NIPS 2016

Marek & Márius

Stats

5680 registered people, 4500+ in mobile app

Industry + academia

10 tutorials, multiple parallel tracks, 173+198+198 posters, 20 demos, 3 symposia, 28+24 workshops

Majority focused on neural networks, some PGMs


Disney Research
Our Partners:
LUCASFILM
Disneyland
Walt Disney World
DisneyToon
Disneyland
Disney
ESPN
MARVEL
PIXAR
abc
Walt Disney Parks and Resorts
Disney Music Group
Disney Theatrical Group
Publishing Worldwide
Disney

Nuts and bolts of buildings AI systems

Bias vs variance

Automatization of PhD study

Slides in references

http://www.computervisionblog.com/2016/12/nuts-and-bolts-of-building-deep.html


AI can help save the world
from efficiency savings... ... to radical breakthroughs
DeepMind
@DeepMindAI
@drewpurves
Intelligent Biosphere – Drew Purves
5


89
Dialog-based Language Learning
Jason Weston
Facebook AI Research
Memory Networks, LSTMs with attention etc. in NLP
(Original) bAbI tasks
The Dialog-based Supervision Tasks
1) Imitation Learning
Related Work
Forward Prediction: how the hell does it work?
Conclusion
Google


TensorFlow w/XLA:
TensorFlow, Compiled!
Expressiveness with performance

Highlights

Lots of new ideas & architectures, but few state-of-the-art methods

RL, open platforms - universe.openai.com, project Malmo

GANs

Word2vec-like unsupervised learning of image representation

Video prediction, Image superresolution

TorontoCity dataset

Using Fast Weights to Attend to the Recent Past

https://arxiv.org/pdf/1610.06258.pdf

Value Iteration Networks

https://papers.nips.cc/paper/6046-value-iteration-networks.pdf

Project Malmo & Universe

https://www.microsoft.com/en-us/research/project/project-malmo/

https://universe.openai.com/

Learning What and Where to Draw

https://arxiv.org/pdf/1610.02454v1.pdf

Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks

https://arxiv.org/pdf/1602.07868.pdf

Attend, Infer, Repeat: Fast Scene Understanding with Generative Models

http://arkitus.com/files/arxiv-attend-infer-repeat.pdf

R-FCN: Object Detection via Region-based Fully Convolutional Networks

https://arxiv.org/pdf/1605.06409v2.pdf

Self Supervised Learning of Visual Representations

https://arxiv.org/pdf/1603.09246v2.pdf

Scaling-up: Image Super-resolution and Compression for the masses

https://arxiv.org/pdf/1609.04802.pdf

https://arxiv.org/pdf/1609.05158v2.pdf

Baidu Research @BaiduResearch · Dec 5

The long lines at #nips2016 speak volumes of the enthusiasm for #AI #machinelearning #deeplearning

1 22 42

Alex Champandard @alexjc · Dec 7

Me: #nips2016 is content-packed, but it's pre-published online and very random.
Petra: Why do people go?
Me: Trying to hire each other! ;-)

4 2 24

Naomi Saphra and 5 others Retweeted

Smerity @Smerity · Dec 8

Celebrating the 20 yr anniversary of LSTM ... being rejected from #nips1996 (credit for perseverance by Hochreiter & Schmidhuber!) #nips2016

