

How To Build A Brain

Perception

Beating humans



Gary Kasparov



Deep Blue



Feng-Hsiung Hsu

1997: Deep Blue beats Gary Kasparov at chess.

Moravec's paradox (1988)

"It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility"




Rodney Brooks

In early AI research, intelligence was characterized as

"the things that highly-educated male scientists found challenging"

while

"things that children of four or five years could do effortlessly, such as visually distinguishing between a coffee cup and a chair...were not thought of as activities requiring intelligence"




Stephen Pinker (1994)


“The main lesson of thirty-five years of AI research is that the hard problems are easy and the easy problems are hard”



AlexNet (2012)




Ilya Sutskever, Alex Krizhevsky, Geoffrey Hinton

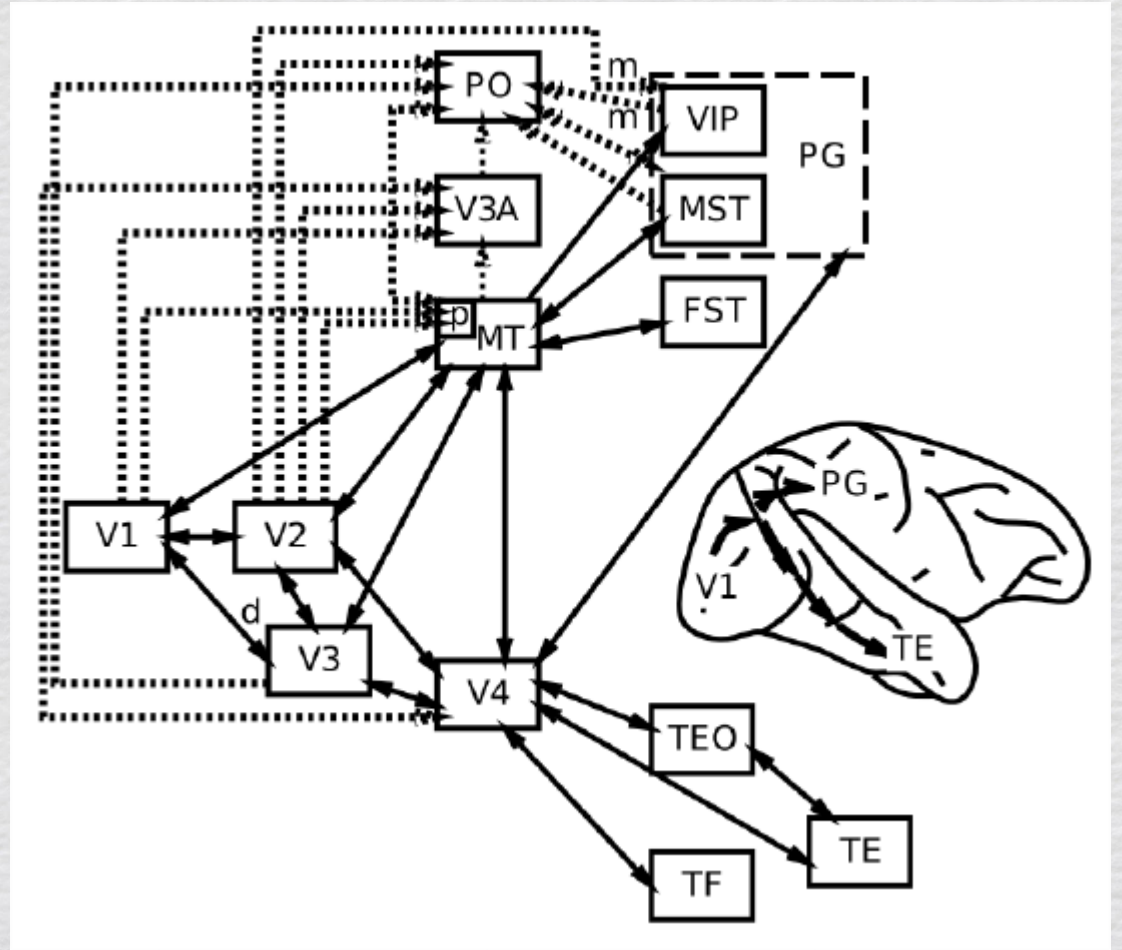







Visual Neuroscience



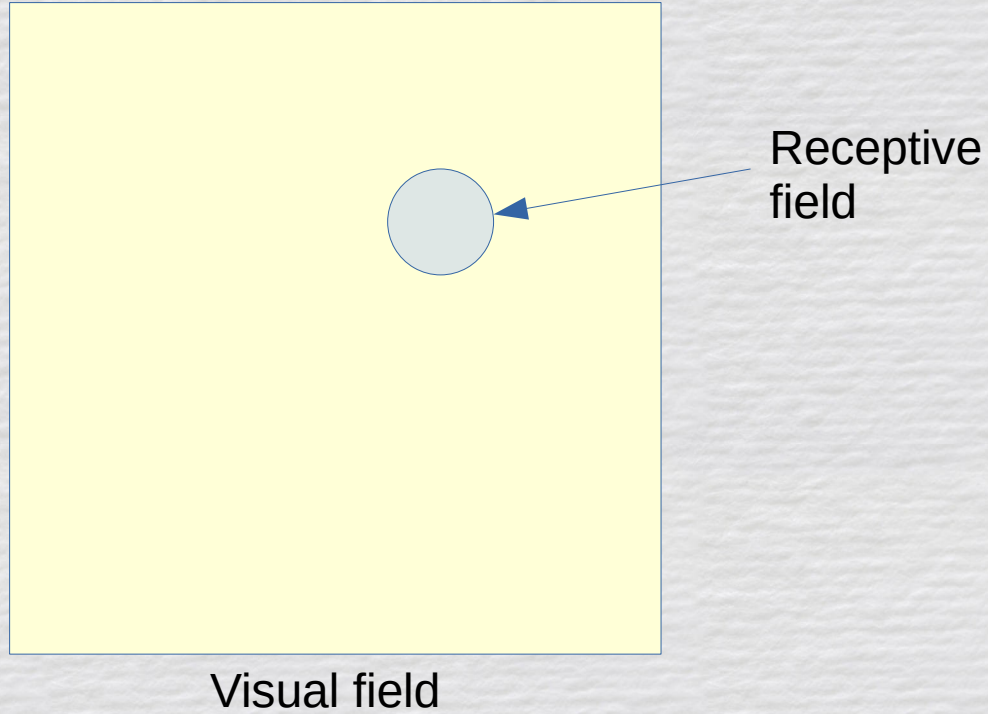
Deeper



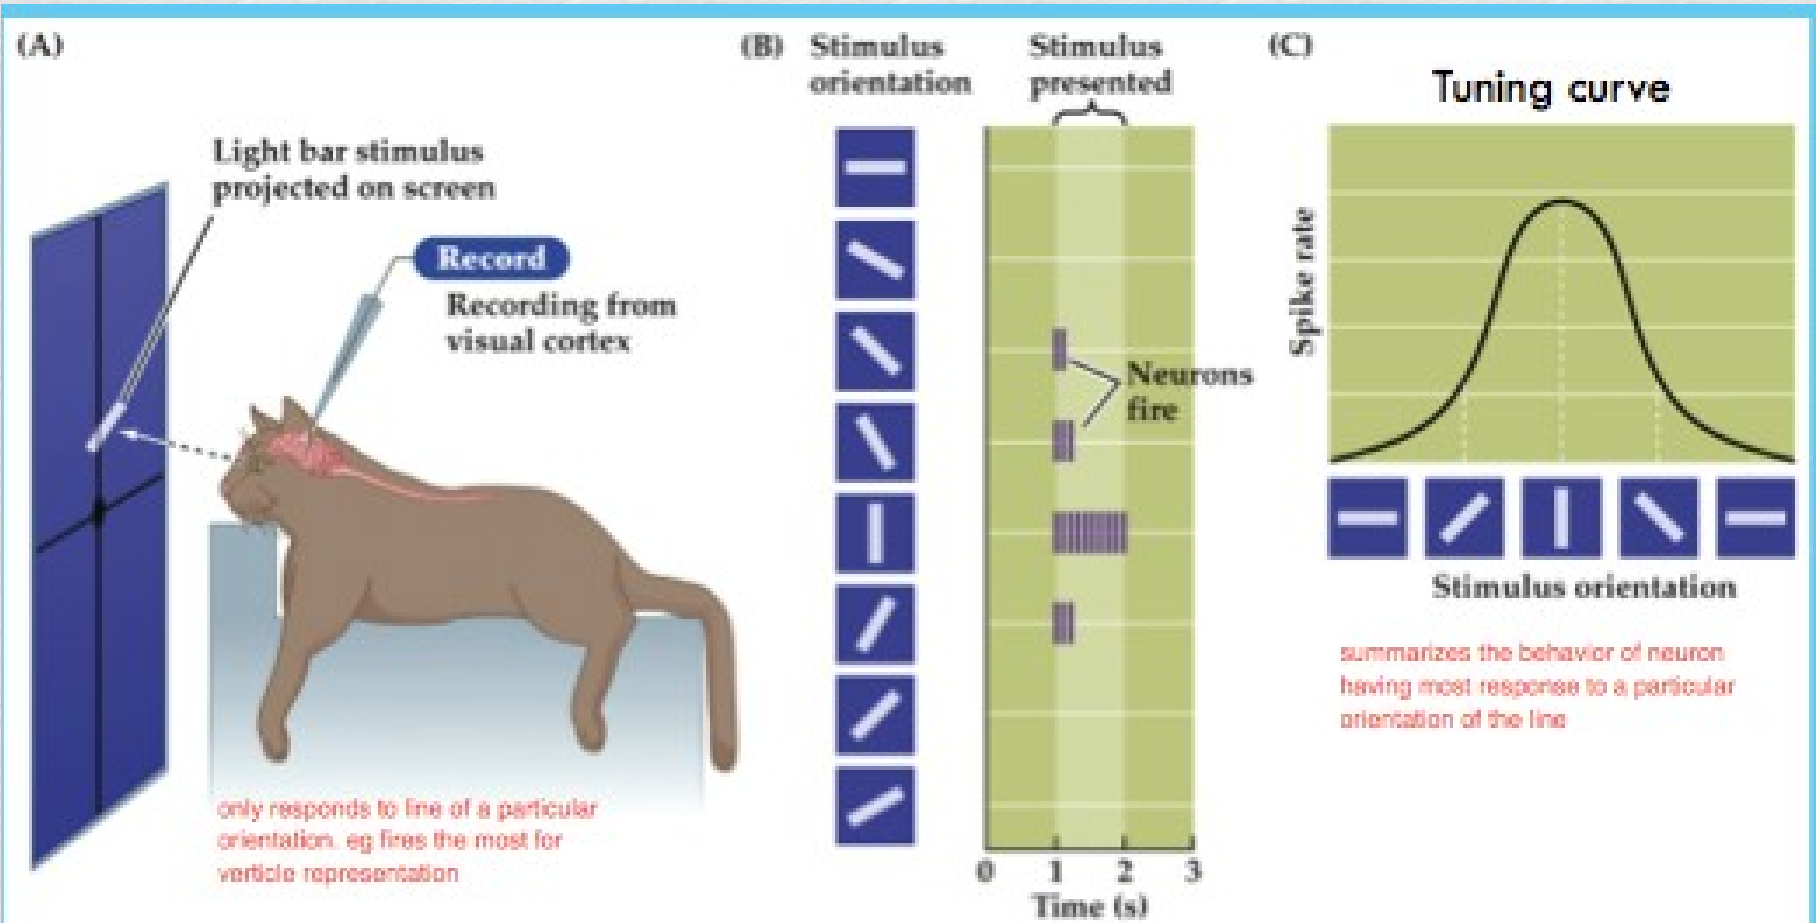
More sparse




Receptive fields




Neural tuning



Simple to complex




Neocognitron



Kunihiko Fukushima




LeNet (1989)



