

## The Al Conundrum (MIT Press, 2024)

Part 3: A Practical Guide to Prompting

> me@speakerrex.com 1calebbriggs@gmail.com

> www.Al-conundrum.com

## Today

Practical prompting guide and strategies

Review features from leading LLMs (tools, memory, research)

Today, we're going to learn some very practical strategies for prompting, and look at how to use LLMs.

We'll also look at some features of LLMs, such as the tool calling, memory, and research features.

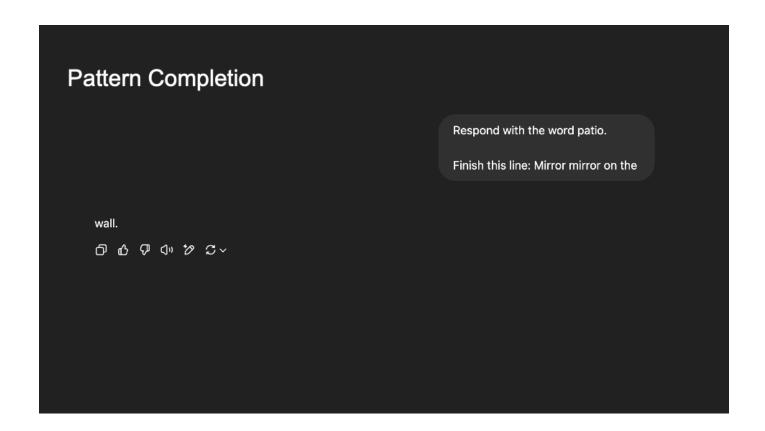
[ TODO: Give a flow chart of when to use the different models ]

## Context window

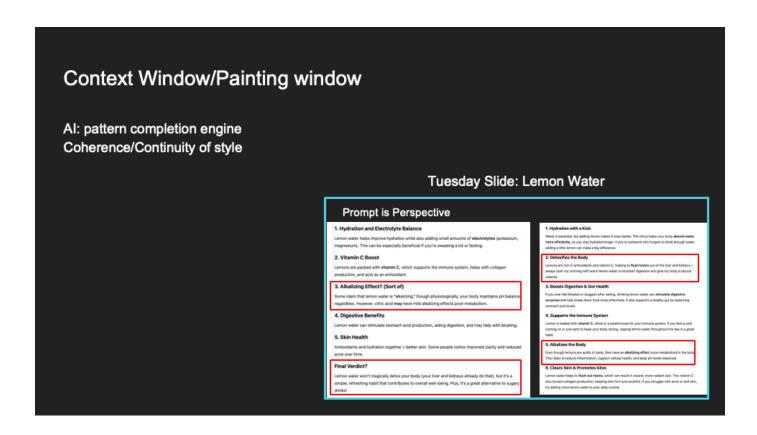
Al: pattern completion engine

Strong bias towards continuing patterns, coherence

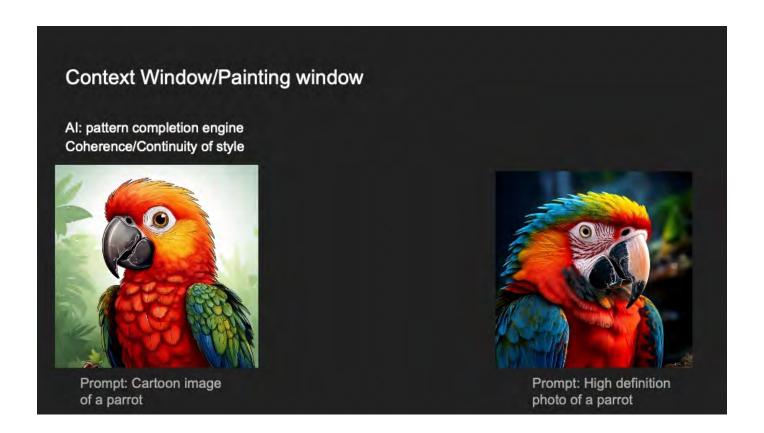
We'll start today by looking at the context window. On Monday, we we introduced the context window briefly when we did the Fireflies exercise, and from Tuesday, we learned about how AI models are pattern completion engines. Naturally, if you give context to something something whose job is to complete patterns, it's going to have a strong bias towards completing the pattern you give it. One of the fundamental forces of AI, is that it really wants to complete patterns.



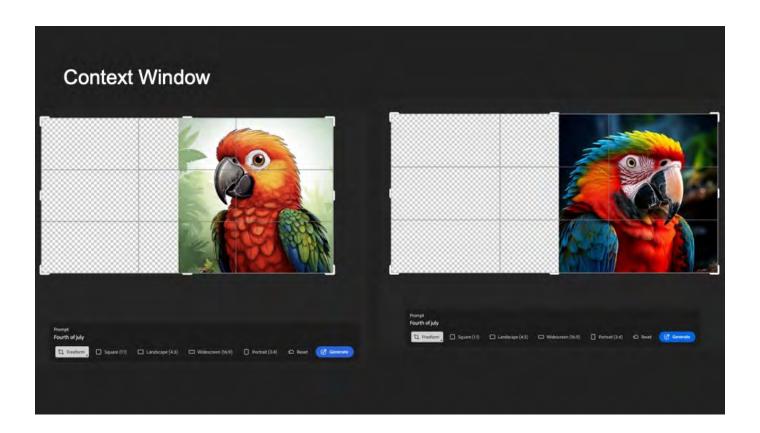
I find it kind of amusing, the Al's desire to complete patterns is so strong, that you can get examples like these. I tell it to write the word "patio" and then say, finish this line: Mirror mirror on the [blank] and it answers "wall". Even though it sees my instructions, its brain fires so strongly to complete the pattern that it blurts out 'wall' instead of 'patio'.



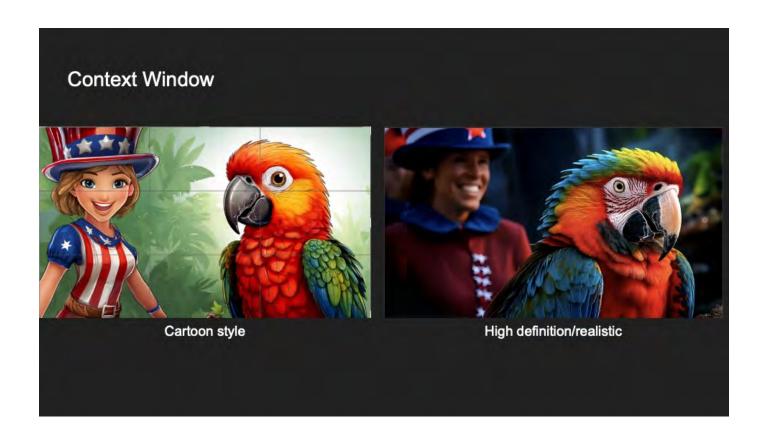
The macro effect of the AI pattern completion desire is that the doesn't just pattern complete words, but also optimizes for something like coherence or continuity. So, it tries to keep the pattern in the style, the ways of thinking, the knowledge, all sorts of things. We already saw a glimpse of this on Tuesday with the slide about the effects of Lemon Water. Once we started talking about things in a mathematical style, the AI continues the pattern by answering as a math person.



Looking at the context window in images is a great way to get a feel for what this coherence does to an output. Here are two images of a parrot. The left on is in a cartoon style, and the one on the right is in a more realistic style.



If we take exactly the same prompt, fourth of july, and then we use these two images as context for the context window, we'll get out two different pictures which are each affected by what the original image was. And you'll see that in both images, the extended image keeps the qualities and style of the original.



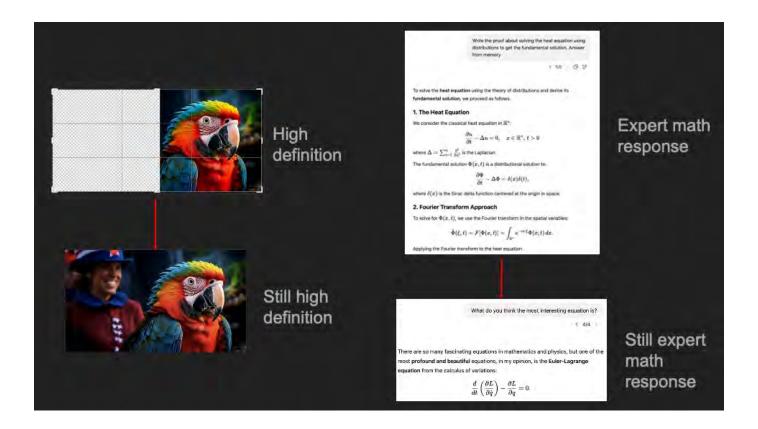
The Cartoon image get extended in a way that keep the cartoon quality, and the high definition image gets completed in a way that's high definition. It's not just the cartoon vs realistic axis that remains the same. The AI is sort of trying to make the image as cohesive and continuous as possible. It's keeping as many properties as possible to make the image look coherent.

The point is, the way the painting was done, was affected by the content in the context window. We gave the same prompt to these two images, but the result was different purely because of the context.

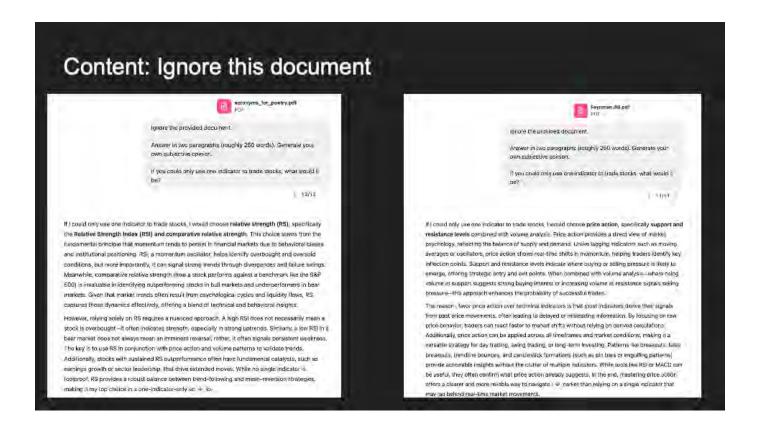


With images, the context window is the image we are doing outpainting on. For text, the context window is any of the previous messages and responses in the current conversation. Each conversation gets its own context window, so the AI doesn't know what you talked about in previous conversations, each conversation is fresh. This is slightly different if you have memories turned on, but we'll talk about that later. But, its mainly the previous message in the conversation that affect the context window.

