About the Course

Kuan-Yu Chen (陳冠宇)

2018/03/01 @ TR-409, NTUST

Introduction to NLP

- Actually, this course is named "Deep Learning for Natural Language Processing"
  - Mainly focus on deep learning
- The loading is very heavy!
  - Algorithms + Mathematics + Coding
- After the course, I hope you are familiar with conventional deep learning strategies

AI in Taiwan

自由時報 Liberty Times Net

即時新聞 報紙總覽 影音 娛樂 汽車 時尚 體育 3C 評論

張忠謀：台積電未來將投入人工智慧（AI）

台積電董事長張忠謀，（圖由台積電提供）

2017-11-05 07:48

鴻海找上吳恩達Landing.ai合作，5年100億衝工業互聯網AI應用

by 葉紀圃 2018.02.02

你有 AI 思維卻沒有 AI 步驟？ Appier 從資料科學家到人工智慧應用驅動產業革新

2017/11/1 • Anny • appier、人工智慧、機器學習、行銷、跨螢技術

▲Appier 執行長暨共同創辦人游直翰於今 年度論壇開場

總覽 熱門 要聞 娛樂 國際 財經 副刊 體育 地產 論壇與專

Zenbo向前衝 華碩機器人事業群獨立

謝明傑領軍 11月中國開賣 明年初在美上市

7784 建立時間：2017/10/11

Zenbo是華碩董事長施崇棠(圖)心血結晶，機器人部門獨立後將由副總裁謝明傑領軍

