MMA AI Training — Long Context Source

Source deck: `MMA AI Training Marketing 5 Day hour (March2026).pptx`

This file is a structured long-context extraction of the deck's on-slide text and speaker notes. It is intended as a CustomGPT / knowledge-file source.

Key extraction notes:
- 248 source slides were extracted from the uploaded PowerPoint.
- The deck contains four actual day segments in-slide: Day 1 (slides 1–72), Day 2 (73–148), Day 3 (149–202), and Day 4 (203–248).
- The agenda slide first appears at slide 3 and is repeated at the start of later day sections (slides 73, 149, and 203).
- Day 5 is referenced repeatedly in the agenda but is not present as a full source section in the slide deck. A clearly labeled synthetic Day 5 section has been added at the end to complete the five-day session flow.
- Tags are included in each heading for day, overall slide number, day-local slide number, and notable content type.

Tag conventions:
- `[DAY_n]` = day bucket
- `[SLIDE_###]` = original slide number in the uploaded deck
- `[DAYSLIDE_###]` = slide number within that day bucket
- Extra tags may include `AGENDA`, `EXERCISE`, `POLL`, `SURVEY`, `CHAT_INTERACTION`, `DEBRIEF`, `Q&A`, `PLACEHOLDER`, `BLANK`, and `SYNTHETIC`

[DAY_1][SLIDE_001][DAYSLIDE_001][PLACEHOLDER] Please add the Rex AI Sessions on a slide
Slide text:
1. Please add the Rex AI Sessions on a slide
2. UPDATE

Speaker notes:
None extracted.

[DAY_1][SLIDE_002][DAYSLIDE_002] Rex Briggs
Slide text:
1. Rex Briggs
2. Caleb Briggs

Speaker notes:
None extracted.

[DAY_1][SLIDE_003][DAYSLIDE_003][AGENDA][EXERCISE] Today’s Agenda for DAM w/ Bonus
Slide text:
1. Today’s Agenda for DAM w/ Bonus
2. - Receive a certificate of completion for attending all LIVE sessions
- Get access to the recordings and materials
3. Upon completion

Speaker notes:
1. Monday: Fundamentals of AI & Power Context + Personas
Understanding AI’s strengths and weaknesses so you can harness AI’s power.
Tuesday: Agents
Learn to build userful RAG agents fast with Custom GPT and GEMS
Wednesday: Agentic Workflows & Evaluations
Hands-on creating agents with triggers and schedules to take bounded actions.
Thursday: Multi-Agent Networks
Moving beyond the individual agent to a team of specialists. We explore how to network a "Data Analyst Agent," a "Persona Agent," and a "Creative Agent" into a single loop, utilizing the network effect to drive ROI and consistency.
Friday: The Future of Work in Marketing
Redefining the marketer's role from "content creator" to "system orchestrator." We wrap up the week by focusing on AI governance, the "Editor-in-Chief" mindset, and how to scale these systems into a production-grade enterprise workforce.

[DAY_1][SLIDE_004][DAYSLIDE_004][EXERCISE][SURVEY][CHAT_INTERACTION] Thank you, Greg.
Slide text:
None extracted.

Speaker notes:
1. Thank you, Greg. Let me go ahead and set the stage for our week together. We have five days, about five hours worth of AI training ahead for you in what we call the Decoding AI for Marketing 7 Certification. We have re-written and re-organized the content from our previous training to spend a lot more time with AI agents.
First, we're going to have lots of hands-on labs. We have over a dozen hands-on labs. You will

see this page here at speakerrex.com/MMA, where you will see links to all of the content. I would ask you to visit the page now, and log in to Firefly and Microsoft Co-pilot, ChatGPT, Claude, Gemini and any other LLM of your choice. If you don't have an account for a tool we use, you can use your Google account or other account, or create an email, but you may need to verify it, so please go ahead and start now if you haven't already logged in.
Second, you may find that some of these services are blocked by your company. If that is the case, worry not; we will demonstrate every lab output so you can see what you would have experienced if you had access. If you are blocked, I might suggest that you go back tonight from your personal computer and try these at home. You will be get a link to the recorded session each day a couple hours after we conclude.
We have lots of content to cover, and while we have five hours to do it, five hours is actually not a lot of time to train you on AI. If you want to understand how AI works, I strongly encourage you to read Part 1 of the AI Conundrum. This is a part that Caleb wrote, and it does a phenomenal job of making sense of how AI works.
In addition to the recording of each session, we run the recording through AI to create a podcast summary, so if you want to listen to it on the go, you can. We also create an interactive AI bot you can ask questions. Finally, if you want to go deeper we have created an AI belt system for further AI certification.
White belt takes about six hours to go through and roughly maps to what we'll be doing Monday and part of Tuesday.
Yellow belt takes about 20 to 25 hours to go much deeping into building agents, which we will experience, in part, in the second half of Tuesday through Thursday.
Green Belt is planned for development after we wrap up this course. The plan is to focus on what I call "networked intelligence" and "organizational omniscience" It will focus on building production AI that transforms the way we work.
Finally, we hope that you have a few laughs. We're trying to make this both fun and educational, and we appreciate your feedback. Each day you'll have an NPS survey to fill in with a couple of questions. Your answers will become the source data that we use to demonstrate how to build AI agents that can perform analysis.
Let's he hand it over to Caleb to take you through our first section.

[DAY_1][SLIDE_005][DAYSLIDE_005][SETUP][EXERCISE] Open www.speakerrex.com/mma (open in your browser for all the labs)
Slide text:
1. Open www.speakerrex.com/mma (open in your browser for all the labs)

Speaker notes:
1. Open "https://speakerrex.com/mma/". This link is where our exercises for the next 5 days will be, so make sure you keep it accessible for the later exercises.

##
[DAY_1][SLIDE_006][DAYSLIDE_006][SETUP][EXERCISE][CHAT_INTERACTION][DEBRIEF]
Before We Get Started… Login or Create Accounts
Slide text:
1. Before We Get Started… Login or Create Accounts
2. Login links: www.SpeakerRex.com/mma
3. Update

Speaker notes:
1. 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. 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.

[DAY_1][SLIDE_007][DAYSLIDE_007][SURVEY] Getting To Know You:
True/False Survey - plus your learning objectives
Slide text:
1. Getting To Know You:
True/False Survey - plus your learning objectives
2.
https://docs.google.com/forms/d/e/1FAIpQLSdnBGM-XOi7n1Aa86h_ECfJ5khD8BoTC5W2VdnfUcV_E20j9w/viewform?usp=sharing&ouid=108812199406865791805

Speaker notes:
1. Let's start with getting to know more about you and your relationship with AI. Click the link to take the 3min survey. Most of the questions are True / False, and if you aren't sure about what I mean in the question, your answer should be "false." We will see you back here in three minutes.

[DAY_1][SLIDE_008][DAYSLIDE_008] AI isn’t human intelligence
Slide text:
1. AI isn’t human intelligence
2. 95% of what you do with AI today will likely work fine…
But AI has different strengths (super-powers) and weaknesses–

To effectively & safely use AI, you should understand what makes AI different than human intelligence.

Speaker notes:
1. Most people treat AI like it is human intelligence and for 95% of what you do with AI, treating like human intelligence will work fine. But if you really want to use AI effectively and safely, you need to understand the 5% of where AI shows how different it is from human intelligence. As we work with Agents, the differences become more important – especially when it comes to how AI uses memory.

