Panel Session
Impact of Deep Learning on Biometrics and Trends

BRIAN LOVELL (THE UNIVERSITY OF QUEENSLAND, AUSTRALIA) (CHAIR)

TERRY BOULT (UNIVERSITY OF COLORADO AT COLORADO SPRINGS, USA)

GANG HUA (MICROSOFT, USA)

XIAOMING LIU (MICHIGAN STATE UNIVERSITY, USA)

MAYANK VASTA (IIIT, DELHI, INDIA)

Why is it so popular?

- 2012-2014 Imagenet results:

CNN
non-CNN

2012 Teams	%error
Supervision (Toronto)	15.3
ISI (Tokyo)	26.1
VGG (Oxford)	26.9
XRCE/INRIA	27.0
UvA (Amsterdam)	29.6
INRIA/LEAR	33.4

2013 Teams	%error
Clarifai (NYU spinoff)	11.7
NUS (singapore)	12.9
Zeiler-Fergus (NYU)	13.5
A. Howard	13.5
OverFeat (NYU)	14.1
UvA (Amsterdam)	14.2
Adobe	15.2
VGG (Oxford)	15.2
VGG (Oxford)	23.0

2014 Teams	%error
GoogLeNet	6.6
VGG (Oxford)	7.3
MSRA	8.0
A. Howard	8.1
DeeperVision	9.5
NUS-BST	9.7
TTIC-ECP	10.2
XYZ	11.2
UvA	12.1

- 2015 results: MSRA under **3.5%** error. (using a CNN with 150 layers!)

figures from Yann LeCun's CVPR'15 plenary

Similar results for Biometrics.
Virtually all new algorithms
Exploit Deep learning
Because the performance
is simply so good.

Either you write a DL paper or
the reviewers will request
you compare your method
to DL. When you compare,
DL invariably wins with ease.

~~Hand Crafted Features~~ Learned Features

- Convolution Layer
  - A set of convolutional kernels

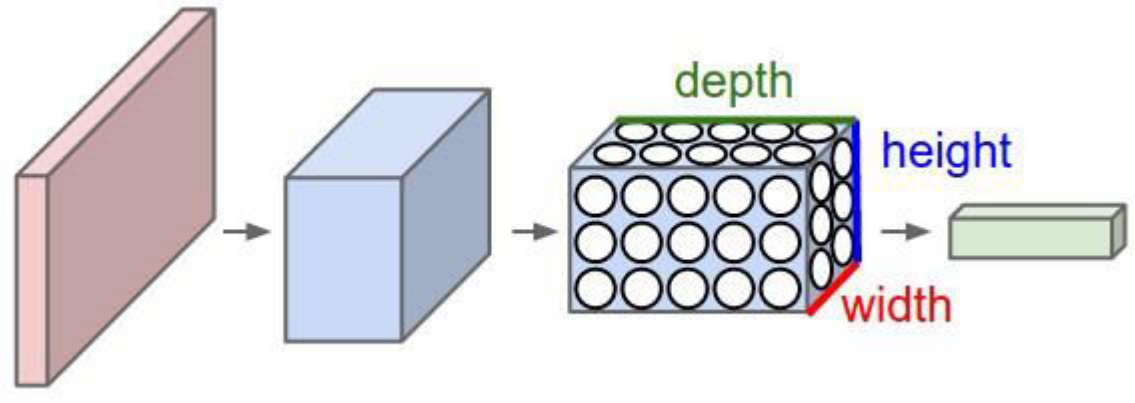

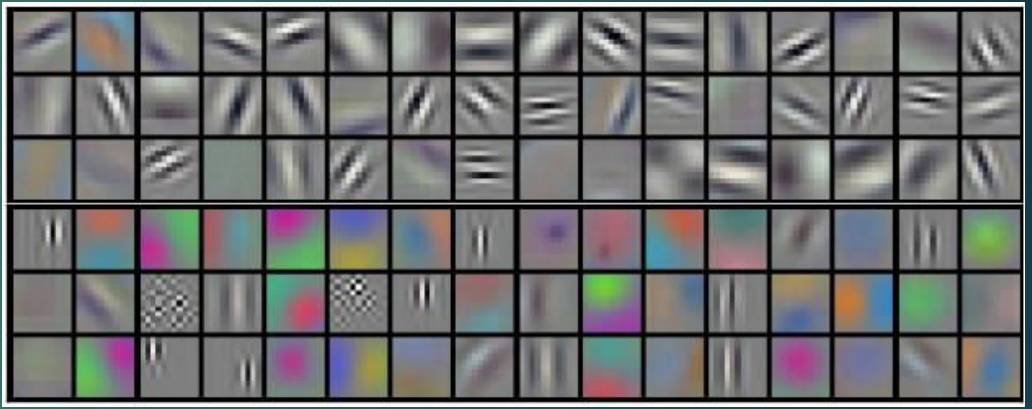

CVPR hot topics (Arnold Wiliem)

