Slide 1: Introduction

- Hello everyone, my name is Bushra and I study Data Science at Millikin University. Today, I'm here to talk to you all about a topic that I'm sure you're at least a little bit familiar with — Artificial Intelligence. For the rest of this presentation, I will be using the acronym AI when I'm addressing Artificial Intelligence.

Slide 2: Survey

- For a little self-reflection, let's do a quick activity.
- (pass out slips of surveys)
- On the paper you've just received, look through each of the different technologies mentioned and check the box that you think is true — either yes, it exists today, or no, it does not exist today. I'll give you a few minutes to do this.
- (5 minutes or until everyone is done)
- Alright, now I will reveal the correct answers, and I want you to check your sheets as I do.
- (Answers appear on the slide as I click, and I explain each example - and then I note that the ones that do not yet exist are in development)

Slide 3: Science Fiction

- In a society that is obsessed with science fiction, it may confuse most people where to draw the line between it and reality.
- What do you first think of when I say "artificial intelligence" or "robot"?
- (audience responds, hopefully someone mentions "terminator", "the matrix", "avengers", or "ex Machina", and the corresponding images appears on my click)

- As awesome as this all looks, it is far more fiction than fact.

Slide 4: Reality

- So let's talk about reality then. I'm going to now provide a few examples.
- The AlphaGo program, a game-playing program that has successfully defeated the human world champion — a feat thought impossible until recently. Funnily enough, AlphaGo was defeated by its descendent AlphaGo Zero, with a score of 100-1.
- Next, we have a terrain robot, named Big Dog. This is only one of many Boston Dynamics creations that are being taught to walk on different kinds of surfaces and maintain their balance.
- Now, before the last example, I'm going to distinguish the difference between Narrow AI and General AI. Narrow AI is what we have today, AI with specific tasks that are coded into it to be completed, like walking, playing a game, etc. The closest we have to a General AI, the human-like, more intelligent kind, is Sophia, a social robot. She has been known to hold amusing conversations with those she interacts with, and is even an honorary citizen of Saudi Arabia.

Slide 5: Development

- This seems all good and dandy until we start looking into the implications of these developments.
- Take Atlas here, for example. He seems harmless enough, struggling to walk and not nearly as good as a human. Now, what if I told you that this is being considered for military use? That it's owned by DARPA, the Defense Advanced Research Projects

Agency. You begin to wonder, what could they possibly have in mind for Atlas? Is it going to start replacing humans in armies? Hold guns? Kill people?

- Not to mention, Google is buying out robotics companies like these and it begs the question of why it is one of the most powerful companies in the world, with most of the information in the world, is interested in military robots.

Slide 6: Drones

- Remember the self-flying war drones we talked about? Here it is.
- The Pegasus x47b, a military drone. Yet again, AI is being used in programs meant for defense. This creation has cameras on the top that record the sky and project it onto the LED screen below, completely camouflaging it. The Pegasus can fly into enemy territory, deliver an attack, and return all on its own. I question whether a machine should be given this much power because with great power comes great responsibility, and is it really wise to give so much power to something that doesn't yet comprehend that?

Slide 7: History

- To provide some context, let's follow the development of technology a little. The first electro-mechanical computer is created by Alan Turing, which then allows research programs to pop up in the 1950s. Over time, computers that took up entire rooms shrink to the size of our palms, and the processing power multiplies exponentially. Exponential growth is something humans have trouble comprehending because it is so vast and grows so fast, but that is exactly the case with computers.
- We're past creating AI, now what? Deep neural networks are being created, called Deep Learning, and the scientists monitoring this don't even fully know what is going on in

these networks. Who knows what will happen at this point (the thinking machine)? Who knows what will happen afterward?

Slide 8: Singularity

- That's what we call the Singularity — the point in time in which artificial intelligence surpasses that of humans, being able to perform all tasks better than we can.
- Imagine that. A being that is smarter than the smartest being on Earth. Think about how we treat animals. We are smarter than them, and so we choose which ones to slaughter for food, which ones to domesticate, which ones to breed, which ones are pets and so on. We drove several animal species into extinction, and this is before we had all the technology and tools that we do now. We've never experienced the short end of the stick because we were always the smarter animal, but with a superintelligence… imagine that.
- Now with such an intelligence, there is a possibility for a good outcome. The superintelligence could find solutions to large problems such as world hunger, global warming, overpopulation, coronavirus, and cancer that humans never would have thought of. But this is wishful thinking, and based on the way humans use their heightened intelligence, our guess for what the superintelligence will do is far more negative.

