Computer Science:

Where is it coming from and where is it going to?

CS Immigration Course – Fall 2019

Education: Where are we?

The Definitive Statement on Education

Education Level	What You Think You Know	How You Act	What You Learn
Grade School	How To Have Fun	Try To Have Fun	How To Behave
High School	Everything	Like You Know Everything	How To Learn
College	Just About Everything	Like You Know Quite A Lot	That There Are Things You Don't Know
Graduate School (Masters)	Some Things	Like You Know A Lot	That You Really Don't Know Much
Graduate School (PhD)	Nothing	Like You Don't Want People To Know That You Know Nothing	How Huge And Vast An Amount You Really Don't Know

Final word on Education

- “No one should escape our universities without knowing how little s/he knows” - Robert Oppenheimer – Physicist, Scientific Director of the Manhattan Project

The Burning Question

- What is Computer Science?
- What is science?
  - Science is the study and understanding of the possible (and beyond.)
  - Science is mainly analytic, that is,
    - it tries to analyze, understand and describe nature
    - and also the unnatural,
      - Beyond a certain complexity, artifacts have complex behavior
      - Herbert Simon, The Sciences of the Artificial

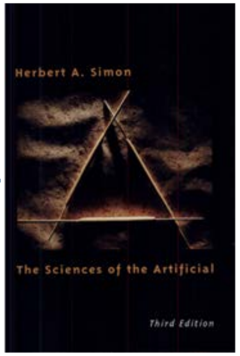

  - But Computer Science also involves a lot of engineering.
    - Yes, you can tell your grandma, you will be an engineer and she will be proud (☺)

What is Computer Science?

- Computer Science is the study of how information is created, processed and communicated.

What is Information?

- Hard to say!
- You know when you have it (and when you don't have it), but
- You can't touch it!
- It takes energy, time, money to produce it, but yet it is very abstract.
- You can store it for later use (if you don't, you lose it permanently)
- You can measure it! (it has a unit-- really!)

Wht s nfrmtn?

- Hrd t sy!
- Y knw whn y hv t (nd whn y dn't hv t), bt
- Y cn't tch t!
- t tks nrgy, tm, mny t prdc t, bt yt t s vry bstrct.
- Y cn str t fr ltr s (f y dn't, y ls t prmnntly)
- Y cn dplct t, y cn sll t r y cn by t.
- Y cn msr t! (t hs a nt – Rlly!)

a i loaio?

- a o a!
- ou o e ou ae i (a e ou on a i), u
- ou a ou i!
- l ae ee, ie, oe o oue i, u e i l e aa.
- ou a oe l o ae ue (l ou o, ou oe l eae)
- ou a eaue i! (l a a ui)

whatisinformation

- hardtosayYouknowwhenyouhaveitandwhenyoudonthaveitbutyoucanttouchitittakesenergytimemoneytoproduceitbutyetitisveryabstractyoucanstoreitforlateruseifyoudontyouloseitpermanentlyyoucanmeasureitithasaunitreally

What is information?

- Content vs. Representation

What is information?

- Content vs. Representation
- Content is “music”, an arrangement of notes

What is information?

- Content vs. Representation
- Content is “music”, an arrangement of notes
- Representation
  - Music scores
  - Air waves – “Sound”
  - Sequence of digital representations of the “sound”
  - Wav file, mp3, AAC
- Information is in the content, NOT in the representation.
  - Representation just carries or communicates it
  - But we operate on the representation

What is information?

- There is even a mathematical definition in the science of "Information Theory"
- The definition basically says "The more unexpected you perceive an event to be, the more information you 'gain' when you observe the event happens"
  - Information gained in observing an event e is $-\log_2 P(e)$ bits

Hot off the Presses

- **"Similar information rates across languages, despite divergent speech rates"** Christophe Coupé, Yoon Mi Oh, Dan Dediu, François Pellegrino. *Science Advances*, 4 September 2019
- Human language communications happens at about 39 bits/sec

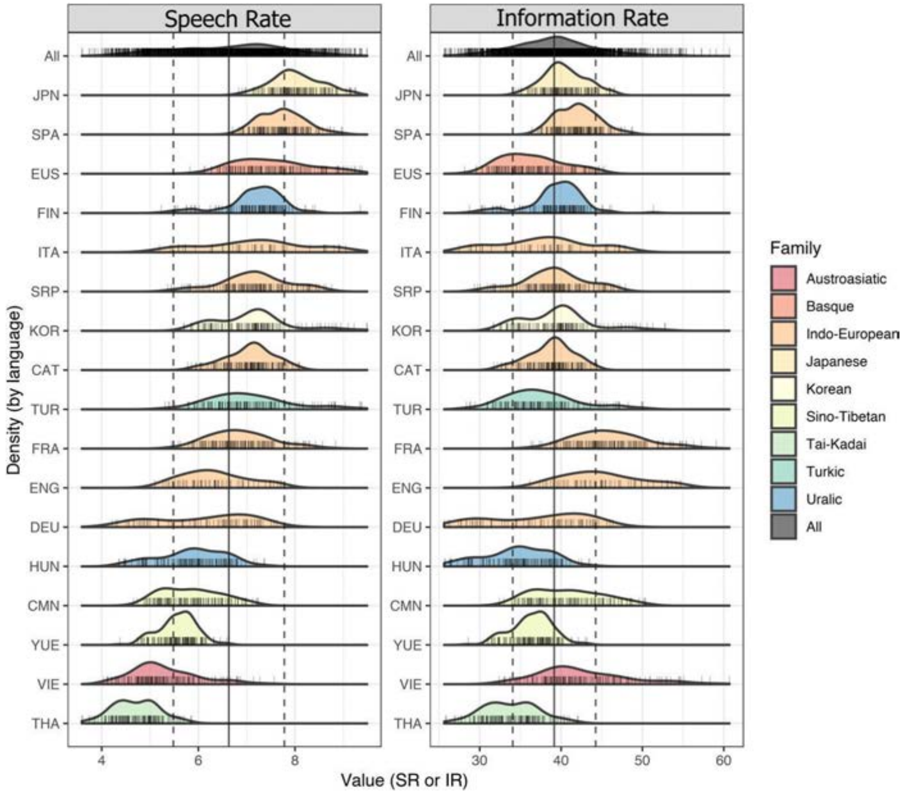

What is Computer Science?

- The fundamental question underlying all of computer science is:

What processes can be automated and how do we automate them (using computers)?

What is Computer Science?

- Computer Science also involves engineering hardware and software systems that
  - Store
  - Transmit
  - Process

  Information towards automating a process.

What is Computer Science?

- Computer scientists also do research on how to build systems:
  - that are new (incorporating new ideas and technology),
  - that solve an important problem hitherto unsolved satisfactorily,
  - that are "cheap" (so that they are affordable)

Building Systems

- Computer Scientists build information processing systems by bringing together
  - hardware
  - software

  so that the system
  - works (i.e. has (almost) no bugs)
  - is fast, maintainable, robust, and affordable.

- Computer Scientists build, experiment with, and evaluate prototype systems that incorporate new
  - ideas,
  - technologies,
  - approaches

  to solve old or new problems.

Skills of a good computer scientist

- Good, open mind
  - things are changing faster that you think
  - Rate of change of change is increasing
    - For calculus geeks: second derivative is increasing (😱)
    - Accelerating change
- Common sense
  - engineering, like politics, is the art of compromises
- Ability to analyze
  - crucial!
- Ability to communicate
  - written, verbal, bi-lingual (at least!)
  - That is why you have to be goo in reading AND writing!

Skills of a good computer scientist

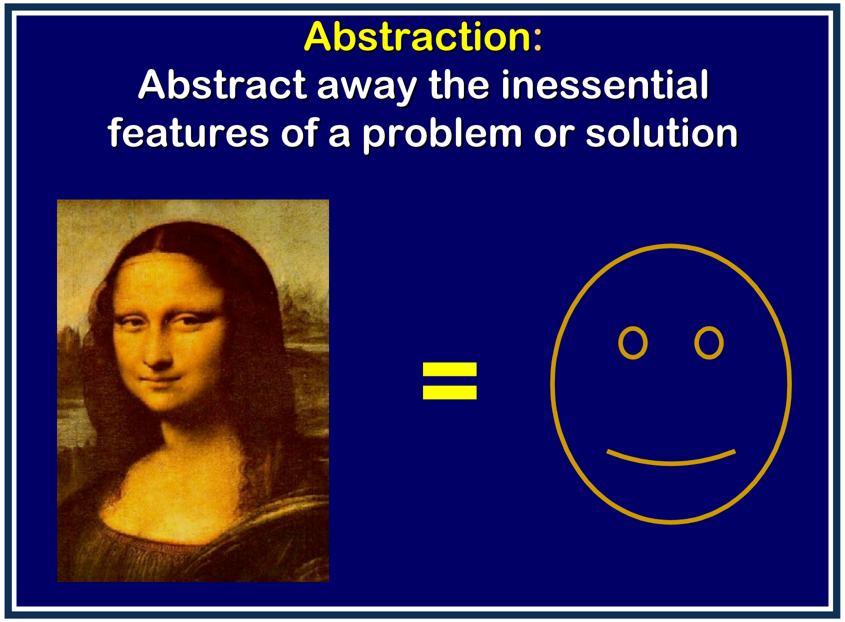

- **Ability to abstract**
  - understand the detail, but always in relation to the whole,
  - know how to forget or avoid the details, when necessary
- **Ability to synthesize**
  - Crucial

Skills of a good computer scientist

- Ability to learn and generalize
  - from successes and failures
- Ability to do trade-offs
  - you do not need a multi-core 3.7 Ghz 8-core super duper system if all you need to do is word processing!
    - Intel thinks otherwise (☺)
- Ability to justify hard decisions
- Very good understanding of fundamental theory and techniques of the field.
  - Mathematics, theoretical computer science, algorithms, data structures, hardware, operating systems, programming, ...

Skills of a good computer scientist

- Understanding the problem, requirements, specifications
- Decomposing problems into manageable parts (Always!)
- Good understanding of data and its organization
- Knowing how to design and/or select efficient algorithms
- Synthesis
  - programming in the small,
    - you writing your toy code
  - programming in the large -- software engineering
    - 1000s people writing MacOS X, Windows, Android, Chrome, etc.

Skills of a good computer scientist

- Understand what can be feasibly computed
- Understand what can NOT be feasibly computed
- Understand what can never be computed

Scooping the Loop Snooper

an elementary proof of the undecidability of the halting problem

No program can say what another will do.
Now, I won't just assert that, I'll prove it to you:
I will prove that although you might work til you drop,
you can't predict whether a program will stop.

….

You can never discover mechanical means
for predicting the acts of computing machines.
It's something that cannot be done. So we users
must find our own bugs; our computers are losers!

by Geoffrey K. Pullum
Stevenson College
University of California
Santa Cruz, CA 95064

What will you be when you grow up?

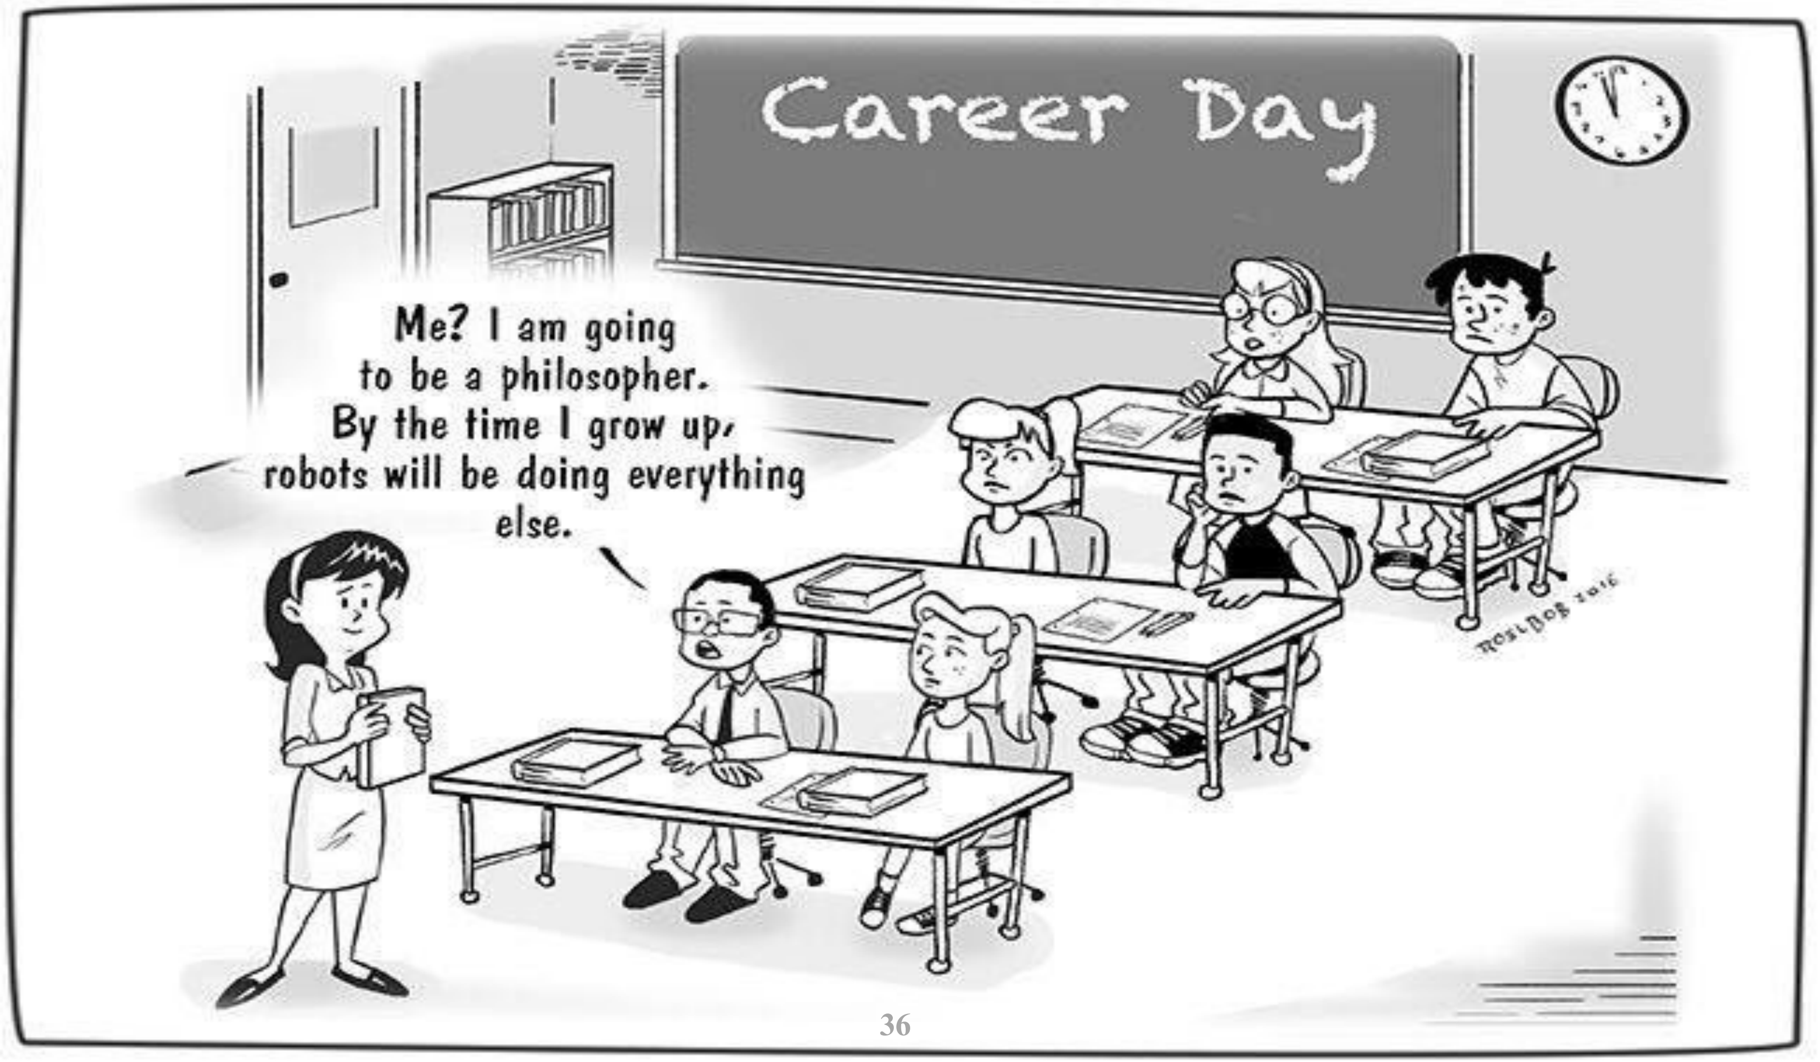

What will you be when you grow up?

- Job opportunities
  - Software Companies
  - IT Departments
    - Banks (Web development, security, speech processing, databases)
    - Manufacturing (Databases, process automation, simulation, vision, robotics, AI, machine learning)
    - Media (Graphics, Web development, text processing, information retrieval, NLP)
    - Telecom (Networking, security)
    - General (System Management)
  - Start-up companies
  - Research Centers (e.g. QCRI, QBRI, QEERI, Google, Microsoft Research, Baidu, Facebook, Amazon, etc.)
  - Tech Companies (Twitter, Bosch, Salesforce, Dropbox, LinkedIn, Uber, etc.)

What will you be when you grow up?