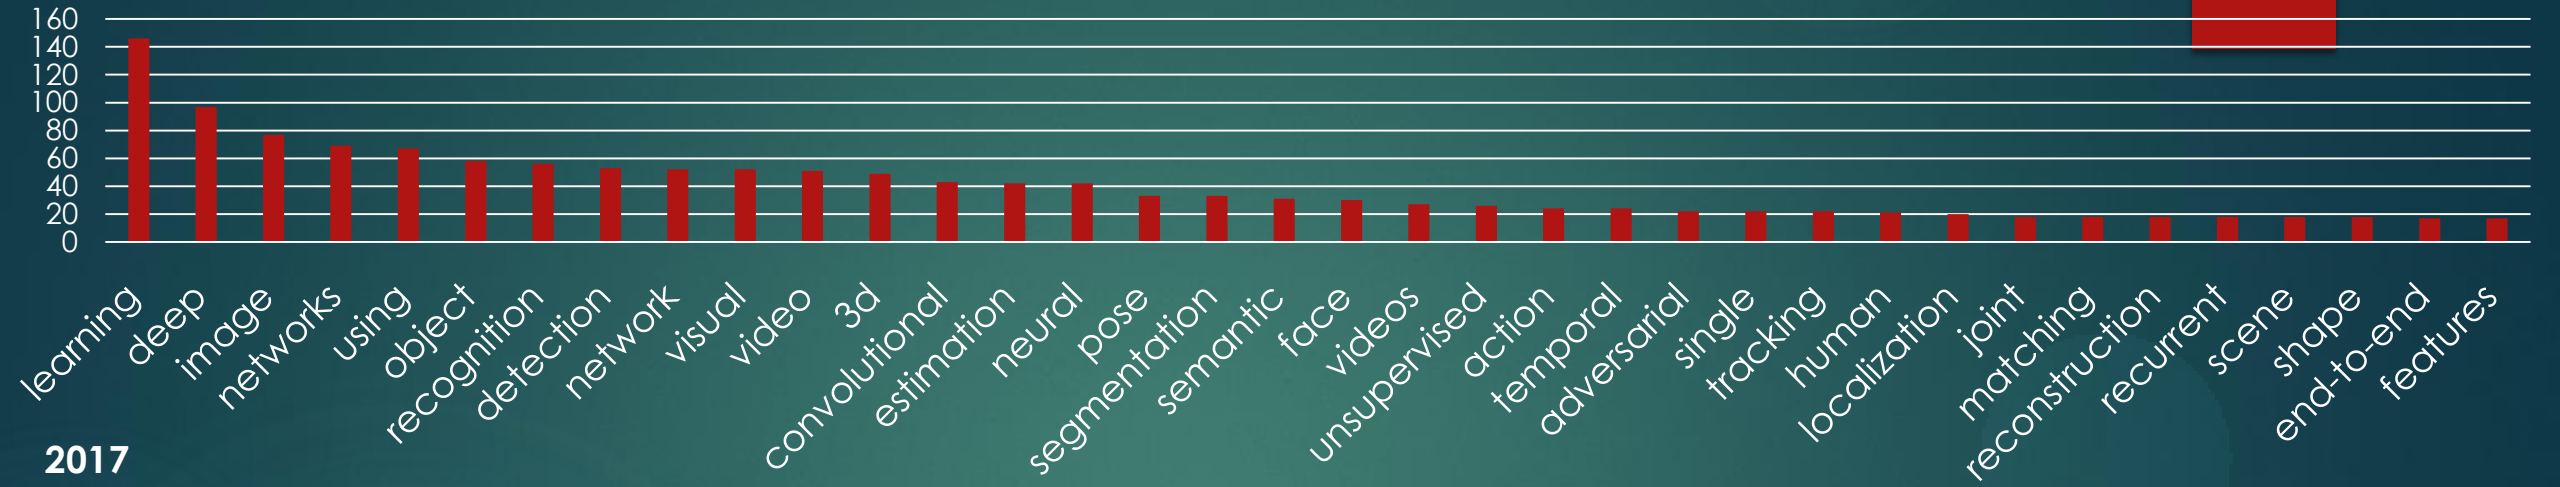

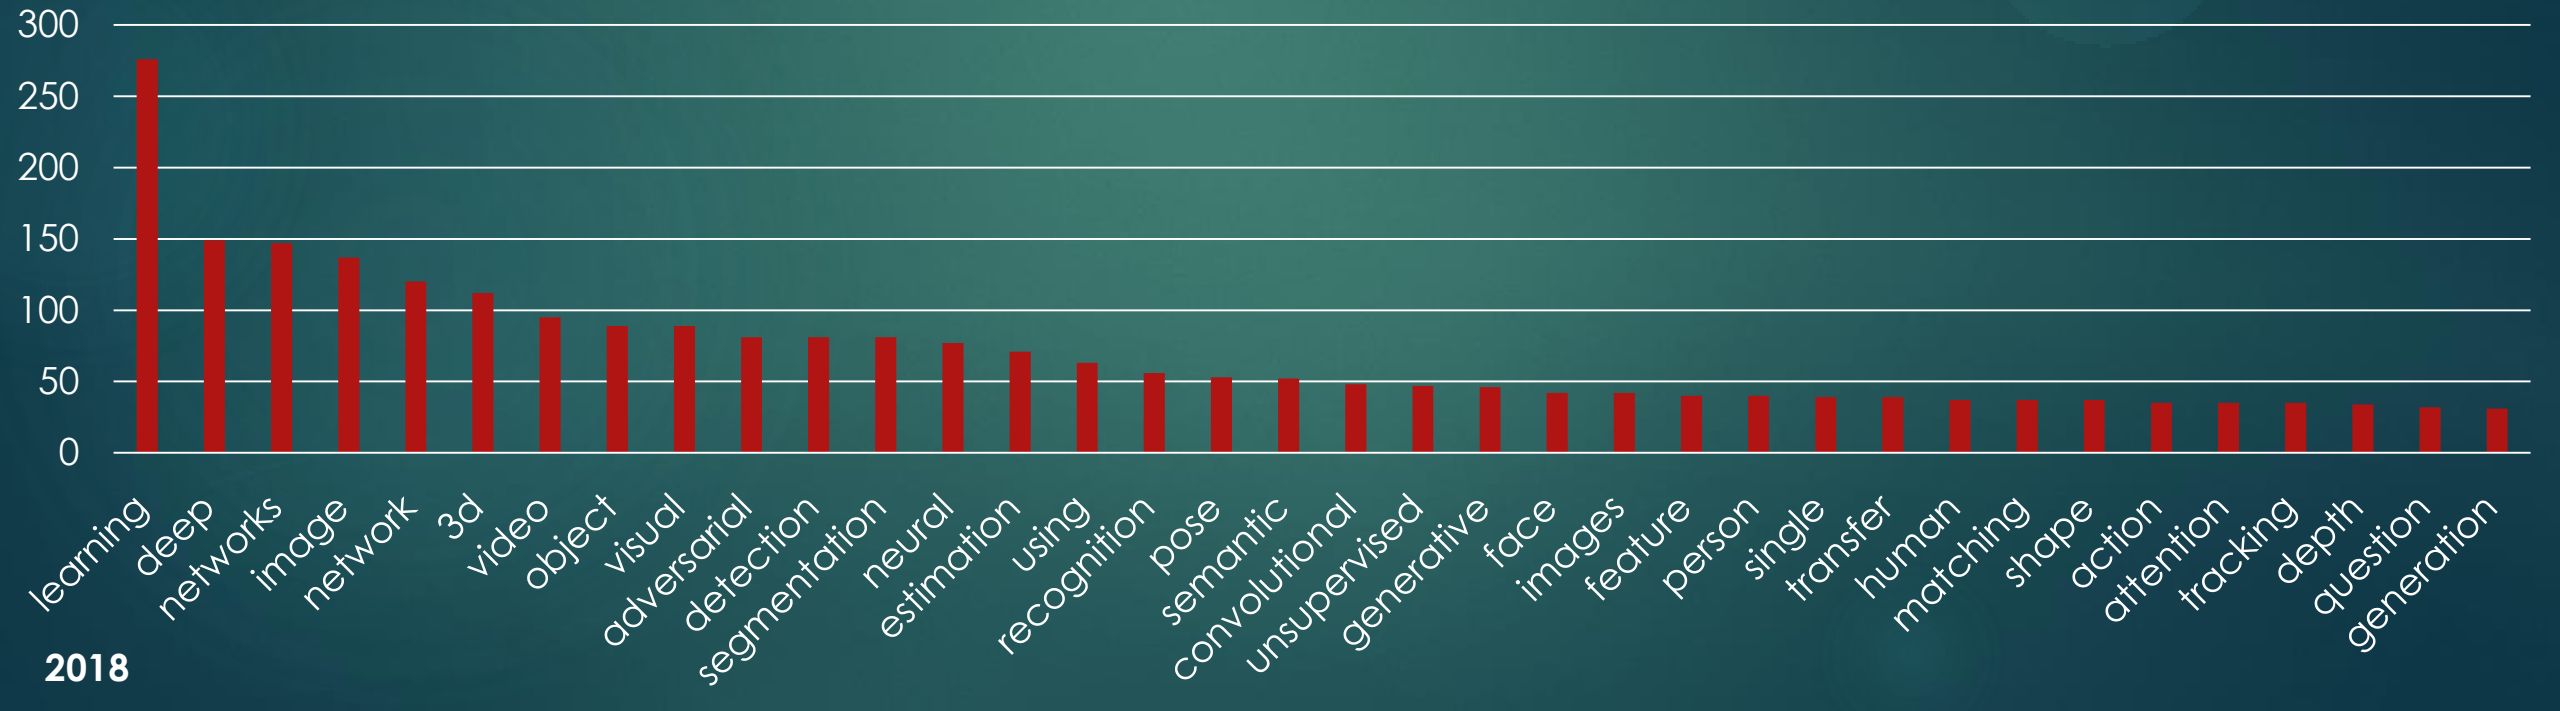

Any non-deep learning paper in CVPR2018?

- Fairly difficult to spot
- BUT found one non-DL paper!

Sensor paper

This paper has been accepted for publication at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, 2018. ©IEEE

A Unifying Contrast Maximization Framework for Event Cameras, with Applications to Motion, Depth, and Optical Flow Estimation

Guillermo Gallego, Henri Rebecq, Davide Scaramuzza
Dept. of Informatics and Neuroinformatics, University of Zurich and ETH Zurich

Abstract

We present a unifying framework to solve several computer vision problems with event cameras: motion, depth and optical flow estimation. The main idea of our framework is to find the point trajectories on the image plane that are best