PLEIADES

PoLynomial Expansion In Adaptive
Distributed Event-based Systems

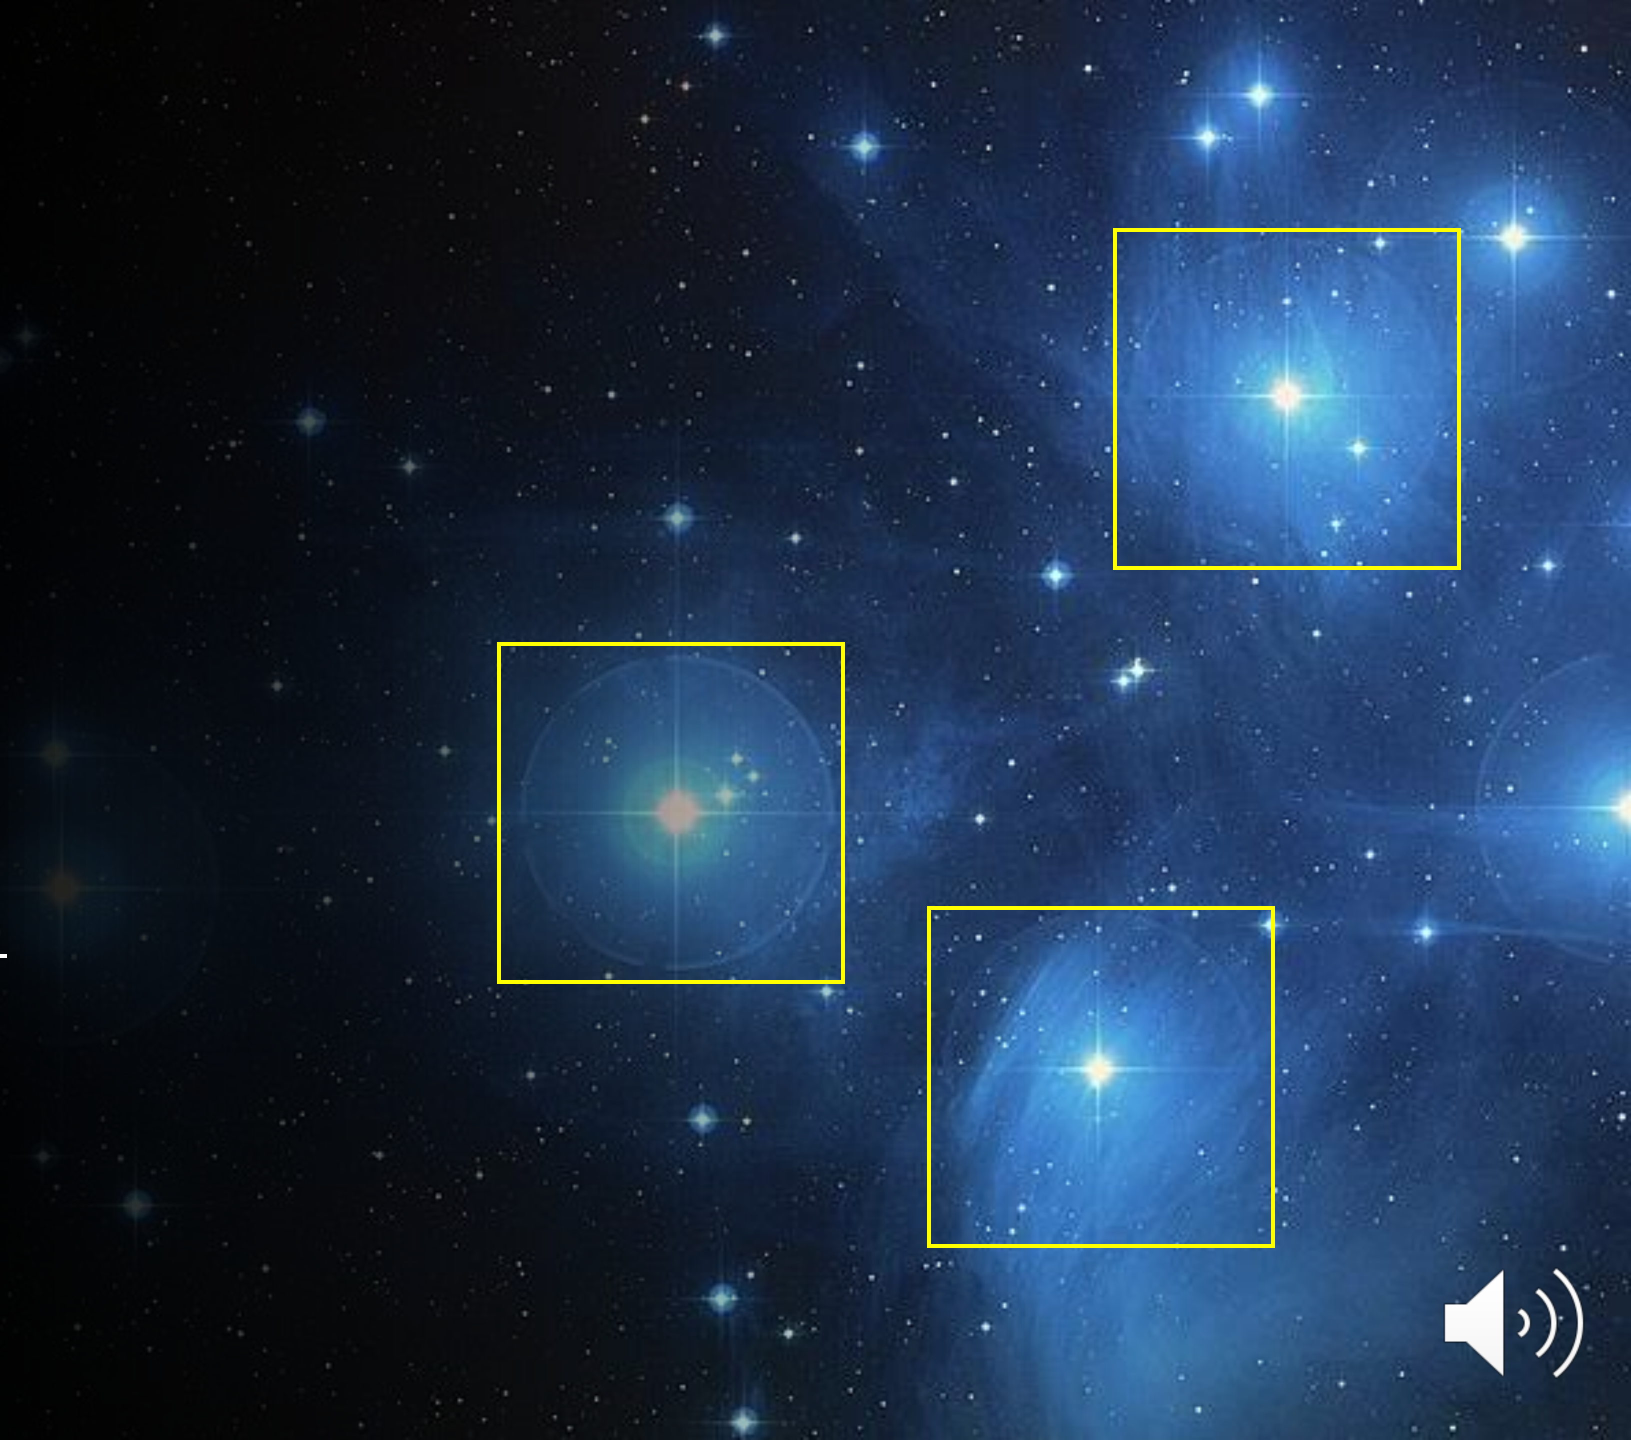

Motivation

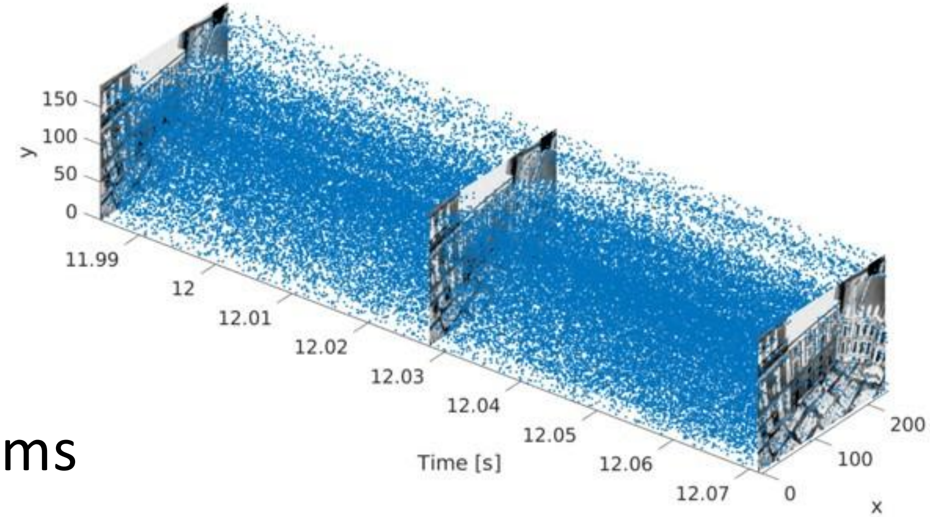

- We want long temporal kernels for event-based systems
- Long temporal kernels
  - Parameterized with coefficients of orthogonal polynomials
  - Fewer parameters = less memory and training stability
- Freely resampled
  - FPS / step-size of input event-frames is flexible
- Optimized computation
  - Via tensor contractions (see CENTAURUS)

