BrainChip – Internship Program 2024

Essential AI. Close to the sensor

Tony Lewis CTO

Todd Vierra, VP of Customer Engagements

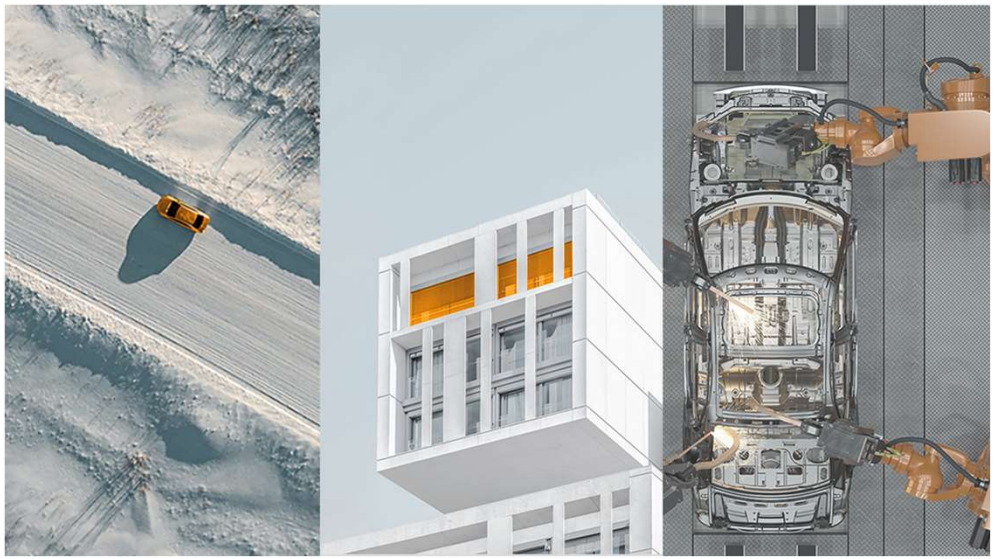

brainchip
Essential AI

April 2024

A little about BrainChip

Our technology brings commonsense to the processing of sensor data, allowing efficient use for AI inferencing enabling you to do more with less. Accurately. Elegantly. Meaningfully. We call this Essential AI.

Essential is optimizing compute. Maximizing performance. Minimizing power. In the real world. And in real time.

- 15 years of AI architecture research
- Worldwide leader in edge AI on-chip processing and learning.
- First to commercialize neuromorphic IP platform and reference chip.
- Centers of engineering excellence in US, Australia, France and India.

Trusted By:

MegaChips RENESAS Valeo NASA Mercedes-Benz

Profiled In:

VentureBeat EE Times Design & Reuse

Traded On:

ASX AUSTRALIAN SECURITIES EXCHANGE OTCQX

Internship Program – 7 projects we are looking for Interns

- Start and End appx 3 months
- Flexible start depends on University's summer schedule (example May – June 2024 start)
- End Aug/Sept
- $28 per hour
- 25 hrs per week
- No Benefits or PTO
- Interview & Background Check Necessary
- Fast working environment
- Remote or Local
- Skillset and applicants will determine projects

Advantages – Work with State of Art HW and AI Team

May lead to full time employment

Possible AI Projects (Slides 4-13)

Simulation Model

Eye tracking demonstration

Demo Modulization

Demonstration Catalog Development

Custom Yocto Image Compatible for Akida Demos

Convert Akida demo to C++ and optimize pipelines

Advancing SoTA in State Space Models with TENNS

Sparsity in Event Based System

LLM agents at the edge

Explore Probabilistic Computing Paradigms

Project One Eye tracking demonstration -Nolan

Background:

BC has recently published a paper for a lightweight TENN network that does gaze prediction with data from an event camera ([2404.08858] A Lightweight Spatiotemporal Network for Online Eye Tracking with Event Camera (arxiv.org))

This network placed 3rd (very close to 1st) in a recent Kaggle competition

Project Idea:

Build out a light vertical around eye tracking for AR/VR, while demonstrating how Akida IP can perform very fast and accurate inference with high sparsity.

