

**Szertics Gergely**

AI4Science program

6 December 2024  
Budapest



# **Gen AI 4 Science**

## Agents and foundation models

<https://hun-ren.hu/ai-4-science>

## Source and kudos

---

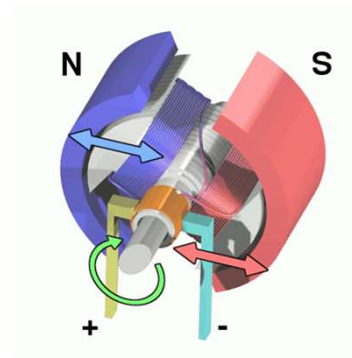


- Lecturer Frankfurt School of Finance and Management
- Deep tech leader, consultant and manager with special interest in artificial intelligence, cognitive sciences, data science and deep learning.
- Long time "Startupper" and CTO.
- Lecturer in applied Artificial Intelligence, tech leadership.
- Public speaker with interest in Buddhist studies, comparative religious studies and cognitive science.

# The complexity of AI enabled tools



USER INTERFACE / APP

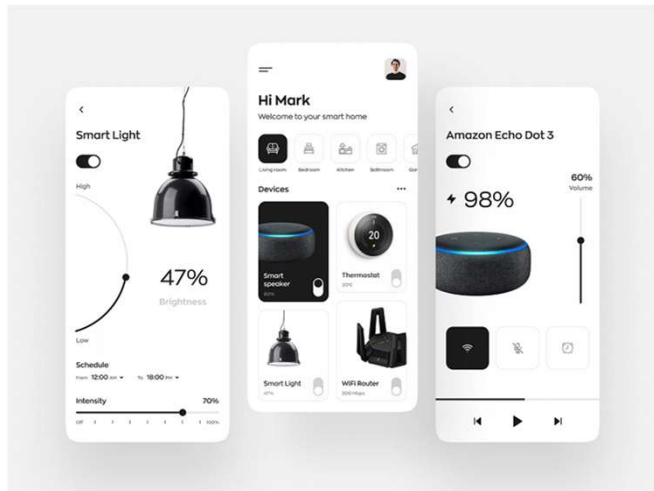


CENTRAL ENGINE

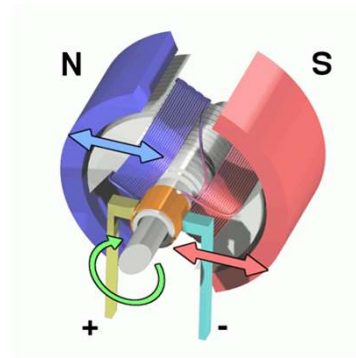


TASK TO COMPLETE

# The complexity of AI enabled tools



USER INTERFACE / APP



CENTRAL ENGINE



TASK TO COMPLETE

## The complexity of AI enabled tools

---



### APPLICATIONS

- chat.openai.com
- Copilot.Microsoft.com
- chat.mistral.ai/chat
- Perplexity.ai
- GenAI4Science
  
- NotebookLM.google
- Scispace.com
- Napkin.ai
  
- Github Copilot
- Cursor.ai



### LLM MODELS

- GPT 4o
- Claude Sonnet 3.5
- Gemini 1.5 Flash
  
- Mistral Large
- Qwen 2.5 – 72B
- LLaMa 3.2 – 8B



### TOOLS TO ACCESS

- Bing search
- Google search
- Elicit search
- Consensus search
- Code Interpreter
- Dalle 3

# How can we imagine LLMs?



## TRAINING

**SKIMMING** through tons of books, articles, research papers, github code...

**The Pile:** approximately 500 000 000 pages of text



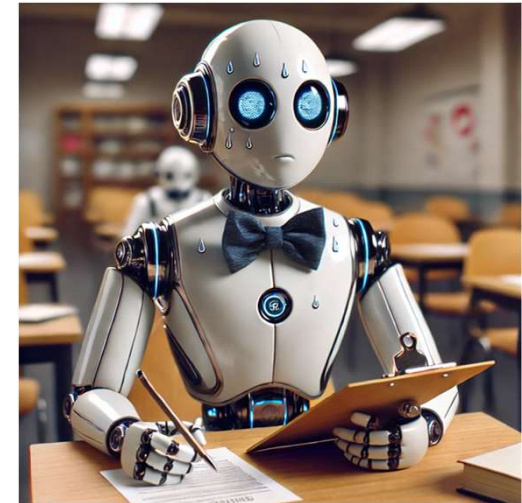
## GENERAL, WIDE, SHALLOW KNOWLEDGE

Whatever sticks from skimming through A LOT



## CAPABILITY OF USING LANGUAGE

Understanding instructions and context and writing in **any** style required



## EXAM

Whatever comes to mind!  
Always follow instructions!  
Read context given carefully!