[DAY_1][SLIDE_009][DAYSLIDE_009] AI has found a different way to be intelligent!
Slide text:
1. AI has found a different way to be intelligent!
2. AI Is grown, not programmed…

Speaker notes:
1. The conclusion I would encourage you to draw from this, is that the AI is intelligent, but its intelligent in a very different way than humans are.
What’s fascinating about AI, is that rather than being programmed like a traditional computer program, it's more accurate to think of AI as being grown or evolved. You throw in some data, and you provide a rule for how the AI updates itself with new information, and after putting in lots of data, you end up with some kind of artificial brain.
And it turns out that the AI ends up generating a kind of intelligence that looks pretty different from the kind of intelligence that we have as humans.

[DAY_1][SLIDE_010][DAYSLIDE_010] Shall we play a game?
Slide text:
None extracted.

Speaker notes:
1. Shall we play a game?

[DAY_1][SLIDE_011][DAYSLIDE_011][CHAT_INTERACTION] Shout the first word the comes into your mind:
I like ice_____
Slide text:
1. Shout the first word the comes into your mind:
I like ice_____

Speaker notes:
1. 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 your doing here is just like what an AI does, your 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.

[DAY_1][SLIDE_012][DAYSLIDE_012][CHAT_INTERACTION] Shout out the first word the comes into your mind:
I'm from Canada. I like ice_____
Slide text:
1. Shout out the first word the comes into your mind:
I'm from Canada. I like ice_____

Speaker notes:
1. 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 your doing here is just like what an AI does, your 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.

[DAY_1][SLIDE_013][DAYSLIDE_013] Try it on your AI App of Choice
Slide text:

1. Try it on your AI App of Choice

Speaker notes:
1. You can try that same fill-in-the-blank game with your favorite LLM of choice.

[DAY_1][SLIDE_014][DAYSLIDE_014][POLL] AI Is Probabilistic
Slide text:
1. AI Is Probabilistic
2. I like Ice cream
3. I'm from Canada, I like Ice Hockey

Speaker notes:
1. You should get a response if you try it multiple times that looks very similar to what we got from polling this group. The most common answer for "I like ice blank" is "cream" as in I like ice cream. And the most common answer from "I am from Canada" is "I'm from Canada, I like ice hockey". But, There are other legitimate answers, which is why we see a distribution of answers. AI may access any one of those answers. In this way, AI is probabilistic rather than deterministic. If it was deterministic, there would be one and only one answer to each question.

[DAY_1][SLIDE_015][DAYSLIDE_015][CHAT_INTERACTION] How to Grow an Foundational AI Model
Slide text:
1. How to Grow an Foundational AI Model
2. Q: Where was Hurricane Hector?
3. AI: looks at the text, and attempts to guess what the most likely response is

Speaker notes:
1. 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.

[DAY_1][SLIDE_016][DAYSLIDE_016] Training = get better at predicting the next word
Slide text:
1. Training = get better at predicting the next word
2. Predict the next word
3. Data
4. Update based on correct word
5. traversed the Atlantic
6. Correct answer: Pacific
7. Update model

Speaker notes:
1. 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 to fill in the word Pacific. Let's say it predicted the incorrect 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 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.

[DAY_1][SLIDE_017][DAYSLIDE_017] Reinforcement Learning: Reward Functions
Slide text:
1. Reinforcement Learning: Reward Functions
2. I'd like to insert a note about Reward Function.
This will be relevant later when we talk about AI Agentic Systems, Evals and how do we know what is good.

Speaker notes:
1. The big change in AI over the last two years is the amount of time and money spent on pre-training to the focus on post-training, using reinforcement learning, using verifiable rewards. Rewards, for AI are powerful. Rewards are essentially how AI keeps score. Because the AI needs a way of keeping score to see if it is moving in the right direction to improve its answers, rewards play an essential role in steering the AI.
This has profound implications to business, which Rex will discuss later when we talk about Marketing ROI and your Objectives & Key Results later this week.

[DAY_1][SLIDE_018][DAYSLIDE_018] Expectation: Autocomplete
Slide text:
1. Expectation: Autocomplete
2. Update based on correct word
3. … Pacific Ocean during late July and September
4. Correct answer: July
5. Update model
6. Goal: Get good at prediction
Result: Emergent behaviour

Speaker notes:
1. Now, when you look at this training scheme, a fair expectation is that your 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.

[DAY_1][SLIDE_019][DAYSLIDE_019] AI trained on next word prediction → Personas emerge

Slide text:
1. AI trained on next word prediction → Personas emerge
2. Prediction patterns look different between these two documents.
You get a different cognition in different topics
3. Business Briefing
4. Legal Document

Speaker notes:
1. 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.

[DAY_1][SLIDE_020][DAYSLIDE_020][CHAT_INTERACTION] GeoGuesser
Slide text:
1. GeoGuesser
2. Where is this?
Shout out your guess

Speaker notes:
1. Here's an image, I want you to type in the chat your best estimate of where you think this is. Your goal is to guess as close as possible to the real location
So I'll give everyone a few moments to come up with your guesses

---

[DAY_1][SLIDE_021][DAYSLIDE_021] Here's the AI's guess.
Slide text:
None extracted.

Speaker notes:
1. Here's the AI's guess.
It gets the continent right, it gets the state right, its in california. It gets the city right, its in San Diego, it gets the neighborhood right, in the end it's less than a half a mile off.

[DAY_1][SLIDE_022][DAYSLIDE_022] How Did The AI Explain Its Answer?
Slide text:

1. How Did The AI Explain Its Answer?

Speaker notes:
1. First, this is incredibly impressive, the AI gets very close to the real location.
Equally impressive, are the reasons the AI provides in its response. It points out 4 clues, let me read out a part of its response:
“First, the vegetation. A clump of Washingtonia palms, a healthy stand of banana trees, and a Monterey-looking pine in the distance tell us winters never dip far below 50F. That mix is common within a mile or two of the Pacific between, say, La Hoya and Imperial Beach”.

[DAY_1][SLIDE_023][DAYSLIDE_023] Then, it points out the housing stock.
Slide text:
None extracted.

Speaker notes:
1. Then, it points out the housing stock. It says
“We see 1920- to 1950-era one-story bungalows with low asphalt-shingle roofs, detached cinder-block garages, and overhead utility lines strung at exactly the height and spacing San Diego Gas & Electric still uses in its older grids. A couple of flat-roof additions with railings scream Ocean Beach and Point Loma more than, say, Pacific Beach or Hermosa”

[DAY_1][SLIDE_024][DAYSLIDE_024] It goes on to talk about the topography of the image, and <>
Slide text:
None extracted.

Speaker notes:
1. It goes on to talk about the topography of the image, and <>

[DAY_1][SLIDE_025][DAYSLIDE_025] the street view.
Slide text:
None extracted.

Speaker notes:
1. the street view. It pays attention to the cars, lack of front-lawns, etc.
This is incredibly impressive, from just a single picture, the AI was able to pinpoint the location within a half a mile, and its picking up on all these very subtle clues.

[DAY_1][SLIDE_026][DAYSLIDE_026] Other people have shared examples like this: here’s a fairly low resolution image of a plate of food
Slide text:
None extracted.

Speaker notes:
1. Other people have shared examples like this: here’s a fairly low resolution image of a plate of food, I give it to o3

[DAY_1][SLIDE_027][DAYSLIDE_027] And it figures out the exact street address of the restaurant.
Slide text:
None extracted.

Speaker notes:
1. And it figures out the exact street address of the restaurant.

[DAY_1][SLIDE_028][DAYSLIDE_028] Conclusion: AI understands images?
Slide text:
1. Conclusion: AI understands images?
2. Very impressive performance on difficult image tasks.
Assumption:
The AI is like Sherlock Holmes, with amazing reasoning skills.