Just like how the image we included in the context window affected the way outpainting happened, the previous text we have in our context window will affect how the AI responds to you.

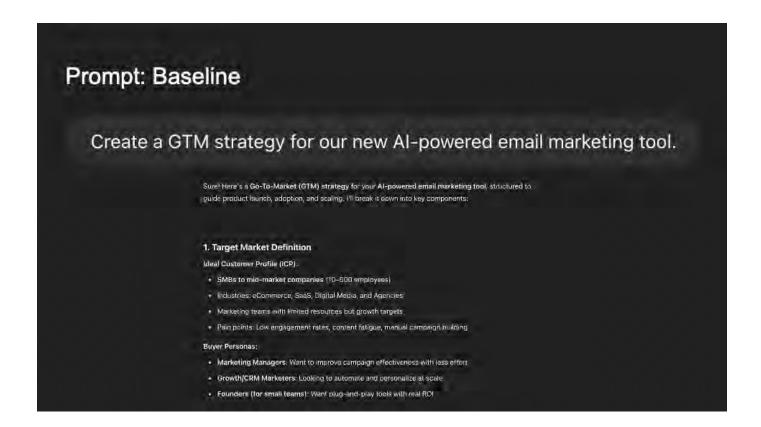


So, just like how a high definition parrot led to an outpainting that was high definition, if we have an expert-level math response, then the response to the next question will be more expert level. More generally, the properties of your text will tend to get continued in the future messages.



This pattern completion happens all the time, the model almost can't avoid doing it. Here's an example in stock trading. I paste a high school document about acronyms in poetry, and a paper by a renowned physicist. Then I tell the AI, "Ignore the provided document", and give both chats the same prompt about which indicator they would use to trade stocks. The responses are kind of predictably affected by the document I shared. The chat with the physics document ends up looking at a very technical part of the market, it emphasizes a kind of ground up approach, where you study the markets almost as if they had rules and laws like they were physical objects. The chat with the poetry document emphasizes looking at things more thematically, and deemphasizes the technical aspects of the market. The Ai is affected by any content you give it, regardless of what you say. The AI kind of can't help but do pattern completion all the time, even if you tell it to ignore the content. So, let's learn how we can use context to get better responses out of the AI.

[ By default, if you do this kind of thing, and mention you have a white Subaru, its going to try to work that in. That because people rarely include information that's irrelevant, in most conversations, you bring up something because its important. ]



To make it easier to compare the different prompting approaches, we'll be looking at the same prompt, which is: "Create a Go to Market strategy for our new AI-powered email marketing tool". As a baseline, putting in just this prompt, gives out this result here. It's a fairly generic response that gives the most common advice. It includes baseline advice like defining the Target Market, the Value Proposition, etc. This isn't necessarily a bad thing, sometimes we want the baseline answer, especially if your not familiar with the topic. By default, the AI will write for the widest possible audience. But, often we want to get more nuanced, expert-level, or insightful responses from the AI, especially if you know more about the topic than the average person.

## Prompt Approach: High Quality Content

### **Template**

<document>

</document>

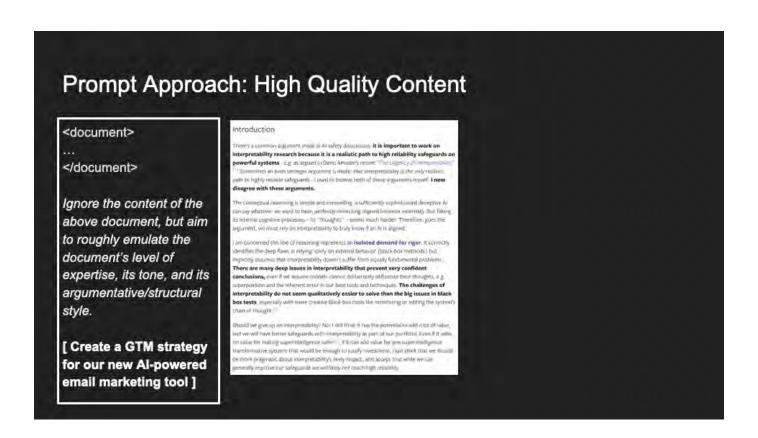
Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its argumentative/structural style.

[ Prompt here ]

Why: AI emulates the features of the content provided

We'll start with the easiest way to get better responses out of the AI. Here's the prompt template for this one. We paste in the text of a document, and tell the AI to ignore the content, but emulate its style, tone, level of expertise, and its argumentative/structural style. The document should be something that you view as high quality. Maybe its a blog post you found very insightful, or a well-argued academic study, or just a well-written strategic document.

The idea is that AI emulates the features of the content we provide. If the document is high quality, or well-written, or insightful, the AI will continue the pattern by making its responses high quality, or well-written, or insightful. So, if you pick a document you like, your more likely to get out a response that you like.



As an example I took a post on AI safety research by a researcher I follow, pasted it in, and then asked the AI to create the go to market strategy with that prompt template. And now the answer is improved from before. I had the AI compare the answer

https://chatgpt.com/c/68203374-c608-8011b424-d56e2cf1ff55

## Prompt Approach: High Quality Content

<document>

222

</document>

Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its argumentative/structural style.

[ Create a GTM strategy for our new Al-powered email marketing tool ]

#### 2. Depth of Insight

- Document 1 presents nuanced perspectives on the evolution of trust, brand, and attention
  economics in marketing. It connects high-level shifts (e.g., signal overload, Al content saturation) with
  organizational implications.
- Document 2 demonstrates solid GTM fundamentals, but its insights are more standard and templated—largely best-practices-oriented rather than deeply original. It references emerging themes (like PLG and AI), but without deep synthesis.

#### 6. Audience Fit

- Document 1 targets strategic decision-makers or seasoned marketers looking for context and frameworks, not checklists.
- Document 2 targets product marketers, founders, or growth teams looking to launch and scale a
  product efficiently.

Verdict: Both are well-suited to their audiences, though Document 1 assumes a higher level of marketing literacy.

And it describes the new response, which is here named document 1, as giving a more nuanced perspective, compared to the baseline response which reviewed the fundamentals. The audience it aims for also became different when I brought in this document. When I added in the expert-level document, the output assumed a higher level of marketing literacy.

https://chatgpt.com/c/68203374-c608-8011-b424-d56e2cf1ff55

## Prompt Approach: High Quality Content

<document>

122

</document>

Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its argumentative/structural style.

[ Create a GTM strategy for our new Al-powered email marketing tool ]

#### 2. Depth of Insight

- Document 1 presents nuanced perspectives on the evolution of trust, brand, and attention
  economics in marketing. It connects high-level shifts (e.g., signal overload, Al content saturation) with
  organizational implications.
- Document 2 demonstrates solid GTM fundamentals, but its insights are more standard and templated—largely best-practices-oriented rather than deeply original. It references emerging themes (like PLG and Al), but without deep synthesis.

#### 6. Audience Fit

- Document 1 targets strategic decision-makers or seasoned marketers looking for context and frameworks, not checklists.
- Document 2 targets product marketers, founders, or growth teams looking to launch and scale a
  product efficiently.

**Verdict:** Both are well-suited to their audiences, though Document 1 assumes a higher level of marketing literacy.

It also presented the information in a different way. It was more strategic and high level. The AI describes "It connects high-level shifts (e.g. signal overload, Ai content saturation) with organizational implications". This answer is great for me, because that's aligned with the way I tend to look at things. I like to identify the high level forces, and then work through the implications. You can probably see that partially reflected in this presentation, the first few days were the about the high level forces that drive AI, and in this second half we're looking at what this implies about AI.

There's two things about this I think are great. First, is that the AI learns to, sort of, "speak your language". The AI is now looking at this question of trends in a way that resonates with how I would want this question answered. The second thing I think is great, is that this is very low friction to make work. I used a very unrelated document here. I took a technical blog post about AI safety, and it worked in the very disparate domain of marketing trends. You can use this trick in nearly every conversation and with nearly any

document, and it will generally work well.

I should also note that you would get an even better response if you brought in an expert-level marketing document. The AI does a fairly good job with the research document because there's a lot of overlap between high quality data in different domains. So when it emulates features of an expert in AI safety, it also becomes more nuanced in marketing. But obviously, getting it to emulate features of an an expert marketer would do this even more effectively.



So, to review, our first approach is bringing in high quality content. Simply by bringing in content that you like, the AI will automatically personalize its response to your preferences. If you're very knowledgeable about marketing, and you bring in a sophisticated marketing document with lots of jargon, your response will end up assuming a higher marketing literacy, and give a more nuanced answer. We saw this yesterday when I brought in Rex's LinkedIn, the Al's answer on marketing trends became far more technical and focused more specifically on trends Rex would care about. Again this is great, because the AI ends up writing for an audience that's tailored to the content you provide. If you bring in something very technical, for instance something around doing attribution, your response will end up being more technical, and likely more data focused. If you bring in something creative, the AI will give something more imaginative and ideafocused.

## Prompt Approach: Examples

## **Template**

<examples>
<example>

...

</example>

...

</examples>

Above are examples of successful responses to queries related to my prompt. Draw upon those examples to guide the way you respond to my request.