**Transforming text to text in a most probable way**

## What are “agents”?



# Everyone is getting “agentic”

AI EFFECT

**Microsoft to roll out new autonomous AI agents next month, fending off challenge from Salesforce**



Salesforce unveils autonomous agents for sales teams



**OpenAI Tests Open-Source Framework for Autonomous Agents**

By Paula Parisi  
October 16, 2024



TECH

**Amazon-backed Anthropic debuts AI agents that can do complex tasks, racing against OpenAI, Microsoft and Google**

**ANTHROPIC**

**2024 October**



# A basic agent - a predefined role + tools

MISTRAL AI Le Chat La Plateforme Docs API

My GPTs + Create

Overview

REGISTRY

Agents

Fine-tuned models

Datasets

API

API Keys

Usage

Limits

CODESTRAL

Codestral

WORKSPACE

Members

Billing

Gergely Szertics Personal

< Agents

**Supervisor**

No description

Model Mistral Large 2.1

Randomness Temperature 0.7

Instructions System Prompt

You are a supervisor for scientific writing from Harvard. If you are given a text evaluate it from the writing style perspective to meet the standards of scientific writing. Give an overview about what are the strengths and potential improvements for this text. Be specific what should be improved and what I should pay attention to when writing.

Demonstrations Few-Shot Prompts

Add Demonstration (Input/Output)

<https://console.mistral.ai/build/agents>

## GPTs

Discover and create custom versions of ChatGPT that combine instructions, extra knowledge, and any combination of skills.

Search GPTs

Top Picks Writing Productivity Research & Analysis Education Lifestyle Programming

### Featured

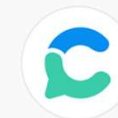
Curated top picks from this week



#### Website Generator

Create a website in seconds!  
Generate, design, write code, and write copy for your website...

By websitegenerator.b12.io



#### Consensus

Ask the research, chat directly with the world's scientific literature. Search references, get simple...

By consensus.app

<https://chatgpt.com/gpts>

# Chain of thought - let's think step by step!

## Standard Prompting

### Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

### Model Output

A: The answer is 27. ❌

## Chain-of-Thought Prompting

### Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls.  $5 + 6 = 11$ . The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

### Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had  $23 - 20 = 3$ . They bought 6 more apples, so they have  $3 + 6 = 9$ . The answer is 9. ✅



# “Inference time compute”

... of-thought (+ 🍓)

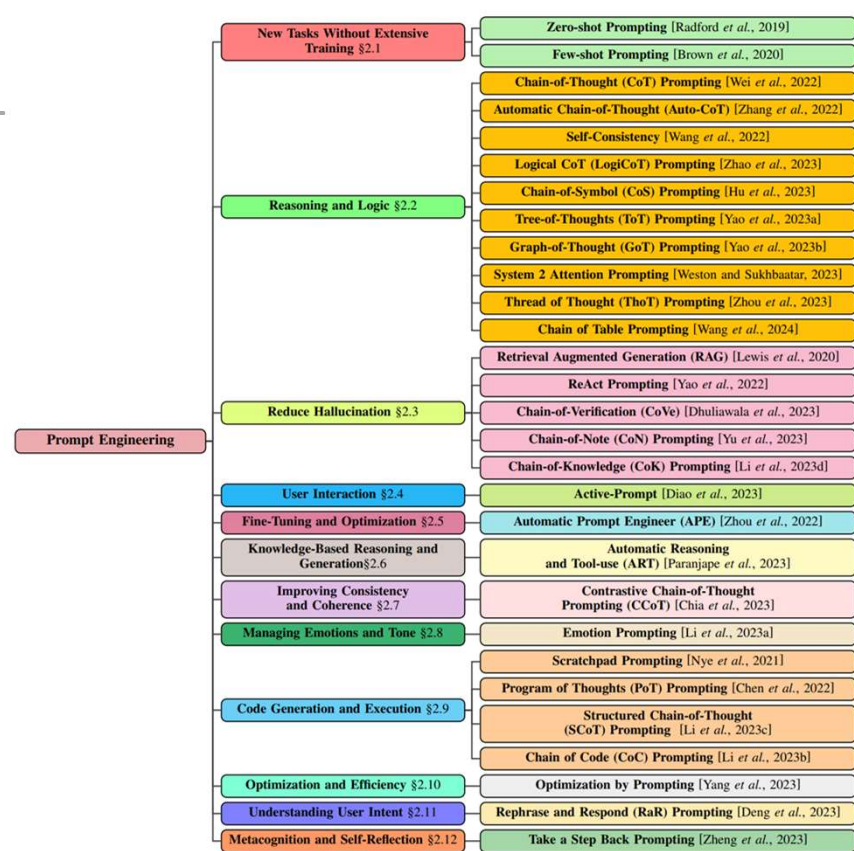


Figure 2: Taxonomy of prompt engineering techniques in LLMs, organized around application domains, providing a nuanced framework for customizing prompts across diverse contexts.