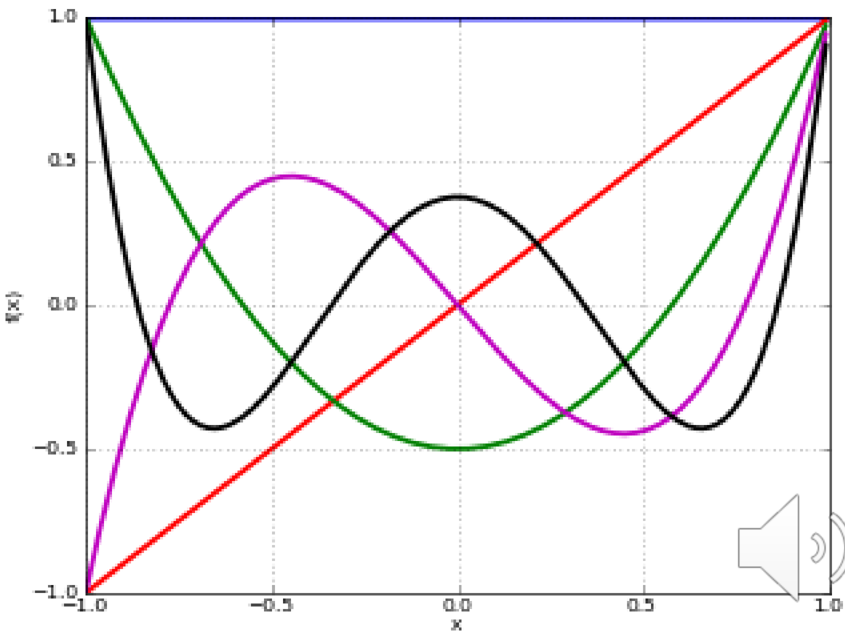

Temporal Kernel from Orthogonal Polynomials

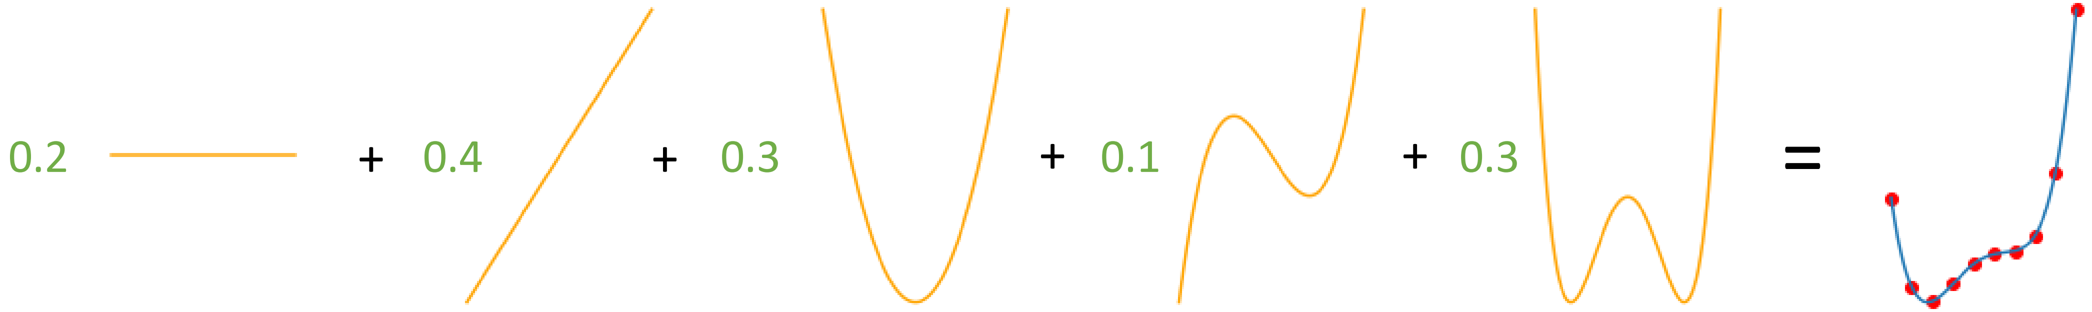

- Express temporal kernel as the sum of a set of coefficients multiplied to the basis functions (orthogonal polynomials), which is then uniformly sampled
- Note: **coefficients are trainable**, but basis functions are untrainable (fixed)
- 💡 : a linear mapping: coefficient vector $\mapsto$ temporal kernel values (In this example: $\mathbb{R}^5 \mapsto \mathbb{R}^{10}$)
- Different ways to discretize: pointwise sampling, trapezoid, analytic integration