2 108 223

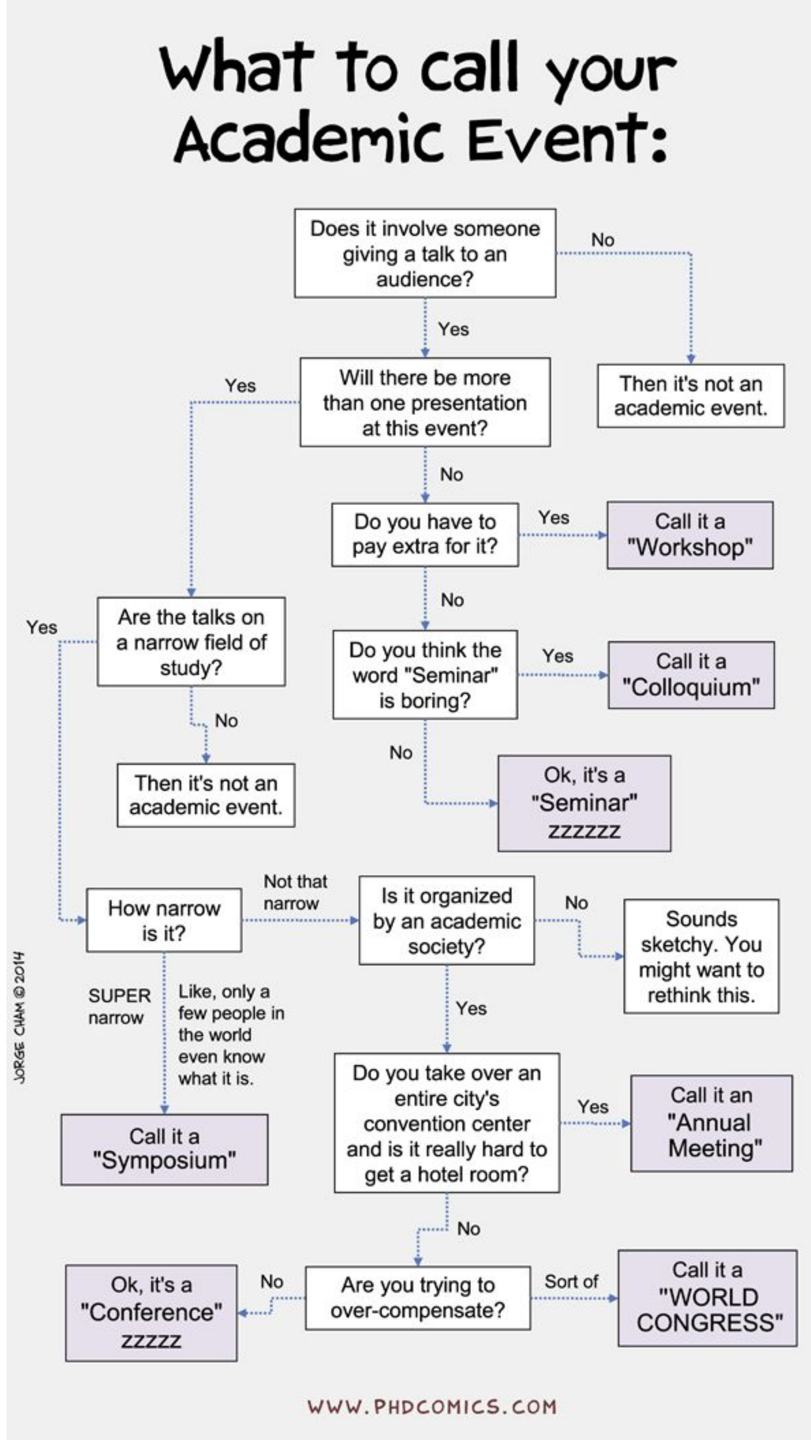
What to call your Academic Event:
Does it involve someone giving a talk to an audience?
No
Then it's not an academic event.
Yes
Will there be more than one presentation at this event?
No
Do you have to pay extra for it?
Yes
Call it a "Workshop"
No
Do you think the word "Seminar" is boring?
Yes
Call it a "Colloquium"
No
Ok, it's a "Seminar" zzzzzz
Yes
Are the talks on a narrow field of study?
No
Then it's not an academic event.
Yes
How narrow is it?
SUPER narrow
Like, only a few people in the world even know what it is.
Call it a "Symposium"
Not that narrow
Is it organized by an academic society?
No
Sounds sketchy. You might want to rethink this.
Yes
Do you take over an entire city's convention center and is it really hard to get a hotel room?
Yes
Call it an "Annual Meeting"
No
Are you trying to over-compensate?
No
Ok, it's a "Conference" zzzzz
Sort of
Call it a "WORLD CONGRESS"
JORGE CHAM © 2014
WWW.PHDCOMICS.COM

Other quick takeaways

- Andrew Ng:
  - Almost all of the value is created by Supervised Learning.
  - For some products at Baidu, the main purpose is to acquire data from users, not revenue.
  - Current machine learning techniques work best when training data and real data come from the same distribution. Real systems rarely exhibit this.
  - If you design a ML system for the real world, comparing against human baseline needs to be the norm
- Want to win Kaggle competitions?
  - XGBoost for tabular data
  - CNNs for images
  - RNNs for time series predictions
- Nando
  - “If your job is to come up with algos that learn sigmoid based networks, you are out of job. Not so much for ReLus”
- RNN Panel
  - "You know, my wife is a local minimum, but you can live with a local minimum -- you have finite time".
  - “I do not think there is a disagreement, I see it as a difference in tactis.”
- GANs
  - “The best idea in ML in the past 20 years” -- LeCun
  - Simple to put together, extra finicky to train

