WESTERN SYDNEY UNIVERSITY

The MARCS Institute for Brain, Behaviour and Development

MARCS Internal Conference Series

Internal ICNS Neuromorphic Processing Conference

15 NOV 2022 | WERRINGTON SOUTH

The International Centre for Neuromorphic Systems Western Sydney
Western Sydney University
University Locked Bag 1797. Penrith NSW 2751. Australia

PAPER TALKS

Neuromorphic Hardware 3
Dr Ying Xu

Neuromorphic Engineering: Bio-inspired Engineering Applied to Neurons 4
Dr Travis Monk

Experimental Theoretical Neuroscience 5
Dr Alexandre Marcireau

POSTERS

Neuromorphic event-based sensors for space situational awareness 6
Andrew Jolley

Real-time Event-based Unsupervised Feature Consolidation and Tracking for Space Situational Awareness 7
Nicholas Ralph

Learning Spatiotemporal Patterns with Low-precision Recurrent Spiking Neural Networks 8
Pablo Urbizagastegui

Hardware Architecture for Event-Driven Feature Extraction Algorithms 9
Philip C Jose

One-bit-per-weight: A single precision network for edge devices 10
Priyanka Bharadwaj

Event-based Atmospherical Mapping and Localisation from the ISS 11
Sami Arja

An FPGA-based Emulation of an Event-based Vision Sensor Using Commercially Available Camera: An Initial Study 12
Samalika Perera

Neuromorphic Hardware

Speaker **Dr Ying Xu**

Abstract

Neuromorphic electronic systems emulate the biological sensory and perceptual functions on hardware platforms, which have great potential to be used in practical applications to achieve real-time, low-power and robust performance. We present the Field-programmable gate array (FPGA) and Application-specific integrated circuit (ASIC) implementations of cochlea, retina, and brain-scale spiking neural network architecture here and demonstrate such neuromorphic electronic sensing systems.

Neuromorphic Engineering: Bio-inspired Engineering Applied to Neurons

Speaker **Dr Travis Monk**

Abstract

Moore's law is approaching saturation. Cramming more transistors on chips is becoming increasingly difficult. To continue improving computational abilities (without sacrificing the size, weight, and power consumption of our devices), we might need to explore new paradigms. Neuromorphic engineering proposes that we look to the animal kingdom for inspiration. Animals solve a different class of problems with fundamentally different hardware (i.e. nervous systems) than conventional computers. Their nervous systems are orders of magnitude lighter, smaller, and more power-efficient than our everyday devices. Neuromorphic engineering seeks to reverse-engineer nervous system design for applications. We will briefly review different ways that animals sense their environment and how they can influence our design of sensory devices.

Experimental Theoretical Neuroscience

Speaker **Dr Alexandre Marcireau**

Abstract

Theoretical - or computational - Neuroscience has produced many neuron models, microcircuits, and large-scale neural networks. However, assessing the behavioural response of such models in real-world scenarios remains challenging. Emulating densely connected and recurrent neural networks with conventional computers requires tremendous resources, while simulating realistic sensory inputs is still out of our reach in many cases.

The silicon sensors and computers developed by Neuromorphic Engineers are inspired by Neuroscience models. We use these devices to solve real-world problems, which lets us test whether they outperform conventional engineering approaches in terms of power consumption, latency, and robustness to noise.

This scientific endeavour aligns itself well with the search for applications of Neuromorphic devices, as the properties we aim to validate are critical to many applications in remote environments. This includes wildlife monitoring, space, and underwater sensing.

