Slide 1: Introduction

- Hello everyone, my name is Bushra and I'm a Data Science major at Millikin University. Today, I'm here to talk about Artificial Intelligence, and namely, why research on AI shouldn't necessarily come to an end but instead, head in a different direction.
- Engage (they're probably super busy, so make this presentation worth their time): For starters, let's talk about you. Each of you are in different respective fields in the tech industry, and I'm curious to hear your opinions before I begin so I can see if I can change your mind by the end.
- (Expert audience offers opinions, some of their findings, etc.) (I respond to each) (In the meantime, the video background in my title slide continues to play) That was all excellent input! Now allow me to show you my perspective on AI.

Slide 2: Agent 57

- It's no surprise to you that Artificial Intelligence is a threat, with its potential sentience and the speed with which it learns. Right now, the most concerning advancements are being done in a vacuum, where the seed for AI's learning is planted but everyone is unaware of what exactly is happening, code by code, in its neural networks.
- DeepMind is at the frontier of such advancements, and it seems innocent enough so far, allowing the AI to iterate through trials and trials of games to find out how to play, how to win, and how to win most efficiently. But you have to consider that you are already introducing the AI to the concept of winning and losing, and a way of keeping score.
- In 2014, Demis Hassabis said that the conceptual layer of intelligence was missing for Agent 57 here, and yet we continue to surge forward without considering that this is the

most important layer of intelligence for something that could one day surpass us. Its understanding of concepts is crucial to safe development and a controlled path to general intelligence.

Slide 3: Technical Safety

- DeepMind claims to maintain a safe methodology of development, and yet some of the games that Agent 57 has played involve guns and shooting, and knowing how fast it learns and how efficiently, we should be concerned as to what 57 would do in a non-game setting.

Slide 4: Drone

- Did we forget about the Pegasus X-47B? Imagine Agent 57 driving contraption, providing it with independence and destructive power. That is far too much to give to the responsibility of something that has yet to be receptive to conceptual intelligence.

Slide 5: Atlas

- Or did we forget about Boston Dynamics? All these mobile robots are getting lighter in weight, stronger, quieter, and more agile, able to cross rougher and rougher terrains far better than humans. Agent 57 wearing one of these, immune to bullets, no longer dependent on human interaction; that is dangerous, and we are blindly inching closer to such a possibility.

Slide 6: More Examples

- AlphaGo, Sophia, Spot, Self-driving cars, and more… technology will not wait for us, we are being pulled on a leash with it, into an uncertain future. What could possibly lay in such a future?

Slide 7: The Technological Singularity

- The Singularity. A point in time in which machine intelligence will surpass our own in almost every domain. I understand the skepticism some of you have, as to if this is even a real threat and if it will ever occur, but let's consider a couple of possibilities.

Slide 8: Technological Utopia

- Say you're a technological utopist. The Singularity occurs, but it was the best thing to ever happen to humanity. We've cured cancer, poverty, world hunger.. Global warming has been dealt with, and all the math problems in the world have been solved. Overpopulation is under control, we never run out of food, and people spend their days feeling nothing but happiness.
- If the Singularity doesn't occur, we might end up somewhere close to this dream future anyways, just without a superintelligence to solve everything, and instead, we solve them ourselves. Humans develop the cures and the solutions to all the world's problems, and while we still use machines and continue to innovate, no Singularity has occurred.
- Both of these seem reasonable enough, and they sound pleasant compared to the world we live in now.

Slide 9: Technological Dystopia

- Let's consider the dystopian outlook. All those sci-fi movies, like the Matrix, and Terminator have come true. We barely stand a chance against our AI overlords and robot soldiers, and the human population is dwindling. We are of no import to beings with far superior intelligences, just like how animals are only second to humans in all our considerations. Due to our arrogance today, we suffer tomorrow.

Slide 10: Chart

- Out of the three, which do think is more likely? (collect audience responses and counter)
- This chart was made around 2015, so it's five years old at this point. We haven't crossed into thinking machines quite yet, as it predicted, but how long do you think it will take for us to reach that inevitably? True Artificial Intelligence exists today, we have created it already, there's no going back. All we can do is shift our research and advancement efforts to keeping it under control, and knowing that control will only last so long, developing an embedded respect for human life, as well as our values and goals.
- This solution is the only way I see all of our hard work and time being worth anything. If we continue to progress for the sake of progress, we may be dealing with the end of life as we know it and the extinction of our species.
- As terrifying and detrimental that all sounds, it's only the Singularity that I believe is inevitable, not a destructive future. What I mean by that is that we can at least attempt, starting now, to reach the utopia that we want, and while that begins with safety measures, controlled databases, and rigorous coding, we also have to take into account public interaction and the diverse array of applications.

Slide 11: Applications

- Based on the results of my research, there is a lot of room for error right now in AI, and that is the problem we must fix first. Racist bots, inaccurate photo recognition, crashing self-driving cars, all of these mistakes are big blows in public acceptance of AI. We have to begin troubleshooting and perfecting simple functions like appropriate social media interactions, correct photo recognition, and safe driving so we can begin to work towards

our utopia. AI learns from humans, and knowing that we are not perfect, we should be mindful of what we feed AI. Humans should have at least some level of trust in technology for the relationship to be successful, but that trust should not mutate into dependence either, so as the developers of AI, we have to maintain the balance.

Slide 12: Conclusion 1

- Having addressed AI's threat, but also its unlimited potential, I hope I've changed the minds of those of you who may not have thought of AI as a substantial entity. We are playing with the inner workings of something revolutionary, and we should not take such a task lightly, but I don't expect us to shut down AI development and technology production due to the fear of the unknown either. We have evolved so much and so far in the last decade than mankind has in the past century, and using our human intelligence, ingenuity, and conceptual reception, we can come up with a solution to the threat of AI without having to lose its benefits as well.

Slide 13: Conclusion 2

- Time isn't slowing down for us, the world won't wait. If the frontrunners of AI technology do not take this responsibility, then someone in their basement might stumble upon a Singularity breakthrough with no idea what to expect or what to do. We don't want to risk such an event, and that is why I'm striving to convince you all to implement safety and some level of understanding of the human life in your respective fields.
- Together, we can reign the intelligence we created and step into a cooperative and safe future.

Slide 14: End - Thank you! (take questions now)