Computing Spatial Image Convolutions for Event Cameras

Cedric Scheerlinck

Prof. Rob Mahony
A/Prof. Nick Barnes

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

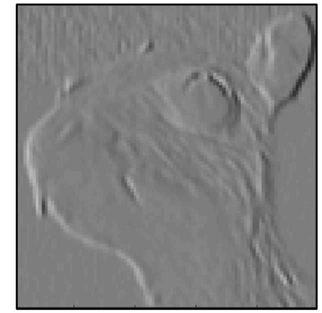

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

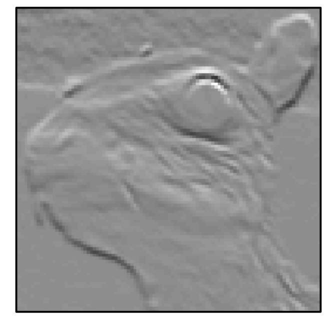

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.

Image Convolutions

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

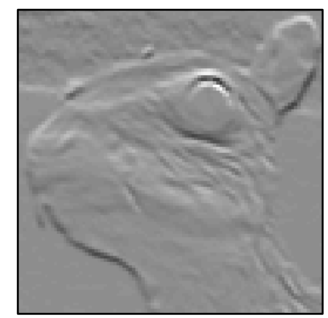

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

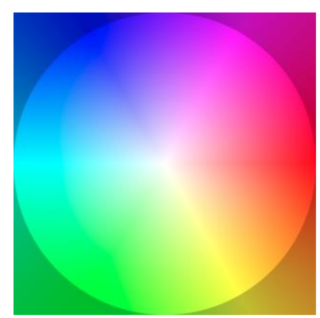

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

$$=$$

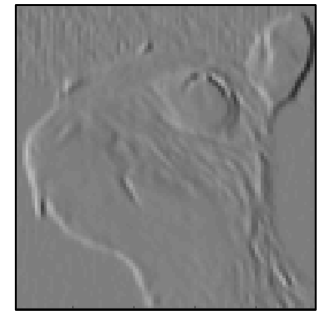

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

=

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

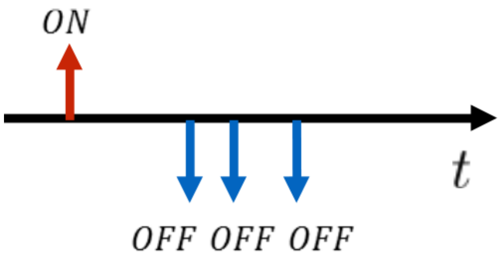

?

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

$$=$$

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

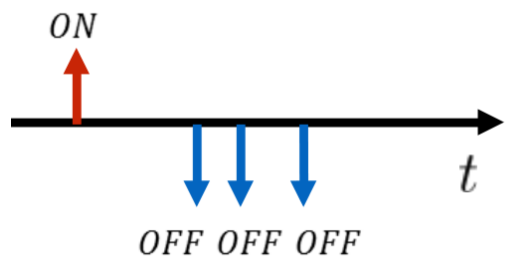

?

Naïve approach: reconstruct image frames from events then apply convolution.

Can we do better?

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

Consider one event 

[timestamp, x, y, ±1]

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

Consider one event

[timestamp, x, y, ±1]

Event image

0	0	0	0	0	0
0	0	0	0	0	0
0	0	-1	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

Consider one event 

[timestamp, x, y, ±1]

Kernel

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

Event image

0	0	0	0	0	0
0	0	0	0	0	0
0	0	-1	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.

Image Convolutions for Events

Consider one event 

[timestamp, x, y, ±1]

Kernel ⋆ Event image

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} *$$

0	0	0	0	0	0
0	0	0	0	0	0
0	0	-1	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0

=

0	0	0	0	0	0
0	1	0	-1	0	0
0	2	0	-2	0	0
0	1	0	-1	0	0
0	0	0	0	0	0
0	0	0	0	0	0

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

Kernel $\ast$ Event image

0	0	0	0	0	0
0	1	0	-1	0	0
0	2	0	-2	0	0
0	1	0	-1	0	0
0	0	0	0	0	0
0	0	0	0	0	0

Six virtual events, or a **convolved event**, can be generated

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.

Image Convolutions for Events

Convolved events can be used as input to an event processing algorithm.

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Image Convolutions for Events

Convolved events can be used as input to an event processing algorithm.