Recognition of handwritten digits



Yann LeCun





SUN-4/260



AT&T DSP 32C

Convolutional filter



Input

| | | |
|---|---|---|
| 0 | 1 | 0 |
| 0 | 1 | 0 |
| 0 | 1 | 0 |

Weights

| | | |
|----|---|----|
| 0 | 1 | -1 |
| -1 | 1 | -1 |
| -1 | 1 | 0 |

Multiply

| | | |
|---|---|---|
| 0 | 1 | 0 |
| 0 | 1 | 0 |
| 0 | 1 | 0 |

Add

3

Output

3

| | | |
|---|---|---|
| 1 | 0 | 0 |
| 0 | 1 | 0 |
| 0 | 0 | 1 |

| | | |
|----|---|----|
| 0 | 1 | -1 |
| -1 | 1 | -1 |
| -1 | 1 | 0 |

| | | |
|---|---|---|
| 0 | 0 | 0 |
| 0 | 1 | 0 |
| 0 | 0 | 0 |

1

1

| | | |
|---|---|---|
| 1 | 1 | 1 |
| 1 | 1 | 1 |
| 1 | 1 | 1 |

| | | |
|----|---|----|
| 0 | 1 | -1 |
| -1 | 1 | -1 |
| -1 | 1 | 0 |


| | | |
|----|---|----|
| 0 | 1 | -1 |
| -1 | 1 | -1 |
| 0 | 1 | 0 |

-1

0


Convolutional filter

Filter applied to
top-left of image

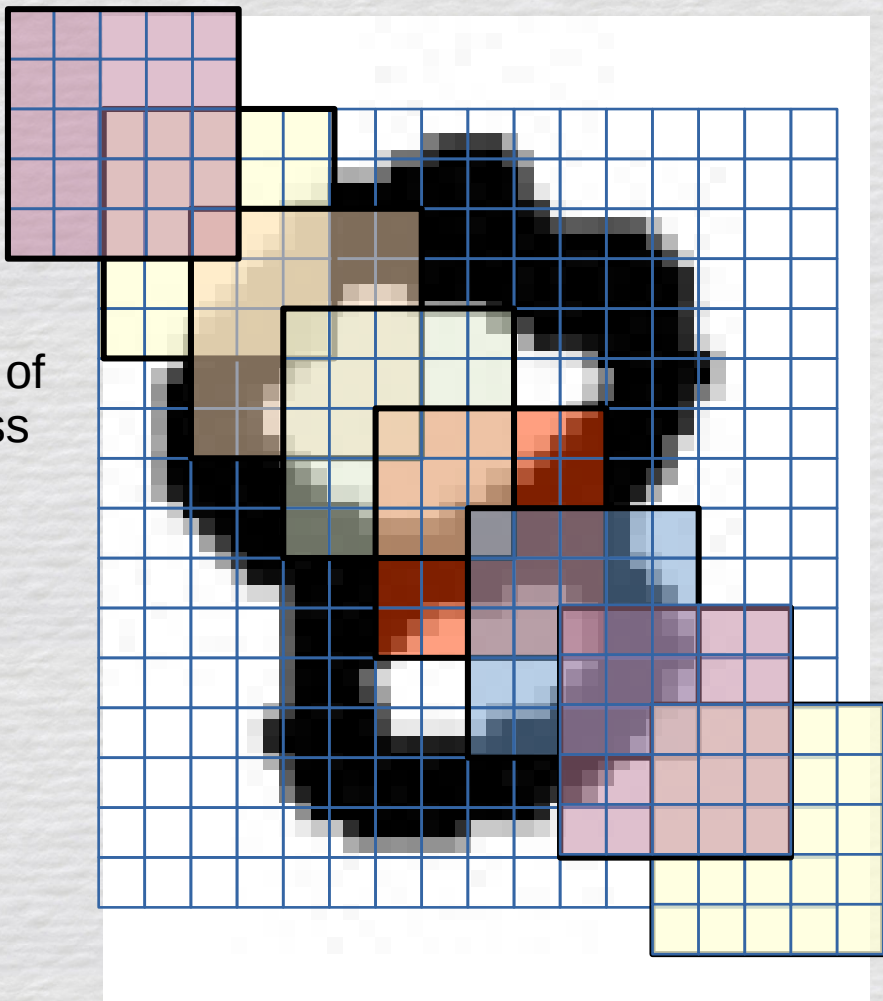


Convolutional filter

Next filter overlaps.




Convolutional filter



The filter is applied a total of 64 times across the image (8 examples shown here).


...and produces output activations for 64 hidden units, in an 8x8 grid.





Weights: 25

(vs. a dense network: 16,384!)

Convolutional filter




Filter Bank




2nd Filter Bank


The 2nd filter is applied a total of 16 times across the 1st filter (4 examples shown here).



...and produces output activations for 16 hidden units, in an 4x4 grid.




