My Concerns About Artificial Intelligence

Abstract

There is growing alarm about the potential dangers of Artificial Intelligence (AI). Giants of the commercial and scientific world have expressed the concern that AI will eventually make people superfluous. A group of activists has argued that AI should not be used in weapon systems, explaining their fear that killer robots might start to fight wars against people and without morality. A Microsoft researcher recently made headlines saying, “As artificial intelligence becomes more powerful, people need to make sure it is not used by authoritarian regimes to centralize power and target certain populations.”

My concerns are very different. I believe that the use of AI techniques will make software intensive systems even less trustworthy than they are now. I see a lot of unscientific thinking and a lack of engineering discipline in AI work.

This talk explains my views. It is designed to stimulate discussion.

Contents

Two Instructive Jokes

What's the difference between a horse-breeder and a Computer Scientist?

- **When the horse breeder talks about "AI", he can explain what is meant.**

What's the difference between a cattle-farmer and an AI fan?

- **A farmer knows BS when she sees it.**

A commandment for scientists:

Thou shalt not use terms that you cannot define.

- **"Intelligence" is one of those terms.**
- **"Artificial Intelligence" is worse.**

Definitions in science must be based on measurement

- **Define terms before you use them in a talk or paper.**

Three Types of AI Research

- Building programs that imitate human behaviour to understand human thinking better (psychology research)
- Building programs that play games well (challenging and fun)
- Demonstrating that practical computerized products can use the same methods that humans use (risky and often naive).

It is only the third that concerns me.

Responsibilities of Engineers

- *Do no harm!*
- *Make products that are fit for their intended use*
- *Make products that can be trusted*
- *Make sustainable/maintainable products (design for change).*

I doubt that AI meets these responsibilities.

Commandment:

Thou shalt always remember that when building a device for others to use, it is your responsibility

- *to know when it is trustworthy/not trustworthy,*
- *to tell the users the circumstances in which it cannot be trusted.*

Surprises are not a "feature".

What Does "AI" Really Mean?

One of the following:

1. **Automation - when arguing that AI is useful.**
2. **A program that appears intelligent — illusionists[1].**
3. **A program that uses methods intended to mimic the thinking methods of people (hopefully intelligent ones)**
4. **A program that uses heuristics[1] (doesn't always work — also like people).**
5. **A lazy way to solve a problem. Don't bother to analyze the problem, just let the program learn[1] on the job.**

Commandment: ***Thou shalt say what you mean*.**

[1] To be discussed later.

Heuristic Programming

Some AI researchers describe the AI approach as “heuristic programming.”

Heuristic programs are based on “rules of thumb”.

- Claimed to be methods that humans use.
- These rules are based on experience but do not always get the right answer.

Computers have different capabilities and can often do better.

Example: Finding a route through a maze.

- Early AI researcher studied how humans find their way through a maze.
- Floyd/Warshall/Tarjan/Dijkstra algorithm outperforms the resulting program. He had not heard of it.
- Computers can use it. People in a maze cannot.

Example: facial recognition

- Human methods and practical computer methods very different.
- Problem is different too.

When Are Heuristics OK?

Heuristics can be safely used in a program if:

- **The specification allows several acceptable solutions and the heuristic is used either to select one of them or to determine the presentation order.**
- **The heuristic is intended to speed up a program that conducts a search that will either find a solution or establish that there is no solution and will not change the result.**

In other situations, heuristic programming is untrustworthy programming.

Commandment:

Thou shalt use heuristics only when the correctness of the program's results does not require the correctness of heuristic choices.

The "Turing Test"

Turing is often cited incorrectly.

- **Turing did not claim his test was a test for artificial intelligence!**
- **Turing said, "Can Machines Think" was an unscientific question.**
- **His test was an example of a scientific question, not a reformulation of a question that was not scientific.**

Commandment

Thou shalt not cite Turing without reading Turing's paper, "Computing Machines and Intelligence" itself, not just articles about Turing's paper. Accept no substitutes

Warning: Paper is not particularly well written.

Example: Character Recognition

Popular topic in early AI research and courses was character recognition.

- **The OCR software that to recognize characters on a scanned printed page frequently errs.**
- **Captchas work.**

Early AI experts recommended interviewing human readers.

- **For example, ask how to distinguish an "8" from a "B."**

People could do the job but could not explain how.

