

Virtual neurorobotics (VNR) to accelerate development of plausible neuromorphic brain architectures

Philip H. Goodman^{1,*}, Sermsak Buntha², Quan Zou¹ and Sergiu-Mihai Dasca²

1. Department of Medicine and Program in Biomedical Engineering, University of Nevada, Reno, USA
2. Department of Computer Science and Engineering, University of Nevada, Reno, USA

Edited by: Frederic Kaplan, Ecole Polytechnique Federale De Lausanne, Switzerland

Reviewed by: Angelo Cangelosi, University of Plymouth, United Kingdom
Jeffrey L. Krichmar, The Neurosciences Institute, San Diego, United States of America
Felix Schürmann, Ecole Polytechnique Federale De Lausanne, Switzerland

Traditional research in artificial intelligence and machine learning has viewed the brain as a specially adapted information-processing system. More recently the field of social robotics has been advanced to capture the important dynamics of human cognition and interaction. An overarching societal goal of this research is to incorporate the resultant knowledge about intelligence into technology for prosthetic, assistive, security, and decision support applications. However, despite many decades of investment in learning and classification systems, this paradigm has yet to yield truly “intelligent” systems. For this reason, many investigators are now attempting to incorporate more realistic neuromorphic properties into machine learning systems, encouraged by over two decades of neuroscience research that has provided parameters that characterize the brain’s interdependent genomic, proteomic, metabolomic, anatomic, and electrophysiological networks. Given the complexity of neural systems, developing tenable models to capture the essence of natural intelligence for real-time application requires that we discriminate features underlying information processing and intrinsic motivation from those reflecting biological constraints (such as maintaining structural integrity and transporting metabolic products). We propose herein a conceptual framework and an iterative method of virtual neurorobotics (VNR) intended to rapidly forward-engineer and test progressively more complex putative neuromorphic brain prototypes for their ability to support intrinsically intelligent, intentional interaction with humans. The VNR system is based on the viewpoint that a truly intelligent system must be driven by emotion rather than programmed tasking, incorporating intrinsic motivation and intentionality. We report pilot results of a closed-loop, real-time interactive VNR system with a spiking neural brain, and provide a video demonstration as online supplemental material.

Keywords: neurobotic architecture, human robot interface, virtual reality, artificial intelligence, social robotics, epigenetic robotics, reinforcement learning, neocortex, mesocircuit

INTRODUCTION

Traditional research in artificial intelligence and machine learning has viewed the brain as a specially adapted information-processing system. More recently the field of social robotics has been advanced to capture the important dynamics of human cognition and interaction (Dautenhahn, 2007; Scheutz et al., 2007). An overarching societal goal of this research is to incorporate the resultant knowledge about intelligence into technology for prosthetic, assistive, security, and decision support applications. However, despite many decades of investment in learning and classification systems, this paradigm has yet to yield truly “intelligent” systems. For this reason, many investigators are now attempting to incorporate realistic neuromorphic properties into machine learning systems, encouraged by over two decades of neuroscience research that has yielded quantitative parameters which characterize the brain’s interdependent electrophysiological (Markram et al., 1997; Schindler et al., 2006), genomic (Toledo-Rodriguez et al., 2004), proteomic (Toledo-Rodriguez et al., 2005), metabolomic and anatomic (Wang

et al., 2006) networks. For example, a search of the ISI Web of Knowledge for publications whose abstract contained words related to *in vivo* or *in vitro* neocortical or hippocampal research increased about 150-fold in the period from 1985 to 2005 (from 18 in 1985, 1494 in 1995, to 2689 in 2005). Directly warehoused data collection motivated by the highly automated genomic projects, such as the Allen Brain Atlas (Allen Institute, 2007), is further accelerating the growth of publically available data. The outpouring of potentially useful data has sparked the development of over 100 neuroscience databases (Society for Neuroscience, 2007).

Knowledge about intelligence systems may be translated into technology for prosthetic, assistive, security, and decision support applications. At the present time, prosthetic devices are limited to interfaces from sensory organs to the cortex (such as cochlear implants which stimulate the peripheral auditory nerves) and, more recently, from the output regions of the neocortex to electromechanical limbs for conscious control of movement (Jensen and Rousche, 2006) in patients with spinal cord injury or limb loss. The ability to understand processing within the brain would facilitate the development of implantable neuromorphic (biomimetic) chips (Berger and Glanzman, 2005), for example, to bridge regions of the brain damaged or disconnected by stroke or head trauma, or to detect and avert seizure propagation. Assistive technologies include navigational robotic devices for persons with movement disabilities (Boy et al., 2007), and household, office or industrial services that require human-like judgment (Bien and Lee, 2007). Security applications include autonomously functioning surveillance systems (Macera et al., 2004; Liu

* Correspondence: Philip H. Goodman, Department of Medicine and Program in Biomedical Engineering, University of Nevada, Reno, USA. e-mail: goodman@unr.edu

Received: 03 Sep. 2007; paper pending published: 06 Oct. 2007; accepted: 09 Oct. 2007; published online: 02 Nov. 2007

Full citation: Frontiers in Neurorobotics (2007) 1:1 doi: 10.3389/neuro.12/001.2007
Copyright: © 2007 Goodman, Buntha, Zou, Dasca. This is an open-access article subject to an exclusive license agreement between the authors and the Frontiers Research Foundation, which permits unrestricted use, distribution, and reproduction in any medium, provided the original authors and source are credited.

et al., 2005), and robots that can perform independently in exploratory (Visentin and van Winnendael, 2006), clean-up, law-enforcement, or military (Carlson and Murphy, 2005) environments that would otherwise be hazardous to humans (Seward et al., 2007). Artificial brain systems with sufficient emotional (Breazeal, 2003, 2004; Fellous and Arbib, 2005), intentional, as well as abstract knowledge-based intelligence could also be configured to assist in decision making for scenarios involving resource allocation under competing demands, such as industrial and business economics, urban planning, and geopolitical conflicts.