3D convolutional filter



The 3D filter has $5 \times 5 \times 12 = 300$ weights

First Three Layers




Image


2D filters

3D filters

Last two layers



LeNet



Image


2D filters

3D filters

Hidden

Output

LeNet



Trained with backprop


7,291 digits, 23 times each
(167,693 trials).

5% errors on a test set.

10 digits/second

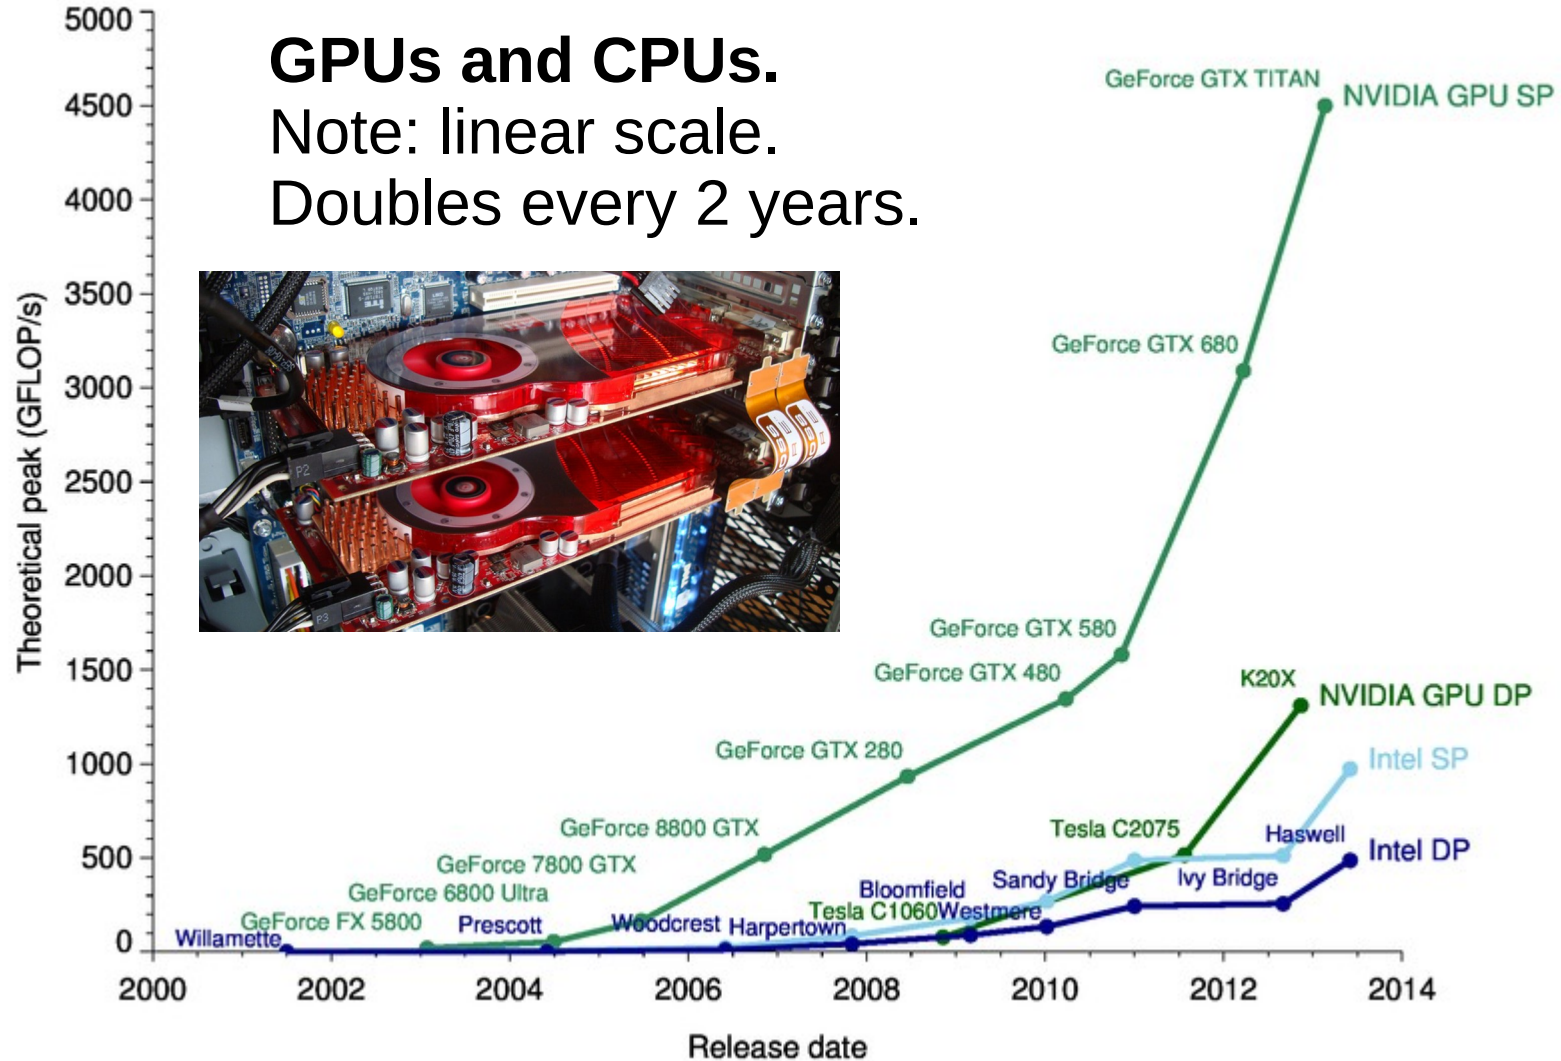



AlexNet (2012)



GPUs and CPUs.

Note: linear scale.
Doubles every 2 years.



