AI, ML, and the AIML

Anton van den Hengel

Director, The Australian Institute for machine Learning

Professor, School of Computer Science, University of Adelaide

Australian Centre for Visual Technologies

Innovation and education in visual information processing

Who we are

- The Australian Institute for Machine Learning
  - 65+ researchers (academics, post-docs + students)
  - Originally a Computer Vision group, but now working in Machine Learning more broadly
  - A major node of the Australian Centre for Visual Robotics
  - Collaborations with Facebook, Google, Microsoft, Canon, Bayer, LBT Innovations, BHP Billiton …
- Recent results:
  - Number 2 in ImageNet Scene Understanding Challenge 2016
    - Beating Samsung, Google, CMU, Oxford …
  - Number 1 on Street View House Numbers
    - Auxiliary outputs overcome vanishing supervision signal problem
  - World's best pedestrian detection
    - Managed to get Deep pedestrian detection to work
  - State of the art Pedestrian Re-ID
    - Paired RNNs as a video autoencoder
  - Number 1 in VQA Synthetic Scenes Challenge
    - Generalised application of CNNs to graphs
  - Won the VQA 2.0 Challenge 2017
    - The primary benchmark in VQA, funded by Facebook

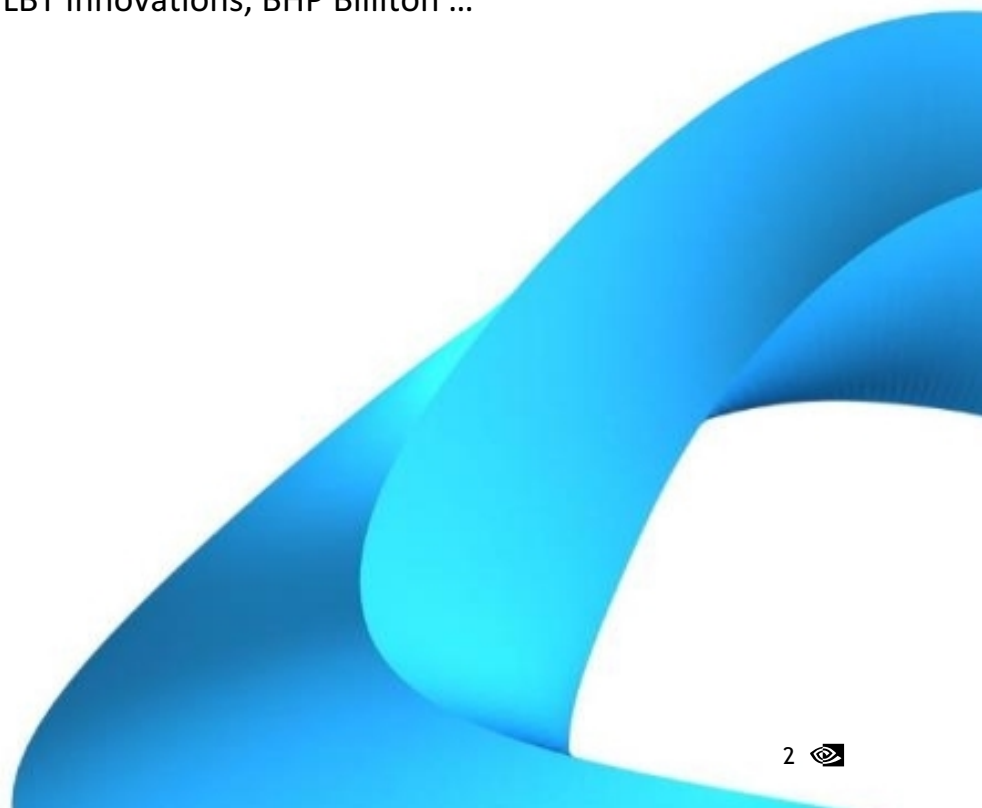

Commercial Research & Development

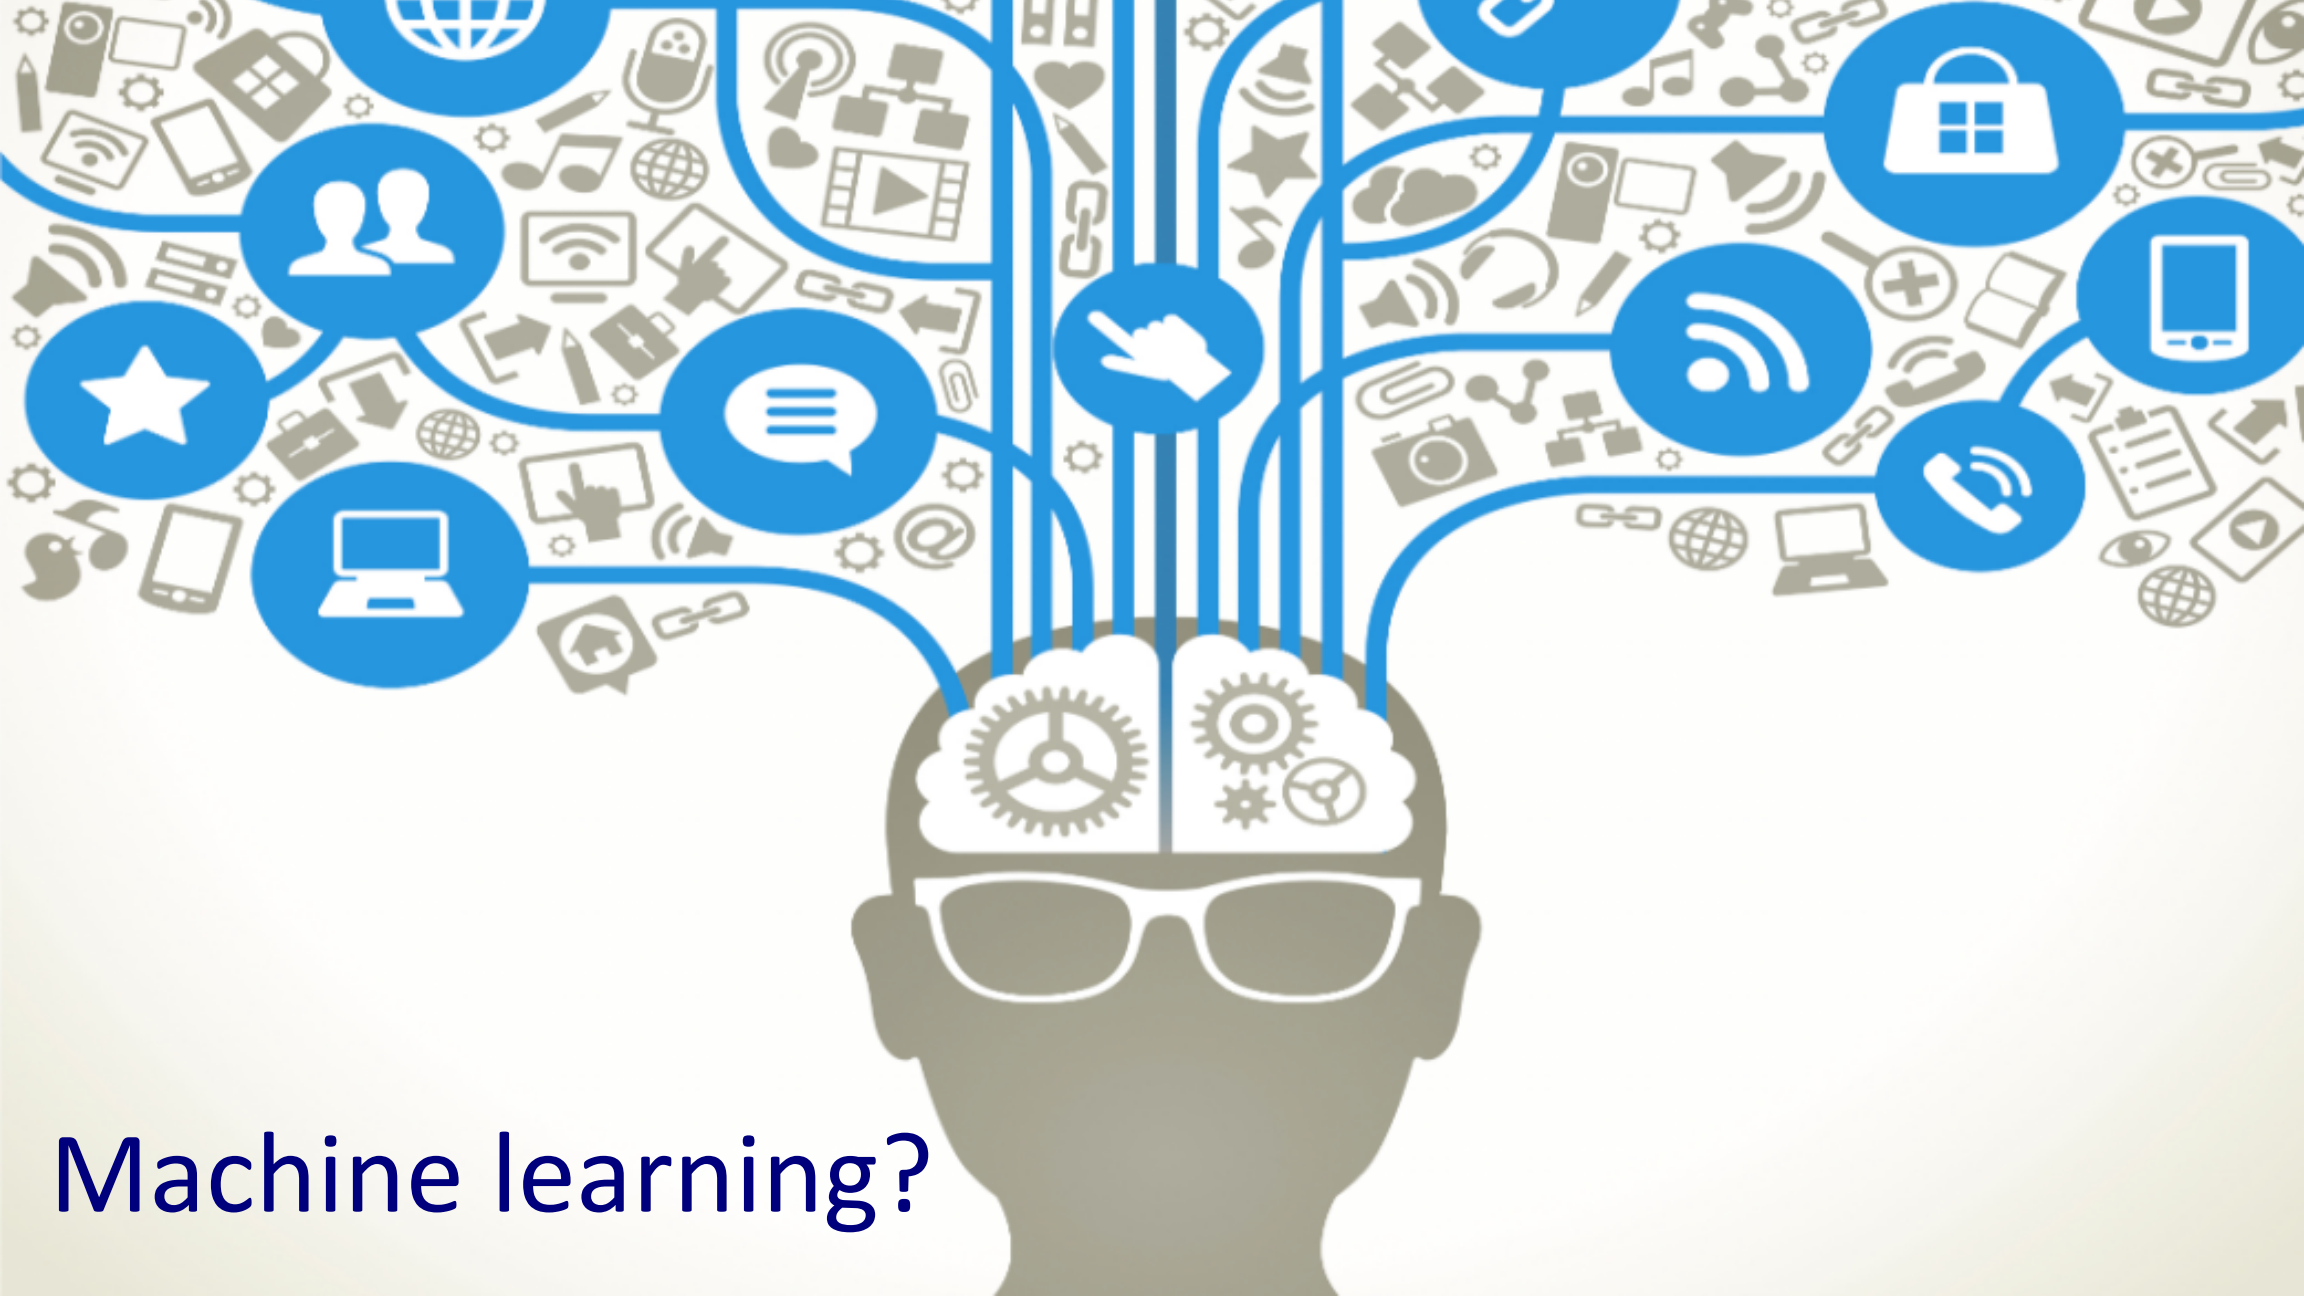

Machine learning?

