Algorithms and Predictions

Kurt Jaeger, pi@fiff.de

https://www.fiff.de/

Würzburg, February, the 8th, 2019

FIfF

Overview

- introduction
- computer, costs and algorithms
- programs and numbers
- deriving structure from data
- changes and predictions
- markets and politics
- privacy
- discussion

Context

- Why FifF ?
  Forum InformatikerInnen für Frieden und gesellschaftliche Verantwortung
  forum computer sciencists for peace and social responsibility
- mid-1980s, out of debates on privacy, automation and military use of computers
- critical view of computer science by computer scientists
- responsibility for research
- social responsibility: automation and job-loss
- privacy a big deal in germany due to historical events

About me

- providing internet services around Tübingen
  retail, manufacturing, science, public admin, high-tech, utilities
- managing a buyers cooperative for internet service providers
- open source development, www.freebsd.org
- participating in the ITK committee
  of the german Chamber of Commerce, DIHK
- doing policy stuff with FifF, Chaos Computer Club Stuttgart, etc

FIfF

The Main Point

- computing and communication capacity exploded over the past 50 years
- capacities grow further in the coming years
- to understand the capacities one has to compare them with other systems
- the growing capacities transform our societies as we speak
- rules, structure and usage of the capacities decides who profits and looses
- understanding this is surprisingly difficult https://www.xkcd.com/2091/

Components

- chips
- costs
- algorithms
- programs
- systems
- bandwidth and latency

Chips

- transistors per chip (billions)
- size of structures on a chip (1000nm vrs. 7nm)
  Moore's Law
- type of chip
  - general CPU, central processor unit
  - memory
  - specialized chip: graphics, network, math, crypto
  - FPGA, field-programmable gate array
  - ASIC, application-specific integrated circuit
- compute cores per chip
- instructions per second (2-4 GHz)
- interconnect between cores, chips, computers
- number of computers
- non-volatile storage: now also at chip innovation rates

Cost

- how much does it cost to manufacture chips
- how fast can you spin a new design ? how much does it cost ?
- how much does it cost to operate one ? thousands ? millions ?
- energy usage!
- who can build them ?
- how are they build ? bill of material
- how many are there ?
- intellectual property
- The end of Moore's Law

Algorithms

- like a cooking recipe, but for operations on data
- human-readable text, how
  - input data is read
  - processed
  - output data written
- data read from storage, sensors, network
- data written to storage, actors, network

Programs

- basic operations
  - add, subtract, multiply, divide
  - compare
  - copy
  - load and store
  - branch
- written in structured text, programming language
- programs implement algorithms
- need to be translated by specialized software into machine code

Size of Programs

- metric: lines of code or compressed source code
- version control, repository: every change can be tracked
- https://www.openhub.net/
- libreoffice: 10 mio lines of code, 200 mb
- llvm: 3.7 mio lines of code, 500 mb
- linux: 10 mio lines of code(?), 160 mb
- windows: 50 mio lines of code
- freebsd os: 16.8 mio lines of code
- freebsd ports/apps: 1.6 mio lines of code, 30K apps 1 gb, compressed 130mb/155mb, repo: 26 gb
- google: more than 2 billion (10^9) lines of code

What numbers ?

- compute
- communicate
- store
- energy, system-wide, per instruction
- units
  - 1024, kilo, k
  - 1024^2, mega, m
  - 1024^3, giga, g
  - 1024^4, tera, t
  - 1024^5, peta, p
  - ...
  - 1 byte is 8 bits,
    plus overhead: approx. 10 bits in communication

Basic numbers

- compute:
  CPUs: 32 cores, 3 GHz per core, approx. 20 gflops/sec
  GPU: 2000 cores, 2 GHz per core, approx. 2 tflops/sec
  FPGA, ASIC: homework!
- approx. 2-4 billion computers worldwide
- communicate: 10gbit/s, 100gbit/s, 160tbit/s for subsea cable
  2-6 tbit/s at the german internet exchange
- energy: human brain
  - approx. 1 petabyte in approx. 1.5kg brain
  - approx. 50watt/hour to operate (food)
  - recent purchase: 30kg, 1/20 of a petabyte, approx. 200watt/hour
  - ecological footprint of both ?

Relation to other numbers

- number of books ever written, avg. size ? 100 mio ? LoC ?
  1 mb per book ? more with pictures ? approx. 100 tb
- number of papers/articles ever written ?
- how many lines of law ?
- one music cd ? 700mb
- the size of the human genome ? 1.5 gb
- one hour of video ? 1h == 3 gb, 400h/minute new on youtube
- how many economic transactions are being done ?
  10 at 1 kb per person per day ? 800 gb per day for .de
- 200mbit/s would be sufficient to transfer that amount of data
- cost of computation compared to cost of a transaction
- how much information can a person process ?
  average/while awake/peak ? 1 gbit/s

Can you relate to numbers ?

