University of Zurich[UZH] | ETH zürich
Institute of Informatics – Institute of Neuroinformatics
ROBOTICS & PERCEPTION GROUP

Event Cameras

Davide Scaramuzza

Slides, Publications, Videos, Code:
http://rpg.ifi.uzh.ch/

Research Overview

Research Overview

Real-time, Onboard Computer Vision and Control for **Autonomous, Agile** Drone **Flight**

Falanga et al., **The Foldable Drone: A Morphing Quadrotor that can Squeeze and Fly**, RAL'19. PDF. Videos.
Featured in IEEE Spectrum.

What does it take to fly as **good as or better** than human pilots?

WARNING! This drone flown is NOT autonomous; it is operated by a human pilot.
Human pilots take years to acquire the skills shown in this video.
Can we use drone racing as a proxy to learn agile flight?

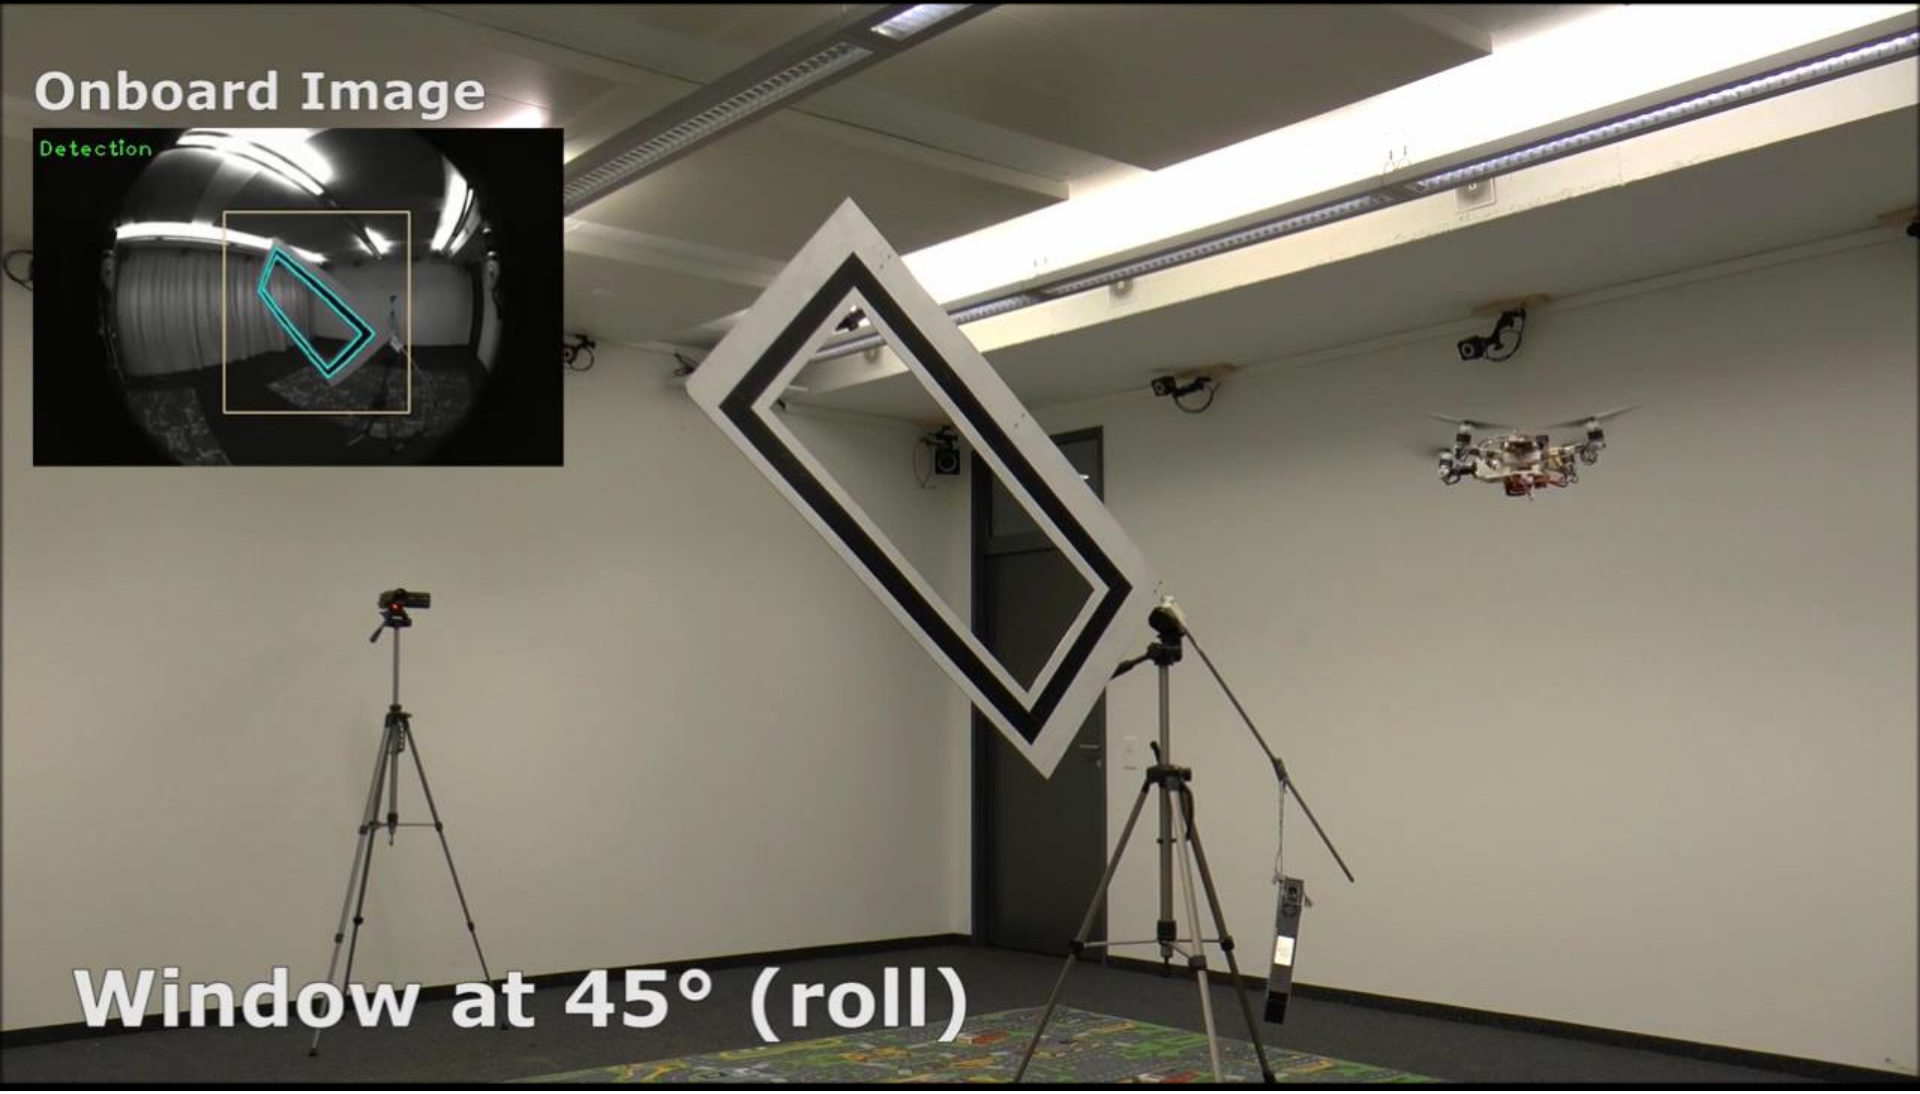

Falanga et al., Aggressive Quadrotor Flight through Narrow Gaps with Onboard Sensing and Computing using Active Vision, ICRA'17. PDF. Video. Featured in IEEE Spectrum.

Autonomous Drone Racing from a single Flight Demonstration

Kaufmann et al., *Deep Drone Racing*, CORL’18, **Best System Paper Award**. PDF. Video.
Kaufmann et al., *Beauty and the Beast: Optimal Methods Meet Learning for Drone Racing*, ICRA’19. PDF. Video
Deployed to win the IROS Autonomous Drone Racing Competition, *IROS’18*. Video.

Event Cameras

Open Challenges in Computer Vision

The past 60 years of research have been devoted to frame-based cameras ...but they are not good enough!

Latency

Motion blur

Dynamic Range

Event cameras do not suffer from these problems!

What is an event camera?