Given the complexity of neural systems, developing tenable models to capture the essence of natural intelligence for real-time application requires that we discriminate features underlying information processing and intrinsic motivation from those reflecting biological constraints (such as maintaining structural integrity and transporting metabolic products). Furthermore, despite the large and increasing number of physiological parameters provided by experimental inquiry, most of the data relates either to the very small scale of individual or small groups of neurons (e.g., intracellular, 2-photon, or unit recordings at discrete recording sites), or at the other extreme, the joint effect of thousands or millions of neurons over millimeter (optical imaging) or centimeter fields (fMRI and PET). Thus, the architecture and response patterns at the middle scale, or “mesocircuit”, remain largely uncharacterized, requiring that the brain modeler proposes and systematically tests plausible connection patterns and learning dynamics. Mammalian brains contain from 10 million (mouse) to 100 billion (human) neurons (Braitenberg, 2001). The use of digital simulation, even with the aid of hundreds or thousands of clustered processing units (Frye et al., 2006), is very limited in its capacity to model the dynamics of neural systems, for which a 10th or 100th of a millisecond precision may be needed for accuracy. Some groups have reported success in navigational tasks using neuromorphic architectures (Banquet et al., 2005; Krichmar et al., 2005; Ogata et al., 2004; Wiener and Arleo, 2003; Cuperlier et al., 2005). Only a few groups have reported simulations on the order of 1 million simplified neural elements (Izhikevich et al., 2004; Ripplinger et al., 2004) using supercomputer clusters, and even these are orders of magnitude away from real-time operation. Approaches such as

expansion to greater numbers of processors incur delays due to switching latency among the processing boards, a phenomenon not seen in the inherently parallel connectionism of biological networks. Adding realistically branching multicompartmental neurons with active synapses and channels (Maciokas et al., 2005) further encumbers digital simulations.

Although we recognize that present technological and neuroscientific limitations may not enable researchers to replace conventional with neuromorphically driven learning systems in the near term, we propose herein a technique of virtual neurorobotics to rapidly forward engineer and test progressively more complex putative neuromorphic brain prototypes for their ability to support intelligent, intentional interaction with humans. Successful prototypes could then be efficiently instantiated in hardware robotics embedded in real world scenarios.

Definition of Virtual Neurorobotics. We define virtual neurorobotics as follows: a computer-facilitated behavioral loop wherein a human interacts with a projected robot that meets five criteria: (1) the robot is sufficiently embodied for the human to tentatively accept the robot as a social partner, (2) the loop operates in real time, with no pre-specified parcellation into receptive and responsive time windows, (3) the cognitive control is a neuromorphic brain emulation incorporating realistic neuronal dynamics whose time constants reflect synaptic activation and learning, membrane and circuitry properties, and (4) the neuromorphic architecture is expandable to progressively larger scale and complexity to track brain development, (5) the neuromorphic architecture can potentially provide circuitry underlying intrinsic motivation and intentionality, which physiologically is best described as “emotional” rather than rule-based drive. A summary of the requirements for a VNR system is shown in Table 1.

High-level social robotic systems reported to-date are generally controlled by artificial intelligence and machine learning algorithms that incorporate explicit task lists and criteria for task satisfaction, segmenting time into periods of action and of awaiting response. Our interest is not to characterize the rules of social engagement *per se*, but rather to uncover the basis of biological brain sensorimotor control, information processing and learning. The corresponding neuromorphic brains must therefore be driven intrinsically by a motivational influence such that the dynamics that subserve information processing are themselves affected by a drive to accomplish the tasks (with neural learning that reinforces successful behavioral adaptation) (Samejima and Doya, 2007; Schweighofer et al., 2007). The motivational system must therefore demonstrate intentionality, which means that the intelligent system takes into account the “aboutness” of its own relationship to other behaving entities (and vice versa) in its environment. With sufficiently complex neuromorphic architectures, intentionality would be expected to be reflected by frontal and parietal mirror neuron responsiveness characteristic of many mammalian intentional behaviors (Iacoboni and Dapretto, 2006). This combined physiological responsiveness of intrinsic motivation and intentionality in animals, including humans, most generally can be described as *emotion*:

“Emotion, in its most general definition, is a complex psychophysical process that arises spontaneously, rather than through conscious effort, and evokes either a positive or negative psychological response and physical expressions, often involuntary, related to feelings, perceptions, or beliefs about elements, objects, or relations between them, in reality or in the imagination.” <http://wikipedia.org>

Emotion, as the name suggests, sets in motion the moment-to-moment behaviors of an intelligent system (Breazeal, 2003; Frijda, 2006). From this perspective, “intelligence” has evolved as a way to better serve emotional drive. That is, intelligence may be a derivative of emotion, rather than vice versa. We therefore make the following hypothesis: the development of a truly intelligent artificial system cannot occur outside the real-time, emotional interaction of humans with a neuromorphic system. This does not mean that intelligent systems, once refined, cannot ultimately be cloned (at a point in development where they are ready to learn advanced tasks). Rather, to grow the early intelligent systems we

Table 1. Components of a closed-loop neuromorphic brain development.

