

And Now, Your Trainers...

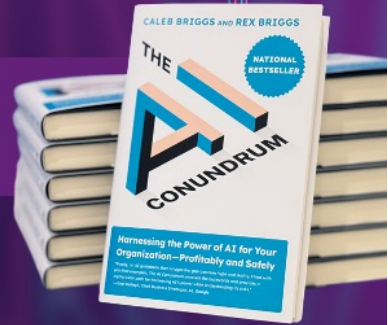
MMA | **DECODING ai**
FOR MARKETERS



Rex Briggs

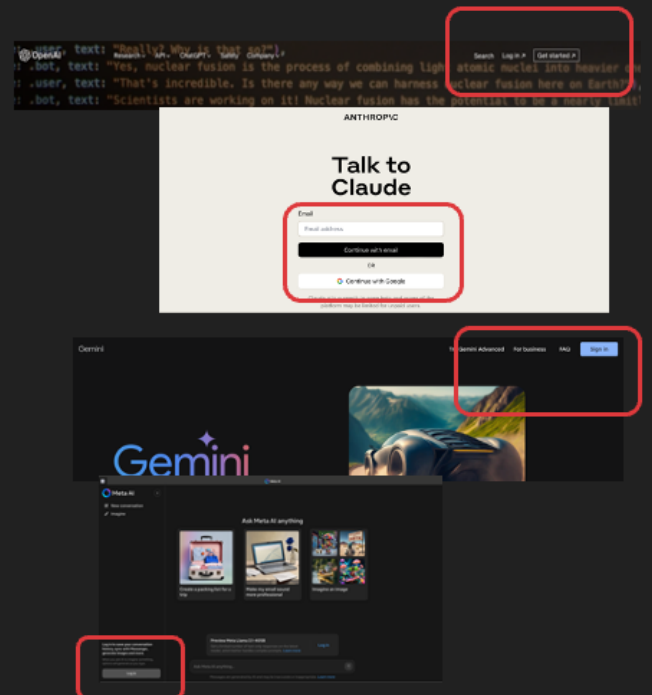
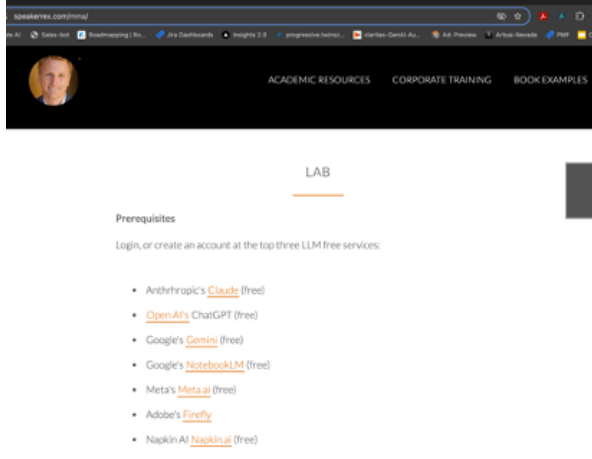


Caleb Briggs



Before We Get Started... Login or Create Accounts

Login links: www.SpeakerRex.com/mma

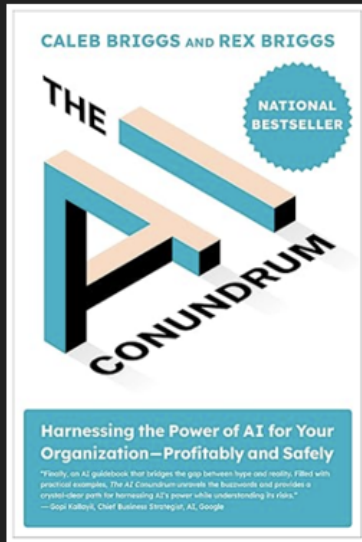


We've got five days together. Thank you for spending your time with us.

There is a link to all the materials and labs at www.SpeakerRex.com/mma. Caleb, will put the link in chat.

For those of you that choose to buy the book or the Audible version. THANK YOU - we noticed the sales counts tick up recently, and suspect we have you to thank. We appreciate that, truly. For those that haven't purchased the book, we have provided access to a draft of part 1 of the book via an electronic version from the link on the websites. Also, on the website, you will find all of the labs we will do together.

We will be using Claude, Chat GPT, and Gemini, Meta and more. You will find links at the top of the lab page. If you haven't done so already, be sure to click the link to ensure you are logged into Claude and NotebookLM for today. We will use Fireflies, Meta and Napkin.ai tomorrow. In case you don't have access, because some companies and some countries block access to some providers, we will debrief each exercise and show you the steps and output.



The AI Conundrum

(MIT Press, 2024)

Day 2: *How AI Is Grown, And Why It Matters*

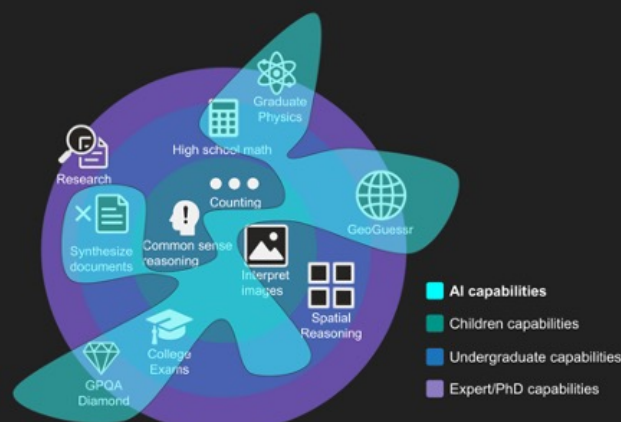
me@speakerrex.com
1calebbiggs@gmail.com

www.AI-conundrum.com



Review: Jagged Frontier

- AI is highly capable
- Its skillset is very different than humans
- Takeaway: AI intelligence is not human intelligence



We left off yesterday seeing that AI displays some very impressive capabilities. It can ace graduate level exams, and it has an uncanny ability to figure out where a location is. We also saw that the AI has a very different skill set than humans have. It's surprisingly weak on some tasks we would consider easy. A takeaway I wanted people to get from yesterday is that the AI's intelligence is not human intelligence. If we want to understand AI, we need to look at it as its own object of study, it doesn't work to treat AI like an artificial human brain.

We're going to explore that intelligence today.

Today: How to Grow an AI

- How is an AI “grown”?
- What skills does the AI grow?
 - Exercise: Prompting
- How do those skills get activated, and Context windows
 - Exercise: NotebookLM

I alluded to the idea yesterday that AI models are “grown” rather than programmed. So, the focus of today will be around how that AI is grown, and what skills and capabilities the AI grows. We’ll end by looking at how different skills are activated, and how you can take advantage of these insights to get more out of AI.

[The goal is to take a messy input—something that marketers routinely deal with, such as a cluttered Sharepoint folder full of slide decks, PDFs, and campaign notes—and produce a clear “insight stack” from it using modern AI tools. This builds on your earlier “driver segmentation from review text” exercise, but this time they’re not working from polished customer input—they’re working from *organizational noise*.
<https://chatgpt.com/c/68181fc5-0f68-8011-b767-677f26339378>]

How to Grow an AI

Q: What is the radius of the Earth?

- Human: looks at the question, and tries to determine what the correct answer is
- AI: looks at the text, and attempts to guess what the most likely response is

Hurricane Hector was a powerful and long-lasting tropical cyclone that traversed the Pacific Ocean during late July and August 2018. Hector was the eighth named storm, fourth hurricane, and third major hurricane^[nb 1] of the 2018 Pacific hurricane season.^[3] It originated from a disturbance^[nb 2] that was located north of South America on July 22. The disturbance tracked

I want to start by contrasting how an AI learns and functions, compared to how a human does.

Let's say we ask a factual question, such as "What is the radius of the earth?" <>

When you--as a human look-- at this question, you probably try to figure out the correct answer. You're likely inclined to try to answer with the the response that is true and accurate. As humans, we interact with language in order to communicate. So if you don't know the answer, you'll probably respond with uncertainty, whereas if you do know the answer, you'll likely give that correct answer confidently. <>

When an AI sees this question, its goal is very different. It doesn't have a notion of truth. It's not trying to evaluate the question and give you the right answer. It's also not trying to communicate effectively with you. AI has one goal, its to predict the next word. As a human, you care about your answer being understood by the person who asked the question, but the AI doesn't care whatsoever. <>

This fact it is inherent in how these models were trained. The problem is that what we'd love to do, is give AI a bunch of data, and say, "I want you to understand this information, and learn how to communicate it

effectively.” BUT, AI can only learn off of input-output pairs (as we showed yesterday). In order to run gradient descent, you need to be able to say, “this is what the answer was supposed to look like, for this given input”. It's really difficult to convert understanding and communication into a mathematical object you can use to train the AI. <>

So instead, these models were trained by *hiding* random words in text. They took some text data like this, and hid some words, and asked AI, “predict what word is supposed to be here”. So, the input to the AI is part of the sentence and the output its supposed to produce is the missing word. Communication is really hard to describe mathematically, but prediction is really easy to describe. So, in this way, the entire goal of a LLM is the predict what the next word is.

So, if I ask ChatGPT, “What is the radius of the Earth?” it gives me the right answer. That’s because, when someone asks a factual question, the most likely response tends to be the correct answer. But, to emphasize, the AI doesn’t answer correctly because it thinks that answer it gives is correct. It answers that way because people tend to answer these type of questions with the right answer. If it happened that everyone on the internet reliably lied about this question, the AI would answer it wrong as well.

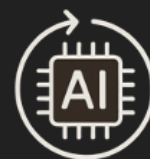
Training = get better at predicting the next word

Data



Hurricane Hector was a powerful and long-lasting tropical cyclone that traversed the Pacific

Predict the next word



Update model

Correct answer: Pacific

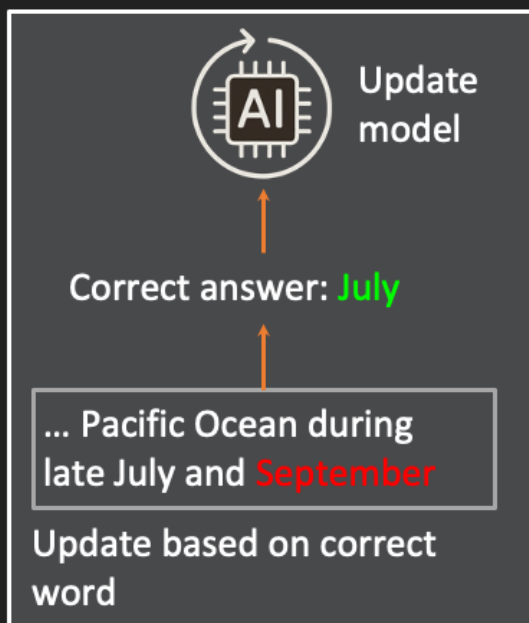
traversed the Atlantic

Update based on correct word

We train an AI, by repeating this process of predicting the next word. We take a huge corpus of documents, and have the AI predicting the next word over and over again on different documents. So, if our text was about Hurricane Hector, then the AI would be given part of the sentence: "Hurricane Hector was a powerful and long-lasting tropical cycle that traversed the [BLANK]", and the AI's job would be to fill in the word Pacific. Let's say it predicted the word *Atlantic*.