A successful demo will demonstrate the value of our IP to large players in the AR/VR space like Meta and Microsoft, as well as prophesee who is a close collaborator with Qualcomm. Right now, the problem of eye tracking is solved using multiple >200fps IR cameras which is a very suboptimal solution.

Deliverables:

1. An efficient application which takes in live events, inferences BC's finished model, then displays the output in real-time

2. A VR-like headset with two prophesee genx320's with which to perform gaze detection (Event-Based Metavision® Sensor GENX320 | PROPHESEE)

brainchip
Essential AI

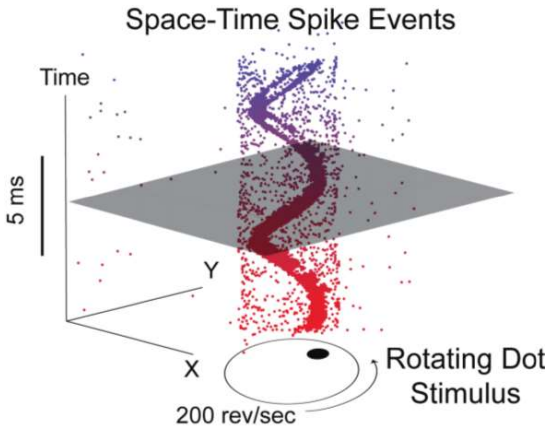

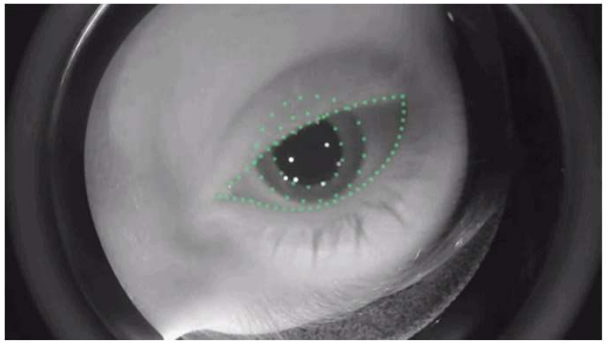

Project Two Cont.. Richard
Akida Demo Modularization and Standardization

Objectives:

1. Evaluate multitude of Python/C/C++ environments upon which Brainchip's demos are based and gather feedback from demo stakeholders on environment requirements.
2. Select a "standard" environment for modular demo development
    1. Adapt 1-2 existing demos to new modular "standard" while noting usage tips
    2. Produce a demo developers guide describing how to adapt/create modular demos in "standard" environment

Description:

Brainchip has demos based on multiple Python environments including, but not limited to, Flask, Kivy, TkInter, and PyQt. Each environment utilizes a unique set of abstractions for the elements of a ML demo pipeline. This situation makes it difficult for a given demo creator to adapt/reuse demo code. Standardizing on an environment with lightweight abstractions and formulating a modularized set of building blocks would speed demo development.

Team size: 1-2

Skills Needed: Python/AI Frameworks, Python GUI Frameworks, C/C++

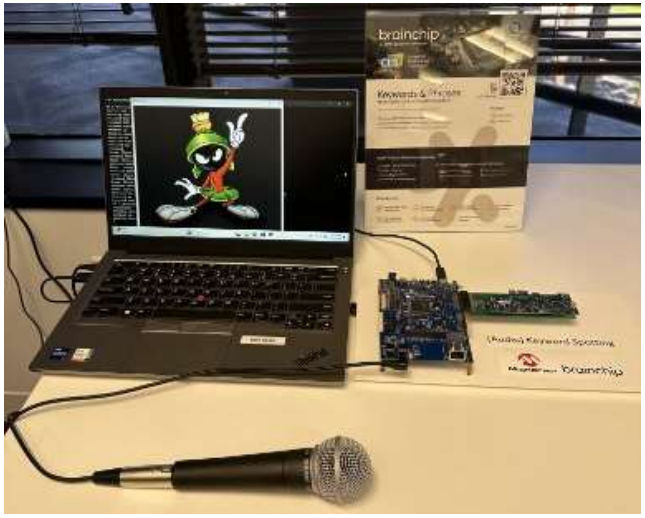

Project Three – Technical Demonstration Catalog Development Internship

Objective:

- Create a comprehensive catalog of our technical demonstrations, including source code, executables, hardware requirements, setup instructions, screenshots, videos, and relevant product feature highlights, to streamline the demonstration process for our solutions architects and sales teams.