Conventional Programming

```
2
3   ass Room(object):
4       def __init__(self, inventory, desc, short_desc):
5           self.__inventory = inventory
6           self.__n = None
7           self.__s = None
8           self.__e = None
9           self.__w = None
10          self.__desc = desc
11          self.__short_desc = short_desc
12          self.__gate_n = None
13          self.__gate_s = None
14          self.__gate_e = None
15          self.__gate_w = None
16
17          if not isinstance(desc, str):
18              raise TypeError ("the input provided is not a string.")
19          elif not isinstance(short_desc, str):
20              raise ValueError ("the string provided is empty.")
21
22      # these set the gates
23      # they set the opposite gates, with checks to avoid recursion loops
24      def set_n(self, other):
25          if not isinstance(other, Room) or not other:
```

Machine Learning

Cat

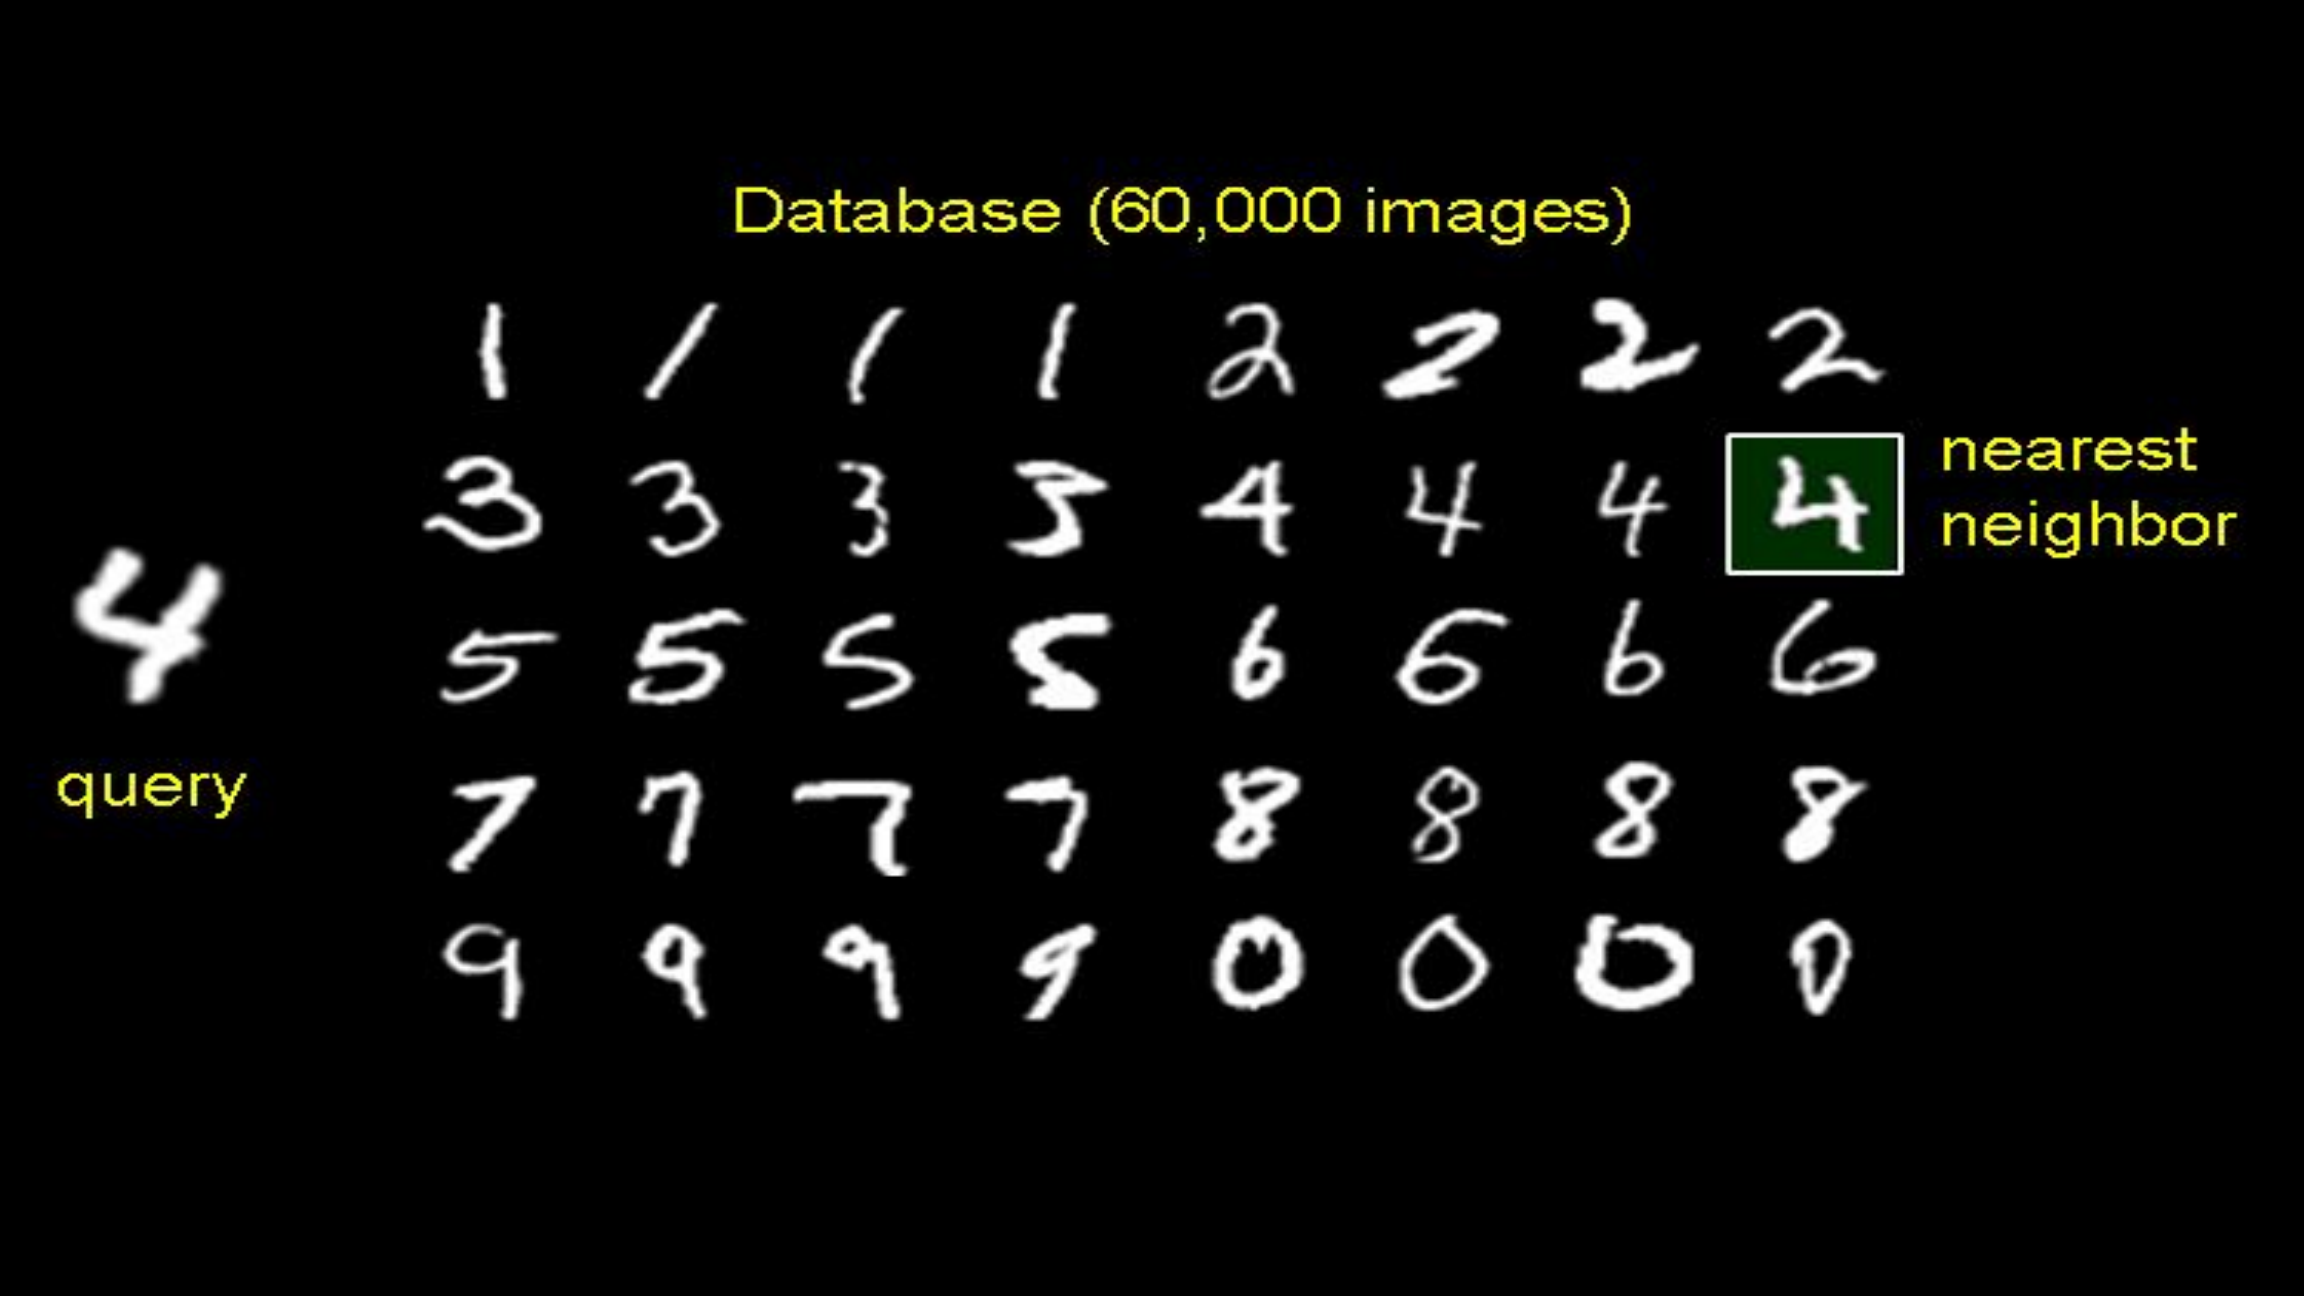
Database (60,000 images)
nearest neighbor
query

