

This presentation is a part of Aalto University's webinar series on

# Research Data Management & Open Science

—  
Autumn 2023



Aalto-yliopisto  
Aalto-universitetet  
Aalto University



All slides in this presentation are licensed with [CC-BY](https://creativecommons.org/licenses/by/4.0/) and can be reused with attribution to [doi:10.5281/zenodo.8413164](https://doi.org/10.5281/zenodo.8413164), unless explicitly mentioned otherwise.

# AI & Research work

---

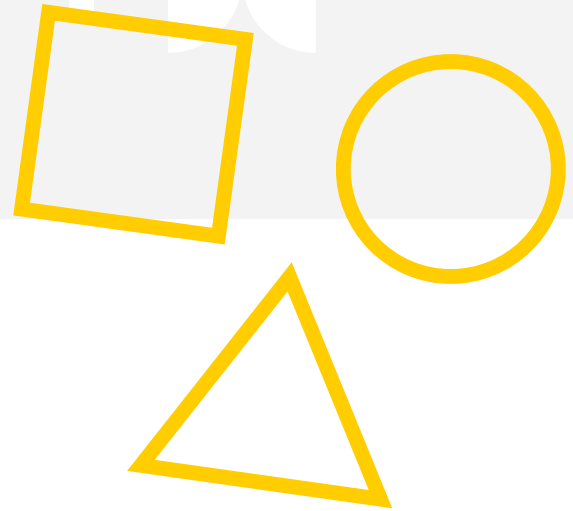
**Dr. Enrico Glerean, Staff Scientist at SCI and Data Agent**

**Dr. Pedro Silva, Postdoc at CHEM and Data Agent**

**10.10.2023**



**Aalto-yliopisto  
Aalto-universitetet  
Aalto University**



# Outline: AI in research work

1. From AI to generative AI
2. Prompt engineering for research
3. AI in the academic workflow
4. Ethical and legal considerations
5. Other tools and advanced topics

## Learning outcomes

- How to use generative AI tools in your research work
- Evaluating the outputs and understanding limitations
- Awareness of the larger picture of AI tools today

# DISCLAIMER

THIS IS NOT A COMPUTER SCIENCE COURSE ON GENERATIVE ARTIFICIAL INTELLIGENCE. THE MATERIALS ARE TAILORED FOR RESEARCHERS WHO HAVE NEVER USED THESE TECHNOLOGIES

SINCE WE ONLY HAVE 1 HOUR, ENRICO AND PEDRO WILL FOCUS ON INTUITIVE EXPLANATIONS AND PRACTICAL HINTS ON HOW TO USE THE MOST POPULAR TOOLS AND WHAT TO BE AWARE OF WHEN USING THEM

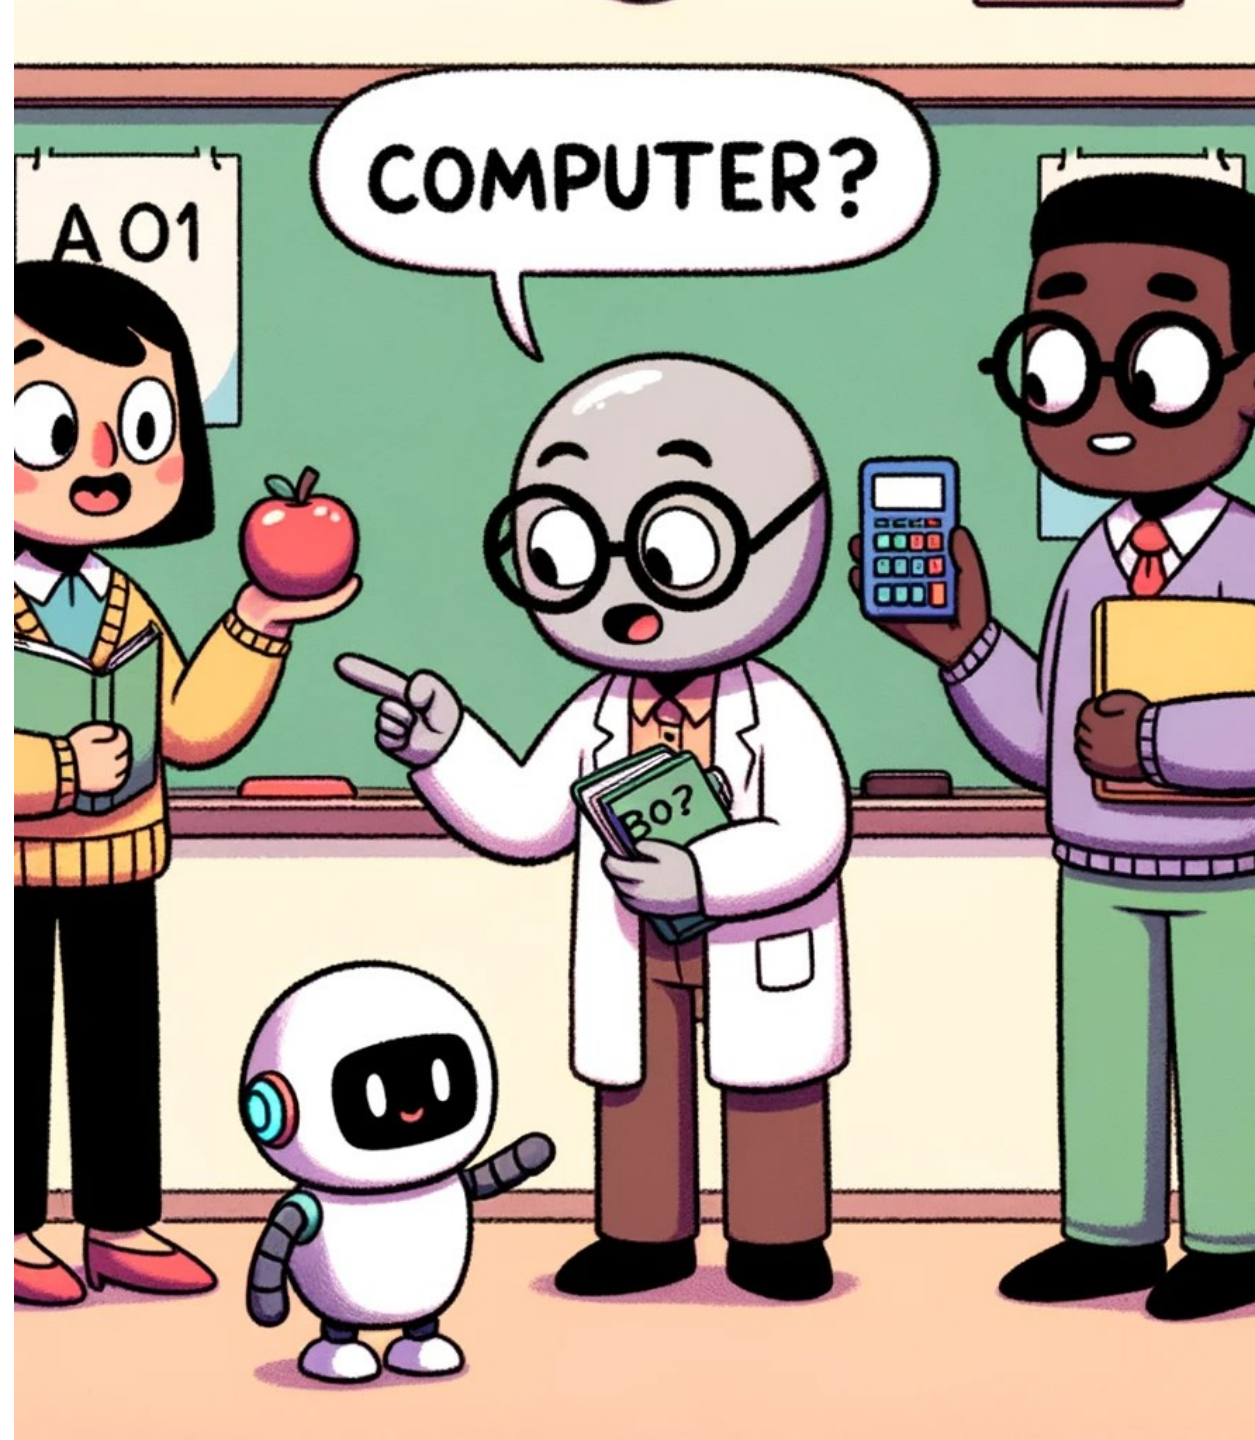
SHARE YOUR EXPERIENCE AND HINTS WITH OTHERS SO THAT WE CAN BUILD AN OPEN GUIDE FOR RESEARCHERS

# The Research Process aided by AI

*with Responsible Conduct of Research at its core and some examples*



# 1. What is (generative) AI?



# What is Artificial Intelligence?

Multiple definitions since 1956, let's pick one from law:

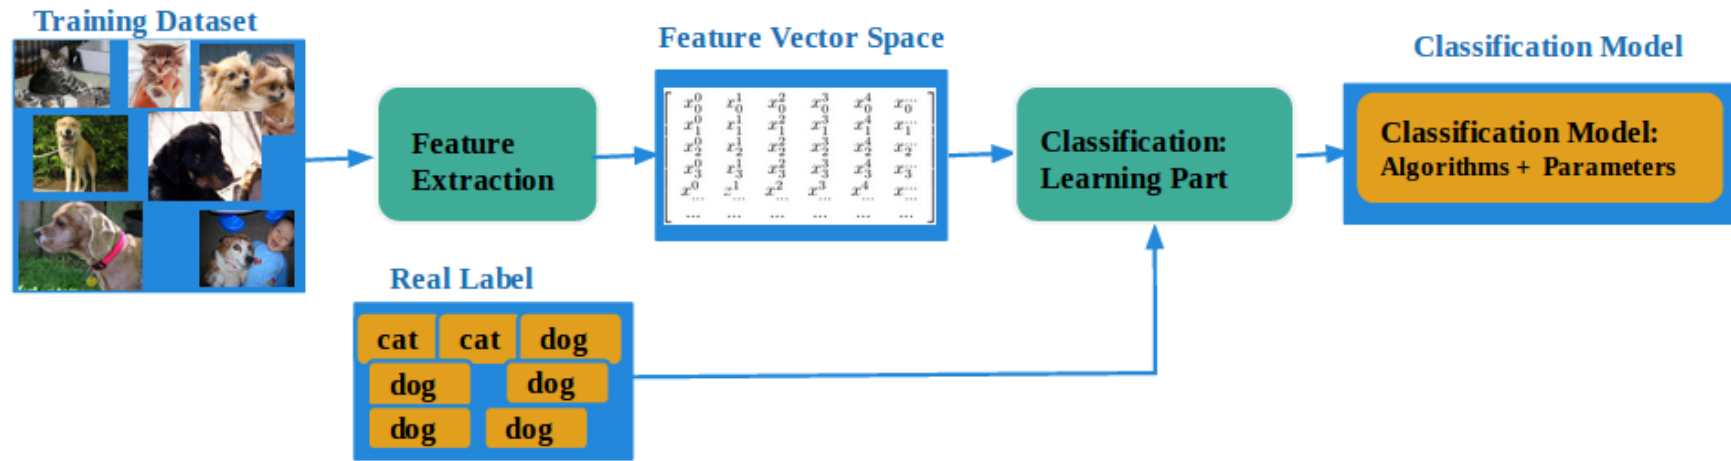
*“‘Artificial Intelligence system’ (AI system) means software that is developed with one or more of the techniques and approaches listed in Annex I (e.g. machine learning and statistics) and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with”*

From “Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS”

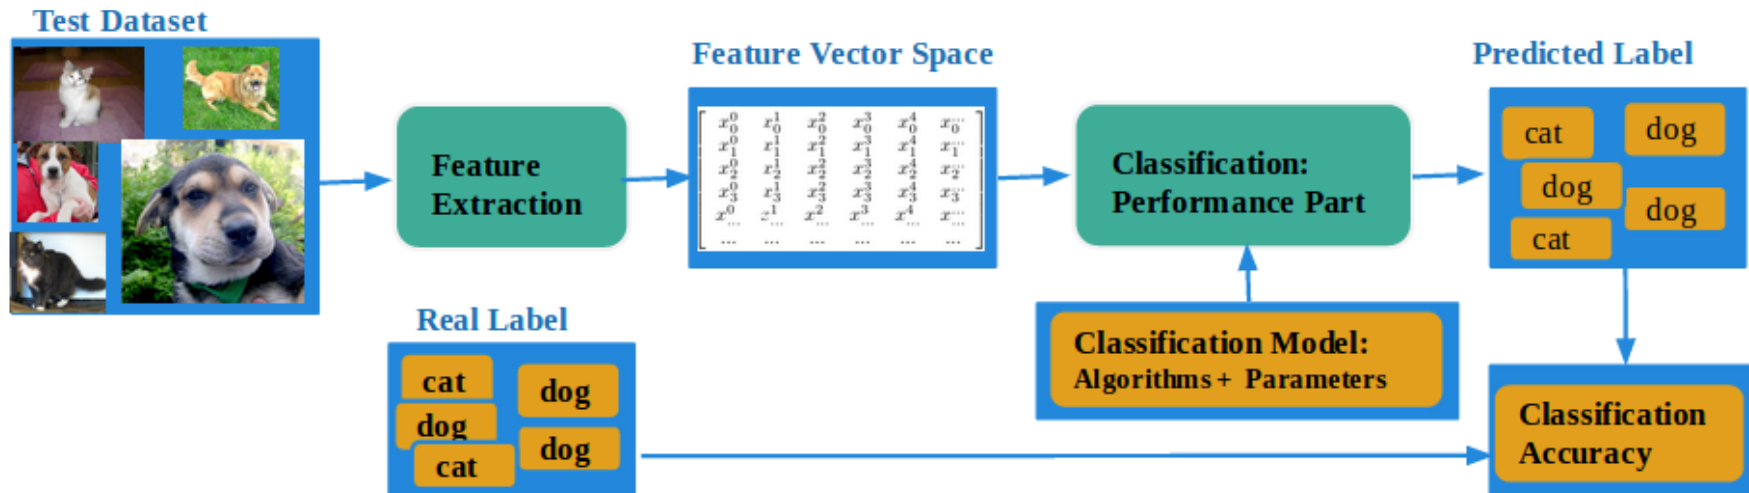
Other AI legal definitions collected by IAPP.

# Learning cats and dogs: Machine learning

## 1) Learning stage



## 2) Testing stage





# Generative AI: an intuitive explanation

- The key task in machine learning is trying to **predict**
- The **classifier** from the previous example is able to **predict labels from new images**: given a new picture of a dog, it can tell if there's a cat or a dog
- There are other types in AI called “**generative**”: the program learns patterns from the training set and is able to generate a picture, or text, or sound following the rules of what it has learned

We are not even scratching the surface here, if you are interested to know more, ask in the chat and we can provide references. See [this link](#) for a great collection of references

# Focus on text generation: large language models (LLMs)

Why focus on language?

- Large amount of data for training these **foundational models**
- Language is possibly the most versatile modality as it allows to describe (almost) anything: concepts, facts, relationships, processes, software code, ...
- Language (as text via a keyboard) is an efficient way to interact with a computer program via instructions

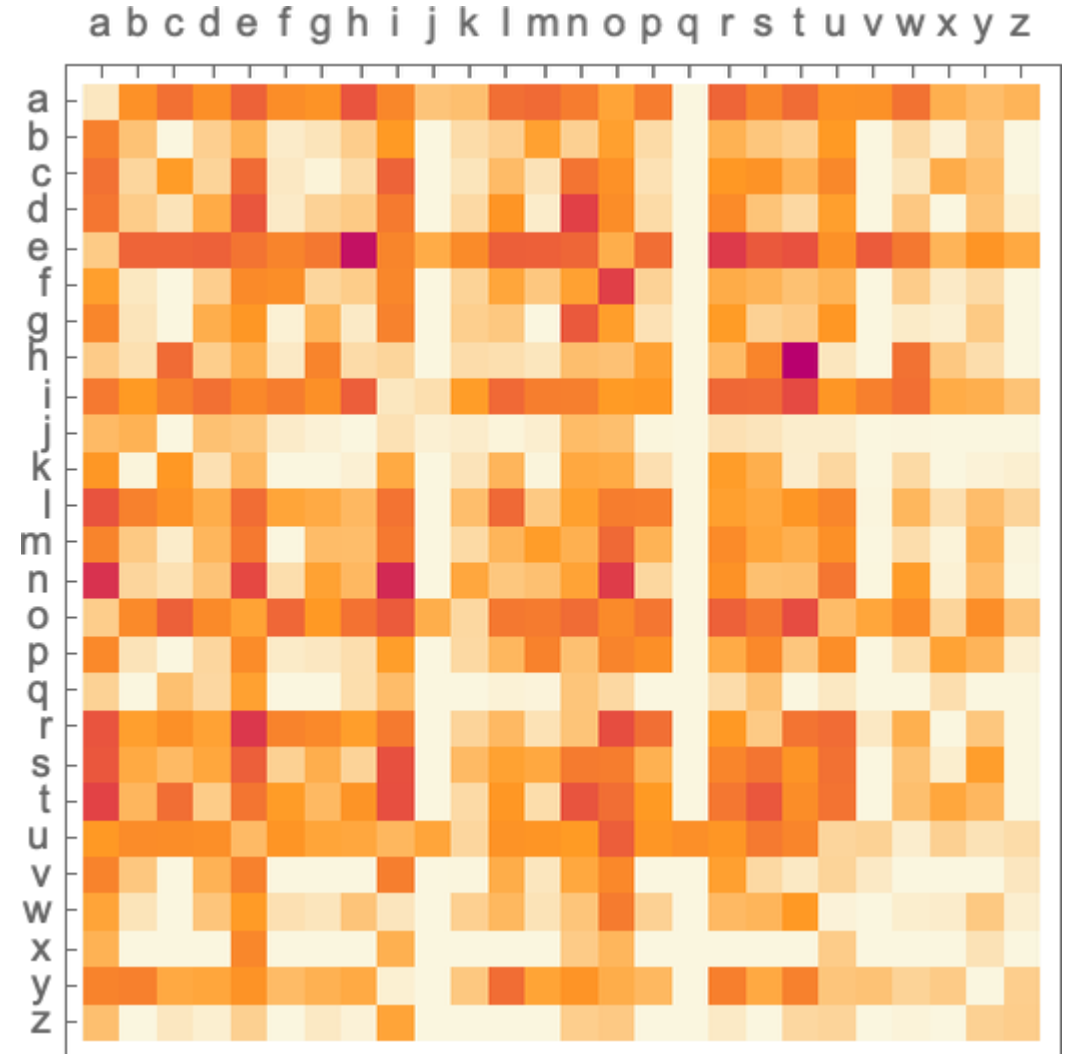
# Guess the letter

*It is an English word starting with the letter 'q': which letter comes after that?*

# Guess the next letter

*If I am thinking an English word that starts with q and I ask you to guess the next letter, what is your most likely guess?*

*Image from What is ChatGPT doing, and why does it work? by Stephen Wolfram*



*How often a letter from the columns is followed by a letter on the rows? Bright yellow = never, dark red = very often*

# Guess the next word, sentence, paragraph, chapter, book...

Large language models estimate the probability between **words in the context of sentences, paragraphs, chapters, books...** Given a **prompt** sentence, the language model can be asked to continue it based on the most likely expected following words, with a given randomness (**temperature**)

