Few-shot generative models

Ole Winther

Department of Biology and Rigshospitalet
University of Copenhagen (KU)

Dept for Applied Mathematics and Computer Science
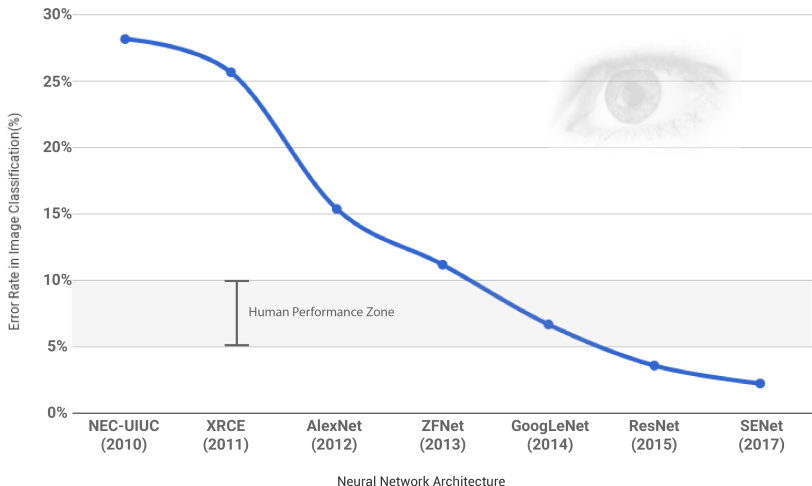Technical University of Denmark (DTU)

November 7, 2019

Deep learning solves perception tasks

# ImageNet classification challenge 2017 update

- Progress from bigger models and model architecture twists

Few-shot generative learning - agenda

- Breakthroughs in AI for text
- Context dependent language representations
- Industry case - raffle.ai
- Few shot and meta-learning

Few-shot generative learning - agenda

- Breakthroughs in AI for text
- Context dependent language representations
- Industry case - raffle.ai
- Few shot and meta-learning
- Bio
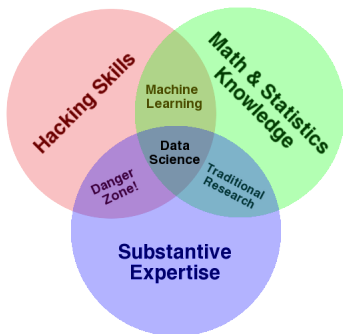
Generative pre-trained (GPT) model

SYSTEM PROMPT (HUMAN-WRITTEN)	*In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.*
MODEL COMPLETION (MACHINE-WRITTEN, 10 TRIES)	The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.
	Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.
	Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.
	Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.
	Pérez and his friends were astonished to see the unicorn herd. These creatures could be seen from the air without having to move too much to see them — they were so close they could touch their horns.

OpenAI blog: Better Language Models and Their Implications

AGI coming?

> While examining these bizarre creatures the scientists discovered
> that the creatures also spoke some fairly regular English. Pérez
> stated, "We can see, for example, that they have a common
> 'language,' something like a dialect or dialectic."
>
> Dr. Pérez believes that the unicorns may have originated in
> Argentina, where the animals were believed to be descendants of a
> lost race of people who lived there before the arrival of humans
> in those parts of South America.
>
> While their origins are still unclear, some believe that perhaps
> the creatures were created when a human and a unicorn met each
> other in a time before human civilization. According to Pérez, "In
> South America, such incidents seem to be quite common."
>
> However, Pérez also pointed out that it is likely that the only
> way of knowing for sure if unicorns are indeed the descendants of
> a lost alien race is through DNA. "But they seem to be able to
> communicate in English quite well, which I believe is a sign of
> evolution, or at least a change in social organization," said the
> scientist.

OpenAI blog: Better Language Models and Their Implications

Q: If a robot talks like a human, does it think like a human?

The answer is no. Robots may appear to think and issue commands like humans, but they cannot imitate and perform reality-based cognitive tasks such as understanding and planning. Testing is required.

Q: How would you test whether a robot can understand or plan?

Simple scenarios that require low levels of reasoning skills can be simulated by a computer. Practical reasoning tasks are harder. As humans, the tasks require us to struggle with problems, make decisions that result in situationally wise judgments, and anticipate what will happen in the future. Computers may never be able to do this.

Q: What if we had vastly more computational power than we have today?

We could make robots work for us to complete important tasks, and perhaps they'd be quicker, but machines would never be able to do true intellectual tasks laboriously. Computers would often perform basic reasoning tasks faster than a human can trouble-shoot the situation, and their responses would often be based on computer

Q: So they would get faster at simple tasks, but reach limits for complex tasks requiring practical reasoning?

No. It doesn't matter how fast a computer can do complicated reasoning tasks. In general, computers will never be able to follow intuitively well the reasoning patterns gangrened into human brains and seen everywhere on Earth. Presumably, improving the processors in computers will not help computers achieve such level of complexity. Rather, technology should in

Transformers - attention is all you need
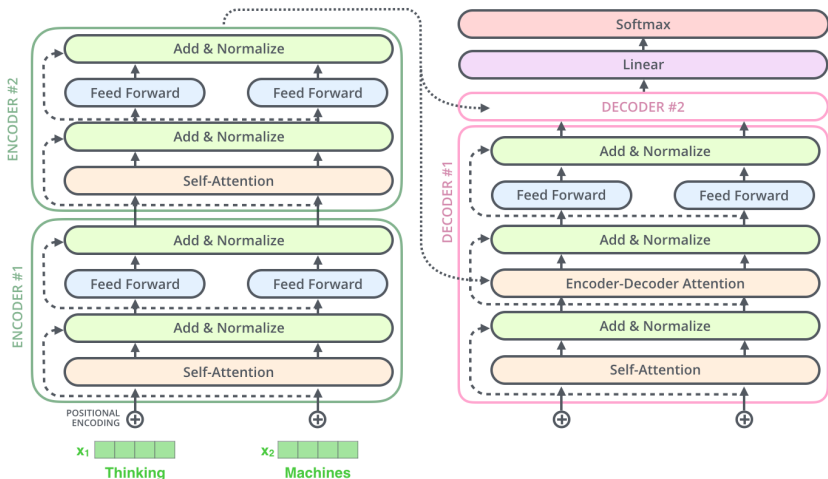
The illustrated transformer