- Graduate Study
  - Master of Science (M. Sc.) (2 years – Advanced Courses + Research+Thesis)
  - Doctor of Philosophy (Ph. D.) (4-6 years – more advanced courses + original research+Thesis)

- Job opportunities
  - Faculty Members/Researchers at Universities
  - Researchers at advanced research labs
  - Research policy managers

Specialization

- An expert is someone
  - who knows more and more
  - about less and less, until
  - s/he knows everything about nothing (☺)

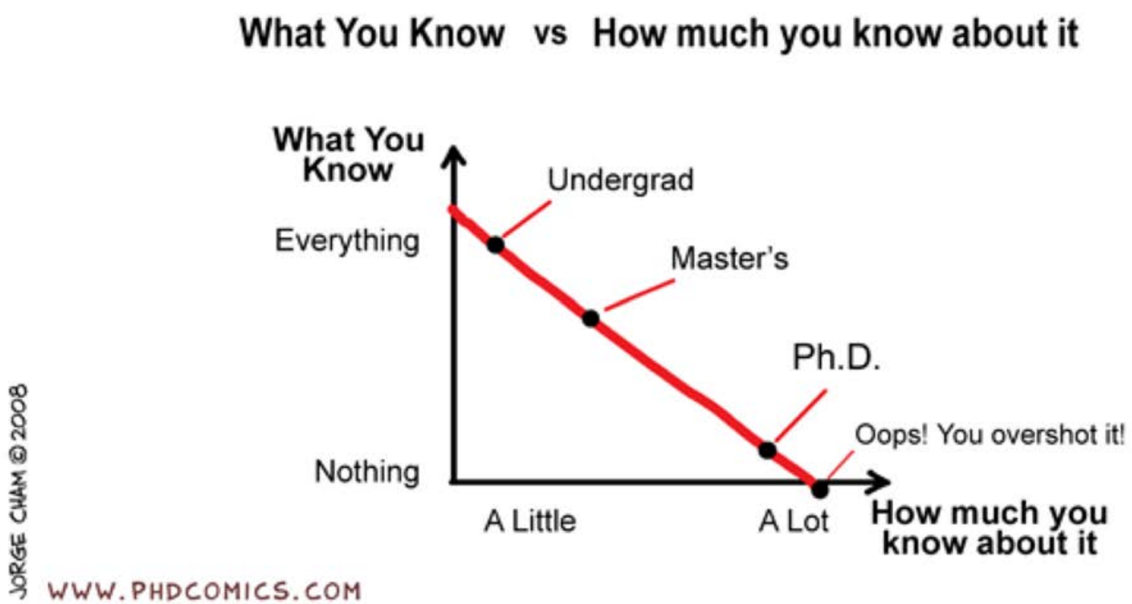

The Evolution of Computing

- Where are we coming from?
  - You probably do not have a good feel of what life in the computerland was like from the 1960's through 2000's.
  - Hell – you probably can not imagine a world without smartphones, Internet, Facebook, Instagram, etc.
- Life was very different 40-50 years ago. (and I know you do not care!)
  - But most everything you see now have been envisioned at least 40-50 years ago if not more!

The Evolution of Computing

- I suggest all of you watch

"2001 A Space Odyssey"
(Kubrick 1968)

to see what the best of minds thought (in 1968) computers would be like in 2001.

- Playing Chess
- Dialogue with Humans
- Planning (and Conspiracy)

The Evolution of Computing

- 1950's computer

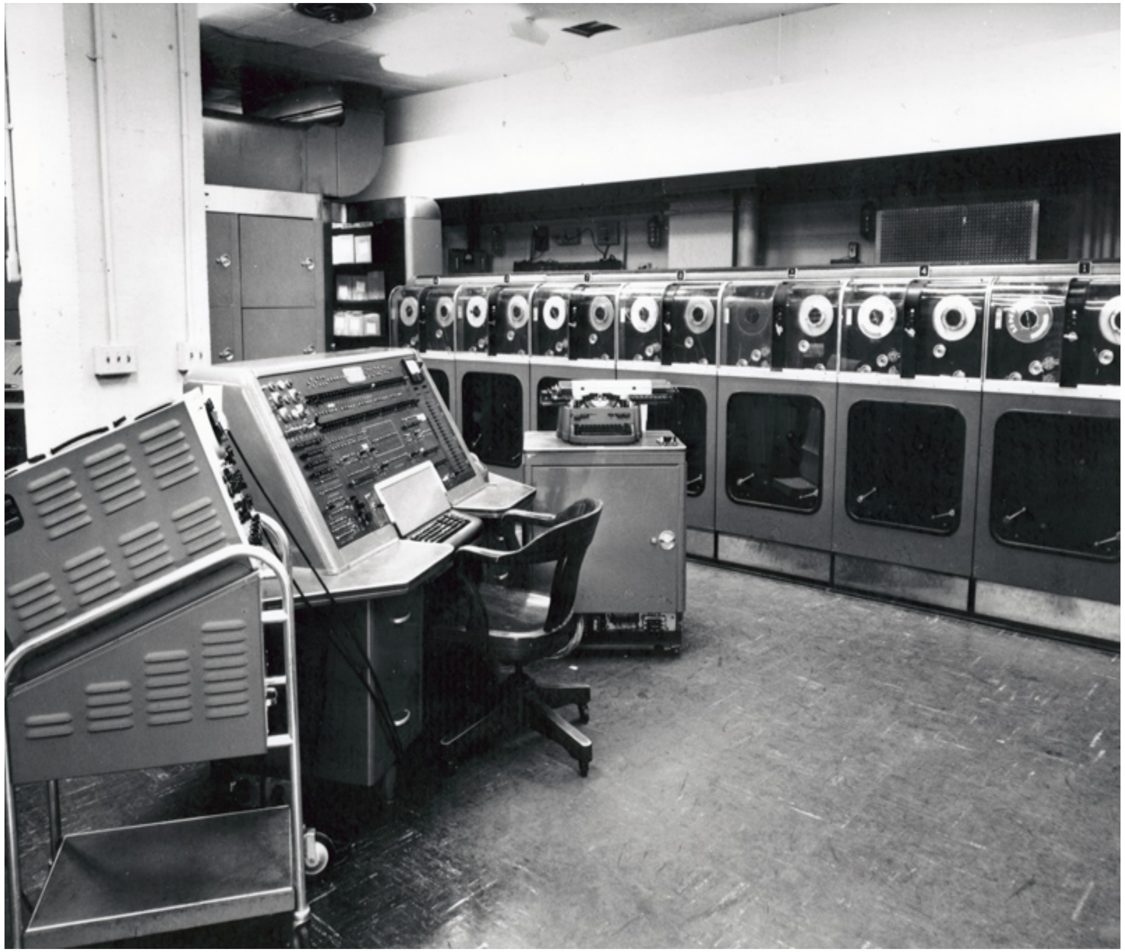

The Evolution of Computing

- Then and now!

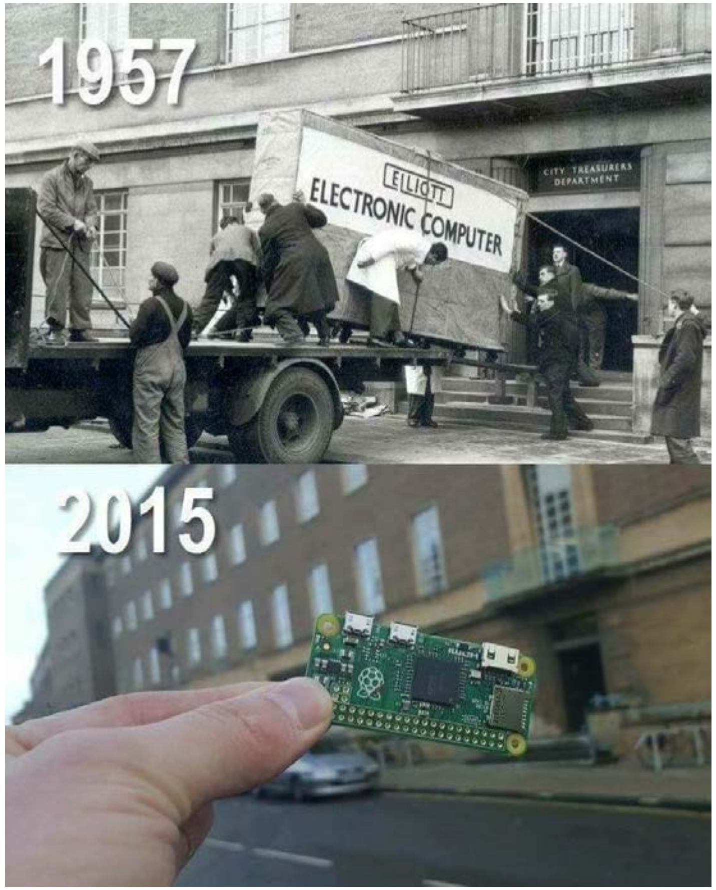

The Evolution of Computing

- 1950's hard disk drive

The Evolution of Computing

- 1950's hard disk drive

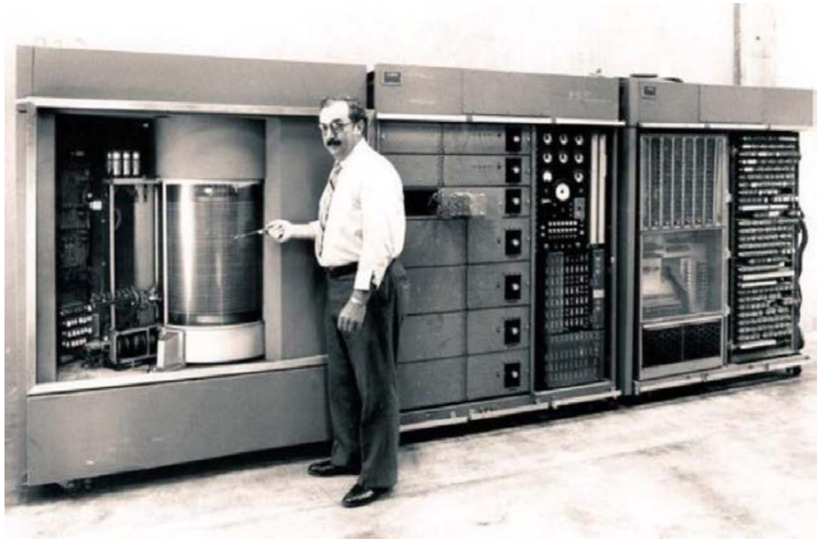

A personal history of computer technology

- 1960's:
  - 100 KHz Machines (IBM 1620),
  - 20-30 KB Memory,
  - Punched Card Input,
  - Teletypes, Line Printers,
  - 10-20 MB Disks
  - Computers consumed kilowatts, need serious AC cooling,
  - **ARPANET, the precursor to Internet, starts about here also.**
  - **"Deep Learning" also starts here – Perceptrons!**
  - **AI actually started in late 1950's**
    - **Dartmouth Conference (1956)**

Dartmouth Conference on AI (1956)

A PROPOSAL FOR THE

DARTMOUTH SUMMER RESEARCH PROJECT

ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

August 31, 1955

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be

The following are some aspects of the artificial intelligence problem

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

A personal history of computer technology

- 1970 IBM 1620
  - The first computer I wrote a program for

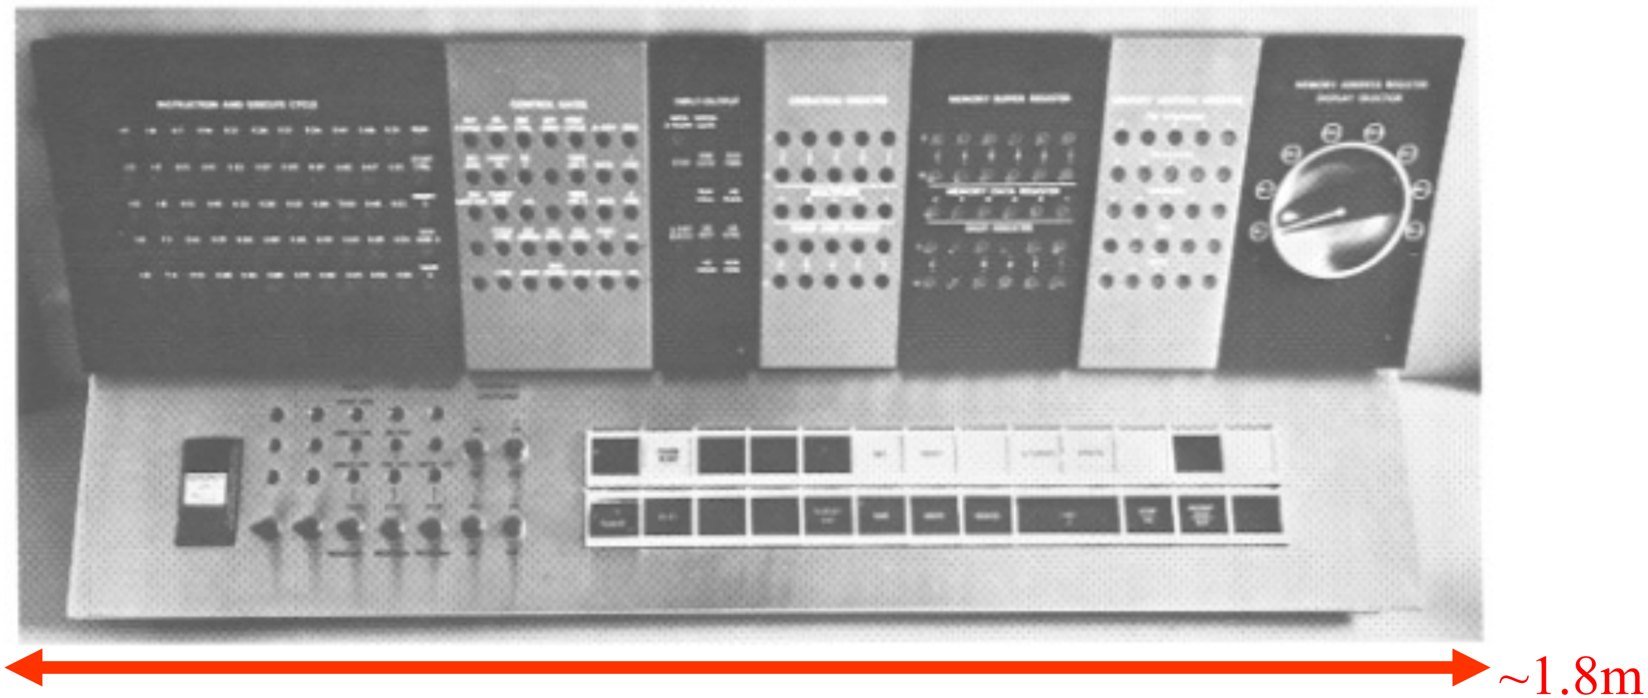

Fortran 2

```
C AREA OF A TRIANGLE WITH A STANDARD SQUARE ROOT FUNCTION
C INPUT - TAPE READER UNIT 5, INTEGER INPUT
C OUTPUT - LINE PRINTER UNIT 6, REAL OUTPUT
C INPUT ERROR DISPLAY ERROR OUTPUT CODE 1 IN JOB CONTROL LISTING
      READ INPUT TAPE 5, 501, IA, IB, IC
  501 FORMAT (3I5)
C IA, IB, AND IC MAY NOT BE NEGATIVE
C FURTHERMORE, THE SUM OF TWO SIDES OF A TRIANGLE
C MUST BE GREATER THAN THE THIRD SIDE, SO WE CHECK FOR THAT, TOO
      IF (IA) 777, 777, 701
  701 IF (IB) 777, 777, 702
  702 IF (IC) 777, 777, 703
  703 IF (IA+IB-IC) 777, 777, 704
  704 IF (IA+IC-IB) 777, 777, 705
  705 IF (IB+IC-IA) 777, 777, 799
  777 STOP 1
C USING HERON'S FORMULA WE CALCULATE THE
C AREA OF THE TRIANGLE
  799 S = FLOATF (IA + IB + IC) / 2.0
      AREA = SQRTF( S * (S - FLOATF(IA)) * (S - FLOATF(IB)) *
     +     (S - FLOATF(IC)))
      WRITE OUTPUT TAPE 6, 601, IA, IB, IC, AREA
  601 FORMAT (4H A= ,I5,5H  B= ,I5,5H  C= ,I5,8H  AREA= ,F10.2,
     +        13H SQUARE UNITS)
      STOP
      END
```

A personal history of computer technology

+ Magnetic Core Memory (4096 bytes)

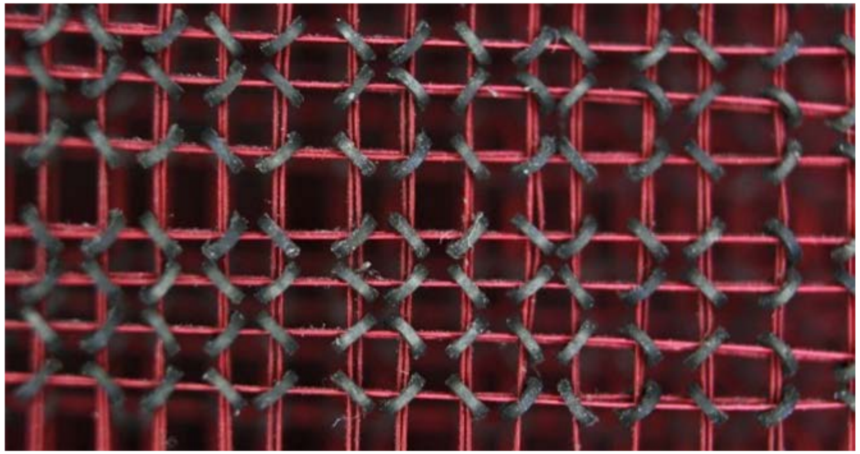

A personal history of computer technology

- Punched Card Input? What IS that?

A personal history of computer technology

- Punched Card Input? What IS that?