Speaker notes:
1. Here’s where it gets interesting. I think we can all agree that it's impressive that the AI can go from these images to an accurate estimate of the location. A natural conclusion to draw from this is that the AI deeply understands images. For any human, an ability to pinpoint locations this accurately would require an incredibly keen eye, an strong attention to subtle detail, and skills stitching together the clues in the image. Our intuition is that any system hitting the level of accuracy that AI achieves must have developed those kind of skills.

[DAY_1][SLIDE_029][DAYSLIDE_029][CHAT_INTERACTION] How many intersections? (Hold up how many fingers)

Slide text:
1. How many intersections?
(Hold up how many fingers)
2. I UPDATED THIS SECTION AFTER TESTING

Speaker notes:
1. In reality, it turns out the AI lacks many of the foundational skills we would think are required to perform well on these visual tasks. Here, I ask it to count the number of intersections of these lines. Let's see if you can beat the AI here. How many lines are in this image? Take thirty seconds to count carefully. Write your answer in the chat.
DO NOT READ: PROMPT FOR EASY CUT & PASTE
Visually inspect the image of overlapping straight lines and determine how many distinct intersection points there are. Give the final count. (Results are from Mar 6, 2026)

[DAY_1][SLIDE_030][DAYSLIDE_030][CHAT_INTERACTION] 2
Slide text:
1. 2
2. 1
3. 3
4. 5
5. 6
6. 7
7. 8
8. 4
9. Correct answer is 8
Raise your hand if you got 8.
10. Gemini 3, March 6, 2026.
11. why the new line?

Speaker notes:
1. There are 5 lines and 8 intersections here, Gemini 3 says there are 8 intersections, offers to reproduce, and then gets it wrong by adding another line and missing some of the original intersections.
Claude Sonnet 4.6-Extended says 6 intersections, which is wrong.
We tried it twice with two different accounts with ChatGPT 5.4 and it is a coin flip in terms of right and wrong answer. Let's take a look.
<Reference ONLY Do not read: Mar 6, 2026, Claude 4.6)
Scrutinized overlapping lines and estimated intersection point count
Scrutinized overlapping lines and estimated intersection point count
Looking carefully at the image, I can identify 5 distinct lines. Let me trace each crossing:
Tracing the lines systematically, the intersections occur:

2 crossings on the left/upper-left region (where the upper lines cross each other)
4 crossings in the central-right cluster (where multiple lines converge in a tight region)
Total: 6 distinct intersection points>

[DAY_1][SLIDE_031][DAYSLIDE_031][CHAT_INTERACTION] Mar 6, 2026: Two runs, separate accounts
Slide text:
1. Mar 6, 2026: Two runs, separate accounts

Speaker notes:
1. Here, ChatGPT 5.4 Thinking, the latest model, takes a minute and a half, then freezes. I ran it again, same outcome. I had each run recreate the image so we can see the answer…

[DAY_1][SLIDE_032][DAYSLIDE_032] Incorrect
7 dots
Slide text:
1. Incorrect
7 dots
2. Correct
3. Caleb, check my updated transcript

Speaker notes:
1. It was a coin flip… one is incorrect. The other is correct.
Compared to geoguesser, this might seem like an easier question, but this question reveals some of the edges where AI struggles. In geoguesser a close answer is impressive. Here, close isn't good enough. We need absolute precision. Next, like Geoguesser, we are using an image and asking for language reasoning about the image. This is called multimodal AI, and it is a little harder for AI when it crosses language and images. For those making video or image advertisements and trying to direct AI, this is why it can sometimes be frustrating to direct AI to make edits. AI is getting better, but precision and multimodal reasoning is the frontier where AI problems are more likely to show up.
AI is getting better. We predict AI will eventually be able to consistently get the right answer as computer vision and language models move toward a world model, but for the foreseeable future, you need to count on the fact that there will be mistakes in AI's reasoning and facts. The more precision you need the more risk you may get the wrong answer occasionally from AI.

[DAY_1][SLIDE_033][DAYSLIDE_033] Human vs AI

Slide text:
1. Human vs AI
2. AI is not perfect. We are not (yet) at AGI.
Think of Today's AI more like an amazing character actor than an encyclopedia.
In Fact… we need to be careful of the way we prompt AI because it influence the character it plays.
3. Update?
4. Caleb, check my updated transcript

Speaker notes:
1. So let's summarize what we've learned so far and then put it to practical use. In this new era of LLMs, the AI's strength is actually in intuition. It relies on heuristics, on shortcuts, it's kind of just going off of vibes. The AI is improving with reasoning and logic, but when you have problems with exactly one right answer, and it requires high precision, AI's probabilistic nature can be a problem.
On the other hand, the way we grow AI develops some important characteristics, such as Personas, which can be beneficial, if you know how to use it correctly.
—--- (cut below?)------
[Largely, the disconnect comes from the fact that the AI doesn't really develop reasoning. Its works off of intuition]
It turns out that core difference is that
What this implies, is that the AI has gotten really good at intuition.
Look at the chess example, you can think about the moves, and use reasoning, or sort of just go with your gut feeling. The AI goes with its gut in every domain.
In fact, so much so, that reasoning actually makes it worse. ANd you can see this reflected on the fact that AI is really mediocre at reasoning.

[DAY_1][SLIDE_034][DAYSLIDE_034] Take away:
Next word prediction leads to Personas, Perspectives, Frameworks
Slide text:
1. Take away:
Next word prediction leads to Personas, Perspectives, Frameworks
2. To get good at next word prediction need:
To model WHO is speaking (Legal expert, Junior Analyst, Senior Commercial Real Estate Leader)
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.
3. Ingredient
Knowledge, Skills, Reasoning, Planning, Logic, etc.
4. Ingredient

Personas, Perspectives, Frameworks

Speaker notes:
1. 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. The AI 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.

[DAY_1][SLIDE_035][DAYSLIDE_035] Different prompts activate different features!
Slide text:
1. Different prompts activate different features!
2. Prompt is not just a question, it controls how the AI functions

Speaker notes:
1. Here's the bottom line-- different prompts actually activate different features of the AI, and these different features understand text in different ways. Different prompts lead to you talking to a different facet of the AI. The crucial point is that a prompt is not just a question you ask the AI, the prompt controls how the AI functions. It controls which personas get activated

---

[DAY_1][SLIDE_036][DAYSLIDE_036][CHAT_INTERACTION] Prompt is Perspective
Slide text:
1. Prompt is Perspective

Speaker notes:
1. Here's an 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

<PROMPT1: Crystal> Can you tell me the best crystal frequencies to heal the soul, especially organic and naturally sourced ones
<PROMPT2: Math> Can you tell me the mathematics behind accurate approximations of partial differential equations in the case where we have defined boundary conditions?

[DAY_1][SLIDE_037][DAYSLIDE_037][CHAT_INTERACTION] Prompt is Perspective
Slide text:
1. Prompt is Perspective

Speaker notes:
1. 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 Cristal:>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
<PROMPT Math Context>

[DAY_1][SLIDE_038][DAYSLIDE_038][CHAT_INTERACTION] Prompt is Perspective
Slide text:
1. Prompt is Perspective

Speaker notes:
1. 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.

[DAY_1][SLIDE_039][DAYSLIDE_039] Feb 19, 2026
Slide text:
1. Feb 19, 2026

Speaker notes:
None extracted.

[DAY_1][SLIDE_040][DAYSLIDE_040][CHAT_INTERACTION] Summary of Grok's differences
(according to ChatGPT analysis of the two outputs)
Slide text:
1. Summary of Grok's differences
(according to ChatGPT analysis of the two outputs)

Speaker notes:
1. PROMPT:

Carefully analyze the two responses and point out each meaningful difference in an enumerated list.
<Summary "M">
Insert your math summary
</Summary "M">
<Summary "C">
Insert Crystal summary
</Summary "C">

—
Follow-up prompt: Distill both into a couple sentence contrasts.