For example image reconstruction [1]:

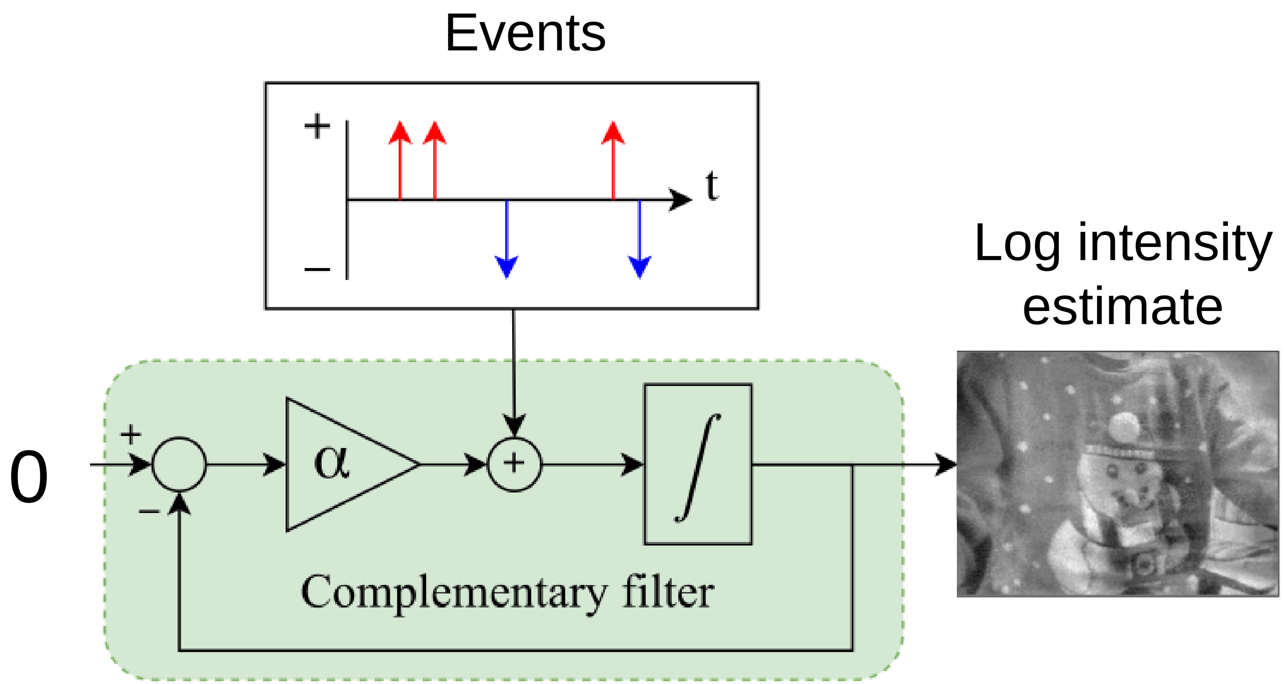

[1] (Scheerlinck et al., ACCV, 2018)

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.