- Novel sensor that measures only **motion in the scene**
- **First commercialized in 2008** by T. Delbruck (UZH&ETH) under the name of Dynamic Vision Sensor (DVS)
- **Low-latency** (~ 1 μs)
- **No motion blur**
- **High dynamic range** (140 dB instead of 60 dB)
- **Ultra-low power** (mean: 1mW vs 1W)

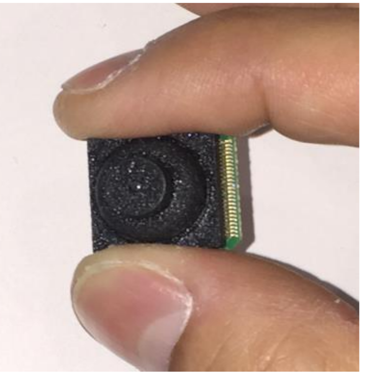

Mini DVS sensor from IniVation.com

Traditional vision algorithms cannot be used because:

- **Asynchronous** pixels
- **No intensity information** (only binary intensity changes)

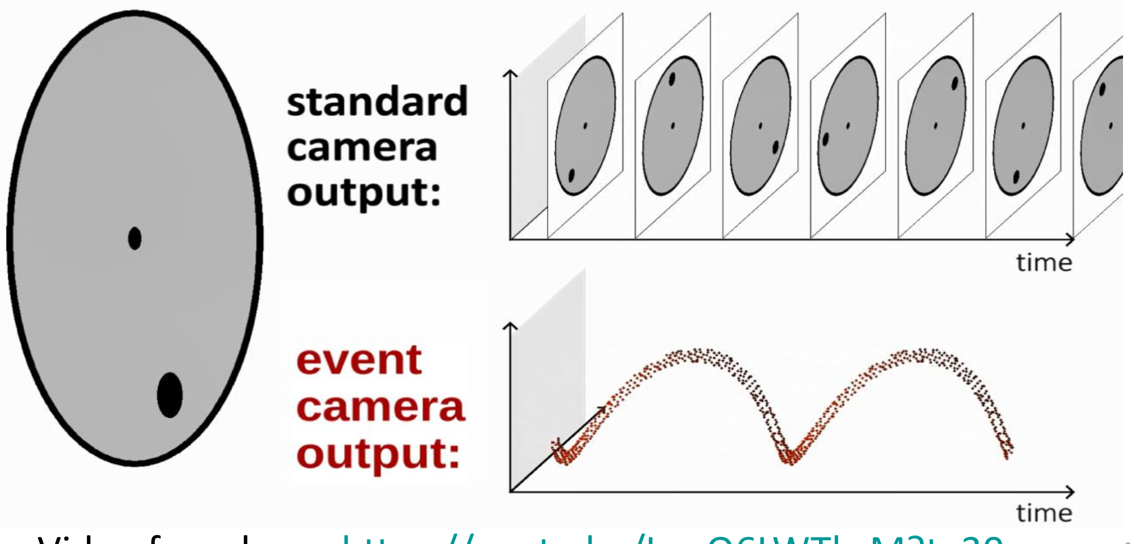

Video from here: https://youtu.be/LauQ6LWTkxM?t=30

What is an event camera?

- Novel sensor that measures only **motion in the scen**
- **First commercialized in 2008** by T. Delbruck (UZH&ET under the name of Dynamic Vision Sensor (DVS)
- **Low-latency** (~ 1 μs)
- **No motion blur**
- **High dynamic range** (140 dB instead of 60 dB)
- **Ultra-low power** (mean: 1mW vs 1W)

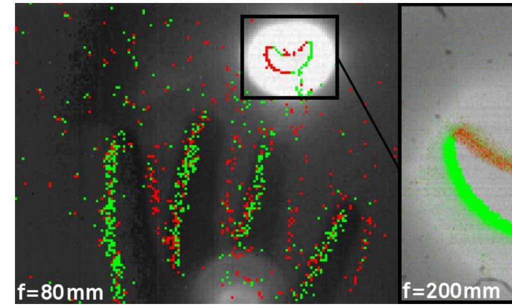

Image of the solar eclipse captured by a DVS

Traditional vision algorithms cannot be used because:

- **Asynchronous** pixels
- **No intensity information** (only binary intensity changes)

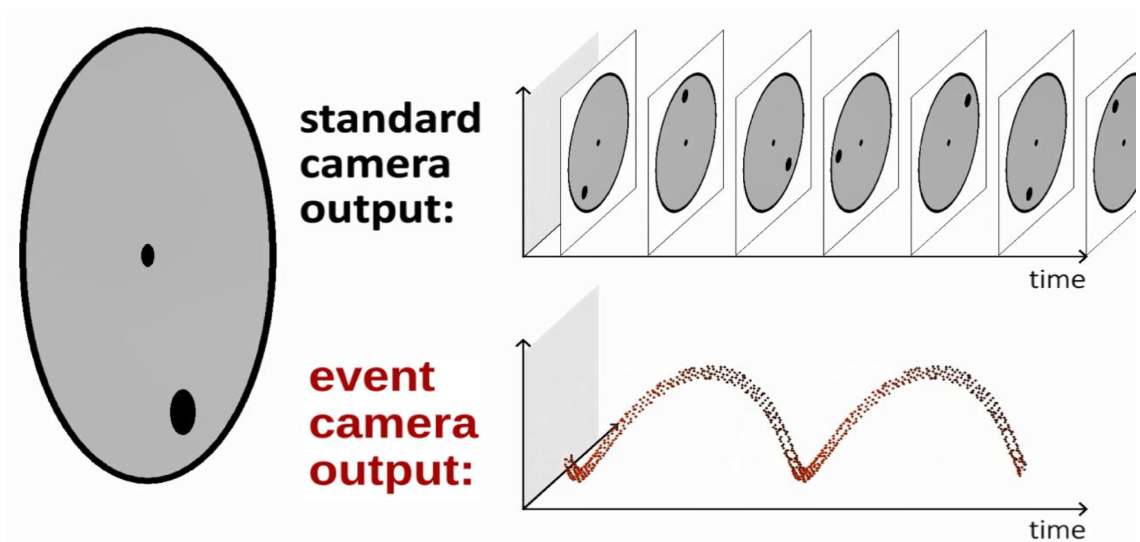

Video from here: https://youtu.be/LauQ6LWTkxM?t=30

Sequence: Fan and Coin

Conventional Frames

Events

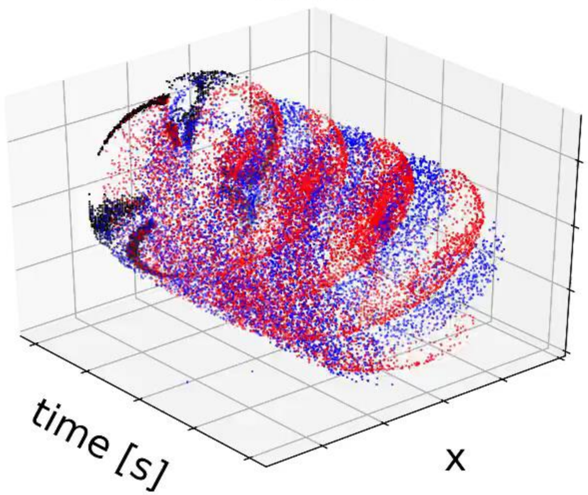

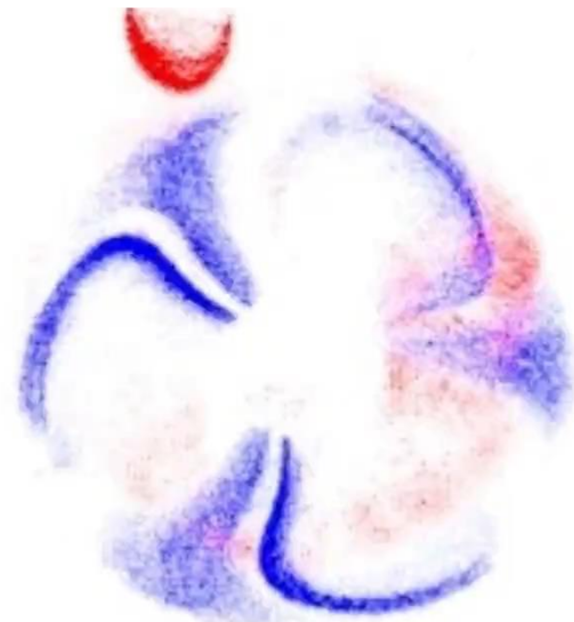

Events in the **space-time** domain (x, y, t)

Events in the **image domain** (x, y)
Integration time: $\Delta T = 10$ ms

Stoffregen et al., Motion Segmentation by Motion Compensation, ICCV'19. [PDF](). [Video]().

