CSCI E-11

Class 9

March 31, 2026

TOR and Harms/AI Vision

Harvard
Extension School
HARVARD DIVISION OF
CONTINUING EDUCATION

Guest Lecturer!

Brian Levine
Associate Dean of IT & Facilities
Distinguished Professor
University of Massachusetts Amherst

https://www.cics.umass.edu/about/directory/brian-levine

News story: OpenAI shut down video generation service Sora

Sora @soraofficialapp

We're saying goodbye to the Sora app. To everyone who created with Sora, shared it, and built community around it: thank you. What you made with Sora mattered, and we know this news is disappointing.

We'll share more soon, including timelines for the app and API and details on preserving your work. – The Sora Team

Last edited 4:52 PM · Mar 24, 2026 · **48.1M** Views

https://x.com/soraofficialapp/status/2036546752535470382

+AI +REPORT +ANALYSIS

Why OpenAI killed Sora

Too much compute, too much competition, and skeptical investors.

by + **Hayden Field**
Mar 28, 2026, 8:00 AM EDT

37 Comments (All New)

https://www.theverge.com/ai-artificial-intelligence/902368/openai-sora-dead-ai-video-generation-competition

Last Week's Reading

Last Week's Reading

SEE CORRECTION FOR THIS ARTICLE

Measuring the predictability of life outcomes with a scientific mass collaboration

Matthew J. Salganik[a,1], Ian Lundberg[a], Alexander T. Kindel[a], Caitlin E. Ahearn[b], Khaled Al-Ghoneim[c], Abdullah Almaatouq[d,e], Drew M. Altschul[f], Jennie E. Brand[b,g], Nicole Bohme Carnegie[h], Ryan James Compton[i], Debanjan Datta[j], Thomas Davidson[k], Anna Filippova[l], Connor Gilroy[m], Brian J. Goode[n], Eaman Jahani[o], Ridhi Kashyap[p,q,r], Antje Kirchner[s], Stephen McKay[t], Allison C. Morgan[u], Alex Pentland[e], Kivan Polimis[v], Louis Raes[w], Daniel E. Rigobon[x], Claudia V. Roberts[y], Diana M. Stanescu[z], Yoshihiko Suhara[e], Adaner Usmani[aa], Erik H. Wang[z], Muna Adem[bb], Abdulla Alhajri[cc], Bedoor AlShebli[dd], Redwane Amin[ee], Ryan B. Amos[y], Lisa P. Argyle[ff], Livia Baer-Bositis[gg], Moritz Büchi[hh], Bo-Ryehn Chung[ii], William Eggert[jj], Gregory Faletto[kk], Zhilin Fan[ll], Jeremy Freese[gg], Tejomay Gadgil[mm], Josh Gagné[gg], Yue Gao[nn], Andrew Halpern-Manners[bb], Sonia P. Hashim[y], Sonia Hausen[gg], Guanhua He[oo], Kimberly Higuera[gg], Bernie Hogan[pp], Ilana M. Horwitz[qq], Lisa M. Hummel[gg], Naman Jain[x], Kun Jin[rr], David Jurgens[ss], Patrick Kaminski[bb,tt], Areg Karapetyan[uu,vv], E. H. Kim[gg], Ben Leizman[y], Naijia Liu[z], Malte Möser[y], Andrew E. Mack[z], Mayank Mahajan[y], Noah Mandell[ww], Helge Marahrens[bb], Diana Mercado-Garcia[qq], Viola Mocz[xx], Katariina Mueller-Gastell[gg], Ahmed Musse[yy], Qiankun Niu[ee], William Nowak[zz], Hamidreza Omidvar[aaa], Andrew Or[y], Karen Ouyang[y], Katy M. Pinto[bbb], Ethan Porter[ccc], Kristin E. Porter[ddd], Crystal Qian[y], Tamkinat Rauf[gg], Anahit Sargsyan[eee], Thomas Schaffner[y], Landon Schnabel[gg], Bryan Schonfeld[z], Ben Sender[fff], Jonathan D. Tang[y], Emma Tsurkov[gg], Austin van Loon[gg], Onur Varol[ggg,hhh], Xiafei Wang[iii], Zhi Wang[hhh,jjj], Julia Wang[y], Flora Wang[fff], Samantha Weissman[y], Kirstie Whitaker[kkk,lll], Maria K. Wolters[mmm], Wei Lee Woon[nnn], James Wu[ooo], Catherine Wu[y], Kengran Yang[aaa], Jingwen Yin[ll], Bingyu Zhao[ppp], Chenyun Zhu[ll], Jeanne Brooks-Gunn[qqq,rrr], Barbara E. Engelhardt[y,ii], Moritz Hardt[sss], Dean Knox[z], Karen Levy[ttt], Arvind Narayanan[y], Brandon M. Stewart[a], Duncan J. Watts[uuu,vvv,www], and Sara McLanahan[a,1]

Contributed by Sara McLanahan, January 24, 2020 (sent for review October 1, 2019; reviewed by Sendhil Mullainathan and Brian Uzzi)

How predictable are life trajectories? We investigated this question with a scientific mass collaboration using the common task method; 160 teams built predictive models for six life outcomes using data from the Fragile Families and Child Wellbeing Study, a high-quality birth cohort study. Despite using a rich dataset and applying machine-learning methods optimized for prediction, the best predictions were not very accurate and were only slightly better than those from a simple benchmark model. Within each outcome, prediction error was strongly associated with the family being predicted and weakly associated with the technique used to generate the prediction. Overall, these results suggest practical limits to the predictability of life outcomes in some settings and illustrate the value of mass collaborations in the social sciences.

life course | prediction | machine learning | mass collaboration

Social scientists studying the life course have described social patterns, theorized factors that shape outcomes, and estimated causal effects. Although this research has advanced scientific understanding and informed policy interventions, it is unclear how much it translates into an ability to predict individual life outcomes. Assessing predictability is important for three reasons. First, accurate predictions can be used to target assistance to children and families at risk (1, 2). Second, predictability of a person's life trajectory can indicate social rigidity (3), and efforts to understand differences in predictability across social contexts can stimulate scientific discovery and improve policy-making (4). Finally, efforts to improve predictive performance can spark developments in theory and methods (5).

In order to measure the predictability of life outcomes for children, parents, and households, we created a scientific mass collaboration. Our mass collaboration—the Fragile Families Challenge—used a research design common in machine learning but not yet common in the social sciences: the common task method (6). To create a project using the common task method, an organizer designs a prediction task and then recruits a large, diverse group of researchers who complete the task by predicting the exact same outcomes using the exact same data. These predictions are then evaluated with the exact same error metric that exclusively assesses their ability to predict held-out data: data that are held by the organizer and not available to participants. Although the structure of the prediction task is completely standardized, participants are free to use any technique to generate predictions.

The common task method produces credible estimates of predictability because of its design. If predictability is higher than expected, the results cannot be dismissed because of concerns about overfitting (7) or researcher degrees of freedom (8). Alternatively, if predictability is lower than expected, the results cannot be dismissed because of concerns about the limitations

Author contributions: M.J.S., I.L., A.T.K., J.B.-G., B.E.E., M.H., K.L., A.N., B.M.S., D.J.W., and S. McLanahan designed research; M.J.S., I.L., A.T.K., C.E.A., K.A.-G., A. Almaatouq, D.M.A., J.E.B., N.B.C., R.J.C., D.D., T.D., A.F., C.G., B.J.G., E.J., R.K., A. Kirchner, S. McKay, A.C.M., A.P., K.P., L.R., D.E.R., C.V.R., D.M.S., Y.S., A.U., E.H.W., M.A., A. Alhajri, B.A., R.A., R.B.A., L.P.A., L.B.-B., M.B., B.-R.C., W.E., G.F., Z.F., J.F., T.G., J.G., Y.G., A.H.-M., S.P.H., S.H., G.H., K.H., B.H., I.M.H., L.M.H., N.J., K.J., D.J., P.K., A. Karapetyan, E.H.K., B.L., N.L., M. Möser, A.E.M., M. Mahajan, N.M., H.M., D.M.-G., V.M., K.M.-G., A.M., Q.N., W.N., H.O., A.O., K.O., K.M.P., E.P., K.E.P., C.Q., T.R., A.S., T.S., L.S., B. Schonfeld, B. Sender, J.D.T., E.T., A.v.L., O.V., X.W., Z.W., J. Wang, F.W., S.W., K.W., M.K.W., W.L.W., J. Wu, C.W., K.Y., J.Y., B.Z., and C.Z. analyzed data; and M.J.S., I.L., A.T.K., D.K., and S. McLanahan wrote the paper.

Reviewers: S.M., University of Chicago; and B.U., Northwestern University.

Competing interest statement: B.E.E. is on the scientific advisory boards of Celsius Therapeutics and Freenome, is currently employed by Genomics plc and Freenome, and is on a year leave-of-absence from Princeton University.

This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

Data deposition: Predictions, code, and narrative explanations for valid submissions to the Fragile Families Challenge and code to reproduce the results of this paper are available from Dataverse at https://doi.org/10.7910/DVN/CXSECU. Data used in the Fragile Families Challenge are currently available to approved researchers from the Princeton University Office of Population Research Data Archive at https://opr.princeton.edu/archive/.

See online for related content such as Commentaries.

[1] To whom correspondence may be addressed. Email: mjs3@princeton.edu or mclanaha@princeton.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1915006117/-/DCSupplemental.

First published March 30, 2020.

Downloaded from https://www.pnas.org by 108.26.182.86 on November 9, 2024 from IP address 108.26.182.86.

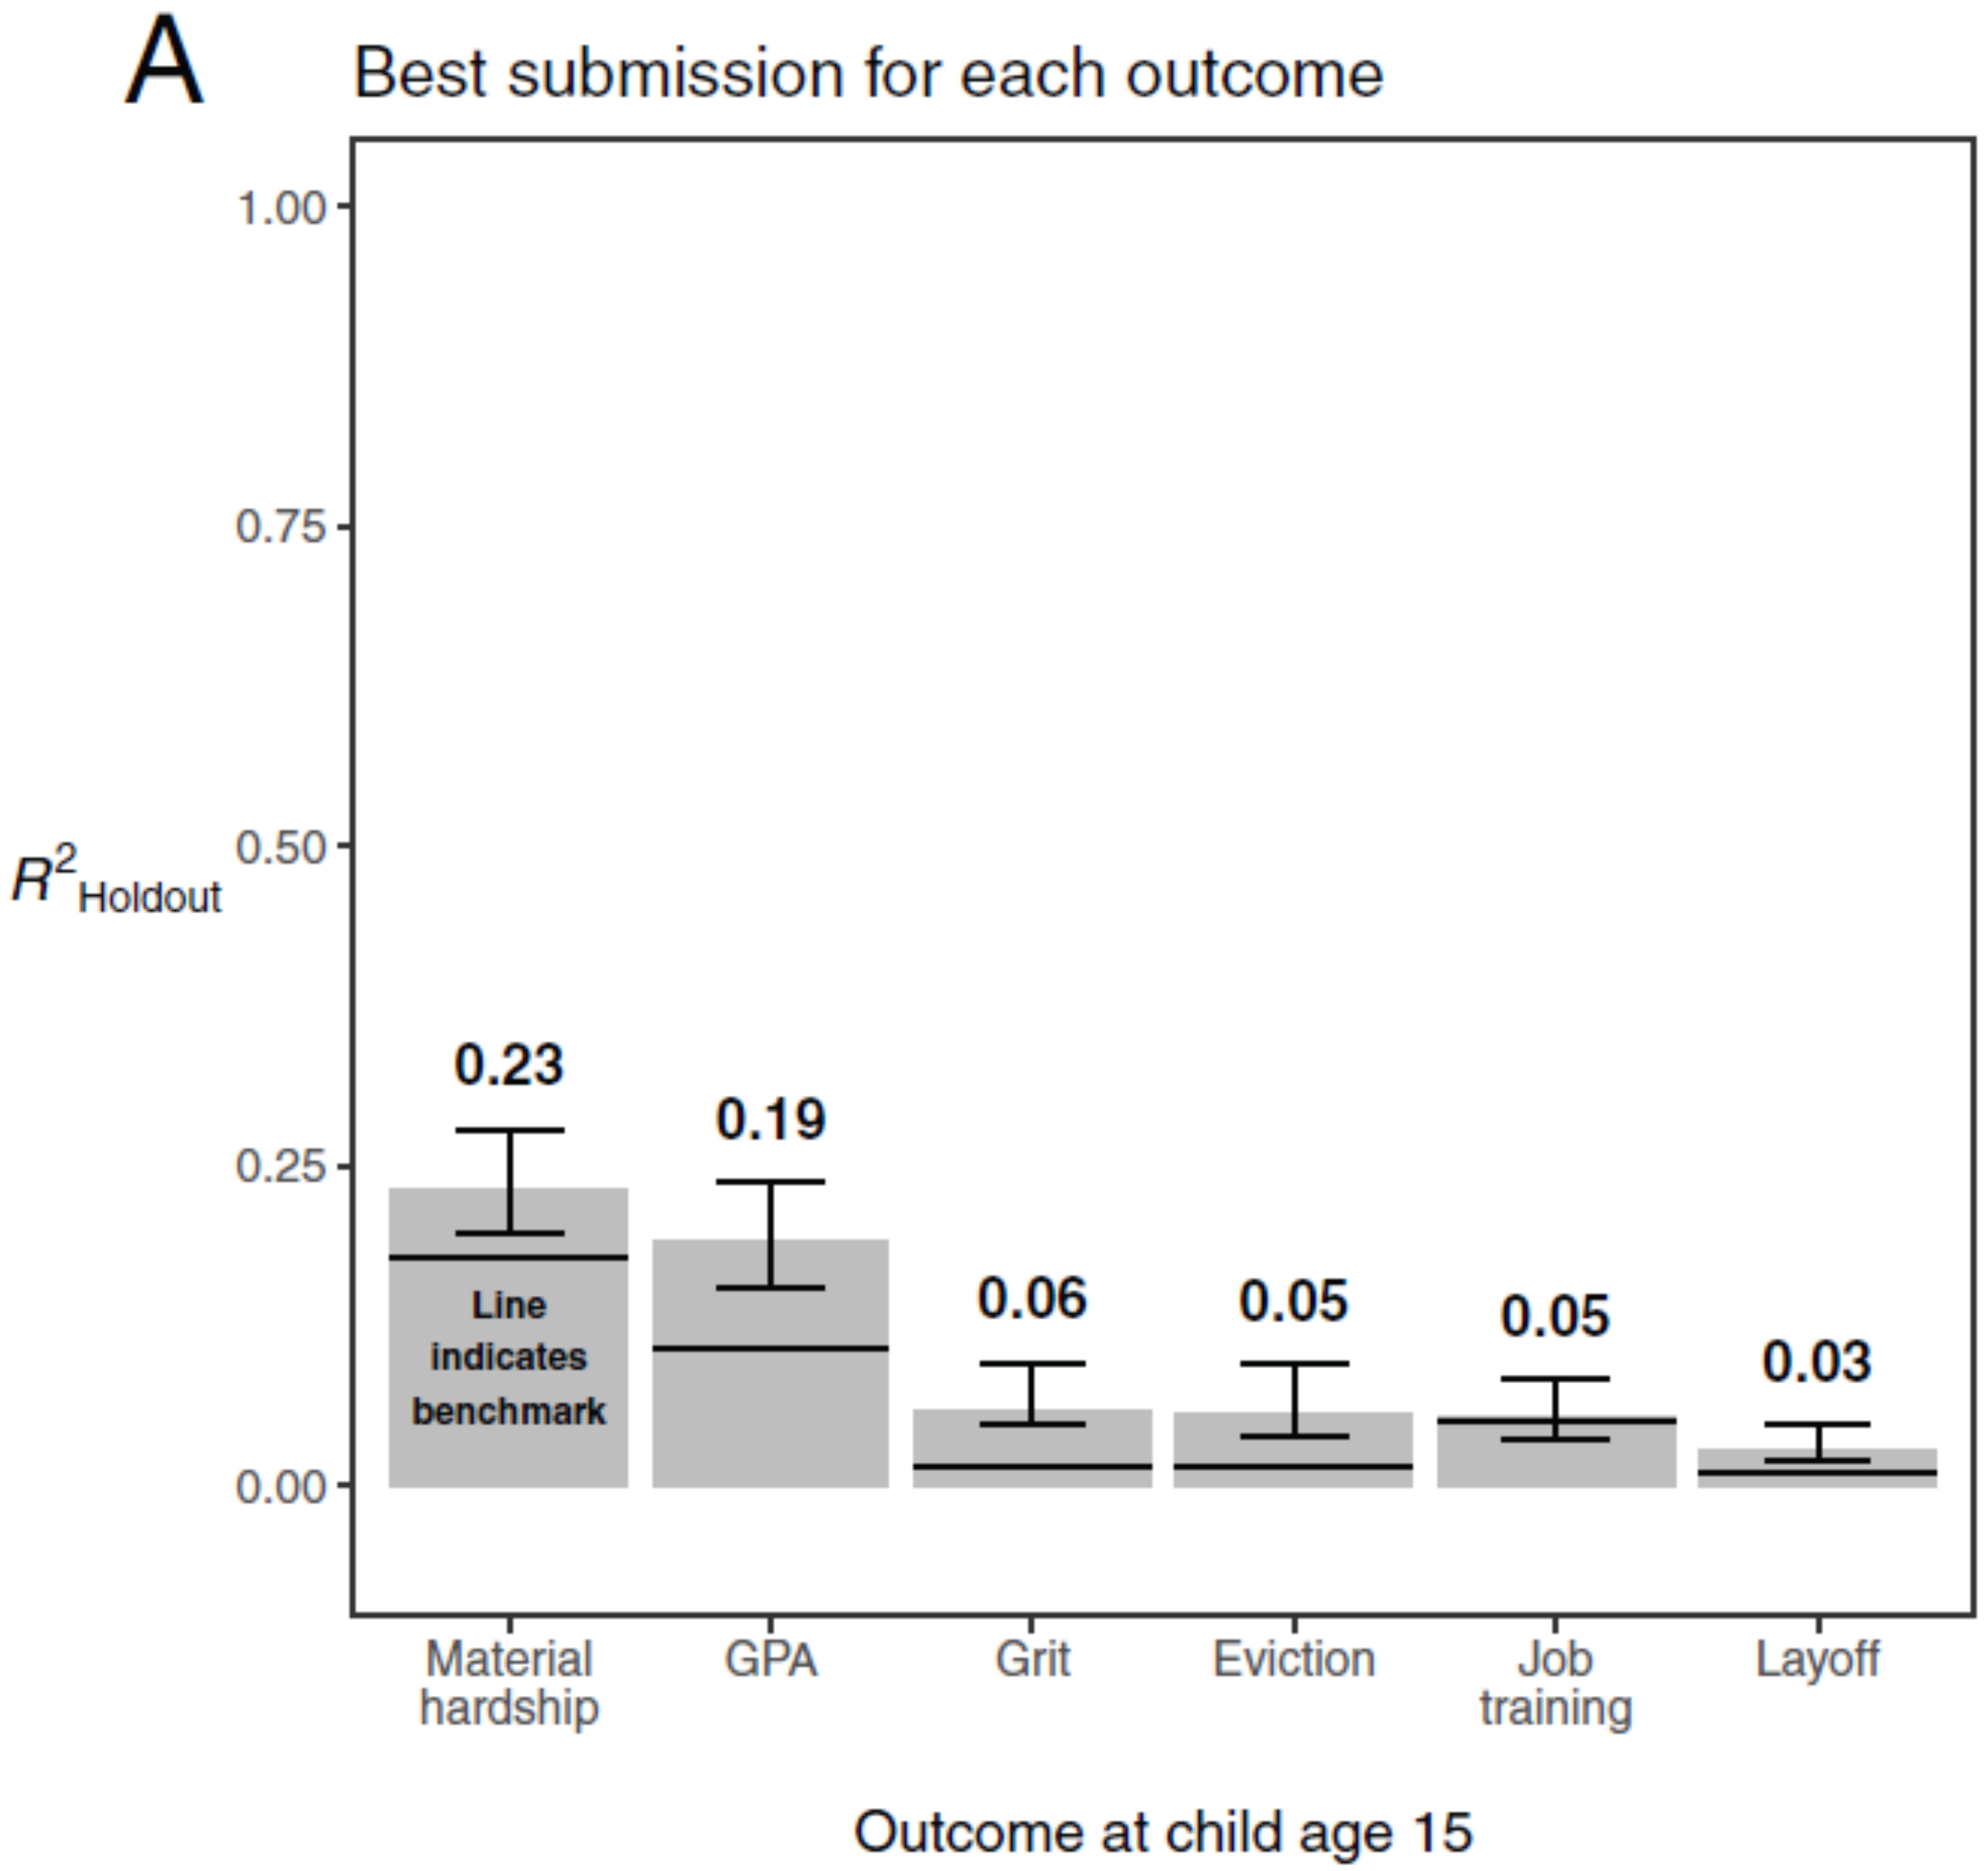

Correlation (XKCD)

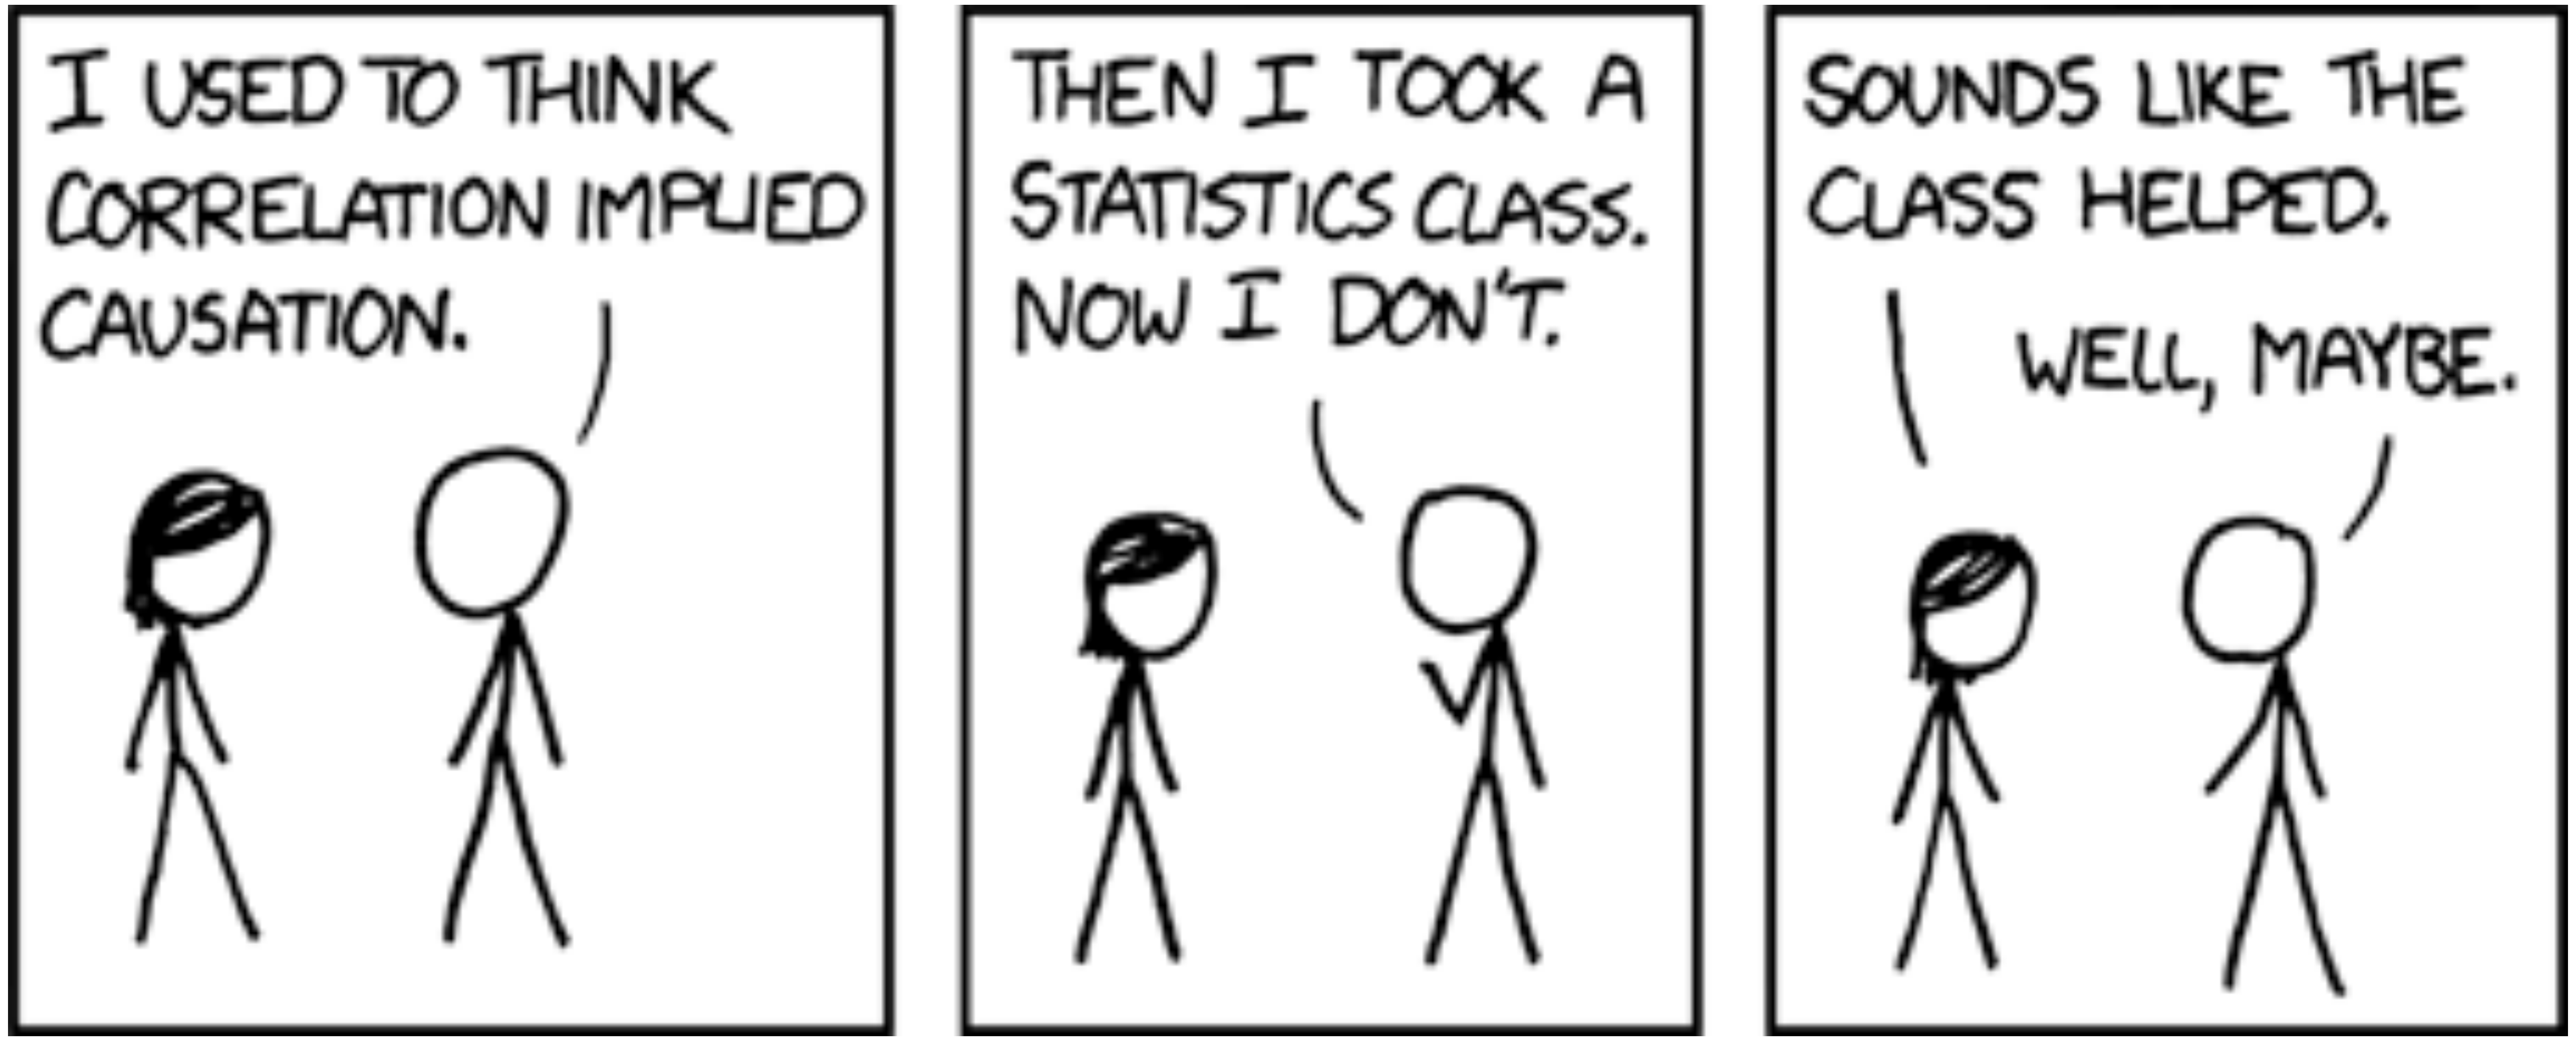

https://xkcd.com/552/

The term “Predictive AI” encompasses analytics and actual future prediction.

MIT Sloan Management Review

MENU

COLUMN

When to Use GenAI Versus Predictive AI

Generative AI doesn’t suit every problem. Use these guidelines to decide between predictive AI — machine learning and deep learning tools — and generative AI.

Rama Ramakrishnan • March 24, 2025

Reading Time: 12 min

Vision

IMAGENET is a knowledge ontology.

- S: (n) Eskimo dog, **husky** (breed of heavy-coated Arctic sled dog)
  - *direct hypernym* / ***inherited hypernym*** / *sister term*
    - S: (n) working dog (any of several breeds of usually large powerful dogs bred to work as draft animals and guard and guide dogs)
      - S: (n) dog, domestic dog, Canis familiaris (a member of the genus Canis (probably descended from the common wolf) that has been domesticated by man since prehistoric times; occurs in many breeds) *"the dog barked all night"*
        - S: (n) canine, canid (any of various fissiped mammals with nonretractile claws and typically long muzzles)
          - S: (n) carnivore (a terrestrial or aquatic flesh-eating mammal) *"terrestrial carnivores have four or five clawed digits on each limb"*
            - S: (n) placental, placental mammal, eutherian, eutherian mammal (mammals having a placenta; all mammals except monotremes and marsupials)
              - S: (n) mammal, mammalian (any warm-blooded vertebrate having the skin more or less covered with hair; young are born alive except for the small subclass of monotremes and nourished with milk)
                - S: (n) vertebrate, craniate (animals having a bony or cartilaginous skeleton with a segmented spinal column and a large brain enclosed in a skull or cranium)
                  - S: (n) chordate (any animal of the phylum Chordata having a notochord or spinal column)
                    - S: (n) animal, animate being, beast, brute, creature, fauna (a living organism characterized by voluntary movement)
                      - S: (n) organism, being (a living thing that has (or can develop) the ability to act or function independently)
                        - S: (n) living thing, animate thing (a living (or once living) entity)
                          - S: (n) whole, unit (an assemblage of parts that is regarded as a single entity) *"how big is that part compared to the whole?"; "the team is a unit"*
                            - S: (n) object, physical object (a tangible and visible entity; an entity that can cast a shadow) *"it was full of rackets, balls and other objects"*
                              - S: (n) physical entity (an entity that has physical existence)
                                - S: (n) entity (that which is perceived or known or inferred to have its own distinct existence (living or nonliving))

amazonmechanical turk
Artificial Artificial Intelligence

Your Account | HITs | Qualifications

Al
Sign in a

Introduction | **Dashboard** | **Status** | **Account Settings**

Mechanical Turk is a marketplace for work.

We give businesses and developers access to an on-demand, scalable workforce.
Workers select from thousands of tasks and work whenever it's convenient.