Vanilla GAN

Vanilla GAN
(Goodfellow, et al., 2014)

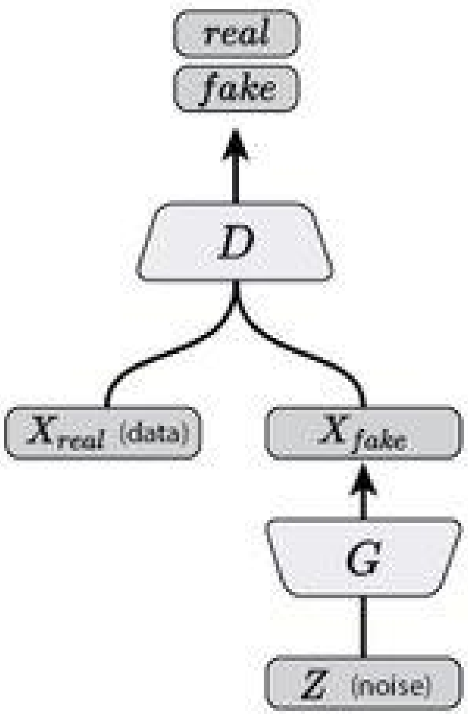

Discriminator Looks at Latent Variables

Conditional GAN
(Mirza & Osindero, 2014)

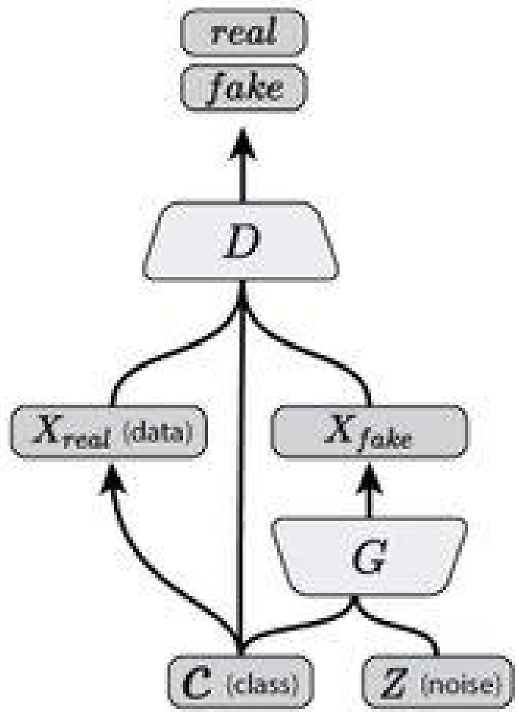

Bidirectional GAN
(Donahue, et al., 2016; Dumoulin, et al., 2016)

