

Dear Professor Belli and readers,

The subject of advanced artificial intelligence is one that has fascinated me since I was ten years old, poring over publicity material for Terminator 2, a film I was not allowed to see. I marvelled at the special effects in the trailers, collected the trading cards featuring images of scenes from the movie, and, most importantly, memorized the plot points and went through them over and over in my mind. What was this “Skynet” thing? How could it be so powerful? Did people really create another kind of life that was now going to destroy them?

This fascination with the creation of a new, alien form of life stuck with me through the decades. In the mid 2000s, I discovered the works of Oxford professor and philosopher Nick Bostrom. Bostrom’s early studies had been on the subject of existential risk; he developed theories and calculated statistical likelihoods of events that could threaten the continued existence of humanity. One of the greatest risks he discovered -- greater than a massive asteroid hit or a nuclear armageddon -- was the development of an advanced artificial intelligence.

I was enthralled.

For the last fifteen-or-so years, I’ve pored over Bostrom’s work, as well as that of his associates and other leaders in both academia and industry. Many have arrived at the same series of conclusions:

- Artificial intelligence, despite its fits and starts, has markedly improved over the decades.
- In time, artificial intelligence, whether conscious or not, will be greater than that of a human.
- Without proper guidelines and safeguards, artificial intelligence has a better likelihood of killing us all than not.

It is that final point that caused me to choose this topic for my presentation. In retrospect, it might have been a poor one because of its breadth.

As I’ve prepared this presentation, I kept getting hung up on the need to detail every little thing I introduced. While this would be fine for an unconstrained academic treatise, it’s a killer for a time-bound presentation. There are parts I’ve had to skip or just briefly touch on that would ideally need ten minutes of presentation time on their own. To make matters worse, I, personally, think this subject is a matter of significant importance -- but I don’t know if my audience will all come away thinking that’s the case.

While this sounds dire, it is far better now than it was when I started. The editing, revision, and peer review process allowed me to find and integrate better examples into the presentation. The first bit about the paper clips brings the much-needed “so what” aspect into play right away, without the audience having to sit there waiting for it to come up.

The “so what” part was hard for me to articulate at first, and that’s a problem for a topic I claimed is a matter of significant importance. How can I think something is important and not even be able to explain why? The individual conference I had with Professor Belli helped get that notion into my head, and while it was unpolished in my lightning presentation, I believe I’ve captured it (or at least most of it) in the final version.

I will be honest and state that this presentation is not my best work. It is not due to a lack of effort or time, but, perhaps, it is a demonstration of my unpolished ability to articulate importance. In my head, and I’m sure this is not unique to me, I assume most things I’m passionate about just make sense. Once that starts getting poured onto a page, though, those ironclad connections bridging ideas fray. It’s hard to connect one thing to another in practice.

Ultimately, however, this made me a better thinker, as well as a better presenter. By forcing me to examine ideas I have and reinforce them with real research and synthesize them into a presentation, I feel I am able to grasp the ideas better than I had before. I hope my presentation can reflect that.

Sincerely,

Max Lobdell