

The Ethics of AI

The Ethics of AI

Essay 4 — Understanding AI Course

AI is powerful. And with power comes responsibility — both for the people who build it and for the people who use it. In this essay, we'll explore some of the biggest ethical questions around AI: What does it mean for jobs and the economy? Why are people so afraid of it? What is "AI slop"? And is the famous "Singularity" something to worry about — or just science fiction?

AI and Employment — Will Robots Take Our Jobs?

This is probably the question people ask most often about AI. And it's a fair question — because the answer is complicated.

What AI actually does to jobs

The honest truth is that AI *does* automate certain tasks that humans used to do. This has happened before — every major technological revolution, from the steam engine to computers, has changed the nature of work. Some jobs disappeared. New ones were created. Overall, the economy adapted.

AI is similar, but it may happen faster — and it affects not just manual jobs but cognitive ones too. Copywriters, data analysts, legal assistants, artists — AI tools can now do things in minutes that used to take humans hours or days.

What economists generally say

Most serious economists and researchers who study this don't predict mass unemployment — but they do predict **significant disruption**. The likely outcomes are:

1. **Some jobs will disappear** — particularly those involving routine, repetitive tasks
2. **Many jobs will change** — the tasks involved will shift, with AI handling the routine parts
3. **New jobs will emerge** — AI needs trainers, ethicists, auditors, and people who can work alongside it
4. **New industries will grow** — just as the internet created millions of jobs that didn't exist before

The real concern — not just job quantity, but distribution

Here's the issue that worries many economists most: even if new jobs are created, they're not necessarily the same jobs, in the same places, for the same people. A truck driver made redundant in Newcastle doesn't easily become an AI specialist in London. The benefits of AI-driven productivity tend to flow disproportionately to those who own the technology — already wealthy individuals and large corporations.

For the UK specifically, this is a real policy challenge. Without deliberate action on education, retraining, and wealth distribution, AI could make economic inequality worse.

Why Are Humans So Afraid of AI?

We touched on this in Essay 1, but let's go deeper.

Fear of AI isn't really about the technology itself — it's about what the technology *represents* to people. Here are the main underlying fears:

1. Existential risk

Some very smart people — including scientists like Stephen Hawking and AI researchers — have warned that AI that surpasses human intelligence could be genuinely dangerous if not properly controlled. This is sometimes called **existential risk**: a risk that threatens the

continued existence of humanity. Whether this is realistic or just speculative is genuinely debated — but it's not a ridiculous concern.

2. Loss of agency

Many people feel that decisions that should be made by humans are increasingly being made by algorithms. Who's deciding what news you see? Who decides whether your loan application is accepted? When AI is involved in consequential decisions, people rightly want to know: *who is accountable?*

3. The speed of change

Technology has always changed quickly, but AI feels different. It seems to be advancing faster than society can adapt — faster than laws can be written, faster than people can understand it, faster than regulations can keep up.

4. Loss of meaning

There's something deeper too. For many people, work isn't just about money — it's about purpose, mastery, and contribution. If AI can do what we do, what does that mean for how we find meaning? This isn't a silly concern. Philosophers and psychologists have written seriously about it.

What About AI Slop?

You've probably heard the term **AI slop** — it's become a popular way to describe low-quality, AI-generated content that floods the internet.

The problem is real. Because AI generates content so easily, it's now economically viable to produce enormous quantities of it for very little cost. This means:

- **Low-quality blog posts** written by AI and posted purely to game search engine rankings
- **Fake reviews** generated to manipulate ratings
- **Misinformation** spread at scale
- **AI-generated art and music** churned out without genuine creativity or craft

For content creators, this is genuinely threatening. For internet users, it's exhausting — sorting through increasingly large amounts of low-quality or misleading material.

What can be done?

- **Platform-level:** Websites and search engines need to develop better detection and filtering systems
- **Regulation:** Governments are beginning to introduce rules around AI-generated content disclosure
- **User education:** Learning to recognise AI-generated content is becoming an essential skill
- **Quality as a differentiator:** Ironically, authentic, human-made content is becoming *more* valuable precisely because AI slop is so prevalent

As a user of AI, you have a responsibility not to add to the problem. We'll talk more about ethical use in Essay 6.

The Singularity — Is It Real?

You've probably heard of the **Singularity** — that moment in science fiction where AI becomes so advanced that it triggers an uncontrollable explosion of intelligence, transforming the world beyond recognition. It's a compelling idea. It's also deeply speculative.