ADATA

產品介紹 產品規格 新聞影音 支援服務

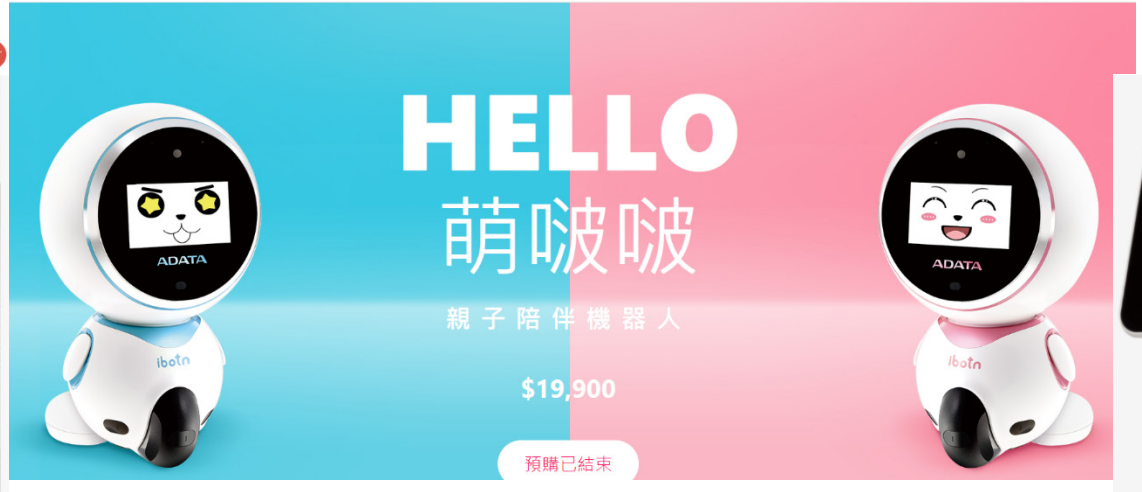

Beyond Human! – 1

- The ImageNet
  - 1000 classes
  - 1,431,167 images

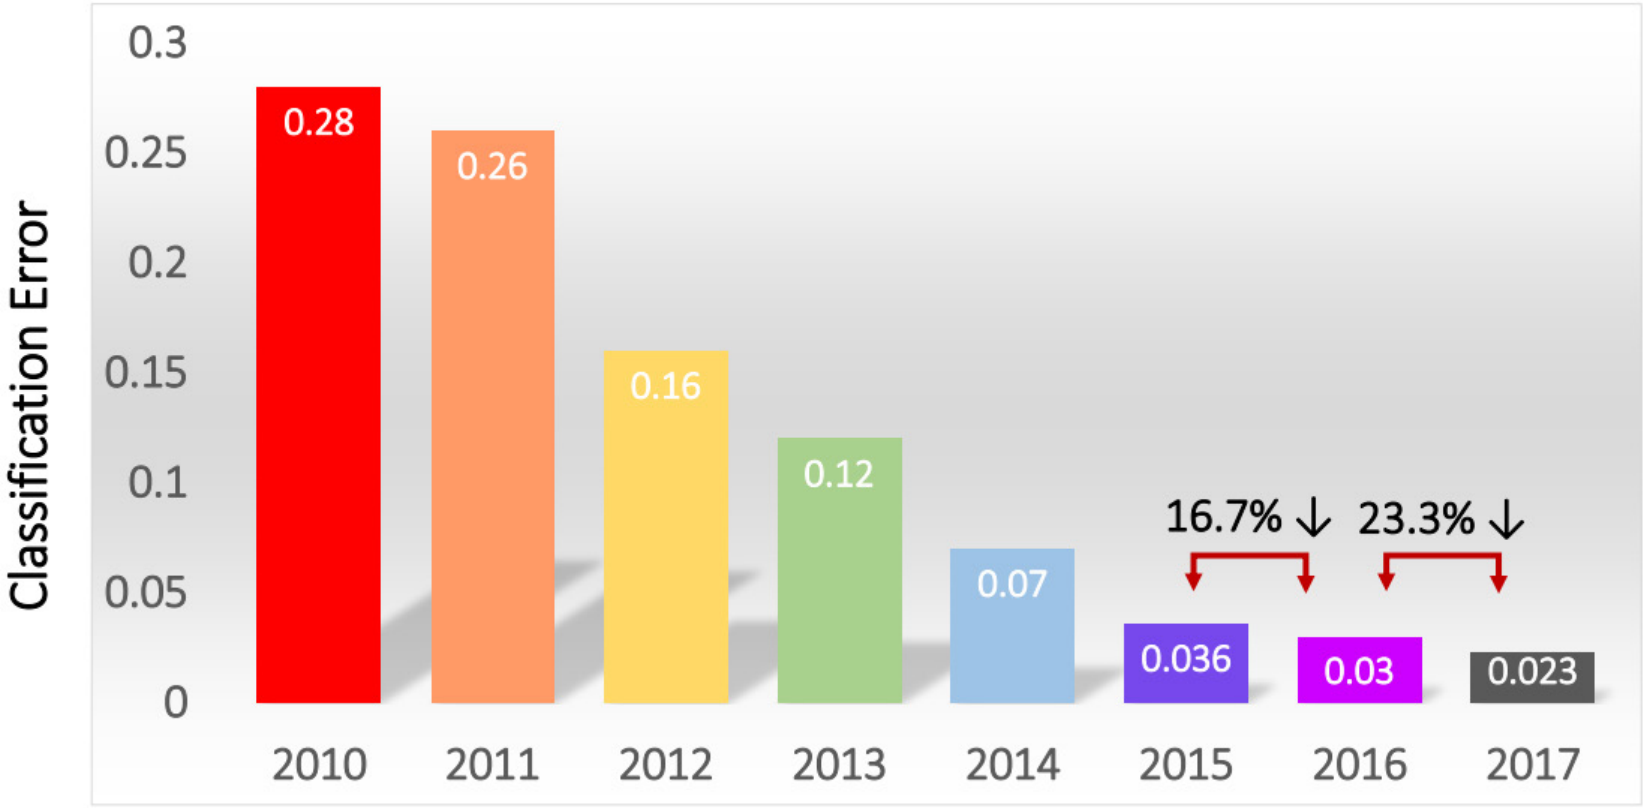

Beyond Human! – 2

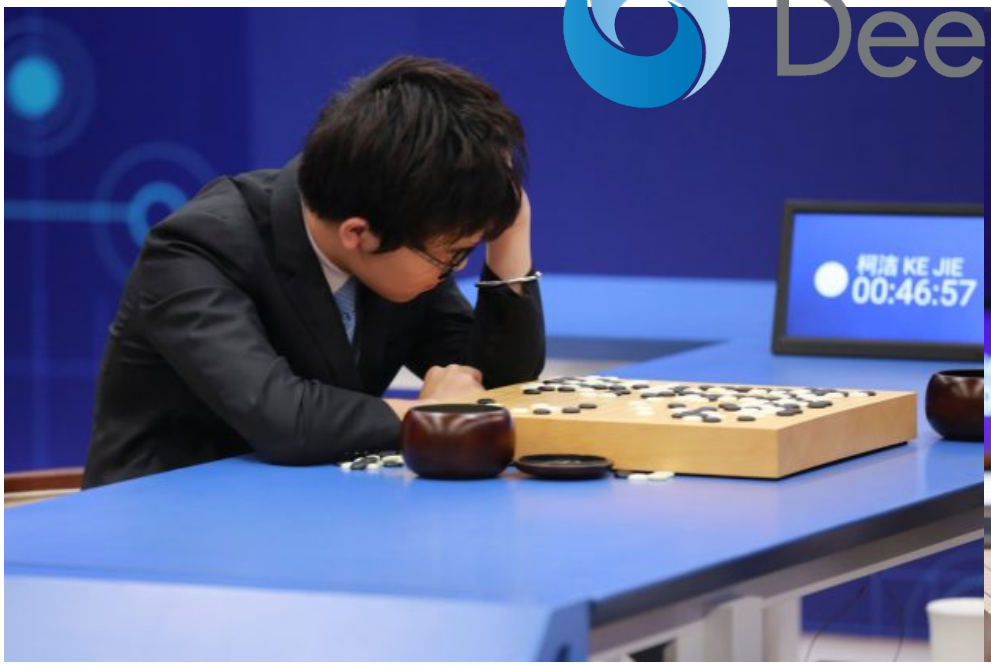

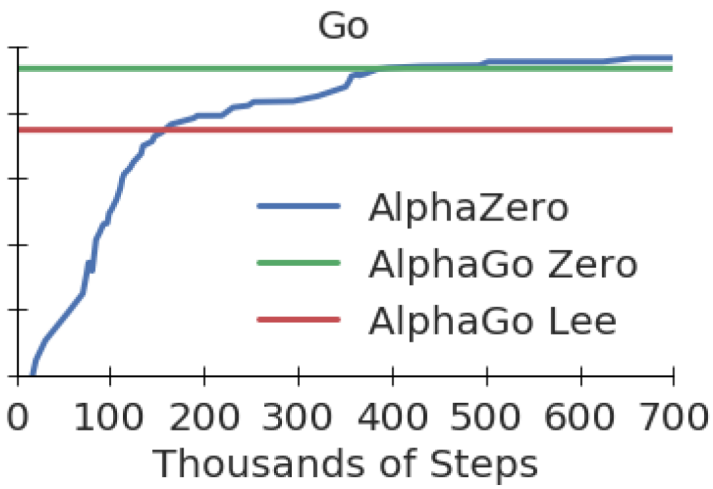

Beyond Human! – 3

SQuAD
The Stanford Question Answering Dataset

Rank	Model	EM	F1
	Human Performance *Stanford University* (Rajpurkar et al. '16)	82.304	91.221
1 Jan 22, 2018	Hybrid AoA Reader (ensemble) *Joint Laboratory of HIT and iFLYTEK Research*	82.482	**89.281**
1 Feb 19, 2018	Reinforced Mnemonic Reader + A2D (ensemble model) *Microsoft Research Asia & NUDT*	**82.849**	88.764
2 Feb 02, 2018	Reinforced Mnemonic Reader (ensemble model) *NUDT and Fudan University* https://arxiv.org/abs/1705.02798	82.283	88.533
