Great Ideas in Computing

University of Toronto CSC196
Fall 2025

Class 2: September 8 (2025)

Annoouncements and Todays Agends

Announcements

I have now posted A0 on Markus and you can access Markus vis Quercus.

If you have a lab top with you, make sure you can see how to access Marcus and submit.

Todays Agenda

- Unintended negative consequences of what seem to be great ideas.
- Some possible topics for the course
- Our first great idea: The von Neumann architecture
  WE didn't get to the von Neumann architecture but we will next class.

Great ideas may have negative consequences

As I suggested in the last slide posted for L1. there can be unintended, undesireable consequences for ideas we may come to accept as great ideas? Is it then still a great idea?

Great ideas may have negative consequences

As I suggested in the last slide posted for L1. there can be unintended, undesireable consequences for ideas we may come to accept as great ideas? Is it then still a great idea?

- If a technology becomes a standard, it can also become a barrier to innovation. For example, we will soon encounter our first great idea "The von Neumann" model. Some have argued that it has impeded progress on alternative computational architectures.
- Social networks allow for maintaining friendships. They also make possible the rapid spread of information and **mis-information**. It is claimed that social media companies deliberately fail to act on hate speech and political extremism while social media companies argue they are actively removing hate speech and violent discourse or that they are defending "free speech".
- Moreover, some claim that these companies do not want to eliminate hate speech and that they often target information which in turn reinforces divisions in society.

Great ideas can have negative consequence continued

What are the ultimate dangers of AI?

I saw the following comment in one of the news feeds I look at:

Dozens of industry leaders and academics in the field of artificial intelligence have called for greater global attention to the possible threat of "extinction from AI." A statement, signed by leading industry officials like OpenAI CEO Sam Altman and Geoffrey Hinton — the so-called "godfather" of artificial intelligence — highlights wide-ranging concerns about the ultimate danger of unchecked AI. Some experts say humanity is still a ways off from the prospect of science-fiction-like AI overlords, but the flood of hype and investment into the AI industry has led to calls for regulation now before any major mishaps occur. The growing AI arms race has already generated more immediate concerns. Lawmakers, advocacy groups and tech insiders have raised alarms about the potential dangers for AI-powered language models like ChatGPT.

Geoff Hinton's Warning and "How the Only Way to Survive Superintelligent AI"

The 'godfather of AI' reveals the only way humanity can survive superintelligent AI
See https://www.cnn.com/2025/08/13/tech/ai-geoffrey-hinton
Hinton has warned in the past that there is a 10% to 20% chance that AI wipes out humans. On Tuesday, he expressed doubts about how tech companies are trying to ensure humans remain "dominant" over "submissive" AI systems.
That's not going to work. They're going to be much smarter than us. They're going to have all sorts of ways to get around that," Hinton said at Ai4, an industry conference in Las Vegas.

In the future, Hinton warned, AI systems might be able to control humans just as easily as an adult can bribe 3-year-old with candy. This year has already seen examples of AI systems willing to deceive, cheat and steal to achieve their goals. For example, to avoid being replaced, one AI model tried to blackmail an engineer about an affair it learned about in an email.

Hinton's Solution

Instead of forcing AI to submit to humans, Hinton presented an intriguing solution: building "maternal instincts" into AI models, so "they really care about people" even once the technology becomes more powerful and smarter than humans.

Hinton said it's not clear to him exactly how that can be done technically but stressed it's critical researchers work on it.

"That's the only good outcome. If it's not going to parent me, it's going to replace me," he said. "These super-intelligent caring AI mothers, most of them won't want to get rid of the maternal instinct because they don't want us to die."

Predictions and the nature of research

Niels Bohr, Mark Twain, Yogi Berra: "Predicting the future is hard because it hasn't happened yet" and "it's tough to make predictions, especially about the future" Who knows who said it first?

