IMAGE RESOLUTION AND ENERGY COST IN EMBEDDED COMPUTER VISION FOR RIVER STAGE MONITORING: CHARACTERIZATION TOWARD FIELD DEPLOYMENT AT THE RÍO GRANDE DE AÑASCO

Anton Wallentin Delgado and Wilfredo Lugo-Beauchamp

Department of Computer Science and Engineering
University of Puerto Rico at Mayagüez, Mayagüez, PR 00681 USA

anton.wallentin@upr.edu

River stage in Puerto Rico is measured by a small network of gauging stations that are expensive to install and maintain, exist in only a few watersheds, and are damaged by the very events that make them most valuable, such as hurricanes María and Fiona, which produce flash flooding in rural, mountainous basins.

This project explores a different way of measuring: letting a camera do the reading. A small, battery-powered station on the riverbank photographs the river at regular intervals, and a vision model running on the station separates, in each image, water from everything else. The model does not report a water level; the level follows from geometry, from where the water–land boundary falls against fixed reference marks whose elevation was surveyed beforehand. Because the image is interpreted on board, on microcontroller-class hardware, instead of being sent elsewhere for processing, the station is cheap to install and frugal enough to run for long periods on a small battery. Every step down in resolution saves energy, so the question that organizes the work is how far it can be lowered before the station stops resolving the changes in water level that matter.

The first stage of that study is complete, and it is what this presentation reports. In a controlled setup, with the camera fixed, the water level was raised in stages to eight known levels, and at each level the station photographed the scene repeatedly at each resolution under study. Every pair of levels is therefore a change of known size, from half a centimeter upward, and the question put to the data is how often the station detected it at each resolution, which identifies the smallest change it resolves reliably. Even the smallest images tested, one sixteenth the size of the largest, detect changes below one centimeter, provided the water line is located more finely than the whole-pixel scale. A separate experiment addressed energy: the station captured bursts of images at each resolution while a current sensor and an independent logger recorded consumption, and the smallest images proved several times cheaper per capture — a smaller battery, or a longer unattended deployment. The detection result is a first estimate of the lowest usable resolution, and should be read as a floor rather than a verdict: the experiment was closed, with stable lighting and clear water, and the conditions a river imposes — turbidity, rain, glare, shadows, vegetation — were absent, so the usable resolution in the field may be coarser.

That is what the next phase will establish: the station is being deployed at a reach of the Río Grande de Añasco to photograph the river unattended for about a week. The resulting images will be used to train and compare lightweight water-segmentation models at each resolution, measuring, for each combination, both the smallest change in level the station resolves and the energy each inference consumes.