Event Camera output with Motion

Standard Camera

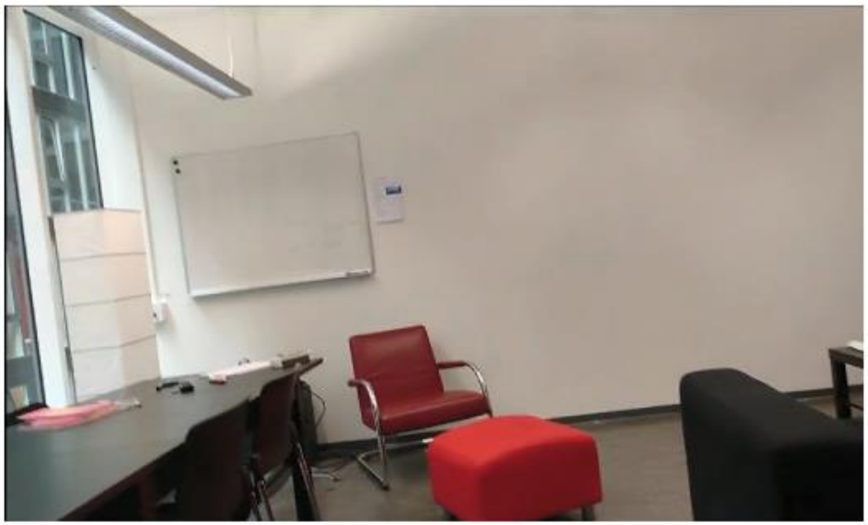

Event Camera (ON, OFF events)

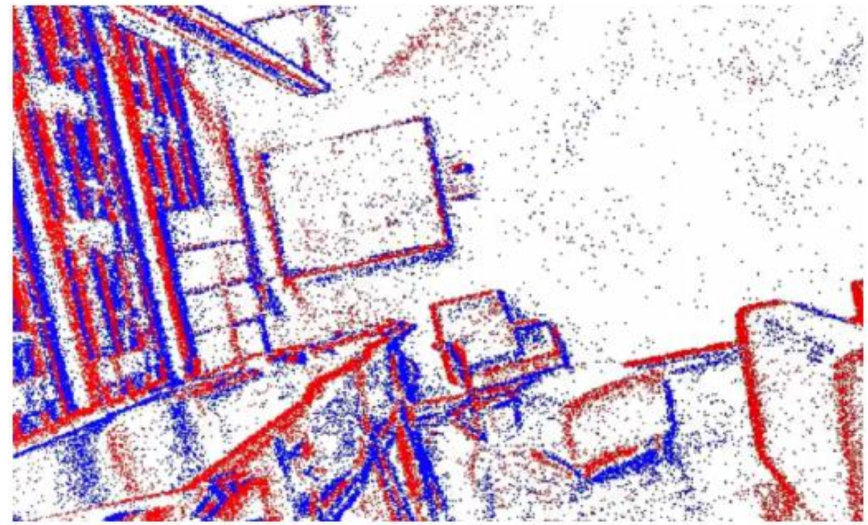

ΔT = 40 ms

Application 1: Event-based Lucas-Kanade Tracking (E-KLT)

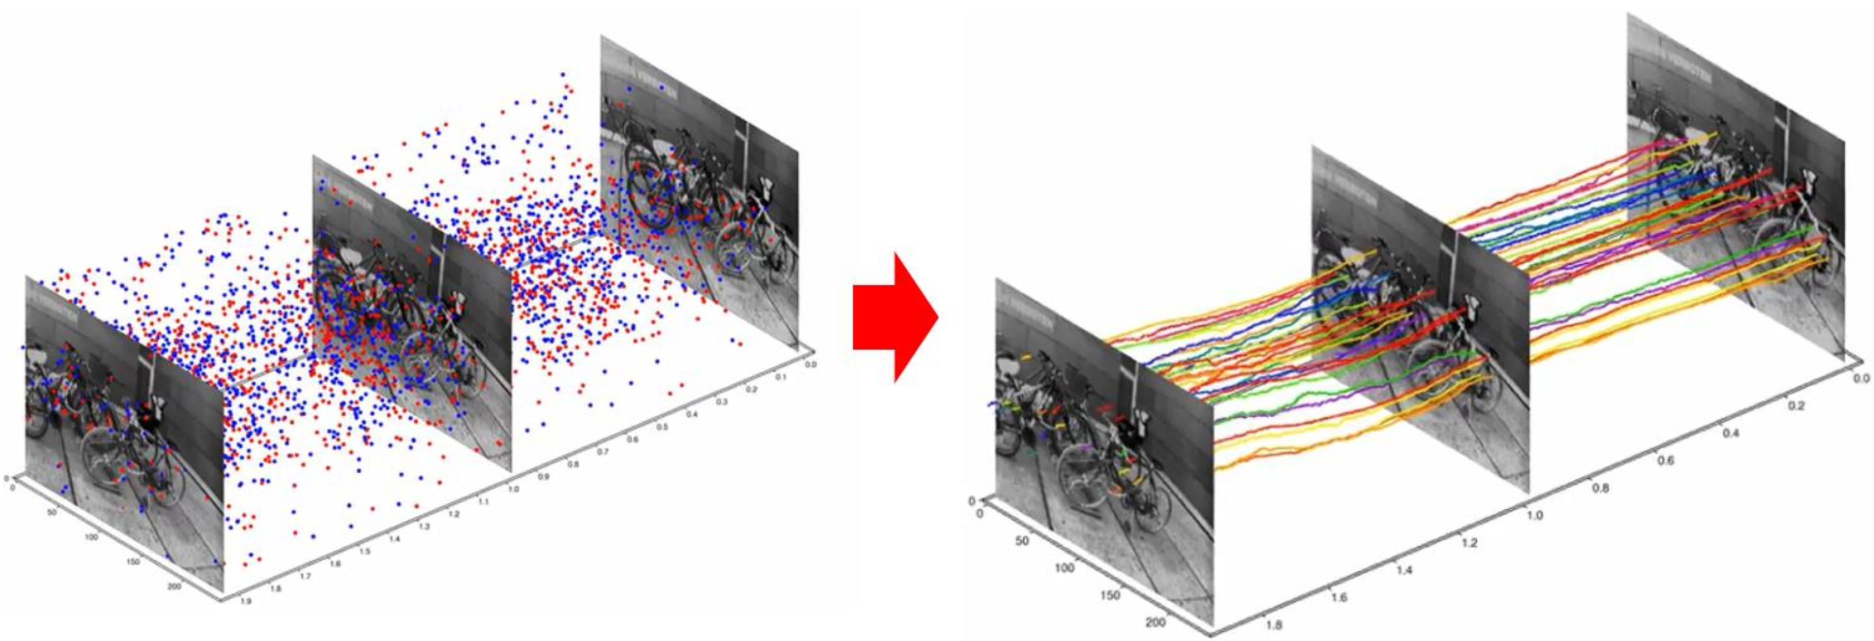

Gehrig et al., EKLT: Asynchronous, Photometric Feature Tracking using Events and Frames, IJCV'19. PDF. Video.

UltimateSLAM: Frames + Events + IMU

85% accuracy gain over standard Visual-Inertial SLAM in HDR and high speed scenes