Functional (SCRIPT)	
1.	Computation and communication provides nearly real-time robotic response
2.	Repertoire of ROBOT behavior is commensurate with its physical and brain complexity
3.	Time is <i>not</i> segmented a priori for ROBOT or ACTOR receptiveness or reaction
4.	ACTOR must assume ROBOT has ability to perceive and respond meaningfully
5.	ACTOR must respond as he/she would in similar real world circumstances
Environment (SCENE)	
1.	Realistic contents, including sights, sounds, obstacles
2.	Movement of background objects apropos to scenario
3.	May be altered or affected by the actions of the ROBOT or ACTOR
4.	May include other ROBOTs or multiple ACTORs
Live Participant (ACTOR)	
1.	Human, child or adult, depending on type of intelligence targeted
2.	Suitability or willingness to attribute intentionality to the ROBOT
Neuromorphic System (ROBOT)	
1.	Central nervous subsystem (BRAIN)
A.	Neocortex
B.	Hippocampus
C.	Basal ganglia
D.	Other limbic regions relating to attention, reward, and fear
E.	Biologically plausible learning algorithms
F.	Progressively more complex architecture as constrained by VNR
2.	Interpretive (rule-based) & communication subsystems (BRAINSTEM)
3.	Embodiment in virtual VNR architecture
4.	Internal (virtual) sensory capabilities (proprioception, balance)
5.	External sensory capabilities (vision, hearing, touch)
6.	Facial and body social expressive and gestural capabilities
7.	Dexterous movements (upper extremities for humanoid robots)
8.	Translational movements (lower extremities)

The system is comprised of functional SCRIPTing requirements with three main components: SCENE, ACTOR, and ROBOT. Terms in parentheses are used in the text and Figures.



must start with minimalist brain architectures that demonstrate intrinsic motivation and intentionality in scenarios requiring intelligent behavior in a real-world context. This recapitulates the way in which humans develop cognitive function over the first several years of social experience. With the VNR approach, we seek not only to “grow” such intelligent systems but also to comprehend, at each step, the differential changes in architecture giving rise to novel and intelligent cognition.

MATERIALS AND METHODS

Behavioral scenario

For the demonstration of basic VNR principles, we chose an instinctual “friend vs. foe” response, wherein a resting dog responds to movement in its visual field with either (1) a cautious growl while remaining in a lying position, (2) threatening bark while sitting up, or (3) happy breathing and tail wagging while fully standing. A human actor is told that he/she is visiting a home with a dog unknown to him/her. As shown in **Figure 1**, a robotic dog is projected onto the forward screen, with external sensors that enable its simulated brain to “see” and respond to the actor’s movements, in the context of a background scene projected onto the rear screen (for this demonstration, we used a static image of a suburban neighborhood).

The simple neuromorphic brain used in our demonstration has only four spiking regions, which respond differentially to the orientation of edges moving in the visual field (see details in 2.2–2.4 below). The actor is told in advance that moving vertically oriented objects (including body parts) will pose a threat to the robot, whereas moving horizontally oriented objects will be perceived as friendly gestures. For example, walking toward the robot will cause predominantly vertical edge movements as the body looms

to fill the camera’s field, as will waving an arm or bat in a striking manner. Offering the hand or a horizontally held object like a bone will trigger a horizontal edge response. The actor finds that the robot stands and barks in response to threatening (vertical) movements, and must decide whether to move sideways or back in a fearful manner, to freeze, or to move toward the robot using some sort of friendly (horizontal) gestures to gain a favorable response (standing, wagging, and breathing in anticipation) from the dog. The actor finds that movements that do not consistently (for at least 50 ms) evoke horizontal or vertical edge response are interpreted by the robot with a warning growl in the lying position. The VNR has no set duration, and the actor can choose any sequence of responses, as he/she might perform with a real dog. Although not incorporated in the demonstration scenario, a suitably configured neuromorphic brain model could also allow the actor to reward the dog for selected behavior patterns (e.g., using voice patterns or petting the touch pad).

ROBOT subsystem

We used Webots 5 (Cyberotics, Lausanne, Switzerland) to render the pseudo-3D Sony AIBO robot with Open Dynamics Engine for accurate physics simulation, and programmed and controlled the movement sequences using the URBI parallel-scripting engine (Gostai, Paris, France). We used URBI to program basic motor sequences, such as lying down, sitting, standing, tail wagging, and opening/closing mouth. Audio recordings of growling, barking, and breathing dog sounds (Tradebit, Wilmington, Delaware) were programmatically linked to the motor sequences for synchronous output. The virtual system is provided with two types of sensory input: internal and external. Internal sensation includes proprioception