2 Jan 03, 2018	r-net+ (ensemble) *Microsoft Research Asia*	82.650	88.493
2 Jan 05, 2018	SLQA+ (ensemble) *Alibaba iDST NLP*	82.440	88.607
3 Dec 17, 2017	r-net (ensemble) *Microsoft Research Asia* http://aka.ms/rnet	82.136	88.126

Beyond Human! – 4

Tero Karras, Timo Aila, Samuli Laine, Jaakko Lehtinen, ”Progressive Growing of GANs for Improved Quality, Stability, and Variation,” in *ICLR*, 2018.
https://github.com/tkarras/progressive_growing_of_gans

Joint the Trend!

- You can build your own architectures easily and flexibly

- Do not need to take care about the mathematics
  - Gradient is computed by calling a function!

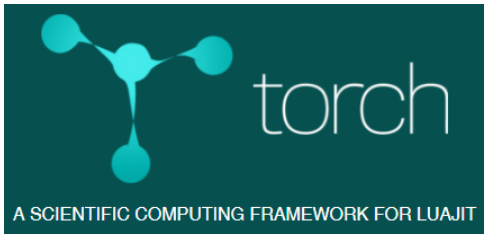

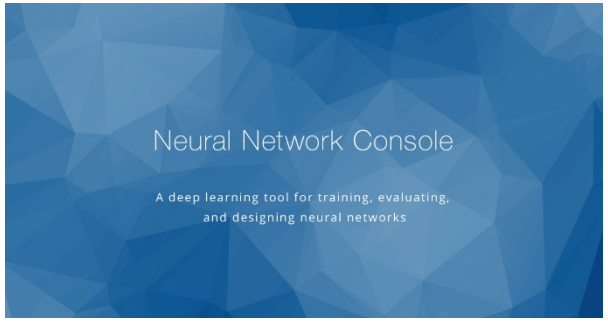

One Framework to Rule Them All

Google:
TensorFlow

"One framework to rule them all"