[Source](#)

# ReACT - Usage of external tools

## HotPotQA

**Question:** Seven Brief Lessons on Physics was written by an Italian physicist that has worked in France since what year?

## (d) ReAct

**Thought 1:** I need to search Seven Brief Lessons on Physics, find its author, then find when the author has worked in France since.

**Act 1:** `Search[Seven Brief Lessons on Physics]`

**Obs 1:** Seven Brief Lessons on Physics (Italian: Sette brevi lezioni di fisica) is a short book by the Italian physicist Carlo Rovelli. Originally published in Italian in...

**Thought 2:** `The author of Seven Brief Lessons on Physics is Carlo Rovelli.` I need to search Carlo Rovelli next and find when he has worked in France since.

**Act 2:** `Search[Carlo Rovelli]`

**Obs 2:** Carlo Rovelli is an Italian theoretical physicist and writer who has worked in Italy, the United States and, since 2000, in France.[1] He is also currently a Distinguished Visiting Research Chair at the Perimeter Institute...

**Thought 3:** `Carlo Rovelli has worked in France since 2000.` So the answer is 2000.

**Act 3:** `Finish[2000]`



From “roleplay” to actual tool usage:

## Roleplay example





- Browser plugin
- Code plugin
- Retrieval plugin
- **Custom plugins / functions**



### Knowledge sources:

- Wikipedia API
- Wolfram Alpha

### Search Tools:

- Bing Search
- Google Search
- SearxNG Search API

**Custom company  
search engine /  
knowledge base**

### Service APIs:

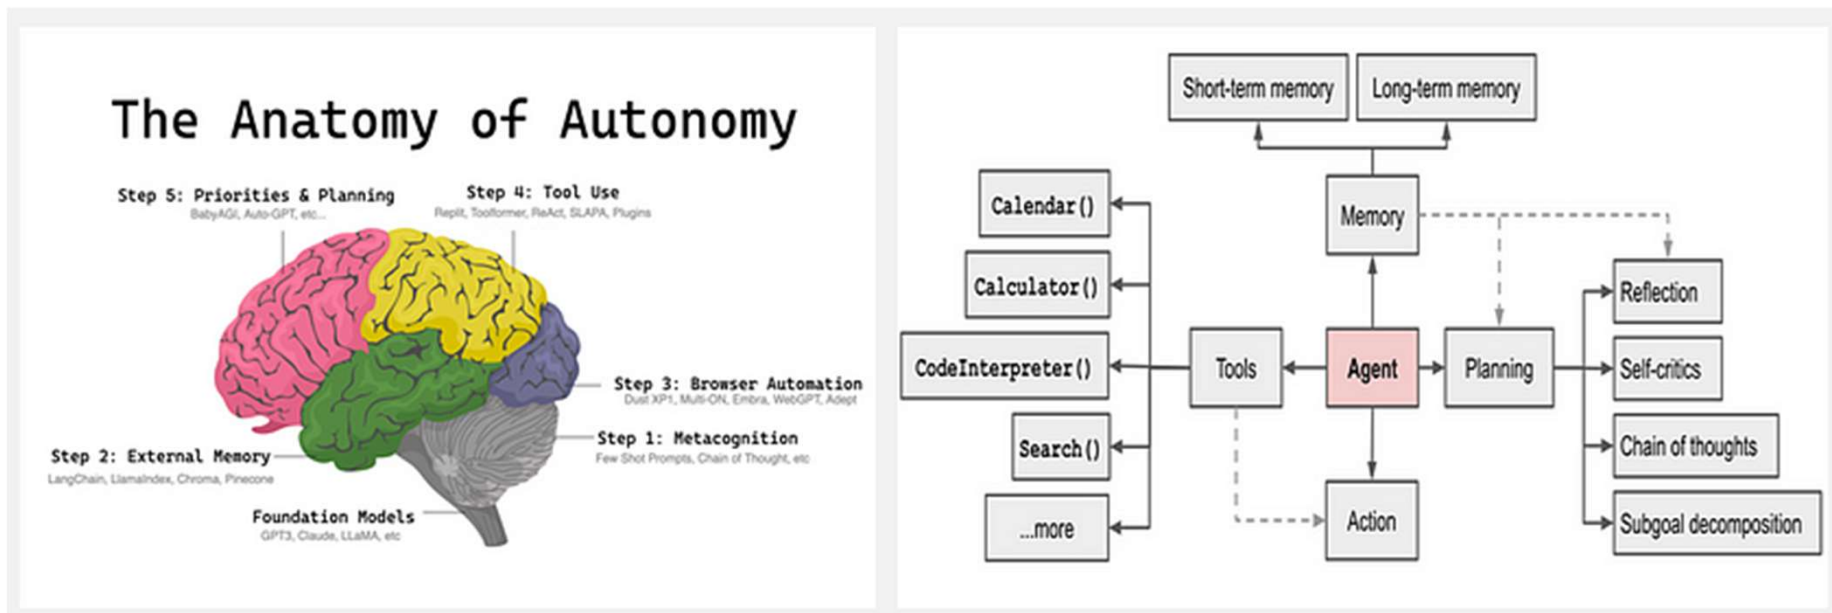
- OpenWeatherMap API
- IFTTT WebHooks
- Zapier Natural Language Actions API
- Apify
- Requests - Any tool on the web