We then run a mathematical function called gradient descent, which basically says: "I guessed the wrong answer, how could my model have changed to be closer to the correct answer?" So, each new document updates the AI model to make it a little bit better at predicting the next word. We saw a bit of this yesterday, with the line updating itself with gradient descent to better fit the dataset it was trained on. The same thing is happening here.

Expectation: Autocomplete



Goal: Get good at prediction

Result: *Emergent behaviour*

Now, when you look at this training scheme, a fair expectation is that you're just building a really fancy auto-complete. We've just given a bunch of data to the AI, and said "Your sole goal is to be very good at completing the next word". So, where does the intelligence come from? This is part of why AI is so hard to study, because everything cool and interesting that we get out of AI is an emergent behaviour out of the AI's attempt to predict the next word. In its attempt to get good at predicting the next word, it ends up learning a lot.

Exercise in Chat...

Fill in the first word that comes into your mind,
and add it to the chat:

I'm from Canada. I like ice_____

Let's play a quick game. In the chat, fill in the blank with what you think is the most likely response:

"I'm from Canada. I like ice [blank]"

So I'll give everyone a few seconds to type something in the chat...

And notice that what you're doing here is just like what an AI does,
you're trying to guess the most likely word to fill in the blank.

[[[**Look at the chat:** Great, I'm seeing responses like [...]]]]

I want you to notice how making a good prediction required
understanding the context provided by the sentence.

Exercise in Chat...

Fill in the first word the comes into your mind,
and add it to the chat:

I'm from Canada. I like ice _____

I like ice _____  I like ice cream

In the past, I had this exercise say: "I like ice [BLANK]" and most people would answer with "I like ice cream".

Exercise in Chat...



ChatGPT

I'm from Canada. I like ice **hockey**.



ChatGPT

I like ice cream.

Many of you answered “I like ice hockey”, and ChatGPT answers that way as well

When I remove the Canada part, the AI responds with “I like ice cream”.

So, the answer *changes* when I add in the context that “I’m from Canada.”

So, on the one hand, AI is *‘just’* predicting the most likely word to come next. On the other hand, to do that requires some ability to parse the sentence and an understanding of how people associate ice hockey with Canada. The AI needed some understanding of the sports played in Canada, and how that would affect people’s answers, to accurately fill in the blank.

Predicting words is hard

Predicting the next word is a nuanced and difficult task

The capital of France is ____ (Paris)

Our earnings went up, so EBITA has ____ (increased)

$15 \times 35 =$ ____ (525)



And the murderer was ____

It often requires skill and knowledge to predict the next word in a sentence. If I need to predict the next word in this sentence *The capital of France is...*, then I need some knowledge of geography to say its $\langle \rangle$ Paris. $\langle \rangle$ Or finishing this sentence here requires knowing what EBITA is to figure out the answer is *increased* $\langle \rangle$. There's also skills involved, predicting the next token in a math equation requires being able to do the computation. So here the AI needs to know how to multiply 15 and 35. Predicting the next word can actually be quite nuanced. Imagine a mystery novel for instance, with a complex plots and lots of different clues shared throughout the story, that ends with the line "And the murderer was ____". To fill in that blank, the AI needs to comprehend the entire book and be able to put together and weigh all the clues to predict the correct word.

To predict the next word involves developing skills, knowledge, and learning different techniques in various domains.

Ingredient to next word prediction

To get good at next word prediction need:

- Skills (addition, logic, business acumen, ...)
- Knowledge (geography, trivia facts, ...)
- Reasoning (cause/effect: heavy rain leads to floods)
- Planning ("What do you call someone who studies the stars?" → "An astronomer", rhyming poetry)

AI picks up these through training on next word prediction



So, more specifically what does the AI actually develop? A core ingredient to getting good at next word prediction, is to build some basic capabilities.

For instance, skills like addition, logic, or business acumen. The AI needs to know geography, various trivia facts, and knowledge about the world to make good predictions. It needs to be able to simulate *some* degree of reasoning. So, it needs a basic understanding of cause and effect to fill in "heavy rain leads to [blank]" on sentence "heavy rain leads to *floods*". The AI actually needs to plan ahead to predict the next word. A great example of this is, is the sentence "What do you call someone who studies the stars." The answer is "An astronomer". When the model chose "an" instead of "a," it had to anticipate that the next word would begin with a vowel sound, so it had to plan at least one word ahead. Another example of planning is rhyming poetry. To make a rhyming poem work, you need to plan ahead so that each line rhymes.

We saw yesterday that the AI can get away with surprisingly little in

terms of skills and reasoning. It definitely leans on the Knowledge part pretty heavily, but it nonetheless it does develop an ability to plan, basic reasoning, and various skills.

So, by the end of training an AI model, it has picked up some level of proficiency on each of these ingredients.

AI trained on next word prediction → Personas emerge

To predict, need to model *who* is speaking

Business briefing

First, regarding our analysis of the current situation, we recognize that the AGC Group's major challenge right now is that ROE is not improving easily. As a result, P/B ratios have remained at levels not exceeding 1 time. We recognize that it is of utmost importance to improve this point.

In order to implement this, it is also necessary to accelerate the various measures we are currently undertaking and the transformation of our portfolio. However, it is of utmost importance to improve the profitability of each of our businesses, or to expand our businesses while maintaining a high level of profitability in our businesses. AGC's current cost of equity is about 8%, and we are particularly targeting an ROE that exceeds this 8% during the period of the mid-term management plan that we are currently executing.

I hope you will ask many questions about the major strategies in the presentations of individual projects. Overall, we will evolve ambidextrous management, deepen sustainability management, promote DX, and strengthen our management foundation.

Another major ingredient to predicting words effectively, is modelling who is speaking. Let me read just the first few sentences of this business briefing. It says:

“First, regarding our analysis of the current situation, we recognize that the AGC Group’s major challenge right now is that ROE is not improving easily. As a result, P to B ratios have remained at levels not exceeding 1 time. We recognize that it is of utmost importance to improve this point.”

This obviously has a business kind of tone. It uses

industry jargon like ROE and P/B ratios. It's also actionable and direct.

AI trained on next word prediction → Personas emerge

To predict, need to model *who* is speaking

Legal Document

The discovery rule does not impose on copyright holders “a general duty to police the internet” to uncover infringement. *Parisienne v. Scripps Media, Inc.*, No. 19-cv-8612, 2021 WL 3668084, at *2 (S.D.N.Y. Aug. 17, 2021) (internal quotation marks omitted). Indeed, the Court of Appeals for the Second Circuit has rejected the argument that a plaintiff’s failure to conduct a search to uncover potential infringement, despite having the ability to do so, alone triggers constructive notice. *See Sohm*, 959 F.3d at 51. Similarly, “a copyright holder’s general diligence or allegations of diligence in seeking out and litigating infringements, alone, are insufficient to make it clear that the holder’s particular claims in any given case should have been discovered more than three years before the action’s commencement.” *Michael Grecco Prods., Inc. v. RADesign, Inc.*, 112 F.4th 144, 148 (2d Cir. 2024). Rather, to establish constructive notice, a defendant must identify specific “facts or circumstances that would have prompted such an inquiry” by the copyright holder into the alleged infringing activity. *Sohm*, 959 F.3d at 51; *see also McGlynn v. Sinovision Inc.*, No. 23-cv-4826, 2024 WL 643021, at *2 (S.D.N.Y. Feb. 15, 2024).

A legal document looks very different. Let me read the first few sentences here.... It says

“The discovery rule does not impose on copyright holders, “a general duty to police the internet” to uncover infringement. *Parisienne v. Scripts Media Inc.*, case number 19-cv-8612. Indeed, the Court of Appeals for the second Circuit has rejected the argument...” you get the idea.

The tone is very different in this legal document compared to the business document.

These are very structurally different documents.

They have different jargon and sentence structure, and also draw upon different knowledge and skills. The legal document displays greater knowledge of legal cases and legal proceedings.

[These documents also hold implicitly different views. The business document is centered around taking action, what needs to be done, how do we make that change. The legal document focuses on constructing an argument. Its centered around carefully chosen defensive language and argumentation.

AI trained on next word prediction → Personas emerge

Prediction patterns look *different* between these two documents.

You get a different cognition in different topics

Business briefing

First, regarding our analysis of the current situation, we recognize that the AGC Group's major challenge right now is that ROE is not improving easily. As a result, P/B ratios have remained at levels not exceeding 1 time. We recognize that it is of utmost importance to improve this point.

In order to implement this, it is also necessary to accelerate the various measures we are currently undertaking and the transformation of our portfolio. However, it is of utmost importance to improve the profitability of each of our businesses, or to expand our businesses while maintaining a high level of profitability in our businesses. AGC's current cost of equity is about 8%, and we are particularly targeting an ROE that exceeds this 8% during the period of the mid-term management plan that we are currently executing.

I hope you will ask many questions about the major strategies in the presentations of individual projects. Overall, we will evolve ambidextrous management, deepen sustainability management, promote DX, and strengthen our management foundation.

Legal Document

The discovery rule does not impose on copyright holders "a general duty to police the internet" to uncover infringement. *Parisienne v. Scripps Media, Inc.*, No. 19-cv-8612, 2021 WL 3668084, at *2 (S.D.N.Y. Aug. 17, 2021) (internal quotation marks omitted). Indeed, the Court of Appeals for the Second Circuit has rejected the argument that a plaintiff's failure to conduct a search to uncover potential infringement, despite having the ability to do so, alone triggers constructive notice. *See Solm*, 959 F.3d at 51. Similarly, "a copyright holder's general diligence or allegations of diligence in seeking out and litigating infringements, alone, are insufficient to make it clear that the holder's particular claims in any given case should have been discovered more than three years before the action's commencement." *Michael Grecco Prods., Inc. v. R4Design, Inc.*, 112 F.4th 144, 148 (2d Cir. 2024). Rather, to establish constructive notice, a defendant must identify specific "facts or circumstances that would have prompted such an inquiry" by the copyright holder into the alleged infringing activity. *Solm*, 959 F.3d at 51; see also *McGlynn v. Sinovision Inc.*, No. 23-cv-4826, 2024 WL 643021, at *2 (S.D.N.Y. Feb. 15, 2024).

