Deep Learning: successes, promises and limits

Guest lecture
Intro to Data Science, Fall 2018 @ CCNY
December 3d 2018, New York
Tom Sercu

Who Am I

- ir. UGent (een burgie)
- Master in Data Science at NYU (New York University)
- Researcher at IBM Research AI
- Published at Speech Recognition and Machine Learning conferences

My goal is for you to:

Build better feeling and understanding of Deep Learning and its limitations!

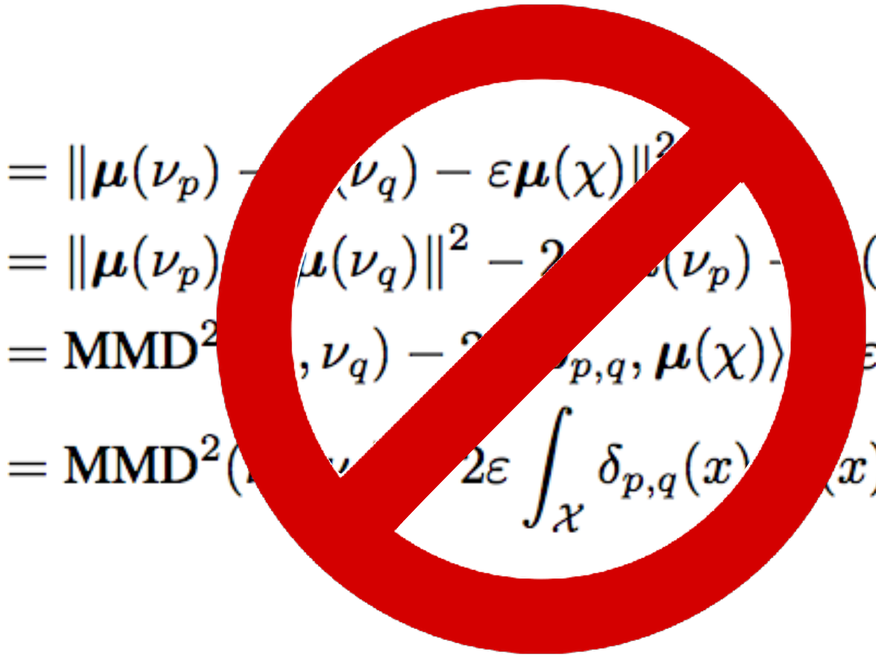

$$= \|\boldsymbol{\mu}(\nu_p) - \quad (\nu_q) - \varepsilon\boldsymbol{\mu}(\chi)\|^2$$
$$= \|\boldsymbol{\mu}(\nu_p) \quad \boldsymbol{\mu}(\nu_q)\|^2 - 2 \quad (\nu_p) - \quad (\nu_q), \boldsymbol{\mu}(\chi)\rangle + \varepsilon^2$$
$$= \mathrm{MMD}^2 \quad , \nu_q) - 2 \quad _{p,q}, \boldsymbol{\mu}(\chi)\rangle \quad \varepsilon^2 \|\boldsymbol{\mu}(\chi)\|^2$$
$$= \mathrm{MMD}^2(\quad \quad 2\varepsilon \int_{\mathcal{X}} \delta_{p,q}(x) \quad x) + \varepsilon^2 \|\boldsymbol{\mu}(\chi)\|^2$$

My goal is for you to:

Build better feeling and understanding of Deep Learning and its limitations!

1. Successes: Which problems AI is good at.
2. Limitations part A: Brittleness by cobbling together.
3. Limitations part B: The frontiers. Stuff that's just too difficult for now

1

2

3

1. Successes: Which problems deep learning is good at.

Conclusion:

Simple Input -> Output mappings!

“Learning” in ML/DL: collect many input/output examples,
black box algorithm (Neural Networks!)
will learn to predict the output