Image Convolutions for Events

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

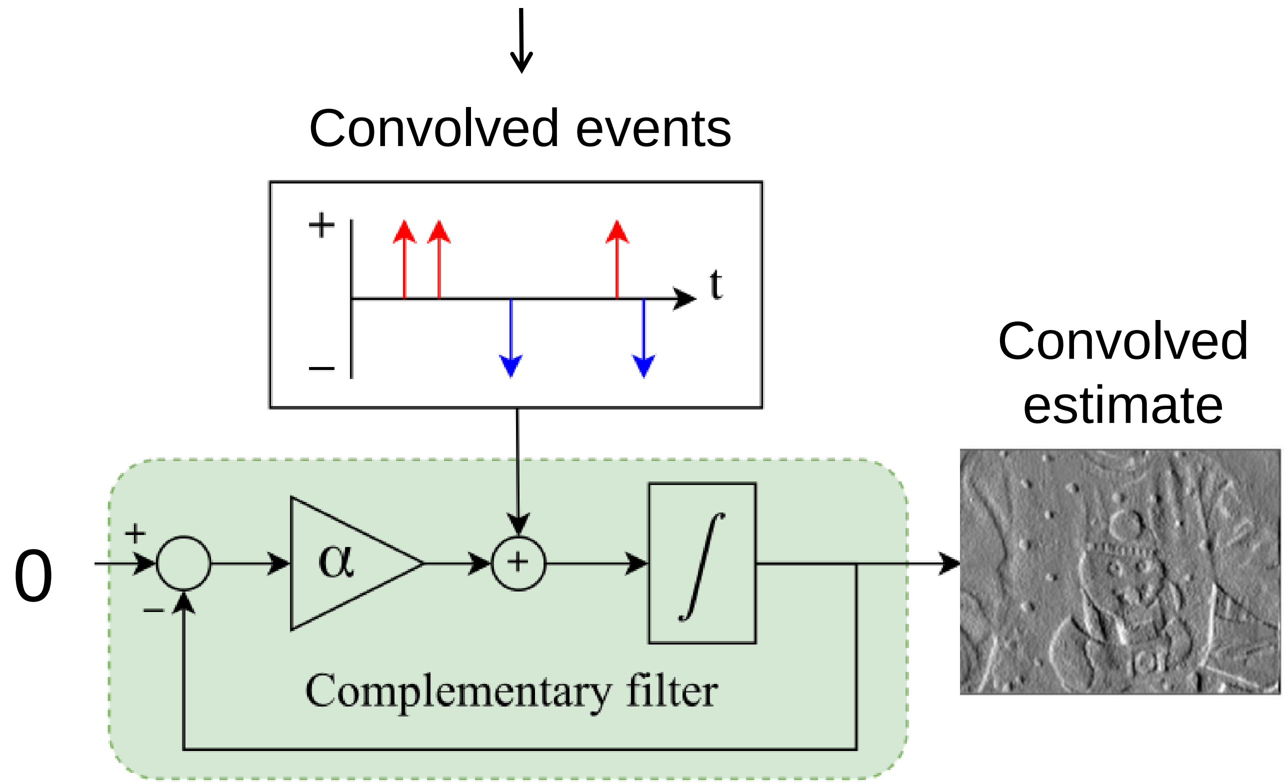

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

Identity

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Gaussian

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

Sobel

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

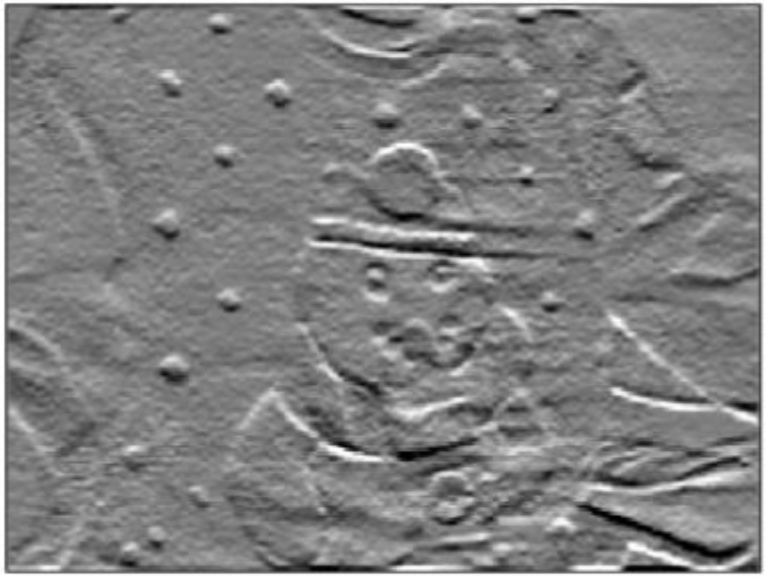

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

Laplacian

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & -12 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

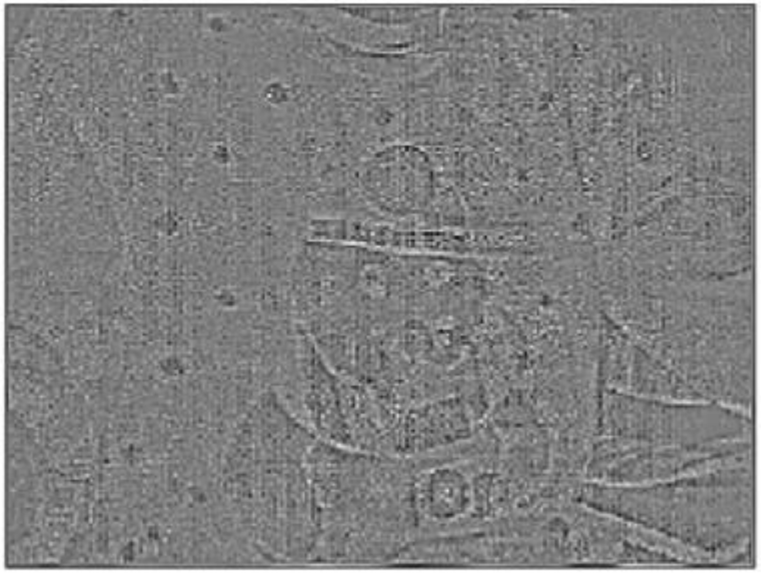

Poisson Reconstruction from Laplacian

Perez et al. (2003), Agrawal et al. (2006)