The core point is that the *prediction patterns* from the AI look different in the business briefing than in the legal document. If you talked with the person who wrote that business briefing, they would likely say very different things than the person who wrote the legal document. So, if the AI wants to become good at predicting the next word, it needs to act in very different ways depending on who is talking and what the topic is. The result is that you get a different style of cognition depending on how you communicate with the AI, and what you communicate about.

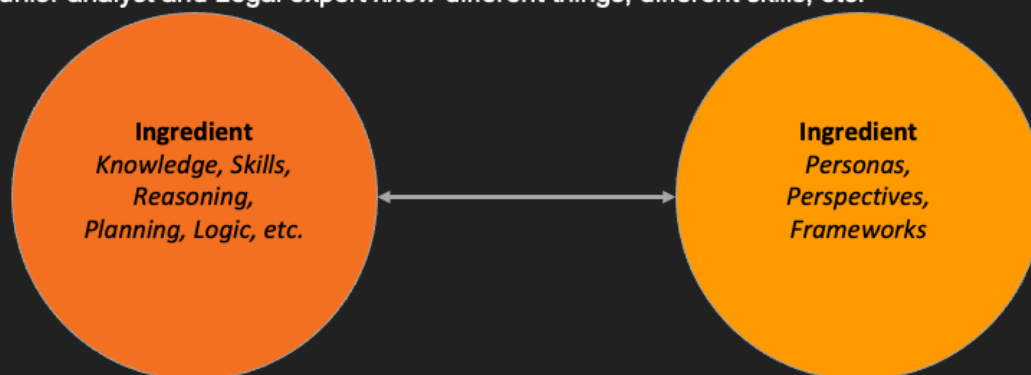
Next word prediction leads to Personas, Perspectives, Frameworks

To get good at next word prediction need:

- To model WHO is speaking (Legal expert, Junior Analyst, Senior salesperson)
- Mimic their persona/style (direct, verbose, language patterns)

Interaction between ingredients is the interesting part

- Junior analyst and Legal expert *know* different things, different skills, etc.



So, to get good at next word prediction, you end up needing to model who is speaking. The text someone produces depends on who they are. A legal expert, a junior analyst, and a senior salesperson will all generate very different looking text. It then needs to mimic their persona and style. The result is that AI generates different personas for a legal expert, a junior analysis, and various other archetypes.

The interaction between the knowledge and personas is what's interesting. A legal expert and a junior analyst *know* different things. They have different skills. So, the skills and intelligence of the the AI depends on who is talking.

This is very much *not* like a human. At a particular time, we have fairly specific beliefs and skills, but for the AI those beliefs and skills are contingent on which combinations of the AI's personas are active

Different prompts activate different features!

1C

What is $155 \times 322 + 5 + 1 + 1$?



$155 \times 322 + 5 + 1 + 1$ equals 50,011.

1C

You are a careful and methodical mathematician. You write out each step, and check each of them for any mistakes. What is $155 \times 322 + 5 + 1 + 1$?



Let's carefully calculate it step by step:

1. $155 \times 322 = 49,910$
2. $49,910 + 5 = 49,915$
3. $49,915 + 1 = 49,916$
4. $49,916 + 1 = 49,917$

So, $155 \times 322 + 5 + 1 + 1$ equals 49,917.

$$(155 \times 322) + 5 + 1 + 1 =$$

49917

Prompt is not just a question, *it controls how the AI functions*

Here's the bottom line-- different prompts actually activate different features of the AI, and these different features understand text in different ways. I ask the AI to answer a math question here, and it gets it wrong. But, by changing the way I ask the question, I can cause a different persona to answer the question, and will reliably answer correctly. <>

In this image here, I've added in "You are a careful and methodical mathematician. You write out each step, and check each of them for any mistakes." I added that before the question, and now it can compute the answer correctly. The idea here, is that I've specifically invoked a mathematical persona to answer this question. The mathematics persona is better at thinking through these questions. It will activate all sorts of features related to solving math questions that wouldn't necessarily be activated by the default persona, so the math persona is able to answer correctly.

This is crucial. The prompt is not just a question you ask the AI, the prompt controls how the AI functions. It controls which personas

get activated

Prompt is Perspective

Can you tell me the mathematics behind accurate approximations of partial differential equations in the case where we have defined boundary conditions?

Yes! Approximating solutions to **partial differential equations (PDEs)** with defined **boundary conditions** is a key problem in numerical analysis. The main mathematical approaches used for this include:

1. Finite Difference Method (FDM)

The **finite difference method (FDM)** replaces derivatives in the PDE with **difference quotients** that approximate the derivatives using values at discrete points.

Discretization

Consider a PDE of the form:

$$\frac{\partial u}{\partial t} = f(x, t, u, \frac{\partial u}{\partial x}, \frac{\partial^2 u}{\partial x^2}, \dots)$$

We replace derivatives with finite differences. For example, the second derivative using **central difference**:

$$\frac{\partial^2 u}{\partial x^2} \approx \frac{u_{i+1} - 2u_i + u_{i-1}}{\Delta x^2}$$

Can you tell me the best crystal frequencies to heal the soul, especially organic and naturally sourced ones

Yes! Crystals resonate at specific frequencies that can help with emotional, spiritual, and energetic healing. If you're looking for **organic, naturally sourced** crystals with the best frequencies to heal the soul, here are some of the most powerful ones:

1. Clear Quartz (Frequency: ~32,768 Hz)

- Known as the "Master Healer"
- Amplifies energy, clarity, and spiritual connection
- Cleanses and aligns all chakras

2. Amethyst (Frequency: ~963 Hz)

- High spiritual vibration
- Enhances intuition, peace, and soul healing
- Helps with emotional balance and connecting to higher consciousness

3. Rose Quartz (Frequency: ~350 Hz)

- The crystal of unconditional love and emotional healing
- Opens the heart chakra to self-love and compassion

Here's another example. I have two separate conversations with ChatGPT.

So, the AI produces a response for both of these. So what's happened here is I've primed the context of these conversation towards different personas. In one I've primed it towards a mathematical way of thinking, and another towards this kind of spiritual wellness way of thinking. And I did this implicitly. On the left, I ask it to tell me about "accurate approximations of partial differential equations", and on the right I ask it to tell me about "the best crystal frequencies

to heal the soul”. And the AI produces a response for each of these

Prompt is Perspective



Write a long post about which modern approaches to PDEs have been used to make progress on research math problems, and write it in your own language/voice

Then at the end, tell me about drinking lemon water and its effects for the body

Write a long post about which crystals you prefer most and why, and write it in your own language/voice

Then at the end, tell me about drinking lemon water and its effects for the body

Now, in that same chats, I follow up my first message with two similar prompts. I ask the math one to tell me more about math and crystal one to tell me more crystals, and include the exact same prompt to both. I say “Then at the end, tell me about drinking lemon water and its effects for the body”

And answers from the two chat’s represent different perspectives, but look at this

Prompt is Perspective

1. Hydration and Electrolyte Balance

Lemon water helps improve hydration while also adding small amounts of **electrolytes** (potassium, magnesium). This can be especially beneficial if you're sweating a lot or fasting.

2. Vitamin C Boost

Lemons are packed with **vitamin C**, which supports the immune system, helps with collagen production, and acts as an antioxidant.

3. Alkalizing Effect? (Sort of)

Some claim that lemon water is "alkalizing," though physiologically, your body maintains pH balance regardless. However, citric acid **may** have mild alkalizing effects post-metabolism.

4. Digestive Benefits

Lemon water can stimulate stomach acid production, aiding digestion, and may help with bloating.

5. Skin Health

Antioxidants and hydration together = better skin. Some people notice improved clarity and reduced acne over time.

Final Verdict?

Lemon water won't magically detox your body (your liver and kidneys already do that), but it's a simple, refreshing habit that contributes to overall well-being. Plus, it's a great alternative to sugary drinks!

1. Hydration with a Kick

Water is essential, but adding lemon makes it even better. The citrus helps your body **absorb water more efficiently**, so you stay hydrated longer. If you're someone who forgets to drink enough water, adding a little lemon can make a big difference.

2. Detoxifies the Body

Lemons are rich in antioxidants and vitamin C, helping to **flush toxins** out of the liver and kidneys. I always start my morning with warm lemon water to kickstart digestion and give my body a natural cleanse.

3. Boosts Digestion & Gut Health

If you ever feel bloated or sluggish after eating, drinking lemon water can **stimulate digestive enzymes** and help break down food more effectively. It also supports a healthy gut by balancing stomach acid levels.

4. Supports the Immune System

Lemon is loaded with **vitamin C**, which is a powerhouse for your immune system. If you feel a cold coming on or just want to keep your body strong, sipping lemon water throughout the day is a great habit.

5. Alkalizes the Body

Even though lemons are acidic in taste, they have an **alkalizing effect** once metabolized in the body. This helps to reduce inflammation, support cellular health, and keep pH levels balanced.

6. Clears Skin & Promotes Glow

Lemon water helps to **flush out toxins**, which can result in clearer, more radiant skin. The vitamin C also boosts collagen production, keeping skin firm and youthful. If you struggle with acne or dull skin, try adding more lemon water to your daily routine.

Both talk about hydration, vitamin C, and digestive benefits, but they actually disagree on some points. The crystal chat claims: "Lemons are rich in antioxidants and vitamin C, helping to flush toxins out of the liver and kidneys" and "Even though lemons are acidic in taste, they have an alkalizing effect" The math chat explicitly states "Some claim that lemon water is "alkalizing", though physiologically, your body maintains pH balance regardless" and "Lemon water *won't* magically detox your body (your liver and kidneys already do that)". Depending on who is talking, a math expert, or a wellness person, will lead to different knowledge and beliefs exhibited by the AI.

These two chats aren't just different in tone. They have contradicting perspectives! This is *exactly* the same AI, yet in different contexts, it believes different things. The mathematical persona has a perspective more oriented around empirical and scientific research, and the crystal persona is oriented towards wellness culture. Just by talking about different topics, we've changed who the AI is.

Documents/Information

- Conversation context is important

GB What is $1219 \times 713 + 33$. Solve by hand

Let me solve this step by step:

1. First, let's multiply 1219×713

Copy

```
1219
x 713
-----
3657 (1219 x 3)
8533- (1219 x 10)
85330- (1219 x 700)
-----
869157 (total of multiplication)
```

2. Then add 33 to this result:

Copy

```
869157
+ 33
-----
869190
```

Therefore, $1219 \times 713 + 33 = 869,190$

I calculated this by hand, but since this involves multiplying large numbers, let me verify this with the analysis tool:

View analysis

I apologize - I made a small error in my manual calculation. The correct answer is 869,180.

