Event-based Neuromorphic Perception and Computation: The Future of Sensing and AI

R.B. Benosman,
McGowan Institute,
BST-3, Rm 2046, 3501 Fifth Avenue
Pittsburgh, PA 15213
benosman@pitt.edu

Event-based ~~Neuromorphic~~ Perception and Computation: The Future of Sensing and AI

R.B. Benosman,
McGowan Institute,
BST-3, Rm 2046, 3501 Fifth Avenue
Pittsburgh, PA 15213
benosman@pitt.edu

Historic Timeline

Franck Rosenblatt: Perceptron

1958

Russell's Infant Son: 5cm by 5cm (176x176 array) Portland Art Museum.

1959

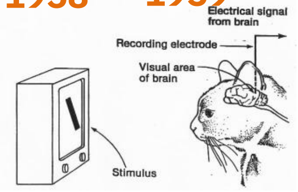

David Hubel and Torsten Wiesel— in 1959. Their publication, entitled "*Receptive fields of single neurons in the cat's striate cortex*"

1963

Lawrence Roberts' "Machine perception of three-dimensional solids". Process 2D photographs to build up 3D representations from lines.

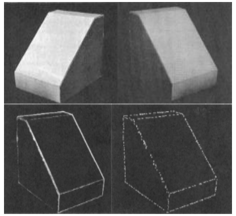

1966

MIT summer Vision Project, Seymour Papert, automatically, background/foreground segmentation, extract non-overlapping objects

1969

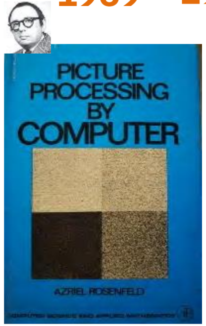

Azriel Rosenfeld
Early applications of image analysis

1982

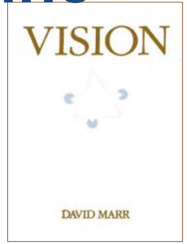

David Marr,
"Vision a computational investigation into the human representtion and processing of visual information" vision is hierarchical

1990's

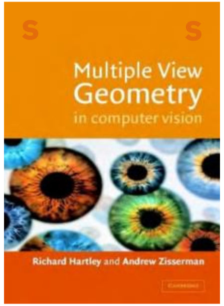

Vision is ruled by Geometry, (projective) 3D reconstruction, fundamental matrix, RANSAC, bundle,

2000's

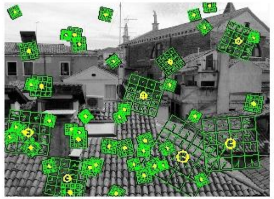

Feature based Pattern recognition.
Keypoints
3D reconstruction gets "solved", generic object recognition

Historic Timeline

k
nblatt:
ptron

Russell's Infant Son: 5cm by 5cm (176x176 array) Portland Art Museum.

MIT summer Vision Project, Seymour Papert, automatically, background/foreground segmentation, extract non-overlapping objects

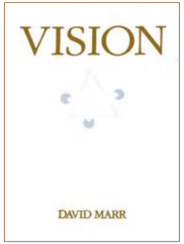

David Marr,
"Vision a computational investigation into the human representtion and processing of visual information" vision is hierarchical

Feature based Pattern recognition.
Keypoints
3D reconstruction gets "solved", generic object recognition

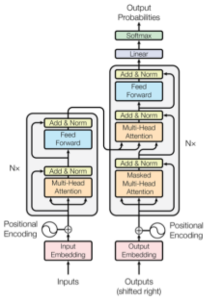

?

958 | 1959 | 1963 | 1966 | 1969 | 1982 | 1990's | 2000's | 2000's

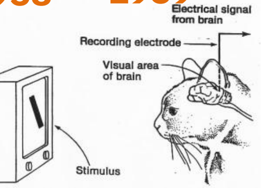

David Hubel and Torsten Wiesel— in 1959. Their publication, entitled "*Receptive fields of single neurons in the cat's striate cortex*"

Lawrence Roberts' "Machine perception of three-dimensional solids". Process 2D photographs to build up 3D representations from lines.

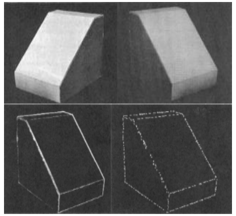

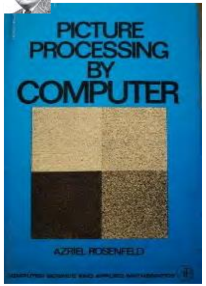

Azriel Rosenfeld
Early applications of image analysis

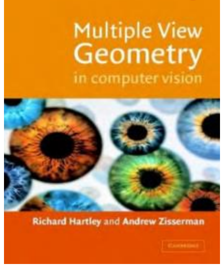

Vision is ruled by Geometry, (projective) 3D reconstruction, fundamental matrix, RANSAC, bundle,

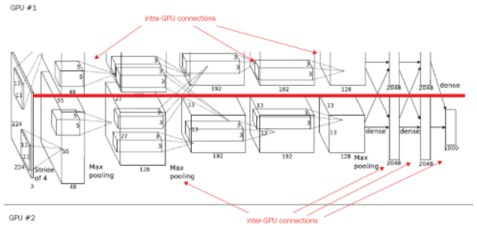

What Went Wrong?

- Computer vision has been reinvented at least three times.

- Too close to the market: **applications based research**

Tendency to **resist novelty** choosing applications over potentially more promising methods that could not yet deliver

- **Not idea driven**

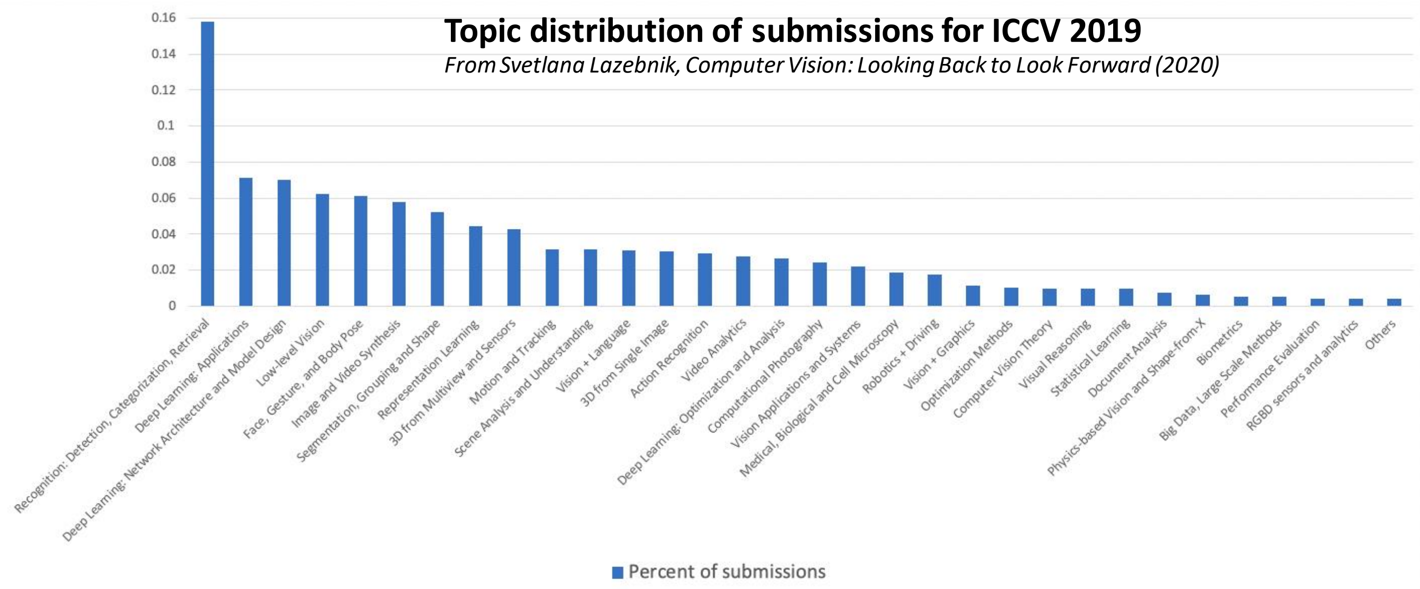

What Went Wrong?

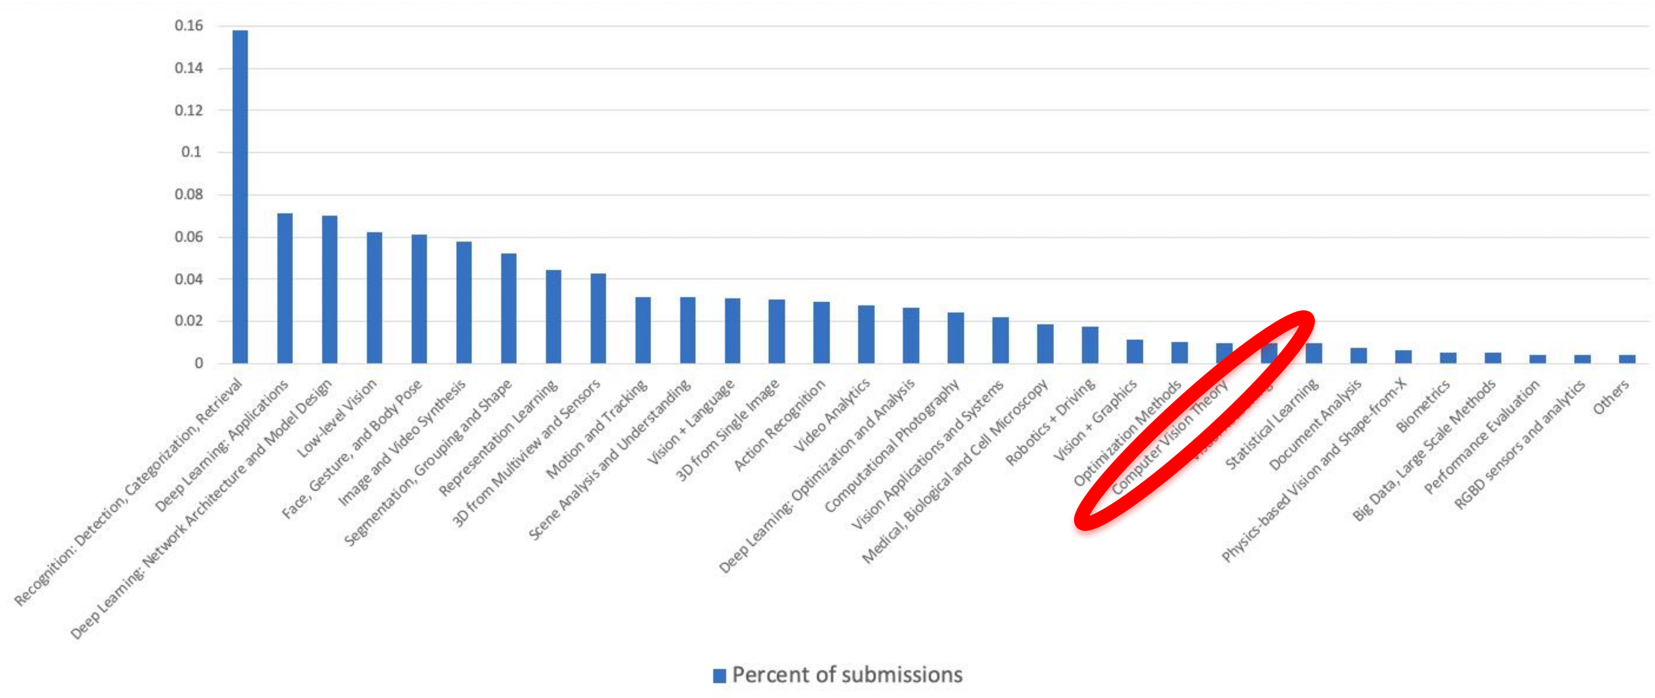

Why Are We Using Images?

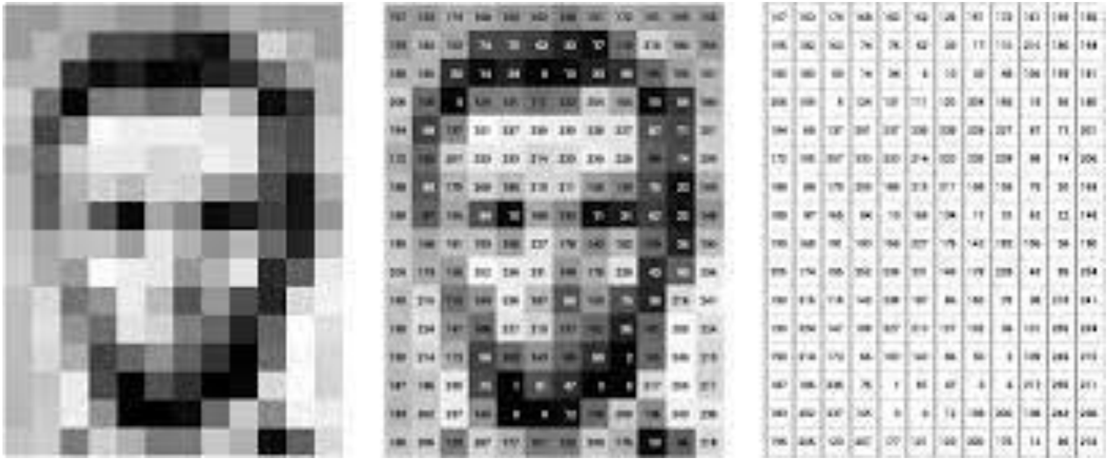

- **Images** are the optimal **structure of data**
- **Grey Levels** as **source of information**

Computer Vision: a Heritage from Art!

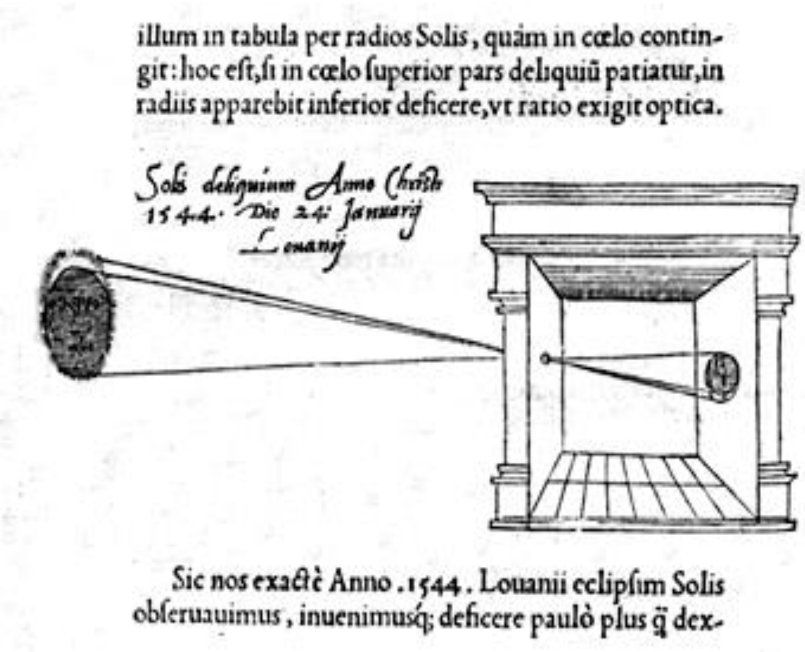
illum in tabula per radios Solis, quàm in cœlo contin-
git: hoc eſt, ſi in cœlo ſuperior pars deliquiũ patiatur, in
radiis apparebit inferior deficere, vt ratio exigit optica.