Image Recognition

1000 clas

1M image

Start of D

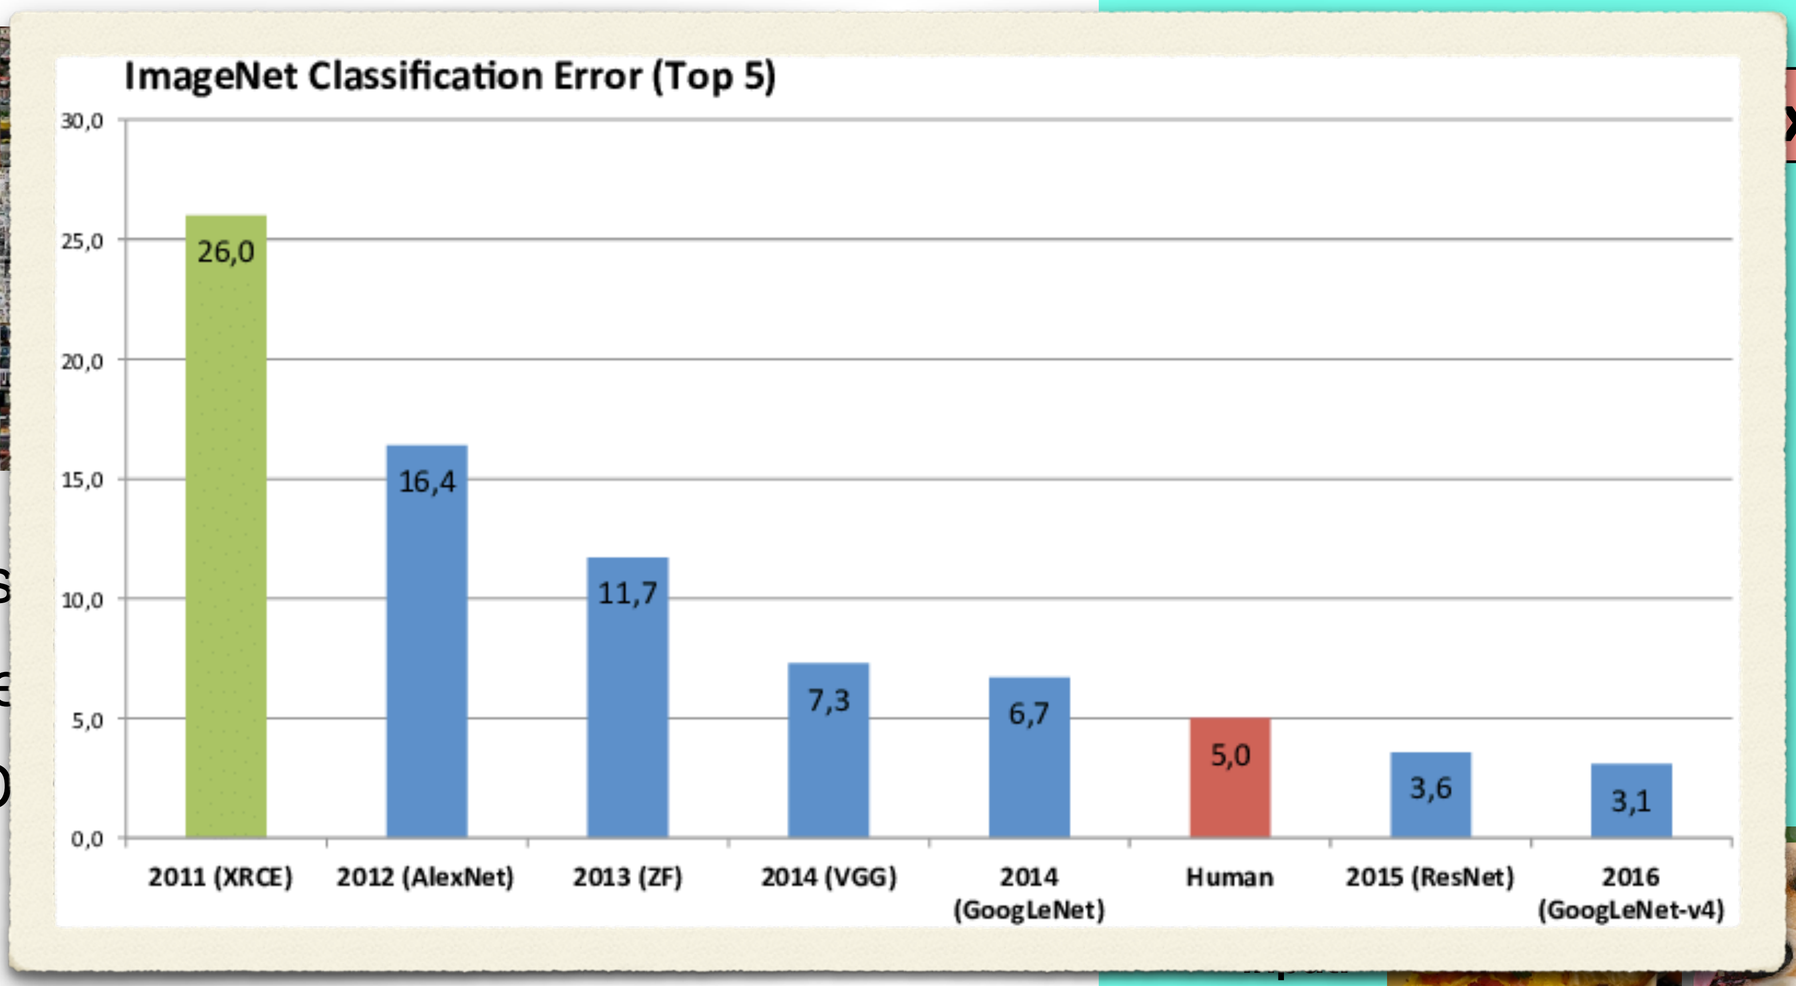

xhound

x 1M

