

Machine Learning: successes, promises and limits

AI Academy Seminar
13 November 2018, Kortrijk
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



I am NOT

- Representing IBM (opinions are my own)
- Selling you anything
 - Except truth!
- Funny.
 - but I'll try

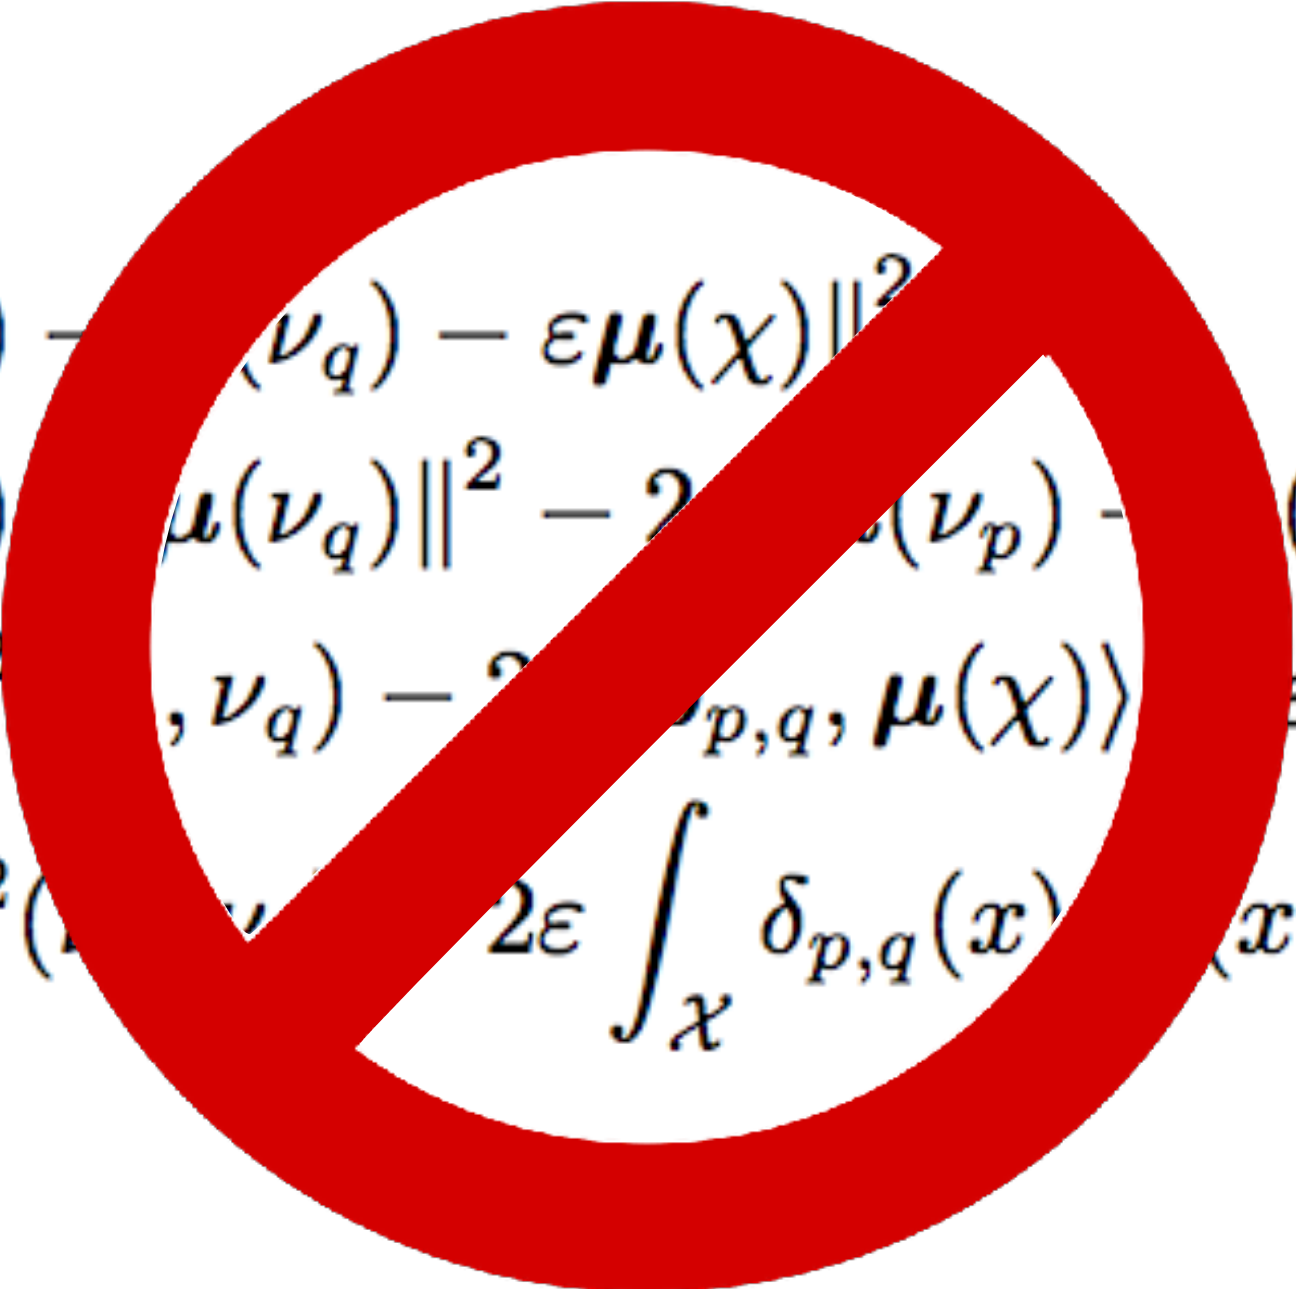


Many companies throw money at AI, but few get real returns.



My goal is for you to:

Build better feeling and understanding of AI and its limitations!


$$\begin{aligned} &= \|\mu(\nu_p) - \mu(\nu_q) - \varepsilon\mu(\chi)\|^2 \\ &= \|\mu(\nu_p) - \mu(\nu_q)\|^2 - 2\varepsilon\langle\mu(\nu_p) - \mu(\nu_q), \mu(\chi)\rangle + \varepsilon^2 \\ &= \text{MMD}^2(\nu_p, \nu_q) - 2\varepsilon\langle\mu(\nu_p) - \mu(\nu_q), \mu(\chi)\rangle + \varepsilon^2\|\mu(\chi)\|^2 \\ &= \text{MMD}^2(\nu_p, \nu_q) + 2\varepsilon\int_{\mathcal{X}}\delta_{p,q}(x)\mu(x) + \varepsilon^2\|\mu(\chi)\|^2 \end{aligned}$$



My goal is for you to:

Build better feeling and understanding of AI 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 AI (machine learning / deep learning) is good at.



Conclusion:

Simple Input -> Output mappings!

