The Art of Asking

Blue J's guide to crafting better prompts for better tax answers

blue J

What to expect

- Foundations of Generative AI
- Effective Prompting Basics
- Advanced Prompting Techniques
- Avoiding Common Pitfalls
- Prompting Glossary

Foundations of Generative AI

blueJ

What is generative AI?

At its core, generative AI is a type of artificial intelligence that creates content based on a user's instructions—or prompts. Content produced by generative AI can include everything from audio and images, to video, code, and text. In Blue J's case, this content comes in the form of answers to users' tax questions, which are generated from a highly curated database containing only the most trusted tax resources.

Leveraging Blue J's expertise as the leading generative AI solution for tax research, we've created this guide on key prompting techniques, demonstrated using real-world examples from Blue J. But the tips here aren't just useful for tax research. With the right prompting skills and mindset, you can consistently generate better outputs from your inputs, no matter what you're prompting for.

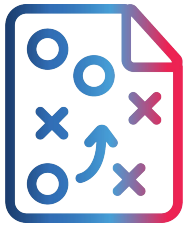

Adopting generative AI

Adopting generative AI usually requires an adjustment period. While this may sound a bit intimidating, simply understanding how these systems work can do a lot to make the transition easier.

Here's the basic process behind generative AI: after receiving a prompt, the system retrieves the relevant information and crafts a cohesive response. With Blue J, prompting the system by asking a tax question results in it generating an answer that summarizes the information you asked for, along with citations and a source list so you know where the answer is coming from.

Seems simple enough, right? It can be, but if you want to get the highest possible quality of responses from generative AI, you need high-quality prompts, prompts written with an understanding of how generative AI will process them.

Effective Prompting Basics

blueJ

Ensure clarity and precision

When prompting any generative AI system, clarity and precision are key. The more ambiguity you leave in your prompt, the harder you make it for the system to retrieve the right information and generate an accurate response.

Blue J Example

Ambiguous prompt:

"Gifting partnership interest."

Clear prompt:

"Explain the key provisions of the Tax Cuts and Jobs Act of 2017."

Provide sufficient context

Like clarity, context works to guide generative AI towards the right information. Think of it like directing the system to the correct section of a library before it begins its search. If you're used to keyword-based research, including context might feel like an extra step—but it's one that leads to more targeted and relevant responses with generative AI.

Blue J Example

Non-contextual prompt:

"What happened with the new tax law?"

Contextual prompt:

"What changes did the Tax Cuts and Jobs Act of 2017 introduce to corporate tax rates?"

Be specific

The large language models (or LLMs) powering many generative AI systems are trained on truly massive amounts of content. The more specific your prompt, the easier it is for the system to generate a response that matches your intent. By focusing on specific aspects in your prompt, the system can then deliver more targeted information in return.

Blue J Example

Vague prompt:

“Eligibility criteria.”

Specific prompt:

"What are the eligibility criteria for the Earned Income Tax Credit?"

Keep prompts relevant

Generative AI systems will do their best to deliver an output that answers your input, but they can't read your mind. If your prompt doesn't quite align with what you're actually looking for, the response you get will likely miss the mark. In order to avoid this mismatch of expectations and outcomes, make sure each prompt is *directly* relevant to your desired outcome.

Blue J Example

Irrelevant prompt:

"What is the best way to invest money?"

Relevant prompt:

"How does Section 179 affect small business equipment deductions?"

Include complete requirements

If your question has multiple components, it’s typically best to include everything in your initial prompt. The more information you provide upfront, the more complete and accurate of a response you'll receive. Including all necessary details also helps to prevent the system from making incorrect assumptions about what your question actually means.

Blue J Example

Incomplete prompt:

"Explain tax deductions.”

Complete prompt:

"Explain tax deductions available for home office expenses."

Use the right terms

The right terminology can act as an extra layer of context for your prompt, since domain-specific language helps generative AI to narrow its retrieval and generation processes. In Blue J, incorporating relevant tax terms can improve accuracy and relevance, especially if those terms frequently appear in the source materials.

Blue J Example

Non-specific language:

"What are the rules for business expenses?"

Specific language:

"What are the IRS guidelines for deductible business expenses under Section 162?"

Advanced Prompting Techniques

blueJ

Break down complex queries

While it's important to make sure your prompts are contextual and complete, the best responses sometimes come from giving the system a chance to answer layered questions one at a time. Try breaking down complex queries into smaller, manageable parts. This way, you can ensure each part is answered in full.

Blue J Example

Original prompt:
"Explain the tax implications of international income for U.S. citizens, focusing on foreign tax credits and the Foreign Earned Income Exclusion."

First prompt:
"What are the criteria for claiming the Foreign Earned Income Exclusion?"

Second prompt:
"How does the foreign tax credit work for U.S. citizens with income from foreign sources?"

Prompt iteratively

Similar to breaking down complex queries, iterative prompting involves splitting your work across multiple inputs. With this technique, you refine and adjust your prompts based on the system's initial responses, allowing you to gather more detailed and accurate information over time.

Blue J Example

First prompt:

"What are the tax benefits for small businesses?"

