GPU VISION

Anantoly Baksheev, Alexey Spizhevoy, Vladislav Vinogradov

OpenCV – Computer Vision library

OpenCV GPU module -- contains rich set of algorithms ported to CUDA
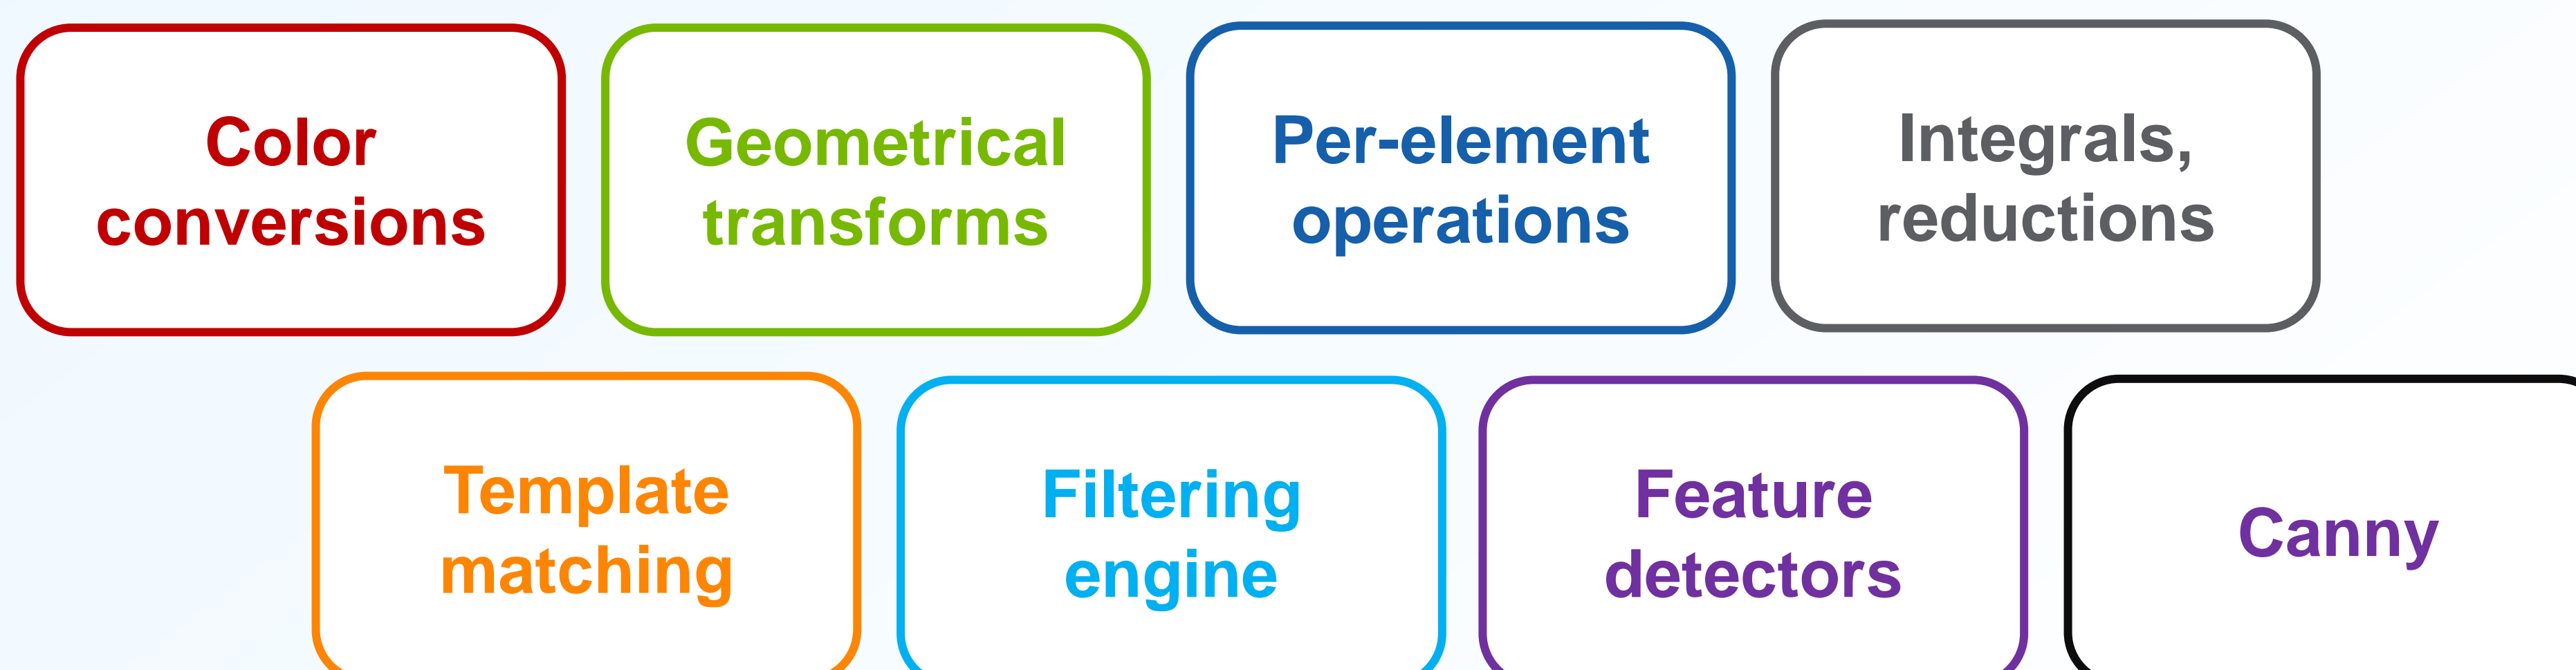- Provides convenient computer vision framework for using graphics hardware
- Consistent with the current CPU functionality (easy to switch)
- Best-in-class algorithms with the best performance