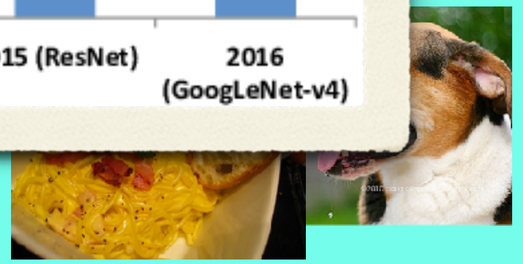

x 1M

Image search

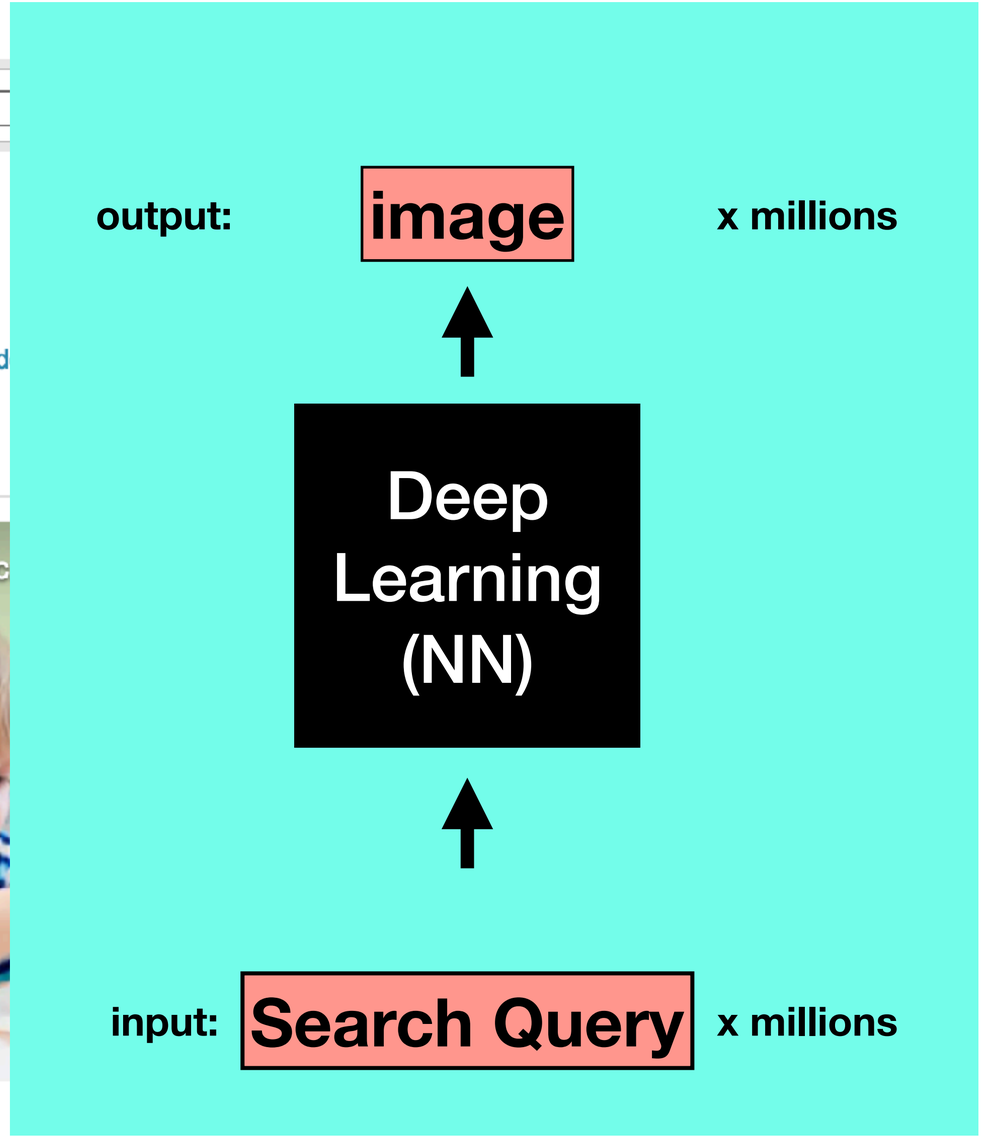

Speech Recognition