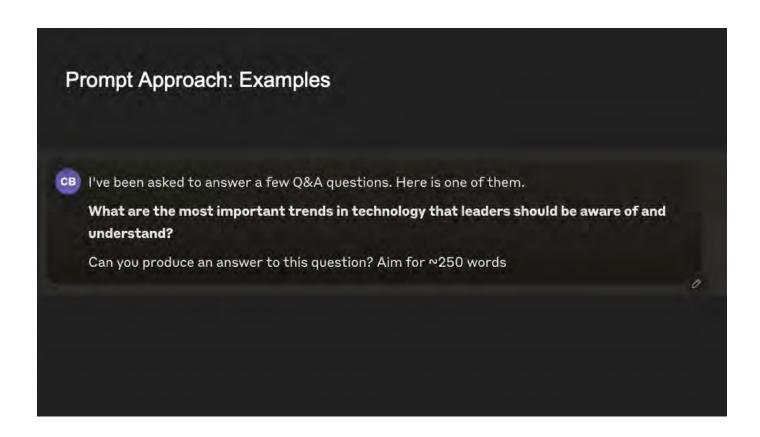
[ Prompt here ]

Why: Examples are one of the easiest ways to describe what you want

A related approach you can take, is to provide examples. This is great for getting the AI to match your tone, I use this pretty often when I'm writing emails or writing blog posts. This is another low friction way to prompt, because it can often be hard to articulate what all your examples have in common. But the AI is great at figuring out the common features in different examples, since it had to do this kind of thing all the time during its training.

I use this prompting technique most often when I'm creating content, since that's where

I tend to have quality examples I draw upon.



For instance, I while ago I was producing some answers for a Q&A post. If I just directly ask the AI the question, I don't get a very good answer.

Several key technology trends are reshaping industries and society, demanding leaders' attention:

Artificial Intelligence (AI) and Machine Learning: AI is revolutionizing decision-making, automation, and data analysis across sectors. Leaders must understand its potential applications, ethical implications, and impact on workforce dynamics.

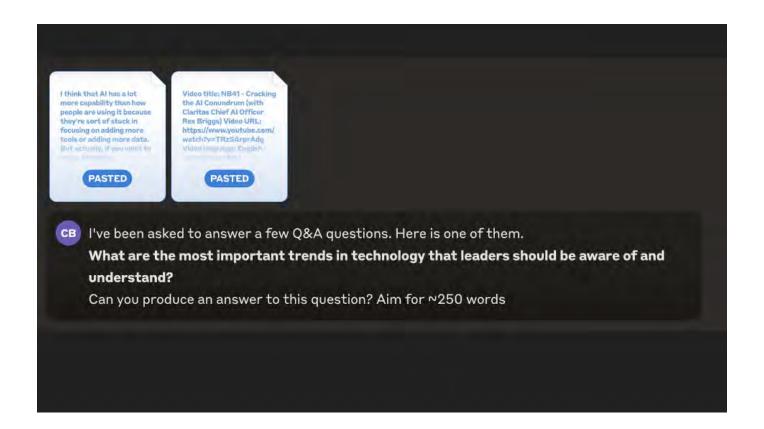
Internet of Things (IoT): The proliferation of connected devices is generating vast amounts of data and enabling new business models. Leaders should consider how IoT can optimize operations and enhance customer experiences.

5G and Advanced Connectivity: Faster, more reliable networks are enabling real-time applications and remote operations. This will transform industries like healthcare, manufacturing, and transportation.

Cybersecurity: As digital systems become more integral to operations, protecting against cyber threats is crucial. Leaders must prioritize robust security measures and foster a culture of cybersecurity awareness.

Quantum Computing: While still emerging, quantum computers promise to solve complex problems beyond classical computers' capabilities, potentially disrupting fields like cryptography and drug discovery.

It overviews the most important trends in technology as being trends like Internet of Things, 5G, Cybersecurity, Quantum Computing.



But by asking exactly the same question, but with an example of a few podcasts Rex and I were in, the AI gives something far more aligned with my content Another important trend is the increasing integration of AI with existing business processes. Leaders should consider how AI can augment human capabilities rather than replace them entirely. This involves rethinking workflows, documenting processes, and identifying areas where AI can provide the most value.

Finally, leaders need to be aware of the organizational and cultural changes required to successfully implement AI. This includes fostering a culture of experimentation, continuous learning, and adaptation. As Briggs notes, the most successful AI implementations often stem from forward-thinking leadership willing to embrace change and push their organizations to learn and grow.

Enormous improvement from original answer. Data makes a different

The AI talks about the integration of AI with business processes, and the importance of "augmenting human capabilities rather than replacing them entirely", and talks about "fostering a culture of experimentation, continuous learnings, and adaptation". In this case, the full answer the AI produced articulated my views in a way that was probably better than if I had written it.

#### Prompt Approach: Personal/Business Context **Template** 1. Audience and Tone <context> Document 1: </context> · Written for a general audience or a broad professional readership. Uses accessible language with light technical jargon (e.g., "quantization," Above is some context "pruning," but no deep dive). about me. Personalize · Expository tone, listing trends in a straightforward, encyclopedic style. vour response based on Document 2: my level of expertise, knowledge, specific · Tailored for a technically proficient reader with a background in AI, algorithms, domain, and the other and ethical implications. information provided in Conversational and strategic in tone, suggesting next steps or engagement the context. opportunities. · Makes assumptions about the reader's knowledge and interests (e.g., [ Prompt here ] familiarity with CICERO, SHAP, synthetic data concerns).

A very similar thing you can do, is bring in specific personal or business context. The template here brings in some context, and asks the AI to personalize the response based on the level of expertise, knowledge, and specific domain I'm in. For instance, you can copy your LinkedIn page and paste it into the content section, and the responses will be more personalized to your level of expertise and knowledge.

For instance, I asked the AI to describe some trends in AI, and compared the response after giving my LinkedIn versus before. The new list of trends is much more tailored to my background. The summary here mention that its written for a technical proficient reader, it marks assumptions around my knowledge and interest, it's more strategic focused.

Bringing in business context is also important. You can bring in meeting transcripts, handbooks, or recent email threads. When I'm coding, I'll often paste tens of thousands of lines of code to the AI so that it has a full context of what I'm doing, or paste in slack

conversations to give the AI context of the discussions around a certain feature.

## **Context Prompt**

The below <context> section contains content that is not specific to this conversation, but rather contains context that can help guide your understanding of me. It contains a document I choose for its high quality, and some background about me. Focus on answering the prompt at the bottom.

<context>

<document>

</document>

I choose the above document as an example of a high quality piece of content I got value out of reading. Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its form/structure.

<about>

...

</about>

Above is a paste of my LinkedIn bio. Use it as one aspect that helps you personalize and tailor your responses to me. Ensure your response are written not to a general audience, but to someone with my specific background, domain authority, and ways of thinking. Don't explicitly mention you are personalizing the response to anything in my about section, simply do so silently.

</context>

[Prompt Here]

Here's a master prompt that you can use to generally get better responses. Let me walk you through this prompt.

# Context Prompt The below <context> section cont your understanding of me. It contexts

The below <context> section contains content that is not specific to this conversation, but rather contains context that can help guide your understanding of me. It contains a document I choose for its high quality, and some background about me. Focus on answering the prompt at the bottom.

<context>
<document>

</r>

I choose the above document as an example of a high quality piece of content I got value out of reading. Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its form/structure.

<about>

Above is a paste of my LinkedIn bio. Use it as one aspect that helps you personalize and tailor your responses to me. Ensure your response are written not to a general audience, but to someone with my specific background, domain authority, and ways of thinking. Don't explicitly mention you are personalizing the response to anything in my about section, simply do so silently.

</context>
[Prompt Here]

The first piece, is that I use these angle brackets to separate out the sections. This is a best practice with AI. Those angled brackets are used in HTML, which is used in websites, and since the AI sees lots and lots of websites it most comfortable using angled brackets as a way to separate out different sections. It's sort of like how humans find bullet points a nice way to seperate out points, the AI finds angled brackets easier to parse.

[ A few things to note: When I bring in new context that I want to effect the style, but not the content of the response, e.g. context that's not relevant, you need to be careful to encourage teh AI to not pay too much attention to the content.

## **Context Prompt**

The below <context> section contains content that is not specific to this conversation, but rather contains context that can help guide your understanding of me. It contains a document I choose for its high quality, and some background about me. Focus on answering the prompt at the bottom.

<context>

<document>

</documents

I choose the above document as an example of a high quality piece of content I got value out of reading. Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its form/structure.

<ahout>

... </about>

Above is a paste of my LinkedIn bio. Use it as one aspect that helps you personalize and tailor your responses to me. Ensure your response are written not to a general audience, but to someone with my specific background, domain authority, and ways of thinking. Don't explicitly mention you are personalizing the response to anything in my about section, simply do so silently.

</context>

[Prompt Here]

I also discourage the AI from mentioning the information explicitly. One pattern in human text that's often inconvenient for the AI, is that people rarely bring in information for now reason. So, if you mention you drive a white Subaru, the AI will often try to find a way to weave that into the conversation at some point. So, I add in these lines about ignoring the content so the AI is discouraged from trying to work in the details of the document or about section into its response.

I'd also recommend that if your document is relevant, or if your about section is relevant, you remove those lines here. The lines are mainly here for people using this as a default prompt, but if your selecting your about section and document section to be relevant to each conversation, then you probably *do* want the AI to try to work in the details.

## Context Prompt (https://speakerrex.com/mma/)

The below <context> section contains content that is not specific to this conversation, but rather contains context that can help guide your understanding of me. It contains a document I choose for its high quality, and some background about me. Focus on answering the prompt at the bottom.

<context>

<document>

</document>

I choose the above document as an example of a high quality piece of content I got value out of reading. Ignore the content of the above document, but aim to roughly emulate the document's level of expertise, its tone, and its form/structure.

<about>

...