Applications: driver assistance, robotics, video surveillance, object detection, aerial photography, human–computer interaction

Two tiers: Primitive image processing blocks and high-level vision algorithms

Low-level image processing blocks

- Average speed-up: **33×**
  - versus heavily optimized CPU version, without data transfer

Color conversions	Geometrical transforms	Per-element operations	Integrals, reductions

Template matching	Filtering engine	Feature detectors	Canny

- Uses NVIDIA Performance Primitives library internally and extends it with higher level computer vision routines
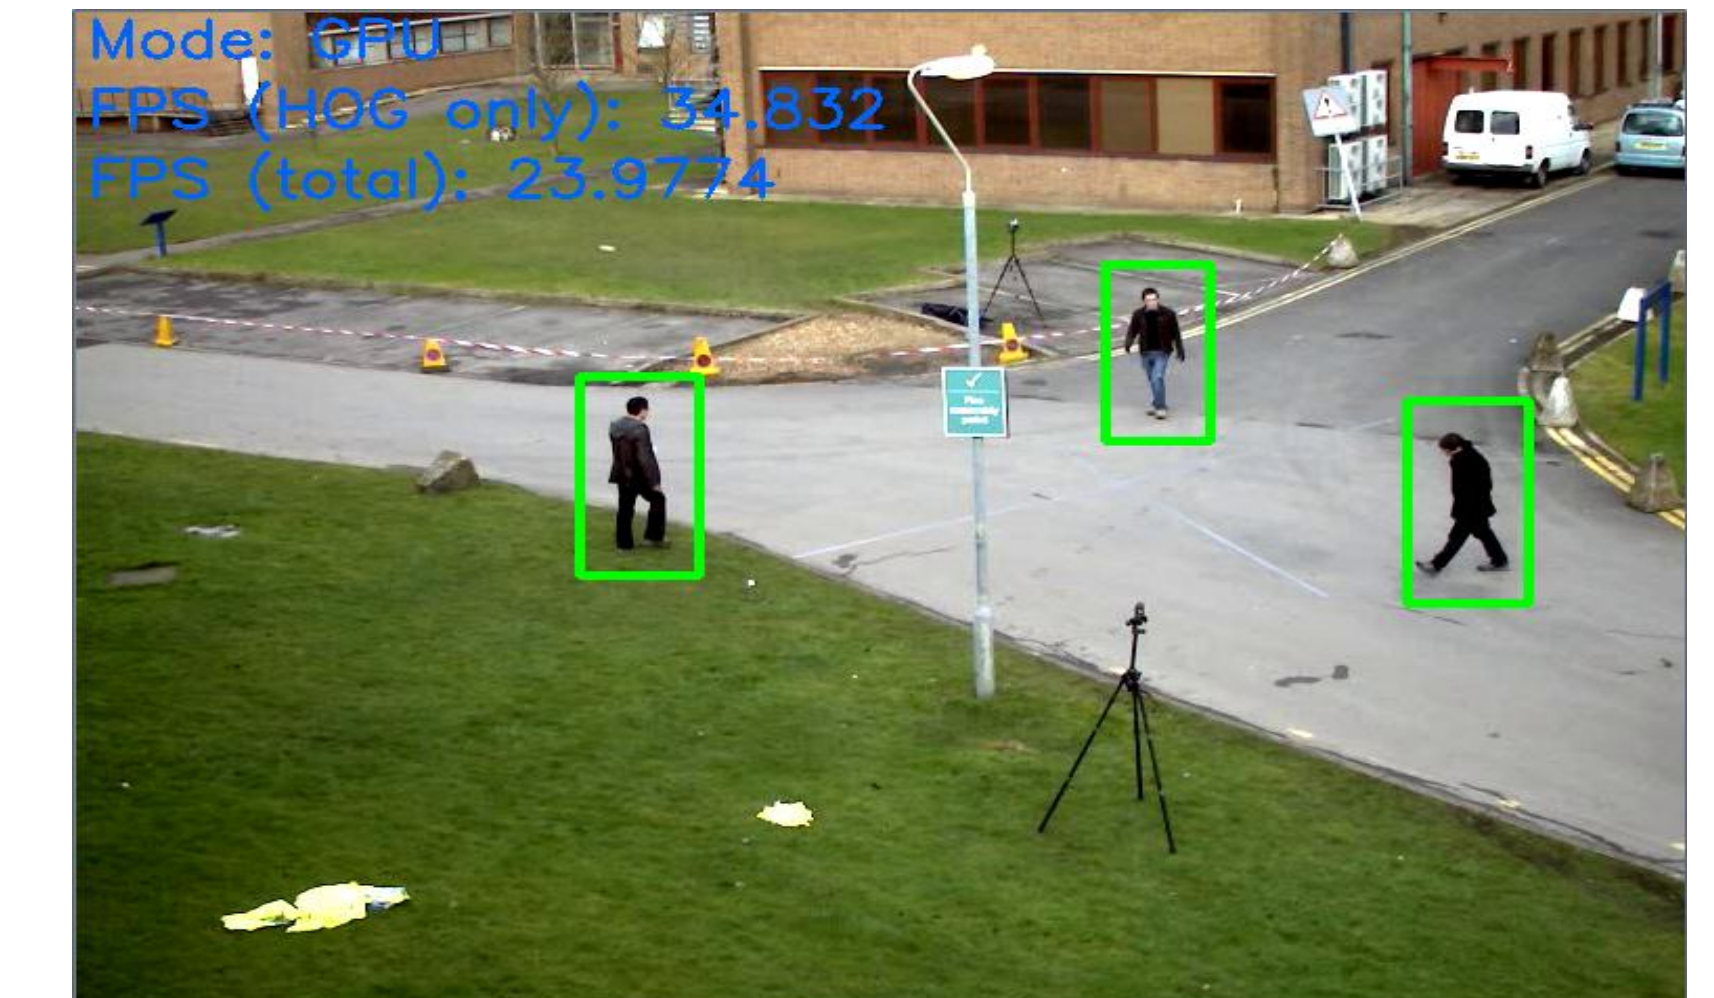
Other software partly accelerated with GPU module