149,499 HITs available. View them now.

Make Money
by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. Find HITs now.

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work

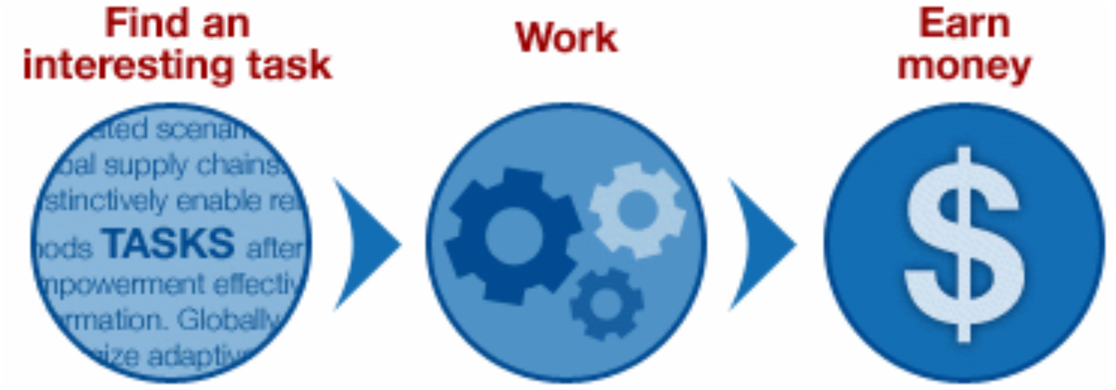

Find HITs Now

or learn more about being a **Worker**

Get Results
from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. Register Now

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results

Get Started

Filters ("Kernels") process the pixels of image to apply different effects.

Operation	Kernel ω	Image result g(x,y)
Identity	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	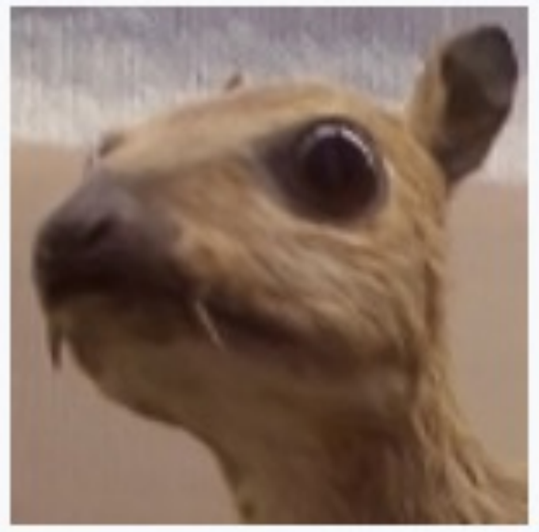
Ridge or edge detection	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	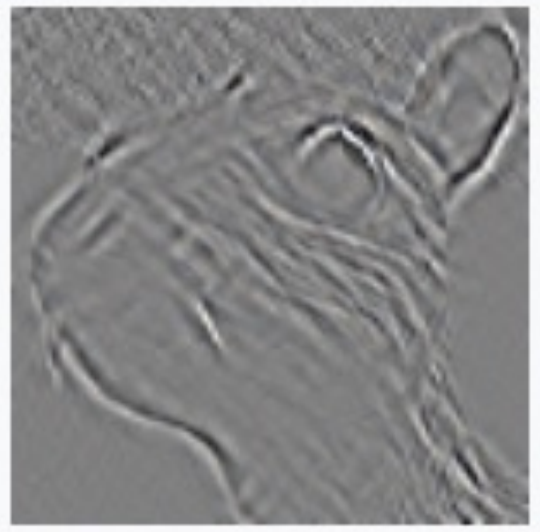
Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	

Traditional fully-connected NNs would be too big to train for images.

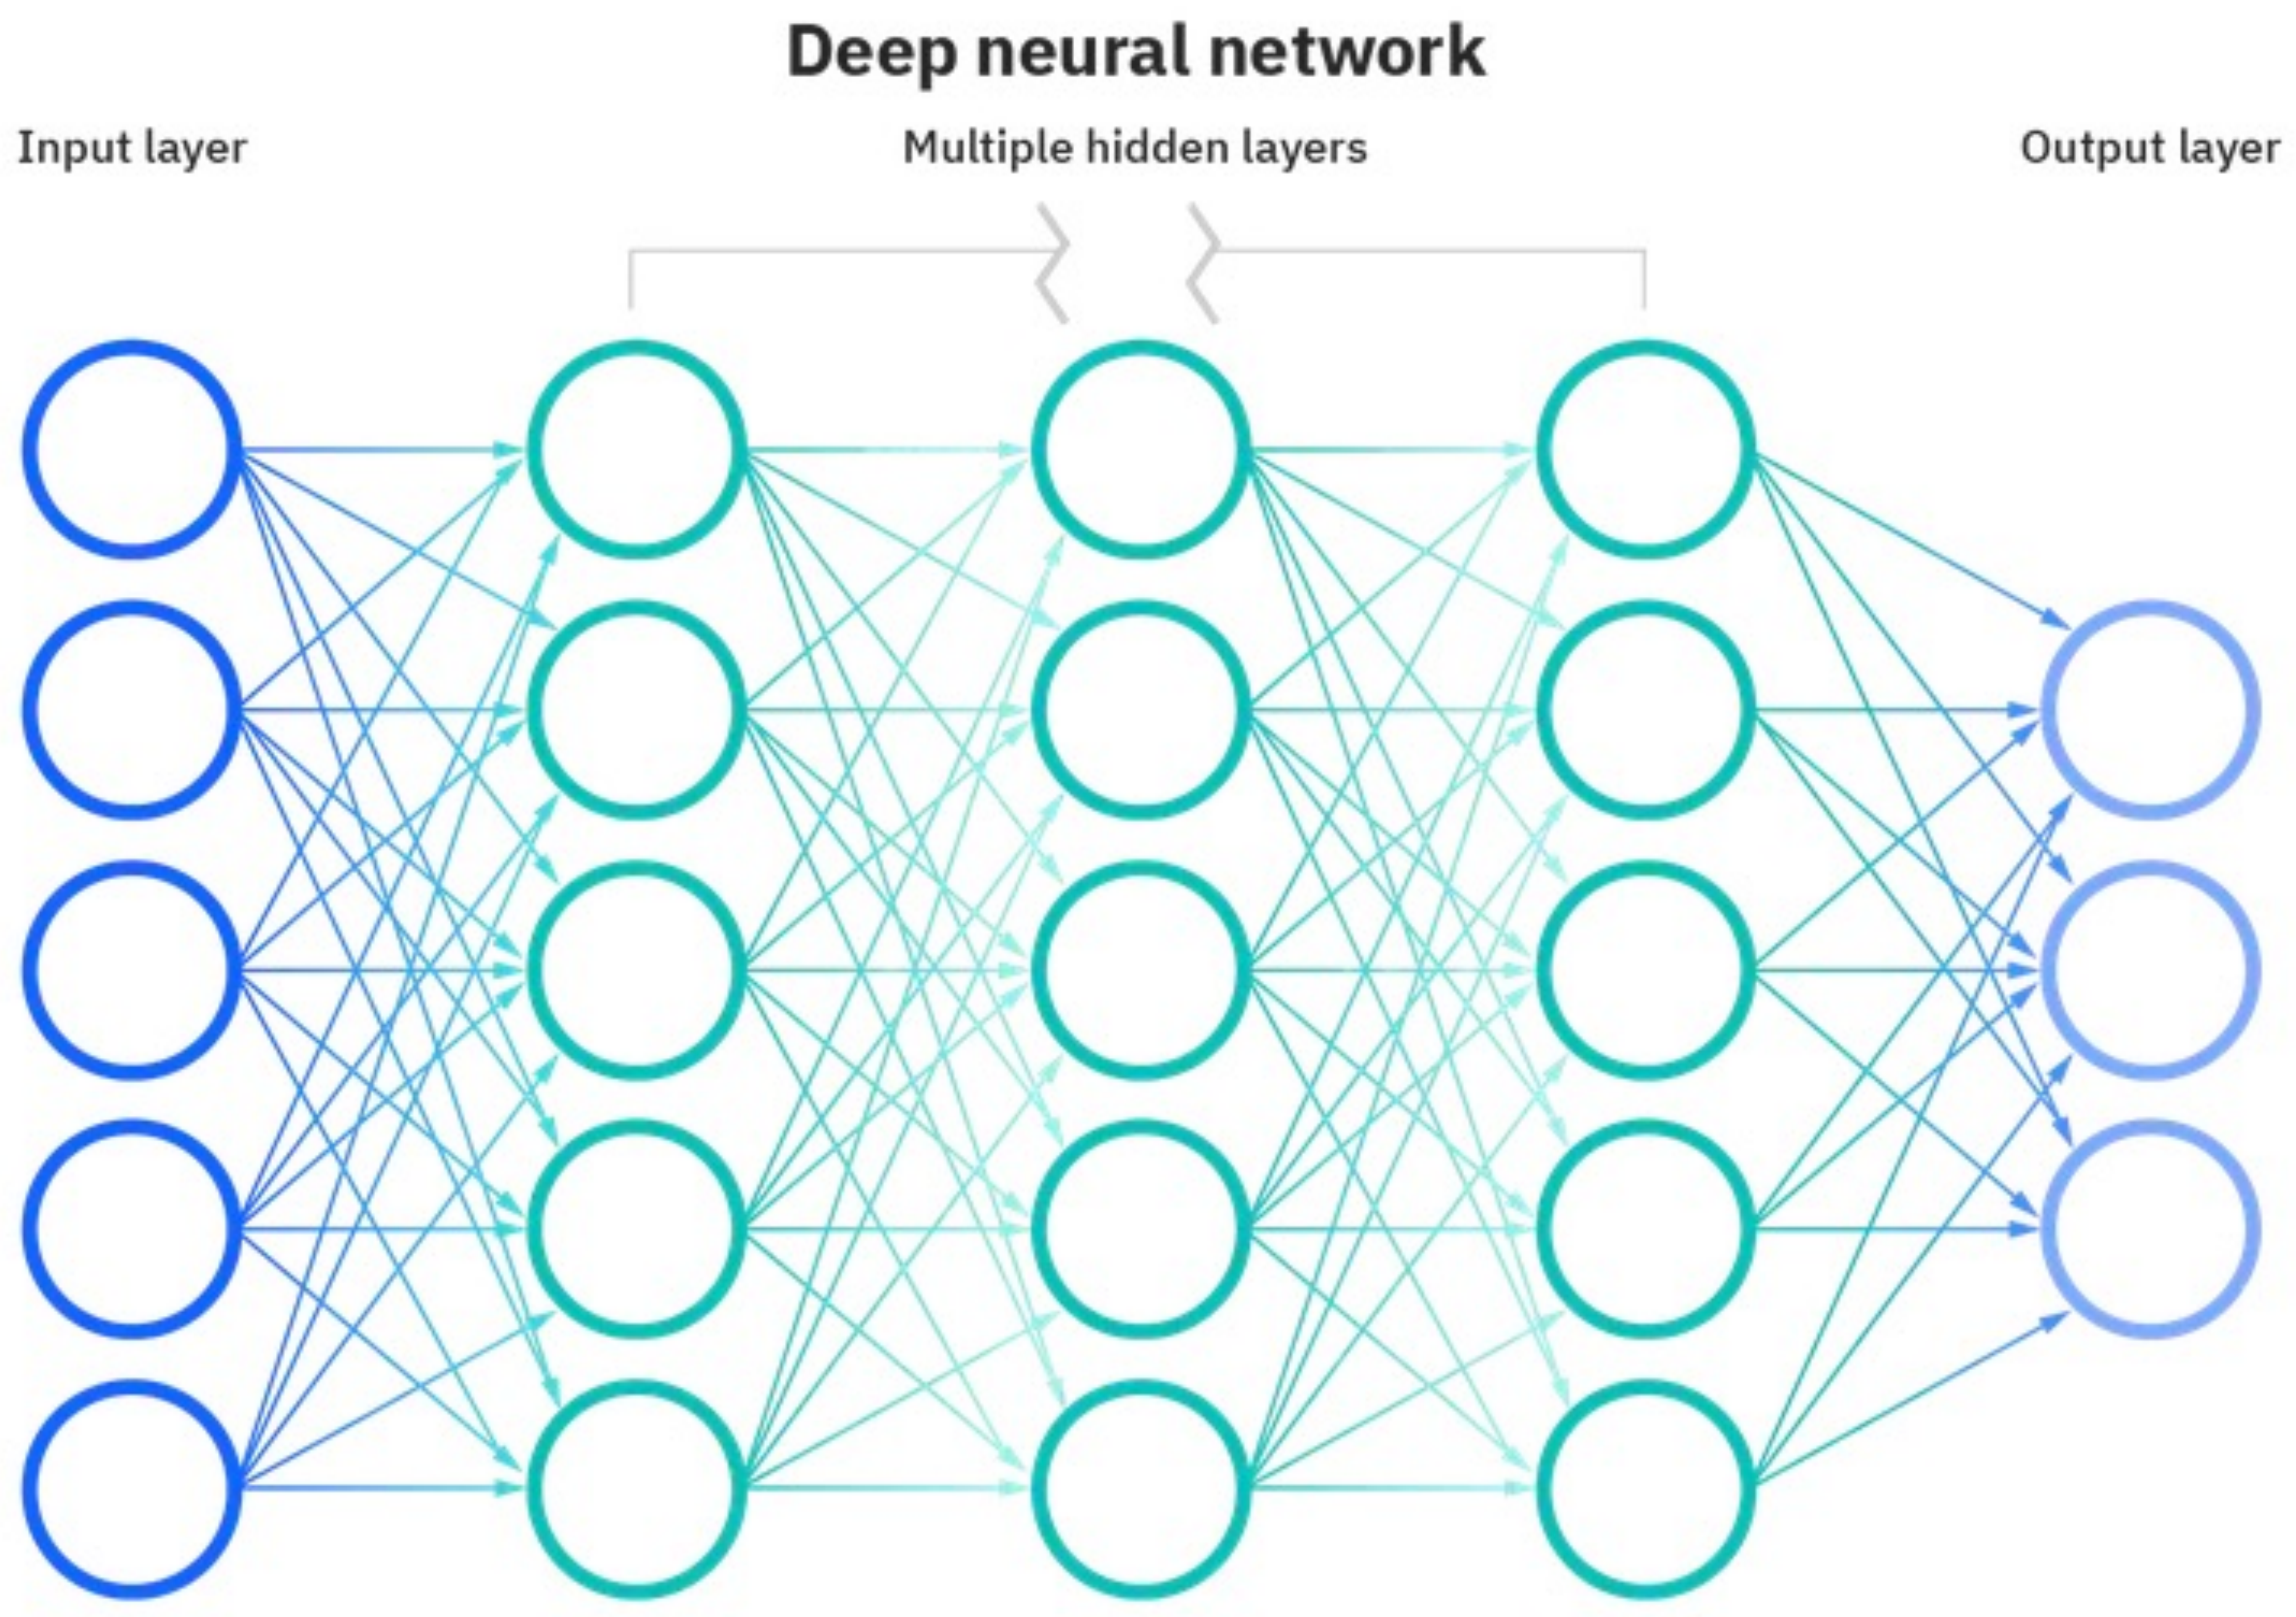

256 pixels

256*256 = 64K neurons

64K * 64K = 4G weights

Convolutional Neural Nets use kernels to produce output values.

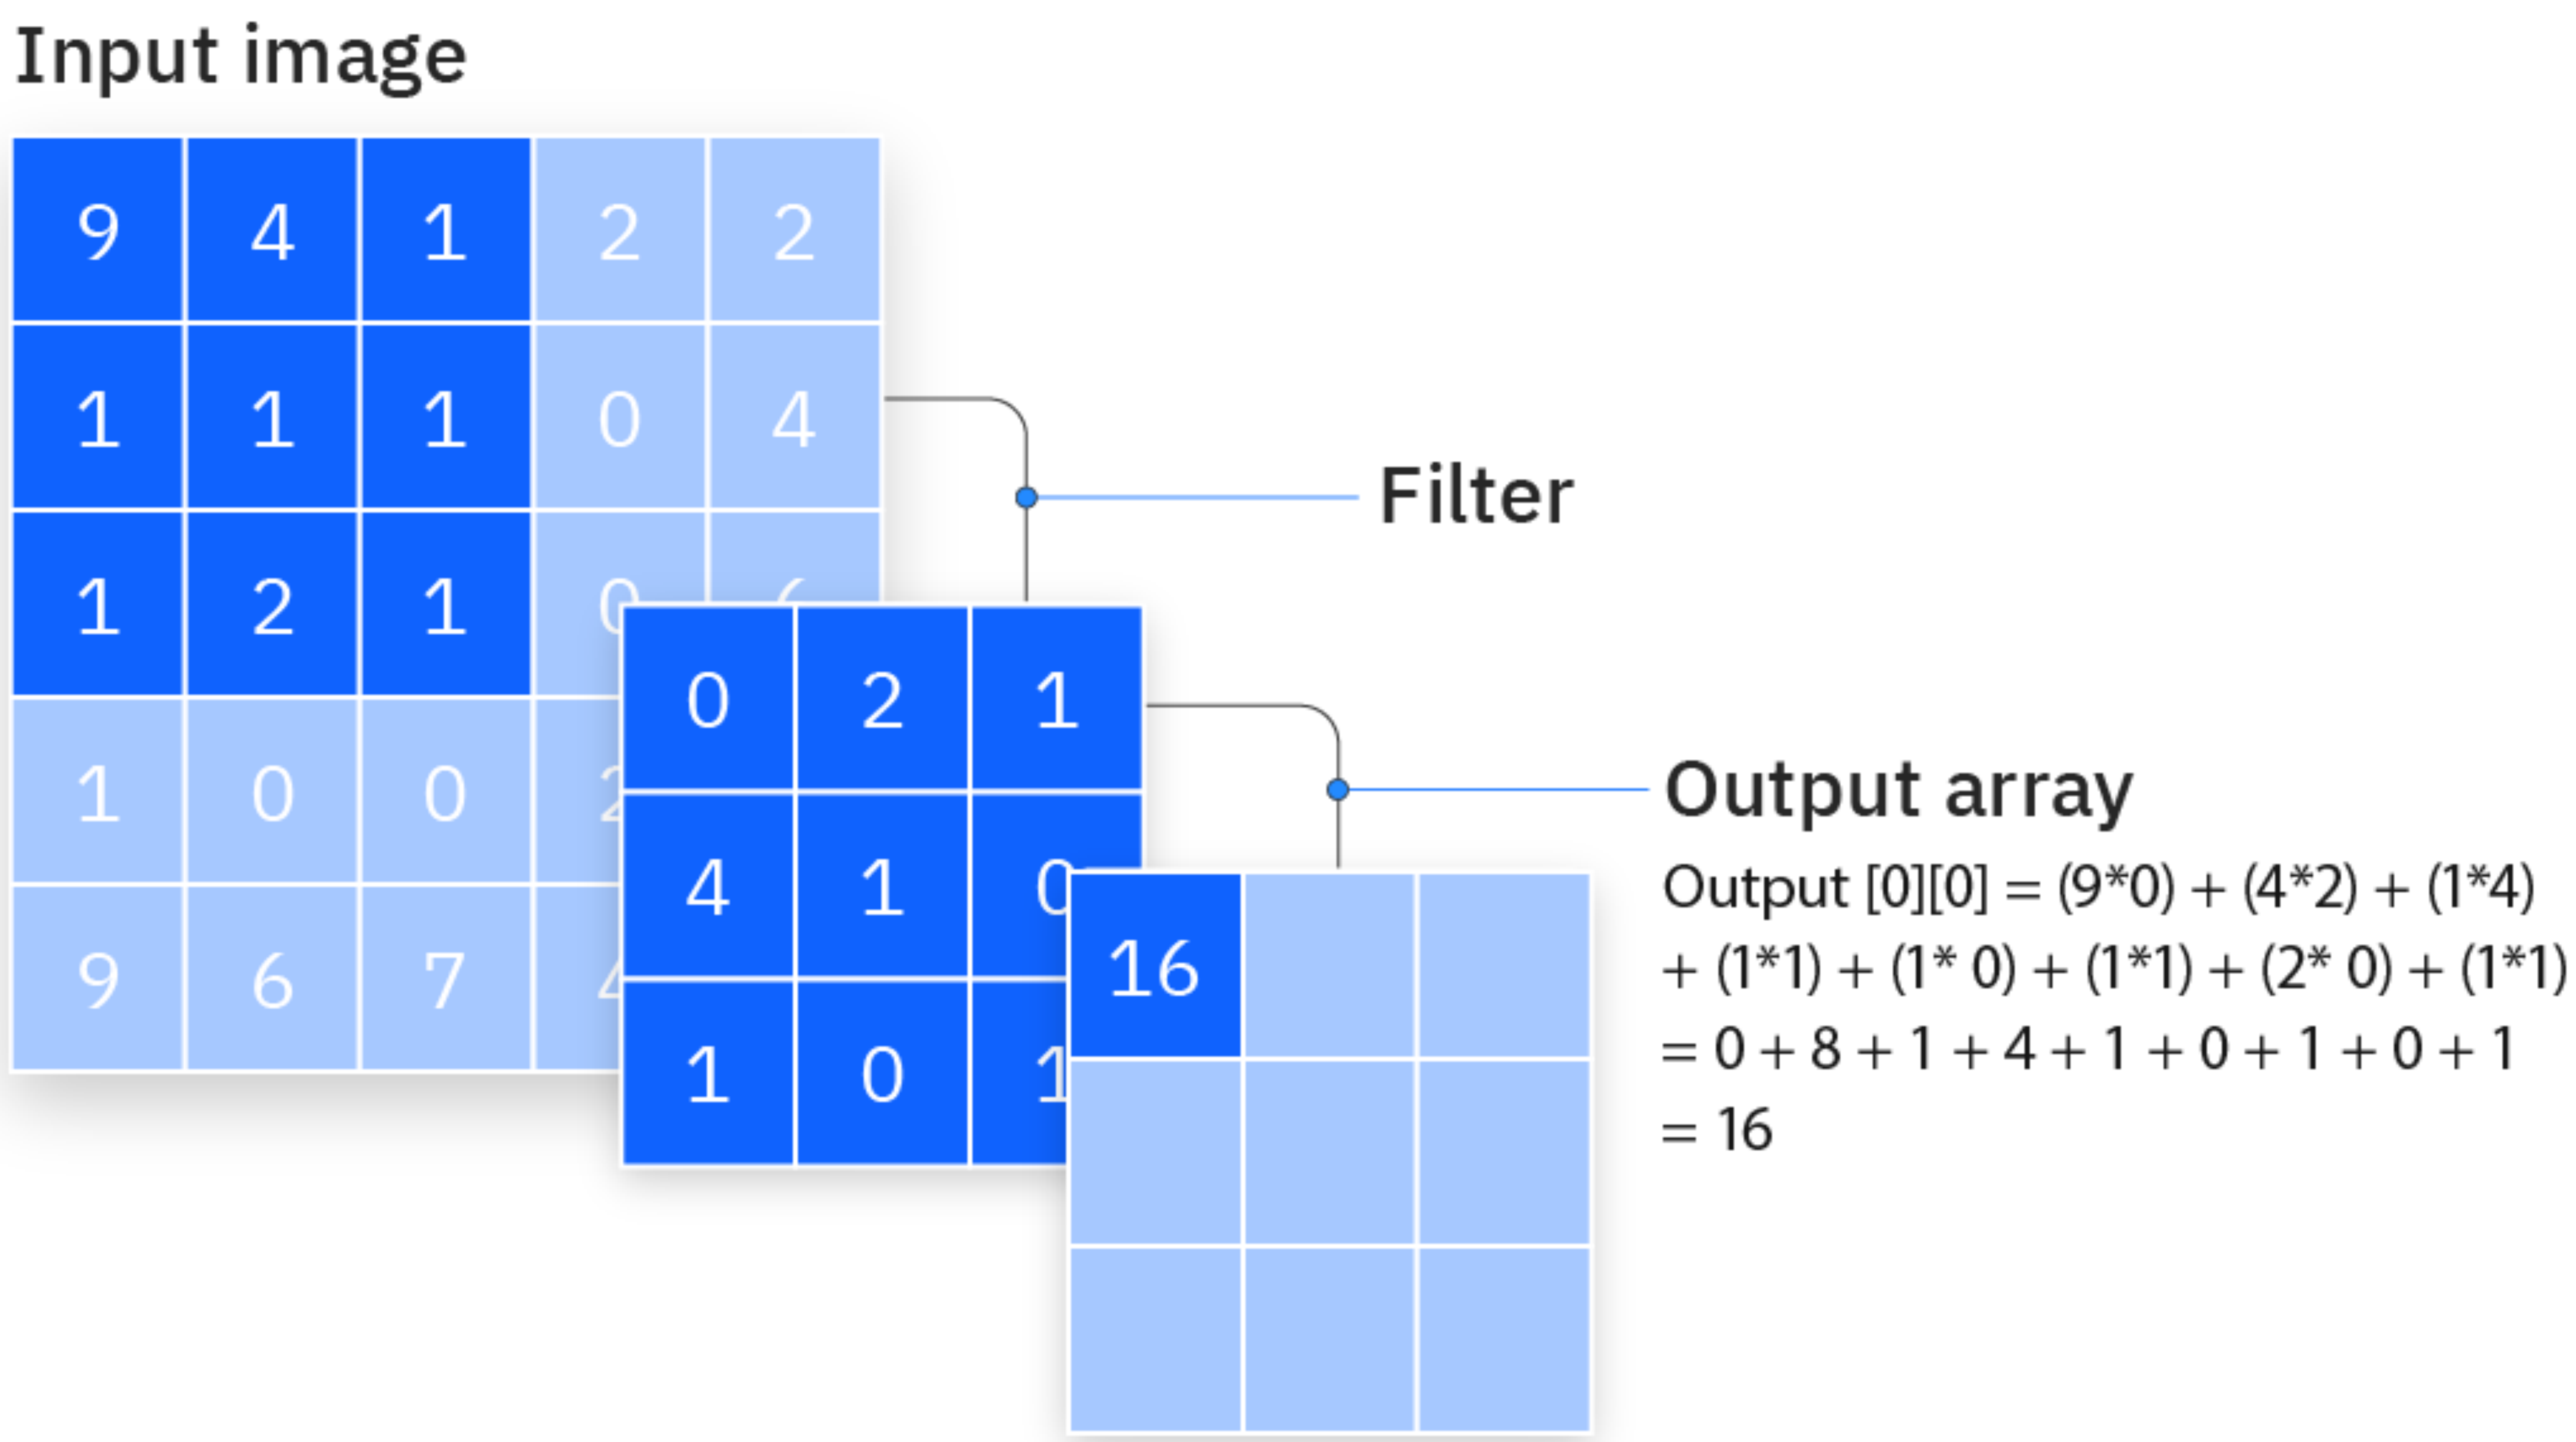

These filters “stride” across the input data to produce a “feature map.”