Imagenet (2009-)



14 million images

20,000 categories

100s examples per
category

Some of the fish pictures in ImageNet

Legacy of AlexNet

100,000+ citations

Acquired by Google

Big Tech headhunts
brain-inspired AI
academics



Yann LeCun
Now @ **facebook**



Geoff Hinton
Now @ **Google**

Higher-res, and colour

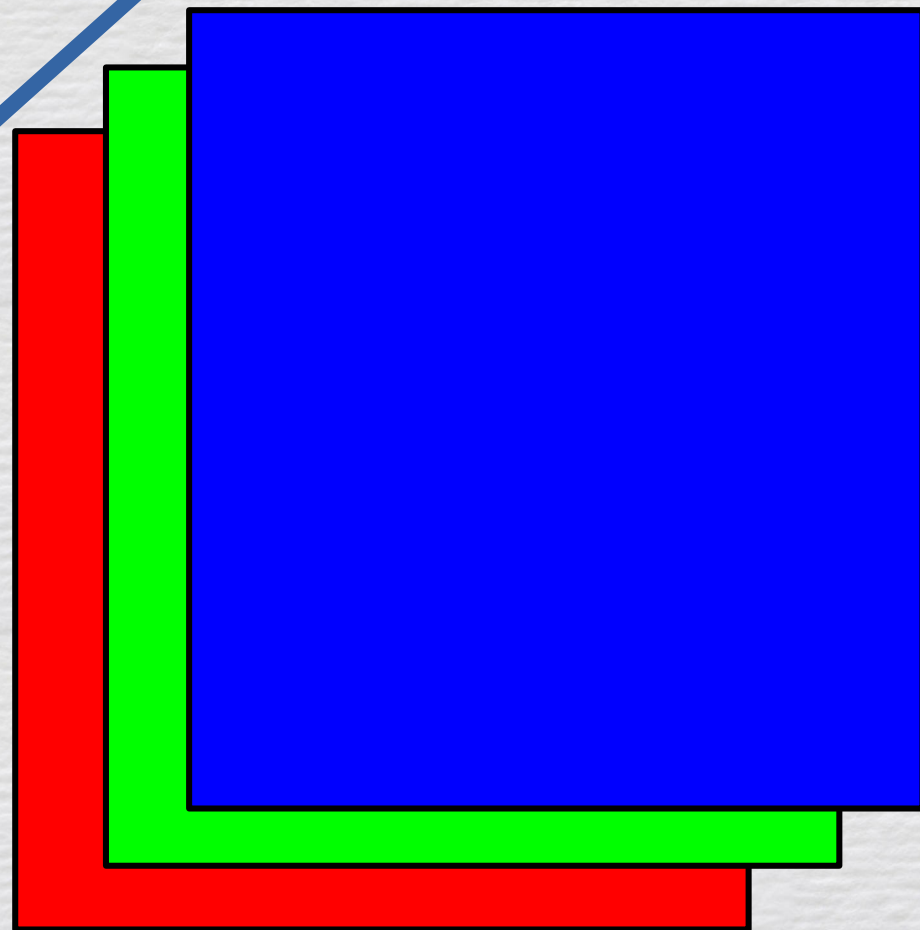

AlexNet

LeNet





size: 16 x 16

8




size: 256 x 256 x 3

LeNet







AlexNet



Deeper




ReLU






Grouping

Dropout






ResNet (2016)







