

Introduction to EE807: Recent Advances in Deep Learning

Jinwoo Shin, KAIST EE

Course Information

- Goal: Cover a very partial subset of recent advances in deep learning
- Course homepage: http://alinlab.kaist.ac.kr/ee807_2018.html
 - Slides are made by students in Algorithmic Intelligence Laboratory
 - Reference papers will be uploaded for each class (we have no textbook)
- Office hours: Every Tuesday, 11am-1pm, N1-914
- Instructor: Jinwoo Shin (N1-914)
 - Associate professor, KAIST EE
 - Cf. I am currently working on deep learning, but I was not a PhD in machine/deep learning or related fields. Hence, my views might be different from other deep learners (do not blame me).
 - Email: jinwoos@kaist.ac.kr
- TA: Insu Han (N1-917)
 - 2nd year PhD student, KAIST EE
 - Email: insu.han@kaist.ac.kr

Course Information

- Assignments
 - Option A (1 project) or Option B (2 paper summary)
 - You have to choose A or B in this week (send your decision to TA and if you do not send, I will assume that you will drop the class)
 - No written exam

- Option A: 1 Project
 - Each student should choose a separate topic decided on your own.
 - I will help the progress of your project (e.g., use the office hours).
 - Send the short description of your project (title and abstract) to TA in this week
 - Any topic is OK as long as it is (weakly) related to deep learning.
 - It is also OK if you choose some topic currently working with your advisor.
 - Send the intermediate progress report (of pdf form) by Oct. 21th and the final report by Dec. 16th to TA.
 - You have to present your project at the end of this class.

Course Information

- Option B: 2 paper summary
 - Each student should choose two deep learning papers published at NIPS, ICML or ICLR in last 3 years.
 - You can decide which papers to study later, but you have to send an email to TA in this week for notifying that you choose Option B.
 - I will help for deciding which papers to study (e.g., use the office hours to ask or send emails to me, including your generic interests and backgrounds)
 - Once you choose papers, try implementing the algorithms on your own using TensorFlow or PyTorch and applying to other datasets
 - Send the report on the first paper by Oct. 21th and the report on the second paper by Dec. 16th to TA. You also have to send your source-code files with the reports.
 - No presentation in the class


Course Information: Q&A

- How much backgrounds do I need on deep learning?
 - This course is not an introductory course to deep learning
 - I will cover some backgrounds quickly, but not spend too much time
 - For example, I will not teach how to use TensorFlow or PyTorch
 - Very sorry, but if you worry about this, please drop the class and be a listener (not for credit).

- How I can ask help for project or paper summary?
 - Use the office hours
 - Or, you can send emails to me for additional meeting requests

Recent Trends in Artificial Intelligence

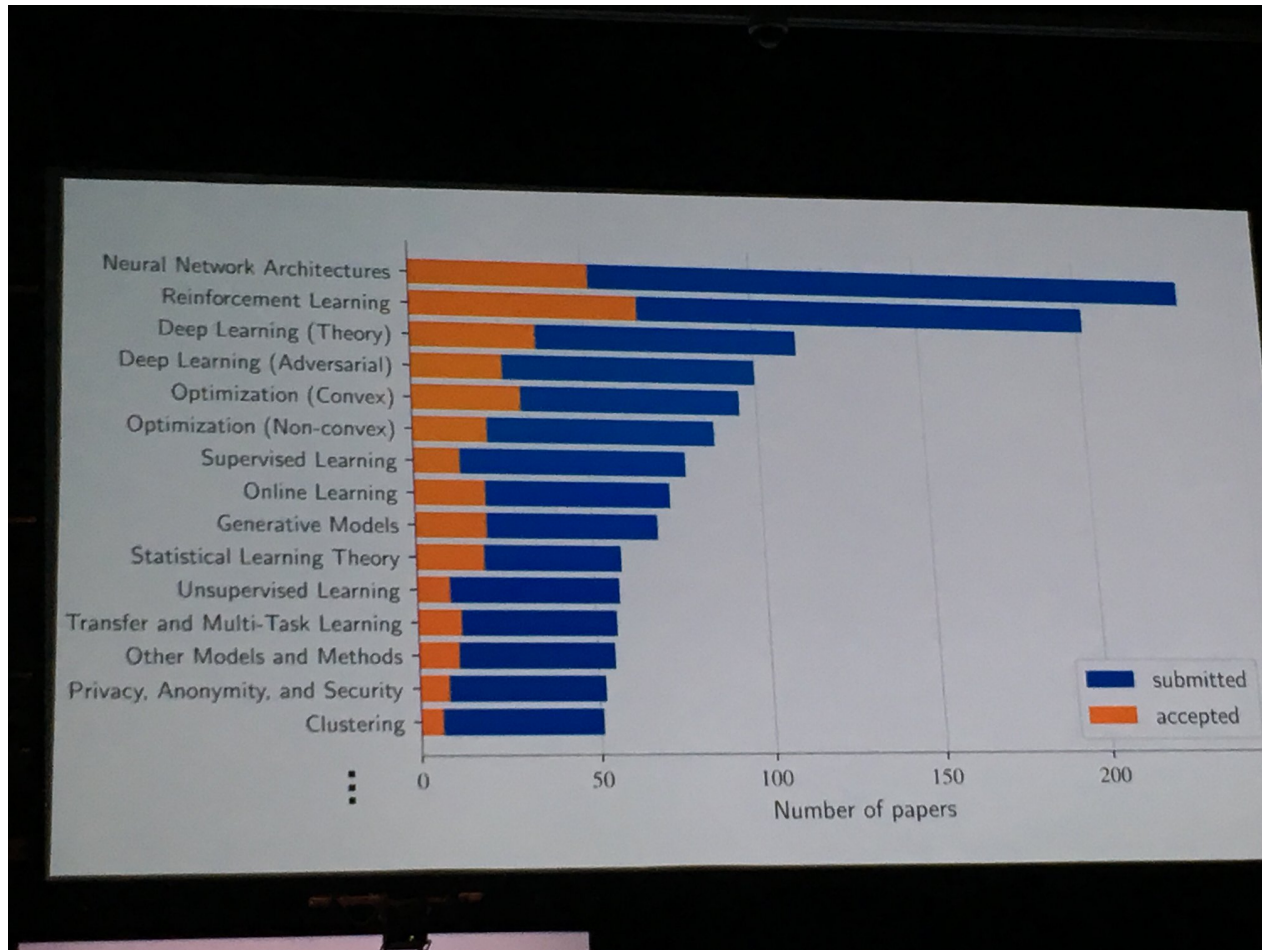
- Deep Learning (DL) is a subfield of Machine Learning (ML)
 - Machine Learning is a subfield of Artificial Intelligence (AI)