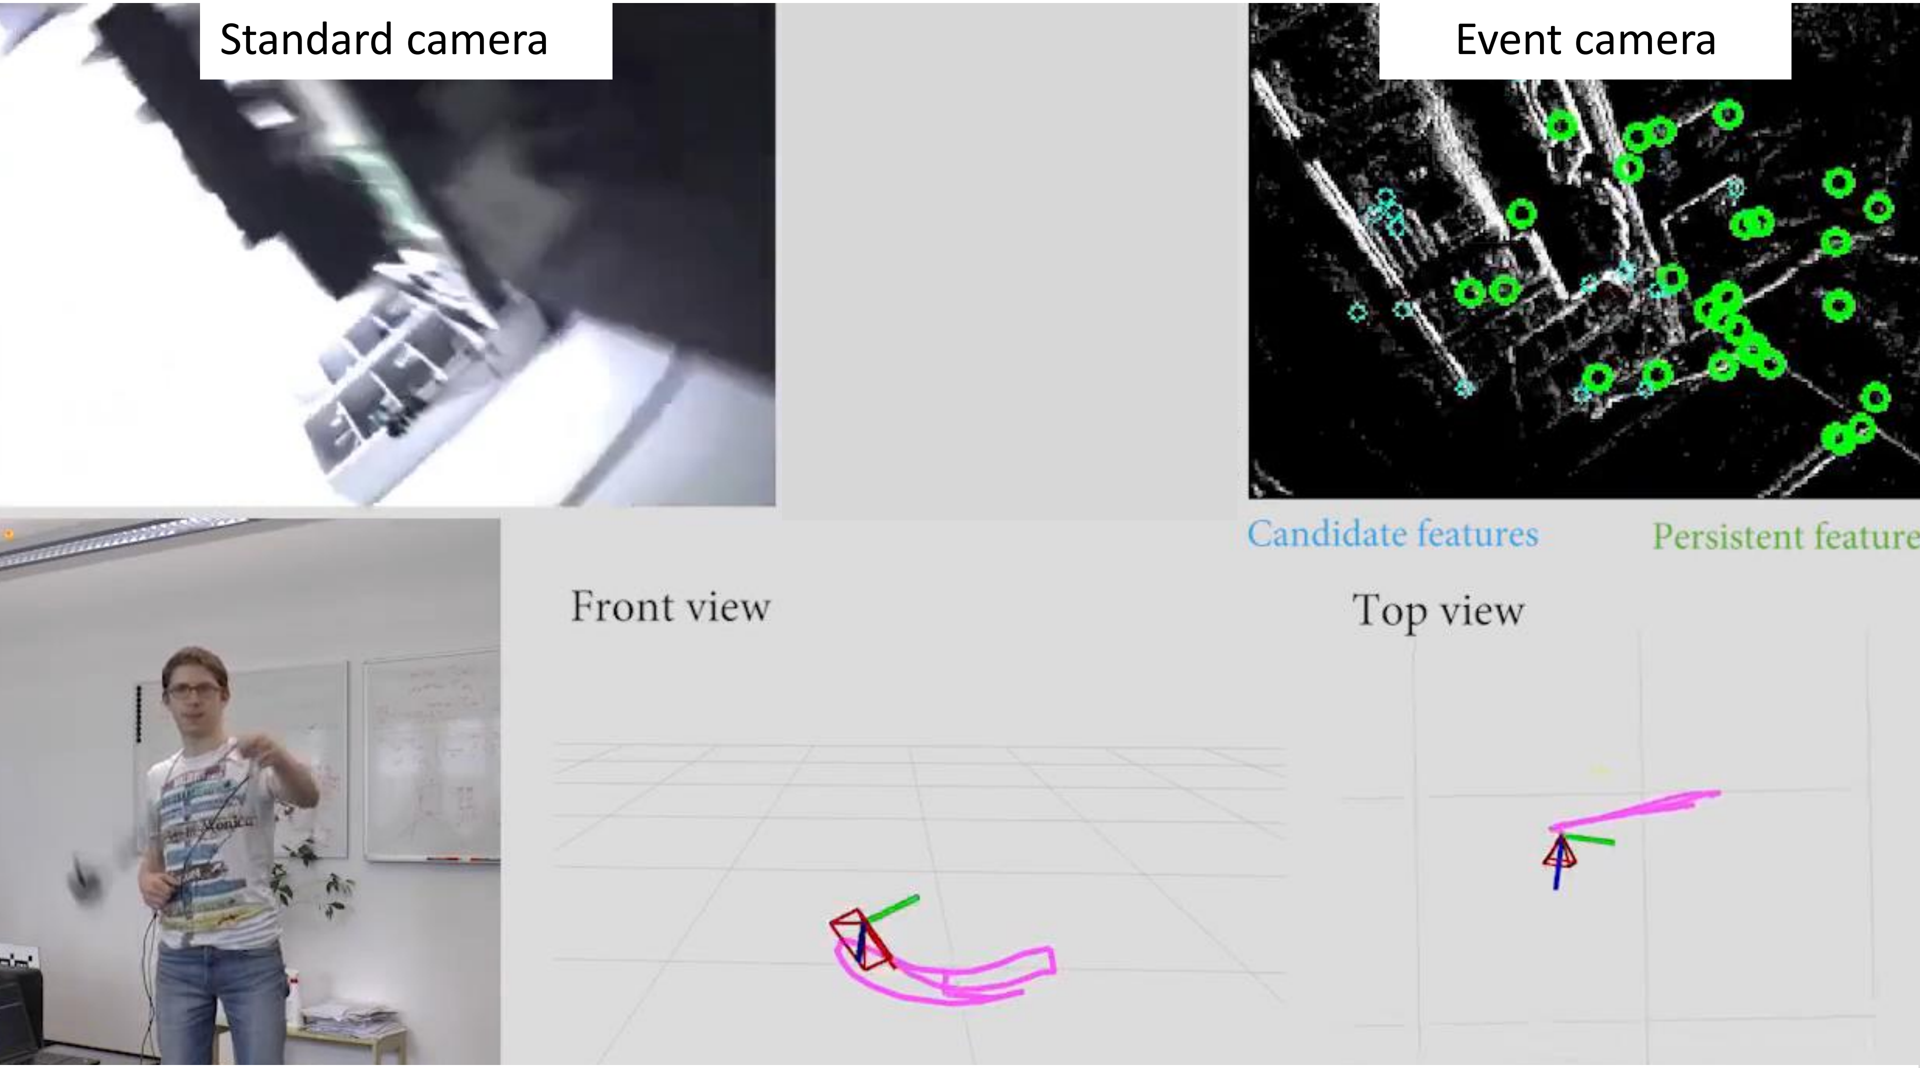

Rosinol et al., Ultimate SLAM? **RAL'18 – Best RAL'18 Paper Award Honorable Mention** [PDF](). [Video](). [IEEE Spectrum]().
Mueggler et al., Continuous-Time Visual-Inertial Odometry for Event Cameras, **TRO'18**. [PDF]()

Application 3: Autonomous Drone Navigation in Low Light

UltimateSLAM running on board (CPU: Odroid XU4)