E.g. the sentence “***The best thing about AI is its ability to***” has these as the 5 most likely following words: “learn”(4.5%), “predict” (3.5%), “make”, (3.2%), “understand” (3.1%), “do” (2.9%)

# Guess the next word, sentence, paragraph, chapter, book... (II)

{The best thing about AI is its ability to,

The best thing about AI is its ability to create,

The best thing about AI is its ability to create worlds,

The best thing about AI is its ability to create worlds that,

The best thing about AI is its ability to create worlds that are,

The best thing about AI is its ability to create worlds that are both,

The best thing about AI is its ability to create worlds that are both exciting,}

The best thing about AI is its ability to learn. I've always liked the

The best thing about AI is its ability to really come into your world and just

The best thing about AI is its ability to examine human behavior and the way it

The best thing about AI is its ability to do a great job of teaching us

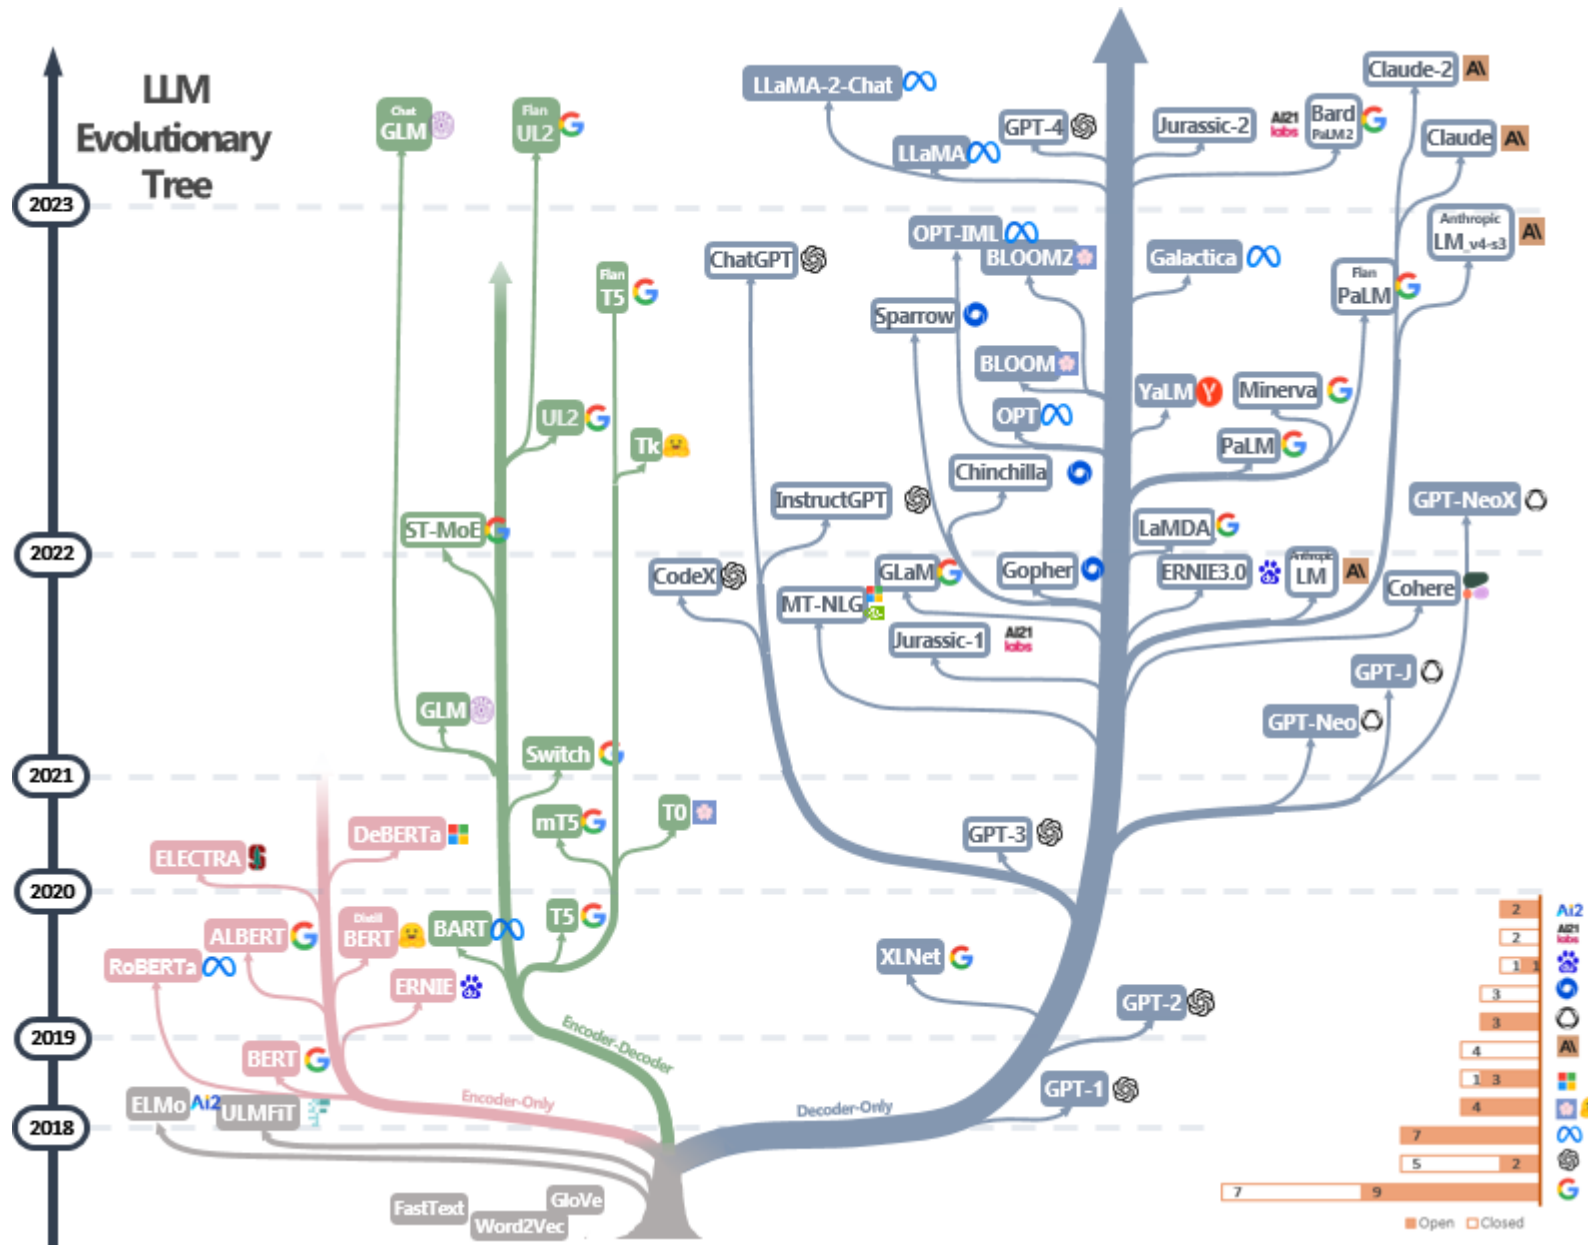
The best thing about AI is its ability to create real tasks, but you can

At each generated word, we re-ask the model to pick another one from the top N words that would follow **from all the previous words seen so far.**

**Different runs, produce different outputs** (goodbye reproducibility!)

*From What is ChatGPT doing, and why does it work?  
by Stephen Wolfram*

# History of LLMs



Intuitive explanations:

- **Encoder** models: they “understand” (= map it to a multidimensional space of words) the prompt
- **Encoder-decoder** models: they “understand” the prompt and produce a new response
- **Decoder** models: they are usually *autoregressive* and only generate based on what has been given so far.

Yang, J., Jin, H., Tang, R., Han, X., Feng, Q., Jiang, H., ... & Hu, X. (2023). Harnessing the power of llms in practice: A survey on chatgpt and beyond. *arXiv preprint arXiv:2304.13712*.

# How to use them?

## Platforms currently available

Historically the way to query generative models was to ask “**continue the sentence**”. Later **OpenAI** made the **chat interface** successful, but they were not the first (ref)

Most current platform require you to **sign-up to a service** since the actual model and its inference software **are not running on your local computer**.