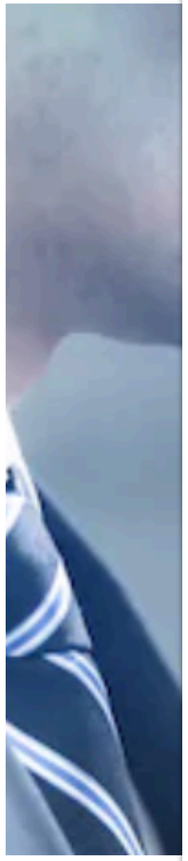

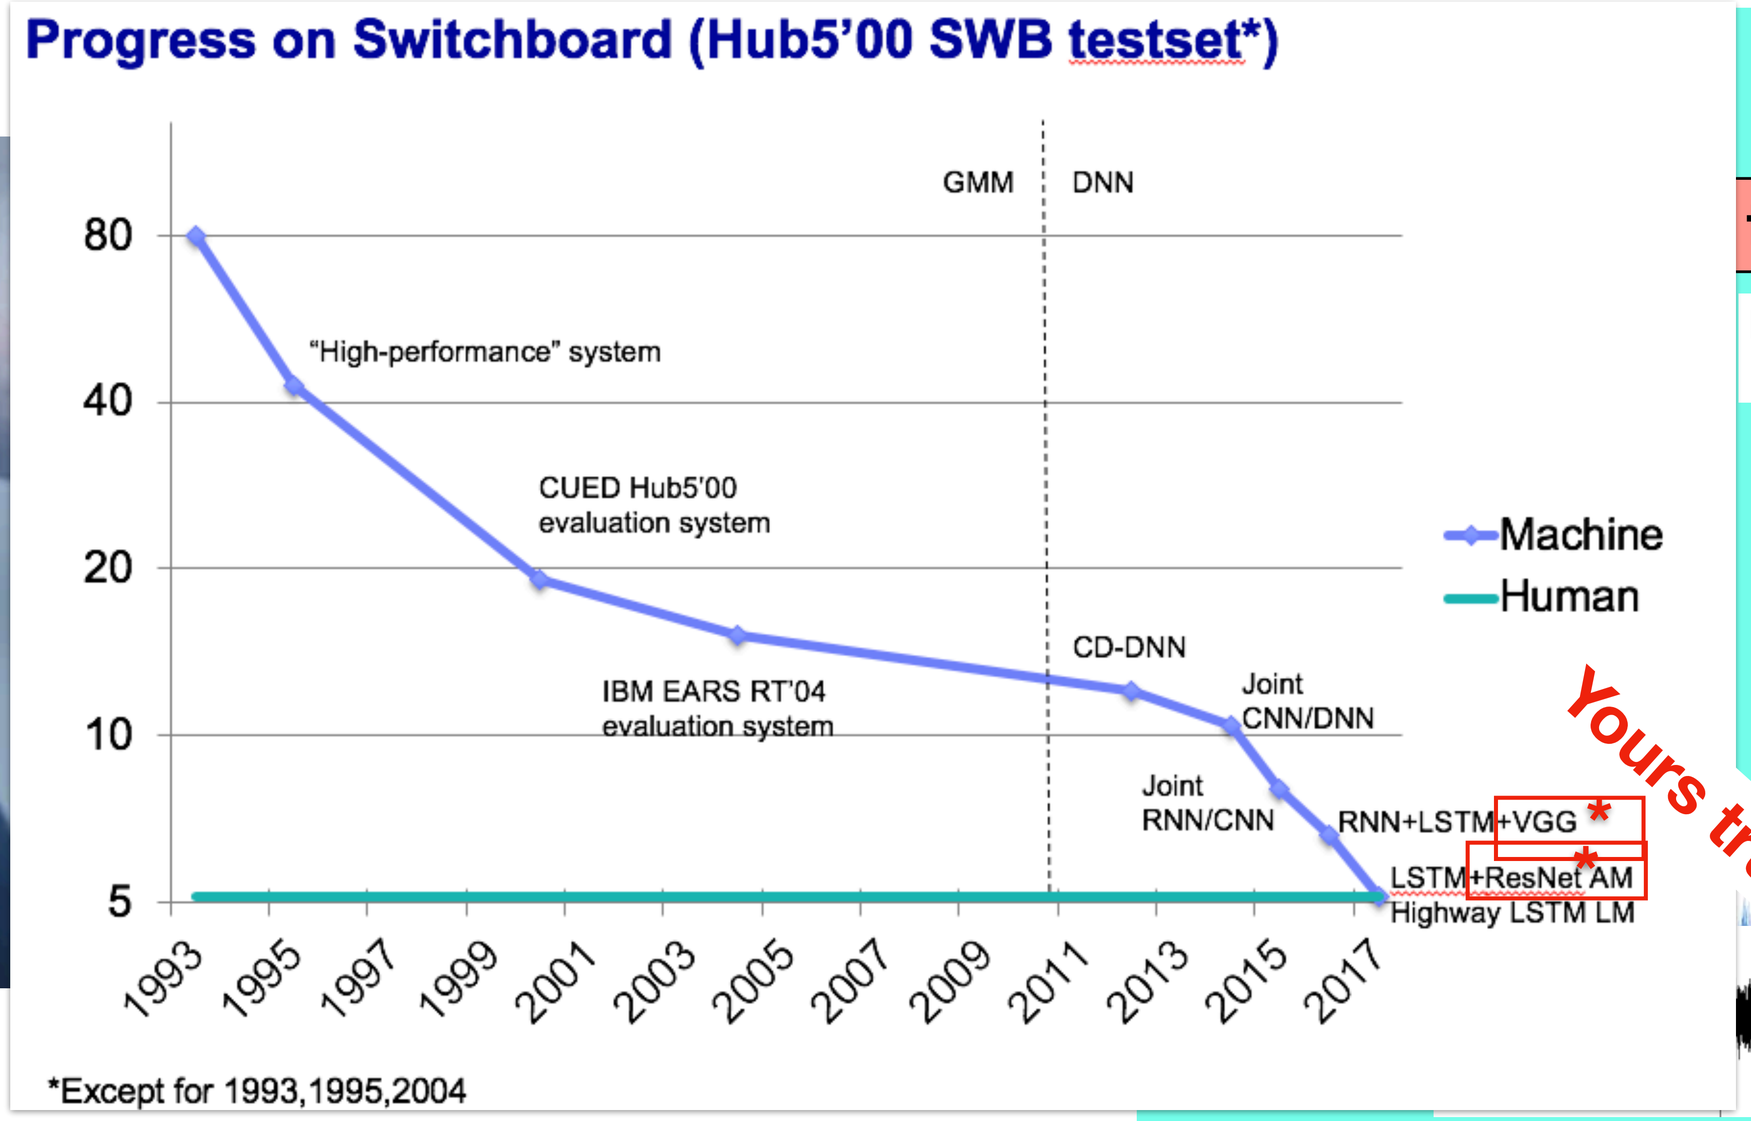

"whut"

the mat"