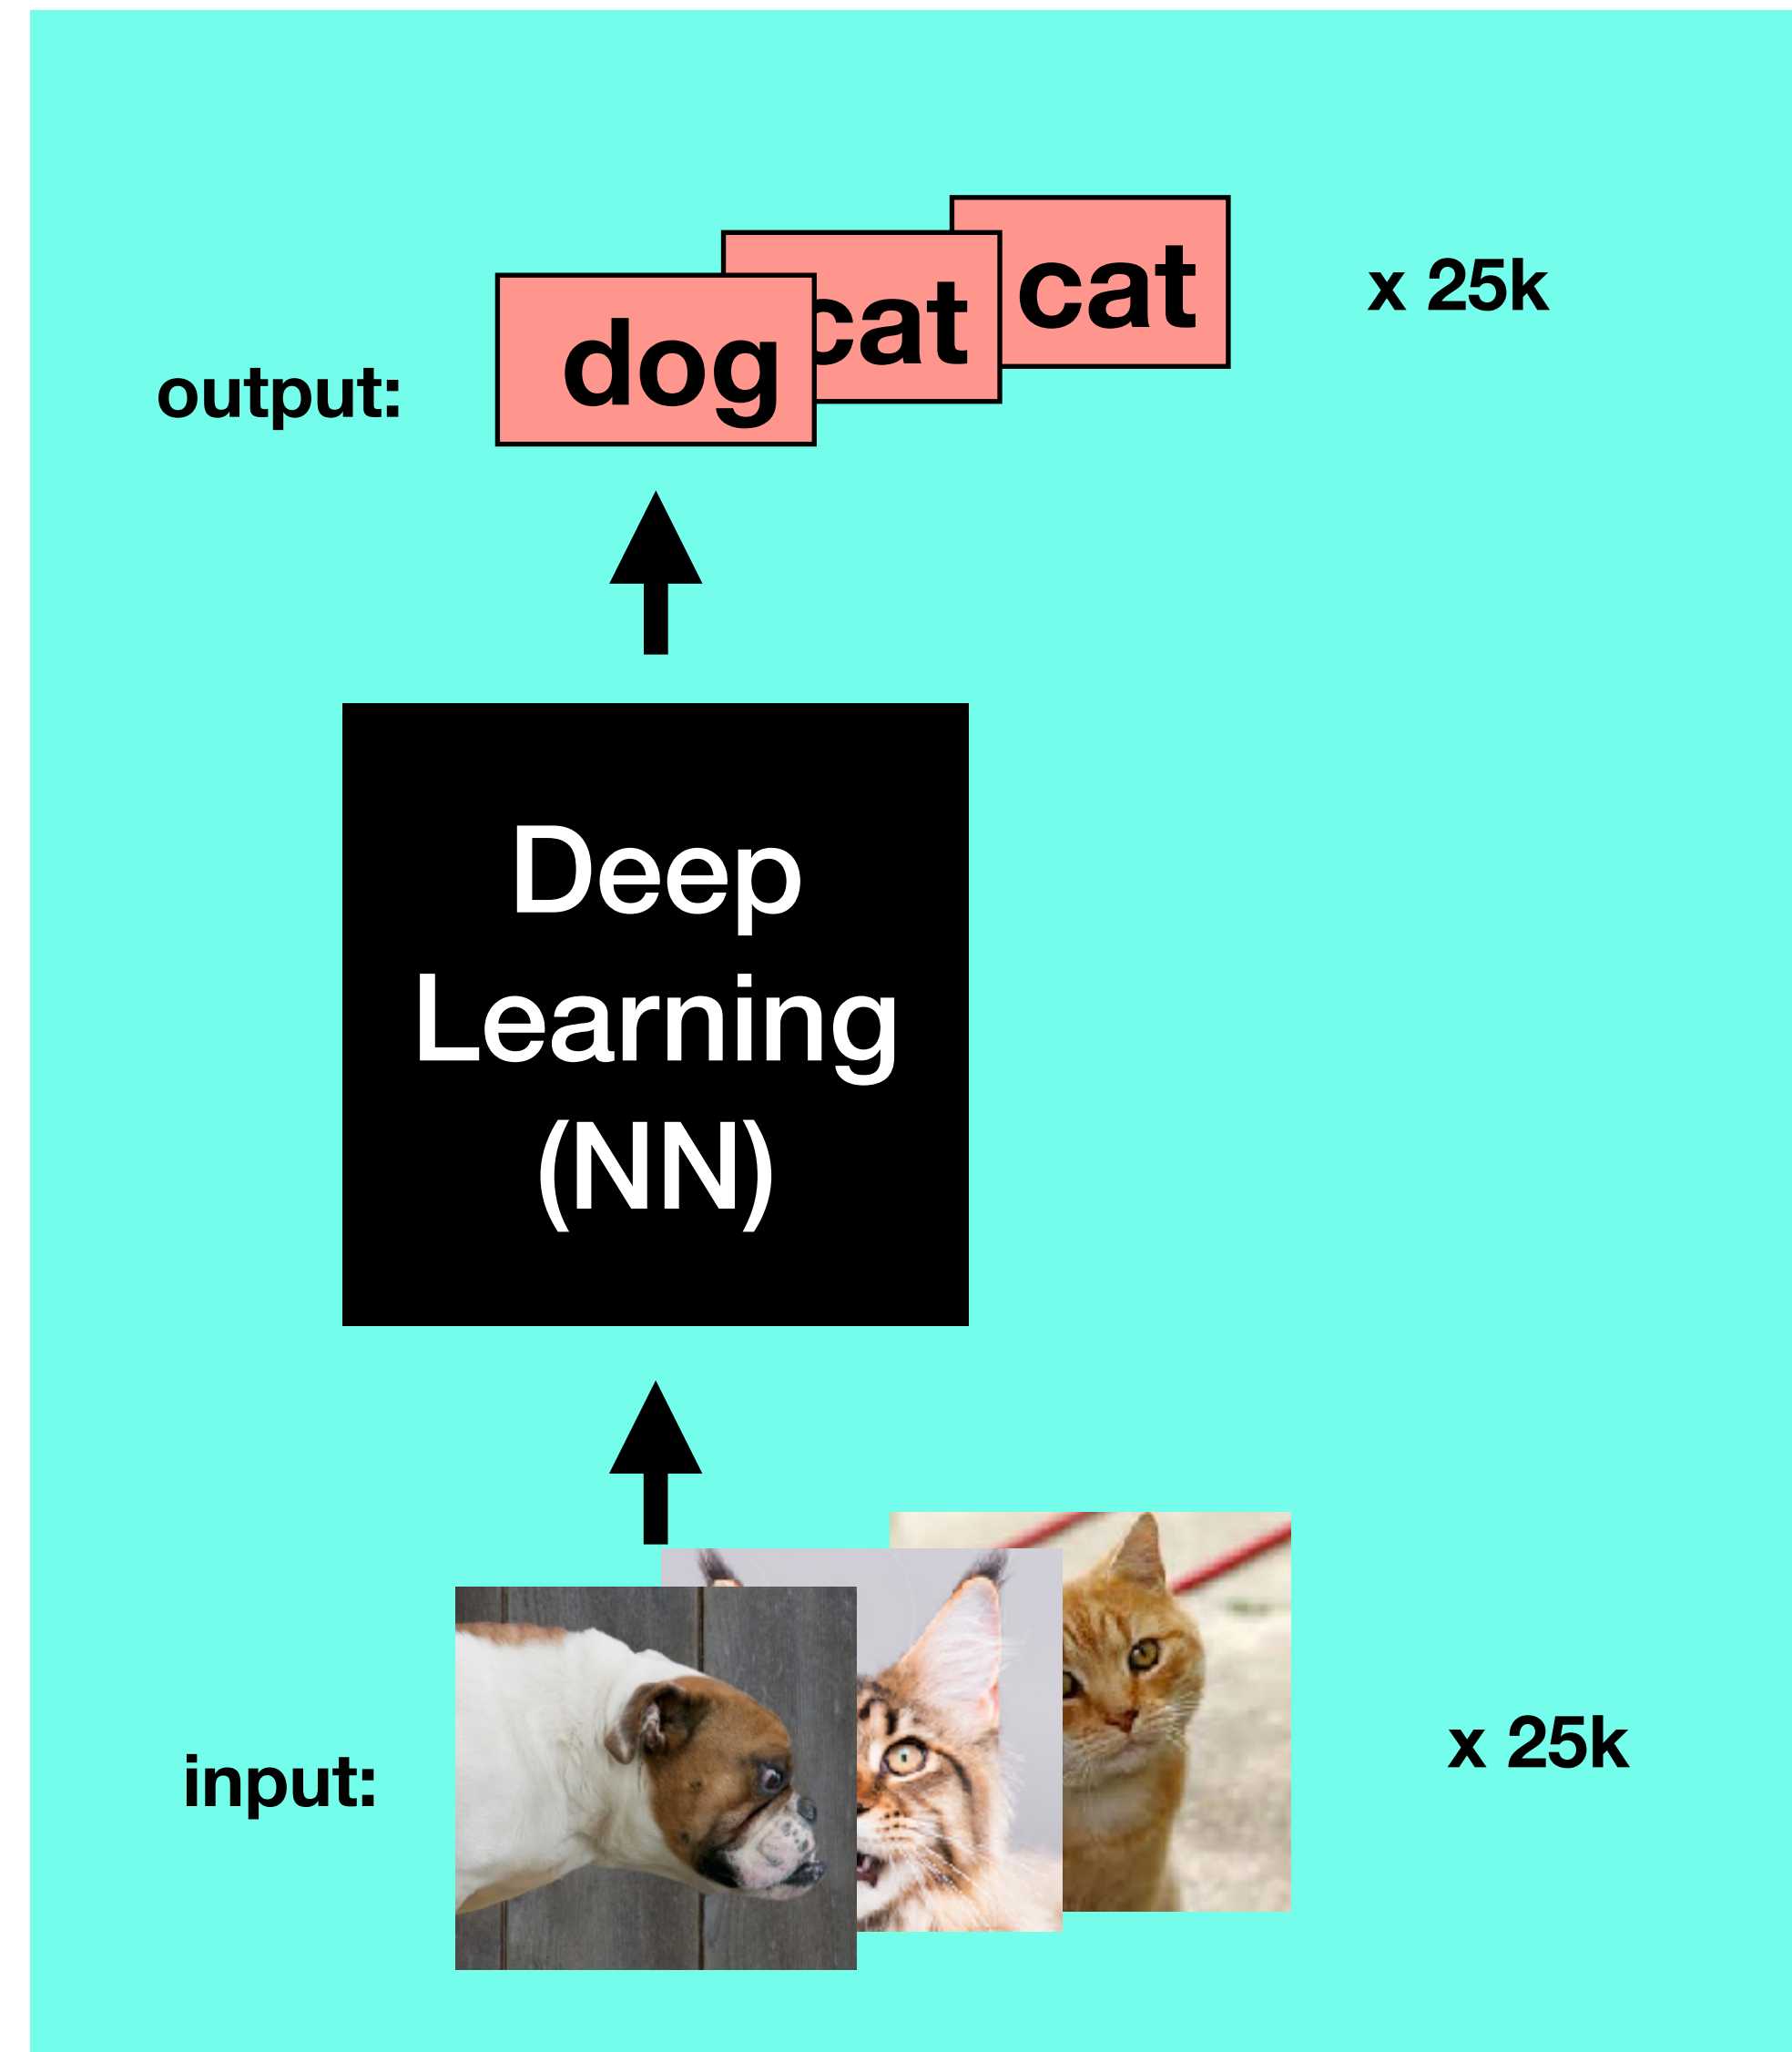
“Learning” in ML/DL: collect many input/output examples,
black box algorithm (Neural Networks!)
will learn to predict the output