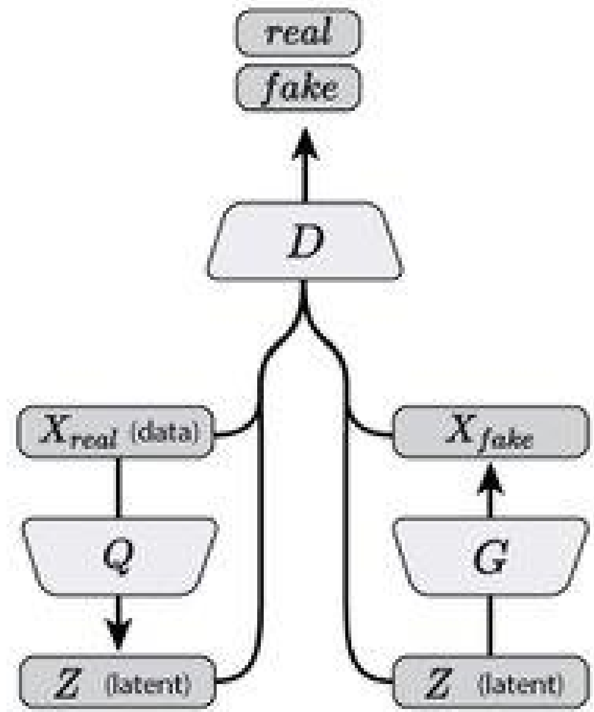

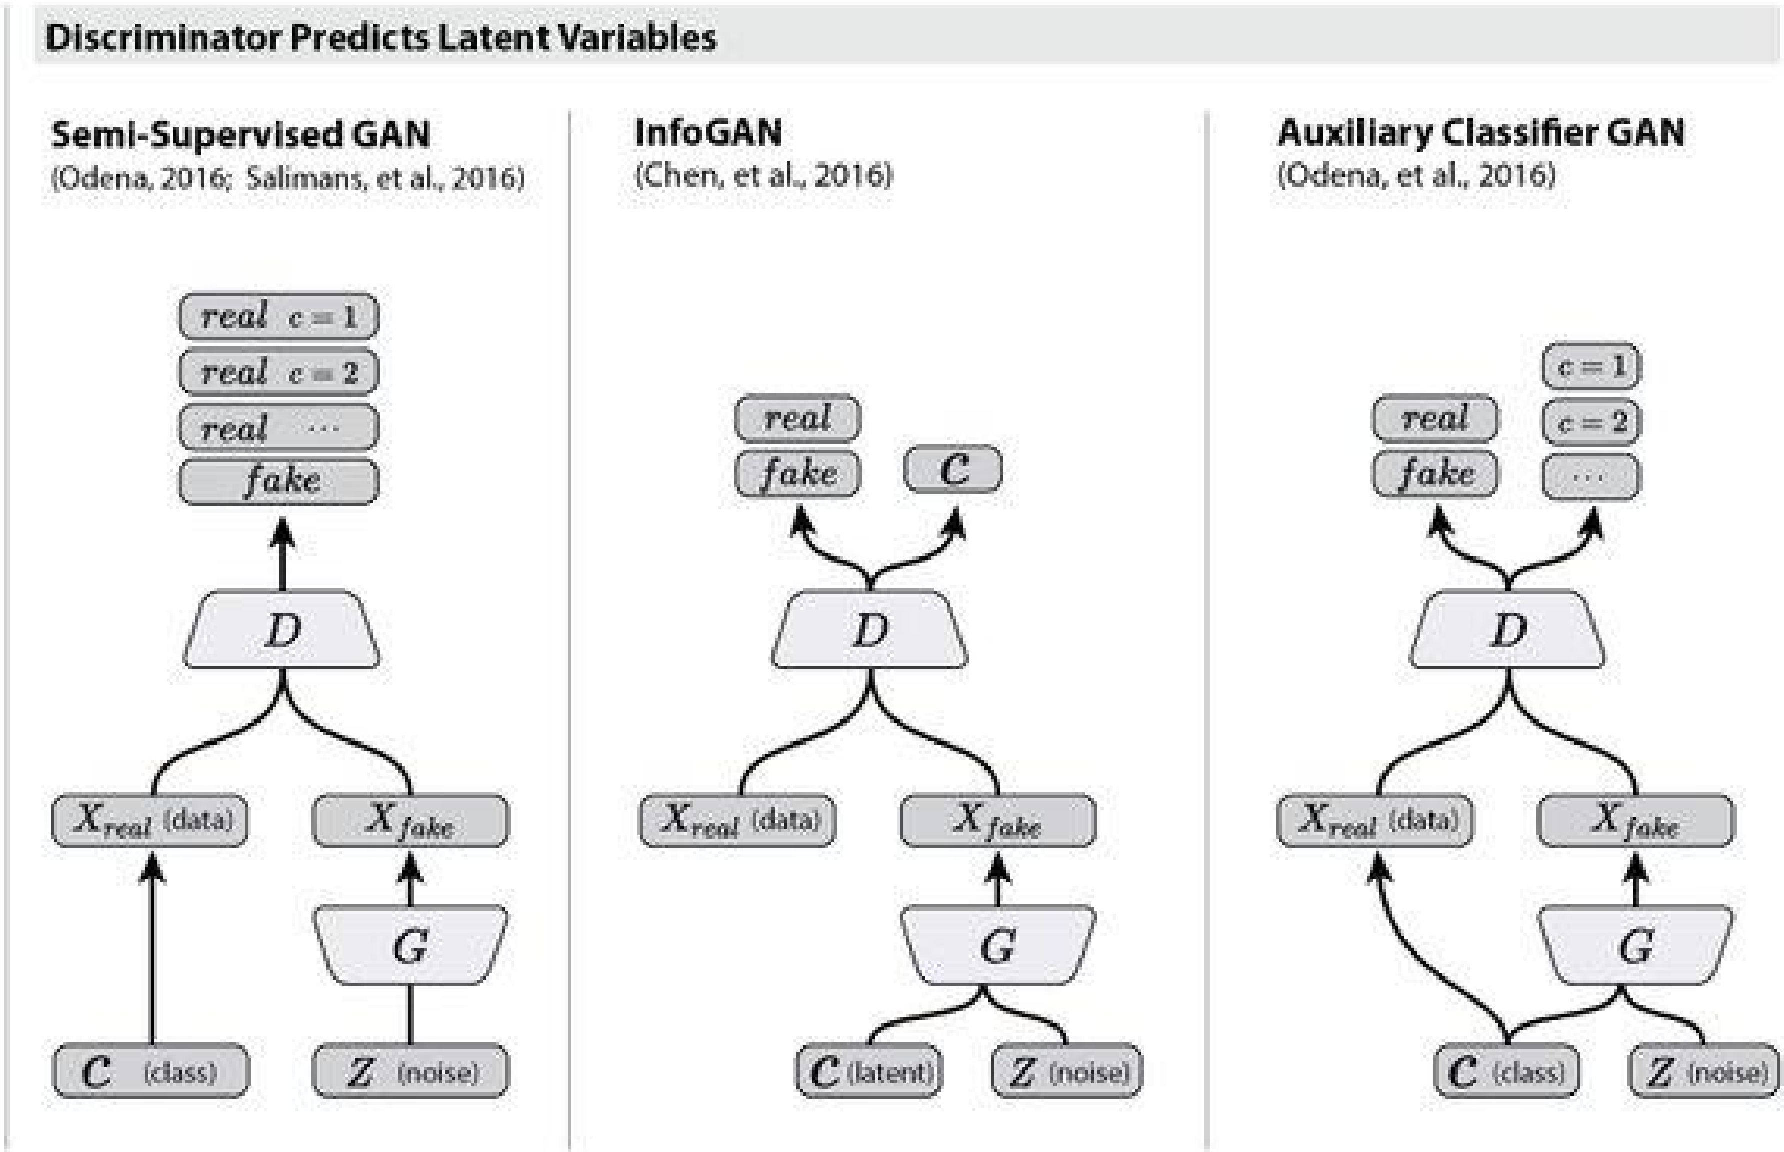
Discriminator Predicts Latent Variables
Semi-Supervised GAN
(Odena, 2016; Salimans, et al., 2016)
real c = 1
real c = 2
real ...
fake
D
X real (data)
X fake
G
C (class)
Z (noise)
InfoGAN
(Chen, et al., 2016)
real
fake
C
D
X real (data)
X fake
G
C (latent)
Z (noise)
Auxiliary Classifier GAN
(Odena, et al., 2016)
real
fake
c = 1
c = 2
...
D
X real (data)
X fake
G
C (class)
Z (noise)