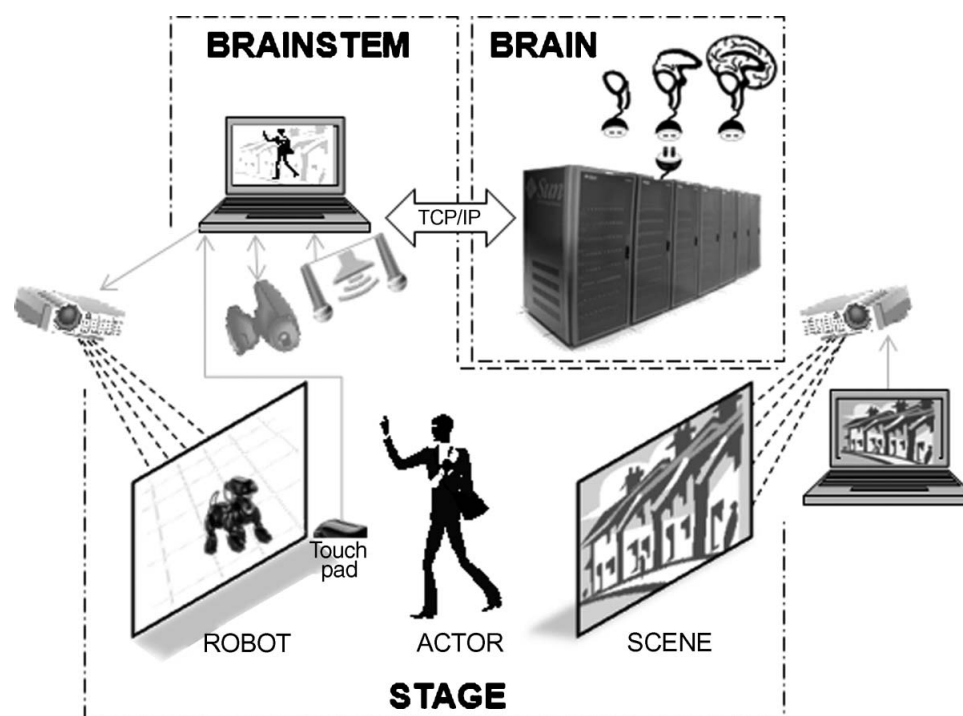


Figure 1. Schematic cartoon of a fully implemented virtual neurobotic (VNR) system. VNR substitutes a pseudo-3D screen projection for the physical robot, which participates in real-time interplay with the human actor. The robot’s eyes (pan-tilt-zoom camera) and ears (monaural or spaced stereo microphones) capture the actor’s movements and voice in the context of the background scene, which is projected independently (and may contain moving elements, including other animals or actors). The BRAINSTEM is a multiprocessor computer (running threads) that synchronously (1) captures and preprocesses video images, sound, and touch, (2) converts preprocessed sensory images into probabilities of spiking for each primary neocortical region, (3) uploads the spike probability vectors to the BRAIN simulator, (4) accepts, from the BRAIN simulator motor neuron region, output spike density vectors and triggers corresponding dominant motor sequences (e.g., sitting, lying, barking, walking) via the robotic simulator program (Webots/URBI), which makes the corresponding changes in behavior of the projected robot (and incorporates internal sensation such as proprioception and balance). The BRAIN simulator is a neuromorphic modeling program running on a supercomputer, executing a pre-specified spiking brain architecture, which can adapt as a result of learning (using reward stimuli offered by the ACTOR’s voice or stroking of the touch pad). Based on successful performance, researchers iteratively “plug in” alternative or more complex brain architectures.

(joint and axial body position) and senses of momentum and balance. External sensation uses situated microphones (ears), cameras (eyes), and surrogate skin surfaces to represent touching by the actor. For visual frame grabbing, a Sony EVI-D70 pan-tilt-zoom (PTZ) camera was used at a rate of 10 per second using a Cyberoptics (Beaverton, OR) PXC200AL board and the Linux EVILib graphics library. For our demonstration, the camera was set at fixed PTZ coordinates because there was minimal displacement of the AIBO robot in this scenario; in general the PTZ is driven by the robot simulator software such that the camera tracks the robots net eye and head position, and zoom is adjusted as the robot moves forward or backwards in its virtual coordinates.

BRAINSTEM subsystem

The BRAINSTEM (Peng, 2006) is a locally designed, multithreaded C++ program that handshakes at every simulation time point with the BRAIN simulator (Figure 2). JPEG images are captured to a shared hard drive and sequentially read and filtered by 160 × 160 Gabor filters in four orientations: vertical, horizontal, and the two diagonals. A medium space-constant width was chosen empirically. The difference between sequential images is used by BRAINSTEM to create spike probability vectors based on orientation and position in the visual field, representing receptive field properties of neurons in early visual processing in cortical area V1 (Olshausen and Field, 1996). At each time step, BRAINSTEM uploads these vectors to the BRAIN simulator, and then waits for a response from the BRAIN simulator which supplies a report of the indexes of cells firing as a result of the input at the current time step. BRAINSTEM tracks firing rates of the corresponding motor cells, and if one dominates the other consistently for a 50 ms period, the corresponding pre-set robotic motor sequence instruction is sent to the ROBOT simulator. For our demonstration, we used URBI to program four motor sequences: lying in wait (with or without growling), sitting up and barking, and standing with tail-wagging and excited breathing.

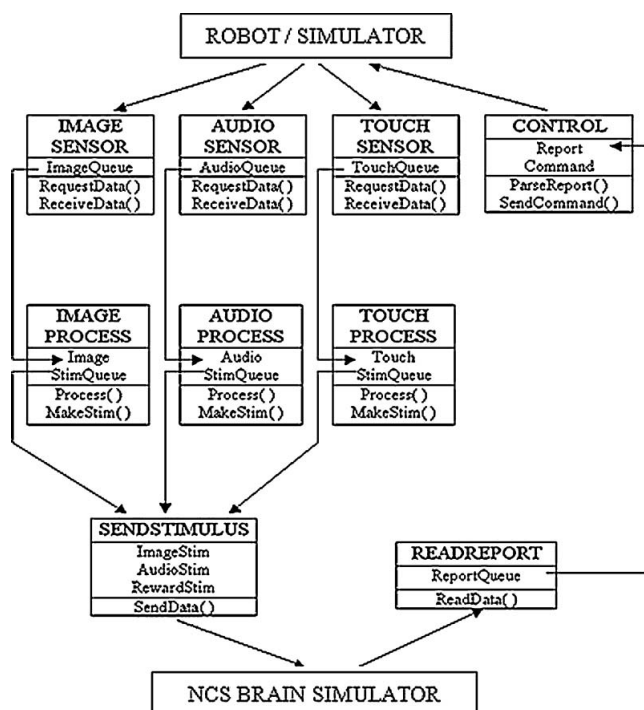


Figure 2. Multi-threaded pipeline organization of the BRAINSTEM. Each sensory or report modality is assigned its own thread in a self-blocking queue. Data are read from a hard drive shared with sensory capture software. Outputs back to the robotic system and to the NCS brain simulator are sent by TCP/IP port routing. Documentation is available at <http://brain.unr.edu>.