Image Recognition

DOG



CAT




1. Successes

Image Recognition



1000 classes
1M images
Start of D

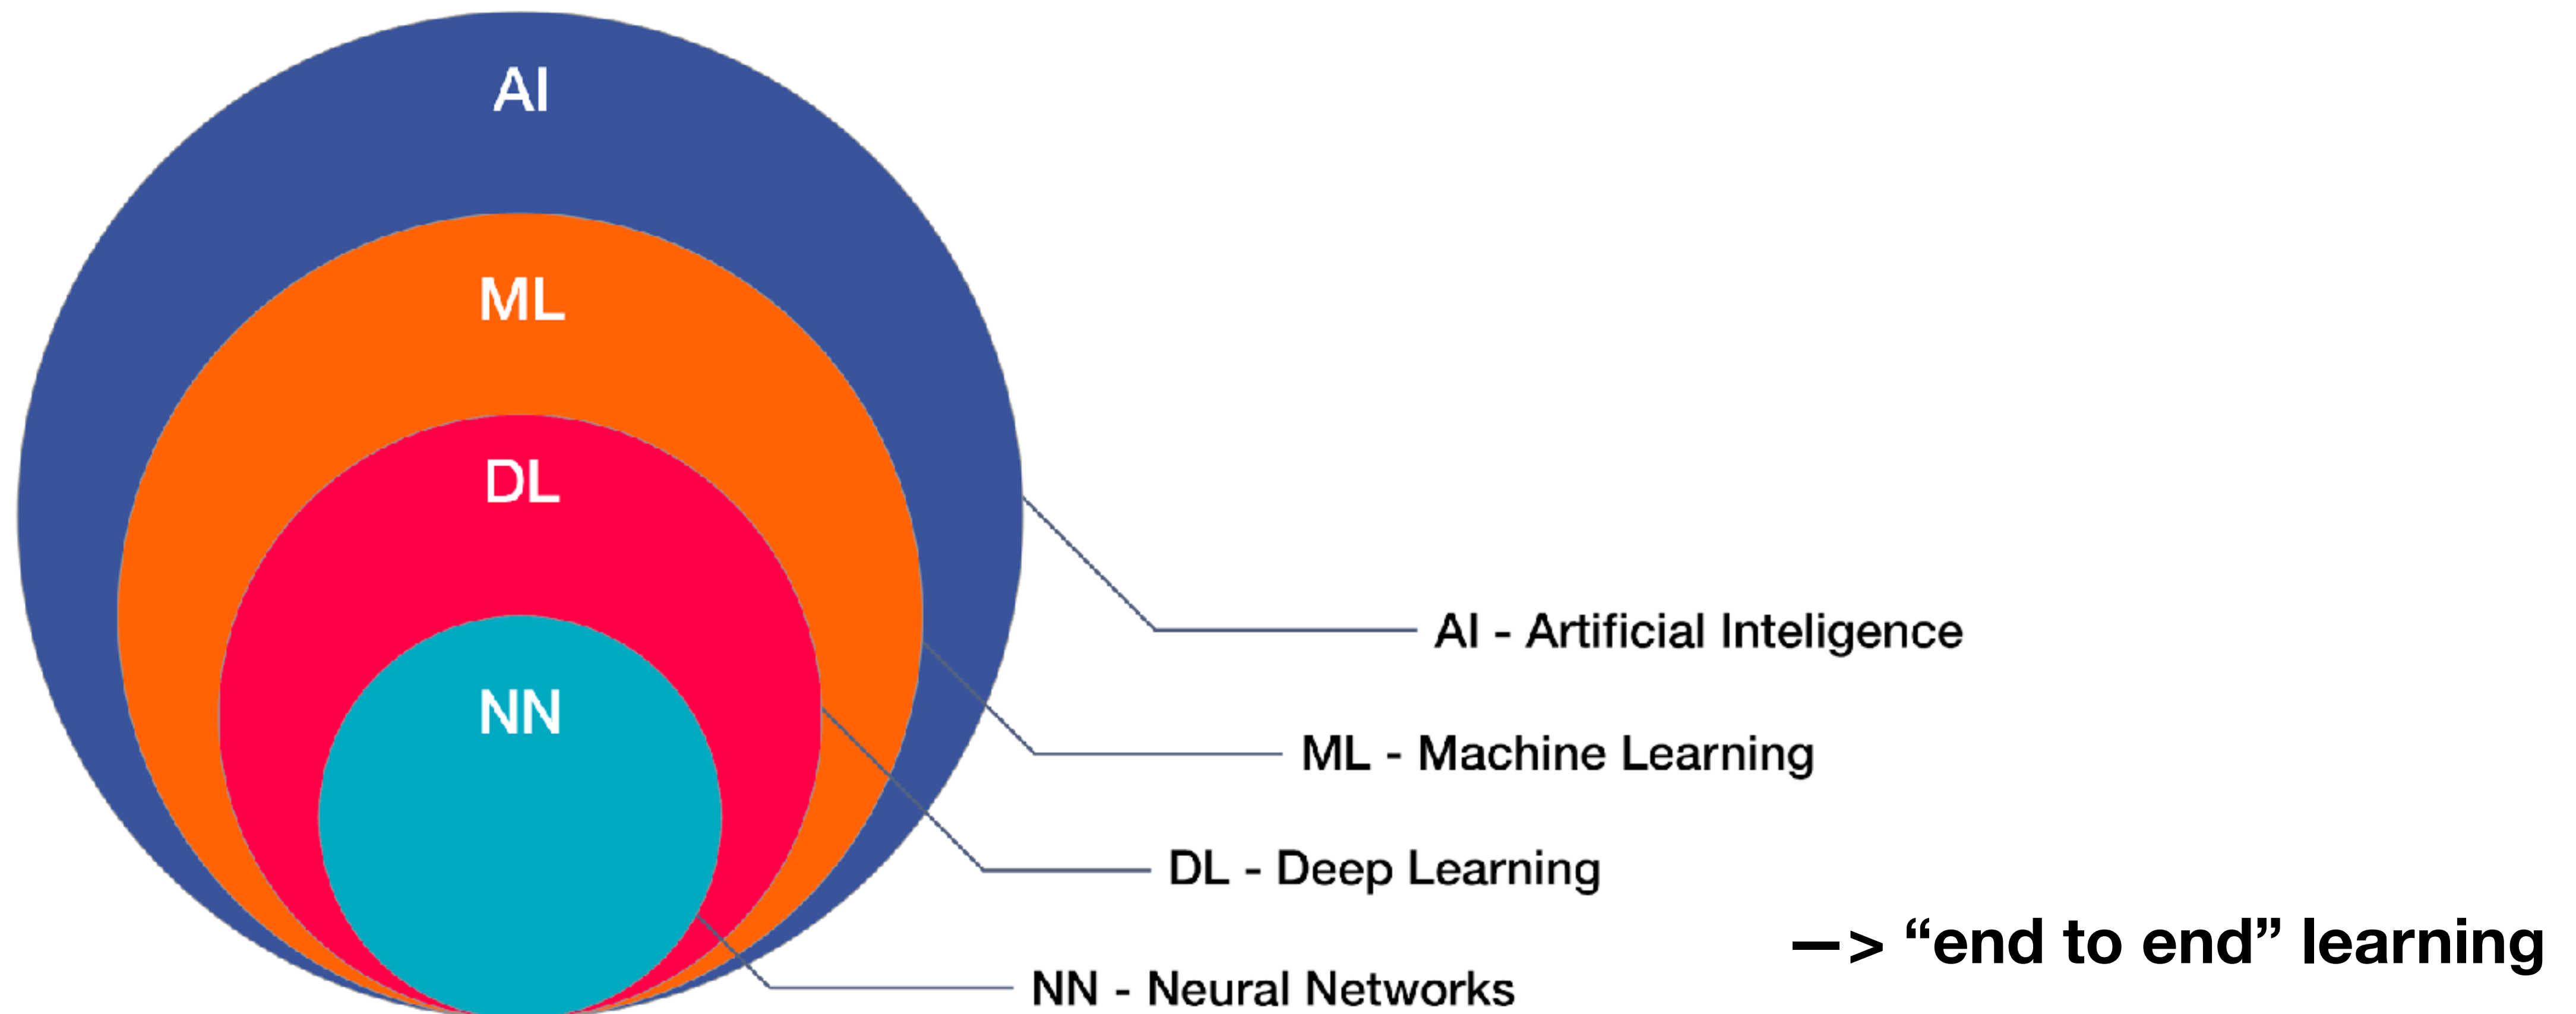


xhound

x 1M

x 1M

Intermezzo: AI > ML > DL



1. Successes

Image search

shutterstock

toddler doing crafts

Toddler doing crafts stock photos