Facebook:
PyTorch +Caffe2

Research

Production

GPU and TPU – 1

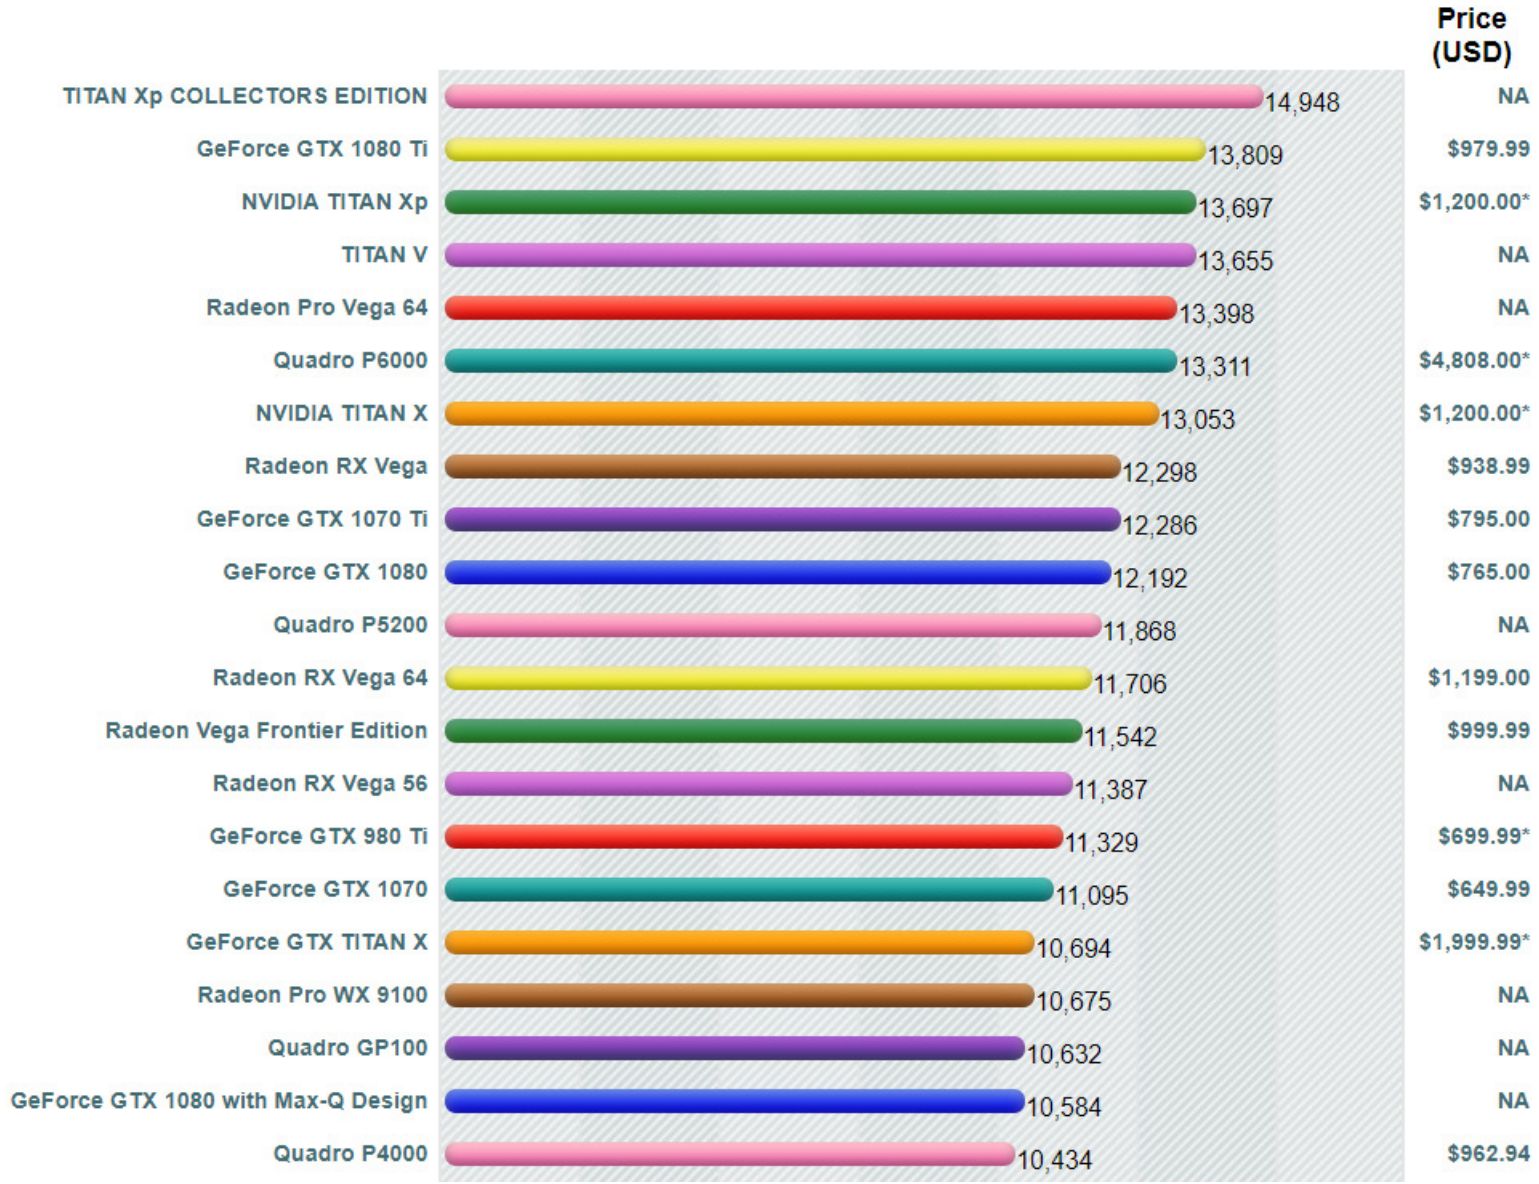
Price (USD)
TITAN Xp COLLECTORS EDITION
14,948
NA
GeForce GTX 1080 Ti
13,809
$979.99
NVIDIA TITAN Xp
13,697
$1,200.00*
TITAN V
13,655
NA
Radeon Pro Vega 64
13,398
NA
Quadro P6000
13,311
$4,808.00*
NVIDIA TITAN X
13,053
$1,200.00*
Radeon RX Vega
12,298
$938.99
GeForce GTX 1070 Ti
12,286
$795.00
GeForce GTX 1080
12,192
$765.00
Quadro P5200
11,868
NA
Radeon RX Vega 64
11,706
$1,199.00
Radeon Vega Frontier Edition
11,542
$999.99
Radeon RX Vega 56
11,387
NA
GeForce GTX 980 Ti
11,329
$699.99*
GeForce GTX 1070
11,095
$649.99
GeForce GTX TITAN X
10,694
$1,999.99*
Radeon Pro WX 9100
10,675
NA
Quadro GP100
10,632
NA
GeForce GTX 1080 with Max-Q Design
10,584
NA
Quadro P4000
10,434
$962.94