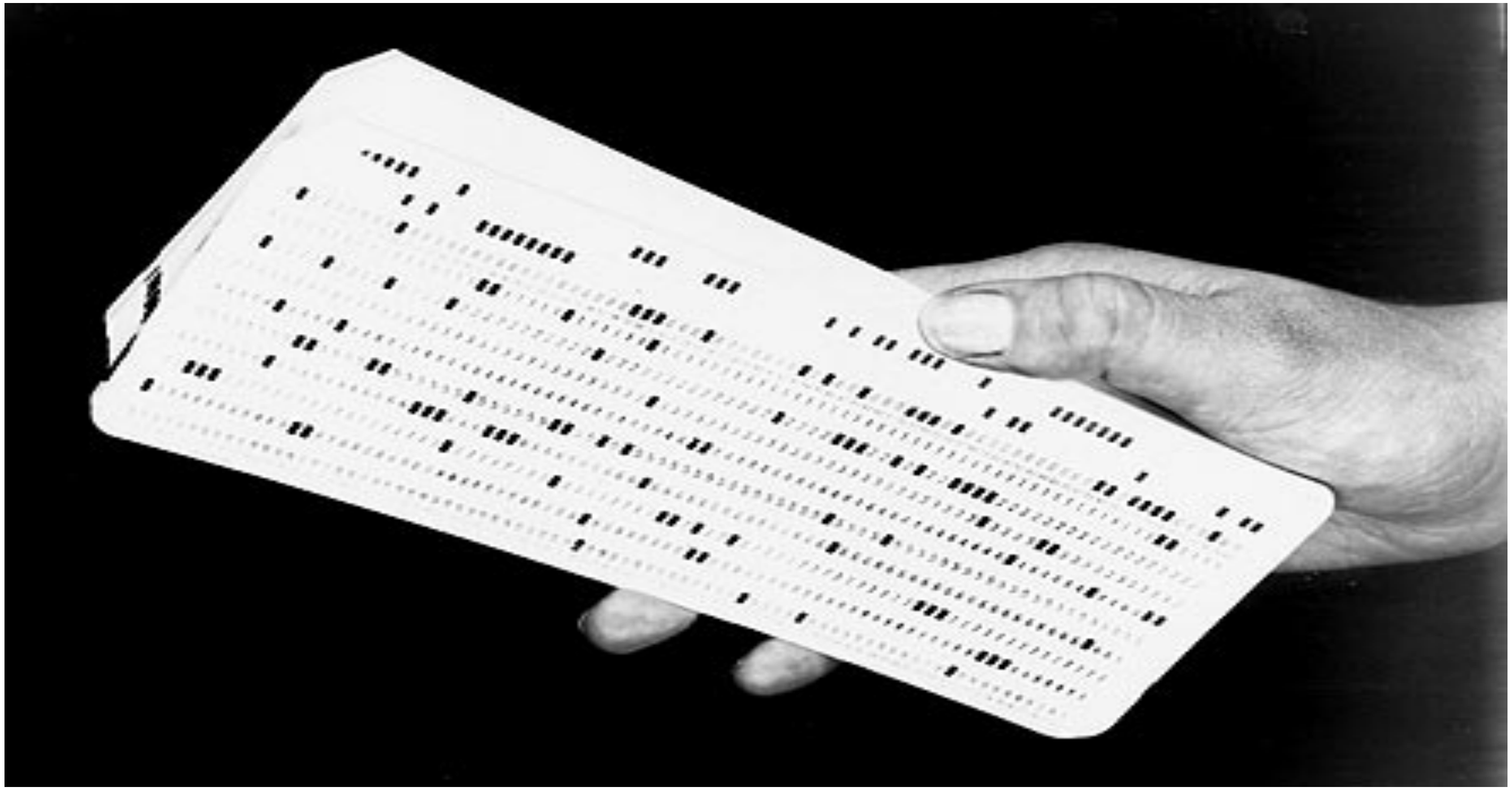

A personal history of computer technology

- Punched Card Input? What IS that?

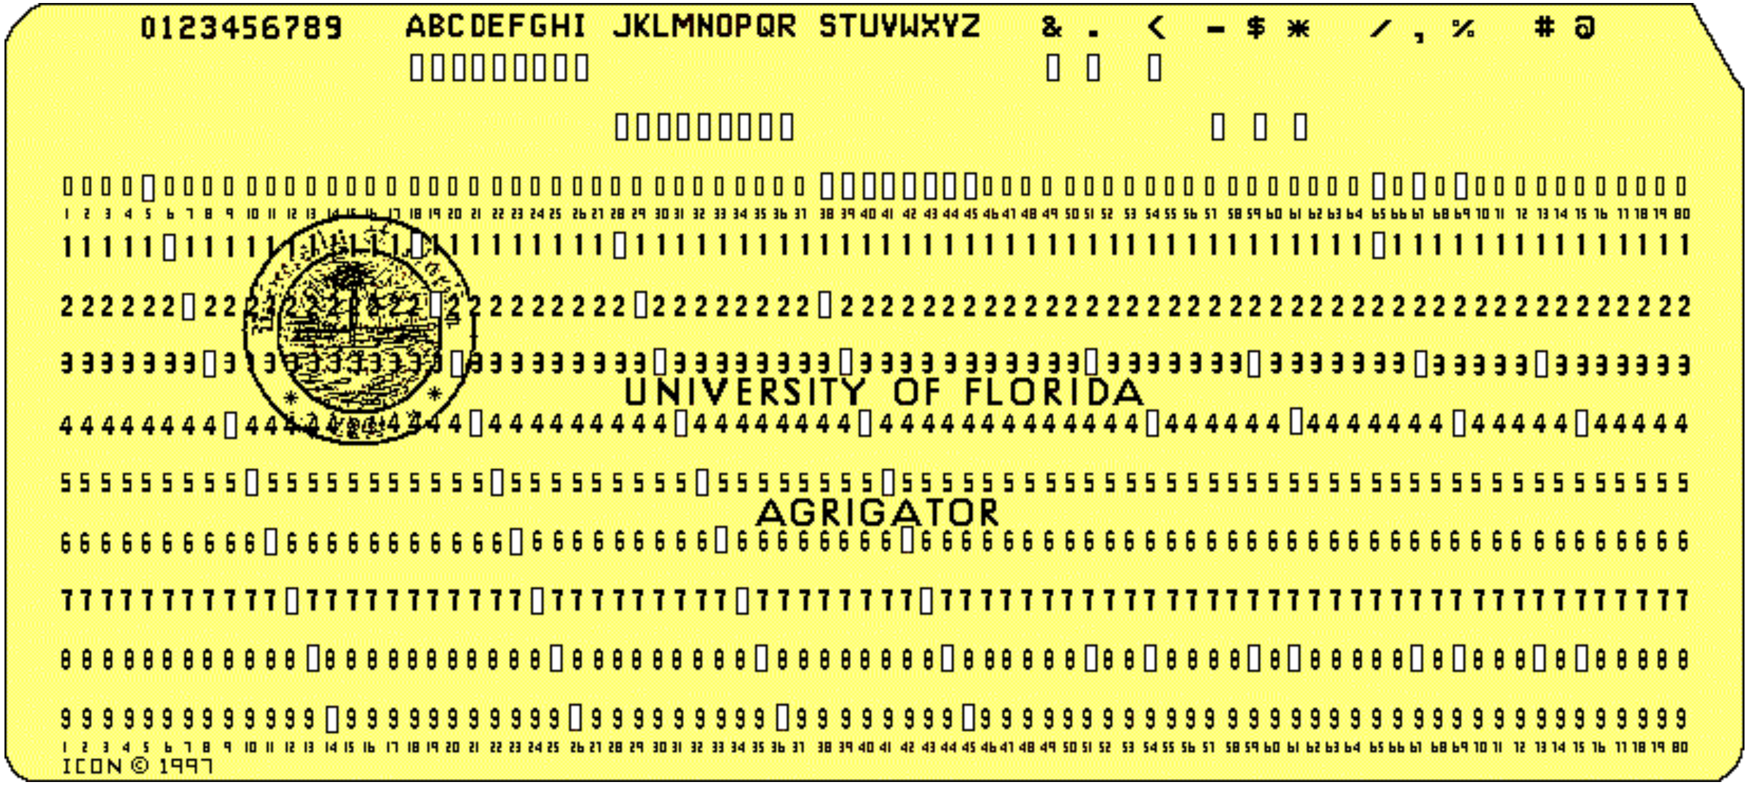

A personal history of computer technology

- Input with punch cards

A personal history of computer technology

+ Input with punch cards

Programs

A personal history of computer technology

- 1970's: 500 KHz Machines (IBM 360/370), 100-500 KB Memory, Video Terminals
- 1975 IBM 370
- 392 Kilobytes, 500 Khz
- Timesharing systems

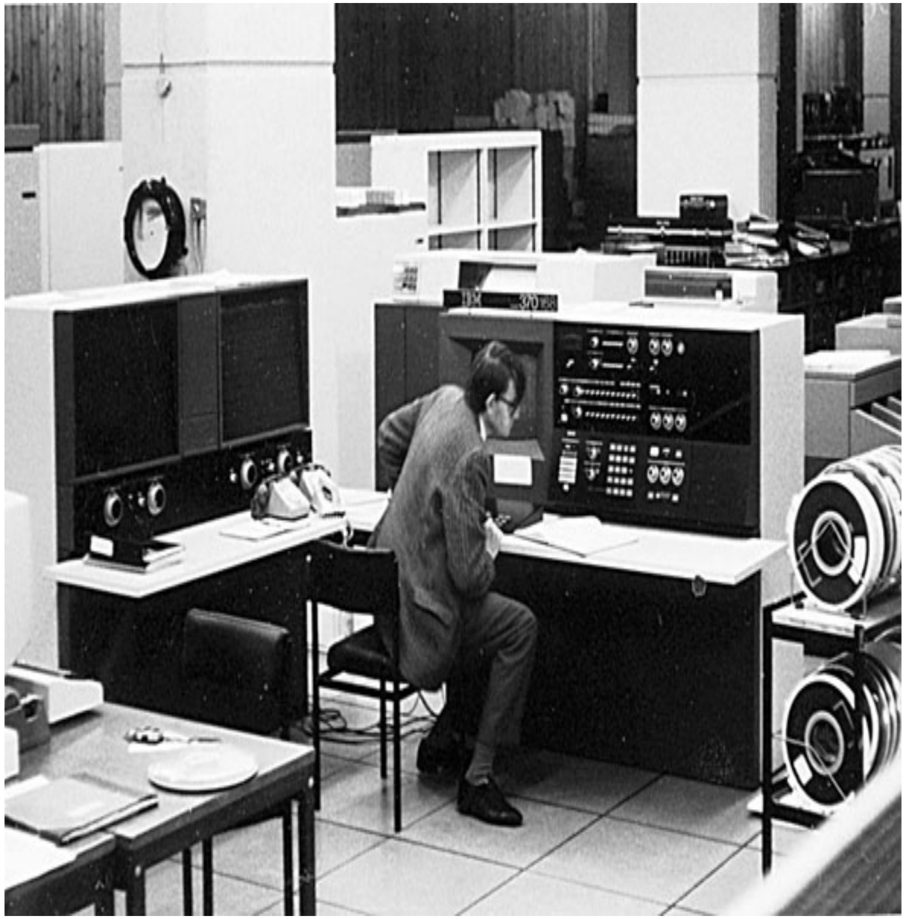

A personal history of computer technology

ARPANET circa 1973

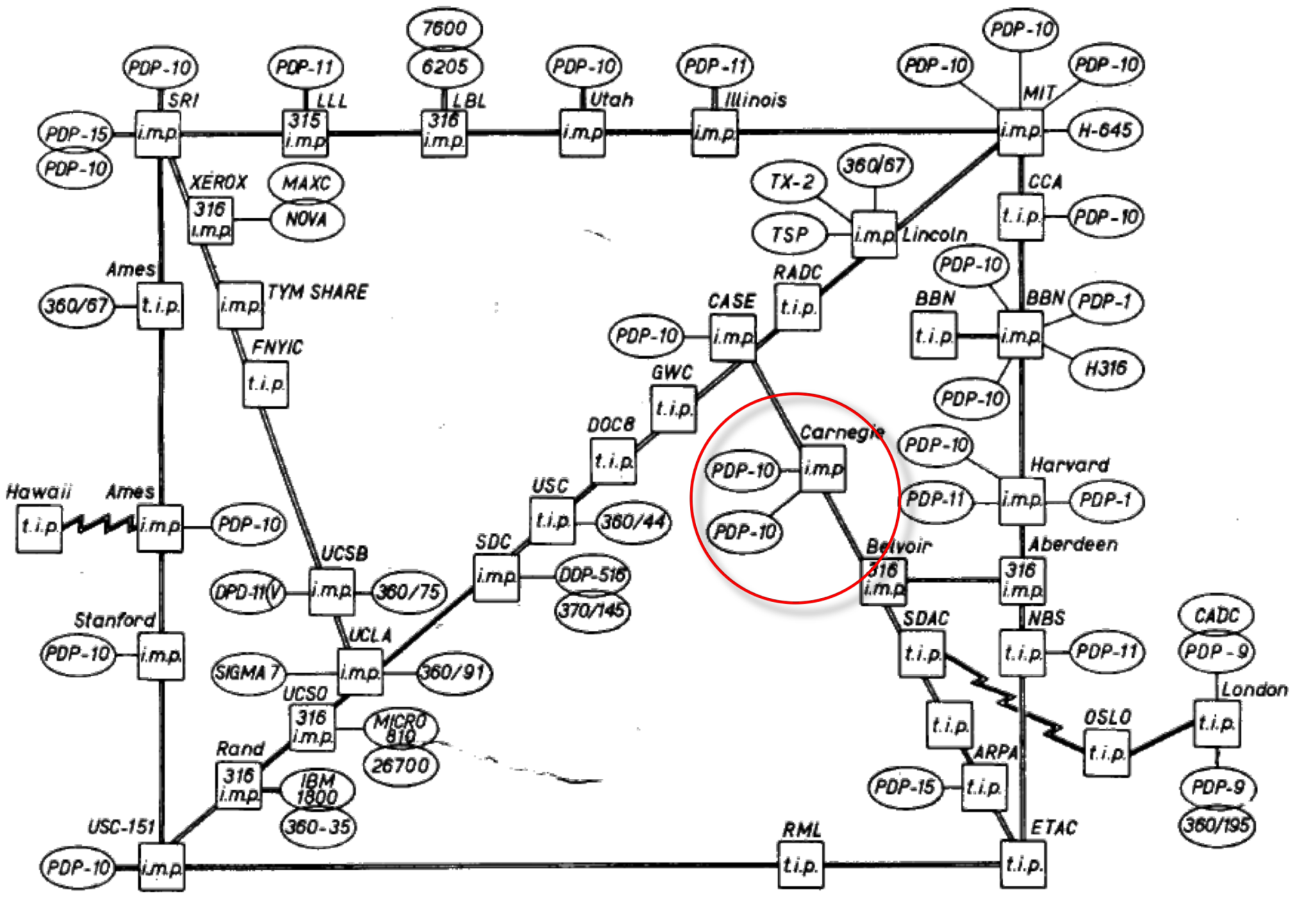

ARPANET circa 1977

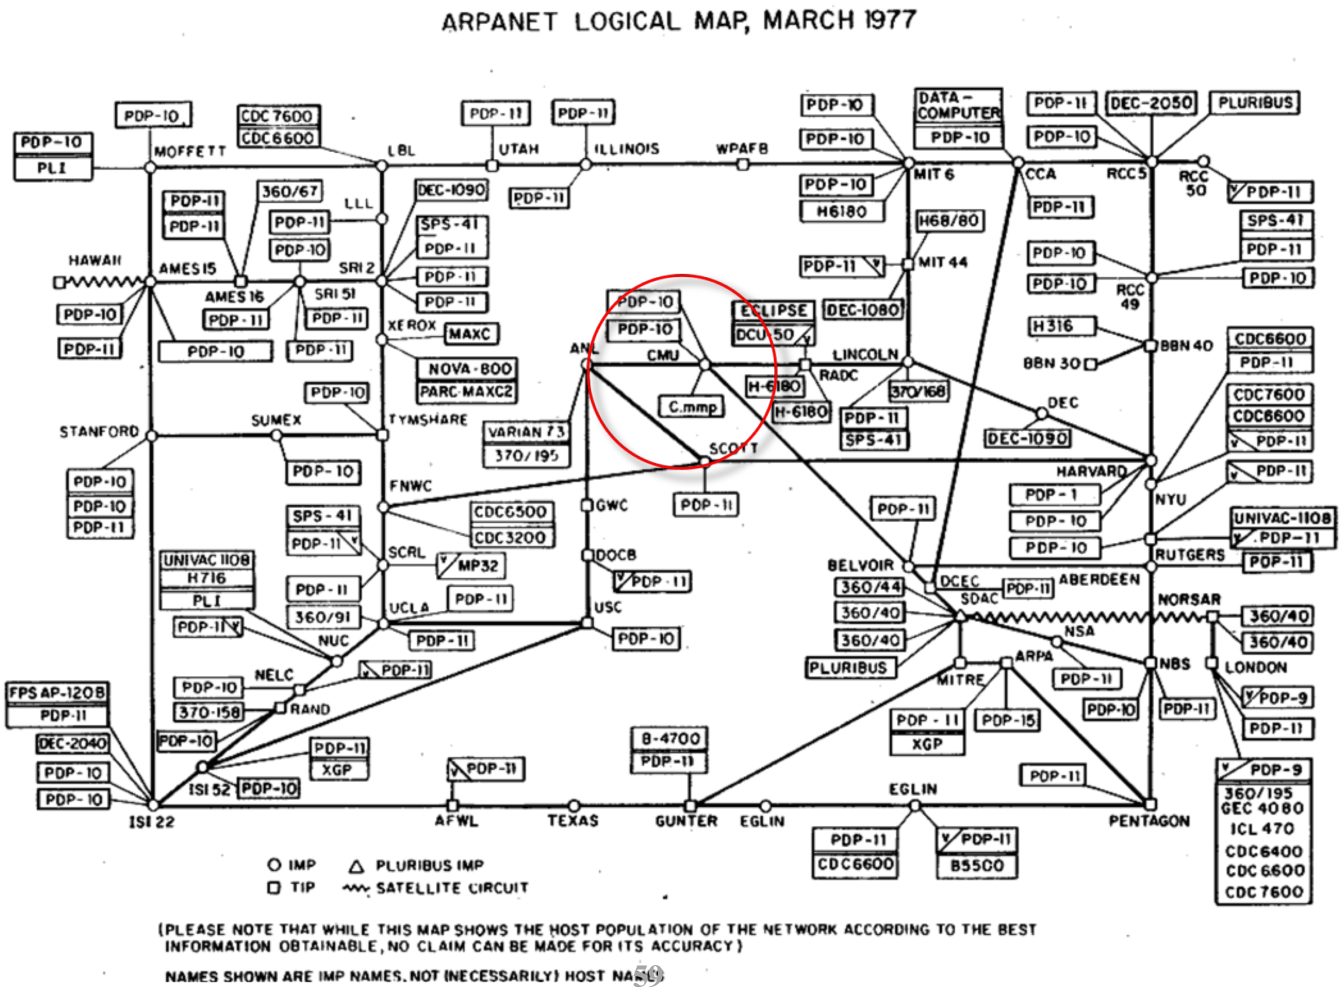

A personal history of computer technology

- Late 70's PDP-8
  - First "personal computer"