505 Toddler doing crafts stock photos, vectors, and illustrations are available royalty-free. See [toddler d](#)

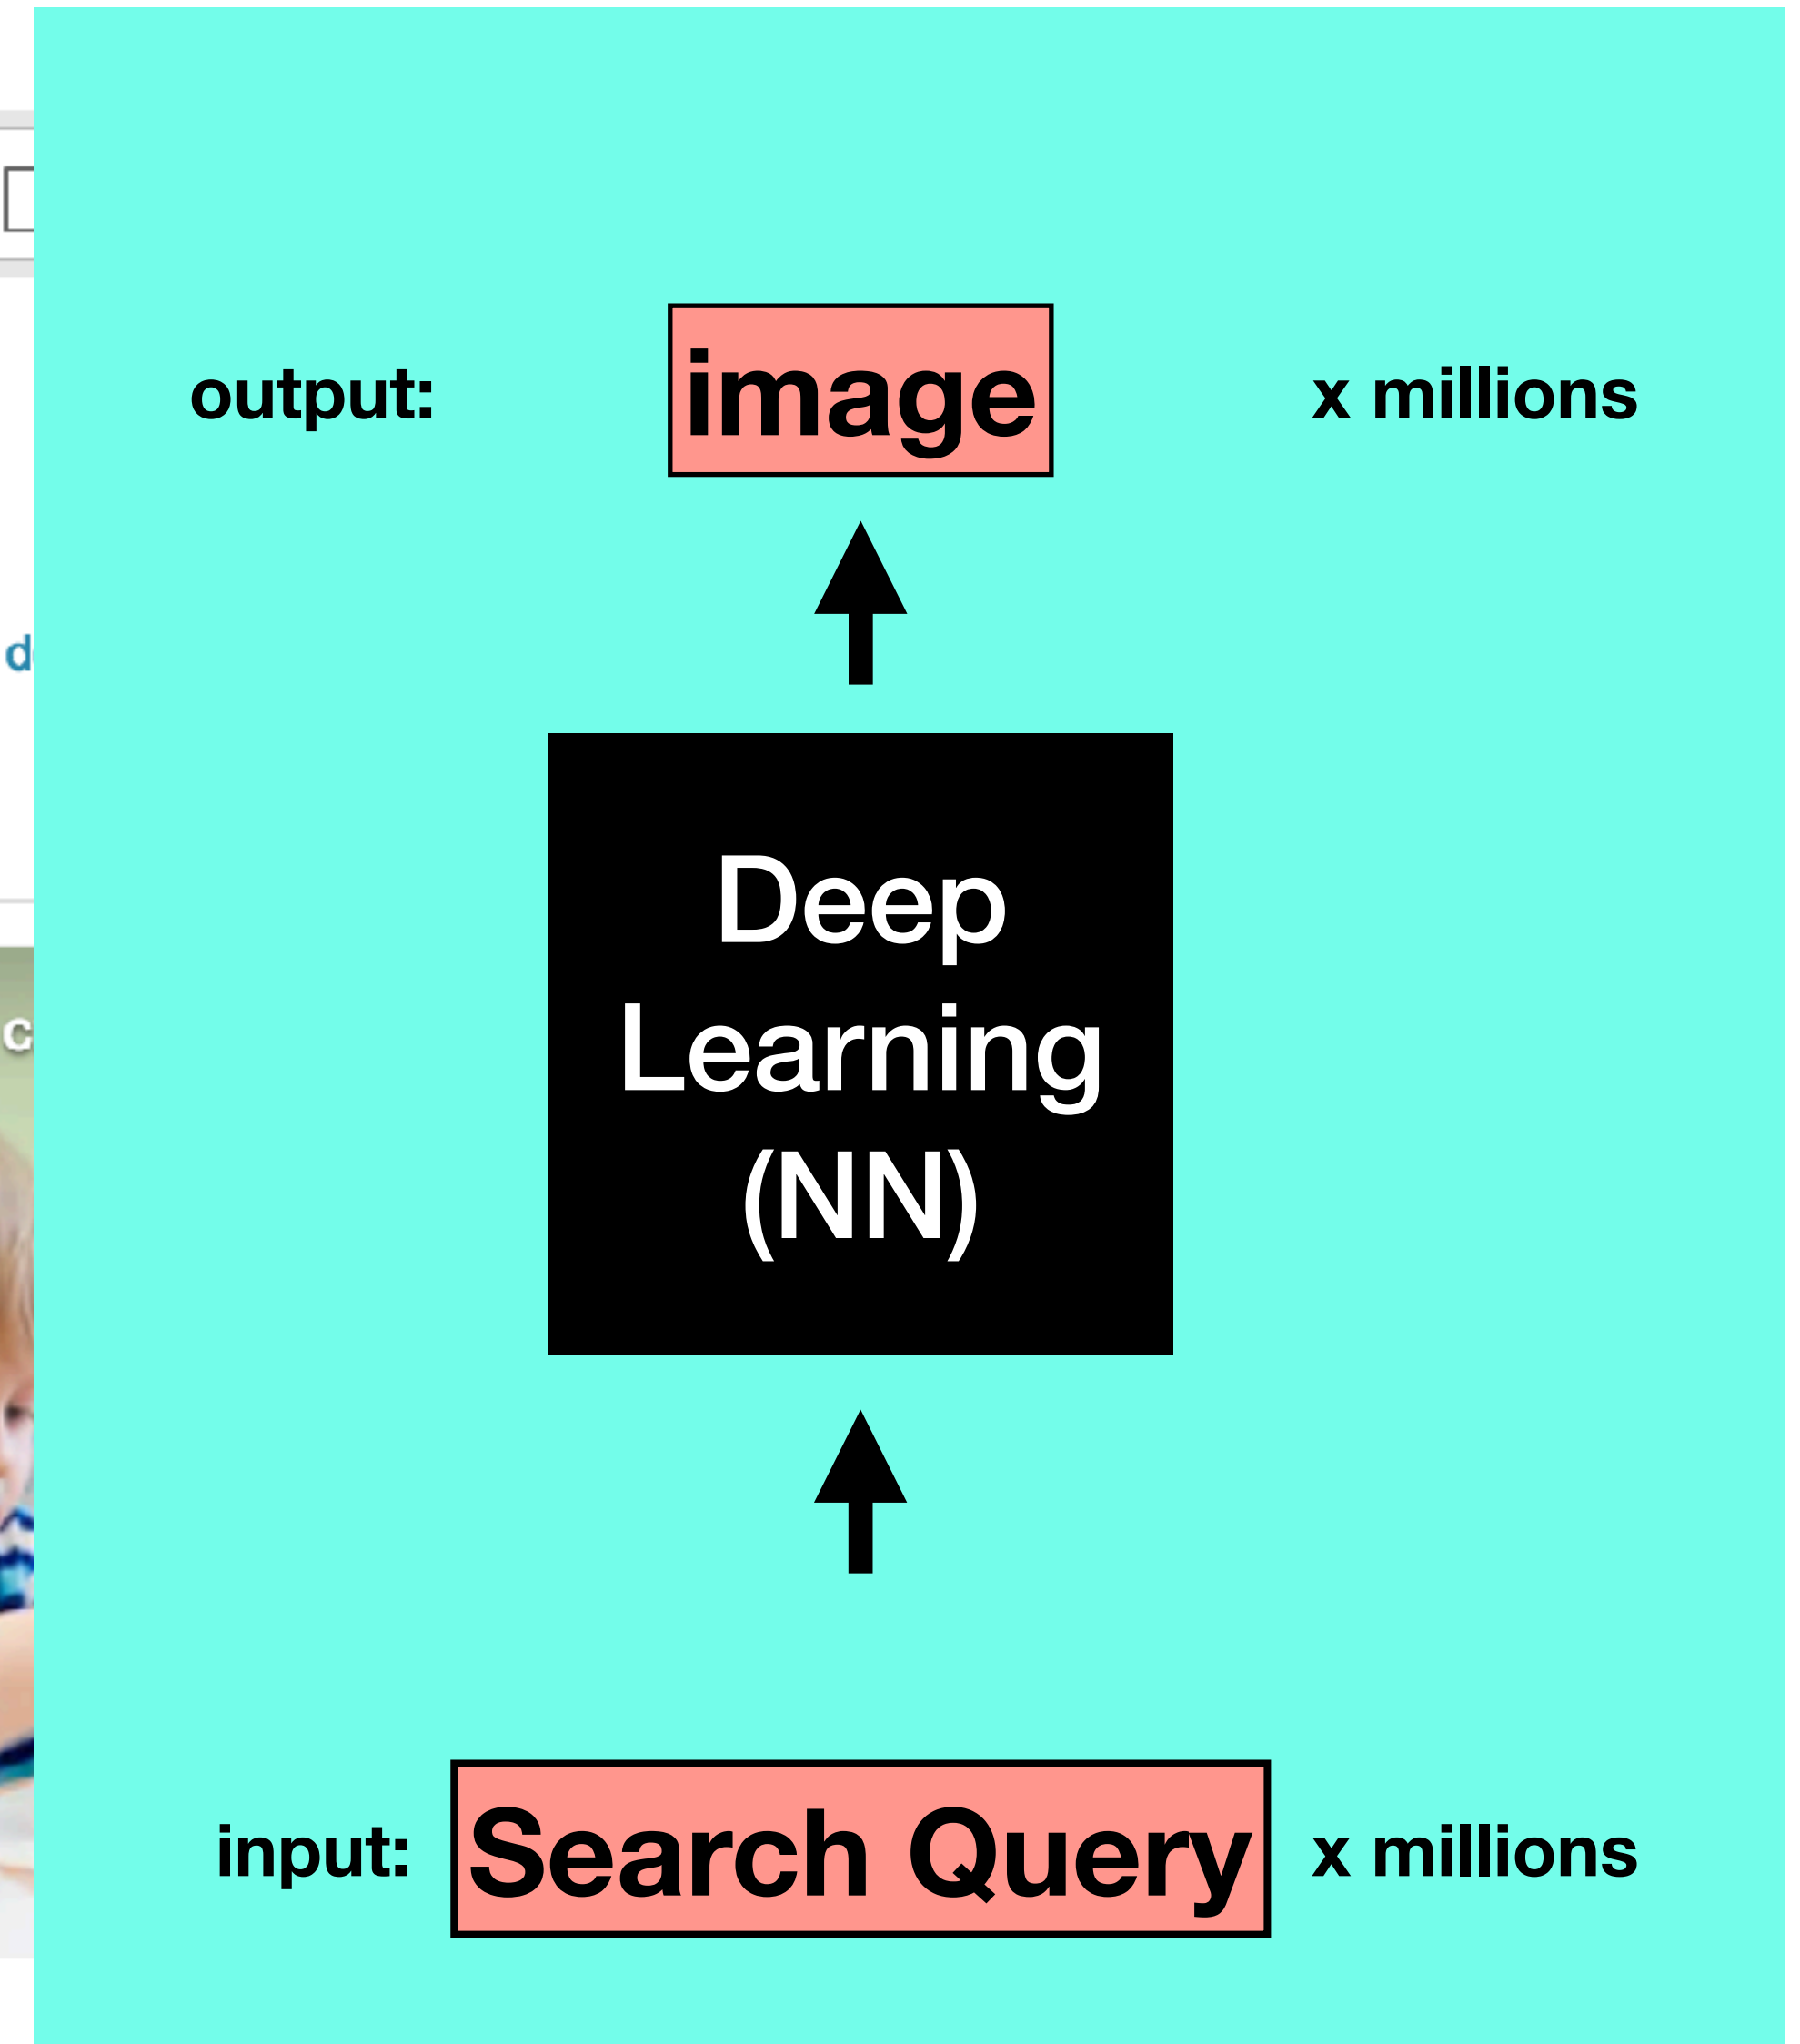
Most Relevant

Fresh Content

Image Type ▾

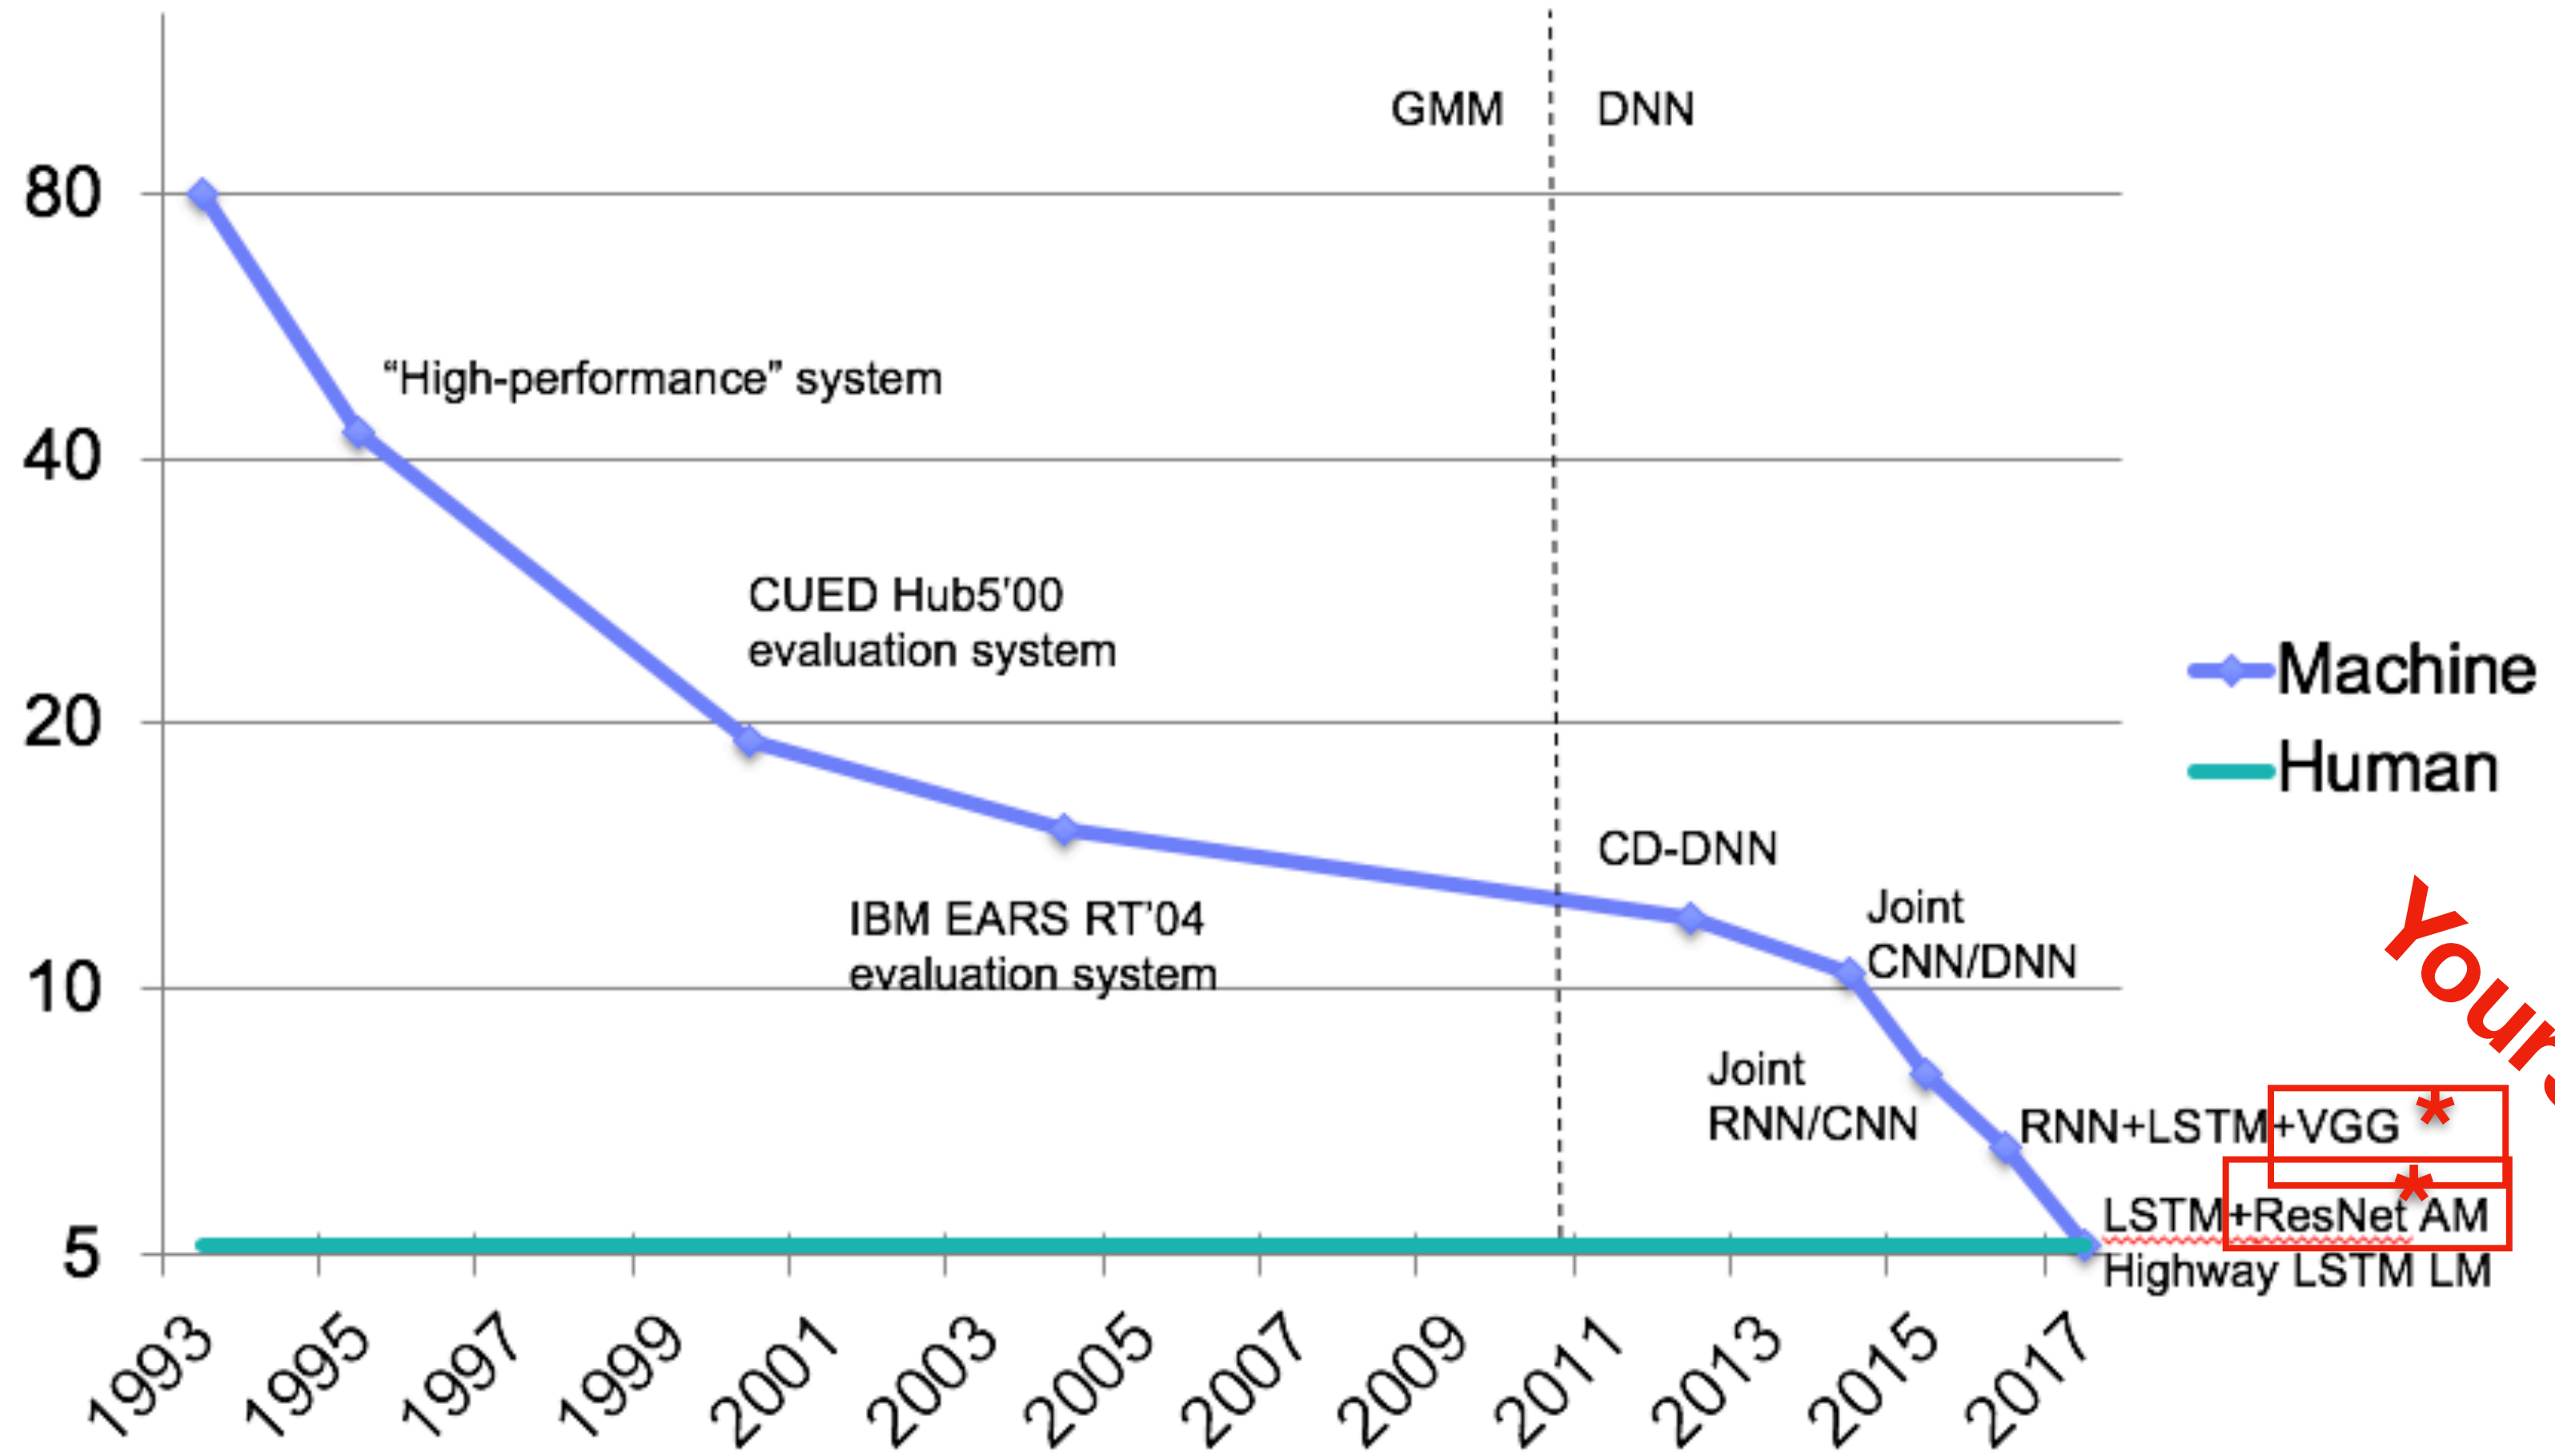
Orientation ▾

Color ▾



Speech Recognition

Progress on Switchboard (Hub5'00 SWB testset*)