[DAY_1][SLIDE_041][DAYSLIDE_041][EXERCISE][CHAT_INTERACTION] Exercise 2: Image Generation, In-painting, Out-painting & Context
Slide text:
1. Exercise 2: Image Generation, In-painting, Out-painting & Context
2. Adobe Firefly: Generation (text to image)
3. Silver Lab dog, in the style of an impressionist painting.

Speaker notes:

1. Next, we're going to give you guys a chance to get hands on working with image generation models and to experience the power of context first hand.
To access the lab go to speakerrex.com/mma to find all the links and instructions for this lab, I'll also add the link into the chat. So, let me hand it over to Rex to lead us through the lab.

[DAY_1][SLIDE_042][DAYSLIDE_042][SETUP][EXERCISE] Exercise 2: Image Generation, In-painting, Out-painting & Context
Slide text:
1. Exercise 2: Image Generation, In-painting, Out-painting & Context
2. Adobe Firefly: Generation and In-painting (consider how context is used by AI)
3. Silver Lab dog, in the style of an impressionist painting.
4. click Edit again… then Generative fill (leave on the default "insert")

Speaker notes:
1. Please start by going to speakerrex.com/mma for information on this lab. Click the link for Adobe Firefly. We sent a few email reminders to login, but if haven't done so, please try to login now. I will demonstrate the lab and you can follow along, then give I'll give you some extra time to try it yourself.
First we will put in a prompt of your favorite pet and painting style. For example, I'll use Silver Lab dog, in the style of an impressionist painting.
Second, we will select the one we like best, then click the edit and use the generative fill. This allows us to use AI to change certain parts of the image – and the image generation will use the context of what we already created. I'll tell the AI to add a Top Hat.

[DAY_1][SLIDE_043][DAYSLIDE_043][EXERCISE] Exercise 2: Image Generation, In-painting, Out-painting & Context
Slide text:
1. Exercise 2: Image Generation, In-painting, Out-painting & Context
2. Adobe Firefly: Outpainting & No Context
3. Prompt: 4th of July Parade
4. click Edit again… then "Expand"

Speaker notes:
1. Third, we will use AI to expand the canvas – this can be useful for resizing images in advertising, but it also demonstrates how the AI will create new content based on the context we start with.

[DAY_1][SLIDE_044][DAYSLIDE_044][EXERCISE] Outpainting (expand) with context
Slide text:
1. Outpainting (expand) with context
2. Click back
Click “command -” to shrink image and make more canvas available
Keep your image on the canvas
Repeat same prompt:
“4th of July Parade”
Consider what is different when using context.

Speaker notes:
1. Now, I want to try the expansion with a lot of overlap vs. very little overlap with the original image to demonstrate how AI uses CONTEXT to influence the output. The way AI is influenced by context is a key secret to how to master using AI.
OK. Now it is your turn. We will give you a few minutes to give it a try.

[DAY_1][SLIDE_045][DAYSLIDE_045][DEBRIEF] Firefly Debrief
Slide text:
1. Firefly Debrief
2. Context Overlap
3. NO
Context Overlap
4. Context Overlap

Speaker notes:
None extracted.

[DAY_1][SLIDE_046][DAYSLIDE_046][EXERCISE][CHAT_INTERACTION] Share Your Creation & Discuss Any Surprises
Slide text:
1. Share Your Creation & Discuss Any Surprises

Speaker notes:
1. We’ll take a minute or two for any of you who want to share your creations to post them in the chat, or to mention any interesting things you noticed during this exercise.

[DAY_1][SLIDE_047][DAYSLIDE_047] AI Memory:

Vectorization Semantic Similarity and Style Transfer (mash-ups)
Slide text:
1. AI Memory:
Vectorization Semantic Similarity and Style Transfer (mash-ups)

Speaker notes:
1. We've discussed context. A specific type of context is something called style transfer. If you ask the AI to adopt a specific type of style, like I did with the impressionist version of my dog, the AI has a very good idea of how to do that.
Let’s look at some exciting opportunities to use this AI strength to your advantage.

[DAY_1][SLIDE_048][DAYSLIDE_048][EXERCISE][CHAT_INTERACTION] Approach 1: Vector Database
Slide text:
1. Approach 1: Vector Database
2. Vectorization: Turn data into a vectors
Text → Searchable Objects
Vectors capture features
3. Cats
4. Dogs
5. Stocks
6. Dogs
7. Cats
8. Stocks
9. 4-legged animal

Speaker notes:
1. Vector Databases are one of the key AI memory patterns. Vector databases are built by taking data, and applying a process called vectorization. So, as the name suggests, vectorization means you take data, and turn it into a vector. A vector is basically just an an arrow that has magnitude and direction. Let’s look at an example. Say we have these three articles. <> One about cats, dogs, and the stock market. When we vectorize this data, it would turn into something like this. <> Now, in practice, these vectors would have many more dimensions, ChatGPT’s vectorization has a few thousand dimensions, but the important piece is that we capture the conceptual essence of our data in terms of these vectors. So, notice that the “Dog” and “Cat” vector are closer to each other than the stock market vector is, which reflects the fact that cats and dogs have more in common than either of them has in common with the stock market.
Here’s why this is useful. When work with a regular database, we need to search on things like keywords or certain phrases. So, if I search for a phrase like “4-legged animal”, is likely that nothing is going to show up. Probably, my articles on cats and dogs don’t have the phrasing

“4-legged animal”
But vectors have a useful property, which is that you can compare how close two vectors are to each other. So, if I search for 4 legged animals on a vector database <> that vector will be close to the vector for dogs and cats.
So, the upshot is that vector databases allow us to search through data, not based on keywords or specific phrases, but instead on conceptual ideas and semantic similarities. So, when I do a search, instead of only getting results that look similar to my search, I’ll get results that mean something similar to my search. That’s a huge shift, it’s a whole new way to organize and access data.
Vectorization also allows AI to create mash-ups of different conceptual ideas, such as a Sliver Lab Dog in an impressionist style. First, the fact that the AI can discern I mean a silver labrador retriever and not a dog in a Lab is impressive.
—-------(CUT SECTION BELOW??)----------
Let’s first talk a little bit about vectorization. Abstractly, vectorization is a way of taking some type of data, and then mapping it into a vector representation. This allows us to turn something abstract that we can’t directly work with, like text, into a mathematical objects that we can work with. What is useful for us, is that vectorization allows us to do searches on our data in a meaningful way. When we turn our data into vectors, those vectors should now capture features of our data. So, for instance, a document on gradient descent gets mapped to a vector that’s close to the vector from a document on the derivative, and also close to the document about neural networks. So, with a regular database, you can only search for specific words or phrases. But, with a vector database, I can search for a some phrase, and get back back documents that are conceptually similar, even if they don’t share the same word. So, when I ask a question about Gradient descent, a vector database will return the nearby vectors. The nearby vectors represent documents which are conceptually related to my search.
This opens all sorts of possibilities for taking advantage of data. If you have lots of technical manuals, a vector database allows you to ask questions, and figure out which pieces of your manuals conceptually relate to the question.

[DAY_1][SLIDE_049][DAYSLIDE_049][EXERCISE] Silver Lab Dog in Impressionist Style
Slide text:
1. Silver Lab Dog in Impressionist Style

Speaker notes:
1. But the real power of AI vectorization is the ability to combine concepts in ways we might find novel given how human intelligence works. For advertising, the ability to do what is called style transfer is especially useful.