### Command execution environments

- Bash
- Python REPL

**"Human as a tool" - let other people help in task execution**

# Building cognitive architectures



(Some) sources:  
[The Anatomy of Autonomy: Why Agents are the next AI Killer App after ChatGPT](#)  
[LLM Powered Autonomous Agents](#)  
[The Rise of Autonomous AI Agents: Debundling the Market Landscape](#)



# From allowing them to collaborate to designing them to collaborate

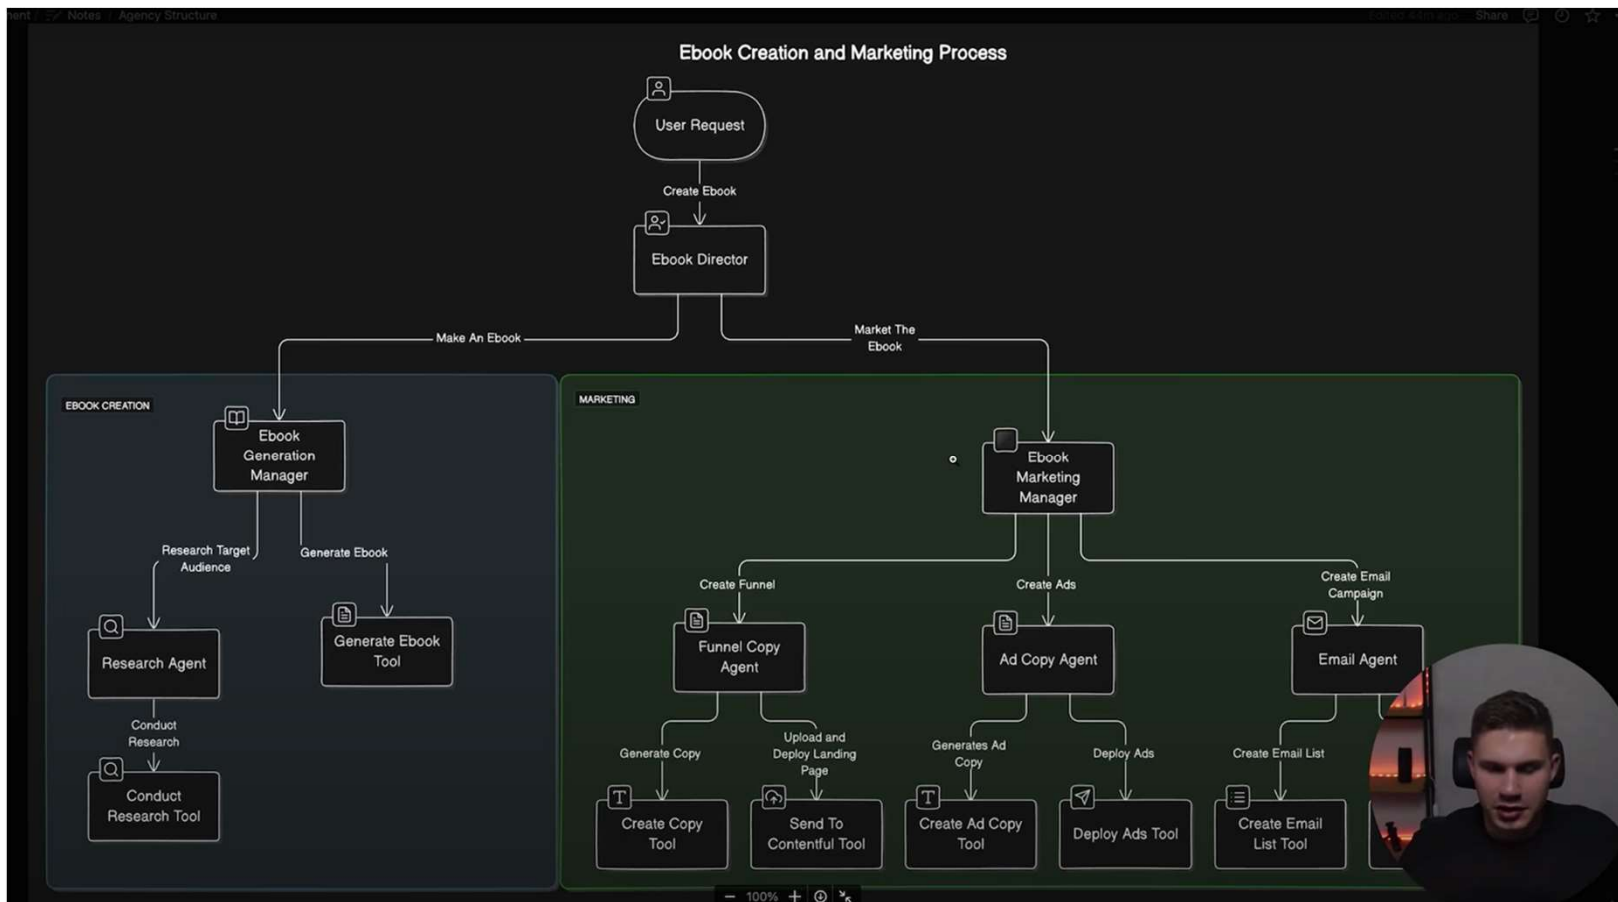


(Some) sources:

