

Deep Learning with Spherical Data


Min Sun

National Tsing Hua University

Vision Science Lab

<http://aliensunmin.github.io/>

Self-Intro



Assistant Professor in Electrical Engineering
at National Tsing Hua University (Sept. 2014)
[CV],[Google Scholar]



Postdoctoral Researcher with Steve Seitz and Ali Farhadi
in CSE at University of Washington (Jan. 2013 - Aug. 2014)



Sep 09-Dec 12



Jul 11-Sep 11



Jul 09-Sep 09



Sep 07-Jul 09



Sep 05-Sep



Sep 99-Jul 03

07
遠見專訪 <https://www.gvm.com.tw/article.html?id=43046>


Large-scale video analysis

Video Title Generation




ECCV 2016

Video QA



AAAI 2017

Visual Forecasting




ICCV 2017

ACCV 2016 Best Paper Award



Car Accident Anticipation


CVPR 2017



Risk Assessment

Handcam

Object Recognition via Handcam




ECCV 2016

Grasp Assistant for BVI



Under Submission 2018


ICCV 2017



Intention Anticipation


360 (Spherical) Vision

Deep Pilot




CVPR 2017

Saliency Prediction




CVPR 2018

AAAI 2018




360 Grounding


Robotic Vision



Learn from Demonstration




Indoor Place Recognition & Navigation
ICRA 2018



Robust Control

Health Care

Utilizing Deep Neural Network to Simulate Radiotherapy Planning in Esophageal Cancer Patients






Outline

- Introduction to Deep Learning
- Spherical data
- Deep Learning for Spherical data

Deep Learning (DL)

- Data
- GPU Computing
- Talents



Data:

IMAGENET

- 開始於 2007 @ Princeton
- 初登場於 2009 @ CVPR
- 照片停止搜集於 2010
 - 總共類別：21841
 - 總共圖片：1千4百萬
- ILSVR Challenge 從2010到現今



Jia Deng



Fei-Fei Li

Info from <http://www.image-net.org/>

1K Image Classification

Label = $f(\text{Image})$

Deep Learning

深度學習






Figure from Olga Russakovsky ECCV'14 workshop

GPU: NVIDIA CUDA

The Difference between a CPU and GPU



CPU



GPU

Tesla P100 With Over 20 TFLOPS Of FP16

Read more: <http://wccfttech.com/nvidia-pascal-gpu-gtc-2016/#ixzz456KT75Jf>

Facebook Machine Learning Server



Big Sur: 8 NVIDIA M40


<https://techcrunch.com/gallery/a-look-inside-facebooks-data-center>

Talents: DNNresearch acquired by Google



Geoffrey Hinton (right: Professor) Alex Krizhevsky (middle; PhD student), and Ilya Sutskever (left; Postdoc)


Deep Learning: Convolutional Neural Network (CNN)



$$\text{Character} = f(\text{Image})$$

What is Deep?


2 Convolution layers



A story in computer Vision!

5 Convolution layers


[LENGUET, LeCun 1980]



AlexNet, 2012, Krizhevsky et al


Automatic Speech Recognition (ASR)

- DNN has become the-state-of-the-art for speech recognition since 2010
 - ✓ Large volume data available
 - ✓ Deep learning is changing everything. AM, LM...
 - ✓ ASR is going to practice. Siri, google, baidu, Tencent...



DBN-HMM

Microsoft: 33% Recognition Error Reduction over discriminatively trained GMM-HMMs



F. Seide, Gang Li, Dong Yu, Conversational Speech Transcription Using Context-Dependent Deep Neural Networks, Interspeech 2011, Florence, Italy

Natural Language Processing (NLP)


- “Deep Learning waves have lapped at the shores of computational linguistics for several years now, but 2015 seems like the year when the full force of the tsunami hit the major Natural Language Processing (NLP) conferences.” -Dr. Christopher D. Manning, Dec 2015

Recurrent Neural Network

Probability of
“arrive” in each slot

Probability of
“**Taipei**” in each slot


Probability of
“on” in each slot



Outline


- Introduction to Deep Learning
- Spherical data
- Deep Learning for Spherical data

Data – Weather




Data – Earthquake


2018/02/6
花蓮大地震



2018/04/13



Data – Air-quality



<https://airbox.edimaxcloud.com/>

Data

- Anything else?
 - **Please let me know**




Outline

- Introduction to Deep Learning
- Spherical data
- Deep Learning for Spherical data

Talent

- Spherical CNNs. Taco Cohen · Mario Geiger · Jonas Koehler · Max Welling. In ICLR 2018 (**Best Paper**)



<https://openreview.net/forum?id=Hkbd5xZRb>




Figure 1: Any planar projection of a spherical signal will result in distortions. Rotation of a spherical signal cannot be emulated by translation of its planar projection.

Talent

- Spherical CNNs. Taco Cohen · Mario Geiger · Jonas Koehler · Max Welling. In ICLR 2018 (**Best Paper**)




Figure 5: The ray is cast from the surface of the sphere towards the origin. The first intersection with the model gives the values of the signal. The two images of the right represent two spherical signals in (α, β) coordinates. They contain respectively the distance from the sphere and the cosine of the ray with the normal of the model. The red dot corresponds to the pixel set by the red line.