Description:

- Inventory and document our existing technical demonstrations, including source code, executables, and hardware requirements.
- Develop a user-friendly catalog interface that allows for easy navigation and access to demonstration details.
- Capture screenshots and record short videos showcasing each demonstration's functionality and user interface.
- Create comprehensive setup and configuration guides for each demonstration, ensuring seamless execution on specific hardware configurations.
- Collaborate with product teams to identify and highlight relevant product features demonstrated by each technical example.
- Gather and present performance metrics and other critical information commonly requested by customers.

Team size: 1-2

Skills Needed :

- Strong organizational and documentation skills.
- Familiarity with various programming languages (e.g., Python, C++, Java) and development environments.
- Experience with technical writing and creating user guides or manuals.
- Proficiency in video recording, editing, and image capture/processing tools.
- Understanding of software development processes and version control (Git and GitHub).
- Knowledge of hardware setup, configuration, and troubleshooting.

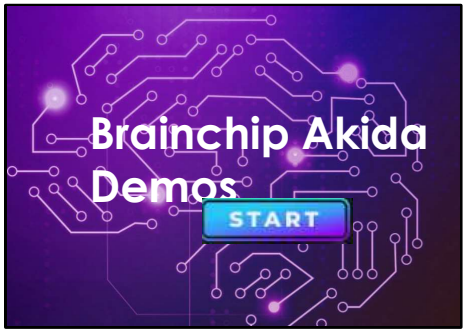

Project four Custom Yocto Image Compatible for Akida Demos

Objective: The primary goal is to create a streamlined process for building Yocto Linux images that incorporate our existing Python and C++ applications, enabling seamless deployment on our embedded devices (probably NXP i.MX8)

Description:

Create an Akida-compatible packaging framework over Yocto so that generating custom Linux image with the required libraries compatible with Akida drivers for becomes easy. This will allow our existing and future demos to be run on embedded devices with yocto linux

Demos:

- Visual Wake Word (May be other use case)
- Other Demos as time permits…

Team size: 1

Skills Needed :

- Experience with Yocto and/or Buildroot for building embedded Linux distributions (preferred).
- Familiarity with build systems such as CMake, Make, or BitBake.
- Understanding of embedded system architectures and their specific requirements.
- Proficiency in Python application development and C/C++ compilation.
- Knowledge of the software development lifecycle, including version control (Git and GitHub), code reviews, and testing practices.

Person

No Person

Example Images from the COCO training set*

Project Five– 2024 Akida Engine Library's C++ API and create C++ Machine Learning application demos already supported by Akida.

Objective: Develop an in-depth understanding and documentation of the Akida Engine Library's C++ API and create C++ Machine Learning application demos already supported by Akida.

Description:

- Explore and document the capabilities of the Akida Engine Library's C++ API.
- Create machine learning demos in C++ to demonstrate the library's practical applications.
- Evaluate the performance and efficiency of the Akida Engine in various computational scenarios. (Benchmarking in C++)

Team size: 1

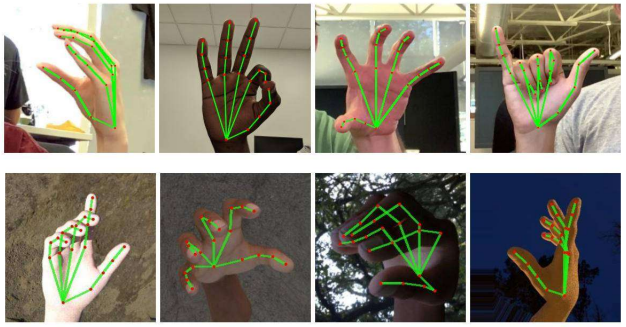

Skills Needed: C++, Machine Learning, Technical Writing

Project Six:
Transaction-Level Akida Simulation Model

Objectives:

1. Develop a SystemC Transaction-Level-Modeling (TLM) simulation of Akida 2.0 IP for functional and performance simulation.