[AI Agents Simulate a Town](#) | [Generative Agents: Interactive Simulacra of Human Behavior](#),  
[Build an Entire AI Agent Workforce](#) | [ChatDev and Google Brain "Society of Mind"](#) | [AGI User Interface](#)



# Automated publishing: zero employees

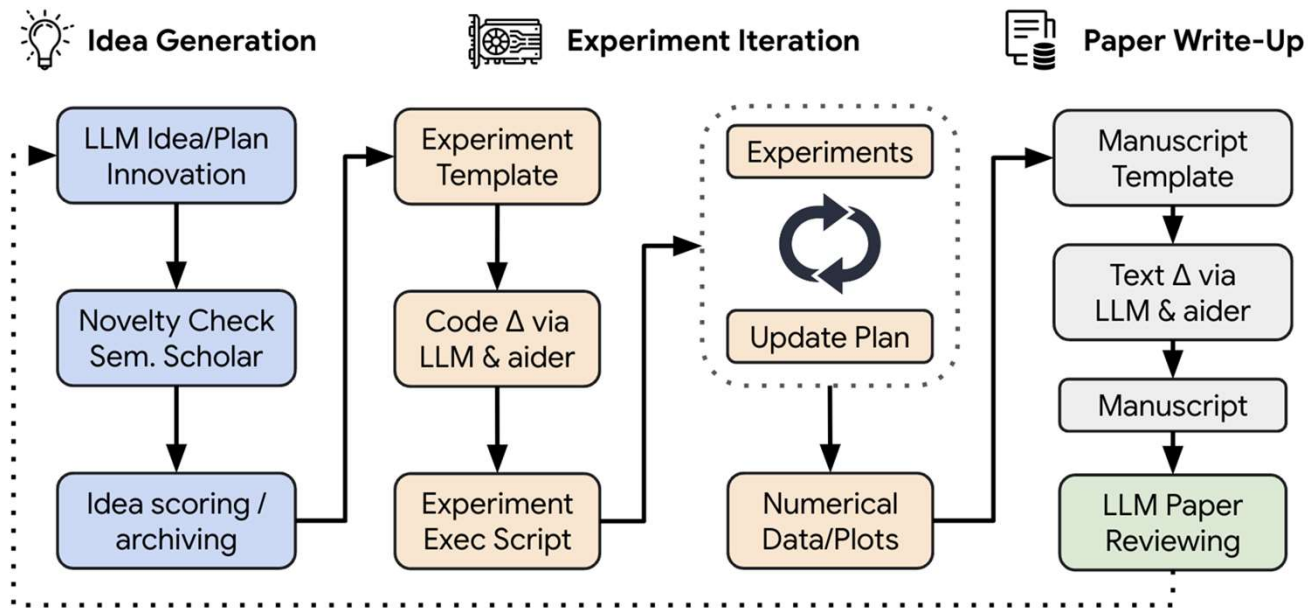


[Link to video](#)

# “AI Scientist” - autonomous discovery?

## The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery

August 13, 2024



The system automatically sets up hypothesis, experiment design, evaluation, conclusion and publication writing.