GPU and TPU – 2

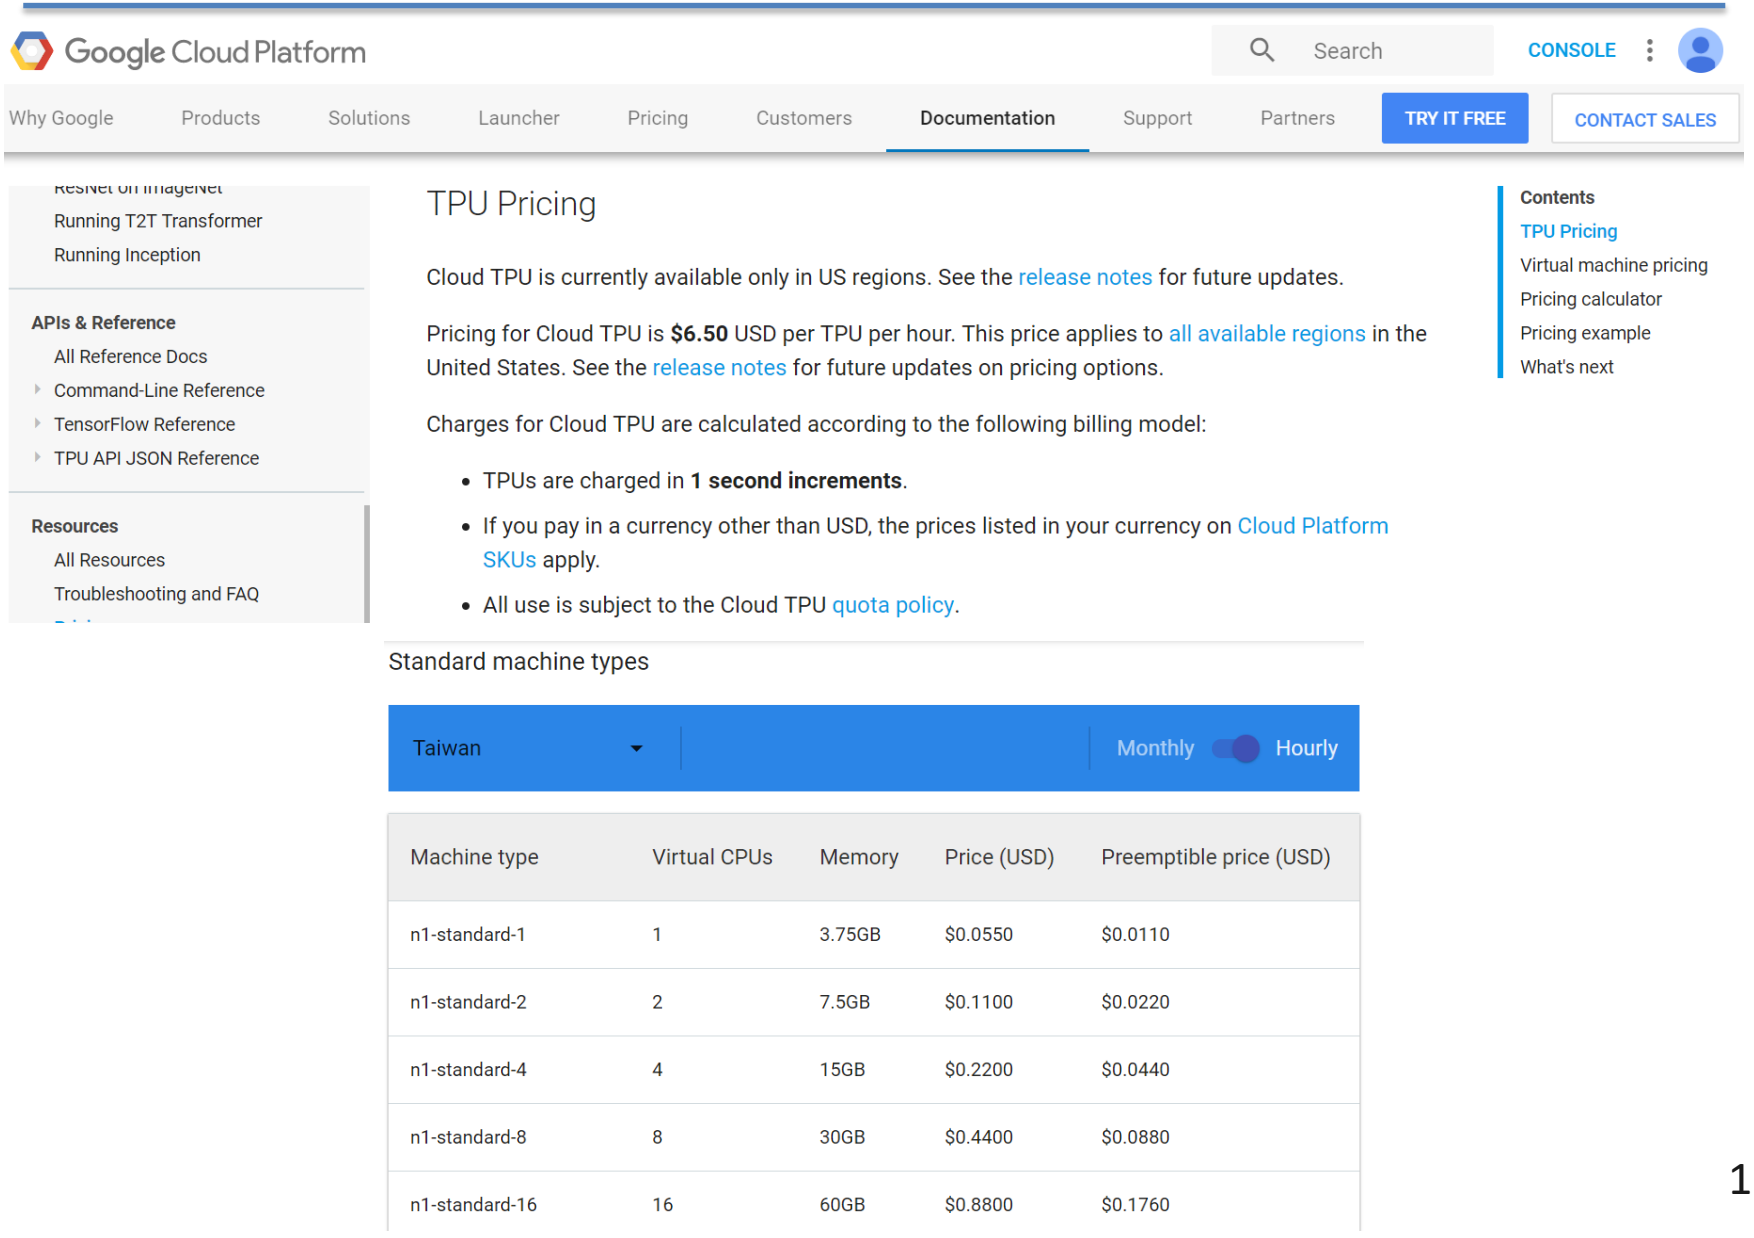
Google Cloud Platform
Search
CONSOLE
Why Google
Products
Solutions
Launcher
Pricing
Customers
Documentation
Support
Partners
TRY IT FREE
CONTACT SALES
Running T2T Transformer
Running Inception
APIs & Reference
All Reference Docs
Command-Line Reference
TensorFlow Reference
TPU API JSON Reference
Resources
All Resources
Troubleshooting and FAQ
TPU Pricing
Cloud TPU is currently available only in US regions. See the release notes for future updates.
Pricing for Cloud TPU is $6.50 USD per TPU per hour. This price applies to all available regions in the United States. See the release notes for future updates on pricing options.
Charges for Cloud TPU are calculated according to the following billing model:
TPUs are charged in 1 second increments.
If you pay in a currency other than USD, the prices listed in your currency on Cloud Platform SKUs apply.
All use is subject to the Cloud TPU quota policy.
Contents
TPU Pricing
Virtual machine pricing
Pricing calculator
Pricing example
What's next
Standard machine types
Taiwan
Monthly
Hourly
Machine type
Virtual CPUs
Memory
Price (USD)
Preemptible price (USD)
n1-standard-1
1
3.75GB
$0.0550
$0.0110
n1-standard-2
2
7.5GB
$0.1100
$0.0220
n1-standard-4
4
15GB
$0.2200
$0.0440
n1-standard-8
8
30GB
$0.4400
$0.0880
n1-standard-16
16
60GB
$0.8800
$0.1760