With multiple channels

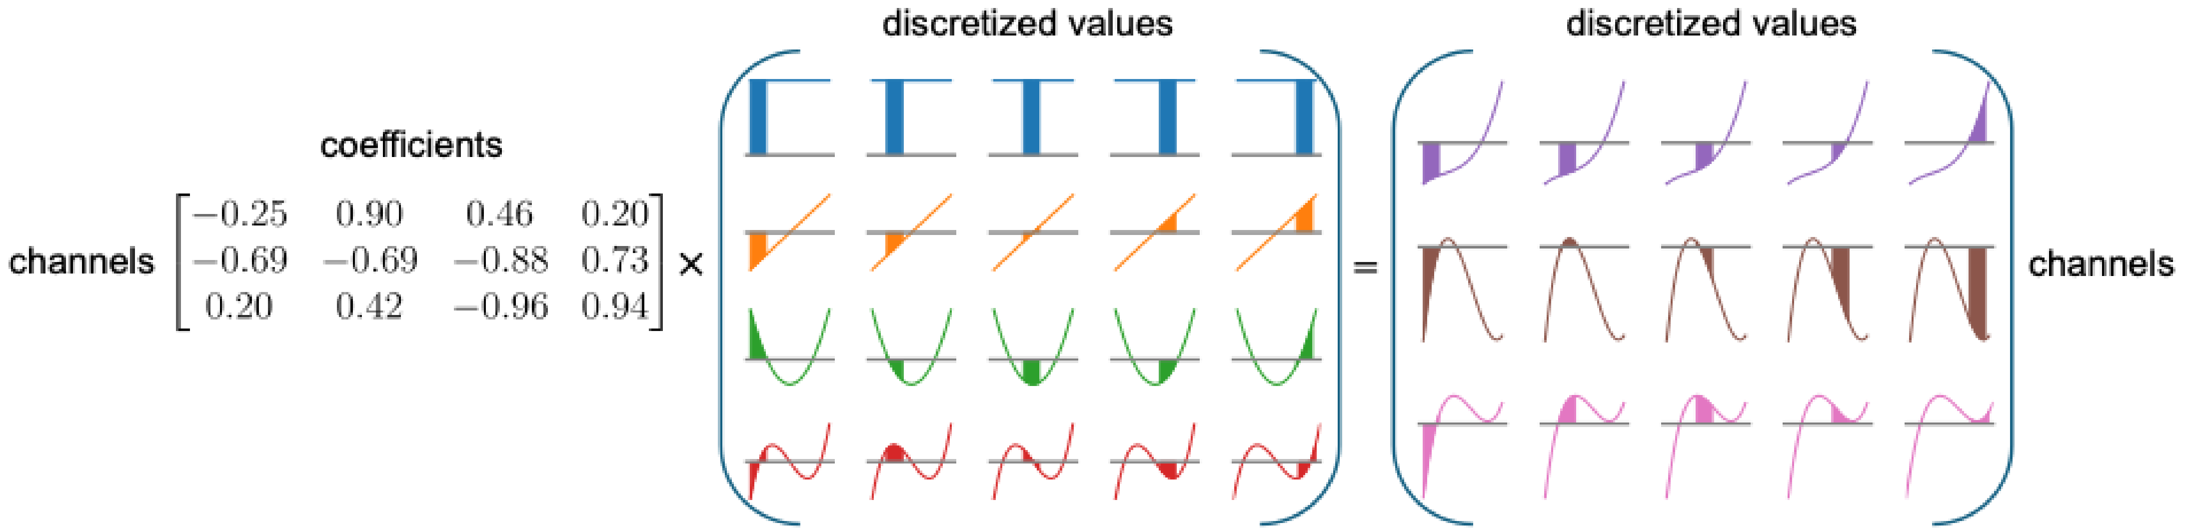

This is just one way of doing it. You generate the temporal kernels, then convolve with inputs.

Can also reproject the inputs into coefficient space and convolve directly there (SSM origins).

Optimal order of operations depend on the input and kernel dimensions

See our sister work, Centaurus: https://arxiv.org/abs/2501.13230

Network architecture (auxiliary)