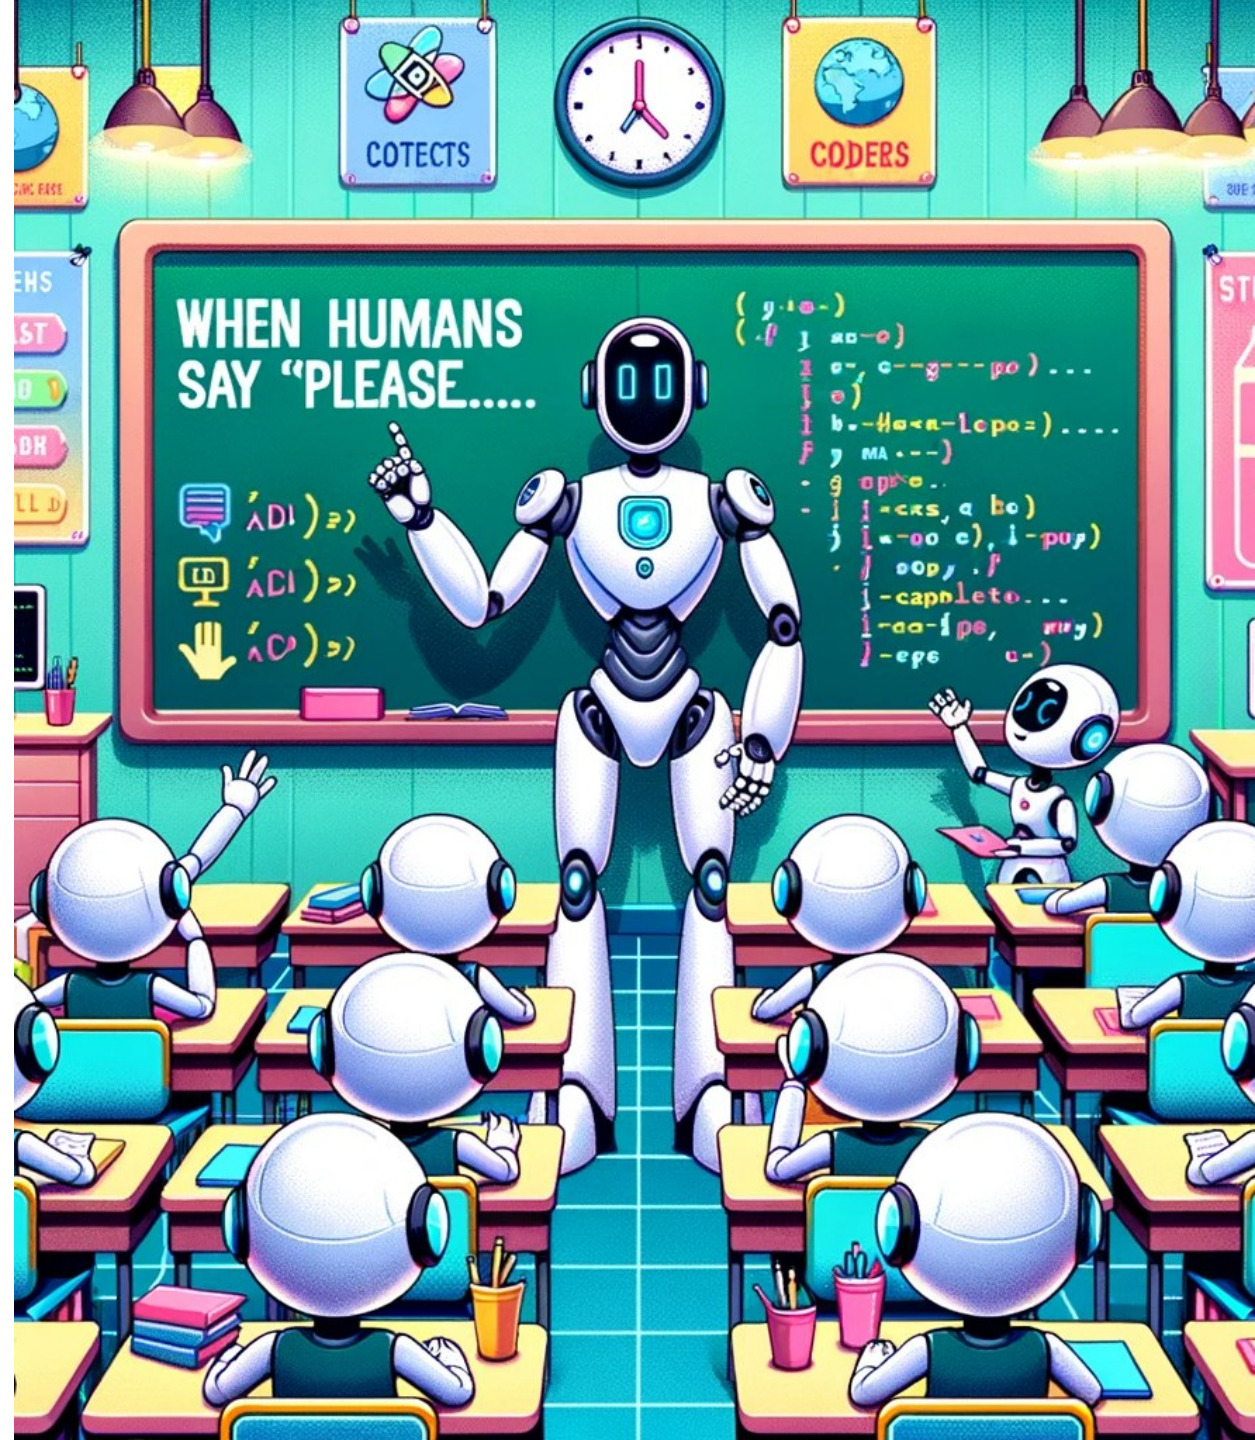
Current platforms and technologies: OpenAI with free **ChatGPT 3.5** and paid **ChatGPT 4** (also Microsoft Bing search engine), **Anthropic Claude 2** (requires VPN), **Google Bard**, Huggingface chat (various open-source models available)

If privacy and protection of your data is important, **open-source models can also be run locally**.



## 2. Prompt engineering

*“Prompt engineering is the process of constructing queries or inputs for AI language models so as to elicit the most precise, coherent, and pertinent responses.”*

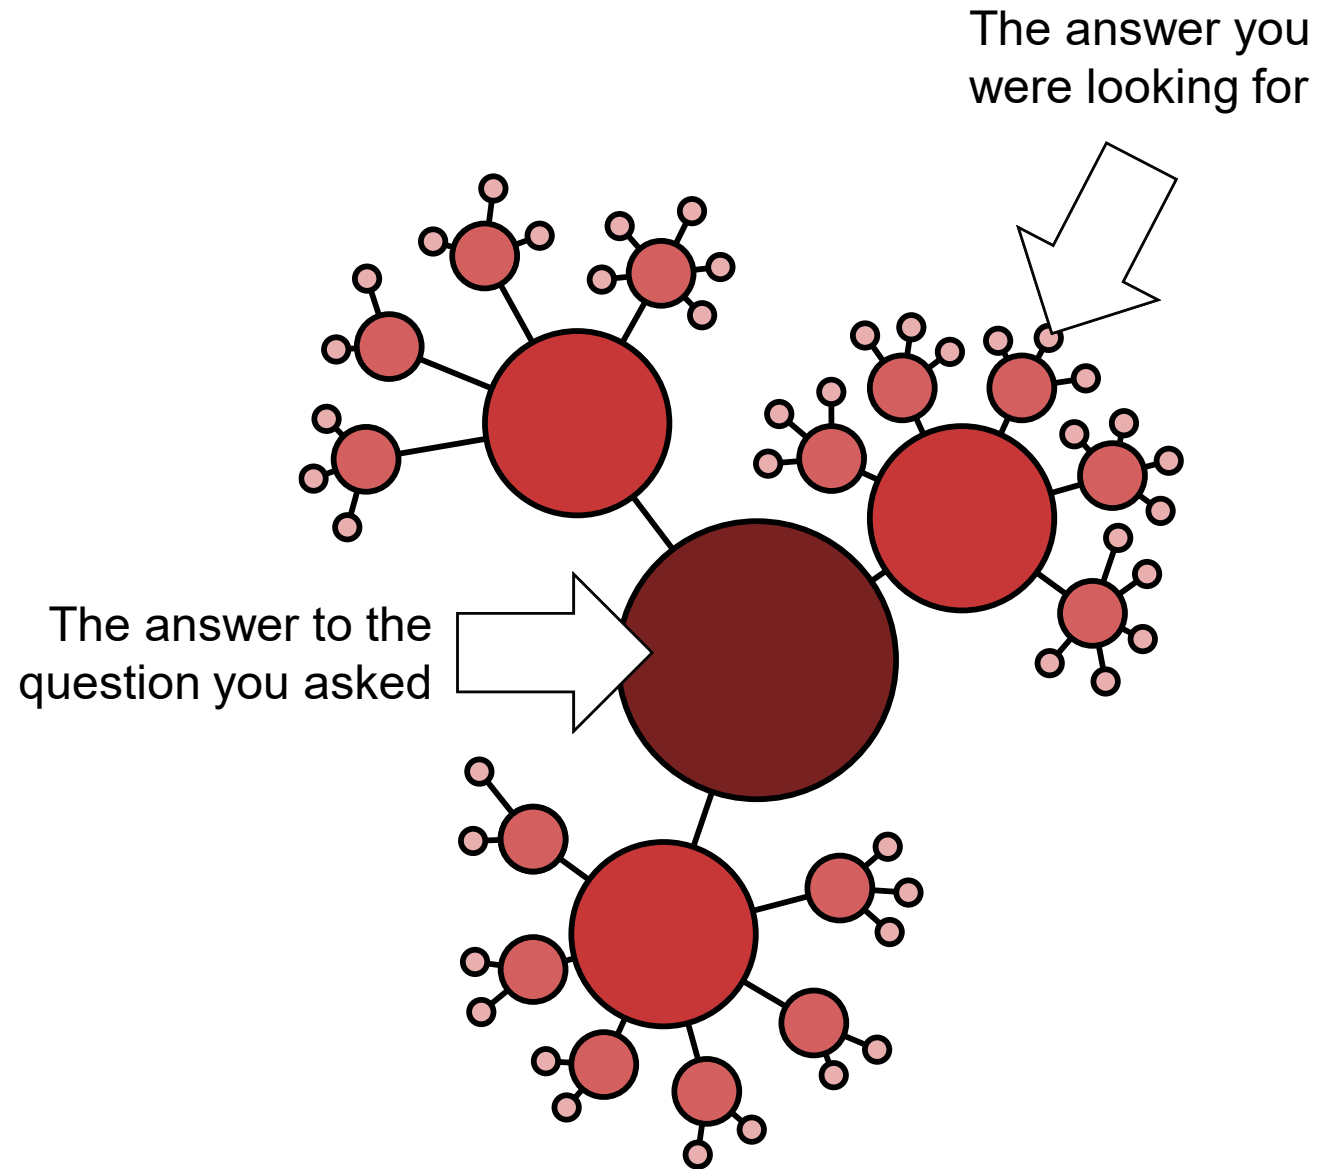


# The art of asking



Is Herlinda asking:

1. to meet with Tiyna?
2. if Tiyna is available for leisure after work?
3. what is the current time?



# Structuring the prompt

## 1. Start with a concise question:

Direct: What is the capital of France?

Instructional: Translate the following text to French: “(...)”

Exploratory: Discuss the implications of climate change on beetles.

Contextual: Given the economic downturn in 2008, what (...)

Comparative: Compare the economic systems of Finland and Sweden.

Scenario-based: Imagine you are a researcher, how would you defend funding in fundamental research?

...

- 2. Specify tone and style:** ex. write in an academic tone, explain like a 5<sup>th</sup> grader, explain to a 5<sup>th</sup> grader, translate to Finnish, ...
- 3. Add examples as input:** this can help the AI towards the answer
- 4. Iterative refinement:** after getting the first answer, you can engage in an iterative dialogue

# CLEAR Principles

1. **Concise:** brevity and clarity in prompts
2. **Logical:** structured and coherent prompts
3. **Explicit:** clear output specifications
4. **Adaptive:** flexibility and customization in prompts
5. **Reflective:** continuous evaluation and improvement of prompts

Explain

~~“Can you provide me with a detailed explanation of the process of photosynthesis and its significance?”~~

“(1) List the steps to write a research paper, (2) beginning with selecting a topic and (3) ending with proofreading the final draft.”

Instead of “What are some renewable energy sources?”, ask for “Identify five renewable energy sources and explain how each works.”

If asking, “What are some ways to conserve water?” leads to generic responses, try a more targeted and adaptive prompt like, “List household practices for conserving water and their potential impact.”

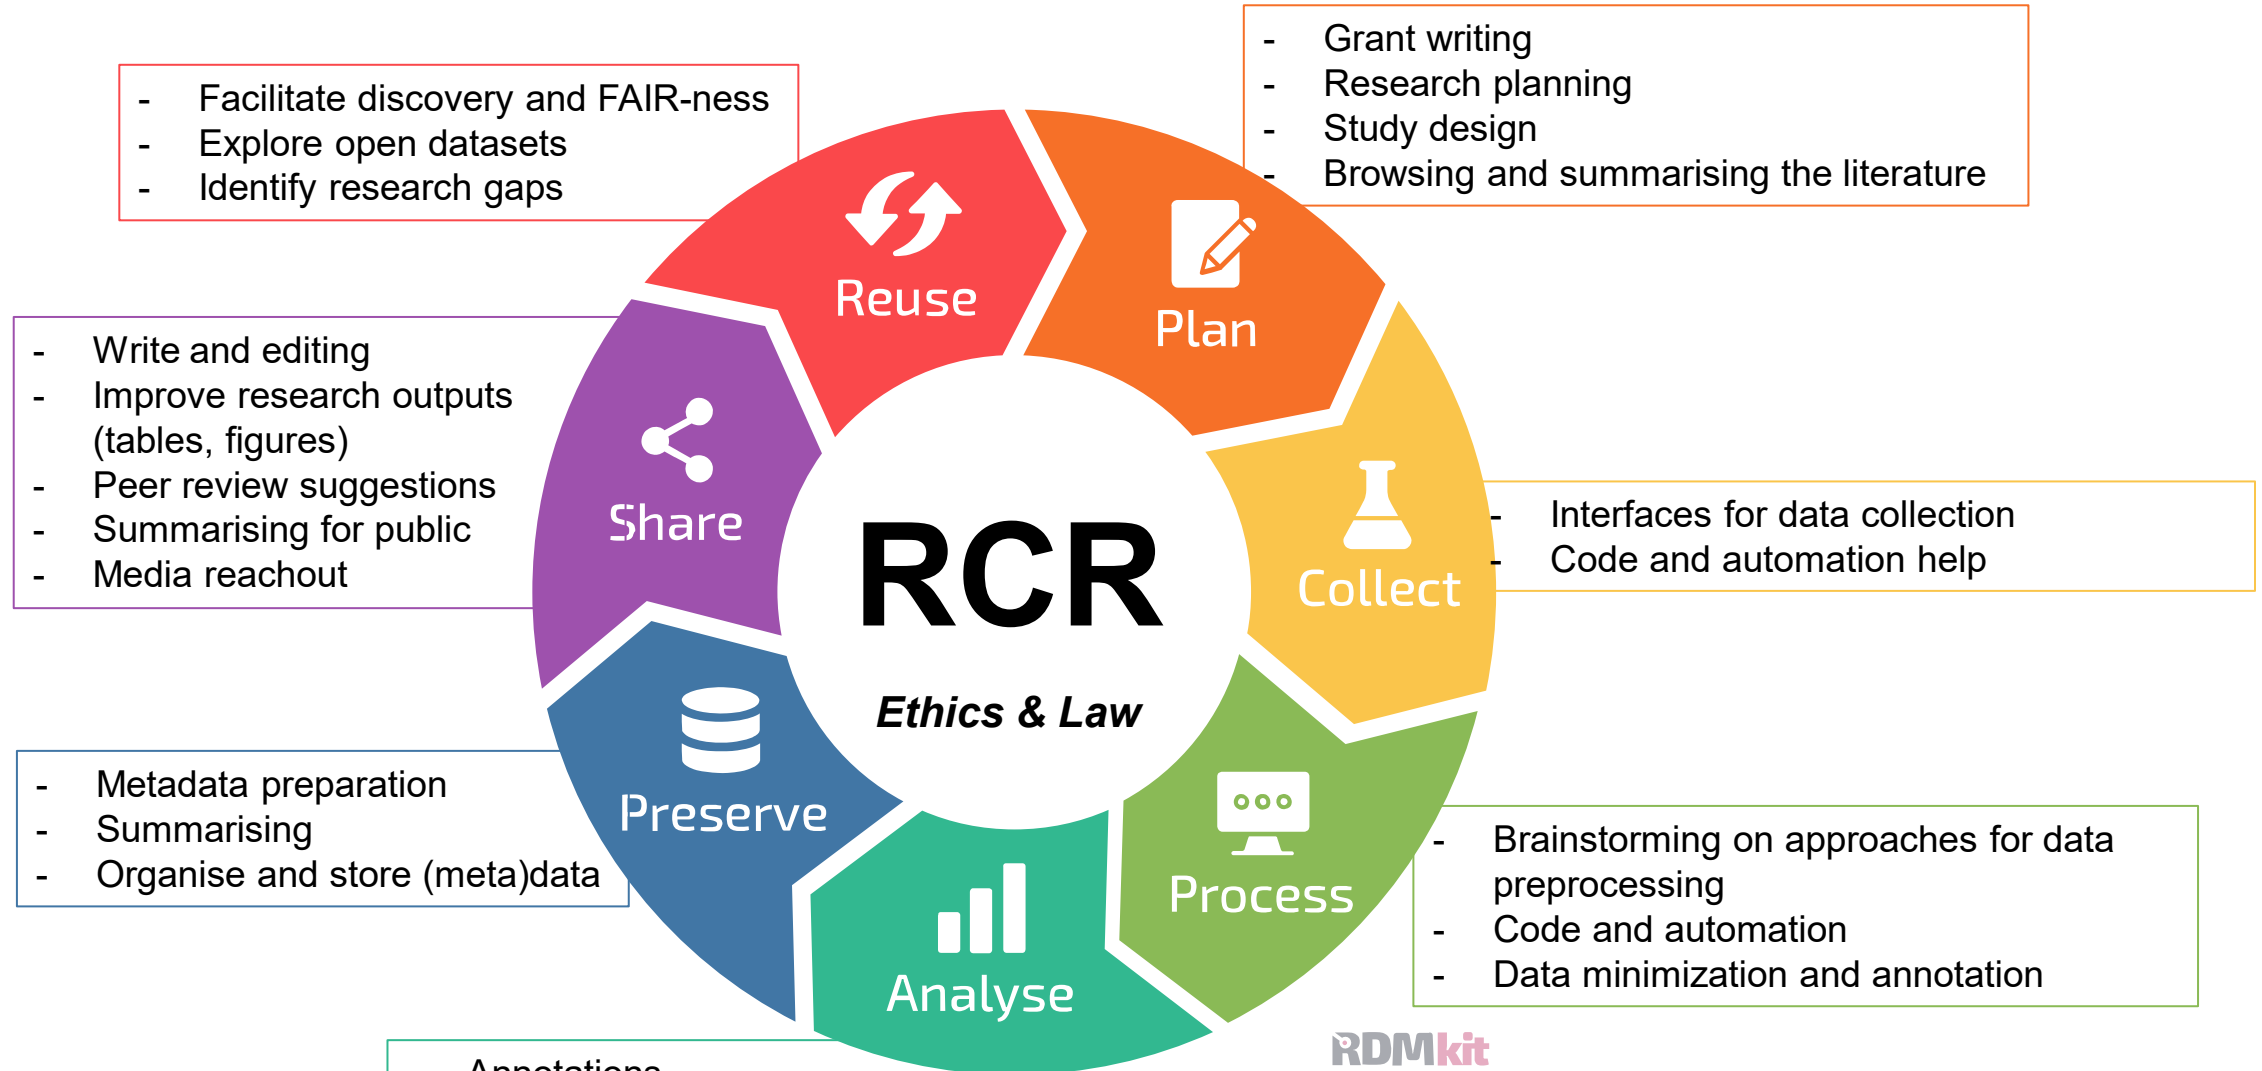
*Make a strategy request(...). Evaluate also the relevance and applicability of each strategy. Consider the target audience's needs and use this information to tailor future prompts to generate content that better addresses specific challenges or contexts.*

### 3. Use of LLMs in the academic workflow

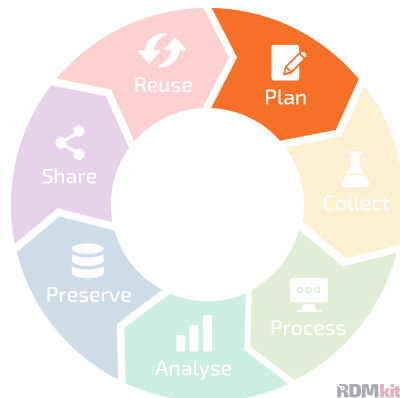


# The Research Process aided by AI (II)

*with Responsible Conduct of Research at its core and some examples*



# What approach should I use?



I'm unfamiliar with different simulations techniques to study the dynamics of elastic structures. Can you briefly mention which are the main approaches used in science and what advantages each approach present in comparison to the others?



Certainly! There are several simulation techniques used to study the dynamics of elastic structures in science and engineering. Each approach has its own advantages and is suitable for different types of problems. Here are some of the main simulation techniques:

## 1. Finite Element Analysis (FEA):

- **Advantages:** FEA is a versatile and widely used method for simulating the behavior of complex elastic structures. It can handle a wide range of geometries and boundary conditions, making it suitable for a variety of engineering problems.
- **Use Cases:** FEA is commonly used for structural analysis, heat transfer analysis, and fluid flow analysis, among other applications.

## 2. Molecular Dynamics (MD):

- **Advantages:** MD simulations are used to model the behavior of atoms and molecules in materials. They provide atomic-level insights into the dynamics and properties of elastic materials.
- **Use Cases:** MD is commonly employed in materials science to study the mechanical properties of materials at the atomic scale, such as the elasticity of polymers and metals.

## 3. Finite Difference Method (FDM):

- **Advantages:** FDM is a straightforward numerical technique that discretizes differential equations to model elastic structures. It is simple to implement and suitable for many linear problems.
- **Use Cases:** FDM is often used for solving partial differential equations governing wave propagation, heat conduction, and other elastic phenomena.

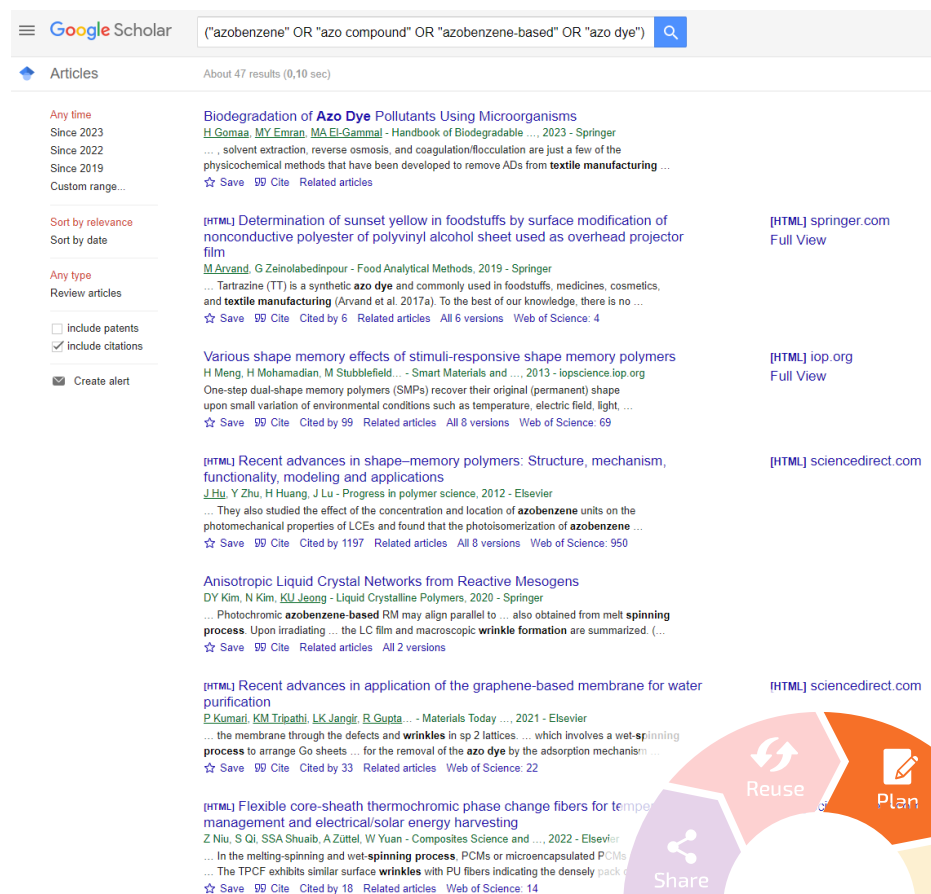
## 4. Boundary Element Method (BEM):

- **Advantages:** BEM is particularly useful for problems with open boundaries, like acoustic and electromagnetic wave propagation. It reduces the dimensionality of the problem by

# Prompt a prompt



Make a prompt for a google scholar search to find articles about azobenzene and spinning processes to produce yarns for textiles. Consider variations about compounds names. Also the query should return articles with wrinkles.



The screenshot shows a Google Scholar search interface with the query "(azobenzene" OR "azo compound" OR "azobenzene-based" OR "azo dye)". The results page lists several articles, including "Biodegradation of Azo Dye Pollutants Using Microorganisms", "Determination of sunset yellow in foodstuffs by surface modification of nonconductive polyester of polyvinyl alcohol sheet used as overhead projector film", "Recent advances in shape-memory polymers: Structure, mechanism, functionality, modeling and applications", "Various shape memory effects of stimuli-responsive shape memory polymers", "Recent advances in shape-memory polymers: Structure, mechanism, functionality, modeling and applications", "Anisotropic Liquid Crystal Networks from Reactive Mesogens", and "Recent advances in application of the graphene-based membrane for water purification".



To find articles about azobenzene and spinning processes to produce yarns for textiles, while also considering variations in compound names and articles discussing wrinkles, you can use the following Google Scholar search prompt:

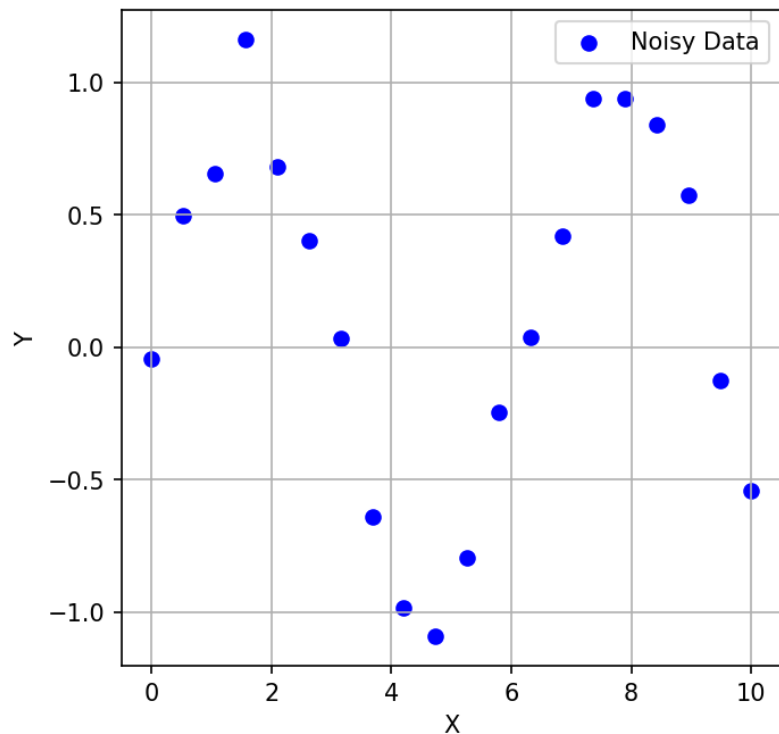
```
arduino Copy code  
  
("azobenzene" OR "azo compound" OR "azobenzene-based" OR "azo dye")  
AND ("spinning process" OR "yarn production" OR "textile manufacturing")  
AND ("wrinkles" OR "wrinkle formation" OR "wrinkle analysis")
```

This prompt includes variations of the term "azobenzene" to ensure you capture a broad range of relevant articles. It also includes keywords related to spinning processes, yarn production, textile manufacturing, and wrinkles to provide a comprehensive search for your topic.





# Coding (I)



Step-by-step  
Response  
(this answer type  
can be prompted)



I have a set of points that contain some noise. In python, how can I smooth the data so the derivative in the boundary is zero?



To smooth a set of points with noise in Python while ensuring that the derivative at the boundaries is zero, you can use various smoothing techniques. One common method is to apply a cubic spline interpolation with boundary conditions. Here's a step-by-step guide on how to do this using the `scipy` library:

1. First, you need to import the necessary libraries:

python

Copy code

```
import numpy as np
from scipy.interpolate import CubicSpline
import matplotlib.pyplot as plt
```

1. Next, you can create your noisy data points. For demonstration purposes, let's generate some noisy data:

python

Copy code

```
# Generate example noisy data
x = np.linspace(0, 10, 20)
y = np.sin(x) + np.random.normal(0, 0.2, len(x))
```

1. Now, you can create the cubic spline interpolation with zero derivative boundary conditions:

python

Copy code

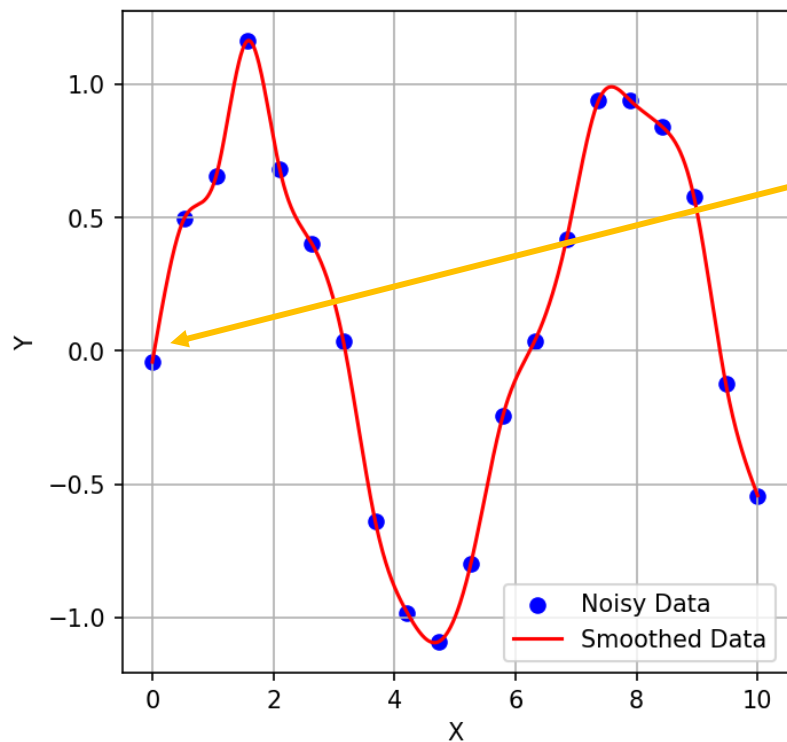
```
# Create a cubic spline interpolation with zero derivative boundary conditions
cs = CubicSpline(x, y, bc_type=((2, 0.0), (2, 0.0)))
```



# Coding (II)

## Second derivative?

In the `bc_type` parameter, `(2, 0.0)` specifies that the second derivative (i.e., curvature) should be zero at both ends, ensuring that the derivative at the boundaries is zero.



The slope does not look like as requested!

1. To obtain smoothed data points, you can evaluate the cubic spline at a higher resolution of `x` values:

```
python  
  
x_smooth = np.linspace(0, 10, 1000)  
y_smooth = cs(x_smooth)
```

1. Finally, you can visualize the original noisy data and the smoothed data using a plot:

```
python  
  
# Plot noisy data and smoothed data  
plt.figure(figsize=(10, 6))  
plt.scatter(x, y, label='Noisy Data', color='blue')  
plt.plot(x_smooth, y_smooth, label='Smoothed Data', color='red')  
plt.xlabel('X')  
plt.ylabel('Y')  
plt.legend()  
plt.grid()  
plt.show()
```

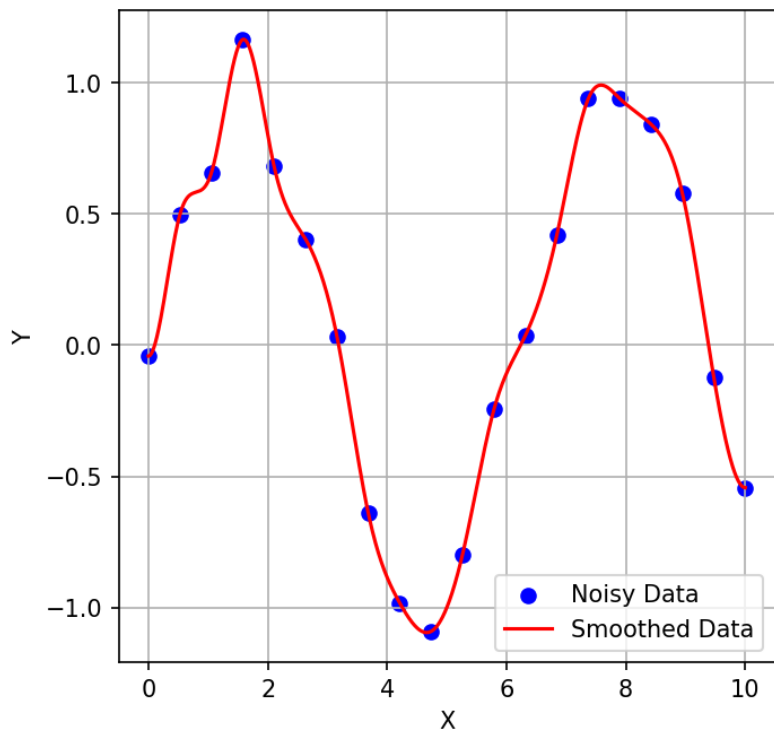
This code will generate a plot with your noisy data and the smoothed data, ensuring that the derivative at the boundaries is zero.

You can adjust the parameters and customize the smoothing further if needed, depending on your specific requirements and dataset.



# Coding (III)

Ask for review!



I've asked for the derivative to be zero in the extremes, but you provided a code with the second derivative being zero. What was the reason you choose the second derivative instead of the first?



I apologize for the misunderstanding. You are correct; you initially asked for the derivative to be zero at the boundaries, which implies setting the first derivative to zero. The code I provided with the second derivative being zero was not in line with your original request.

To ensure that the first derivative (slope) is zero at the boundaries while smoothing the data, you can modify the boundary conditions when creating the cubic spline. Here's the corrected code:

```
python Copy code

import numpy as np
from scipy.interpolate import CubicSpline
import matplotlib.pyplot as plt

# Generate example noisy data
x = np.linspace(0, 10, 20)
y = np.sin(x) + np.random.normal(0, 0.2, len(x))

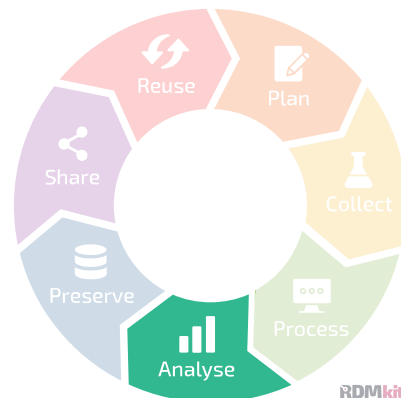
# Create a cubic spline interpolation with zero first derivative boundary conditions
cs = CubicSpline(x, y, bc_type='clamped') # 'clamped' enforces zero first derivative

# Evaluate the cubic spline at a higher resolution of x values
x_smooth = np.linspace(0, 10, 1000)
y_smooth = cs(x_smooth)

# Plot noisy data and smoothed data
plt.figure(figsize=(10, 6))
plt.scatter(x, y, label='Noisy Data', color='blue')
plt.plot(x_smooth, y_smooth, label='Smoothed Data', color='red')
plt.xlabel('X')
plt.ylabel('Y')
plt.legend()
plt.grid()
plt.show()
```

# Long forgotten formulas

Name	Grade1	Grade2	Grade3	Grade4	Average
-----	-----	-----	-----	-----	-----
Alice	4.2	3.5	1.0	4.7	3.8
Bob	2.9	Text	4.2	5.1	#N/A
Charlie	3.7	4.0	3.2	6.3	#N/A
David	1.5	2.8	3.9	4.1	3.075
Emma	4.8	3.6	2.7	4.9	4.25
Frank	3.3	2.4	NaN	3.9	#N/A
Grace	4.7	2.1	3.9	5.2	3.975
Hannah	2.6	3.9	1.8	4.5	3.2
Ian	5.3	3.7	3.4	4.0	#N/A
Jack	3.1	4.5	3.6	2.7	3.475



I have an excel sheet that contains the grades of students. I want to color the cells with the following rules: if between 0 and 1 cell is red; if between 1 and 3 cell is orange; if between 3 and 5 cell is green; if out of the range 0 and 5 or is not a number then it is black.



You can use Excel's conditional formatting feature to apply cell colors based on the rules you mentioned. Here's how you can set up conditional formatting for your grade range:

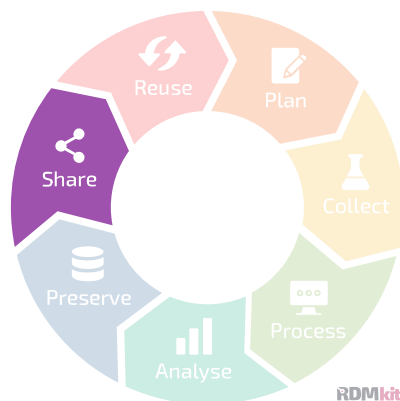
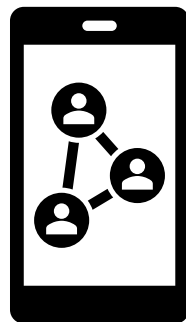
1. Select the cells where you want to apply conditional formatting. For example, if your grades are in column A from row 2 downwards, select all those cells.
2. Go to the "Home" tab on the Excel ribbon.
3. In the "Styles" group, click on "Conditional Formatting."
4. Choose "New Rule" from the dropdown menu.
5. In the "New Formatting Rule" dialog box, select "Use a formula to determine which cells to format."
6. In the "Format values where this formula is true" field, enter the following formula for red (grades between 0 and 1):

```
excel Copy code
=AND(A2>=0, A2<=1, ISNUMBER(A2))
```

Note that A2 refers to the active cell. If you started with a different cell, adjust it accordingly.

7. Click on the "Format" button to choose the red color for the cells. Go to the "Fill" tab and select the red color you want, then click "OK."
8. Click "OK" in the "New Formatting Rule" dialog box to apply the rule.
9. Now, repeat steps 4 to 8 for the other conditions:
  - For orange (grades between 1 and 3):
    - Formula: `=AND(A2>1, A2<=3, ISNUMBER(A2))`
    - Choose an orange fill color.
  - For green (grades between 3 and 5):
    - Formula: `=AND(A2>3, A2<=5, ISNUMBER(A2))`
    - Choose a green fill color.
  - For black (out of the range or not a number):
    - Formula: `=OR(A2<0, A2>5, ISNUMBER(A2)=FALSE)`
    - Choose a black fill color.
10. Click "OK" to apply each rule.

# Text review (I)



Use the following abstract and make a description (max characters as allowed by Twitter and use hashtags in the end) for a picture of the work to be shared on Instagram:

Active fabrics, responding autonomously to environmental changes, are the “Holy Grail” of the development of smart textiles. Liquid crystal elastomers (LCEs) promise to be the base materials for large-stroke reversible actuation. The mechanical behavior of LCEs matches almost exactly the human muscle. Yet, it has not been possible to produce filaments from LCEs that will be suitable for standard textile production methods, such as weaving. Based on the recent development of LCE fibers, here, the crafting of active fabrics incorporating LCE yarn, woven on a standard loom, giving control over the weave density and structure, is presented. Two types of LCE yarns (soft and stiff) and their incorporation into several weaving patterns are tested, and the “champions” identified: the twill pattern with stiffer LCE yarn that shows the greatest blocking force of 1–2 N cm<sup>-1</sup>, and the weft rib pattern with over 10% reversible actuation strain on repeated heating cycles. Reversible 3D shape changes of active fabric utilize the circular weaving patterns that lead to cone shapes upon heating. The seamless combination of active LCE yarns into the rich portfolio of existing passive yarns can be transformative in creating new stimuli-responsive actuating textiles.



"🔪 Crafting Smart Fabrics: Weaving the Future! ✨

Unveiling the "Holy Grail" of smart textiles – Active Fabrics! 💡 Liquid Crystal Elastomers (#LCEs) are the game-changers, mimicking human muscle behavior. 🤖 But wait, there's more! We've woven LCE yarns into standard looms, giving us control over density and structure. 🧵

Discover the "Champions" 🏆: Stiffer LCE yarns in twill pattern for robust blocking force, and weft rib pattern with over 10% reversible actuation strain. 🌐

Watch as our fabric transforms into 3D shapes upon heating 🔥 – cone shapes galore! 🚀

Join us in revolutionizing textiles with seamless integration of active LCE yarns! 🌐  
#SmartTextiles #FutureFabrics #Innovation #TechFashion #MaterialScience"

# Text review (II)



Rewrite the following text and make it formal and suitable for a first contact by email:  
" Hey dude Smith. I saw your ppt in Munich and our work cross like a lot. I am also intereseted in the brain and stuff, so I usually apply some electrode in the brain and then I extract the data with machine learning. Would you be interested in having a chat soon? it can be f2f or in zoom.  
Kiitos tosi paljon!

Pedro Silva



Subject: Introduction and Expression of Interest

Dear Mr. Smith,

I trust this email finds you well. My name is Pedro Silva, and I recently had the pleasure of coming across your presentation during an event in Munich. I was greatly impressed by your work, and I noticed significant commonalities between our respective areas of interest.

I share a strong enthusiasm for the intricacies of the human brain and related fields. My own research involves the application of electrodes to the brain, followed by the extraction and analysis of data using advanced machine learning techniques.

Given our shared interests and the potential for meaningful collaboration, I would like to propose the possibility of arranging a discussion in the near future. This conversation could take place face-to-face, should our schedules align, or via a virtual meeting platform such as Zoom.

I genuinely believe that our exchange of ideas and experiences could prove to be mutually beneficial. Please let me know if you would be open to this idea and, if so, your availability in the coming weeks.

Thank you very much for your time, and I look forward to the opportunity of connecting with you.

Warm regards,

Pedro Silva

## 4. Limitations, ethical and legal considerations



# Limitations & risks of (generative) AI

- **Misinformation**
  - You will experience ~~hallucinations~~ confabulations
- LLMs **cannot reason**, but they can produce output that can be meaningful to you (Enrico's attempt with the classic example of the goat, wolf, and cabbage riddle)
- **Reinforcement of biases** (gender, racial, cultural, ...)
  - They are designed to output what has most frequently been seen, not what should be the correct answer to your question
- Lack of reliable sources
  - **LLMs are not search engines**, they are word generators
- Safeguards are in place, but they can be bypassed.



# Responsible conduct of research

- **Allea principles**

- **Reliability** in ensuring the quality of research, reflected in the design, methodology, analysis, and use of resources.
- **Honesty** in developing, undertaking, reviewing, reporting, and communicating research in a transparent, fair, full, and unbiased way.
- **Respect** for colleagues, research participants, research subjects, society, ecosystems, cultural heritage, and the environment.
- **Accountability** for the research from idea to publication, for its management and organisation, for training, supervision, and mentoring, and for its wider societal impacts.

- **TENK definition of research misconduct**

- Fabrication (false data)
- Falsification (false results)
- Plagiarism (reusing without citing the sources)

# Nothing new here, actually!

- AI cannot be given **authorship**
- Use of AI should be **transparent** by describing how AI is used
- **Personal data protection** (GDPR) is to be followed
- Protect your own **unpublished work**
- The output of the AI system can be **sensitive**
- When creating artistic output, **contributor roles must be explained transparently and specifically.**
- **Reproducibility** can be difficult to achieve
- **Accountability:** you are responsible for what you present

# Be aware of the AI Hype

- You have most likely experienced the **AI Hype**
  - grandiose narratives
  - sparks of AGI
  - AI scaremongering
- We are not going to end up like in the movie *The Matrix*
- The low-barrier for creating all sorts of AI-LLMs applications also means lots of uncertain on where your personal and intellectual data will go

**Nothing new here, it is just another technology for researchers and we need to evaluate if it can be useful for our work.**

# AI governance, AI ethics

- European Commission to soon approve the **AI act**: a framework for risk evaluation of AI systems
- **AI governance** is what governments and companies are right now struggling with (see CAIDP's AI governance recommendations for the UN)
- As researchers we have more freedom of experimenting with these new technologies, **within the boundaries of responsible conduct of research**
  - Aalto ethics committee recommends having an **external ethics advisor for AI**. **FCAI** can help you find one.

## 5. Advanced cases and other tools



# Some more advanced uses of AI

1. **Not just chat:** all other sorts of NLP Natural Language Processing
2. **Local open source LLMs** for increased privacy, safety, and reproducibility
3. **Develop** your AI application (paid APIs, local fine-tuning with OSS)
4. Existing **external applications** for academics
5. ChatGPT-4 plugins and other **OpenAI advanced features**
6. **Beyond text:** Midjourney, Dall.e, ...

# 1. Language models beyond chats

- **Not just chat-bots!** Language models are useful for example for **automatic transcription of interviews**.
- **OpenAI Whisper** is available on our **Triton HPC cluster** and **SECDATA** environment, useful when interviews should not leave Aalto storage. See the **speech2text** app developed by **Aalto RSEs**.
- Other uses are **Named Entity Recognition** for automatic anonymisation of text, or to identify when a text is sensitive
- **Neural Machine Translation** with open-source models such as **No Language Left Behind** (by Meta)

# 2. Local open source LLMs

- **Open-source models**
  - You can download them and run them on your laptop
  - Also available on Aalto Triton HPC cluster
  - Personal favourite: **model Llama2** (and CodeLlama) and llamacpp for **inference**.
- **Why local?**
  - **Cloud scenario: you need Internet and data leaves you**
    - Your Computer -> ChatGPT website -> remote inference software -> remote model
  - **Local scenario: no internet needed, data stays with you**
    - Your computer (or Aalto Triton HPC) -> local terminal/website -> local inference software -> local model



# 3. Build your AI app

- Build a chatbot with Azure OpenAI API
- Or **build it locally** using Llamacpp + open-source models + development tools such as LangChain, LangSmith, PineCone
- **Fine tuning**
  - Prepare a **new set of instructions** (e.g. questions and answers)
  - Tweak the model weights so that the model is ***fine-tuned*** with the new content
  - It can be computationally expensive (many GPUs), but there are efficient ways with few compromises (qLORA)

# 4. External AI applications

- Here a list of Apps that can be useful for academics
  - ChatPDF
  - Dante.ai
  - Cursor
  - GitHub copilot
  - ...what is your favourite?

# 5. ChatGPT-4 new features

- **For 20 USD/month** you can access GPT4 with **Beta** features
  - Custom instructions
  - Advanced data analysis
  - Better model
  - **Plugins**
    - Beware with plugins since it is not always clear where your data will go

# 6. Beyond text

1. DALL·E 3 (openai.com)
2. Midjourney
3. Replicate

## 6. Conclusions and references



# Take home message

- **AI is a junior assistant** you can hire: you don't maybe want to assign the most critical tasks to them, you want to assign tedious tasks where you are an expert so that it's easy to verify the output or assign exploratory tasks for you later to evaluate
- Autoregressive LLMs as they are today **are not "intelligent"**. They are great at generating an average of all the things they were exposed to.... And carry all the **biases** and **errors** from the past
- **Responsible Conduct of Research** is at the core of any academic activity that you do: be honest and transparent

# Where to go from here?

- There are great **resources for self learning**, beware of scammers. A reliable list available at [this link](#)
- **Aalto researchers**: are you unsure of the possibilities for your research? Talk with us: [researchdata@aalto.fi](mailto:researchdata@aalto.fi) is a network of networks that works closely with IT services, Ethics, Legal, Software Development Services.
- **Teachers**: there are similar considerations for teachers, but **protecting student data is of absolute importance**. It would require a dedicated training session; most likely you are already discussing this with your colleagues.

# References

- **Unesco guidelines for generative AI**
- **Excellent guide from Aarhus University**
- **Jeremy Howards's Hacker guide to Language Models**
- **Elements of AI open course**
- **Ethics of AI open course**

A list of questions and answer from the live training session is available at <https://www.aalto.fi/en/open-science-and-research/ai-and-research-work-useful-learning-materials> .



**Thank you!**



**[aalto.fi/en/open-science-and-research](https://aalto.fi/en/open-science-and-research)**

**A”**

Aalto-yliopisto  
Aalto-universitetet  
Aalto University