Towards a Non-Fourier Event-Based Opto-Electronic Convolution Accelerator

FOCUS LAB
UF UNIVERSITY of FLORIDA

Hannah Kirkland, Trung H. Le, Piper Taylor, Mehran Keivanimehr, Isaac J. Sledge, and Sanjeev J. Koppal

Office of Naval Research Science & Technology
NSF

Paper Project Page Poster

Abstract

The current advancement of neural networks has led to a resurgence of research in optical computing. Exploiting the physical properties of light to enable optical neural networks is advantageous due to optical computing's higher parallelism, speed, and power consumption compared to conventional compute methods. However, many state-of-the-art methods rely on coherent light and conventional image sensors or photodetectors. For the first time, we demonstrate using an event camera for reconfigurable optical compute, implemented via a neuromorphic camera, light source, two digital micromirror device (DMDs), and a computer. We describe the fundamental advantages and limitations of our device then report initial results using the device for two simple machine learning tasks.

Qualitative Comparison

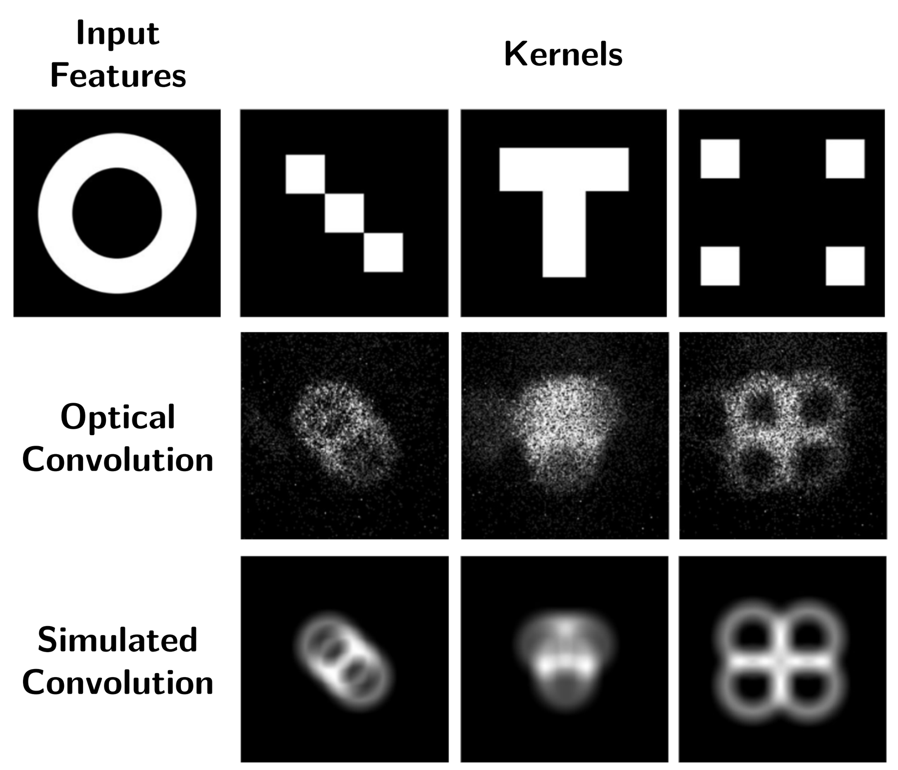

Pictured are example convolutional outputs for each given kernel and the "O"-shaped input features. The real optical convolutions were reconstructed from the captured event stream. The simulated convolutions were performed using the same kernel and input features on a conventional computer.

