Naval Research Laboratory
Washington, DC 20375-5320

NRL/MR/8234--20-10,084

Final Report: Neuromorphic Architectures for Fast Low-Power Robot Perception, Work Unit IT015-09-41-1G25

Dr. Joseph T. Hays

Dynamics & Control Systems Branch
Spacecraft Engineering Division

Dr. Wallace E. Lawson

Dr. Keith M Sullivan

Navy Center for Applied Research in Artificial Intelligence Branch
Information Technology Division

June 18, 2020

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
18-06-2020

2. REPORT TYPE
NRL Memorandum Report

3. DATES COVERED *(From - To)*
Oct 1, 2016 – Sept. 30, 2019

4. TITLE AND SUBTITLE
Final Report: Neuromorphic Architectures for Fast Low-Power Robot Perception, Work Unit IT015-09-41-1G25

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
0601153N

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
1G25

6. AUTHOR(S)
Dr. Joseph T. Hays, Dr. Wallace E. Lawson, and Dr. Keith M Sullivan

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER
NRL/MR/8234--20-10,084

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

10. SPONSOR / MONITOR'S ACRONYM(S)
NRL

11. SPONSOR / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This memorandum summarizes the efforts of work unit 1G25 where new brain-inspired neuromorphic computing technology and deep/convolutional neural networks (CNN) where applied to develop real-time low SWaP scene understanding capabilities for mobile robotic systems. Specifically, we sought understanding of the relationships between the perception task, CNN-based algorithms, and the constraints of neuromorphic systems and to derive principles of CNN design for neuromorphic architectures.

15. SUBJECT TERMS

Radar	Multistatic
Inverse scattering	Sparsity

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
Unclassified Unlimited	Unclassified Unlimited	Unclassified Unlimited	Unclassified Unlimited	8	Dr. Joseph T. Hays
					19b. TELEPHONE NUMBER *(include area code)* (202) 404-4281

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

This page intentionally left blank.

INTRODUCTION

Advances in computer vision and machine learning have demonstrated techniques that can identify hazards, classify objects, detect features of interest, and perform many other autonomous tasks in cluttered, static scenes. Such techniques are suitable for stationary surveillance systems or for robotic vehicles that operate fairly slowly (such as tracked or wheeled vehicles, which typically drive slowly or stop in order to perceive their environment) or that operate at long distances from objects in their environment (such as unmanned aerial vehicles and quadrotors). These systems cannot perform well (or at all) in unknown and highly dynamic scenes.

Additionally, many mobile robot platforms are severely constrained in size, weight, and power (SWaP). Contemporary perception systems are based on very greedy algorithms that are computationally demanding. The real-time mobility requirements of these platforms demand that computation be performed locally, at the edge. As such, significant power is dedicated to the computation resources that exceed the low SWaP needs, therefore, new approaches must be investigated to achieve the needed low SWaP and real-time perception capabilities of future mobile robotic systems.

As the corporate research center for the Navy, the US Naval Research Laboratory (NRL) has a long heritage of investigating solutions to the Navy's need of advancing autonomous robotic systems on the battlespace. NRL's lessons learned from developing aerial, underwater, surface, ground and space robotic systems confirms the need of realizing low SWaP real-time perception capabilities on mobile robotic systems. This is consistent with the vision detailed in the 2018 National Defense Strategy which states,

> **"we cannot expect success fighting tomorrow's conflicts with yesterday's weapons or equipment…we must invest in modernization of key capabilities…**
>
> **The Department will invest broadly in military application of autonomy, artificial intelligence, and machine learning…to gain competitive military advantages."**

As we mature Edge Computing based real-time scene understanding, system autonomy will evolve to reduce human labor for system command and control. Thus humans shift from command and control to supervising the system.

OBJECTIVE

The objective of this work unit was to investigate the application of new brain-inspired neuromorphic computing technology and deep/convolutional neural networks (CNN) to develop real-time low SWaP scene understanding capabilities for mobile robotic systems. Specifically, we sought understanding of the relationships between the perception task, CNN-based algorithms, and the constraints of neuromorphic systems and to derive principles of CNN design for neuromorphic architectures.

Manscript approved May 19, 2020.

DEVELOPMENTS

We chose the use case of a small high-speed ground robot that was tasked with providing surveillance of an area where only coarse waypoints were provided. The environment was assumed to have significant environmental uncertainty between the waypoints. Standard RGB camera data was streamed from the ground robot and sent to the IBM TrueNorth neuromorphic processor to be analyzed.

We investigated a series of CNN-based perception algorithms that were developed for and deploy on to the TrueNorth neuromorphic processor. A short description of each algorithm follows. (Details related to these algorithms can be found in the papers listed in the *Publications Section*.)