*Except for 1993,1995,2004

“whut”
the mat”

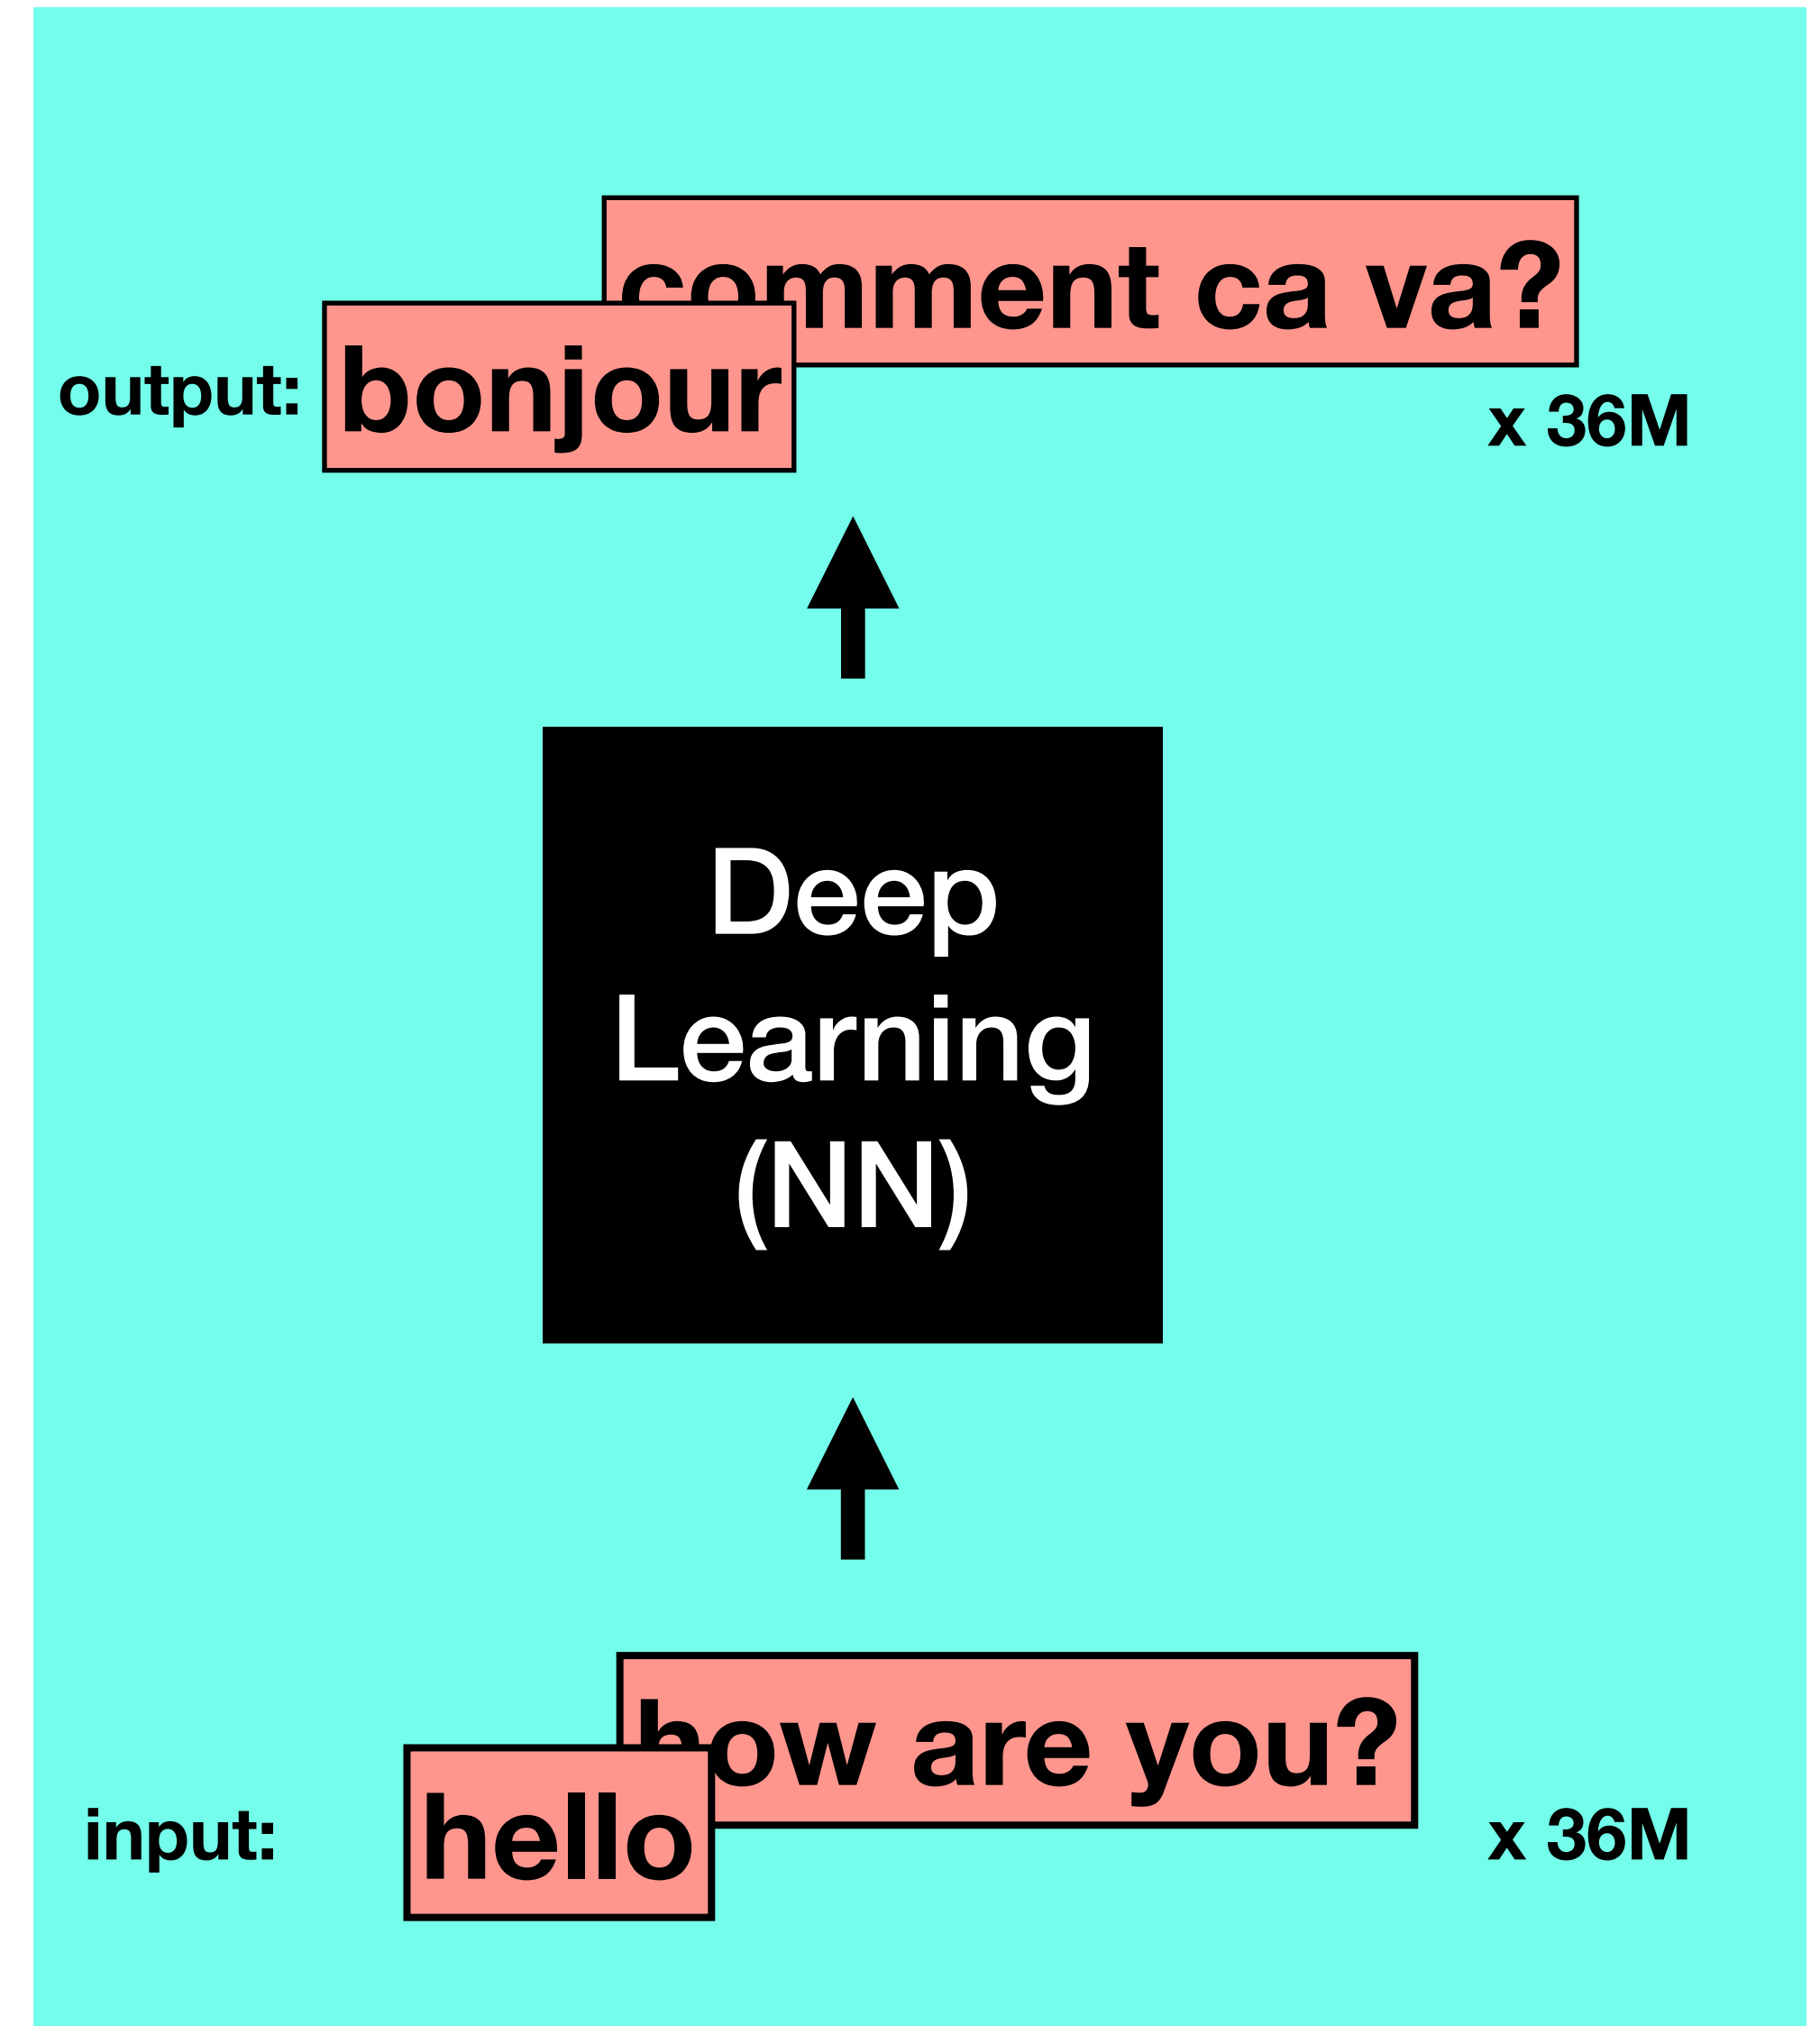
human transcriptions
for every utterance