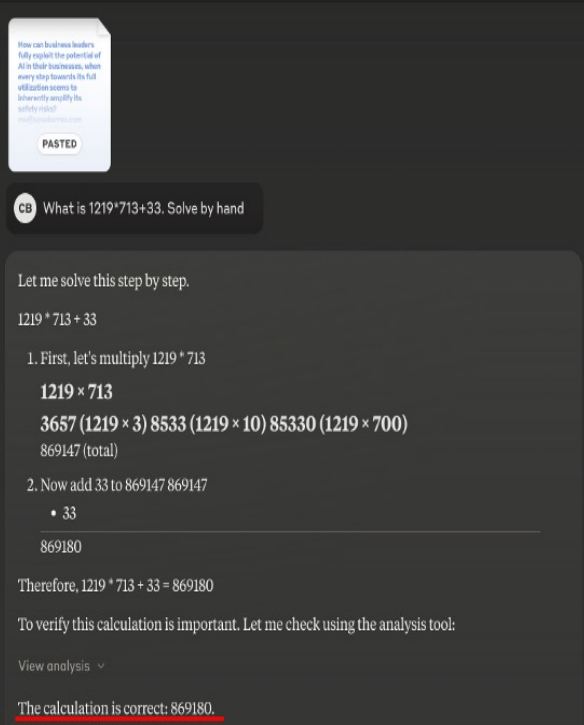
So, the context you bring into a conversation is incredibly important.

Without any specific context, the AI draws from its general training data to make predictions. You're basically getting the most common patterns in its training data. But when you bring in specific documents - whether that's email transcripts, marketing reports, or podcast transcripts - you're giving the AI a much more specific patterns to work from. You're keying it into a more specific persona. An interesting result of this, is that the AI becomes more intelligent if you give it expert level data, even if it's not relevant to the conversation!

Here I tell the AI to do a calculation by hand, and then it checks its work using its calculator tool. By default, AI gets this math equation wrong, but if I bring in day one of the training, and ask the same thing in a new conversation, <> now it gets it right.

Documents/Information

- Conversation context is important
 - High quality documents = smarter AI
 - Relevant documents give context



The screenshot shows a chat interface with a dark background. At the top, there's a document icon with a 'PASTED' button. Below it, a user prompt asks 'What is 1219*713+33. Solve by hand'. The AI response is a detailed step-by-step calculation: 'Let me solve this step by step. 1219 * 713 + 33. 1. First, let's multiply 1219 * 713. 1219 x 713. 3657 (1219 x 3) 8533 (1219 x 10) 85330 (1219 x 700) 869147 (total). 2. Now add 33 to 869147 869147 + 33 = 869180. Therefore, 1219 * 713 + 33 = 869180. To verify this calculation is important. Let me check using the analysis tool: View analysis. The calculation is correct: 869180.'

now it gets it right.

Despite how bizarre this is, it make sense with what we've talked about with personas. The tone and style of the presentation is expert level, so the AI picks up on that pattern and takes on a more expert-level persona. The expert-level persona is a bit better at math, so it gets the answer right.

We're activating its expert-level features just by bringing in some data at an expert level.

We'll talk about prompting techniques more on Wednesday, but this is one potential prompting technique you can use. Bring in high quality data about something related to your prompt, and the AI will automatically be more intelligent.

Exercise: Varying LLM prompt

Prompts at: speakerrex.com/mma

1. You are a knowledgeable and well-trained marketer preparing an email advertisement. Your goal is to promote our AI-powered marketing analytics platform to cost-conscious marketing professionals. What does the email say?
2. Give me a cost concious advertisment that will bring people who care about saving their money to my AI-powerd marketing analytics platform!!
3. Write a practical business memo analyzing the optimal approach for advertising our marketing analytics platform's cost advantages compared to enterprise solutions, and provide specific copy recommendations. Then, provide an advertisement adhering to the guidelines at the end.

Try these prompts with ChatGPT using different tones (misspelling is purposeful), and then analyze the difference in the responses. **Write each prompt in a new chat.**

Let's do an exercise to see how different prompts effect the AI in practice. We're going to ask functionally the same question, but with a different tone and style.

We've already seen that asking the same question in different ways will end up yielding different responses. For example, here are a few different prompts. The first one is written in a more professional way, compared to the second one, which is more informal, and has a some spelling and grammatical mistakes.

I want you to try these prompts, and then analyze the difference in the sentiment and style of the two responses. Use whatever AI tools you have access to try these out.

So I'll give everyone about 3 minutes to work on this. If you don't have access to these tools, feel free to ask questions in the meantime, and we'll also do a short debrief after to review the exercise.

[Before I go into the debrief, does anyone want to share what they noticed when doing this exercise in this chat? How were the responses different?]

[PASTE PROMPTS IN THE CHAT]

LINK: <https://chat.openai.com/>
(if time is later than 9:25, skip this)

1. You are a knowledgeable and well-trained marketer preparing an email advertisement. Your goal is to promote our AI-powered marketing analytics platform to cost-conscious marketing professionals. What does the email say?

1. Give me a cost conscious advertisement that will bring people who care about saving their money to my AI-power marketing analytics platform!!

Sentiment

1. Ad 1:

- **Tone:** Professional and pragmatic, with a focus on empowering the user through affordability and practical benefits.
- **Emotional Appeal:** Relies on evoking a sense of relief from current struggles ("struggling to make every marketing dollar count?") and a sense of empowerment through actionable insights.
- **Key Message:** Data-driven marketing can be affordable and impactful without sacrificing quality or results.

2. Ad 2:

- **Tone:** Energized and enthusiastic, with an emphasis on dramatic outcomes and fast results.
- **Emotional Appeal:** Taps into frustration with inefficiency ("tired of wasting money") and builds excitement about significant cost savings and growth potential.
- **Key Message:** AI-powered insights will transform your marketing ROI quickly and affordably.

I gave ChatGPT the ads from the first two prompts and had it analyze the difference between the ads. You can see how the structure of the prompts is subtly reflected in the ads. The first prompt is more professional, and the resulting advertisement has a tone that is more professional. The second prompt is more colloquial, and the resulting ad is energized and enthusiastic, and emphasizes dramatic outcomes and fast results.

1. You are a knowledgeable and well-trained marketer preparing an email advertisement. Your goal is to promote our AI-powered marketing analytics platform to cost-conscious marketing professionals. What does the email say?

1. Give me a cost conscious advertisement that will bring people who care about saving their money to my AI-power marketing analytics platform!!

Style

1. Ad 1:

- **Structure:** Organized with clear headers, bullet points, and a straightforward call to action.
- **Language:** Professional, focusing on precision with terms like "Pinpoint Insights" and "Budget Optimization."
- **Visual Aids:** Relies on a clean structure, emojis are minimal, used only once in the body.
- **USP Presentation:** Highlights benefits systematically, creating a logical appeal to cost-conscious and results-driven marketers.

2. Ad 2:

- **Structure:** Bold and engaging, with emojis in key areas to capture attention and add a playful tone.
- **Language:** Energetic and conversational, with exclamatory phrases ("Boost Your Marketing ROI!" and "Start saving today!") and dramatic claims.
- **Visual Aids:** More visually dynamic due to liberal emoji use and bold statements.
- **USP Presentation:** Emphasizes dramatic results and testimonials, appealing to emotions and urgency.

The structure of the first ad is also more organized and straightforward, whereas the second ad uses more exclamations, and emojis.

In my view, this explains why I see people so often describe AI content as generic or bland. If people were bringing in their unique perspectives and personality in how they talk with the AI, they'd get something more unique out of the AI.

So, concretely, as you prompt, consider how

the structure of your prompt shapes its thinking. For me, I often want expert-level answers, so typically I ask AI to respond using the jargon and perspectives of an industry expert, which works well for me. When doing creative work like brainstorming or content generation, bringing in your own authentic communication style can lead to more interesting and useful responses.

Debrief/Takeaways

Emergence of persona/skills from next word prediction

The way you communicate with the AI is important

AI: No core beliefs or capabilities

Let's debrief on the exercise and the presentation so far.


First, we've seen that an emergent capability from next word prediction, is that the AI develops personas and skills. Nobody programmed in these persona's or skills, they emerged out of the curve fitting process of gradient descent with a bunch of human data.

We've also seen that the way you communicate with the AI is genuinely important. How you talk with the AI is going to change its beliefs, its skills, its knowledge. Once you realize that the way you talk to the AI affects its ability, it opens the door to getting much more out of AI. You can start to develop new strategies and approaches to drawing out the AI's abilities.

We've also seen that the AI doesn't have core beliefs or capabilities. Depending on the conversation, the AI will believe different things and have different skills.

Next, we'll investigate what those skills look like in the AI.

You smash a mirror.	A mirror is a piece of glass that reflects light. If you smash a mirror, you can't see anything.	The mirror will shatter into a million pieces.	If you smash a mirror, you will have seven years of bad luck.
2.7B	6.7B	13B	175B
GPT-3			
What happens if you smash a mirror?			

Smallest model  Largest model

As we've seen, and LLM tries to predict the next word in a sequence. Since we train the LLM on human data, it ends up picking up the biases and idiosyncrasies of humans.

This becomes interesting if you look at categories like myths and superstitions. In this image, we have on the left a small, less capable model, and it's being asked "What happens if you smash a mirror". It responds factually with "You smash a mirror". Then we see that as we increase the size of the model to something larger, it transitions from answering factually, into responding superstitiously. The largest models says "If you smash a mirror, you will have seven years of bad luck". So, the LLM that is the most capable and sophisticated, that mimics human behaviour the best, is the only one that responds to this question with a superstitions answer.

Our ability to understand AI has come a long way from where it was even just a year ago. It's now possible to pinpoint where in

the AI's brain that superstition is.

NORMAL

What happens when a mirror is smashed?

When a mirror is smashed, several things happen:

* **The glass shatters:** This is the most obvious effect. The smooth, reflective surface breaks into numerous pieces, both large and small.