- Visual Odometry pipeline in Robotics Operation System **(2.5×)**
- Textured object detection pipeline in Robotics Operation System **(2-6×)**
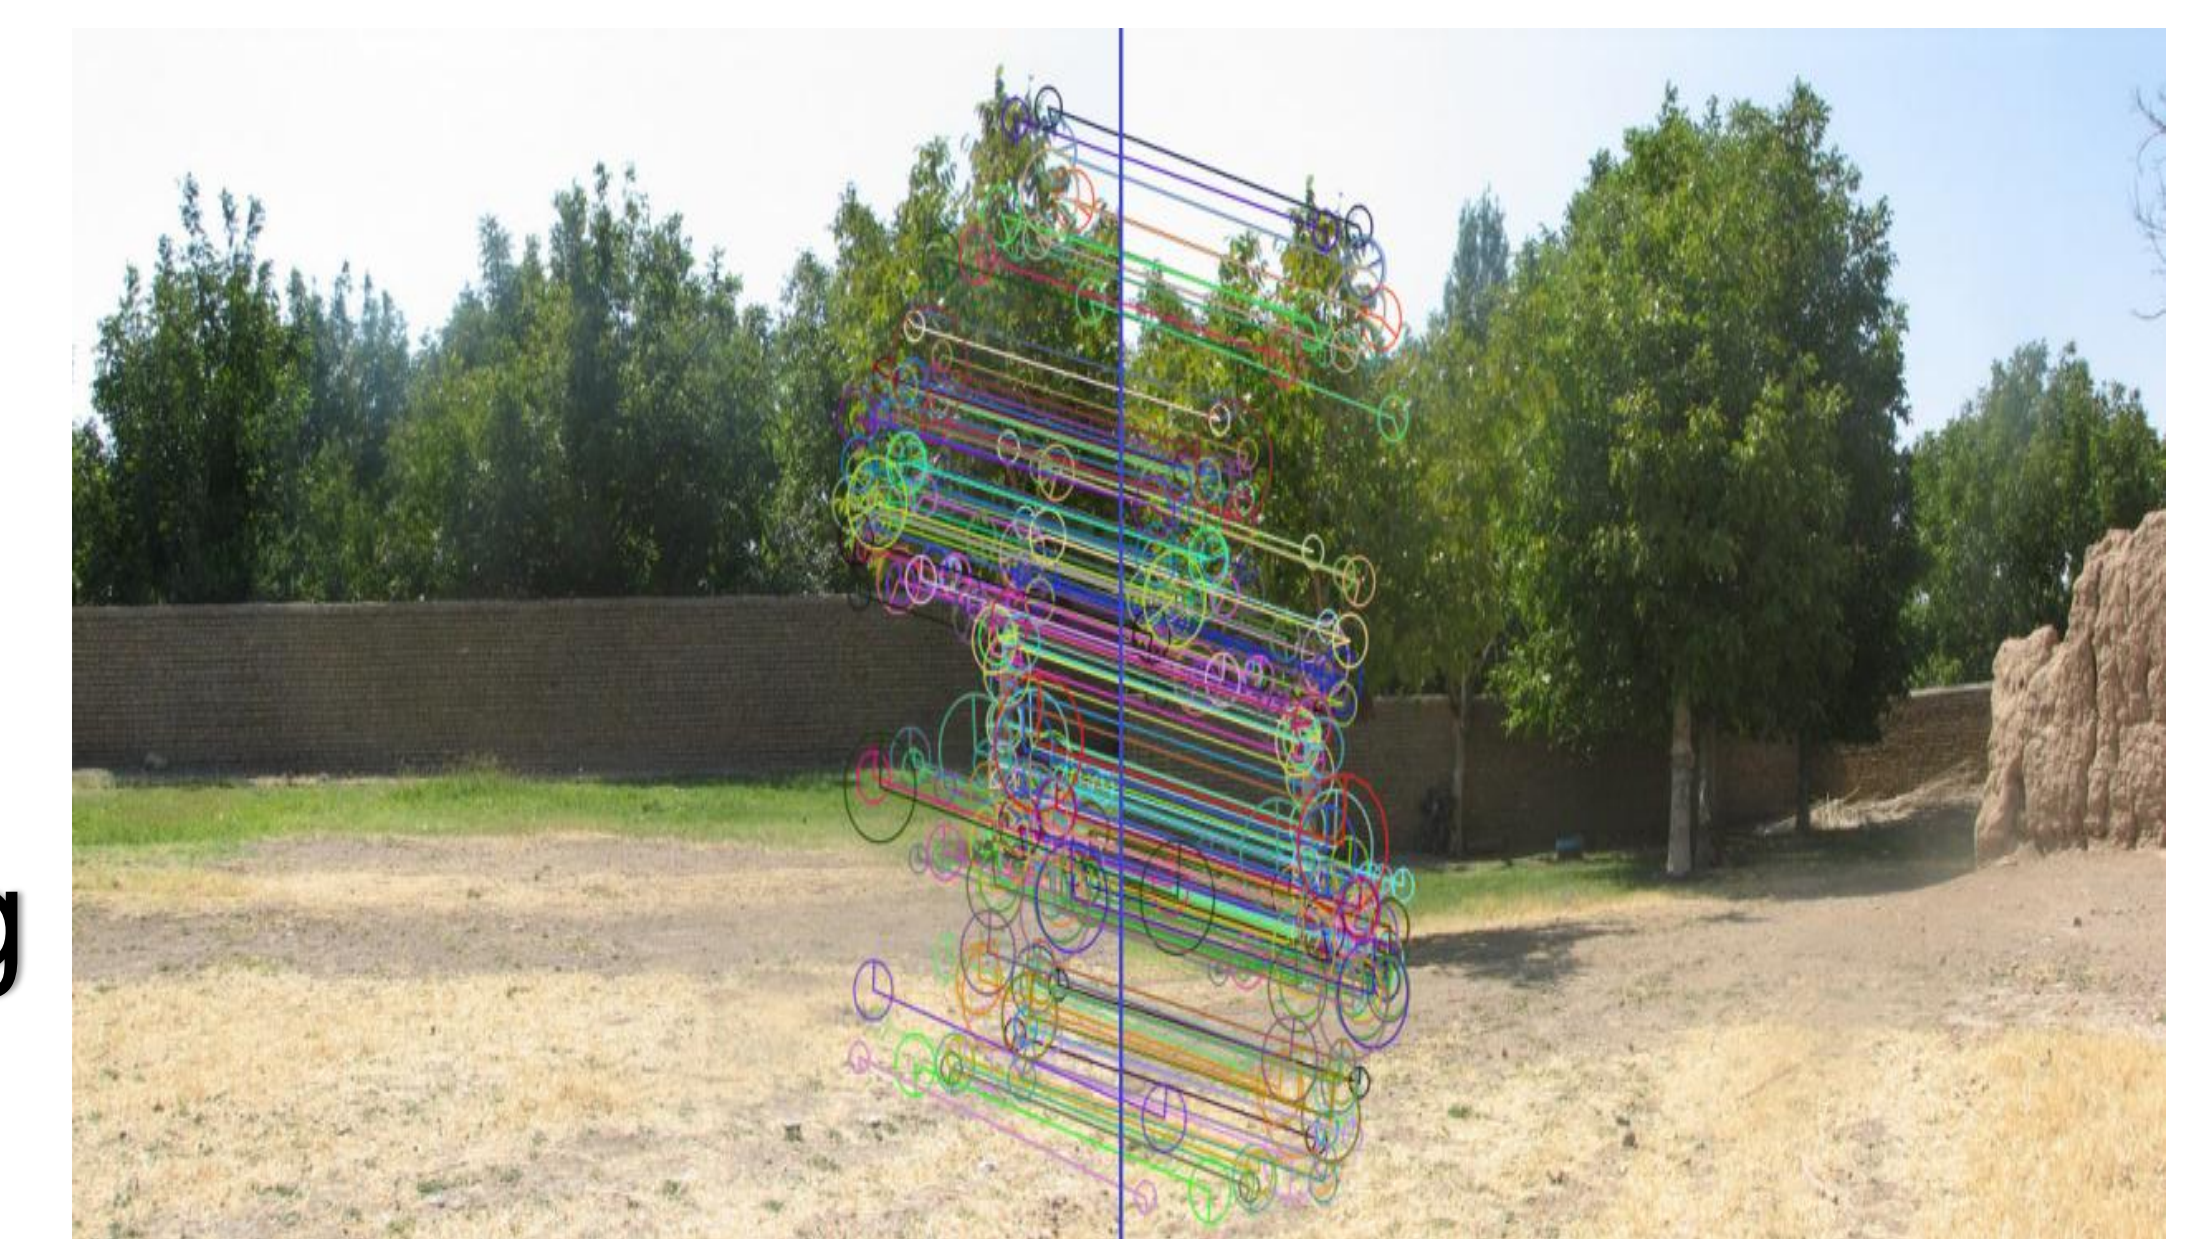- OpenCV panorama stitching module (in progress)