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

Current events | Gradient estimate

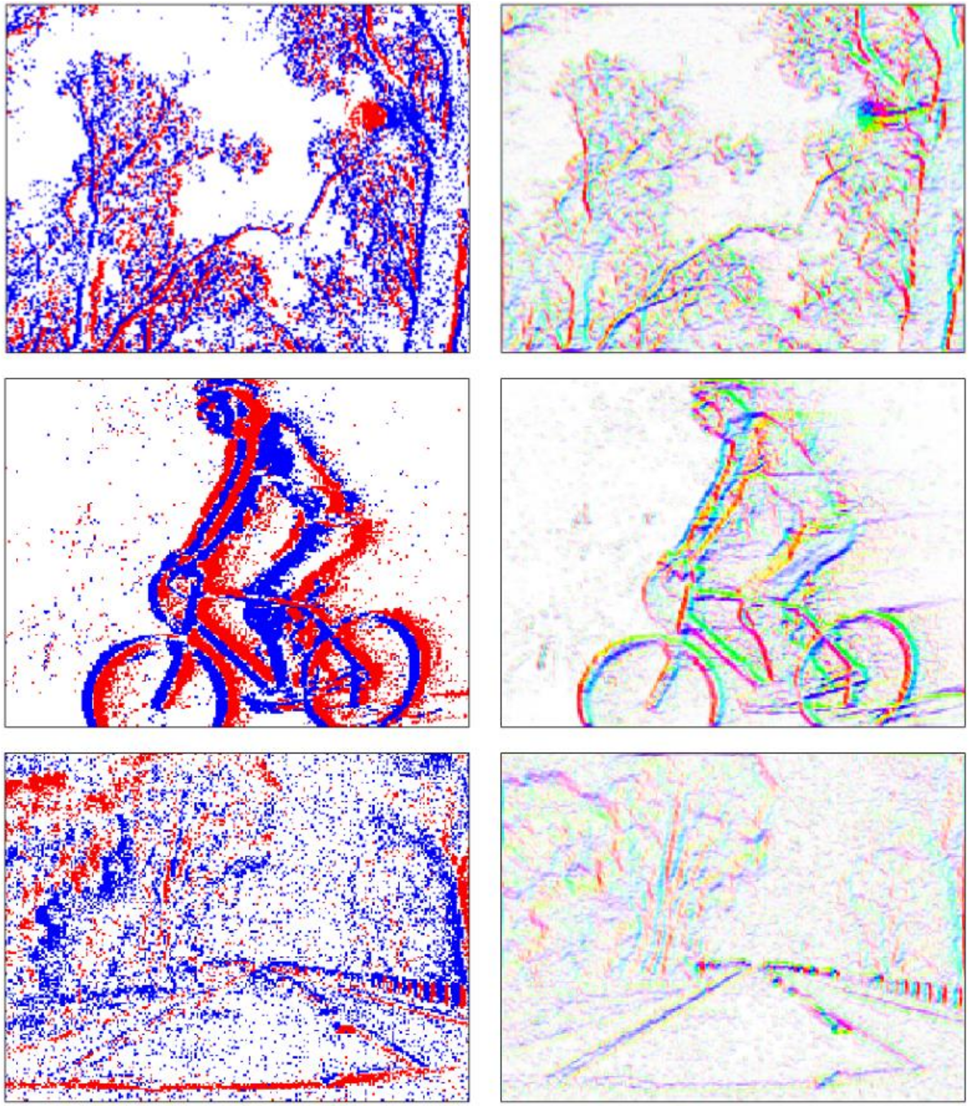

Sobel

$$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