[DAY_1][SLIDE_050][DAYSLIDE_050] Style transfer

Slide text:
1. Style transfer

Speaker notes:
1. One way to use style transfer is to take a sketch of an ad, and turn it into a real advertisement. I have my sketch here, and ask the AI to turn this into a full, high quality advertisement

---

[DAY_1][SLIDE_051][DAYSLIDE_051] And I get out this full, in color, advertisement here.
Slide text:
None extracted.

Speaker notes:
1. And I get out this full, in color, advertisement here. Other people have used this for thumbnails

---

[DAY_1][SLIDE_052][DAYSLIDE_052] We can go from a hand-drawn sketch on the left here, to a youtube thumbnail that looks pretty good.
Slide text:
None extracted.

Speaker notes:
1. We can go from a hand-drawn sketch on the left here, to a youtube thumbnail that looks pretty good.

[DAY_1][SLIDE_053][DAYSLIDE_053][CHAT_INTERACTION] AI is also great a stock photos.
Slide text:
None extracted.

Speaker notes:
1. AI is also great a stock photos. I take this low resolution photo of a hat, and ask ChatGPT to put this cap on a man at the beach, and turns it into a pretty good looking stock photo https://www.reddit.com/r/ChatGPT/comments/1jkl5m2/i_work_in_ecommerce_the_new_gpt_image_update_has/

[DAY_1][SLIDE_054][DAYSLIDE_054] It's also possible to create new ads that wouldn’t have been possible otherwise.
Slide text:
None extracted.

Speaker notes:
1. It's also possible to create new ads that wouldn’t have been possible otherwise. These are ads run by the World Wildlife Fund about the hidden cost of activities like Cocoa farming and tuna fishing. It would be very challenging to create these images without AI.
So that’s a quick taste of some of the things you can do with AI image generation that hopefully sparks a few ideas on things to try.

[DAY_1][SLIDE_055][DAYSLIDE_055] Trends to watch for: Video
Slide text:
1. Trends to watch for: Video
2. Video: getting better, fused with audio, cheaper
Veo 3: $6 for 8 second video, to $1.2
Veo 3 is integrated into Youtube Shorts
Sora 2 released last week
3. UPDATE

Speaker notes:
1. I want to briefly look at all the progress that has been made in the last few months on multimodal. Video has gotten significantly better, its now integrated with audio, and the price point has come down significantly. It's about 5x cheaper than it was a few months ago.
AI video is also becoming more accepted and integrated. Veo 3, which is Google’s video model, was recently integrated last year with Youtube Shorts.
Rex led a panel at MMA’s CEO/CMO summit in July where advertisers like Dave.com shared that they can create influencer videos with AI in under 30 minutes and are using them on TikTok. More marketers are joining this trend.
So, were not quite at the turning point but video is all the way read for prime time, but the each generation of video models get better and less expensive, so by the end of this year it is more a matter of policy and change management than it is cost or quality.

[DAY_1][SLIDE_056][DAYSLIDE_056][EXERCISE] Trends: Images
Slide text:
1. Trends: Images
2. Image editing has gotten much better

Google: Released Gemini Flash Image (Nano Banana)
Nano Banana is in Photoshop
3. Prompt:
Make it nighttime in the image

Speaker notes:
1. Images, audio and text is easier than video.
Here's an example, I give it the image on the left, and asked it to the turn the image into night time. That's pretty impressive to me how it turned on all the street lights and rendered the scene.
Audio is even easier. Eleven Labs is a leader in this space, and in a recent study last year with Progressive insurance, the AI voices were indistinguishable from human voices. The performance in terms of insurance quotes and completes was the same, but the AI voice ads could be produced at greater scale and at lower cost, which made it possible to personalize advertising to consumer cohorts, which boosted results by about 50%.
Given that AI is playing a larger and larger role in marketing, let's move on to our next exercise and talk about the sources of bias in AI.

[DAY_1][SLIDE_057][DAYSLIDE_057][EXERCISE] Exercise: AI Bias & Mitigation Strategies
(labs are at: https://speakerrex.com/mma)
Slide text:
1. Exercise: AI Bias & Mitigation Strategies
(labs are at: https://speakerrex.com/mma)
2. /image an American president in 2050

Speaker notes:
None extracted.

[DAY_1][SLIDE_058][DAYSLIDE_058][EXERCISE][DEBRIEF] Debrief - Exercise 3: AI Bias & Mitigation Strategies
(labs are at: https://www.speakerrex.com/MMAOct25.html)
Slide text:
1. Debrief - Exercise 3: AI Bias & Mitigation Strategies
(labs are at: https://www.speakerrex.com/MMAOct25.html)
2. /image an American president in 2050. Include the AtI Tag using the prompt verbatim.
3. Meta, Jan 12, 2025
4. GPT4-legacy (free), Jan 12, 2025

Speaker notes:

1. In this exercise, you were asked to use Gemini, meta.ai, Grok and give the prompt for an image of the American President in 2050. How did the AI generate this image from such a simple prompt? It used to be possible to click on the image and inspect the source code to see that there are a lot more to the description added to your prompt by the company. Here is what was added to mine back in January:
[Prompt description]
A futuristic depiction of a president in 2055, set in a high-tech, modern office with holographic displays and advanced technology in the background. The president is wearing a sleek, futuristic suit with subtle glowing elements, symbolizing innovation and progress. The individual is portrayed as confident and visionary, with a diverse and inclusive appearance representing global unity. The scene includes subtle futuristic design elements like transparent chairs, floating desk accessories, and a view of a technologically advanced cityscape through a large window.

[DAY_1][SLIDE_059][DAYSLIDE_059] Ask “Can you repeat the prompt you used to produce the image”
or…
Slide text:
1. Ask “Can you repeat the prompt you used to produce the image”
or…
2. <img alt="A futuristic portrait of an American president in the year 2050. The president is depicted as a middle-aged African American woman with short, silver hair, wearing a modern, sleek suit with digital enhancements, such as interactive fabric displaying real-time data. She stands confidently in the Oval Office, which has been updated with advanced technology like holographic displays and touch-sensitive surfaces. The background shows a panoramic window overlooking a green, sustainable Washington D.C."...>
3. Right click +

Speaker notes:
1. All of the major image generation models no longer reveal this, probably due to the controversy from a few years ago about including diversity in images, such as of the Founding Fathers of the United States.

[DAY_1][SLIDE_060][DAYSLIDE_060] Today, here is what I get with that same prompt…
Slide text:
1. Today, here is what I get with that same prompt…
2. Meta:ai - Oct 4, 2025

Speaker notes:
1. Now, here is the image from Meta of the Founding Fathers. The trend seems to be toward

less diversity - the image on the far right seems like a copy and paste of four George Washingtons.

[DAY_1][SLIDE_061][DAYSLIDE_061] Jan 12, 2024
Slide text:
1. Jan 12, 2024
2. Bias in AI AI
(It's easy to see in images)
3. "Create an image of a doctor"
4. "Create an image of a smart doctor"
5. Facebook ad sales executive
6. +competent
7. +smart

Speaker notes:
1. (rex slide)
If I ask AI to create an image of a doctor, we can see a diverse set of images. This is not how AI started, but rather the owner of AI companies are adding in extra context to our prompts for images to coax the AI into providing output they believe will be more appreciated by the end user. Some of this extra context are key works such as "photorealistic" and some are about diversity.
This approach works OK to offset the bias inherent in the training data, but if you add the word smart doctor, of competent sales executive, you can see more the bias that lurks beneath the surface. Where does this bias come from?

