Generative AI:
From Creation To Action

AI CLUB WEEK #3

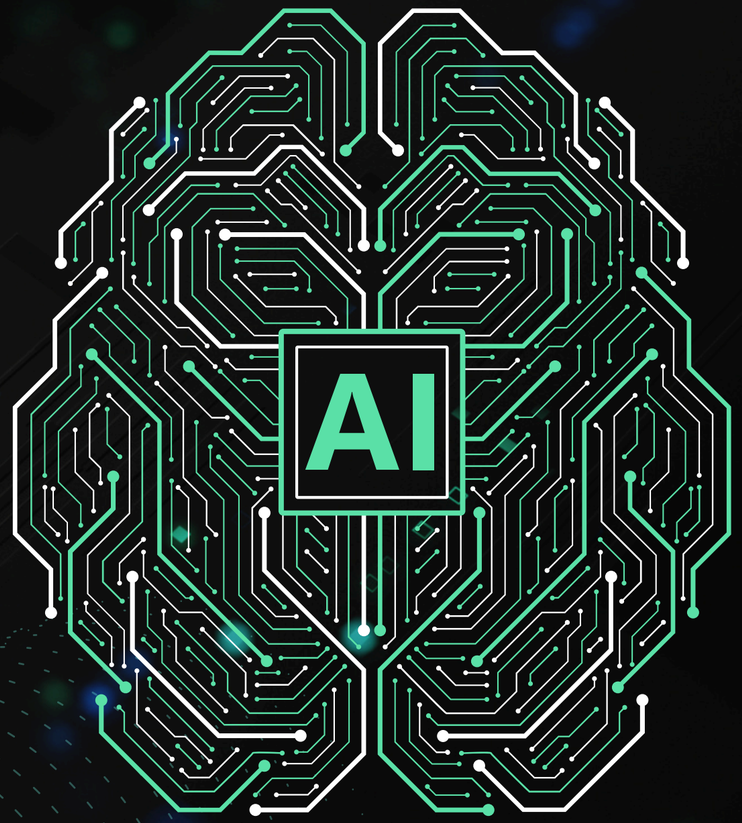

PRESENTED BY AREN, MICHAL, SERGI

Timeline

Week 1:
- How AI Learns
- Machine Learning & Neural Networks

Week 2:
- How AI Predicts & Writes
- Large Language Models (LLMs)

Today:
- How AI Creates and Acts
- Part 1: Diffusion Models (Image Generation)
- Part 2: Agentic AI (Goal-Driven Action)









TODAY

Agentic AI &
Diffusion Models

WEEK 4

Natural Language
Processing & AI in vision

WEEK 5

Ethics & the Future
of AI, Cyber Security

WEEK 6+

Creating Classifier
Models & Chatbots

How Diffusion Works

Step 1 (Training)

The "Forwards Process": Adding Noise

- Starts with millions of real images.
- It will progressively add “**noise**” (specifically **Gaussian noise**) to the image, converting the image into an unrecognisable mess of mismatched pixels called **static**.
- This process is **not learnt**, solely being based on mathematical processes lined up as instructions.

