

## TOPICAL SEMINARS

Discussion of Artificial Intelligence – April 10, 2024

BY Pete Stoehr ([peterstoehr1946@gmail.com](mailto:peterstoehr1946@gmail.com))

“John Henry said to his Captain:

‘A man ain’t nothin’ but a man.

But before I let that steam drill beat me down,

I’ll die with a hammer in my hand.’ ”

John Henry beat the drill. He then laid down and died, with his hammer in his hand.

No one has beaten the drill since.

The printing press, vehemently opposed by those who could read, eventually spawned, or at least hastened the development of, by centuries, the steam engine, which led to the internal combustion engine, mass production of electricity, atomic power, radio, TV, the internet, and artificial intelligence.

Each of these innovative technologies was deeply criticized.

Had the printing press not been invented when it was, Friederich Drumpf, a newly trained 16-year-old apprentice barber in Kallstadt, Germany may not have heard of America and would not have left home to avoid the military draft. He would not have changed his name to Frederick Trump, become a citizen legally in 1892 (**because he had entered the country under the age of 18**). His grandson, Donald, would probably be a barber (or rat catcher) in Kallstadt, or more likely not have been born at all.

My conclusion is that the printing press led to both good and bad things. Its invention was inevitable, (The Chinese had moveable type by the end of the first millennium.), and the world is generally a better place to live in today than it was 1000 years ago. The same will be true of the development of artificial intelligence, regarding societal impact.

The YouTube interview by Economist editor Zanny with Deep Mind co-founder Mustafa and AI skeptic Yuval illuminates many of the potential dangers and benefits of AI. If we think of society as separate from humanity, the benefits are obvious to anyone, I think.

Yuval's assertion that a successful, autonomous AI would mean "**THE END OF HUMAN HISTORY.**" I think if he had added "**AS WE KNOW IT,**" he may have been correct, but that has happened before, and we're still here. Mustapha counters that the goal is universal abundance, not jobs, and that jobs are just a production tool to achieve that goal.

I tend to agree with Mustapha, although see the article by Peter Coy that I've included at the end of this paper,

There is a risk, especially to the political system, of loss of trust in what we see and hear. If trust collapses, so will democracy.

Regulation is obviously needed to prevent AI from achieving "Agency," or the right to self-regulate. The problem is "By Whom?"

The possibilities seem to be:

- **Self-regulation by the industry-** This has begun, to a degree, and it is in the industry's self-interest in the long term. In actual practice, as we have seen with OpenAI, short term profits and investment return will rule.
- **Voters-** My notes say "Ha,Ha2."
- **Congress (or the EU, China, Great Britain, etc.)** – Congress has a great record on Gun Control. The EU has been working on its "AI Act" for 3 years, but the Act is attacked by all sides, China has regulated effectively at the citizen level, but won't restrict use by its leadership or the military. Great Britain has already announced it won't regulate AI for now. Germany, France, and Italy, like Great Britain, fear that innovation would be stifled, and investment would go elsewhere. In other words, "**Show me the money!**"
- **An international Agency,** like the UN- Not in the real world
- **Executive Action-** The Biden administration released an order in October titled "On the Safe, Secure and Trustworthy Development and Use of Artificial Intelligence." It lays out a simple framework to regulate AI at some point in the future. This is the most practical path, although there are many problems. One is the risk of Authoritarianism, as in China. Another is the lack of permanence- executive orders are easy to repeal.

We must also consider the risk of the lone, angry teenager sitting in his bedroom with a laptop. That is beyond the scope of this paper.

Another ethical consideration I'll mention without going further is "Who decides what the safety norms are regarding morality, DEI, etc.?" AI might conclude that the answer to a declining birthrate is Free Love, coupled with free housing and freedom from other needs. It might have a cost-benefit analysis proving its point. Does that make it right for society, and the people involved?

Along with that is the question "Who decides the end goal for an AI project? Mustafa mentioned in the video that Americans might not like being taxed to send food to Africa.

Peter Coy is a non-economist who writes as an observer of economics in the New York Times. His April 5 Opinion column looks at a seldom discussed potential problem:

**"What would life be like if artificial intelligence solved all your problems?"**

You wouldn't have to work anymore. You could get any luxury you wanted by raising an eyebrow. Feelings of pleasure would wash over you constantly. Death would become almost optional because you could take on digital form and keep going for a billion years.

It sounds awful to me."

**Mr. Coy's entire column appears on the next page.**

**Ezra Klein** wrote an excellent commentary titled "The Lesson of the Open AI Mess." (New York Times, Sunday November 26, 2023) He concludes "But if the capabilities of these systems continue to rise exponentially... there is no off switch."