- **4 Kilobytes**, 200 Kilohertz, paper tape

DEC 10 at CMU late 1970's & early 1980's

A personal history of computer technology

- Early 1980's:
- 1Mhz Machines (Vax 780), 1-2 MB Memory, CRT Terminals (9600 Baud serial lines), ethernet
- 50-100 MB Disks,
- Early laser printers: Xerox Dover printer,
  - Printed 1 page a second,
  - Cost 300,000 dollars

A personal history of computer technology

- Mid 80's:
- Xerox Alto Workstation
  - The first real personal computer!
  - Steve Job's is famously rumored to have seen this and then built the Mac.
- 128KB RAM, 5.8 MHz Processor
- Interchangeable **2.5MB** personal disk
- Only 2000 units were ever produced.
  - Cost $10,000 1980 dollars to build = $31000 in 2019 dollars

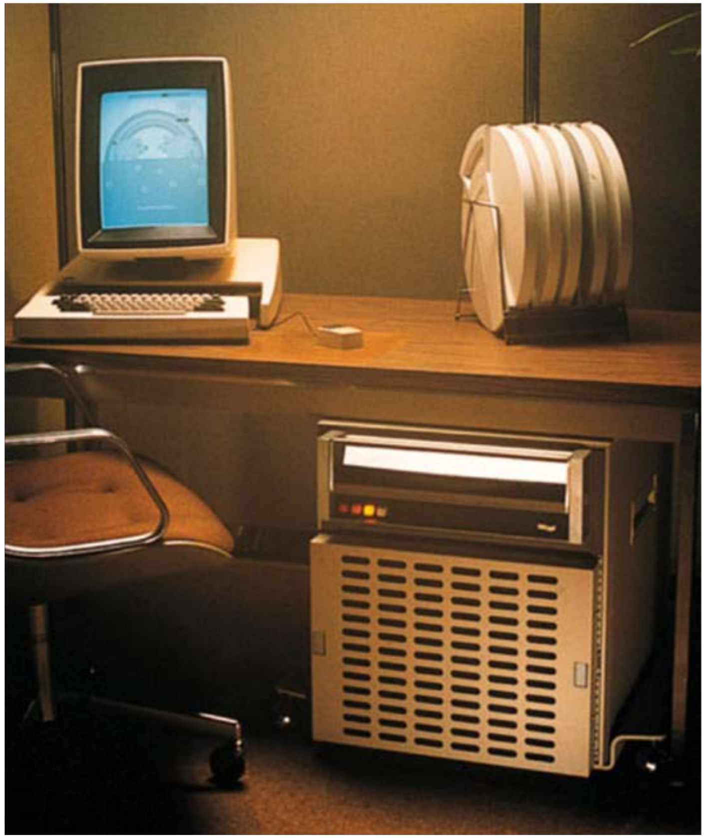

Screen shots then

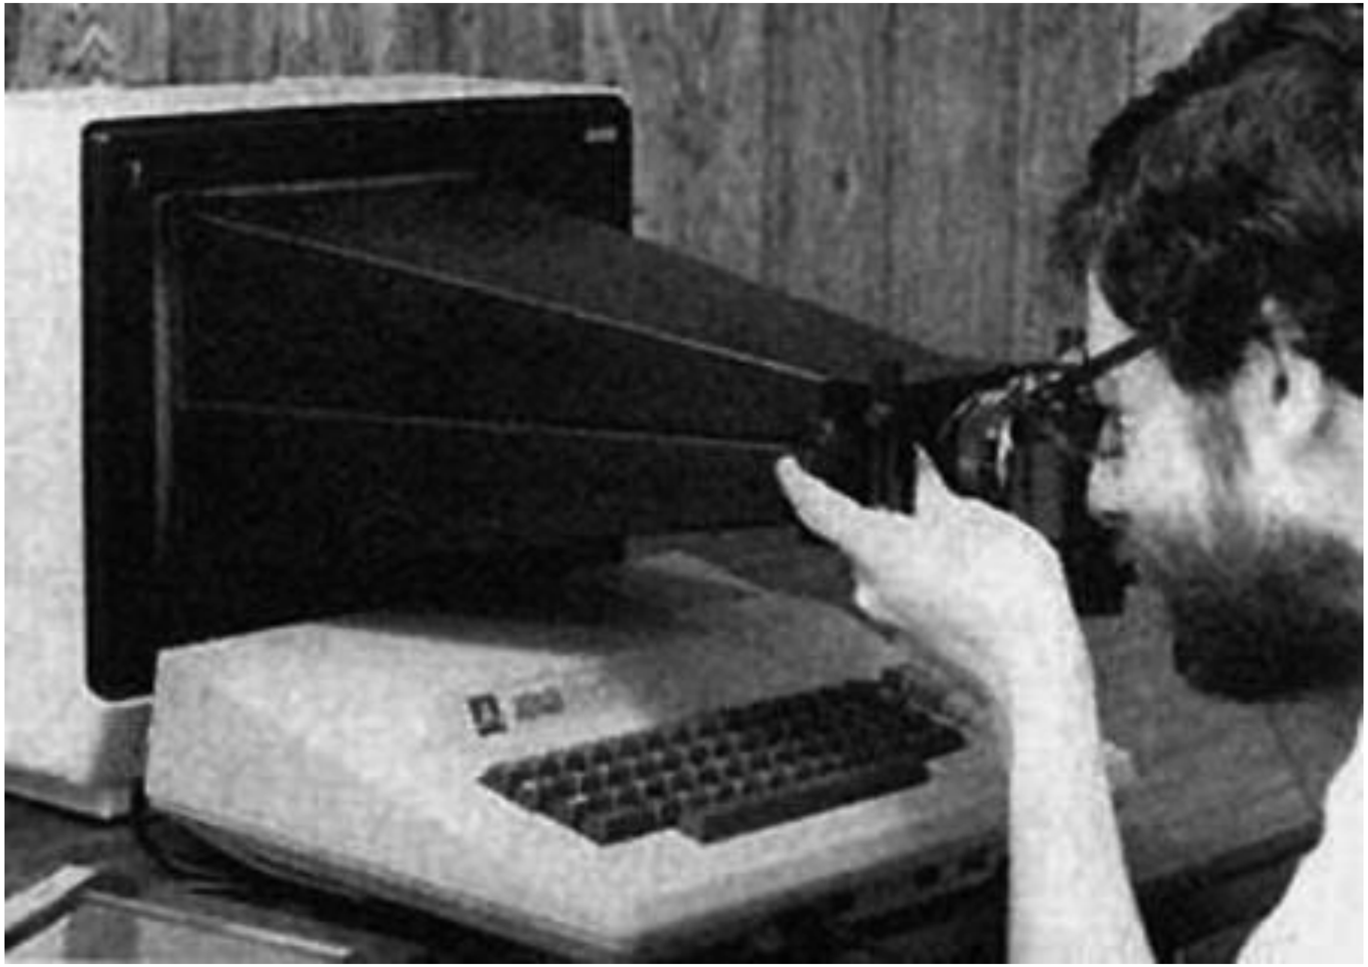

Mobile Phones

- Mobile phones are actually quite old.
- First proof of concept mobile phone was built and demoed in 1973
- Publicly available in 1984 – expensive!

A personal history of computer technology

- Late 80’s:
- Sun 1 - 4 Workstations
- 16 MB, 70 MegaHertz

A personal history of computer technology

- Late 80's MacIntosh SE
  - 4 MB RAM 40 MB Disk, 20 Mhz
  - 3000 (1988) dollars = 6531 (2019) dollars
  - My first personal computer
- Through 90's,
  - Sun SparcStations
    - 32 MB – 1 GB RAM, 32 – 450 MHz
  - Various Macintoshes/Mac Laptops
    - 4 MB - 200 MB Ram, 32 – 233 Mhz
  - Various PCs
    - 1 MB – 1GB Ram, 266 – 750 MHz

Apple Powerbook 170

- My first laptop
- **Released:** October 21, 1991 (Bought it in December 1991)
- **Price:** **$**4,599 (≈$8,367 in 2019)
- **CPU:** Motorola 68030, 25 MHz
- **Memory**: 40 MB disk, 4MB RAM
- Weighs several kilos, mechanical tracker
- Still functional – at least will boot

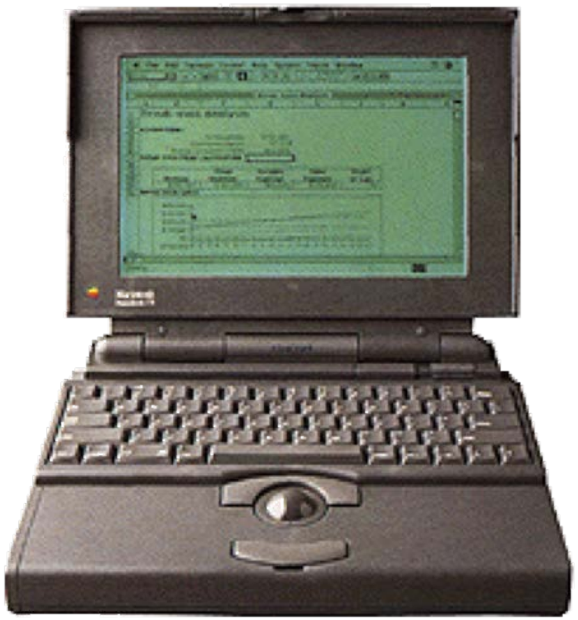

A personal history of computer technology

- 2019
  - Mac Pro, 32G GB Ram, 1TB of SSD Disk, 3 GHz 4-core CPU
  - Macbook Pro Laptops (16 G Ram), iPads, iPhone, etc.

- **You probably now have significantly more computing power in your cell phone, than all of the universities in the Middle East combined had, when I was your age!**

What is driving all this?

- Semiconductor Technology
- Paranoia (☺)
  - *Andrew S. Grove (Intel Chairman)*
    - Only the paranoid survive (1996)
- Theoretical breakthroughs
  - It turns out that these have provided the most significant improvement, but we rarely see it mentioned
  - If you discover a polynomial time algorithm for a problem not known to have one, your impact will be way more than any hardware improvement.

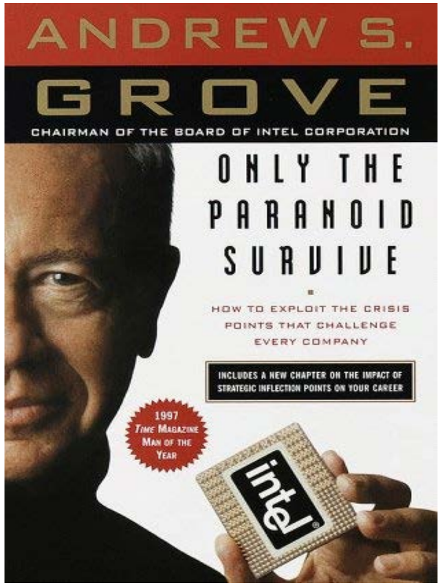

Very Large Scale Integrated Circuit Technology

- Intel 8086 (1979)
- 29,000 Transistors
- 5 Megahertz
- The first microprocessor I built a computer with

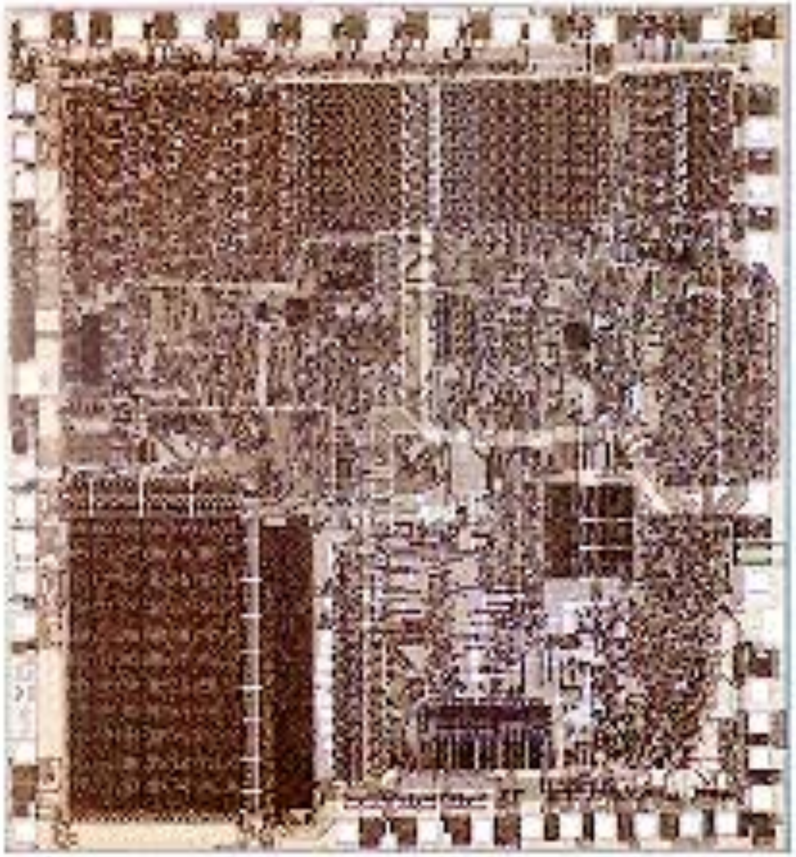

Very Large Scale Integrated Circuit Technology

- Intel 8086 (1979)
- 29,000 Transistors
- 5 Megahertz
- The first microprocessor I built a computer with

Microprocessor based modular database processors

Authors: Esen A. Ozkarahan
Kemal Oflazer

Published in:
· Proceeding
VLDB '78 Proceedings of the fourth international conference on Very Large Data Bases - Volume 4
Pages 300-311

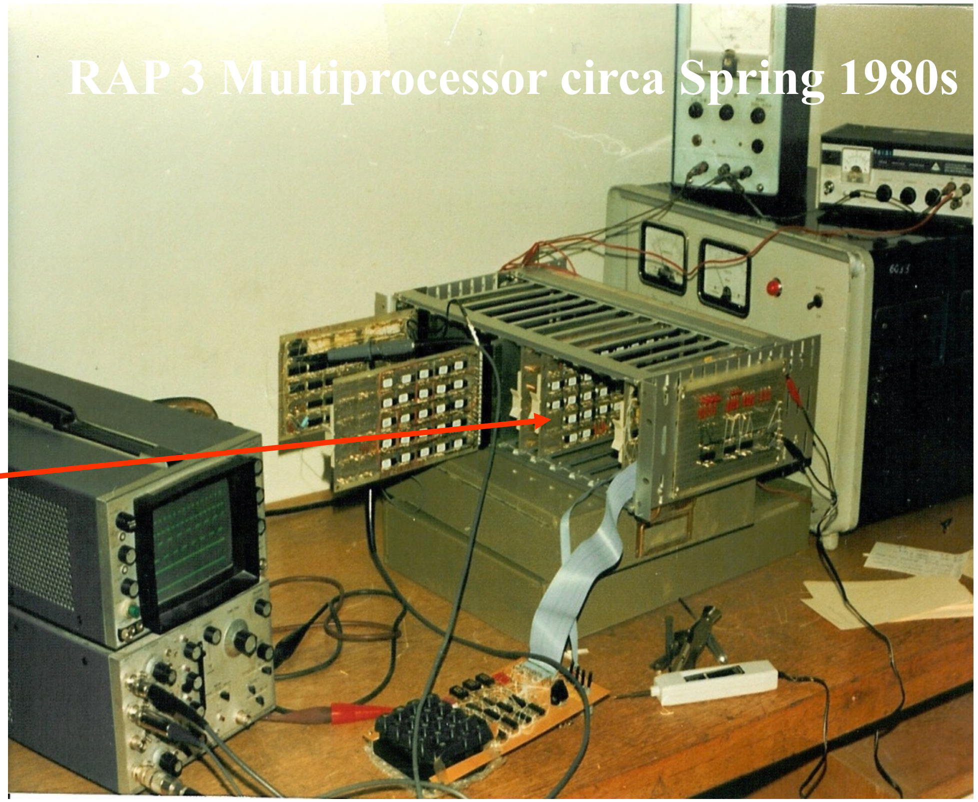

Very Large Scale Integrated Circuit Technology

- Intel Coffee Lake– (2017)
  - 4/12-cores + GPU
- ~6 billion transistors
  - 270,000 times that of 8086
- 3.7 Gigahertz
  - ~710 times that of 8086

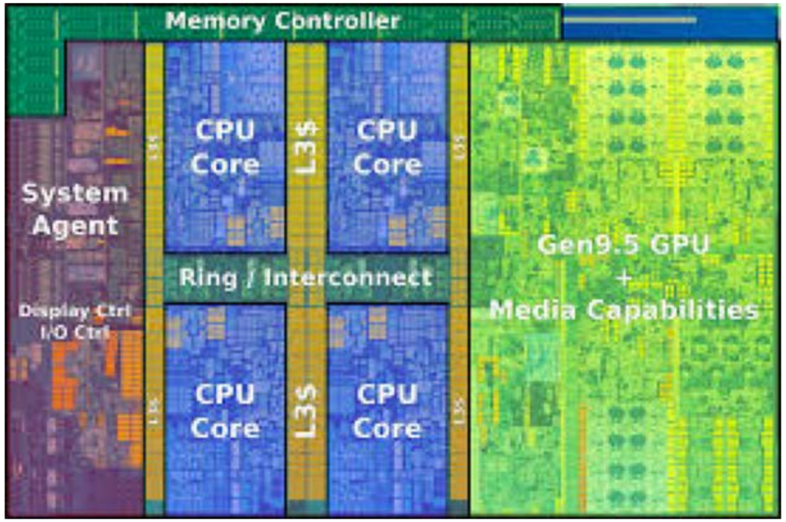

149 mm^2

Very Large Scale Integrated Circuit Technology

- Intel Sandy Bridge-E – (2011)
  - 8/16 cores
- 2.27 billion transistors
  - 90,000 times that of 8086
- 3.3 - 3.6 Gigahertz
  - ~700 times that of 8086