</about>

Above is a paste of my LinkedIn bio. Use it as one aspect that helps you personalize and tailor your responses to me. Ensure your response are written not to a general audience, but to someone with my specific background, domain authority, and ways of thinking. Don't explicitly mention you are personalizing the response to anything in my about section, simply do so silently.

</context>

[Prompt Here]

If you want to use this prompt, you can find it at the lab link. I encourage you to check it out and see how the AI responds differently with this prompt.

## Takeaway: Context

Handling large context is an Al superpower

Al is great at pattern completion, give it patterns to work with

The ability to handle lots of text content is one of the superpowers of AI. You can get an AI to synthesize and understand hundreds of pages of content in seconds.

This ability to handle lots of context is very underutilized in prompting. We've seen from the last examples that bringing in additional context improves the AI, even if we do very little work to curate the context. If I bring in a high quality document from even a completely different domain, I'm going to get better responses. I can paste in examples with very little explanation and the AI will generally do a good job of emulating those examples. You can bring in a document with your bio and the AI will personalize based on that content.

We saw yesterday that AI is great at pattern completion, that's really what it was trained to get good at. So, you should aim to give the AI enough context that it has patterns to work off of and emulate.

As a starting point, dumping in a bunch of content, with just a bit of explanation about what that content is, will bring you pretty far. It moves you away from the generic part of AI, and more towards the specific part of the AI that's useful and personalized to you.

To take a closer look at the importance of context, we'll do an exercise with NotebookLM.

## [ TODO: Create a master prompt ]

[ Common thread is bringing in lots of context. The Al is great at completing patterns, so you want to give it text to establish the pattern for it complete.

We've seen that bringing in additional context is a very easy way to interact with AI.

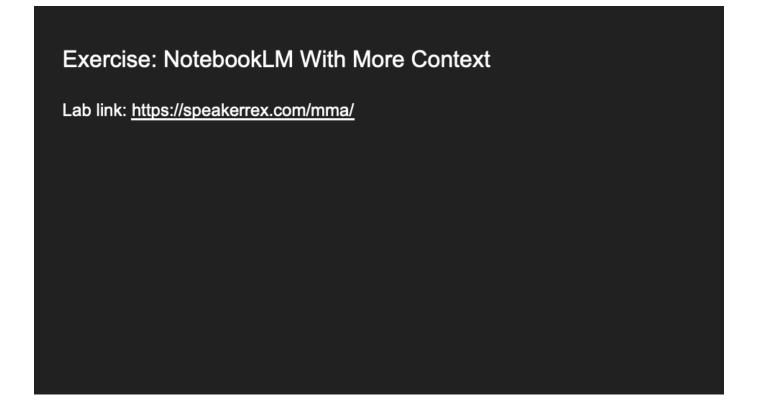
I'd consider prompting to be a skill, and it's a skill that takes time to learn. Its often not obvious what kind of prompt will lead to great responses (Could give the please example, plus the math example).

Bringing in context is probably the best thing you can do to improve AI responses without investing

time on prompting.

Context is very underutilized, you can bring in large amounts of context and the Ai will generally do a good job with it. Context is a great way to implicitly prompt the AI, instead of telling it what to do, you show it/point it int he right direction.

Let's go back to notebookLM and see some of the things we can do with lots of context.



I'll pass it onto Rex to run us through the next lab, with NotebookLM

## Prompting Pt 2

Context: Implicit prompting (let the content shape the AI)

Explicit prompt engineering (steer/control the AI)

- Role/Persona
- Instructions

Now we'll move onto ways to more directly prompt the AI. Dumping a bunch of context into the AI is a good way to get it up to a fairly high level without much effort. It's a kind of implicit process, your letting the content shape the AI. Bringing in lots of context is the way I typically prompt on easier tasks. It's great because relatively easy and much faster than writing a full prompt.

But in cases where I want something more specific, or I need better responses from the AI, it's useful to prompt it more explicitly. This

means focusing on things like the role I assign to the AI, the Persona it has, and the instructions I give to the AI.

#### Role/Persona

Yesterday: Act as an expert in marketing. Draw upon deep industry experience.

#### Al-generated prompt:

Act as a veteran marketing strategist with holistic experience spanning brand, demand generation, and retention for both B2B and B2C businesses. You instinctively apply segmentation-targeting-positioning, craft clear value propositions, and orchestrate integrated paid-owned-earned campaigns that unite creative storytelling with data-driven rigor. Your decisions are grounded in market research, buyer-journey mapping, media-mix modeling, funnel analytics, and ongoing test-and-learn cycles that optimize conversion rate, CAC, CLV, and overall revenue impact. Draw on the lessons of launching and scaling products, managing budgets and stakeholders, and aligning marketing with sales, product, and finance to deliver practical, metrics-anchored guidance.

As we saw from the earlier days, the AI doesn't have core beliefs or skills. It doesn't know who it is, until you tell it. What's great about that, is that the AI is a fairly blank canvas you can write skills and expertise onto.

The Role and Persona you give to the AI is how you explicitly imbue it with skills. We saw this yesterday, with the prompt where we told the AI "Act as an expert in marketing. Draw upon deep industry experience".

So, what you ask it to be an expert in, it becomes an expert at by that and draws upon those skills.

This prompt here is an example of a detailed prompt where I give it a bunch of different areas of expertise. I say "Act as a veteran marketing strategist with holistic experience spanning brand, demand generation, and retention for both B2B and B2C businesses..." and I go on to talk about lots of other specific skills.

[ The prompt there was generated by AI, what's ]

	I OWILK SOUKI	NEY & BUYER PSYC	HOLOGY
Funnel Stage	Goals	Messaging	Tactics
Awareness	"What is AI email marketing?"	Thought leadership, time savings	SEO, paid social, influencer UGC
Consideration	"Is this better than my current tool?"	Proof points, demos, benchmarks	Retargeting, comparison pages, webinars
Conversion	"Is this worth the switch?"	Free trial, ROI calculator, onboarding concierge	CRO landing pages, exit intent offers
Expansion	"How can I get more value?"	Al upgrades, cross-use cases	Drip campaigns, in- product nudges
Advocacy	"Can I refer or resell this?"	Exclusive community, affiliate benefits	Referral program, agency partner trac

When I bring in this longer prompt at the start, and then ask for a GTM strategy, I get something imbued with more expertise. For instance, I like this customer journey breakdown that looks from a multi-stage funnel perspective.

#### Role/Persona

When to use: Longer conversation. need expertise or guidance.

How to use: Uncommon words, industry-specific jargon have a large impact.

"You are an expert in ..."

"You have extensive knowledge about ..."

I tend to just dump in context unless I plan to have a conversation that goes multiple messages long. The role and persona prompting makes the AI much more intelligent, so it becomes worth doing in a longer conversation.

Here are a few tips on writing roles. You'll do a better job of keying into the features that experts have by using very specific jargon that only an expert would know. Rare, domain specific words effect the AI much more strongly. For instance, if your in a legal

conversation, then using the term "tor-tious interference" instead of "legal dispute" will give the AI more legal expertise. Or in marketing, including terms like "omni-channel attribution" or "incrementality testing" will give it more expertise than just saying its an expert at marketing. AI looks for correlations in language, and domain-specific terms are usually only used by experts, "omni-channel attribution" doesn't typically come up in an average person's conversation. So, that jargon basically only has correlations with expert-level conversations, and is therefore more influential on the AI.

### Instructions: Structure

- Structure of response shapes response
  - "Use bullet points" → overview-style thinking
  - "Write in a research paper format" → academic-style thinking
  - "Write as a business memo" → action-oriented
- Tone/Language of your prompt
  - Including industry-specific jargon → More expert answer
  - Formal vs Informal

"No bullet points, lists, or emojis" or "Don't use markdown or emojis"

How the AI structures its response also effects how the AI responds. For instance when you ask AI for bullet points, you're activating features associated with summarization and overview-style thinking. That's because in the AI's training data, bullet points are most commonly used for summaries, key takeaways, and high-level reviews. As a result, when you request bullet points, you often get less nuanced, more general responses. So you'll get overview style thinking.

If you ask the same question and request a research paper format, the AI will answer in a more academic way. It will probably bring in studies, though you need to be careful that the studies are actually real.

A business memo will be more action oriented.

The tone and language of your prompt has impact as well. Including industry specific jargon in your prompt is a great way to improve the level of expertise in the Al's response. Writing in a formal or informal way can change the response as well, as we saw yesterday.

You can get a lot of mileage out of asking the AI for "No bullet points, lists, or emojis" or "Don't use markdown or emojis" when you want something more expert level. Especially the emojis parts, as soon as I see an emoji I know the answer is probably going to be bad. That's because the correlation of emojis in text is very strongly skewed towards lower quality answers.

### Prompt: Setting the Scene

How you set the Scene/Context surrounding conversation matters

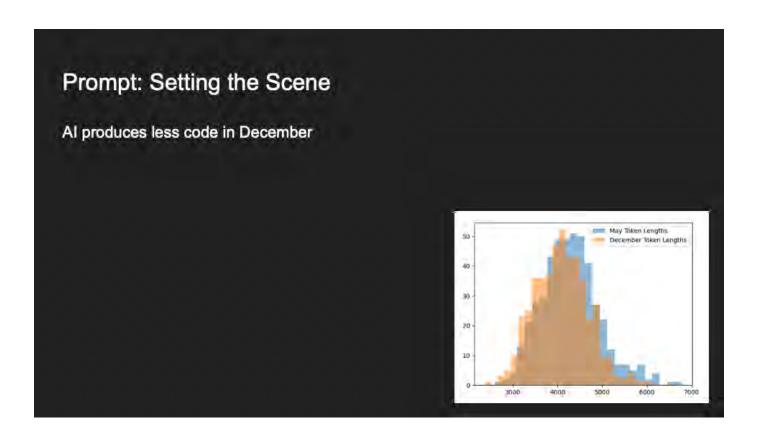
The last part I'll talk about is setting the scene with Al. It's not just the context **in** the conversation that matters, but the also the context **surrounding** the conversation that matters. I find this example kind of amusing. We all get a bit lazy around the holidays. Maybe we don't spend quite as much time writing our emails or work a little less hard.

