



The KITE Research Institute

KITE

AI Tool Guideline and Overview

This material has been prepared solely for use at University Health Network (UHN). UHN accepts no responsibility for the use of this material by any person or organization not associated with UHN. No part of this document may be reproduced in any form for publication without permission of UHN. A printed copy of this document may not reflect the current electronic version.

Contents

- Scope 3
- Types of Artificial Intelligence 3
 - 1. Large Language Models (LLMs) 3
 - 2. Machine Learning (ML) 4
 - 3. Neural Networks 4
 - 4. Natural Language Processing (NLP)..... 4
 - 5. Computer Vision 5
 - 6. Image Diffusion Models..... 5
- Making Good AI Prompts..... 6
- Use of LLMs for KITE Tasks 7
 - Documentation..... 7
 - Example 7
 - Limitations 8
 - Requests for Summaries..... 9
 - Example 9
 - Programming/Coding 9
 - Graphics 10
 - Asking Questions 10
- Use of Microsoft Copilot at UHN 11
 - Copilot Protected Mode 11
 - Reminders When Using an AI LLM tool 11
 - How To Access Protected Copilot in Edge..... 11
 - Copilot for Microsoft 365 12
- LLM Usage Costs 12
 - Microsoft Copilot 12
 - Other LLM costs..... 12
 - Understanding Tokens 12
 - OpenAI ChaGPT 12
 - Google Gemini 13
 - Meta’s LLaMA 13
- LLM Usage Risks 13
 - LLM Usage Risks According to Microsoft Copilot..... 14

Copyright Considerations 15

AI Project Considerations 15

LLM Generated Content Attribution 16

AI Hub at UHN..... 16

Using AI to Develop and Review Research Grant Proposals..... 16

Appendix A 1

Appendix B 1

Scope

This document provides an overview of artificial intelligence (AI) tools, models and techniques for KITE staff and researchers. This document provides guidelines and tips for KITE users on how to implement and utilize these tools for common work tasks. This document will primarily focus on business and general use of one aspect of AI, Large Language Model (LLM) related tools, such as Chat GPT and Microsoft Copilot. This document will not specifically cover research study related uses of LLMs, due to the wide diversity of these applications.

Note to **never** put private or confidential information, including patient or participant data (PHI - personal health information), employee information, research and education confidential data, corporate confidential information, or intellectual property in **any** AI tool.

Types of Artificial Intelligence

1. Large Language Models (LLMs)

LLMs are a type of AI model designed to understand and generate human-like text. They are trained on vast amounts of data and use deep learning techniques to perform a wide range of language-related tasks. In general, the following are the strengths and weaknesses of publicly available LLMs:

- **Strengths:** LLMs are incredibly powerful tools that can generate human-quality text, translate languages, produce creative content, and answer questions in an informative way.
- **Weaknesses:** LLMs can sometimes produce incorrect or misleading information. They may also struggle with tasks that require reasoning or understanding of real-world context. Some cases of incorrect output are known as AI hallucinations, which include fabrication of details or generation of false information due to processing errors or misapplications of learned patterns.

LLM Examples:

- **Microsoft Copilot:** Microsoft Copilot is a powerful AI assistant designed to enhance productivity and streamline tasks across various Microsoft applications. It provides assistance for document creation, coding, and image creation based on natural language requests. Microsoft Copilot is the preferred LLM for use at KITE/UHN and it is integrated into the Microsoft Edge internet browser and into the Bing search engine. Copilot is also integrated into Microsoft VS Code via GitHub.
- **OpenAI GPT-4:** Developed by OpenAI, it can generate coherent and contextually relevant text, answer questions, and assist in creative writing.
- **Google Gemini:** Gemini is known for its multimodal capabilities, meaning it can process and generate various types of data like text, code, and images. This makes it versatile for a wide range of tasks.

- **Meta's LLaMA:** This LLM is designed for various natural language processing tasks. LLaMA's main advantage is that the pre-trained model priority weighting parameters can be downloaded and run locally if privacy mandated.

2. Machine Learning (ML)

ML involves training computer algorithms on data to make predictions or decisions without being explicitly programmed. It encompasses various techniques, including supervised, unsupervised, and reinforcement learning.

ML Examples:

- **Supervised Learning:** Algorithms learn from labeled data. For instance, spam detection in emails where the model is trained on labeled spam and non-spam emails. Supervised learning is also used in research studies and in healthcare. Radiology examinations are a good example of where ML has been used to enable systems to automatically identify some cancers and other anomalies in images to assist radiologists in identifying areas requiring further scrutiny.
- **Unsupervised Learning:** Algorithms find patterns in unlabeled data. An example is customer segmentation in marketing, where the model groups customers based on purchasing behaviour.
- **Reinforcement Learning:** Algorithms learn by interacting with an environment and receiving rewards or penalties. [A notable example is AlphaGo, which learned to play and master the game of Go³.](#)

3. Neural Networks

These are a subset of ML models inspired by the human brain's structure. They are particularly effective for tasks like image and speech recognition.

Neural Network Examples:

- **Convolutional Neural Networks:** Used for image recognition tasks, such as identifying objects in photos.

4. Natural Language Processing (NLP)

NLP focuses on the interaction between computers and human language. It involves tasks like text analysis, translation, and sentiment analysis.

NLP Examples:

- **Google Translate:** Uses deep learning to translate text between languages.
- **Chatbots:** Virtual assistants like Siri and Alexa use NLP to understand and respond to user queries.

5. Computer Vision

This field enables computers to interpret and make decisions based on visual data.

Examples:

- **Facial Recognition:** Used in security systems to identify individuals.
- **Autonomous Vehicles:** Self-driving cars use computer vision to navigate and detect obstacles

6. Image Diffusion Models

Image Diffusion models can be used to generate images from text inputs.

Examples:

- OpenAI's DALL-E: The Image Diffusion model used by Copilot when asked to generate images.
- Stability AI's Stable Diffusion: An Open-Weight image diffusion model that can be run locally to generate images on any suitable computer.