aligned with the event data by maximizing an objective function: the contrast of an image of warped events. Our method implicitly handles data association between the events, and therefore, does not rely on additional appearance information about the scene. In addition to accurately recovering the motion parameters of the problem, our framework produces motion-corrected edge-like images with high dynamic range that can be used for further scene analysis. The proposed method is not only simple, but more importantly, it is, to the best of our knowledge, the first

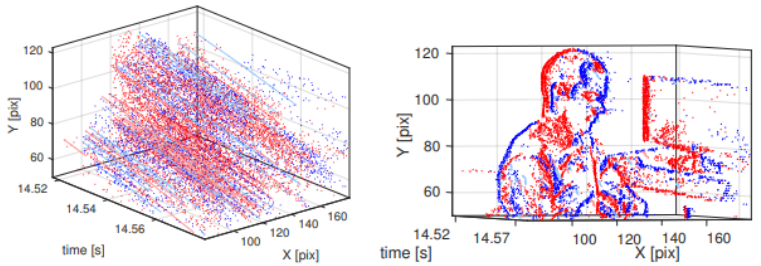

(a) (b)

Figure 1: (a) Events (dots) caused by a moving edge pattern and point trajectories in a space-time region of the image plane, colored according to event polarity (blue: positive event, i.e., brightness increase; red: negative event, i.e., brightness decrease). (b) Visualization of the events along the direction of the *point trajectories* highlighted in (a); corresponding events line up, revealing the edge pattern that caused them. Our approach works by maximiz-

What is the future for Biometrics Research?

- Deep learning has eaten our research field
- CVPR, ICML, NIPS, BTAS, and ICB are becoming indistinguishable
- Accuracy of Biometrics is improving in leaps and bounds through DL
- How do we write a really good paper that is not simply an application of DL?
- Math was King and Data was the Queen
- Now Data is King and Math is the Queen
- Companies like Google, Nvidia, and Microsoft are dominating scientific conferences
  - E.g. NVIDIA has 14 scientific papers in CVPR2018 – like a major University