Master Thesis Description

Topic: Brain-Inspired Visual Attention with Event Cameras and Spiking Neural Networks

Description: The aim of this work is to develop a brain-inspired visual attention model that does not process all the information arriving from every pixel. Instead, it selects the most salient location at each moment and analyses only that location further. You will begin by building an event-based vision pipeline in which large receptive fields pool the visual input into a coarse representation. From this representation, you will extract different features, such as centre-surround and orientation responses, and combine them into a saliency map. A winner-take-all stage then selects the most important spot in the visual field. A key part of the thesis is the top-down attention signal, which tunes the saliency map according to the current goal. For example, when the system is looking for edges, it can place more emphasis on the edge feature layer. Once a location is selected, the receptive fields in and around it are analysed at a finer resolution to extract the information of interest. You will implement this coarse-to-fine model using spiking neural networks and evaluate it against processing the full visual field at full resolution, comparing computational cost, latency and the quality of the selected information. As a possible application, the attention model can guide where a humanoid robot such as iCub looks, touches and moves, for example while it learns about the space around its body.

Motivation: Processing every incoming pixel at full detail is expensive and largely unnecessary, because only a small part of a scene matters at any moment. The brain addresses this through attention, using both stimulus-driven salience and current goals to decide where to spend its resources. Event cameras report only pixel changes, offering fast, low-latency sensing, while spiking neural networks on neuromorphic chips communicate through spikes as neurons do. Together they provide a natural basis for fast, low-power attention that adds detail only where it is needed.

Reference papers:

- A Model of Saliency-Based Visual Attention for Rapid Scene Analysis.
- Event Driven Bio-Inspired Attentive System for the iCub Humanoid Robot on SpiNNaker.

Requirements:

- Programming: Strong proficiency in Python and experience with SNN libraries
- Spiking Neural Networks: Basic understanding of SNN architectures and neuron models.
- Computational Neuroscience: Basic knowledge of the visual system and attention.