Yours truly



x 1000s of hours

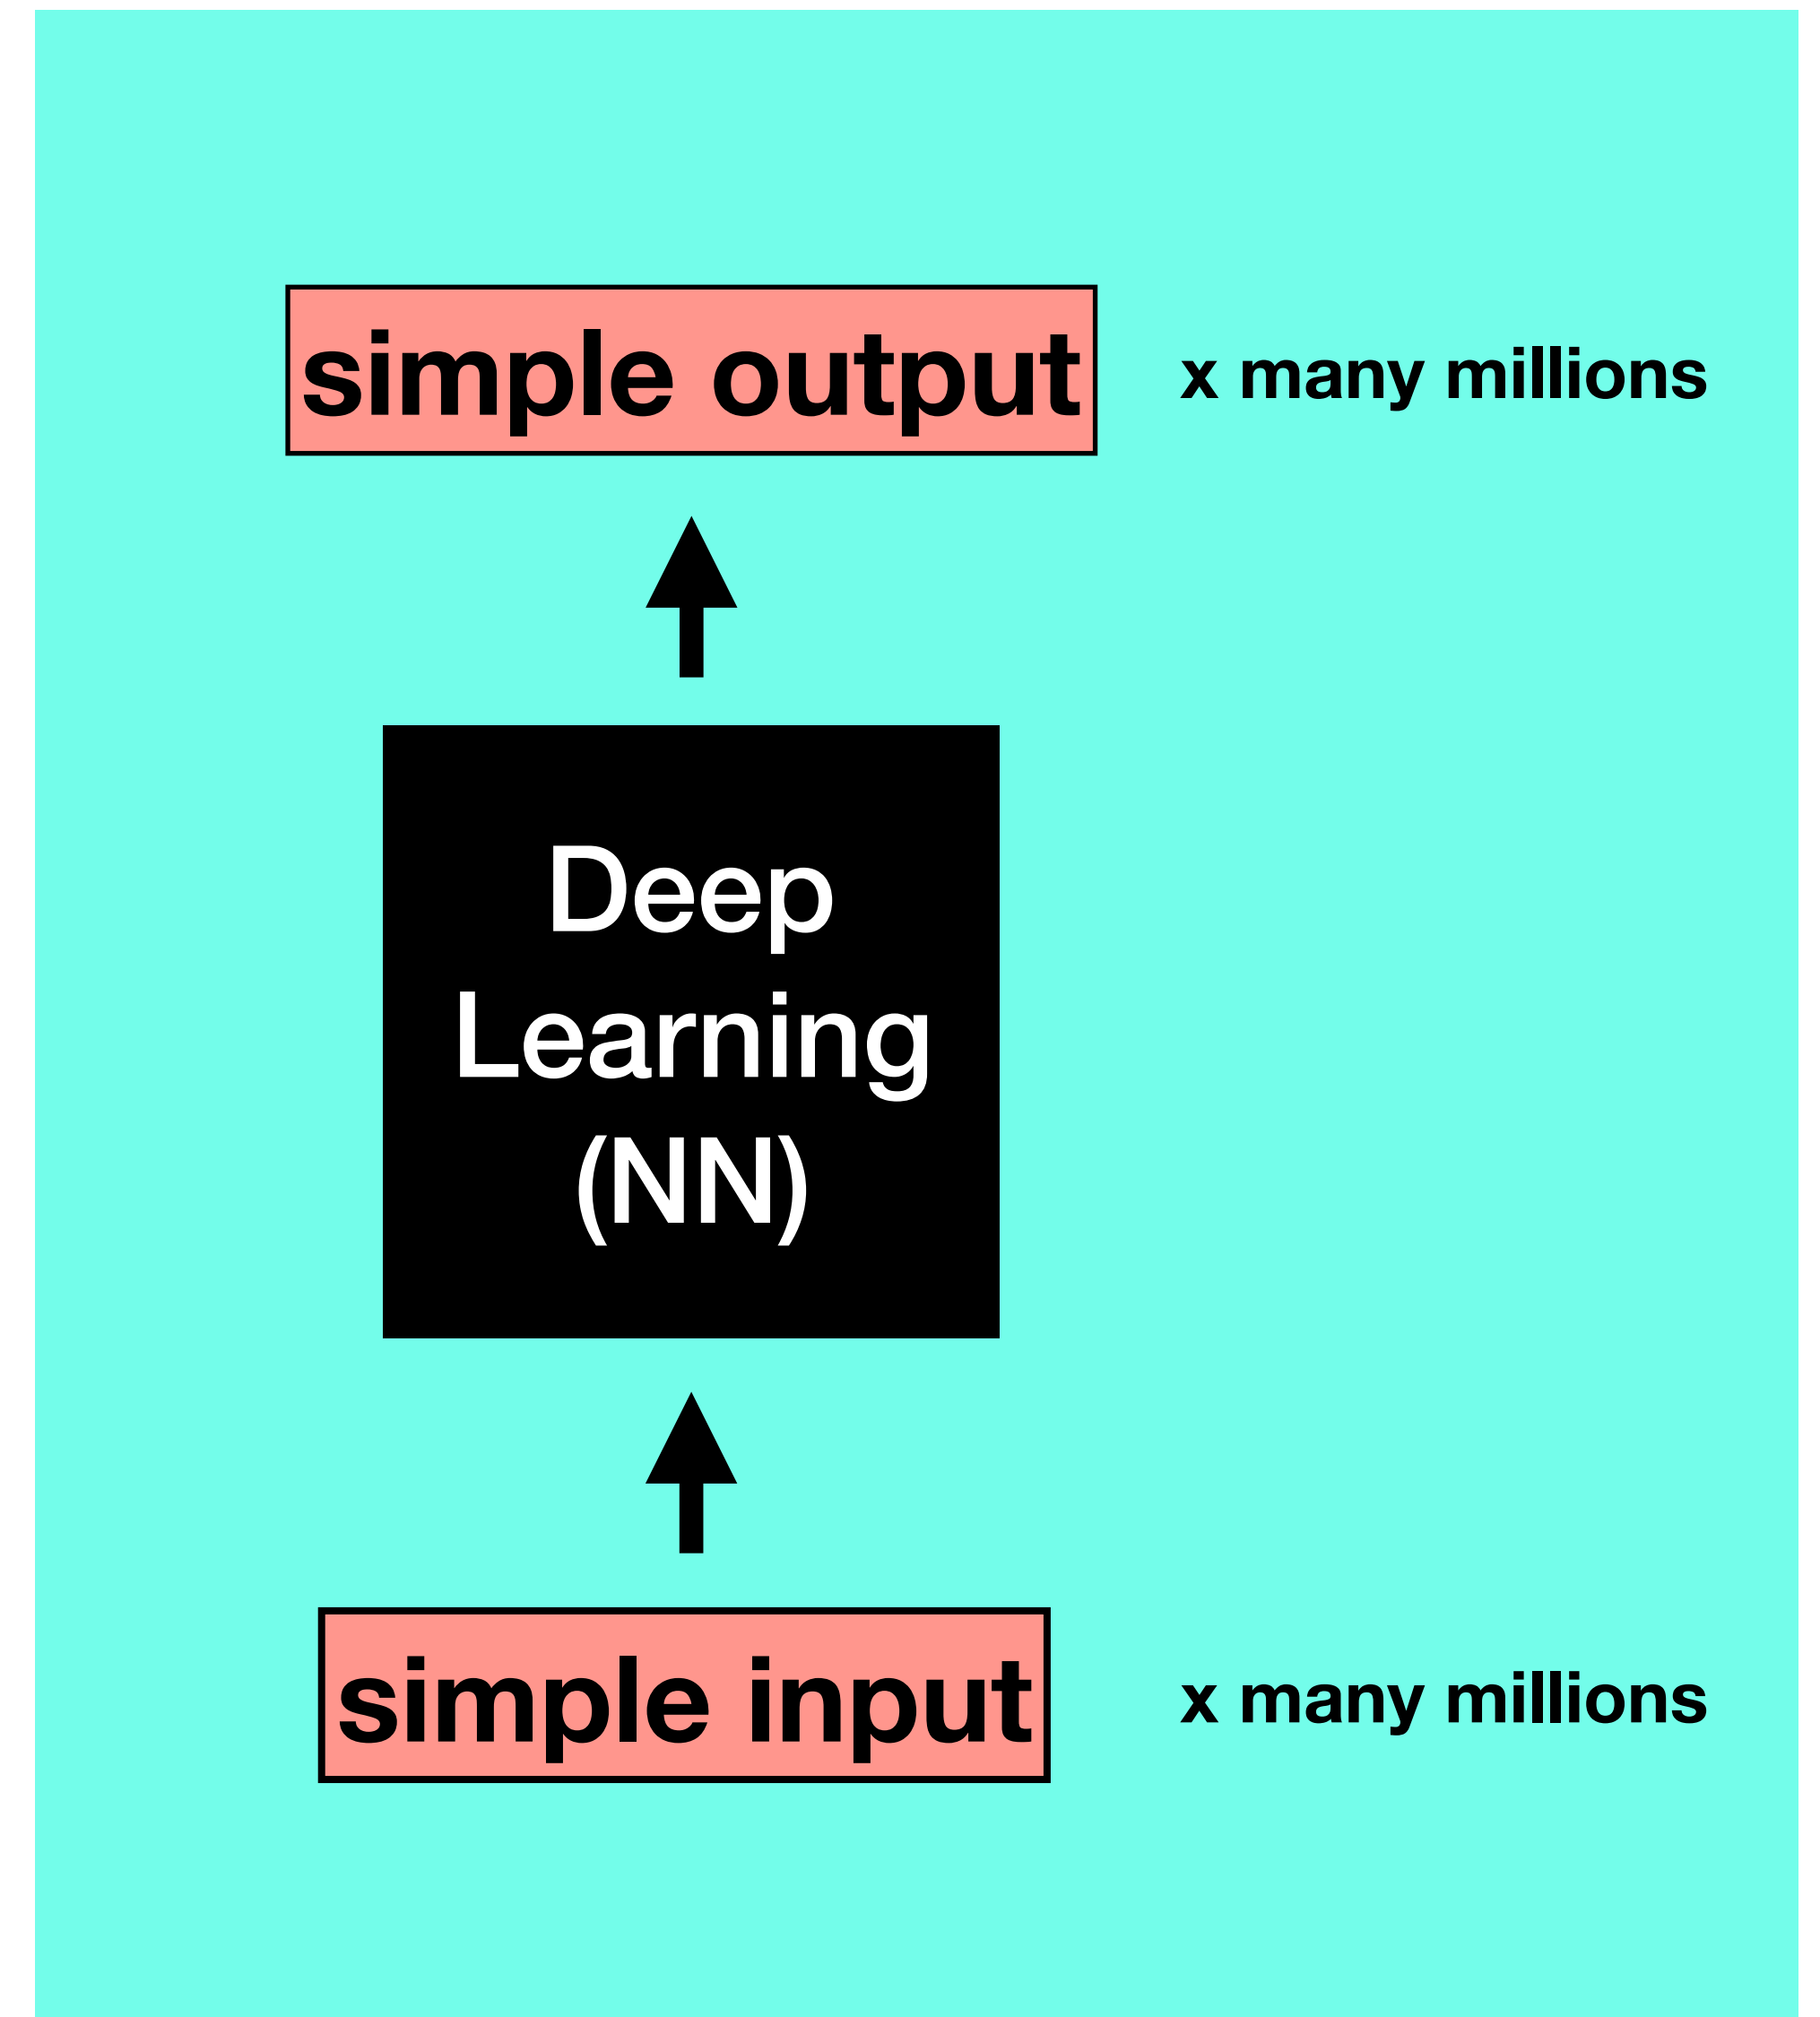
Machine Translation



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

2. Limitations A: patchwork ML

Simple I->O problem with lots of examples?



input: front camera
output: steer L/R

yes but..