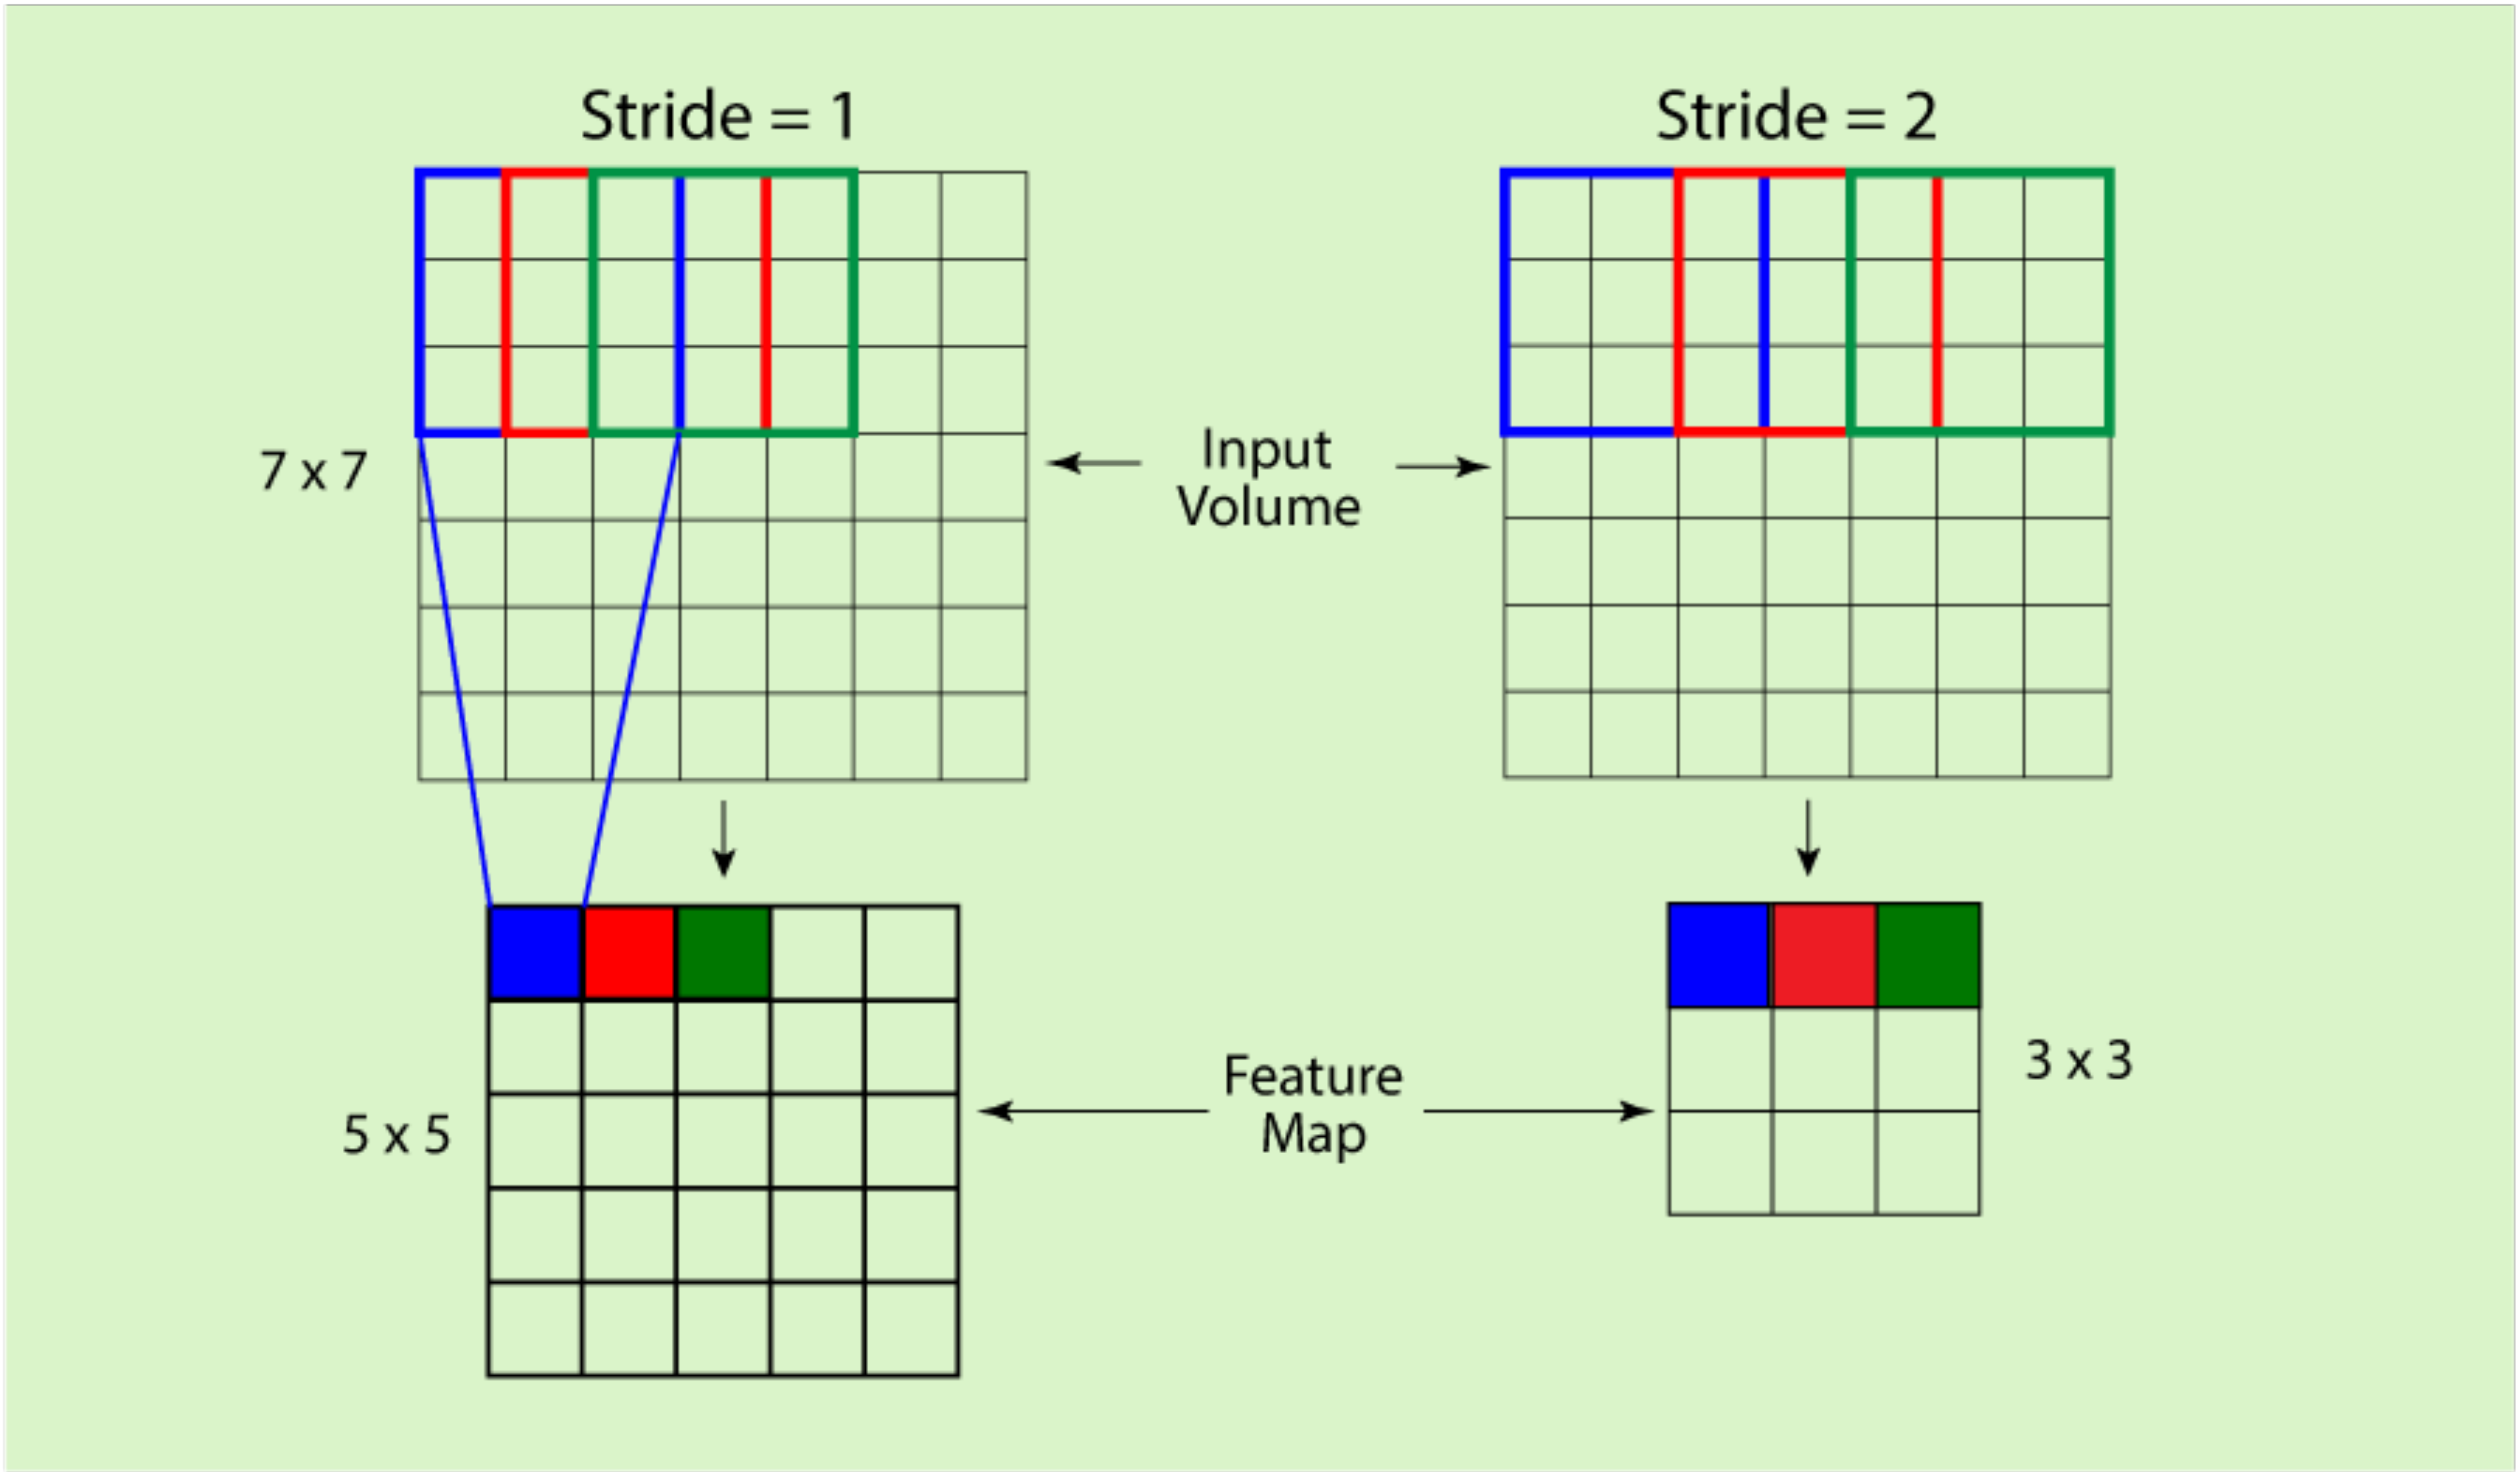

Each filter applied at a layer produces another output, called “depth.”

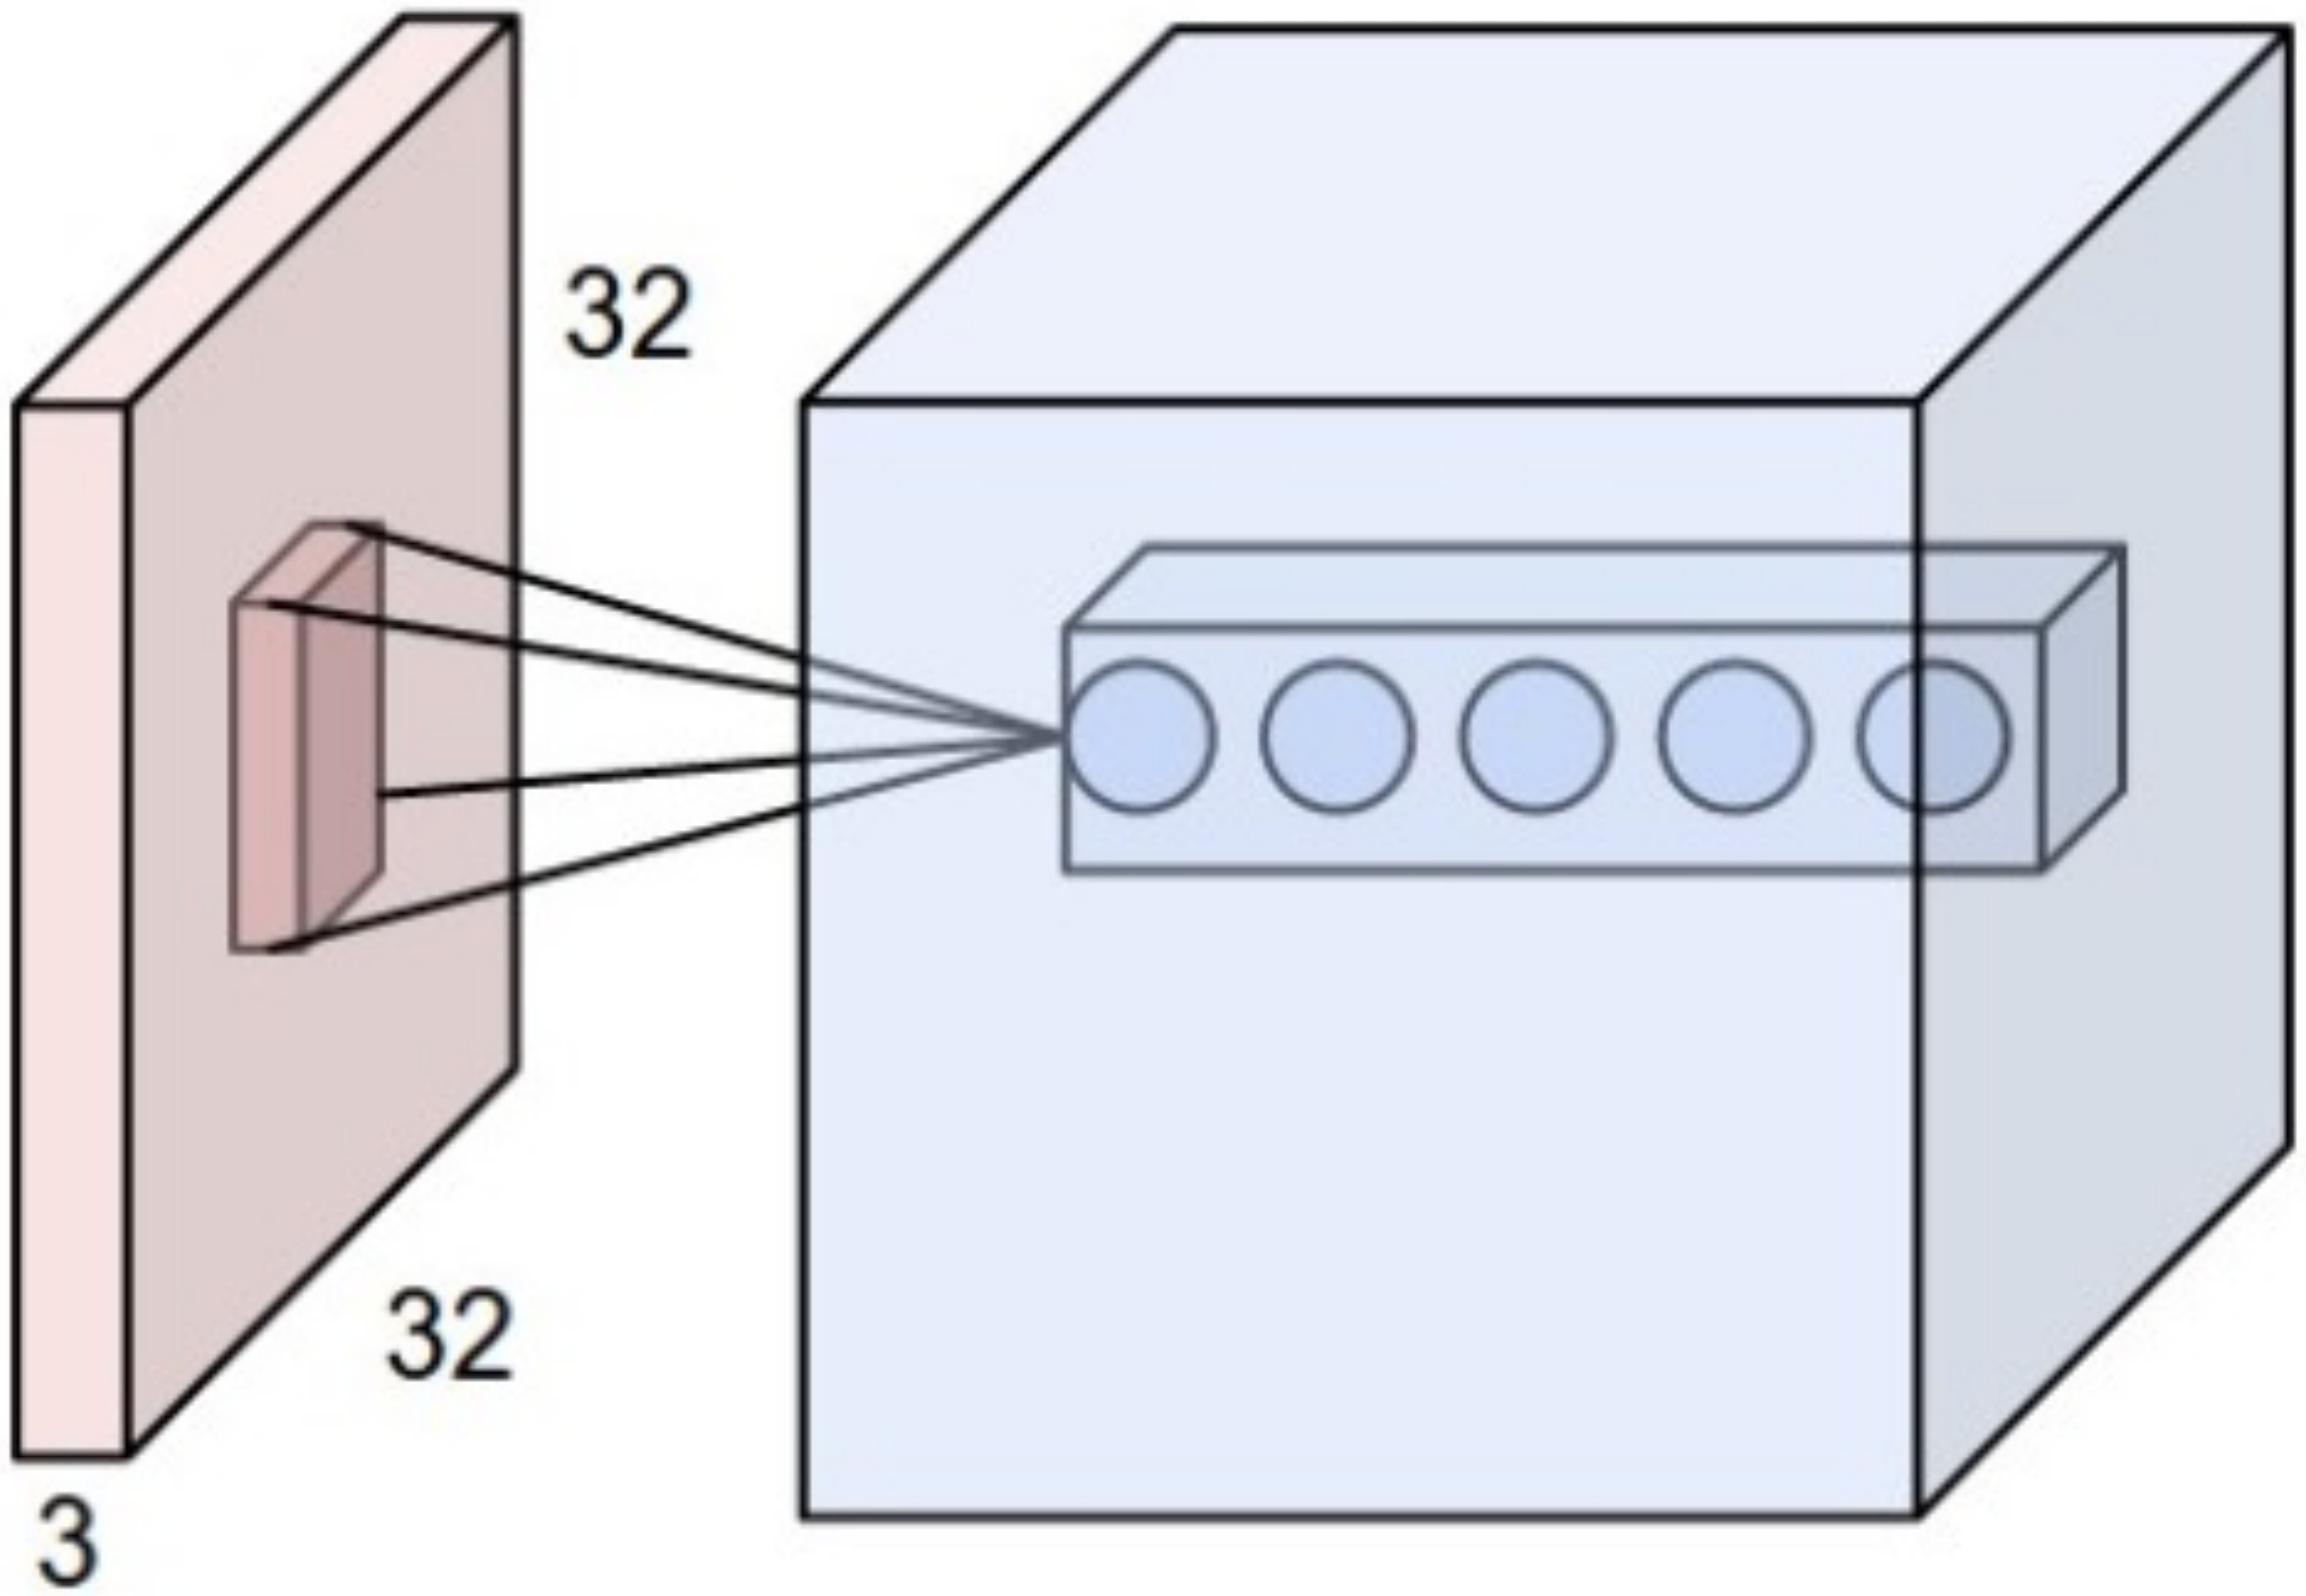

Max pooling reduces the dimensions of the data.

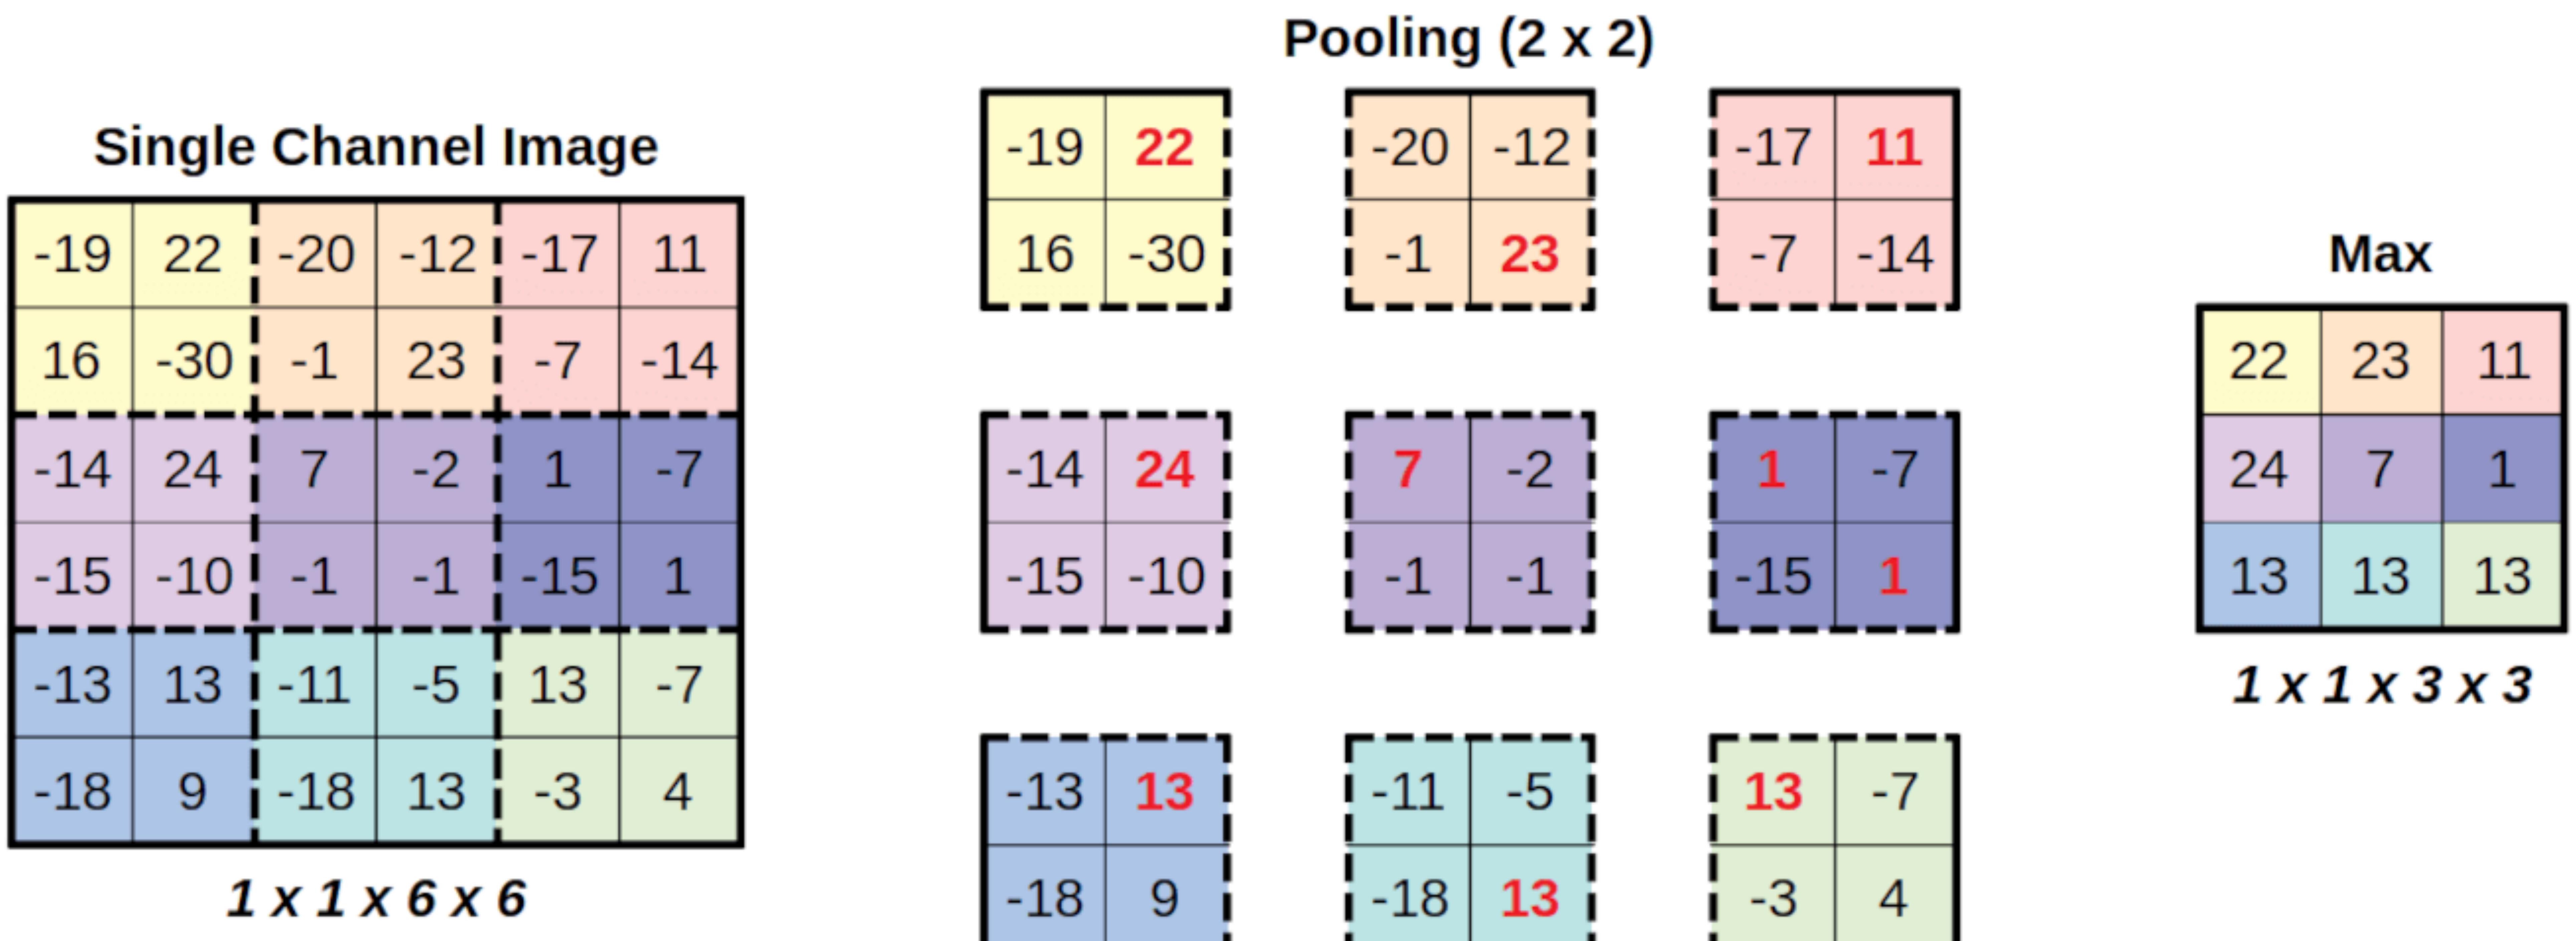

CNN architecture from Krizhevsky combines all of these approaches.