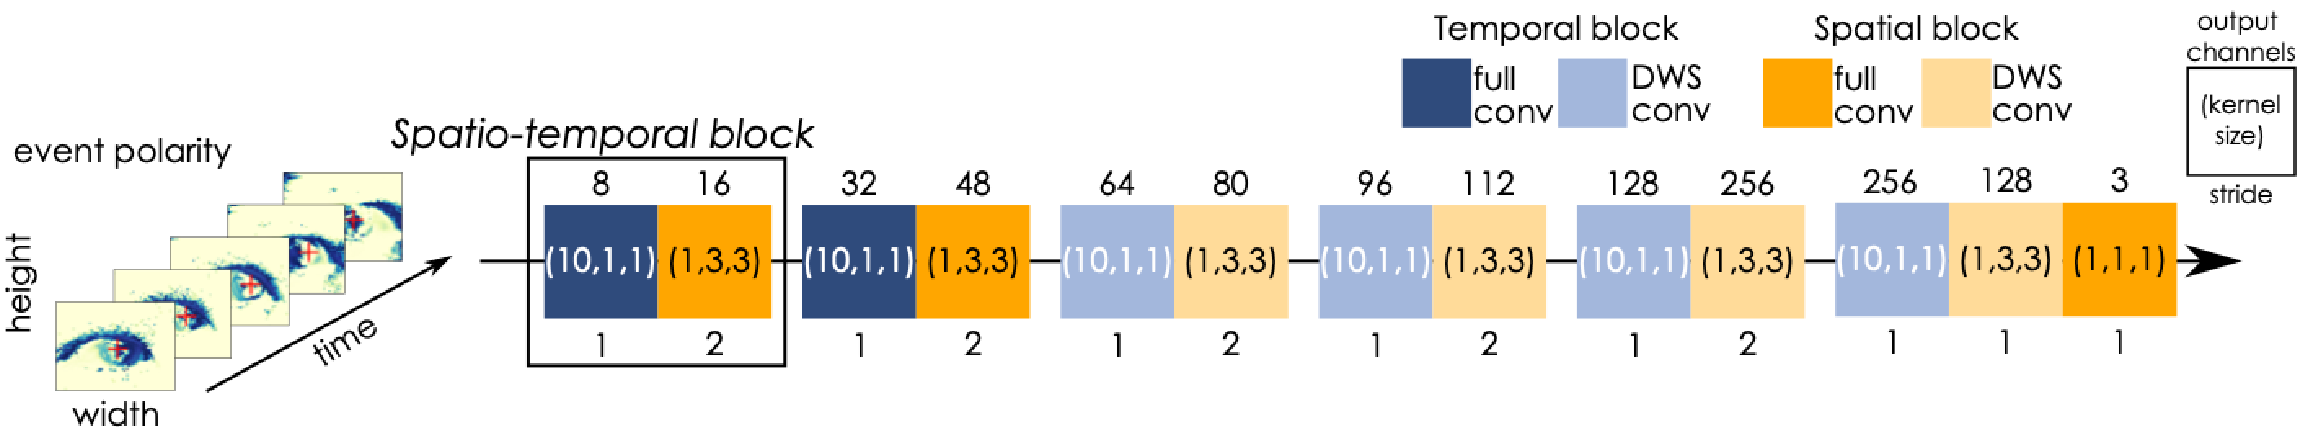

Separated temporal and spatial conv blocks for compute reduction: (1+2)D conv net