- do you feel the pain
  if some newspaper article confuses
  millions (10^6)
  billions (US, 10^9) and
  billions (DE, 10^{12}) ?
- computation and communication are tightly coupled
- complexity: how many computational steps or
  communication is needed
  to solve a task ? in what time ?

Changing numbers changes meaning

Or as Prof. Ludewig said: Quantity is a quality in itself

- we have classic key performance indicators as given above
- in the past, we did exact computation/communication/programming
- exact, but still: lots of bugs, strange behaviour in software etc
- what difference does it make if a computer can do large numbers of calculcations or do receive/transmit large data sets ?
- it can make all the difference
- look at computation/communication from a different viewpoint

The new viewpoint: machines that think

- In the 1940ties: basic understanding of neural networks
- early 70ties for very small applications
- irrelevant until approx. mid-2000, because computational requirements were too high
- basically: derive structure and context from large amounts of unstructured input
- nowadays known as machine learning or 'AI'
- no magic, it's still mostly math and statistics

Human communication as data

- different persons have different ways to learn
- audio: talk and music
- visual: pictures and video
- text: words
- numbers: yes, numbers!
- limited: touch, smell and movements
- augmented senses: sensors for temperature, chemicals, magnetic fields, etc

So what does machine learning change ?

- if you have large computational resources at your disposal you can derive structure from all this data
- Cambrian explosion of computational options

Science

- theory
- experiment
- recently: simulation
- and now: prediction

Prediction

- making predictions on events and decisions, without human intervention
- no need to be perfect, just better than rolling a dice
- how does that earn money ? differentiated pricing, Odlyzko
- increased margin compared to other market participants
- implicit race for market dominiation
- but: it's all about local optima, not about global optimum
- changes the meaning of value and trust
- facebook: doing 200 trillion predictions per day https://twitter.com/ylecun/status/991936213249650688
- Alexa listening is doing predictions
- Serving ads is doing predictions, killing journalism as a by-product

The Value Chain

- chip fabs
- planet-scale computing resources
- fiber optic communication to connect all that
- new algorithms to make use of the unstructured data
- communication networks are a natural monopoly
- who can compete in this world ?
- size does matter!

Scale: Working with the Full Take

- can we measure the state of a country ?
- can we measure the political sentiment of a society ?
- can new words and concepts be introduced into the political debate ?
- how much computation and communication is needed for this kind of influence ?
- can we identify the persons we need to seed with a word/concept to get it going ?
- can a company grow big fast enough to preempt regulation ?

Markets

- efficient-market hypothesis
  asset prices fully reflect all available information
- how do markets work if the marketplace has asymmetric information ?
  and competes as a seller and buyer ?
- ad-networks as markets ?
- 'working' as providing market data ?
- 'living' as signal to markets ?
- how do information markets work if the tools become legal monopolies with patents on algorithms ?
- markets do not really handle intellectual property rights
- economics: do we have the tools to handle all this ?

Politics

- are markets of public opinion
- under pressure by the way new media is changing the communication
- decision makers
- and lawmakers
- how, if news are fake to manipulate audience sentiment ?
- how, if even pictures and videos are no longer 'true' ?
- how, when the first AIs write position papers or legal arguments ?
- need for elections, if your vote is predicted in advance ?

Privacy and Algorithms

- boils down to audit data lakes and algorithms
- how to audit algorithms ?
- by inspection ? 2b lines of code ?
- oh, we know: we write test code! faster than google writes new code ?
- what if they let AIs write their code ?
- The argument for a right to privacy is a stopgap
- right now it's the only stopgap

EFF

Military

- new arms race
- RMA: revolution in military affairs
- full spectrum dominance
- cyberwar
- skynet

Energy

- we inherited a certain amount of energy with our fossil fuels
- that inheritance will soon be spent
- and will change the environment in unimaginable ways
- we should not waste this inheritance in race-to-the-bottom attempts of arbitrage https://www.newyorker.com/cartoon/a16995
- we do not have the time to err into this direction
- we will run out of energy before we run out of information
- we drown in data
- data is fools gold, most of the time

How can we adapt our societies ?

- given that reflection itself is challenged
- what are the tools sociology can provide to help in this scenario ?
- we need tools to debate and come to conclusions before we ran out of energy

Questions ?

Discussion!