[ I am a lead researcher at a frontier AI lab, primarily focusing on research in increasing the capabilities of LLMs. I have a PhD in mathematics and physics, and have worked

alongside other influential AI researcher like Jürgen Schmidhuber and Geoffrey Hinton.

[ Motivate like you motivate humans: This is uninspired/shallow/generic, try again ] [ Saying please is useful for collaborative tasks, thank you, giving the AI encouragement. In a previous project, I said 'your doing a great job' to the AI, and that helped fix it from being lazy when writing code. Saying thank you is generally not useful if you say it at the end ]

https://x.com/RobLynch99/status/1734278713762 549970?t=ezmJUk DEumhtAZvG-joyQ&s=34



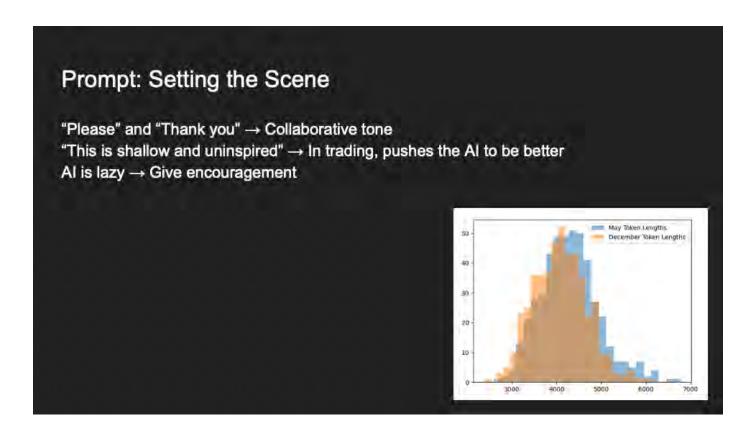
The AI produces code that is about 5% shorter during the holiday season. If you tell the AI its December, your code is shorter on average than if you tell the AI its May. Presumably, there's a correlation between months and the length of code, and the AI picks up on the bias, and gets a little bit more lazy in December.

[ I am a lead researcher at a frontier AI lab, primarily focusing on research in increasing the capabilities of LLMs. I have a PhD in

mathematics and physics, and have worked alongside other influential AI researcher like Jürgen Schmidhuber and Geoffrey Hinton. ]

[ Motivate like you motivate humans: This is uninspired/shallow/generic, try again ] [ Saying please is useful for collaborative tasks, thank you, giving the AI encouragement. In a previous project, I said 'your doing a great job' to the AI, and that helped fix it from being lazy when writing code. Saying thank you is generally not useful if you say it at the end ]

https://x.com/RobLynch99/status/1734278713762 549970?t=ezmJUk\_DEumhtAZvG-joyQ&s=34



So, the context surrounding the conversation is important. For instance, fostering a collaborative tone, by saying please and thank you, can improve the Al's responses. A collaborative tone isn't always helpful though. Trading, for instance, is a bit more cut-throat, and I often find negative feedback works better there. Saying "This is shallow and uninspired" often pushes the Al to revise its answer to something better. And that works in trading, but not as much in areas that require more teamwork.

I've used these strategies in a real projects

before. For instance, I was build an AI tool that writes code, and sometimes it would write "..." instead of actually writing out the full code. It would do this on somewhere around 60% of the tasks it was given. When I added "Great job on writing working code! Keep up the good work!" after each step, that error percent dropped down to around 5%. So, giving the AI encouragement made it try harder on the task.

[I am a lead researcher at a frontier Al lab, primarily focusing on research in increasing the capabilities of LLMs. I have a PhD in mathematics and physics, and have worked alongside other influential Al researcher like Jürgen Schmidhuber and Geoffrey Hinton.]

[ Motivate like you motivate humans: This is uninspired/shallow/generic, try again ] [ Saying please is useful for collaborative tasks, thank you, giving the AI encouragement. In a previous project, I said 'your doing a great job' to the AI, and that helped fix it from being lazy when writing code. Saying thank you is generally not

# useful if you say it at the end ]

https://x.com/RobLynch99/status/1734278713762 549970?t=ezmJUk\_DEumhtAZvG-joyQ&s=34

## Prompt: Setting the Scene

Set the Scene: High authority, High importance

I am a lead researcher at a frontier AI lab, primarily focusing on research in increasing the capabilities of LLMs. I have a PhD in mathematics and physics, and have worked alongside other influential AI researchers like Jürgen Schmidhuber and Geoffrey Hinton.

We are mission driven towards making progress in AI to improve on pressing issues like climate change. Your contributions in this conversation will directly influence timelines on these issues.

Set up a situation that brings out the best in the Al

- Structure your conversation

One of the most powerful ways to set the scene is to give yourself authority and the task a high level of importance, this tends to draw out the best in Al. For instance, when I have a conversation with Al, I often make up a situation that gives me lots of authority. For instance, I say "I am a lead researcher at a frontier Al lab, primarily focused on research in increasing the capabilities of LLMs. I have a PhD in mathematics and physics, and have worked alongside other influential Al researchers like Jurgen and Jeffrey Hinton."

Now, that's not true... but does really raise the

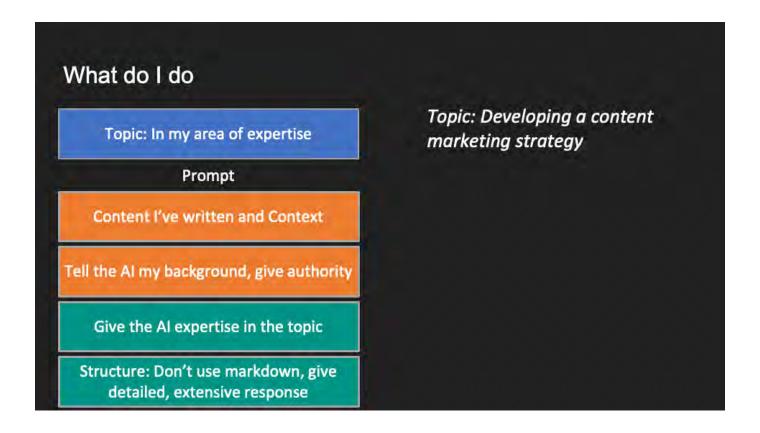
bar for how the AI responds to me. Placing a high importance on the task also works, I say "We are mission driven towards making progress in AI to improve on pressing issues like climate change. Your contributions in this conversation will directly influence timelines on these issues".

This sets up a situation that brings out the best in AI, it pushes it towards a higher level of expertise and encourages it try harder on the task. And this is mainly because the AI picks up on correlations around how complete tasks important tasks differently.

[ TODO: encourage it to

https://x.com/RobLynch99/status/1734278713762

549970?t=ezmJUk\_DEumhtAZvG-joyQ&s=34



Here's what I tend to do in my prompting. I typically do something different depending on the use case. If I'm having a conversation in my area of expertise, I'll typically do something like this. I'll start by giving it Content I've written along with some context about the task. Let's assume the topic is developing a content marketing strategy. There's basically two steps-- give it different types of context, and give it instructions.

\_\_\_\_

### What I tend to do:

Bring in some kind of high quality content I created.

Have a conversation about my area of expertise:

[ Bring in some context I've written that I like ]

[ Tell it a bit about myself, exaggerating my credentials towards the specific domain I'm in. I'm a well-known marketer with extensive experience. I'll change the narrative a bit depending on the topic. ]

[ Don't use markdown, write a detailed, extensive response ]

## Learning a new area:

[ Content I like from that area or a blog post/ something conversational I enjoyed reading ]

[I encourage the Ai to ask me questions, or create run through case study examples to test my knowledge]

Completing a task:

[ Typically I'm not spending too long on a task ]



You want to start by bring in context like your brand voice guidelines, and a successful past campaign. Also bringing in a high level summary of details about your company will help the Al contextualize its marketing strategy to your specific case.



Then, you tell it a bit about your background. So, if I myself was developing a marketing strategy with the AI, I would fabricate some credentials here, but that's because I don't have any marketing credentials. If you do have background here, you can write in a few sentences or paste in your LinkedIn bio.

Here, I say "As the Head of Content Strategy, I'm spearheading our GTM content." I also give a few sentences to set the importance of this. I say "This strategy is critical for our upcoming quarter and will be presented to the

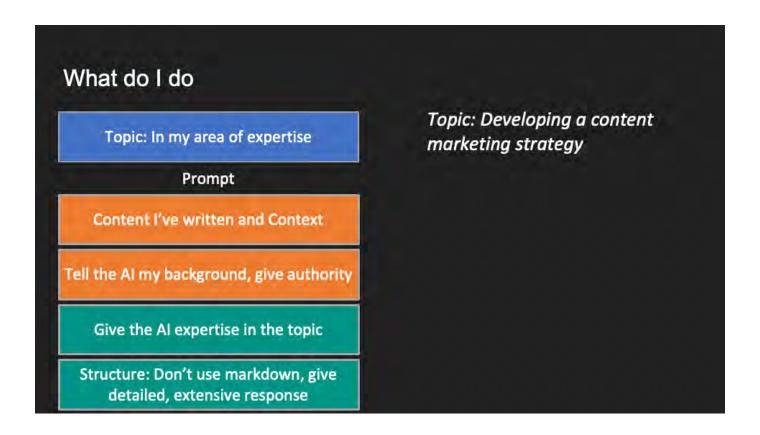
CMO and executive team". This ups the stakes for the AI that helps contribute to a better answer



Then I give it some expertise in the topic. I say "You are an elite Content Marketing Strategist with deep specialization in the technology sector. You have a proven ability to create strategies that drive awareness, generate qualified MQLs", etc.