Each temporal and spatial block can be further made depthwise-separable

Causal temporal convolutions, needed for real-time processing

Lightweight activations, ReLU, good for sparsity, and edge deployments

DVS Hand Gesture recognition: IBM DVS128 Dataset

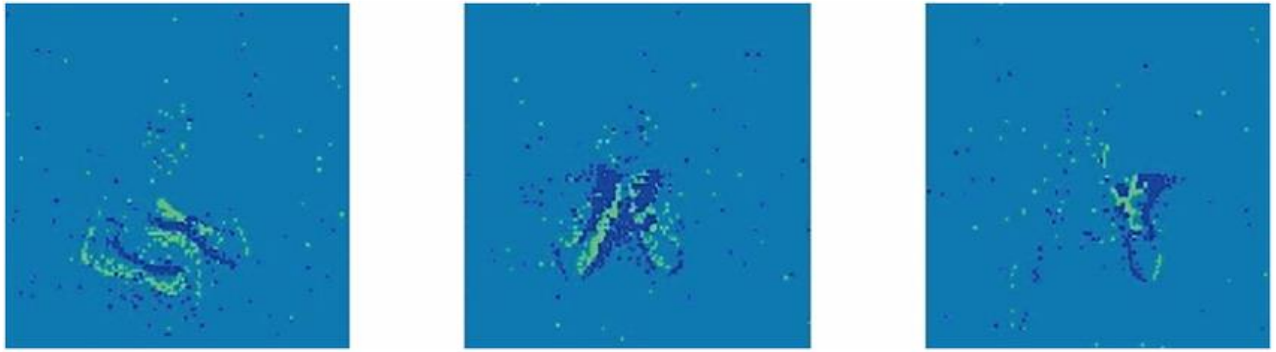

Model	Accuracy	Parameters	MACs / sec
Conv(1+2)D	99.17	196 K	0.429 B
ANN-Rollouts (Kugele et al., 2020)	97.16	500 K	10.4 B
TrueNorth CNN* (Amir et al., 2017)	96.59	18 M	
SLAYER (Shrestha and Orchard, 2018)	93.64		
PLEIADES	99.59 (0.02)	**192 K**	**0.499 B**
PLEIADES + filtering*	**100.00 (0.00)**	**192 K**	**0.499 B**