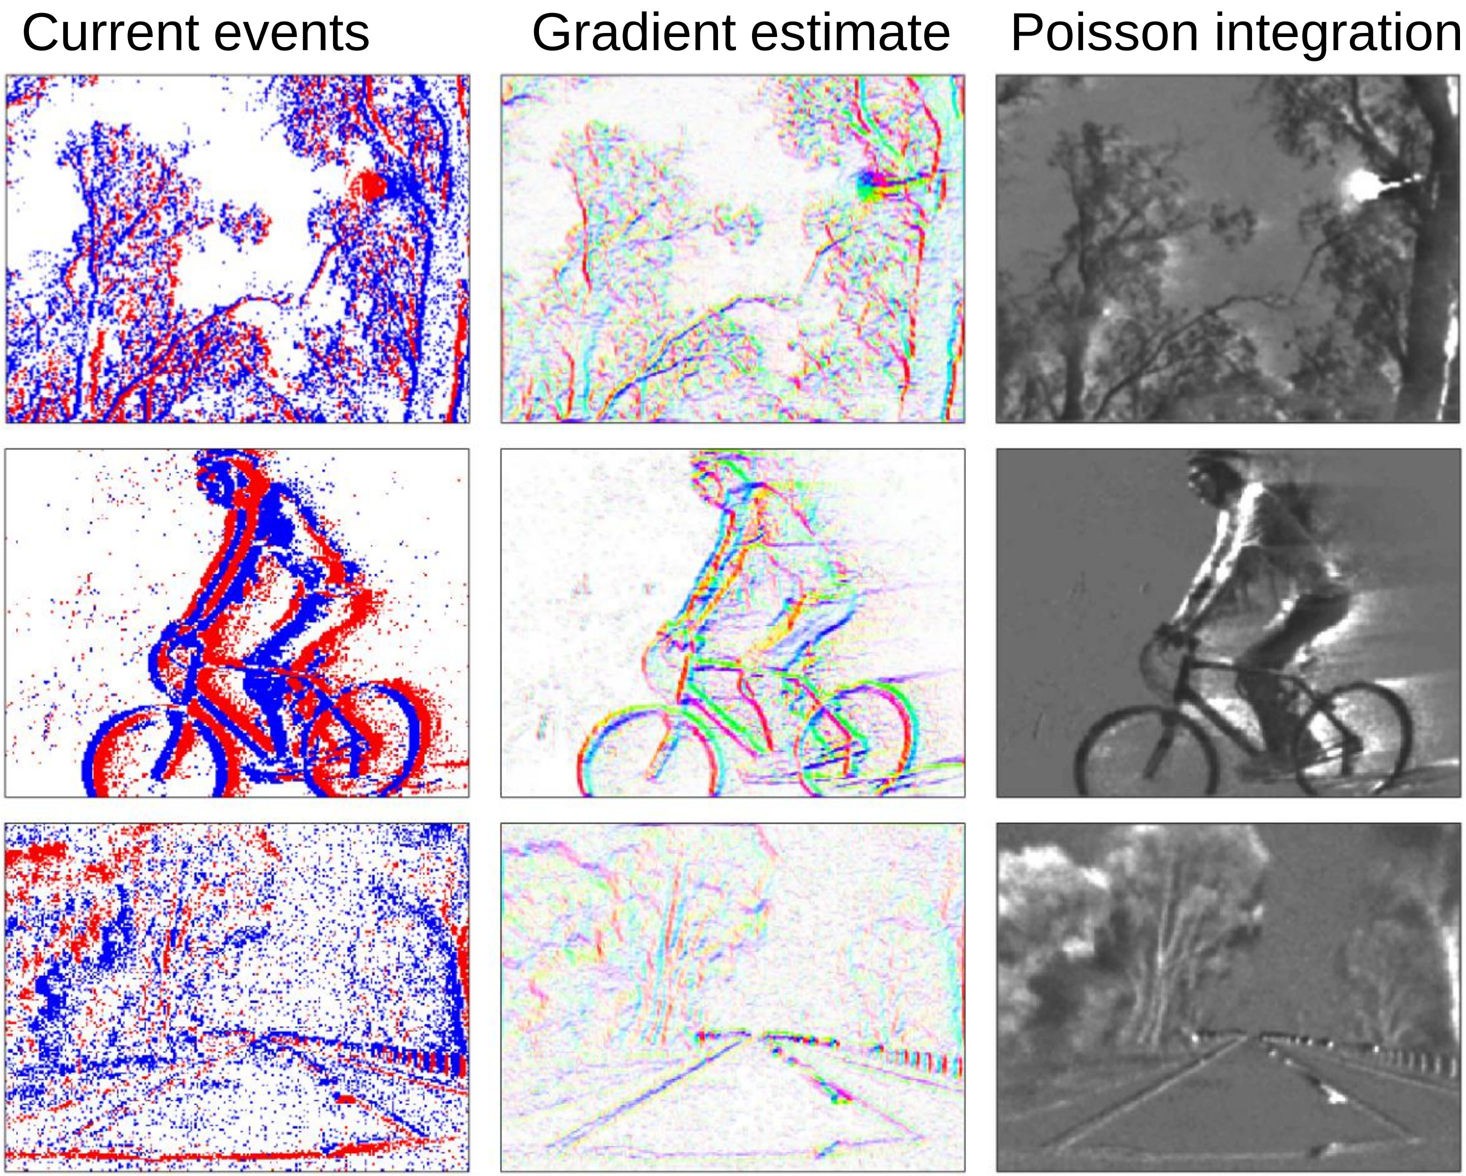

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

Gradient can be used as input to asynchronous Harris corner detector.