Neuromorphic event-based sensors for space situational awareness

HDR Candidate Andrew Jolley

Supervisors A/Prof Gregory Cohen
Prof André van Schaik
A/Prof Andrew Lambert

Abstract

Neuromorphic event-based optical sensors are attracting attention for their utility in detecting objects in Earth orbit, and for their advantages over traditional optical sensors for space situational awareness applications. With temporal resolution measured in microseconds, very high dynamic range, and data rates driven by changes in a scene rather than arbitrary frame readout rates, these relatively new sensors offer novel means for characterising satellites. Because the pixels in the sensor pixel array operate independently and asynchronously, when changes in a scene are detected, only data relating to those changes are output by the sensor. Rather than producing synchronous image frames, neuromorphic sensors output spikes, known as events, whenever a pixel detects an increase or decrease in brightness. By calculating the rate at which events associated with a satellite are recorded, event-rate curves can be produced. Event-rate curves are similar to lightcurves produced using traditional frame-based optical sensors, however event-rate curves do not specify the brightness of a satellite. Instead, event-rate curves provide information relating to changes in a satellite's brightness, which can be used to derive satellite rotation rates, axes of rotation, and reflecting surface characteristics. Using two event-based cameras with different broadband optical filters fitted, multicolour event-rate data were obtained simultaneously, for a variety of satellites. Using these data, colour event-rate ratio curves were produced, which provide high-temporal-resolution information on colour changes in the light reflected from satellites. Such colour changes can be used to infer the type of material responsible for specular reflections, as well as satellite orientation, configuration, and operational status

Real-time Event-based Unsupervised Feature Consolidation and Tracking for Space Situational Awareness

HDR Candidate Nicholas Ralph

Supervisors Dr Saeed Afshar
A/Prof Gregory Cohen
Dr Damien Joubert
Andrew Jolley
Dr Nicholas Tothill
Prof André van Schaik

Abstract

Earth orbit is a limited natural resource that hosts a vast range of vital space-based systems that support the international community's national, commercial and defence interests. This resource is rapidly becoming depleted with over-crowding in high demand orbital slots and a growing presence of space debris. We propose the Fast Iterative Extraction of Salient targets for Tracking Asynchronously (FIESTA) algorithm as a robust, real-time and reactive approach to optical Space Situational Awareness (SSA) using Event-based Cameras (EBC) to detect, localise and track Resident Space Objects (RSO) accurately and timely. We address the challenges of the asynchronous nature and high temporal resolution output of the EBC accurately, unsupervised and with few tune-able parameters using concepts established in the neuromorphic and conventional tracking literature. We show this algorithm is capable of highly accurate in-frame RSO velocity estimation and average sub-pixel localisation in a simulated test environment to distinguish the capabilities of the EBC and optical setup from the proposed tracking system. This work is a fundamental step toward accurate end-to-end real-time optical event-based SSA and developing the foundation for robust closed-form tracking evaluated using standardised tracking metrics.

Learning Spatiotemporal Patterns with Low-precision Recurrent Spiking Neural Networks

HDR Candidate Pablo Urbizagastegui

Supervisors Dr. Mark Wang
Prof. Andre van Schaik

Abstract

Emulating spiking neuron networks in neuromorphic hardware has highlighted promising prospects for solving real-world problems. These types of devices are particularly appealing to applications that require low power consumption, asynchronous operations, and short latency. To make full use of its potential, however, we need to understand how spiking neurons cooperate to compute, and what are the components needed to succeed in a given task. Researchers on neuromorphic systems often approach these questions, but there is always a compromise: On the one hand, we need to reduce the complexity and memory footprint of neural systems. On the other hand, multiple mechanisms may be necessary to introduce stable, efficient learning in silico. In this work, we show how recurrent spiking neurons with low bit resolution can learn a simple spatiotemporal sequence with different mechanisms inspired by neurobiology. We demonstrate that these neurons respond selectively to features conveyed by the input and encode temporal information on their recurrent weights. Additionally, they exhibit brief direction selectivity and extrapolate the next possible elements of the sequence. These results are a necessary step before scaling up to very-large cortical models and answering questions that traditional artificial neural network research has failed to answer.

Hardware Architecture for Event-Driven Feature Extraction Algorithms

HDR Candidate Philip C Jose

Supervisors Dr Runchun Mark Wang
Dr Ying Xu
Prof André van Schaik

Abstract

Dynamic Vision Sensors (DVS) are developed by taking inspiration from mammalian retinas that can detect changes in luminosity to produce a stream of events. This continuous stream of events from the sensor is used to extract the complex spatio-temporal patterns of a scene with the aid of event-based (EB) feature extraction algorithms. However, vast variants of EB feature extraction algorithms reported in the literature use time surface computation to create a spatio- temporal event context with a predefined radius around an incoming event from the sensor at a specific time history. Feature Extraction using Adaptive Selection Threshold (FEAST) is a novel EB feature extraction algorithm that uses time surface-based computation to extract the most commonly occurring patterns or features in an event stream. This work introduces a novel hardware architecture for accelerating the FEAST network in real-time. More importantly, the EB computations within the architecture are based on integers where the time surface computations are approximated with 8-bit integers, and neuron weights are normalised to a 9- bit integer within a range of [-1:1]. A software model of the architecture was tested with the Poker DVS dataset, obtaining an accuracy of 95% for a 9-bit precision of neuron weights with 0% accuracy loss in comparison to the full-precision FEAST network.