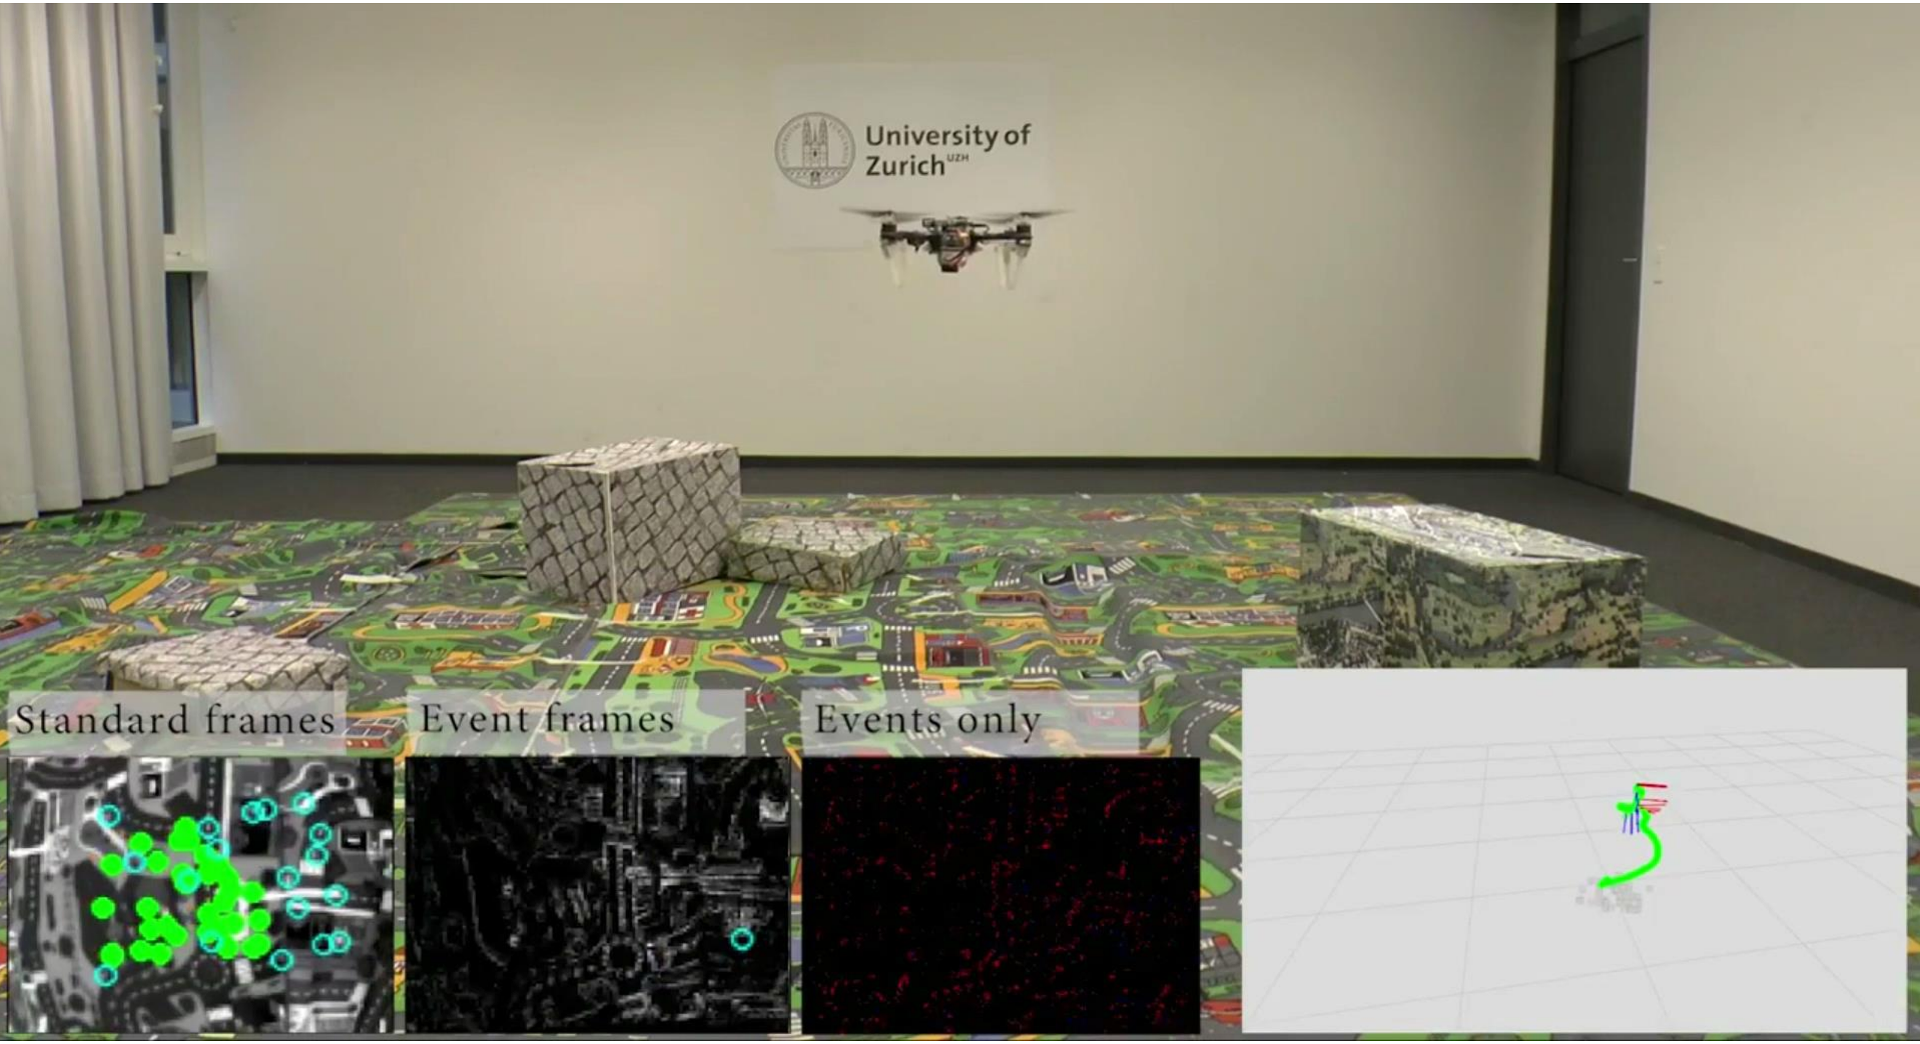

Rosinol et al., Ultimate SLAM? **RAL'18 – Best RAL'18 Paper Award Honorable Mention** PDF. Video. IEEE Spectrum.
Mueggler et al., Continuous-Time Visual-Inertial Odometry for Event Cameras, **TRO'18**. PDF

Application 4: Motion Segmentation

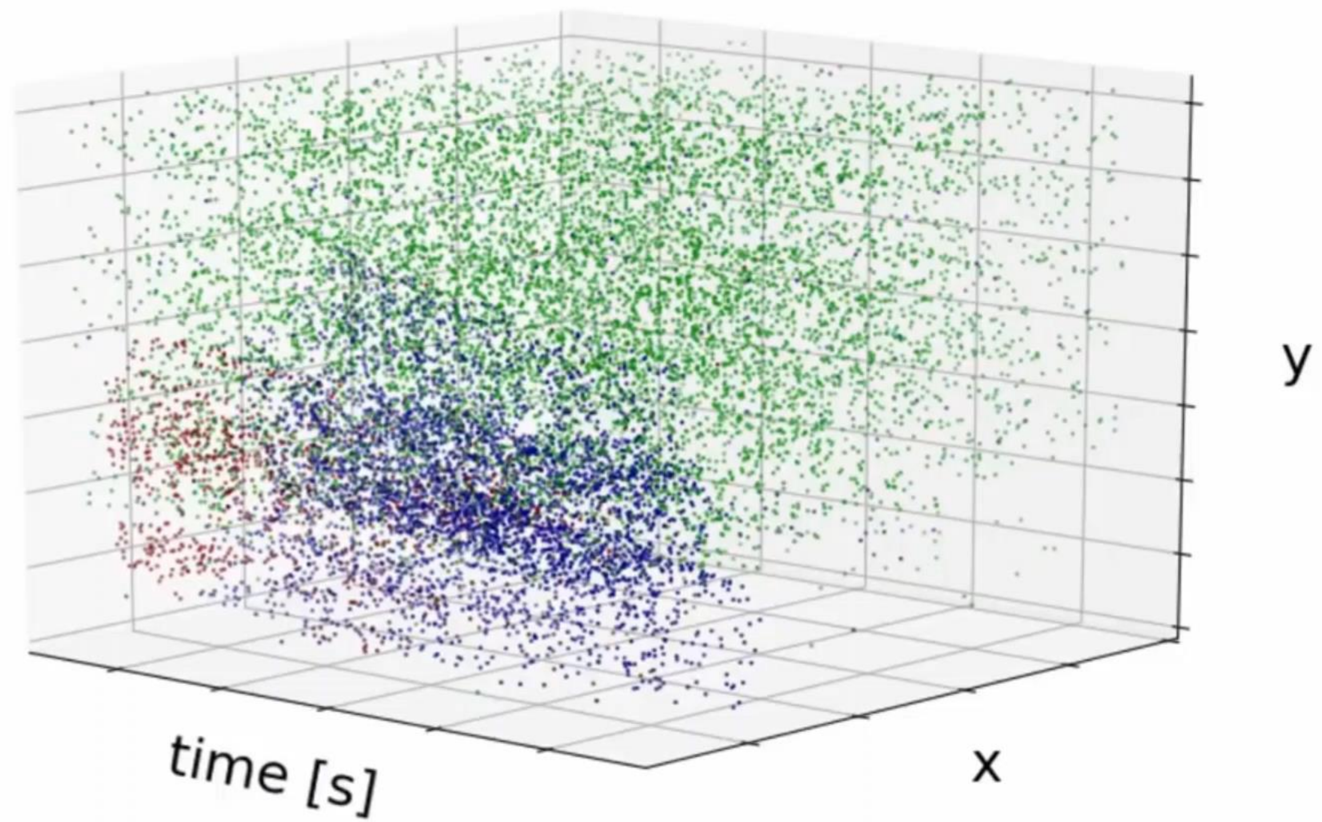

Stoffregen et al., Motion Segmentation by Motion Compensation, ICCV’19. [PDF](). [Video]().