combination of
many subsystems

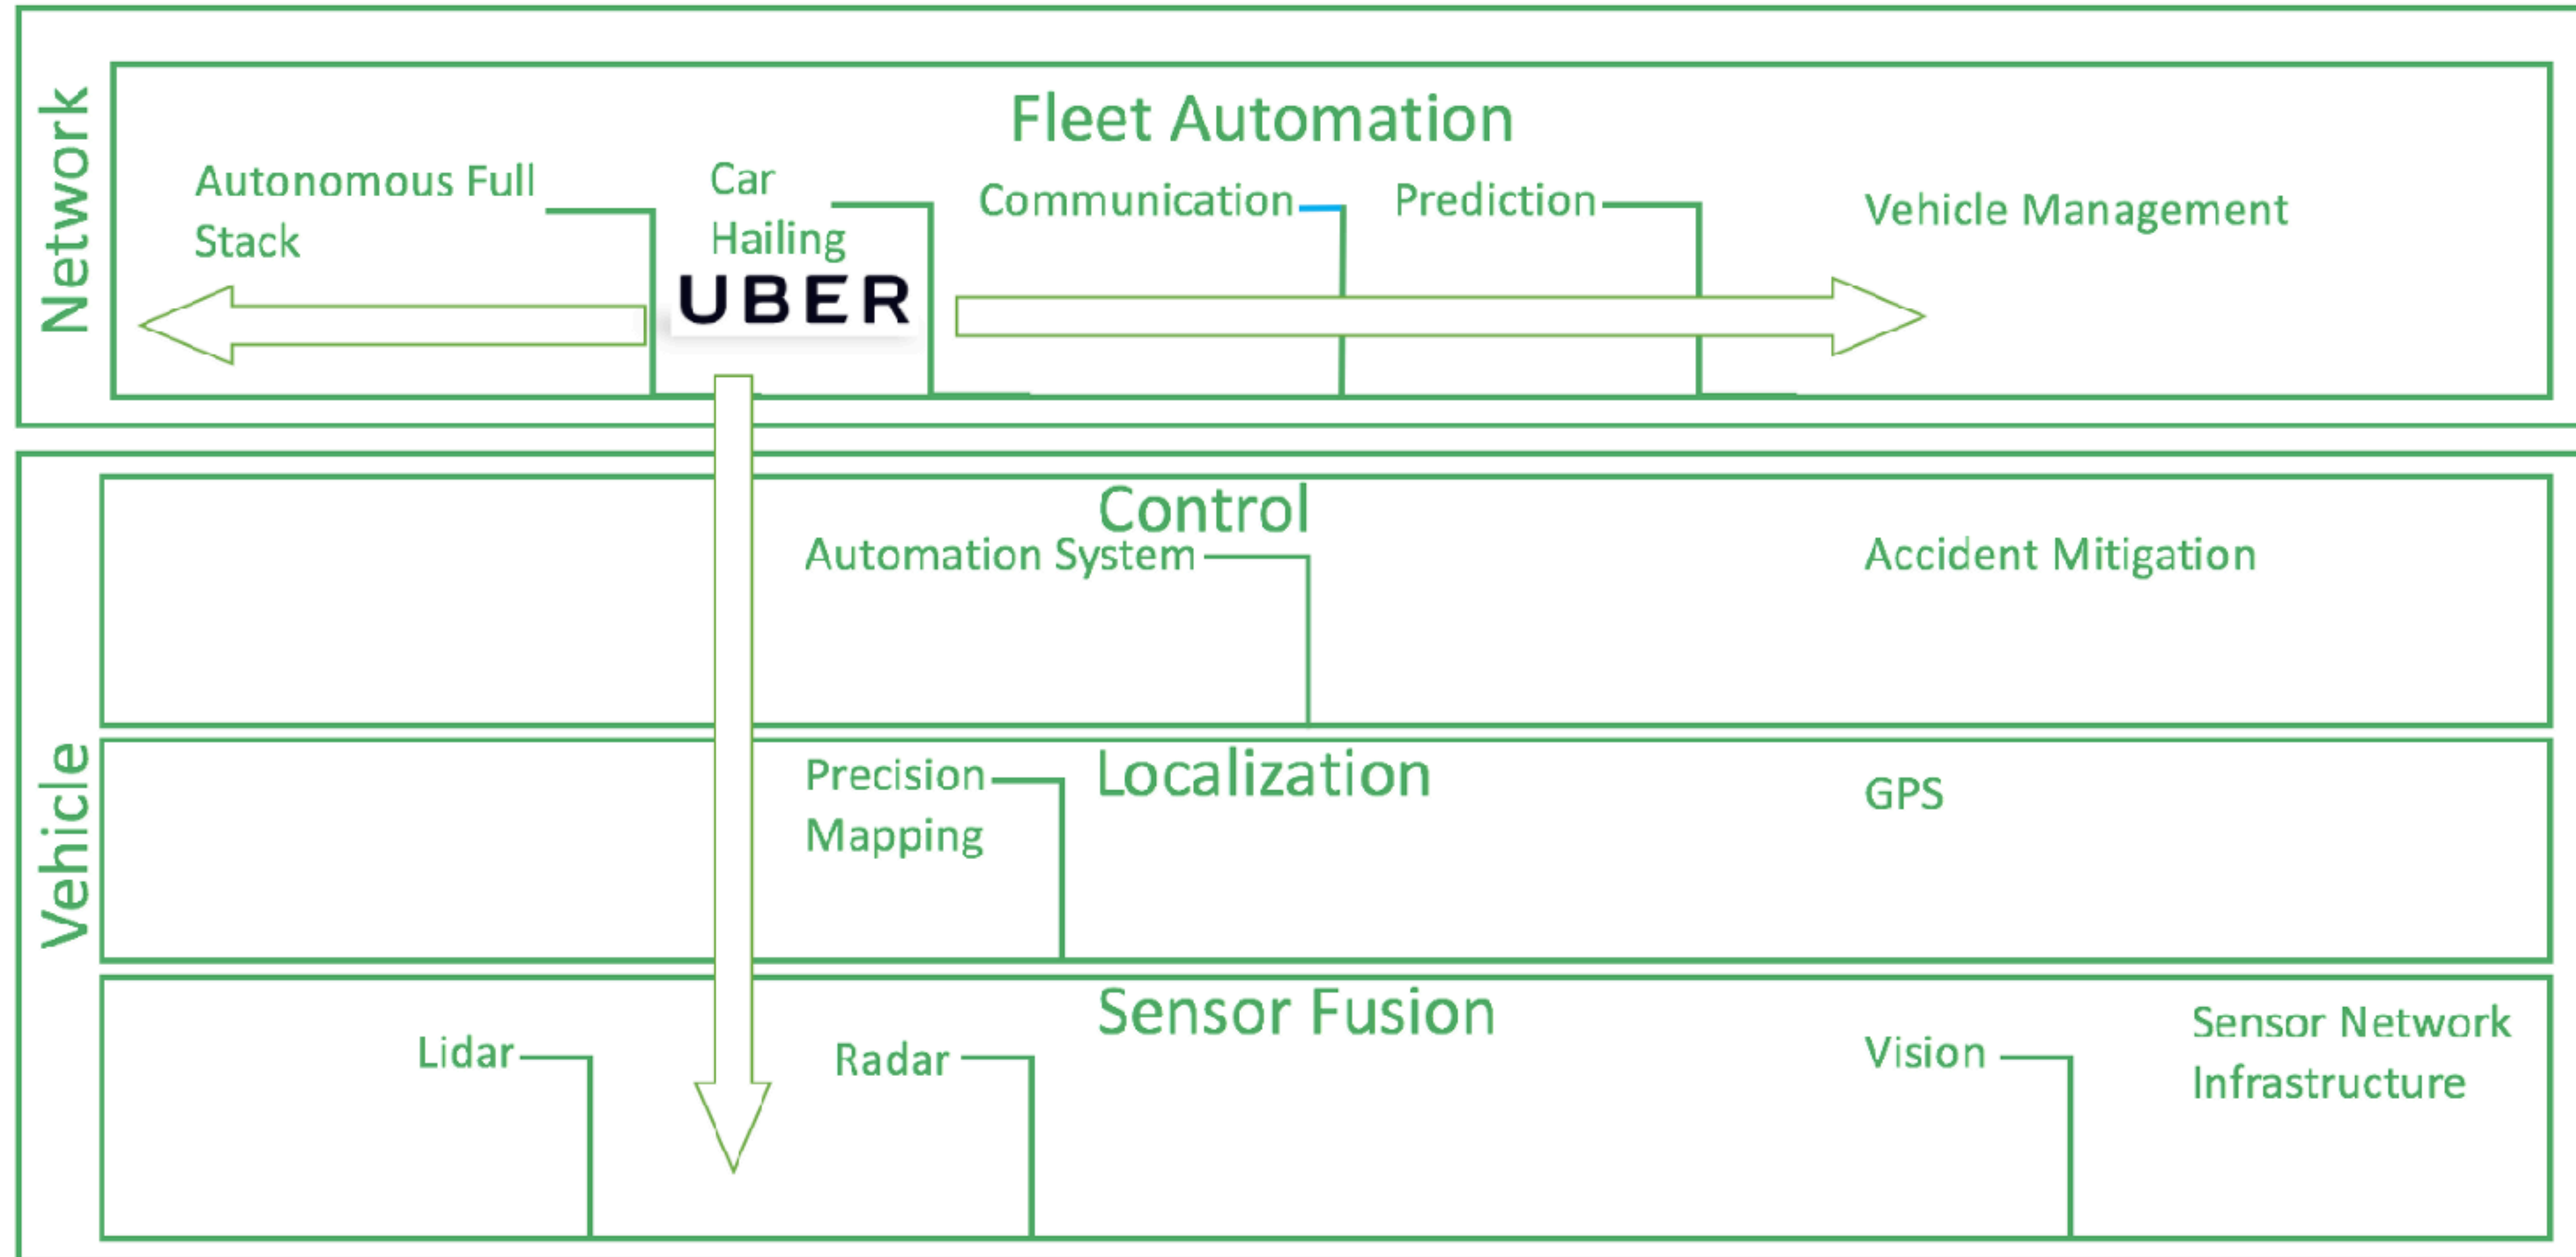


VS



high quality data


Patchwork solution




ML for real estate price prediction



Zestimate [?]
\$144,672

 **ZESTIMATE RANGE** [?]
\$127,000 - \$158,000

 **LAST 30 DAY CHANGE**
-\$4,627 (-3.1%)

Inside the Zestimate

The Zestimate is Zillow's best estimate of a home's value. It is based on a blend of valuation methods, each of which may produce a different estimate depending on the available data.

ESTIMATE BASED ON

Comparable homes	\$140,145	∨
Local tax assessments	\$145,039	∨
Market appreciation	\$143,163	∨
Local sale prices	\$135,874	∨

Intermezzo: computer programs

- Sequence of instructions
- Instructions need to be exact, no ambiguity
- + Correct input: 100% guaranteed output is correct
- Incorrect input: will freak out and stop
- All ML models live in context of classical computer programs
- Combining ML models happens through classical programming: set of precise instructions

Dogs and cats revisited

Robot on a mission: catch only cats

Classification




CAT

Classification
+ Localization




CAT



Patchwork solution

Pros

- + Get to a solution cheaper & faster
- + Use others' solutions (OS / MLaaS APIs)



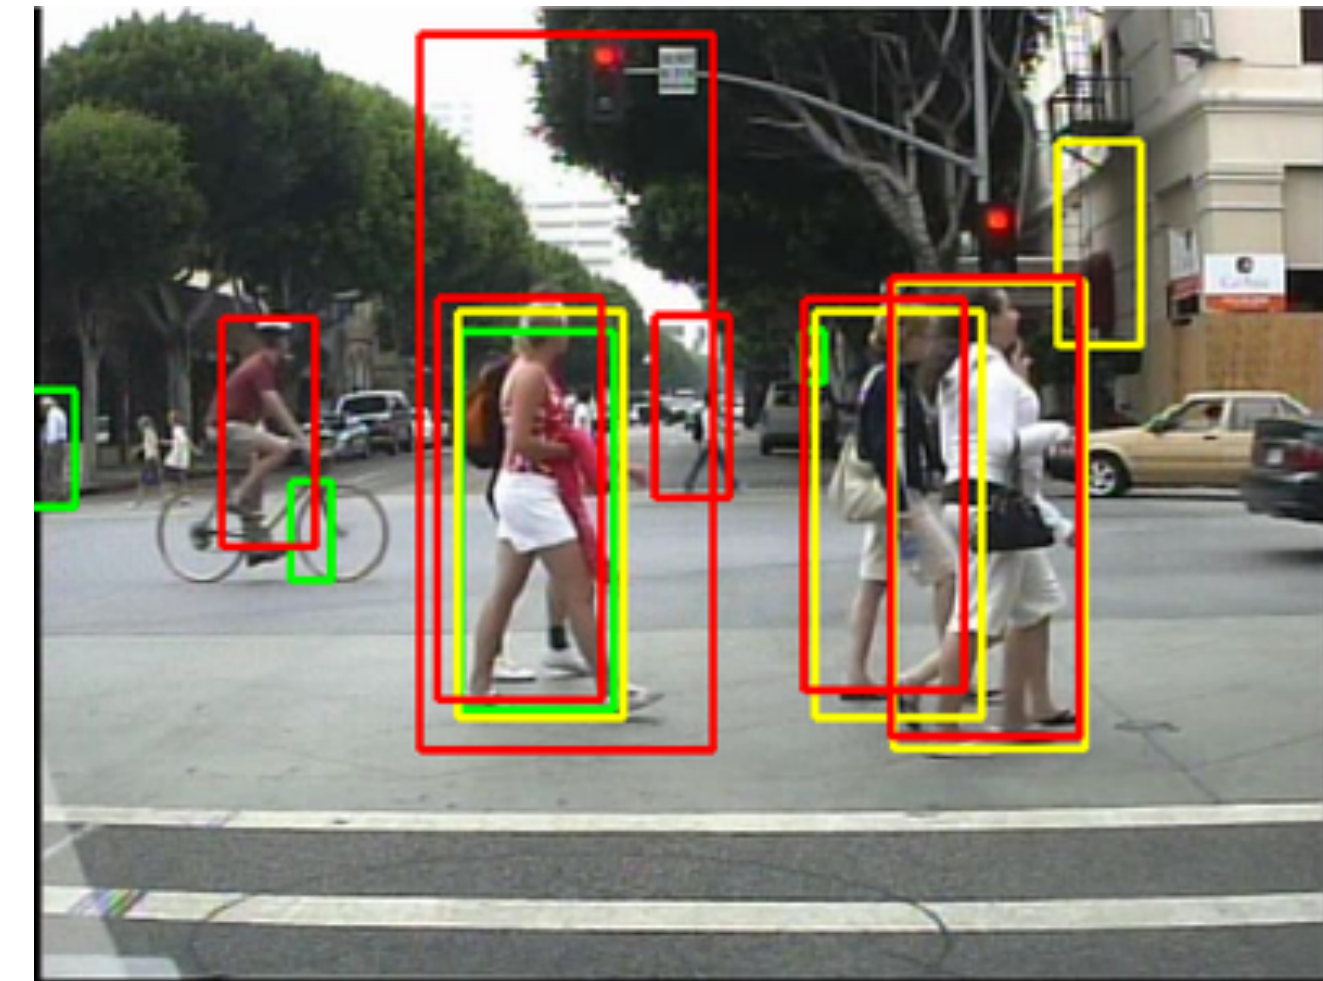
Cons

- More engineering to get good performance & you may never get there
- More cost to maintain
- No easy way to generalize
vs end to end machine learning: just collect new data