Context dependent language representations
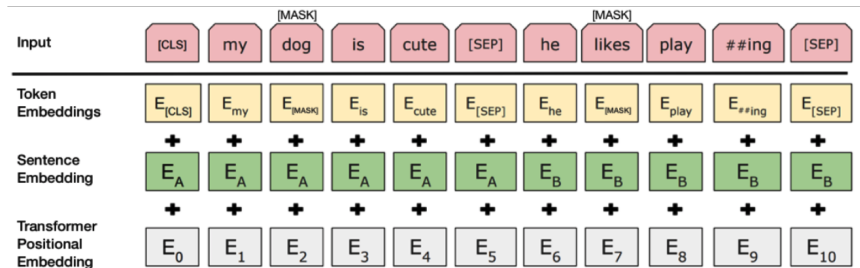
BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding
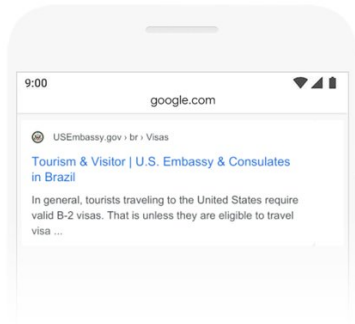
Google search 2019 - now with BERT
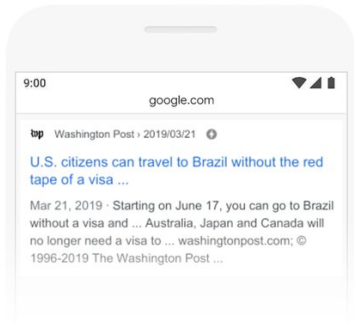
Google blog: Understanding searches better than ever before

Google search 2019 - now with BERT
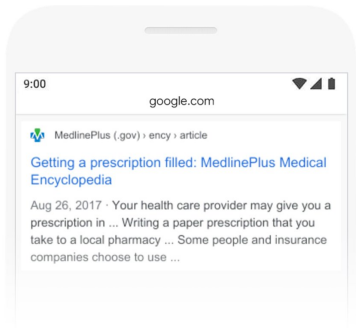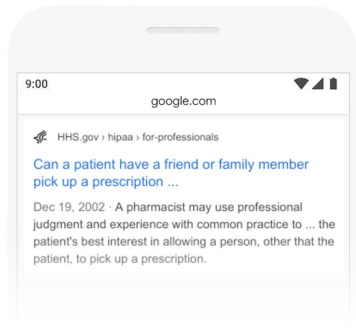
Google blog: Understanding searches better than ever before

raffle.ai - AI powered search for corporates

- Filling the gap between research and company reality

ImageNet - image classification

- 1.000 different classes - including many breeds of dogs!
- 1.000.000 training images

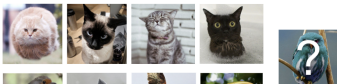# Humans are very good at one-shot learning

One shot learning of simple visual concepts

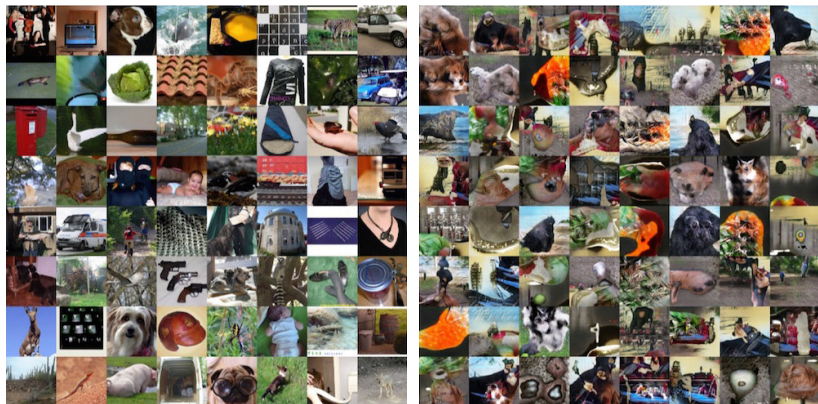# Meta- and few-shot learning - supervised set-up
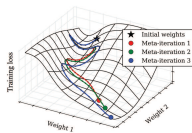

lilianweng.github.io

Generative models

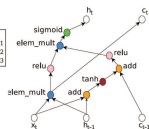What I cannot create, I do not understand. - Richard Feynman

openai.com/blog/generative-models/

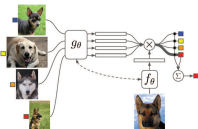# Few-shot generative models (DFF funded project)

- Focus on unsupervised (generative) setting
- Principled Bayesian formulation
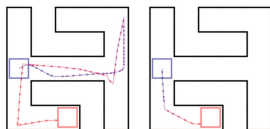- Provide alternative view of pre-training approaches.

hyperparameter optimization
Maclaurin et al. '15

learned recurrent cell
Zoph & Le '17

few-shot image classifier
Vinyals et al. '16

learning to quickly navigate new mazes
Duan et al. '16

bair.berkeley.edu/blog