Kernel Recovery Task

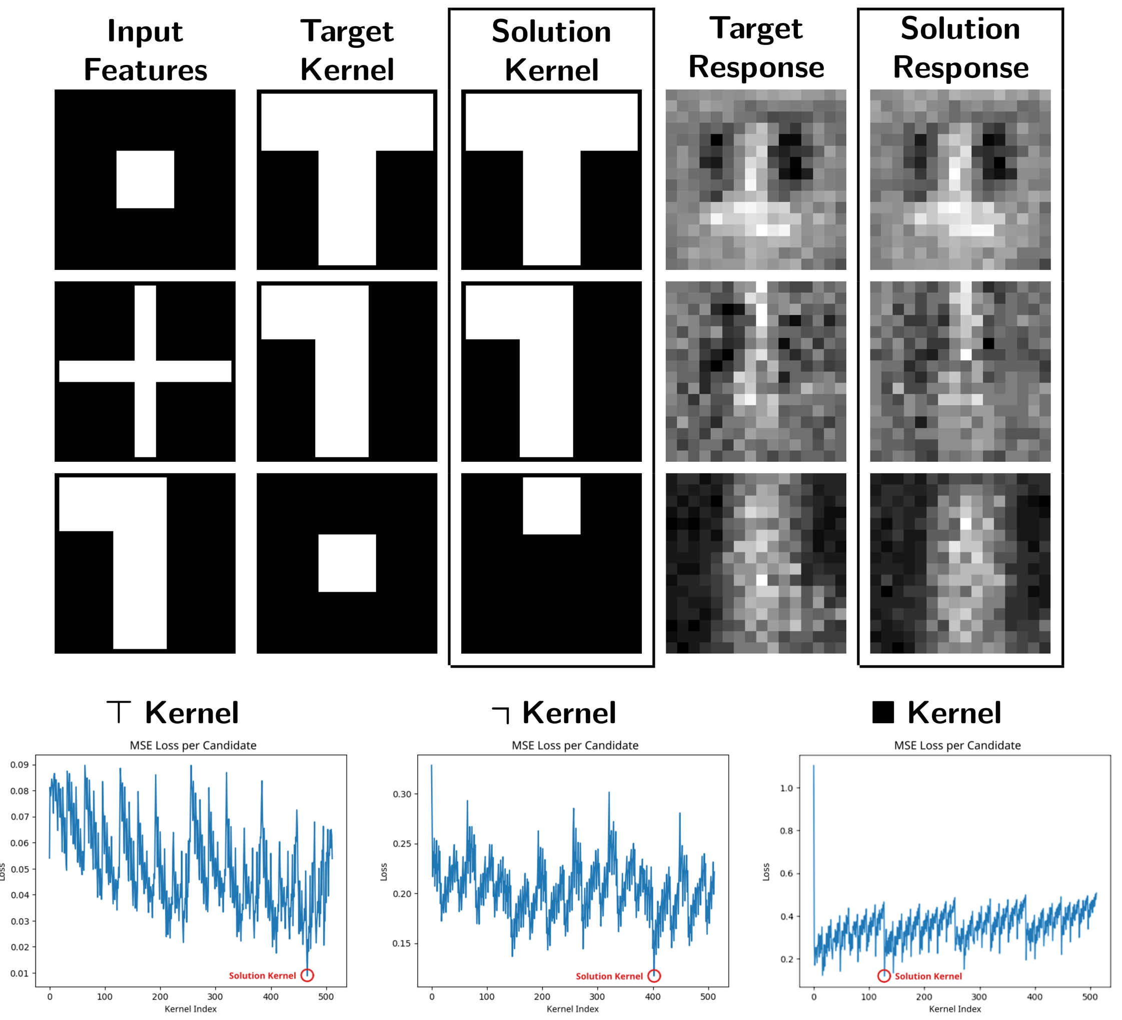

For the opto-electronic convolution accelerator to be useful, the captured response needs to be distinct for different inputs. To determine the functional ability of our system, we implemented a task whose goal of was to recover the kernel used to generate a chosen convolutional response.

First, we chose a 3×3 binary pattern for the kernel and a 16×16 binary pattern for the input features. The input features remained fixed, and each possible kernel was displayed. The resulting convolutional response was captured and compared to the target response. The lower graphs show the MSE loss for every candidate kernel for the task. In all but the last case, the correct kernel was recovered. The last case recovered a shifted result.