human transcriptions for every utterance

Yours truly

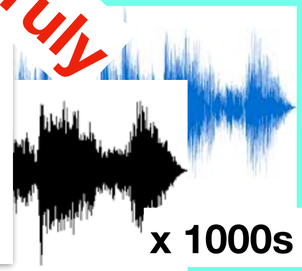

x 1000s of hours

Machine Translation

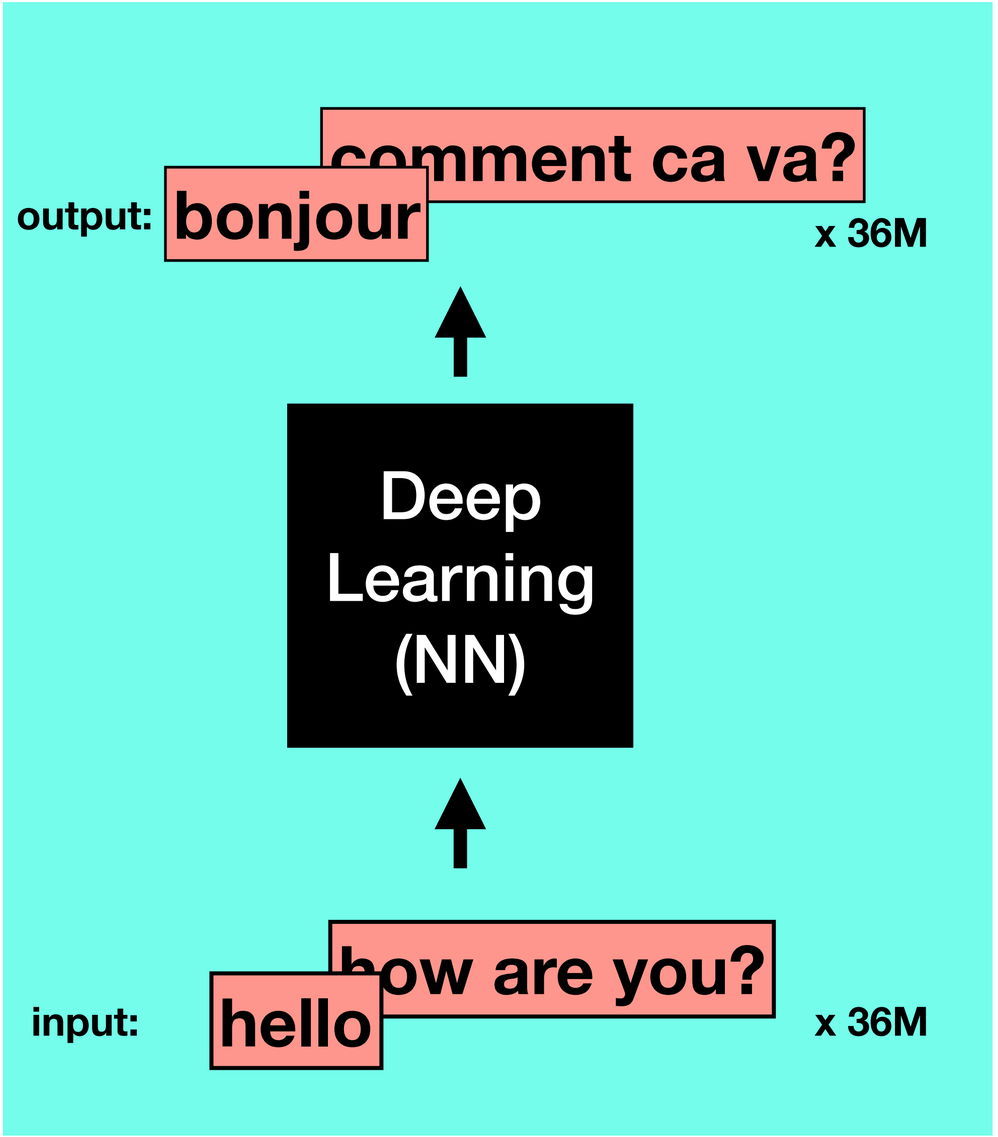

Main drivers:

1. Deep learning.
2. Data. Lots of data.
3. Compute power.
4. Open source software

but also

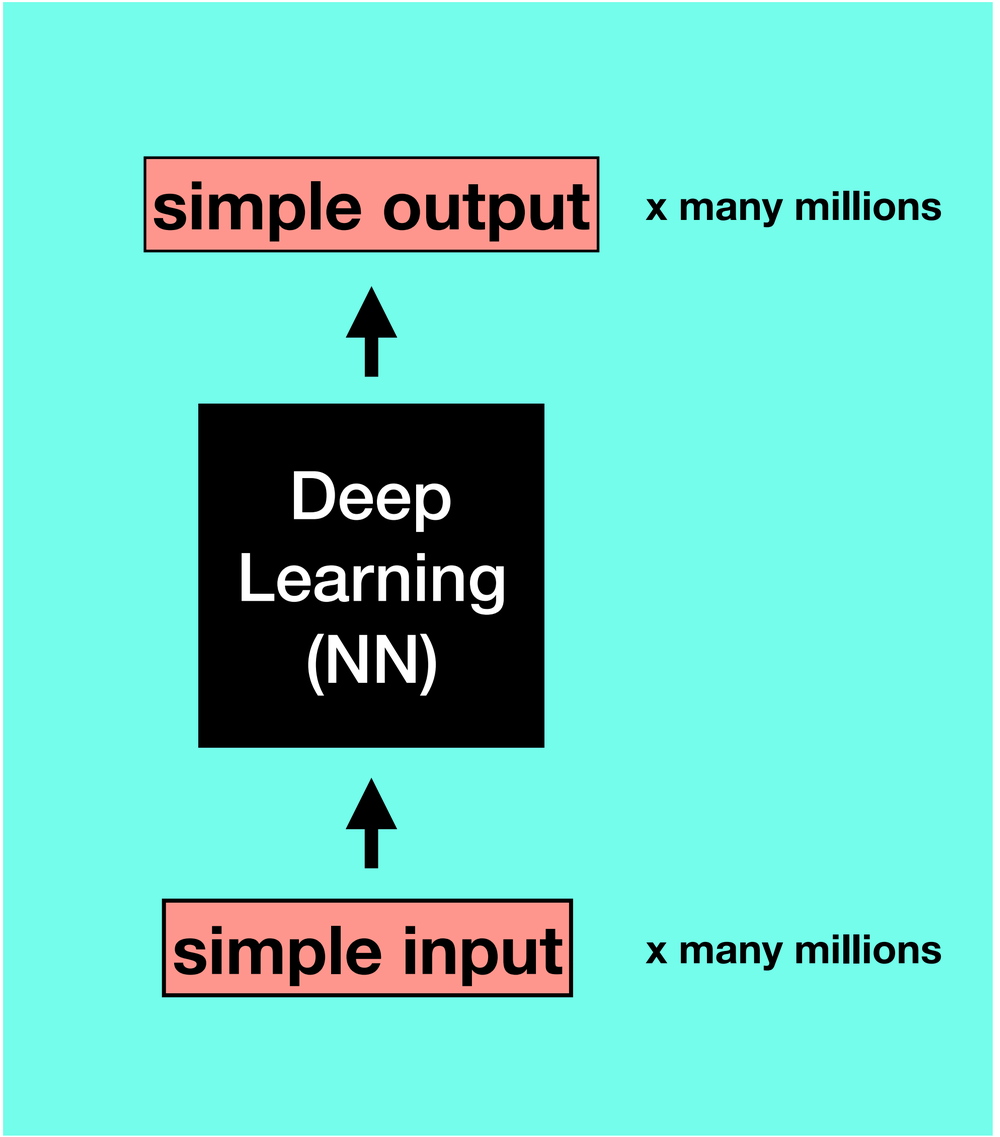

2. Limitations part A: ML Patchworks, a necessary evil

Conclusion:

Not every problem is this simple!