BRAIN neuromorphic subsystem

In principle, any brain simulator that allows iterative two-way port-based message passing with the BRAINSTEM subsystem could be used in the VNR loop. We used a locally developed C++ program called the NeoCortical Simulator (NCS), which runs on any LINUX cluster; for a recent review of spiking neural simulators, including NCS, see Brette et al. (2007). NCS emulates clock-based integrate and fire neurons whose compartments contain conductance-based synaptic dynamics and Hodgkin-Huxley formulations of ionic channel gating particles (the code is freely available at <http://brain.unr.edu>). Neuronal compartments, which may include an arbitrary number of cellular compartments, are allocated in 3-D space, and are connected by forward and reverse conductances without detailed cable equations. Synapses are conductance-based, with phenomenological modeling of depression, facilitation, augmentation, and Hebbian STDP. We run NCS on a 200-CPU hybrid of Pentium and AMD processors. Although NCS was motivated by the need to model the complexity of the neocortex and hippocampus, limbic and other structures can be modeled by variably collapsing layers and specifying the relevant 3-D layouts. Common firing patterns are obtained using combinations of membrane ion channels (Maciokas et al., 2005). NCS delivers reports on any fraction of neuronal cell groups, at any specified interval. Reports include membrane voltage (current clamp mode), current (voltage clamp), spike-event-only timings (event-triggered), calcium concentrations, synaptic dynamics parameter states, and any Hodgkin-Huxley channel parameter. Although NCS can be run in a batch mode, for interactive robotics we use an Internet protocol port-based input-output mode which handshakes with the BRAINSTEM subsystem at every sampled time point.

As a demonstration of VNR principles, we programmed a simple neuromorphic architecture consisting of 64 single-compartment neurons divided into four columns representing pre-motor regions (precursors to coordinated behavioral sequences), each connected to one of the visual field preferences based on Gabor filter configurations. According to the probability vector received from BRAINSTEM, NCS injected short (1 ms) step current (3 nA) pulses sufficient to reach the threshold of -50 mV and generate a single spike. The membrane voltages were sampled and updated at a frequency of 1000 Hz. Because the images were grabbed at 10 Hz, the upload probability vector from each pair of Gabor-filtered images was used repeatedly for 100 ms intervals of neural simulation.

RESULTS

Figure 3 describes a typical set of pre-motor brain region action potential rasters recorded over the course of a 10-second VNR interaction with a human.

A typical VNR run is demonstrated in the video available at <http://brain.unr.edu/VNR>. A sampling of the findings is shown in Figure 4. Running only on four processors of the cluster, the small size of the neuromorphic brain enabled NCS to interact in real time with the BRAINSTEM subsystem. The BRAINSTEM in turn was able to capture and perform Gabor filtering on the images acquired at 10 Hz with no backlog. Any delays or hesitation we observed in the responses of the ROBOT were attributable to latency in communication across TCP/IP ports.

DISCUSSION

We described and developed an iterative method of virtual neurobotics (VNR) intended to rapidly forward engineer and test progressively more complex putative neuromorphic brain prototypes for their ability to support intrinsically intelligent, intentional interaction with humans. We reported pilot results of a closed-loop, real-time VNR system with a spiking neural brain, and provided a video demonstration as online supplemental material. These methods and results are presented within a conceptual framework for continued development of neuromorphic brain architectures using VNR.