Solis deliquium Anno Christi
1544. Die 24: Januarij
Louanij

Sic nos exactè Anno .1544. Louanii eclipſim Solis
obſeruauimus, inuenimusq; deficere paulò plus q̃ dex-

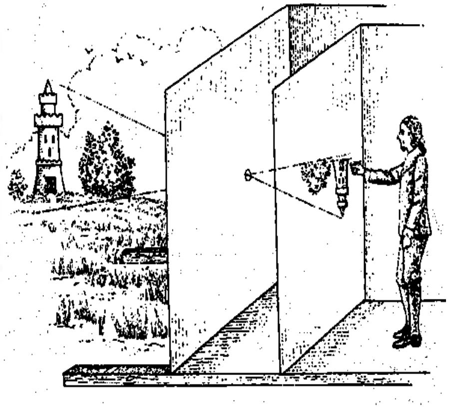

- Invention of the camera obscura in 1544 (L. Da Vinci?)
- The mother of all cameras

Origins of Imaging

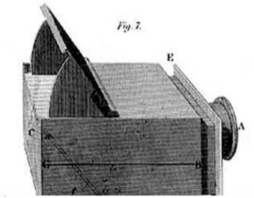

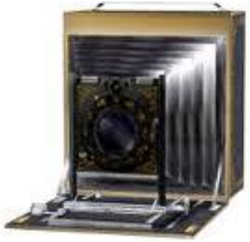

- **Increasing profits**: painting faster
- Evolution from portable models for travellers to current digital cameras
- Evolving **from canvas**, to paper, to glass, to celluloid, **to pixels**

Origins of Video: Motion Picture

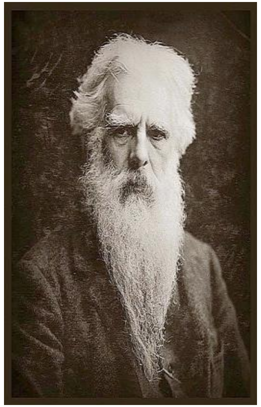

Eadweard Muybridge
(1830-1904)

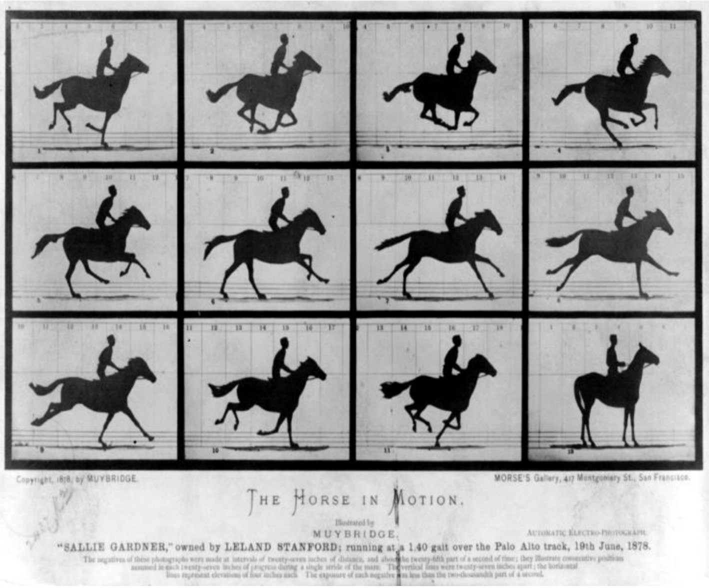

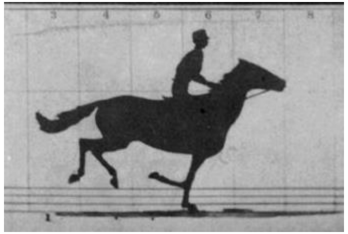

- Early work in **motion-picture** projection
- Pioneering work on **animal locomotion** in 1877 and 1878
- Used **multiple cameras** to capture motion in **stop-motion** photographs

Computer vision, the Impossible Trade Off!

power vs frame rate

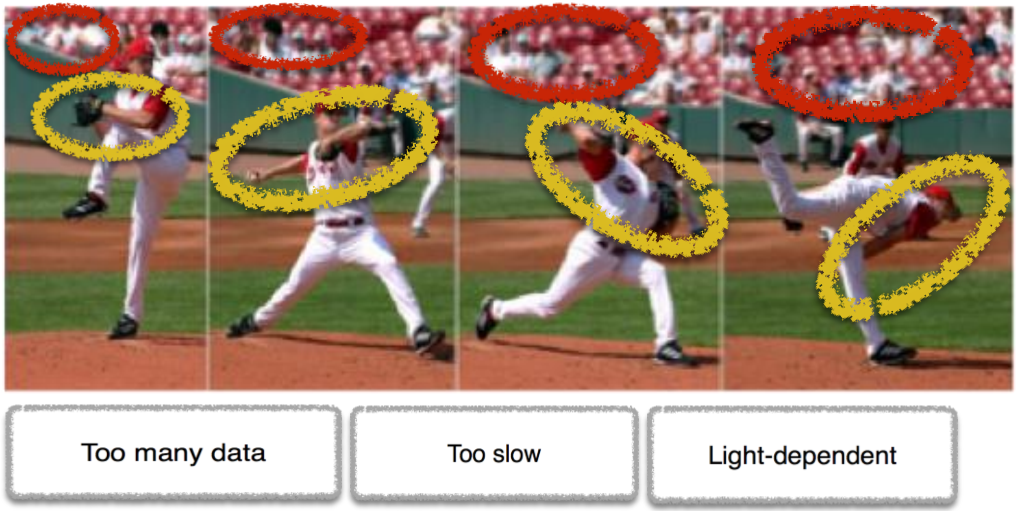

over-sampling

- redundant useless data
- power and resource hungry: need to acquire/transmit/store/process

under-sampling

- motion blur
- displacement between frames

High Power & High Latency

Event Acquisition

Scopes:

- Reduce Data Load and only Detect “meaningful” events, at the time they happen!
- Avoid burning energy to acquire, transmit and store information that ends up being trashed

Solutions:

- No generic solution,
- There are almost an infinite number of solutions to extract events
- Need to be adapted to the dynamics and nature of the data

Event acquisition

Popular solution: Sample on the amplitude axis of signals

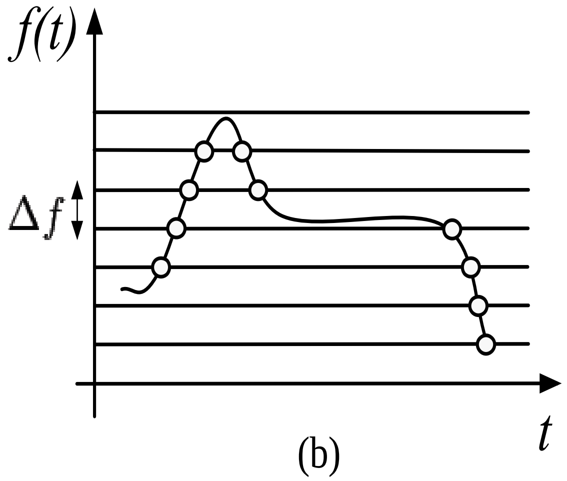

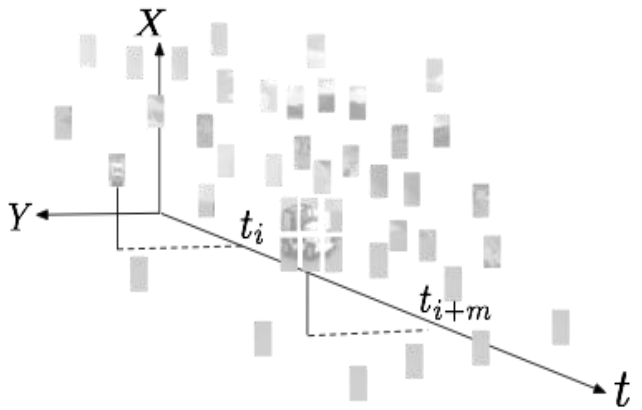

- New Information is detected when it happens
- When nothing happens, nothing is sent or processed
- Sparse information coding

Time is the most valuable information

Event acquisition

<u>Popular solution</u>: Sample on the amplitude axis of signals

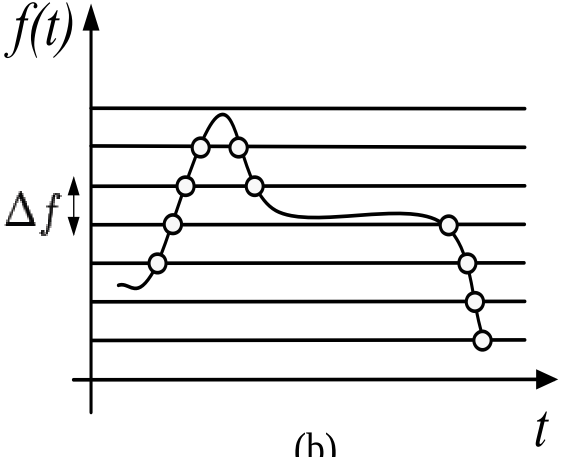

(b)

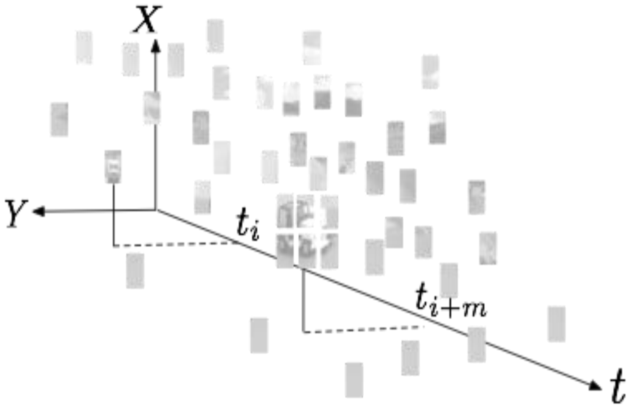

<u>Time is the most valuable information</u>

A 128×128 120dB 15us Latency Asynchronous Temporal Contrast Vision Sensor

Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck, *Member, IEEE*

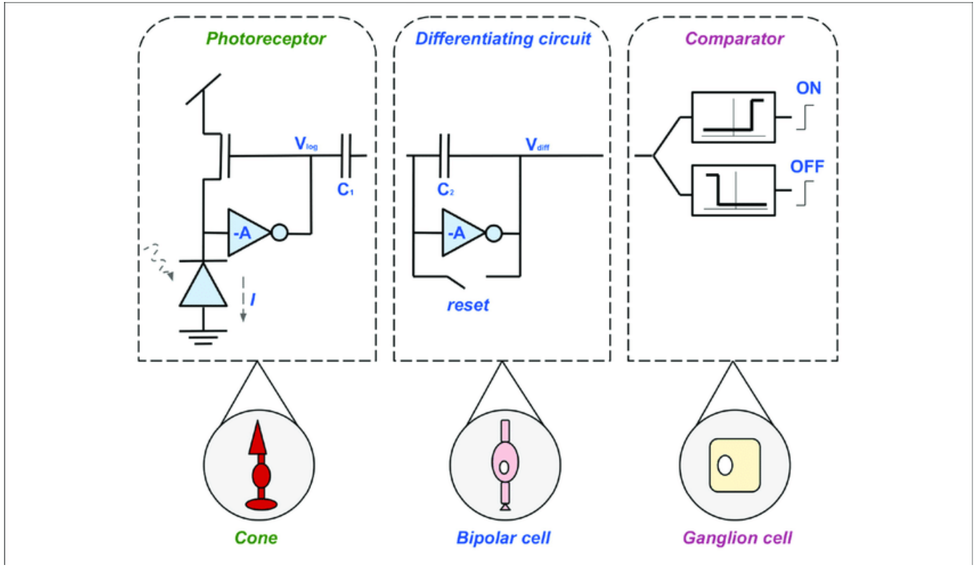

$$|\Delta \log I| > T$$

A QVGA 143 dB Dynamic Range Frame-Free PWM Image Sensor With Lossless Pixel-Level Video Compression and Time-Domain CDS

Christoph Posch, *Member, IEEE*, Daniel Matolin, and Rainer Wohlgenannt

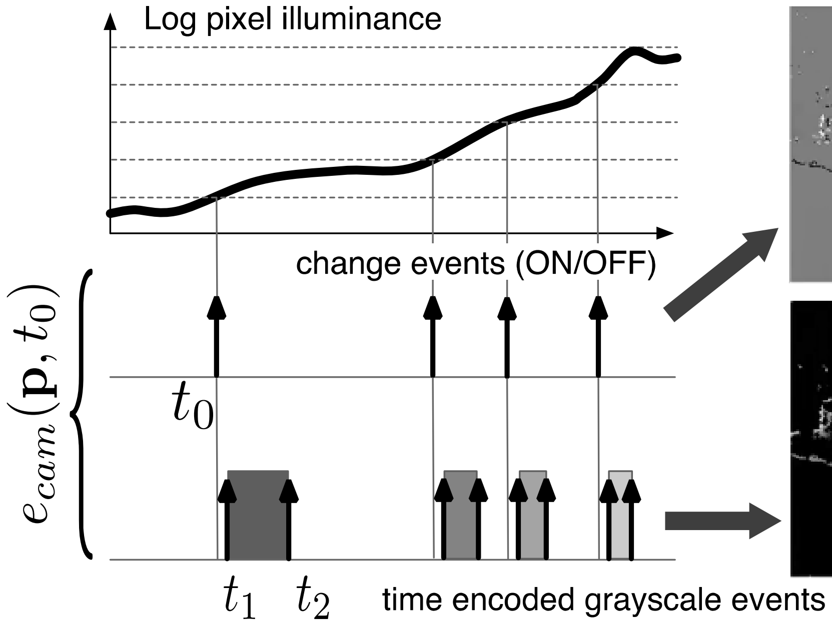

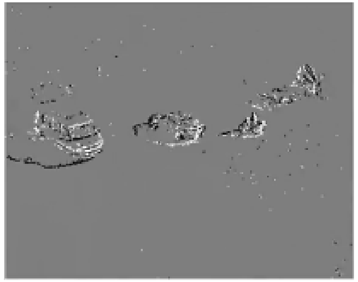

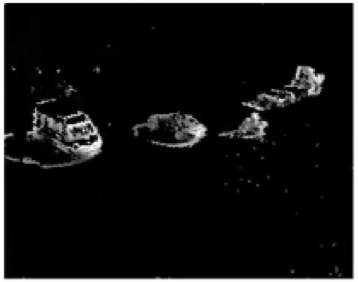

Temporal events and absolute light measurement

Frames vs Events

conventional
frame-based camera

event-based camera

Why Event Based Sensors?

Chaotic pendulum tracking

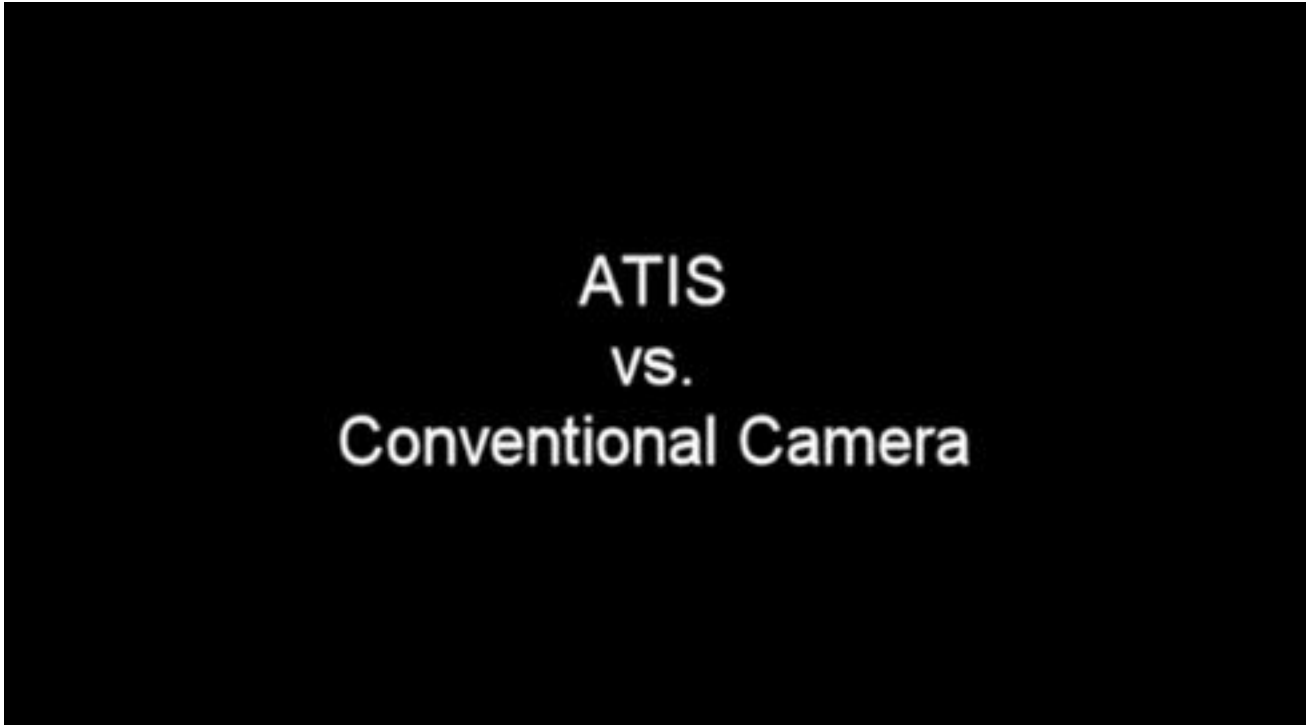

Event Time-based Sensor: Grey Levels Events

Why Event Based Sensors?