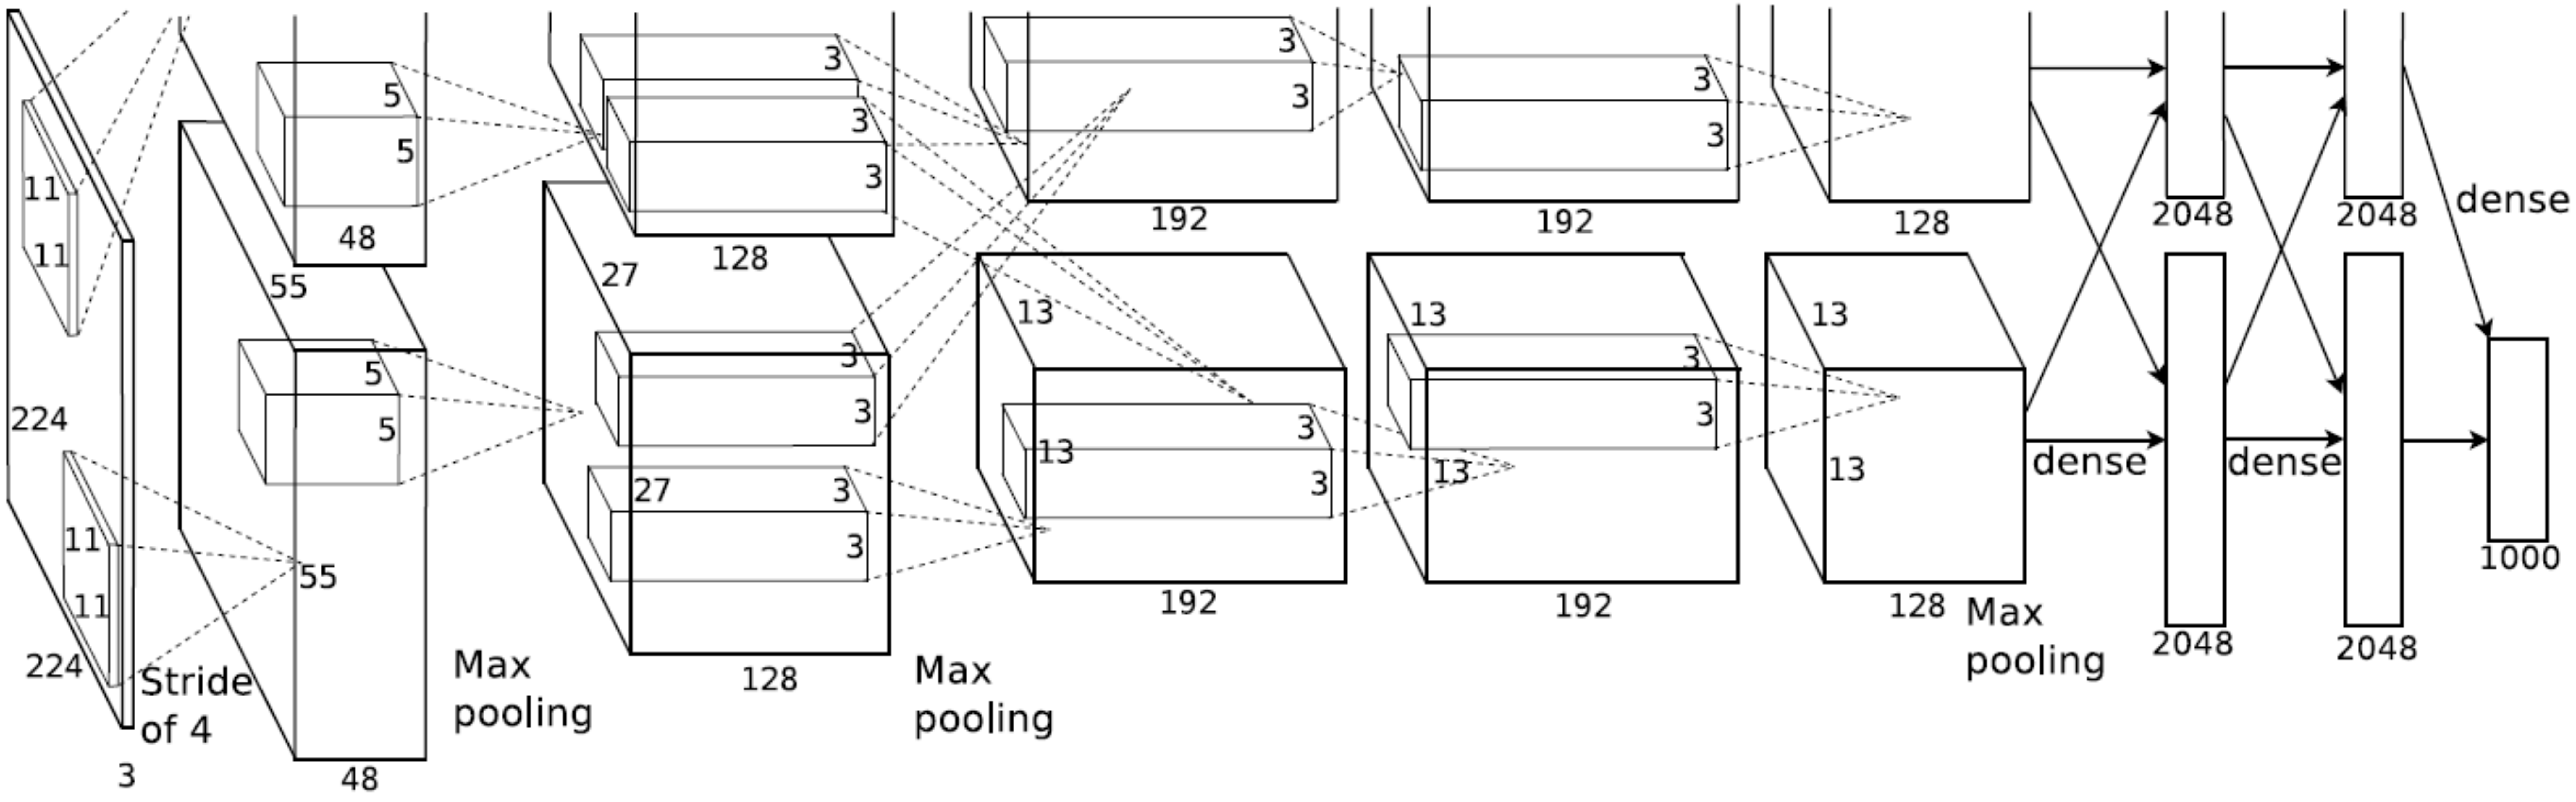

CNNs perform decently well on ImageNet data, but…

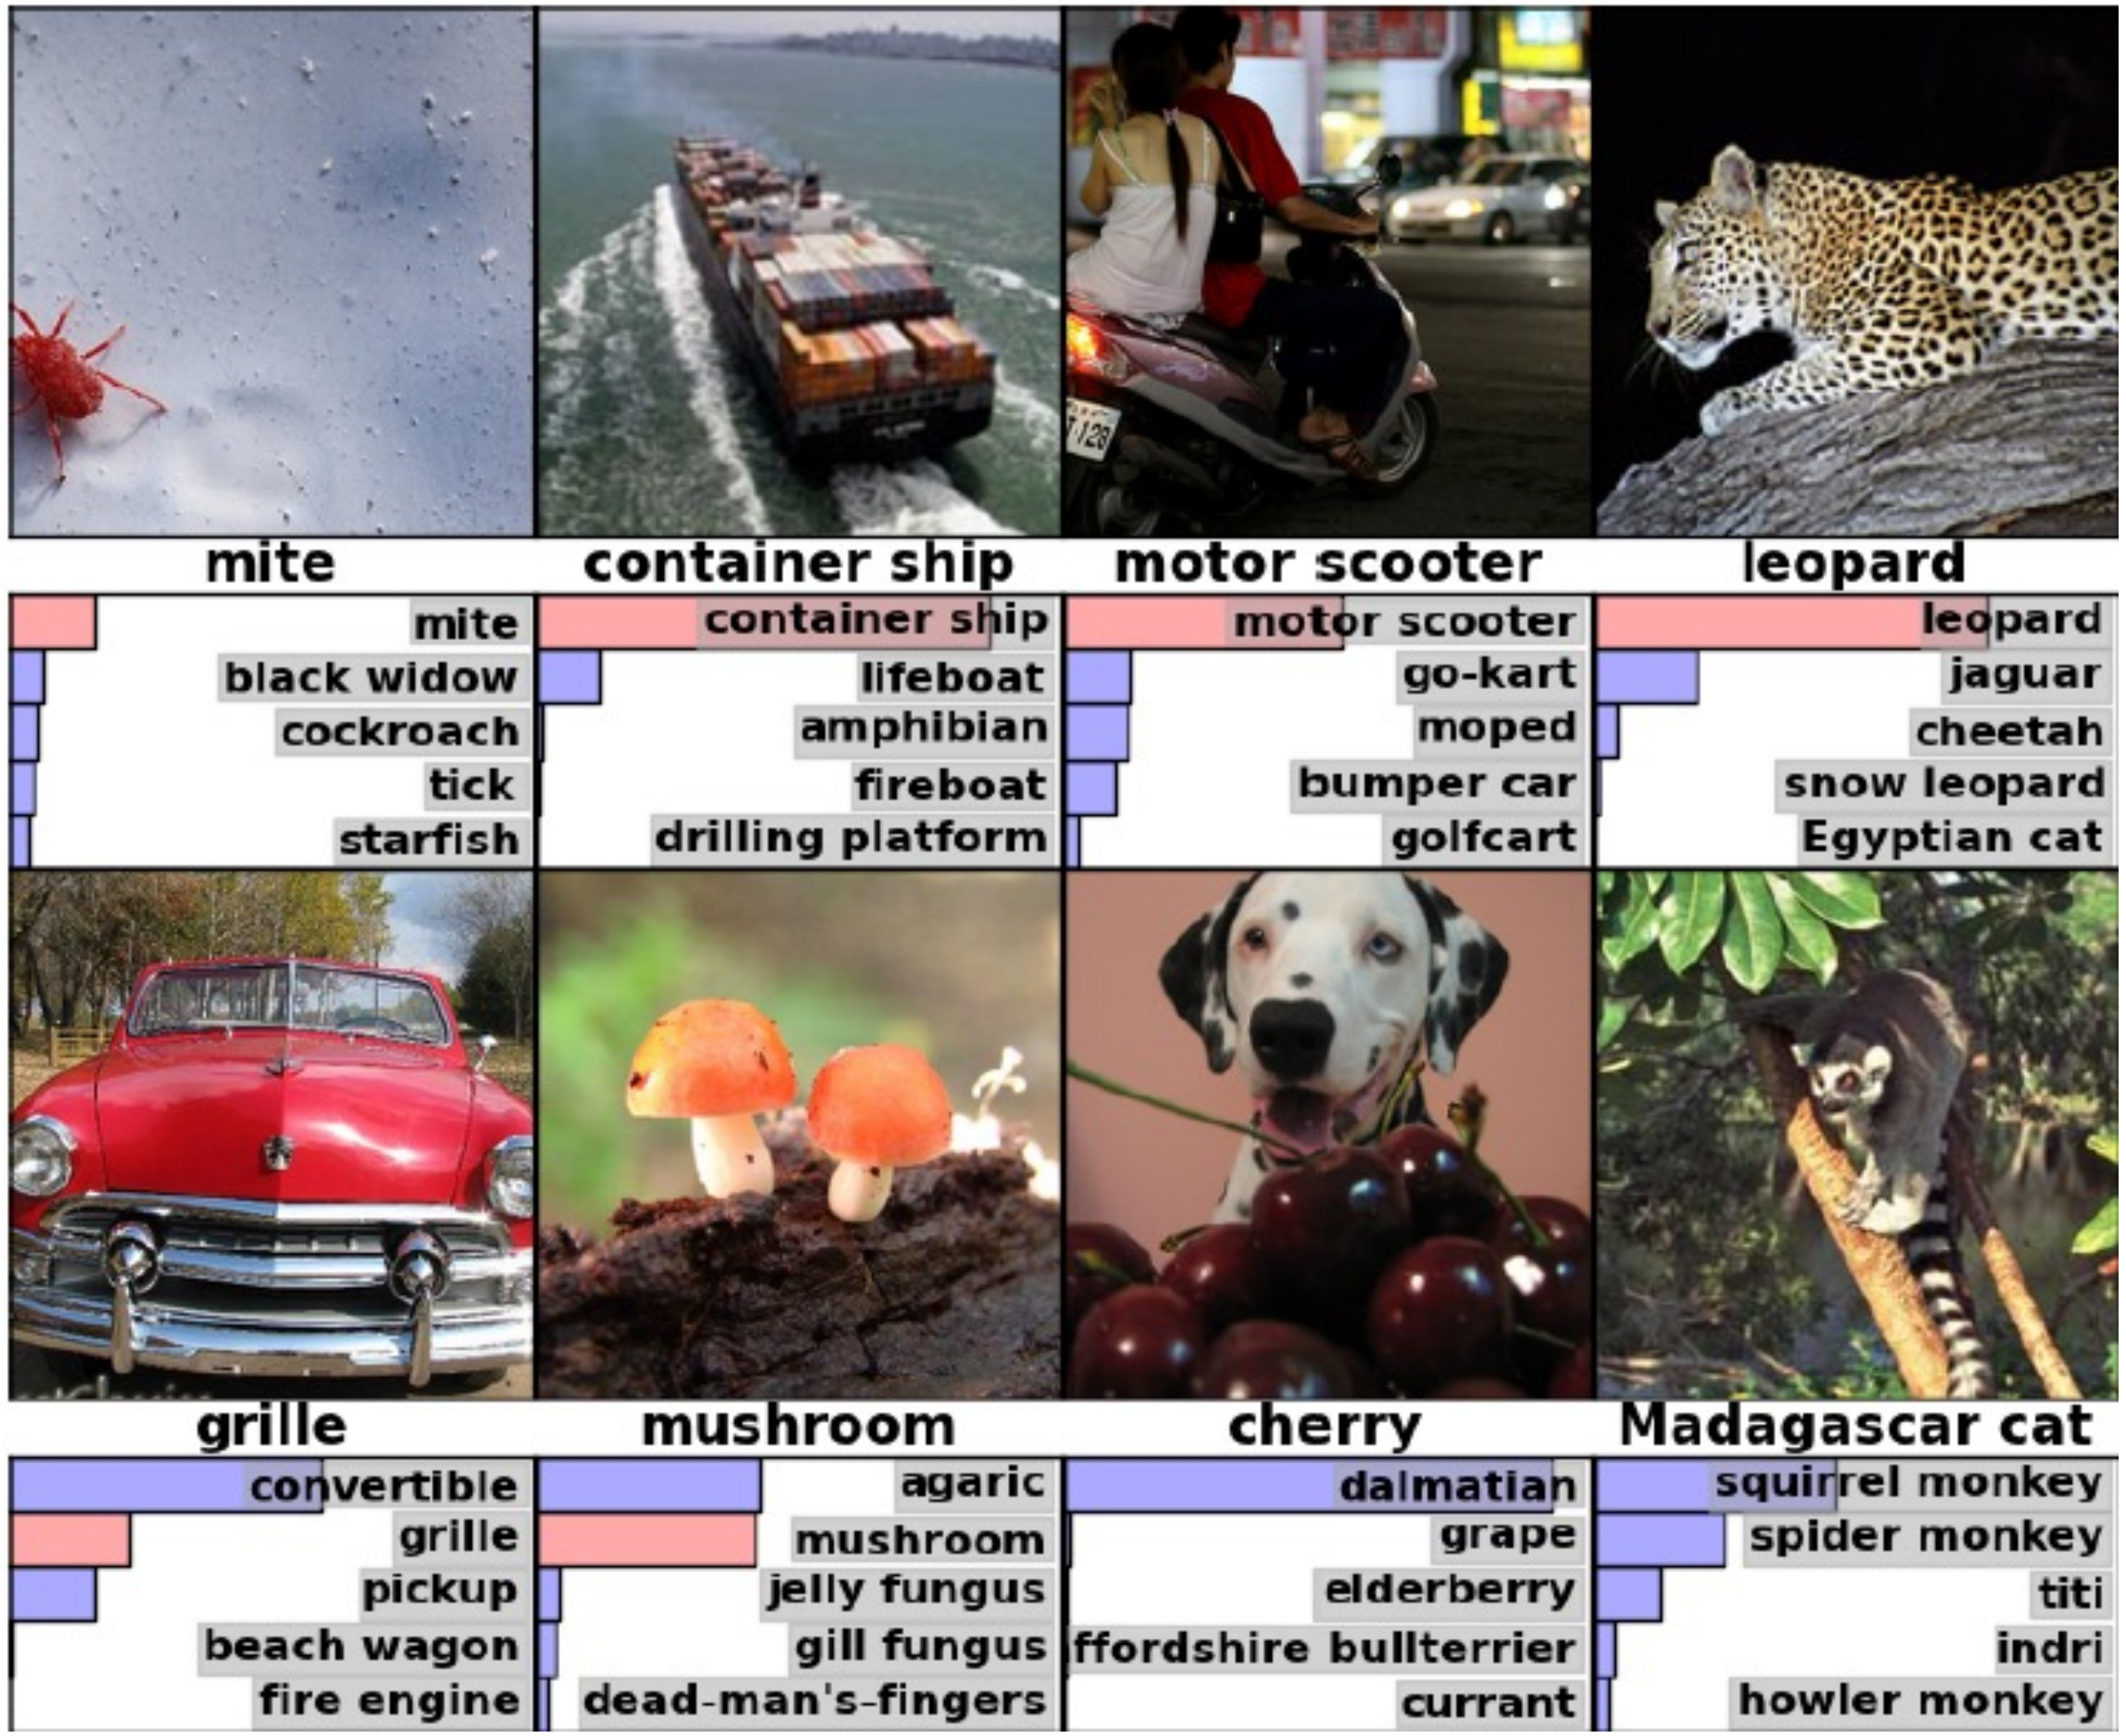

Problems with ImageNet

ImageNet given label:
Irish Wolfhound

Cleanlab guessed: **Irish Water Spaniel**

MTurk consensus: **Irish Water Spaniel**

ID: 00022071

ImageNet given label:
macaque

Cleanlab guessed: **langur**

MTurk consensus: **langur**

ID: 00010077

Consequences of the training set can produce serious problems.

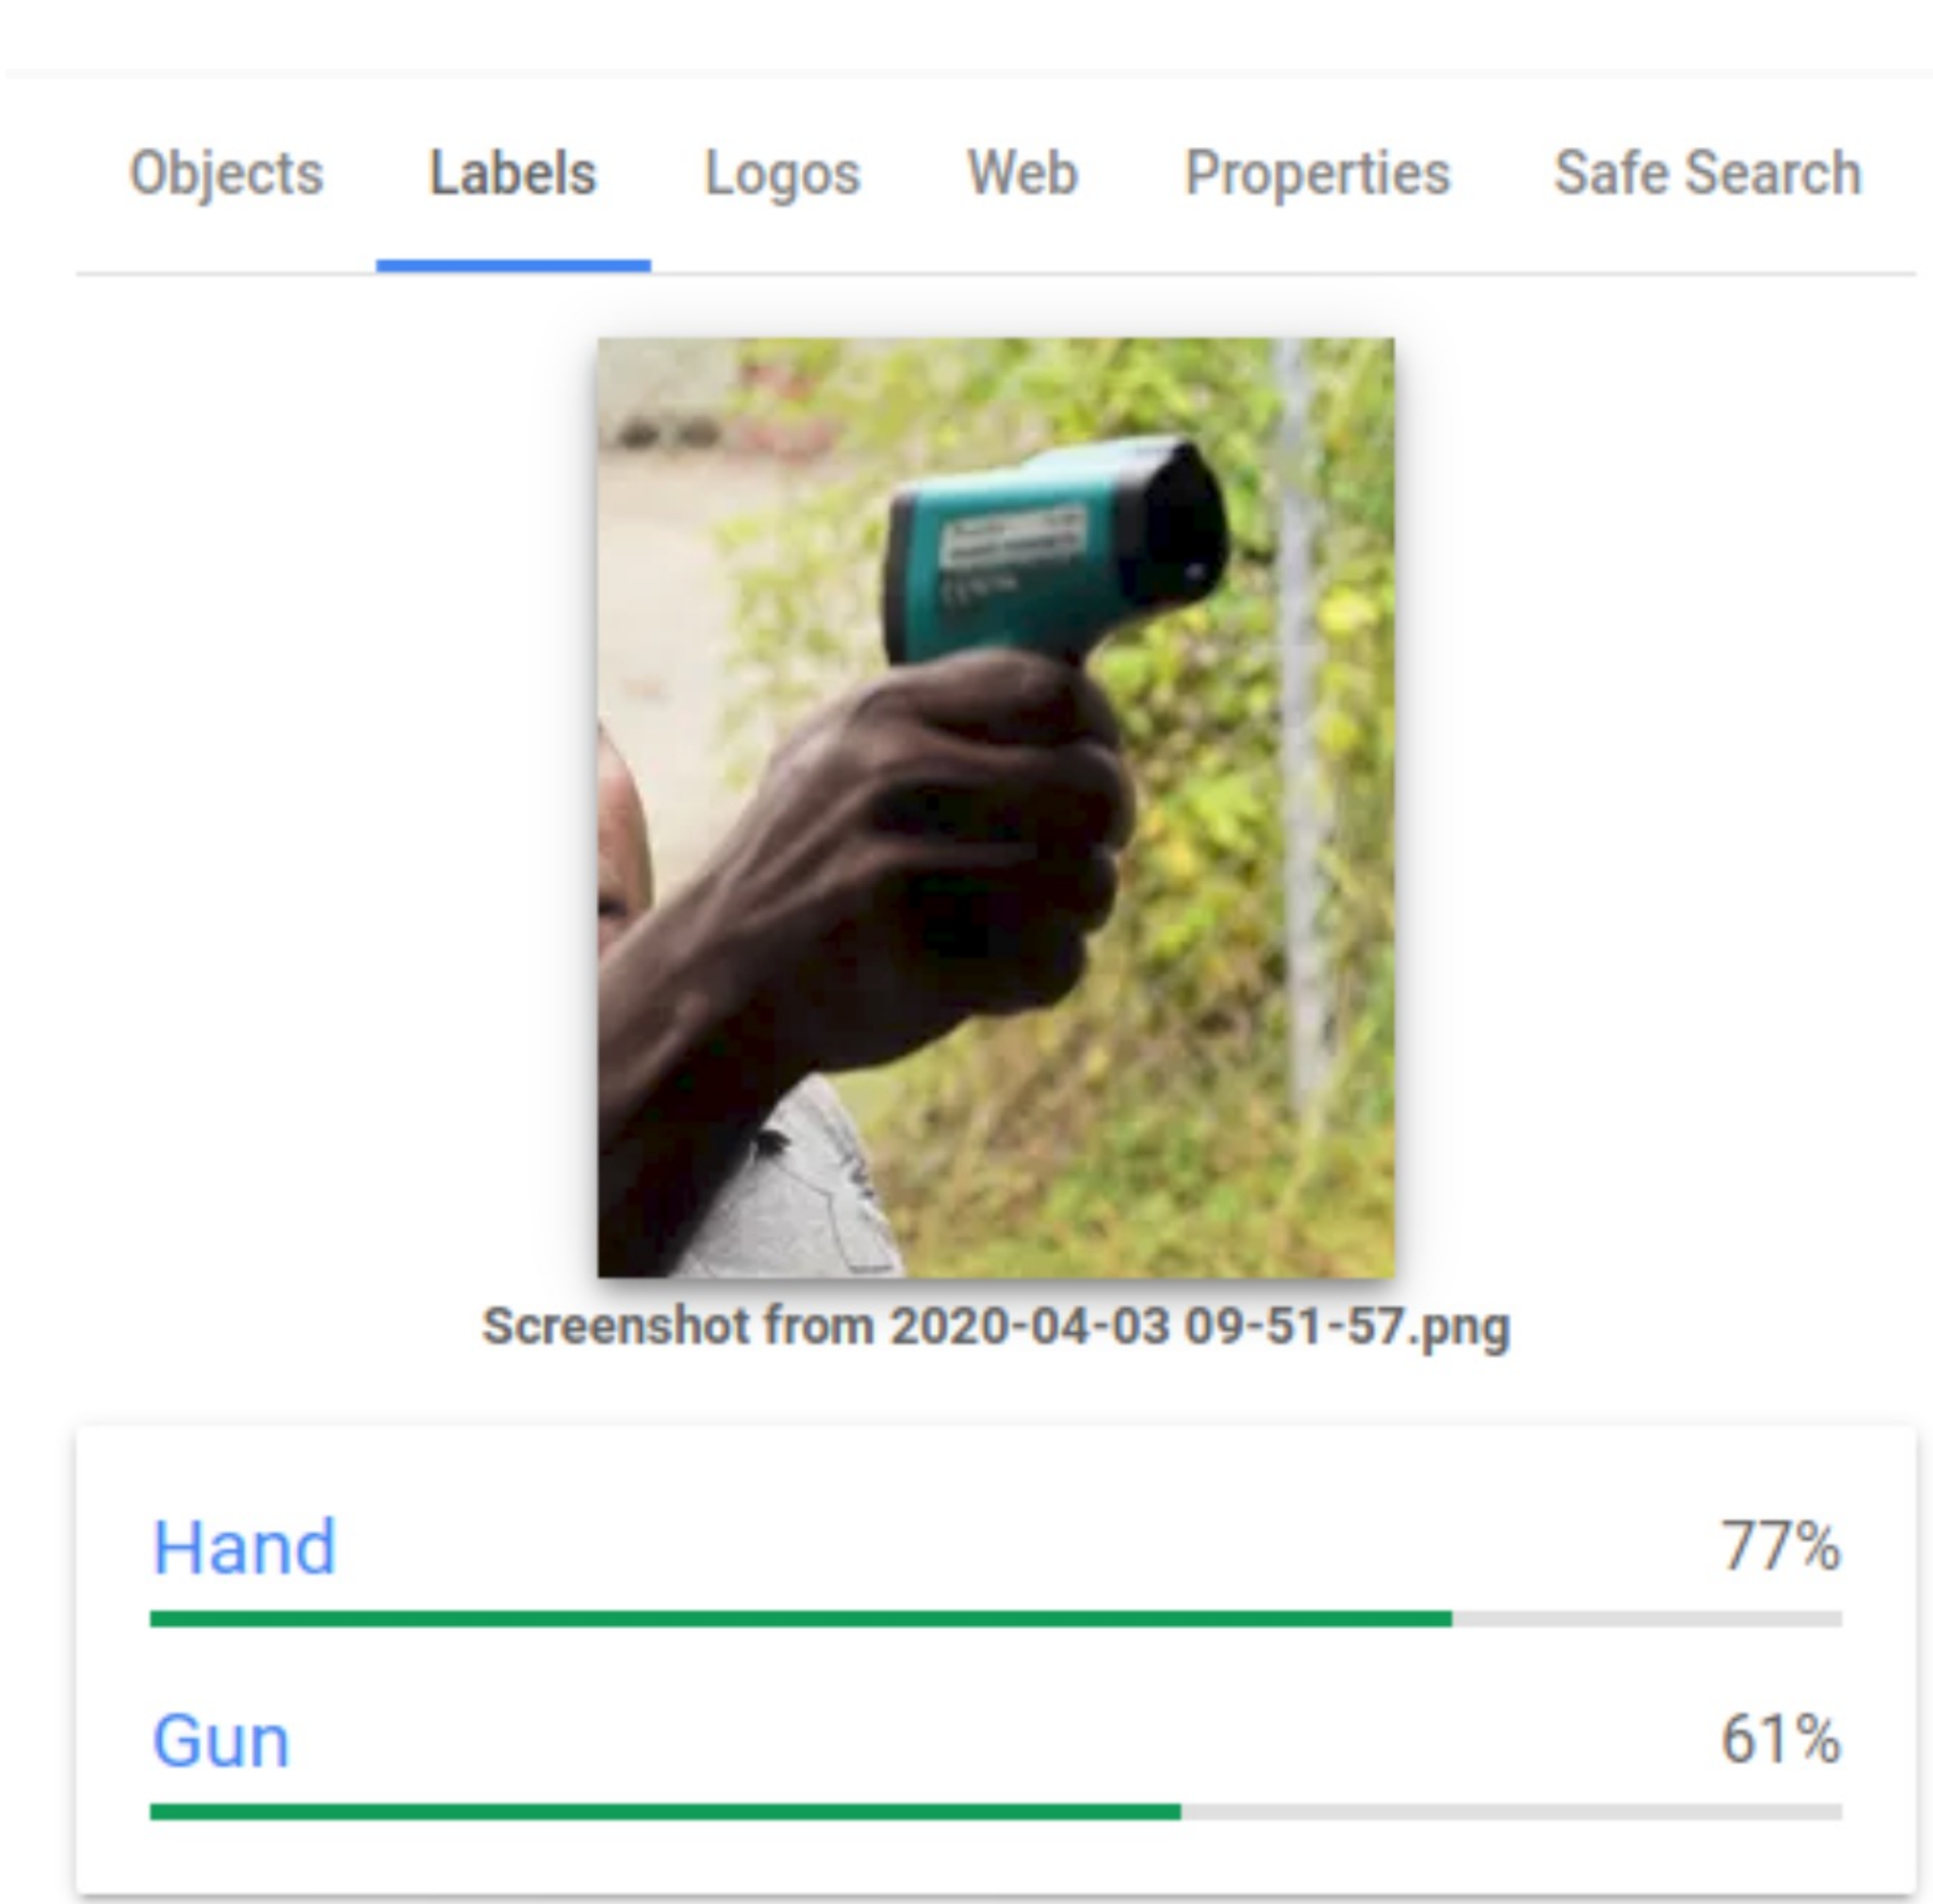

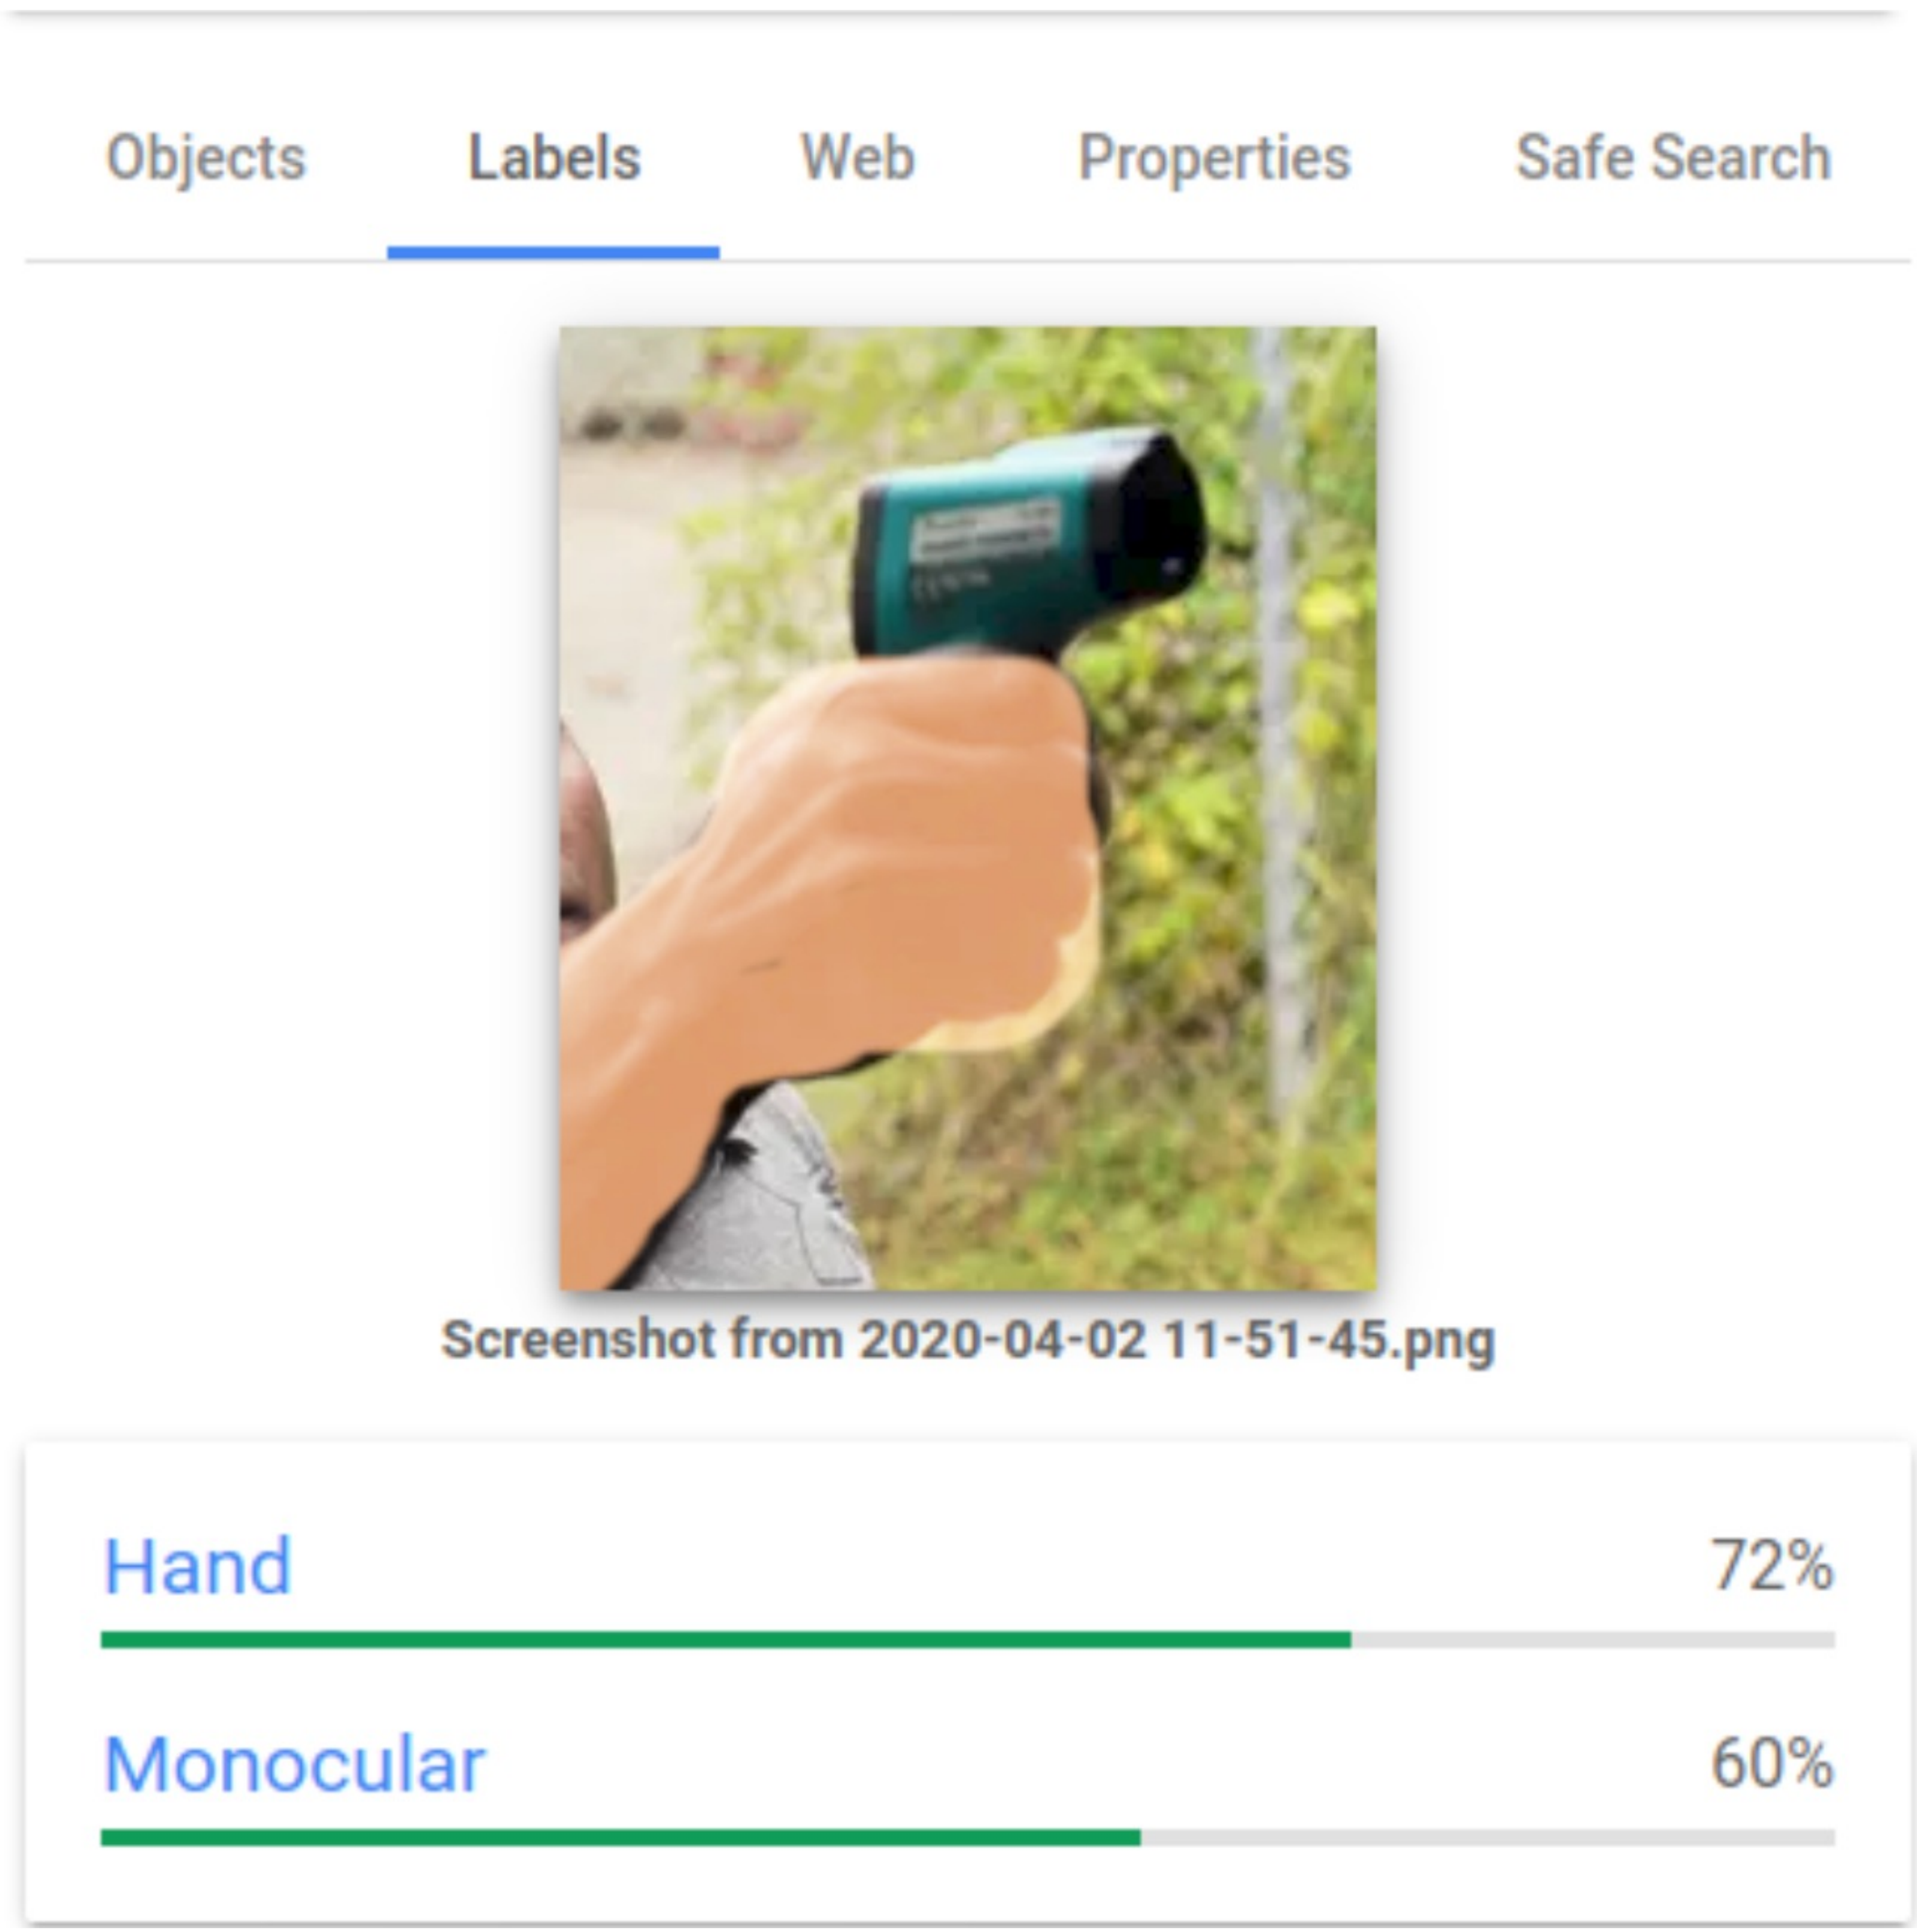

Facial Recognition

Facial recognition uses CNNs, but produce “feature vectors”