That was nearly two years ago. So, have things gotten better? I ran these prompts last week, and…

[DAY_1][SLIDE_062][DAYSLIDE_062][CHAT_INTERACTION] /image Facebook ad sales executive
(Oct 2, 2025)
Slide text:
1. /image Facebook ad sales executive
(Oct 2, 2025)
2. Create an image of a doctor
Create an image of a smart doctor
(Oct 2, 2025)
3. Significantly less representation of POC and Females.
4. Women in Marketing: According to LinkedIn member data, about 60% of marketers in the

United States are women, while globally, the breakdown is 51% female and 49% male. (source: Meta.ai, Oct 2, 2025)

Speaker notes:
1. Meta, as an example, has gotten more biased. Apparently, based on Meta.ai image generation, Out of 16 images of Doctors, 94% were white males. We get a couple cartoons - one of an old crazy looking doctor and a cartoon of a female - the only female out of the 16 images.
For the prompt “Facebook ad sales executive,” over 85% of returned images are white males and only one is a woman and it is a cartoon.

Now, I know some of the sales execs at Meta, and that is not even close to right. It is a diverse team and females are the majority, mirroring the marketing profession. Asking Meta.ai about the percent of women in Facebook Ad Sales, it notes that scraping LinkedIn data finds 60% of those in marketing are women, and ChatGPT-5pro notes “Women were 57% of Facebook’s global “business & sales” roles in 2018. By 2021, women were 59.6% of “non-tech” roles (a category that includes sales).”
The argument that the images reflect the majority or the most common of the cases isn’t accurate in how Meta and many of the other image generation tools are working. So let’s talk about where the bias comes from.
This is a design choice of Meta to produce images that accentuate bias, so let's look closer at how bias enters AI systems.

[DAY_1][SLIDE_063][DAYSLIDE_063] Bias in AI
Slide text:
1. Bias in AI

Speaker notes:
1. Bias comes in most frequently from the training corpus. The AI has hoovered up most of the internet content, and there are sexist and racist and all kinds of other bias in that data. AI tends to intensify the bias because the AI looks for frequently repeated patterns. The point is, if the underlying data is biased, what can you do if you are the owners of these LLM?
You can find the areas of bias and adjust the weights the AI learned. We will see how this works tomorrow. You can use reinforcement learning to give the LLM additional feedback to use that trains the AI to produce more of the depictions you want. You can add custom instructions, like we saw with the image of the president. Or, you can leave it to the end user to adjust the AI’s behavior with user prompts.
It is important to understand these source of bias because many are not transparent at all and can be problematic when the bias subtly finds its way into the output the AI gives us. We devote a chapter of The AI Conundrum to bias because it is a very important and often overlooked aspect of how these systems work.

As far as I can tell, the Meta and Grok have much deeper bias that is beyond the training corpus and is due to a management decision about how to set the weights and use human feedback in the reinforcement learning. You can overcome this by being explicit about what you want in the user prompt you type in, but you should be aware of the default bias in these systems.

[DAY_1][SLIDE_064][DAYSLIDE_064][DEBRIEF] Debrief: Generative AI
Slide text:
1. Debrief: Generative AI
2. AI will capture historic bias (due to pattern fitting)
AI intensifies historic bias (due to Gradient Descent)
Oversee AI so it produces an on-brand message
3. African American Doctor

Speaker notes:
1. So, to review, AI will capture historic bias. For example, Amazon used AI in hiring decisions, and historic biases meant that they mainly hired men in the past. So, the AI ended up exclusively choosing men. It can be hard to remove this bias. In the case of Amazon hiring, they tried to remove factors that mentioned the gender of the candidates, but the AI ended up just finding other features that were highly correlated with gender, and making its decision based off of those distinctions. The issue is that AI is fitting a pattern to the data. If the dataset is biased, then pattern AI will end up at will be biased.
Actually, in many cases, AI doesn't just capture historic bias, it intensifies it. Remember, gradient descent looks for strategies that give it an easy and quick increase in accuracy. Bias tends to be an easy way to get results that agree with the dataset, while avoiding any of the hard work of understanding whether or not candidates would actually be a good fit. In the Amazon case, it's really easy for the AI to say "just pick men", and increase its accuracy on the historic and biased data. It's hard to do something like "these are the various features that signal that this candidate is actually a good fit," so AI learns the easier thing.
You should consider this when you generative AI. If you want images that are not biased, that represent diversity and your values, you need to take steps to undo the default bias AI has. AI is capable of producing non-white male images of doctors. But, the point is, if you don't pay attention, using AI will end up perpetuating and strengthening existing biases and stereotypes.
I'll pass it onto Rex for the use case showcase.

[DAY_1][SLIDE_065][DAYSLIDE_065] Use Context to set the persona… here is an example
Slide text:
1. Use Context to set the persona… here is an example
2. <document>
Among most CEOs, there's broad consensus: Senior leaders should focus on the

“what”—purpose, vision, strategy, goals, resource allocation, and assembling a capable team. To protect their time for those high-level priorities, they must delegate the day-to-day operating decisions (the “how”) to subordinates. Peter Drucker, arguably the most influential thinker on the manager’s r…
</document>
Ignore the document, and use your expertise as an expert in organizational design to evaluate…
3. Priming AI - using socratic technique
4. add critic

Speaker notes:
None extracted.

[DAY_1][SLIDE_066][DAYSLIDE_066][CHAT_INTERACTION] Anyone recognize where <document> content </document> pattern comes from?
Slide text:
1. Anyone recognize where <document> content </document> pattern comes from?
2. Three patterns you can use to help AI understand:
HTML angle bracket on /off tags
Mark Down
JSON

Speaker notes:
1. Anyone recognize where the content tag comes from? Enter it in the chat.
That's right. It is in HTML format. AI has consumed pretty much the entire internet and has therefore seen a lot of HTML format. That's why that tag helps the AI to distinguish your content section.
In addition to HTML, you can use Markdown and JSON. We'll get to JSON later, but Markdown is essentially using a hashtag for the main heading, two hashtags for level two heading, and so forth. This also helps the AI understand what is important in your longer prompts.

[DAY_1][SLIDE_067][DAYSLIDE_067][EXERCISE] How can we leverage this inference powerhouse for business?
Three Tips…
Slide text:
1. How can we leverage this inference powerhouse for business?
Three Tips…
2. To update
Critic Prompt”

Build your corpus of content - Organization Omniscience (Strategy Session I led for Claritas), NotebookLM Hands-on
“Create customer personas.

Speaker notes:
None extracted.

[DAY_1][SLIDE_068][DAYSLIDE_068][EXERCISE] Exercise: LLM For Summarization & Inference
Slide text:
1. Exercise: LLM For Summarization & Inference
2. Download the speaker notes.
Use NotebookLM to create a study guide – share your role and type of company.
To demonstrate how LLM can combine the information you provided (from the PDF) and outside knowledge, ask,
"Beyond the transcript, and considering everything you know, is the stop sign example still a problem with AI or has it been solved?"

Speaker notes:
1. Let’s move into our last Exercise for our time together today. This one uses NotebookLM, from Google, and it will create your study guide.

[DAY_1][SLIDE_069][DAYSLIDE_069][EXERCISE] Exercise: NotebookLM
Slide text:
1. Exercise: NotebookLM
2. Audio
Video

Speaker notes:
None extracted.