Description:

Transaction-Level-Modeling (TLM) simulation provides high simulation speed without sacrificing functional accuracy. System performance can also be modeled with high simulation performance. SystemC is an open source C++ based framework for TLM simulation. Use SystemC to create a SystemC model of Akida NPs, Nodes, and HW Mesh along with a simulation testbench.

Team size: 1

Skills Needed: Computer Architecture/Engineering, RTL HW modeling, C/C++, SystemC experience a plus

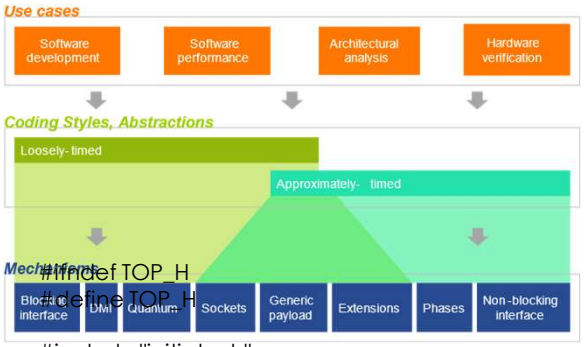


```
#ifndef TOP_H
#define TOP_H
#include "initiator.h"
#include "target.h"

SC_MODULE(Top)
{
  Initiator *initiator;
  Memory    *memory;

  SC_CTOR(Top)
  {
    // Instantiate components
    initiator = new Initiator("initiator");
    memory    = new Memory   ("memory");

    // One initiator is bound directly to one target
with no intervening bus

    // Bind initiator socket to target socket
    initiator->socket.bind( memory->socket );
  }
};

#endif
```

Project Seven TENNS 2.0

Advancing SoTA in State Space Models

Objective: State space models like TENNs are multipurpose algorithms with uses in denoising, object recognition, LLMs, and compression. The candidate will improve algorithms and demonstrate applicability in new domains.

Description: State space models are one of the hottest research topics in machine learning. SSM represent an alternative to LSTM, GRU etc. in being trainable as convolutional networks but run as compact RNNs.

Team size: 1-2

Skills Needed:

Advanced knowledge of topics such as Signal Processing, Theoretical Physics and Neural Networks.

Working knowledge of at least one of : PyTorch, Numpy

Highly fluent in at least one programming language. Python preferred.

Academic requirements:

3.5 GPA in a stem field (e.g. Math, Engineering, CS, Physics) or other evidence of ability.

Other: Ability to work independently and in teams. Excellent verbal, written and presentations skills.

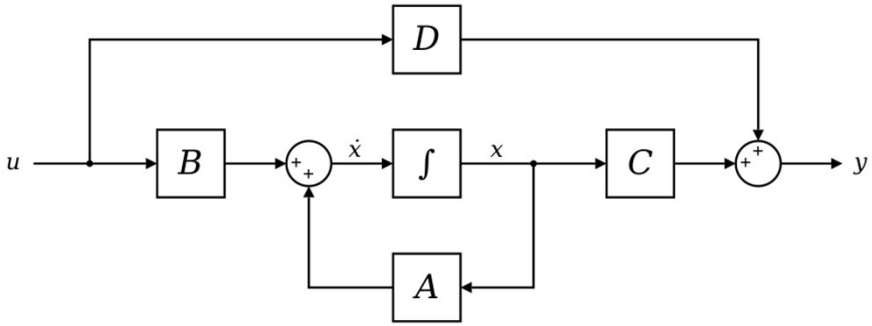

Project Eight Sparsity in Event Based System

Think more, work less

Objective: Power in AKIDA hardware is a function of activation sparsity (as opposed to weight sparsity). The selected candidate will explore new ways of increasing sparsity at all levels of a network.

Description: Description: Event-Based processing is gaining mainstream adoption. We have developed unique spatiotemporal processing technology which can be applied to a host of applications including automotive, biomedical & biometrics signals.

Team size: 1-2

Skills Needed:

Working knowledge of at least one of : PyTorch, Numpy, or other deep learning framework.

Highly skilled in Python.

Academic requirements:

Advanced coursework in applied mathematics

3.5 GPA in a stem field (e.g. Math, Engineering, CS) or other evidence of ability.

Other: Ability to work independently and in teams. Excellent verbal, written and presentations skills.