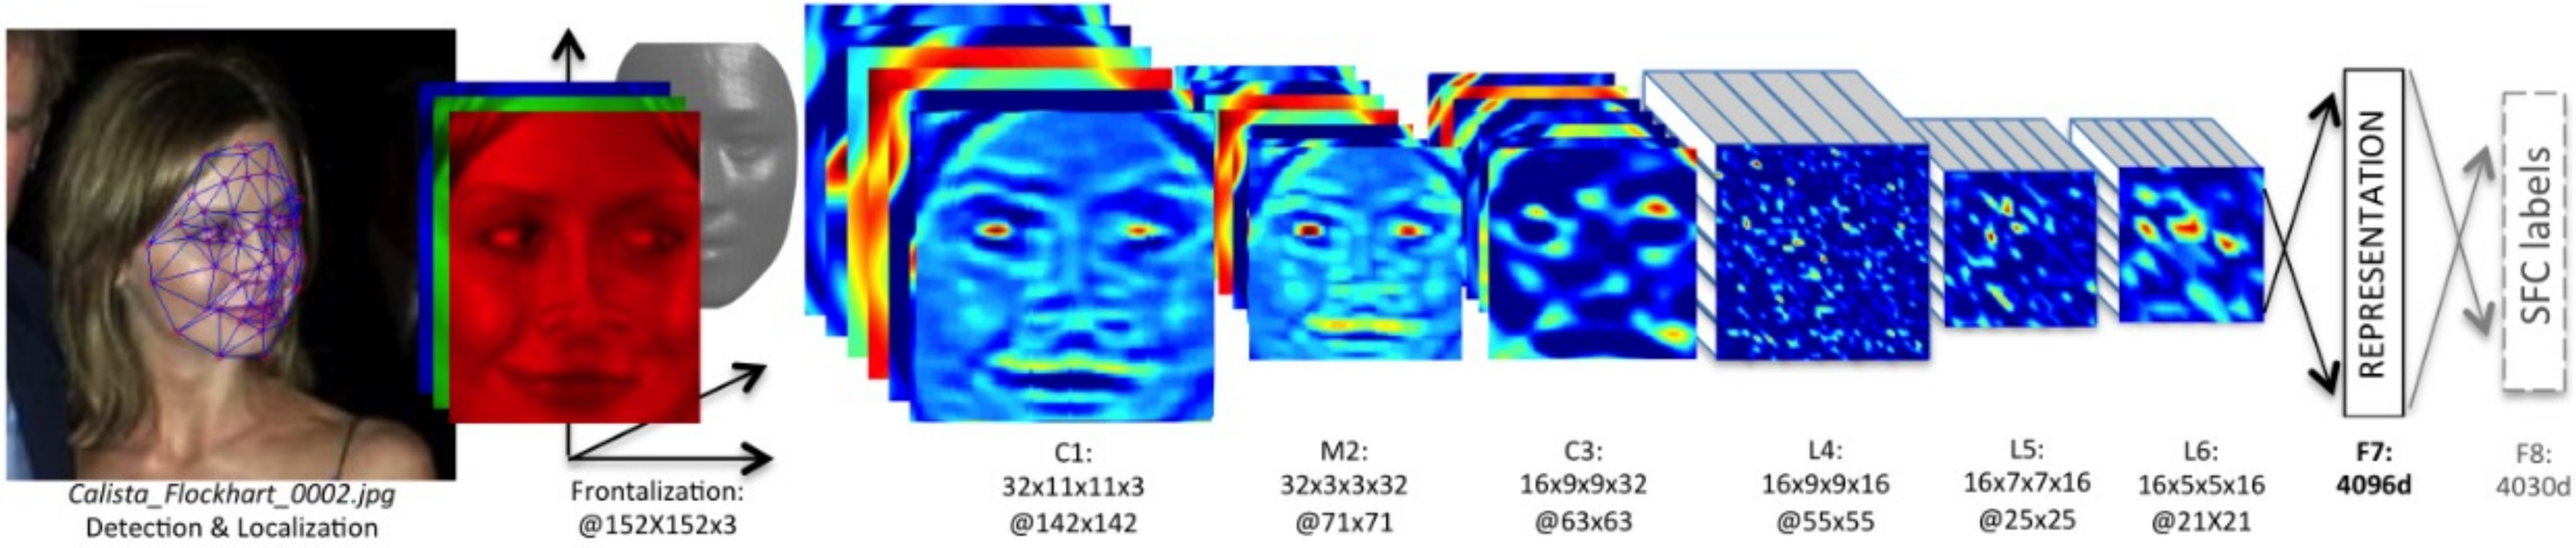

DeepFace used a complex set of 2D and 3D transformations on faces.

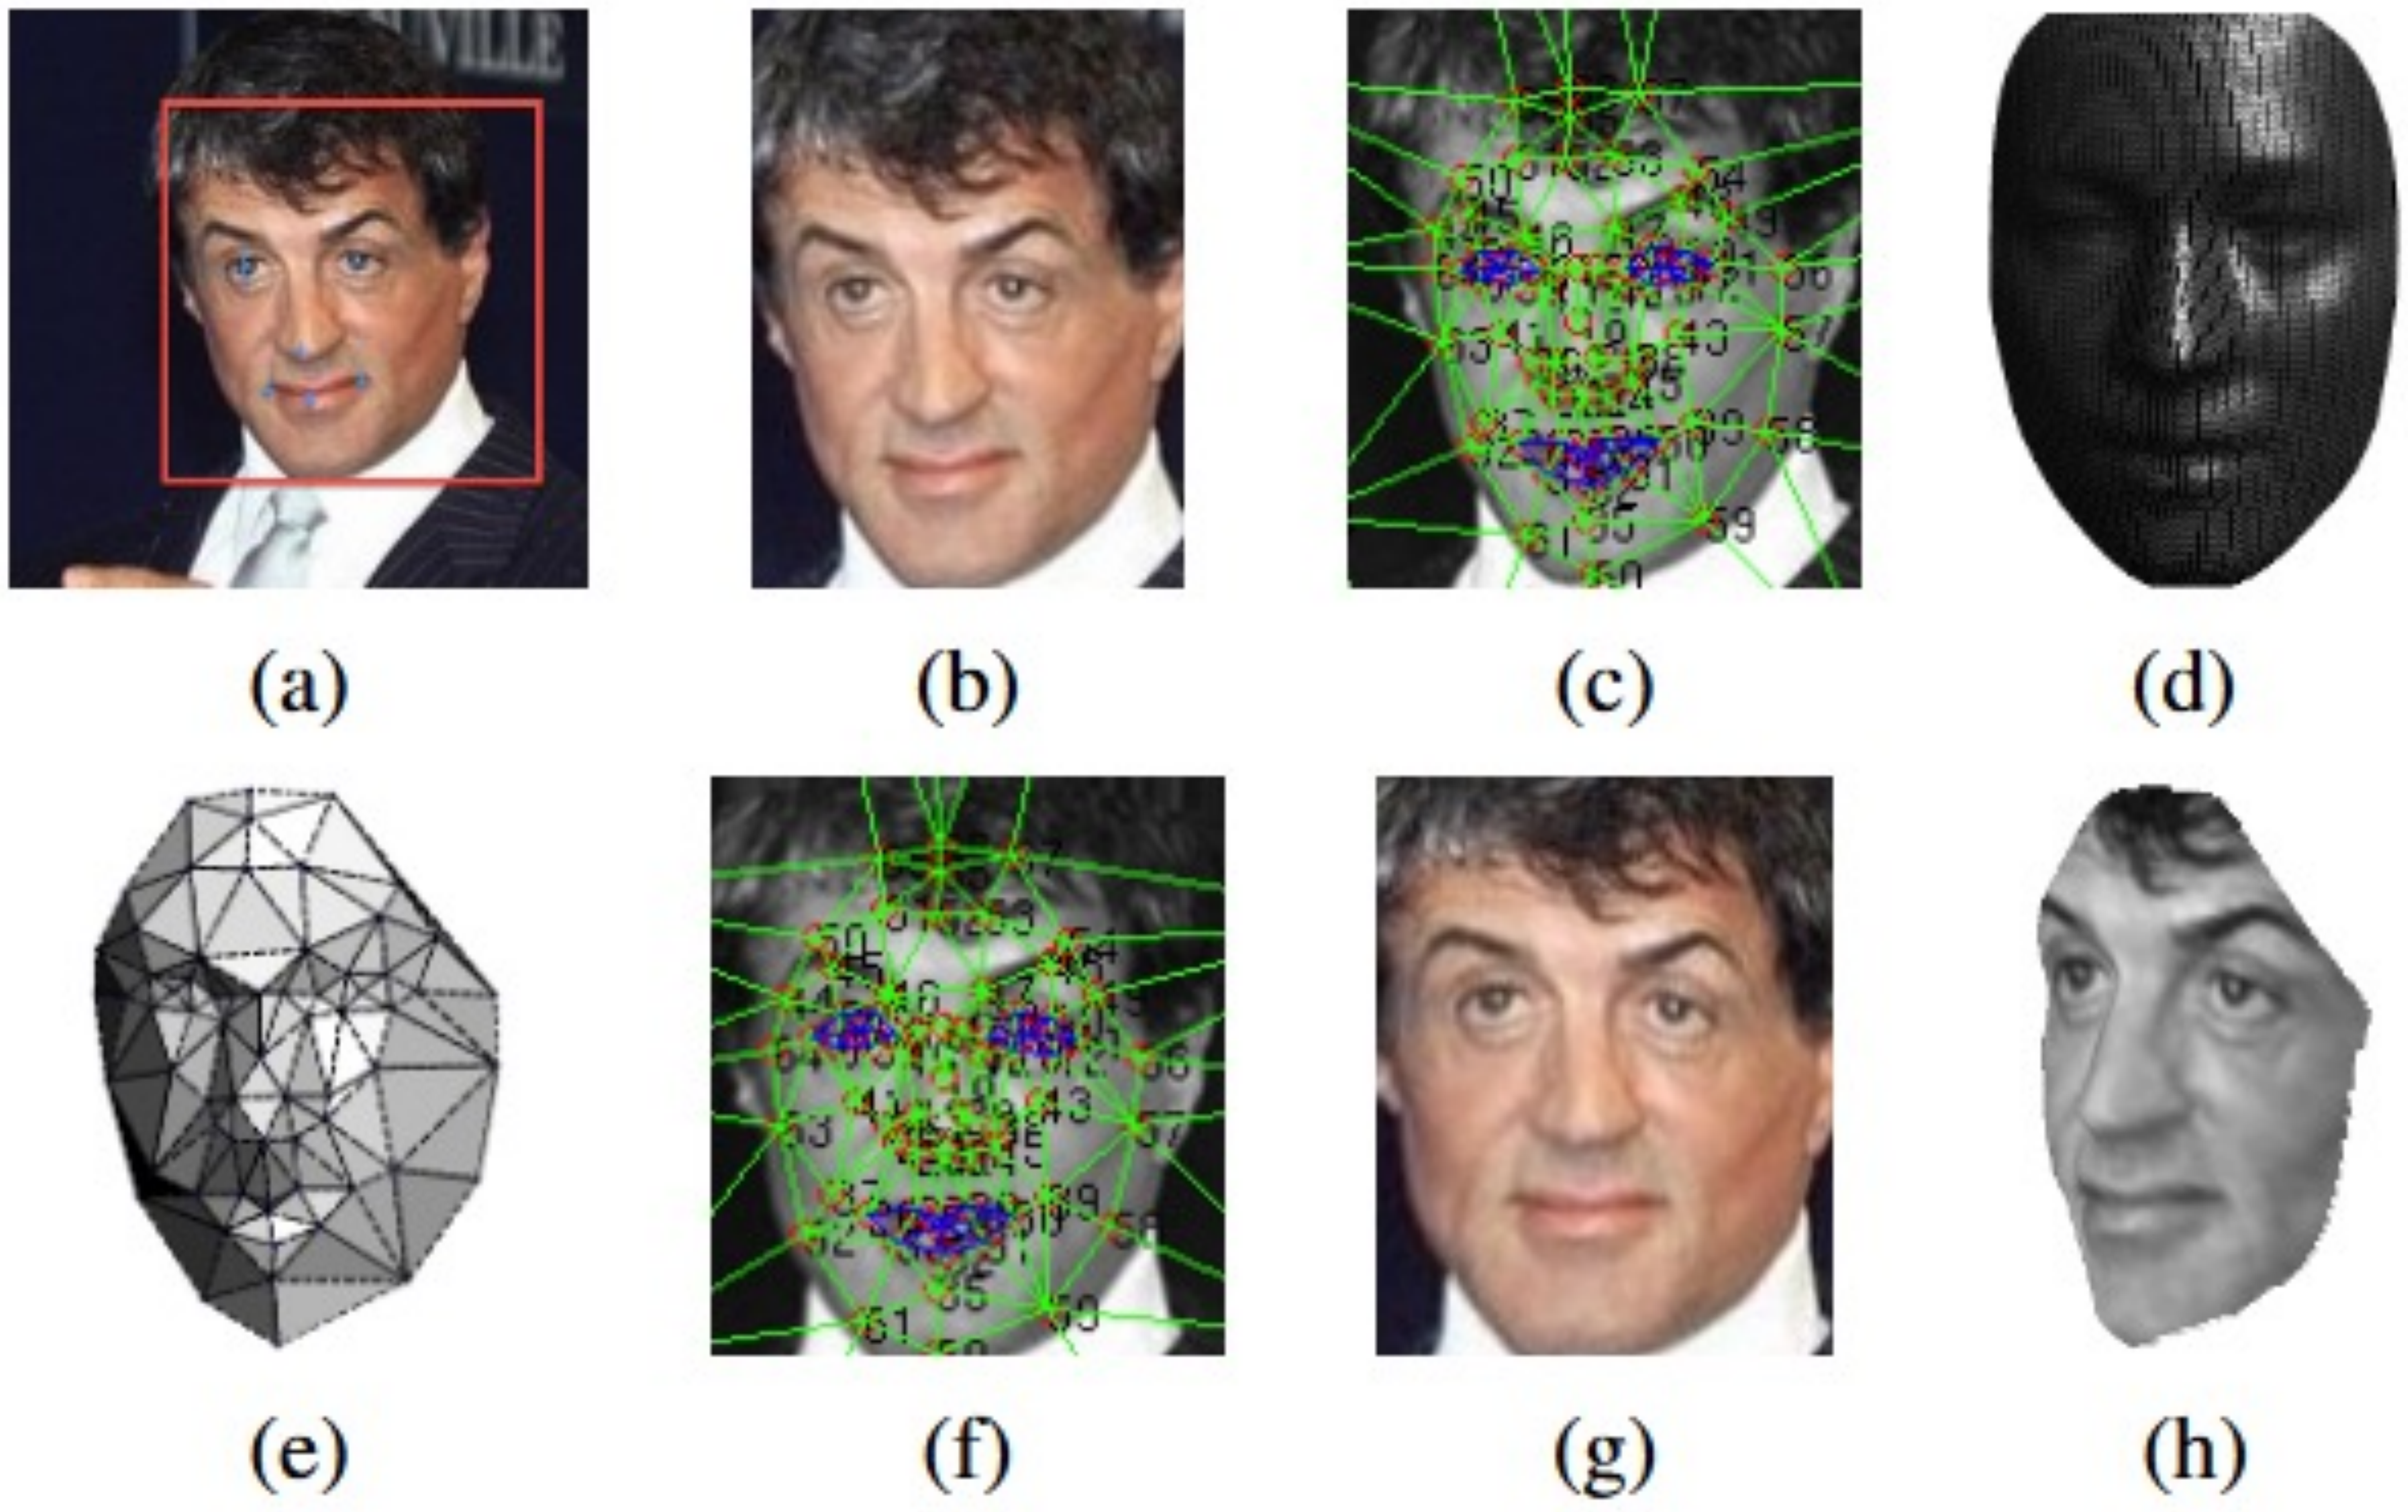

FRT can perform authentication (1:1) or identification (1:N).

- Face Authentication/Verification (1:1 matching)

- Face Identification/Recognition (1:N matching)

1:N

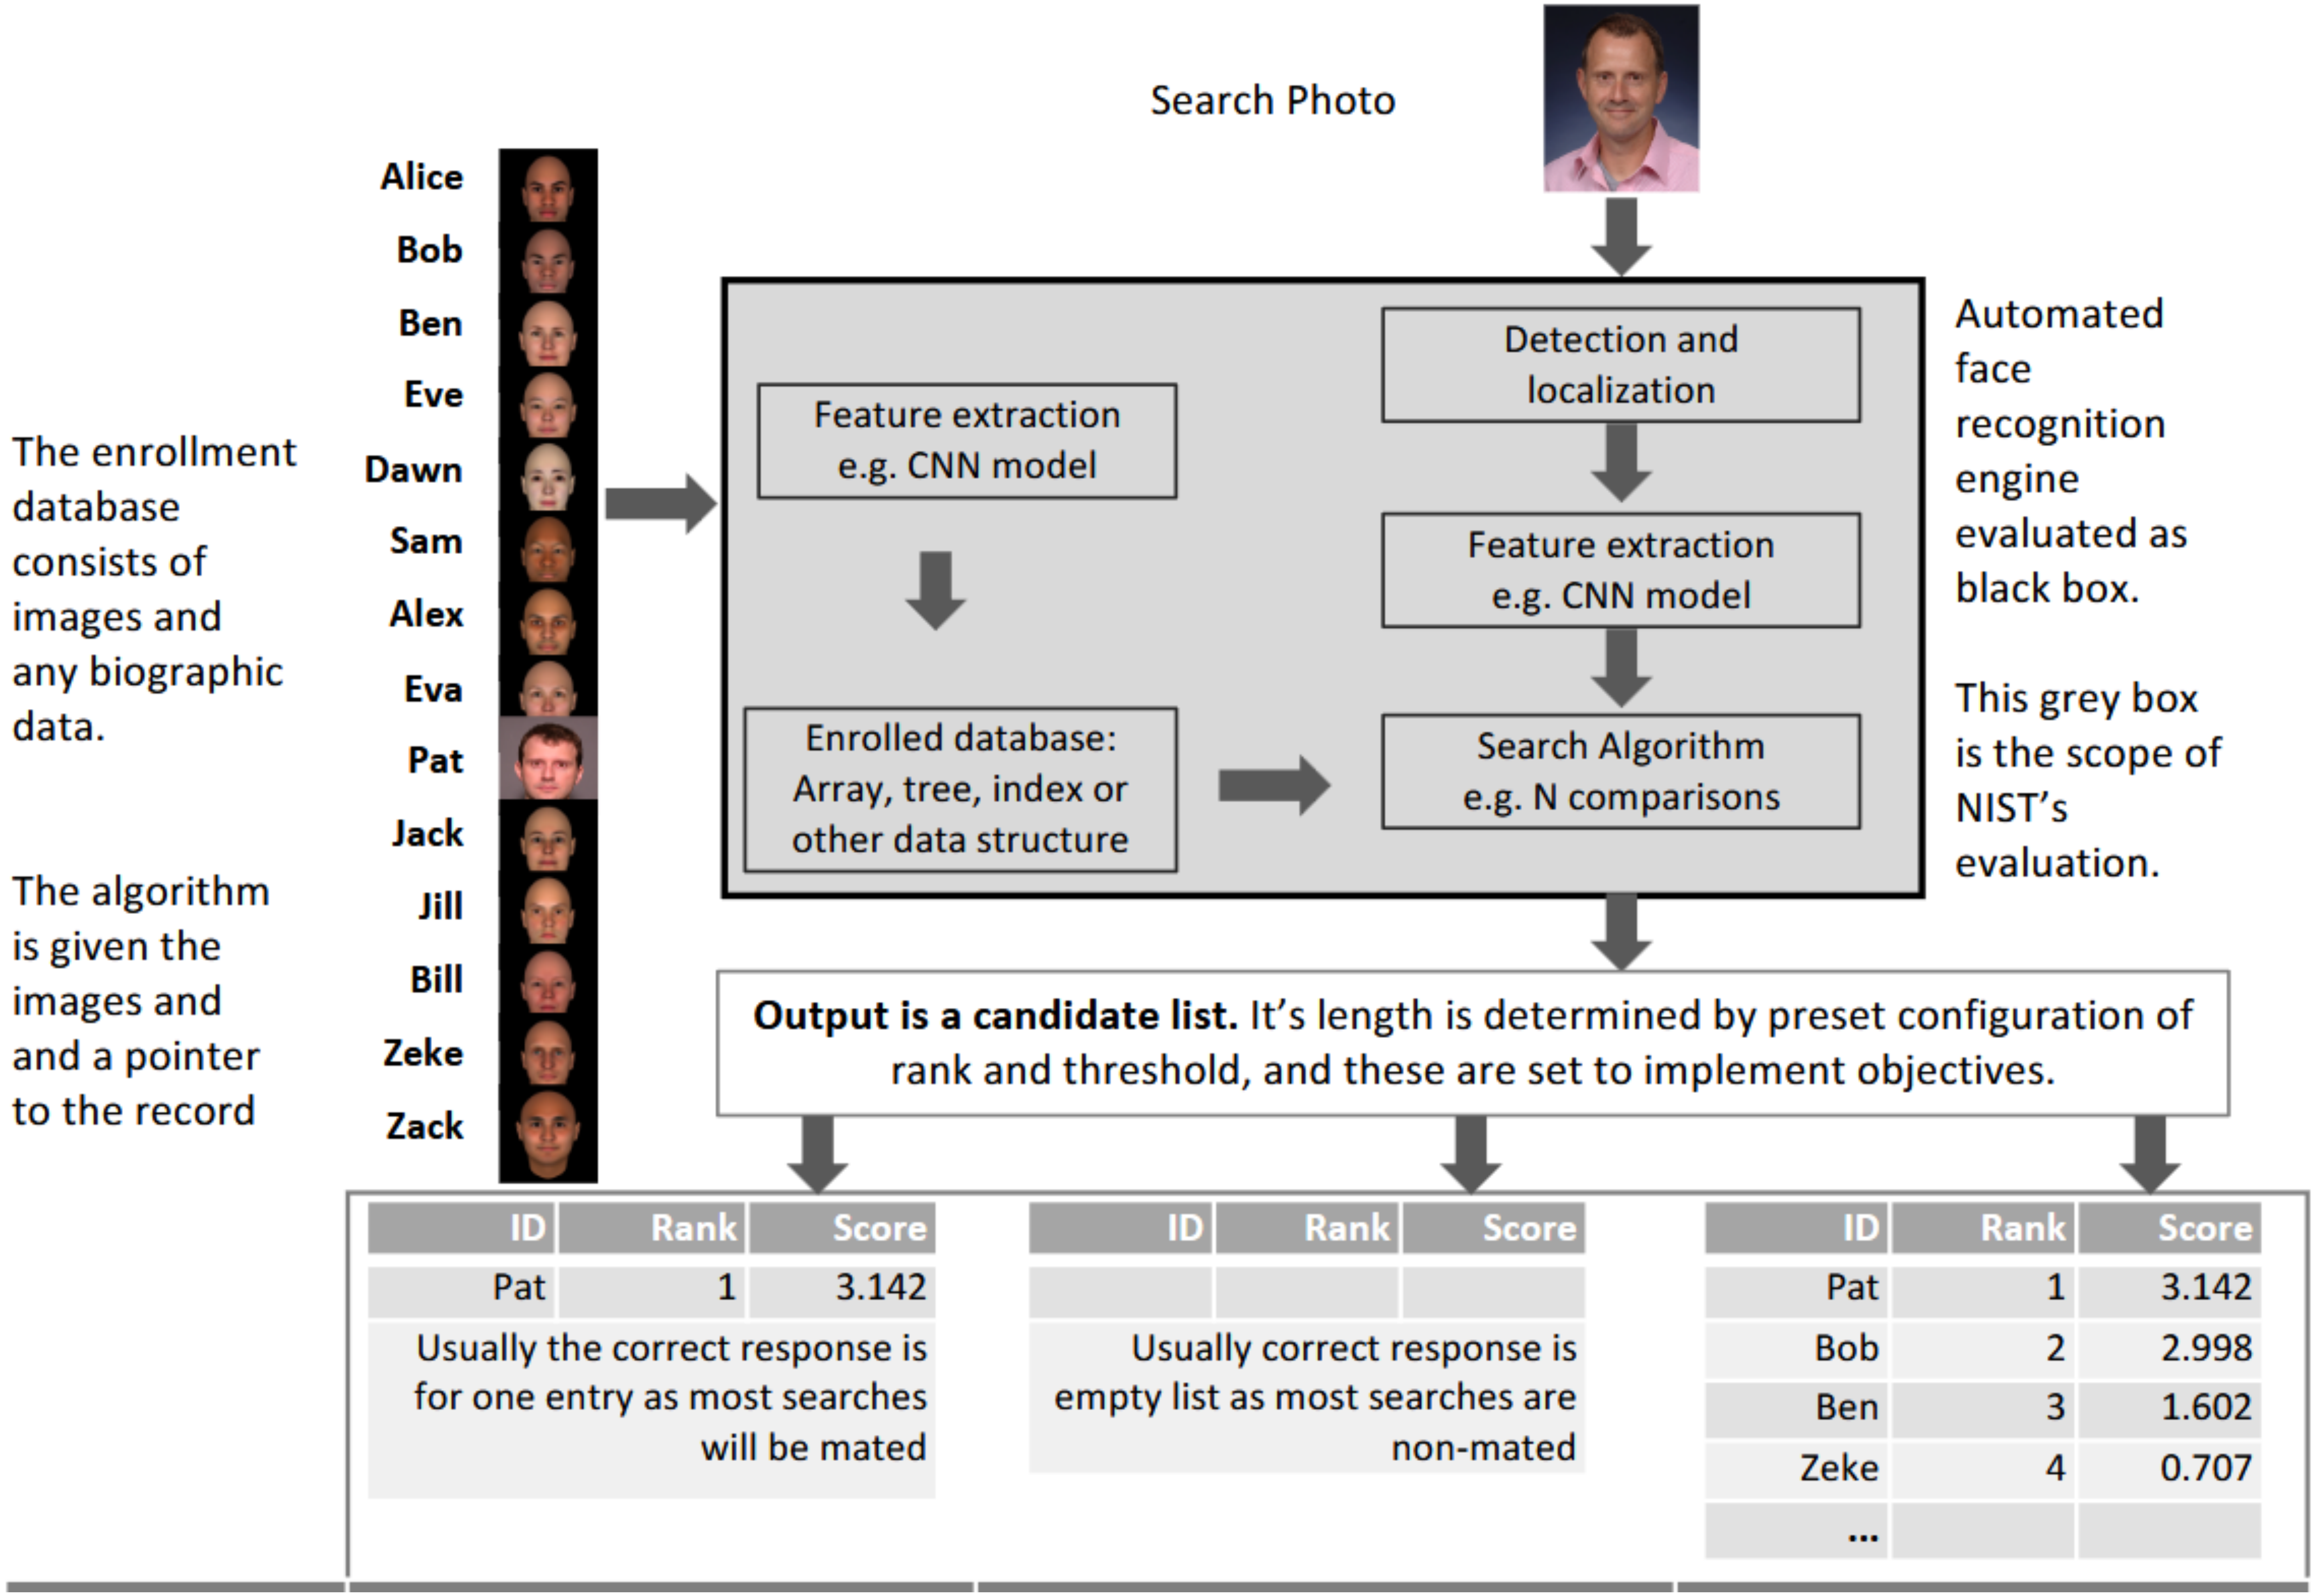

Search Photo
Alice
Bob
Ben
Eve
Dawn
Sam
Alex
Eva
Pat
Jack
Jill
Bill
Zeke
Zack
The enrollment database consists of images and any biographic data.
The algorithm is given the images and and a pointer to the record
Feature extraction e.g. CNN model
Detection and localization
Feature extraction e.g. CNN model
Enrolled database: Array, tree, index or other data structure
Search Algorithm e.g. N comparisons
Automated face recognition engine evaluated as black box.
This grey box is the scope of NIST's evaluation.
Output is a candidate list. It's length is determined by preset configuration of rank and threshold, and these are set to implement objectives.
ID Rank Score
Pat 1 3.142
Usually the correct response is for one entry as most searches will be mated
ID Rank Score
Usually correct response is empty list as most searches are non-mated
ID Rank Score
Pat 1 3.142
Bob 2 2.998
Ben 3 1.602
Zeke 4 0.707
...

There is a trade-off between false positives and false negatives.