Conventional Convolutions

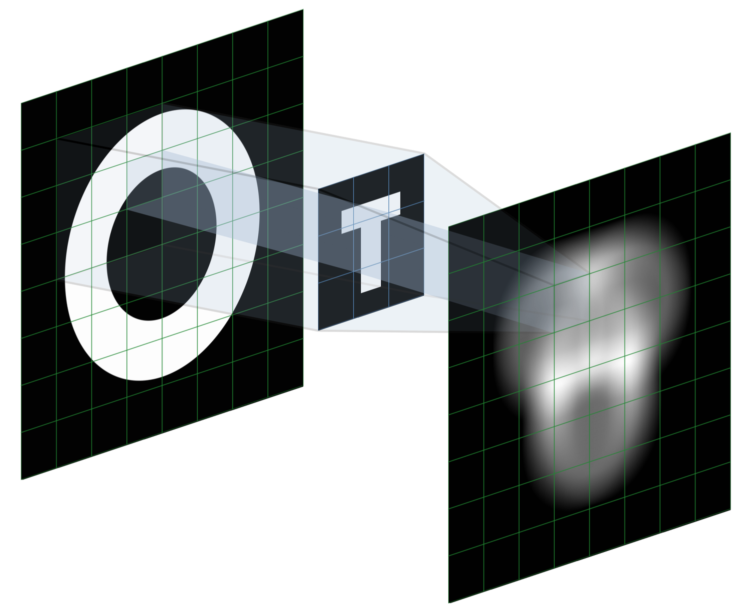

$ Resource Intensive · Compute Limitations · Conventional Computer

Optical Convolutions

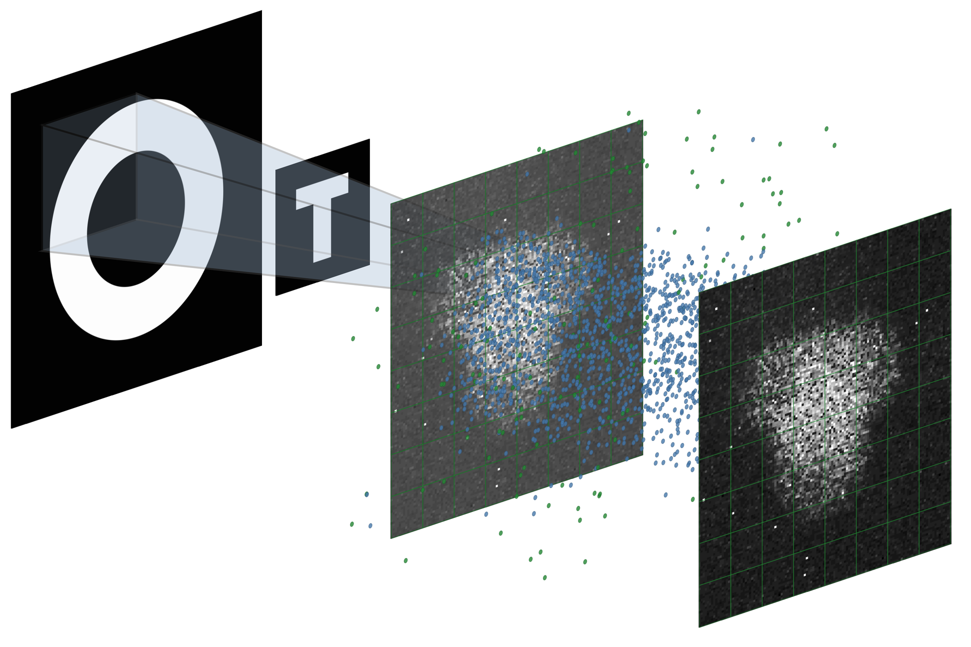

Incoherent · Reconfigurable · Speed of Light · Neuromorphic Camera

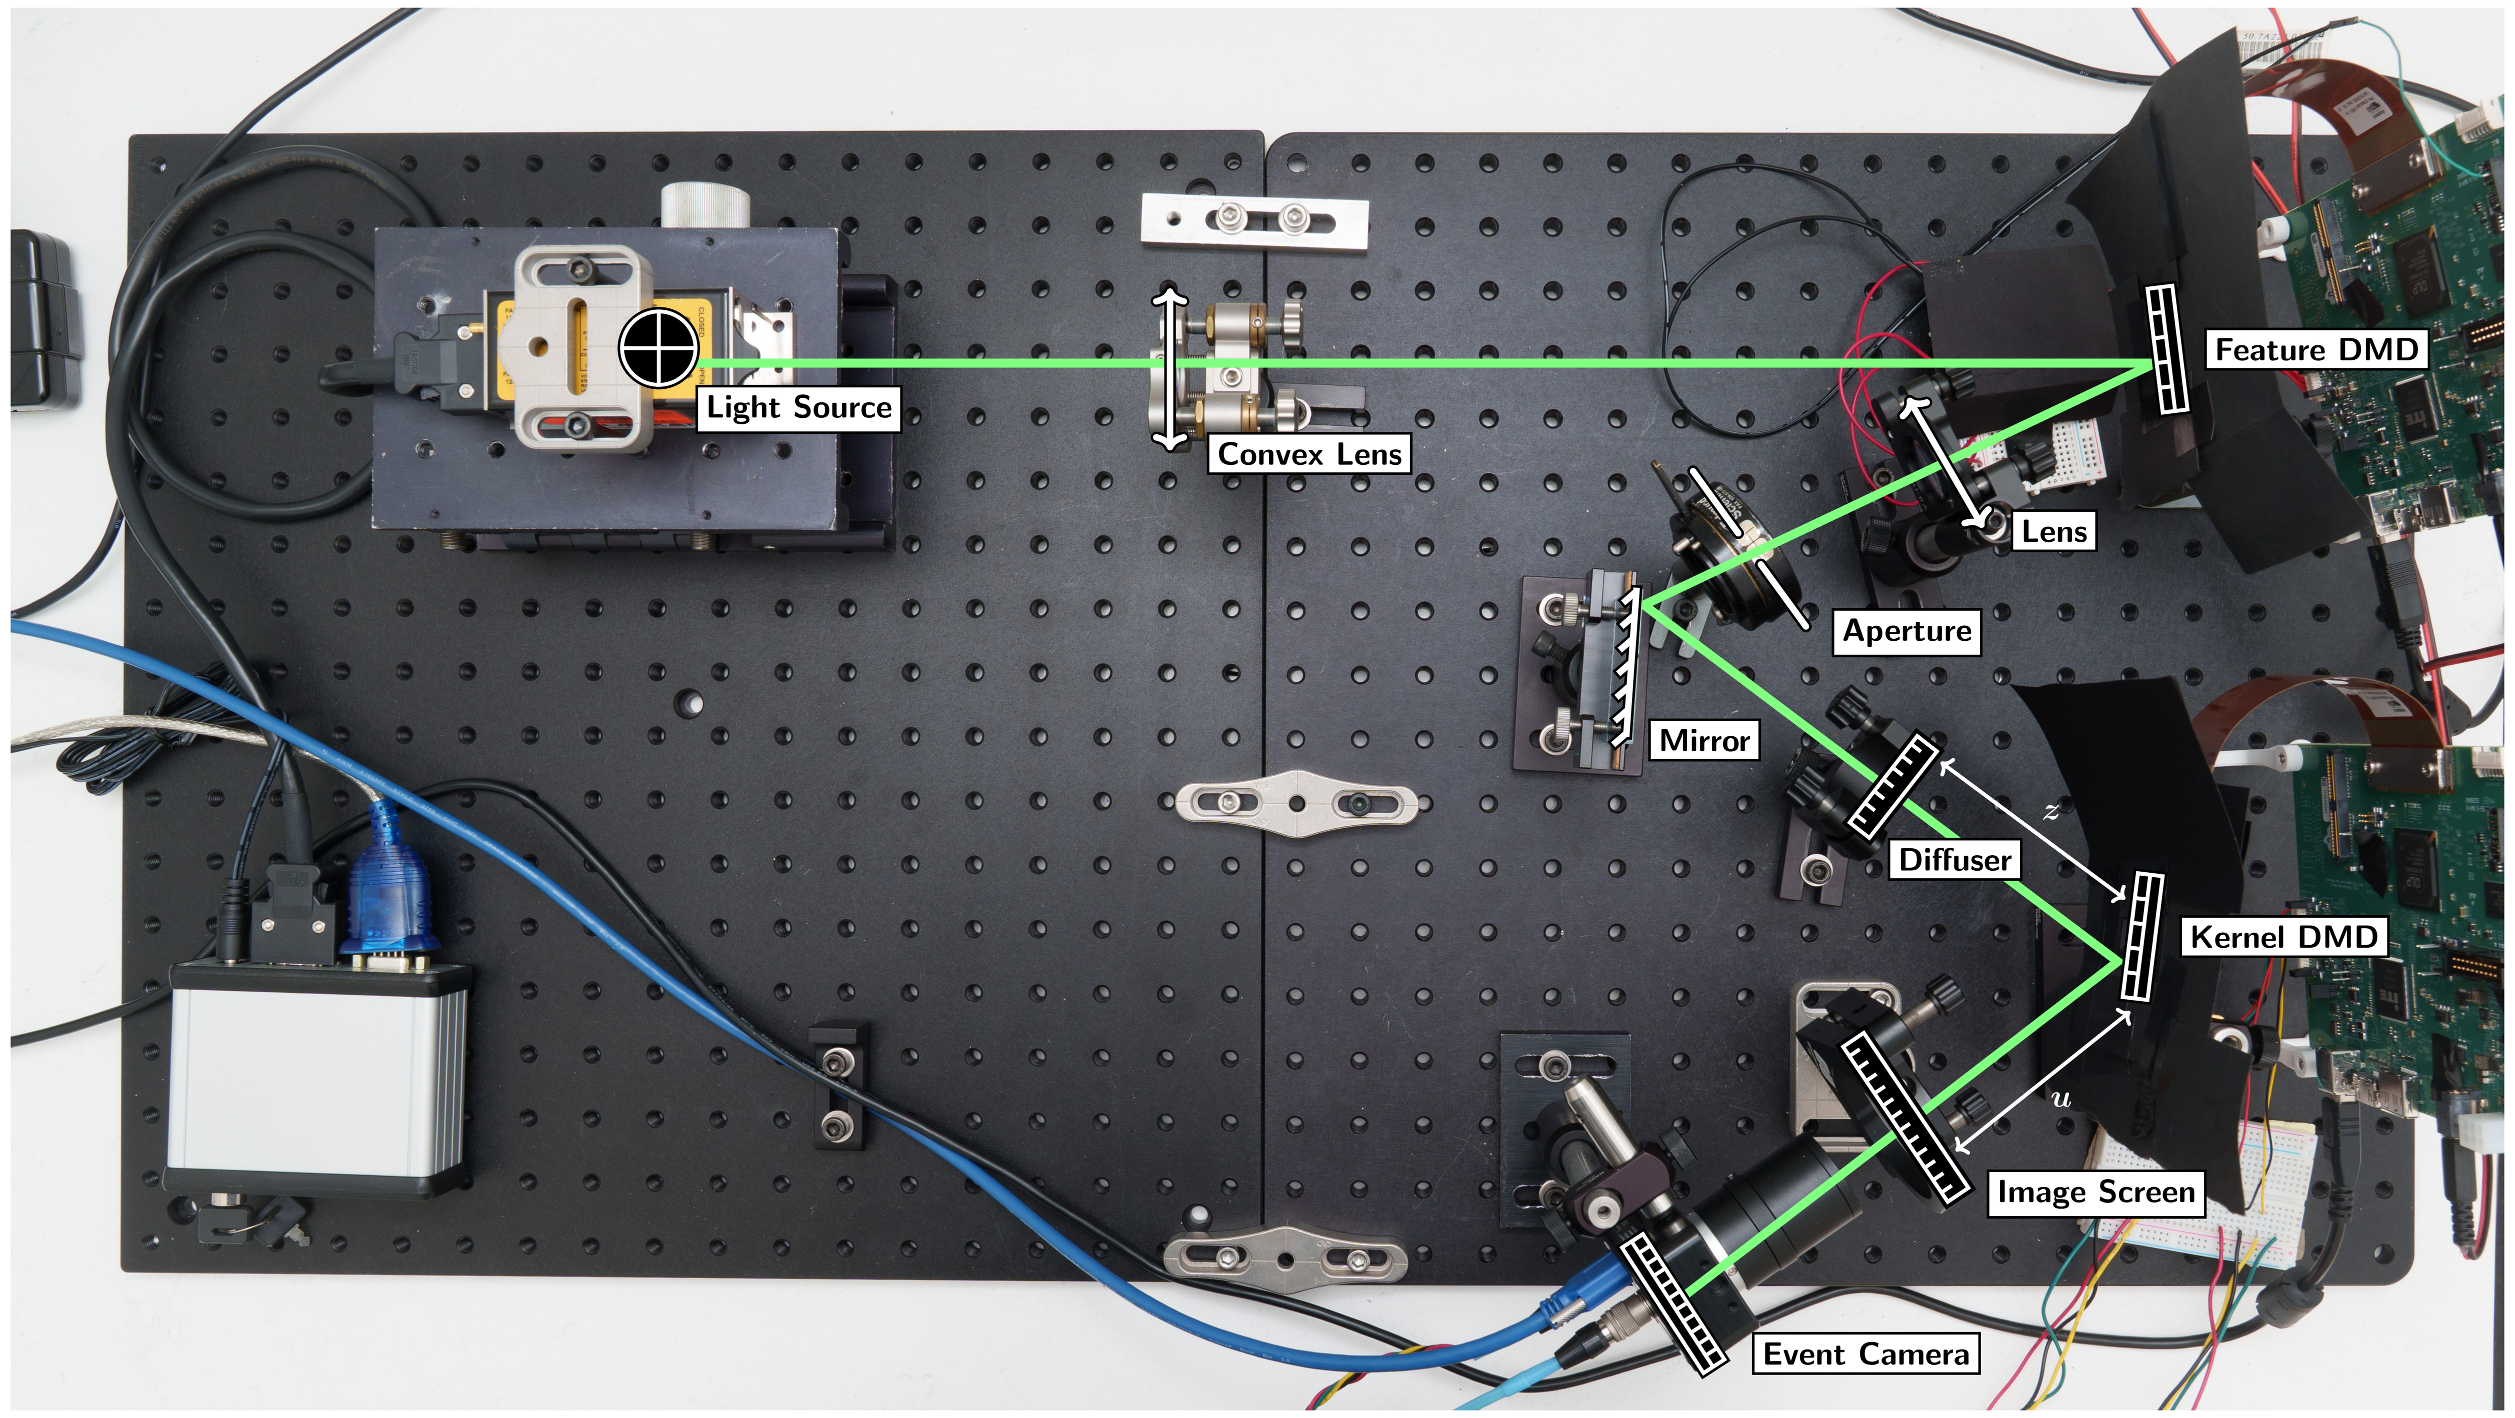

Binary Classification

Statistic	*Simulated Annealing* Conventional Computer	*Simulated Annealing* Opto-Electronic Device		*Random Walk*
Max	72.8%	66.7%	63.0%	56.4%
Min	30.3%	52.4%	46.7%	48.9%
Mean	49.9%	55.1%	54.6%	51.5%
St.D.	9.6%	12.1%	5.9%	3.2%
	temp = 10	*temp = 10*	*temp = 5*	

The real and simulated test accuracies for the three layer BCNN model optimized on the Face-Not-Face dataset. The model was optimized in an iterative, layerwise fashion using simulated annealing or random walk. For two classes, the random baseline is 50%.

Statistic	***Conventional Computer***
Mean	58.5±0.83%
Max	58.9%
	Opto-Electronic Device
Mean	55.8±2.24%
Max	58.6%
	Temperature = 10

The real and simulated test accuracies for the same model optimized on the EMNIST (letters vs numbers) dataset. The model was optimized in an iterative, layerwise fashion using simulated annealing. For two classes, the random baseline is 50%.