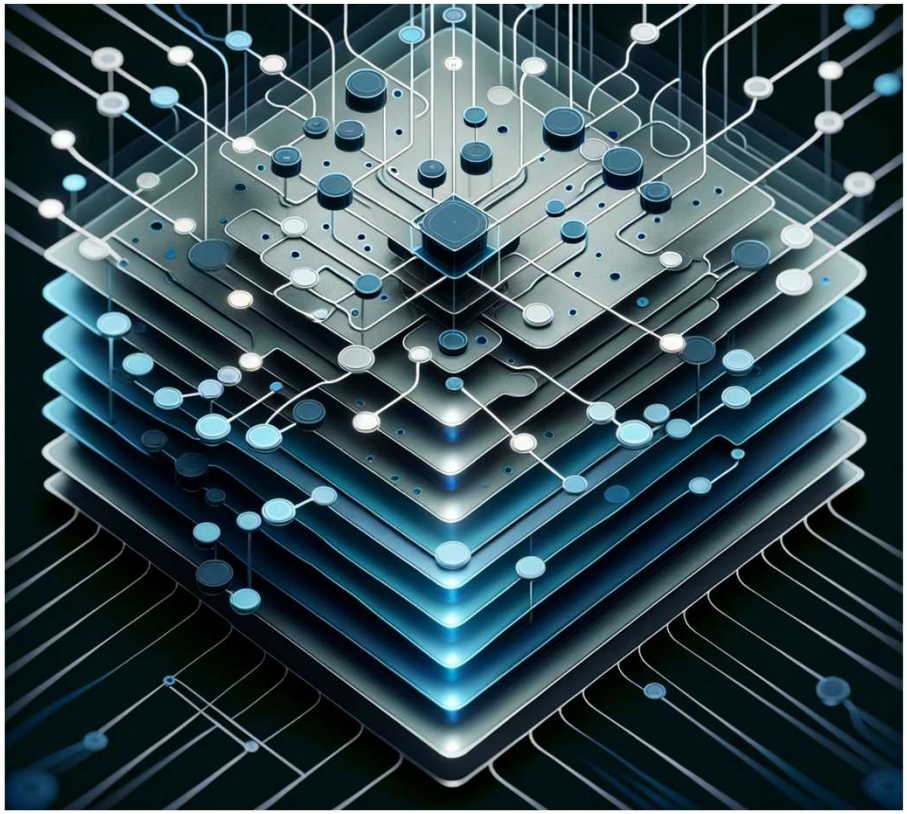

brainchip
Essential AI

Project Nine LLM agents at the edge

Retrieval-Augmented Generation (RAG) to improve LLM performance

Objective: Using off-the shelf tools and specialized prototype LLM modules develop at BrainChip, develop and demonstrate a RAG based agent for consumer use at the edge (i.e. not in the cloud)

Description: Edge agents will preserve privacy and security and at the same time are energy saving and green. In this project you will explore how to integrate BrainChip's small language models (< 1 billion parameters) in an on-edge device while augmenting the LLM with factual information to reduce hallucinations.

Team size: 1-2

Skills Needed : Python. A deep learning framework such as PyTorch. Familiarity with RAG, Vector Databases and principles of NLP a plus.

Academic requirements:

Strong mathematical background

3.5 GPA in a stem field (e.g. Math, Engineering, CS, Physics) or other evidence of ability.

Other: Ability to work independently and in teams. Excellent verbal, written and presentations skills.