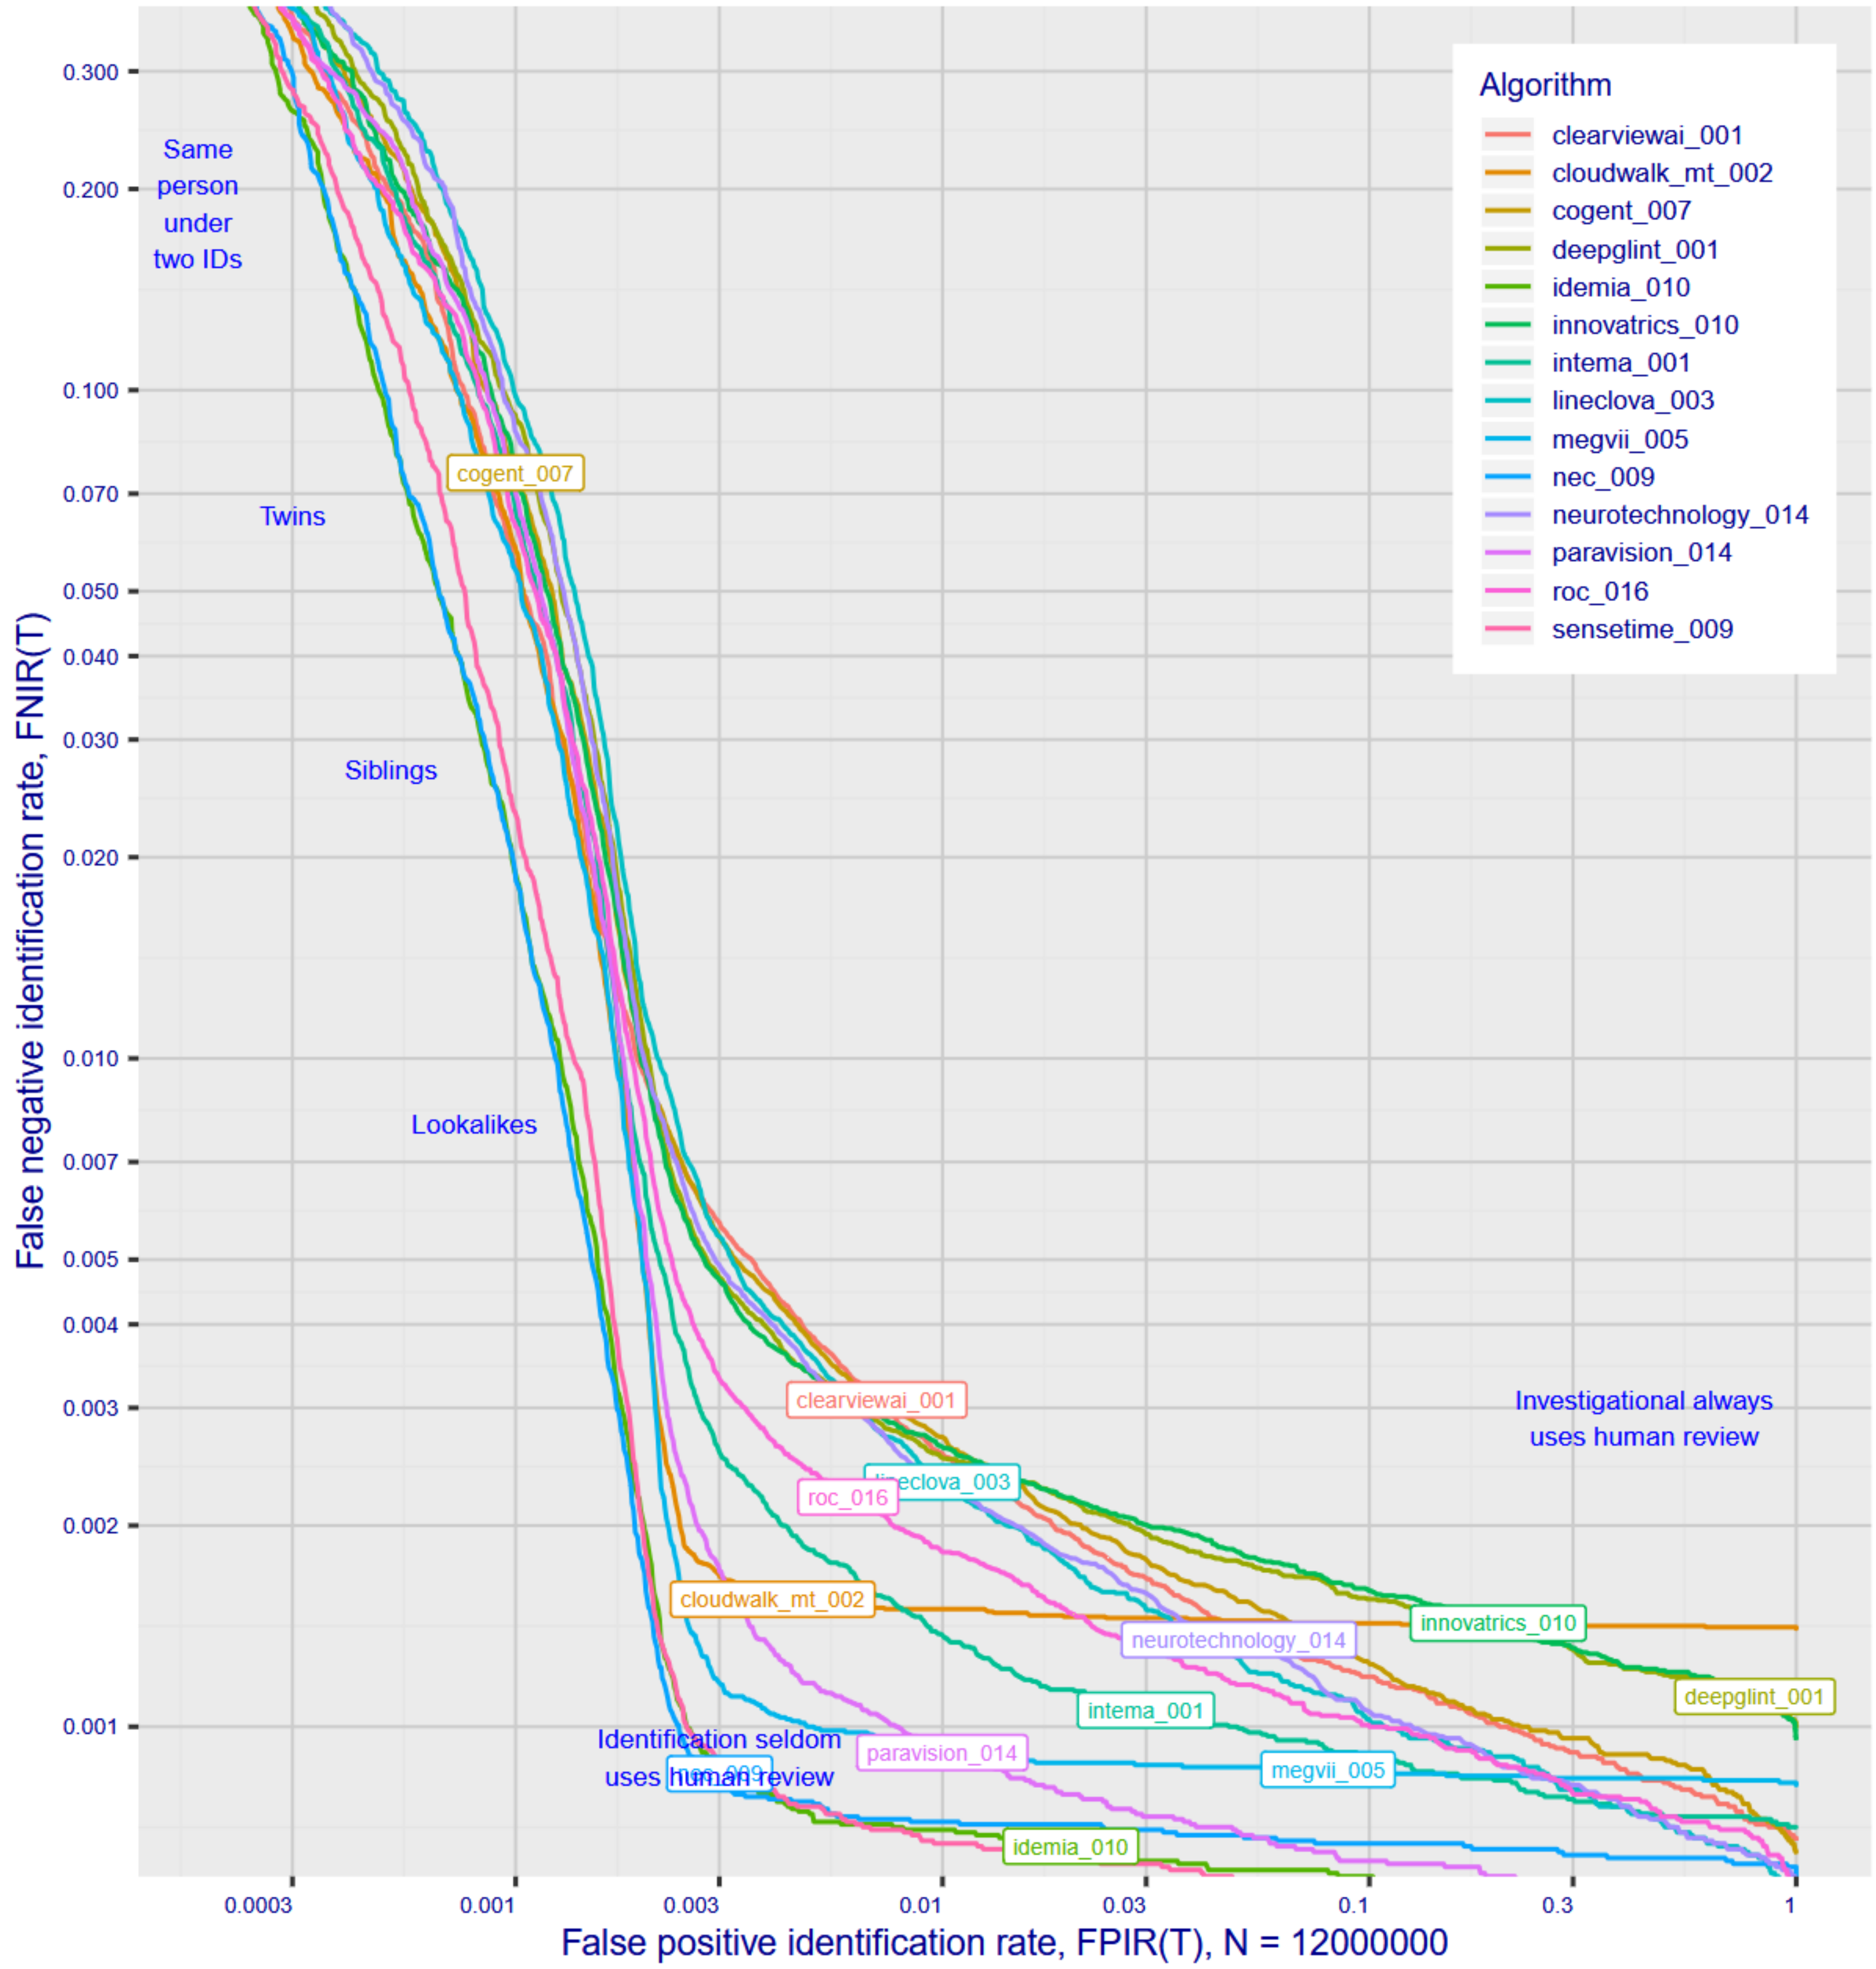

Figure 1: Identification miss rates across the false positive range. N = 12 million individuals are enrolled with one recent image.

Open Set v Closed Set Galleries

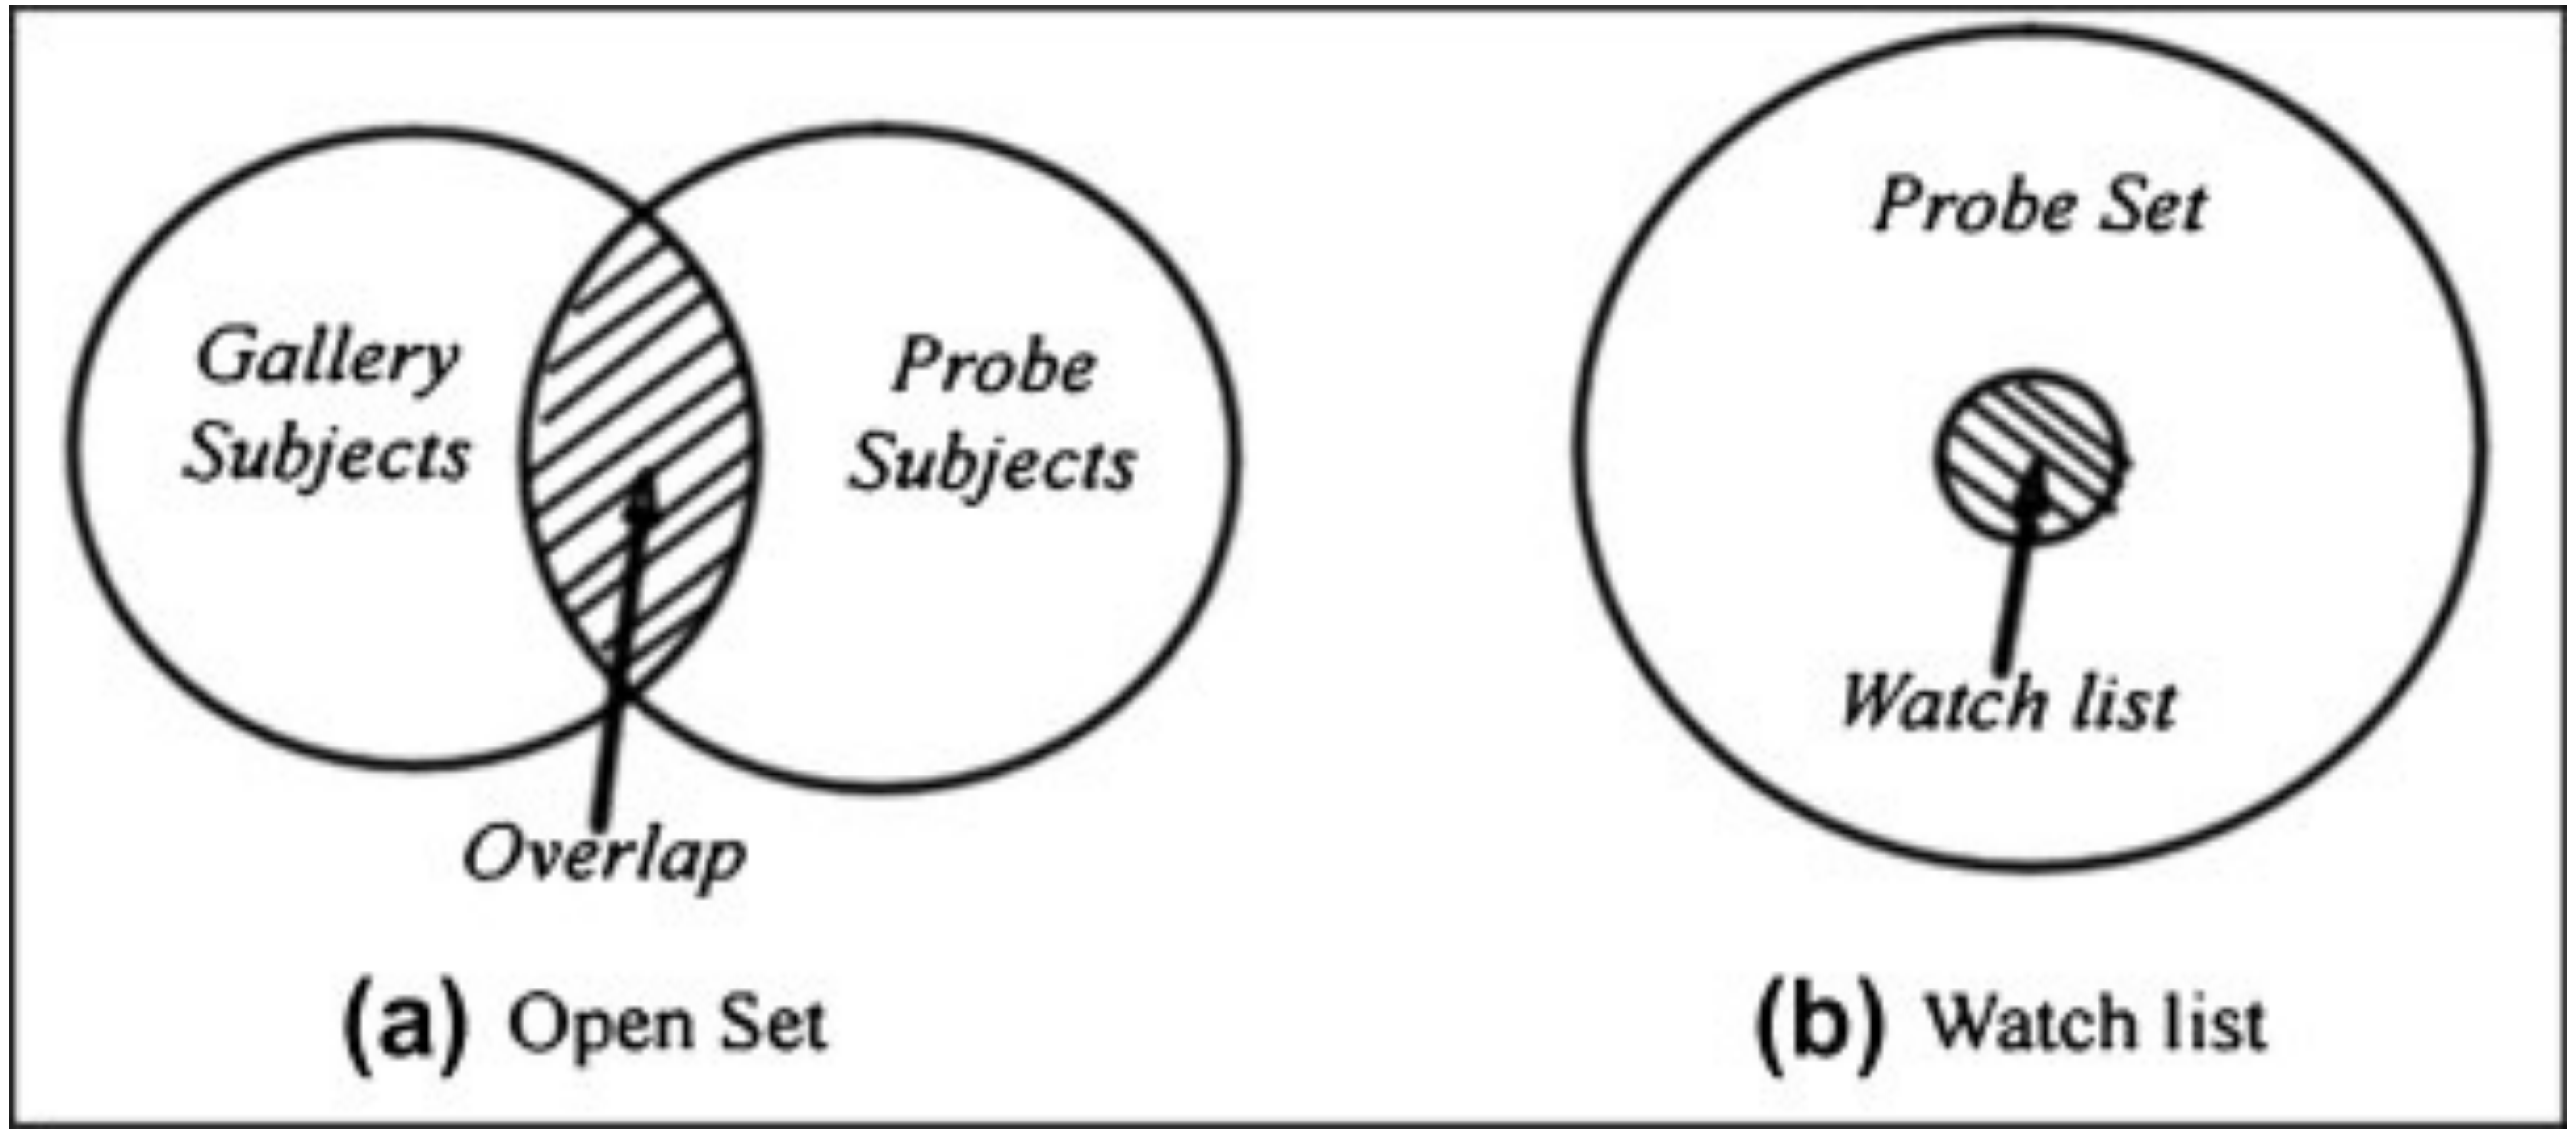

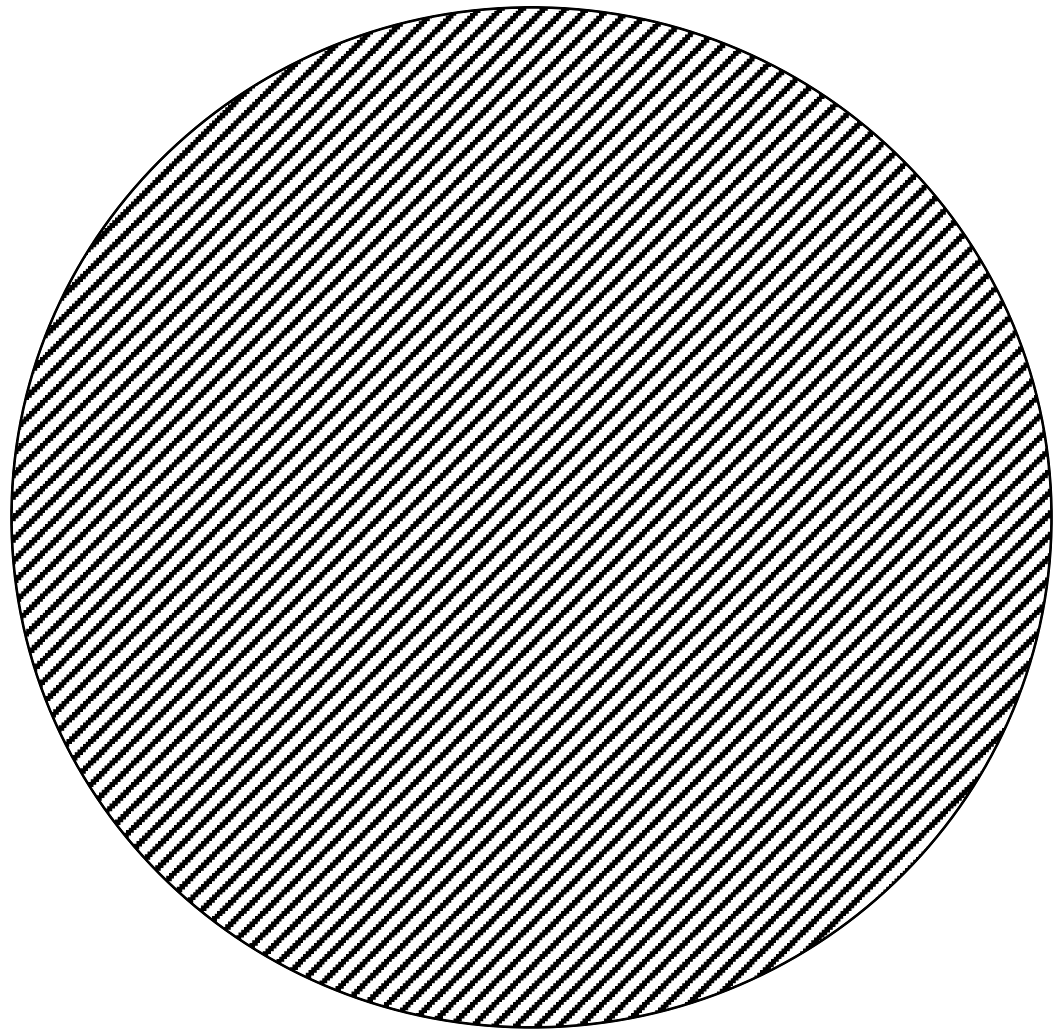

Closed Set

FRT false matches increase with demographic similarity.

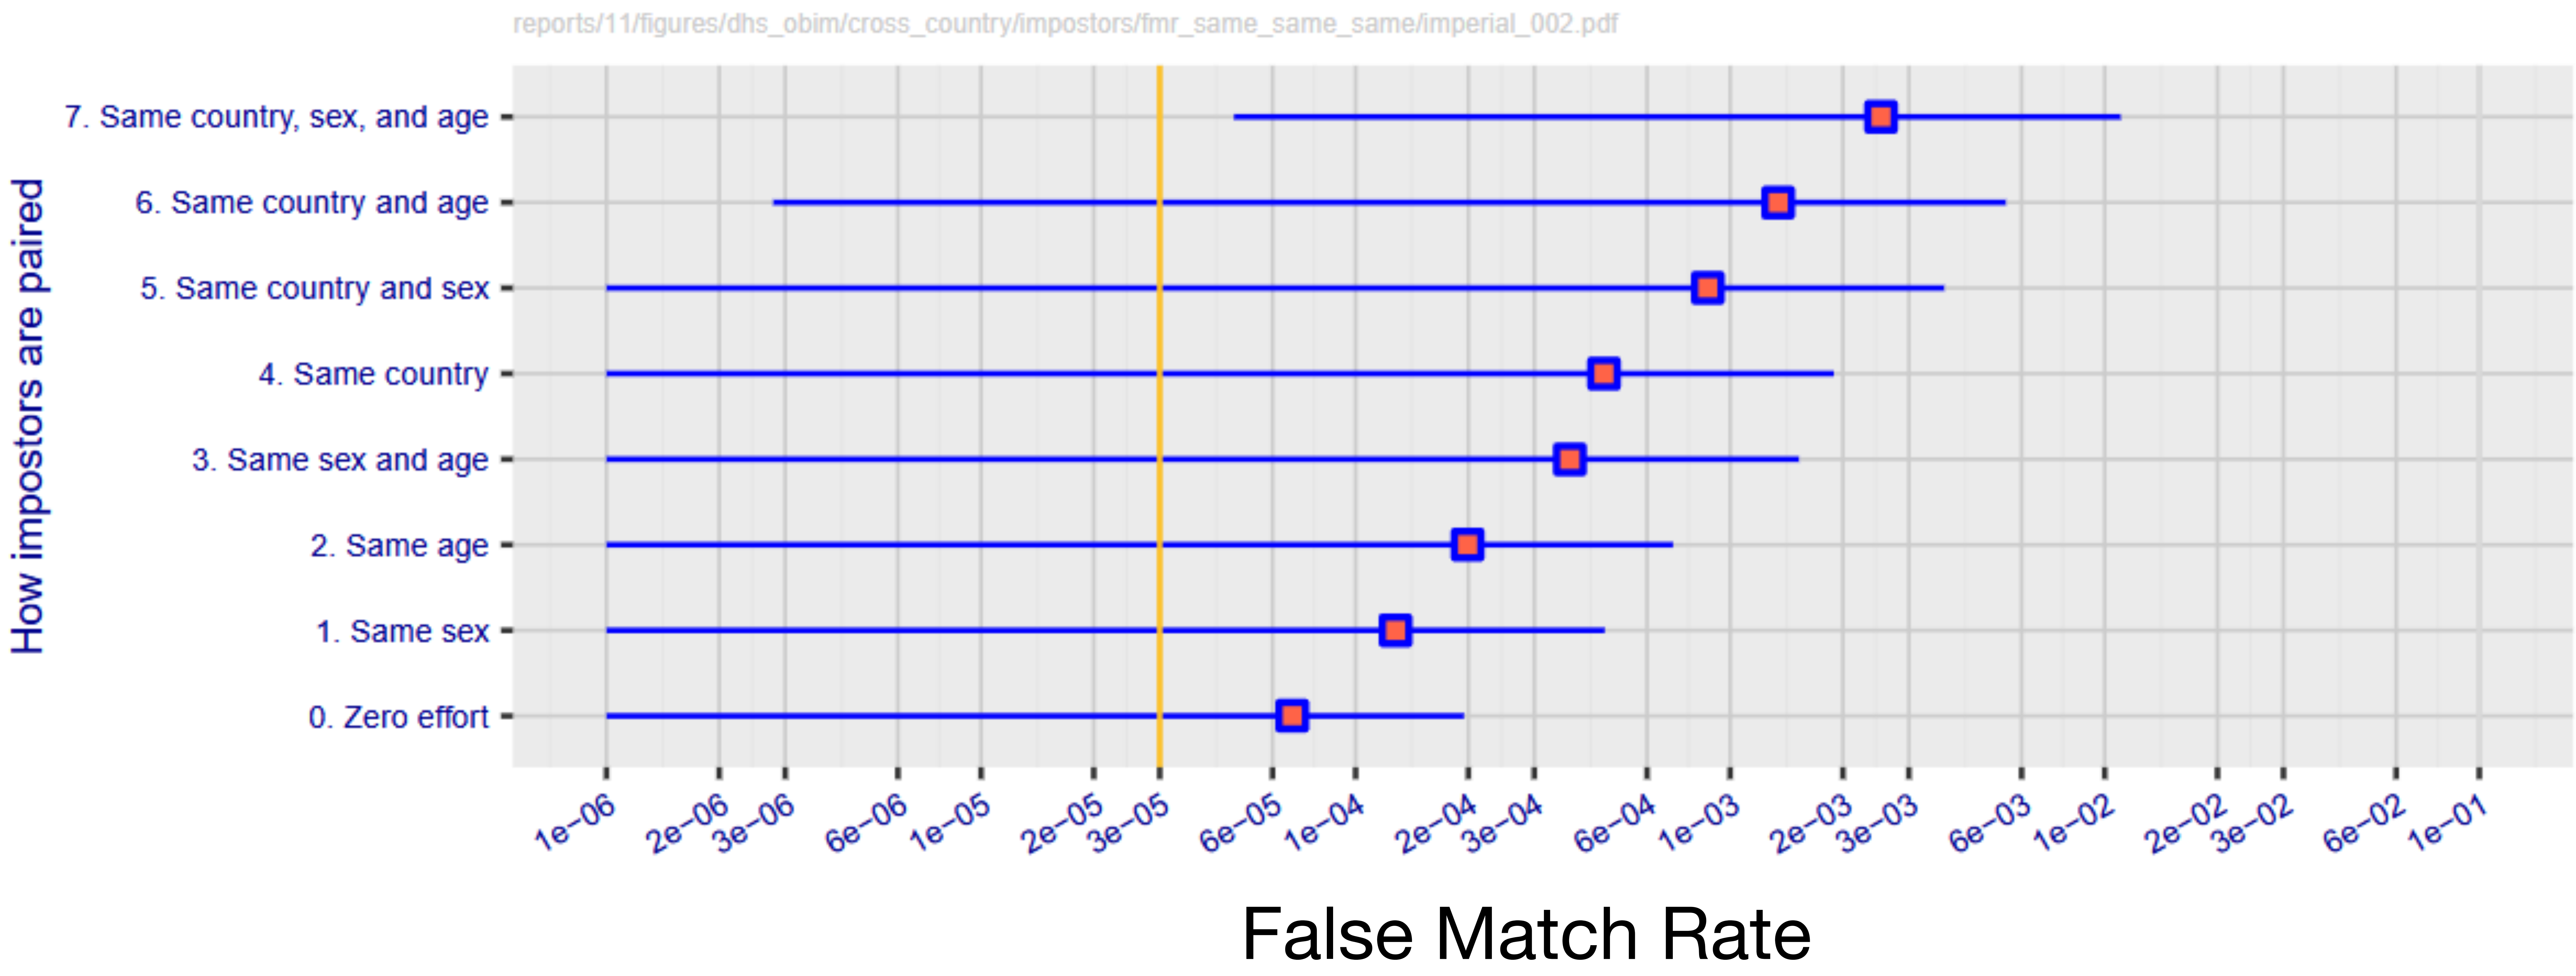

FRT performance can vary by sex and race.