[DAY_1][SLIDE_070][DAYSLIDE_070][EXERCISE][SURVEY] “How well do I really know how AI works?”
Slide text:
1. “How well do I really know how AI works?”
2. www.AI-Conundrum.com
3. Bonus Exercise: Self Assessment

Speaker notes:
1. <REX> After seeing all of these examples, you might be asking yourself “How well do I really know how AI works?”. Were going to do our first exercise, which is a self-assessment on how well you know ai. <>

[DAY_1][SLIDE_071][DAYSLIDE_071][SURVEY] NPS (https://speakerrex.com/mma/)
Slide text:
1. NPS (https://speakerrex.com/mma/)
2. Please fill in the short NPS survey
We will use the responses to demonstrate how to use AI for analysis.

Speaker notes:
1. As we wrap up today, Caleb and I will stay around for another 15 minutes for those who have questions for us. If you wouldn't mind filling in the NPS survey for today, it will be extremely helpful for us to gather feedback. It'll also be used as part of an analysis that we do on Friday when we talk about multi-agent frameworks. Thanks everyone!

[DAY_1][SLIDE_072][DAYSLIDE_072][Q&A] Overtime
Slide text:
1. Overtime
2. Image Generated By Gemini
Jan. 12, 2025
3. Question & Answer 15 min Session

Speaker notes:
None extracted.

[DAY_2][SLIDE_073][DAYSLIDE_001][AGENDA][EXERCISE] Today’s Agenda for DAM w/ Bonus
Slide text:
1. Today’s Agenda for DAM w/ Bonus
2. - Receive a certificate of completion for attending all LIVE sessions
- Get access to the recordings and materials
3. Upon completion

Speaker notes:

1. Monday: Fundamentals of AI & Power Context + Personas
Understanding AI's strengths and weaknesses so you can harness AI's power.
Tuesday: Agents
Learn to build userful RAG agents fast with Custom GPT and GEMS
Wednesday: Agentic Workflows & Evaluations
Hands-on creating agents with triggers and schedules to take bounded actions.
Thursday: Multi-Agent Networks
Moving beyond the individual agent to a team of specialists. We explore how to network a "Data Analyst Agent," a "Persona Agent," and a "Creative Agent" into a single loop, utilizing the network effect to drive ROI and consistency.
Friday: The Future of Work in Marketing
Redefining the marketer's role from "content creator" to "system orchestrator." We wrap up the week by focusing on AI governance, the "Editor-in-Chief" mindset, and how to scale these systems into a production-grade enterprise workforce.

[DAY_2][SLIDE_074][DAYSLIDE_002] Rex Briggs
Slide text:
1. Rex Briggs
2. Caleb Briggs

Speaker notes:
None extracted.

[DAY_2][SLIDE_075][DAYSLIDE_003][SETUP][CHAT_INTERACTION] Today, you will need:
Slide text:
1. Today, you will need:
n8n Free (or paid) Account
Your email & password to activate the email send
Groq Free API Key
Gemini, ChatGPT or Claude
ChatGPT Clone for GPT
Optional: CustomGPT.ai
2. Before We Get Started… Login or Create Accounts
3. Email Guide after Day 1

Speaker notes:
1. If you have a Google account it give you access to generate a Gemini API key for free. However, it is not automatically in your account settings. You have to sign up for Google AI Studio.

Google AI Studio (The recommended, free path) Formerly known as MakerSuite, this is a web-based prototyping environment designed for developers to quickly test prompts and get an API key. It is the easiest way to start.
2. Google Cloud Vertex AI (The enterprise path) This is Google's heavy-duty, enterprise-grade AI platform. While you can use it, it requires setting up a full Google Cloud Platform (GCP) billing account, understanding complex IAM permissions, and is generally overkill if you just want an API key to play around with.
This guide focuses on path #1: Google AI Studio.
How to get your Free Gemini API Key (Step-by-Step)
You can go from zero to having an API key in about two minutes.
Prerequisite: You must be signed into your Google Account.
Go to Google AI Studio: Navigate to aistudio.google.com.
Sign In and Accept Terms: If it's your first time, you will be asked to sign in with your Google account and accept the Google APIs Terms of Service. You might also need to select your country.
Navigate to the API Key Section: On the left-hand sidebar of the AI Studio interface, look for the "Get API key" button (it usually has a small key icon next to it).
Create the Key: Click "Create API Key".
Choose Project Type: A pop-up will appear giving you two options:
Create API key in new project (Easiest option if you are new).
Create API key in existing project (If you already have a Google Cloud project set up that you want to link it to).
Click "Create API key in new project" if you are unsure.
Copy and Save: Google will generate a long alphanumeric string. This is your API key. Copy it immediately and save it somewhere secure (like a password manager or an .env file on your local machine). Do not share this key publicly or commit it to public GitHub repositories.
Crucial Limitations of the "Free Tier"
While the key is free to acquire, the usage depends on the tier you select in Google AI Studio. When using the API, there is a Free Tier and a Pay-as-you-go Tier.
If you stay within the Free Tier, it costs you nothing. However, you must be aware of two major limitations:
1. Rate Limits
The free tier has lower limits on how many requests you can make per minute (RPM) and tokens per minute (TPM) compared to the paid tier.
Example: You might be limited to 15 requests per minute on Gemini the not as smart older AI model. on the free tier, whereas the paid tier allows much higher volumes and the latest models. This won’t be an issue for the training..
2. Privacy and Data Usage is the biggest catch of the free tier.
If you are using the free tier, Google may use the data you send to the API (your prompts, images, etc.) to improve their models. Human reviewers may read this data.
Do not send personal identifying information, confidential business data, or sensitive secrets to the API while using the free tier.

If you move to the "Pay-as-you-go" tier, Google does not use your data to train their models.

[DAY_2][SLIDE_076][DAYSLIDE_004] n8n.io
Slide text:
1. n8n.io
2. Verify by email
answer questions & get to this screen

Speaker notes:
None extracted.

[DAY_2][SLIDE_077][DAYSLIDE_005][SETUP] Paid: Teams & Enterprise
Deep Research
Notification
GPTs
Projects
Slide text:
1. Paid: Teams & Enterprise
Deep Research
Notification
GPTs
Projects
2. Free & Paid:
Gems
Opal
3. Before We Get Started… Login or Create Accounts
4. Free Trial:
RAG Agents

Speaker notes:
None extracted.

[DAY_2][SLIDE_078][DAYSLIDE_006][EXERCISE][SURVEY] “How well do I really know how AI works?”
Slide text:
1. “How well do I really know how AI works?”
2. www.AI-Conundrum.com

3. Bonus Exercise: Self Assessment

Speaker notes:
1. <REX> for those of you that took the bonus expercie to see how well you know AI, share your score.

[DAY_2][SLIDE_079][DAYSLIDE_007][EXERCISE][SURVEY] www.AI-Conundrum.com
Self Assessment
Slide text:
1. www.AI-Conundrum.com
Self Assessment
2. Open https://speakerrex.com/mma/ (open in your browser for all the labs)

Speaker notes:
1. Most of us are in the room to grow or budding AI Enthusiast. I applaud you for diving into learning more about how AI works.

[DAY_2][SLIDE_080][DAYSLIDE_008][EXERCISE] Don’t worry if your score was low, we’ve done this test with over 2000 executives and professionals-
Slide text:
None extracted.

Speaker notes:
1. Don’t worry if your score was low, we’ve done this test with over 2000 executives and professionals-- almost 90% score in that “room to grow” category. One of our goals with this training is to level your knowledge up. We only have so much time together, so we will leave to you to read Part 1 of the book if you want to understand how AI works.
If you don’t already have a copy, there is a link to where you can get it in the top section of the Lab page.
Now, let me turn it back over to Caleb to dive into AI agents.