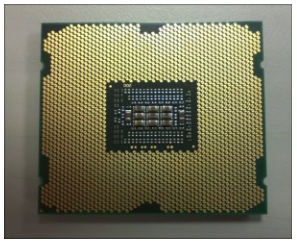

2011 pins

180 Watts

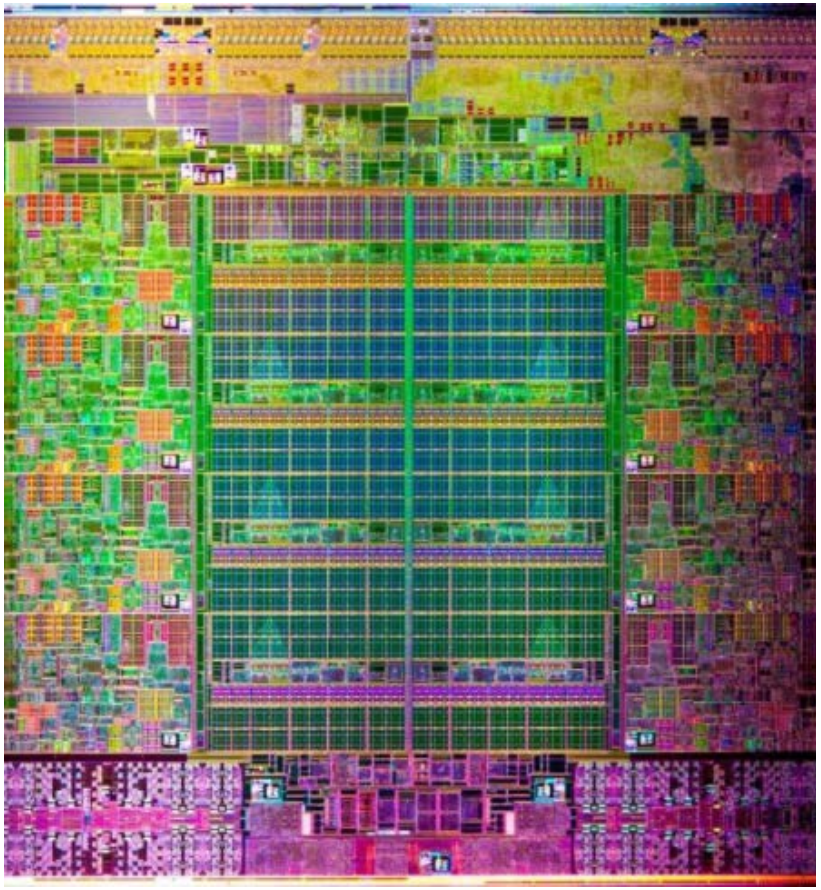

435 mm^2

Very Large Scale Integrated Circuit Technology

9th generation Core/Coffee Lake Refresh [edit]

Model	Price (USD)	Cores/Threads	Base frequency (GHz)	Max turbo frequency (GHz)	GPU	Maximum GPU clock rate (MHz)	L3 cache (MB)	TDP (W)	Socket	Release
i9-9900K	$529	8/16	3.6	5.0	UHD 630	1200	16	95	LGA 1151	Q4 2018
i7-9700K	$399	8/8	3.7	4.9	UHD 630	1200	12	95	LGA 1151	Q4 2018
i5-9600K	$229	6/6	3.7	4.6	UHD 630	1150	9	95	LGA 1151	Q4 2018
i5-9400F	$149	6/6	2.9	4.1	N/A	N/A	9	65	LGA 1151	Q1 2019
i3-9350KF	$172	4/4	4.0	4.6	N/A	N/A	8	91	LGA 1151	Q1 2019

Very Large Scale Integrated Circuit Technology

8th generation Core/Coffee Lake/Kaby Lake Refresh/Whiskey Lake [edit]

Desktop [edit]

Model	Price (USD)	Cores/Threads	Base frequency (GHz)	Max turbo frequency (GHz)	GPU	Maximum GPU clock rate (MHz)	L3 cache (MB)	TDP (W)	Socket	Release
i7-8086K	$425	6/12	4.0	5.0	UHD 630	1200	12	95	LGA 1151	Q2 2018
i7-8700K	$359	6/12	3.7	4.7	UHD 630	1200	12	95	LGA 1151	Q4 2017
i7-8700	$303	6/12	3.2	4.6	UHD 630	1200	12	65	LGA 1151	Q4 2017
i5-8600K	$257	6/6	3.6	4.3	UHD 630	1150	9	95	LGA 1151	Q4 2017
i5-8500	$202	6/6	3.0	4.1	UHD 630	1100	9	65	LGA 1151	Q2 2018
i5-8400	$182	6/6	2.8	4.0	UHD 630	1050	9	65	LGA 1151	Q4 2017
i3-8350K	$168	4/4	4.0	N/A	UHD 630	1150	8	91	LGA 1151	Q4 2017
i3-8100	$117	4/4	3.6	N/A	UHD 630	1100	6	65	LGA 1151	Q4 2017

Source: https://en.wikipedia.org/wiki/List_of_Intel_microprocessors

Very Large Scale Integrated Circuit Technology

8th generation Core/Coffee Lake/Kaby Lake Refresh/Whiskey Lake [edit]

Desktop [edit]

Model	Price (USD)	Cores/Threads	Base frequency (GHz)	Max turbo frequency (GHz)	GPU	Maximum GPU clock rate (MHz)	L3 cache (MB)	TDP (W)	Socket	Release
i7-8086K	$425	6/12	4.0	5.0	UHD 630	1200	12	95	LGA 1151	Q2 2018
i7-8700K	$359	6/12	3.7	4.7	UHD 630	1200	12	95	LGA 1151	Q4 2017
i7-8700	$303	6/12	3.2	4.6	UHD 630	1200	12	65	LGA 1151	Q4 2017
i5-8600K	$257	6/6	3.6	4.3	UHD 630	1150	9	95	LGA 1151	Q4 2017
i5-8500	$202	6/6	3.0	4.1	UHD 630	1100	9	65	LGA 1151	Q2 2018
i5-8400	$182	6/6	2.8	4.0	UHD 630	1050	9	65	LGA 1151	Q4 2017
i3-8350K	$168	4/4	4.0	N/A	UHD 630	1150	8	91	LGA 1151	Q4 2017
i3-8100	$117	4/4	3.6	N/A	UHD 630	1100	6	65	LGA 1151	Q4 2017

The processors do not seem to be getting any faster!

Then this happens!

2019 – 1.2 Trillion Transistor Deep Learning Processor

- 462.25 cm^2 on a "wafer"
  - Compare with 435 mm^2 earlier
  - ~100 times larger
- 56 times larger than the largest GPU today
- Power consumption is not publicized
- But apparently water-cooled

Very Large Scale Integrated Circuit Technology

- CMU in the 1980's was a hotbed of VLSI chip design research
- Undergraduate and Graduate Courses on VLSI chip design
  - I actually TA'ed a VLSI design course in 1982
- Design – Send by Email – Get 100 copies of your chip in the mail - Test

Carnegie Mellon University
Research Showcase

Computer Science Department — School of Computer Science

1-1-1982

Design and implementation of a single-chip 1-d median filter

Kemal Oflazer
Carnegie Mellon University

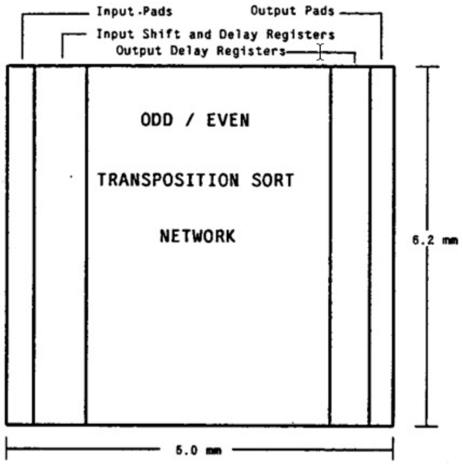

Figure 4-1: The floor plan of the chip

Very Large Scale Integrated Circuit Technology

- Moore's "Law"
  - Transistors double every 18-24 months

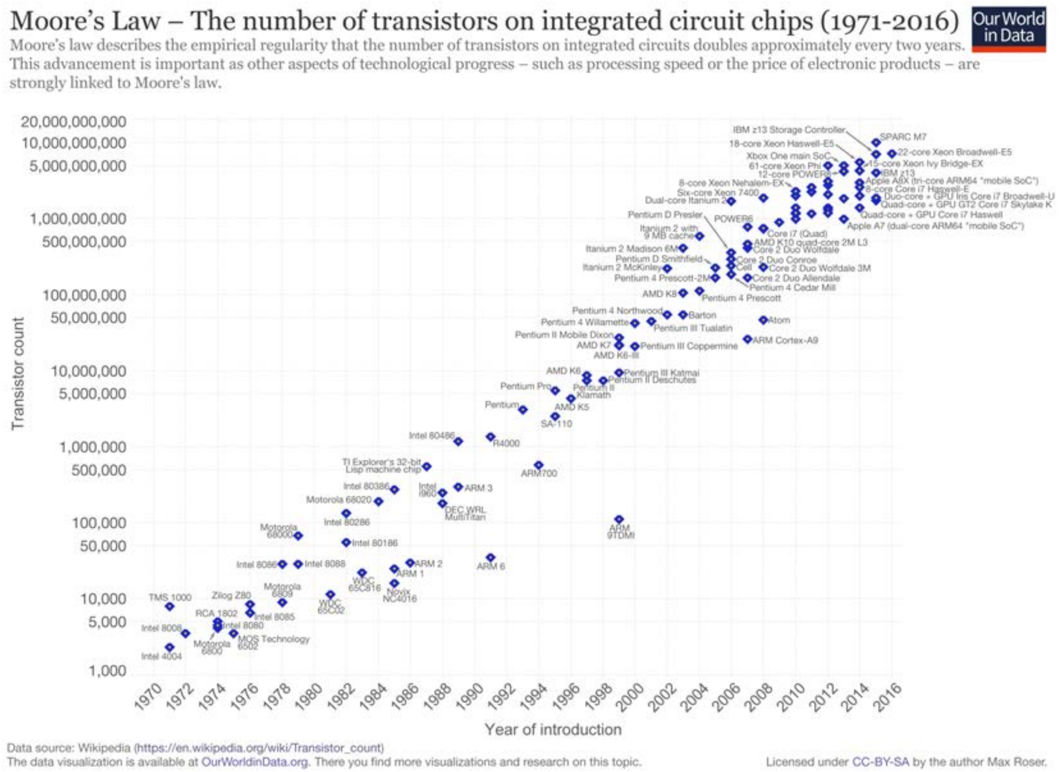

Moore's Law

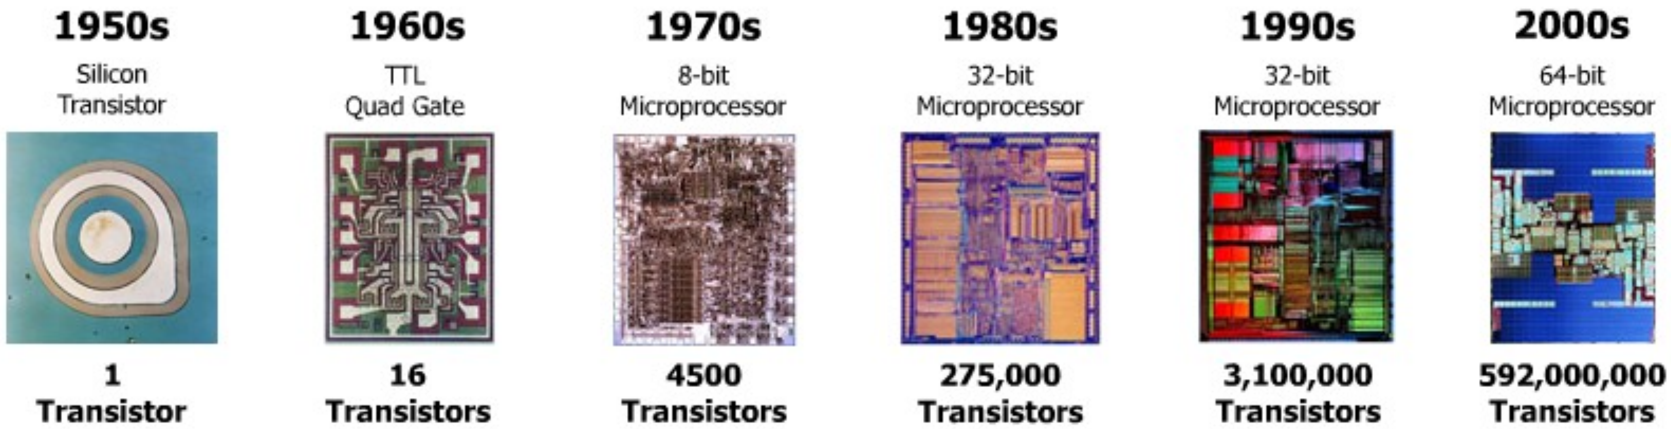
1950s
Silicon Transistor
1 Transistor
1960s
TTL Quad Gate
16 Transistors
1970s
8-bit Microprocessor
4500 Transistors
1980s
32-bit Microprocessor
275,000 Transistors
1990s
32-bit Microprocessor
3,100,000 Transistors
2000s
64-bit Microprocessor
592,000,000 Transistors

Other trends

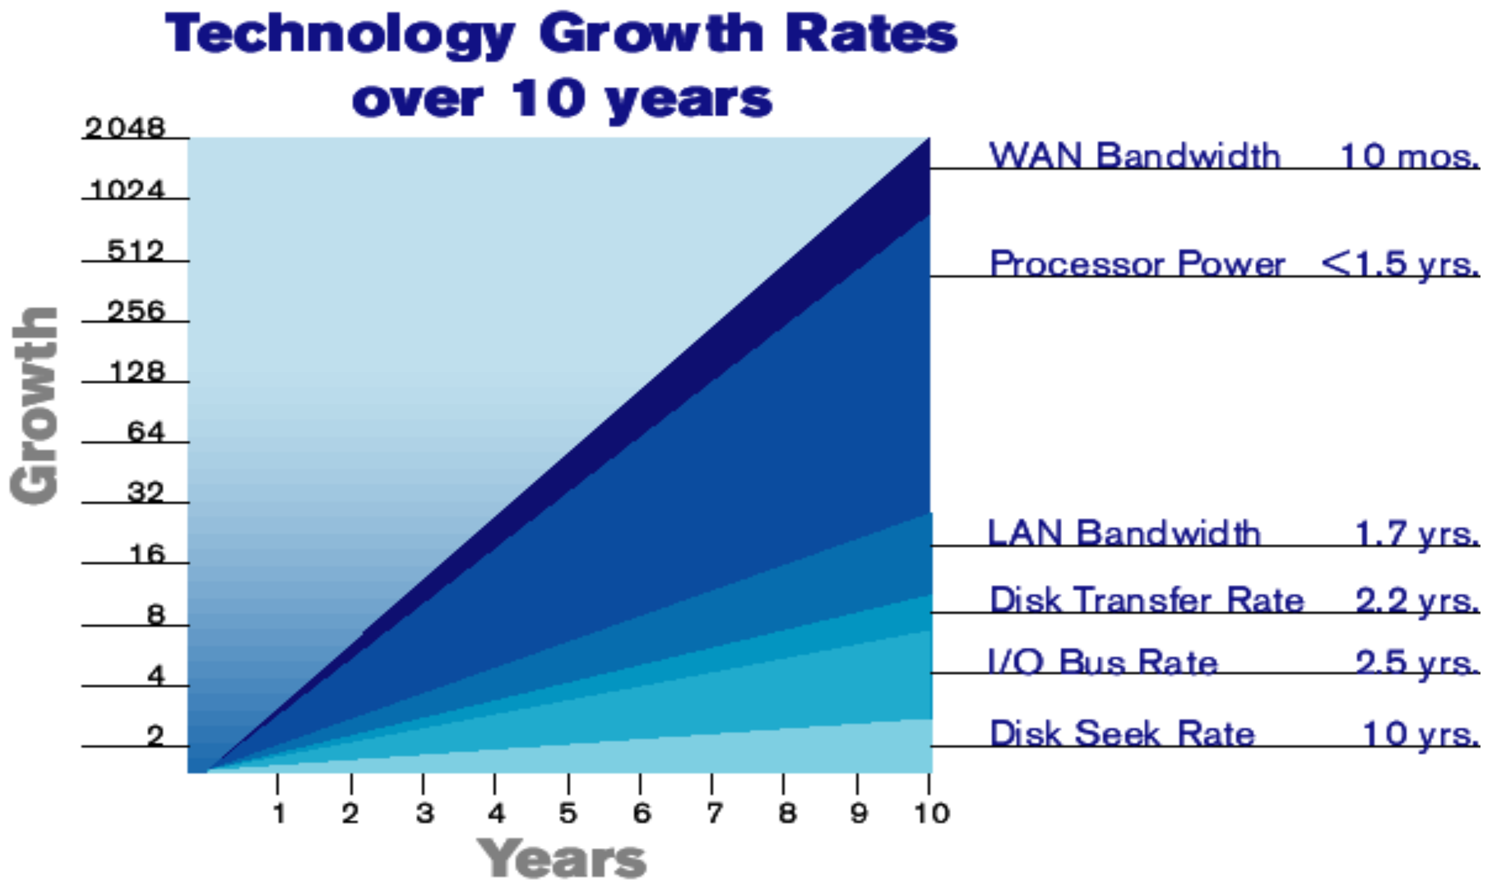

Note that the Y axis is logarithmic, so these are exponential growths

Exponential Growth

- It is good when it applies technological advances.
- Not so good when it applies to problem complexity.
  - There are legitimate problems that will take ~10^{1000} centuries to solve with the fastest computers and algorithms available.
    - Universe has been around for about $14 * 10^7$ centuries!
  - Wait for Prof. Kapoutsis to explain you that!