Patchwork solution: Reasons Why

- Not enough data for the whole
 - but data for the sub problems!
- Need for “guarantees” or explainability
- Legacy and history
- Spread of teams / skills
- No need for DL expertise



Tough decision

Patchwork solution

- make it work fast
- need to fix it forever

End to end ML solution

- spend more time and money upfront
- more expertise needed to “design it right”
- better performance
- more robust

Why a patchwork is more brittle

than end to end ML solution

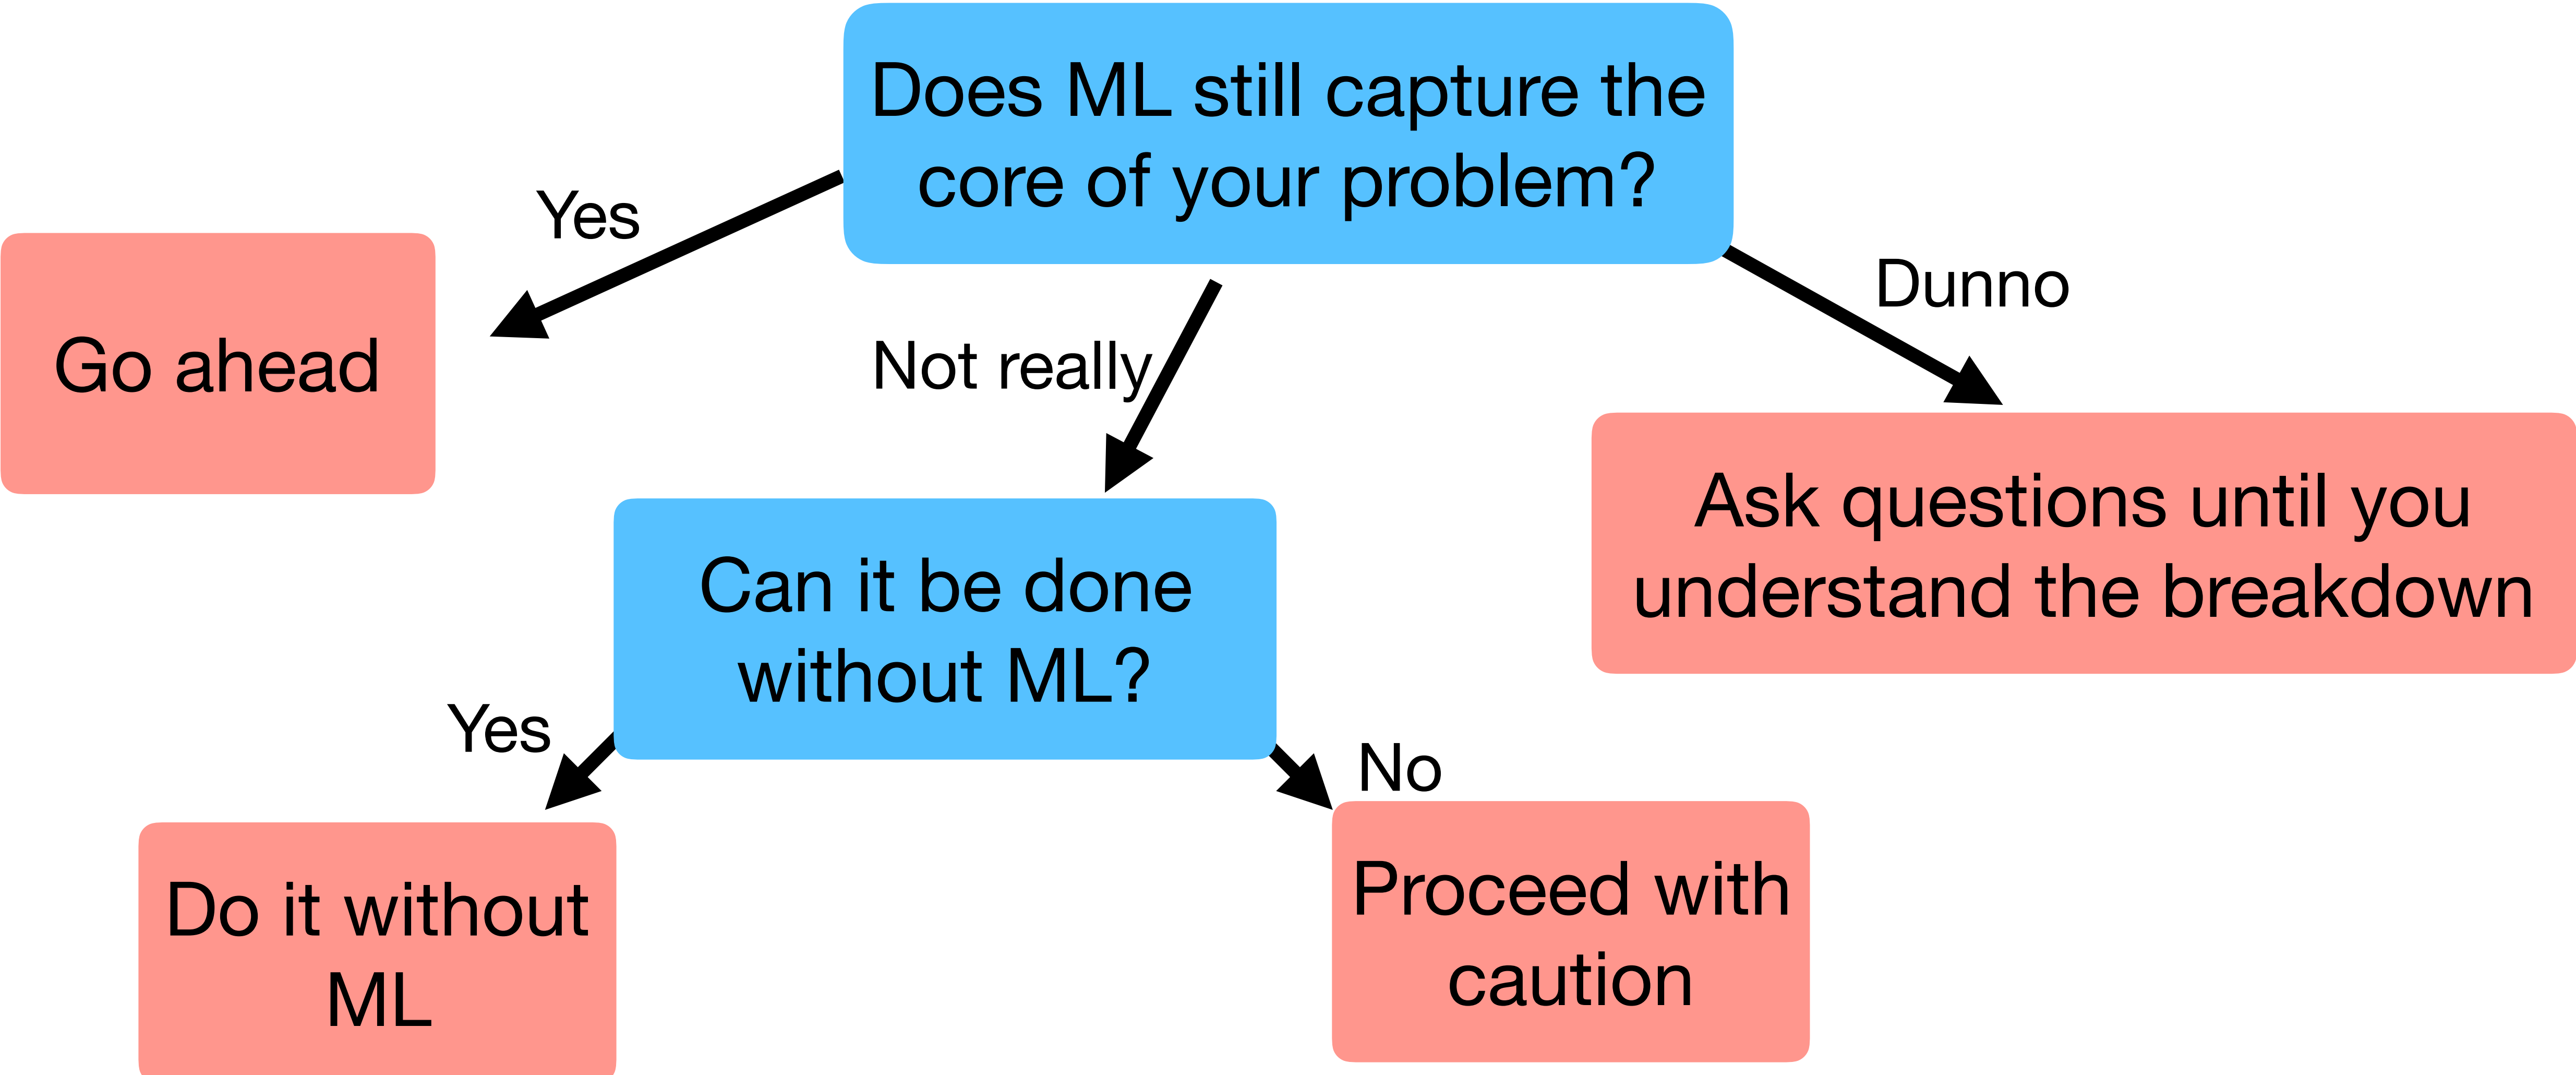
- Often subsystem data doesn't exactly match the actual problem
- Errors accumulate
- Model + heuristics combination
Limited by engineers' imagination
- Needs active human intervention
- More places for bugs to hide



vs



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



FACEBOOK'S ARTIFICIAL INTELLIGENCE ROBOTS SHUT DOWN AFTER THEY START TALKING TO EACH OTHER IN THEIR OWN LANGUAGE

'you i i i everything else'



Dan Joyce @dan_w_joyce · Sep 10
Replying to @zacharylipton @elonmusk
Doctors become irrelevant in a crumbling health service... But AI can save us



Computers could replace doctors, Jeremy Hunt says
At the Expo conference in Manchester, Mr Hunt said: 'The changes in medical innovation are likely to transform humanity by as much in the next
dailymail.co.uk

Elon Musk @elonmusk
This is nothing. In a few years, that bot will move so fast you'll need a strobe light to see it. Sweet dreams...

alex medina @mrmedina
we dead



PM - Nov 26, 2017

Natural Language Understanding

disambiguation:



(EN: bank)

winograd schemes:

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

What is too [small/large]?

dialogue



3. Limitations B: Frontiers

Reasoning



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



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

**Visual
Question
Answering
VQA**



Multi agent games

**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 aka gezond verstand

Imagine what happens if you open your hand:



vs



Common Sense aka gezond verstand

- 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!

1



2. Limitations part A: Brittleness by cobbling together.

2



3. Limitations part B: The frontiers. Stuff that's just too
difficult for now

3





you

Thank you

AI Academy Seminar
13 November 2018, Kortrijk
Tom Sercu