ControlNet

Our initial IBM TrueNorth based CNN classifier enabled reactive behavioral control of a Pioneer robot navigating around obstacles as it moved through an area. This was tested in an indoor playpen area as shown in Figure 1. This neuromorphic reactive controller successfully identified obstacles and generated motor commands to avoid them.

Figure 1: Pioneer robot and playpen test area for the neuromorphic reactive controller, ControlNet.

OluNet/TerrainNet

Our second-generation perception algorithm classified outdoor image patches into one of three classes, *{concrete, grass, asphalt}*. Higher level logic then determined appropriate motor control behavior to keep the robot on concrete surfaces. The controller was again deployed to the TrueNorth to control a Pioneer robot. Figure 2 illustrates part of the path that the Pioneer followed around the NRL campus using OluNet. The controller failed at a few points illustrated in red (i.e. at concrete crossroads and when a concrete sidewall was present). However, the robot successfully navigated along the green regions of the path. Details of this effort are documents in [Sullivan, February 2019].

Figure 2: The path followed by the Pioneer robot using OluNet running on the IBM TrueNorth neuromorphic processor (left). An image showing the sub-patches passed to the TrueNorth for classification. The green patches are identified as grass, while the blue patch is classified as concrete.

Foveated Input to Deep Networks

Based on our experiences with OluNet/TerrainNet, a biologically inspired foveation approach was investigated to reduce the size of the input image to (1) reduce the number of neuromorphic cores required and (2) improve the amount of time needed to train. This was experimentally demonstrated on a deep convolutional network suitable to be run on the IBM TrueNorth. Figure 3 illustrates how a traditional images can be foviated such that more "attention" or priority is given to the central region of the viewing plan. This initial study, documented in [Lawson, 2018], found that the log-polar foviation strategy significantly increased the accuracy of the terrain being classified by OluNet.

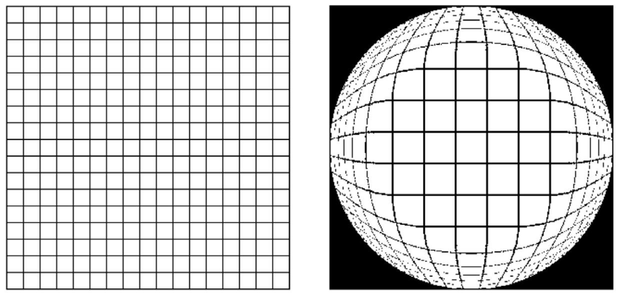

Figure 3: Log-polar foviation of an image.

AriNet

In an effort to improve upon the OluNet navigation controller, the AriNet controller was developed. This network took a different approach. Terrain was now classified as {passable, non-passable}. The neuromorphic terrain classifier constituted the "low level" controller and was integrated with a "high level" GPS waypoint following controller. Since the GPS waypoints were intentionally defined as very course, many obstacles and non-passable terrain was experienced by the low-level controller between the waypoints. At a waypoint, the higher level GPS controller would take over and orient the low level controller toward the next waypoint. A significant speed increase was observed when the robot was under neuromorphic navigation control. Figure 4 illustrates the Packbot robotic system (left) and the path around NRL campus (right). Analysis and results of this study are documented in [Sullivan, March 2019]

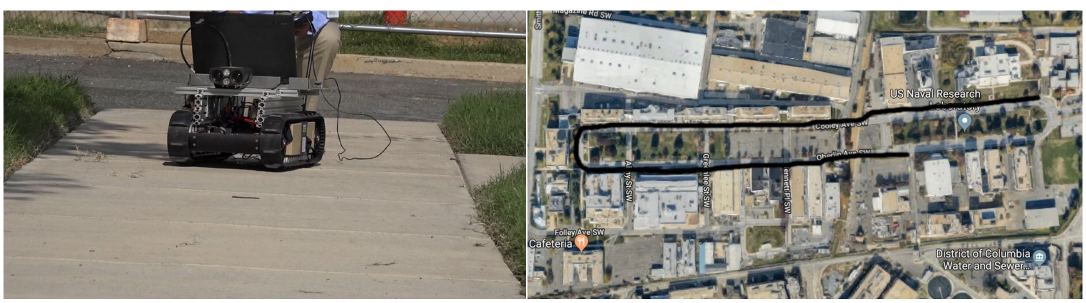

Figure 4: Packbot being controlled by AriNet (left), path to be traversed (right)

PUBLICATIONS

W. Lawson, K. Sullivan, C. Bradel, O. Roy, E. Bekele, "*Biologically Inspired Foveate Robot Vision*", Association for the Advancement of Artificial Intelligence (AAAI) Fall Symposia, 2018, Washington, DC.