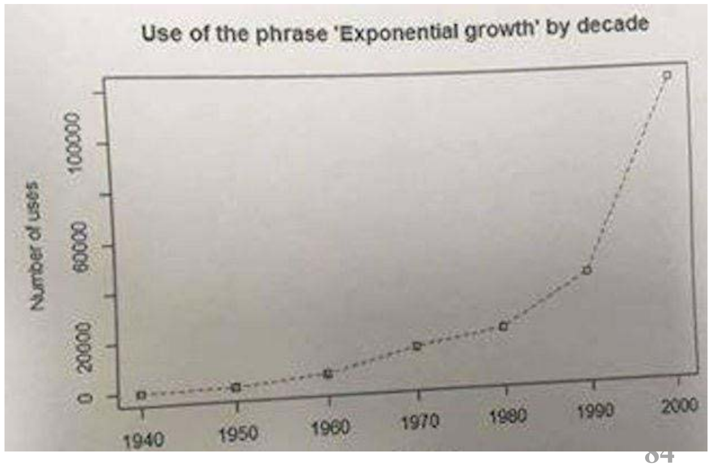

Where are we heading?

- **“Prediction is hard, especially about the future.”**
  Niels Bohr

Where we are heading?

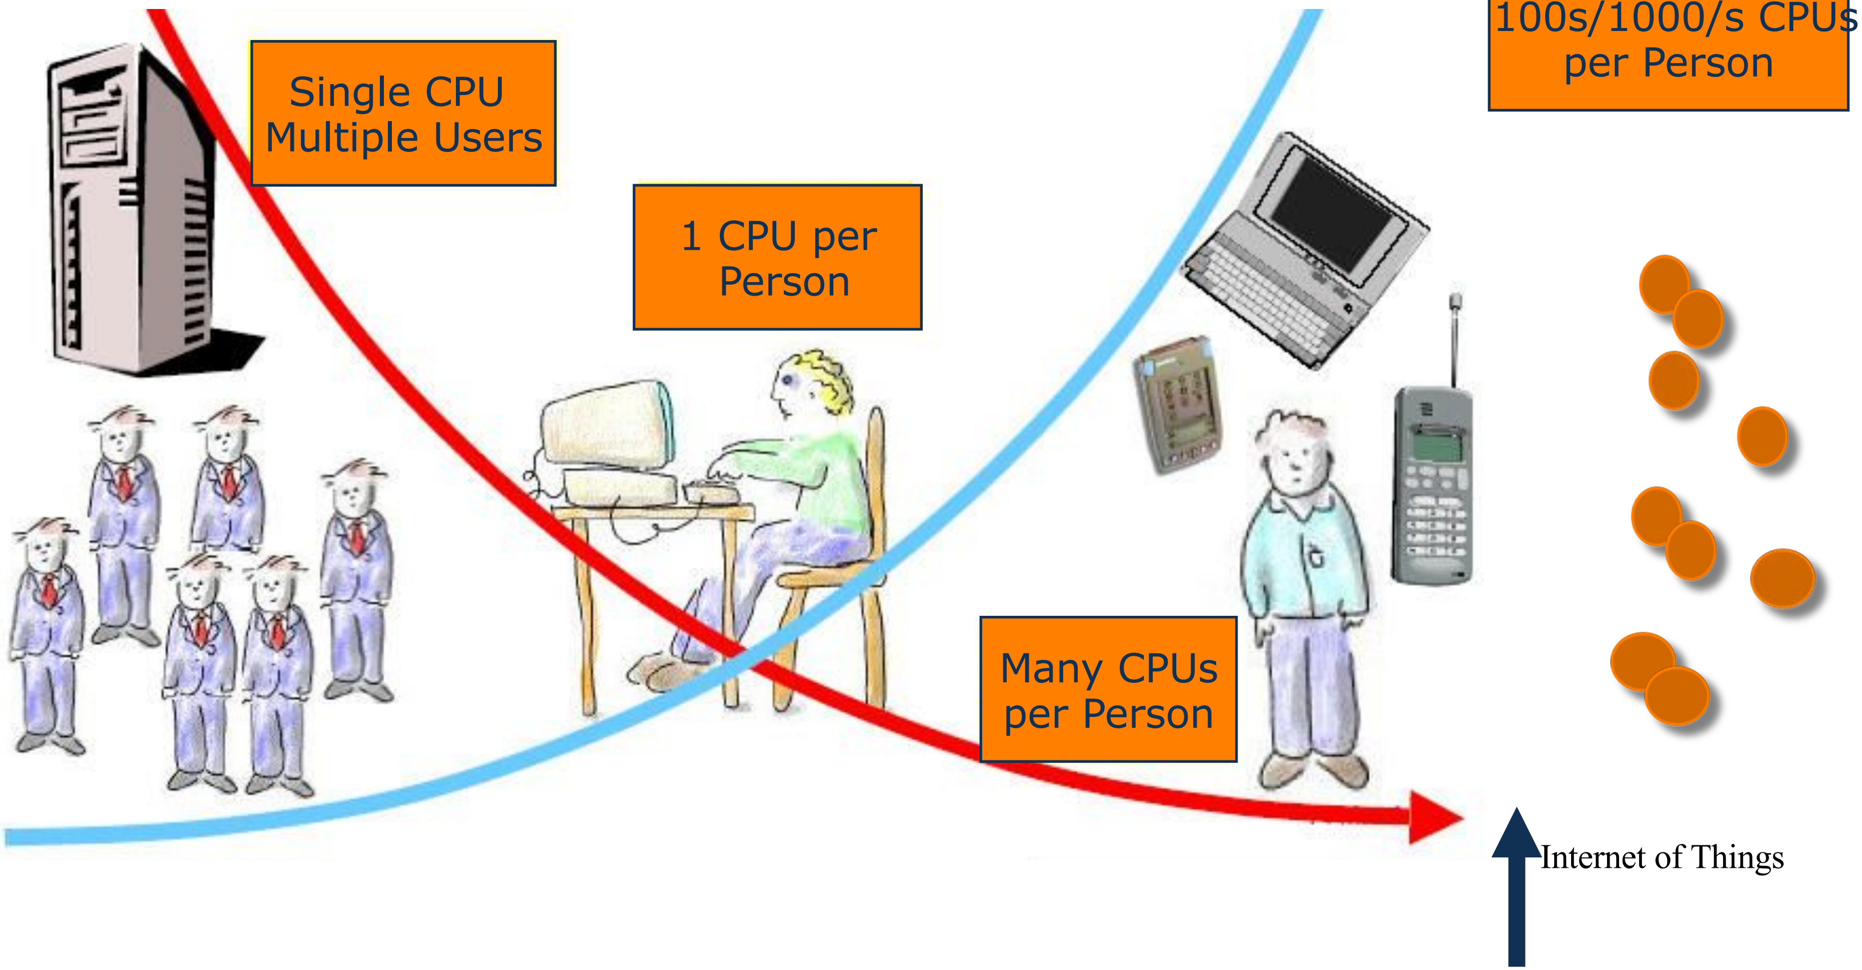

Broadband Speed

+ World Average is about 7.1 Mbit/s (2017)

Ranking position	Country	Mean Download Speed
1	Singapore	60.39
2	Sweden	46.00
3	Denmark	43.99
4	Norway	40.12
5	Romania	38.60
6	Belgium	36.71
7	Netherlands	35.95
8	Luxembourg	35.14
9	Hungary	34.01
10	Jersey	30.90
11	Switzerland	29.92
12	Japan	28.94
13	Latvia	28.63
14	Taiwan	28.09
15	Estonia	27.91
16	Spain	27.19
17	Republic of Lithuania	27.17
18	Andorra	27.14
19	Hong Kong	26.45
20	United States	25.86
21	Slovakia	25.30
22	Madagascar	24.87
23	France	24.23
24	Finland	24.00
25	Germany	24.00

Typical home internet in SFO in 2019

Wireless

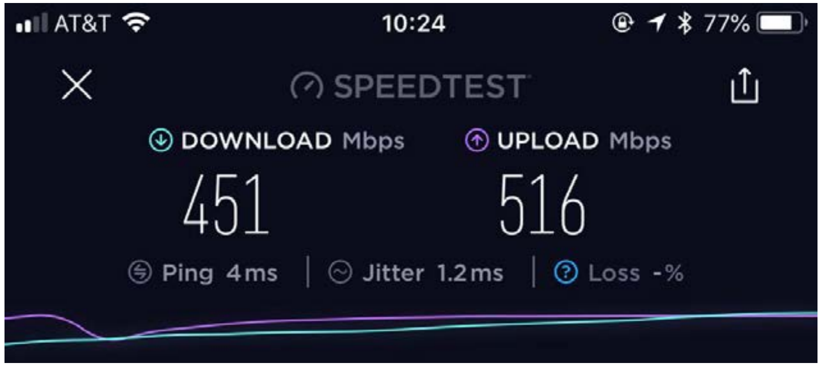

Wired Ethernet

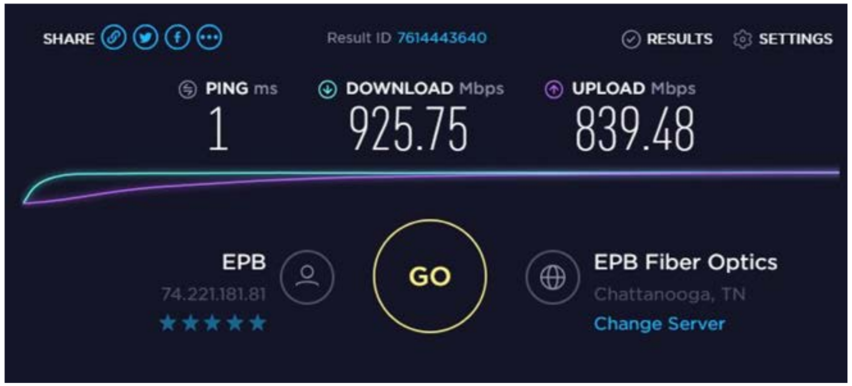

Happening Now

- Late 2010s:
  - Cars, appliances, telephones, everything have lots computing power somewhere
    - Car service is really SW upgrade mostly
  - Essentially infinite bandwidth (and wireless too),
  - Speech and visual interfaces, coming
    - Siri, Kinect, Google Now, Alexa
    - Dictation is widely available on mobile platforms
  - Wearable computers
    - Smart Watches
    - Virtual Reality
    - Google Glass (?)
  - Trackers
    - Fitbit, Apple Watch, Smart Implants (e.g. Cont. Glucose Monitors)
  - Deep Learning/Machine Learning/AI Tsunami

Wearable Computers

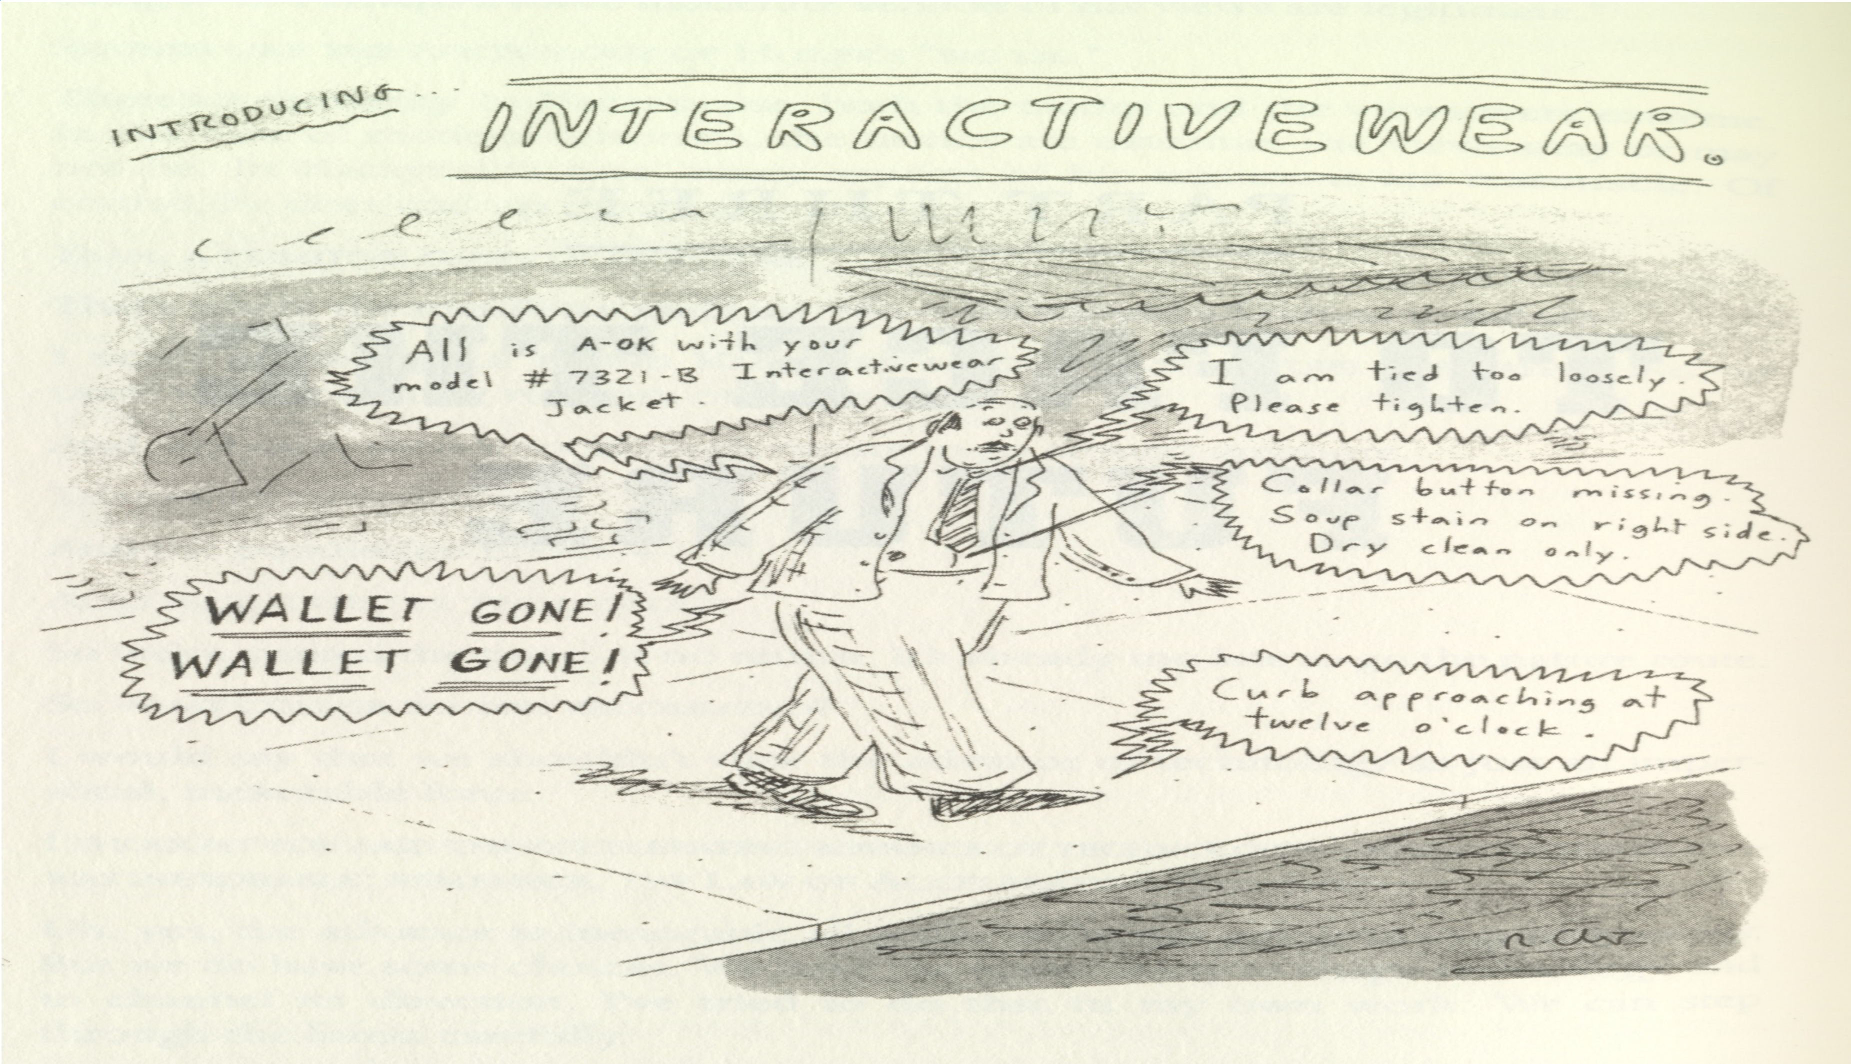

Computational Journalism!