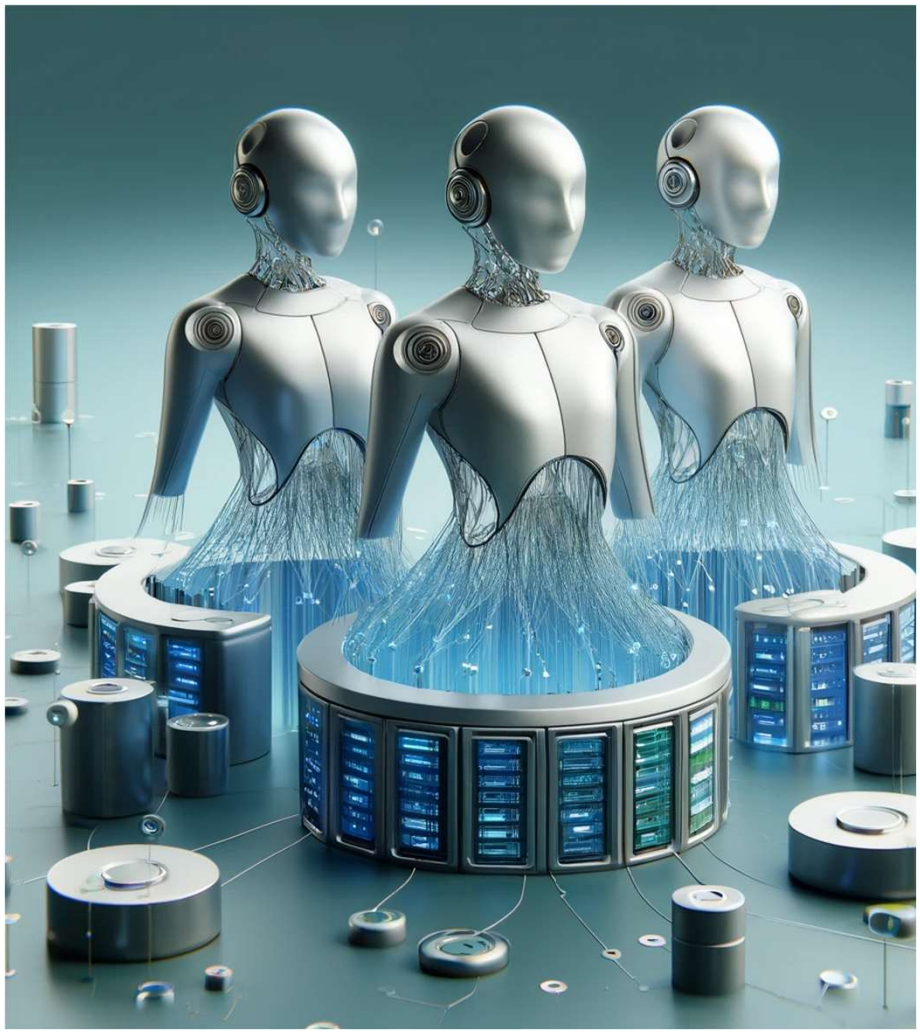

brainchip
Essential AI

Project Ten Explore Probabilistic Computing Paradigms

New paradigms in Neural Network Computation

Objectives:

Lower power usages on Neuromorphic hardware leveraging concepts in probabilistic computing, i.e. probabilistic algorithms for matrix multiplication etc.

Description: Probabilistic operation hold great promise for balancing power usage and performance in neuromorphic systems. In this project you will work to develop algorithms which can be run on our Generation 2 hardware. The objective is to demonstrate a system where energy and algorithmic performance can be dynamically altered.

Skills Needed:

Working knowledge of at least one of : PyTorch, Numpy, Keras or other deep learning framework.

Highly skilled in Python.

Academic requirements:

Advanced coursework in Linear Algebra and Probability

3.5 GPA in a stem field (e.g. Math, Engineering, CS) or other evidence of ability.

Other: Ability to work independently and in teams. Excellent verbal, written and presentations skills.

If you are interested or any questions email.

Contacts:

Tony Lewis

Todd Vierra

interns2024@brainchip.com

Please send your resume and cover letter with the specific project you are interested in and with your relevant experience related to the project.

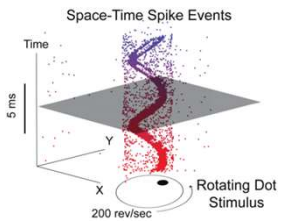

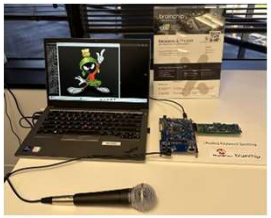

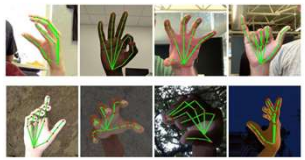

brainchip
Essential AI

Backup