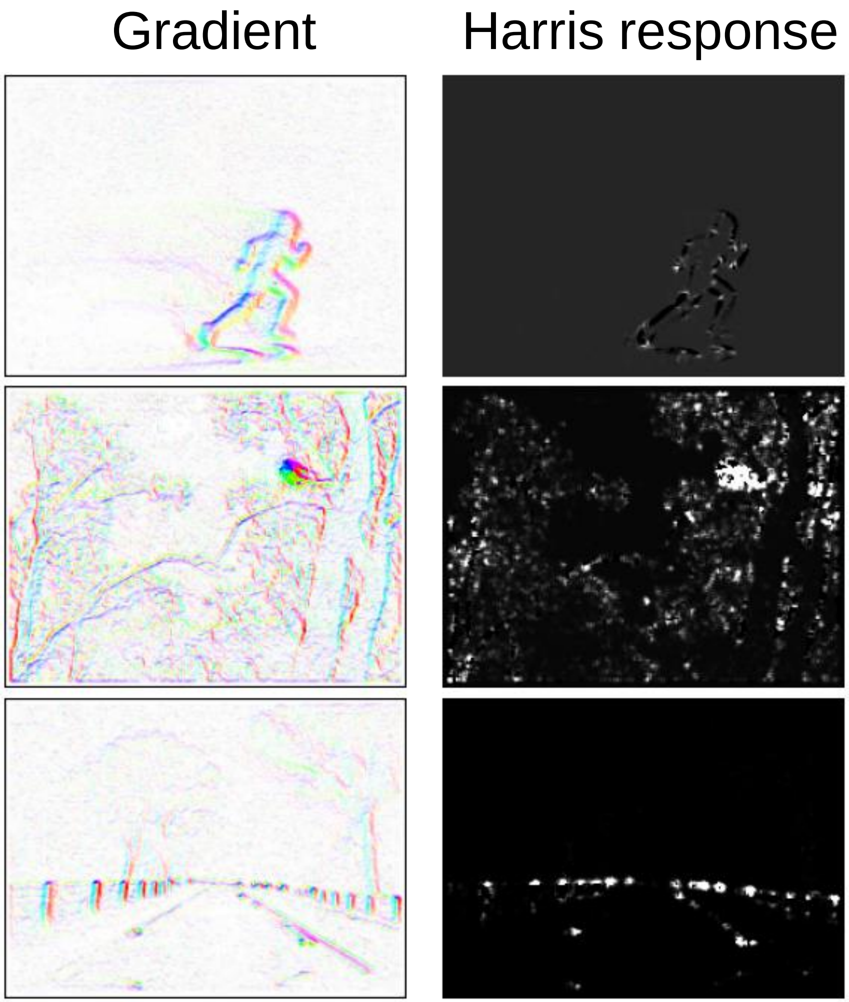

When an event arrives, the Harris response is only updated in a local neighbourhood.

C. Scheerlinck, N. Barnes, R. Mahony, “Computing Spatial Image Convolutions for Event Cameras,” arXiv, 2018.

Results

Gradient can be used as input to asynchronous Harris corner detector.