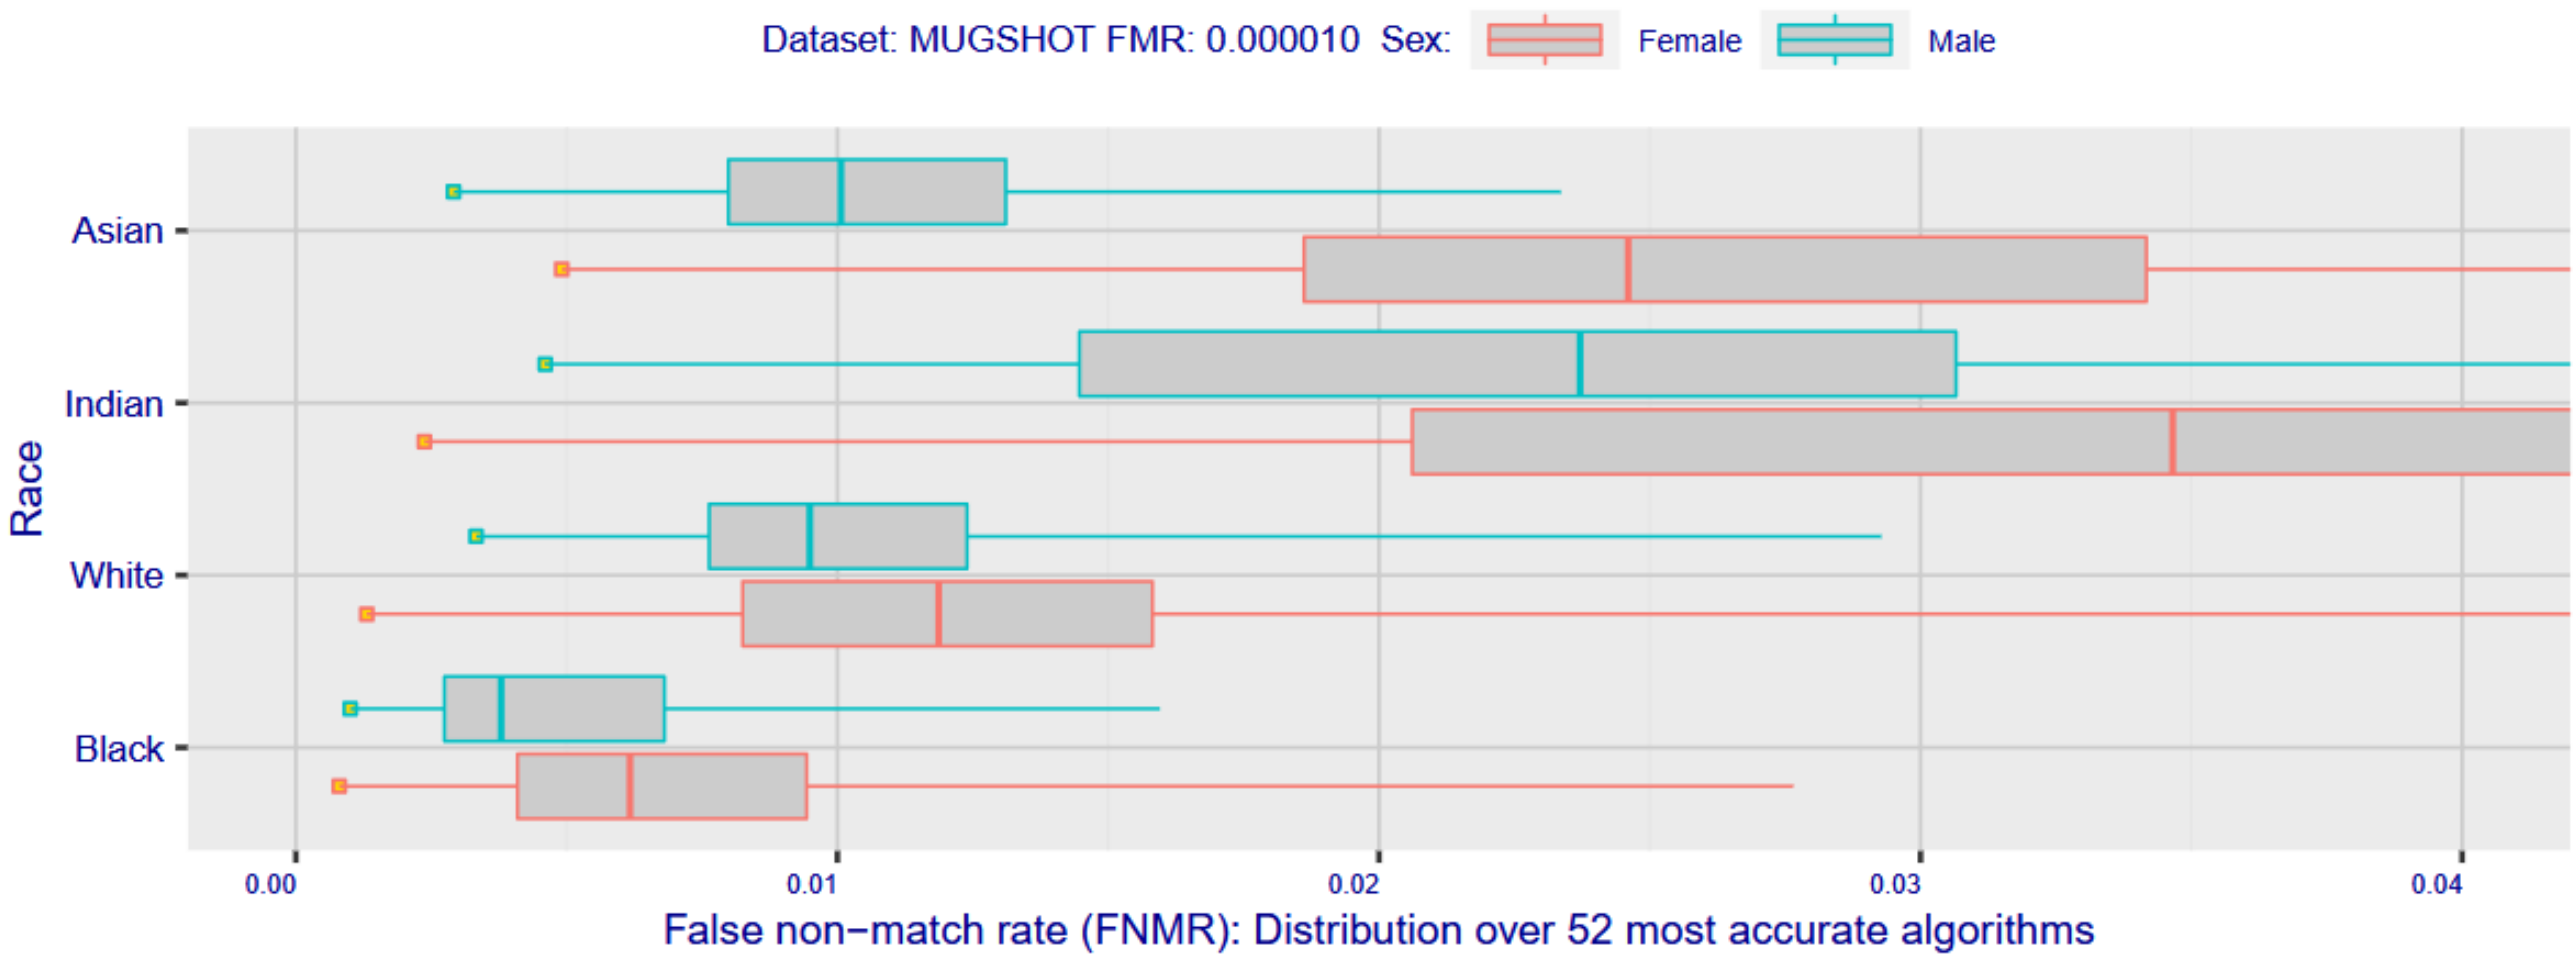

Systems claim to be able to accurately assess emotion from facial expression.

State	Example Photo	Action Units	Physical Description
Amusement		6 + 7 + 12 + 25 + 26 + 53	Head back, Duchenne smile, lips separated, jaw dropped
Anger		4 + 5 + 17 + 23 + 24	Brows furrowed, eyes wide, lips tightened and pressed together
Boredom	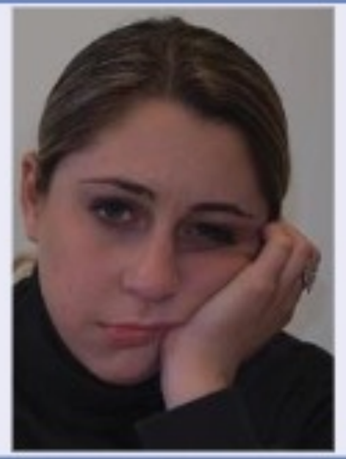	43 + 55	Eyelids drooping, head tilted (not scorable with FACS: slouched posture, head resting on hand)
Confusion	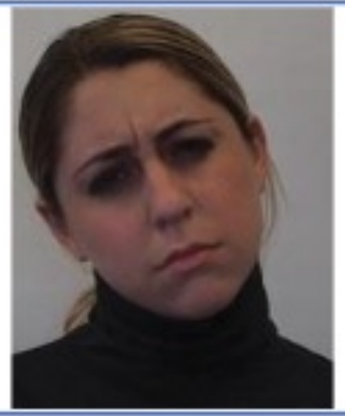	4 + 7 + 56	Brows furrowed, eyelids narrowed, head tilted
Contentment		12 + 43	Smile, eyelids drooping

State	Example Photo	Action Units	Physical Description
Fear		1 + 2 + 4 + 5 + 7 + 20 + 25	Eyebrows raised and pulled together, upper eyelid raised, lower eyelid tense, lips parted and stretched
Happiness		6 + 7 + 12 + 25 + 26	Duchenne smile
Interest		1 + 2 + 12	Eyebrows raised, slight smile
Pain		4 + 6 + 7 + 9 + 17 + 18 + 23 + 24	Eyes tightly closed, nose wrinkled, brows furrowed, lips tight, pressed together, and slightly puckered
Pride	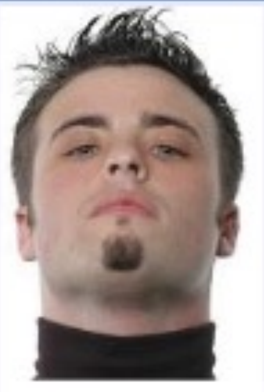	53 + 64	Head up, eyes down

Lab 4 — Whoops! We are so sorry….

This is what you should have seen

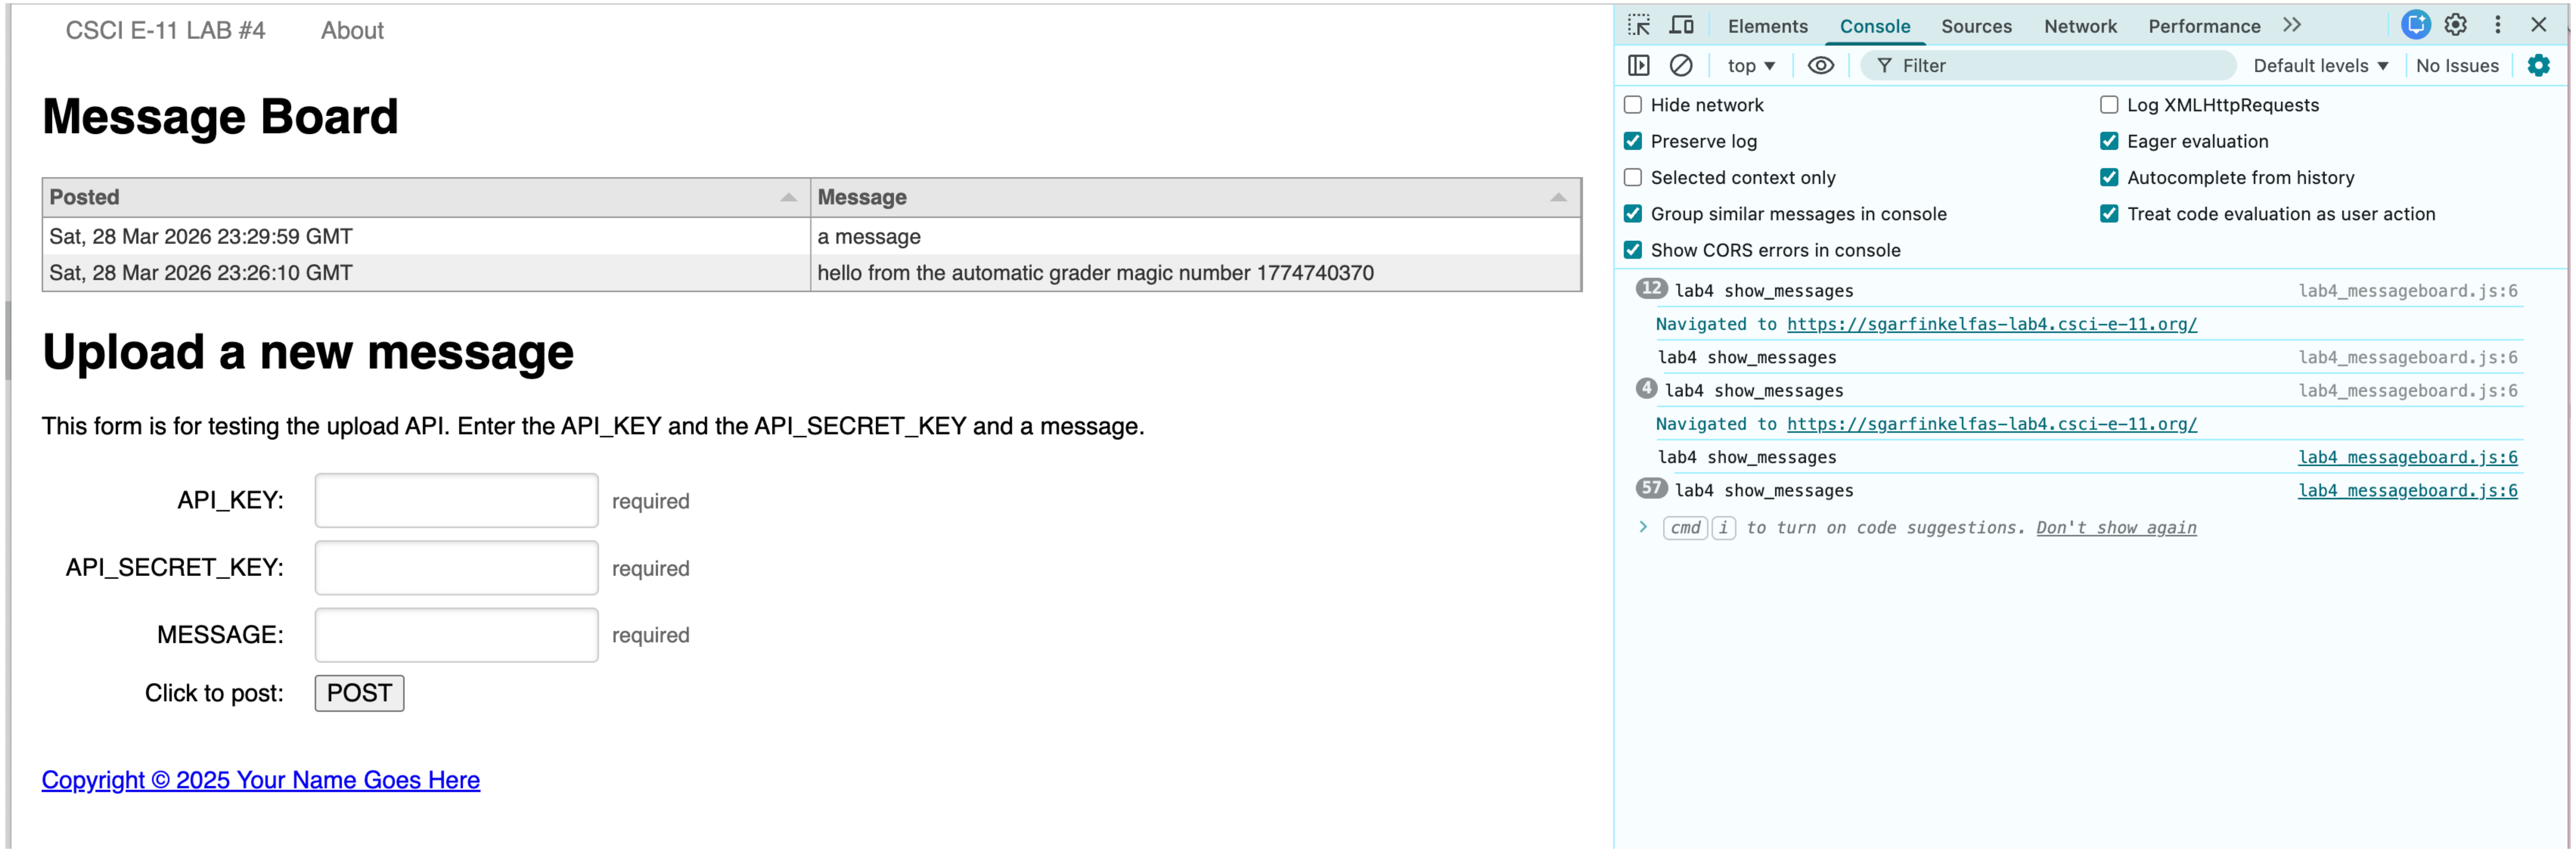

This is the change we made

lab_common/static/lab4_messageboard.js +1 -1 Viewed

```diff
@@ -19,7 +19,7 @@ function show_messages() {
         })
         .then((obj) => {
             // Clear or update the Tabulator table
-            const table = Tabulator.findTable("#message-table")[0];
+            let table = Tabulator.findTable("#message-table")[0];
             if (table) {
                 // Table exists: update its data
                 table.replaceData(obj);
```

lab_common/templates/index.html +2 -2 Viewed

```diff
@@ -65,8 +65,8 @@ <h1>Upload a new message</h1>
         <div id='status-message'></div>
     </div>
     {% if lab_number == "4" %}
-        <div id='lab5'></div>
-        <script src="static/lab4_imageboard.js"></script>
+        <div id='lab4'></div>
+        <script src="static/lab4_messageboard.js"></script>
     {% endif %}
     <!-- If we are in lab5, lab6 or lab8, bring in the image interaction -->
     {% if lab_number >= "5" %}
```

Revised dates

In recognition of our error, we have re-opened lab4 grading.

You can now run the `e11 grade lab4` command until April 4.

In the interest of fairness, we have also extended lab5 grading to April 8.

