Adaptive Video Sampling for Energy-Efficient Object Detection

Dr. Suren Jayasuriya, Assistant Professor, Arizona State University

A critical performance requirement for embedded computer vision is energy-efficiency in order to preserve battery life for mobile and autonomous platforms. In particular, the image sensor and readout can take up a significant amount of energy in a computer vision pipeline, particularly if the sensor is capturing and processing video data in real-time. The primary mechanism in which image sensors can save energy is to limit their readout to portions of the array known as regions-of-interest (ROIs). This is a form of spatial subsampling, and can be achieved using windowing, column/row skipping, or binning in the image sensor. Such a content-driven approach can result in significant energy savings. However, this comes at the cost of potential loss of visual detail for objects, that may be necessary for end-task performance. This will certainly be the case if the subsampling approach is agnostic to the semantic information in the frames. Thus, there is an opportunity to design smart sampling approaches, which can determine the sampling pattern based on scene content, to save energy while preserving computer vision task performance.

In this project, we discuss the effectiveness of video subsampling for object detection. Object detection for videos is a critical application for self-driving cars, surveillance, and autonomous robotics. We propose adaptive algorithms to subsample video frames that utilize semantic information from a previous key frame in the video to determine sub-sampling patterns for future frames. Further, our project analyzes the computational footprint and energy cost for implementing these adaptive subsampling strategies that will be invaluable for future design of software-defined image sensors.