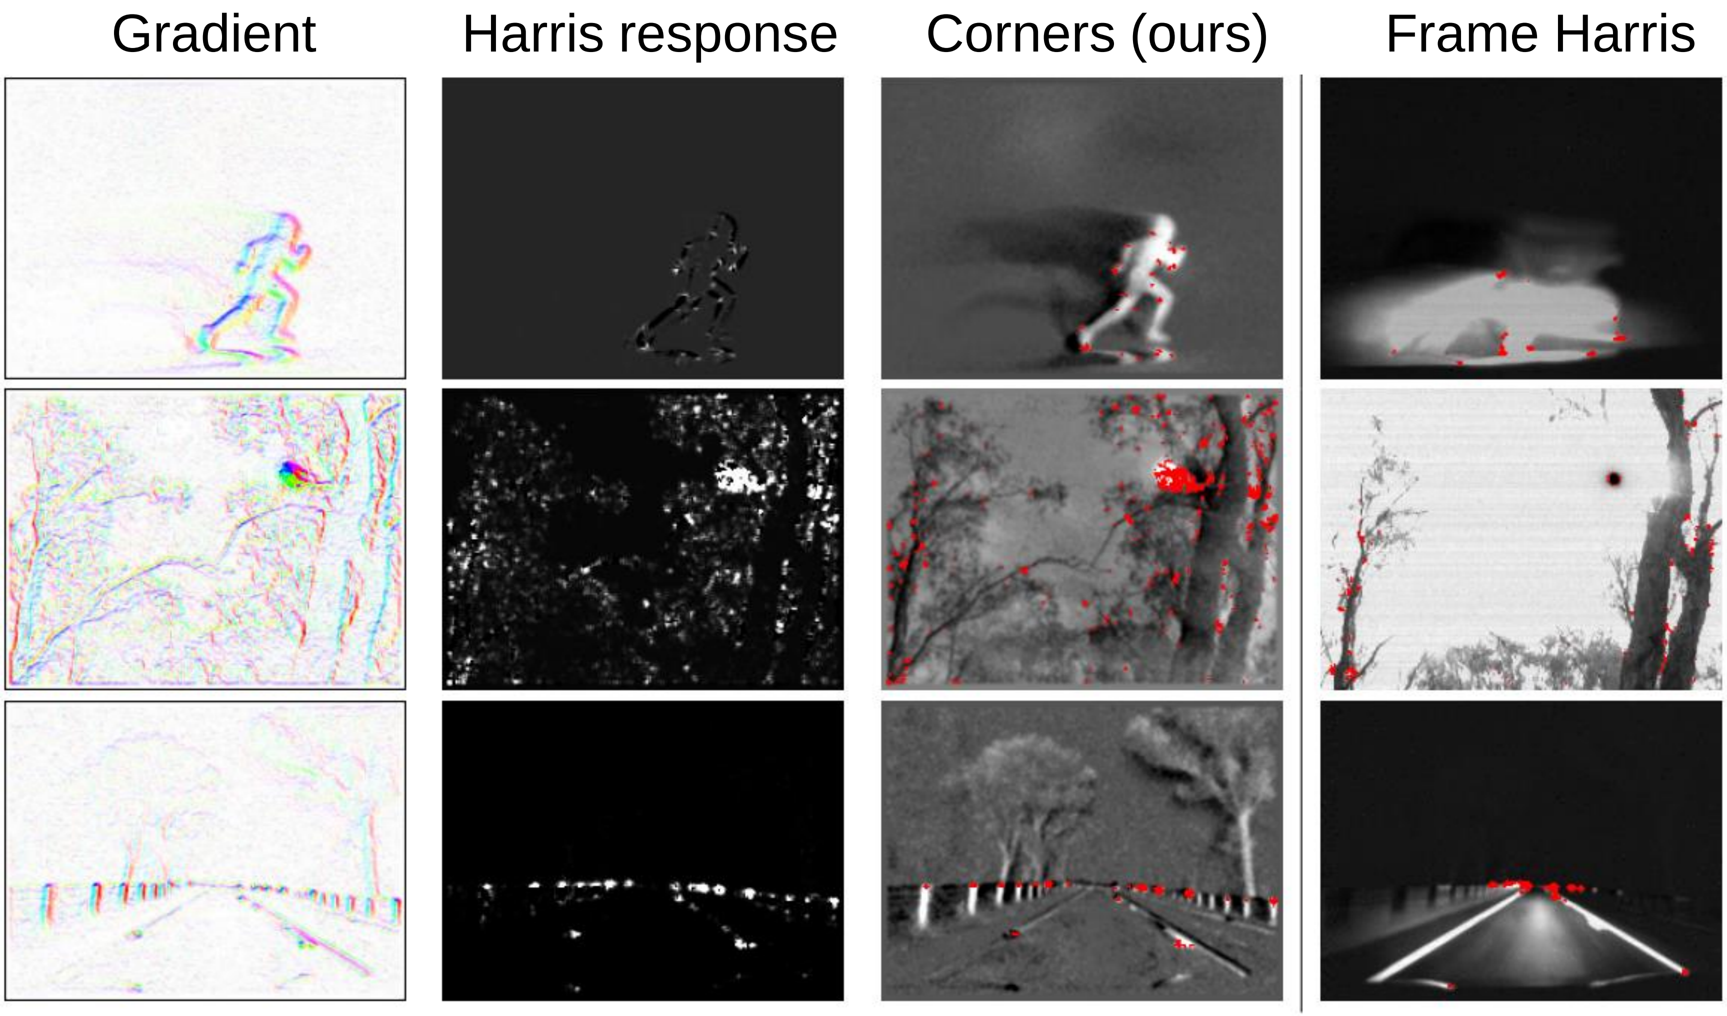

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.

Results

Comparison to state-of-the-art event-based corner detection.