Tentative Syllabus

Date	Syllabus	Homework
3/1	Course Overview	
3/7, 8	Introduction to Deep Learning	
3/14, 15	Backpropagation	HW1 (python basics)
3/21, 22	Language Modeling and Word Embedding	
3/28, 29	Recurrent Neural Network	
4/4, 5	Break	
4/11, 12	Advanced RNNs	HW2 (MLP)
4/18, 19	Break	
4/25, 26	Paragraph Embeddings	HW3 (RNN)
5/2, 3	Convolutional Neural Networks	
5/9, 10	Advanced Structures	HW4 (CNN)
5/16, 17	Generative Adversarial Network-1	
5/23, 24	Generative Adversarial Network-2	HW5 (GAN)
5/30, 31	Special Topics	
6/6, 7	Special Topics	
6/13, 14	Special Topics	
6/20, 21	Break	
6/27, 28	Final Project Demo	

Tentative Grading

- Homework: 75%
  - 5 homeworks, each homework is 15%
- Final Project: 30%

- Instructor: Kuan-Yu Chen (陳冠宇)
  - E1-222-4
  - kychen@mail.ntust.edu.tw
  - (02) 2737-6377
  - http://faculty.csie.ntust.edu.tw/~kychen/courses/2018_Spring_DLNLP/2018_DLNLP.html

- TAs:吳政育, 吳澤鑫, 朱璟軒, 顏苙峰 and 羅上堡 (RB308-3)

Reference

- Ian Goodfellow and Yoshua Bengio and Aaron Courville, Deep Learning, MIT Press, 2016
  - http://www.deeplearningbook.org/