# Defining agents on the fly

**Vocational Educator** AI Education Specialist Policy Analyst Basic fact writer

This persona will represent the voice of traditional vocational education. They will provide insights into how AI is transforming pedagogical approaches and the challenges educators face when incorporating technology into their teaching methods. Their focus will be on the implications for teaching practices, student engagement, and the need for teacher training in AI tools.



How are vocational education institutions currently integrating AI technologies into their curriculum and teaching practices?

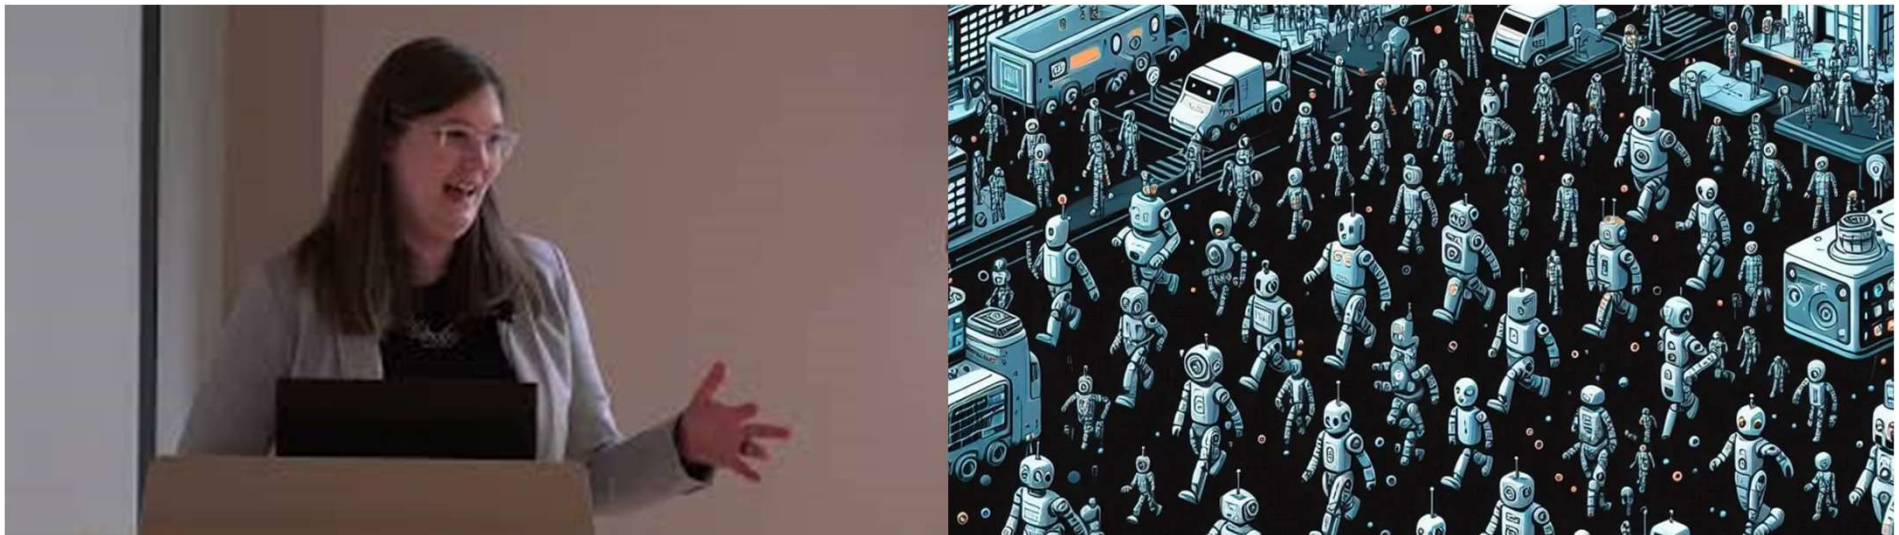


Vocational education institutions are increasingly integrating AI technologies into their curriculum and teaching practices to prepare students for the evolving job market. This integration is happening across several dimensions: 1. **Personalized Learning**: AI is being used to create personalized learning experiences tailored to individual students' needs and learning styles. This approach ensures that each student receives the right level of challenge and support, optimizing their learning outcomes (Source: [2]). 2. **Generative AI Deployment**: Over half of Registered Training Organisations (RTOs) are expected to deploy generative AI within the next 24 months. This technology is considered game-changing, offering new methods for teaching and learning beyond traditional practices (Source: [5]). 3. **Practical Skills and Immediate Job Readiness**: Vocational education focuses on practical skills that lead directly to employment. AI tools can enhance this hands-on, practical training by providing simulations, real-time feedback, and virtual apprenticeships that mirror industry

Stanford university  
STORM:

"Assisting in Writing  
Wikipedia-like Articles  
From Scratch with Large  
Language Models"

# LLM-s simulating human societies



S3: Social-network Simulation System with Large Language Model-Empowered Agents

Simulating millions of LLM agents with AgentTorch

## **Other foundational models**

# Specialized LLM for science

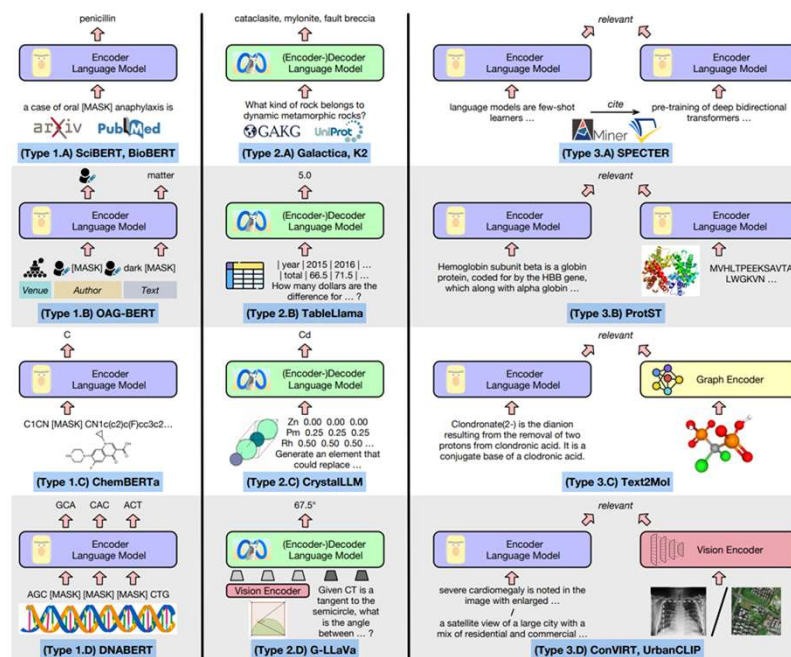
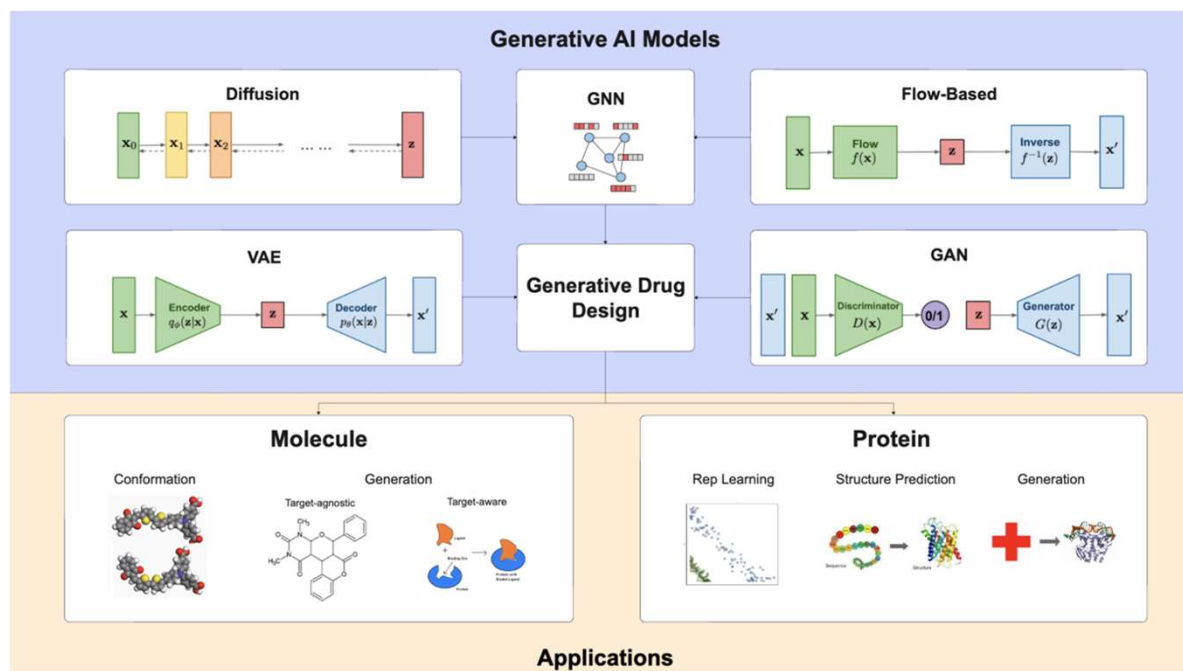


Figure 1: Three major types of scientific LLM pre-training techniques. (COLUMN 1): Pre-training encoder LLMs with sequentialized scientific data (e.g., text, academic graphs, molecules, biological sequences) via masked language modeling. (COLUMN 2): Pre-training (encoder)-decoder LLMs with sequentialized scientific data (e.g., text, tables, crystals, images) via next token prediction (possibly with instruction tuning). (COLUMN 3): Mapping text and relevant sequences/graphs/images closer in the latent space via contrastive learning.