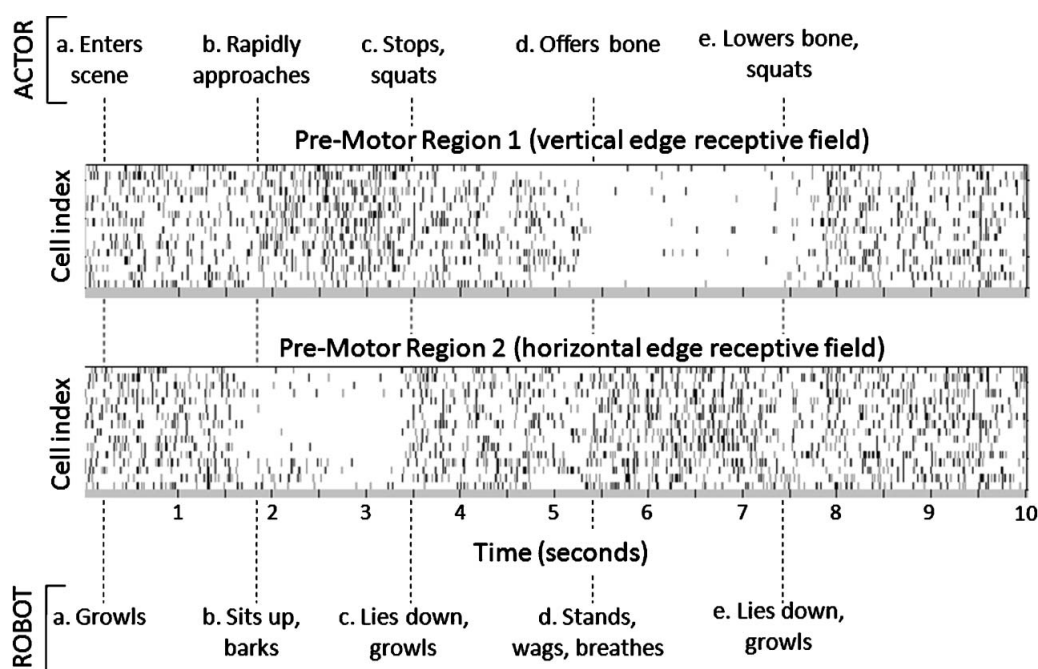


Figure 3. ACTOR-BRAIN-ROBOT interplay. Ten-second behavior scenario indicating timing of ACTOR (upper row) and ROBOT (lower row) events. The ACTOR in this scenario was free to choose any sequence of movements in response to perceived intent of the ROBOT. ROBOT behavioral sequences are triggered when the neuromorphic BRAIN output to BRAINSTEM has 50 ms of consistent spiking in one pre-motor region compared with another. Periods without domination of one pre-motor region over another trigger the ROBOT to lie down and growl. In cell rasters, each row represents the timing of action potentials (spikes) of a single neuron; darker gray markers indicate clustered bursts of spikes.

Comparison with other social robotic approaches

The VNR approach echoes the importance of social embeddedness in the stepwise, ontological development of robotic cognition, as popularized by researchers starting in the mid-1990s (Breazeal and Scassellati, 2000;

Brooks et al., 1998), and continuing with epigenetic robotics emphasis on the interplay of cognitive and perceptual brain systems (Lungarella and Berthouze, 2002; Schlesinger, 2003). Almost all of these models use a combination of psychological production rules, fitness functions, and hier-

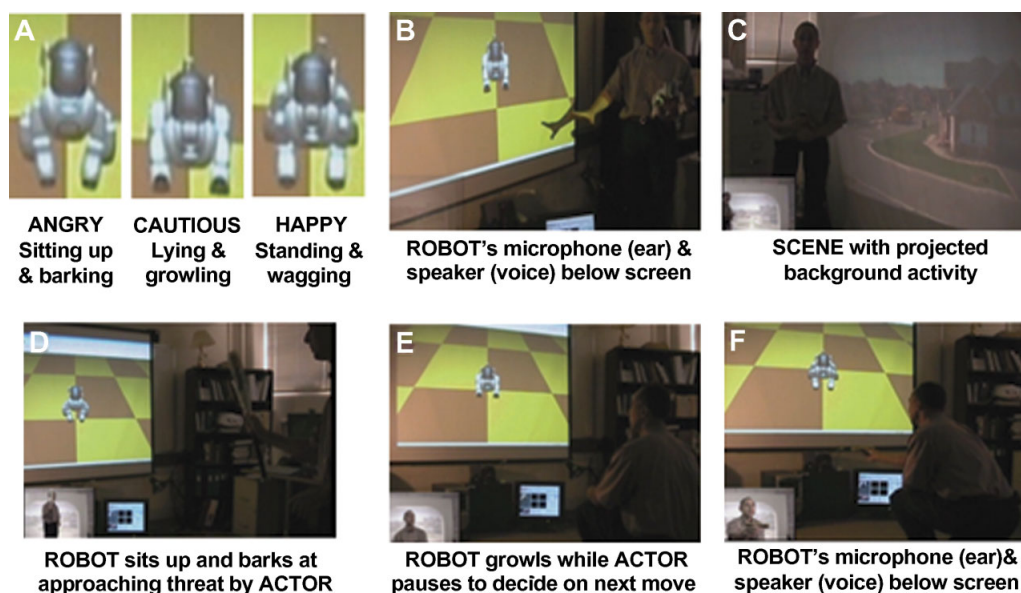


Figure 4. Demonstration of VNR interaction. The system is comprised of three main components, SCENE, ACTOR, and ROBOT, as viewed by an external observer. (A) Three major behaviors of the ROBOT. (B) Positioning of ROBOT's external sensory devices. (C) Background scene consisting of suburban neighborhood. (D–F) VNR loop (ROBOT-eye view superimposed in lower left corner): (D) ACTOR approaches with threatening behavior (bat) and ROBOT responds by sitting up and barking. (E) ACTOR responds by lowering bat and squatting down, then ROBOT responds by lying and growling a warning. (F) ACTOR offers dog bone and ROBOT responds by standing and wagging tail. Online video, <http://brain.unr.edu/VNR/VNRdemo.avi>.

archical machine learning algorithms to map behavior to robotic cognition. This is the case even for models intended to capture neuronal epiphenomena such as mirror neuronal activity (Triesch et al., 2007). Our approach differs in three ways. First, our focus is to understand brain physiology at the “mesocircuit” level, using social robotics to constrain the many possible architectures bridging the well-characterized measurements at the cellular level (e.g., patch clamp and unit recordings) and those at the scales of millions of cells (e.g., optical and fMRI imaging). By maintaining biological plausibility, we hope ultimately to facilitate technology transfer to the fields of neuroprosthetics, assistive and security robotics, and human-like support systems for geopolitical decision-making. Second, the stipulation of neuromorphic architecture precludes us from using production rules or hierarchical algorithms as psychological models; any assumptions on behavioral triggering, motivation, and intentionality must arise from the tissue models themselves, and learning from behavioral reinforcement must manifest as synaptic change. Third, the simulation must respect the actual distribution of physiological time constants that characterize membranes, channels, and synapses; otherwise, there would not be temporal coherence (i.e., social interaction) between the simulated robotic brain and that of the human actor. Operating progressively large neural simulations in real time presents a difficult computational challenge, which VNR ameliorates by allowing flexible prototyping of robotic embodiment and scripting of specific social interactions in variable contexts. The rationale for the virtual paradigm is similar to that put forward by Krichmar and Edelman (2005) for robotic instantiation of brain-based devices.