- **Modern software for character recognition based on restricting fonts.**
- **Solve the problem by producing papers on a computer! OCR is not needed.**

Commandment:

When thou encounters a problem that needs AI, try to avoid it.

Example: Assembly Line Assistant

Tool-handling assistants sped up assembly line.

- **Assemblers tossed tool in a box when done with it.**
- **Assistant retrieved tool when needed.**

Researchers tried to replace the human assistants with robots.

"Computer vision" algorithms could not find the desired tool in the heap.

- **The problem was changed. Robot put tool in the box.**
- **Robot remembered where the tool was and could retrieve it easily.**

"AI" could not imitate the human but simple machine could do more.

Commandment: *Wherever possible, thou shalt avoid AI by expanding the problem*.

Machine Learning

Programs collect data during use; use data to improve performance.

- This is now called machine learning. (Sounds more magical)
- Not new: Turing speculated about a "child program" to be trained like a child.

Requires no "intelligence."

- Robert Dupchak's 1960 penny-matching machine did it.
- A small black box could "outsmart" us by remembering our past behaviour.

"Machine Learning" seems easier than problem analysis but,

- The resulting programs may be untrustworthy.
- ML may exhibits the weaknesses of "hill-climbing" algorithms; miss the peak.
- ML may also err because of incomplete or biased experience.
- Learning is a restricted form of statistical classification.

Machine-learning algorithms are heuristic; may fail in unusual situations.

Commandment: *Thou shalt use physics and math to guarantee correctness; use collected data to optimize performance.*

Neural Nets

Discrete model of brain structure.

No theoretical increase in power.

Brain structure is constrained to allow growth.

Computers do not have that constraint.

Other highly parallel processor structures available.

Tell me why we should build things that way.

Creating Illusions

Alan Perlis (first Turing Prize winner) called AI researchers "illusionists".

- **They try to create the illusion of intelligence.**
- **Should be considered stage magicians not scientists.**
- **Dupchak (penny matcher) and Weizenbaum (Eliza) showed this to be easy.**

We do not want computer systems that perform tricks

- **We need trustworthy tools.**
- **Trustworthy systems must be based on sound mathematics and science;**
- **They must not be based on heuristics or illusionist's tricks.**
- **Scientists shed light on facts; they are not illusionists.**
- **Engineers build trustworthy products, not chimera.**

Commandment: ***Thou shalt act as scientist or engineer — not an illusionist*.**

The Ethics of AI

Whatever it is, AI is part of a machine.

Machines don't have to have ethics.

Their users/designers/purveyors all have to follow ethical principles.

The question is not, "Is the Robot 'moral'?".

The right question is, "Is the use of this machine, in this situation, moral?".

AI Has Become a Clique

Cliques are commonly formed among researchers.

- They read, write, and meet only within the clique. They have their journal.
- They are often unaware of results, tools, and techniques from outside.
- A huge amount of duplication results.
- Bad solutions are used where good ones exist.
- Vocabulary differences make communication impossible.

CS has become a clique. We are isolated from computer users.

CS is being divided into cliques. We don't talk across disciplines enough.

AI has become a clique on its own.

- Example: Vision researcher who was unaware of signal processing progress.
- Example: Neural net not compared with least squares method.
- Example: Researchers not familiar with algorithm and software design results.

Some Conclusions

When someone says, "This product uses AI", ask questions.

- **Ask for a definition of AI. It might be good; you can't assume its bad.**
- **Ask the method. "AI" usually hides sloppy work but not always.**

AI methods least risky if getting an incorrect result or no result at all is OK.

- **If you are prepared to accept "I don't understand" or an irrelevant answer from your "personal assistant," AI is harmless.**
- **If the response is important, be hesitant about using AI.**

Programs that almost always work can be the most dangerous.

Demonstrations of AI can mislead. Test it yourself.

Trying to win the "Imitation Game" is a waste of time.

Always use science and mathematics fully before turning to heuristics.

Be afraid, very afraid, of AI solutions. Fear sloppiness, not power.

Get the Whole Story

Parnas, D. L., “The Real Risks of Artificial Intelligence”,

Inside Risks Viewpoint,

Communications of the ACM, Vol 60 Issue 10, October 2017, pp. 27 - 31.

DOI>10.1145/3132724