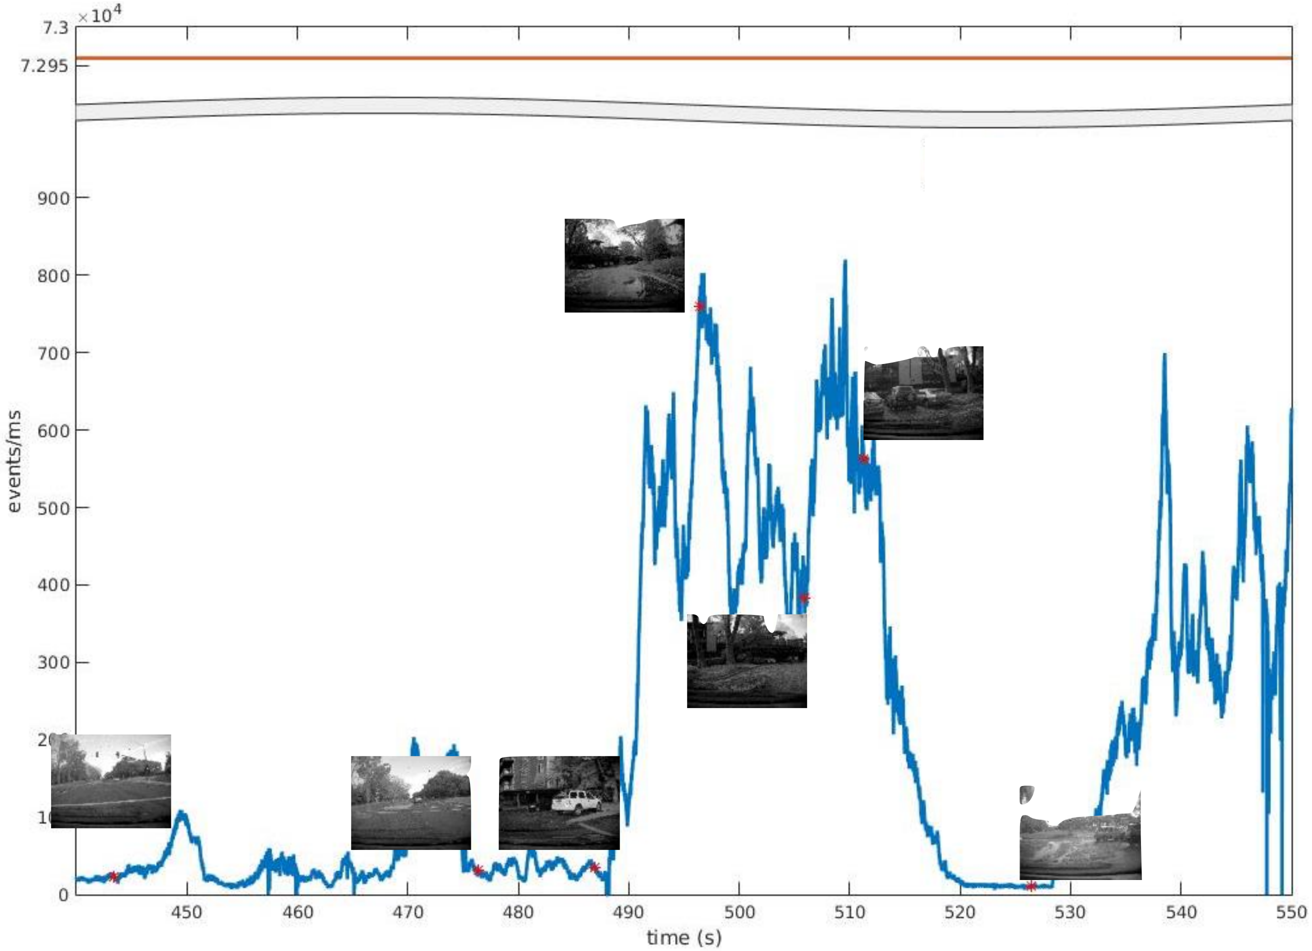

Event Cameras

iniLabs

PROPHESEE
METAVISION FOR MACHINES

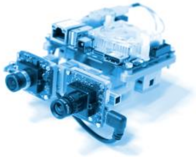

insightness

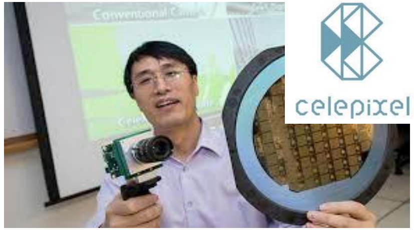

Event Cameras

intel

SONY

SAMSUNG

iniLabs

PROPHESEE

METAVISION FOR MACHINES

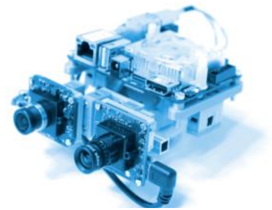

SONY

OmniVision.

elepixel

Event cameras have become <u>a commodity</u>

How to Process Events?

How to Process Events?

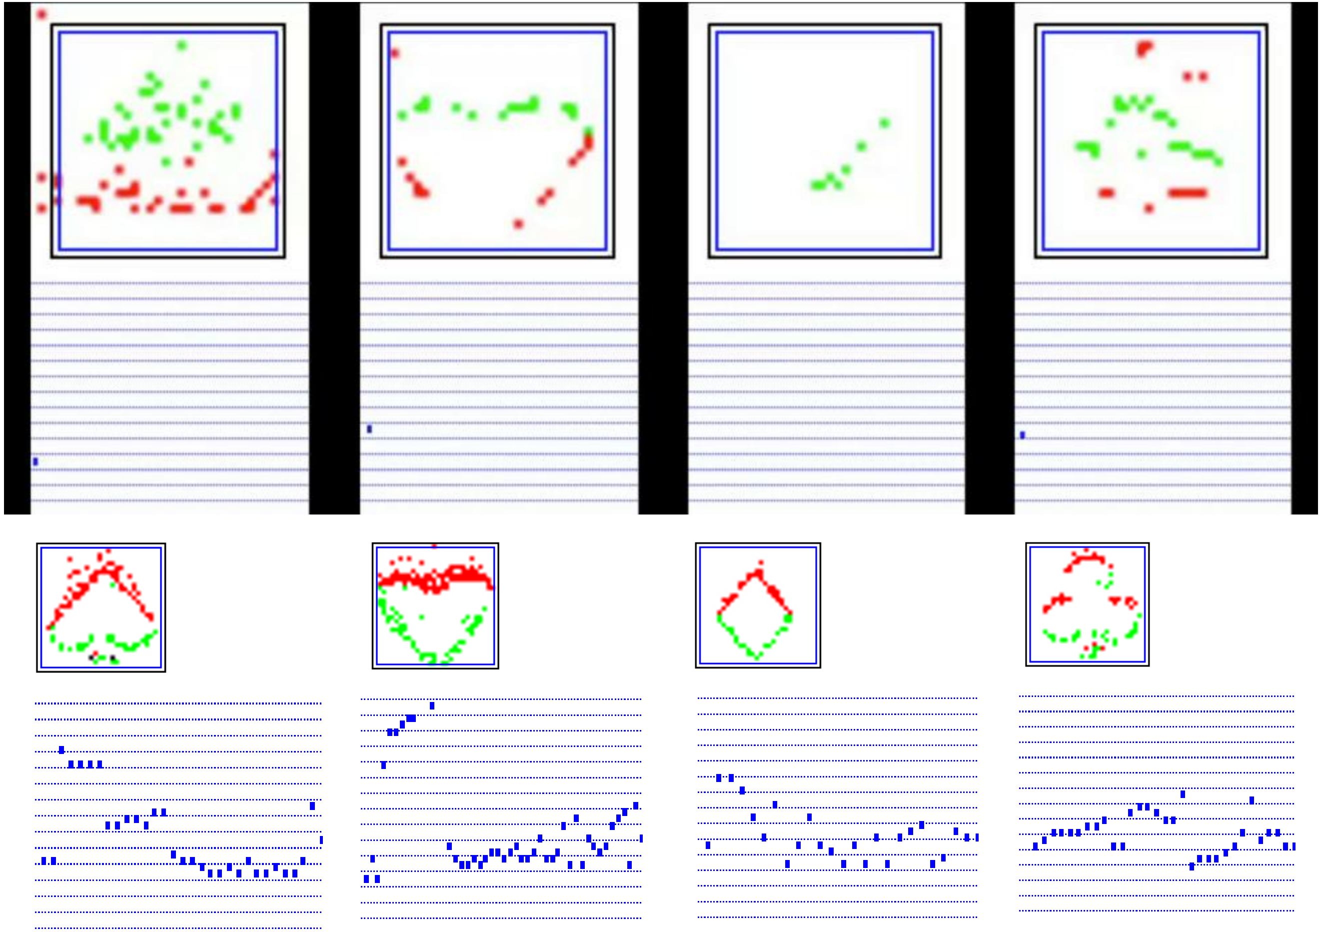

Event Computation

Incremental

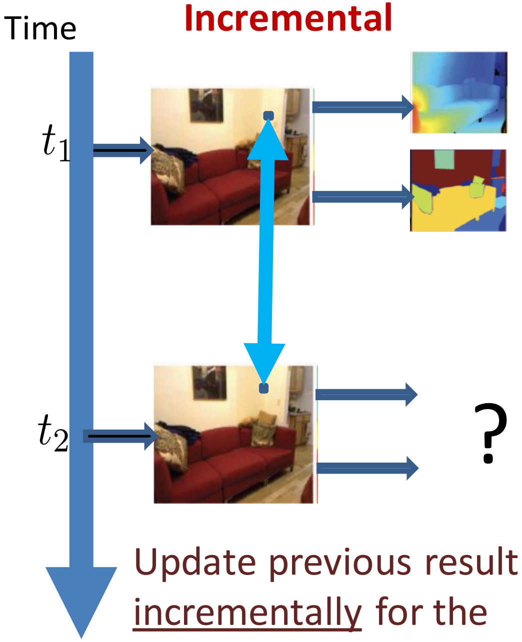

Update previous result <u>incrementally</u> for the single pixel change

VS

Batch or Frame

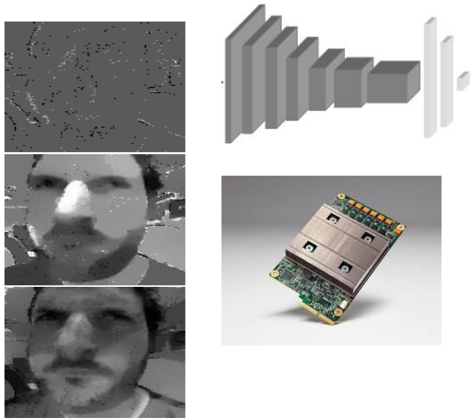

Compute the new 1 pixel-change frame and compute result using <u>the whole frame</u>

Applications: Event Stereovision

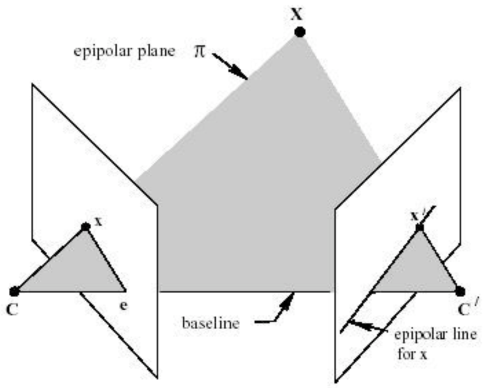

$$uu'f_{11} + uv'f_{21} + uf_{31} + vu'f_{12} + vv'f_{22} + vf_{32} + u'f_{31} + v'f_{23} + f_{33} = 0,$$

$$A\mathbf{f} = 0$$

- Matching pixels is hard
- Changing lighting conditions, occlusions, motion blur….

Event Stereovision

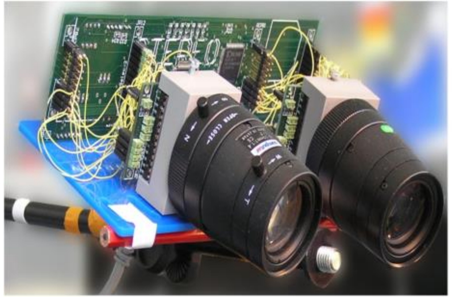

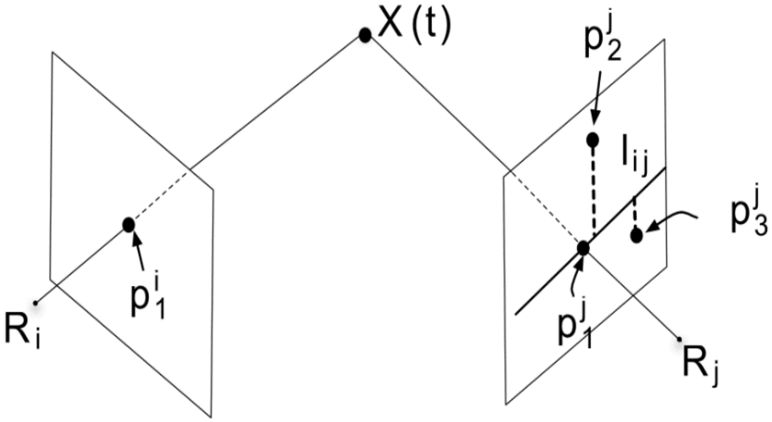

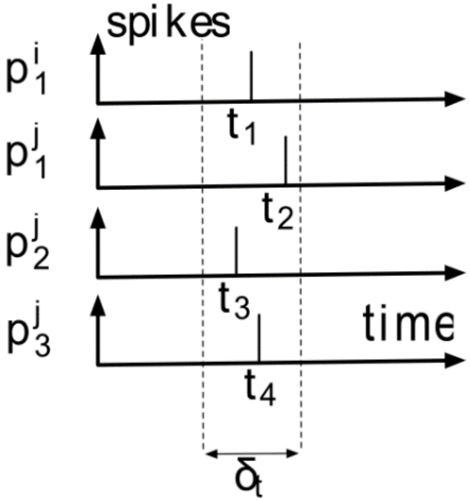

- Matching binocular events only using the time of events
- Two events arriving at the same time and fulfilling geometric constraints are matched

P.Rogister, R.B. Benosman, S-H. Ieng, P. Lichtsteiner, T. Delbruck, Asynchronous Event-based Binocu- lar Stereo Matching. 2011 December 23, IEEE Transactions on Neural Networks 23(2) : 347-353, DOI : 10.1109/TNNLS.2011.2180025

Event Stereovision

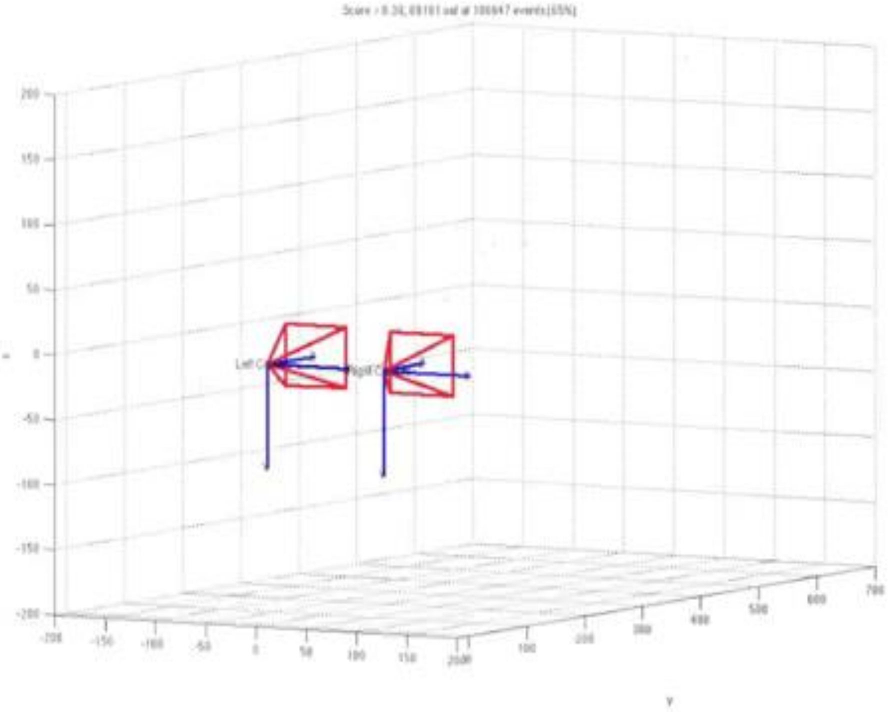

S.H. Ieng, J. Carneiro, M. Osswald, R. B.Benosman, *Neuromorphic Event-Based Generalized Time-based Stereovision*, 2018 July 18, Frontiers in Neuroscience, 12(442), DOI :10.3389/fnins.2018.00442

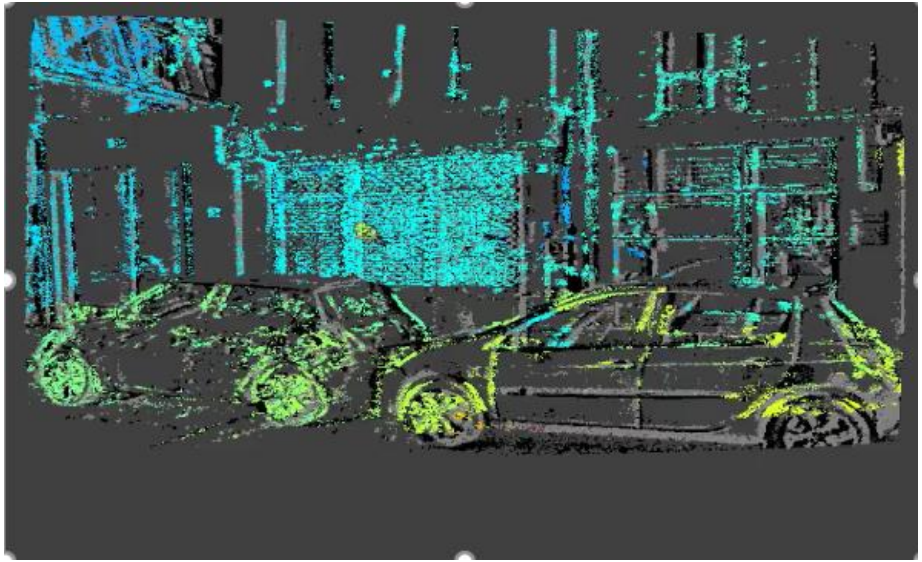

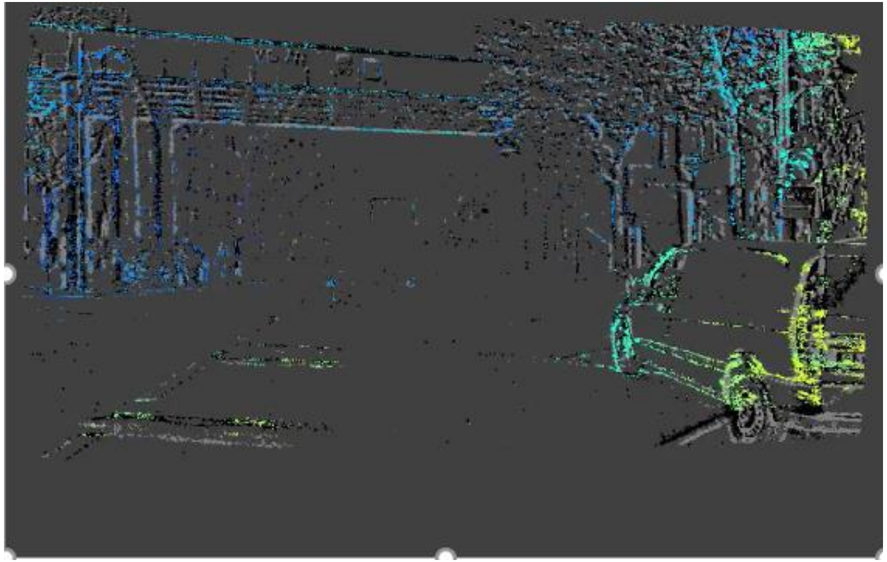

Courtesy Chronocam 2016

Visual Odometry

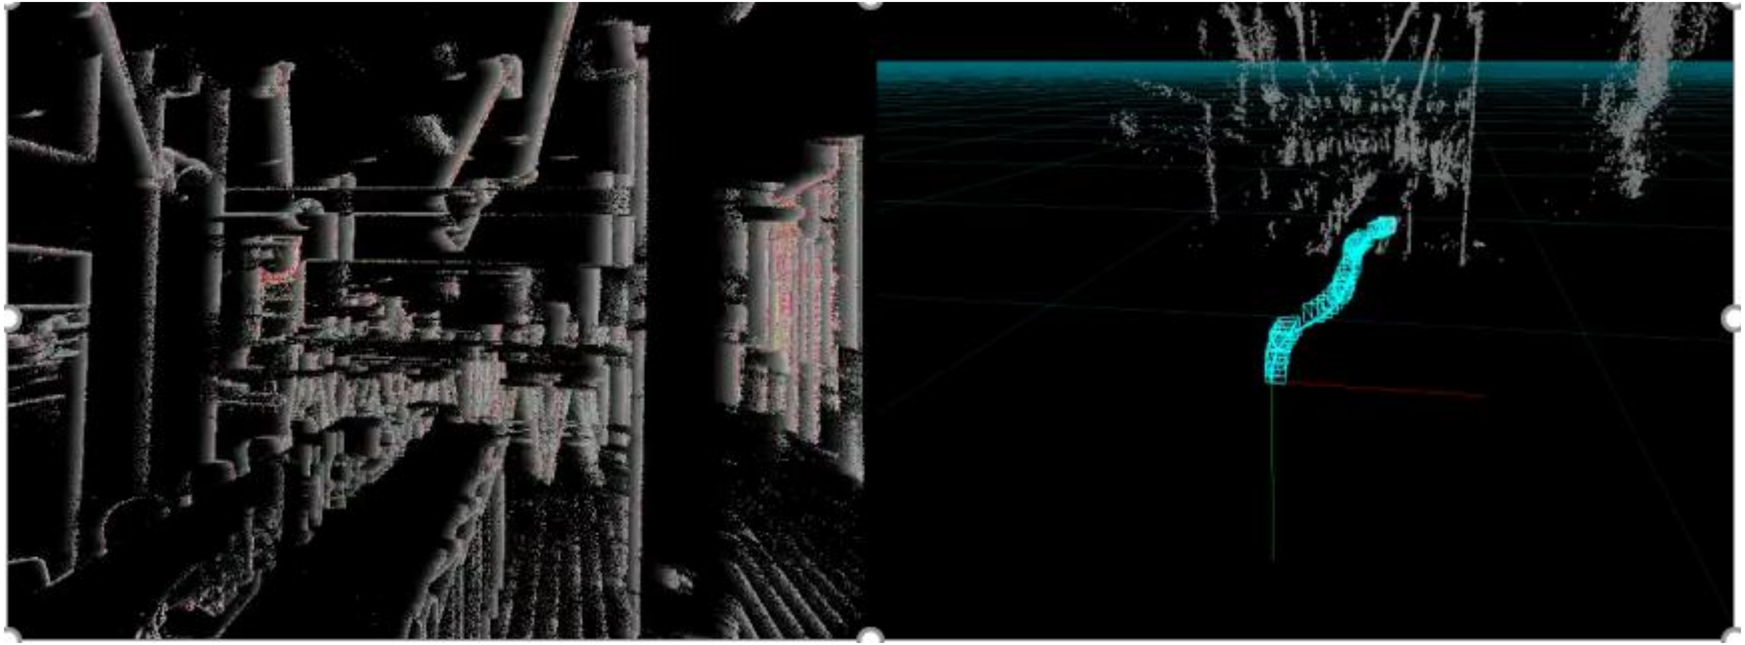

Courtesy Chronocam 2016

IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS, VOL. 25, NO. 2, FEBRUARY 2014

Event-Based Visual Flow

Ryad Benosman, Charles Clercq, Xavier Lagorce, Sio-Hoi Ieng, and Chiara Bartolozzi

$$I(x, y, t) = I(x + \Delta x, y + \Delta y, t + \Delta t)$$

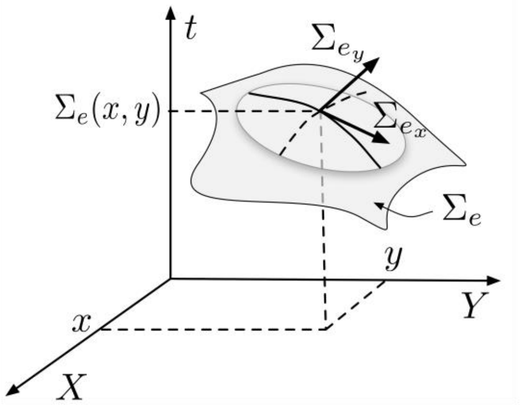

For an incoming event :

$e(\mathrm{p}, t) = (\mathrm{p}, t)^T$

Form the surface (event times):

$\Sigma_e : \mathbb{R}^2 \rightarrow \mathbb{R}^3$

$\mathrm{p} \mapsto t = \Sigma_e.$

We then have:

$$\frac{\partial \Sigma_e}{\partial x}(x, y_0) = \frac{\mathrm{d}\Sigma_e|_{y=y_0}}{\mathrm{d}x}(x) = \frac{1}{v_x(x, y_0)},$$

$$\frac{\partial \Sigma_e}{\partial y}(x_0, y) = \frac{\mathrm{d}\Sigma_e|_{x=x_0}}{\mathrm{d}y}(y) = \frac{1}{v_y(x_0, y)},$$

$$\nabla \Sigma_e = \left(\frac{1}{v_x}, \frac{1}{v_y}\right)^T,$$

Event Flow

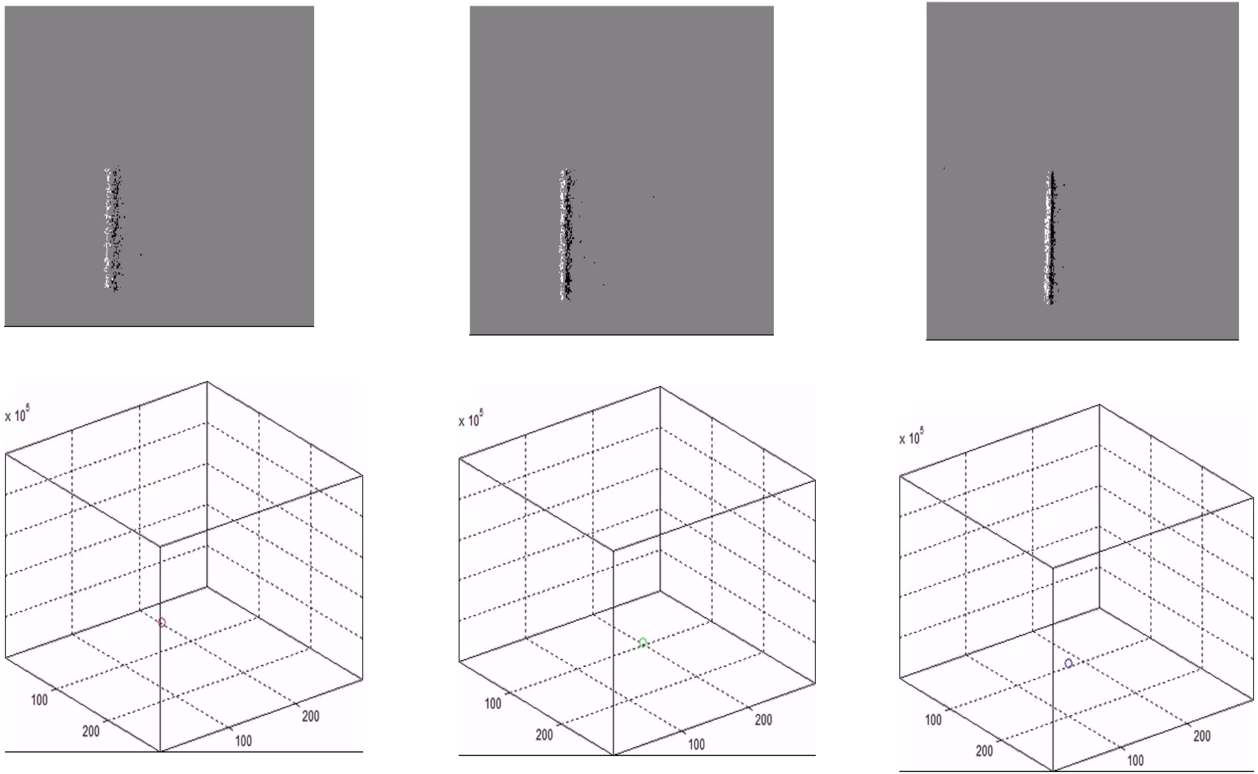

Event Flow

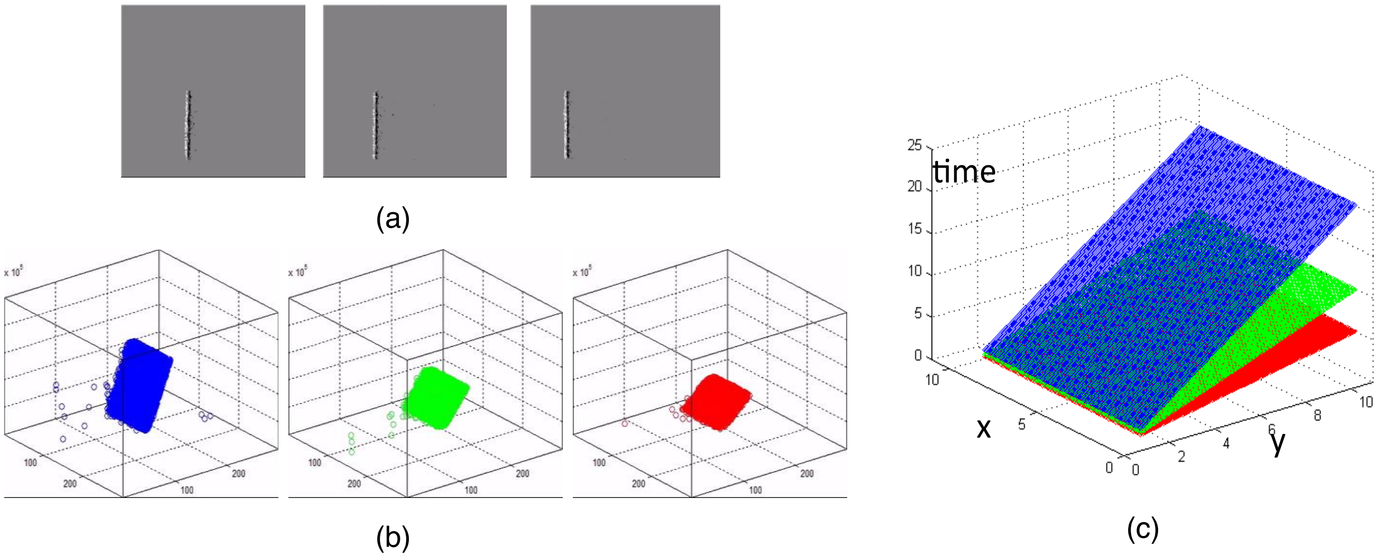

- High temporal resolution generates smooth space-time surfaces
- The slope of the local surface contains the orientation and amplitude of the optical flow

Event-Based Visual Flow

Ryad Benosman, Charles Clercq, Xavier Lagorce, Sio-Hoi Ieng, and Chiara Bartolozzi

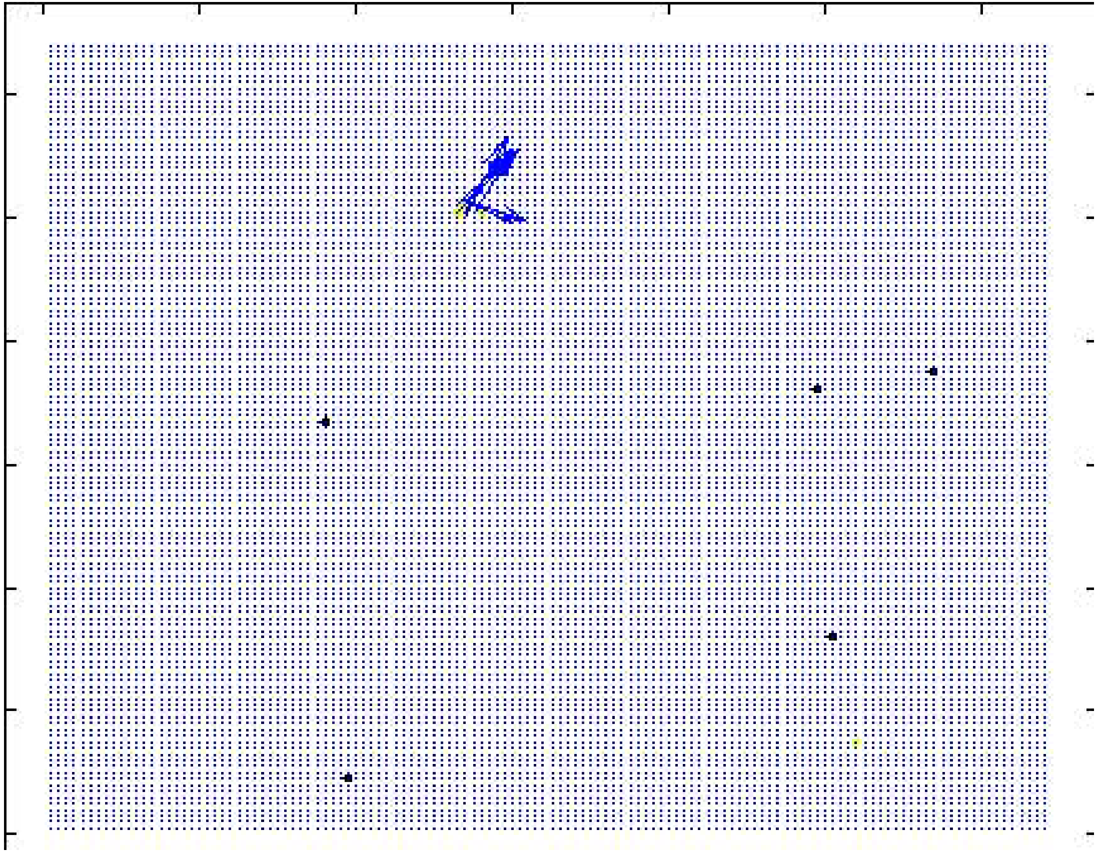

Tracking Real-Time Outdoor Scenes

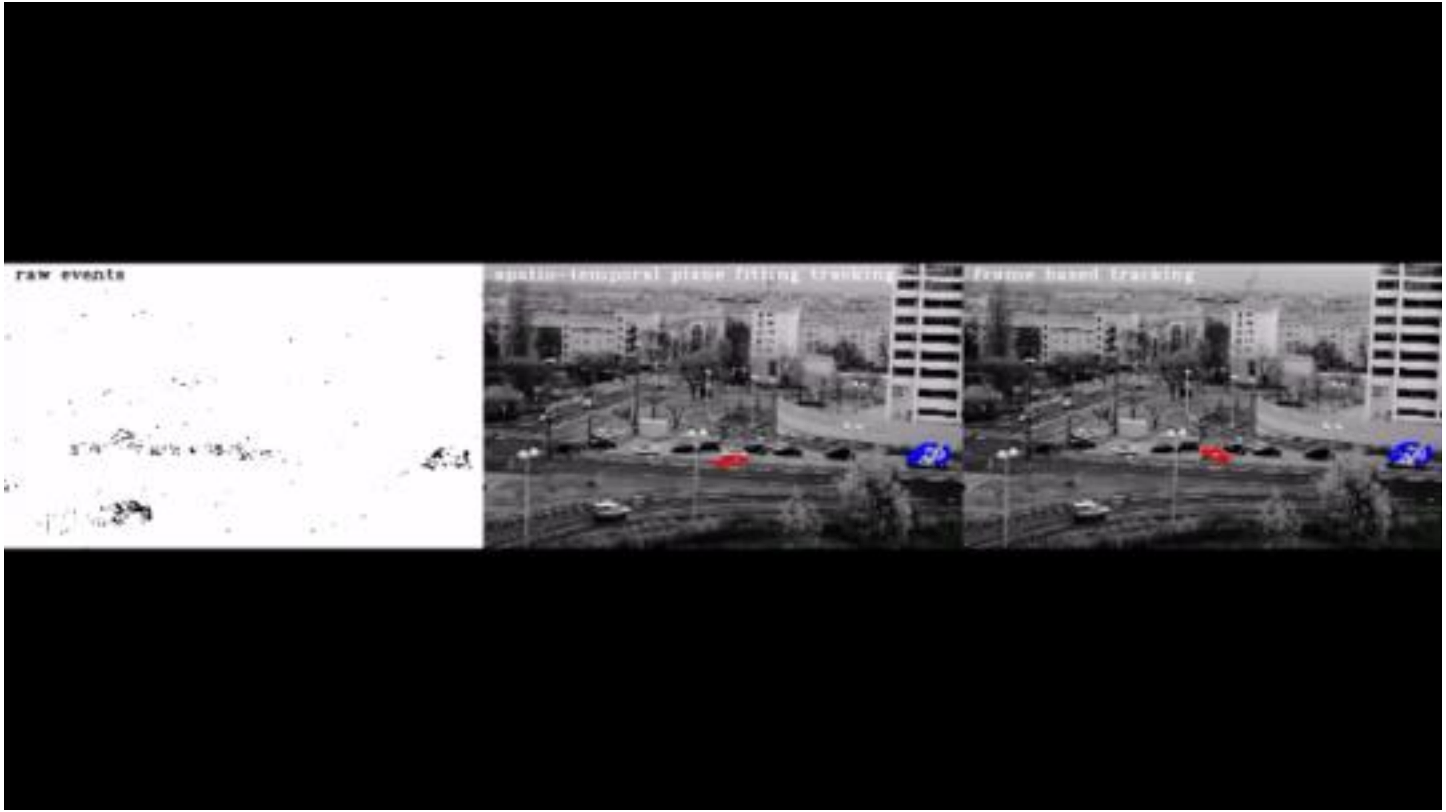

Z. Ni, S.H. Ieng, C. Posch, S. Regnier, R.B. Benosman, Visual Tracking using Neuromorphic Asynchronous Event-based Cameras, 24 February 2015 Neural Computation 27(4) :925-53, DOI : 10.1162/NECO-a-00720

Tracking Real-Time Outdoor Scenes

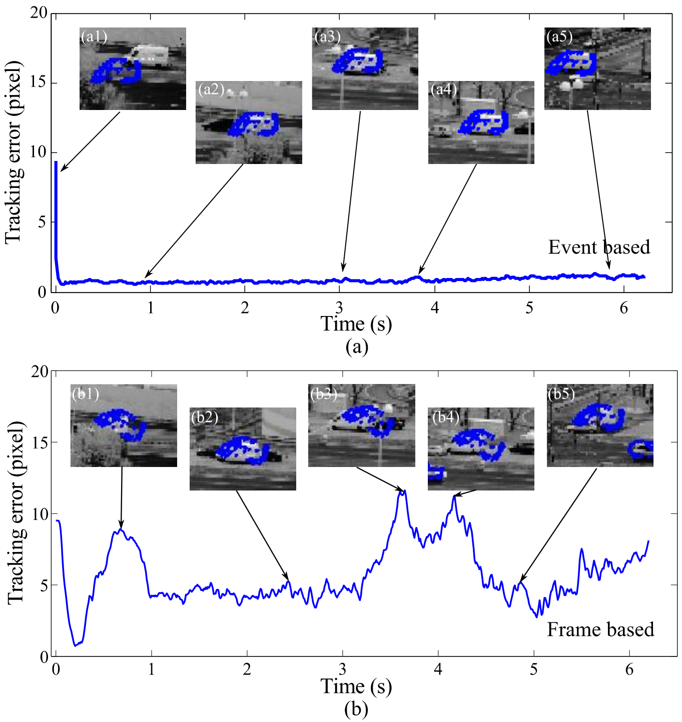

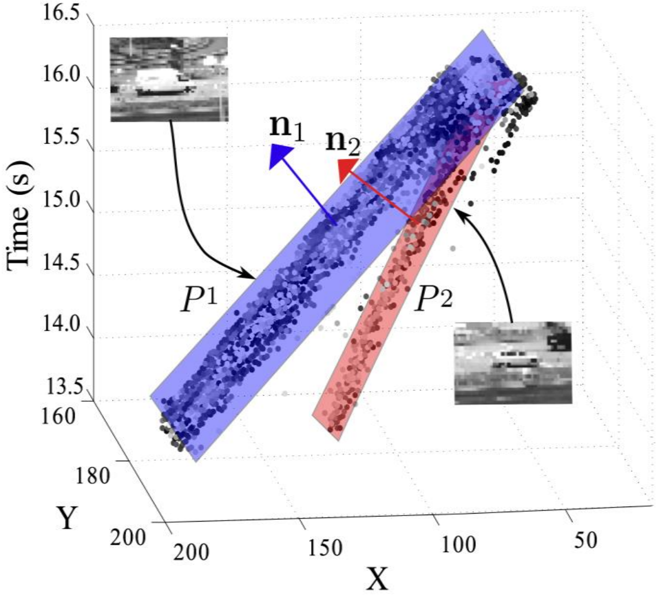

Z. Ni, S.H. Ieng, C. Posch, S. Regnier, R.B. Benosman, Visual Tracking using Neuromorphic Asynchronous Event-based Cameras, 24 February 2015 Neural Computation 27(4) :925-53, DOI : 10.1162/NECO-a-00720

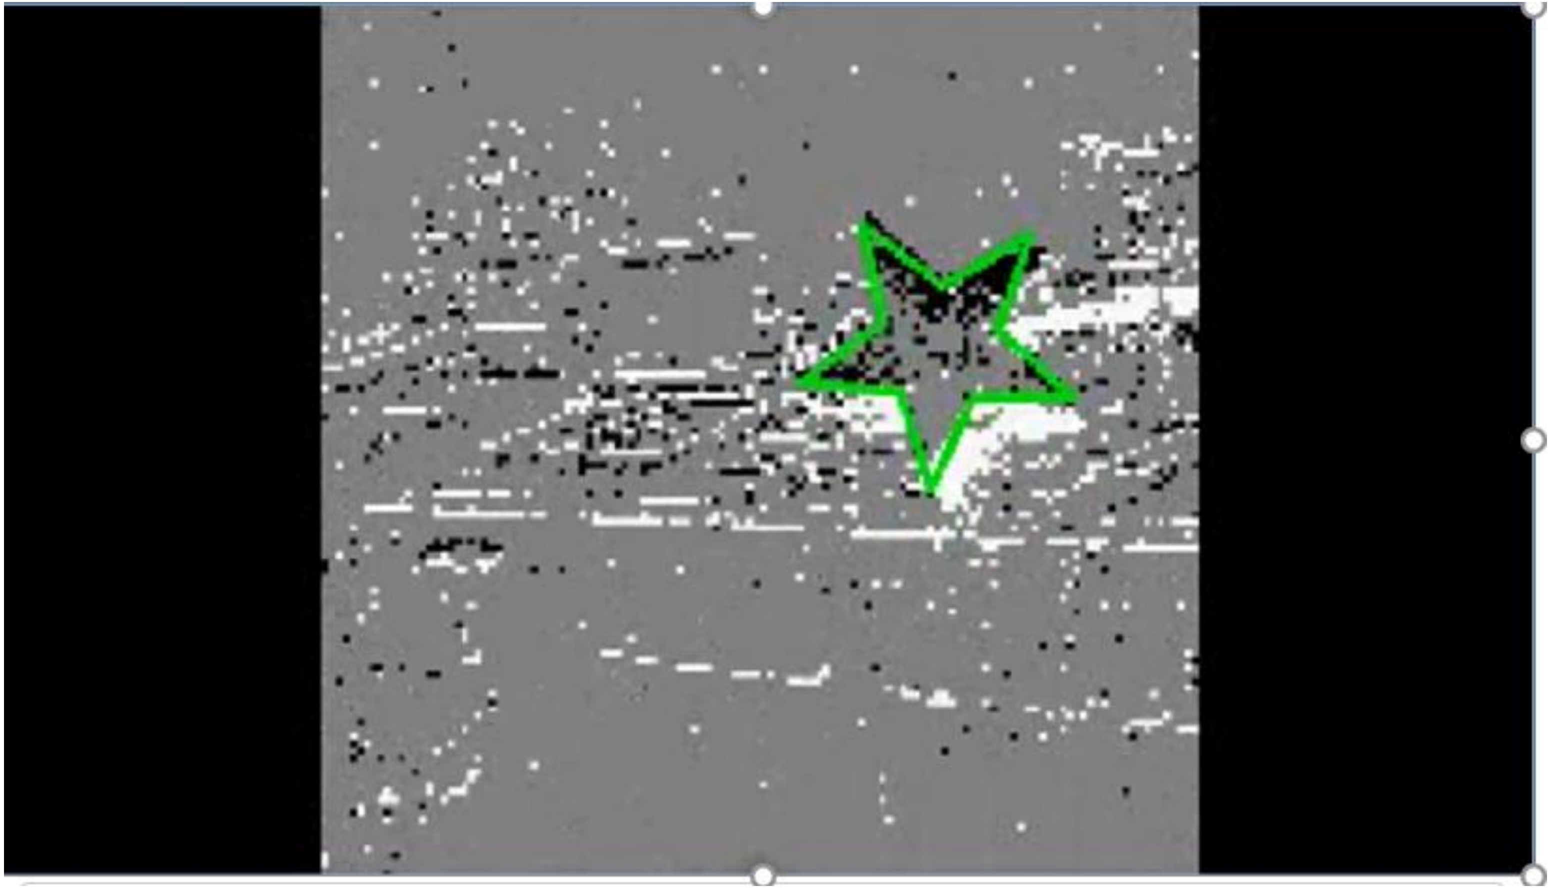

Z. Ni, S.H. Ieng, C. Posch, S. Regnier, R.B. Benosman, Visual Tracking using Neuromorphic Asynchronous Event-based Cameras, 24 February 2015 Neural Computation 27(4) :925-53, DOI : 10.1162/NECO-a-00720

Event-Based 3D Tracking and Pose Estimation

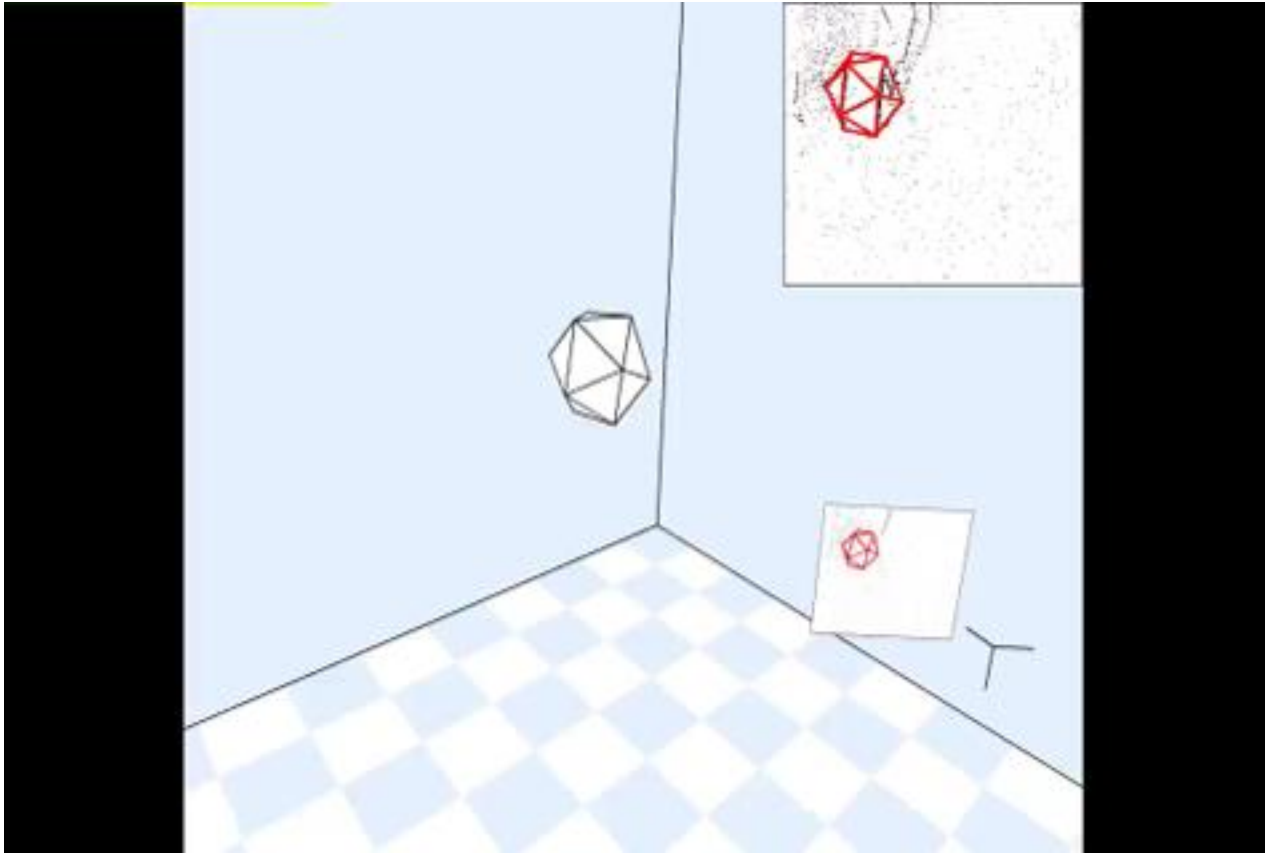

D. Reverter-Valeiras, G. Orchard, S.H. Ieng, R.B. Benosman,Neuromorphic Event-Based 3D Pose Estimation, 2016 January 22, Frontiers in Neuromorphic 9(522), DOI : 10.3389/fnins.2015.00522

Low Power and Latency Streaming

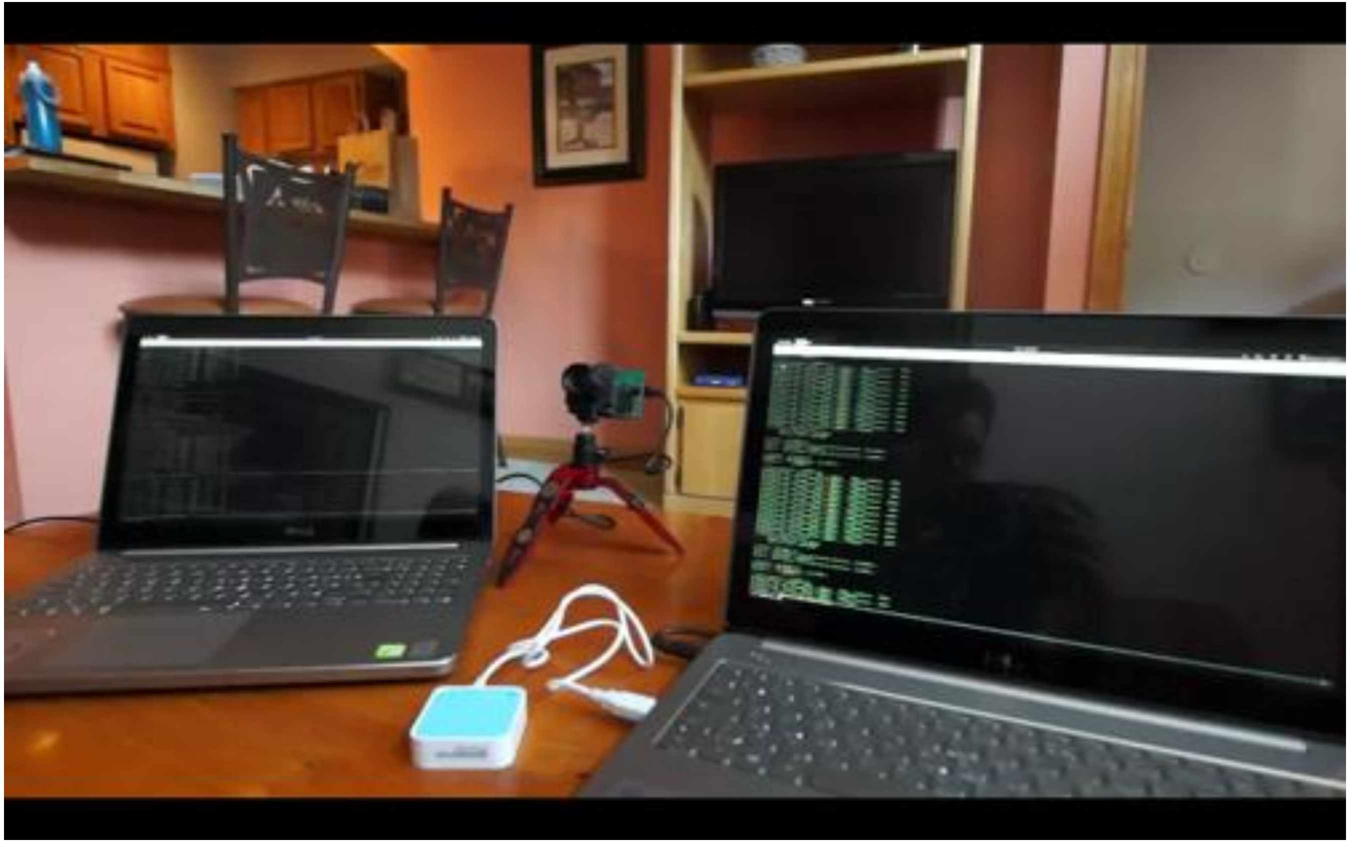

Asynchronous Event-Based Fourrier Analysis

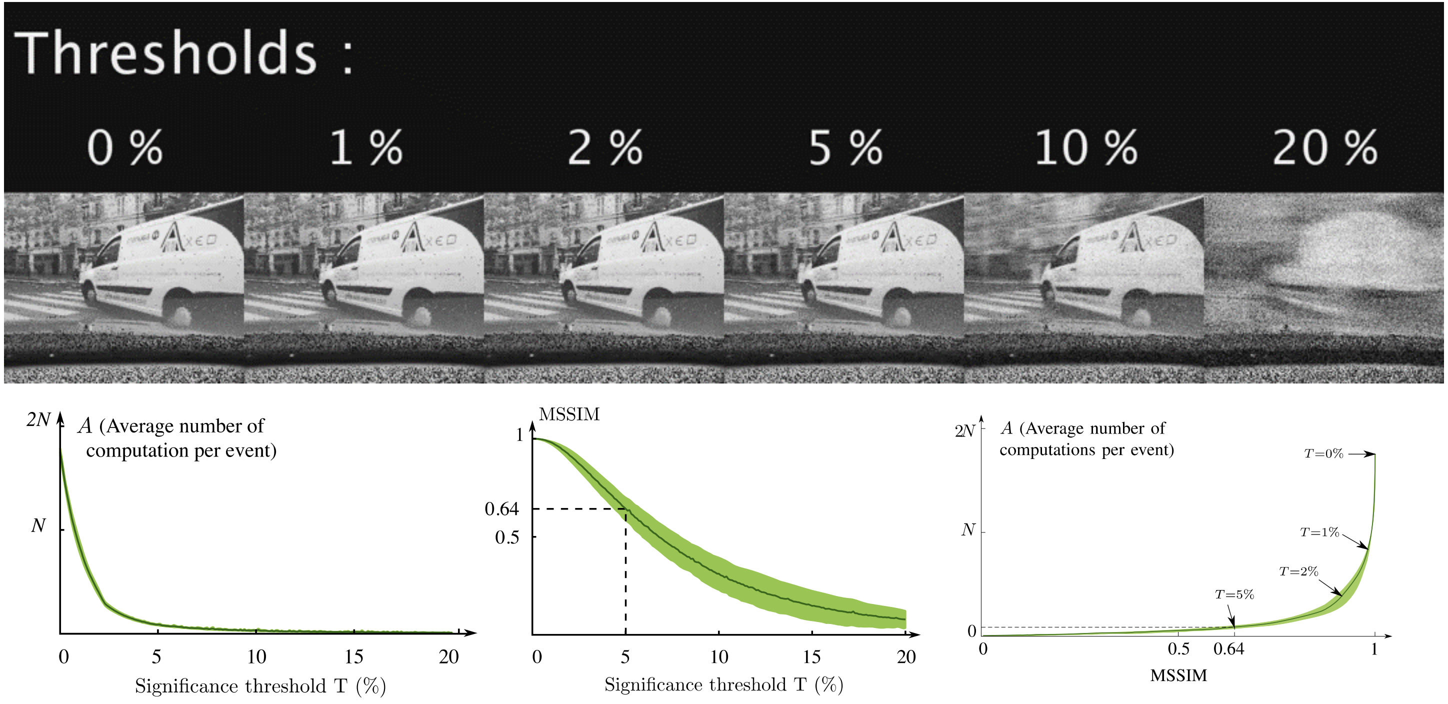

Q. Sabatier, S.H. Ieng, R.B. Benosman, Asynchronous Event-based Fourier Analysis, 2017 February 6, IEEE Transactions on Image Processing 26 (5) :2192-2202, DOI : 10.1109/TIP.2017.2661702

Last Two Decades: Rethinking Computer Vision in The Time Domain

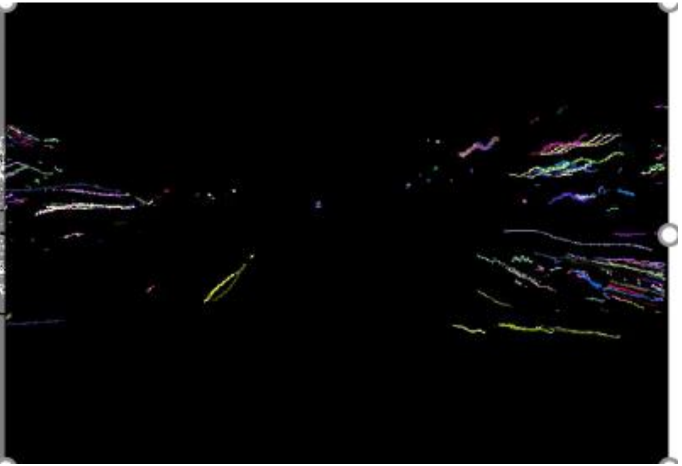

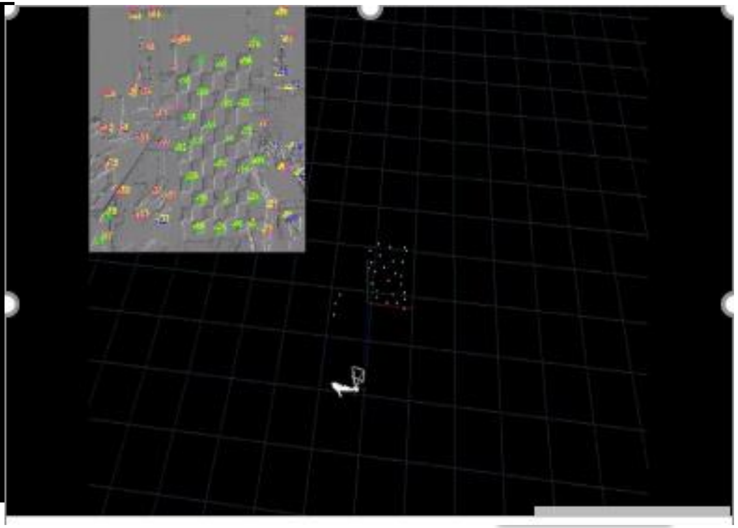

Deep Temporal Learning: Time Surfaces

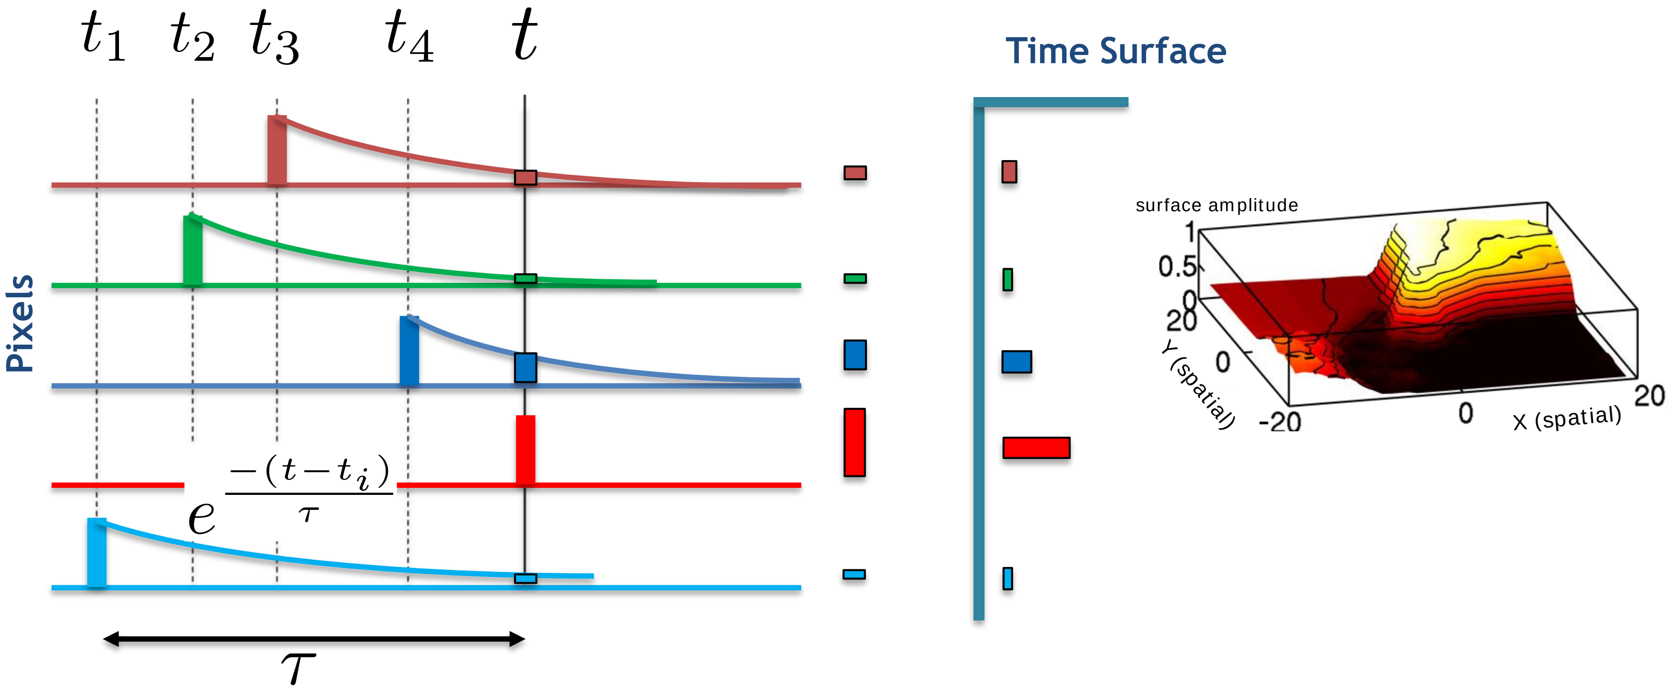

X. Lagorce, G. Orchard, F. Galluppi, B. E. Shi, R.B. Benosman, HOTS : A Hierarchy Of event- based Time-Surfaces for pattern recognition, 2016 July 11, IEEE Transaction on Pattern Analysis and Machine Intelligence 39(7) :1346-1359, doi :10.1109/TPAMI.2016.2574707

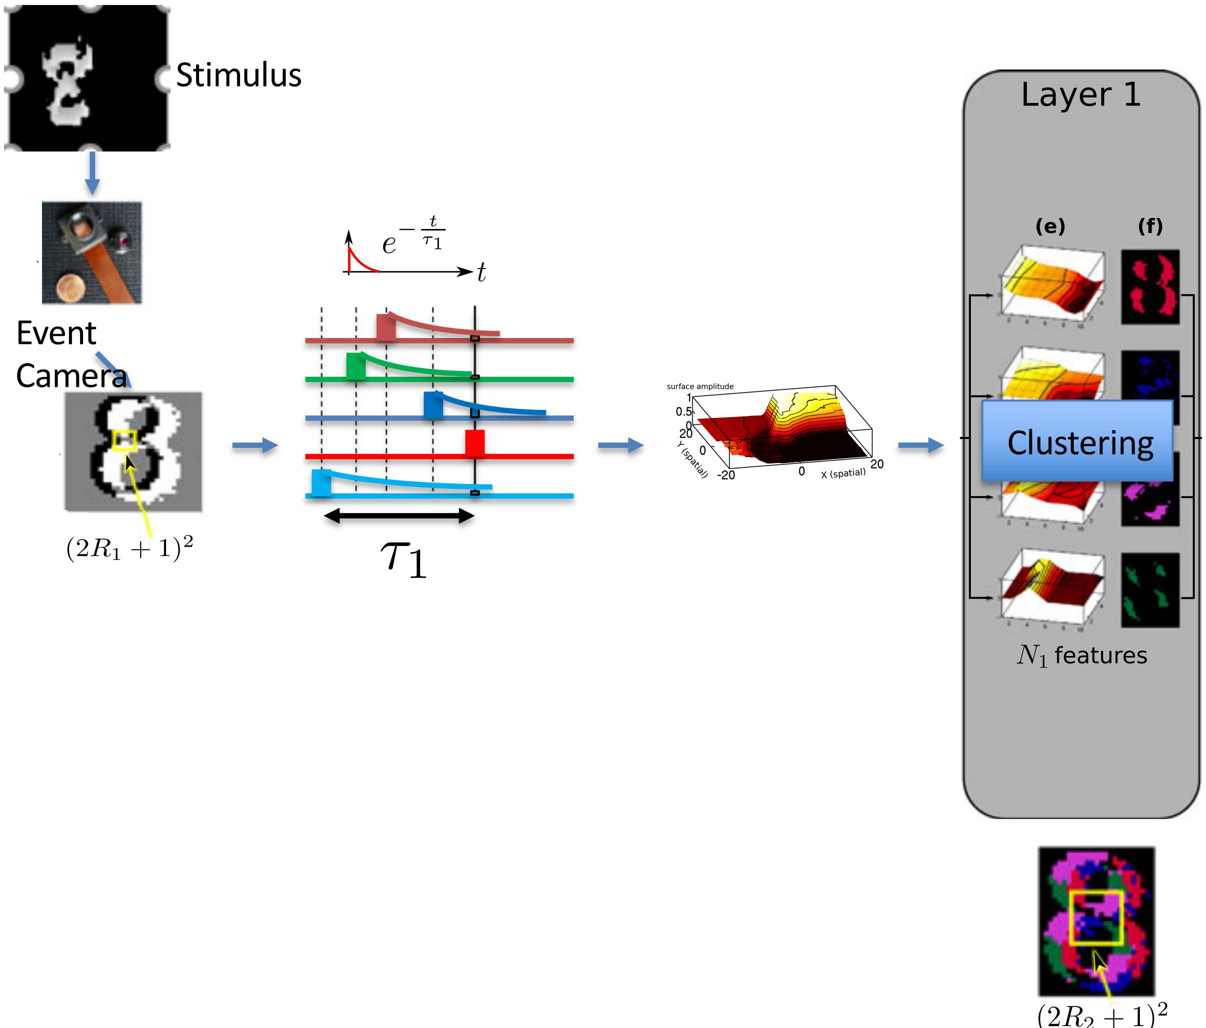
Stimulus
Event
Camera
$(2R_1+1)^2$
$e^{-\frac{t}{\tau_1}}$
t
τ_1
surface amplitude
X (spatial)
Y (spatial)
Layer 1
(e)
(f)
Clustering
N_1 features
$(2R_2+1)^2$

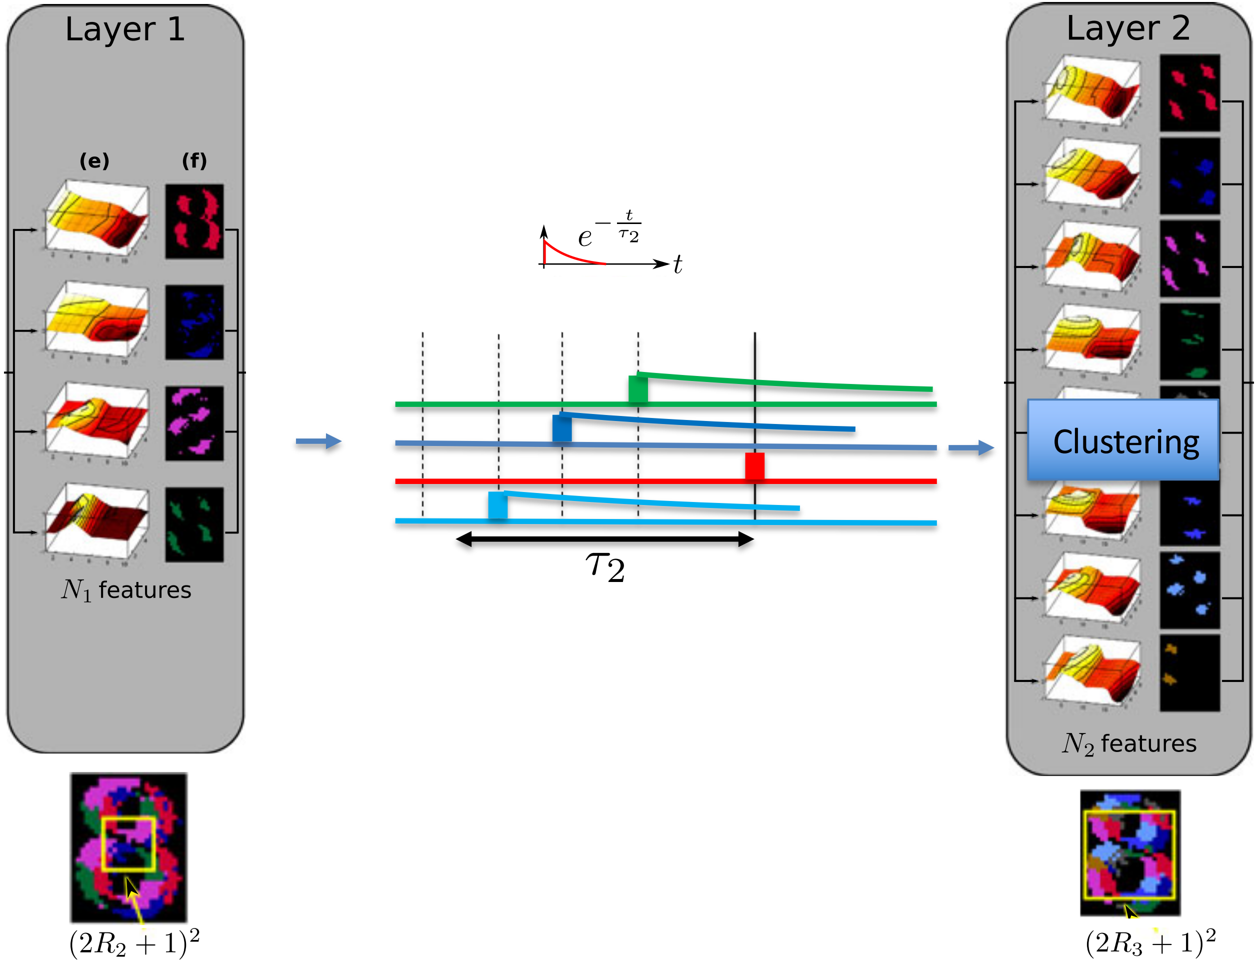
Layer 1
(e)
(f)
N_1 features
$(2R_2+1)^2$
$e^{-\frac{t}{\tau_2}}$
t
τ_2
Layer 2
Clustering
N_2 features
$(2R_3+1)^2$

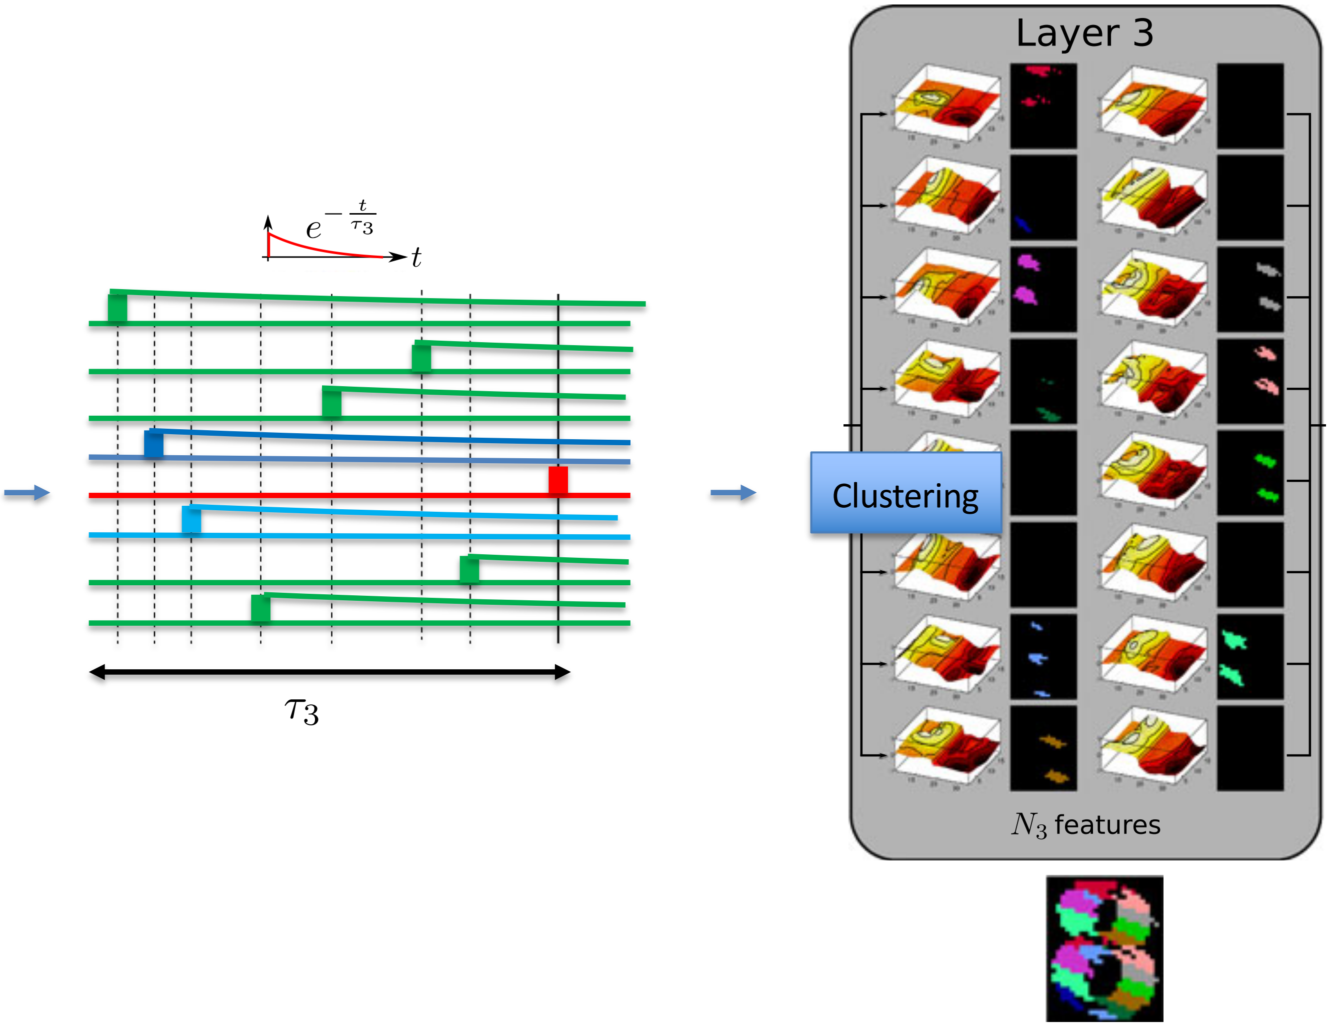

$e^{-\frac{t}{\tau_3}}$
t
τ_3
Layer 3
Clustering
N_3 features

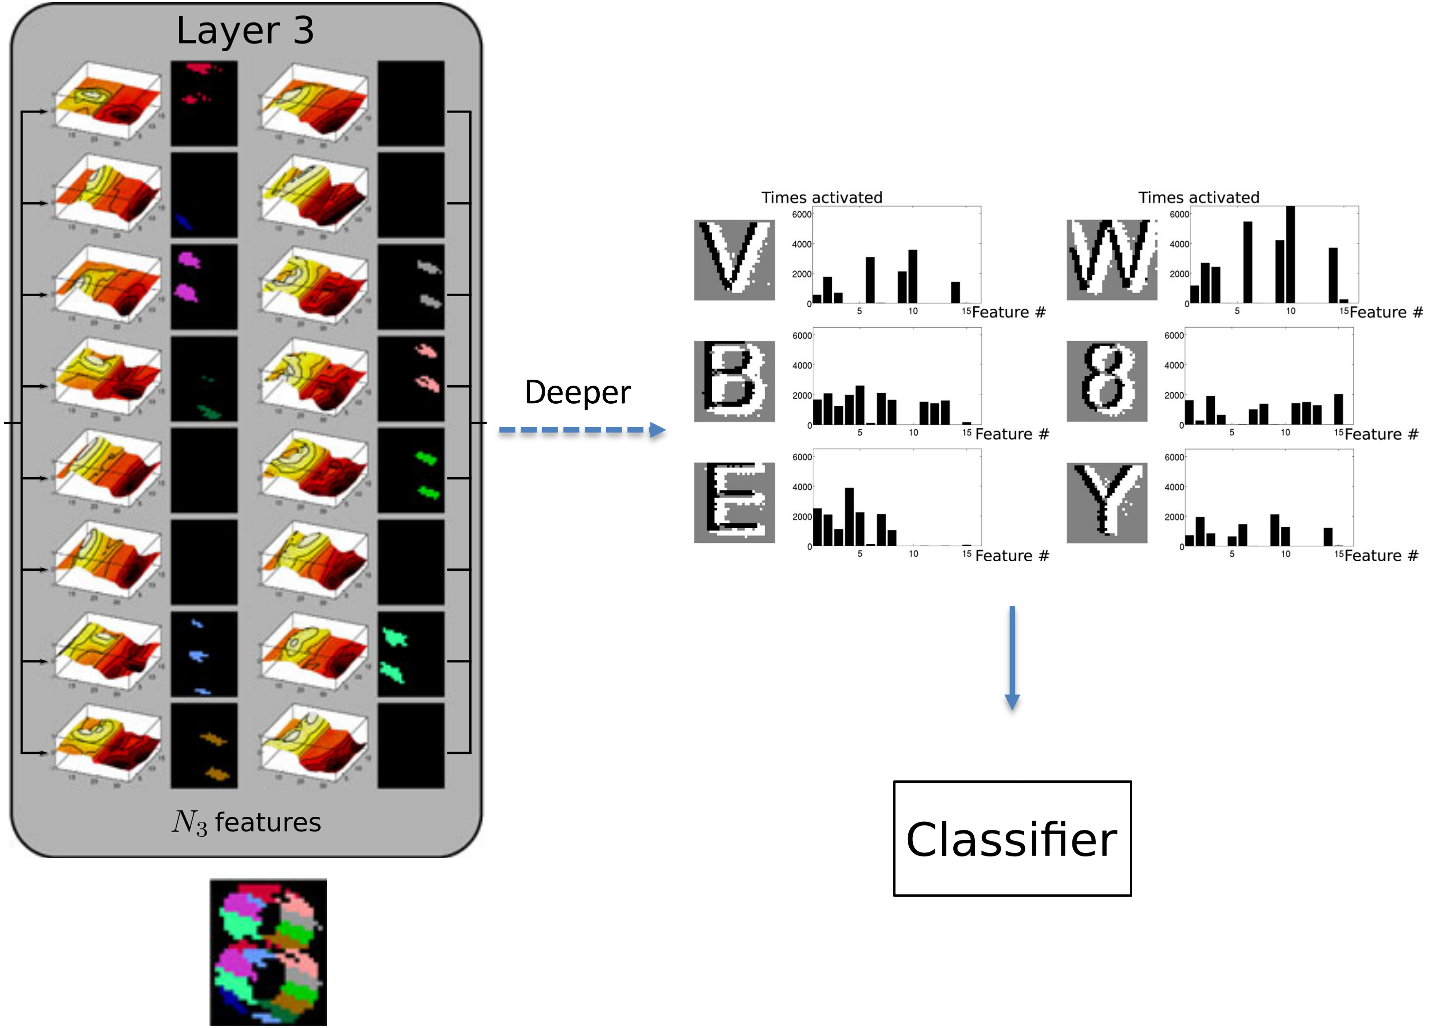
Layer 3
N_3 features
Deeper
Times activated
Feature #
Classifier

Deep Temporal Learning with Adaptive Temporal Feedback: Temporal Surfaces

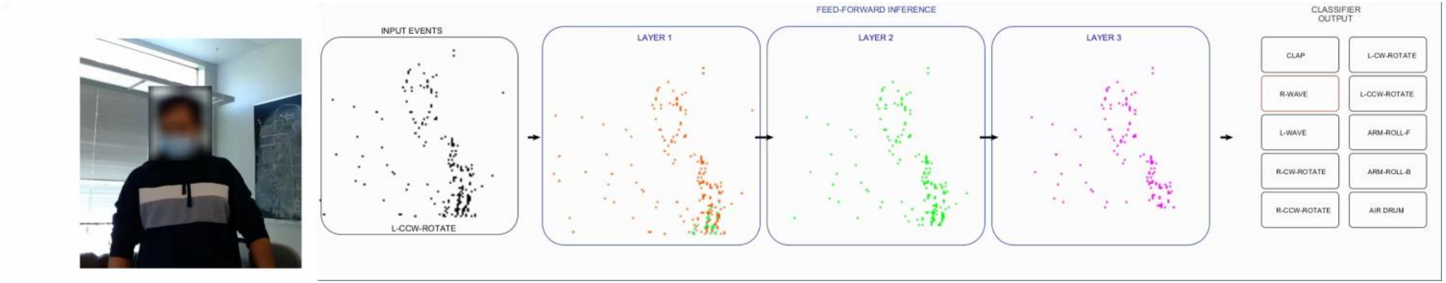

X. Lagorce, G. Orchard, F. Galluppi, B. E. Shi, R.B. Benosman, HOTS : A Hierarchy Of event- based Time-Surfaces for pattern recognition, 2016 July 11, IEEE Transaction on Pattern Analysis and Machine Intelligence 39(7) :1346-1359, doi :10.1109/TPAMI.2016.2574707

Deep Temporal Learning with Adaptive Temporal Feedback: Temporal Surfaces

X. Lagorce, G. Orchard, F. Galluppi, B. E. Shi, R.B. Benosman, HOTS : A Hierarchy Of event- based Time-Surfaces for pattern recognition, 2016 July 11, IEEE Transaction on Pattern Analysis and Machine Intelligence 39(7) :1346-1359, doi :10.1109/TPAMI.2016.2574707

Computation Platfoms?

Two tendencies

Replicate

understand

Biomimetism

Understand and Accelerate

Computation Platfoms?

Replicate

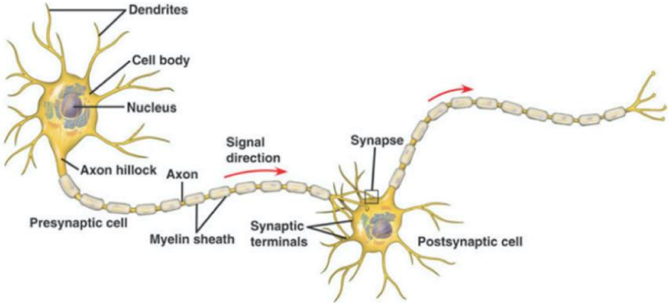

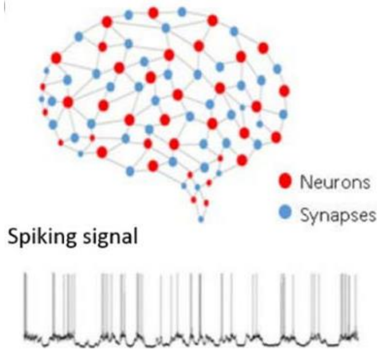

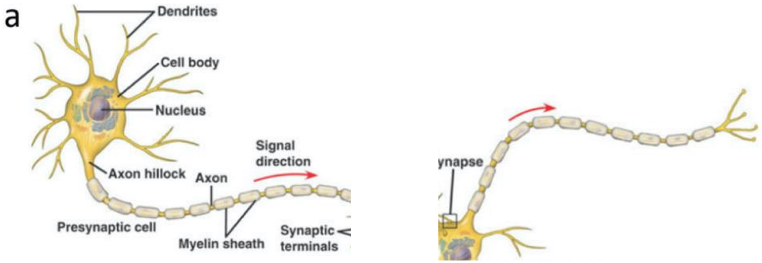

*From: Goi, E., Zhang, Q., Chen, X. et al. Perspective on photonic memristive neuromorphic computing. PhotoniX **1,** 3 (2020). https://doi.org/10.1186/s43074-020-0001-6*

Analog vs Digital

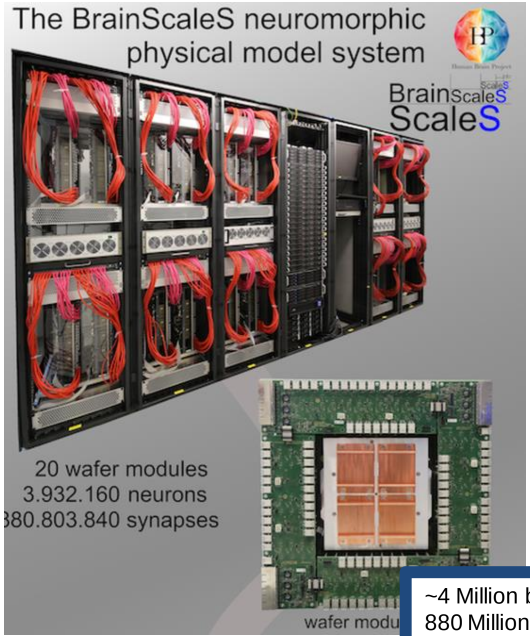

20 wafer modules
3.932.160 neurons
880.803.840 synapses

~4 Million bio-realistic neurons
880 Million learning synapses
105 faster than real time
164 Giga-events/s (normal)
1 Tera-event/s (burst)
several hundreds of kW [2010]

https://wiki.ebrains.eu/bin/view/Collabs/neuromorphic/BrainScaleS/

Neuromorphic Computing

Today's digital neuromorphic hardware

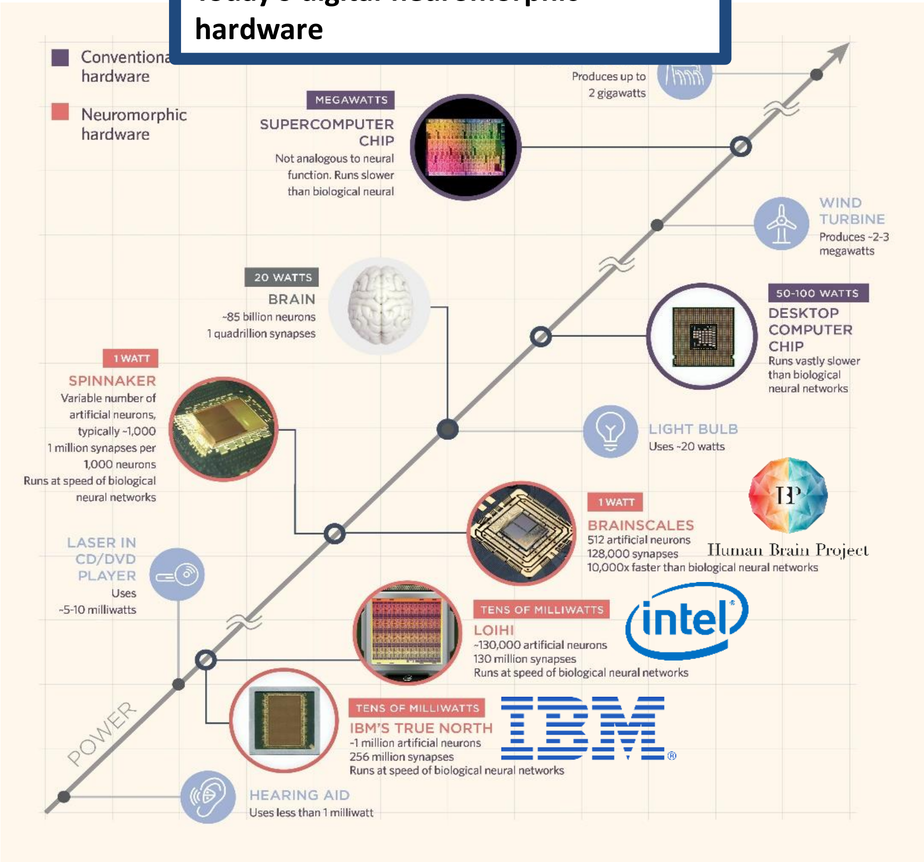

[from The Scientist, 2019]

A Processing Solution Adapted to Event Data

understand

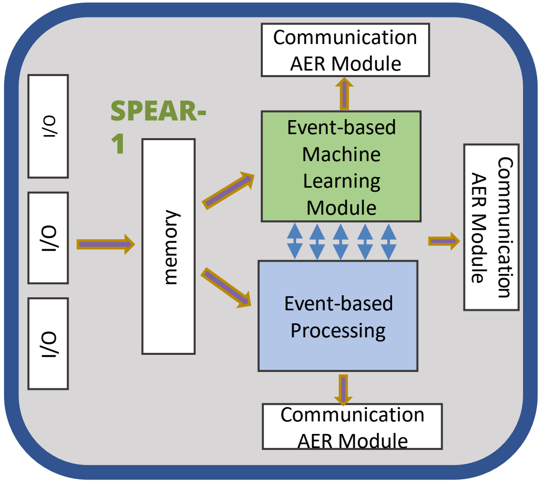

- Architecture for **Event per Event**: **programmable event-based computation** and **machine learning**
- Industry standard I/O and programmability
- **Scalable**, enabling fast and cost-effective derivatives

<10-100mW and up to 20 GigaEvents/s processing

Radar and LIDAR and much more…

Improving the Accuracy of Spiking Neural Networks for Radar Gesture Recognition Through Preprocessing

Ali Safa, *Graduate Student Member, IEEE*, Federico Corradi, *Member, IEEE*, Lars Keuninckx, Ilja Ocket, *Member, IEEE*, André Bourdoux, *Senior Member, IEEE*, Francky Catthoor, *Fellow, IEEE*, and Georges G. E. Gielen, *Fellow, IEEE*

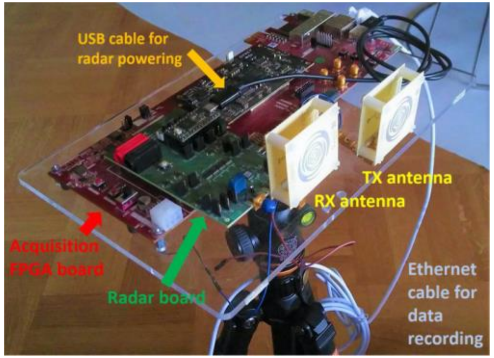

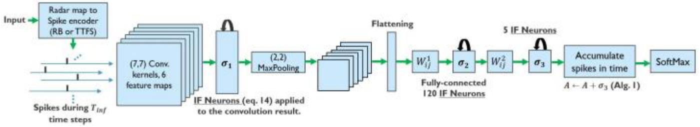

Event-Based Structured Light for Depth Reconstruction using Frequency Tagged Light Patterns

T. Leroux, S.-H. Ieng and R. Benosman
University of Pittburgh, Carnegie Mellon Univeristy, Sorbonne Universitas
benosman@pitt.edu

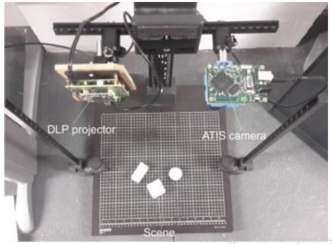

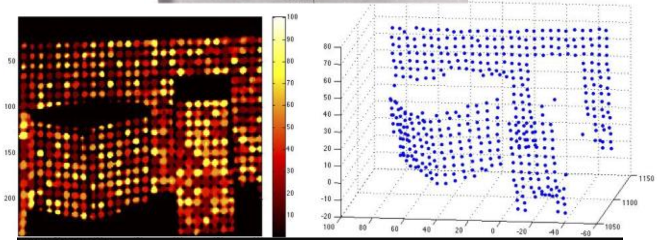

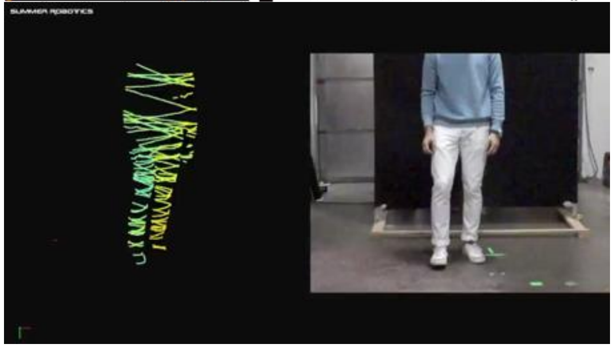

https://www.summerrobotics.ai

Sight Restoration: Prosthetics and Optogenetics

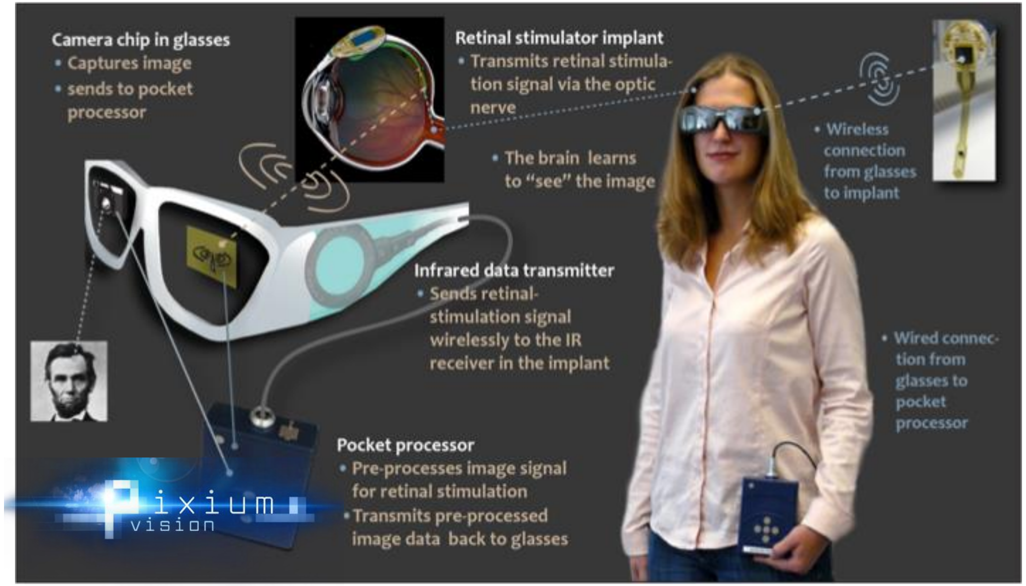

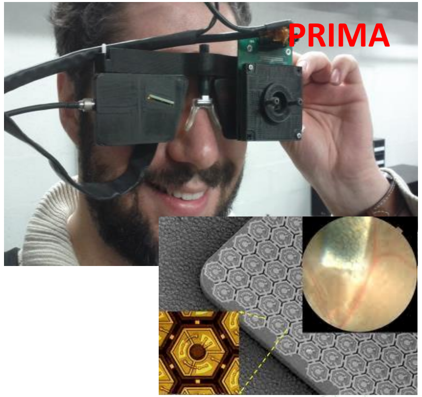

GenSight

- Development of Retina Stimulation Goggles
- 3 generations of Retina Prosthetics
- Asynchronous Retina Stimulation: Prosthetics and Optogenetics

Sight Restoration: Prosthetics and Optogenetics

& much more..


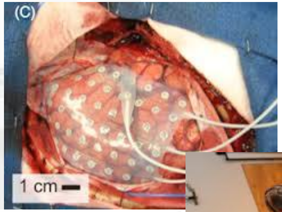


Space awareness

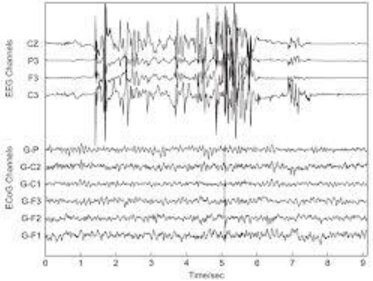
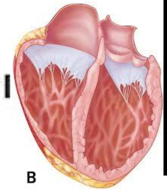


Low power Online decoding and classification


Décision making: game theory stock Market

Robotics

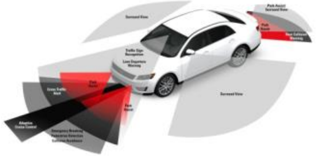

ADAS


Sensory Substitution

Always on sensing

Conclusions

A Chance to ground Perception in Basic Science!

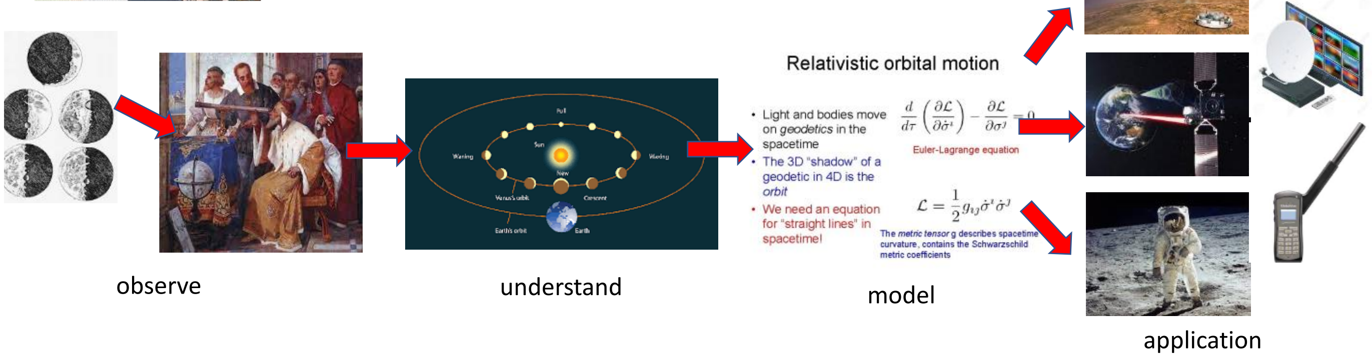

- A whole new world to explore
- A deep paradigm shift for Sensing & AI
- Novel sensors to build
- New adapted processing architectures to design

Conclusions

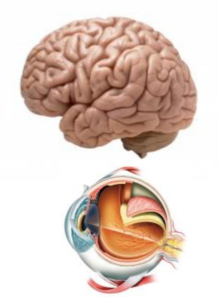

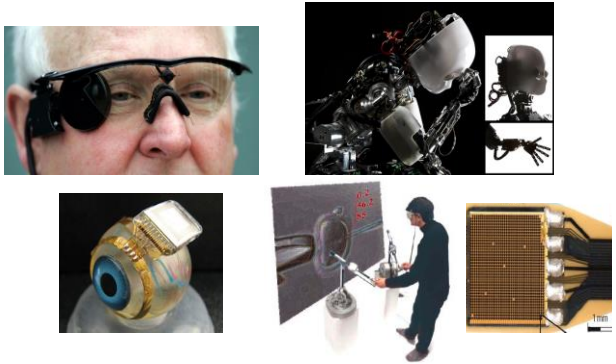

Physiology →
- Models
- Hardware

→
- Prosthetics
- Robotics
- Computation

- A whole new world to explore
- A deep <u>paradigm shift for Sensing & AI</u>
- Novel sensors to build
- New adapted processing architectures to design