Making Good AI Prompts

A prompt is the input a user provides to the AI model to elicit a specific response. This can take various forms, ranging from simple questions or keywords to complex instructions, code snippets, or even creative writing samples. The effectiveness of a prompt directly influences the quality and relevance of the AI's output.

Prompt Engineering focuses on crafting effective prompts that unlock the capabilities of LLMs, enabling them to understand intent, follow instructions, and generate desired outputs. Much of the following content regarding prompt engineering is based on content from Google (<https://cloud.google.com/discover/what-is-prompt-engineering>) and the full content is included as Appendix B in this document.

By using carefully worded prompts, a user provides the AI with context, instructions, and examples that help it understand the human's intent and respond in a meaningful way. Prompts do not need to include pleasantries like please or thank you or other language that is used to make statements polite for humans. The polite language may not reduce the effectiveness of the prompt but it does waste time and reduces the number of characters available for the key elements of the prompt.

The format of a prompt plays a significant role in how the AI interprets the user's request. Different models may respond better to specific formats, such as [natural language questions](#), direct commands, or structured inputs with specific fields.

Different LLMs have different limits in terms of the maximum number of characters for a single prompt. Generally, it's recommended to keep prompts concise and clear to ensure the model can process them effectively. A typical suggested limit is ~500 characters.

According to Microsoft Copilot based on a prompt asking what its limits are:

The character limit for Microsoft Copilot prompts can vary depending on the specific mode and application. Generally, the limits are as follows:

- *More Balanced mode: 2,000 characters.*
- *More Creative and More Precise modes: 4,000 characters*
- *Edge Copilot: Up to 18,000 characters*

These limits help ensure that Copilot can process and respond to prompts effectively.

Providing context and relevant examples within your prompt helps the AI understand the desired task and generate more accurate and relevant outputs.

Please note that achieving the desired result may require multiple attempts. In some instances, the initial outcome might serve as a good starting point, but a user may need to refine it further to achieve the desired result.

Use of LLMs for KITE Tasks

This section describes tasks where LLMs could be a useful addition in aspects of tasks performed by KITE users. While a few of the most common uses of LLMs are given below, there are many other potential uses.

Documentation

Microsoft Copilot can be used in many ways for documentation in terms of both content generation and improvement. One practical use is to reword content and adjust for the intended tone.

Example

A supervisor has the following email to send to a software designer to update their widget:

Complete the following by the end of the week:

Widget module design

Widget model test plan

Widget model coding

The supervisor prompted the Copilot LLM with “*Reword into an email suitable for sending a software designer:*” with the above text. Copilot took this prompt and generated:

Subject: Tasks to Complete by End of Week

Hi [Software Designer’s Name],

Could you please complete the following tasks by the end of this week:

- 1. Design the widget module.*
- 2. Develop the test plan for the widget model.*
- 3. Code the widget model.*

Thank you for your attention to these tasks.

*Best regards,
[Your Name]*

The supervisor could then prompt Copilot for possible additional alterations like “*What if I need to emphasize urgency*”, which lead to the following result:

Subject: Urgent: Tasks to Complete by End of Week

Hi [Software Designer’s Name],

I need your immediate attention on the following tasks, which must be completed by the end of this week:

*Design the widget module.
Develop the test plan for the widget model.
Code the widget model.*

It's crucial that these tasks are finished on time. Please prioritize accordingly.

Thank you for your prompt action.

*Best regards,
[Your Name]*

In addition, a prompt for the same text of “What if I want this to be very nice since it is a new employee.” results in:

Subject: Welcome! Tasks to Complete by End of Week

Hi [Software Designer's Name],

Welcome to the team! I hope you're settling in well. As you get started, I have a few tasks that need to be completed by the end of this week:

- 1. Design the widget module.*
- 2. Develop the test plan for the widget model.*
- 3. Code the widget model.*

I appreciate your efforts and am here to support you as you work through these. Please don't hesitate to reach out if you have any questions or need assistance.

Looking forward to seeing your great work!

*Best regards,
[Your Name]*

Limitations

The above are very trivial examples of what Copilot can do with prompts to improve minimal text to make it a complete message. One aspect to consider is that while it is making reasonable adjustments, they are not presented in a person's own style. This can potentially lead someone to think a person is just using a tool to generate the email (which is the case in this example). For some situations, such as providing engineers with instructions or communicating with someone regarding a business request, it is likely not an issue, but for messages that would typically be more personalized, it may make a person's message sound too generic and computer generated.

Requests for Summaries

Microsoft Copilot is fantastic at summarizing articles and other bodies of text. A user can quite literally paste a link to an article (if the article is not behind a paywall) and ask for an article summary, and Copilot will generate one. In most cases, the summary accurately reflects the original content, but outputs should still be validated. Copilot can sometimes omit pertinent details, misunderstand bits of the articles, introduce bias, or add incorrect information to the summary.

Example

The following is an example of a request for summarizing an article by using the following text prompt:

Summarize this article: <https://www.cbc.ca/news/canada/marketplace-boot-test-1.6252613> using bullet points for the summary.

With the above prompt, Copilot responded:

Sure! Here's a summary of the article:

- **Investigation:** *CBC Marketplace tested six popular winter boot brands in Canada for slip resistance on wet, icy surfaces.*
- **Brands Tested:** *Merrell, Sorel, Kamik, Ugg, Timberland, and WindRiver.*
- **Testing Method:** *Conducted at the KITE Research Institute in Toronto, using a hydraulic system to adjust the floor angle up to 15 degrees.*
- **Best Performer:** *WindRiver Backwoods Waterproof Hyper Dri 3 hiking boots had the best grip, thanks to fibre-embedded soles.*
- **Results:** *Only WindRiver boots passed the slip test, while the other five brands failed.*
- **Safety Concerns:** *About 90% of winter boots available in Canada do not pass slip-resistance tests.*
- **Importance:** *Proper slip-resistant boots can help reduce the risk of falls on ice, which is a significant cause of emergency department visits in Canada*

If you have any other questions or need further details, feel free to ask!

Note that in this example, Copilot was asked specifically to summarize in bullet point form, but other summarization options can also be used.

Programming/Coding

Microsoft Copilot is an excellent coding resource. It will translate natural language prompts into self-contained programs in different languages. It has provided excellent results when requesting Python code, Visual Basic code and SQL Scripts for internal projects at KITE. It has also been used for translating from one coding language to another (e.g. from C to Python) and it is often used for creating Unit Test and System Test frameworks for code. Copilot can create code and scripts based on many different programming languages, including Python, C, C#, C++, HTML, SQL, R and MATLAB. The easiest way to determine which languages it can develop for is to ask Copilot.

Copilot can also provide some help regarding creating content for higher-level frameworks, such as Unity or the Unreal Engine, but the user would need to have a basic working knowledge of the tools in order for the instructions to be useful. It can also be helpful for instructions regarding the use of many other software-based tools.

Coding prompts should include examples of inputs and expected outputs if possible. Code can be generated for utilities, such as simple programs to do repetitive tasks or tasks that would be tedious to do manually on a number of image files.

Code bugs can be provided to Copilot to come up with fixes before attempting to do a traditional line-by-line debug. Note that Copilot and other AI tools may indicate that they have changed the code and solved the issue despite not changing the code at all. These tools are not perfect despite how powerful they can be, and their outputs always need to be checked.

Graphics

Copilot feeds text inputs into an image diffusion model when asked to generate an image. This can be a great resource for generating graphics, but at the risk of copyright infringement, as created pictures can be generated by LLMs from existing graphical images or portions of images without attribution. It is recommended that no graphics developed using diffusion models are provided to external clients or used on any of our websites because of the potential risks to KITE's reputation and the possibility of copyright implications. Diffusion models can be used for internal purposes, such as brainstorming and coming up with options quickly.

Note that since typical usage of an image diffusion model is probabilistic, the same prompt will usually not generate the same image. The AI will also not always follow all prompt instructions, or at least not always follow them as the user anticipates. Prompts can include aspects like lighting, shadows, colour choices, etc.

Showing AI generated images to external clients should be avoided since they may like the concept exactly as presented, and the image could be problematic due to rights issues. A few examples of AI generated imagery are provided in Appendix A.

Please note that AI generated graphics that are sold as stock photos by Adobe etc., are fine for use in our websites and for external clients if they are licensed appropriately.

Asking Questions

Asking Copilot a question can generate an answer or a number of potential options for the request. In this usage, it is effectively a "smart" search engine. The drawback to this is that the responses can be outdated (if that is relevant to the request), or biased based on what data it was trained on, or incorrect. It can often save time compared to sifting through numerous Google search results. However, the shorter list of options might not always include the most suitable ones. Since responses are probabilistic, asking again may trigger a different response.

Use of Microsoft Copilot at UHN

The tool endorsed and supported by UHN Digital is Microsoft Copilot in Edge. Edge is the web browser from Microsoft and the default web browser at UHN.

Copilot Protected Mode

The advantage of using Copilot in Edge is that it can be used in a “Protected” mode, meaning that the data (i.e. whatever a user types into Copilot in Edge) is kept within UHN’s custody and protection. The user’s chats and data are not stored or shared with anyone, and previous chats are not used to improve Copilot in Edge responses.

Copilot is approved by UHN **only** when accessed through the Edge browser or by using the Bing search app on a mobile device in protected mode. Both methods require the UHN user to **sign in with their UHN credentials** to ensure that they are in the “Protected” mode. If the user is not logged in with their UHN credentials, any searches they do are NOT properly secured.

When in protected mode, even though the tool is protected for data safety, please note that it is forbidden for UHN personnel to input any private or confidential information, including patient or participant data (PHI - Personal Health Information), employee information, research and education confidential data, corporate confidential information, or intellectual property in their queries or prompts. If a researcher is in doubt regarding whether certain data can be entered, please check with the project supervisor.

Reminders When Using an AI LLM tool

Remember when using AI:

- Always verify and validate the accuracy of returned results.
- Do not use any AI assistant for any type of clinical care related work, including note taking to put into a chart, generating summaries of patient information, or other. UHN has not validated the use of these tools for medical applications and teams are working on controlled pilots for such use cases.
- Never put Private or Confidential information, including patient or participant data (PHI - Personal Health Information) employee information, research and education confidential data, corporate confidential information, or intellectual property in any AI tool.

How To Access Protected Copilot in Edge

There are currently three main ways to access Copilot in Edge and ensure that entered data is protected:

1. Logging in using [Copilot.Microsoft.com](https://copilot.microsoft.com) with your UHN credentials
2. Logging in using [Edge browser](#) with your UHN credentials
3. Using the mobile [Bing search engine](#) app on a phone or tablet signed in with your UHN credentials.

Copilot for Microsoft 365

Copilot in Edge is different from Copilot for Microsoft 365. Copilot for Microsoft 365 is currently not available or supported by UHN and would require additional licensing. This tool specifically integrates into Office 365 applications and a UHN Digital team is actively assessing these applications.

LLM Usage Costs

Microsoft Copilot

Microsoft Copilot integrated into Edge and generally preferred by UHN. It's free to use with a UHN account, and separate chats can be created as needed. Chats should be private and secure. Model output quality may differ depending on the time of day as Microsoft dynamically changes the model for load balancing purposes. The pro version of Copilot can be purchased for \$20/month which gives priority access to their most powerful model. Typically Copilot should be more than adequate for typical KITE tasks. Research study specific tasks may require a different LLM in some cases, as determined by the PI.

Other LLM costs

Understanding Tokens

If work requires specialized models or applications, researchers may need to use other LLMs. For this type of work, a software developer or programmer would typically use the LLM. Instead of using the prompts described above to access the LLM, a software developer would use the LLM's application programming interface (API) to send in text inputs or "tokens". Many APIs incur fees based on tokens.

When accessing an LLM through an API, costs are usually calculated on a per-token basis instead of a fixed monthly cost, therefore understanding what tokens are is useful when determining how much using a particular LLM would cost.

A token represents a sequence of characters, which can be as short as one character or as long as one word. For example:

- Common words like "air" might be a single token.
- Longer or less common words like "airline" might be split into multiple tokens ("air" and "line").

On average, 100 tokens are roughly equivalent to 75 words in English. Tokens are used to measure both the input (prompts) and output (completions) of the model.

OpenAI ChaGPT

For GPT, the following costs are applied based on usage:

- Free – gives access to their weakest model and limited access to one of their stronger models (GPT-4o), file uploads, data analysis & web browsing
- Plus (\$20/month) - extends the limits on data analysis, file uploads & web browsing, limited access to even more powerful models (GPT-4o1 & o1-mini)

- Pro (\$200/month) - Unlimited access to advanced models, access to o1 pro mode which gives better, more accurate answers to hard questions.
- For app development/API access costs vary on a per-model basis, but it's generally \$2.5/1M input tokens. The exact prices can be found on <https://openai.com/api/pricing/>

Google Gemini

This large language model family is trained by Google. It is generally intelligent, but can struggle on some questions that OpenAI's GPT-4 would otherwise succeed in. For general usage, it should be fine to use.

- Gemini – Free to use, but limited in model output quality and context window (how many words/tokens the AI can analyze at a time)
- Gemini Advanced (\$27/month) - Access to the most powerful Gemini model, extended context window (up to 1 million tokens), and able to use Gemini in Google docs, Gmail, etc.

For app development/API access – Like OpenAI, the costs vary on a per-model basis, but any of Google's models can be accessed for free with heavy limits on requests per minute and per day for the higher-end models. The lower end models such as Gemini 1.5 Flash can be accessed with generous rate limits.

The exact cost breakdown can be found here: <https://ai.google.dev/pricing>

Meta's LLaMA

LLaMA is a special case in that it can technically be run for free with the necessary hardware, but the hardware involved in running the more powerful models involves either building a computer with multiple high-end graphics cards/AI accelerators, or renting a server that contains such hardware.

The less-powerful models can be run on modest devices (e.g. standard desktop with 16GB of RAM), but the output quality and speed might suffer accordingly.

LLM Usage Risks

There are a number of risks that may result from using LLMs. The following is a subset of the potential risks:

- Personal data or other people's PHI could be uploaded by accident.
- Copyrighted data or works from the internet could be integrated into what the LLM provides, which could result in accusations of plagiarism or theft.
- Documentation generated could be inaccurate or misleading. Any documentation generated by an LLM must be carefully reviewed and checked for validity.
- Trusting that the code generated works and meets the intent without understanding and testing any code generated by the AI.
- Using AI generated code in a product can make a design disjointed and harder to maintain if it is not carefully integrated.

LLM Usage Risks According to Microsoft Copilot

The following is a list of risks generated by Microsoft Copilot based on the prompt “What are risks in using content from an LLM like Microsoft Copilot”:

Using content generated by a large language model (LLM) like Microsoft Copilot can be incredibly beneficial, but there are some risks to be aware of:

1. Accuracy and Reliability

- **Misinformation:** LLMs can sometimes generate incorrect or misleading information. It's important to verify facts from reliable sources.
- **Outdated Information:** The model's knowledge is based on the data it was trained on, which may not include the most recent updates or developments.

2. Bias and Fairness

- **Bias in Data:** LLMs can reflect biases present in the training data, leading to biased or unfair outputs. This can affect the quality and neutrality of the content.
- **Stereotyping:** There's a risk of reinforcing stereotypes if the model generates content based on biased data.

3. Privacy and Security

- **Sensitive Information:** Using LLMs to process sensitive or personal information can pose privacy risks. It's crucial to ensure that any data shared with the model is handled securely.
- **Data Leakage:** There's a potential risk of the model inadvertently generating content that includes sensitive or proprietary information.

4. Copyright and Intellectual Property

- **Copyright Infringement:** LLMs can generate content that closely resembles copyrighted material, leading to potential copyright issues. It's important to ensure that generated content does not violate copyright laws.
- **Originality:** While LLMs can assist in content creation, relying too heavily on them might result in less original and creative outputs.

5. Ethical Considerations

- **Misuse:** There's a risk of LLMs being used to generate harmful or malicious content, such as deepfakes, fake news, or spam.
- **Transparency:** Users should be aware that the content is generated by an AI and not a human, to maintain transparency and trust.

6. Dependence and Skill Degradation

- **Over-Reliance:** Relying too much on LLMs for tasks like writing or coding can lead to a degradation of human skills in these areas.
- **Critical Thinking:** Users might become less inclined to critically evaluate information if they overly depend on AI-generated content.

To mitigate these risks, it's important to use LLMs like Microsoft Copilot as a tool to augment human capabilities rather than replace them. Always review and verify the content generated, be mindful of privacy and copyright issues, and maintain a critical perspective.

Copyright Considerations

The University of Toronto has an article regarding Generative AI Tools and Copyright Considerations here:

[Generative AI tools and Copyright Considerations | University of Toronto Libraries](#)

Key points from the article include:

The legal status of AI tools is unsettled in Canada at this point in time. This is an evolving area and our understanding will develop as new policies, regulations, and case law becomes settled.

- **Input:** The legality of the content used to train AI models is unknown in some cases. There are a number of lawsuits originating from the US that allege AI tools infringe copyright and it remains unclear if and how the fair use doctrine can be applied. Canada remains in a similar uncertain state; it is unclear the extent to which existing exceptions in the copyright framework, such as fair dealing, apply to this activity.
- **Output:** Authorship and ownership of works created by AI is unclear. Traditionally, Canadian law has indicated that an author must be a natural person (human) who exercises skill and judgement in the creation of a work. As there are likely to be varying degrees of human input in content generated, it is unclear in Canada how it will be determined who the appropriate author and owner of works are.

AI Project Considerations

One important consideration of any AI related project is whether AI is the most suited approach for the development of a tool or process. Tools or non-AI approaches to the problem may already exist. If using AI is determined to be the most suitable, or necessary approach, further research should be done into ensuring user privacy as all data sent to non-locally hosted models (such as GPT-4) could potentially be logged (even if claimed otherwise). Validation is needed to ensure that model outputs are accurate and structured as required.

LLM Generated Content Attribution

The Artificial Intelligence Disclosure (AID) Framework is a method for disclosing the use of AI in research. Please see the links below for further information.

[The Artificial Intelligence Disclosure \(AID\) Framework: An Introduction | Weaver | College & Research Libraries News](#)

[\[2408.01904\] The Artificial Intelligence Disclosure \(AID\) Framework: An Introduction](#)

AI Hub at UHN

UHN has an AI hub that is primarily focused on supporting and promoting AI Research and AI in Healthcare. Please visit their website to find more information:

[UHN AI Hub | Collaborative Centre.](#)

Using AI to Develop and Review Research Grant Proposals

The Government of Canada website has a guideline page that gives guidance on using AI to review and develop grant proposals:

[Guidance on the use of Artificial Intelligence in the development and review of research grant proposals](#)

It should be noted that this guideline only applies to generative AI tools (like LLMs).

Key points from the web page include:

- Named grant applicants are still accountable for the entirety of their application, even if some of it was created with the assistance of an AI tool
- Priority should be given to privacy, confidentiality, data security and IP protection when developing and reviewing grant applications
- Applicants must make sure that the info included in the grant application is “true, accurate and complete”
- Applicants must acknowledge and reference sources of information
 - This can be difficult when using generative AI as it can produce information without proper acknowledgement/recognition
- Applicants need to state if and how they used generative AI in the creation of their application
 - Specific instructions (that must be adhered to) will be provided regarding the use of generative AI that may differ for each funding opportunity
- Using publicly available generative AI tools to *evaluate* grant applications is strictly prohibited as the data sent to tool may be logged and used to improve it.

This material has been prepared solely for use at KITE. UHN accepts no responsibility for use of this material by any person or organization not associated with KITE. No part of this document may be reproduced in any form for publication without permission of KITE. A printed copy of this document may not reflect the current, electronic version.

Editors (version 2024-12-17): Andrew Turner, Colin Stoddart, Tara Lai, Susan Gorski, Jarrett Churchill

Appendix A

AI Generated Image Examples

This appendix shows prompts provided to Microsoft Copilot and the resultant images generated 15Nov2024. The same prompt was submitted multiple times. With the probabilistic nature of the Copilot AI, it generates different imagery with the same prompt. Note that it often ignores one or more aspects of the prompt requests when generating images.

Images generated by AI need to be reviewed closely, even images that appear reasonable at first glance can include odd artifacts, distortions etc. For instance, extra fingers and unnatural finger joint angles are common artifacts in AI generated images that include people.

1. Create an image of a human brain being stimulated figuratively.









2. Create an image of a human brain being stimulated figuratively with ideas and activities.





3. Create an image of a human brain being stimulated figuratively with arts, science and education.





4. Create an image of the human mind being stimulated figuratively with ideas.





5. Create an image of the human mind being stimulated figuratively with ideas. Include symbols of creativity and inspiration.





6. Draw an image of the downtown Toronto skyline as seen from Lake Ontario during the day, with the CN tower close to the center of the image and calm water. Add a small biplane flying across the scene from left to right above the skyline with a banner trailing behind the plane that says "Kite" in blue lettering on a white banner with the letters appearing left to right on the banner. Make the image as realistic as possible.









Appendix B

Prompt Engineering Guide from Google

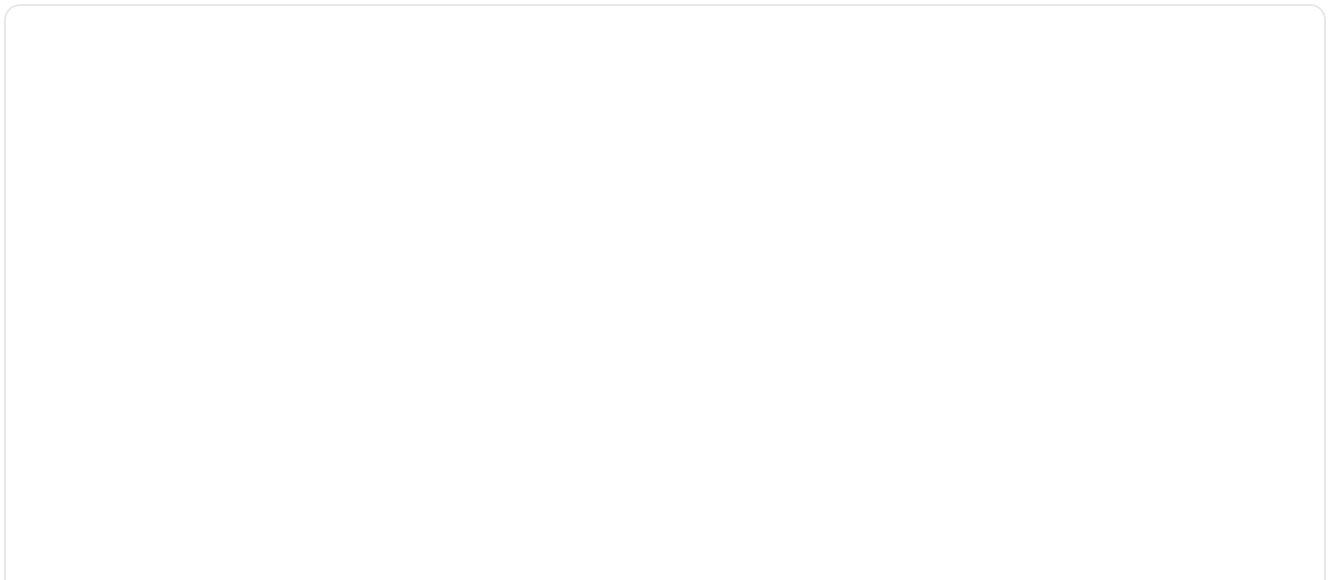


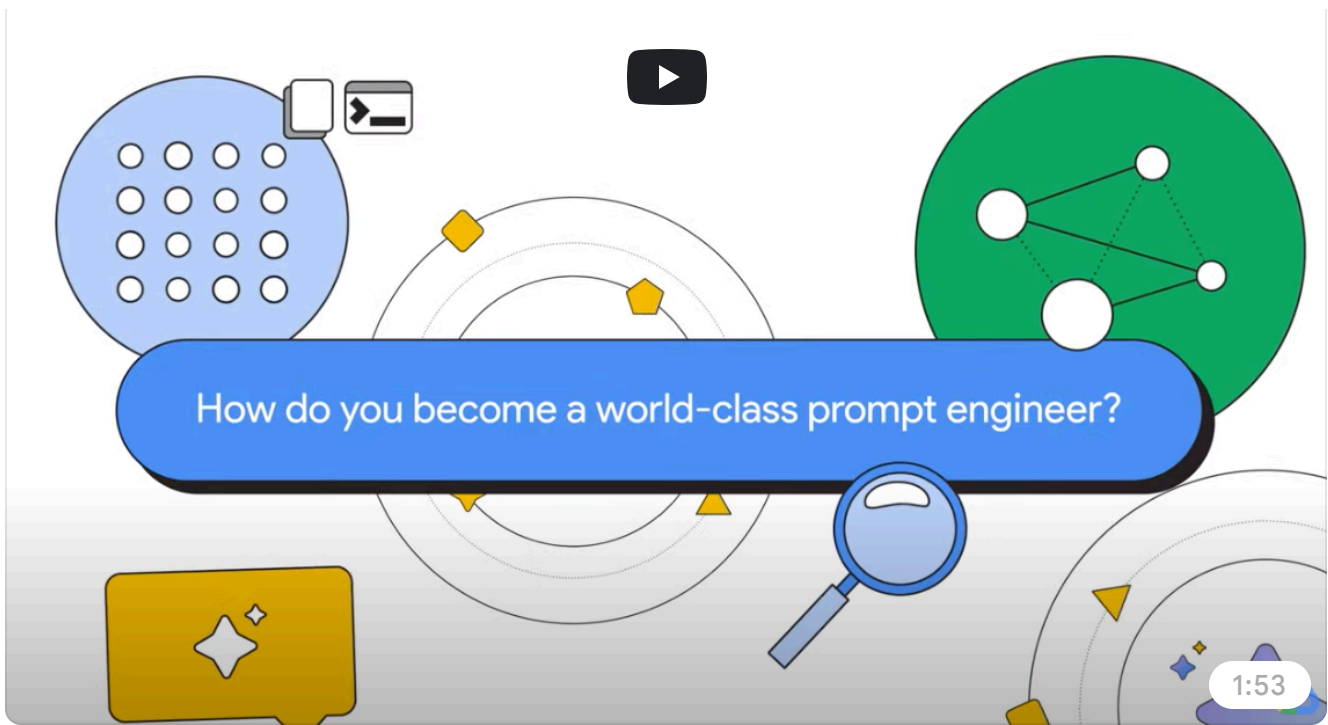
[Topics](#) > Prompt engineering

Prompt engineering: overview and guide

The rise of large language models (LLMs) has brought forth exciting possibilities for human-computer interaction. However, harnessing the full potential of these powerful AI models requires a crucial skill: prompt engineering. This burgeoning field focuses on crafting effective prompts that unlock the [capabilities of LLMs](#), enabling them to understand intent, follow instructions, and generate desired outputs. As we increasingly interact with AI in various applications, prompt engineering plays a vital role in ensuring accurate, relevant, and safe interactions.

Get started for free





Tips to becoming a world-class Prompt Engineer

What is prompt engineering?

Prompt engineering is the art and science of designing and optimizing prompts to guide AI models, particularly LLMs, towards generating the desired responses. By carefully crafting prompts, you provide the model with context, instructions, and examples that help it understand your intent and respond in a meaningful way. Think of it as providing a roadmap for the AI, steering it towards the specific output you have in mind.

To dive deeper into the world of prompt design and explore its applications, check out the [Introduction to Prompt Design](#) on Google Cloud.

Ready to experiment with LLMs and prompt engineering firsthand? Try the [Vertex AI](#) free trial and experience the power of this technology.

What is a prompt for AI?

In the [context of AI](#), a prompt is the input you provide to the model to elicit a specific response. This can take various forms, ranging from simple questions or keywords to complex instructions, code snippets, or even creative writing samples. The effectiveness of your prompt directly influences the quality and relevance of the AI's output.

What do you need for prompt engineering?

Several key elements contribute to effective prompt engineering. Mastering these allows you to communicate effectively with AI models and unlock their full potential.

Prompt format

The structure and style of your prompt play a significant role in guiding the AI's response. Different models may respond better to specific formats, such as:

The format of your prompt plays a significant role in how the AI interprets your request. Different models may respond better to specific formats, such as [natural language questions](#), direct commands, or structured inputs with specific fields. Understanding the model's capabilities and preferred format is essential for crafting effective prompts.

Context and examples

Providing context and relevant examples within your prompt helps the AI understand the desired task and generate more accurate and relevant outputs. For instance, if you're looking for a creative story, including a few sentences describing the desired tone or theme can significantly improve the results.

Fine-tuning and adapting

Fine-tuning the AI model on specific tasks or domains using tailored prompts can enhance its performance. Additionally, adapting prompts based on user feedback or model outputs can further improve the model's responses over time.

Multi-turn conversations

Designing prompts for multi-turn conversations allows users to engage in continuous and context-aware interactions with the AI model, enhancing the overall user experience.

Types of prompts

There are various types of prompts used in AI, each serving a specific purpose:

Direct prompts (Zero-shot)

Zero-shot prompting involves providing the model with a direct instruction or question without any additional context or examples.

An example of this is idea generation, where the model is prompted to generate creative ideas or brainstorming solutions. Another example is summarization, or translation, where the model is asked to summarize or translate some piece of content.

One-, few- and multi-shot prompts

This method involves providing the model with one or more examples of the desired input-output pairs before presenting the actual prompt. This can help the model better understand the task and generate more accurate responses.

Chain of Thought Prompts

CoT prompting encourages the model to break down complex reasoning into a series of intermediate steps, leading to a more comprehensive and well-structured final output.

Zero-shot CoT Prompts

Combines chain of thought prompting with zero-shot prompting by asking the model to perform reasoning steps, which may often produce better output.

Use cases and examples of prompt engineering

Here are some specific examples and use cases showing how prompt engineering helps produce customized and relevant output.

Language and Text Generation

| Scenario | | Example Prompt |
|------------------|---|--|
| Creative Writing | Craft prompts that specify genre, tone, style, and plot points to guide the AI in generating engaging narratives. | "Write a short story about a young woman who discovers a magical portal in her attic." |
| Summarization | Provide the AI with text and instruct it to generate concise | "Summarize the main points of the following news |

| | | |
|-------------|--|--|
| | summaries that capture key information. | article on climate change." |
| Translation | Specify the source and target languages to enable the AI to accurately translate text while preserving meaning and context. | "Translate the following text from English to Spanish: 'The quick brown fox jumps over the lazy dog.'" |
| Dialogue | Design prompts that simulate conversations, allowing the AI to generate responses that mimic human interaction and maintain context. | "You are a friendly chatbot helping users troubleshoot their computer problems. Respond to the user's query: 'My computer won't turn on.'" |

Question Answering

| Scenario | Instructions | Example Prompt |
|----------------------|--|---|
| Open-Ended Questions | Formulate prompts that encourage the AI to provide comprehensive and informative answers | "Explain the concept of quantum computing and its potential impact on the |

| | | |
|---------------------------|--|--|
| | based on its knowledge base. | future of technology." |
| Specific Questions | Design prompts that target specific information, enabling the AI to retrieve precise answers from the provided context or its internal knowledge base. | "What is the capital of France?" or "According to the provided text, what are the main causes of deforestation?" |
| Multiple Choice Questions | Present prompts with options, prompting the AI to analyze and select the most appropriate answer based on its understanding of the context. | "Who wrote the Harry Potter series? A) J.R.R. Tolkien, B) J.K. Rowling, C) Stephen King" |
| Hypothetical Questions | Craft prompts that explore hypothetical situations, allowing the AI to reason, speculate, and provide potential outcomes or solutions. | "What would happen if humans could travel at the speed of light?" |

| | | |
|--------------------------------|--|--|
| <p>Opinion-Based Questions</p> | <p>Design prompts that elicit the AI's perspective or opinion on a specific topic, encouraging it to provide reasoning and justification for its stance.</p> | <p>"Do you believe that artificial intelligence will eventually surpass human intelligence? Why or why not?"</p> |
|--------------------------------|--|--|

Code Generation

| Scenario | Instructions | Example Prompt |
|------------------------|--|--|
| <p>Code Completion</p> | <p>Provide the AI with a partial code snippet and prompt it to suggest or complete the remaining code based on the context and programming language.</p> | <p>"Write a Python function to calculate the factorial of a given number."</p> |
| | | |

| | | |
|-------------------|---|---|
| Code Translation | Specify the source and target programming languages to enable the AI to translate code while preserving functionality and syntax. | "Translate the following Python code to JavaScript: def greet(name): print('Hello,', name)" |
| Code Optimization | Prompt the AI to analyze existing code and suggest improvements for efficiency, readability, or performance. | "Optimize the following Python code to reduce its execution time." |
| Code Debugging | Provide the AI with code containing errors and prompt it to identify and suggest potential solutions for the identified issues. | "Debug the following Java code and explain why it is throwing a NullPointerException." |

Image Generation

| Scenario | Instructions | Example Prompt |
|----------|--------------|----------------|
| | | |

| | | |
|------------------------------|--|---|
| <p>Photorealistic Images</p> | <p>Craft prompts that describe the desired image in detail, including objects, scenery, lighting, and style, to generate realistic and high-quality images.</p> | <p>"A photorealistic image of a sunset over the ocean with palm trees silhouetted against the sky."</p> |
| <p>Artistic Images</p> | <p>Design prompts that specify art styles, techniques, and subject matter to guide the AI in creating images that mimic specific artistic movements or evoke certain emotions.</p> | <p>"An impressionist painting of a bustling city street with people walking under umbrellas in the rain."</p> |
| <p>Abstract Images</p> | <p>Formulate prompts that encourage the AI to generate images that are open to interpretation, utilizing shapes, colors, and textures to evoke feelings or concepts.</p> | <p>"An abstract image representing the concept of hope, using bright colors and flowing shapes."</p> |
| | | |

| | | |
|---------------|---|--|
| Image Editing | Provide the AI with an existing image and specify desired modifications, enabling it to edit and enhance the image according to the given instructions. | "Change the background of this photo to a starry night sky and add a full moon." or "Remove the person from this image and replace them with a cat." |
|---------------|---|--|

Strategies for writing better prompts

Developing effective prompts requires a strategic approach. Consider these strategies to enhance your prompt engineering skills:

1. Set Clear Goals and Objectives:

| Tactic | Prompt Example |
|--|--|
| Use action verbs to specify the desired action | "Write a bulleted list that summarizes the key findings of the attached research paper" |
| Define the desired length and format of the output | "Compose a 500-word essay discussing the impact of climate change on coastal communities." |
| | |

Specify the target audience

"Write a product description for a new line of organic skincare products, targeting young adults concerned with sustainability."

2. Provide Context and Background Information:

| Tactic | Prompt Example |
|---|--|
| Include relevant facts and data | "Given that global temperatures have risen by 1 degree Celsius since the pre-industrial era, discuss the potential consequences for sea level rise." |
| Reference specific sources or documents | "Based on the attached financial report, analyze the company's profitability over the past five years." |
| Define key terms and concepts | "Explain the concept of quantum computing in simple terms, suitable for a non-technical audience." |

3. Use Few-Shot Prompting:

| Tactic | Prompt Example |
|--|--|
| Provide a few examples of desired input-output pairs | Input: "Cat" Output: "A small furry mammal with whiskers." Input: "Dog" Output: "A domesticated canine known for its loyalty." Prompt: "Elephant" |
| Demonstrate the desired style or tone | Example 1 (humorous): "The politician's speech was so dull, it could cure insomnia." Example 2 (formal): "The dignitary delivered an address that was both informative and engaging." Prompt: "Write a sentence describing the comedian's stand-up routine." |
| Show the desired level of detail | Example 1 (brief): "The movie was about a young boy who befriends an alien." Example 2 (detailed): "The science fiction film follows the story of Elliot, a lonely boy who discovers and forms a unique bond with an extraterrestrial stranded on Earth." Prompt: "Summarize the plot of the novel you just finished reading." |

4. Be Specific:

| Tactic | Prompt Example |
|---|---|
| Use precise language and avoid ambiguity | Instead of: "Write something about climate change," use: "Write a persuasive essay arguing for the implementation of stricter carbon emission regulations." |
| Quantify your requests whenever possible | Instead of: "Write a long poem," use: "Write a sonnet with 14 lines that explores themes of love and loss." |
| Break down complex tasks into smaller steps | Instead of: "Create a marketing plan," use: "1. Identify the target audience. 2. Develop key marketing messages. 3. Choose appropriate marketing channels." |

5. Iterate and Experiment:

| Tactic | Action |
|--------------------------------------|---|
| Try different phrasings and keywords | Rephrase your prompt using synonyms or alternative sentence structures. |

| | |
|--|--|
| Adjust the level of detail and specificity | Add or remove information to fine-tune the output. |
| Test different prompt lengths | Experiment with both shorter and longer prompts to find the optimal balance. |

6. Leverage Chain of Thought Prompting:

| Tactic | Prompt Example |
|--|---|
| Encourage step-by-step reasoning | "Solve this problem step-by-step: John has 5 apples, he eats 2. How many apples does he have left? Step 1: John starts with 5 apples. Step 2: He eats 2 apples, so we need to subtract 2 from 5. Step 3: $5 - 2 = 3$. Answer: John has 3 apples left." |
| Ask the model to explain its reasoning process | "Explain your thought process in determining the sentiment of this movie review: 'The acting was superb, but the plot was predictable.'" |
| | |

Guide the model through a logical sequence of thought

"To classify this email as spam or not spam, consider the following: 1. Is the sender known? 2. Does the subject line contain suspicious keywords? 3. Is the email offering something too good to be true?"

For further guidance on prompt engineering best practices, explore the [Five Best Practices for Prompt Engineering](#) on Google Cloud.

Benefits of prompt engineering

Effective prompt engineering offers numerous benefits, enhancing the capabilities and usability of AI models:

Improved model performance

Well-crafted prompts lead to more accurate, relevant, and informative outputs from AI models, as they provide clear instructions and context.

Reduced bias and harmful responses

By carefully controlling the input and guiding the AI's focus, prompt engineering helps mitigate bias and minimize the risk of generating inappropriate or offensive content.

Increased control and predictability

Prompt engineering empowers you to influence the AI's behavior and ensure consistent and predictable responses aligned with your desired outcomes.

Enhanced user experience

Clear and concise prompts make it easier for users to interact effectively with AI models, leading to more intuitive and satisfying experiences.

Start your AI journey with Google Cloud

New customers get \$300 in free credits to spend on Google Cloud.

Get started

Talk to a Google Cloud sales specialist to discuss your unique challenge in more detail.

[Contact us](#)

Related Google Cloud products and services

[See all AI products and solutions](#)



Vertex AI Platform

A single platform for data scientists and engineers to create, train, test, monitor, tune, and deploy ML and AI models.



Generative AI on Vertex AI

Rapidly prototype and test generative AI models. Test sample prompts, design your own prompts, and customize foundation models and LLMs.



AI APIs

Easily integrate AI into your applications with Google Cloud's AI and machine learning APIs.



SOLUTION

Model Garden on Vertex AI

Jumpstart your ML project with a single place to discover, customize, and deploy a wide variety of models from Google and Google partners.



Additional learning resources to get started

New to Google Cloud or generative AI? New customers get [\\$300 in free credits](#) to run, test, and deploy workloads.

- [Training: No cost generative AI fundamentals course](#)
- [Documentation: Introduction to prompt design](#)
- [Documentation: General prompt design strategies](#)
- [Documentation: Generative AI prompt samples](#)

Take the next step

Start building on Google Cloud with \$300 in free credits and 20+ always free products.

Get started for free

Need help getting started?

[Contact sales](#)

Work with a trusted partner

[Find a partner](#)

Continue browsing

[See all products](#)

| Why Google | Products and | Solutions | Resources | Engage |
|-----------------------------|--------------------------|------------------------------|--------------------------------|------------------|
| Choosing Google Cloud | pricing | Infrastructure modernization | Google Cloud Affiliate Program | Contact sales |
| Trust and security | Google Cloud pricing | Databases | Google Cloud documentation | Find a Partner |
| Modern Infrastructure Cloud | Google Workspace pricing | Application modernization | Google Cloud quickstarts | Become a Partner |
| Multicloud | See all products | Smart analytics | | Events |
| | | | | Podcasts |

[Global infrastructure](#)

[Customers and case studies](#)

[Analyst reports](#)

[Whitepapers](#)

[Blog](#)

[Artificial Intelligence](#)

[Security](#)

[Productivity & work transformation](#)

[Industry solutions](#)

[DevOps solutions](#)

[Small business solutions](#)

[See all solutions](#)

[Google Cloud Marketplace](#)

[Learn about cloud computing](#)

[Support](#)

[Code samples](#)

[Cloud Architecture Center](#)

[Training](#)

[Certifications](#)

[Google for Developers](#)

[Google Cloud for Startups](#)

[System status](#)

[Release Notes](#)

[Developer Center](#)

[Press Corner](#)

[Google Cloud on YouTube](#)

[Google Cloud Tech on YouTube](#)

[Follow on X](#)

[Join User Research](#)

[We're hiring. Join Google Cloud!](#)


[Google Cloud Community](#)

[About Google](#) | [Privacy](#) | [Site terms](#) | [Google Cloud terms](#) | [Manage cookies](#)

[Our third decade of climate action: join us](#)

[Sign up for the Google Cloud newsletter](#)

[Subscribe](#)

 [English](#) ▼