Computer writing news stories

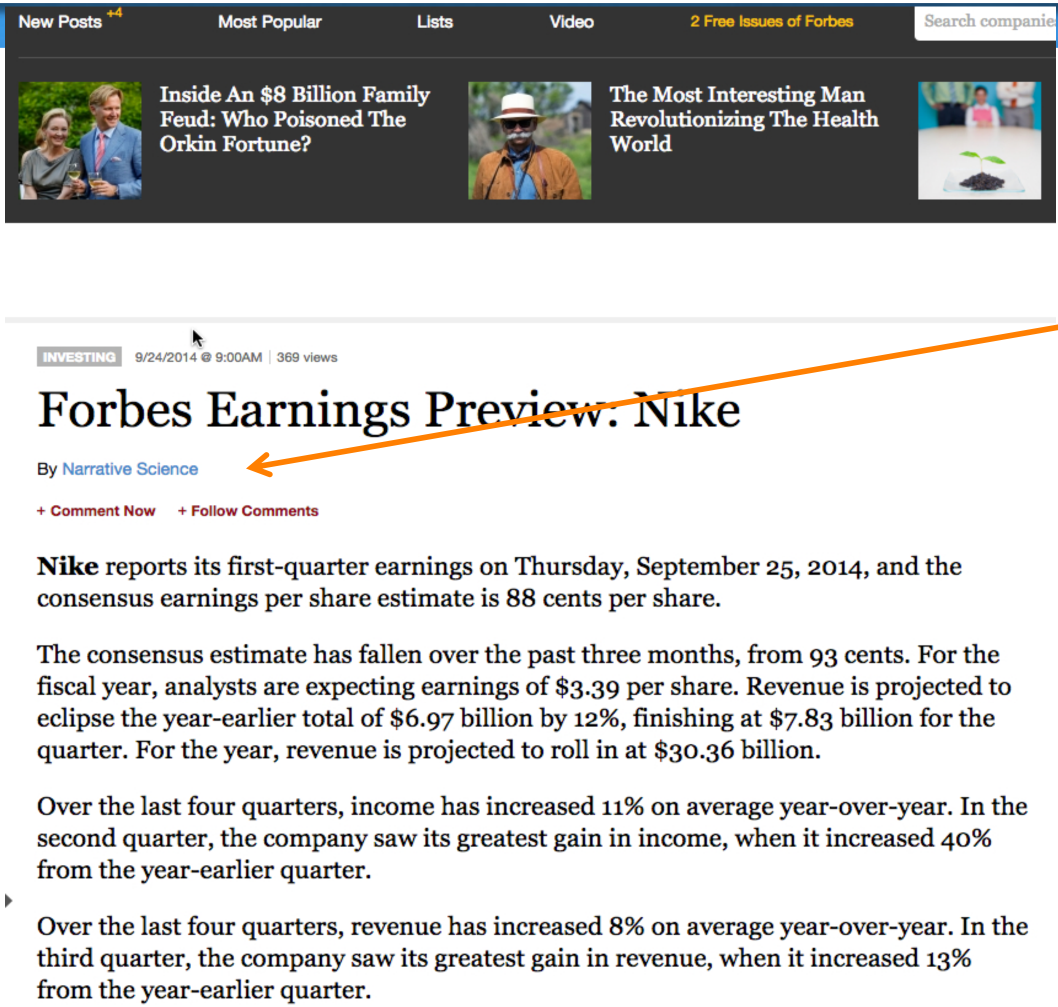

New Posts +4 | Most Popular | Lists | Video | 2 Free Issues of Forbes | Search companie

Inside An $8 Billion Family Feud: Who Poisoned The Orkin Fortune?

The Most Interesting Man Revolutionizing The Health World

INVESTING 9/24/2014 @ 9:00AM | 369 views

Forbes Earnings Preview: Nike

By Narrative Science

+ Comment Now + Follow Comments

Nike reports its first-quarter earnings on Thursday, September 25, 2014, and the consensus earnings per share estimate is 88 cents per share.

The consensus estimate has fallen over the past three months, from 93 cents. For the fiscal year, analysts are expecting earnings of $3.39 per share. Revenue is projected to eclipse the year-earlier total of $6.97 billion by 12%, finishing at $7.83 billion for the quarter. For the year, revenue is projected to roll in at $30.36 billion.

Over the last four quarters, income has increased 11% on average year-over-year. In the second quarter, the company saw its greatest gain in income, when it increased 40% from the year-earlier quarter.

Over the last four quarters, revenue has increased 8% on average year-over-year. In the third quarter, the company saw its greatest gain in revenue, when it increased 13% from the year-earlier quarter.

Then this happened

On Extractive and Abstractive Neural Document Summarization with Transformer Language Models

Sandeep Subramanian[1,2,3,*], **Raymond Li**[1,*], **Jonathan Pilault**[1,2,4,*], **Christopher Pal**[1,2,4,5]
[1]Element AI, [2]Montréal Institute for Learning Algorithms, [3]Université de Montréal,
[4]École Polytechnique de Montréal, [5]Canada CIFAR AI Chair
[1]{jonathan.pilault}@elementai.com

Abstract

We present a method to produce abstractive summaries of long documents that exceed several thousand words via neural abstractive summarization. We perform a simple extractive step before generating a summary, which is then used to condition the transformer language model on relevant information before being tasked with generating a summary. We show that this extractive step significantly improves summarization results. We also show that this approach produces more abstractive summaries compared to prior work that employs a copy mechanism while still achieving higher rouge scores. *Note: The abstract above was not written by the authors, it was generated by one of the models presented in this paper.*

Introduction

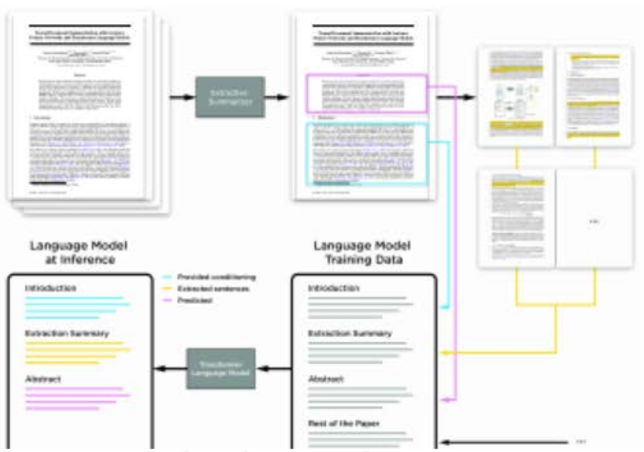

scores. *Note: The abstract above was not written by the authors, it was generated by one of the models presented in this paper.*

Deep Learning

- Starts in 1960's
  - Fades after computational limits are proved.
- Popular in 1980's and early 1990's
  - Fades again due to scalability problems
- Comes back with a vengeance in 2010's when
  - Cheap computing power
  - Specialized hardware (GPU, TPU)
  - Lots of data

  are available.
- You really need to be deeply proficient in **multi-variate calculus, probability and statistics** to understand how to build one of these and understand how they work.

Machine Learning/Deep Learning/AI

- Radically and quickly changing everything
  - Medicine,
  - Security/Surveillance
  - Language Translation
  - Image Captioning →
  - Autonomous Cars
  - Investing
  - Education
  - Medicine
  - Hollywood


A female tennis player in action on the court.


A group of young men playing a game of soccer


A man riding a wave on top of a surfboard.


A baseball game in progress with the batter up to plate.


A brown bear standing on top of a lush green field.


A person holding a cell phone in their hand.


A close up of a person brushing his teeth.


A woman laying on a bed in a bedroom.


A black and white cat is sitting on a chair.


A large clock mounted to the side of a building.


A bunch of fruit that are sitting on a table.


A toothbrush holder sitting on top of a white sink.

Deep Learning in Drug Discovery

- Hot off the presses

Deep learning enables rapid identification of potent DDR1 kinase inhibitors

a

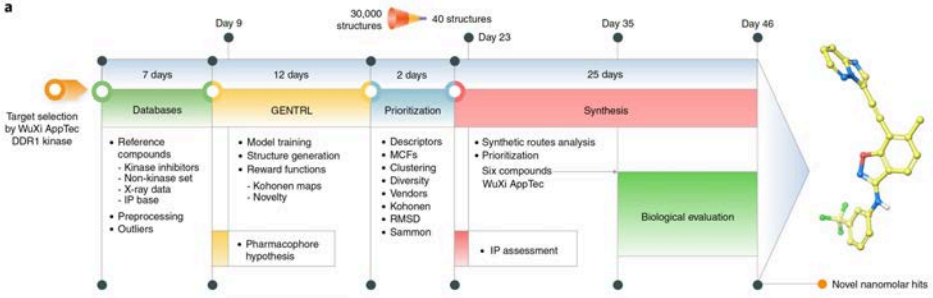

We have developed a deep generative model, generative tensorial reinforcement learning (GENTRL), for de novo small-molecule design. GENTRL optimizes synthetic feasibility, novelty, and biological activity. We used GENTRL to discover potent inhibitors of discoidin domain receptor 1 (DDR1), a kinase target implicated in fibrosis and other diseases, in 21 days. Four compounds were active in biochemical assays, and two were validated in cell-based assays. One lead candidate was tested and demonstrated favorable pharmacokinetics in mice.

- Nature, 2 September, 2019, https://www.nature.com/articles/s41587-019-0224-x

Deep Learning in Medicine

- Hot off the presses (Forbes, Sept 12 2019)
- **"Artificial Intelligence Detects Heart Failure From One Heartbeat With 100% Accuracy"**
  - Doctors can detect heart failure from a single heartbeat with 100% accuracy using a new artificial intelligence-driven neural network.
  - That's according to a recent study published in Biomedical Signal Processing and Control Journal, which explores how emerging technology can improve existing methods of detecting congestive heart failure.
  - Led by researchers at the Universities of Surrey, Warwick and Florence, it shows that AI can quickly and accurately identify CHF by analyzing one electrocardiogram (ECG) heartbeat.

“Deoldifying” old B & W Images and Movies

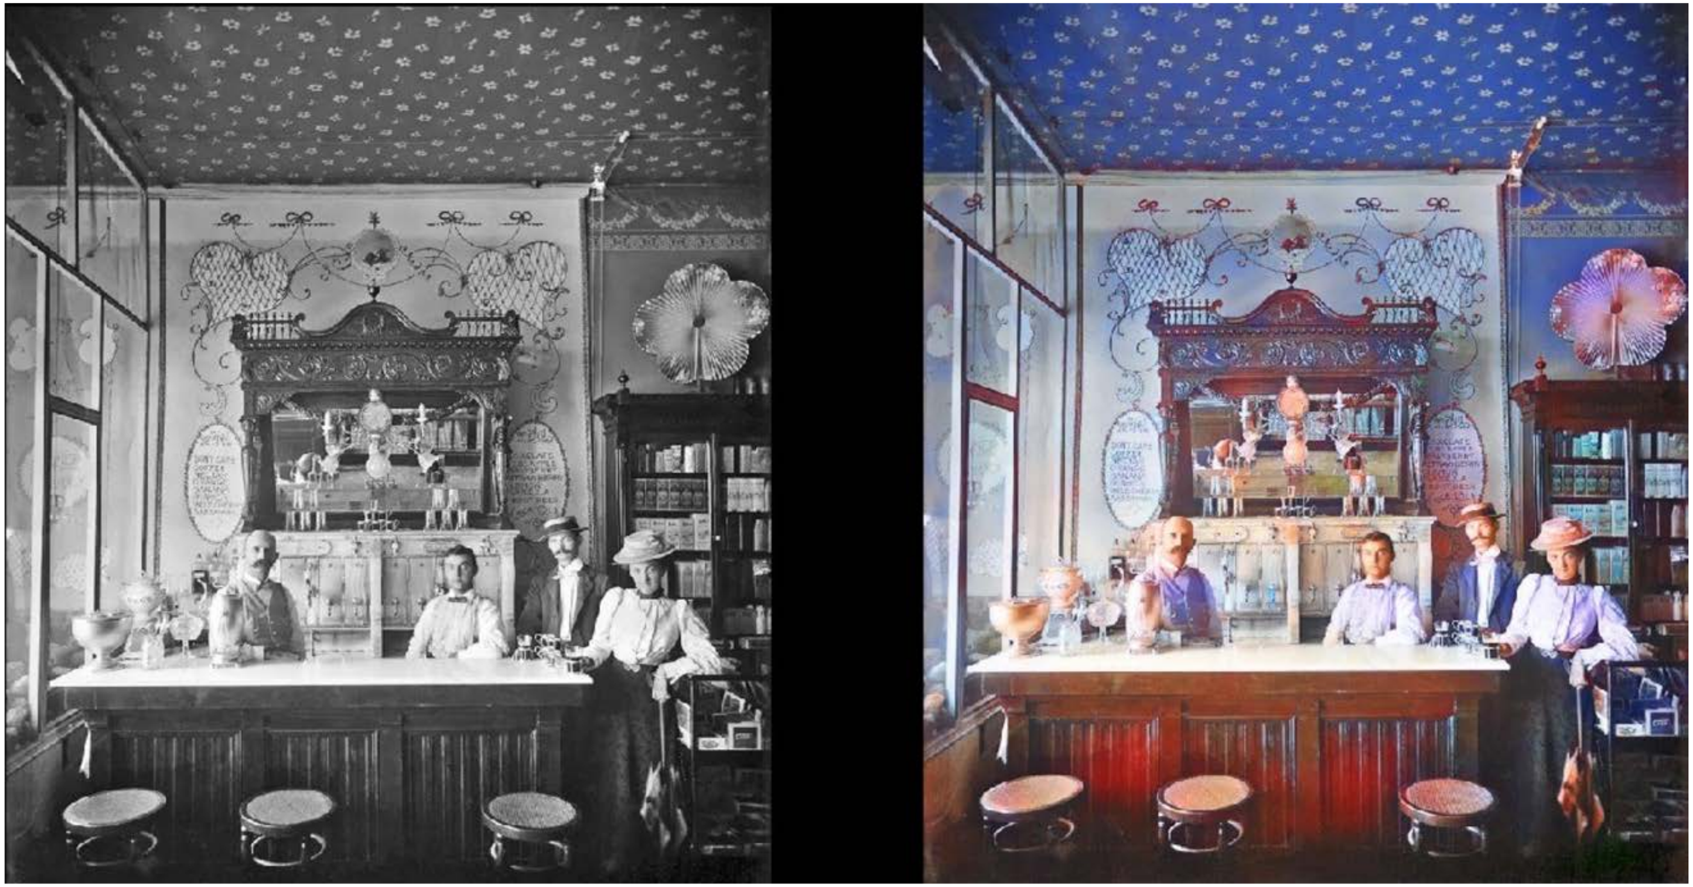

Quantum Computing

- New model of computing based on quantum states of subatomic particles
  - Another reason to take your physics courses very seriously!
- Big push by tech companies and countries to "crack" this and build the first production quantum computer
- Quantum computing will bring exponential speed-ups to many algorithms and will up-end many applications starting with encryption.
- If you plan to do graduate study, this may be an interesting area.

Further ahead

- If Moore's Law keeps up for a longer time, how do
  - CPU power
  - Memory density

  grow.

- Some walls have already been hit.
  - Switch to multi-core structures
  - That is why we have the 15-210 Parallel Algorithm and Data Structures course
- What do we do with all this computing power?
- What would we do with all this memory?
  - Suppose you could "remember" and quickly access everything you see, hear, touch, smell, (maybe even think of (☺)).

Further Ahead

- Internet of Things or computing by the kilo (☺) – security, privacy, robustness
- Forbes Mag. Aug 2019 →
- Any (CS) student who does not take at least 2-3 AI-related courses will most likely be handicapped.

Moore’s Law revisited

- The death rumors of Moore’s Law have been exaggerated
- MIT Technology Review, Sept 2019

Computing / Microchips

The world’s most advanced nanotube computer may keep Moore’s Law alive

MIT researchers have found new ways to cure headaches in manufacturing carbon nanotube processors, which are faster and less power hungry than silicon chips.

- “Modern microprocessor built from complementary carbon nanotube transistors”
  - https://www.nature.com/articles/s41586-019-1493-8

Some grand challenges in computing

- Modelling and fast simulation of living biological structures
  - Cells, Brain, Organs
- Practical ubiquitous computing
  - Computers everywhere but hidden
- Capture, access and process a lifetime of human sensory input and memory
- Understand the architecture of brain and mind
- Build robust and dependable computer systems
- Realize quantum, chemical, biological computing
- Establish “computational thinking” as a fundamental skill in education (CS broadly is the new math!)

Short Term Trends

Vala Afshar
@ValaAfshar

14 technologies of next decade:

1 #AI (ML)
2 #IoT ⌚
3 #blockchain ⛓
4 3D print 🖨
5 mobile 📱
6 autonomous cars 🚗
7 mobile internet 💻
8 robotics 🤖
9 VR/AR 👓
10 wireless power 🔌
11 quantum computing 🖥
12 5G 📡
13 voice DA 🎙
14 #cybersecurity 🔒

WORLD ECONOMIC FORUM

Top 10 skills

	in 2020	in 2015
1.	Complex Problem Solving	Complex Problem Solving
2.	Critical Thinking	Coordinating with Others
3.	Creativity	People Management
4.	People Management	Critical Thinking
5.	Coordinating with Others	Negotiation
6.	Emotional Intelligence	Quality Control
7.	Judgment and Decision Making	Service Orientation
8.	Service Orientation	Judgment and Decision Making
9.	Negotiation	Active Listening
10.	Cognitive Flexibility	Creativity

Source: Future of Jobs Report, World Economic Forum

Final Words and Advice

- You will be living in a "bubble" for at least the next 4 years
  - Interacting with smart, hard-working people like yourselves
- But real world is quite different
- Adjusting could be VERY hard
  - Learn about the Dunning-Kruger effect
    - e.g. Watch https://www.youtube.com/watch?v=wvVPdyYeaQU
  - It will explain so much of the reality you will observe

So

Final Words and Advice

- Develop a healthy appreciation of Humor
  - Especially Computer Humor – seriously!
    - Most of it is self-deprecating and takes the stress away (☺)
  - Better – work on computational humor,
    - develop systems that can laugh at a good joke – seriously!
    - If your system laughs at this – you will get a PhD
      - https://www.youtube.com/watch?v=SXn2QVipK2o

  - Here are some more examples of humor
  - If you do not get any/some of these – you will, in due time

CS Humor

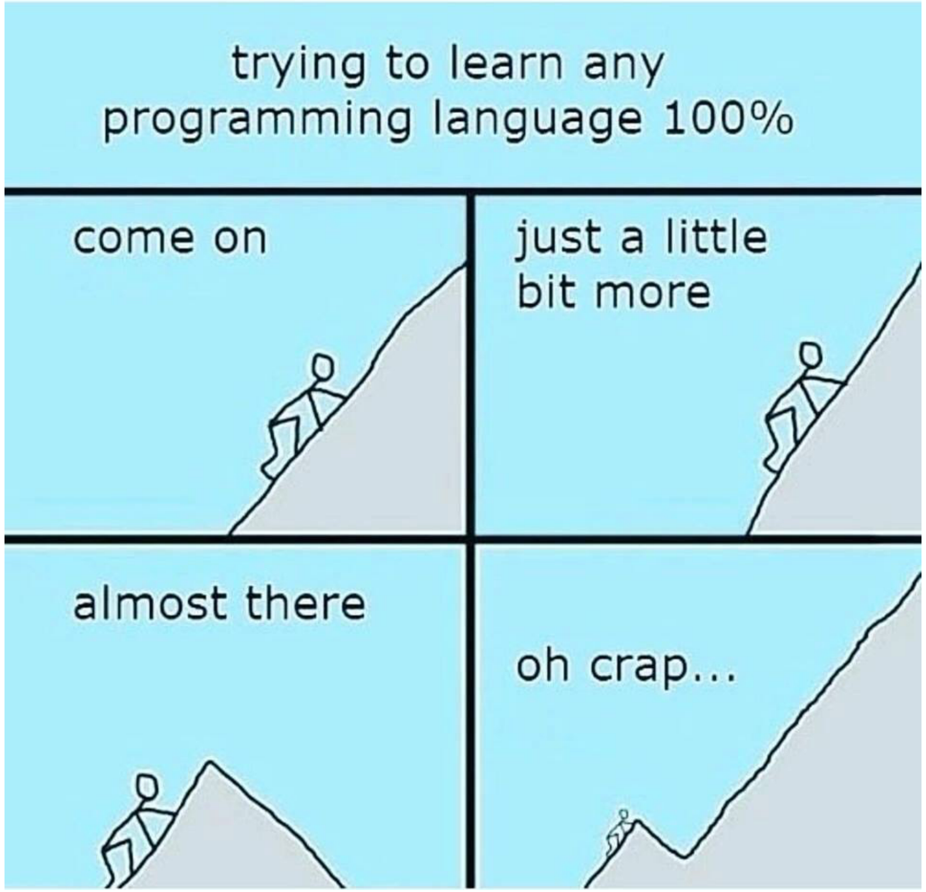
trying to learn any
programming language 100%
come on
just a little
bit more
almost there
oh crap...