AAAI, IJCAI = AI, NIPS, ICML = ML, CVPR = computer vision, ICRA = robotics, ACL = natural language processing

Recent Trends in Machine Learning

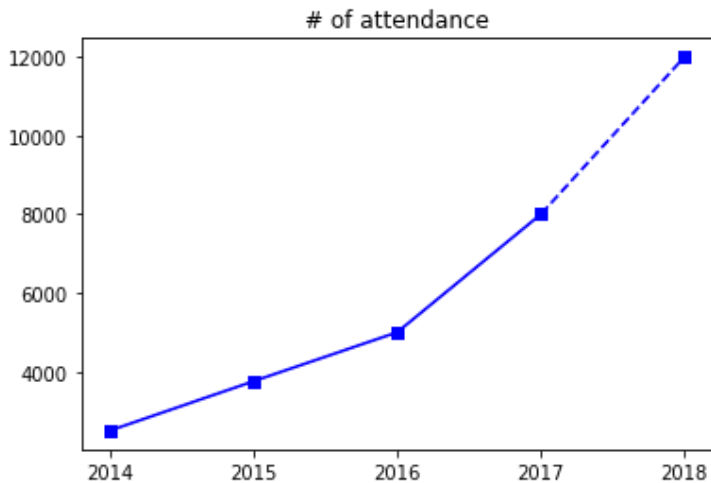
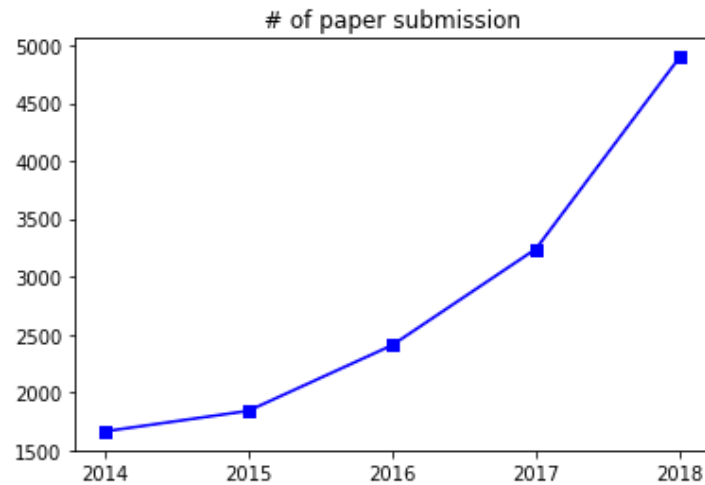
- Statistics of ICML 2018



The ratio of deep learning papers is roughly 30%

Why top AI conference publications matters?

of submissions and attendances at NIPS



Nvidia CEO reveals the company's latest GPU at A.I. conference

- Nvidia's Titan V desktop for PCs is specifically aimed at artificial intelligence, an area where the company has made gains in recent years.
- Nvidia says the new graphics card provides nine times the performance of its predecessor.