Machine Learning vs Artificial Intelligence

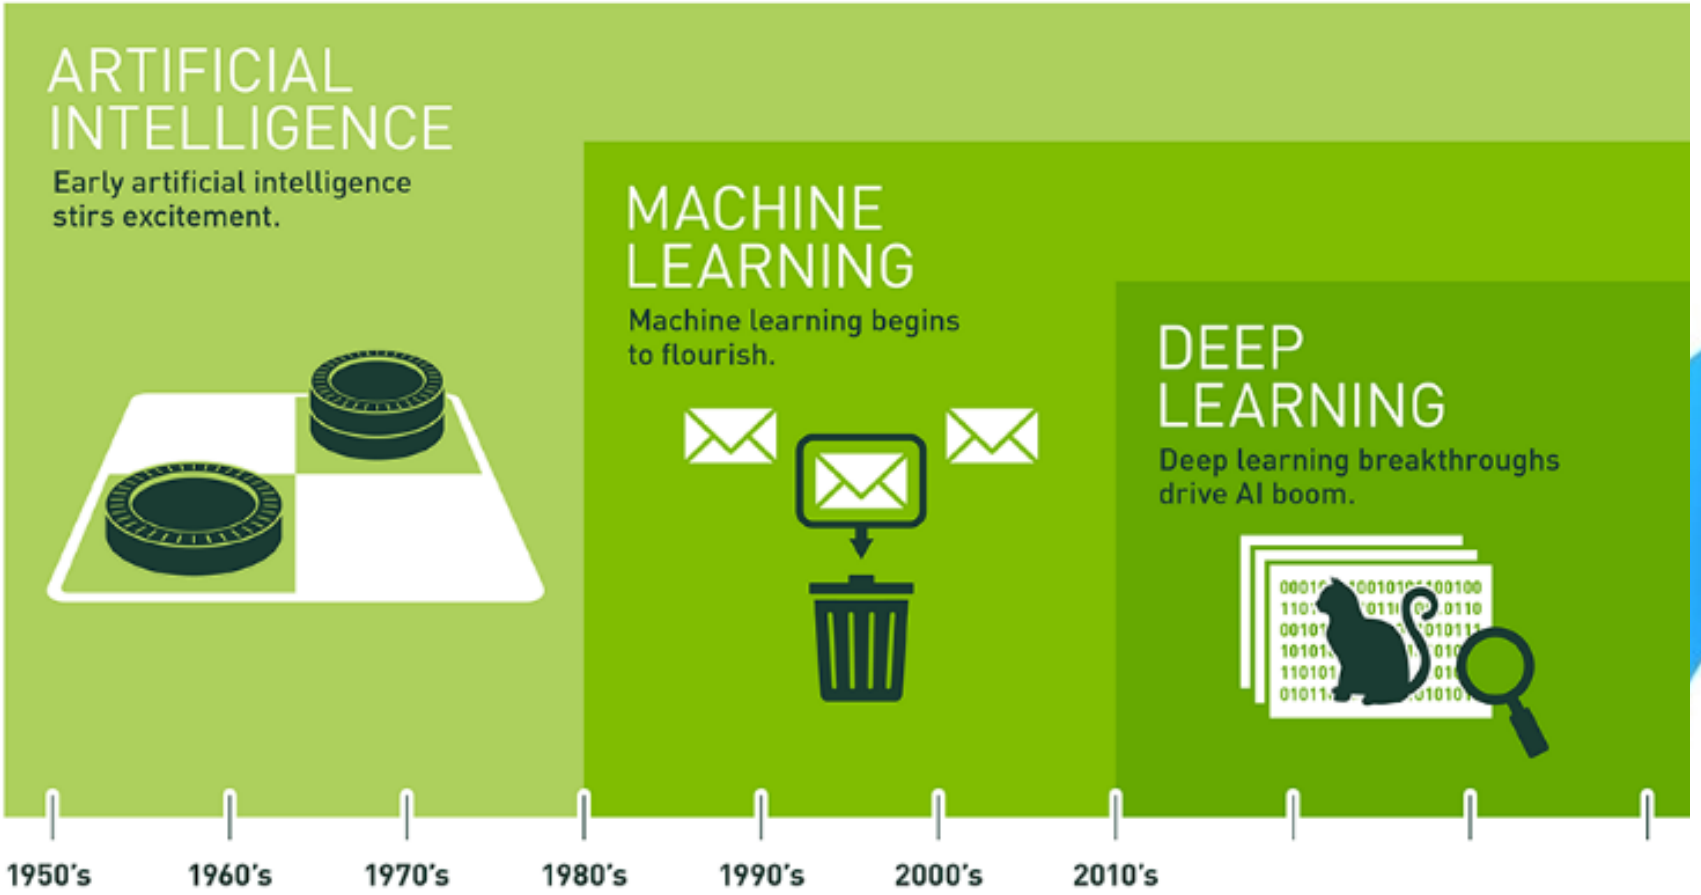

Image by NVIDIA

AI was supposed to be solved by now

- The idea was to start simple and slowly add complexity
  - 1965, H. A. Simon: "machines will be capable, within twenty years, of doing any work a man can do."
  - 1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."
  - 1970, Marvin Minsky: "In from three to eight years we will have a machine with the general intelligence of an average human being."
- It didn’t work

