Plate Reader

xcore.ai License Plate Camera and Reader

ELECOMP Capstone Design Project 2020-2021

Sponsoring Company:

XMOS Ltd.
Queens Quay
33 – 35 Queen Square
Bristol, BS1 4LU
United Kingdom
http://www.xmos.com

Company Overview:

Our relationship and reliance on technology is changing rapidly. In this increasingly complex, connected world, we want a more seamless connection with the technology around us. We want technology that can talk, empathize, relate, and react – an interface between human and machine that feels … human. One that comes with no instructions needed, delivered through artificial intelligence and connectivity.

XMOS is at the forefront of the far-field voice interface market, with the most qualified, most comprehensive range. Adopted by leading brands, our high-performance, easily integrated solutions are transforming consumers' interactions with everyday devices.

With more than 78 hardware and algorithm patents, XMOS silicon and software solutions deliver best in class voice capture for the smart environment.

We are headquartered in the UK, with offices across Asia and the United States.

Technical Director:

Dr. Andrew Cavanaugh
Algorithm Engineer
andrewc@xmos.com
https://www.linkedin.com/in/andrew-c-7805ba8/

Project Motivation:

As consumer and industrial electronics become more 'intelligent' there is a growing demand to move that intelligence from datacenters in the cloud to devices themselves. This demand solves issues of privacy, limited connectivity, and any case where the cost (in Watts or dollars) of transmitting sensor data is higher than the cost of processing it locally.

XMOS' newest chip, xcore.ai, is perfectly suited for sensing, DSP, neural network inference, and wireless connectivity. Customers buy products to solve problems, and a product with a track record of solving problems is easier to sell. Since this project uses a camera, DSP, and machine learning (with optional radar sensor and WiFi connectivity) the result should be a prototype that solves several complex problems with a single chip.

The market for Artificial Intelligence Internet of Things (AIoT) devices is extremely large and growing extremely fast. By providing a technology that enables these devices to operate with low energy consumption, and regard for end-user privacy the xcore.ai platform will be well positioned to be the first choice for AIoT manufacturers in the coming years.

Anticipated Best Outcome:

The Anticipated Best Outcome (ABO) of this project is a fully functional prototype that can accurately report license plate numbers in a variety of lighting and motion conditions. An enclosure with an Explorer board, any supporting boards, camera, and power supply that can survive damp conditions would be ideal. It is important to protect the sensors and board to the maximum extent possible without significantly obscuring sensors such as a camera or radar. The images captured by the camera need to be as clear or clearer than those used in ML training data sets.

Firmware for the xcore.ai chip, which should be able to start in parallel with the electrical design, should include basic data acquisition, image processing, and a neural network for interpreting the images. The results of the inference should be made available via some backhaul method (I2C, USB, WiFi, etc.)

Additional features that go beyond the ABO would be:

- Using radar sensor (will be provided) to improve accuracy and/or wake device
- Reporting wirelessly via WiFi (included on explorer board)
- Low-light performance

Project Details:

While designing and training a neural network is a necessary part of this project and underpins everything else, it will be the implementation on xcore.ai and integration with sensors that should constitute the bulk of the work. Leverage existing networks that do the same or similar things, as well as open source training data whenever possible. A few weeks of research on what is currently available for free may save a few months of scraping and labeling data.

A camera with basic driver will be provided, but the team is free to change the camera / modify the driver based on the requirements of the final product.

A trained neural network should be frozen early in the design process. Ideally this network would be validated by images captured on the camera selected for the final product. This network will then be quantized and ported (with the help of provided tools) to xcore.ai. A good guideline would be to have a trained model with >85% accuracy and quantized model with > 80% accuracy. Higher accuracy is better, but this is not the overall focus of the project. You may also be able to filter the outputs from multiple frames to get a better final estimate.

The level of integration we are looking for is a project box that holds all components and can tolerate some external moisture. We do not think custom PCBs will be required, but that decision is left to the team.

The Product:
If the ABO is achieved your team will have created a demonstrator box that can:

- Sit on a sidewalk for an hour
- Report unique license plate #s that pass the camera
- With all image processing and inferencing done on xcore.ai
- Report accuracy of trained neural network, and accuracy loss due to quantization
- Energy efficiency estimates (i.e. the average cost in Joules of reporting a plate#)

Block Diagram:

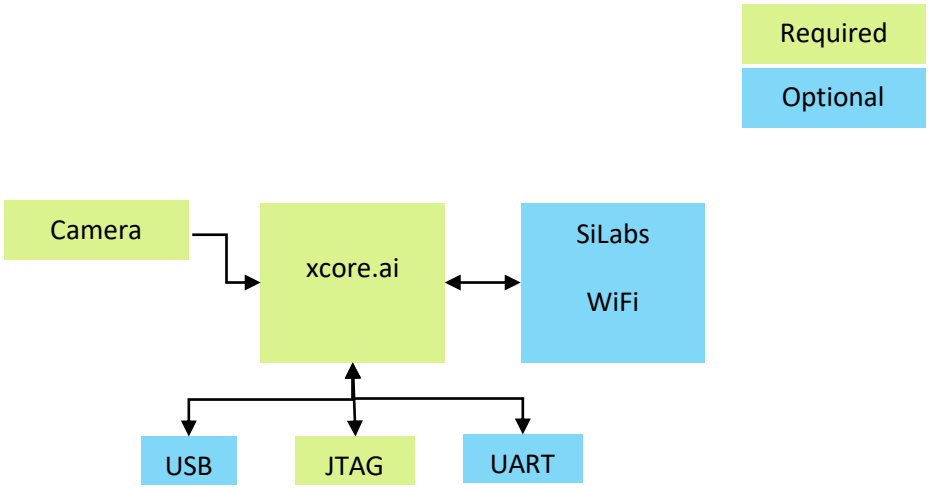

Hardware/Electrical Tasks:

- Responsible for selecting what processing should be done on raw image data before feeding into neural network (can be prototyped in a non-embedded implementation, but keep memory and processing limitations in mind)
- Determine layout for boards, sensors, etc.
- Design protective enclosure that does not interfere with camera or other sensors / WiFi
- Provide method for measuring overall system power while deployed

Firmware/Software/Computer Tasks:

- Firmware will be written in C or C++ on either bare metal or FreeRTOS
    - Modular code will be essential
    - Leverage existing libraries and tools as much as possible
    - Learn tradeoffs between fixed point and floating-point arithmetic
- You will be using development tools that might not be released publicly
    - Report any suspected bugs to technical director
    - Give feedback on what you like and what could be improved
    - While most coding will be in C/C++ there is additional “linker script” to determine how different software and hardware components should be “glued” together

- Host interface
    - Create a basic tool for logging and displaying data from the board
    - Ultimately this will just need basic data like plate#, but capturing raw images may be useful for debugging the end-to-end system
    - Use whatever languages/libraries you feel comfortable with
    - Implement some sort of versioning to ensure that host and device apps are either compatible or they will complain loudly

Joint Tasks:

- The neural network design and training will require a great deal of effort up front. Both team members should familiarize themselves with:
    - Neural network basics (simple math wrapped in jargon)
    - Convolutional neural networks
    - NN framework (Tensorflow, PyTorch, etc.)
    - Post-training quantization
- Select a network architecture
    - Total # weights should be under 300k
    - Maximize use of operators that have optimized implementations
    - Remember to look for existing networks that solve this or similar problems!
- Training
    - Find a trained network
    - Or determine how to best get training data

Composition of Team:

1 Electrical Engineer & 1 Computer Engineer

Skills Required:

Electrical Engineering Skills Required:

- Strong signal processing background
- Linear Algebra a plus
- Familiarity with embedded programming

Computer Engineering Skills Required:

- Familiarity with concurrent / real time computing
- Embedded C/C++
- Basic Signal Processing knowledge
- Experience moving data off chip: UART/SPI/I2C/USB/etc.

Anticipated Best Outcome's Impact on Company's Business, and Economic Impact

This will be one of the first projects even implemented on this new platform. As such it will set expectations both internally and with prospective customers about what xcore.ai is best suited for. We have specified the problem to be solved but have not specified how to solve it in too much detail. This allows the team to have maximum flexibility and creativity which will allow them to teach us something new about our own platform.

Integrating sensors and 'intelligence' on embedded devices without sending high bandwidth data to the cloud is a widely shared business need. Achieving the ABO will show the market that xcore.ai is well suited to filling this need.

Broader Implications of the Best Outcome on the Company's Industry:

In just the last decade, machine learning has gone from a niche research toping to a commodity product deployed in everybody's pocket supercomputer. Commercialization has been largely met with excitement on the part of both companies and end users – early adopters tend to, by definition, meet new products with excitement. Because the most computationally intense part of neural networks is training, much of the hardware development has been in making training faster or more energy efficient, doing on-line inference is more of an afterthought (luckily, since training requires running millions of inferences, this hardware was also well suited to it). Unfortunately, having all the relevant computational power in the cloud has resulted in applications where most of the time and energy of an inference is in transmitting data and waiting for an answer.

The move to doing inference on deeply embedded (much lighter weight than a cell phone processor) hardware has just started, and a "killer app" has yet to emerge. We need to explore the limits of what can be done on these types of processors and we need to ask ourselves what benefits they are providing over a traditional cloud-based inference in terms of cost, power, privacy, etc.