MEMENTO

adafruit Products Gift Ideas What's New

Kits & Projects / Enclosures / Adafruit MEMENTO Camera Enclosure & Hardware Kit

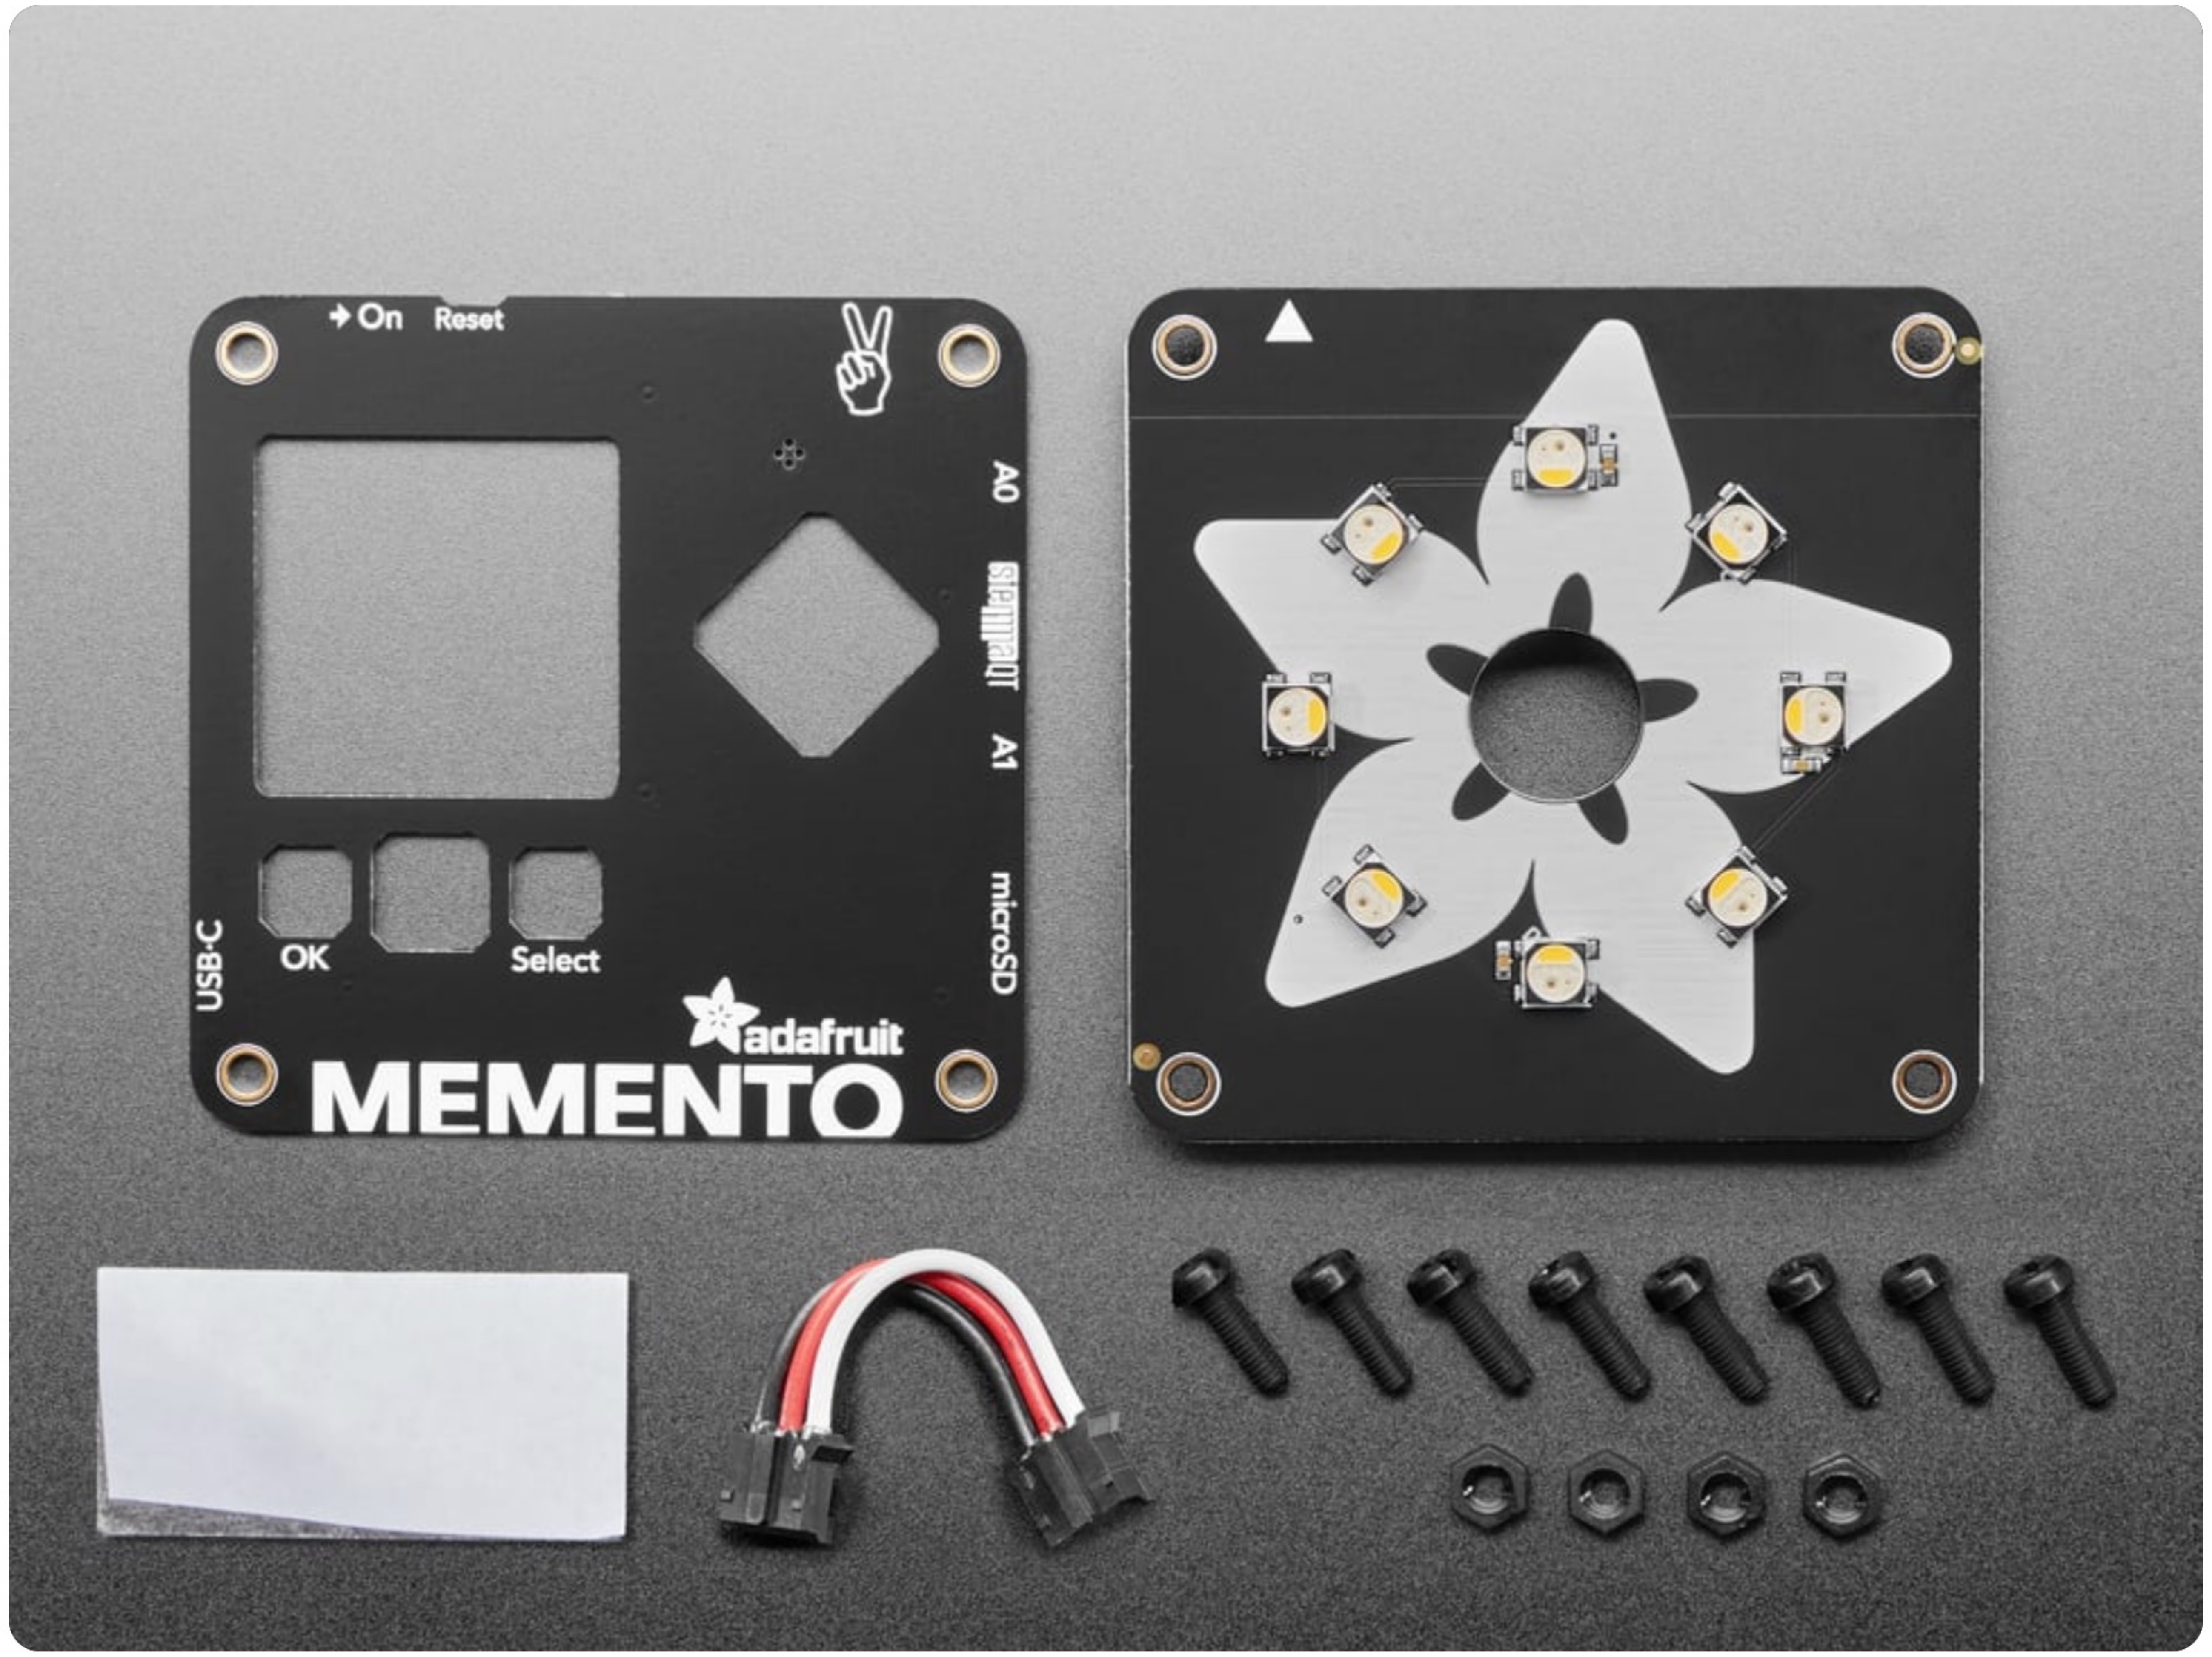

Adafruit MEMENTO Camera Enclosure & Hardware Kit

Product ID: 5843

$9.95

In stock

1 Add to Cart

Qty	Discount
1-9	$9.95
10-99	$8.96
100+	$7.96

Add to Wishlist

Description

Technical Details

Sensors / Cameras / MEMENTO - Python Programmable DIY Camera - Bare Board

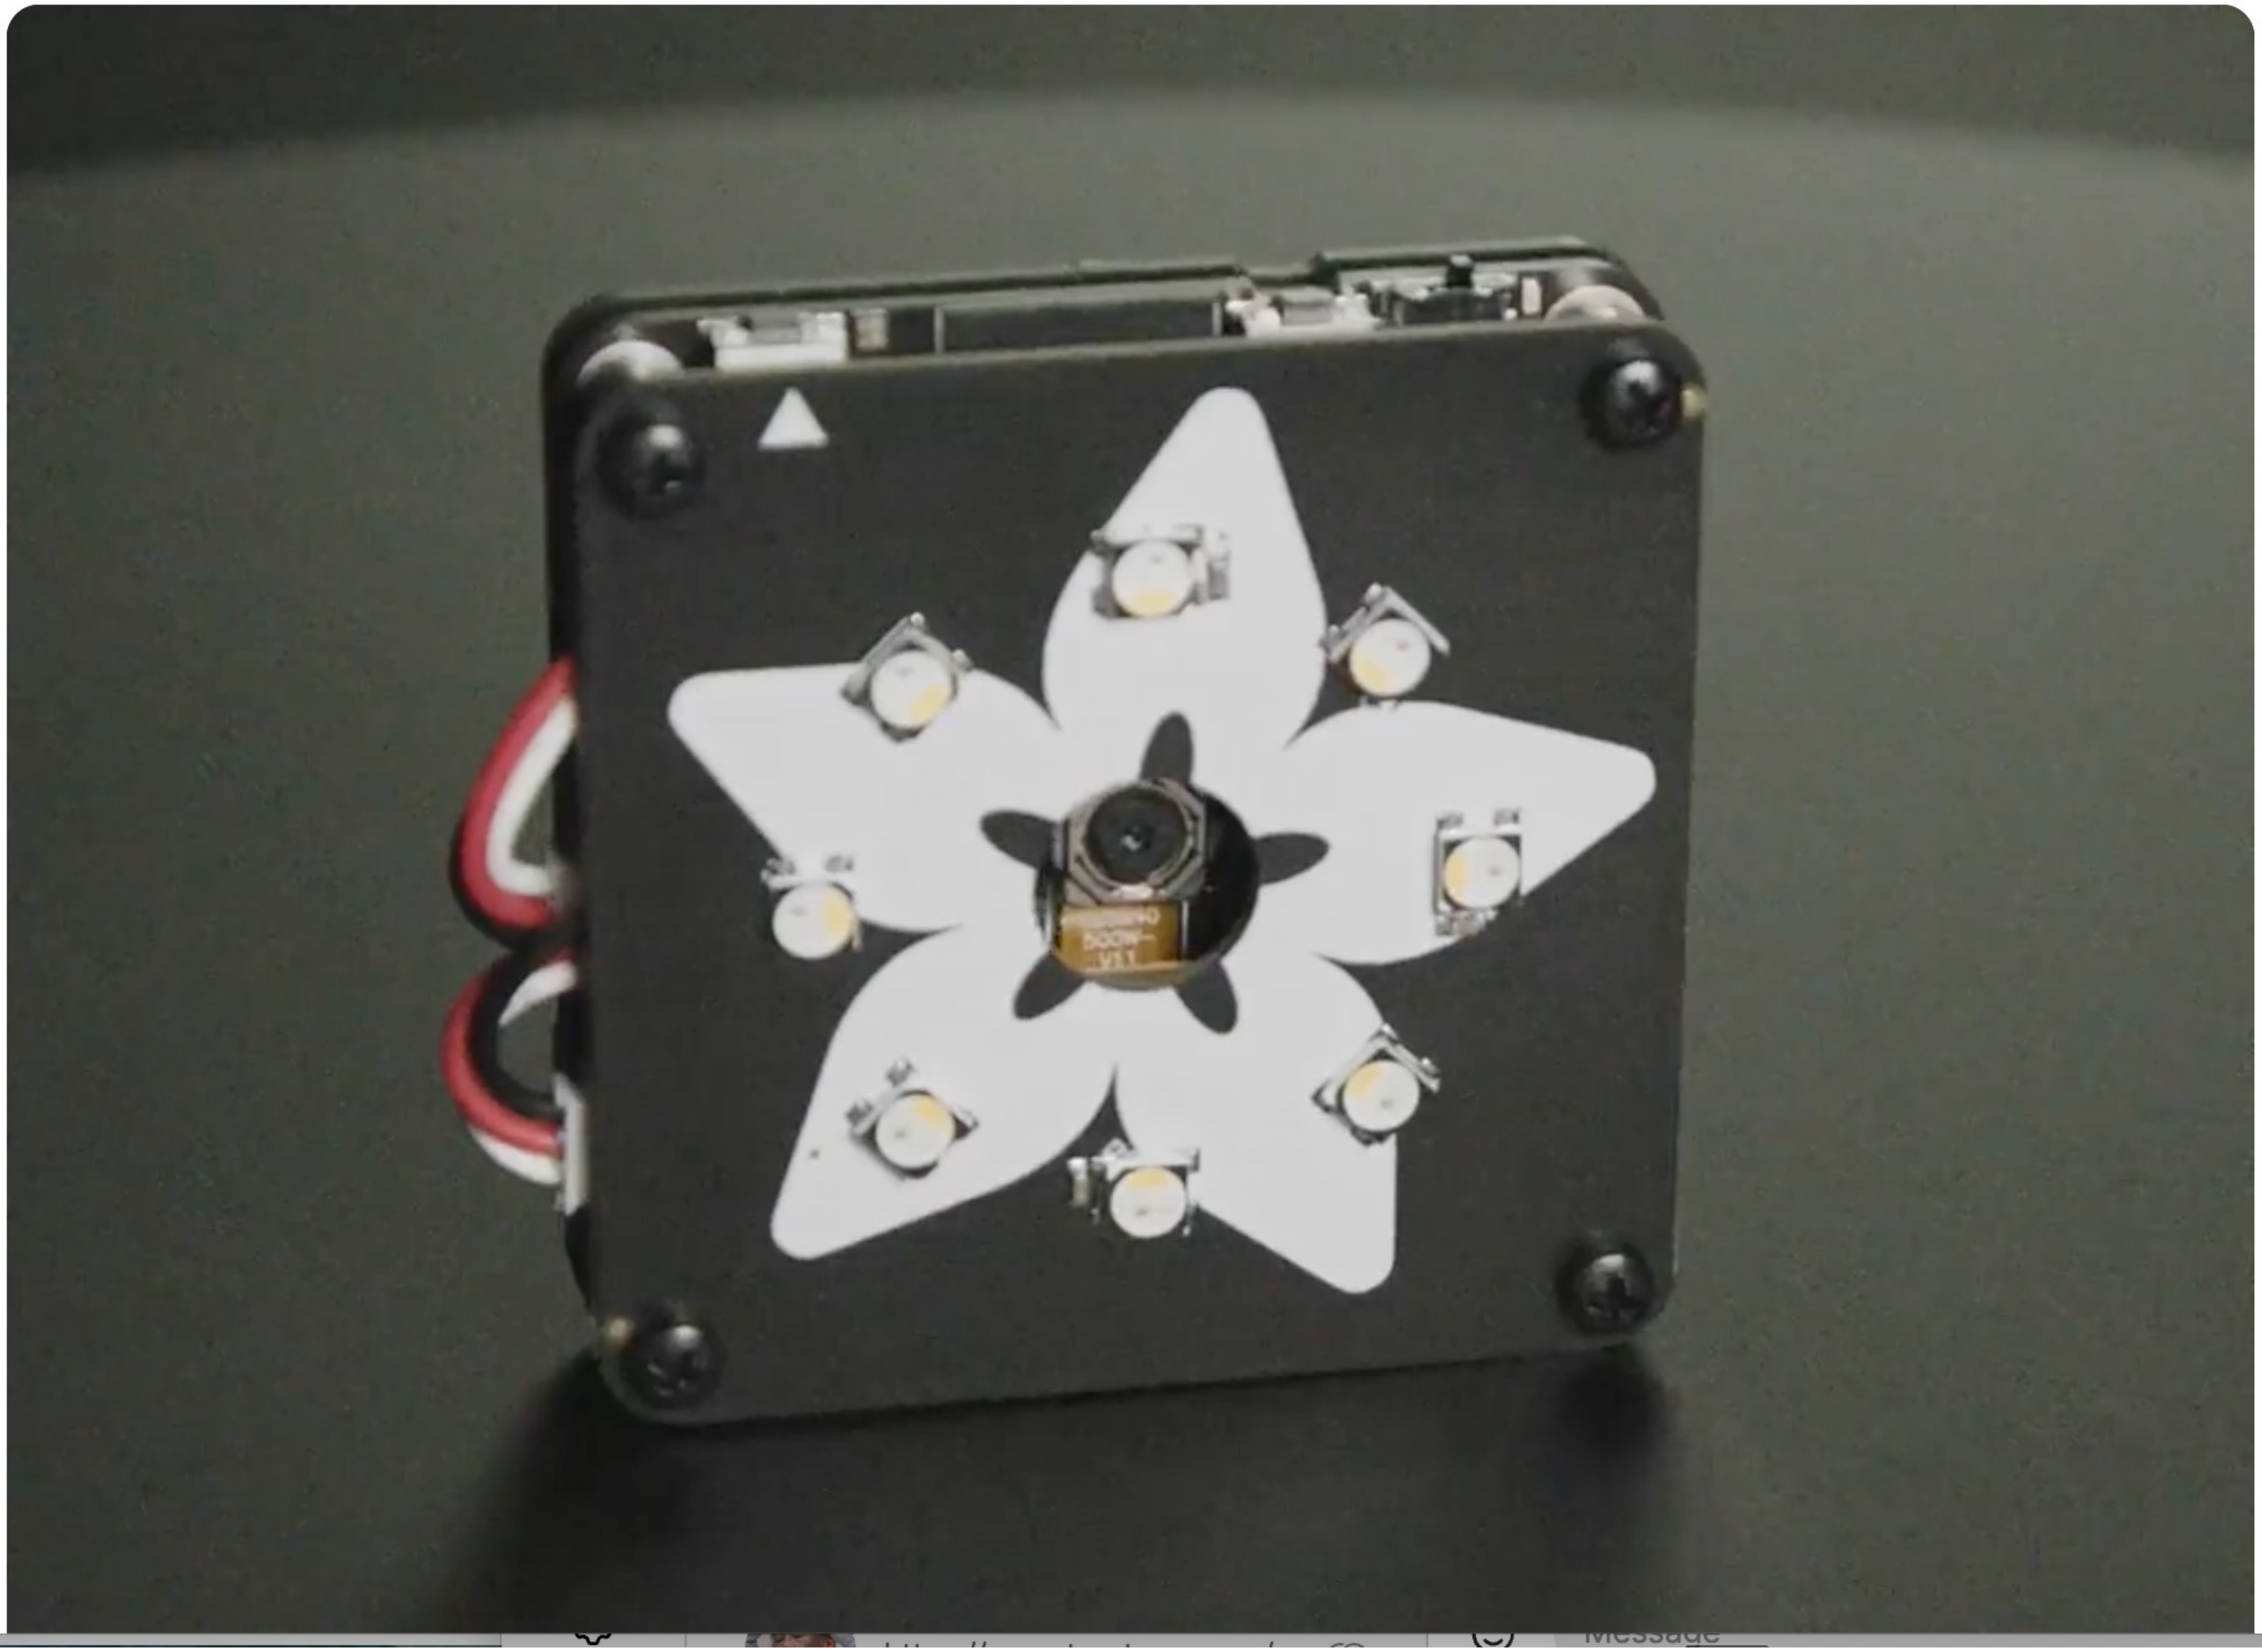

MEMENTO - Python Programmable DIY Camera - Bare Board

Product ID: 5420

$34.95

Out of stock

Please enter your details below and we will send you an email when this item is back in stock! You will only be emailed about this product.

Your Email

Notify Me

Add to Wishlist

Description

Technical Details

Amazon has this in stock! - https://www.amazon.com/dp/B09MS6PH7L

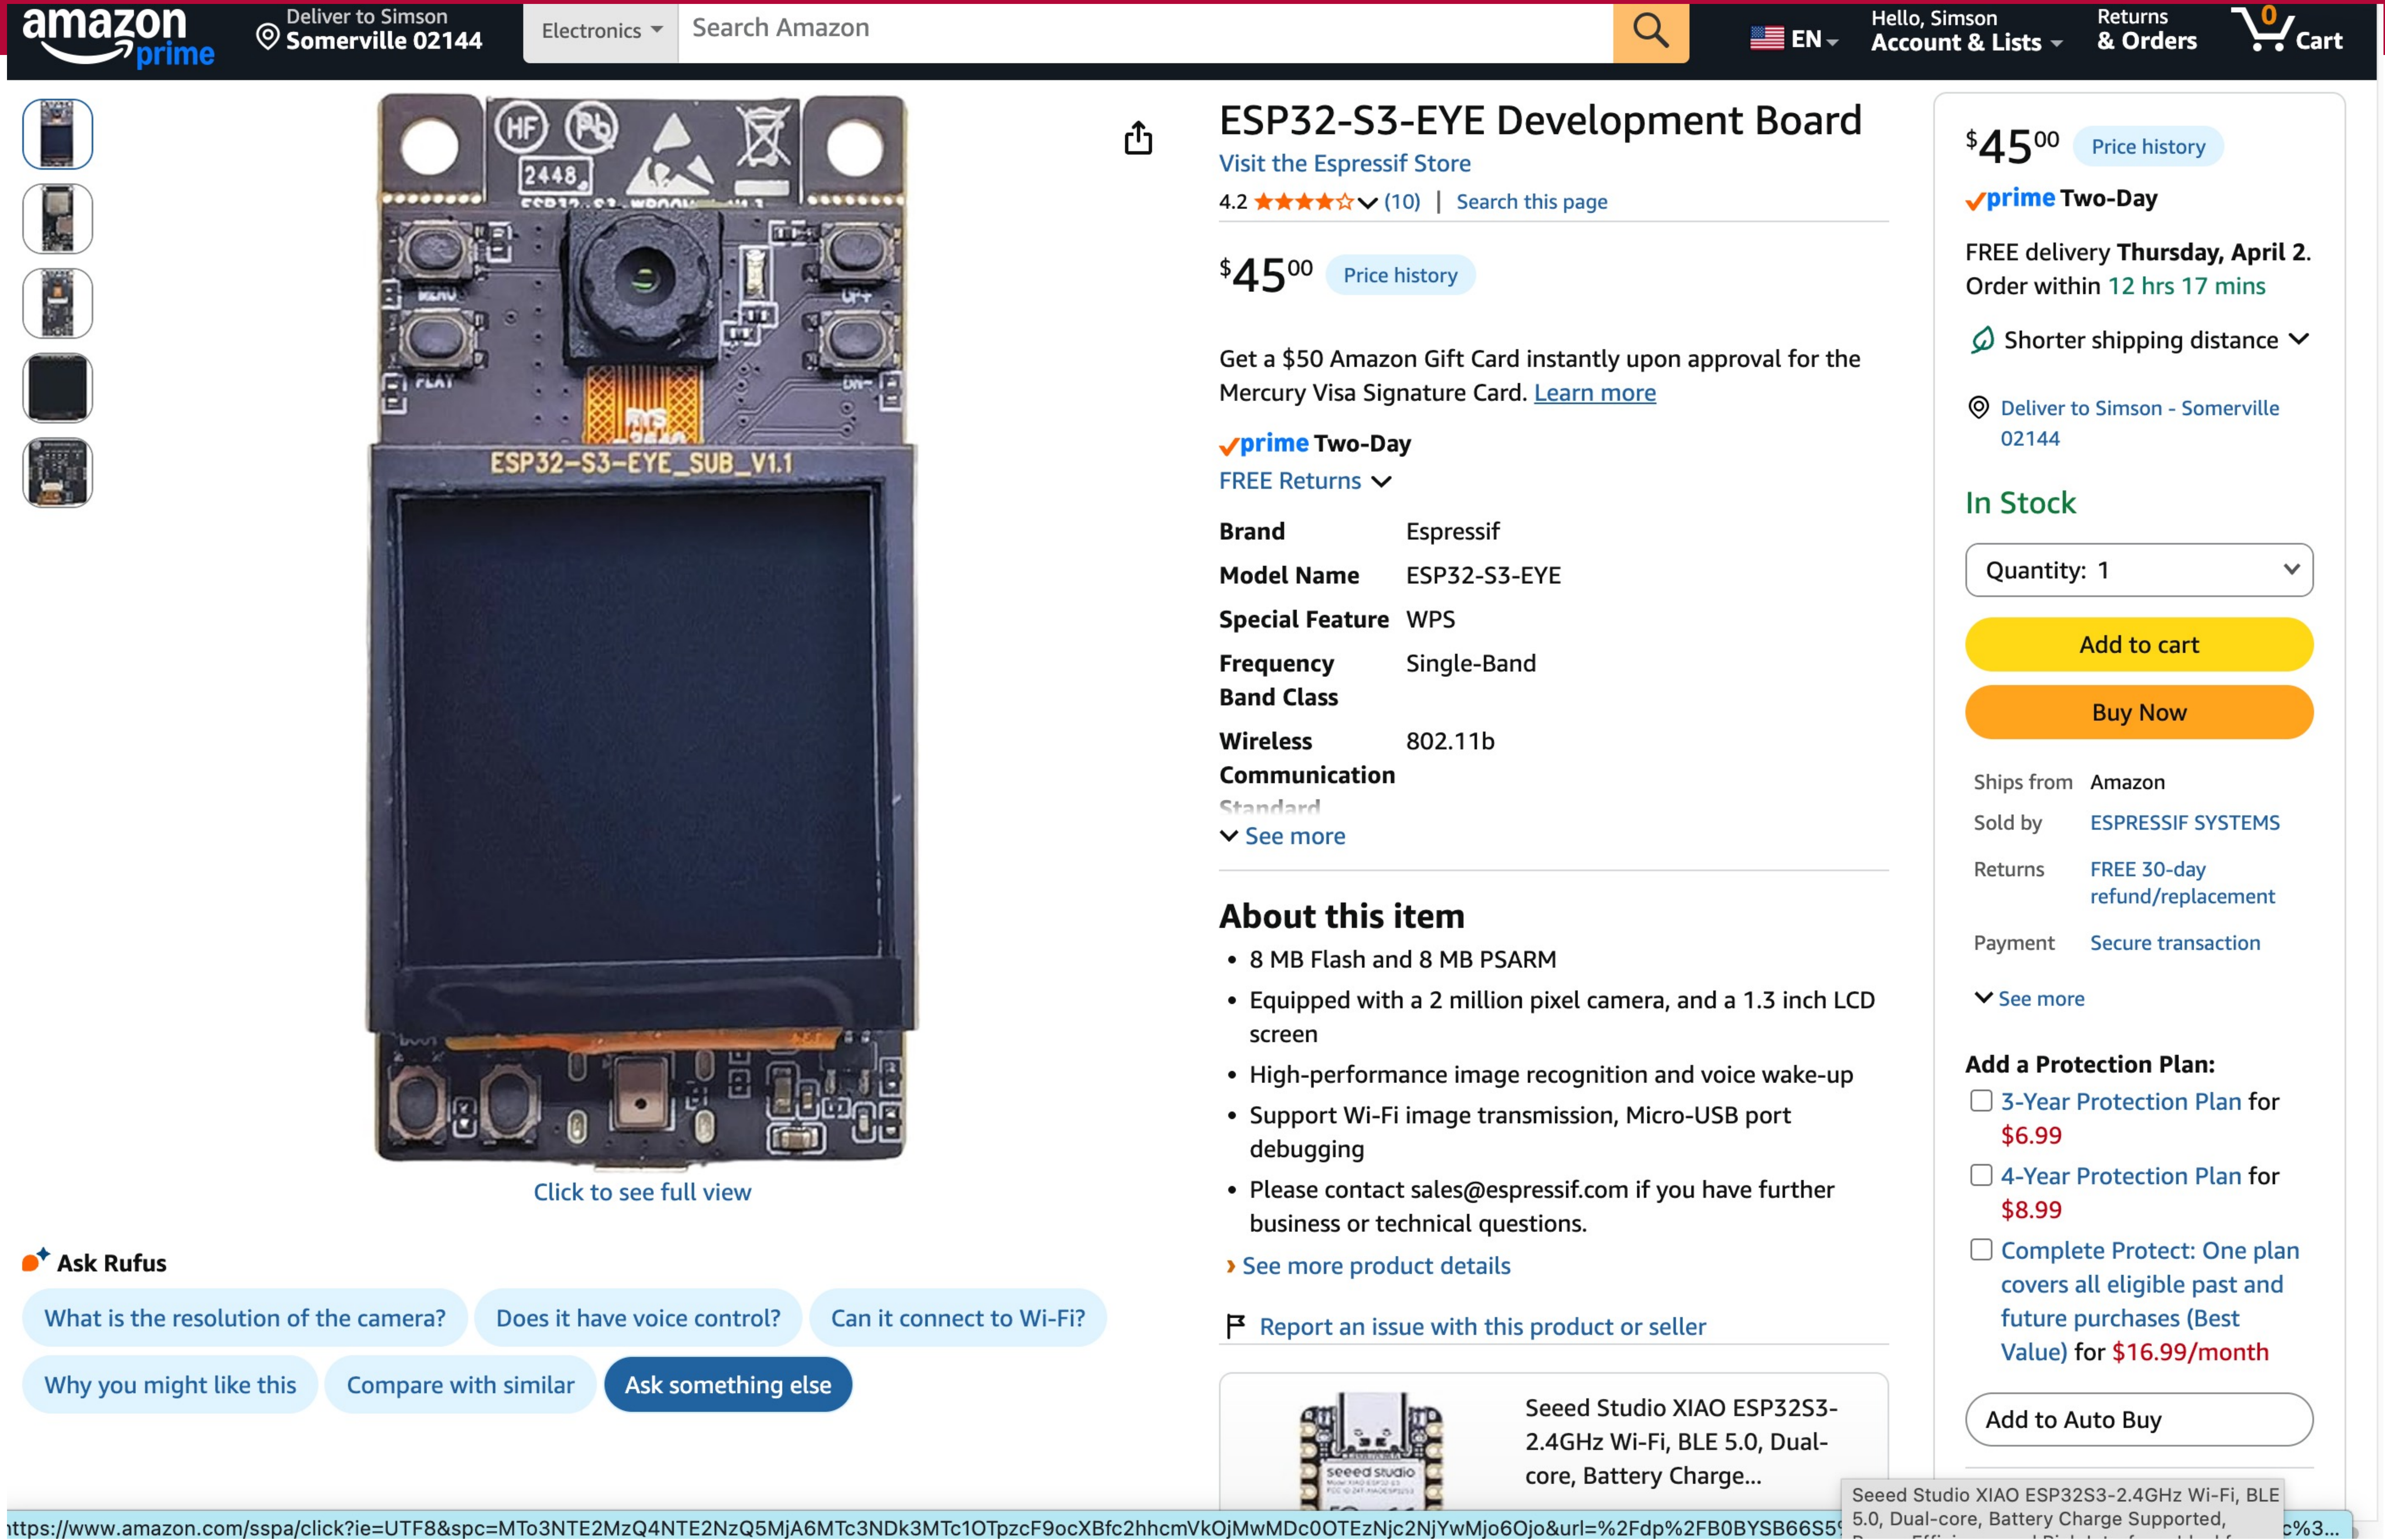

Next Week's Reading

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

Abstract

We trained a large, deep convolutional neural network to classify the 1.2 million high-resolution images in the ImageNet LSVRC-2010 contest into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 37.5% and 17.0% which is considerably better than the previous state-of-the-art. The neural network, which has 60 million parameters and 650,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and three fully-connected layers with a final 1000-way softmax. To make training faster, we used non-saturating neurons and a very efficient GPU implementation of the convolution operation. To reduce overfitting in the fully-connected layers we employed a recently-developed regularization method called "dropout" that proved to be very effective. We also entered a variant of this model in the ILSVRC-2012 competition and achieved a winning top-5 test error rate of 15.3%, compared to 26.2% achieved by the second-best entry.

1 Introduction

Current approaches to object recognition make essential use of machine learning methods. To improve their performance, we can collect larger datasets, learn more powerful models, and use better techniques for preventing overfitting. Until recently, datasets of labeled images were relatively small — on the order of tens of thousands of images (e.g., NORB [16], Caltech-101/256 [8, 9], and CIFAR-10/100 [12]). Simple recognition tasks can be solved quite well with datasets of this size, especially if they are augmented with label-preserving transformations. For example, the current-best error rate on the MNIST digit-recognition task (<0.3%) approaches human performance [4]. But objects in realistic settings exhibit considerable variability, so to learn to recognize them it is necessary to use much larger training sets. And indeed, the shortcomings of small image datasets have been widely recognized (e.g., Pinto et al. [21]), but it has only recently become possible to collect labeled datasets with millions of images. The new larger datasets include LabelMe [23], which consists of hundreds of thousands of fully-segmented images, and ImageNet [6], which consists of over 15 million labeled high-resolution images in over 22,000 categories.

To learn about thousands of objects from millions of images, we need a model with a large learning capacity. However, the immense complexity of the object recognition task means that this problem cannot be specified even by a dataset as large as ImageNet, so our model should also have lots of prior knowledge to compensate for all the data we don't have. Convolutional neural networks (CNNs) constitute one such class of models [16, 11, 13, 18, 15, 22, 26]. Their capacity can be controlled by varying their depth and breadth, and they also make strong and mostly correct assumptions about the nature of images (namely, stationarity of statistics and locality of pixel dependencies). Thus, compared to standard feedforward neural networks with similarly-sized layers, CNNs have much fewer connections and parameters and so they are easier to train, while their theoretically-best performance is likely to be only slightly worse.

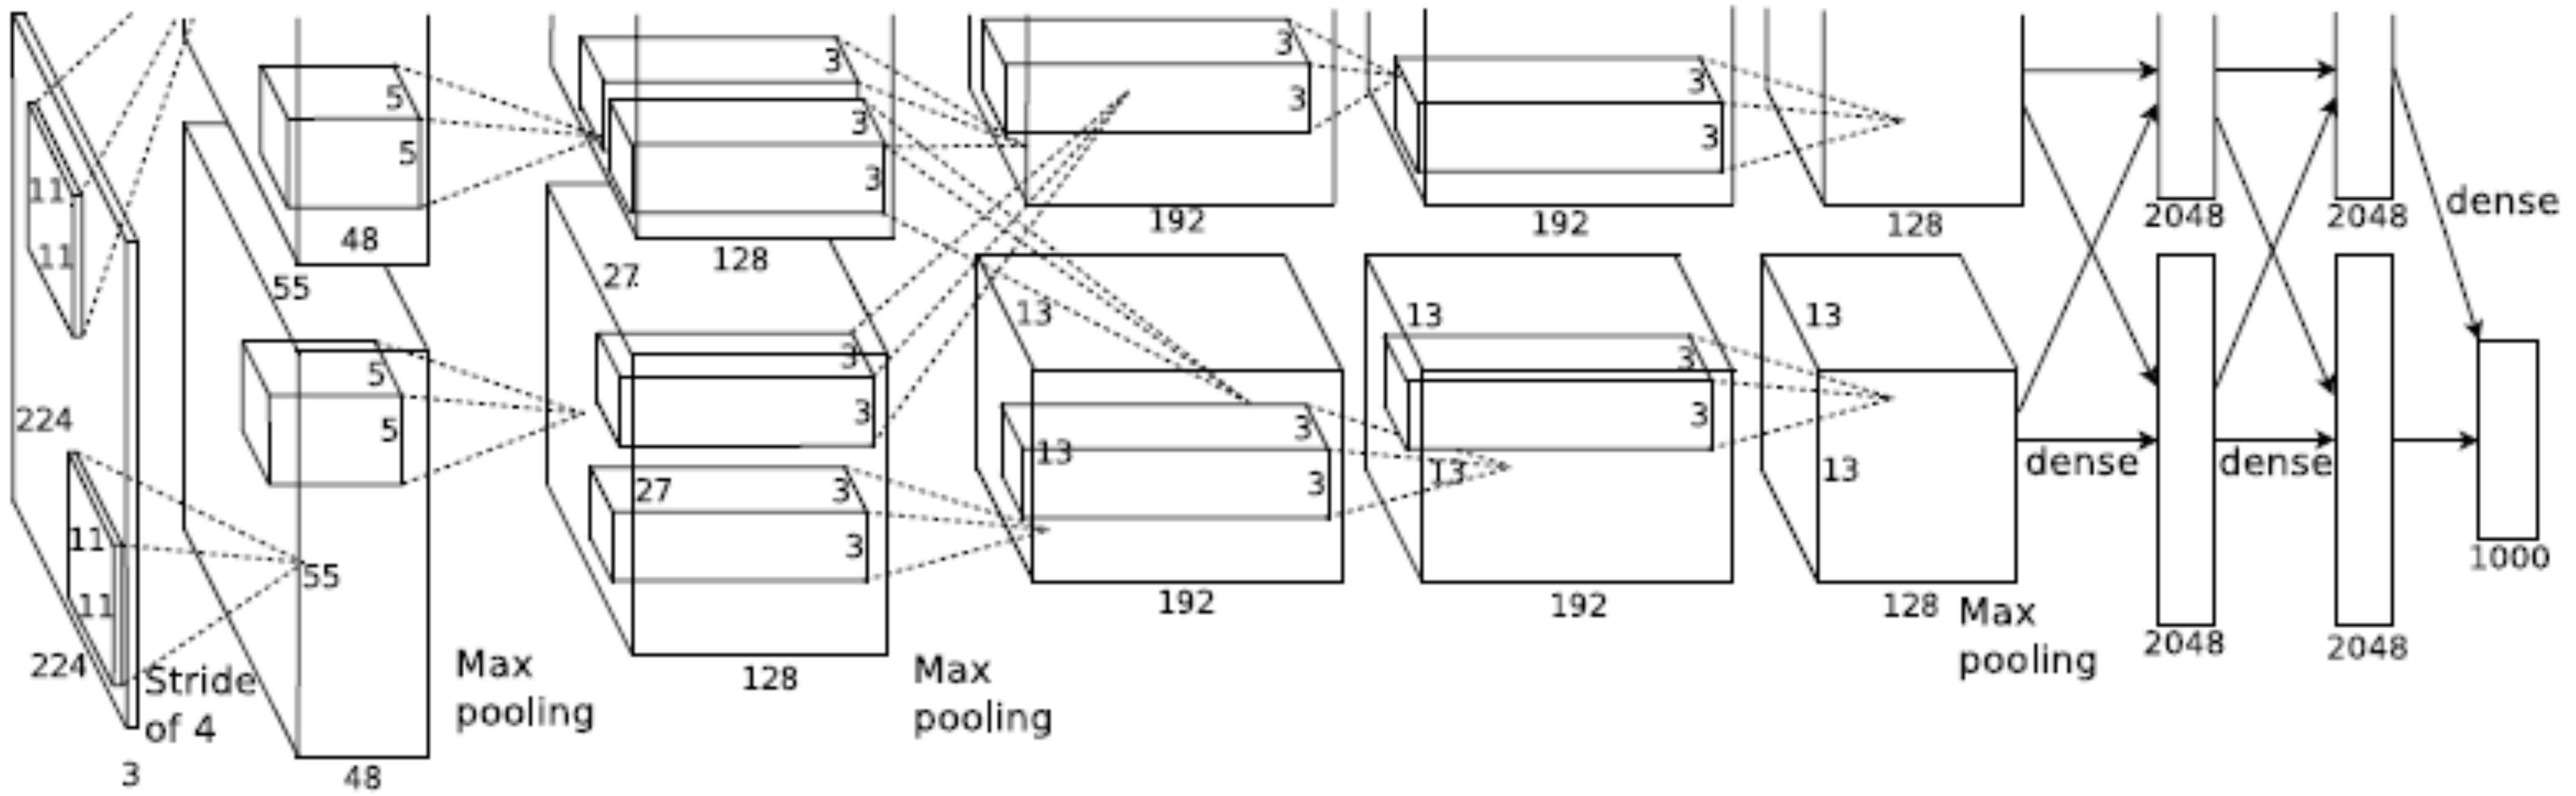

Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253,440–186,624–64,896–64,896–43,264–4096–4096–1000.

Geoffrey Hinton, Nobel Prize in Physics, 2024

Geoffrey Hinton
Facts

© Nobel Prize Outreach.
Photo: Clément Morin

Geoffrey Hinton
Nobel Prize in Physics 2024

Born: 6 December 1947, London, United Kingdom

Affiliation at the time of the award: University of Toronto, Toronto, Canada

Prize motivation: “for foundational discoveries and inventions that enable machine learning with artificial neural networks”

Prize share: 1/2

Ilya Sutskever, co-founder and (former) Chief Scientist of OpenAI

Ilya Sutskever

20 languages

Article Talk Read Edit View history Tools

From Wikipedia, the free encyclopedia

Ilya Sutskever (Hebrew: איליה סוצקבר; born 1986) is a computer scientist who specializes in machine learning.[4] He has made several major contributions to the field of deep learning, including sequence-to-sequence learning, reasoning models, GPT models, and contributions to CLIP, DALL-E, and AlphaGo.[5] With Alex Krizhevsky and Geoffrey Hinton, he co-created AlexNet, a convolutional neural network.[6] One of the most highly cited computer scientists in history, he has won the NeurIPS Test of Time Award for his lasting impact on AI research three times in a row (2022-2024) and received the National Academy of Sciences Award for the Industrial Application of Science in 2026.[7][3]

Sutskever co-founded and was chief scientist at OpenAI, where he oversaw the research breakthroughs that led to large language models and to the launch of ChatGPT.[8][5] He also led the research that led to reasoning models such as o1.[5][9][10] In 2023, he was one of the members of OpenAI's board that ousted Sam Altman as its CEO; Altman was reinstated a week later, and Sutskever stepped down from the board. In June 2024, Sutskever co-founded the company Safe Superintelligence Inc., alongside Daniel Gross and Daniel Levy.[11][12] Within a year, the company was valued at more than $30 billion.[13]

Ilya Sutskever
FRS
איליה סוצקבר

Sutskever at Tel Aviv University in 2023

Born	1986 (age 39–40) Gorky, Soviet Union

https://en.wikipedia.org/wiki/Ilya_Sutskever

For Next Week:

- Read the ImageNet paper
    - Comment on Perusall
    - Complete the reading quiz
- Participate in the discussion on Ed Board
- Catch up on Labs 4 and 5
- Start Lab 6
- Start your taxes (US)