA	Exc.	2.2	0.0	2.2	1.1	0	0.65	1	1	0	0	0	
				WMA	Exc.	3.6	0.0	2.8	1.4	0	0.65	1	0	1	0	0	
				MA	Exc.	20.6	0.0	9.2	4.5	0	15	2	1	0.5	0	0	
Exc.	3136	672	MA	Inh.	49.5	1.4	14.0	4.7	0	180	3	0	0	1	0.1		
			VFA	Exc.	15	0.0	31.0	215.0	0	10	2	1	0.5	0	0		
Exc.	3136	672	WMA	Exc.	14.9	0.0	40.3	280.0	0	10	2	1	0.5	0	0	0	

Supplementary Table 2. Anatomical and synaptic parameters used in DMS simulations.

Bibliography

- Albright, T. D., Desimone, R., and Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *J Neurophysiol* 51, 16–31.
- Amit, D. J., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex* 7, 237–252.
- Baddeley, A. (2012). Working memory: theories, models, and controversies. *Annual Review of Psychology* 63, 1–29.
- Bechara, A., Damasio, H., Tranel, D., and Anderson, S. W. (1998). Dissociation Of working memory from decision making within the human prefrontal cortex. *J Neurosci* 18, 428–437.
- Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18, 10464–10472.
- Bisley, J. W., Zaksas, D., Droll, J. A., and Pasternak, T. (2004). Activity of neurons in cortical area MT during a memory for motion task. *J Neurophysiol* 91, 286–300.
- Buschman, T. J., Siegel, M., Roy, J. E., and Miller, E. K. (2011). Neural substrates of cognitive capacity limitations. *Proc Natl Acad Sci USA* 108, 11252–11255.
- Chen, Y., McKinstry, J. L., and Edelman, G. M. (2013). Versatile networks of simulated spiking neurons displaying winner-take-all behavior. *Frontiers in Computational Neuroscience* 7, 16.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X. J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex* 10, 910–923.
- Compte, A., Sanchez-Vives, M. V., McCormick, D. A., and Wang, X.-J. (2003). Cellular and network mechanisms of slow oscillatory activity. *J Neurophysiol* 89, 2707–2725.
- Cortés, J. M., Greve, A., Barrett, A. B., and van Rossum, M. C. W. (2010). Dynamics and robustness of familiarity memory. *Neural Computation* 22, 448–466.
- Cortés, J. M., Marinazzo, D., Series, P., Oram, M. W., Sejnowski, T. J., and van Rossum, M. C. W. (2012). The effect of neural adaptation on population coding accuracy. *J Comput Neurosci* 32, 387–402.
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci* 24, 87–114; discussion 114–85.
- Curtis, C. E., and Lee, D. (2010). Beyond working memory: the role of persistent activity in decision making. *Trends Cogn Sci* 14, 216–222.
- Dempere-Marco, L., Melcher, D. P., and Deco, G. (2012). Effective Visual Working Memory Capacity: An Emergent Effect from the Neural Dynamics in an Attractor Network. *PLoS ONE* 7, e42719.
- Derdikman, D., Hildesheim, R., Ahissar, E., Arieli, A., and Grinvald, A. (2003). Imaging spatiotemporal dynamics of surround inhibition in the barrels somatosensory cortex. *J Neurosci* 23, 3100–3105.