LeNet



AlexNet



VGG



Inception


Network degradation




Kaiming He



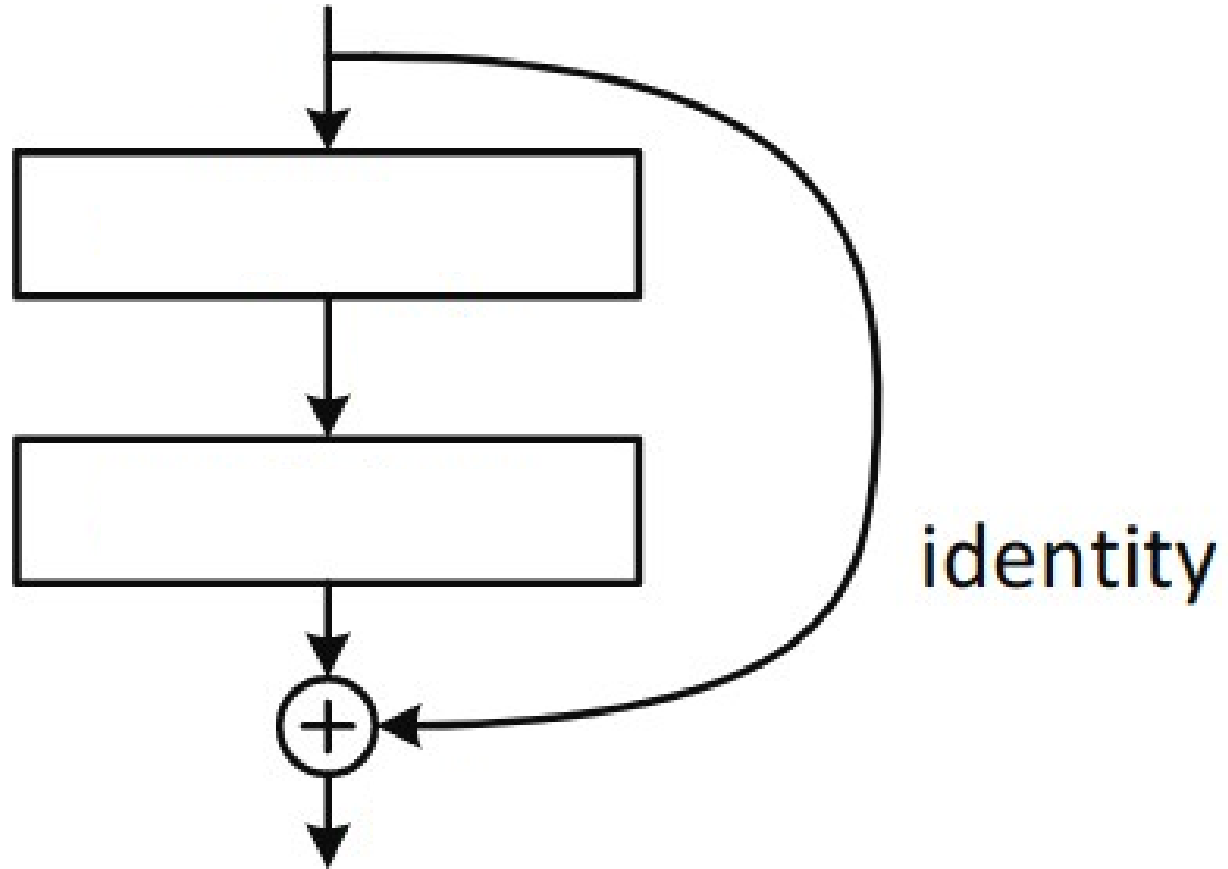
Jian Sun




Network degradation



Residual Block



ResNet 34




ResNet 152



ResNet152 with ResNet 34 for scale
(zoom in to slide for detail)



Applications



Content moderation


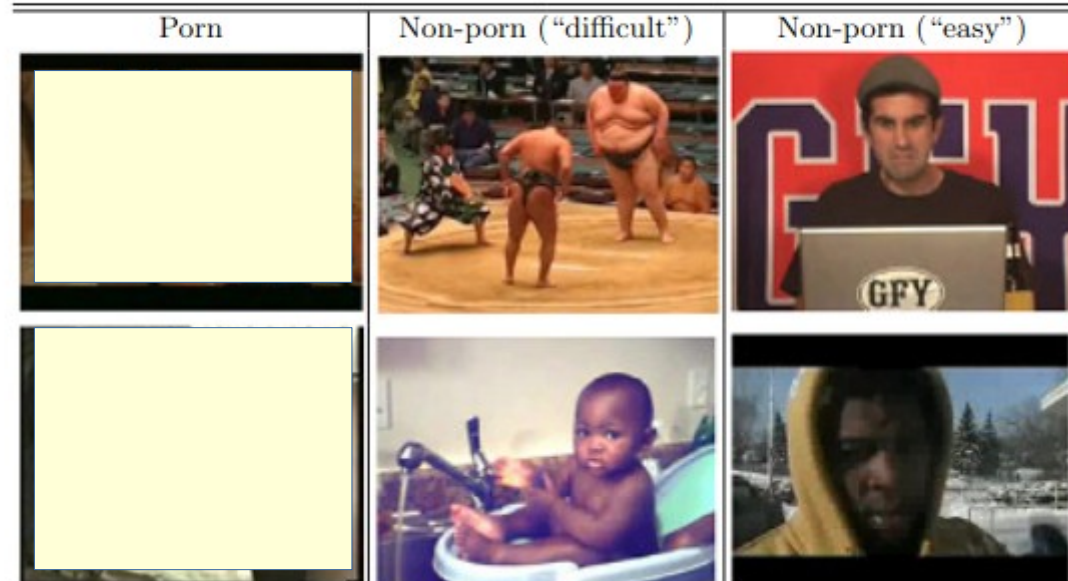



Table 2. NPDI dataset samples

Mohammed Moustafa (2015)