**DISCLAIMER:** I don't like the paper I've written. I haven't read any of the others yet, I hope they did better. I did not utilize Artificial Intelligence in writing this. Unfortunately, I also didn't use enough Native Intelligence.

Instead of following facts through a funnel to a clear conclusion, I feel that I've paddled down a little creek into a bay and crossed the bay into the ocean.

More concise conclusions:

- The growth of AI is inevitable and largely uncontrollable (see Twitter, etc.)
- In 50 years, it will be judged mostly positive, with problems (smart phones)
- Not much obvious change in society for 10-20 years
- In 20 years, I will have used up all or most of my remaining life. In 50 years, my children will have, and my grandchildren will be well past their primes.

# Life without problems to solve sounds great until it starts sounding awful



By Peter Coy

What would life be like if artificial intelligence solved all your problems? You wouldn't have to work anymore. You could get any luxury you wanted by raising an eyebrow. Feelings of pleasure would wash over you constantly. Death would become almost optional because you could take on digital form and keep going for a billion years.

It sounds awful to me.

We human beings dislike our problems, naturally, but if we had no problems to solve, what meaning would life have? How long would we be happy sitting on top of our big rock candy mountains? For me, one long weekend would about do it. A billion years of perfect happiness would be perfect misery.

Nick Bostrom, a philosophy professor at Oxford, wrestles with this question in a fascinating book that was published on March 27, "Deep Utopia: Life and Meaning in a Solved World." Solved here means solved the way the game of tic-tac-toe has been solved: If you move first, there is no reason you should ever lose.

A lot has changed in A.I. in the 10 years since Bostrom's last book on the topic, "Superintelligence: Paths, Dangers, Strategies." A.I. is folding proteins, generating art and blowing minds. The idea that it will change the world has gone from a nerdy obsession to conventional wisdom.

What we still don't know is how A.I. will change the world. It could enslave or kill us all. But for "Deep Utopia," Bostrom smartly chose to imagine the opposite: That A.I. does exactly what we ask it to do.

It's not a pretty picture. "Basically," he writes, "we're unfit to inhabit a perfect world."

People could still “work” at “jobs” in “offices” even if A.I. becomes able to do everything better than people can — a scenario examined by the economist Pascual Restrepo of Boston University in [one of my newsletters](#) last month. But the jobs will really be more like harmless hobbies, Bostrom writes.

Fine, you might say, people can fulfill themselves in other ways, such as shopping or gardening or child care. Bostrom punctures that balloon as well. One reason people find satisfaction in such activities is that they pride themselves on being good at them — knowing how to find the right throw pillow or grow gardenias or dry a child’s tears.

But A.I. will be able to do all of those things better than you can, Bostrom says. “It is quite conceivable that, at technological maturity, you would in fact be slightly harming your child whenever you indulged in some D.I.Y. parenting,” he writes.

I predict that prediction will infuriate a lot of people. I’m appalled by it myself. Bostrom, though, is not aiming to please. He clearly prides himself on following trends in A.I. to their logical conclusions, even when that’s discomfiting. He’s a jokey, quirky polymath who has a doctorate in philosophy but has also studied computational neuroscience, mathematical logic, astrophysics and general relativity.

Although he says we’re not cut out for perfection, Bostrom does explore ways that life in utopia could be made bearable after all. For instance, we could ask the computers to give us absorbing challenges. He cites a science fiction novel, “Permutation City” by Greg Egan, about a man who has achieved immortality. To avoid boredom, he programs himself to be seized, at random intervals, by fresh passions. At the moment of the novel it’s making table legs, of which he has manufactured 162,329.

That feels more like hell than heaven. Another vision of the future, even more repulsing, is that we become aimless blobs, perhaps hooked up to electrodes for jolts of pleasure.

Bostrom asks us to suspend judgment, though. Maybe we’d like it. “The question before us here is a very different one: *not* how interesting a future is *to look at*, but how good it is *to live in*.”

Bostrom told me on Friday that he sees the book as an exploration of the A.I. future, not a conclusion about it. “It’s going to be ultimately a very deep question what we choose,” he said. There is, he said, “a hunger for people to think about these questions.”

He likened the book to a particle accelerator that smashes atoms together to study their parts, such as quarks. In “Deep Utopia,” he said, he smashed values into one another to study their composition.

Toward the end of the book, a professor named Bostrom who very much resembles the author promises to reveal the secret of life. Just as he gets to that point in the lecture, the dean tells him to wrap up because the venue has been booked only until 6.

“Well,” the Bostrom character says, “I guess that’s that.”