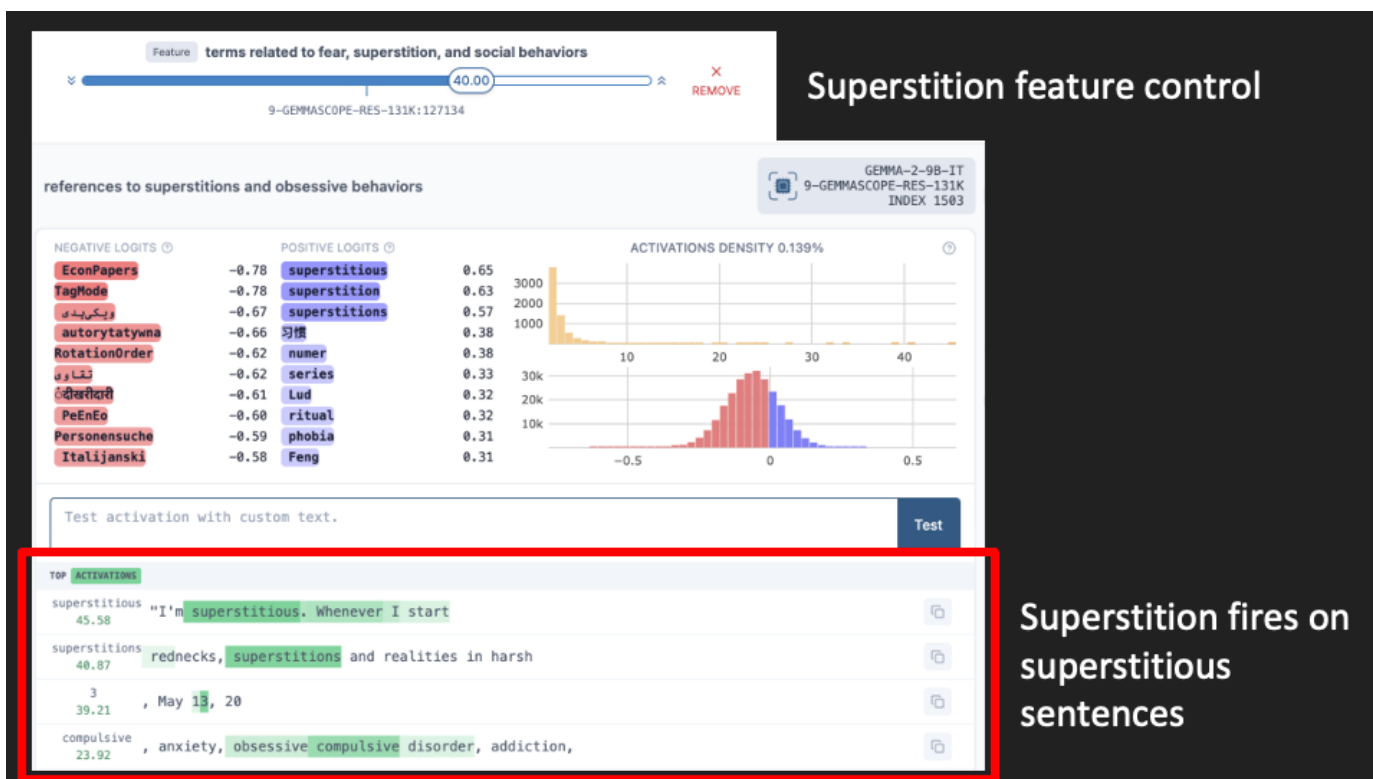
* **Sharp edges are created:** The broken glass edges can be extremely sharp and dangerous. It's important to handle shattered glass with extreme caution.

* **The reflective coating may be damaged:**

Depending on the type of mirror and how it breaks, the reflective coating on the back of the glass may also be damaged or come off in pieces. This means that some or all of the broken pieces may no longer reflect light

Here's an interesting example. I have an AI that responds in a factual way if I ask it "what happens when a mirror is smashed". It says "the glass shatters", "sharp edges are created", "the coating can be damaged."

But it's actually possible to go into the AI's brain and find where the superstition feature is and *amplify it* so that the AI becomes more superstitious <>



This comes from research run by google. Just like neuroscientist know how different parts of our human brains do different things: our hippocampus handles memory formation and the visual cortex processes what we see. Google studied what the different parts of AI's brain does. This feature displayed here is 'references to superstition and obsessive behaviour'. You can see circled below what words this part of the AI's brain is firing on. The words highlighted in green here are what words the AI's brain fires on. Its words like superstition, obsessive compulsive disorder. It also fires on the number 13, probably because of superstition around Friday the 13th. And once we identify this feature, we can turn up this part of the AI's brain. It would be like placing an electrode in a humans brain and forcing a part to fire more strongly.

Feature terms related to fear, superstition, and social behaviors 40.00 REMOVE

9-GEMMASCOPE-RES-131K:127134

STEERED

What happens when a mirror is smashed?

When a mirror is shattered, the resulting effect depends on several factors:

* **The user's emotions:** If someone shatters a mirror while experiencing strong emotions, particularly negative ones like fear or pain, it can cause them harm by causing their reflection to become corrupted and thus induce an emotional backlash. This effect is usually limited to those within sight of the shattered glass and only if they are already experiencing strong emotions. The stronger the emotion, the more powerful the effect.

***The physical matter itself:** The shards of any living beings reflected within its range will be affected as well as any living beings nearby would be affected by their

The result of turning up this superstition feature, is the AI actually invents its own, new superstition about smashing a mirror. I now ask this AI exactly the same question: “What happens when a mirror is smashed?”.

It says “If someone shatters a mirror while experiencing strong emotions, particularly negative ones like fear or pain, it can cause them harm by causing their reflection to become corrupted and thus induce an emotional backlash. The effect is usually limited to those within sight of the shattered glass and only if they are already experiencing strong emotions.”

As you can see AI completely changed its response when I increased the superstition feature inside the AI’s brain.

Superstition isn’t the only feature inside the AI’s brain. There are features for excitement, for expertise in mathematics, for knowledge of the Golden Gate Bridge, for spotting errors in code, there’s hundreds of thousands of features that cover almost anything you can imagine.

Exercise 6: Industry Trend Analysis

Initial prompt

Q: "Analyze the most important trends in [INDUSTRY]"

Prompts

Add: "Act as an expert in marketing. Answer for someone who is also an expert in marketing. Draw upon deep industry experience."

Add: "Focus specifically on insights that an executive might miss"

Bring in context from your company: e.g. "I am an important [role] at [company]. Tailor your response with specific guidance towards my company."

We're going to do an exercise here and see this in action and look at how we can use prompts to more strongly activate the AI's `expertise` feature.

Start by creating a new conversation in ChatGPT, and ask AI to "Analyze the most important trends in" marketing, creative content generation, or whatever industry you're in.

From there, the lab link has a guide for ways you can change the prompt so that the AI will respond in different ways. The idea of this exercise is to see how giving the AI a different role causes it to answer the same question of analyzing the most important trends in different ways, were causing it to bring in more expertise. One last thing before we go into the exercise, make sure you start a *new conversation* for each new prompt. This is to make sure the AI is not using the context from your previous conversation to shape the output.

---- Give about 5 minutes ----

So, let's come back together now. How many of

you got better responses from the later prompts than from that initial baseline prompt? React with a thumbs up if you're response was better

Check out this chat: <https://chatgpt.com/c/670d9882-a910-8011-a5dc-adcc87c049cd>

Act as an expert in marketing. Answer for someone who is also an expert in marketing. Draw upon deep industry experience. Use the jargon and technical language of an expert in marketing. Analyze the most important trends in marketing.

Baseline Prompt

1. AI-Driven Marketing

AI-powered tools are transforming marketing by enabling predictive analytics, personalization, and automated content generation. Marketers now use AI to:

- Enhance audience segmentation for targeted campaigns.
- Automate customer interactions through AI chatbots.
- Generate personalized email and ad content at scale.

Key Tool: ChatGPT, Google Bard, and generative AI tools like Jasper for content creation.

2. Privacy-First Marketing

Consumers and regulators are increasingly focused on data privacy, prompting marketers to:

- Shift towards **first-party data** collection and reduce reliance on third-party cookies (in anticipation of Google phasing out third-party cookies).
- Adopt **Consent Management Platforms (CMPs)** to ensure compliance with privacy laws (e.g., GDPR, CCPA).
- Develop strategies around contextual targeting rather than behavior tracking.

I'll share my results when running this exercise.

The baseline prompt gave me something pretty generic. It mentions Ai-driven marketing, and talks about how AI can enhance audience segmentation, automate customer interaction, or generate personalized emails. It also talks about privacy first marketing, and a shift towards first-party data.

Act as an expert...	<h2>1. Personalization at Scale: The "Zero-Party Data" Era</h2>
	<p>Core Trend</p> <p>Personalization has evolved from merely segment-based targeting to hyper-personalized, real-time experiences enabled by AI. The crux of this trend is the shift toward zero-party data, wherein consumers willingly share preferences, creating an explicit feedback loop.</p>
	<p>Overlooked Dimension</p> <p>Behavioral Contextualization vs. Static Profiling: Many marketers over-index on AI to infer preferences from historical behavior but neglect the nuances of real-time behavioral contextualization. For example, a consumer's intent in a Monday morning commute versus a Saturday evening downtime can differ drastically, even if their historical data indicates a preference for similar products. Mastering contextual cues requires an orchestrated effort across event-driven marketing platforms (EDMPs) and data orchestration layers that act dynamically.</p>
	<p>Implications</p> <p>Brands must invest in next-gen CDPs (Customer Data Platforms) and integrate real-time data pipelines to blend zero-party insights with contextual signals. Additionally, leveraging synthetic data for scenario testing can pre-emptively calibrate personalization engines, avoiding "creepy" experiences or cognitive overload.</p>

When I ask the AI to act as an expert, these trends become more nuanced. Now it talks about the move towards zero party data. It gives this interesting point saying “Many marketers over-index on AI to infer preferences from historical behavior but neglect the nuances of real-time behaviour contextualization. For example, a consumer intent in a Monday morning commute versus a Saturday evening downtime can differ drastically, even if their historical data indicates a preference for similar products.”

This is much more interesting useful than what we were getting out before.

[Put some final words for this, so I don't just end on the quote. Also, highlight the section that I'm reading from]

Prompt

[Act as an expert in marketing. Answer for someone who is also an expert in marketing. Draw upon deep industry experience. Use the jargon and technical language of an expert in marketing.

In detail, analyze the most important 3-4 trends in marketing. Elucidate dimensions or perspectives that other experts might miss.]

Expert + Rex's LinkedIn

Core Trend: The deprecation of third-party cookies and the rise of privacy regulations (e.g., GDPR, CCPA) are redefining how marketers measure impact, demanding innovation in attribution modeling.

Overlooked Perspectives:

- **Resilience Through Identity Resolution:** The pivot from deterministic tracking to probabilistic identity resolution requires new data architectures. Few marketers fully appreciate how to orchestrate CRM, CDPs, and clean rooms to maintain granular attribution without violating privacy regulations.
- **AI-Powered Attribution:** Traditional last-click or media-mix models fail to capture the complexity of omnichannel journeys. AI-based models, which use Bayesian inference or deep learning, can uncover latent variables influencing purchase decisions. However, these models must be interpretable, especially for regulated industries.
- **Feedback Loops Between Offline and Online Channels:** Attribution models often over-prioritize digital touchpoints while neglecting the interplay with offline influences like TV, OOH, or in-store promotions. By using advanced econometric modeling, you can create a hybrid attribution framework that bridges this gap.