Second prompt:

"How does Section 179 allow small businesses to deduct equipment purchases?"

Learn from suggested prompts

In some generative AI-powered applications, like Blue J, you'll receive suggested prompts based on the system's initial responses. These suggestions can be used to guide your follow-up prompts toward deeper research, helping you explore avenues you might not have otherwise considered.

Blue J Example

Original prompt:

"What are the rules for determining the allocation of consideration received in an applicable asset acquisition?"

Suggested follow-up prompt:

"What are some potential consequences if the allocation on Form 8594 is incorrect?"

Avoiding Common Pitfalls

blueJ

Using overly broad prompts

If you're only interested in getting a high-level summary, broad prompts can be very useful. However, if you're after the finer details, these kinds of prompts can result in vague or general responses, leaving out important specifics that would have otherwise been helpful to your search.

Blue J Example

Broad prompt:
"Eligibility criteria for tax credits."

Narrow the scope:
"What are the requirements for qualifying as a head of household?"

Focus on specific provisions:
"What are the depreciation rules under Section 179?"

Asking leading questions

Biased prompts can result in biased responses—responses that may overlook relevant information. When the way you've phrased your prompt suggests a preferred result, the system will then try to deliver that result. Essentially, the more bias you introduce into your input, the more you can expect in your output.

Blue J Example

Leading question:
"Why is the IRS penalty for late filing unfair?"

Use neutral language:
"What are the implications of the IRS penalties for late tax filings?"

Ask for factual information:
"What are the IRS guidelines for abating penalties for reasonable cause?"

Ignoring context

While some contextual information might seem obvious or irrelevant to you, it's important to remember how much content generative AI has to handle. Never assume the system fully understands what you're referring to or wouldn't benefit from a bit of additional context.

Blue J Example

Ignoring context:
"What are some tax deductions I can use?"

Include background information:
"What are the tax implications for small businesses receiving PPP loans?"

Define the context:
"In the context of estate planning, what are the tax benefits of setting up a trust?"

Specify relevant code sections:
"How does Section 199A apply to sole proprietors?"

Incorporating ambiguous language

Ambiguous prompts often lead to unclear or irrelevant responses. If it's not clear what you're asking for, you can't expect the system to produce a clear response. Ambiguity forces generative AI to make assumptions—assumptions that may not always be accurate.

Blue J Example

Ambiguous language:
"What are the rules for travel?"

Use precise terms and language:
"What are the eligibility requirements for the First-Time Homebuyer Credit?"

Specify exact information needed:
"What are the IRS guidelines for classifying workers as independent contractors versus employees?"

Clarify the query:
"What documentation is required to substantiate charitable contributions on a tax return?"

Being redundant and repetitive

While specificity and context are key for generating better responses, don't forget that artificial intelligence has some real intelligence behind it. For generative AI built on large language models, the system can often read between the lines, so there's no need to include redundant information or repeat yourself when prompting.

Blue J Example

Redundancy and repetition:
"Explain the benefits of the Child Tax Credit. What are the advantages of the Child Tax Credit?"

Review for repetitive elements:
"What are the tax benefits and potential drawbacks of Roth IRA conversions?"

Be concise, but complete:
"What are the IRS guidelines for home office deductions for self-employed individuals?"

Adopting a prompting mindset

Prompting effectively is about more than just mastering techniques. It's also about adopting the right mindset.

A big part of crafting better prompts is understanding that generative AI essentially acts as a mirror. The clarity, relevance, and completeness of your prompt directly influence the clarity, relevance, and completeness of the response. Similarly, if your prompt is broad, biased, or ambiguous, your output likely will be too.

By crafting your prompts with the responses you want in mind, you're way more likely to receive those responses. Armed with the techniques in this guide, you're ready to write better prompts and generate better outputs—to make the most of generative AI.

Prompting Glossary

blueJ

Prompt

A prompt is a specific input or query provided to a generative AI system to elicit a relevant response. It is a carefully crafted question or statement designed to guide the system in retrieving and generating the desired information.

Context

> Context refers to the background information relevant to a particular prompt or query, providing additional details that help clarify the meaning and scope of the request. Context helps to align the retrieval and generation processes with the specific needs and nuances of the user's query.

Retrieval

> Retrieval is the process where relevant documents are identified from a large database in response to a given prompt. This phase leverages information retrieval techniques to find the most pertinent sources that can inform the generative response.

Generation

> Generation is the process where the system synthesizes the information retrieved and produces a coherent, contextually appropriate response to the prompt. This step uses language models to generate text that explains, summarizes, or elaborates on the retrieved information.

Query Refinement

Query refinement involves modifying and improving a query to obtain more accurate and relevant results. This may include adding context, specifying terms, or breaking down complex questions.

NATP Membership Has Benefits

NATP
National Association of Tax Professionals

NATP members receive:

Professional members:
Enjoy 7 consecutive days of unlimited questions, once per membership year.

Premium members:
Unlock 14 consecutive days of unlimited questions, once per membership year.

All members:
60% off your annual Blue J license fee

Access your benefits →