Jordan Novet | @jordannovet

Published 11:00 PM ET Thu, 7 Dec 2017 | Updated 1:53 PM ET Fri, 8 Dec 2017



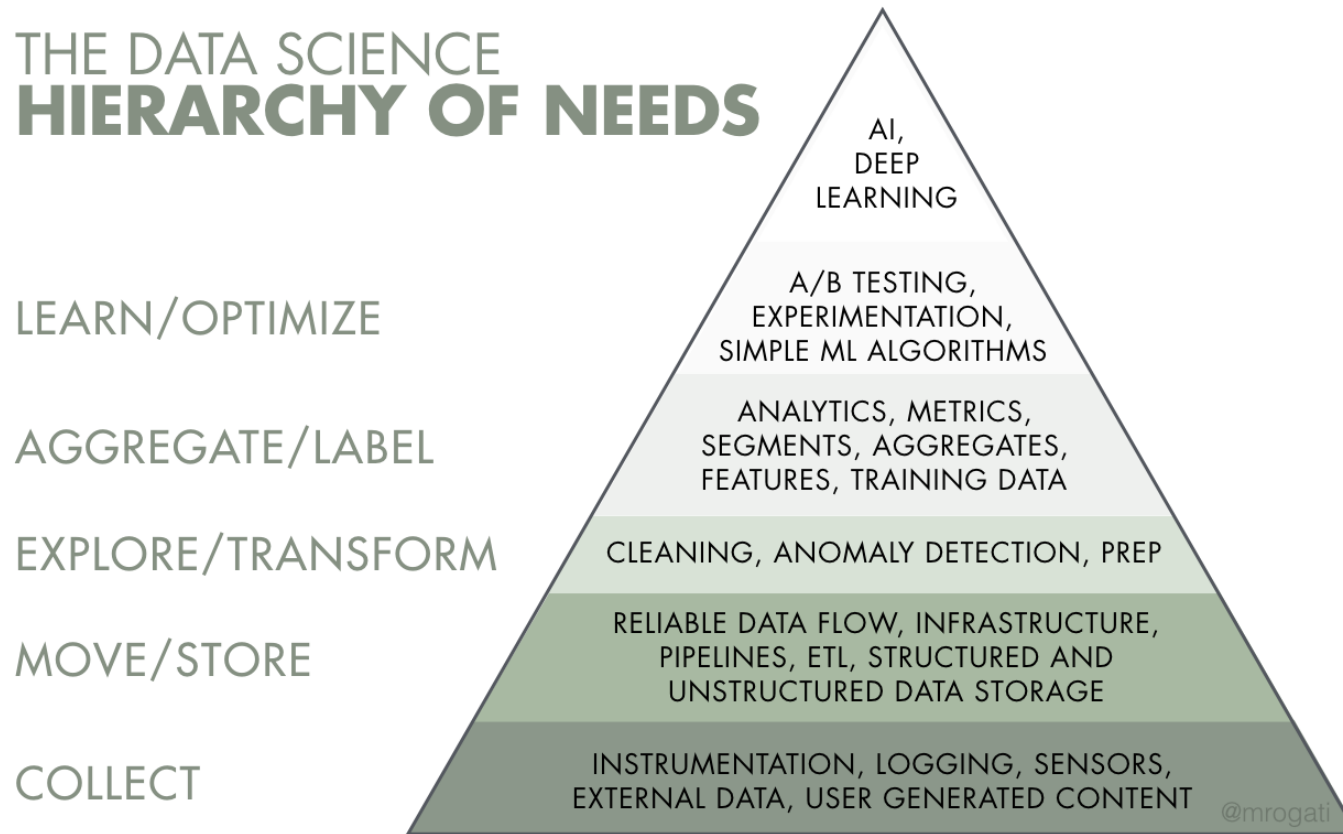
Nvidia on Thursday announced the launch of its latest graphics card for artificial intelligence.

CEO Jensen Huang showed off the Titan V graphics processing unit for desktop PCs, which employs Nvidia's seventh-generation Volta architecture, for the first time at a major AI conference in Long Beach, Calif.

The new GPU delivers 110 teraflops of performance, or nine times that of the comparative model from Nvidia's previous architecture, the Titan Xp, which was announced in April. But like Nvidia's Tesla V100 GPU for data centers, which was unveiled in May, the new Titan V is more clearly oriented toward AI than previous generations.

AI Hierarchy of Needs?

THE DATA SCIENCE HIERARCHY OF NEEDS




From <https://hackernoon.com/the-ai-hierarchy-of-needs-18f111fcc007>

For example, Ilya Sutskever (Research Director of OpenAI) is paid more than \$ 1.9 million USD in 2016.

From <https://www.nytimes.com/2018/04/19/technology/artificial-intelligence-salaries-openai.html>

Recent hype on deep learning ends soon?

Hype Cycle for Emerging Technologies, 2018



As of July 2018

gartner.com/SmarterWithGartner

Source: Gartner (August 2018)
© 2018 Gartner, Inc. and/or its affiliates. All rights reserved.

Gartner.