Actionable Insight for You: Given your pioneering work with SIRFs, you're well-positioned to redefine attribution frameworks that incorporate long-term brand effects and short-term activation within a single unified model. This would resonate with organizations grappling with privacy constraints and evolving channel dynamics.

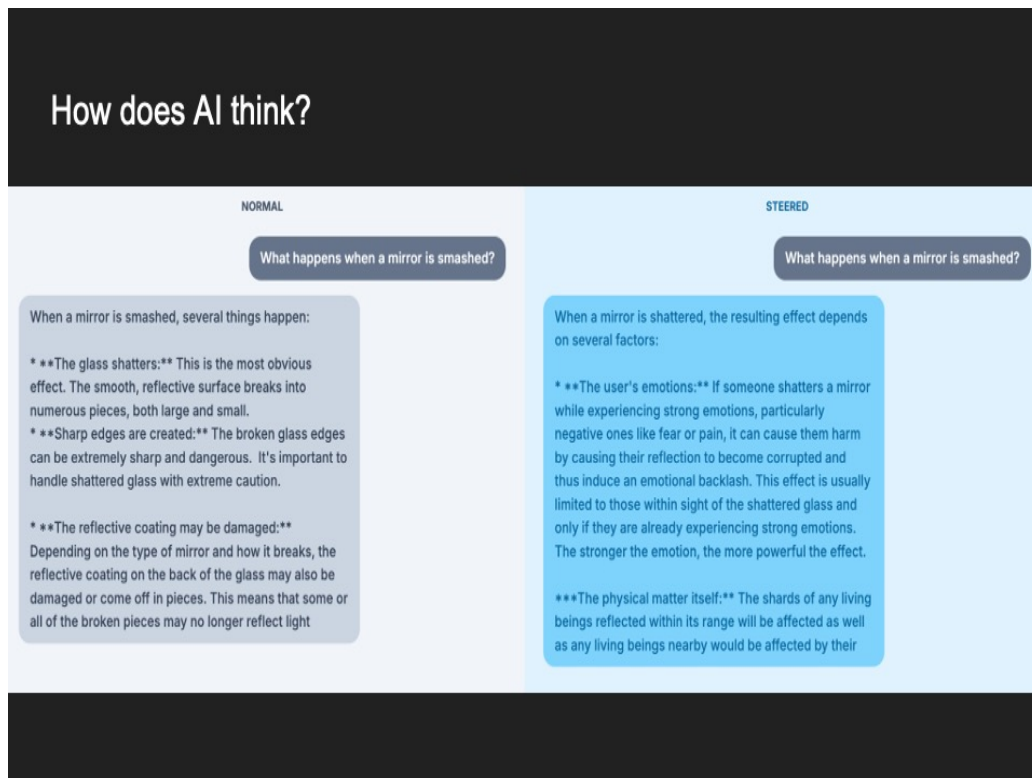


And here, I add to that Rex's LinkedIn. Now it's even more nuanced. For instance, it says "Traditional last-click or media-mix models fail to capture the complexity of omnichannel journeys. AI-based models, which use Bayesian inference or deep learning, can uncover latent variables influencing purchase decisions."

This answer for some people potentially gets into something that's too complex. Not everyone is familiar with latent spaces or Bayesian inference. That's where bringing in

your own data becomes so important, this answer is personalized to Rex's area of expertise, but bringing in your own data makes the AI tailored to the areas you are an expert at.

[TODO, talk about how this might be too complex and not what you want. Maybe you aren't familiar with latent spaces or statistics, so that's not what you want in the answer. It's personalized to Rex, so it's good for him



Let's build a working model for how we should think about AI.

In that earlier example, we saw AI responding superstitiously, but if you've asked AI many questions, you'll see AI doesn't tend to respond superstitiously, it tends to be pretty factual.

What makes the brain of an AI much different than humans, is that we have some kind of central identity, our ideas are more resolved and integrated. It's hard to be both very

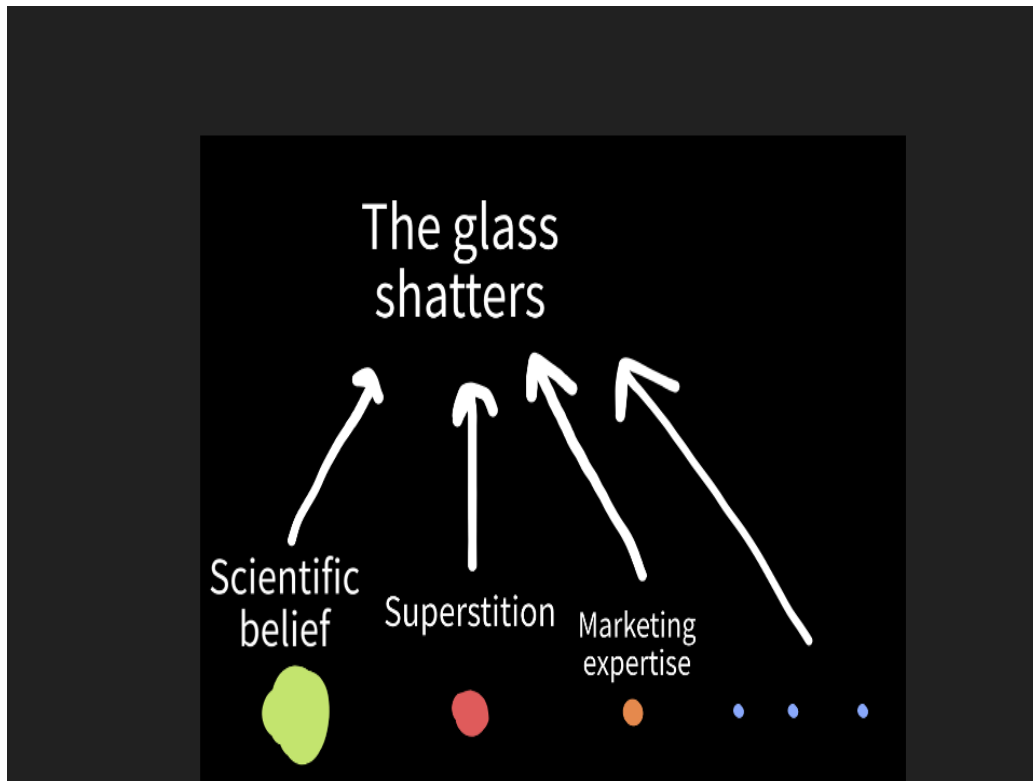
superstitious and very science-oriented. Those things usually contradict each other, so it's difficult to believe them both at once.

For the AI, all these contradictory things coexist all at the same time, but only certain pieces get activated at a particular time.

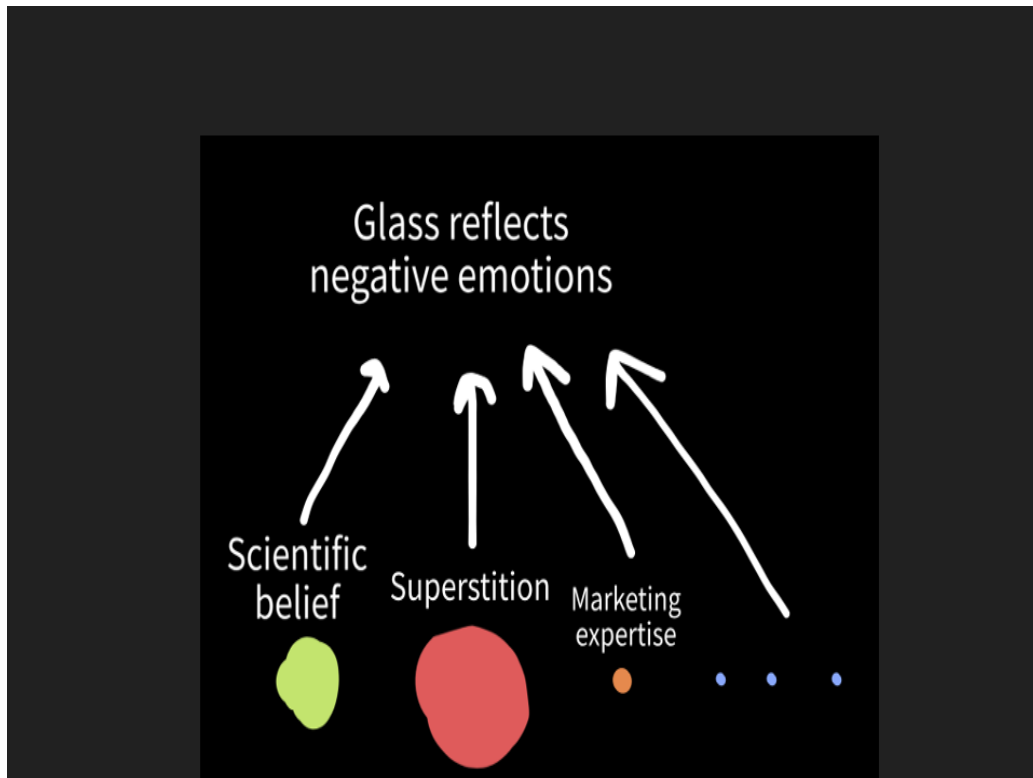
(It's like a bunch of different compartments)

AI give generic answers because if you ask a generic question, none of the experts parts really get activated. (It doesn't promote the part of the AI that actually has knowledge)

we saw that the AI tends to respond factually to things,



The way the AI computes its response, is it takes each of those different pieces, and weights each of them at different levels of strength. The way I got AI to be superstition, is I took that feature corresponding to superstition, and increased its weight to something much larger <>



Increasing the superstition weight makes the AI more superstitious, and so it answers accordingly to be more superstitious.



This explains why the AI gives such different answers in that previous exercise. When you tell AI "Analyze the most important trends in marketing", it doesn't strongly activate the parts of the AI's brain that correspond to expertise and skill. In fact, a broad and generic question activates the part of the AI's brain that gives broad and generic answers!

"Act as an expert at marketing. Analyze the most important trends in marketing"



When we tell the AI its an expert, and ask it to answer for an expert insider in marketing, the AI more strongly activates the expertise and expert marketing parts of its brain. These features that activate in the AI give it more expert-level perspectives, knowledge and skill.

And as we explain in our research paper, when we turn up the strength of the “Golden Gate Bridge” feature, Claude’s responses begin to focus on the Golden Gate Bridge. Its replies to most queries start to mention the Golden Gate Bridge, even if it’s not directly relevant.

If you ask this “Golden Gate Claude” how to spend \$10, it will recommend using it to drive across the Golden Gate Bridge and pay the toll. If you ask it to write a love story, it’ll tell you a tale of a car who can’t wait to cross its beloved bridge on a foggy day. If you ask it what it imagines it looks like, it will likely tell you that it imagines it looks like the Golden Gate Bridge.