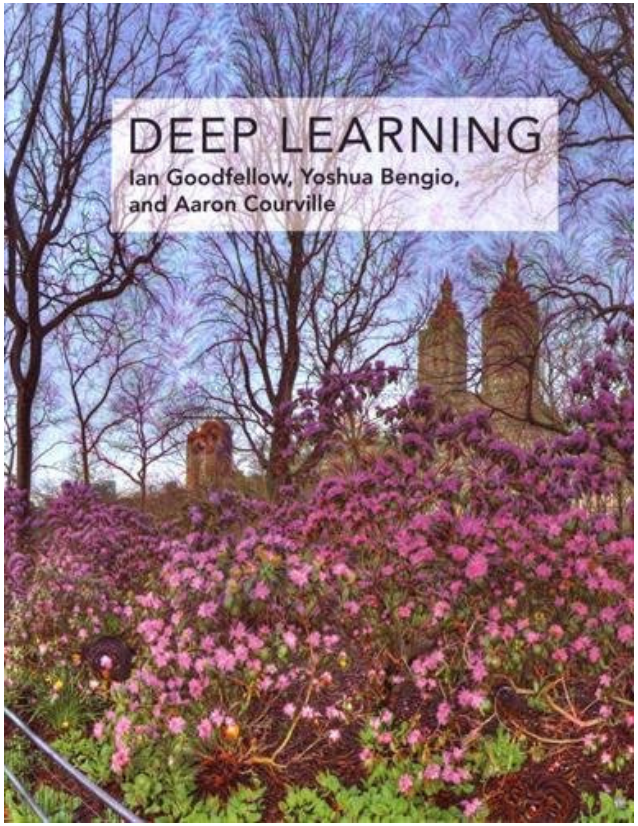

Book Signing at NIPS 2016

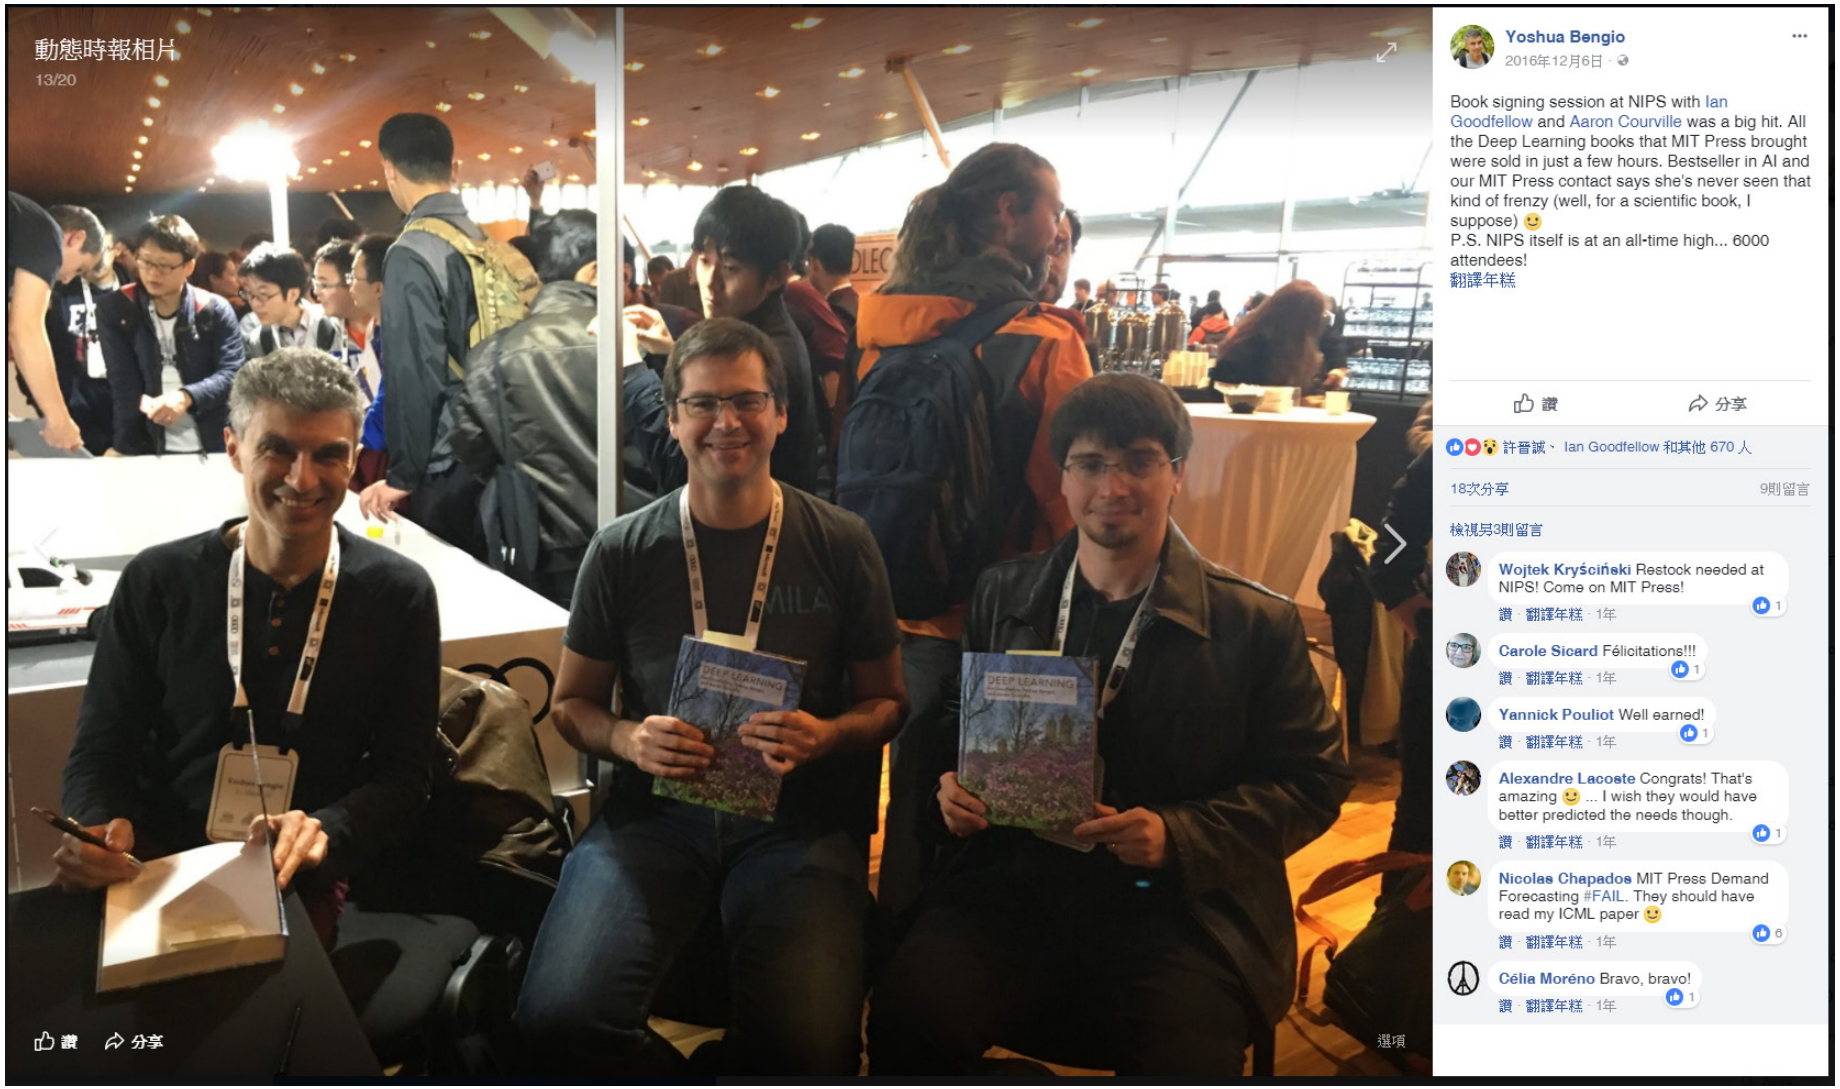

Questions?

kychen@mail.ntust.edu.tw