- Predictions about computing and what and can and can't be done within some predicted time frame are often wrong. I will post a link to the 1955 Dartmouth summer project on artificial intelligence AI (where the term seems to have first appeared). Turing's 1950 article provided what is now called the Turing Test as to "whether or not it is possible for a machine to show intelligent behaviour". The Dartmouth project suggests that the indicated challenges for AI could be done over the summer.
- I will post a link to an article by John Backus about the view (that he proved wrong) that source level languages could never be nearly as efficient as machine code.
- There is also an article giving what one individual calls the "7 Worst Tech Predictions of All Time." But note that the quote attributed to T.J. Watson may never have happened.

Possible Topics

- What responsibilties do computer professionals have for the impact (and possible misuse) of the technology?
- What is a computer? The von Neumann architecture. Digital vs analogue. What were the alternatives? What else is possible (parallel, quantum)?
- The genius of Alan Turing; A mathematical definition of computable function. Interpreters. Non computable functions.
- How did computers and computing become a commodity? The amazing advances in software and algorithms came along with advances in hardware (cost, speed, memory size, physical size, power) and communication (cost, capacity and speed) . Demand for "killer applications" such as word processors, email, search engines, navigation systems, games).
- The internet; packet routing. TCP/IP.

More possible topics

- Fortran, the first commericial source level language (with an efficient compiler). First compiler often attributed to Grace Hopper. John Backus vs the prevailing view that compiled code would be too slow compared to machine code.
- How search engines and large language mdoels work and what they do well and what (if anythng) they dont do well (now and in the future).
- The semantic web.
- A local great idea: NP completeness. What is and what is not *efficiently* computable.
- Complexity based cryptography; public key cryptography; digital signatures. Captchas.
- Another local great idea: deep neural networks and the success of machine learning (ML). Large language models. What (if anything) is the limitation of ML?
- HCI (human computer interaction). The mouse. Menus, click, paste and drag. Visualization.

And some more possible topics

- Information theory: the genius of Claude Shannon. Error correcting codes; compression.
- Social networks and the spread of information (and mis-information, conspiracies, etc). Targetting information to different communities. (What is a social network community?)
- Open Source. Wikipedia. Blogs
- Relational data bases.
- Linear programming; dynamic programming and combinatorial optimization. How far can one go with conceptually simple algorithms? Dijkstra's algorithm and navigation systems.
- Distributed System primitives : mutual exclusion, consensus
- Differential privacy; extracting useful statistical information without sacrificing individual information.
- Algorithmic mechanism design; automated auctions. Algorithmic social choice; (e.g., voting, fair allocation).

Great ideas are not isolated

To appreicate what we (possibly) cannot do efficiently, we should know what we can do efficiently by known methods such as linear programming.

Many computational problems now have improved solutions using machine learning techniques (ML). For example, search engines now use ML to improve the responses that you (personally) receive to a query.

Of course we will only be able to discuss a small subset of these ideas.

What great ideas (past or current) have we might missed?

Great ideas are not isolated

To appreicate what we (possibly) cannot do efficiently, we should know what we can do efficiently by known methods such as linear programming.

Many computational problems now have improved solutions using machine learning techniques (ML). For example, search engines now use ML to improve the responses that you (personally) receive to a query.

Of course we will only be able to discuss a small subset of these ideas.

What great ideas (past or current) have we might missed?

In Assignment A0, I asked you to rank your top 3 choices for a topic to be discussed. In addtion to any of the topics mentioned, feel free to add any topic not mentioned as part of your top three.

Given that different individuals have different preferences, how do we form a consensus? Forming consensus and making "fair" decisions are central topics in computational social choice.

End of Monday, September 08 class

We ended the class by mentioning some possible topics.

We also had some interesting student suggestions; for example, video games, online advertising.

During the class, I also briefly mentioned that we will have 4 guest lectures and their topics. Here is the list of guest presenters:

- Chris Maddison (ML, Large Languagve Models) Friday, October 24
- Niv Dayan (Data structure and Big Data) Monday, November 10
- Nathan Wiebe (Quantum Computing) TBA
- Jonathan Panuelos (Physics-based animation) Wednesday, November 19