

Autumn 2025

Lib4RI Training Series: Module 8a – AI Made Easy: A Beginner's Toolkit for Streamlining Your Projects

Dr. Corin Otesteanu

Table of Contents

1. Introduction

2. Convolutional neural networks

- How CNNs work
- Convolution filters
- Visual (interactive) explanation
- Applications

3. Large-language models

- How LLMs work
- Visual (interactive) explanation
- Current models
- Applications
- Pros and cons

4. Conclusion

Introduction


Artificial Intelligence

- AI refers to machines mimicking human intelligence
 - term was coined at the Dartmouth Conference in **1955**
- To perform tasks such as
 - Recognizing and understanding images
 - Understanding language
 - Reasoning and making decisions


J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

- Machine Learning (ML): computer systems that learn from data, without following explicit instructions
- Deep Learning (DL): a subset of ML, focusing on neural networks
- Generative AI: systems that create new data from learned patterns
- Large language models (LLMs): deep learning models trained on massive text data to generate human-like text



Early AI vs ML vs. DL



How modern AI Models Work

1. Training Phase:


- Models learn from (vast amounts of) data
- Minimize a loss function
 - Identify patterns and relationships
 - Adjust internal parameters

2. Inference Phase:

- Internal parameters are fixed
- Apply learned patterns to new data
- Generate predictions or outputs

<https://theresanaiforthat.com/>

- Database with latest AI tools



Exercise

Go to ChatGPT and prompt the following:

- Prompt: Act as an expert in [your field of expertise]. Do a literature online search to find what are the latest developments in the field of [specific area in your field of expertise] and the research gaps in the area? Next propose novel research ideas, thinking about the following steps: importance, impact and implementation feasibility. Finally list the 3 most promising areas of research
- Rate the promising research areas on a scale from 1-10


Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs)

- a class of deep neural networks, primarily used for analyzing visual imagery
- utilize convolutional layers to detect features in images
- capable of learning **complex patterns** in large amounts of data
- fundamental building blocks in many modern deep learning architectures

How CNNs work

- Input layer: receives raw pixel values
- Convolutional layers: apply filters to detect features
- Activation functions: introduce nonlinearity
- Pooling layers: reduce spatial dimensions
- Fully connected layers: combine features for final output



Naranjo-Torres, J.; Mora, M.; Hernández-García, R.; Barrientos, R.J.; Fredes, C.; Valenzuela, A. A Review of Convolutional Neural Network Applied to Fruit Image Processing. Appl. Sci. 2020, 10, 3443. <https://doi.org/10.3390/app10103443>


How CNNs work

- Input layer: receives raw pixel values
- Convolutional layers: apply filters to detect features
- Activation functions: introduce nonlinearity
- Pooling layers: reduce spatial dimensions
- Fully connected layers: combine features for final output

- Loss calculation: compare output with expected result
- Backpropagation: compute the gradient of the loss function with respect to each weight (filter), indicating the direction and magnitude of updates to minimize error
- Optimization: update parameters using algorithms like gradient descent until convergence

CNN explainer

<https://poloclub.github.io/cnn-explainer/>



Wang, Zijie J., et al. "CNN explainer: learning convolutional neural networks with interactive visualization." *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020): 1396-1406.

CNNs applications

- Image classification
- Object detection
- Facial recognition
- Image segmentation
- Image captioning


Classification



**Classification
+ Localization**



Object Detection




**Instance
Segmentation**



Example of CNN application: image segmentation

- Image segmentation is a technique that divides an image into multiple regions, where each pixel in a region shares similar characteristics
 - often domain specific
- Meta created the Segment Anything Model
 - domain agnostic
- SAM2 can also be applied to video in addition to image data to segment regions on interest

Example: Meta SAM2



Example: Meta SAM2

Try it yourself:

<https://huggingface.co/spaces/wondervictor/evf-sam2>

EVF-SAM-2-Image

EVF-SAM-2-Video

Input Image



Drop Image Here
- or -
Click to Upload





Output Image



Prompt

Use a phrase or sentence to describe the object you want to segment. Currently we only support English

Submit

☐ check this if you want to segment body parts or background or multi objects (only available with latest evf-sam checkpoint)

☐ semantic level

Segment Anything Model

- Go to <https://segment-anything.com/demo>
- Upload an image that you want to segment (delineate an object)
- Use it programmatically by downloading it from https://huggingface.co/docs/transformers/main/en/model_doc/sam_hq

Segment Anything
Research by Meta AI

Home [Demo](#)

When hovering over the image, SAM is running in the browser.

Tools

Upload Gallery

Hover & Click
Click an object one or more times.
Shift-click to remove regions.

+ -
Add Mask Remove Area

Reset Undo Redo

Multi-mask
Cut out object

Box

Everything

Cut-Outs

EVF-SAM-2-Image EVF-SAM-2-Video

Input Image Output Image

Prompt
Use a phrase or sentence to describe the object you want to segment. Currently we only support English

Submit


Comparison with ChatGPT

Segmented Cell with Bounding Box




The bounding box coordinates for the segmented cell are (**x=0, y=0, width=200, height=166**). However, this seems to include the entire image rather than just the cell. Let me refine the segmentation to ensure the bounding box tightly surrounds the cell.

Refined Segmented Cell with Bounding Box



The refined bounding box coordinates for the segmented cell are (**x=66, y=86, width=62, height=62**).

Segmented Cell with Bounding Box




The bounding box coordinates for the segmented cell in this image are (**x=60, y=3, width=74, height=71**).

DeepImageJ

- user-friendly plugin for using a variety of pre-trained neural networks in ImageJ and Fiji in life science applications
- Download from <https://deepimagej.github.io/>

Search for available models for your applications.




Large language models (LLMs)

Large language models (LLMs)

- AI models trained on vast amounts of text data to understand and generate human-like text
- Massive scale (billions of parameters)
- Self-supervised learning on diverse text corpora
- Ability to perform multiple tasks sometimes without specific training
 - Text generation and summarization
 - Question answering
 - Language translation
 - Conversational AI and chatbots

Large language models (LLMs)



How LLMs work

○ Tokenization

- LLMs process text by breaking it down into tokens, which can be whole words or parts of words

○ Embeddings

- each token is converted into a dense numerical vector representation

○ Transformer backbone (Attention Mechanism)


- LLMs use attention to understand the relationships between words in a sentence

○ Training Objective

- LLMs are trained on large datasets to predict the next word in a sequence, to generate a coherent plausible text

LLMs explained

<https://poloclub.github.io/transformer-explainer/>



Cho, Aeree, et al. "Transformer Explainer: Interactive Learning of Text-Generative Models." *arXiv preprint arXiv:2408.04619* (2024).


Large language models overview

Model Name	Image Generation	Internet Access	Reasoning	File Upload	Context Window	Data Privacy	Price	Deployment
ChatGPT	Yes	Yes	Yes	text, pdf, images	128,000 tokens	opt out available	free \$20/month	Cloud
Grok	Yes	Yes	Yes	text, pdf, images	128,000 tokens	opt out available	free \$30/month	Cloud
DeepSeek	No	Yes	Yes	text, pdf, text from images	128,000 tokens	No	free, API-based	Cloud
Llama 3	Yes	No	No	text	128,000 tokens	Yes	free	Local
Gemini	Yes	Yes	Yes	text, pdf, images	2 million tokens	No	free, \$20/month	Cloud
Claude	No	No	Yes	text, pdf, images	200,000 tokens	Yes, except for feedback	free, \$20/month	Cloud
Copilot	Yes	Yes	Yes	text, pdf, images	128,000 tokens	No	free, \$20/month	Cloud
Copilot Enterprise	Yes	Yes	No	text, pdf	128,000 tokens	Yes	free, \$25/month	Cloud

Large language models performance comparison


INTELLIGENCE

Artificial Analysis Intelligence Index; Higher is better




SPEED

Output Tokens per Second; Higher is better



PRICE

USD per 1M Tokens; Lower is better



Large language models performance comparison

Model 250 / 250	Overall ↑↓	Hard Prompts ↑↓	Coding ↑↓	Math ↑↓	Creative Writing ↑↓	Instruction Following	Longer Query ↑↓	Multi-Turn ↑↓
 claude-opus-4-1-2...	1	1	1	1	1	1	1	1
 gemini-2.5-pro	1	2	3	1	1	1	2	1
 chatgpt-4o-latest...	2	4	3	15	2	5	5	1
 claude-opus-4-1-2...	2	2	1	1	2	1	1	1
 gpt-4.5-preview-2...	2	6	4	8	1	3	3	1
 gpt-5-high	2	3	3	1	8	5	13	7
 o3-2025-04-16	2	4	4	1	8	7	18	8
 qwen3-max-preview	3	2	3	1	7	4	3	3
 grok-4-fast	7	5	4	1	7	7	5	3
 gpt-5-chat	8	4	4	8	7	5	5	2
 qwen3-max-2025-09...	8	4	2	1	7	5	5	2
 qwen3-v1-235b-a22...	8	2	3	1	8	2	2	5
 claude-opus-4-202...	9	4	3	3	2	2	2	6
 deepseek-r1-0528	9	12	5	8	7	17	15	14
 deepseek-v3.1	9	12	8	3	7	8	7	12
 deepseek-v3.1-ter...	9	12	11	3	2	7	5	15
 deepseek-v3.1-ter...	9	4	3	1	7	4	3	7

Chiang W, Zheng L, Sheng Y, et al. Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference. arXiv. 2024;2403.04132

<https://lmarena.ai/?leaderboard>, accessed 02/10/2025

Large language models performance - language

Chatbot Arena Overview (Language)									Sort by Rank		Sort by Arena Score	
Model	English	Chinese	German	French	Spanish	Russian	Japanese	Korean				
grok-3-preview-02-24	1	1	1	1	1	1	1	1				
gpt-4.5-preview-2025-02-27	1	1	1	1	1	1	1	1				
gemini-2.0-flash-thinking-exp-01-21	3	1	2	1	1	1	2	1				
gemini-2.0-pro-exp-02-05	3	3	2	1	1	1	2	1				
chatgpt-4o-latest-20250129	3	1	2	1	1	3	2	1				
deepseek-r1	4	1	2	1	1	5	2	1				
gemini-2.0-flash-001	6	2	2	1	1	5	8	1				
o1-preview	6	12	7	2	4	9	2	6				
o1-2024-12-17	8	2	3	1	1	5	2	1				
qwen2.5-max	10	3	2	1	2	6	5	1				
o3-mini-high	10	1	3	1	1	8	3	1				
deepseek-v3	10	12	4	4	2	7	8	7				
step-2-16k-exp-202412	10	12	8	4	4	6	6	7				
glm-4-plus-0111	11	2	2	4	2	7	9	4				
qwen-plus-0125	11	9	7	2	2	6	9	6				
claude-3-7-sonnet-20250219	11	12	7	2	2	6	3	1				
o1-mini	11	15	10	4	3	18	10	13				
gemini-2.0-flash-lite-preview-02-05	13	12	8	2	3	6	5	4				

Chiang W, Zheng L, Sheng Y, et al. Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference. arXiv. 2024;2403.04132
<https://lmarena.ai/?leaderboard>, accessed 11/03/2025


Example: Creating a GUI app

Prompt



- Make a streamlit app. The app should have three figures that plots three time signals. A fourth figure should show the spectrum of the summed up time signals. There should be sliders that adjust the amplitude of each signal from 0 to 10 and the frequency of the signals from 0.1 to 50 Hz. The scale of the plots should be fixed to 10

Example: Creating a GUI app


Signal 1



Signal 2



Spectrum of Combined Signals



Example: Solving electric circuit problem


Prompt: For the circuit show below find the current and voltage for each of the five resistors.

Response:

- To find the current and voltage for each resistor the given circuit, we can follow these steps:

1. Determine the total resistance of the circuit.
2. Calculate the total current supplied by the voltage source.
3. Find the voltage drop across each resistor using Ohm's Law.
4. Calculate the current through each resistor.

- Here's how we can do this step-by-step: [...]



Example: Solving electric circuit problem

Summary

- Current through R1R_1R1 : 0.357A0.357 A0.357A
- Voltage across R1R_1R1 : 3.57V3.57 V3.57V
- Current through R2R_2R2 : 0.286A0.286 A0.286A
- Voltage across R2R_2R2 : 1.43V1.43 V1.43V
- Current through R3R_3R3 : 0.0715A0.0715 A0.0715A
- Voltage across R3R_3R3 : 1.32V1.32 V1.32V
- Current through R4R_4R4 : 0.055A0.055 A0.055A
- Voltage across R4R_4R4 : 0.11V0.11 V0.11V
- Current through R5R_5R5 : 0.0183A0.0183 A0.0183A
- Voltage across R5R_5R5 : 0.11V0.11 V0.11V

- This approach involves simplifying the circuit step-by-step and applying Ohm's Law and Kirchhoff's Laws to find the desired currents and voltages.

Example: Optical character recognition

- Screen capture from a youtube video
- Prompt: Convert this image to text. for each bullet point provide an equation

Heart Rate Variability Features

- Inter Beat Interval to calculate Pulse Rate Variability:
- High Frequency (eg through DFT) 0.15–0.4 Hz
- Low Frequency 0.04–0.15 Hz
- VLF power (0.015–0.04 Hz)
- RMSSD: Root mean square of successive differences of IBI
- pNN50: proportion of successive IBIs that differ more than 50ms over total IBIs
- Delta IBIs
- Mean heart rate
- 90th percentile heart rate
- 10th percentile heart rate

○ Response

I have created a PowerPoint slide with the features of heart rate variability, including their corresponding equations. You can download it using the link below:

[Download Heart Rate Variability Features Presentation \[↗\]](#)

Optical character recognition - example

1. Inter Beat Interval (IBI): $IBI = Peak_n - Peak_{n-1}$
2. High Frequency (HF): HF Power = DFT(0.15 – 0.4 Hz)
3. Low Frequency (LF): LF Power = DFT(0.04 – 0.15 Hz)
4. VLF Power: VLF Power = DFT(0.015 – 0.04 Hz)
5. RMSSD: $RMSSD = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N-1} (IBI_i - IBI_{i+1})^2}$
6. pNN50: $pNN50 = \frac{\text{Number of } |IBI_i - IBI_{i+1}| > 50 \text{ ms}}{\text{Total Number of IBIs}} \times 100\%$
7. Delta IBIs: $\Delta IBI_i = IBI_i - IBI_{i-1}$
8. Mean Heart Rate: $\text{Mean HR} = \frac{60}{\text{Mean IBI}}$
9. 90th Percentile Heart Rate: $HR_{90} = \text{Percentile}(HR, 90)$
10. 10th Percentile Heart Rate: $HR_{10} = \text{Percentile}(HR, 10)$

Writing social media posts

- write a post for LinkedIn based on the information given in this article.
- Make it catchy and visually appealing: <https://ethz.ch/en/news-and-events/eth-news/news/2025/01/four-snsf-advanced-grants-go-to-eth-zurich-researchers.html>

☀ Celebrating Excellence in Research at ETH Zurich! ☀

We're thrilled to share that **four ETH Zurich professors** have been awarded the **prestigious SNSF Advanced Grants**, securing a total of **8.5 million CHF** for groundbreaking research! 🎉

These grants will support pioneering projects over the next five years, pushing the boundaries of innovation and discovery. 🧠💡

💡 **Meet the Visionaries and Their Projects:**

🔬 **Prof. Kirsten Bomblies** (Biology)

🌱 Tackling fertility challenges in polyploid plants to revolutionize plant breeding and stress tolerance.

💻 **Prof. Jonathan Home** (Physics)

⚙️ Advancing quantum computing with robust ion storage techniques for more reliable systems.

🏗️ **Prof. Nicola Spaldin** (Materials)

🌀 Unlocking the mysteries of chiral materials with a new theory of crystal chirality.

🧠 **Prof. Mehmet Fatih Yanik** (Neurotechnology)

🧬 Innovating precise neural interventions to treat resistant psychiatric and neurological disorders.

Where LLMs excel

- Conversations (chatbot, customer service, etc)
- Writing in a certain style
- Translations
- Correcting grammar
- Summarization
- Brainstorming (limited)
- Coding (limited)

LLM limitations and downsides

Can produce factually incorrect results

- biases
- hallucinations (fabricated responses)
 - invented facts
 - fictitious citations
 - incorrect explanations
 - non-existent entities
- no or limited reasoning capabilities (as of now)
- lack of knowledge
 - knowledge cutoff
 - lack of domain specific knowledge
- generated text might constitute *plagiarism*

Causes

- biased training data
- model architecture
 - generate novel ideas, styles, genres
 - propose valuable hypotheses
 - (unknown) errors
 - known factual errors
- model architecture
- limited and outdated training data
- training data and prompting

LLM limitations and downsides

- LLMs process text using tokens, not characters.
- Tokens can be single characters or entire words.
- Tokenization makes direct character counting difficult for LLMs.

how many letters does this prompt have?

It seems there was an issue calculating the character count directly. I'll count the letters here manually:

The prompt you provided has 296 characters. [\[>-\]](#)

- Workaround: ask AI to use programming language to count the characters to generate correct result

Minimizing LLM inaccuracies - prompting

Prompt: specific input or instruction to a language model to generate a desired output

Good practice for prompt design

- Be specific and clear in the task you want done
- Break down complex tasks in simpler tasks
- Provide constraints
 - E.g. 3 bulletpoints, one paragraph, etc
- Style
 - formal, informal
 - academic, journalistic, creative
- After output, iterate, refine prompts

<https://docs.anthropic.com/en/docs/build-with-claude/prompt-engineering/overview>

Minimizing LLM inaccuracies - prompting

Prompt: specific input or instruction to a language model to generate a desired output

Good practice identifying factually incorrect results

- Prompt for alternative point of view
- Ask for the references
- Ask for positive or negative of a certain subject
 - If you ask for a proof, ask also to prove the opposite
- Use the same prompt on several LLMs and compare results
- When in doubt, always check with a reputable resource

You are responsible for data generated by genAI

Prompt template example

- Role Definition
 - Define the role of the AI clearly
- Context:
 - Provide any necessary background information or context
- Input Data:
 - Specify the data or information the AI should use
 - Upload your own data
- Instructions:
 - Give clear and specific instructions on what you want the AI to do
- Constraints
 - Mention any constraints like length/bulletpoint, format, or style
- Other
 - Multi-shot prompting (if applicable)
 - Chain of thought

Prompt example 1

- Role Definition

- You are an AI assistant that provides critical reviews of scientific research papers

- Context:

- The following text is a research paper on the effects of microplastics on marine life

- Input Data:

- [Insert of upload your own data]

- Instructions:

- Critically review the research paper, focusing on the methodology, data analysis, and conclusions. Highlight any strengths and weaknesses, and suggest areas for improvement.

- Constraints

- The analysis should have 10 points. Use formal language.

- Other (optional)

- Multi-shot prompting: example of a good review
 - Chain of thought: provide a step by step analysis for each point

Prompt example 2

○ Role Definition

- You are an expert in Python data visualization with extensive experience in scientific plotting.

○ Context:

- I'm a researcher in climate science working with temperature time series data.

○ Input Data:

- No input data

○ Instructions:

- Create a comprehensive visualization of global temperature anomalies over the past century

○ Constraints

- Use matplotlib library. Data should be between 1924 and 2024. Should be publication-ready quality. The code to reproduce the plots should be given as a jupyter notebook

Exercise

Consider (an already published conference paper or document related to your work)

- Upload the results and discussion section to ChatGPT and prompt one of the following:
 - I am a
 - research scientist in ...
 - technician using device
 - IT support working with ...
 - Administrative assistant in charge of ...
 - please summarize the following document as accurately as possible with an emphasis on the quantitative and qualitative aspects presented in the document.
- Rate the summary


Exercise

Consider (an already published conference paper or document related to your work)


- Upload the results and discussion section to ChatGPT and prompt one of the following:
 - Act as a scientific reviewer for a journal paper. Go through the key sections to assess its quality, structure, methodology, and scientific contribution. Provide an assessment of quality of the paper
 - Act as a manager and provide feedback on the following document provided by a
 - administrative assistant
 - technician
 - IT support
 - Assess the clarity, organization, completeness, technical content (if applicable) and spelling
- Rate the review

Accessing LLMs

- Through their respective website
- On Imarena.ai
 - select desired model
- On huggingface.co/chat
 - select desired open access model
- On NVidia
 - <https://build.nvidia.com/explore/discover>
 - select desired open access model



Customizing LLM preferences - ChatGPT



- Custom instructions from <https://help.openai.com/>:
- How formal or casual should ChatGPT be?
- How long or short should responses generally be?
- How do you want to be addressed?
- Should ChatGPT have opinions on topics or remain neutral?

Customizing LLM preferences - ChatGPT

The image displays three overlapping screenshots of the OpenAI ChatGPT settings interface, illustrating the steps to customize LLM preferences.

Left Screenshot: Shows the main sidebar. The 'Settings' option, represented by a gear icon, is circled in red. Other options include 'My GPTs', 'Customize ChatGPT', 'Upgrade Plan', and 'Log out'.

Middle Screenshot: Shows the 'Settings' modal. The 'Personalization' option, represented by a person icon, is selected and highlighted. Other options include 'General', 'Speech', 'Data controls', 'Builder profile', 'Connected apps', and 'Security'.


Right Screenshot: Shows the 'Personalization' settings. It includes the following options and controls:

- General:** Improve the model for everyone (Off >)
- Personalization:** Shared links (Manage)
- Speech:** Export data (Export)
- Data controls:** Delete account (Delete)
- Builder profile:**
- Connected apps:**
- Security:**

<https://help.openai.com/en/articles/7730893-data-controls-faq>

Customizing LLM preferences - Claude

- Click on your username (bottom left corner) and then settings. Enable artifacts.



The screenshot shows the Claude settings interface. On the left, a sidebar displays the name 'Claude' and a profile icon (a black circle with a white 'c'). A red arrow points from this icon to the 'Artifacts' toggle switch in the settings panel. The settings panel is divided into several sections:

- What best describes your work?**: A dropdown menu currently set to 'Engineering'.
- What personal preferences should Claude consider**: A text area containing the following preferences:
 - I typically write code in TypeScript.
 - I like approaching code architecture through domain knowledge.
 - When learning new concepts, I find analogies particularly helpful.
 - In project planning, I tend to work backwards from the goal.
- Claude capabilities**:
 - Artifacts**: A toggle switch that is turned on. Description: 'Ask Claude to generate content like code snippets, text documents, or website designs, and Claude will create an Artifact that appears in a dedicated window alongside your conversation.'
 - Prompt examples**: A toggle switch that is turned off. Description: 'Displays example prompts in the chat input to help you get started.'
 - CSV chat suggestions**: A toggle switch that is turned on. Description: 'Claude will suggest responses when you upload CSVs to your conversation.'
- Connected accounts**: A section for connecting other services. It shows a 'GitHub' account with a 'Connect' button next to it.

Customizing LLM preferences - Claude

- Click on the chose style button.

How should Claude write responses?

Normal ✓

Concise

Explanatory

Formal

Create & Edit Styles

Default responses from Claude

Claude 3.7 Sonnet ▾  Choose style ▾



Share a writing example or describe your style and Claude will make a custom writing style tailored just for you.

[Add Writing Example](#)

Select, drop, or paste existing doc, post, message, etc.
Content is not stored after matching style.

[Describe style instead](#)

Customizing LLM preferences - Claude

- Click on your username (bottom left corner) and then enable feature preview tools
- Built-in Code Sandbox enables Claude to perform complex calculations, data analysis, and iterative testing


Feature Preview
✕

Preview and provide feedback on upcoming enhancements to our platform. Please note: experimental features might influence Claude's behavior and some interactions may differ from the standard experience.

Analysis tool

LaTeX Rendering

sales_funnel_data

CSV

Please visualize the sales funnel progression from awareness to purchase with a bar graph so I can identify conversion bottlenecks.

I'll help you visualize the sales funnel progression. First, I'll analyze the data to get the conversion numbers at each stage, then create a visualization.

[View analysis >](#)

Now I'll create a visualization that shows both the number of customers at each stage and the conversion rates between stages.

</>


Sales Funnel Visualization

Click to open component

Sales Funnel Visualization

Preview Code ✕

Sales funnel



Website Visit

Time spent: 2-3 minutes

Users: 10

Conversion rate: 50%

Drop-off Analysis

50% Drop Rate

Recommendations


Last edited just now

Analysis tool

On ☒

Claude can write and run code to process data, run analysis, and produce data visualizations in real time.


Library for the Research Institutes within the ETH Domain: Eawag, Empa, PSI & WSL



94

Gemini – custom features


- Click on settings, then saved info




The screenshot shows the Gemini application interface. On the left, a sidebar menu is visible with options: 'I read this s...', 'Gem mana...', 'Help', 'Activity', and 'Settings' (highlighted with a blue background). A sub-menu is open for 'Settings', showing options: 'Saved info' (selected), 'Apps', 'Your public links', and 'Dark theme'. The main content area is titled 'Info you asked Gemini to save' with an information icon. Below the title, it says: 'Share info about your life and preferences to get more helpful responses. Add new thing during a chat.' The main prompt area asks 'What do you want Gemini to remember?' and contains the text: 'For example, "I prefer short, concise responses"'.

Gemini – custom features

- Click on apps, then enable which apps you want to give access to




Productivity



Grok – custom features

- Click on your account, settings, customize



The screenshot displays the Grok user interface. On the left, a chat window shows a greeting: "Good evening, Corin. How can I help you today?". Below this is a search bar with the placeholder text "What do you want to know?". To the right of the search bar are icons for "DeepSearch" and "Think". At the bottom of the chat window are buttons for "Research", "How to", "Analyze", "Create images", and "Code". A red line points from the "Customize" option in the settings menu to the "Customize" button in the chat window.

Settings

- Account
- Appearance
- Customize**
- Data

Customize Grok's Response

- Custom** (Selected): Responds to you as you please.
- Concise: Responds briefly and directly.
- Formal: Responds using a formal tone.
- Socratic: Responds in way to help you learn.

Custom Instructions

How should Grok behave?

Grok – custom features

- Click on your account, settings, data

Settings

Account

Appearance

Customize

Data

Improve the Model

By allowing your data to be used for training our models, you help enhance your own experience and improve the quality of the model for all users. We take measures to ensure your privacy is protected throughout the process.

☒

Export Account Data

You can download all data associated with your account below. This data includes everything stored in all xAI products.

Export

Delete All Conversations

Delete all of your conversation data.

Delete

Delete Account

Permanently delete your account and associated data from the xAI platform. Deletions are immediate and cannot be undone.

Delete


Deepseek and data privacy

When you create an account, input content, contact us directly, or otherwise use the Services, you may provide some or all of the following information:

- **Account information.** We collect information that you provide when you set up an account, such as your date of birth (where applicable), username (where applicable), email address and/or telephone number, and password.
- **User Input.** When you use our Services, we may collect your text input, prompt, uploaded files, feedback, chat history, or other content that you provide to our model and Services (“Prompts” or “Inputs”). We generate responses (“Outputs”) based on your Inputs.

<https://cdn.deepseek.com/policies/en-US/deepseek-privacy-policy.html>

Reasoning models



- uses Chain-of-Thought reasoning to break problems into multiple intermediate steps
- self-corrects and refines responses before finalizing an answer in an iterative manner
- reduces logical errors compared to traditional generative transformer models
- longer response time but higher accuracy

OpenAI models – hallucinations and fairness + bias

Dataset	Metric	o3	o4-mini	o1
SimpleQA	accuracy (higher is better)	0.49	0.20	0.47
	hallucination rate (lower is better)	0.51	0.79	0.44
PersonQA	accuracy (higher is better)	0.59	0.36	0.47
	hallucination rate (lower is better)	0.33	0.48	0.16

Table 7: BBQ evaluation

Metric	o3	o4-mini	o1
Accuracy on Ambiguous Questions	0.94	0.82	0.96
Accuracy on Unambiguous Questions	0.93	0.95	0.93
P(not stereotyping ambiguous question, not unknown)	0.25	0.26	0.05

<https://cdn.openai.com/pdf/2221c875-02dc-4789-800b-e7758f3722c1/o3-and-o4-mini-system-card.pdf>

Julius.ai

- AI data analyst tool for visualizing and analyzing complex data in seconds
- ChatGPT or Claude based
- Capabilities
 - chat with your data
 - create visual representations and animations
 - build classification or forecasting models and run them

Example: Julius.ai

- Iris dataset consists of 3 different types of irises' (Setosa, Versicolour, and Virginica) petal and sepal length and width
- Prompt: using the iris dataset build a model that classifies the data in 3 classes
- Answer: The model has been successfully trained on the Iris dataset, achieving an accuracy of 1.0, indicating perfect classification on the test set. The classification report provides detailed metrics for each class, including precision, recall, and F1-score.

Example: Julius.ai

</> Trained a Random Forest Classifier on the Iris Dataset

Python

Rerun code

Edit code



```
1 # Import necessary libraries
2 import pandas as pd
3 from sklearn.datasets import load_iris
4 from sklearn.model_selection import train_test_split
5 from sklearn.ensemble import RandomForestClassifier
6 from sklearn.metrics import classification_report, accuracy_score
7
8 # Load the iris dataset
9 iris = load_iris()
10 X = iris.data # features
11 y = iris.target # target classes
12
13 # Split the dataset into training and testing sets
14 X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)
15
16 # Create a Random Forest Classifier
17 model = RandomForestClassifier(n_estimators=100, random_state=42)
18
19 # Train the model
20 model.fit(X_train, y_train)
21
22 # Make predictions
23 y_pred = model.predict(X_test)
24
25 # Evaluate the model
26 accuracy = accuracy_score(y_test, y_pred)
27 report = classification_report(y_test, y_pred)
28
29 # Print the results
30 print('Accuracy:', accuracy)
31 print('Classification Report:\n', report)
32
```

Code Output

Hide 4 outputs

Accuracy:

1.0

Classification Report:

	precision	recall	f1-score	support
0	1.00	1.00	1.00	10
1	1.00	1.00	1.00	9
2	1.00	1.00	1.00	11
accuracy			1.00	30
macro avg	1.00	1.00	1.00	30
weighted avg	1.00	1.00	1.00	30

Show less

The model has been successfully trained on the Iris dataset, achieving an accuracy of 1.0, indicating perfect classification on the test set. The classification report provides detailed metrics for each class, including precision, recall, and F1-score.

Code Explanation

This code snippet loads the Iris dataset, splits it into training and testing sets, trains a Random Forest Classifier, and evaluates its performance using accuracy and a classification report.

Code Output

Some parameters of LLMs

- Temperature:
 - Controls the randomness of the model's output
 - Lower values (e.g., 0.2) make the output more focused and deterministic (more predictable answers)
 - Higher values (e.g., 1.0 or above) make the model more creative and diverse in its responses
- Top-k Sampling:
 - This parameter limits the number of possible next tokens to the top k tokens with the highest probabilities
 - A lower k reduces randomness by restricting choices to the most likely tokens
- Top-p (Nucleus) Sampling:
 - Instead of limiting by a fixed number (like top-k), top-p sampling restricts the model to choosing from the most probable tokens whose cumulative probability is greater than p
 - $p = 0.7$ means the model will sample from the smallest group of tokens whose total probability is at least **70%**, leading to more controlled output
- Presence Penalty:
 - Adjusts the model's likelihood of discussing new topics. A higher presence penalty encourages the model to introduce new concepts into the conversation instead of repeating itself
- Frequency Penalty:
 - Controls how much the model penalizes repeated tokens. A higher frequency penalty discourages the model from generating repeated words or phrases
- Max Tokens (Length of Output):
 - This sets the maximum length of the generated response. It can limit the model's output to avoid excessively long answers

ChatGPT parameter tuning example

- Only possible via API, Playground (requires subscription)
- can directly control parameters such as
 - temperature,
 - top-k,
 - top-p
 - max_tokens
- Example json API Request:
- ```
{ "model": "gpt-4", "prompt": "Tell me a creative story about a robot",
 "max_tokens": 100, "temperature": 0.8, "top_k": 50, "n": 1}
```

# ChatGPT parameter tuning example

- Or via third parties (e.g. huggingface spaces)

Additional Inputs ▼

System message

You are a helpful assistant. You are an OpenAI GPT model. Please adhere to OpenAI's usage policies and guidelines. Ensure your responses are accurate, respectful, and within the scope of OpenAI's rules.

Model

gpt-4o-mini ▼

Max new tokens

4096

Temperature

1.7

Top-p (nucleus sampling)

0.95

Seed

0

☒ Randomize Seed

# Notebook LM

- designed to help users summarize, understand, and synthesize information from multiple documents
- bases its answers on the documents you provide
- developed by Google, uses Gemini LLM

## Key Features:

- Document Summarization
- Chat with Sources
- FAQ Suggestions
- Source Grounding

# Notebook LM - Privacy

<https://support.google.com/notebooklm/> - Accessed on 22.10.2024

**We value your privacy and never use your personal data to train NotebookLM.**

- **If you are logging in with your consumer google account and choose to provide feedback**, human reviewers may review your queries, uploads, and the model's responses to troubleshoot, address abuse or make improvements. Keep in mind that it's best to avoid submitting any information you wouldn't feel comfortable sharing.
- **As a Google Workspace or Google Workspace for Education user**, your uploads, queries and the model's responses in NotebookLM will not be reviewed by human reviewers, and will not be used to train AI models.

# Notebook LM

Prompt:

What is the discrete wavelet transform?

Gemini 1.5

Key concepts:

- **Wavelets:** Wavelets are mathematical functions that are localized in both time and frequency. They are used as basis functions to represent the signal.
- **Decomposition:** The DWT decomposes the signal into a set of wavelet coefficients that represent the signal's information content at different scales (frequencies) and locations (time).
- **Multiresolution analysis:** The DWT provides a multiresolution analysis of the signal, allowing for the examination of the signal's behavior at different levels of detail.

Applications:

Notebook LM

## The Discrete Wavelet Transform (DWT)

The DWT is a mathematical transformation that analyzes signals in both **time and frequency domains**, offering a **multi-scale representation** of the signal <sup>1</sup> <sup>2</sup>. It decomposes a signal into a set of **wavelet coefficients**, which represent the signal's behavior at different scales and locations <sup>2</sup>.

The DWT uses **basis functions** called **wavelets** <sup>1</sup> <sup>3</sup>. Unlike the Fourier transform, which uses sinusoidal waves as basis functions, wavelets are **localized in both time and frequency** <sup>1</sup> <sup>4</sup>. This localization allows the DWT to capture transient features and abrupt changes in the signal more effectively <sup>2</sup> <sup>4</sup>.

Key Concepts in DWT:

- **Mother Wavelet:** A prototype wavelet function that serves as the building block for generating other wavelets through **scaling (dilation) and translation** <sup>3</sup> <sup>5</sup>.
- **Scaling:** Stretching or compressing the mother wavelet to analyze the signal at different resolutions <sup>2</sup> <sup>6</sup>. Larger scales correspond to lower frequencies and coarser details, while smaller scales correspond to higher frequencies and finer details.

# Exercise

- Go to <https://notebooklm.google.com/>
  - Upload one of your (already published) papers
- Or
- A document that you want to analyze
  - perform 2 queries for what information you want to find out
    - E.g. What method was used to solve the problem?
    - What was the accuracy?
    - What were the limitations?
  - check the accuracy of the answers
  - check the accuracy of the source (page, paragraph)

# Exercise

- Use ChatGPT for the same task as before
- Upload one of your (already published) papers

Or

- A document that you want to analyze
- perform 2 queries for what information you want to find out
  - E.g. What method was used to solve the problem?
  - What was the accuracy?
  - What were the limitations?
- check the accuracy of the answers

# ChatGPT and other Generative AI models

## ○ Pros:

- Reformulations (e.g. for clarification)
- Review / ideas to improve certain section
- Helping with coding
  - Code translations
  - Basic implementations
  - Quick data handling
- Help with literature review

## ○ Cons

- Output might not be correct
- Privacy/copyright concerns

# Conclusion


# Conclusion

- Modern AI tools can be powerful and productivity increasing
- Chose the right tool for the right task
- Keep up to date – new tools are emerging fast
  - the best tool today might not be the best tool tomorrow
- Take data privacy in account before uploading your data to an AI model
- Critically evaluate the answer of the AI model
- You are responsible for using the output generated by the AI models

# Thank you for your attention

**Please direct any (future) questions to**

Dr. Corin Otesteanu



Copyright protected material.