- Desimone, R., Albright, T. D., Gross, C. G., and Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4, 2051–2062.
- Edelman, G. M. (1987). *Neural Darwinism: The Theory of Neuronal Group Selection*. The MIT Press.
- Edelman, G. M., and Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proc Natl Acad Sci USA* 98, 13763–13768.
- Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegnér, J., and Compte, A. (2009). Mechanism for top-down control of working memory capacity. *Proc Natl Acad Sci USA* 106, 6802–6807.
- Engel, T. A., and Wang, X.-J. (2011). Same or Different? A Neural Circuit Mechanism of Similarity-Based Pattern Match Decision Making. *J Neurosci* 31, 6982–6996.
- Ferrera, V. P., Rudolph, K. K., and Maunsell, J. H. (1994). Responses of neurons in the parietal and temporal visual pathways during a motion task. *J Neurosci* 14, 6171–6186.
- Fino, E., and Yuste, R. (2011). Dense inhibitory connectivity in neocortex. *Neuron* 69, 1188–1203.
- Fuster, J. M., and Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science* 173, 652–654.
- Gibson, B., Wasserman, E., and Luck, S. J. (2011). Qualitative similarities in the visual short-term memory of pigeons and people. *Psychon Bull Rev* 18, 979–984.
- Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron* 14, 477–485.
- Haider, B., Krause, M. R., Duque, A., Yu, Y., Touryan, J., Mazer, J. A., and McCormick, D. A. (2010). Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron* 65, 107–121.
- Hirsch, J. A., and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat's visual cortex. *J Neurosci* 11, 1800–1809.
- Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol (Lond)* 551, 139–153.
- Isaac, J. (2003). Postsynaptic silent synapses: evidence and mechanisms. *Neuropharmacology* 45, 450–460.
- Izhikevich, E. M. (2007). *Dynamical systems in neuroscience*. The MIT Press.
- Izhikevich, E. M., and Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proc Natl Acad Sci USA* 105, 3593–3598.
- Kaschube, M., Schnabel, M., Löwel, S., Coppola, D. M., White, L. E., and Wolf, F. (2010). Universality in the evolution of orientation columns in the visual cortex. *Science* 330, 1113–1116.
- Kisvárdy, Z. F., Crook, J. M., Buzás, P., and Eysel, U. T. (2000). Combined physiological-anatomical approaches to study lateral inhibition. *Journal of Neuroscience Methods* 103, 91–106.
- Lisman, J. E., and Idiart, M. A. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science* 267, 1512–1515.
- Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical

- pyramidal neurons. *Proc Natl Acad Sci USA* 95, 5323–5328.
- Miller, E. K., and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202.
- Miller, E. K., Erickson, C. A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16, 5154–5167.
- Mishkin, M., and Pribram, K. H. (1956). Analysis of the effects of frontal lesions in monkey. II. Variations of delayed response. *J Comp Physiol Psychol* 49, 36–40.
- Miyashita, Y., and Chang, H. S. (1988). Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature* 331, 68–70.
- Monaghan, D. T., and Cotman, C. W. (1985). Distribution of N-methyl-D-aspartate-sensitive L-[3H]glutamate-binding sites in rat brain. *J Neurosci* 5, 2909–2919.
- Mongillo, G., Barak, O., and Tsodyks, M. (2008). Synaptic theory of working memory. *Science* 319, 1543–1546.
- Myme, C. I. O., Sugino, K., Turrigiano, G. G., and Nelson, S. B. (2003). The NMDA-to-AMPA ratio at synapses onto layer 2/3 pyramidal neurons is conserved across prefrontal and visual cortices. *J Neurophysiol* 90, 771–779.
- Nusser, Z. (2000). AMPA and NMDA receptors: similarities and differences in their synaptic distribution. *Curr Opin Neurobiol* 10, 337–341.
- Passingham, R. (1975). Delayed matching after selective prefrontal lesions in monkeys (*Macaca mulatta*). *Brain Res* 92, 89–102.
- Romo, R., and Salinas, E. (2003). Flutter Discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4, 203–218.
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3, 919–926.
- Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., and Richmond, B. J. (2008). Short-Term Memory Trace in Rapidly Adapting Synapses of Inferior Temporal Cortex. *PLoS Comput Biol* 4, e1000073.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896.
- Wang, X. J. (1999). Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J Neurosci* 19, 9587–9603.
- Wohrer, A., and Kornprobst, P. (2009). Virtual Retina: a biological retina model and simulator, with contrast gain control. *J Comput Neurosci* 26, 219–249.
- York, L. C., and van Rossum, M. C. W. (2009). Recurrent networks with short term synaptic depression. *J Comput Neurosci* 27, 607–620.

LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

APE-X	Name given to a humanoid BBD constructed at The Neurosciences Institute
BBD	Brain-based device
CAS	Center-annular-surround (type of network connectivity)
COTS	Commercial off-the-shelf
DMS	Delayed match-to-sample
STDP	Spike-timing-dependent (synaptic) plasticity
UAV	Unmanned aerial vehicle
UAGV	Unmanned aerial-ground vehicle
UGV	Unmanned ground vehicle
WM	Working memory
WTA	Winner-take-all (type of network and population behavior)