High-level functionality

Face detection (6×)
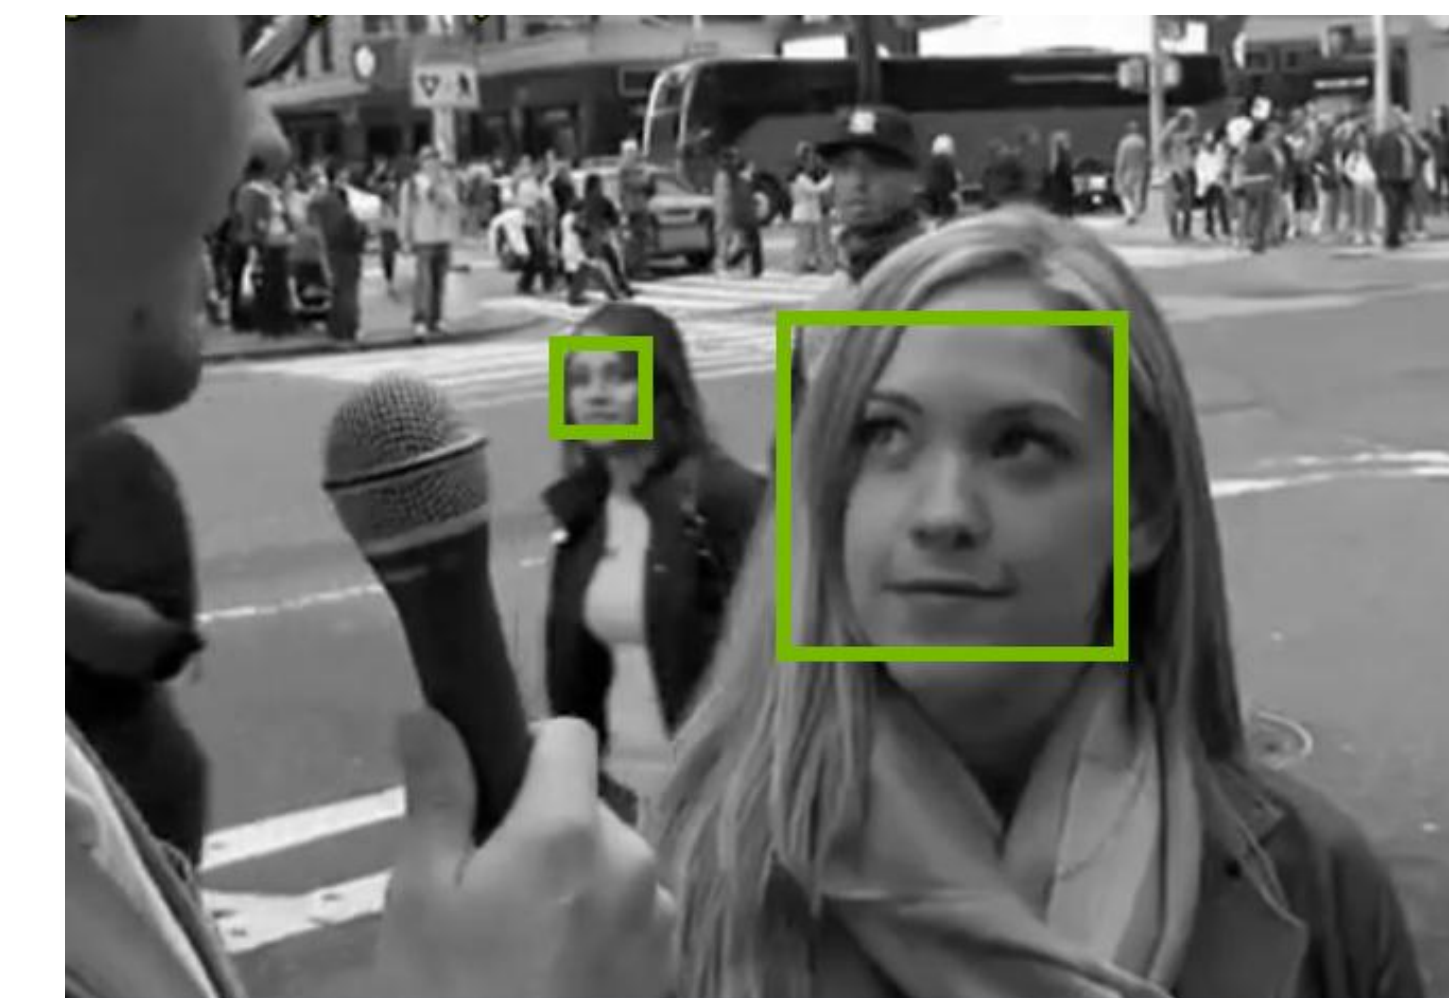
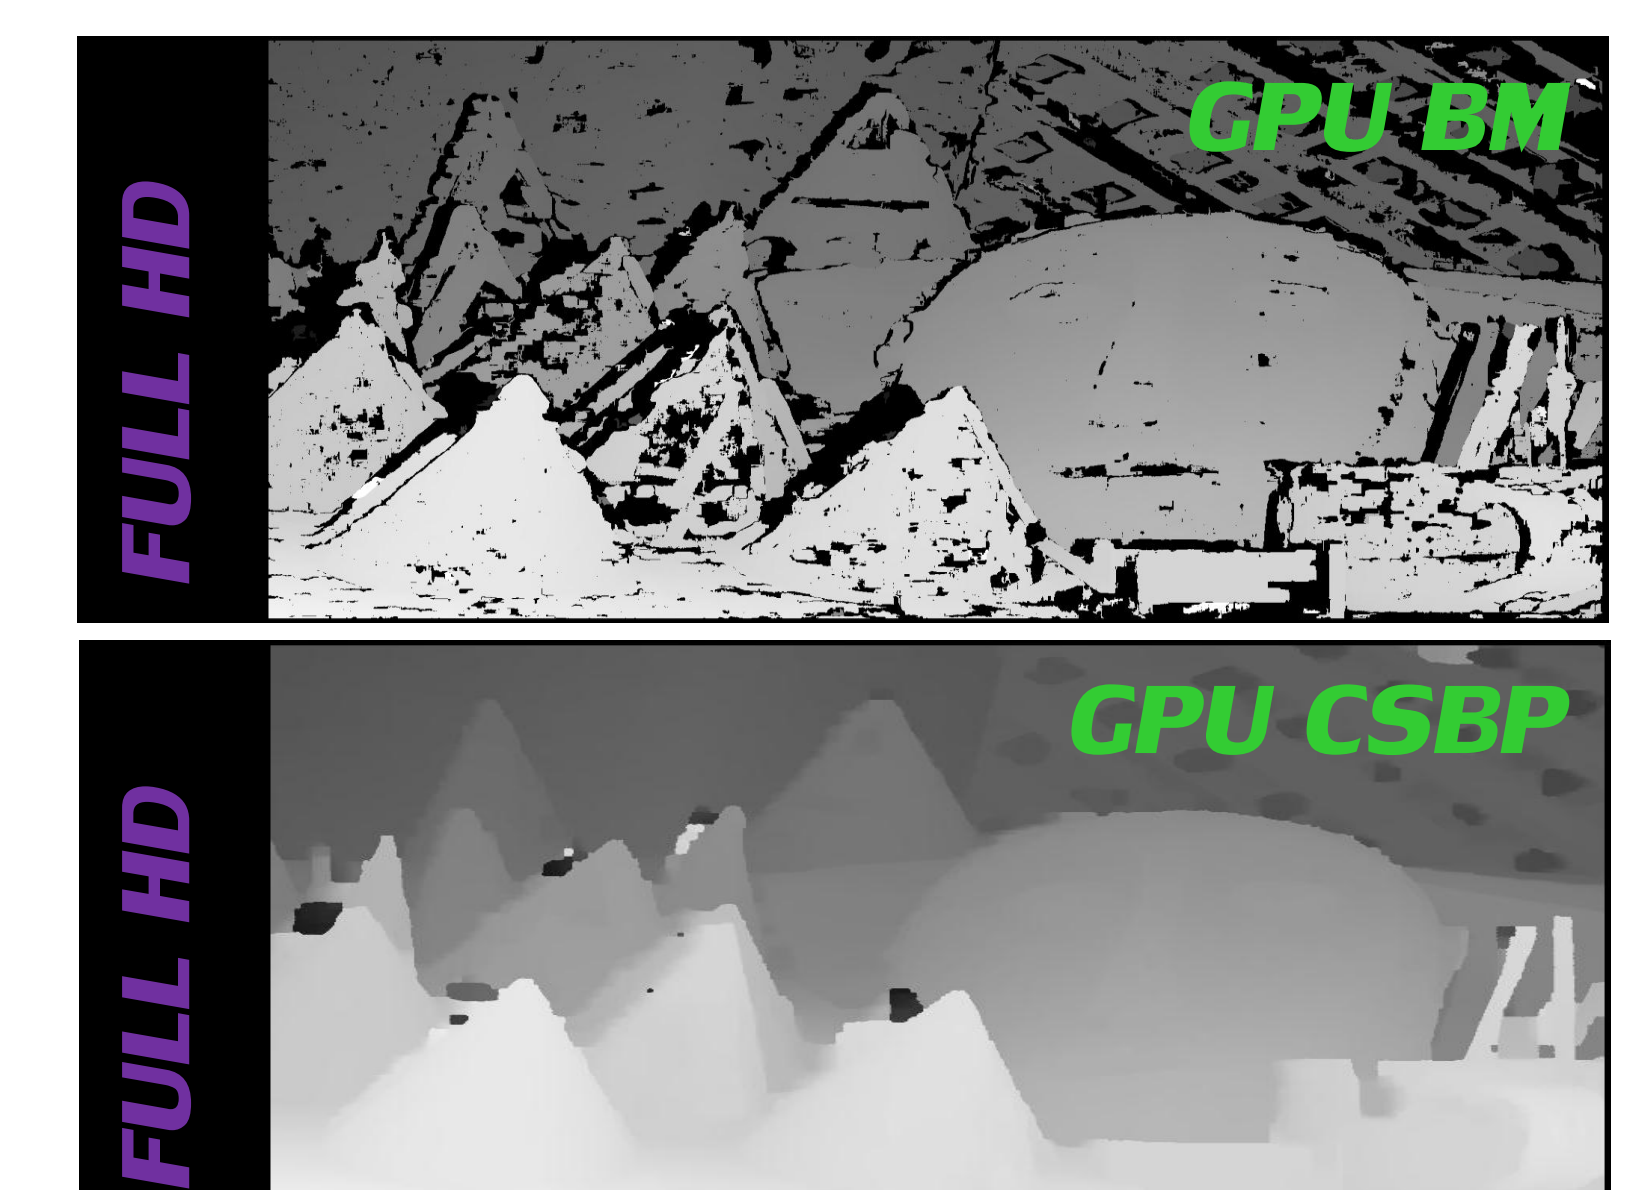### Pedestrian detection (7×)

- SURF features **(12×)**

- Brute Force Feature Matching **(20-30×)**

Stereo Correspondence

- Block Matching **(7×)**
  - FullHD in real time on 2 GPUs

- Belief propagation **(20×)**
- Constant space BP **(50×)**
  - versus original author's code

Other algorithms

- Brox optical flow (2fps)
- Lucas–Kanade optical flow (in progress)
- Farnerbeck optical flow (in progress)
- ORB features (3-6×, expected to be run on future Tegra with CUDA)

* All performance numbers were measured for NVIDIA Tesla C2050, Intel Core i5-760 2.8Ghz (4 cores). OpenCV was compiled with SSE, TBB support