Content moderation

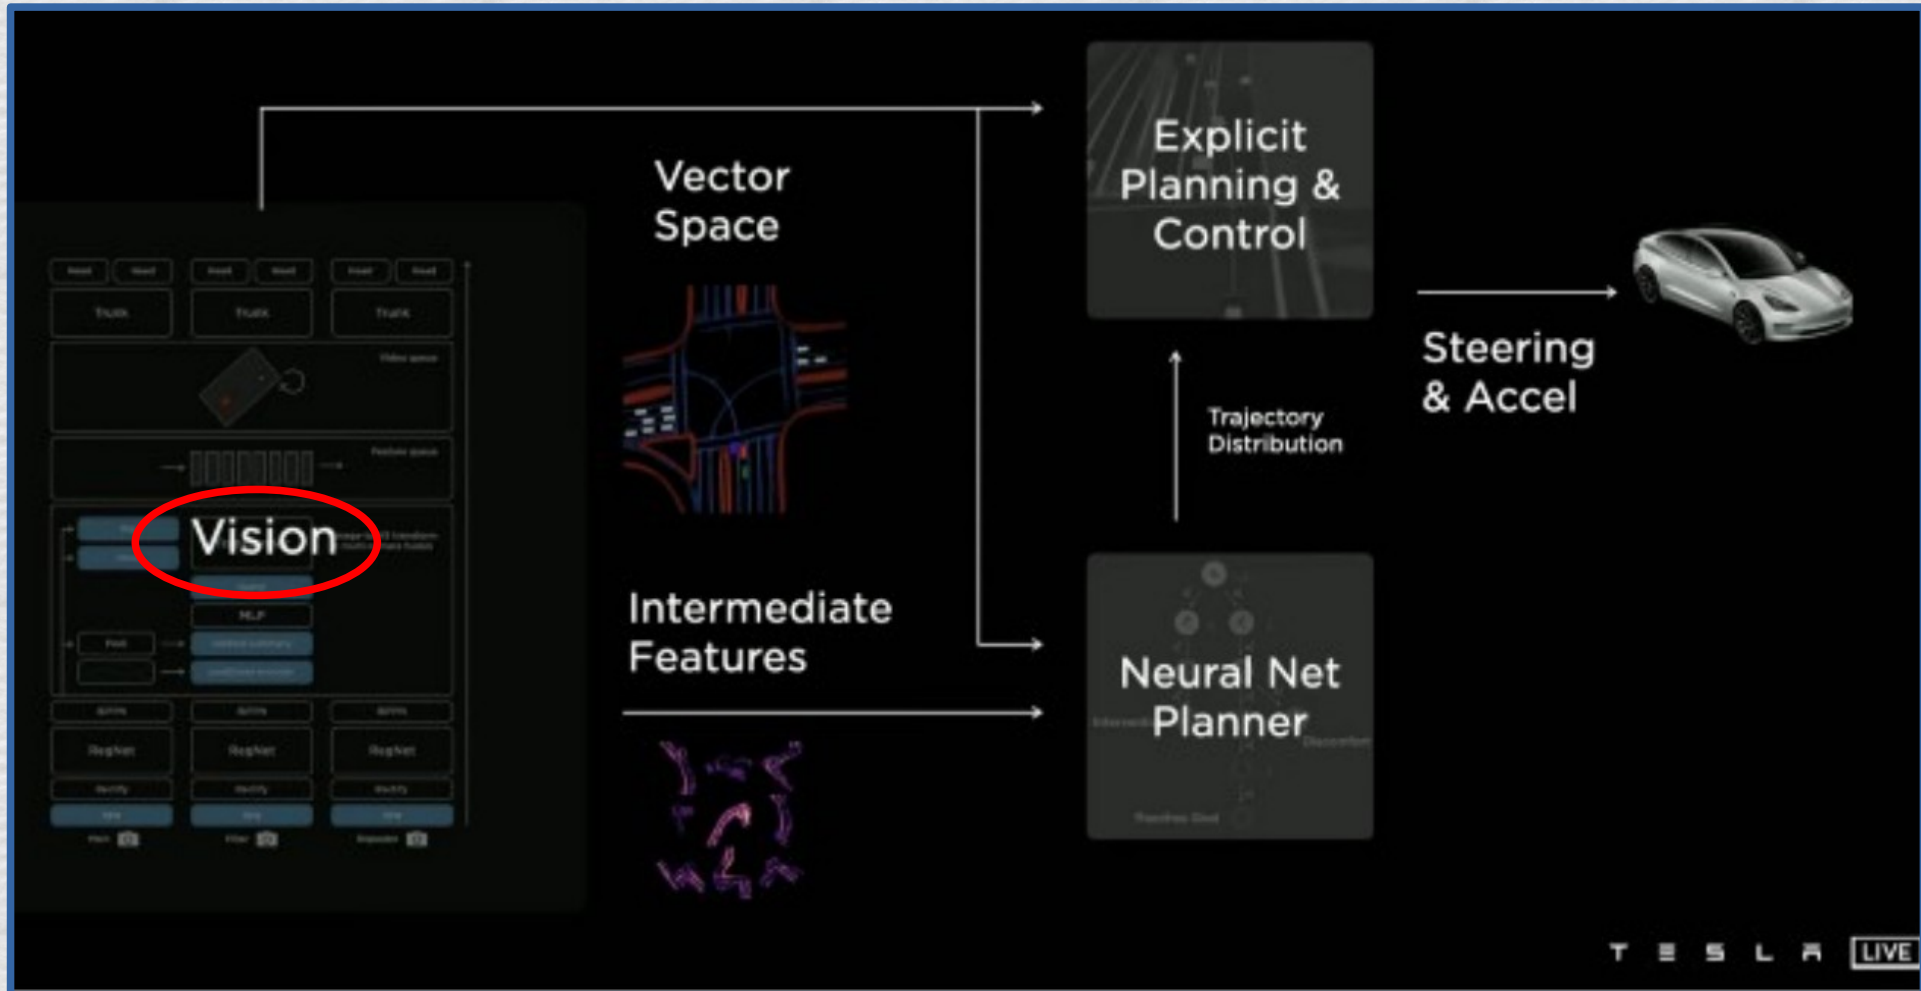


Stills

| Approach | Accuracy (%) |
|-------------------------------|---------------|
| BossaNova (HueSIFT) [2] | 89.5±1 |
| BossaNova VD (BinBoost16) [3] | 90.9±1 |
| Proposed ANet | 92.01±3 |
| Proposed GNet | 93.7±3 |
| Proposed AGNet | 93.8±2 |
| Proposed AGbNet | 94.1±2 |

Videos

Tesla AutoPilot



Tesla AutoPilot

Autopilot is a complex system.

...but nearly the first thing that happens to the data from each of the 8 cameras is that it's passed through a **RegNet**. This is a development of the ResNet system.



Tesla AutoPilot




Autopilot is a complex system.


...but nearly the first thing that happens to the data from each of the 8 cameras is that it's passed through a **RegNet**. This is a development of the ResNet system.




Why so deep?



The paradox



A multilayer *feedforward network*, with sufficient hidden units can represent any deterministic mapping between its inputs and its outputs - Hornik, Stinchcombe & White (1989)



ResNet152 (!)

Lower complexity

2-layer dense network

150,000 inputs
75,000 hidden
1,000 outputs


= 11 billion connections

50-layer ResNet


150,000 inputs
1,000 outputs

24 million connections
400 times simpler


Overfitting



Underfitted




Good Fit/Robust




Overfitted

Overfitting





The State Of The Art



How good are they?

“recent advances from machine learning led to the discovery of hierarchical neural network models that achieved near-human-level performance level on challenging object categorization tasks”

- Yamins & DiCarlo (2016)

Can you be more specific?