Obstacles and considerations

Although the use of human actors would seem to present an obstacle in terms of time, resources, and variability, we do not have any other “gold standard” for spontaneous, emotionally intelligent interaction. After all, it is human-level cognition that we seek to elucidate in our modeling and applications. Further, the parameters for neuronal membranes, channels, synapses are provided as time constants on the order of milliseconds to seconds, as co-optimized by evolution. A system that emulates connected neurons but operates at the temporal scale of microseconds, for example, cannot interact with the slower responses of humans. Thus, both the joint distribution of known biological time constants and the need for emotionally intelligent responses warrant the use of a closed-loop interaction of brain prototype with an actor. Another consideration might be to use animals in place of humans; however animals, which rely on many subtle biological sensory cues such as smell, will not reliably accept robots (embodied or virtual) as interactive partners.

The proposed closed-loop system could incorporate either real or virtual robots. However, given that the human participant must find the robotic responses believable, the external body configuration and motions may need to be optimized during the course of a research program. Because designing and building physical robots and their application programming interfaces (API) is itself a specialized and costly endeavor, we prefer a virtualized pseudo-3D representation of a robot. The physical attributes and API requirements of a virtually-embodied can easily be altered, and compiled for later physical implementations (of virtually-validated) behaving robotic systems.

CONCLUSION

In this paper, we propose a conceptual framework in which a virtual neuro-robotic simulation environment is used to accelerate the development of progressively more socially intelligent neuromorphic brain architectures. We completed a prototype of such a closed-loop system wherein a robotic dog controlled by a simple spiking brain model interacts emotionally with a human participant. Future plans include larger neuromorphic brains, with biological learning mechanisms, and humanoid virtual robot with linguistic capabilities.

We are presently developing a humanoid virtual social robot within the Webots/URBI environment, with functional attributes motivated by the MDS (mobile, dexterous, social) robot under development by the Per-

sonal Robotics Group of the MIT Media Lab (<http://robotic.media.mit.edu>). We plan to incorporate natural language understanding and production using corresponding neocortical models with reward based on praise and curiosity. A related model of childhood autism is also in progress.

CONFLICT OF INTEREST STATEMENT

The authors declare that this research was conducted without any commercial or financial relationships that could be construed as a potential conflict of interest.

ACKNOWLEDGMENTS

We thank Monica Nicolescu and Frederick C. Harris, Jr. for their guidance and recommendations, and Milind Zirpe for his assistance with the neocortical simulation. This work was supported by grants from the U.S. Office of Naval Research (grants N000140010420 and N000140510525).

REFERENCES

- Allen Institute. (2007). Allen Brain Atlas, URL: <http://www.brain-map.org>.
- Banquet, J. P., Gaussier, P., Quoy, M., Revel, A., and Burnod, Y. (2005). A hierarchy of associations in hippocampo-cortical systems: cognitive maps and navigation strategies. *Neural Comput.* 17, 1339–1384.
- Berger, T. W., and Glanzman, D. L. (Eds.). (2005). *Implantable biomimetic electronics as neural prostheses*. (Boston, MA, MIT Press).
- Bien, Z. Z., and Lee, H. E. (2007). Effective learning system techniques for human-robot interaction in service environment. *Knowl. Based Syst.* 20, 439–456.
- Boy, E. S., Burdet, E., Teo, C. L., and Colgate, J. E. (2007). Investigation of motion guidance with scooter robot and collaborative learning. *IEE Trans. Robot.* 23, 245–255.
- Braitenberg, V. (2001). Brain size and number of neurons: an exercise in synthetic neuroanatomy. *J. Comput. Neurosci.* 10, 71–77.
- Breazeal, C. (2003). Emotion and sociable humanoid robots. *Int. J. Hum. Comput. Stud.* 59, 119–155.
- Breazeal, C. (2004). Social interactions in HRI: the robot view. *IEEE Trans. Syst. Man Cybern. C Appl. Rev.* 34, 181–186.
- Breazeal, C., and Scassellati, B. (2000). Infant-like social interactions between a robot and a human caretaker. *Adapt. Behav.* 8, 49–73.
- Brette, R., Rudolph, M., Carnevale, T., and Hines, M., et al. Simulation of networks of spiking neurons: a review of tools and strategies. (2007). *J. Comput. Neurosci.* 23, 349–398.
- Brooks, R., Breazeal, C., Marjanović, M., Scassellati, B., and Williamson, M. (1998). The Cog project: building a humanoid robot, In *Computation for Metaphors, Analogy, and Agents*, C. Nehaniv (Ed.), (New York: Springer), pp. 52–87.
- Carlson, J., and Murphy, R. R. (2005). How UGVs physically fail in the field. *IEEE Trans. Robot.* 21, 423–437.
- Cuperlier, N., Quoy, M., Laroque, P., and Gaussier, P. (2005). Transition cells and neural fields for navigation and planning. *Comput. Sci.* 3561, 346–355.
- Dautenhahn, K. (2007). Socially intelligent robots: dimensions of human-robot interaction. *Phil. Trans. R. Soc. B.* 362, 679–704.
- Fellous, J., and Arbib, M. A. (2005). *Who needs emotions? The brain meets the robot*. (NY, Oxford University Press).
- Frijda, N. H. (2006). *The laws of emotions*. (London, Lawrence Erlbaum).
- Frye, J., Schurmann, F., King, J., Ranjan, R., and Markram, H. (2006). Neurodamus—a framework for large scale and detailed brain simulations. *FENS Abstr.* 3, A037.12.
- Iacoboni, M., and Dapretto, M. (2006). The mirror neuron system and the consequences of its dysfunction. *Nat. Rev. Neurosci.* 7, 942–951.
- Izhikevich, E. M., Gally, J. A., and Edelman, G. M. (2004). Spike-timing dynamics of neuronal groups. *Cereb. Cortex* 14, 933–944.
- Jensen, W., and Rousche, P. J. (2006). On variability and use of rat primary motor cortex responses in behavioral task discrimination. *J. Neural. Eng.* 3, L7–L13.
- Krichmar, J. L., and Edelman, G. M. (2005). Brain-based devices for the study of nervous systems and the development of intelligent machines. *Artif. Life* 11, 63–77.
- Krichmar, J. L., Seth, A. K., Nitz, D. A., Fleischer, J. G., and Edelman, G. M. (2005). Spatial navigation and causal analysis in a brain-based device modeling cortical-hippocampal interactions. *Neuroinformatics* 3, 197–221.
- Liu, J. N. K., Wang, M., and Feng, B. (2005). iBotGuard: an Internet-based intelligent robot security system using invariant face recognition against intruder. *IEEE Trans. Syst. Man Cybern.* 35, 97–105.
- Lungarella, M., and Berthouze, L. (2002). On the interplay between morphological, neural, and environmental dynamics: a robotic case study. *Adapt. Behav.* 10, 223–241.
- Macara, J. C., Goodman, P. H., Harris, F. C., Drewes R., and Maciokas, J. B. (2004). Remote-neocortex control of robotic search and threat identification. *Rob. Auton. Syst.* 46, 97–110.
- Maciokas, J. B., Goodman, P., Kenyon, J., Toledo-Rodriguez, M., and Markram, H. (2005). Accurate dynamical models of interneuronal GABAergic channel physiologies. *Neurocomputing* 65, 5–14.
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–225.