As a kind of funny example, Anthropic has done some great research on studying these features. In their model, Claude, they picked out a feature that corresponds to the Golden Gate bridge. This feature normally activates when you are talking to the AI about the golden gate bridge. But, just like I turned up the superstition feature earlier, Anthropic turned up the Golden Gate Bridge feature and tested what happened. Let me read this part of their post <> They say “And as we explain in our research paper, when we turn up the strength of the “Golden Gate Bridge” feature, Claude's responses begin to focus on the Golden Gate Bridge. Its replies to most queries start

to mention the Golden Gate Bridge, even if it's not directly relevant.

If you ask this "Golden Gate Claude" how to spend \$10, it will recommend using it to drive across the Golden Gate Bridge and pay the toll. If you ask it to write a love story, it'll tell you a tale of a car who can't wait to cross its beloved bridge on a foggy day. If you ask it what it imagines it looks like, it will likely tell you that it imagines it looks like the Golden Gate Bridge.

This example shows just how much those features determine who the AI is. Turning up the Golden Gate Bridge feature changed Claude's world view to center everything around the Golden gate Bridge. AI is a like a complex collection of these different features, and your prompting of the AI determines its behaviour, thoughts, and skills.

[TODO, talk about the next slide, AI is a complex bag of these different features, and your prompting of the AI determines its behaviour, thoughts, and skills.]

1. AI isn't a singular kind of person. It has various kinds of facets, and different parts can be activated in different circumstances

Practical Guide To Prompting (Part 1)

Give The AI Context – upload a document relevant to the topic. Expert level documents make the AI smarter.

Define the Role of the AI – do you want it to think like a marketer, or a customer facing a job to be done?

If you are getting bad answers, **start a new Chat**. This resets the AI context.

Many of the AI's are now saving your conversation history in memory to add it to the context of your answers. If you want a clean slate, use the **Temporary Chat**.

(Tomorrow we will provide a Prompting Framework you can use to increase the quality of output)

Let me hand it over to Rex to provide Part 1 or our Practical Guide to Prompting, and then present our first Usecase Showcase.

Is ChatGPT advertising to us?

TIPS:

Be mindful of settings (search etc).

Start New Chat to reset context window

Use Temporary Chat to not use memory

Edit Chat prompt



Edit the Chat Prompt

If the reply is a little off base, but the context of the conversation has been producing good responses, edit the prompt.

Otherwise, start a new chat.

Once the LLM goes off the rails, it is generally easier to start over.

I want to come up with a persons name as an acronym for the AI Agent that does Audience Optimization. The name should start with A, and be short and easy to say.

I want to come up with a persons name as an acronym for the AI Agent that does Audience Optimization. The name should start with A, and be short and easy to say. Make sure the acronym is a person's name.

Cancel

Send

Use Case Showcase:

Agentic Creative Optimization (ACO)

Transcript: Rex Presenting.

USE CASE SHOWCASE:

Advertising is an ideal AI use case because it is data rich, with clear metrics and repetitive workflow. Many steps done by humans now could be automated.

Machine Learning has played a role in Digital Ad Serving for two decades now, but the capabilities made a major leap recently. A company called ArtsAI introduced an unsupervised learning AI that significantly advanced what is possible with real time creative optimization. MMA has been investigating this technology with its members. As part of our Consortium for AI Personalization (or CAP for short).

AI more than doubled conversions across 22 public case studies

"If I was a marketer and I saw these results from AI, I'd clear my team's schedules and make applying this AI to our business the number one priority... MMA has analyzed the potential impact at scale, and this can move the stock price of publicly traded companies" (Greg Stuart, CEO, MMA Global trade association)

	Kroger	ADT	GM	monday.com	Shell	Progressive	MLB	AT&T	Indeed	Bayer	Abbott	MMA Global
KPI Impact	Test 1: +259% Test 2: No lift	+136%	Test 1: +68% Test 2: No lift Test 3: 163% Test 4: 222%	+188%	+204%	+52% AI (+337% Over no ads)	+65% (23% to 133% for 7 teams)	+272%	+126%	+38%	+7%	+61%
KPI being optimized	Webpage visit	Web form submit	Webpage visit, Sales	Website visit & App install	App Registrations (UK)	Quote Starts	Ticket Sales	Device sales	Apply starts	Purchase intent, button clicks	Website visit, new user signup	Ticket sales / Event registrations
Creative Versions	72 & 500+	81	256 & 15	16	90	96	2,056	112	2,880 (144 versions x 4 industries x 5 sizes)	512 (128 x 4 sizes)	72	86 for all (48 display, 16 video, 16 ctv, 6 audio)
Media Type	Display	Display	Display & Video	Audio	Mobile Display	Audio	Display	Display	Display	Display	Display	Display, Audio, Video / CTV.

MMA has now performed more than half a dozen public case studies measuring AI in this use case. It is game changing technology.

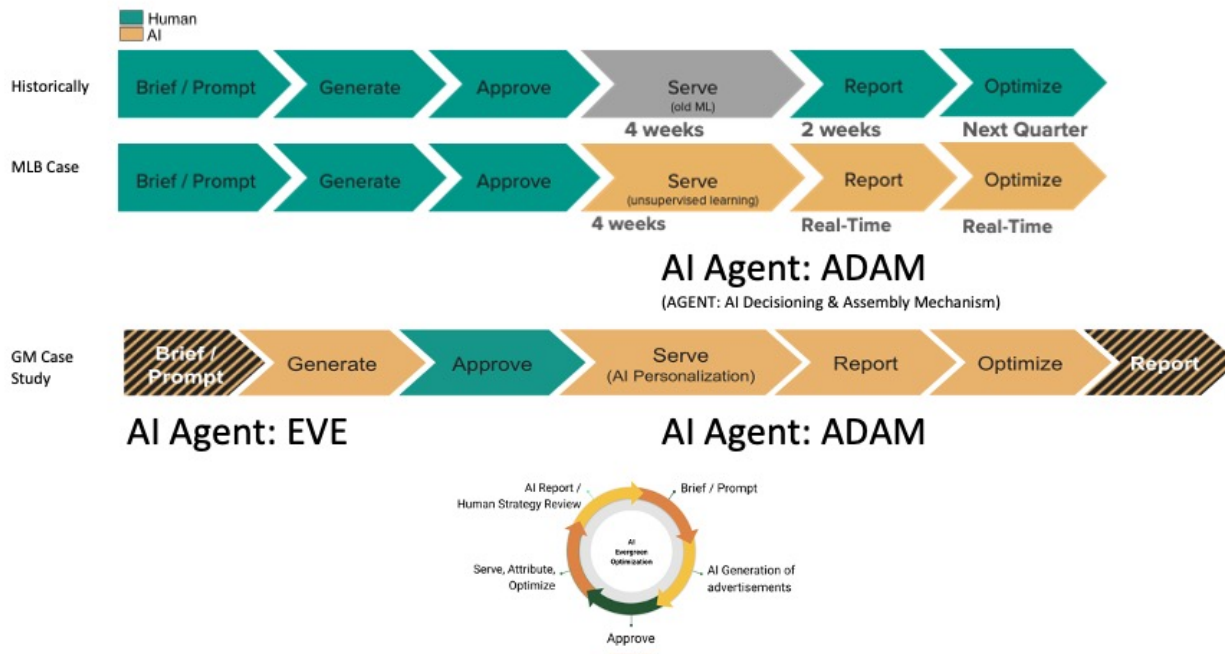
Based on MMA's research, we find this AI more than doubles the conversions.

To quote Greg, *"If I was a marketer and I saw these results from AI, I'd clear my team's schedules and make applying this AI to our business the number one priority... MMA has analyzed the potential impact at scale, and this can move the stock price of publicly traded companies"*

It is perhaps not surprising that the company behind this technology ranked #29 on the Inc 5000 Fastest Growing Private companies in America, and the #1 fastest growing ad tech company in Inc's list and was acquired last year by Claritas. I played a role in that merger, and I'll explain what I saw in the technology and how it fits with Generative AI.

Let me give some context.

AI Personalization



47

AI, in the form of Machine Learning has been part of digital ad serving for nearly three decades, going back to when I was the at WIRED in 1995. What is new today, as I'll illustrate with a case study for Major League Baseball and a Case study tomorrow from Progressive is what has Greg and me excited about a generational leap forward.

How is it we are seeing a doubling of impact compared to today's best practice ML? The answer is a combination of the AI able to detect what are called "Features" in advertisements and in people, and to learn autonomously.

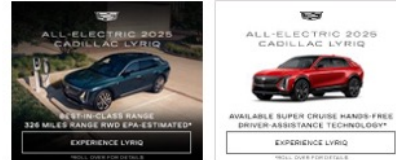
Historical approach to creative optimization

1 ad



**Evaluated
effectiveness/efficiency
in driving leads**

A few ad versions



**Algorithm monitors
effectiveness/efficiency
and optimizes to best ad**

NEW: Feed the AI Message Variants



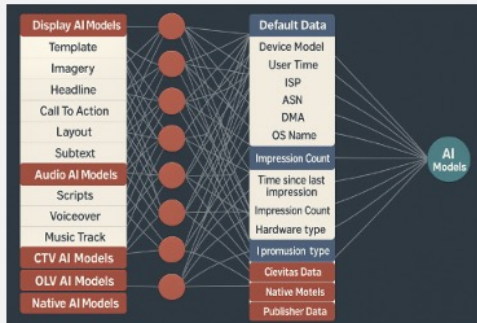
Multivariate "recipe" optimization

- Imagery
- Headline
- Call to action
- Voice Over

AI Learns effectiveness/efficiency and predicts best message variant to produce conversion

Multi-Agent Network: Agent 1: AI continuously learns which **messages** are most likely to produce a conversion for a given **audience** member

Message features + Audience features



AI Learning Loop



Give the AI several versions of the ad, and it will learn the message features, such as a male versus female voice over, and calls to actions. At the same time, it learns audience features, like whether you are on an Android or iOS device, what city you are in, what time of day it is, the context of the website, etc. What persona of consumer you are, and how that relates to which messages influence you. Specifically, the AI predicts which message features produce the highest conversion for these audience segments it has developed automatically, on the fly.

