

Foundations of the AI Era

Course: GENS 176 | **Instructor:** Prof. Hundley | **Date:** Spring 2026

(This document was generated in part with Gemini)

I. The Big Picture

What is Artificial Intelligence (AI)?

The broad field of computer science dedicated to building systems capable of performing tasks that typically require human intelligence—such as recognizing patterns, making decisions, or translating languages.

- **Narrow AI:** Systems designed for one specific task (e.g., facial recognition, Netflix recommendations). This is what exists today.
- **AGI (Artificial General Intelligence):** A theoretical AI that can learn and apply intelligence across any task at a human level. It doesn't just "play chess"; it can reason, plan, and learn new skills autonomously.

What is Machine Learning (ML)?

A subfield of AI. Instead of a human programmer writing a rigid list of "if-then" rules, Machine Learning allows a computer to look at vast amounts of data, find patterns, and improve its performance over time.

The Metaphor: If Traditional Programming is a **recipe** (follow steps A, B, and C), Machine Learning is a **student** (looking at 1,000 solved math problems to learn how to solve the 1,001st).

II. The Mechanics: How it Works

What is a Neural Network?

The technical architecture behind modern AI, inspired by the structure of the human brain.

- It consists of layers of interconnected "nodes" (neurons).
- When the network is "trained," it adjusts the strength of these connections (called **weights**) until it can accurately turn an input (like an image of a cat) into the correct output (the word "Cat").

Large Language Models (LLMs)

A specific type of AI trained on massive datasets of human language (books, code, articles).

- **How they work:** LLMs do not "know" things in a human sense. They are highly advanced **prediction engines**. They analyze the context of your prompt and calculate which word (or "token") is most likely to come next.
- *Examples: GPT-4, Claude, Gemini.*

Reinforcement Learning (RL)

A training method based on trial and error.

- An AI "agent" is placed in an environment and given a goal. It receives **rewards** for good moves and **penalties** for bad ones.
- Over millions of attempts, the AI learns the most efficient strategy to maximize its reward. This is often how AI learns to play games or navigate physical spaces.

III. Logic and Reliability

Deterministic vs. Stochastic Models

Understanding this distinction is crucial for knowing when to trust an AI's output.

- **Deterministic:** A system where the same input **always** produces the exact same output. (Example: A calculator. $\$5 \times \5 is always $\$25$).
- **Stochastic:** A system that involves **probability and randomness**. Even with the same input, the output may vary slightly each time. (Example: A weather forecast or an LLM).

Why "Hallucinations" Happen

Because LLMs are **stochastic** and **predictive**, they sometimes prioritize "sounding likely" over being "factually true." When an AI makes up a fact, it isn't "lying"—it is simply following a statistical path that turned out to be incorrect.

IV. The "Singularity"

The **Technological Singularity** is a hypothetical future point in time when technological growth becomes uncontrollable and irreversible, resulting in unfathomable changes to human civilization.

The term is borrowed from **astrophysics**: In a black hole, a "singularity" is a point where the laws of physics as we know them break down. In AI, it represents a point where our ability to predict the future breaks down because the "intelligence" on the planet is no longer primarily human.

How it connects to AGI:

Most theorists believe AGI is the "trigger" for the Singularity through a process called an **Intelligence Explosion**:

1. **Step 1:** We create an **AGI** that is as smart as a human.
2. **Step 2:** Because it is a computer, it can work 24/7 at incredible speeds. Its first "job" is to design an even better version of itself.
3. **Step 3:** That "better" version (now smarter than any human) creates an even *more* advanced version.
4. **Step 4:** This creates a **feedback loop** where the AI improves itself exponentially in weeks, days, or even hours, leading to **Superintelligence (ASI)**.

V. Summary Table

Concept	Role	Simple Analogy
AI	The Goal	A robot that can think.
Machine Learning	The Method	Learning by observing examples.
Neural Network	The Structure	Digital "brain cells" wired together.
LLM	The Specialist	A super-powered version of "autofill."