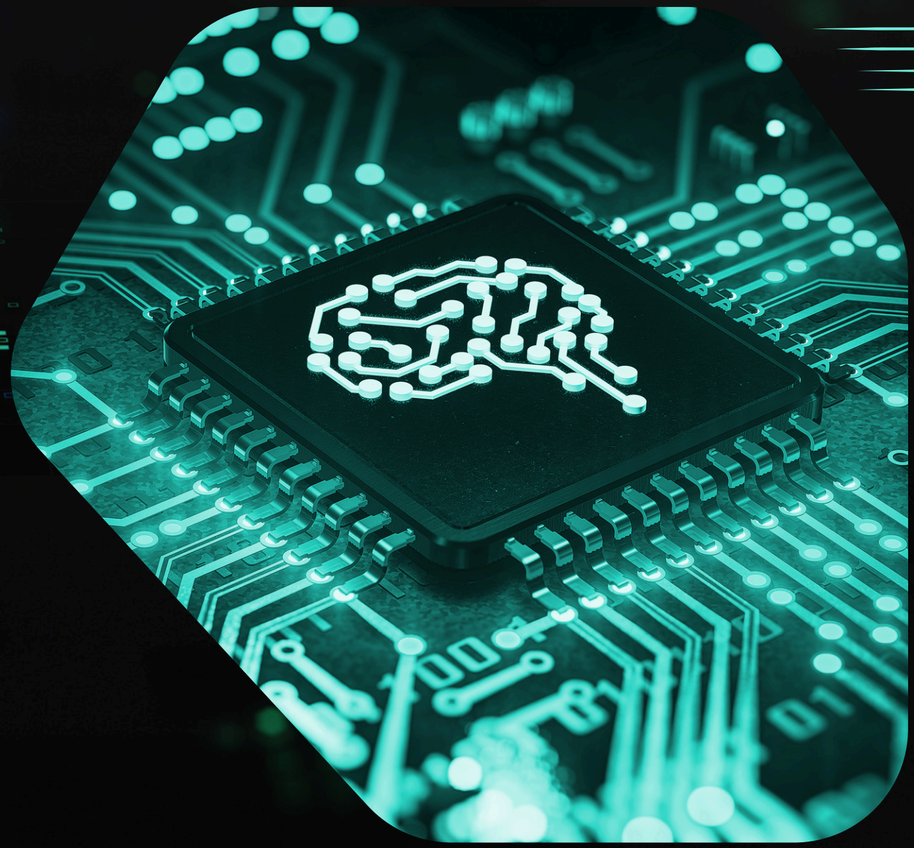

How Diffusion Works

Step 2 (Generating)

The "Reverse Process": Removing Noise

- It starts with pure static, created by the **addition** of noise
- The model **learns to predict** the formation of noise, as it is a **fixed** mathematical process
- Removing the predicted Gaussian noise from the provided image, clearing the image up
- It will then **begin to form** the original image, or more a **similar image** of similar characteristics

...So why does it have to go through all these steps?