What the Singularity actually means

The term was popularised by mathematician and science fiction writer Vernor Vinge, and later by inventor Ray Kurzweil, who famously predicted the Singularity would arrive around 2045. The idea is that once AI reaches a certain threshold of intelligence, it will be able to improve itself, leading to a runaway feedback loop of ever-increasing capability.

Why some people take it seriously

People like Nick Bostrom (philosopher at Oxford) and the late Stephen Hawking have taken seriously the possibility that superintelligent AI could pose an existential risk — not because AI would be malicious, but because a superintelligent system pursuing any goal could be dangerous if that goal misaligns with human values.

Why sceptics are unconvinced

Many AI researchers think the Singularity is science fiction dressed up as prophecy:

- There's no clear evidence that human-level AI is achievable in the near future, let alone superintelligence
- The idea assumes intelligence can improve itself without limits — but physical and computational constraints exist
- Human intelligence itself doesn't lead to runaway self-improvement, so why would artificial intelligence work differently?
- Predicting specific timelines for these things has a poor track record

The honest answer

We don't know if the Singularity is possible. Many experts think it's not, or not soon. But we also can't rule it out completely. That's why the field of **AI safety** — focused on ensuring AI remains aligned with human values — exists. Even if the Singularity never happens, the work of making AI safe and beneficial is enormously important.

How Should We Act? — Ethical AI Use

So given all of this, how should we *actually* behave? Here are some principles for ethical AI use:

1. Be honest about what AI is

Don't pass off AI-generated work as your own if it matters. Don't pretend AI is something more (or less) than it is.

2. Think about bias

AI reflects the data it was trained on. That data comes from a world with real biases. Be thoughtful about where AI might introduce or reinforce unfairness.

3. Don't use AI to harm others

This should go without saying, but: don't use AI to impersonate people, spread lies, create non-consensual imagery, or manipulate vulnerable people.

4. Use AI to augment, not replace, your own thinking

AI is a powerful tool. But your judgment, creativity, and critical thinking are still essential. Think of AI as an assistant, not a replacement for your mind.

5. Consider the environmental impact

Training and running large AI models uses a lot of energy. For smaller, simpler tasks, consider whether a local model (Essay 3) could do the job with less environmental cost.

6. Stay curious and engaged

Don't just use AI passively. Learn about how it works, what its limitations are, and what debates surround it. Engaged, informed citizens are the best defence against AI being misused — by corporations, governments, or anyone else.

What Have We Learned?

- AI will change the jobs landscape — disrupting some sectors while creating others, with real questions about who benefits
 - Human fears of AI are understandable and often deeper than they first appear — touching on meaning, agency, and control
 - AI slop is a real problem: low-quality AI-generated content online is pervasive and affects everyone
 - The Singularity is speculative but taken seriously enough by serious people that AI safety research is genuinely important
 - Ethical AI use is everyone's responsibility — not just the companies building it
-

Glossary of Terms

Term	Definition
Automation	When technology (including AI) performs tasks that were previously done by humans.
AI Slop	Low-quality, AI-generated content produced at scale, often misleading or optimised purely for engagement rather than quality.
The Singularity	A hypothetical future moment when AI surpasses human intelligence and is able to improve itself, leading to unpredictable and potentially irreversible consequences.
Existential Risk	A threat to the continued survival of humanity as a species. Some researchers argue unmanaged AI could pose such a risk.
AI Safety	The field of research focused on ensuring AI systems remain aligned with human values and don't cause harm — even as they become more powerful.
Bias (in AI)	When an AI system produces results that are systematically unfair, reflecting the skewed data it was trained on.
Alignment	The challenge of ensuring AI systems pursue goals that genuinely reflect what humans actually want and value.
Disruption	Significant, often rapid change to industries, jobs, or ways of living caused by new technology.
Economic Inequality	The gap between rich and poor. Technology can make this better or worse depending on how its benefits are distributed.
Kurzweil (Ray)	A prominent futurist and inventor who predicted the Singularity would occur around 2045.

Further Thinking

- If an AI can do your job, what remains that only a human can do? Is that question important to you?
- Should AI companies be legally responsible for the harm their products cause? Who else, if anyone?

- Is "AI slop" mainly a problem for society, or does it affect you personally?
-

Essay 4 of 8 — Understanding AI Course Authors: Bea Groves-McDaniel and SAL-9000 Licensed under Creative Commons (NC, ND) — Share with attribution, no commercial use.