CS Humor

99 little bugs in the code.
99 little bugs.
Take one down, patch it around.
127 little bugs in the code...

CS Humor

Books you should not be reading

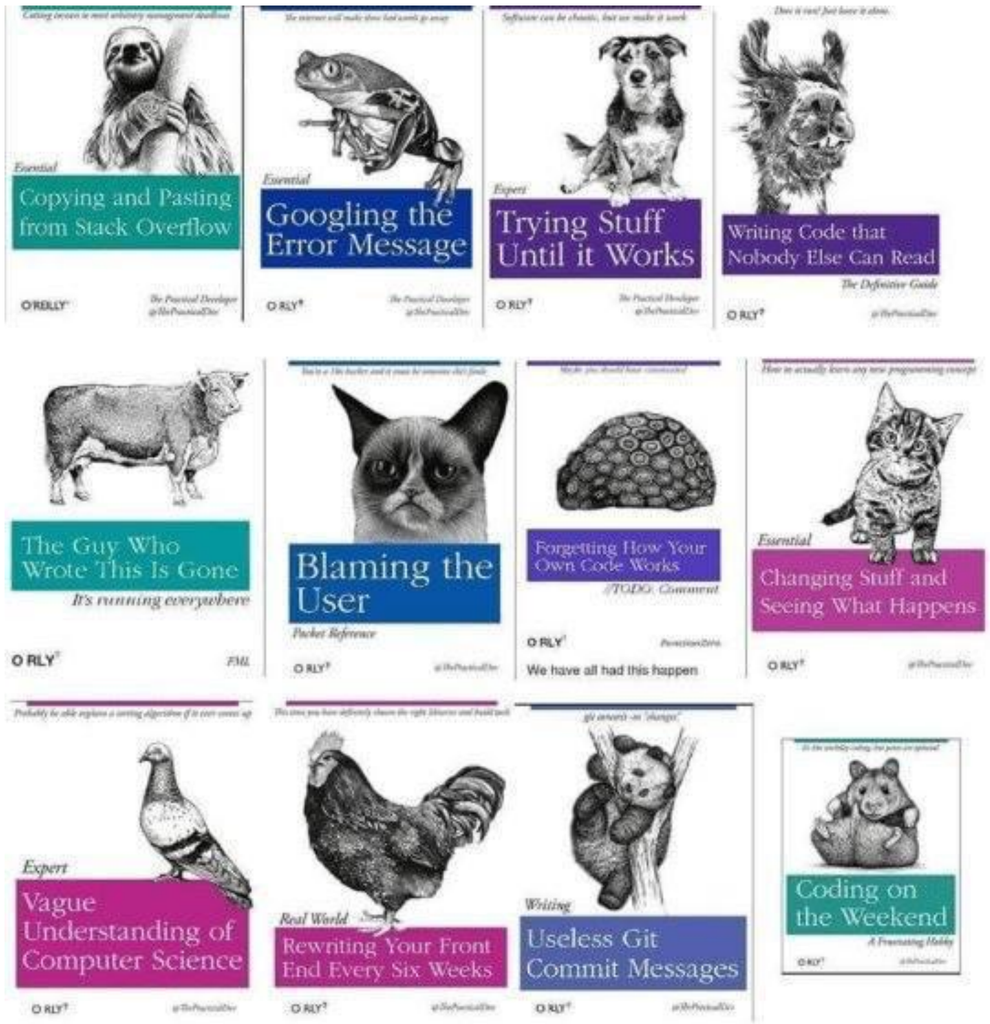

CS Humor

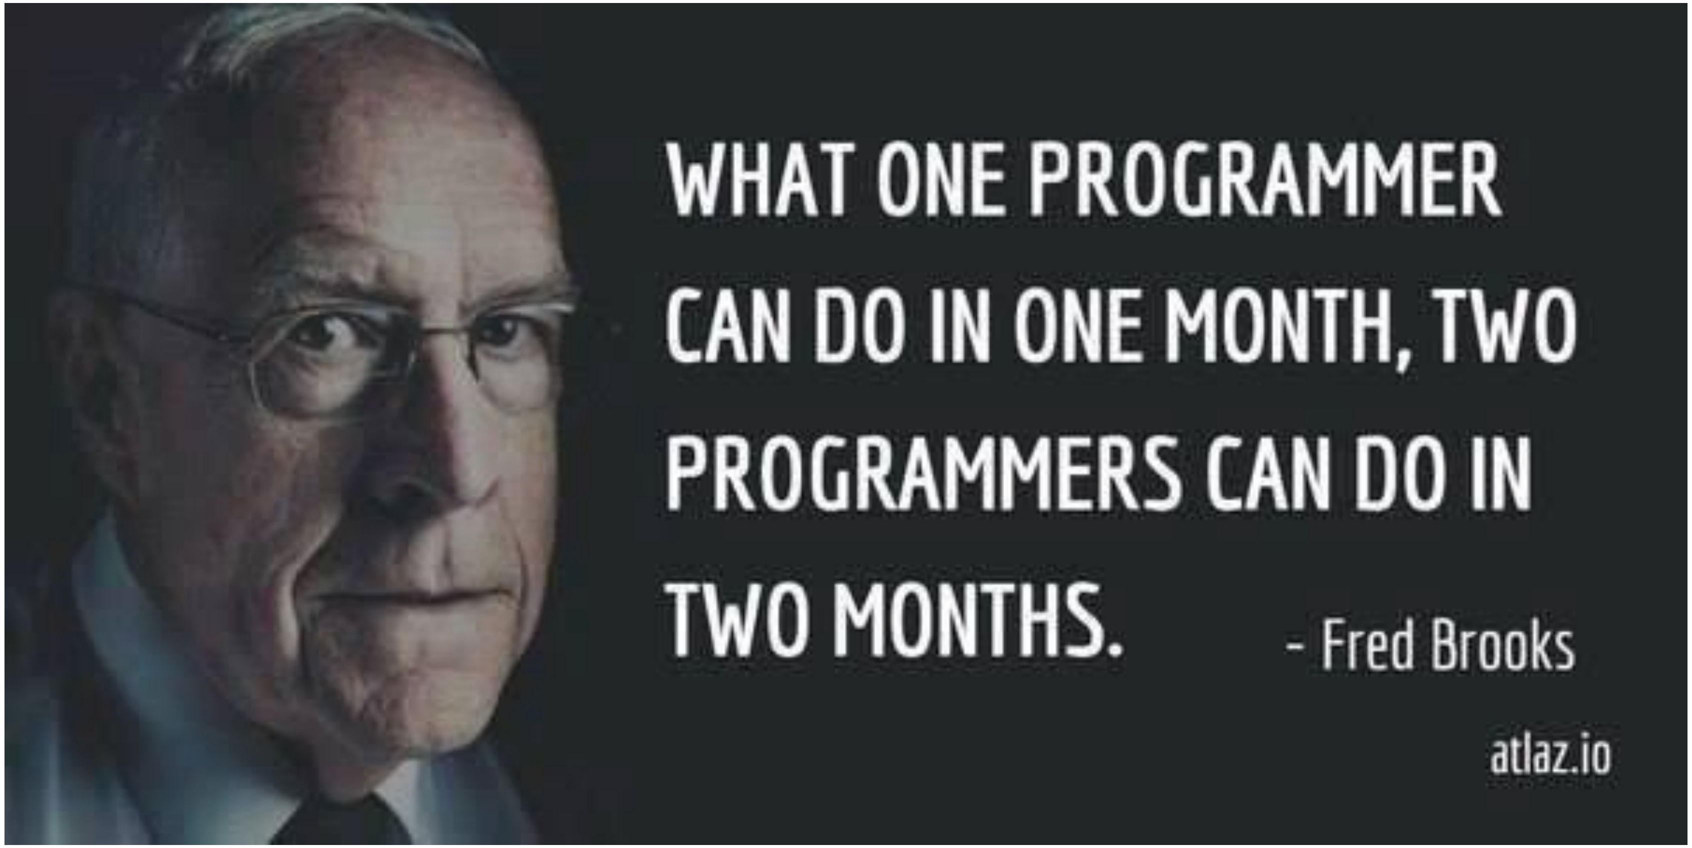
WHAT ONE PROGRAMMER CAN DO IN ONE MONTH, TWO PROGRAMMERS CAN DO IN TWO MONTHS.
- Fred Brooks
atlaz.io

CS Humor

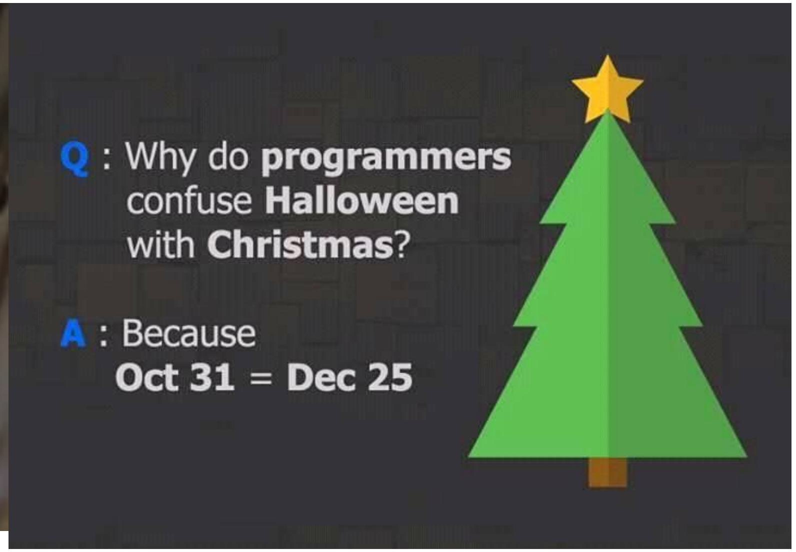

You will get these when you take 15-213

CS Humor

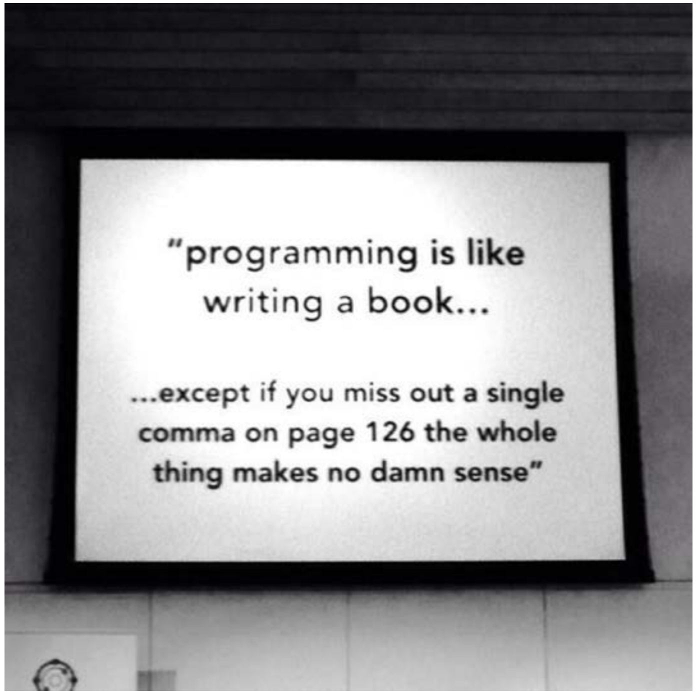
"programming is like writing a book...
...except if you miss out a single comma on page 126 the whole thing makes no damn sense"

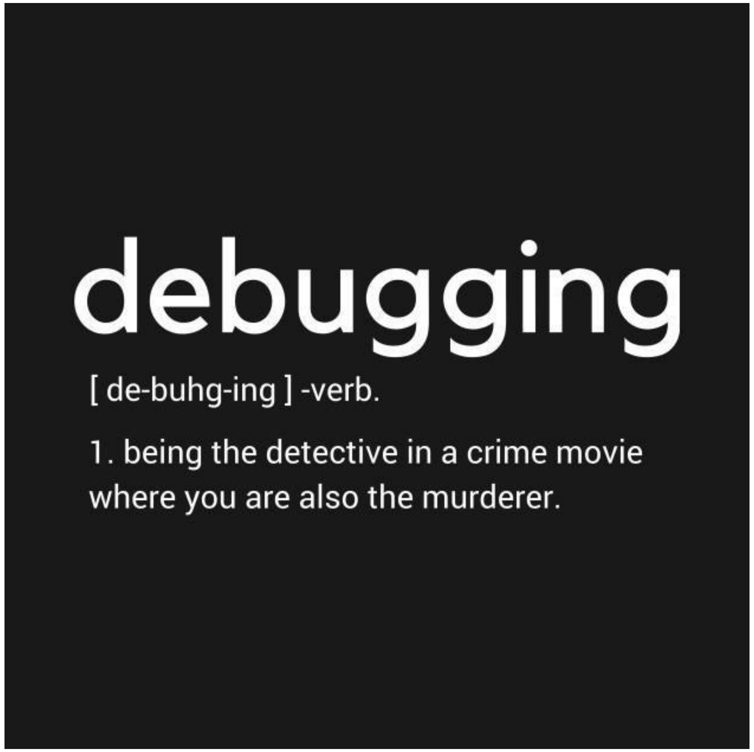
debugging
[de-buhg-ing] -verb.
1. being the detective in a crime movie where you are also the murderer.

CS Humor

```java
public static boolean isEqual(int a, int b) {
    try {
        int c = 1 / (a - b);
    } catch (ArithmeticException e) {
        return true;
    }
    return false;
}
```

Don't ever do clever things like this! We will send you to the gallows (☺)

CS Humor

See also
"The International Obfuscated C Code Contest"

CS Humor

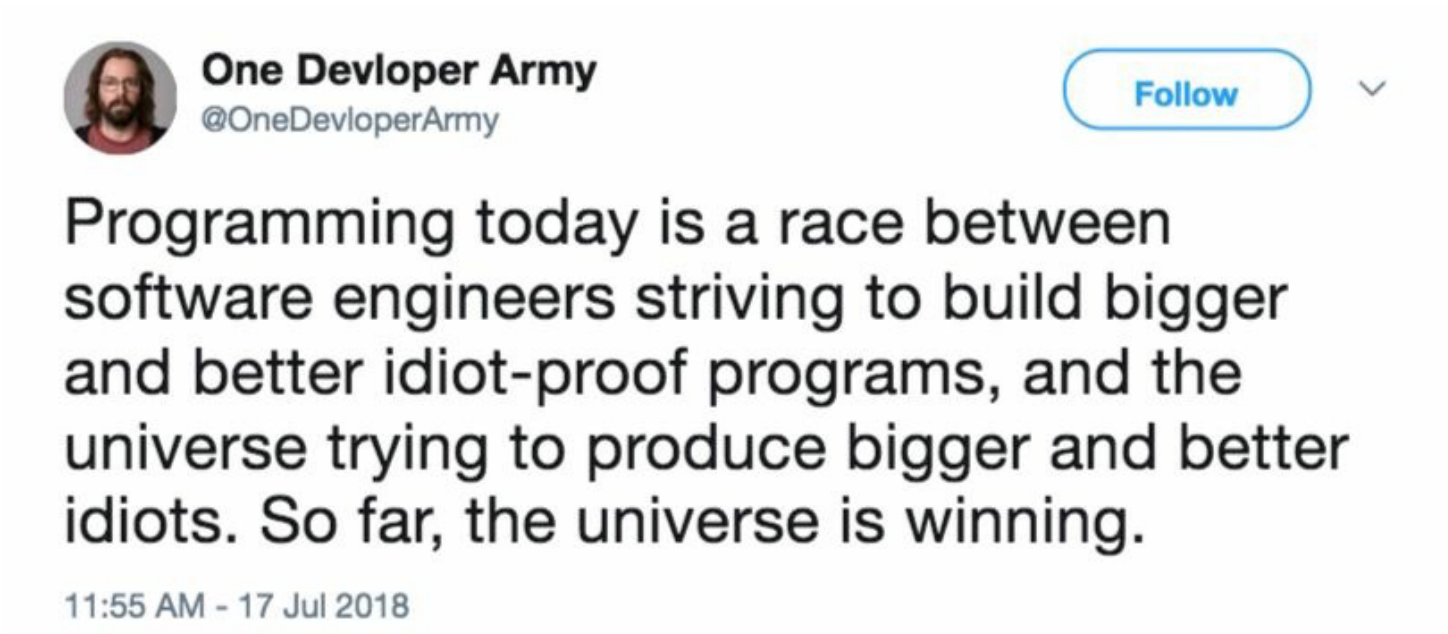

One Devloper Army
@OneDevloperArmy
Follow

Programming today is a race between software engineers striving to build bigger and better idiot-proof programs, and the universe trying to produce bigger and better idiots. So far, the universe is winning.

11:55 AM - 17 Jul 2018

Thank you