Application 5: Dynamic Obstacle Avoidance

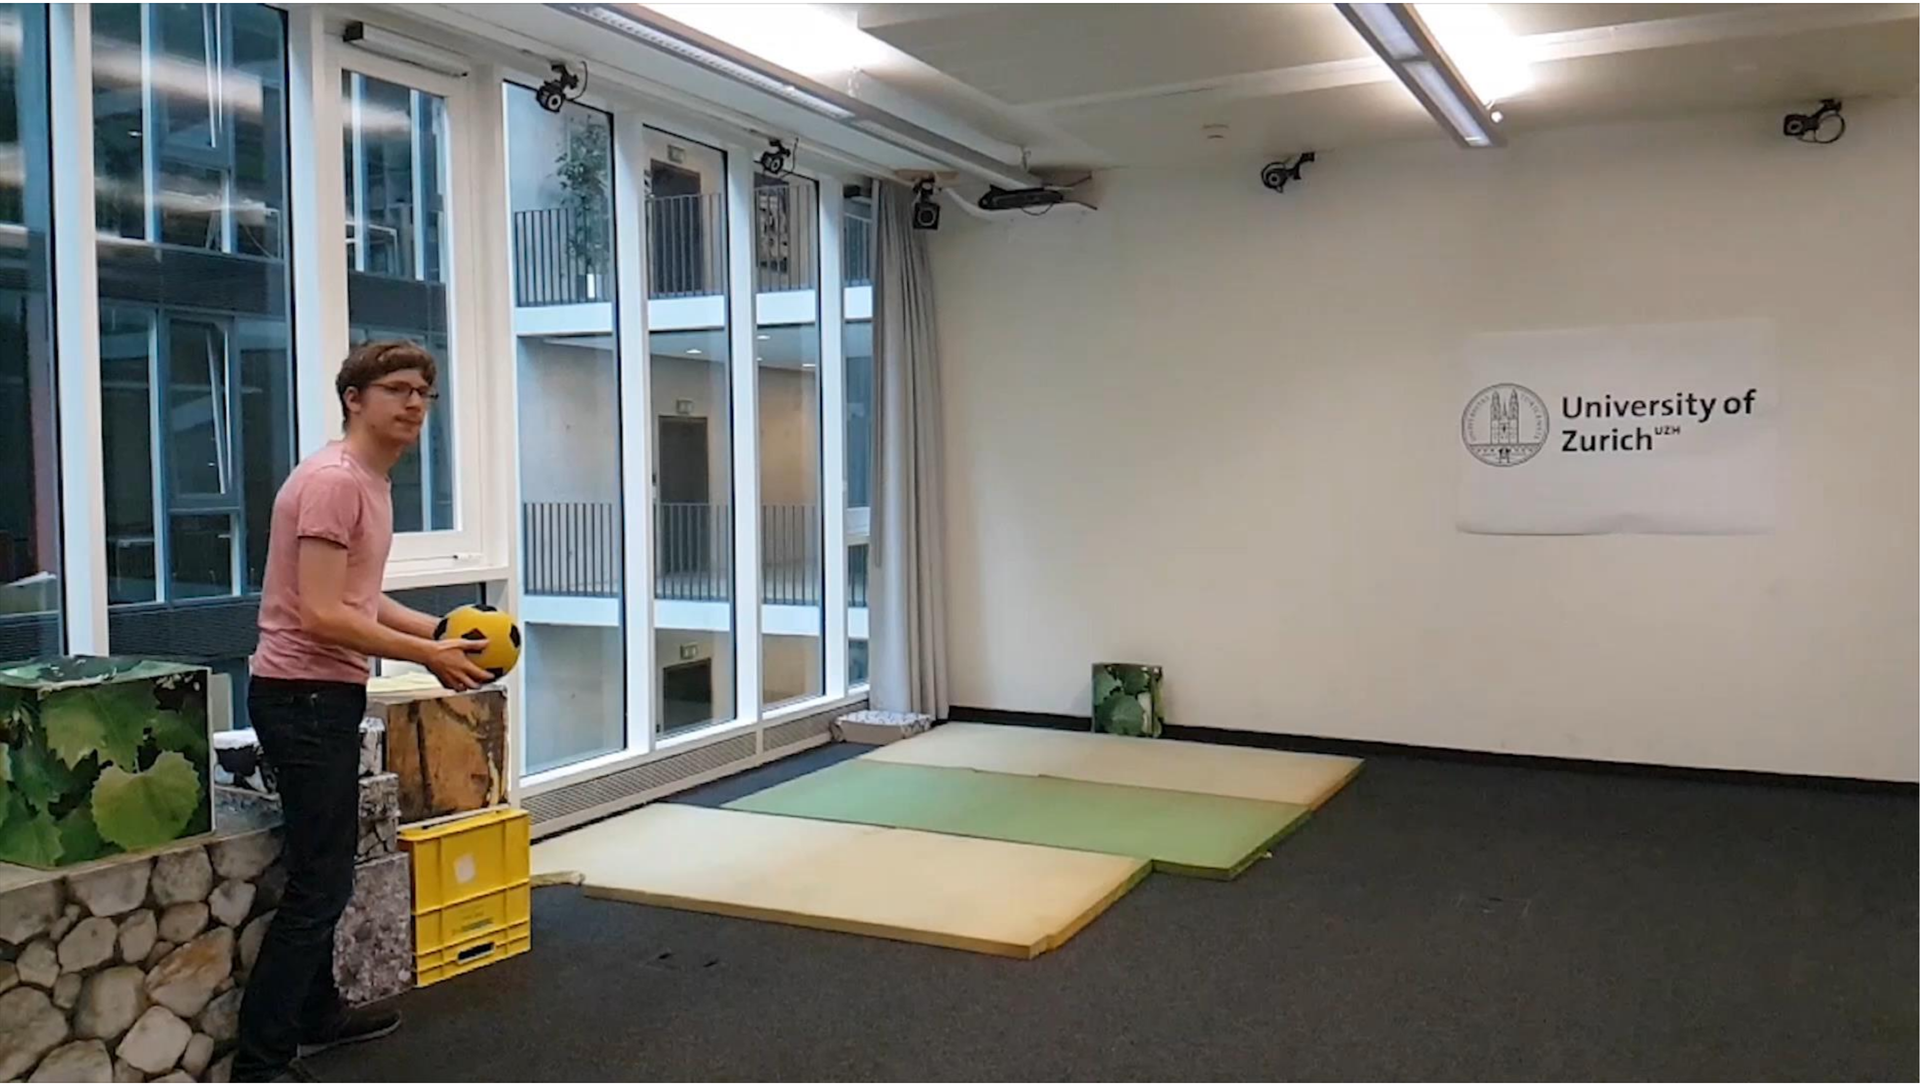

Falanga et al. *How Fast is too fast? The role of perception latency in high speed sense and avoid*, RAL'19.
PDF. Video. Featured in IEEE Spectrum.

Image Reconstruction

Image Reconstruction from Events

Events

Reconstructed image from events

Rebecq et al., “High Speed and High Dynamic Range Video with an Event Camera”, PAMI’19. PDF Video Code

Overview

- **Recurrent neural network** (main module: Unet)
- Input: sequences of *event tensors* (spatio-temporal volumes of events[3])
- **Trained in simulation only**, without seeing a single real image

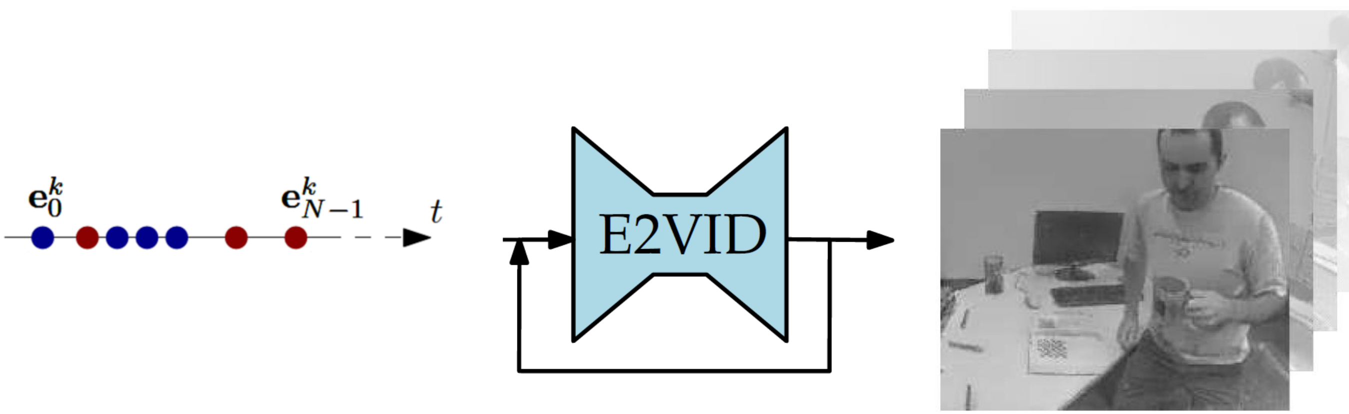

Rebecq et al., "High Speed and High Dynamic Range Video with an Event Camera", PAMI'19. PDF Video Code