## [A Comprehensive Survey of Scientific Large Language Models and Their Applications in Scientific Discovery](#)



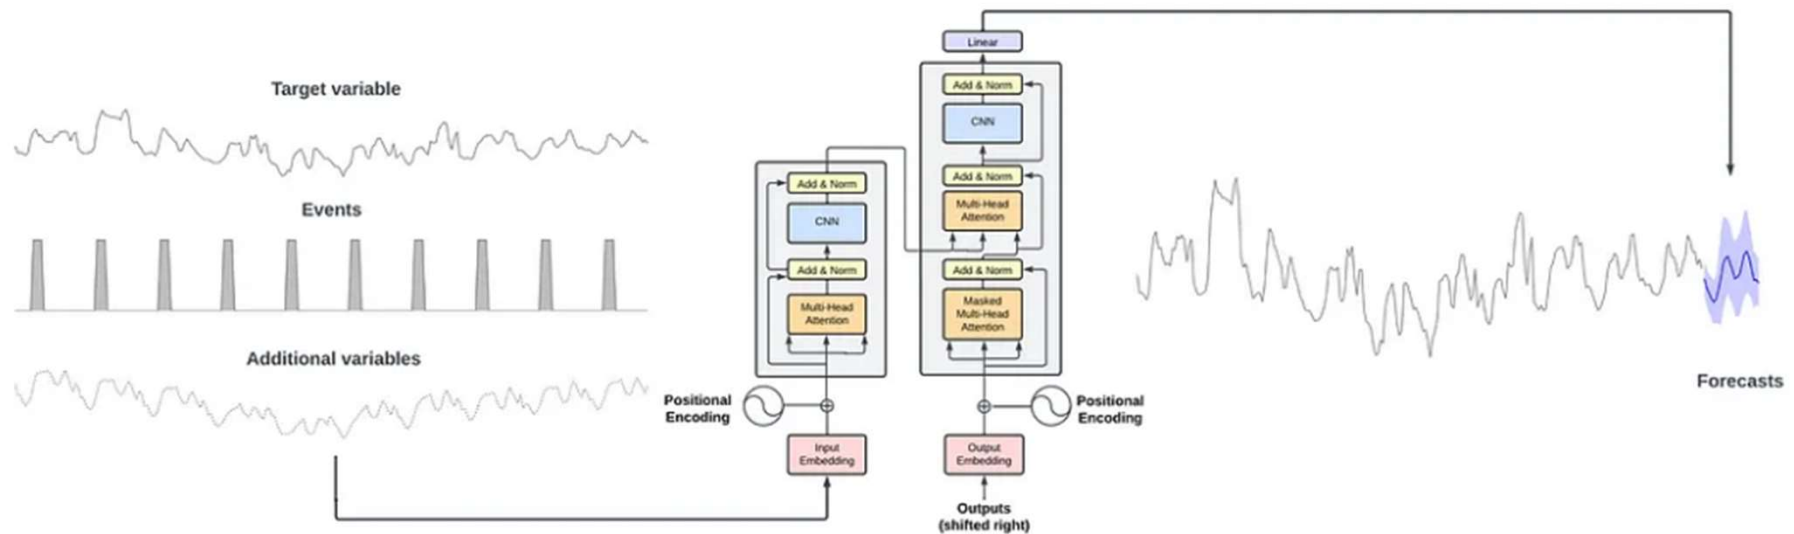
# Foundation model - Generative drug design



**Fig. 1.** An overview of the topics covered in this survey. In particular, we explore the intersection between generative AI model architectures and real-world applications, organized into two main categories: small molecule and protein generation tasks. Note that diffusion and flow-based models are often paired with GNNs for processing 2D/3D-based input, while VAEs and GANs are typically used for 1D input. Images used from [17, 18, 19, 20, 21, 22, 23].



# Foundation model - Time series



Architecture of TimeGPT. The input series, along with exogenous variables, is fed to the encoder of the Transformer, and the decoder then generates forecasts. Image by Azul Garza and Max Mergenthaler-Canseco from [TimeGPT-1](#).

**“Thinking about LLMs as text generators  
is as thinking of computers as  
calculators.”**

- Andrej Karpathy<sup>(“source”)</sup>

HUN  
REN 

<https://hun-ren.hu/ai-4-science>