Was it worth it?

Alex Champandard
@alexjc

Following

These 50 Lessons are particularly useful & well written! But you could have learned >90% without going to #NIPS2016.

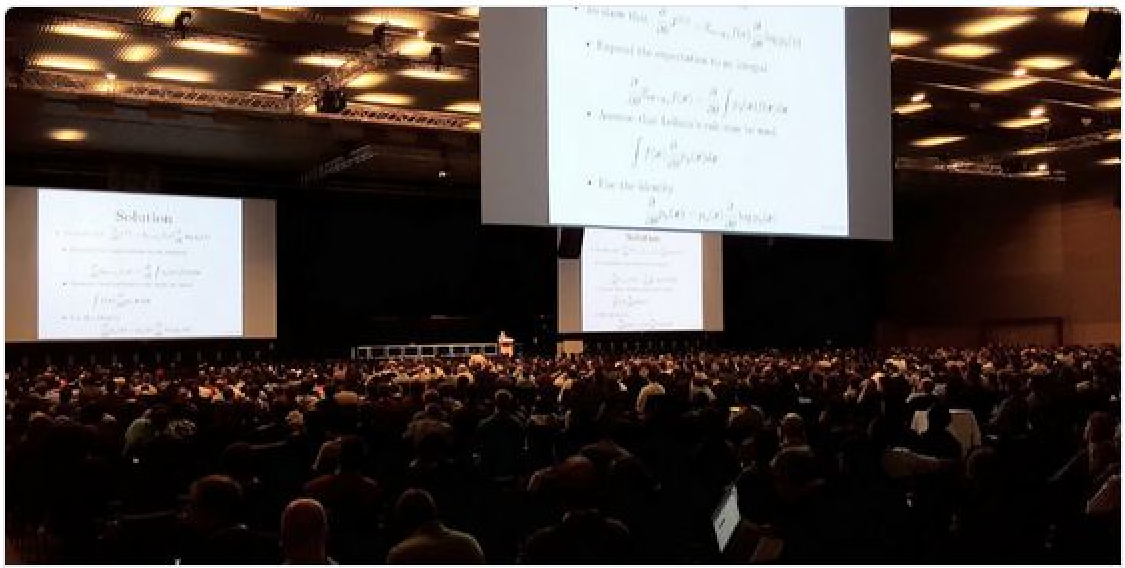

50...
I l...
blog.ought.com

Was it worth it?

Alex Champandard @alexjc · 3h
The field of #ML, thanks to arXiv / reddit / Twitter, is more accessible than ever for anyone regardless of their location or resources.

2 1 5

Alex Champandard @alexjc · 3h
If you're considering NIPS next year, ask yourself if you'd like to meet the authors of papers you read in-depth. If not, just go to arXiv!

1 3

Alex Champandard @alexjc · 3h
Even beginner tutorials are better online, whether MOOC or otherwise. Sessions at NIPS were good but live talks can't match well-prepared!

1 1 3

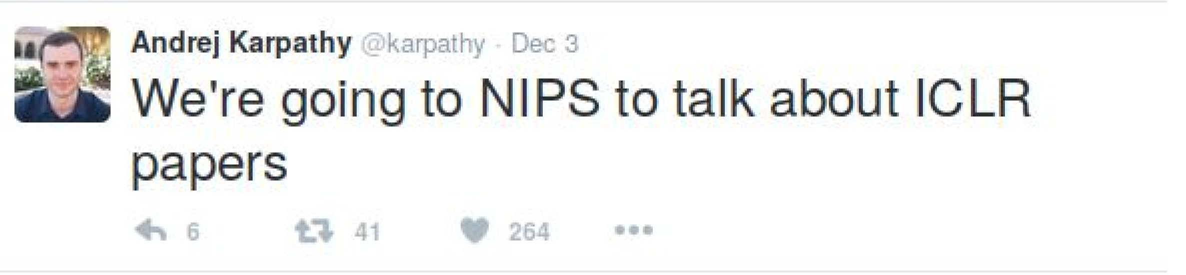

Example of stuff that is already (quite) old: Weight Norm

https://arxiv.org/pdf/1602.07868.pdf

References

- NIPS Review Process: http://www.tml.cs.uni-tuebingen.de/team/luxburg/misc/nips2016/index.php
- Towards biologically plausible Deep Learning -Yoshua Bengio: http://www.iro.umontreal.ca/~bengioy/talks/Brains+Bits-NIPS2016Workshop.pptx.pdf
  - Video from a similar talk: https://archive.org/details/Redwood_Center_2016_09_27_Yoshua_Bengio
- Magenta Demo
  - https://magenta.tensorflow.org/2016/12/16/nips-demo/
- Andrew Ng Tutorial Slides
  - https://www.dropbox.com/s/dyjdq1prjbs8pmc/NIPS2016%20-%20Pages%202-6%20(1).pdf?dl=0
- A list of code for work presented at NIPS
  - https://www.reddit.com/r/MachineLearning/comments/5hwqeb/project_all_code_implementations_for_nips_2016/
  -

Other NIPS reviews/writeups

- http://www.slideshare.net/SebastianRuder/nips-2016-highlights-sebastian-ruder
- http://inverseprobability.com/2016/12/13/nips-highlights.html
- https://beamandrew.github.io/deeplearning/2016/12/12/nips-2016.html
- https://www.linkedin.com/pulse/nips-2016-towards-end-dynamic-dialogue-system-vishal-bhalla
- https://medium.com/@libfun/nips-2016-experience-and-highlights-104e19e4ac95#.ngf9jdu3e
- https://www.linkedin.com/pulse/some-general-take-aways-from-nips2016-igor-carron
- http://www.inference.vc/my-summary-of-adversarial-training-nips-workshop/
- https://tryolabs.com/blog/2016/12/06/major-advancements-deep-learning-2016/
- http://apeiroto.pe/ml/nips-2016.html
- http://www.machinedlearnings.com/2016/12/nips-2016-reflections.html
- http://blog.aylien.com/highlights-nips-2016/
-