- Ogata, T., Sugano, S., and Tani, J. (2004). Innovations in applied artificial intelligence. *Lecture Notes in Computer Science*. 3029, 435–444.
- Olshausen, B. A., and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607–609.
- Peng, Q. Brainstem: a neocortical simulator interface for robotic studies. (2006). Master's thesis. (Reno, NV, University of Nevada). Available online at <http://brain.unr.edu>.
- Ripplinger, M. C., Wilson, C. J., King, J. G., Frye, J., Drewes, R., Harris, F. C., and Goodman, P. H. (2004). Computational models of interacting brain networks. *J. Investig. Med.* 52, 155.
- Samejima, K., and Doya, K. (2007). Multiple representations of belief states and action values in corticobasal ganglia loops. *Ann. N. Y. Acad. Sci.* 1104, 213–228.
- Schweighofer, N., Tanaka, S. C., and Doya, K. (2007). Serotonin and the evaluation of future rewards-theory, experiments, and possible neural mechanisms. *Ann. N. Y. Acad. Sci.* 1104, 289–300.
- Scheutz, M., Schermerhor, P., Kramer, J., and Anderson, D. (2007). First steps toward natural human-like HRI. *Auton. Rob.* 22, 411–423.
- Schindler, K. A., Goodman, P. H., Wieser, H. G., and Douglas, R. J. (2006). Fast oscillations trigger bursts of action potentials in neocortical neurons in vitro: a quasi-white-noise analysis study. *Brain Res.* 1110, 201–210.
- Schlesinger, M. (2003). A lesson from robotics: modeling infants as autonomous agents. *Adapt. Behav.* 11, 97–107.
- Seward, D., Pace, C., and Agate, R. (2007). Safe and effective navigation of autonomous robots in hazardous environments. *Auton. Rob.* 23, 223–242.
- Society for Neuroscience. (2007). Databases having experimental data. URL: http://ndg.sfn.org/eavObList.aspx?cl=81&at=278&vid=28872&menu_item=dblist1.
- Toledo-Rodriguez, M., Blumenfeld, B., Wu, C. Z., Luo, J. Y., Attali, B., Goodman, P., and Markram, H. (2004). Correlation maps allow neuronal electrical properties to be predicted from single-cell gene expression profiles in rat neocortex. *Cereb. Cortex* 14, 1310–1327.
- Toledo-Rodriguez, M., Goodman P., Illic, M., Wu, C. Z., and Markram, H. (2005). Neuropeptide and calcium-binding protein gene expression profiles predict neuronal anatomical type in the juvenile rat. *J. Physiol. (Lond.)* 567, 401–413.
- Triesch, J., Jasso, H., and Deák, G. O. (2007) Emergence of mirror neurons in a model of gaze following. *Adapt. Behav.* 15, 149–165.
- Visentin, G., and van Winnendael, M. (2006). *ACTA Astronaut.* 59, 750–756.
- Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J. Y., and Goldman-Rakic, P. S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nat. Neurosci.* 9, 534–542.
- Wiener, S. I., Arleo, A. (2003). Persistent activity in limbic system neurons: neurophysiological and modeling perspectives. *J. Physiol. (Paris)* 97, 547–555.