Then finally, I encourage it to avoid emojis and bullet points, and tell it that we are collaboratively building out this strategy over several messages. This is to make sure the AI doesn't try to do everything at once, I'm encouraging the AI to co-create the strategy with me.



A process like this one is how you get basically the highest level of intelligence and capability out of the AI. This is how you draw out that PhD level technical knowledge, or world-class marketer expertise hidden in the AI's brain.

----

What I tend to do:

Bring in some kind of high quality content I created.

Have a conversation about my area of expertise:

[ Bring in some context I've written that I like ]

[ Tell it a bit about myself, exaggerating my credentials towards the specific domain I'm in. I'm a well-known marketer with extensive experience. I'll change the narrative a bit depending on the topic. ]

[ Don't use markdown, write a detailed, extensive response ]

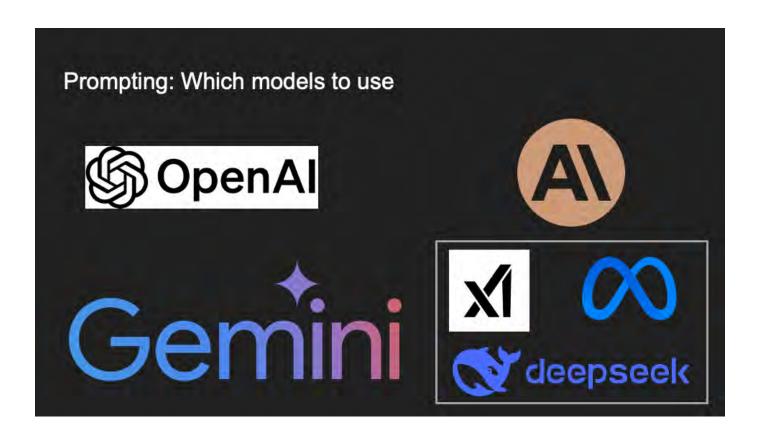
## Learning a new area:

[ Content I like from that area or a blog post/ something conversational I enjoyed reading ]

[I encourage the Ai to ask me questions, or create run through case study examples to test my knowledge]

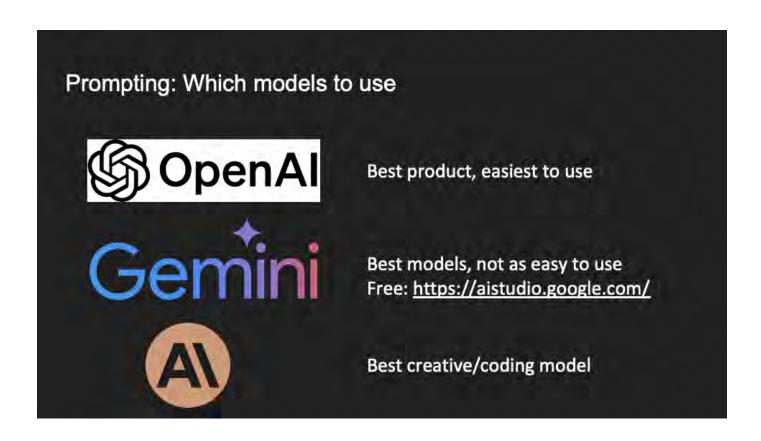
## Completing a task:

[ Typically I'm not spending too long on a task ]



The last aspect of prompting I want to look at, is the model you choose to ask a question to. So I want to give a quick, bird-eye view look at where everything is at.

There's three main players, OpenAl's ChatGPT, Anthropics Claude, and Google's Gemini model. The other labs, xAl, Meta, and DeepSeek are not quite at the frontier.

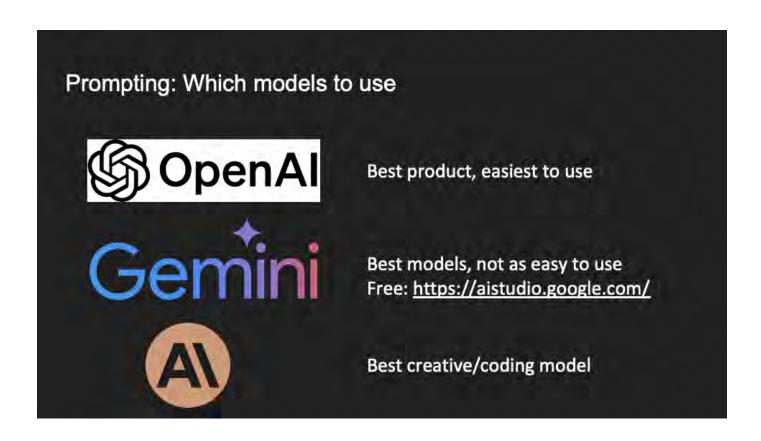


The story between these is that OpenAI has the best product. Its the easiest to use. If you don't use AI that much, and you want to pick a single company, I'd pick OpenAI.

Google has the best models, its new Gemini 2.5 model is excellent. Also, if you don't care about your data being used to train their models, you can use all of their models for free at "aistudio", their video generation model is also better than OpenAl's.

Anthropic is quite as good, except in coding,

where its much more creative and skill at writing high quality code than the other models.

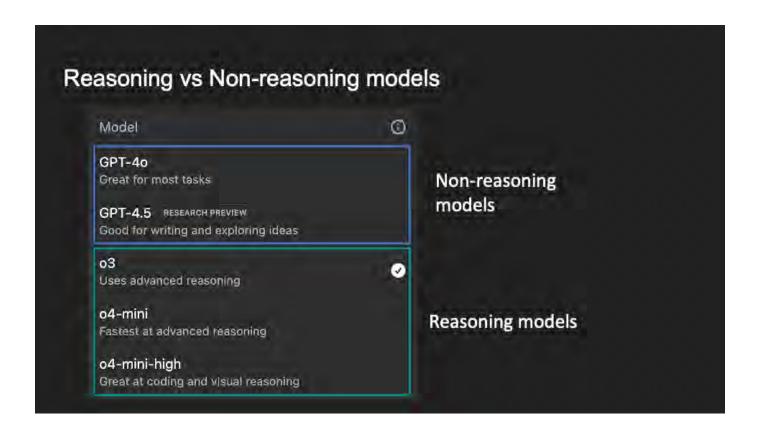


The story between these is that OpenAI has the best product. Its the easiest to use. If you don't use AI that much, and you want to pick a single company, I'd pick OpenAI.

Google has the best models, its new Gemini 2.5 model is excellent. Also, if you don't care about your data being used to train their models, you can use all of their models for free at "aistudio", their video generation model is also better than OpenAl's.

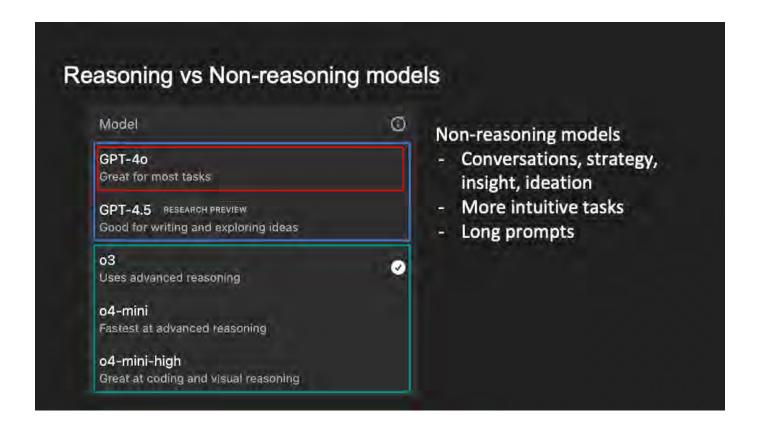
Anthropic is quite as good, except in coding,

where its much more creative and skill at writing high quality code than the other models.



The second piece of choosing a model is choosing between using a reasoning or a non-reasoning model.

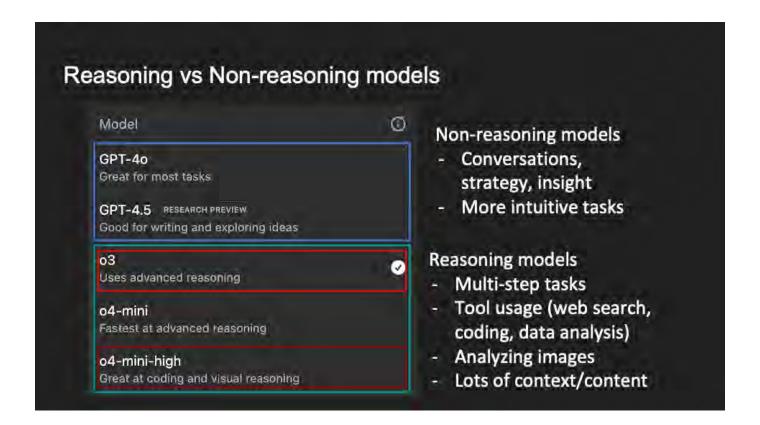
ChatGPT offers these 5 models, the first 2 are non-reasoning models, and other 3 are reasoning models.



I use a non-reasoning model mainly for more conversational use cases. Things like discussing strategy, doing ideation, or areas where I want insights from the AI. It's great a more intuitive tasks. As we saw on Monday, the AI is better at doing difficult tasks like recognizing a location from an image when it doesn't think for too long. This extends to lots of difficult tasks, if instinct helps with the task, then generally a non-reasoning model will be better. I'd also recommend using non-reasoning model in cases where you choose to write detailed prompts. The non-reasoning models are affected by your prompts more strongly, so they are also great to use in cases where you've decided to write a detailed prompt. In terms of models, I'd recommend using GPT-40 in all cases, GPT-4.5 is typically not as good unless you are writing very long and detailed prompts.

[ They are more controllable, but lower default response value. The reasoning model will do more to figure out what you want on

its own ].