K. Sullivan, W. Lawson, J. Hays, "*Neuromorphic Computing Used for Terrain Classification in Order to Guide an Autonomous Robot*", The Third Annual Workshop on Naval Applications of Machine Learning, 11–14 February 2019, San Diego, California.

K. Sullivan, W. Lawson, J. Hays, "*Neuromorphic Robot Perception for Autonomous Control*", Neuro Inspired Computational Elements Workshop, 26–29 March 2019, Albany NY.

TOWARD ONLINE LEARNING

The developments of ControlNet, OluNet and AriNet revealed the need to be able to continue learning online, or, to provide the robot with the ability to recognize when it is in a state that it does not confidently know what to do and then to ask for input. Once user input is received, the intent is to update the proper learned behaviors online and expand the systems navigation capabilities.

Much of this work unit was consumed in: (1) get access to both IBM TrueNorth neuromorphic hardware, (2) working with IBM and Intel to mature their software support to be able to solve problems of the size required by mobile robots. All published developments from this work unit are based on the IBM TrueNorth. This is an online inference-only device, meaning, all learning is carried out offline.

Unlike the IBM TrueNorth, the Inel Loihi neuromorphic processor has online learning cababilities. The team has successfully prototyped an initial Intel Loihi-based AriNet (initially developed for the IBM TrueNorth) but the current software support does not allow the network to span multiple chips. Therefore, the network was decimated down to an unusable size but proved that the Loihi was performing inference. Multichip spanning capabilities are yet to come online for NRL to be able to continue this invesigation.

SUMMARY OF KEY ISSUES/FINDINGS IN NEUROMORPHIC COMPUTING FROM AN APPLICATIONS PERSPECTIVE

The following includes a summary of key issues, or findings, learned during our developments in this work unit. This summary clearly illustrates that the Neuromorphic Computing community still has a number of challenges ahead before it will be more broadly adopted. These issues/findings include:

- Current neuromorphic hardware:
  - Does not allow for large/complex networks
  - Is very limited in types of network layers (doesn't support common layers such as LSTM, skip connections, etc.)
- General availability of neuromorphic hardware is still in question
  - Are any of the current research-grade chips going to be productized and generally available in a year? In three years? This is still a high-risk factor.
- Sensor-to-chip-to-actuator I/O performance improvements are needed
  - Current technology has unacceptably poor latency and bandwidth for many robotic applications
- Algorithmic needs include:
  - Advancements in unsupervised/meta online learning
  - Advancements that enable engineered online learning systems to be verifiably safe
- What is the temporal relationship of I/O data to encoding (e.g. is there a needed "presentation" time when simulating spiking neural networks (SNN))
- Low SWAP neuromorphic technology needs to be developed to expand out to the periphery
  - Too much power-hungry electronics are still required for a deployed neuromorphic system "in the wild".
  - Advancements in neuromorphic sensors (beyond DVS cameras) are needed, for example: strain-gauges, angle encoders, rate gyros, accelerometers, touch sensors, etc… Some work is being investigated for analog-to-digital and digital-to-analog converters to produce/consume spikes directly (similar in spirit to the DVS sensor). This would help produce lower SWAP periphery electronics.
- The neuromorphic community:
  - Requires specialized knowledge and is at risk to be able to transition to the larger Machine Learning and Robotics communities who do not have this specialize knowledge
  - Is struggling to define an undisputed advantage of applying neuromorphic computation over other low SWAP Von Neumann solutions.

ROBUST ARTIFICIAL INTELLIGENCE FOR NEUROROBOTICS (RAI-NR 2019) WORKSHOP

RAI-NR 2019 (https://blogs.ed.ac.uk/rai-nr/) was held on 26-28 August 2019 at the University of Edinburgh, Scotland. This event was a direct outgrowth of this Work Unit (IT015-09-41-1G25) and was seed funded by codes 8200 and 5500 from NRL through the Collaborative Science Program (CSP) of ONR-G (UK office).

RAI-NR 2019 explored issues of reliability, safety and resource efficiency in the context of how emerging neural network and other neuromorphic technologies could address these requirements in realistic applications.

The proceedings of the workshop will appear in a special issue of the Frontiers Journal of Neurorobotics (https://www.frontiersin.org/research-topics/11012/robust-artificial-intelligence-for-neurorobotics) of which the NRL research team associated with Work Unit (IT015-09-41-1G25) are guest content editors.

CONCLUSION

A number of key issues have been itemized in this report reflecting the current state-of-the-art in regards to that application of neuromorphic processing technology to mobile robotic perception tasks. This technology is still very immature but shows promise if the recorded limitations can be overcome. If realized, high performance low-SWaP processing can be realized. DoD mobile robotic systems are in great need of a high performance low-SWaP processing and therefore it is recommended that further research in this area continue.