Event Camera Simulator

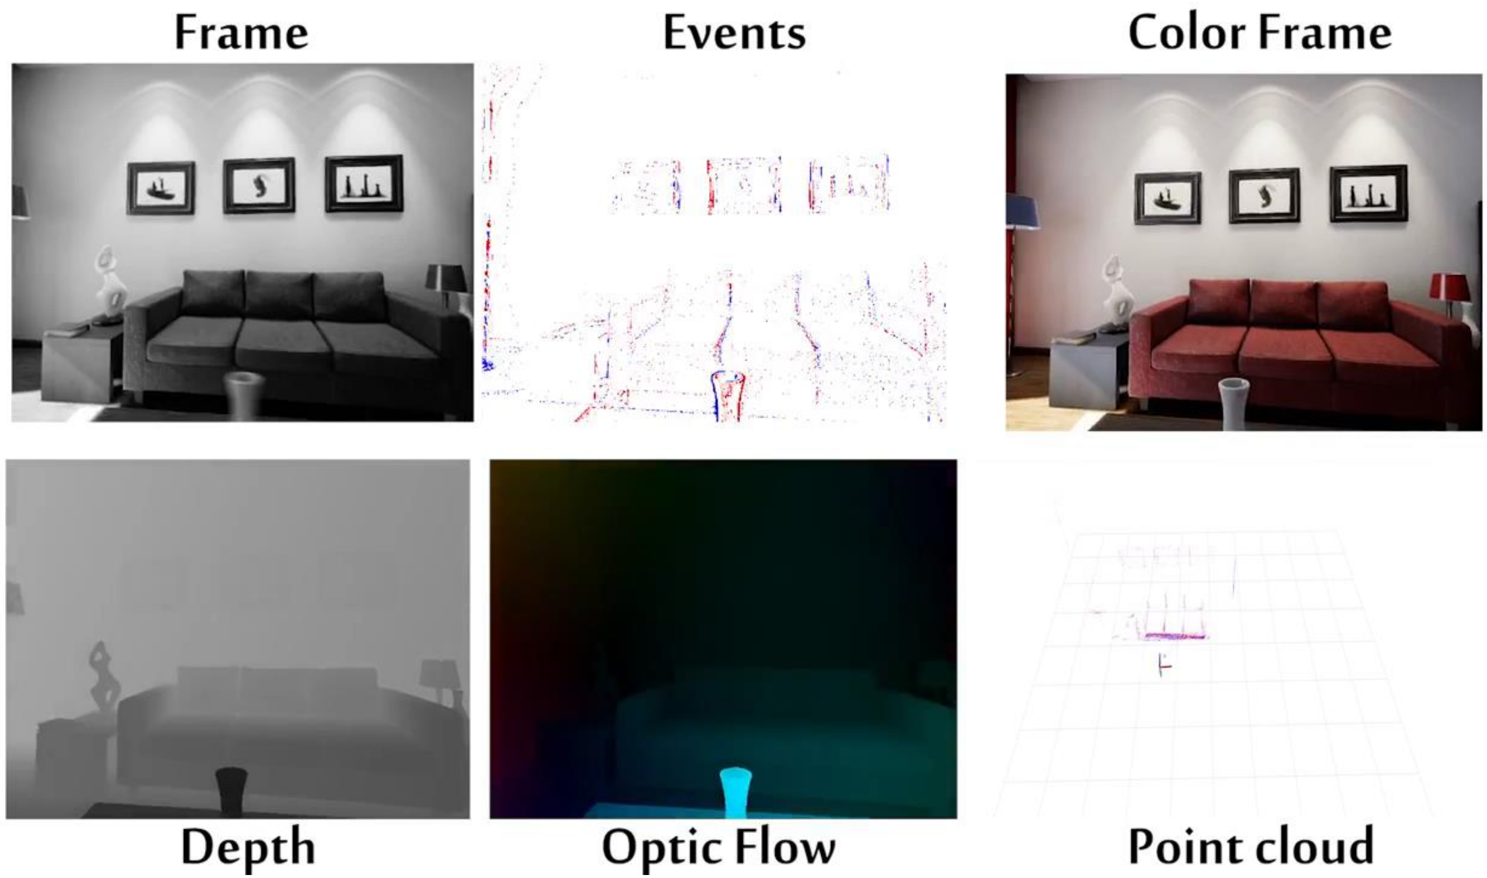

Open Source: http://rpg.ifi.uzh.ch/esim.html

Rebecq, ESIM: an Open Event Camera Simulator, CORL'18. PDF, YouTube, Project page

High Speed Video Reconstruction Results

Popping a water balloon

Recall: trained in simulation only!

https://youtu.be/eomALySSGVU

Huawei P20 Pro (240 FPS)

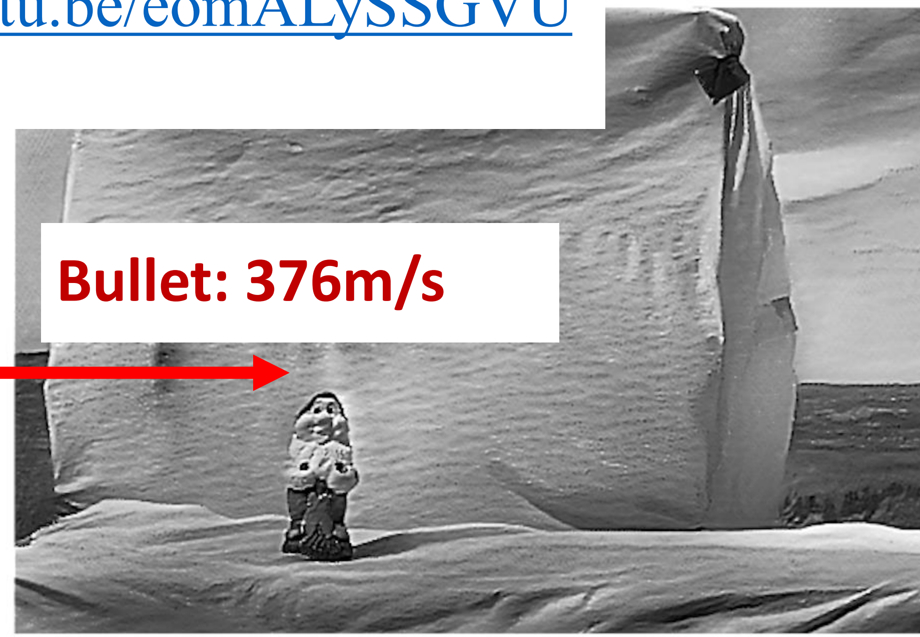

Our reconstruction (5400 FPS)

We used Samsung DVS

Source Code: https://github.com/uzh-rpg/rpg_e2vid

Real time

Rebecq et al., "High Speed and High Dynamic Range Video with an Event Camera", PAMI'19. PDF Video Code

HDR Video: Driving out of a tunnel

Recall: trained in simulation only!

https://youtu.be/eomALySSGVU

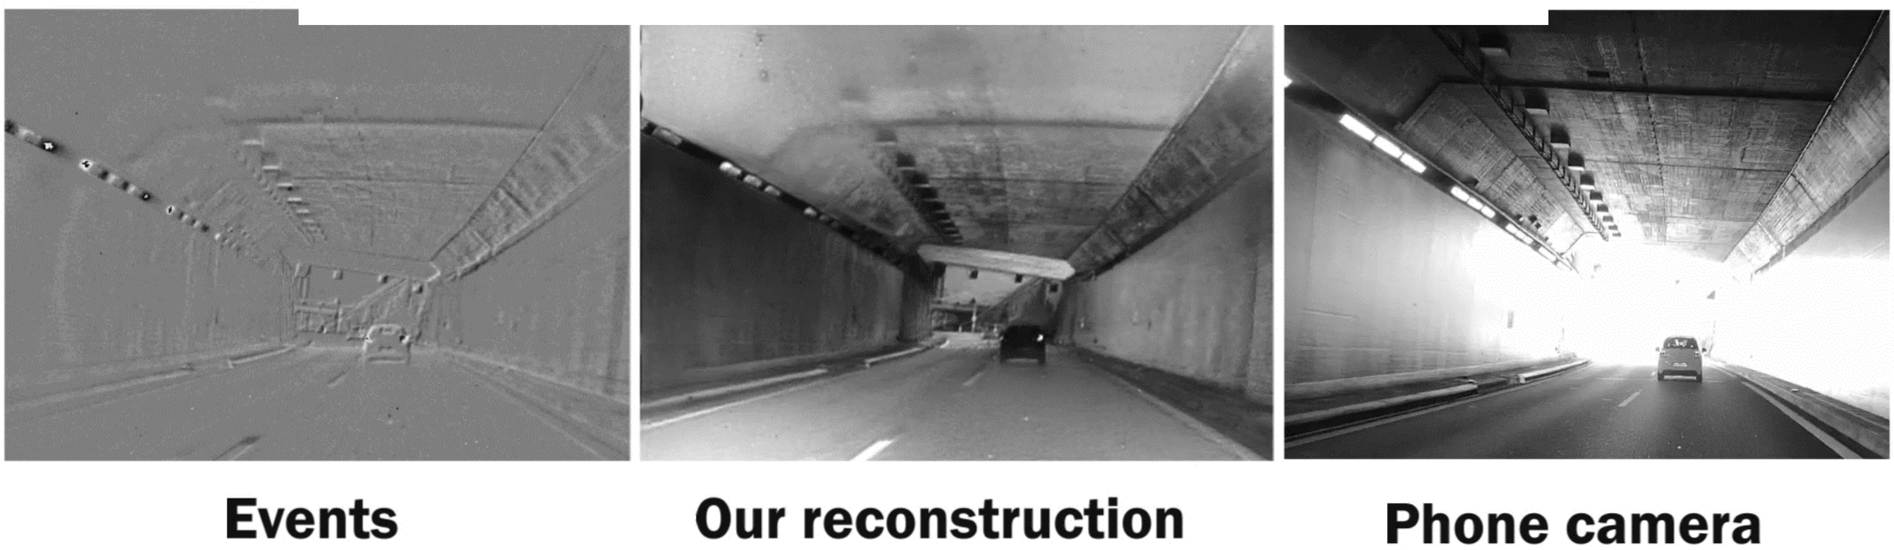

Source Code: https://github.com/uzh-rpg/rpg_e2vid

Rebecq et al., “High Speed and High Dynamic Range Video with an Event Camera”, PAMI’19. PDF Video Code

HDR Video: Night Drive

Recall: trained in simulation only!