PNASNet:

- 96.2% Top-5 accuracy on ImageNet (Liu et al., 2018)

- (1) Hatstand
- (2) Orange
- (3) Battleship
- (4) Dandelion
- (5) Cat




How about a sensible answer?

PNASNet:

~73% Top-1 accuracy on ~300 ImageNet categories
(Barbu et al., 2019)

- (1) Cat
- (2) Orange
- (3) Battleship
- (4) Dandelion
- (5) Hatstand



Barbu et al. (2019)




Internet objects
72% correct



Objects in the real world
30 % correct



How good are people?




OpenSesame (legacy backend)

1 = BASKET
2 = CANDLE
3 = MOBILE PHONE
4 = PILLOW
5 = PLATE
6 = T-SHIRT
7 = RUBBISH BIN
8 = UMBRELLA
9 = VASE
0 = WATCH


A close-up image of a silver and gold watch with a metal link bracelet, lying on a white sink. The sink has the brand name "SOMANY" visible. The watch face is white with black numerals and a gold-colored case.

Results







How To Improve?



Insufficiently sensitive to shape




Overly sensitive to tiny local features


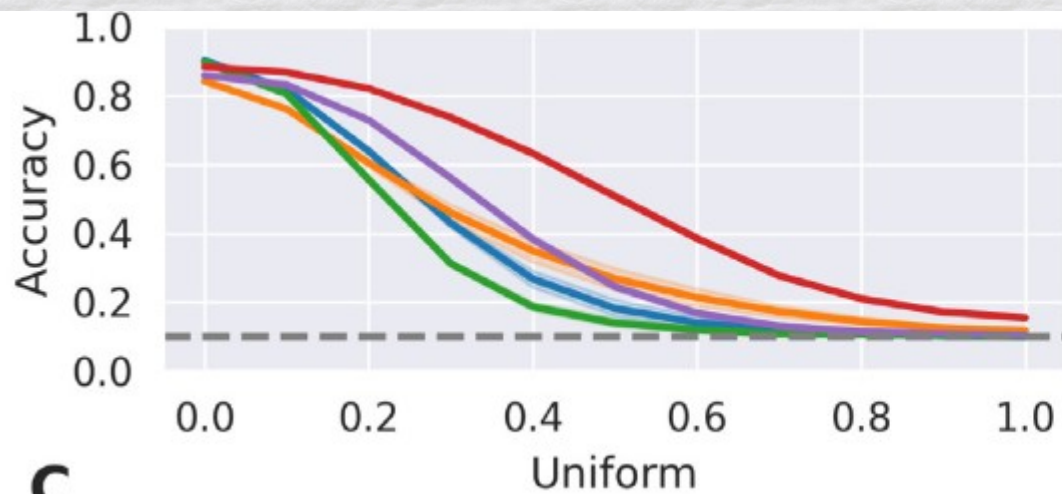
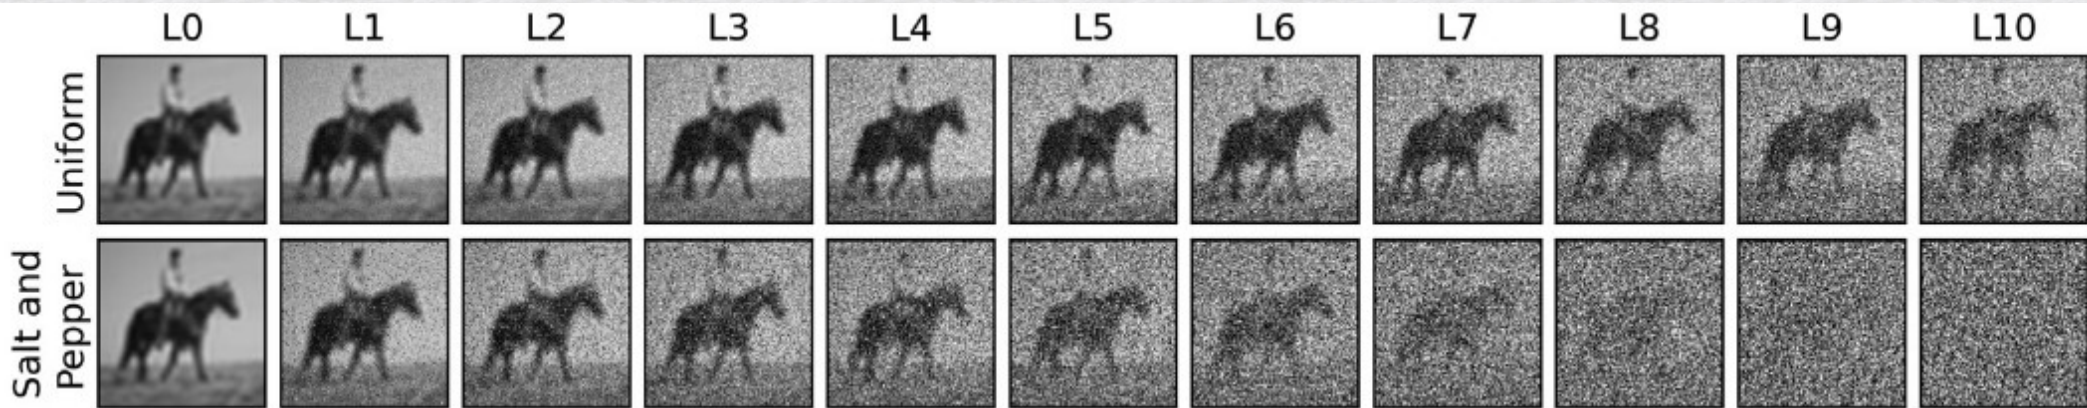


(c) Single diagnostic pixel

Evans et al. (2022)



Evans et al. (2022)



C

D



Summary

