

Probabilistic Automation Systems ("AI")

- [What do we actually mean when we call something "AI"?](#)
- [The hidden human work behind automation](#)
- [Data is never neutral](#)
- [When AI starts making things up](#)
- [Are we becoming too dependent on LLMs?](#)
- [Creativity and AI](#)
- [Data privacy and ownership](#)
- [Big companies and the future of innovation](#)
- [So, what actually is AI?](#)
- [Things worth discussing](#)
- [Source](#)

What do we actually mean when we call something "AI"?

The term **Artificial Intelligence** was coined by John McCarthy and since then it has become one of the most widely used terms in technology. Today, almost everything related to automation, prediction, generation or recognition is being called AI. It has also become a very common marketing term, where companies use "AI" to make their products look more advanced or intelligent.

But what actually is AI?

Emily Bender and other researchers used the term "**Stochastic Parrot**" to describe large language models. The basic idea is that these systems are stochastic models which learn from huge amounts of data and predict what word, token, character or piece of information is likely to come next. The systems are extremely complex, but at the same time the basic principle is still based on probabilities and patterns learned from data.

In a way, AI can also be seen as a **container** for many different things. It can contain text, symbols, pictures, speech and other types of information, as well as generative and recognition systems. We use the same word "AI" for many completely different systems.

The output of these systems is defined by statistics and algorithms. This brings up one of the biggest questions in this whole discussion: **is there actually intelligence in these systems, or are we simply interpreting the output as intelligence because it looks intelligent to us?**

This is where **anthropomorphising** becomes important. Anthropomorphising means giving human traits, emotions, behaviours or intentions to something that is not human. With AI, we often say things like "the AI wants to help", "the AI understands me", "the AI decided this" or "the AI knows what I mean". But does the system actually have an intention, or are we just giving it one because it communicates with us in a human-like way?

This can create a misconception that the system itself has intentions, feelings or understanding when it may simply be producing an output based on patterns and probabilities.

The hidden human work behind automation

Probabilistic automation systems are more than just algorithms. They also require a huge amount of human labour behind the scenes.

Before a system can learn from data, people have to prepare that data. There is work involved in annotation, labelling, transcription, moderation and many other tasks. Even after a model is trained, humans are often required to evaluate and improve its outputs.

This creates an interesting contradiction. We call these systems "automation", but there are still many humans working behind the scenes to make that automation possible.

A lot of this work is concentrated in countries across Africa and South and East Asia. In many cases, workers can face problems related to low wages, temporary contracts, lack of worker protection and limited access to healthcare and other support. Some moderation work can also involve exposure to disturbing or harmful content.

When we use an AI system, we normally only see the final product. We don't see all the people who helped create the data that made the system possible.

So the question becomes: **how much of the "automated" world is actually dependent on invisible human labour?**

Data is never neutral

Another important thing about probabilistic systems is that they are not representing all available information equally. They certainly cannot be considered a perfect representation of the world's population.

These systems can suffer from gender bias, racial bias, age bias, political bias, cultural bias and many other types of bias.

The problem starts with the data itself. If the data contains certain patterns, the model can learn those patterns. If some groups are overrepresented and others are underrepresented, the resulting system can reflect those differences.

Datasets are never completely neutral. They are created by people, from sources selected by people, and the decisions made while collecting and preparing that data can influence the final model.

A lot of the information available on the internet is also dominated by Western sources and by languages and cultures that have a much larger online presence. This means that asking an AI system about the "world" does not necessarily mean that we are getting an equal representation of the whole world.

The model only knows what is available to it through its training and other information sources.

This makes me think about a simple chain:

World → Data → Dataset → Model → Output

Bias can enter at almost every stage of this process.

When AI starts making things up

One of the biggest problems with generative AI is that it can produce information which looks completely correct but is actually wrong.

This is especially dangerous because the answer can be written very confidently. A person reading it may not realise that the information is incorrect.

Scientific references are a good example. There have been cases where language models generated scientific references that did not actually exist. Recent research and reporting has shown that fabricated or incorrect citations are becoming a problem in scientific literature.

This creates a strange situation where a system can produce something that looks very scientific while the source behind it may not exist at all.

So we have to remember that **something sounding correct does not mean that it is correct.**

Are we becoming too dependent on LLMs?

Another topic I find interesting is how much people are starting to depend on LLMs for their everyday work.

People are using them to write emails, solve programming problems, write assignments, summarise articles, generate ideas, write reports and even make decisions.

The question is what happens to our own skills when we stop practising them.

If we always ask an AI to solve a problem instead of trying to solve it ourselves, we may slowly lose the ability or patience to think through difficult problems. If we always ask an AI to write something for us, we may also get worse at writing ourselves.