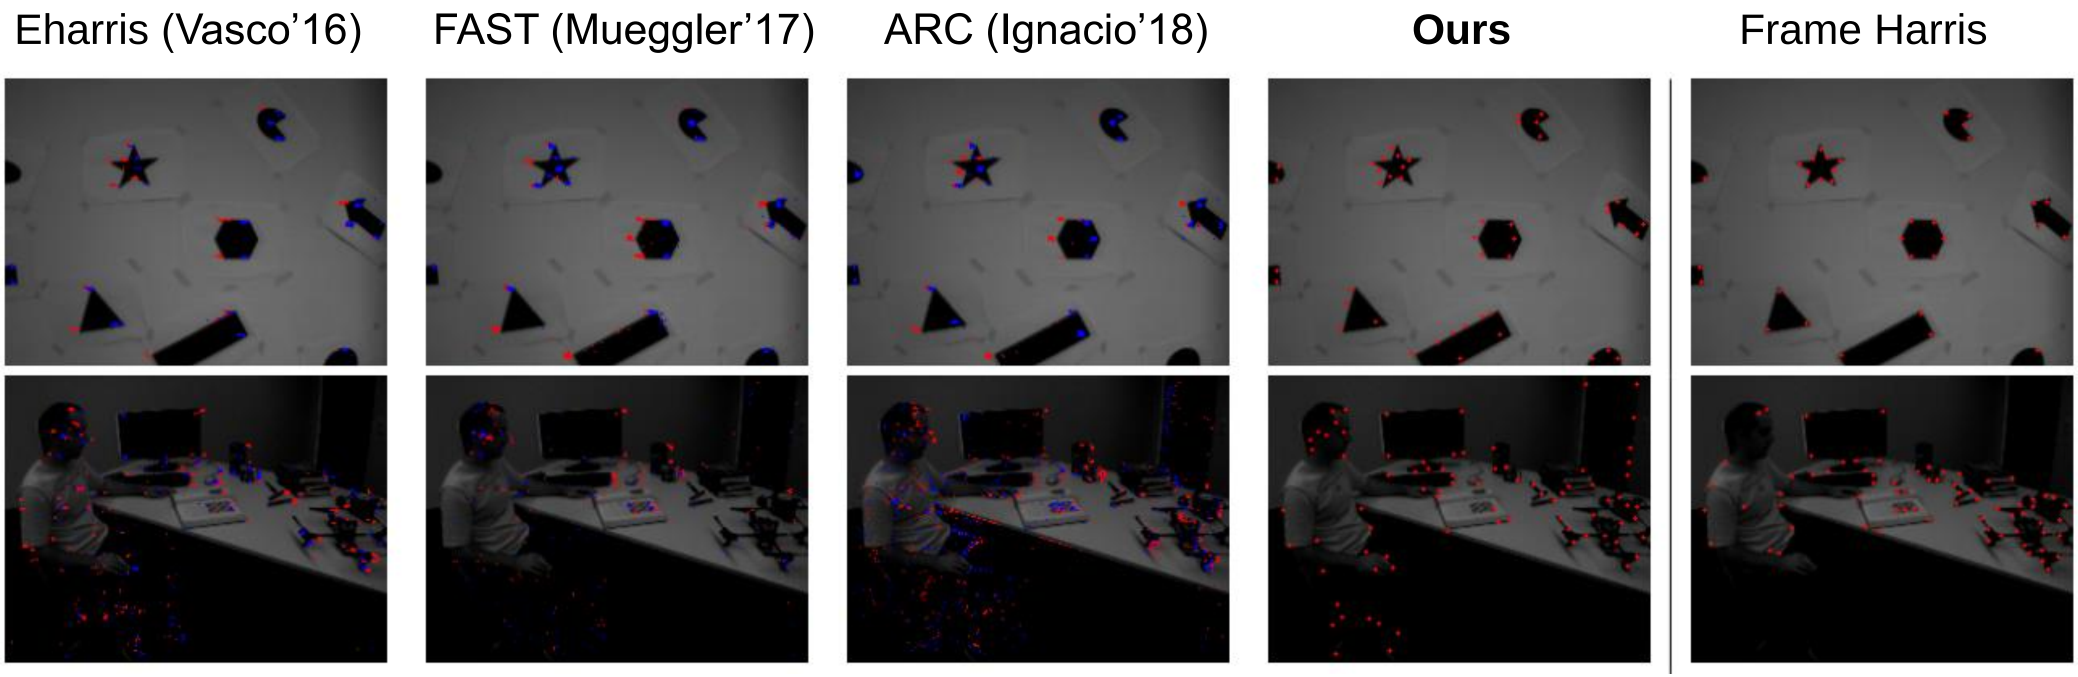

Since we output a continuous-time Harris response state, at any point in time we can apply non-maximum suppression to get clean corners.

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.

Conclusion

- We have introduced a methodology for event-based convolutions.
- Each event is individually convolved, producing a cluster of convolved events.
- Convolved events are fed into an asynchronous image reconstruction algorithm to produce a continuous-time state estimate of the convolved image.
- We introduce an asynchronous Harris corner detector based on gradients produced by our method.

C. Scheerlinck, N. Barnes, R. Mahony, "Computing Spatial Image Convolutions for Event Cameras," arXiv, 2018.