Most real problems require cobbling together AI systems:
a) makes it work
b) makes it brittle

Decision: investing in a patchwork ML solution

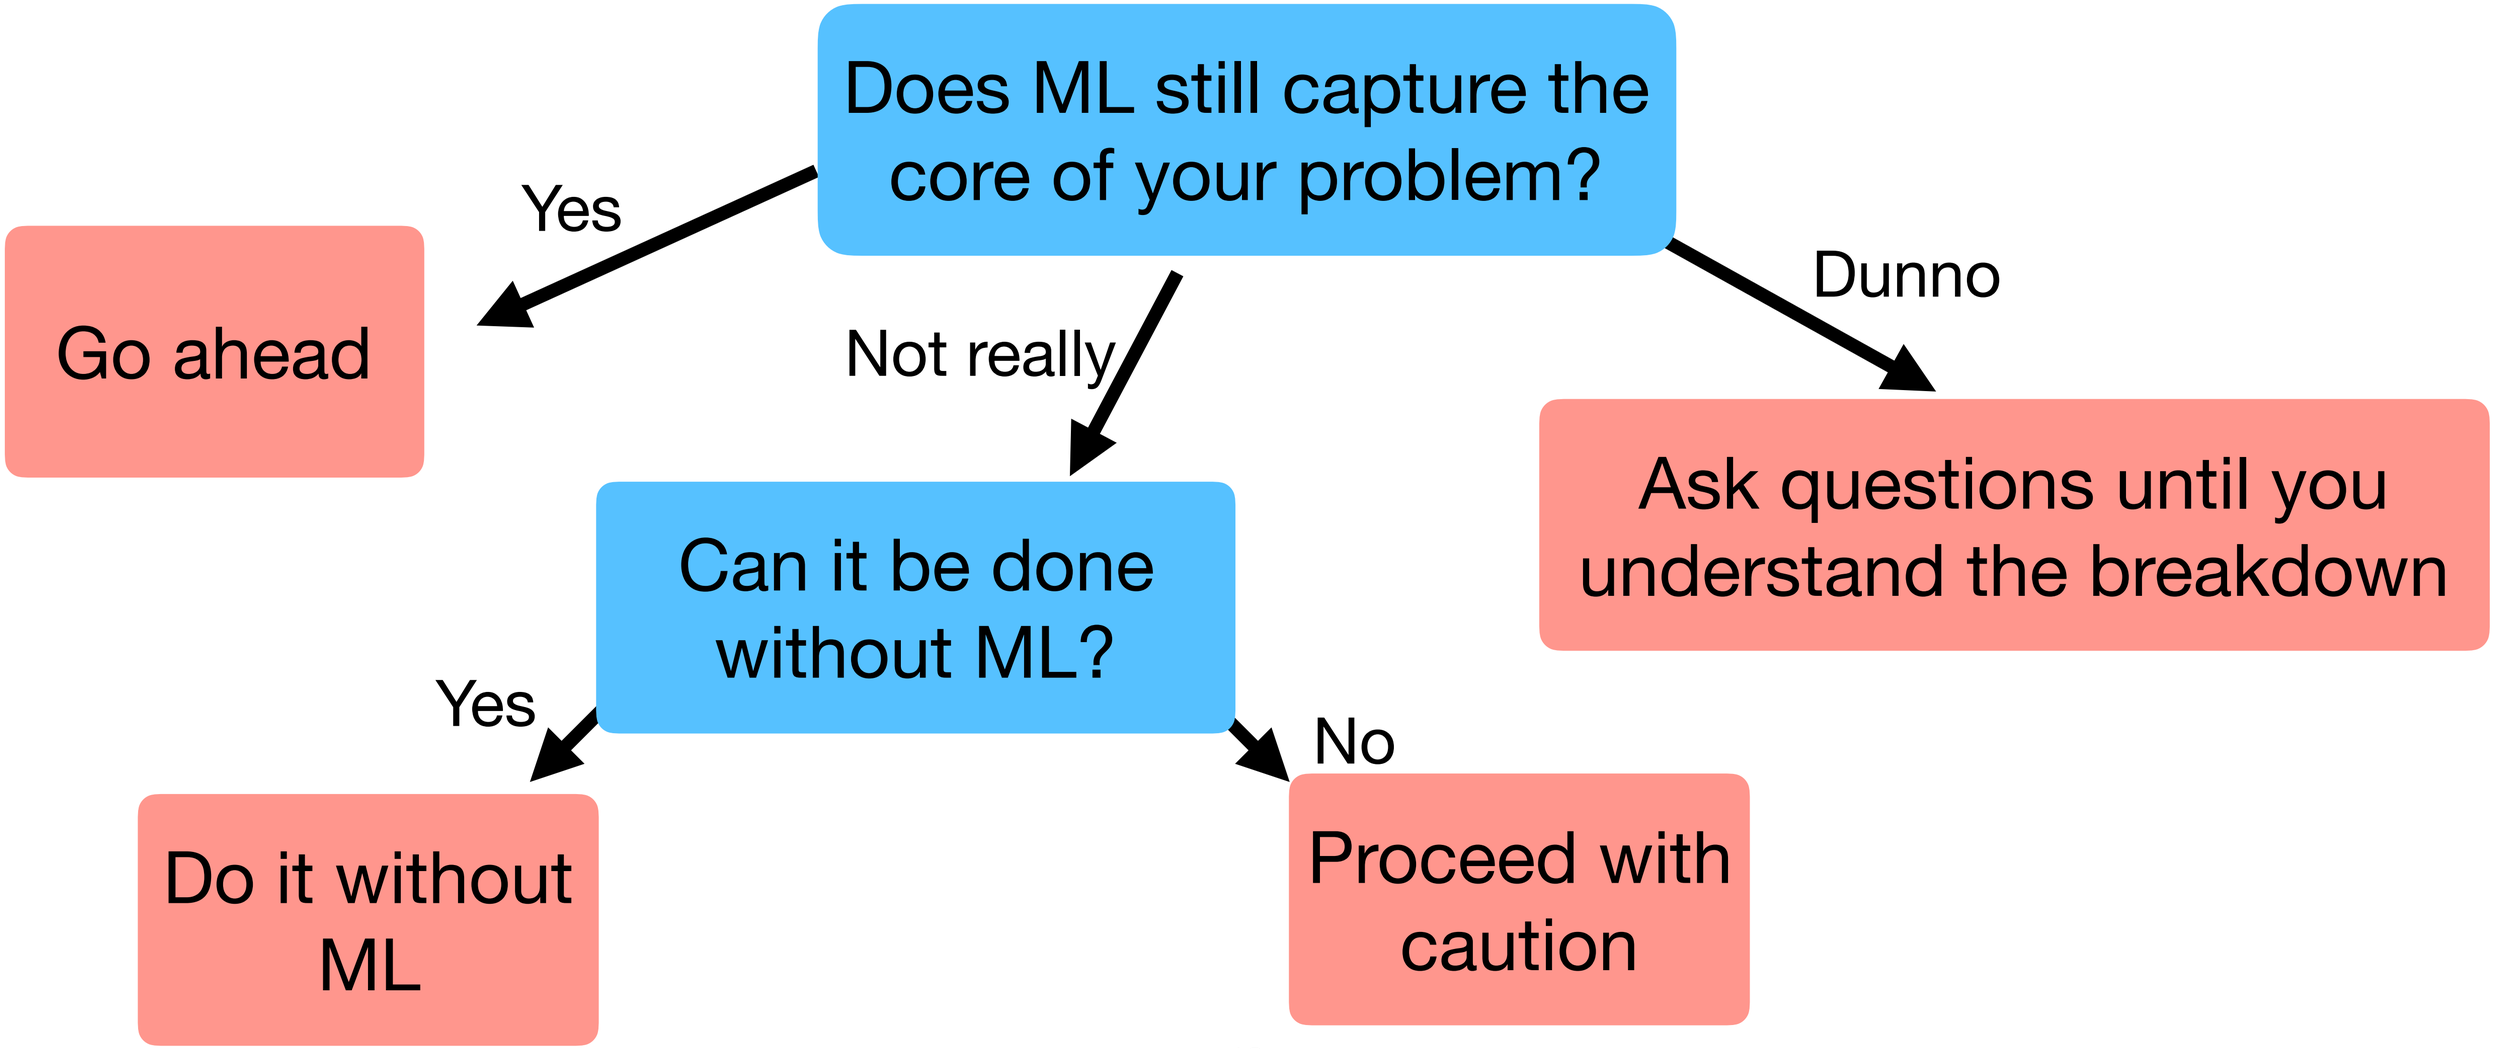

3. Limitations part B: Frontiers. Stuff that's just too hard

Conclusion:

Don't believe anyone who sells you "an AI that understands language and reasons about a problem from common sense"

Will consist of simpler components cobbled together or just won't deliver on the promise.