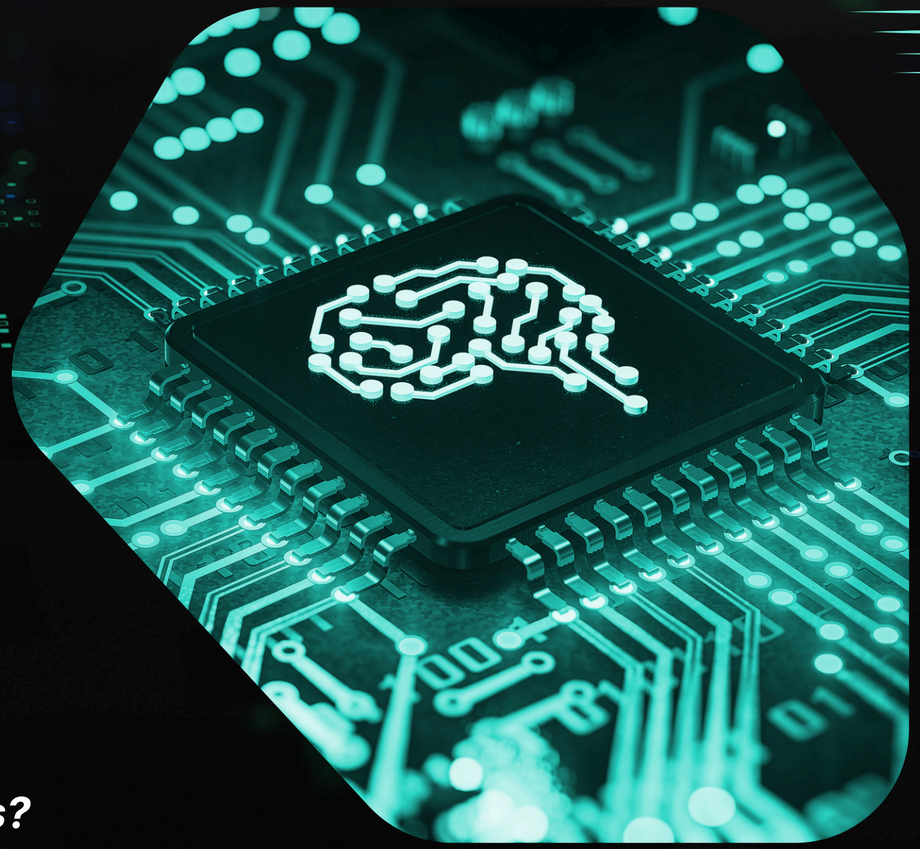

Why Diffusion Matters For You

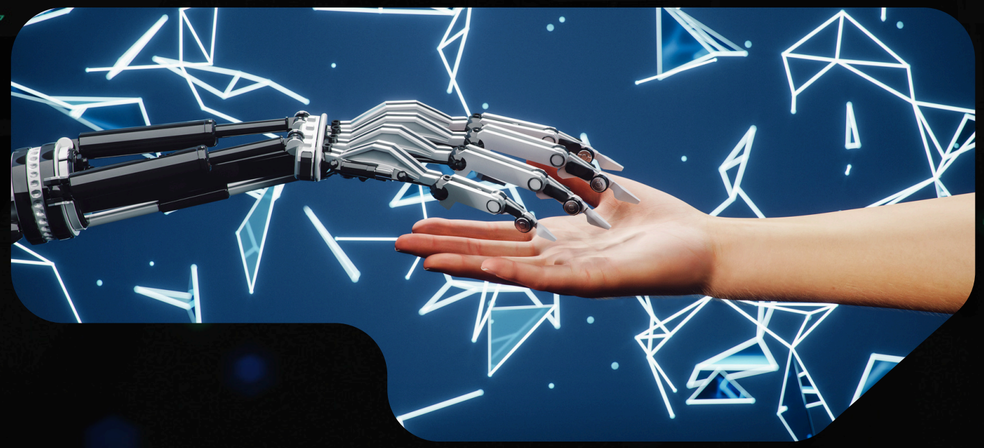

01 **Study:** Instantly create diagrams, historical scenes, or concepts for presentations (e.g. A labelled diagram of a heart for Biology).

02 **Creativity:** Design logos, artwork, or website assets in *seconds* (e.g. Probably ALL of the photos on these slides).

03 **Future Careers:** A core tool in marketing, game design, architecture, and media (e.g. Presenting a project to a team which contains images generated by AI).

From Creating To Acting

Generative AI

- Creates new content (text via LLMs, images via Diffusion).
- Responds to a **single, specific** prompt.

Generative AI

- Output is the **final product** (e.g. an essay, an image, a presentation)
- It's a powerful tool that waits for your command.

Agentic AI

- An autonomous system that takes action.
- Operates in a **multi-step** loop (**Think, Act, Observe**).

Agentic AI

- Output is a **completed goal** (e.g., a booked flight, a finished project).
- It uses the "brain" to achieve a goal on its own.

It's the difference between a *calculator* and a *mathematician*.

The Components of an AI Agent

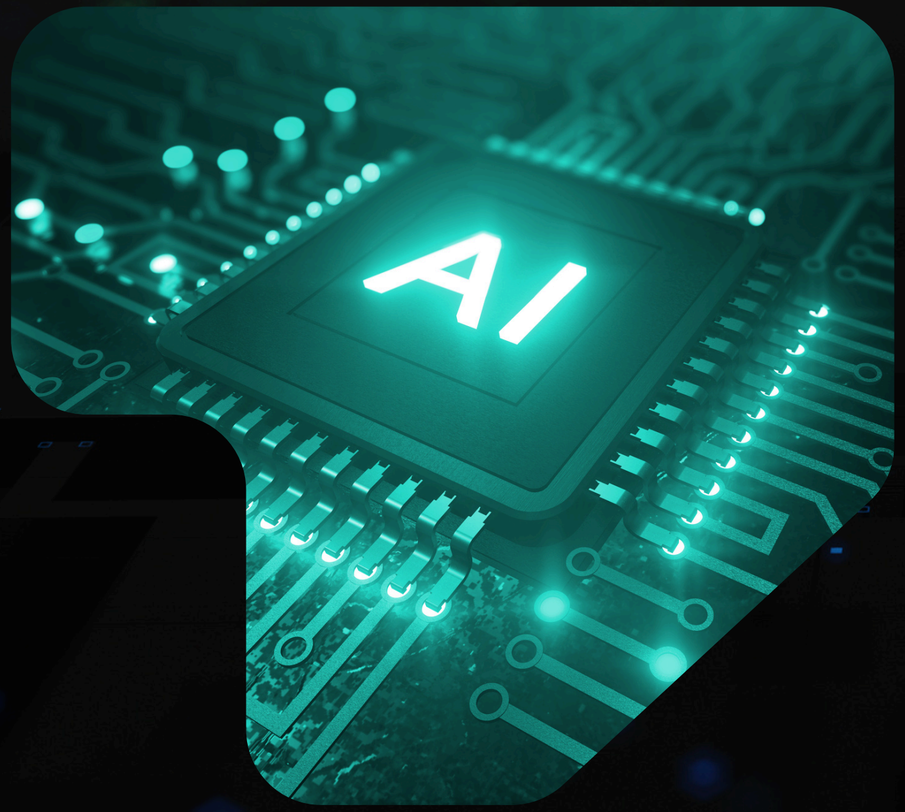

1. Goal (The Mission):

- A high-level, complex objective.
- e.g., "Plan a 3-day trip to Paris for under £400."