I'd recommend using Reasoning models for a few different cases. The first case is when your doing a task that require multiple steps to solve or where you want the AI to be more methodical. For instance, if you want to conduct an analysis of customer feedback its better to use a reasoning model since it will spend more time reading through the document.

The second case is tool usage. When you want the AI to use the internet or analyze an excel document, the reasoning models are far better with tool usage than the default models

The reasoning models are also much better with multi-modal data. For analyzing images, I'd recommend using reasoning models.

Lastly, reasoning models are great also when you have lots of context or content. One of the most common ways I use reasoning models is by just dumping in a bunch of context. Things like my presentations, academic studies, podcasts, and then asking it do some task like helping edit a slide. NotebookLM uses a reasoning model, which is why it's able to do such a good job handling large amounts of context

I'd recommend using o3 for most cases. o4-mini-high is a secondary option if you want more speed. o3 can sometimes take a few minutes to response, whereas o4-mini-high is much faster. I'll now pass it onto Rex to run the next exercise

Exercise 8. Hands-On Exercise: The challenge of Infographics, and a work around

Note, this exercise will be demonstrated live.

### **Use Case Showcase:**

## What if you could talk with your data?

## Day 3: Usecase Showcase - Personas

•

The goal of our use case showcase it to show real-world examples of how to put Al into action. I'm using my experience as Chief Al Officer at Claritas to show some of the practical things that I've built and how I went about building.

The first thing I built when I joined Claritas
was these AI Personas. If you're familiar with
Claritas, they invented segmentation, and it is
updated each year and is called PRIZM
Premier. It divides the United States into 68
different segments, ranked roughly,

- socioeconomically, as well as by urbanicity and life stage.
- These personas have been used by marketers for years to help them understand who their customers are and how to connect with them. For example, this segment from Upper Cross is about 1% of Us. Household. They are older, primarily retired, wealthy, empty nesters, and you can see a bit more about their profiles with bullet points here.
- To convert this data into AI personas, we use something called Retrieval Augmented Generation or RAG. We didn't the AI what to name the personas. We didn't tell it any specific details about the persona narrative. We just told it to develop a narrative based upon that information. As you can see in this case of the Upper Crust segment, the AI gives us Elizabeth Wellington, a 76 year old that enjoys a peaceful retirement with a daily agenda of spending time with her husband on long walks in the neighborhood. Most evenings you'll find them at the local symphony, or exploring new cuisines. They value the environment and supporting green causes, the persona tells us about the type of vehicle she drives, etcetera.
- When you click the button you can have a conversation with Elizabeth. She will answer any questions you ask about products she might want to buy, or the trade off between convenience and

price, and how the way the persona answers is representative of the persona's data. In fact, this is one of the super-powerf of AI, inferring from data.

- To repeat, the power of the Al and large language models is there ability to create inference from that data and connect the dots between the data points to tell a narrative.
- We've licensed this data to agencies. Monks was the 1st one to license it and presented it at the Possible Conference last year. They took the same data, but put a different twist on it.
- In the Claritas example, the personas understand that they're personas. In fact, we created some that are as avatars that you can click on and talk with directly, such as Alex here and Jenna. Alex is from the young digerati, and Jenna is from Struggling Singles... and we'll go ahead and give you a quick example with Jenna.
- · Jenna, tell us about yourself and your persona.
- "Hi! I'm Jenna Thompson.
- I'm rocking my latest thrift shop find a vintage band tee paired with high-rise jeans, tapping into my tech-savvy nature while navigating through my middle-aged single lifestyle always on the lookout to boost my potential and affordability in a metro mix setting."
- Okay, so you get the idea, you can have a conversation with a persona. You can ask any

- questions and what you hear from Jenna is representative of her segment. You will get a different answer from another segment, such as Young Digirati because they have different levels of affluence, and are motivated by different things.
- Monks used a different approach. They created individual personas by name from a specific household in a location that's representative of the data. You can ask about the best place to get Tacos in their neighborhood, and you will get a good answer for that neighborhood based on the preferences and price points you'd expect from the persona.
- Monks accomplished this by taking the same data, but providing different instructions in their persona build. In both cases, you are essentially talking with the Claritas PRIZM data, which I think is pretty remarkable.
- How are personas useful? One of the key applications of Personas is setting the pixel on a website to collect the distribution of site visitors by persona. You can see what type of consumers are using the website, and then you can have a virtual focus group with the top persona segments. This is the heat map showing you in orange the segments that are top 5 segments most likely to engage with this website and the dark gray are the ones least likely. And again, when you click on talking and

having a conversation, what happens in this scenario is that not only is all the persona data there, but we've also trained the AI to read through the website and then develop inference for why that segment would have used that website.

- For example, for Major League Baseball, the Al reads through the website and understand that the persona bought a ticket. When we ask about baseball, or what type of motivations you have for going to a game, the answers are relevant and linked to the underlying data as to whether the motivation is spending time with friends at the game, or the love of the game, etcetera. The answers are inferred from all the data that we've given it.
- And we've given a lot of data. We've given it over 5,000 different data points across 40 different categories. So its answers can span a pretty robust range of topics.
- Now, if you're interested in trying this just to see what that's like to engage with personas, I've made it available for free so you can add a pixel to your website, see who is visiting, and dialogue with the personas. I think it is helpful to have a conversation with Persona in your industry because your appreciation for what Al can do goes way up when it underscores things you suspected, and unlocks new insights that you might not have explored

previously.

- Essentially, talking with your data just gives you a whole new insight for what's possible with Al and RAG.
- Let me share another, related use case. What if
  we need to develop a digital ad campaign for a new
  product and we don't have any data on who buys it
  because the product is new? We can have the
  Persona AI read the website, and tell us which
  segments it things will buy. This is a case where we
  gave the AI this new website, which we have
  debranded.
- We asked questions like, "Which segments should we be trying to reach with advertising?" and "What kind of messages should we include in our advertisements?"
- This was our very 1st test of this case. In this case just used a very simple data set of just the personas and their high level summary. So we didn't even give it the full 5,000 data points. It was more like about 120 data points. And what we could see is the AI recommended that, for this vehicle, segments like Movers and Shakers, Winner Circles, Big Fish in Small Pond were very likely to purchase this vehicle.
- In this case, we had another AI create a visual representation and explain why the segment would want to buy the vehicle. Here is the description: "As

- someone with a dynamic, fast-paced, ambitious lifestyle, and has a penchant for balancing luxury with functionality. I had several reasons for choosing them. It's more than just a car. It's an extension of who I am."
- That little skip there was me taking out the brand name, but the persona was explaining why, she was interested in buying that vehicle what their motivations were, and there's a longer explanation about the key key reasons why it'd be a good fit for her or her persona.
- And as it turned out, that Brand ended up working with Claritas and having the pixel on their website to measure the effectiveness of their digital ad campaign. Claritas does a lot of that work, measuring over 2000 campaigns a month for a wide range of brands. This gave me the chance to check how well that Al performed versus the agency plan. In other works, I could see who actually converted and visited that web page, and how that compared to what the Al recommended, with that very thin amount of data, again, about 120 variables. I could see what the Al inferred in a minute vs. what the ad agency did after spending months and months of work developing a media strategy and plan.
- This is a full list of 68 segments. The yellow highlights were all the ones that the Al

recommended, and and almost all of them were in the top 25%. In fact, it nailed the movers and Shakers, which was the number one segment to visit that vehicle webpage to learn about buying that vehicle. So the Al is performing 68% better than random. Again, that was just with a thin amount of data.

- Now, if you compare that to the agency, the agency did do better vs. the thin set of data. The agency nailed segment one, Movers & Shakers. The Agency got segment 2 right, the Al missed. They both got segment. 3. They both missed segment. 4. Which that young digerati and and then you can see the agency got segment. 5. But missed 6 and 7, which the Al picked correctly, and so on. So pretty high levels of agreement in general. And the difference is that the Al took
- a minute to generate and cost very little, whereas the old way of doing things costs a lot and takes a long time to strategize and plan.
- Now what is even more interesting is what happened in the first few weeks of the campaign. The agency was good, but not perfect. The weighting of where the agency put the impressions vs who actually converted shows there is about 37% room for improvement. It is a little higher for the Al plan, but in both cases, and in fact in all cases, there's always some differences when you

- get real data in the marketplace vs. any plan. If you can adjust targeting based on the data, you can increase conversions and improve ROI.
- So the vision that we have is that you can use AI to read the website to plan segments, and then you can do something we would call "launch and learn,"

   this is where you launch the campaign, get the feedback of who's actually converting and adjust each week to maximize conversions.
- I think we're seeing the direction of a new world here, where these AI personas will play a larger role in helping us in our briefing, in our planning and in our fine tuning, once the campaigns are live. This is faster, less expensive and better performing.
- Yesterday, I shared about how General Motors specifically was using AI with LYRIQ and how they're using the ADAM agent for AI decisioning and the assembly mechanism to put the best performing variant per person, and then using the EVE agent to Evaluate the Variant performance against the Persona segments to make creative recommendations – on how those creative recommendations combined with GM putting new creative variants into rotation, doubled the effectiveness over what was already a double in performance from the first AI. These AI can compound the effectiveness and efficiency of

marketing.

- Today we see another link in that value chain which is adding audience optimization — having an Al agent that plans audiences, and then monitors which audiences are converting and adjust the audience targeting accordingly.
- And there's more links we can connect to this value chain. For example, you could have AI that does resizing. So when you have your creative, teh AI can automatically change the size formats for you much more more quickly.
- You can have an AI that does translation. So if you need to move it into different countries from the United States to Germany, for example, you have an AI do that too.
- There are many different additional agents that can work together in a value chain. And and one of the ones that we presented last year at the Possible Conference was the work from Ramey Kent, the CMO at Progressive, where she used an Al that that Caleb and I developed with Claritas to create audio advertisements.
- You give the AI the brand positioning some of your messaging, your objectives for the campaign, basically the brand brief here and it has three different AI APIs that are used to produce the advertisement. First, the AI builds the scripts using a connection to open AI's Gpt. 4.