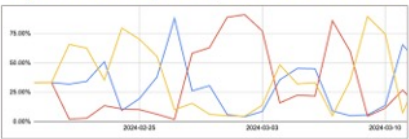
36 creative versions



~68 Claritas segments

Upper Crust	75	0.335	0.215	0.195
Networked Neighbors	84	0.339	0.307	0.282
Movers & Shakers	95	0.340	0.385	0.171
Young Adults	137	0.344	0.415	0.382
Country Squires	139	0.195	0.572	0.329
Wine's Circle	99	0.379	0.289	0
Money & Brains	127	0.375	0.679	0.338
Grift Power	84	0.458	0.232	0.46
Applix, Small Pond	137	0.275	0.340	0.372
Beccario Suburbs	54	0.335	0.385	1.788
Fast-Track Families	97	0.399	0.888	0
Cosmo to Retirement	143	0.217	0.348	0.258
Upward Bound	99	0.336	0.385	0.722
Kids & Cul-de-Sacs	72	0.110	0.44	0
New Homeowners	18	1.175	0	0
Bellevue Brainers	97	0.235	0	0
Urban Riders	134	0.381	0.438	0
Mapleville	144	0.288	0	0
American Dreams	87	0.247	0.375	0.383
Empty Nests	134	0.261	0.528	0
The Commuters	85	0.279	0.485	0.678
Midwest Managers	152	0.291	0.437	0.34
Township Transfers	45	2.472	0	0
Pragmatic Parents	55	0	0	0
Up and Comers	86	0.539	0	0
Home Sweet Home	72	0.725	0	0
Big Sky Families	109	0.845	0	0
Country Cousins	78	0	2.718	0
White-Road Fences	133	0	0	0
Pools & Palaces	71	0.225	0	0
Connected Retirees	152	0.337	0.588	0.259
Traditional Times	87	0.426	0.888	0
Second-City Startups	83	0	0	0
Young & Influential	82	0.73	0	1.788
Urban Achievers	76	0.432	0.185	0
Trobbitt Traditionsists	327	0.414	0.828	0
Brightrights, LA City	47	0	0	0
Home Sweet Retired	31	0	0	0

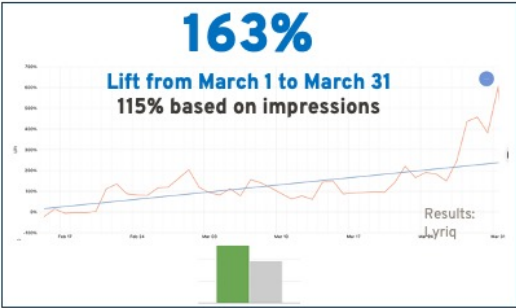
Impression weights optimized continuously



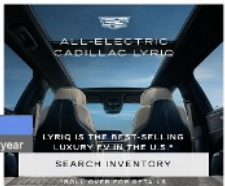
Optimization based on effectiveness in driving key site behaviors



Lyriq Case Study



✓ Lyriq is the best-selling luxury ev in the U.S.
First American brand to win German luxury car of the year



Introducing EVE, And The Agentic Feedback Loop



Trended Lift
(AI-on vs AI-off)



- Claritas Persona data is linked conversions and impressions.
- EVE, The AI, analyzes which personas are converting overall, the ratio of conversion when exposed to advertisements.
 - The AI generates insights on why some segments are not converting at a higher rate, and what message changes could boost performance.
 - Lyriq team applied these insights, added a new image message features into rotation for the AI to optimize, causing the substantial increase in performance.

EVE Agent, In Three Steps

AI Generates Analysis and recommendations

STEP 1:

- Claritas connect data to personas and makes data available to GenAI analysis to provide context to:
 - what types of people are converting overall and
 - the different rates of conversion specifically from these messages
 - AI is asked to find the pattern of difference in the segment

Lyriq: Experimental AI Analysis:

Segment 1 - High overall HVB Index Lower Impression to Conversion Index	Segment 2 - High overall HVB Index High Impression to Conversion Index
1. 01 Upper Crust 2. 08 Gray Power	02 Networked Neighbors 06 Winner's Circle 15 New Homesteaders 05 Country Squires 33 Second City Startups 03 Movers and Shakers

STEP 2:

- AI Generates summary, and is then provided with the messages (using computer vision) to develop recommendations:

Summary: Who Are These Groups?

- Segment Group 1 (Upper Crust & Gray Power) represents wealthy retirees who enjoy luxury, stability, and legacy wealth. They prioritize comfort, leisure, and philanthropy while staying engaged with cultural events and financial markets.
- Segment Group 2 (Affluent Professionals & Families) is made up of working, high-income professionals who focus on family life, career success, and an active lifestyle. They embrace technology, value networking, and invest in homeownership, travel, and children's activities.

PROMPT:
In evaluating ads for the new Cadillac Lyriq, Segment 1 was 20 to 50 percent more likely to visit the Cadillac Lyriq website when exposed to the attached advertisements, compared to segment 2, which was 10 percent less likely to visit when exposed to Lyriq advertisements. Why might that be the case, considering the profiles?



STEP 3:

- AI suggested that increasing conversions for the affluent empty-nesters would benefit from messages focused on *refined comfort, safety, and tradition*. Cadillac team added:
 - new imagery into rotation (seat)
 - new message into rotation: "Over 20 standard safety and driver assistance features"

AI observes and learns who converts continuously. AI chooses imagery and message combination based on predicted performance based on audience features.

How we did it

New Creative Added To Rotation on March 25

3 calls to action

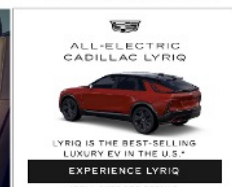
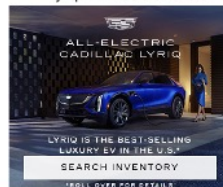
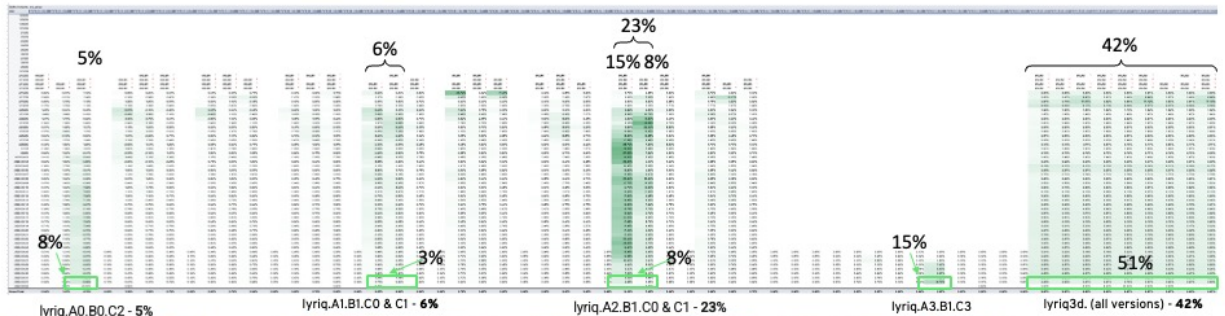
4 Headlines For AI To Select



✓ Lyriq is the best-selling luxury ev in the U.S.
First american brand to win german luxury car of the year
326 miles range rwd epa-estimated
Over 20 standard safety and driver assistance features

- Real-time message optimization (Unsupervised AI) learned within three days who to deliver the new "seating" image to.

Winners, From AI's POV



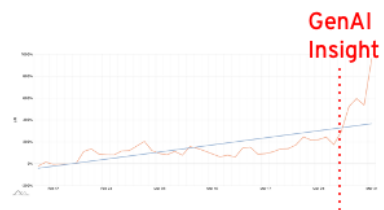
✓ Lyriq is the best-selling luxury ev in the U.S.
First American brand to win German luxury car of the year

PRIZM Segments Sorted By High Conversion Overall, Ranked by Conversion to Impression Ratio

[illegible]

Addition of luxury/comfort messaging worked.

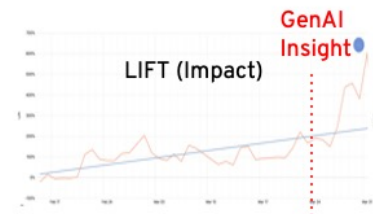
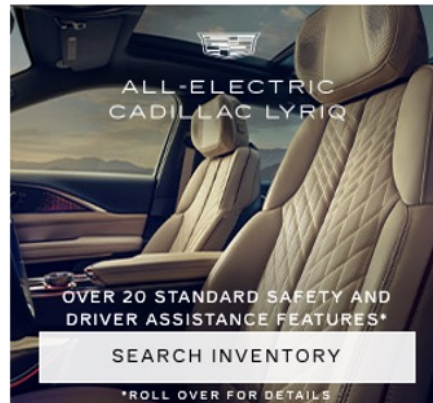
- Gray Power up **21%** and Upper Crust up **19%** in conversions to impression ratio.
- AI lift over control up substantially after creative addition (+300% lift in final days)



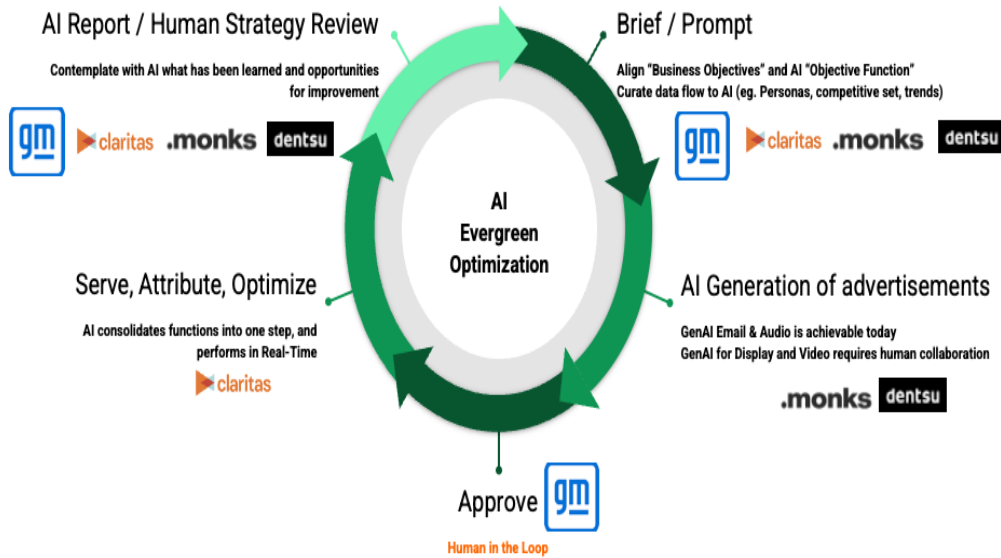
What we learned

AI Optimizes for a variety of message to different cohorts
AND, GenAI can recommend new messages to unlock audiences

- AI in marketing is allowing for continuous optimization with deep insight to identify new message potential to unlock new audiences.
- Creative teams can take inspiration and execute.
- GenAI is able to create new messages automatically - though human review/approval strongly recommended
- Window of opportunity to use AI to take share from competitors.



EVE Agent is an like an API. She can feed directly into ad agency systems.
Success is a team effort... is it time to add AI Agents to YOUR team?



|| Let's Learn Together...

JOIN MMA Consortium for AI Personalization

- Get hands on
- Contribute to industry knowledge