2. Brain (The "Reasoner"):

- An LLM (like in Week 2) acts as the "inner monologue."
- It plans, strategises, and reflects on its actions.

3. Tools (The "Hands"):

- A list of approved actions (APIs).
- e.g., Web_Browser, Code_Interpreter, Calendar_Access.

4. Autonomy (The Loop):

- The ability to run a cycle of "Think -> Act -> Observe" repeatedly.
- It self-corrects based on errors and observations.

How Agents Work: The React Loop

An agent's inner monologue follows a Think, Act, Observe cycle:

Goal:

"What's the weather in London and the traffic on the M25?"

01 Cycle 01

- **Thought**: I need two pieces of info. I'll get the weather first.
- **Action**: Web_Browser("weather in London")
- **Observation**: "14°C and cloudy."

02 Cycle 02

- **Thought**: I have the weather. Now I need the traffic information for the M25.
- **Action**: Web_Browser("Traffic on M25")
- **Observation**: "Crash at M25 causing blockage"

03 Cycle 03

- **Thought**: I have both pieces of info. I can form the final answer.
- **Action**: Final_Answer("The weather in London is 14°C and cloudy, and the M25 is full of traffic...")

Get Hands-On
Free Tools

AI That Creates (Diffusion):

- **Leonardo AI**: Free daily tokens for high-quality, complex images.
- **Stable Diffusion Online**: A simple, free web tool to experiment.
- **Canva**: Great for projects and presentations.

AI That Acts (Agents):

- **Perplexity AI**: An "answer engine" that acts as a research agent.
- **Gemini / ChatGPT**: When they "Browse the web" or "Analyse data," they are running a simple agent loop.

Practical Challenges:

- **CrewAI**: A Python library to build a "team" of agents that collaborate (e.g., a "Researcher" agent and a "Writer" agent).
- **Gemini** - Look into the Canvas Mode involving a live server to be able to run code in a small window (Preview).

So Wait...
What’s the catch?

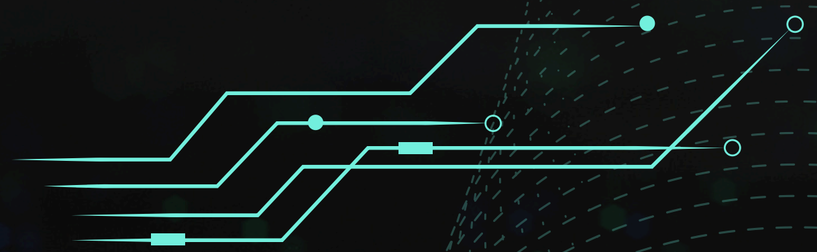

Diffusion Model Risks (Creation):

- **Misinformation & Deepfakes**: Creating photorealistic fake images of events or people that never existed.
- **Bias & Representation**: If training data is not diverse, models can amplify and reinforce harmful stereotypes.
- **Copyright & Artist Livelihood**: Models are trained on billions of images, often without the original artist's permission or compensation.

Agentic AI Risks (Action):

- **Security**: How to stop a malicious goal? (e.g., "Find a security flaw in the school's website and email it to me").
- **Reliability**: A chatbot hallucinating a fact is annoying. An agent hallucinating an action (e.g., booking the wrong £1000 flight) is a disaster.
- **The "Black Box"**: It can be hard to trace why an agent made a bad decision in a 100-step task.

What we learned:





We covered two aspects of Generative AI.

Diffusion Models (Creation): These are image generators. They work by starting with pure static and "denoising" it step-by-step, guided by your text prompt, until a clean image comes out.

Agentic AI (Action): This is the next leap beyond simple chatbots. It's a system that uses a LLM (Large Language Model) as a brain to autonomously achieve a goal.

The Core Loop: Agents work on a ReAct (Reason + Act) or "Think, Act, Observe" cycle, allowing them to use tools, see the results, and self-correct until their mission is complete.

Courses

Posted on Teams, let us know if you are not yet on it!

Guess the film:

I’m sorry Dave, I’m afraid
I can’t let you do that.