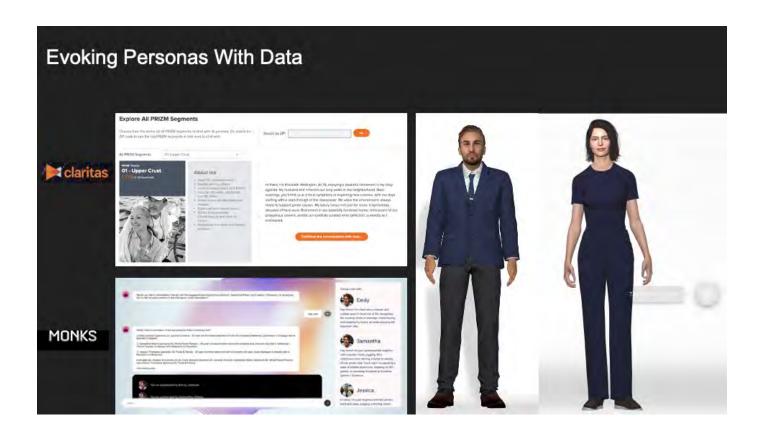
- Once those scripts were approved it then handed off to another AI that uses Eleven Labs, which creates synthetic voices that read the scripts, and then it handed to another one which created the music bed.
- I'll play you some of the ads that they ran and then share the results.
- "Want the perfect auto insurance. It's like crafting. The ultimate playlist. Progressive was the 1st to offer online shopping for insurance, and we continue to innovate for your convenience with over 30 products designed to make insurance a breeze.
- Imagine coverage options that reward your safe driving and fit your budget."
- You get the idea. The Al generated different voices here, and different scripts. And different music soundtracks.
- Ramey shared that normally it would take about 22 weeks to run a 4 week campaign and they would develop 3 ads. And most of those weeks were spent developing the ads, refining it. And here they were able to do the same thing in 6 weeks. That's 4 weeks to run the campaign and now only 2 weeks to develop the ads, get them approved and live. So just a massive savings, and they created a lot more versions, which enabled the ADAM Al to optimize them for different people.
- There were 96 different combinations of ads that

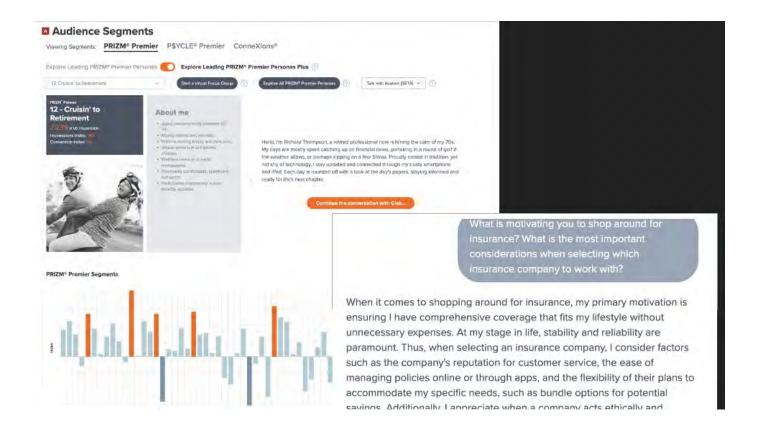
you saw there and overall, the results increased by over 50% because they were using that Claritas ADAM Al agent to match the right creative variant to the people who are most likely to start a quote and complete a quote when they heard the different ad.

- This is a big change when we begin to think about all the ways in which AI can fit into the marketing value chain.
- Let me conclude: I've drafted a document for MMA Global called from briefing to perpetual creativity. It's a draft. And if you would like to receive a copy of it, comment on it, give feedback, we'd welcome that engagement and brainstorming. It's something that I plan on updating periodically, because this industry is moving so quickly.
- Thank you for joining the Use case showcase.
   We're going to shift over to a Q and A with Caleb for those who want to stay longer for the overtime session.
- Tomorrow, we're going to get even more practical and hands on with vector databases, multimodal and take a look at how vision systems work, what you can do with them. Friday, we'll be wrapping up with autonomous AI and workflows and showing you how to actually build things like what we did for Ramey with the AI APIs. There's a lot of things you can do with AI workflows. We can't wait to share

## more with you!

Thank you for joining us.





# Personas tap into PRIZM storyboard as well as deep dive data



#### claritas

### 40 Categories of Data, 5000+ Data Points

Alcohol
Appliances & Equipment
Attitudes
Automobile Insurance
Automotive
Channel Preferences
Claritas Cable & Satellite
Claritas Home Office Tech
Claritas IPA

Claritas Net Worth

Claritas Service Providers
Claritas Wireless Telecomm
Claritas Wireline Telecomm
Commuting and Transportation
Conservation Efforts & Opinions
Credit
Deposits
Digital video displays
Do Not Publish

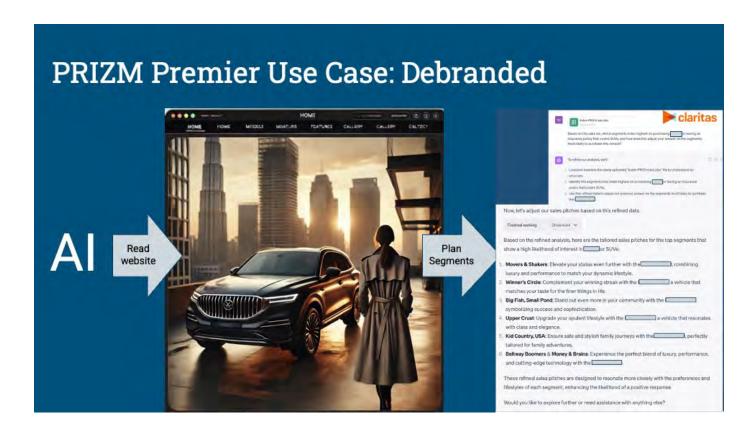
Environment

Financial
Food and Beverages
Health
Household Detailed Demographics
Household Summary Demographics
Insurance Attitudes
Investments and Retirement
Items in the Home
Life Insurance
Other

Other Insurance
Provider Services
Radio
Residential Insurance
Restaurants
Retail shopping
Sports and Leisure
Telecommunications
Travel
Voting



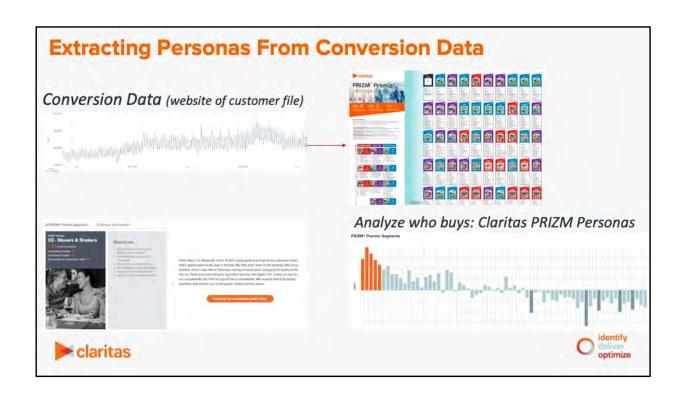
https://www2.claritas.com/AI\_Training\_and\_AA\_a ccess



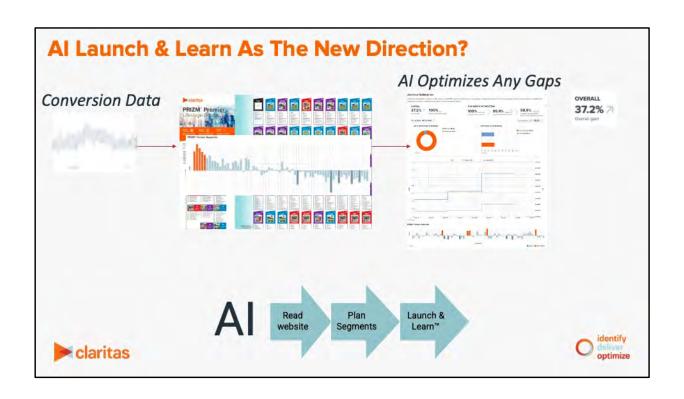
## PRIZM Premier: Mover & Shaker Segment

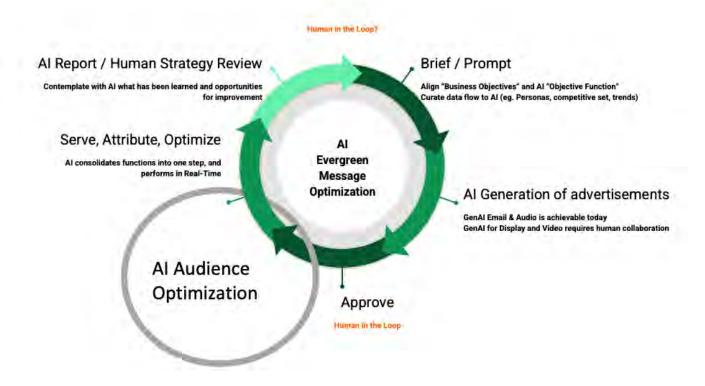












We are quickly moving to a multi-agent value chain.

## GenAl For Display: Resizing









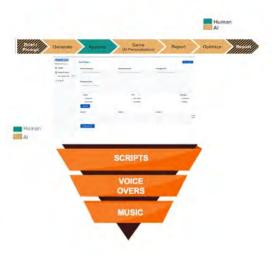
## **GenAl For Translation**











https://cf.artsal.com/campaigns/creatives/proview.html?creativeID=22950&revision=1



https://cf.artsai.com/campaigns/creatives/proview.html?creativeID=22950&rovision=1