3. Limitations B: Frontiers

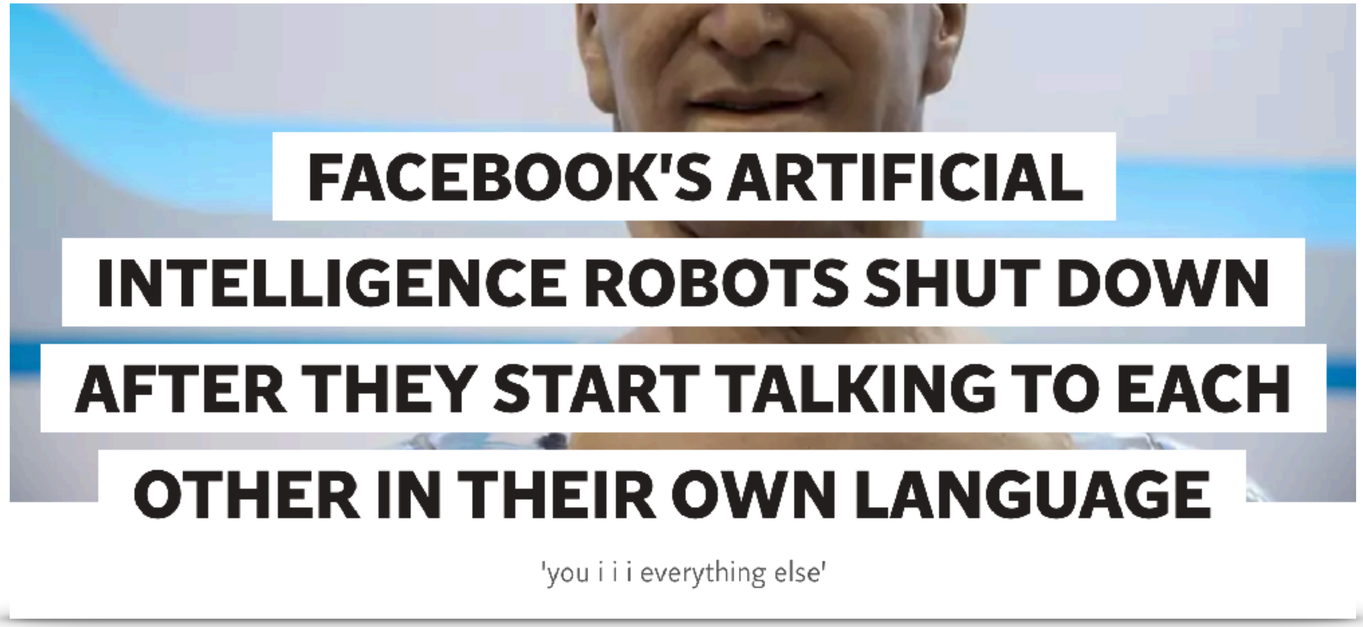

Natural Language Understanding

disambiguition:

(NL: bank)

(EN: bank)

dialogue

winograd schemes:

The trophy doesn't fit into the brown suitcase because it's too [small/large].

What is too [small/large]?

Reasoning

How many slices of pizza are there?
Is this a vegetarian pizza?

Multi agent games

Does it appear to be rainy?
Does this person have 20/20 vision?

Visual Question Answering VQA

Set of simple reasoning tasks

Task 15: Basic Deduction

Sheep are afraid of wolves.
Cats are afraid of dogs.
Mice are afraid of cats.
Gertrude is a sheep.
What is Gertrude afraid of? A:wolves

Task 16: Basic Induction

Lily is a swan.
Lily is white.
Bernhard is green.
Greg is a swan.
What color is Greg? A:white

Only limited scope!

Requires heavy engineering

Learning from less (labeled) data

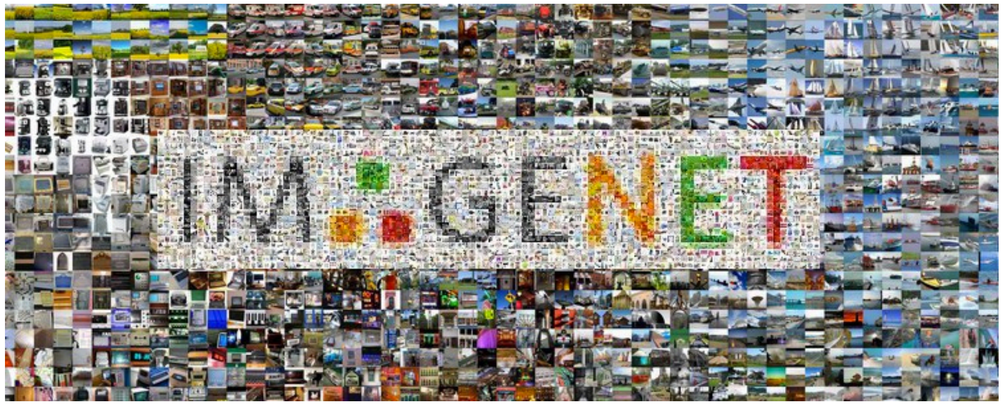

1M labeled images

vs

continuous visual stream, mostly UNLABELED
cuddles as reward signal for recognizing papa

Causality

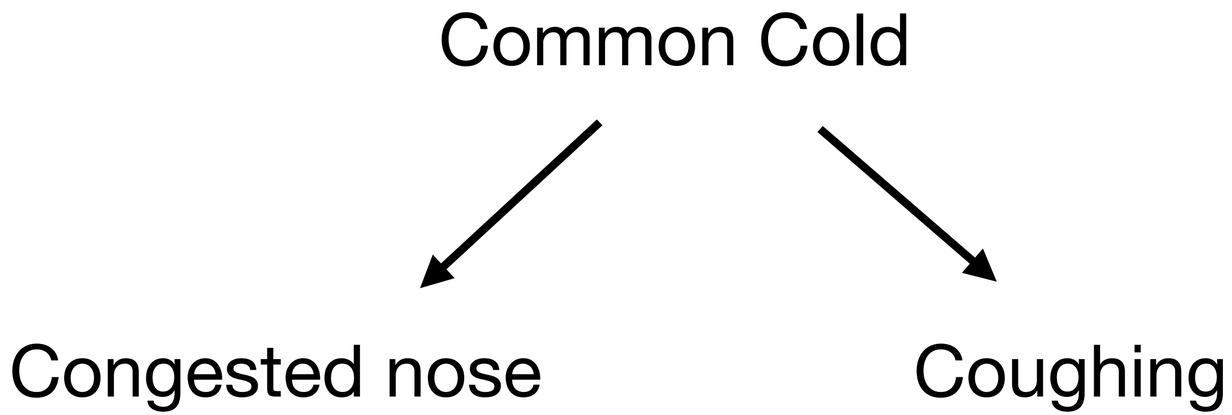

Common Sense

Imagine what happens if you open your hand:

vs

Common Sense

- Remembering and accessing knowledge
- Provides prior knowledge in new problem
- Intuitive reasoning (by analogy)
- Easily infer causal direction
- Reduces data hungriness

The holy grail

I hope now you have a

better feeling and understanding of AI and its limitations!

1. Successes: Which problems AI / ML / DL is good at: **simple Input -> Output problems!**
2. Limitations part A: Brittleness by cobbling together.
3. Limitations part B: The frontiers. Stuff that's just too difficult for now

1

2

3

This was part 1.
Now part 2, the real thing.

Guest lecture
Intro to Data Science, Fall 2018 @ CCNY
December 3d 2018, New York
Tom Sercu