Slide 9: Evidence for negative outcome

- Why can't we hope for the best? Well, humans are flawed beings and artificial intelligence mimics us, and therefore, our flaws as well. Take this AI for example.
- Tay, a Microsoft AI chatbot, became extremely racist, sexist, anti-Semitic within 24 hours of being online. You can see to what extent in these images. Tay learned this

discrimination from humans. It had to learn how to be offensive first by looking through other people's tweets, and then mirrored the same negative behavior when it was active.

- Imagine if a more powerful AI displayed these discriminatory qualities? Adopting human's discriminatory tendencies?

Slide 10: " "

- Another example is Google's first-generation photo-labeling system. It misidentified an image of African Americans as gorillas. Comparing people of color to gorillas in order to make them feel inferior and subhuman was a tactic practiced by European colonists, and that historical detail reflects in technology developing today. I think this says a lot about our people as a society, and how we have still yet to fix our own flaws, much less our machine creations.

Slide 11: Medical

- AI isn't just socially dangerous, but also medically dangerous. You would think that a robot that is trained would perform better than humans at certain medical tasks, and while that may be true in some cases such as surgery, here we have another issue of misidentification. In an article I came across in my research, an AI program mistakenly diagnosed a fetus as having Down Syndrome due to white spots near the fetus' heart. It turned out that the white spots were due to a higher pixelation, and the baby was, in fact, perfectly healthy. The parents almost agreed to an amniocentesis or a process in which amniotic fluid is sampled, which is a risky procedure for the mother and the child. If this misidentification was amplified, then think about how many families could be put at risk.

Slide 12: Cars

- Now, let's talk about another place where misidentification can occur. This video shows Waymo cars, or Google's self-driving car project, in action as cars drive around with no one at the steering wheel.
- Again, there is room for positivity here. Cars with AI drivers have been statistically shown to reduce accidents significantly because they don't make mistakes like humans, but they do make mistakes. In 2016, 2017, and 2018, accidents and loss of life occurred when the AI mistook objects in its surroundings as inanimate. If this could be tweaked to perfection, then maybe we can have safer roads and a significant decrease in road collisions and such, but if these mistakes were made at large or on the interstate, the following loss of life and wreckage could be severe.

Slide 13: Arrive at the main point

- I could go on and on about all the various flaws and mistakes of AI that should make us think a little bit more before we start trusting technology a little too much, but I think these examples provide enough insight for you guys to get the gist of what I'm saying.
- I don't want to be a downer. I want to hope for the best. I mean, if all this advancing technology can lead us to a tech-mech utopia where AI and humans live side by side peacefully, then I'm all for it. It's just that this seems highly unlikely based on the incidents we've seen.

Slide 14: Conclusion

- Why am I telling you all this? Well, humanity is at the dawn of an AI Revolution. AI lives in our homes, in our phones - it walks the Earth alongside us right now. It is our responsibility to embed human values and respect for human life into AI before it

actually matters. My entire project is to prepare for the birth of a superintelligence, and if it doesn't care for human beings, then, to put it lightly, we are screwed. You might think, "oh we can just unplug it, just cut off its power,", but a superintelligence would know to connect to the Internet and there is no off switch for that.

- Of course, not every single one of us can immediately put a stop to what seems to be the inevitable end of humanity at the hands of AI, but we can all help by educating ourselves. Keep up with the advancing technology, know what is being developed and what is flying in the sky, what is cruising the Internet, and what is in your own home. Always be skeptical and careful with technology, and try to keep yourself updated on current events in the world.
- I guess what I'm saying is, don't let AI catch you off guard. Let's support AI regulation and research, so that we can find a way to gain the benefits of AI without risking humanity's future. Let's not let AI be our last invention. Let's create a safe future together.

Slide 14: End

- Thank you!
- (take questions)