Temporal kernels freely resampled

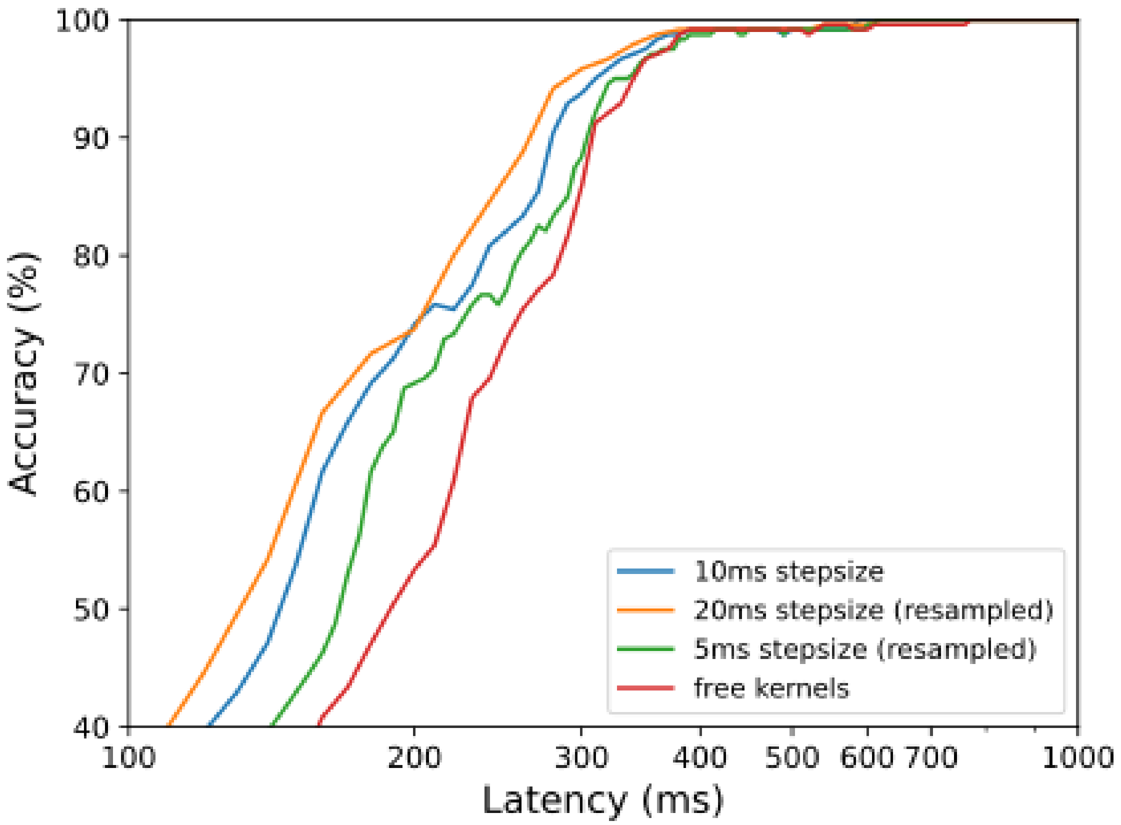

Note: we do not do additional fine-tuning, only re-sampling

Event-based eye tracking: CVPR AIS 2024

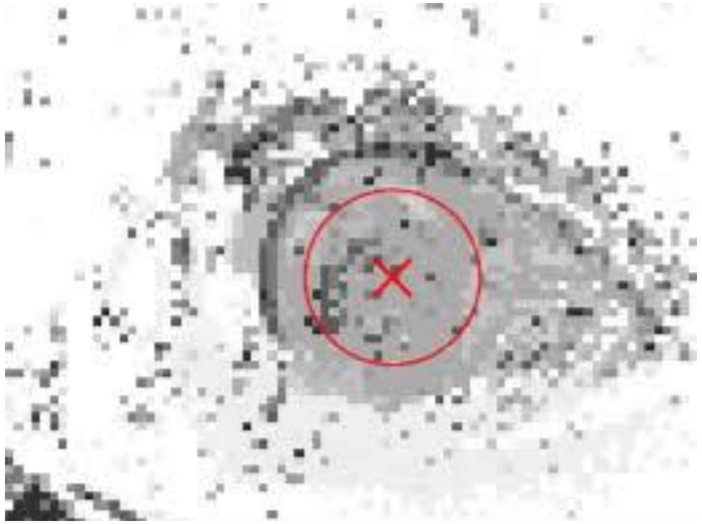

Model	p10	p5	p3	Parameters
MambaPupil	99.42	97.05	90.73	-
CETM	99.26	96.31	83.83	7.1M
Conv(1+2)D	99.00	97.97	94.58	1.1M
ERVT	98,21	94.94	87.26	150K
PEPNet	97.95	80.67	49.08	640K
PLEIADES + CenterNet	**99.58 (0.03)**	**97.95 (0.03)**	**94.94 (0.04)**	**277K**