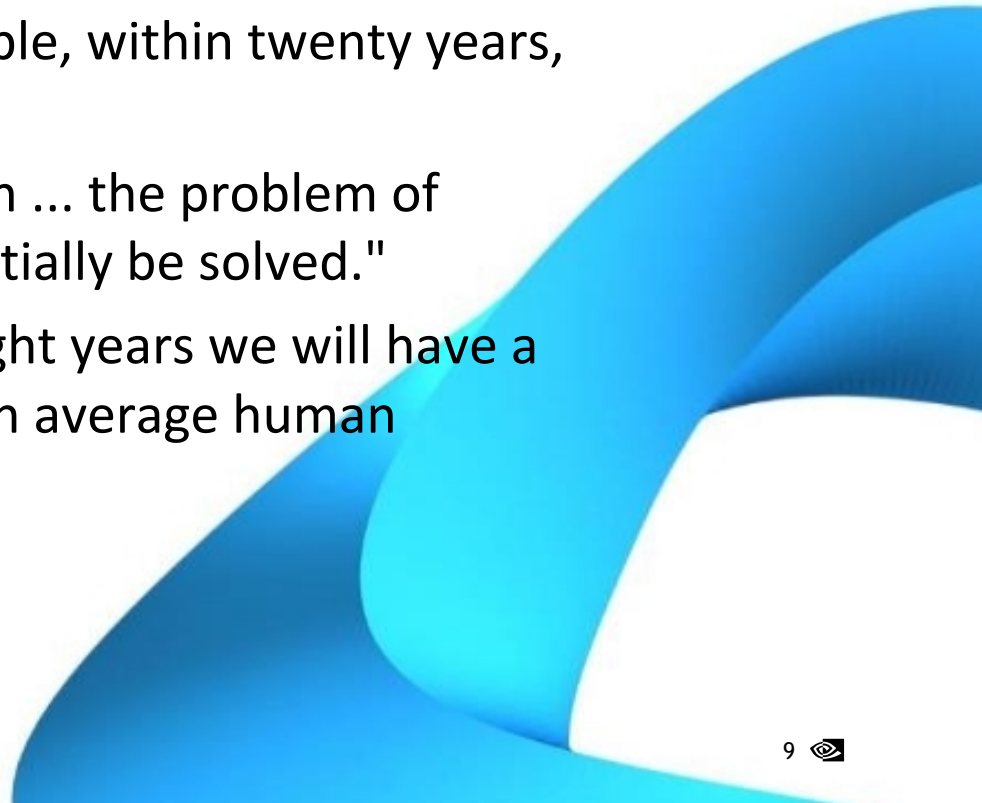

Expert predictions of years until AI

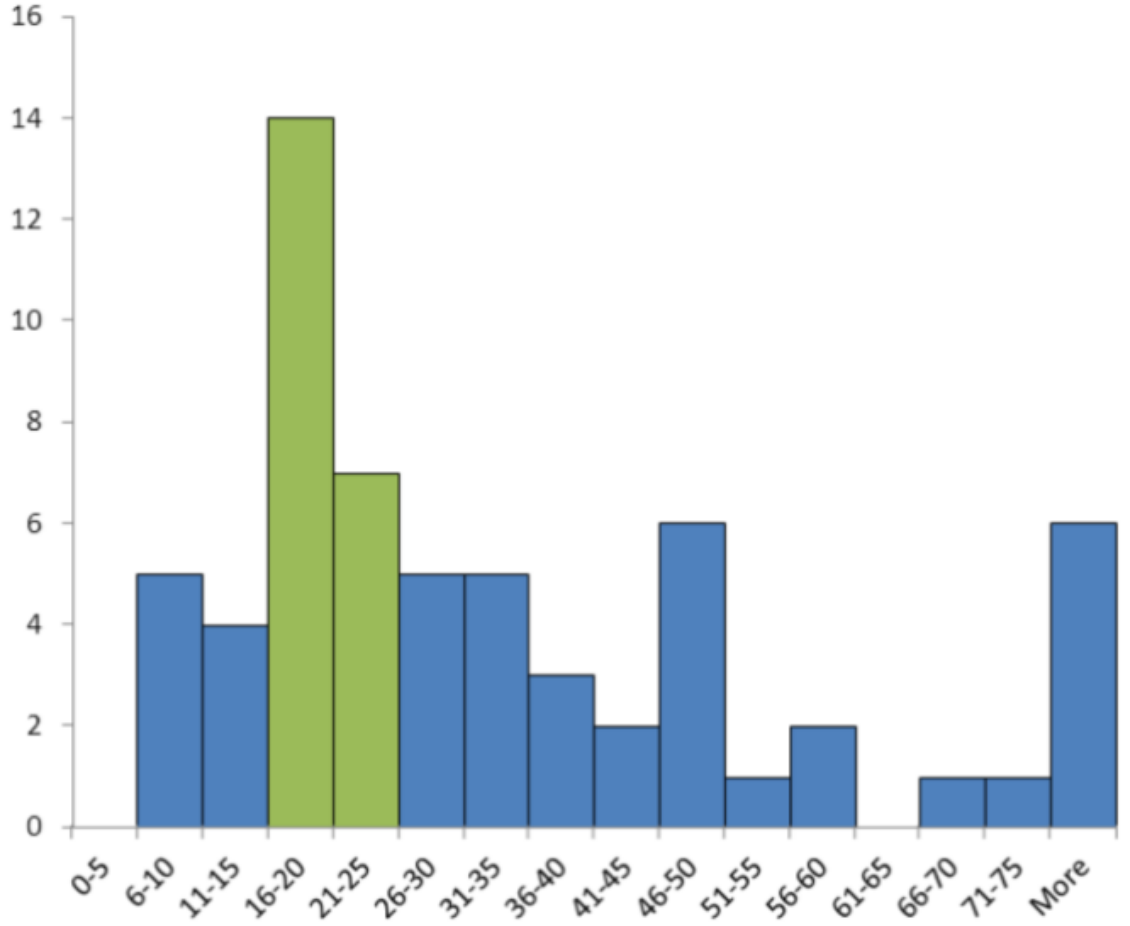

Armstrong, Stuart, and Kaj Sotala. "How we're predicting AI–or failing to." *Beyond Artificial Intelligence*. Springer International Publishing, 2015. 11-29.

"man in black shirt is playing guitar."

"construction worker in orange safety vest is working on road."

"two young girls are playing with lego toy."

"boy is doing backflip on wakeboard."

"girl in pink dress is jumping in air."

"black and white dog jumps over bar."

"young girl in pink shirt is swinging on swing."

"man in blue wetsuit is surfing on wave."


Q: Is this a fish or a bicycle?
A: No


Q: What game are they playing
A: Soccer


But counting is a very complicated idea
Q: How many horses are in the image?
A: 2

ACVT: Medical Devices

- We are developing 5 devices
  - 2 assisting a human expert
    - 1 recently achieved FDA approval
  - 3 allowing non-experts to make semantic decisions
- Mortal risks
- High-level decisions
- Uncontrolled environments


The Singularity is a way off yet

Translation ~= Understanding