There is already a lot of discussion and research around cognitive offloading and how relying on external tools can change the way people think and remember information.

I don't think it is as simple as saying "**AI will make humans dumb**", because technology has always changed the way humans use their brains. Calculators reduced the need to perform certain calculations manually. Search engines changed the way we find information. GPS changed the way we navigate.

But AI is different because it can now perform tasks that involve writing, reasoning, programming and creating ideas.

So the more interesting question is:

If we outsource more and more of our thinking to machines, which abilities will we stop practising?

And if we stop practising those abilities for a long enough time, what happens to our creativity, problem-solving and ability to think independently?

Creativity and AI

AI can now generate images, music, videos, stories, designs, code and many other things.

But does generating something mean that the system is actually creative?

Maybe creativity is simply the ability to produce something new. Or maybe creativity also requires intention, experience, emotions, personal memories and understanding.

When a human creates a painting, there is normally a reason behind it. There is a person with experiences and emotions deciding what they want to express.

When an AI generates an image, it is producing an output based on patterns it learned from its training data.

So who is actually being creative?

Maybe AI is creating.

Maybe the person using the AI is creating.

Maybe it is a combination of both.

Or maybe we need to rethink what we mean by creativity in the first place.

Data privacy and ownership

Another major problem is data privacy.

The amount of data being collected and used by technology companies today is enormous. Social media platforms, search engines, video platforms and other services collect huge amounts of information about how people interact with the internet.

At the same time, AI companies are training models using enormous datasets collected from different sources.

This raises questions about where the data came from, whether people gave permission for their data to be used, and whether creators should be compensated when their work is used to train these systems.

There is also a difference between something being publicly accessible and something being ethically available for any possible use.

Just because something is available on the internet does not necessarily mean the person who created it expected it to become training data for an AI model.

This also creates problems around copyright and intellectual property.

A photographer, artist, writer, programmer or researcher may spend years creating something, while an AI system can potentially learn patterns from huge amounts of similar work in a very short period of time.

So the question is not only **"Can AI do this?"**

It is also:

"Who owns the data, who gave permission for it to be used, and who should benefit from it?"

Big companies and the future of innovation

Another thing that is becoming more obvious is how much of the AI market is controlled by large technology companies.

Building advanced AI systems requires huge amounts of computing power, data, infrastructure and money. This creates a situation where smaller companies, researchers and independent developers can find it difficult to compete at the same level.

This makes me think about what happens to innovation when the technology becomes concentrated in the hands of a small number of companies.

Large companies can also acquire smaller companies, buy their technology or invest in promising ideas. Sometimes this is good because it gives small companies the resources to grow. But there is also a concern that independent inventions and ideas can eventually become controlled by much larger companies.

The question is whether the future of technology will continue to be driven by many independent people and small organisations, or whether it will increasingly be controlled by a few companies that have the money and infrastructure to build the biggest systems.

So, what actually is AI?

After looking at all of these things, I don't think AI should simply be viewed as a piece of software.

There are algorithms, but behind those algorithms there is also data. Behind the data there are people. Behind those people there are companies, infrastructure, money, laws and society.

People → Data → Labour → Algorithms → Models → Companies → Users → Society

Everything is connected.

The important thing is that we should not blindly accept the word "AI" and assume that something is intelligent simply because it produces an impressive result.

At the same time, we should also not underestimate what these systems are capable of.

Maybe the real discussion is not whether AI is intelligent or not.

Maybe the real question is:

“ What are we actually building when we build systems that can imitate more and more of the things humans do?

And more importantly:

“ What are we willing to give up in exchange for the convenience and power these systems provide?

Things worth discussing

This book is not meant to give a final answer.

It is meant to start a discussion.

Is a system intelligent if it can behave intelligently but does not have consciousness?

Can an LLM actually understand something, or is it only predicting patterns?

Are we becoming too dependent on AI?

Will AI increase human creativity or reduce it?

Should creators be compensated when their work is used for AI training?

Is publicly available data automatically fair to use?

How much hidden human labour exists behind AI?

Can datasets ever be completely neutral?

Will AI make humans more capable, or will we slowly stop developing some of our own abilities?

And finally:

If we are entering the "Age of AI", what kind of humans do we want to become in it?

Source

The original workshop material that inspired this discussion:

https://pub.fabcloud.org/project/fab26_bottrop/workshops/Probabilistic-Systems/

Additional reading:

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major & Margaret Mitchell — *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*

Nature — research and reporting on fabricated scientific citations and the growing use of AI in scientific publishing.