Event-based road-scene object detection: Prophesee GEN4

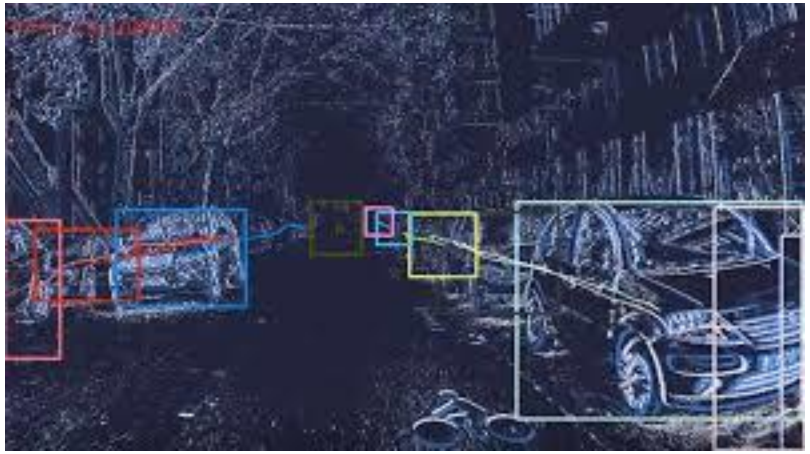

Model	mAP	Parameters	MACs / sec	FPS
RED (Perot et al., 2020)	0.43	24.1 M		20
Gray-RetinaNet (Perot et al., 2020)	0.43	32.8 M	2060 B	20
S5-ViT-B (Zubic et al., 2024)	0.478	18.2 M		20
GET-T (Peng et al., 2023)	0.484	21.9 M		
PLEIADES + CenterNet	**0.556**	**0.576 M**	**122.5 B**	**100**

Check PLEIADES out!

Paper:

Code: