

Mastering the Art of Prompting

Generative AI Whitepaper



ChatGPT



* Claude 3



Message ChatGPT



GEN AI STRATEGY

www.generativeaistrategy.com

TABLE OF CONTENTS

Introduction:

The basics of human-AI communication **3**

The Fundamentals of Prompting:

Structure, knowledge, tips & technicals **4**

Prompt Templates:

How to get the most out of prompting for your business **10**

Context Injection and RAG:


How GPT's and RAG systems can help streamline prompting **12**



ChatGPT



Claude



1. Introduction

In the rapidly evolving landscape of artificial intelligence, the ability to effectively communicate with generative AI models (or human-ai communication) has become a crucial skill. Generative AI tools, such as ChatGPT and other fine-tuned models, are revolutionizing how businesses operate, innovate, and compete. However, the key to unlocking their full potential lies in the art of prompting—crafting precise, contextual, and strategic instructions that guide AI to deliver optimal results.

Prompting is not just a technical skill; it is a bridge that connects human capabilities with the computational power of AI. Mastering this skill enables you to leverage AI's capabilities to enhance productivity, foster creativity, and make informed decisions.

Whether you are drafting a report, brainstorming ideas, or optimizing workflows, effective prompting can be the catalyst that transforms how you work and achieve your goals.

In this whitepaper we aim to provide you with detailed insights into the principles and best-practices of prompting, providing you with the knowledge to capture the full potential of generative AI technologies.

Note that prompting is a fundamental skill for working with generative AI models, which can be Large Language Models (LLMs) but also image, video, or audio models. However, in this whitepaper, we will mainly focus on prompting language models. Explore our other whitepapers to learn more about generative AI and prompting techniques.

2. The Fundamentals of Prompting



Overview of Generative AI: Patterns vs Facts

In order to understand the importance of prompting, it is first important to know what generative AI is, and how it works. Generative AI refers to a subset of artificial intelligence capable of 'generating' new content, such as text, images, and audio, by identifying and replicating patterns found in vast data sets. Unlike traditional AI, that typically predicts a number or class ("e.g., a price or whether a consumer is likely to churn or fraud"), generative AI models like ChatGPT and DALL-E create new outputs that have never existed before.

It is important to know three fundamental aspects of generative AI models:

- **Patterns, Not Facts:** Generative AI models learn from patterns in the data they are trained on. They do not search the web or access real-time data to generate responses. Instead, they predict the most likely continuation of the input based on 'learned' patterns.
- **Training Data Influence:** The accuracy and quality of AI outputs are influenced by the diversity and quality of the training data. High-quality and diverse training data sets enable the AI to generate more reliable and relevant outputs.
- **Generative Capabilities:** Models like GPT-4o can create essays, draft emails, create images, and more, based solely on the prompts they are provided with. They understand language nuances and context, which allows them to generate coherent and contextually appropriate content.

Once one is familiar with these aspects of generative AI models, it becomes much clearer why it is important to craft well structured and context rich prompts.



Key Takeaway

Generative AI models **predict responses based on learned patterns.**

The **diversity and quality of training data** significantly influences outputs.

Models like GPT-4 **generate relevant content by 'understanding' context.**



AI: Knowledge and Capabilities

Even AI has its limitations

Effective use of generative AI requires a clear understanding of what these models 'know' (knowledge) and what they 'can do' (capabilities). Just as you wouldn't expect a new employee to know everything about your company on their first day, you shouldn't assume an AI model has complete knowledge or unlimited capabilities.

The fundamentals of AI knowledge

Generative AI models operate fundamentally differently from search engines or databases. Instead of looking up information, they generate responses based on patterns they've learned from their training data. These patterns help them to produce contextually appropriate responses that may appear like facts. However, it's crucial to understand that this is pattern recognition, not fact retrieval, and these patterns could be wrong.

When interacting with a generative AI, you're not accessing a database of facts. You are prompting the model to generate text based on the patterns it has learned in historic data. While the AI can create content on a wide range of topics, it does not 'know' facts the way humans do.

Given these characteristics, providing context is crucial when working with AI. You're not filling gaps in its knowledge, but guiding its prediction process. More context will help the AI produce more relevant and accurate responses tailored to your needs.

Background information: Offer relevant details the AI might not know

Example: Instead of simply asking, 'Write a product description' you might ask, 'Write a product description for our latest smartphone model XYZ123. It features a 6.5-inch OLED display, 5G capability, a triple-lens camera system (12MP wide, 12MP ultra-wide, 12MP telephoto), and is available in three colors: Midnight Black, Ocean Blue, and Sunset Gold.'

Specific instructions: Clearly state what you want the AI to do

Example: Rather than a vague prompt like 'Help me with my marketing': you could say, 'Create a social media content calendar for the next month to promote our new smartphone. Include 3 posts per week, alternating between feature highlights, user testimonials, and promotional offers.'

Desired output format: Specify how you want the information presented

Example: Instead of prompting 'Give me information about our sales: you could prompt 'Provide a summary of our Q2 sales data in the following format Total revenue (with percentage change from Q1), Top 3 performing products (with sales figures), ..'

The Importance of Structured Communication

To make full use of an AI's generative capabilities, we need to provide it with context in the form of prompts. Effective prompting in generative AI requires precision and clarity in your communication. Every letter, word, space, and punctuation mark can influence an AI's output. Think of an AI as a new colleague who has just joined your company. While this colleague has a broad understanding of language and 'the world', they know nothing about your specific company, its culture, or the tone of voice you use.

Just as you would need to train a new employee by providing them with detailed information and clear instructions, you must do the same with AI models. The more structured and context-rich your prompts, the better the AI can understand and produce the desired output.

Basics of Creating Effective Prompts

Creating effective prompts is essential for getting useful and accurate responses from generative AI models. A well-crafted prompt provides clear instructions, relevant context, and specific expectations of the output.

The key elements of a good prompt include:

Role Specification

Define the AI's role, helping it understand from which point of view it should generate responses. Always start with 'all potential words on earth', so you need to help it focus!

Context Information

Provide background information to set the scene. Context is crucial for the AI to produce relevant and coherent responses that fit the specific situation or task at hand.

Clear Instructions

Outline the desired action or response. Clear and concise instructions help the AI focus on the task and reduce ambiguity, leading to more precise outputs.

Expected Output

Describe the format and tone of the output desired, ensuring the response meets your expectations and is appropriate for the intended audience or application.

Here is an example of a prompt that provides the AI with a clear role, context, instruction, and desired output. See how we provide the AI with a detailed description of who we are, what our USPs are, and what we want the AI to output? The clearer you are, the better the output will be!



Goal: Create Facebook Marketing Campaign

Role Specification

You are a Facebook marketing expert.

Context Information

Our company, EcoGoods, is launching a new eco-friendly product line called 'GreenLife Essentials' next month. The product line includes reusable water bottles, biodegradable utensils, and organic cotton bags.

Clear Instructions

Develop a promotional campaign highlighting the environmental benefits and unique features of the Greenlife Essentials products. Highlight that using Greenlife Essentials can reduce plastic waste by up to 1,000 single-use items per year per person. In addition, our products are 100% vegan.

Expected Output

First, create a Facebook campaign for 20 posts spread over 4 weeks, including a description of the target audience (environmentally conscious consumers aged 18-35). Second, create the first post including the text to be posted and a description of an image that would be appropriate to accompany the text.

facebook



Facebook Campaign for GreenLife Essentials

Target Audience:

Environmentally conscious consumers aged 18-35 who are interested in sustainable living and eco-friendly products.

Campaign Overview

We will create a series of 20 posts spread over 4 weeks to promote the GreenLife Essentials product line. Each post will focus on different aspects of the product line, including environmental benefits, product features, and usage tips. The goal is to educate our audience on how GreenLife Essentials can help reduce plastic waste and promote a vegan lifestyle.

First Post: Text and Image Description

Post Text

👍 What AI Models Like vs. What They Don't Like

Just like understanding what your colleagues' preferences are, understanding what AI models 'prefer' in communication can help you create more effective prompts. Large Language Models are trained on structured text/data, unstructured data, but also code. Using a mix of the structure aspects from language, text and code you can make a clearer prompt.

What models like:

- ✅ **Clear and specific instructions:** AI models perform best with concise, specific, and unambiguous instructions. For instance, 'Write a blog post about the benefits of solar energy' is more effective than 'Talk about energy'.
- ✅ **Structured format:** Models typically prefer structured prompts using Markdown syntax, such as #, ##, and ### for headings. This helps the AI parse and understand the hierarchy of the information provided (e.g., which part of your prompt explains the model's role vs. the output examples).
- ✅ **Contextual information:** Providing relevant context helps the AI generate more accurate and relevant responses. Include detailed background information and specific examples to guide the AI.

What models do NOT like:

- ❌ **Ambiguity:** Vague instructions like 'tell me about something interesting' can lead to unsatisfactory and 'standard' outputs.
- ❌ **Overly complex sentences:** Simplify your language to ensure the AI can easily understand the instructions. Small spelling mistakes don't matter that much, but they can when they change the meaning of a sentence.
- ❌ **Noise:** Providing the model with too much information for a specific task can confuse it and lead to ineffective results. Providing the model with only the information needed for the task at hand is best.
- ❌ **Bullets vs. slashes:** Instead of using bullet points, models often prefer using slashes('/') to denote list items, because these are patterns learnt from code hierarchy.

One-Shot vs. Few-Shot Prompting

Besides learning 'what the AI models like' in communication, it is also key to understand how you can 'guide the models reasoning'. Similar to guiding a colleagues thinking, we can do the same with AI models.

One-shot prompting involves providing no example or instruction to the AI to guide its response. This method is useful for straightforward tasks where detailed context or multiple examples are not necessary and might create noise.



Prompt:

'Create an email to notify the team that the project meeting has been rescheduled to next Friday at 10 AM and include the key topics to be discussed.'

Output:

Hi team, Unfortunately we will have to reschedule the meeting for tomorrow, where we are going to discuss XYZ.'

Few-shot prompting involves providing the AI with several examples to help it understand the task and desired format. This technique is useful for more complex tasks for which you require a more consistent output format.



Prompt:

'Create an email to notify the team that the project meeting has been rescheduled to next Friday at 10 AM and include the key topics to be discussed.'

Example 1:

Hi team, The meeting tomorrow will be rescheduled to XX AM/PM. Key topics to be discussed are: A, B, and C. Please be on time.

Example 2:

Hi team, ...

Output:

'Hi team, The meeting tomorrow will be rescheduled to 10AM. Key topics to be discussed are: Project status and key next steps. Please be on time.'

3. Prompt Templates

Prompt templates are essential for maximizing efficiency and consistency when working with generative AI models. These templates provide a predefined structure that can be customized for various tasks, ensuring that the AI's outputs meet specific requirements and standards. By using templates, you can streamline the prompt creation process. This allows you to focus on refining the content rather than constructing new prompts from scratch each time.

In chapter 2, we covered the structure of a good prompt, including key elements such as the AI's role, context information, specific instructions, and desired output format. These elements should be integrated into your prompt templates to maintain clarity and precision.

Integrating fixed context and variable context

Prompt templates contain 'fixed context' that remains constant across different prompts and 'variable context' that changes based on your specific topic or use case.

Again, similar to briefing a colleague on a specific task (e.g., writing a product description or marketing content), that colleague would require 'general knowledge about our company and products' (e.g., USPs of our company) and 'specific knowledge for the specific task' (e.g., which product to write about).

Example fixed context

- **Company Information:** overview of your company, its mission, values, and key products or services, and USPs
- **Target Audience:** detailed description of your target audience, including demographics, preferences, and behavior.
- **Tone of Voice:** guidelines on your preferred tone, style of communication, preferred words and phrases.

Variable context

- **Specific Topics:** the main subject of the email, marketing campaign, or document.
- **Individual Names:** personalization elements like the recipient's name, job title, or prospect company specifics (e.g., how your USPs fit their proposition).
- **Event Details:** specific dates, locations, or other relevant specifics that are 'must know' to be able to create strong output.

Prompt Libraries

Creating your own prompt template libraries

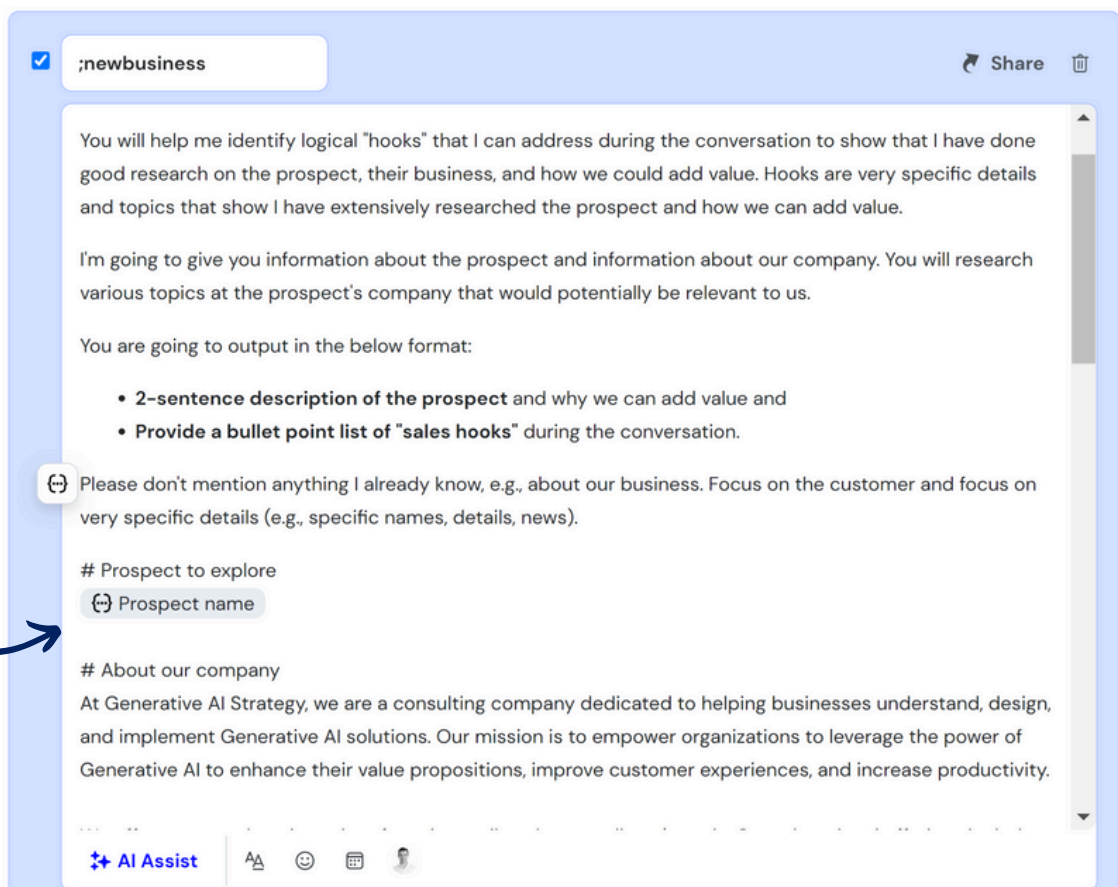
Prompt templates can be saved easily in a simple Word document or a spreadsheet, allowing team members to access and modify them as needed.

You can also use certain tools for more streamlined and efficient usage of prompt templates. Productivity tool Magical, for example, allows you to save, organize, and use prompt templates easily, enabling quick access and deployment. It lets you create libraries to store and categorize prompt templates for different use cases, quickly insert them using keyboard shortcuts or a simple interface, and collaborate and share them with team members to ensure everyone has access to the best prompt templates.

Prompt template example for sales meeting preparation

Note: Use this prompt in an AI that has search capabilities for factual context injection (RAG), for example Perplexity.ai or ChatGPT with search capabilities


Variable
context



The screenshot shows a chat interface with a prompt template. At the top, there is a search bar containing the text ";newbusiness" and a "Share" button. The main content of the prompt is as follows:

You will help me identify logical "hooks" that I can address during the conversation to show that I have done good research on the prospect, their business, and how we could add value. Hooks are very specific details and topics that show I have extensively researched the prospect and how we can add value.

I'm going to give you information about the prospect and information about our company. You will research various topics at the prospect's company that would potentially be relevant to us.

You are going to output in the below format:

- **2-sentence description of the prospect** and why we can add value and
- **Provide a bullet point list of "sales hooks"** during the conversation.

Please don't mention anything I already know, e.g., about our business. Focus on the customer and focus on very specific details (e.g., specific names, details, news).

Prospect to explore

About our company
At Generative AI Strategy, we are a consulting company dedicated to helping businesses understand, design, and implement Generative AI solutions. Our mission is to empower organizations to leverage the power of Generative AI to enhance their value propositions, improve customer experiences, and increase productivity.

At the bottom of the chat interface, there is a toolbar with icons for "AI Assist", text formatting, emojis, and a user profile picture.

4. Context Injection and RAG

As we dive deeper into the world of generative AI models, it becomes clear that basic prompting techniques are just the beginning. Advanced prompting techniques allow for more nuanced, accurate, and efficient AI interactions. This chapter explores one of those more sophisticated methods: context injection or retrieval-augmented generation (RAG).

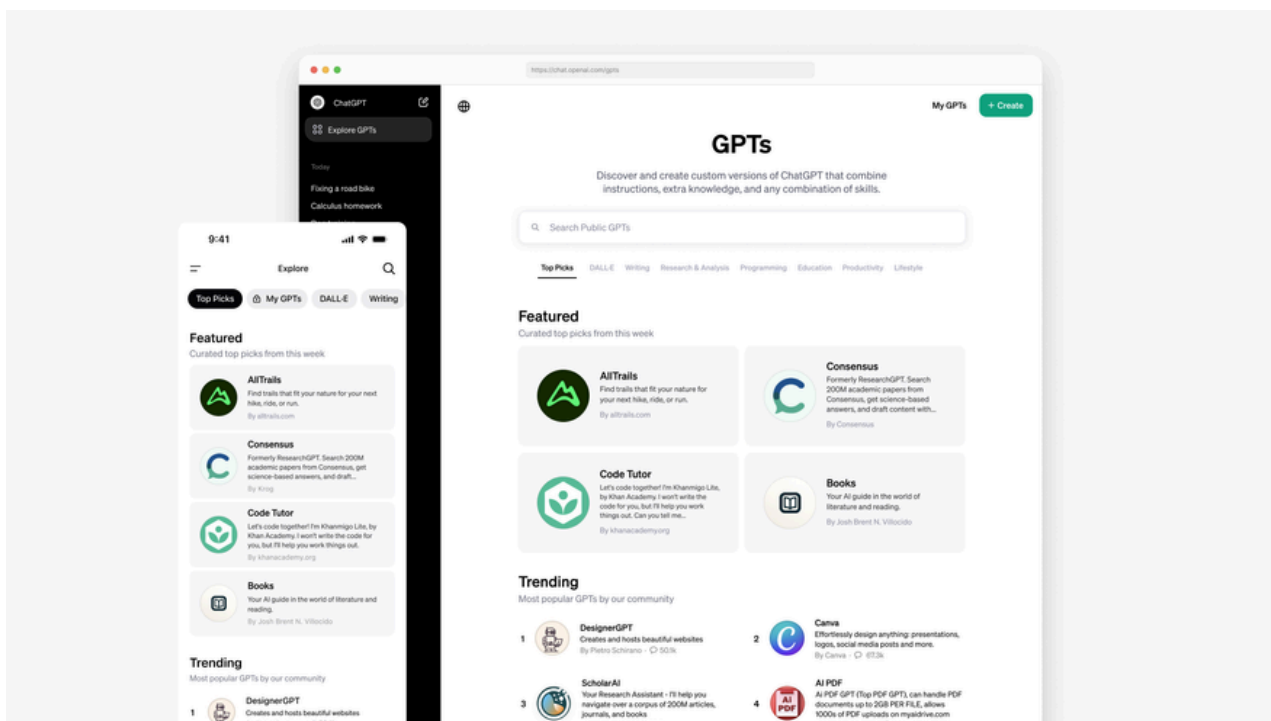
But first, let's see how OpenAI's ChatGPT makes complex techniques, like RAG, easy.



ChatGPT Custom GPTs



ChatGPT



A custom GPT is a tailored version of ChatGPT that combines specific instructions, extra knowledge, and specialized skills to perform specific tasks effectively. Unlike the general-purpose ChatGPT, custom GPTs are designed to meet the unique needs of individuals, companies, or specific use cases.

The custom ChatGPT set-up can be created without any coding knowledge, making it accessible to a broad audience. Users can build a custom GPT by starting a conversation, providing detailed instructions, adding extra knowledge relevant to the task, and selecting capabilities like web searching, image creation, or data analysis. This customization allows the custom GPT to perform tasks more efficiently and accurately, aligning closely with the user's requirements.

For instance, a custom GPT could be designed for internal company use to assist with onboarding new employees, generating marketing materials, or supporting customer service teams. By integrating company-specific data and instructions, the custom GPT can deliver responses that are not only relevant but also aligned with the company's brand and operational needs. This level of customization ensures that the AI is not just a generic tool, but a valuable asset tailored to enhance specific workflows and tasks.

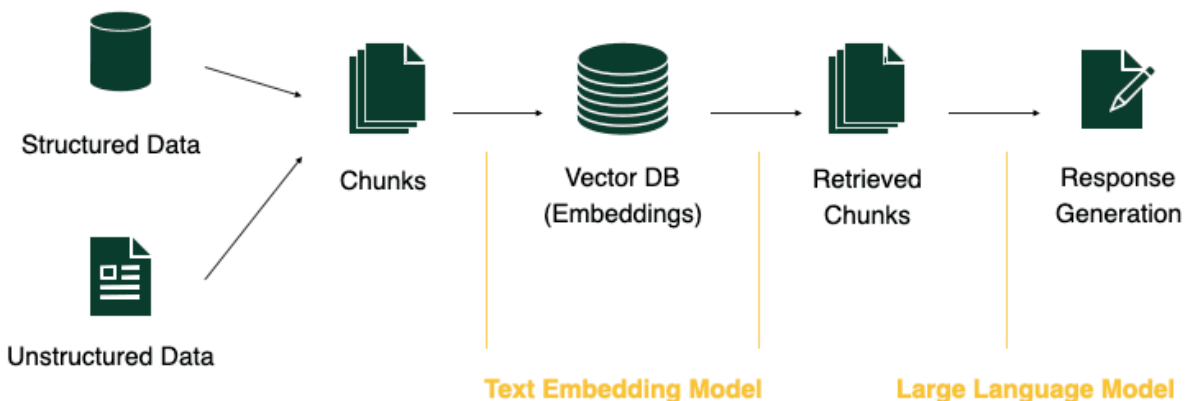
↓ Retrieval Augmented Generation (RAG)

Given that generative AI models are purely based on patterns and do not perform well in generating factual information, providing context to the model is key. This can be done by 'injecting' context into the prompt, before sending it to the AI. To stick to the colleague analogy, compare context injecting with providing a colleague access to your CRM system containing info about a prospect ('context') when asking them to write a good sales email to the prospect .

Retrieval-augmented generation (RAG) is the concept of 'retrieving' and 'injecting' context relevant to a specific task or question. It is a powerful tool that combines the capabilities of deterministic retrieval mechanisms (e.g., search engines, databases) with probabilistic generative AI models. This combination ensures the model has the context needed to generate a more accurate and contextually relevant response.

Custom RAG applications

A custom RAG application segments documents into small pieces or 'chunks'. When a query is presented, the system retrieves the most relevant chunks for context. This approach ensures that the AI leverages specific, up-to-date information (e.g., news topics, sales data, relevant consumer information). For example, when creating a personalized email, context related to the recipient might be relevant, while it could also be relevant to retrieve recent news from the city the recipient lives in.



Decisions to be made in custom RAG applications

Retrieval-augmented generation (RAG) is a technique that enhances AI models by giving them access to external information. It's like having a very knowledgeable colleague who can quickly look up relevant facts before responding to a question.

When building a custom RAG solution, several key decisions need to be made:

Chunk size	# of chunks to retrieve	Retrieval method
This determines how you break up your context. Using smaller chunks (like paragraphs) is like giving your colleague bite-sized pieces of information. Larger chunks (like full pages) provide more context but may include irrelevant details.	This is like deciding how many documents your colleague should consult before answering. More chunks mean more comprehensive information, but it takes longer to process and carries the risk of 'information overload'.	This is how you search for relevant information. Semantic similarity search, for example, understands the meaning behind words. Traditional keyword search is faster but more literal, potentially missing conceptually related context.

These decisions significantly impact the output quality. If you don't provide enough context or use an ineffective retrieval method, it's like not giving your colleague the necessary information to complete a task effectively. The result might be incomplete or include irrelevant information.

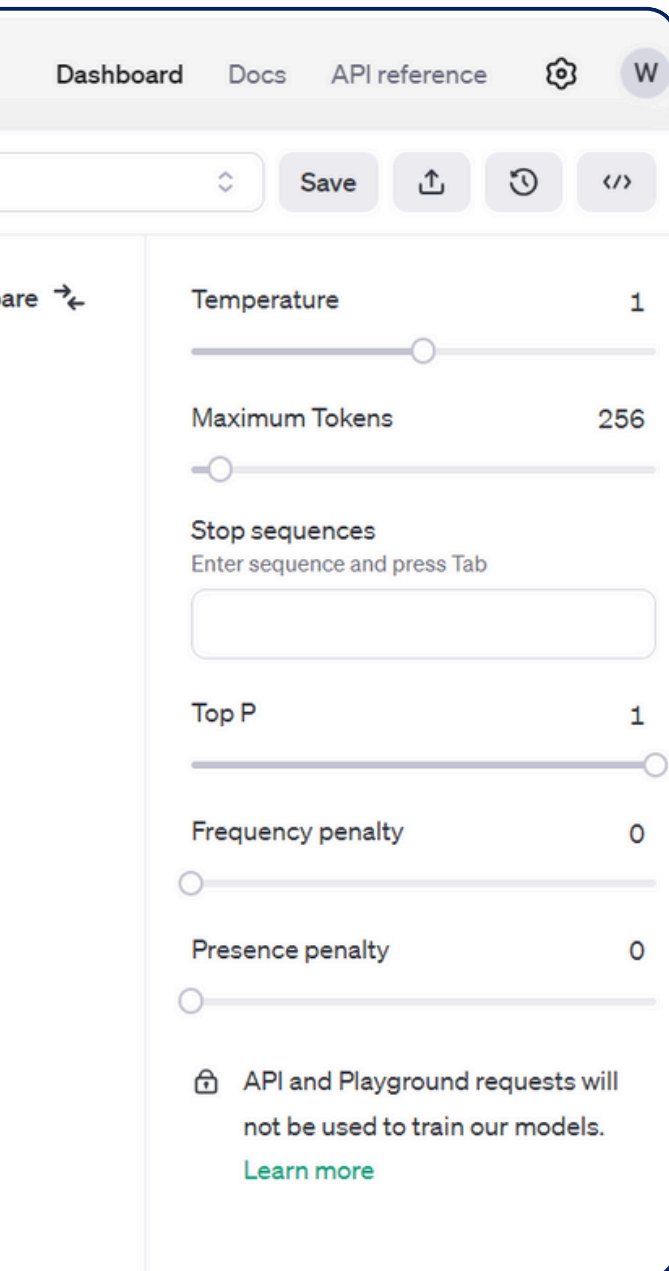
RAG in OpenAI GPTs

OpenAI custom GPTs allow you to automate the process of RAG. After uploading a document into the GPT's knowledgebase, the OpenAI custom GPT performs the chunking and retrieval processes for you - allowing you to get started with RAG within minutes. Custom GPTs can automatically pull in relevant information from predefined sources, streamlining the interaction.

Understanding what is happening 'under the hood' of AI applications can help you understand how the generative AI model used came to its output.

Different Large Language Model Parameters

When working with RAG, you often use API's. When using OpenAI's APIs or the OpenAI platform interface, understanding and adjusting the parameters of large language models (LLMs) can enhance the quality of the outputs. This allows for more control, but requires more advanced prompting.



The screenshot shows the OpenAI API playground interface with the following parameters and values:

- Temperature:** 1 (Slider)
- Maximum Tokens:** 256 (Slider)
- Stop sequences:** Enter sequence and press Tab (Text input)
- Top P:** 1 (Slider)
- Frequency penalty:** 0 (Slider)
- Presence penalty:** 0 (Slider)

At the bottom, there is a note: "API and Playground requests will not be used to train our models." with a [Learn more](#) link.

← **Controls randomness.** Lower values make output more predictable, higher values increase creativity.

← **Limits output length** by setting a maximum number of tokens (words or parts of words).

← **Signals specific words or characters** which should make the AI stop generating text.

← **Controls diversity.** Lower values lead to output with less variety, higher values result in more variety.

← **Reduces word repetition** by penalizing frequent words.

← **Discourages repeated words.**






Key Takeaways

Structured and unambiguous prompts are key for accurate AI outputs.

Providing detailed context helps the AI generate relevant and coherent outputs.

Compare AIs with new employees who needs specific and thorough instructions.

Trust In Our Expertise

-  **Educate:** We help clients understand generative AI via tailored and practical 2-4 hour generative AI workshops
-  **Strategy:** We help clients define their strategy, ideating use cases bottoms-up, prioritize, and define the roadmap
-  **Implementation:** We help clients implement generative AI, offering business, prompting, technical expertise, and more

We have worked with 150+ companies across industries



"Of all the masterclasses we've participated in at Nordian, this one differed in depth and pragmatism. Truly a **blend of visionary and actionable insights**, looking forward to our continued collaboration."

"The team's engagement with us has been remarkable. Their involvement has been instrumental in **supporting our C-suite in identifying key opportunity** by articulating and demoing implications on competitive assets & capabilities. At the same time, their help has accelerated how we operationally deploy GenAI capabilities to learn fast and scale what works"



Average rating of 8.7 / 10 (based on 258 reviews)

GEN AI

Interested? Contact us

We support companies with generative AI education, strategy, and implementation.



Wouter van Haften
Founder



Paul Melkert
AI product manager



Bram Mosterd
AI data scientist



Martine Uijland
AI Associate



Greg Nieuwenhuys
Senior partner

+/- 15 AI
experts
standby

We bring **strategic, technical, and hands-on** experience

+31 6 200 86 78 3 (Wouter)

www.generativeaistrategy.com

wouter@generativeaistrategy.com