One-bit-per-weight: A single precision network for edge devices

HDR Candidate Priyanka Bharadwaj

Supervisors Dr Runchun Mark Wang

Dr Damien Joubert

Prof André van Schaik

Abstract

Deep neural networks (DNNs) trained with single precision are more energy efficient than their full-precision counterparts on resource-constrained embedded devices. Recent works placed a strong emphasis on these architectures, which are now capable of achieving human like performances on CIFAR-10 and CIFAR-100 datasets. More particularly, Hardware friendly models such as 1-bit-per-weight Binary Neural Network (BNN) exhibit optimal performance when Batch Normalization (BN) is replaced with scaling and shifted Rectified Linear Unit (sReLU) layers. However, complex classification datasets such as ImageNet have not been investigated to their full potential using single bit precision DNN. This work presents the results of classification models on the ImageNet dataset (ILSVRC2012 classification ImageNet challenge). CIFAR-10, CIFAR-100 and ImageNet have achieved 91%, 71% and 79% (Top-5 accuracy) accuracies respectively in case of 1-bit-per-weight Wide residual networks which is close to that of a full precision network of the respective datasets. These findings may provide significant optimization opportunities in the arithmetic core by removing the need for computationally expensive multiplication units. Large silicon area required for implementing BN layer can be managed by using smaller silicon area to implement sReLU and thereby reduce the power dissipation associated with the hardware acceleration units.

Event-based Atmospherical Mapping and Localisation from the ISS

HDR Candidate Sami Arja

Supervisors Dr Alexandre Marcireau
Dr Saeed Afshar
A/Prof Gregory Cohen

Abstract

Earth's upper atmosphere is home to several transient luminous events (TLE) such as lightning sprites, or red sprites. Many TLE observations have been made since 1989, however, they are still poorly understood because they are extremely fast and rarely occur. TLEs often occur in tandem with lightning and are short-lived electrical discharges in the order of hundreds of microseconds. TLE observations above thunderstorms are difficult because the light emissions of sprites are superimposed on the intense light emissions of the lightning diffused by clouds. Event cameras make a good candidate for detecting high-speed transient luminous due to their fast temporal response and pixel bandwidth. To predict the sensor's ability to map and localise sprites and lightning effectively, we build a three steps network that maps small cities and vessels at night and reports their precise geographical location from the event stream. The network features precise ISS ephemeris, event-based motion compensation and homography estimation. Each part of the network interacts to ensure a final and accurate localisation value is obtained. Our results indicate that with appropriate camera settings and suitable mapping parameters, the network can capture these transient luminous events when they occur. Finally, we include initial mapping and localisation results of the network from several recordings from the ISS NADIR camera.

An FPGA-based Emulation of an Event-based Vision Sensor Using Commercially Available Camera: An Initial Study

HDR Candidate Samalika Perera

Supervisors Dr Mark Wang

Dr Ying Xu

Prof André van Schaik

Abstract

Event-based vision sensors are a promising approach for bio-inspired vision sensors to mimic the human eye in receiving visual information. It provides a number of advantages over conventional sensors, such as low latency and high temporal resolution. Existing event-based vision sensors use dedicated analogue circuits to generate events. In this paper, we present an FPGA-based framework for an event-based sensor that uses a commercially available camera to generate events digitally. To emulate the behavior of an event-based vision sensor, it is essential to detect the changes in brightness of each pixel with time. The camera module provides a digital value for the brightness of each pixel. These digital values are stored on the on-chip Static Random Access Memory (SRAM) of the FPGA, and the stored values are compared to generate events according to the brightness changes as the event-based vision sensor's behavior. We have illustrated a stable and cost- effective event generation system implemented on a Cyclone V FPGA board with an inexpensive 8- megapixel camera module running up to 60 frames/second.