

What is new in AI and deep learning in 2015?


Tapani Raiko

Aalto University

16 Dec 2015

Who is Tapani Raiko?

- ▶ Assistant professor and academy research fellow
- ▶ Leader of Deep Learning research group at Aalto
- ▶ Lecturing AI (spring) and deep learning (autumn)



Achilles' heel of AI 1955-2014: Perception in natural environment

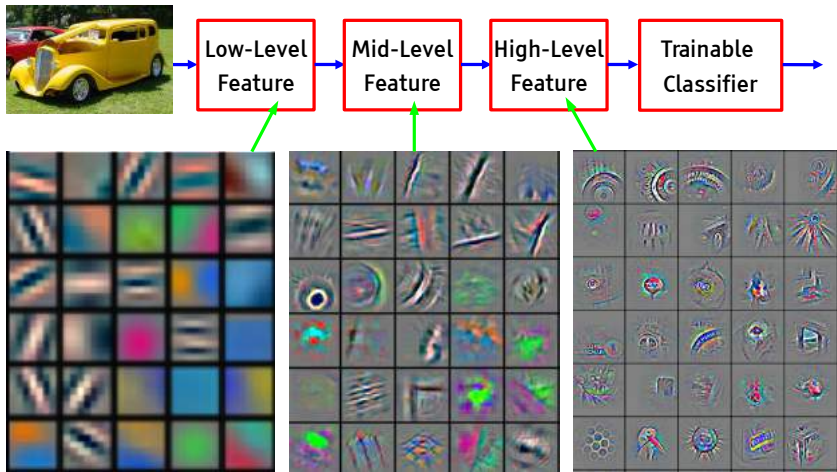


xkcd.com/1425

Deep Learning = Learning Hierarchical Representations

Y LeCun

It's deep if it has more than one stage of non-linear feature transformation




Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Deep learning is **hot** in academy

- ▶ "Deep learning ... dramatically improved the state-of-the-art in speech recognition, visual object recognition..." (LeCun et al., Nature, 2015)
- ▶ "...bridges the divide between high-dimensional sensory inputs and actions, resulting in the first artificial agent..." (Mnih et al., Nature, 2015)
- ▶ "Knowing the sequence specificities of DNA- and RNA-binding proteins is essential ... deep learning outperforms other state-of-the-art methods" (Alipanahi et al., Nature Biotechnology, 2015)


Deep learning is **hot** in industry



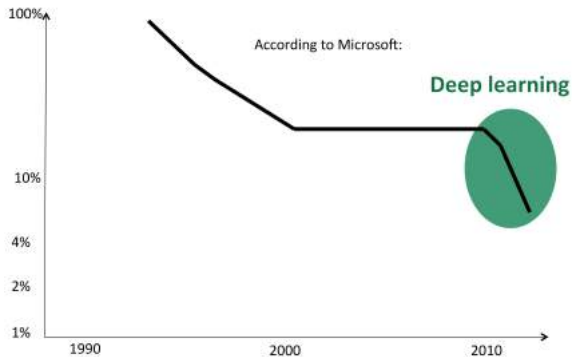
Google acquired startup DeepMind for \$500M in 2014.
Global race: Facebook, Microsoft, Baidu, IBM, Amazon, Samsung, Apple...

NIPS Growth

Total Registrations 3755




Speech to text



2015

Image classification



Images to text (Vinyals et al., 2015)



A woman is throwing a **frisbee** in a park.



A **dog** is standing on a hardwood floor.



A **stop** sign is on a road with a mountain in the background



A little **girl** sitting on a bed with a teddy bear.



A group of **people** sitting on a boat in the water.



A giraffe standing in a forest with **trees** in the background.

► Note: Interdisciplinary. Attention modelling.

Text to images (Mansimov et al., 2015)



A very large commercial plane flying in blue skies.



A very large commercial plane flying in rainy skies.




A herd of elephants walking across a dry grass field.



A herd of elephants walking across a green grass field.

Generating faces torch.ch/blog/2015/11/13/gan.html

Noise $\sim N(0,1)$



Batch normalization (Ioffe and Szegedy, 2015)

Input: Values of x over a mini-batch: $\mathcal{B} = \{x_1 \dots x_m\}$;
Parameters to be learned: γ, β

Output: $\{y_i = \text{BN}_{\gamma, \beta}(x_i)\}$

$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{ mini-batch mean}$$


$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{ mini-batch variance}$$

$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{ normalize}$$

$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma, \beta}(x_i) \quad // \text{ scale and shift}$$

Batch normalization (Ioffe and Szegedy, 2015)


- ▶ Robust to (scale of) initialization
- ▶ Speeds up learning
- ▶ Regularizes



Why now?

- ▶ Cheap computation (GPU)
 - ▶ Good match with deep learning
- ▶ Big data (open data)
- ▶ Open source algorithms
 - ▶ Easy GPU usage
 - ▶ Automatic differentiation
 - ▶ OpenAI started

Cambrian explosion for robotics?





Thanks!
Tapani Raiko, Aalto University

Finnish AI Society

- ▶ Finnish AI Conference every second year since 1984
- ▶ Nowadays local conferences not appreciated
- ▶ Currenty inactive, would require a whole new board and new activities
- ▶ Ask more if interested