Video courtesy of Prophesee

Our reconstruction from events
(we used a Prophesee sensor)

GoPro Hero 6

Source Code: https://github.com/uzh-rpg/rpg_e2vid

Rebecq et al., "High Speed and High Dynamic Range Video with an Event Camera", PAMI'19. PDF Video Code

Downstream Applications

Monocular Depth Estimation

https://youtu.be/eomALySSGVU

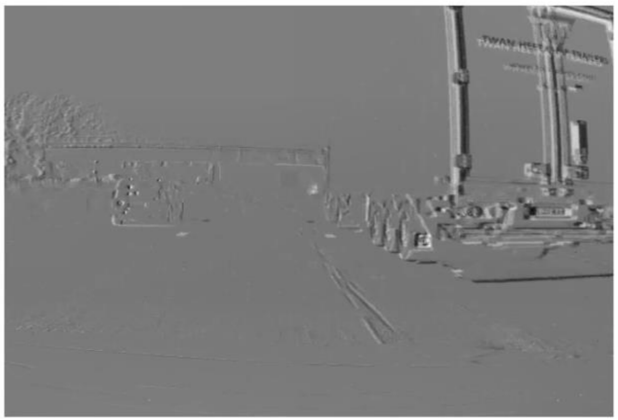

Events

Our reconstruction

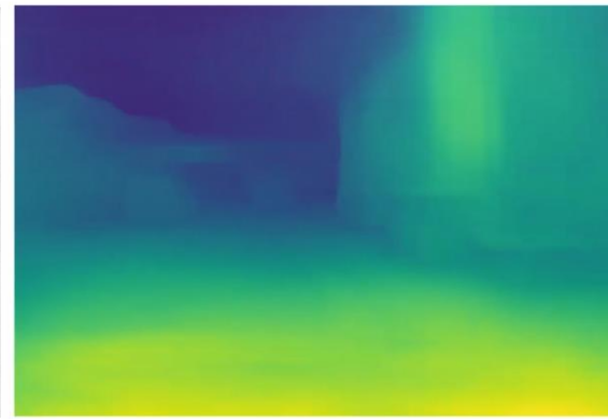

Monocular depth

Monocular depth estimation (Megadepth) applied on the reconstructed frames

Rebecq et al., “High Speed and High Dynamic Range Video with an Event Camera”, PAMI’19. PDF Video Code

Object detection

YouTube https://youtu.be/eomALySSGVU

Events

Our reconstruction + object detections (YOLOv3)

Rebecq et al., "High Speed and High Dynamic Range Video with an Event Camera", PAMI'19. [PDF](#) [Video](#) [Code](#)

Does it mean that in order to use event cameras we must first reconstruct an image?

NO!

These results were only to show that it should be possible to design algorithms that process events **end-to-end without passing through image reconstruction**!

Example: End-to-End Object Classification

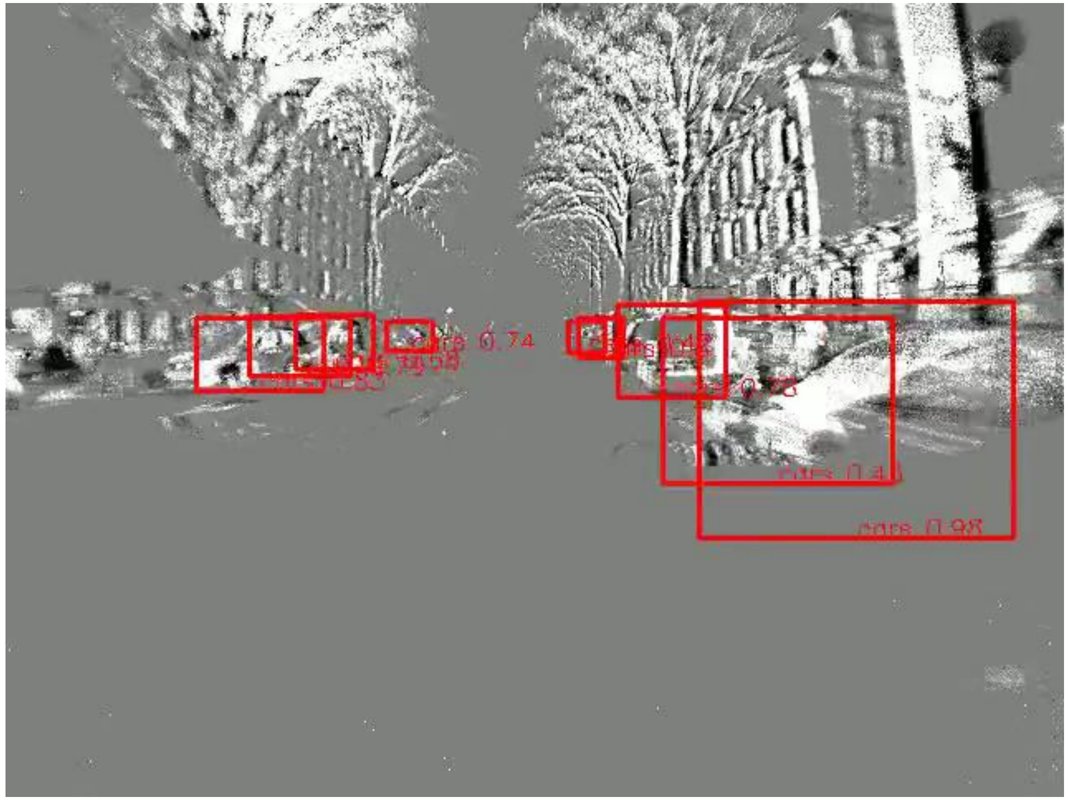

- Dataset: https://www.prophesee.ai/dataset-n-cars/
- Collected by PROPHESEE (largest event-camera company)
- Contains: Event, Images, car and pedestrian annotations

Sironi et al., "HATS: Histograms of Averaged Time Surfaces for Robust Event-based Object Classification". CVPR'18

Thanks!

- Code, Dataset, Simulator, tutorials, resources on event cameras:
  http://rpg.ifi.uzh.ch/research_dvs.html

- Survey paper on event cameras:
  http://rpg.ifi.uzh.ch/docs/EventVisionSurvey.pdf

- Code, datasets, videos, and publications:
  http://rpg.ifi.uzh.ch/

@davsca1 @davidescaramuzza

Reference

Event-based Vision: A Survey

Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andr
Stefan Leutenegger, Andrew Davison, Jörg Conradt, Kostas Daniilidis, Davide Scaramuzza

Abstract— Event cameras are bio-inspired sensors that work radically different from traditional cameras. Instead of capturing images at a fixed rate, they measure per-pixel brightness changes asynchronously. This results in a stream of events, which encode the time, location and sign of the brightness changes. Event cameras posses outstanding properties compared to traditional cameras: very high dynamic range ($140\,\mathrm{dB}$ vs. $60\,\mathrm{dB}$), high temporal resolution (in the order of $\mu\mathrm{s}$), low power consumption, and do not suffer from motion blur. Hence, event cameras have a large potential for robotics and computer vision in challenging scenarios for traditional cameras, such as high speed and high dynamic range. However, novel methods are required to process the unconventional output of these sensors in order to unlock their potential. This paper provides a comprehensive overview of the emerging field of event-based vision, with a focus on the applications and the algorithms developed to unlock the outstanding properties of event cameras. We present event cameras from their working principle, the actual sensors that are available and the tasks that they have been used for, from low-level vision (feature detection and tracking, optic flow, etc.) to high-level vision (reconstruction, segmentation, recognition). We also discuss the techniques developed to process events, including learning-based techniques, as well as specialized processors for these novel sensors, such as spiking neural networks. Additionally, we highlight the challenges that remain to be tackled and the opportunities that lie ahead in the search for a more efficient, bio-inspired way for machines to perceive and interact with the world.

Index Terms—Event Cameras, Bio-Inspired Vision, Asynchronous Sensor, Low Latency, High Dynamic Range, Low Power.

1 Introduction and Applications

"THE *brain is imagination, and that was exciting to me; I wanted to build a chip that could imagine something*[1]." that is how Misha Mahowald, a graduate student at Caltech in 1986, started to work with Prof. Carver Mead on the stereo problem from a joint biological and engineering perspective. A couple of years later, in 1991, the image of a cat in

as well as new computer vision and robotic tasks. Sight is, by far, the dominant sense in humans to perceive the world, and, together with the brain, learn new things. In recent years, this technology has attracted a lot of attention from both academia and industry. This is due to the availability of prototype event cameras and the advantages that these devices offer to tackle problems that are currently unfeasible

http://rpg.ifi.uzh.ch/docs/EventVisionSurvey.pdf