AI, WAR, AND RESPONSIBILITY

35–40 Minute Presentation

---

1. Opening: Why AI Triggers Fear (5 minutes)

- When people hear “AI and war,” the immediate reaction is often fear rather than curiosity.
- That fear is understandable, because most public narratives frame AI as an independent force rather than as a human-built system.
- Popular culture encourages us to imagine AI as a “ghost in the machine” — opaque, autonomous, and potentially hostile.
- When something feels mysterious, we tend to project intention onto it.
- My goal is not to dismiss ethical concerns about AI, but to clarify where those concerns properly belong.
- If we misunderstand what AI is, we will misidentify the source of danger.
- This talk is about shifting the focus from *what AI might do* to *what humans choose to do with it*.

---

2. What AI Actually Is (and Is Not) (7–8 minutes)

- Artificial intelligence is not a single system, but a class of technologies designed to detect patterns in data.
- Modern AI systems are trained by adjusting large numbers of internal parameters to reduce error.
- These parameters do not represent beliefs, memories, or understanding.
- AI does not “know” things; it predicts outputs based on learned structure.
- It does not understand meaning, context, or consequence in the human sense.
- AI has no awareness, no intention, and no internal experience.
- Any sense of intelligence we perceive comes from coherence, not comprehension.
- This matters because responsibility cannot belong to something that does not understand what it is doing.

---

3. Syntax vs. Semantics: Why AI Sounds Intelligent (6–7 minutes)

- Linguistics distinguishes between syntax and semantics.
- Syntax concerns structure: how symbols relate to one another.
- Semantics concerns meaning: what those symbols refer to in the world.
- A sentence can be syntactically correct while being semantically empty.
    - “Colorless green ideas sleep furiously.”
- AI systems operate almost entirely at the level of syntax.
- They model how language is structured, not what it means.
- AI does not connect words to lived experience, intention, or consequence.
- When AI speaks fluently, humans instinctively supply the meaning.
- This creates the illusion of understanding without actual understanding.
- AI has syntax without semantics — and moral reasoning requires semantics.
- Understanding requires grounding.
- Grounding requires interaction with the world.
- Neural networks operate entirely within abstract mathematical space.
- Even when connected to sensors or images, the system still processes data, not experience.
- This is why AI can sound thoughtful while remaining morally blind.
- It can manipulate language about war, ethics, and suffering without grasping any of them.
- This is not a flaw in the technology.
- It is the way the technology works – it simulates understanding but has none.
- And it is why responsibility must remain human.

---

4. Learning vs. Theft: Clearing a Moral Confusion (5 minutes)

- A common claim is that AI “steals” from artists and writers.
- This claim often confuses learning with copying.
- Copying involves storing and reproducing content.
- Learning involves internalizing patterns and relationships.
- Humans learn by exposure and abstraction, not by storing exact replicas.

- AI learning is analogous in structure, though not in experience.
- The real ethical questions concern consent, compensation, and power.
- These are governance issues, not evidence of machine intent.
- Treating learning itself as theft leads to confusion rather than clarity.

---

5. The Practical vs. the Philosophical (4–5 minutes)

- Philosophical questions ask what should be done and who is responsible.
- Engineering questions ask what can be done and how reliably.
- These domains overlap, but they are not the same.
- AI has no internal moral life.
- It does not deliberate, regret, or reflect.
- Moral responsibility does not disappear when decisions are automated.
- It shifts upstream to those who design, deploy, and authorize systems.
- AI intensifies human choices; it does not replace them.

---

6. AI in Military Contexts: Precision and Risk (8–9 minutes)

- Most military AI is used in support roles such as logistics, maintenance, and intelligence filtering.
- These systems reduce human error and cognitive overload.
- Some AI systems provide decision support rather than decisions.
- Humans remain responsible for interpreting and acting on recommendations.
- Ethical risk emerges when humans over-rely on automated outputs.
- Autonomous systems are the most controversial application of AI in war.
- These systems operate within rules defined by humans.
- They do not choose goals, enemies, or justifications.
- The ethical issue is not machine autonomy, but human delegation.
- AI makes war more efficient.
- Efficiency magnifies both good intentions and bad ones.

---

7. Sentience, Survival, and the “Runaway AI” Fear (4–5 minutes)

- A common fear is that AI will one day seek to preserve itself.
- In humans, the drive to survive is biological and evolutionary.
- AI has no evolutionary history, no pain, and no subjective experience.
- AI systems have goals only because humans assign them.
- Optimizing a goal is not the same as wanting it.
- Intelligence does not imply consciousness.
- Treating AI as a future moral agent is a category error.
- Focusing on speculative sentience can distract from present accountability.

---

8. Conclusion: The Human Remains Central (3–4 minutes)

- AI is not a malicious actor.
- It is a force multiplier.
- It amplifies human competence, intention, and error.
- Used carefully, it can reduce harm and improve decision-making.
- Used irresponsibly, it can accelerate violence and injustice.
- The danger is not that AI will develop its own will.
- The danger is that humans will avoid responsibility by blaming machines.
- AI does not remove moral responsibility.
- It concentrates it.

Closing line

> “The ethical challenge of AI is not preventing machines from becoming human — it’s ensuring humans remain accountable for what they build.”

Discussion Questions (Choose 4–6)

Warm-Up

- Which aspect of AI concerns you most: the technology itself, or who controls it?

- Does automation change moral responsibility, or just obscure it?

Deeper Ethical Questions

- Are there decisions that humans should never delegate, even if AI could perform them better?
- If AI operates at the level of syntax rather than meaning, should it ever be trusted in morally loaded contexts?
- Does increased precision in warfare make war more ethical, or merely more efficient?

Reflective Questions

- If AI reflects human goals, what does its military use say about us?
- Are we more afraid of AI becoming powerful, or of humans avoiding accountability?