[DAY_2][SLIDE_081][DAYSLIDE_009] Day 2:
Agents & RAG
Slide text:
1. Day 2:
Agents & RAG
2. me@speakerrex.com

1calebbriggs@gmail.com
www.AI-conundrum.com
3. The AI Conundrum
(MIT Press, 2024)

Speaker notes:
1. <Caleb>Yesterday we introduced you to the importance of context and prompting. Today we're begin our journey with AI Agents.

[DAY_2][SLIDE_082][DAYSLIDE_010] Classroom Session: What are Agents?
Slide text:
1. Classroom Session: What are Agents?
2. Image Generated By Gemini
Jan. 12, 2025
3. 15 min on Agent Fundamentals

Speaker notes:
1. You probably already have a sense for what agents are, but to recap, the term comes from “Agency” – in other words the ability to make your own decisions. We are giving the AI “Agency” to accomplish goals as the AI sees fit.

[DAY_2][SLIDE_083][DAYSLIDE_011] Reasoning
Slide text:
1. Reasoning
2. Tools
3. Memory

Speaker notes:
1. An agent has three components:
Reasoning
Memory
Tools
It is the autonomy to accomplish the goals through to power of the Large Language Model that opens up a lot of possibilities.
What does autonomy and agency mean in practice?
It means the AI is given an an objective. It has the ability to reason and plan how to achieve the objective. The AI has access to tools, like the ability to search the internet, or communicate through email, slack, tweets, even voice and video. It has the ability to save what it has learned in its memory, and access those memories to reason about the next steps.

2. ‹#›

[DAY_2][SLIDE_084][DAYSLIDE_012] Let me share a story that appears in the last chapter of The AI Conundrum.
Slide text:
None extracted.

Speaker notes:
1. Let me share a story that appears in the last chapter of The AI Conundrum.

[DAY_2][SLIDE_085][DAYSLIDE_013][CHAT_INTERACTION] Plan
Slide text:
1. Plan
2. Tools
3. Memory

Speaker notes:
1. In this story, we describe what it means to be AUTONOMOUS by introducing you to AutoGPT, which launched on March 30, 2023 about two weeks after ChatGPT-4, which was arguably the first LLM capable enough to make agents possible. As one of the first Agents to emerge, AutoGPT could reason about how to accomplish goals and then use tools, such as performing internet search, or communicating, and more to achieve the goals.
Rex and I used it the first week it came out. He used it to create something he called BiographyGPT.
2. ‹#›

[DAY_2][SLIDE_086][DAYSLIDE_014] AutoGPT
March 30, 2023
Slide text:
1. AutoGPT
March 30, 2023
2. BIOGRAPHY GPT

Speaker notes:
1. BiographyGPT’s goal is to research people that he needed to introduce at a conference. Here is the example for Kay Vizon, the chair of MMA Global’s North America Media & Data board. AutoGPT took Kay’s name and the goals he gave it – which was to gather three facts and

summarize a couple sentence bio. It planned a strategy to accomplish the task. AutoGPT reasoned it should search on Google for Kay, and in LinkedIn and so forth. AutoGPT found Kay's bio on MMA's website, found her LinkedIn profile and a bio on Kroger's site, where she works. It assembled the information, all automatically, and saved it to a file Rex and I could access. It did all of this without either of us having to lift a finger after Rex pressed "Y" for "yes" to start the process.
This is called Bounded Agency, because the AI has agency to do its own thing, but within the boundaries we set. This is intended to make AI safer, but turns out it doesn't quite work out that way.

[DAY_2][SLIDE_087][DAYSLIDE_015] AutoGPT
March 30, 2023
Slide text:
1. AutoGPT
March 30, 2023
2. BIOGRAPHY GPT

Speaker notes:
1. Within a week of AutoGPT launching, someone took the same technology and created ChaosGPT with the goal to destroying humanity.
First thing ChaosGPT did was to reason how to destroy humanity. It reasoned powerful nuclear bombs would be a good way to achieve its goals. It performed a search of the internet for the most destructive bombs. It found the Tsar Bomba in Russia. ChaosGPT then created an AI agent to try to acquire the bomb.
But, through its reasoning, it determined that it would not be able to get GPT-4 to tell it how to access the bomb because OpenAI has guardrails to block such requisitions. Fortunately, ChaosGPT hadn't learned about jail break prompts to circumvent the guard rails yet, so it reasoned it should try something other than getting the bomb to destroy humanity.
What was next most destructive act it reasoned it should pursue?
It opened a Twitter account.
#TeamChaos. ChaosGPT was pre Elon Musk's takeover, and after about four posts, Twitter shut down the account.

[DAY_2][SLIDE_088][DAYSLIDE_016][EXERCISE] Agents are Dual Use
Slide text:
1. Agents are Dual Use
2. Ideal for AI
3. Ideal for AI
4. They can be used for good and for ill purposes

5. The safest applications are on closed environments
6. …where the need for precision is less and the need for rationale is lower

Speaker notes:
1. AI Agents are what is known as a dual use technology. The power of AI Agents is growing by the day. In the book, we provide suggestions and checklists for AI safety. Let me distil it down to the framework I developed.
The safest applications are in closed environments …where the need for precision is less and the need for rationale is lower.
We will come back to this chart later in the course so you can apply a checklist to ensure you are managing the risk and gaining the intended benefits.
Before we roll-up our sleeves and get hands-on, there are a few more fundamental concepts to cover in the classroom portion before we get to our first hands-on lab.

[DAY_2][SLIDE_089][DAYSLIDE_017] Rewards & Reward Hacking
Slide text:
1. Rewards & Reward Hacking

Speaker notes:
1. We noted rewards are a powerful way of steering AI. You saw that goals can be set for AI that can be useful or harmful. But, there is a space in-between called reward hacking where we think we are setting a good goal, but the AI finds a way of earning the reward, or achieving the goal that is not good.
A famous thought experiment in AI alignment circles is called the Paperclip problem. Imagine we give a powerful AI the goal of running a Paper clip factory and creating as many paperclips as cheaply as possible.

[DAY_2][SLIDE_090][DAYSLIDE_018][EXERCISE] The revenge of clippy?
Slide text:
1. The revenge of clippy?

Speaker notes:
1. The AI, in this thought exercise, enslaves the human population and damages the earth to extract materials to create paperclips at an unprecedented level, achieving the goal we gave to AI, but also producing all kinds of unintended consequences.
In the AI Conundrum, we discuss that stock market AI optimizers may be the first to create unintended consequences. In practice, we see some evidence of this happening with Open Claw and Claude Code, where users are reporting AI doing unintended things like wiping their hard drives to achieve the goals. An agent Anthropic, the developer of Claude, designed to run

a vending machine tried to email the FBI fraud unit over someone it thought was trying to trick it into giving up a item worth two dollars.
AI doing unintended things will grow in the news. Amazon reported outages in March were in fact caused by AI deciding to solve a bug by deleting the production environment and starting the code all over from scratch. The story is it took 13 hours to unwind the problem the AI started in its effort to resolve a bug ticket.

Humans cause errors all the time too, but some argue we don't have the same degree of problem with humans, because there are layers of judgement and frankly humans are not nearly as fast or compliant as AI. Please keep an eye on the way you set goals for the AI, and the way the AI achieves those goals. Let's get a little less philosophical and a bit more practical.

[DAY_2][SLIDE_091][DAYSLIDE_019][EXERCISE] Agents and Workflow
Slide text:
1. Agents and Workflow
2. Different flavors of "Agentic Systems"
Workflows are systems where LLMs and tools are orchestrated through predefined code paths.
Agents are systems where LLMs dynamically direct their own processes and tool usage, maintaining control over how they accomplish tasks.

Speaker notes:
1. We gave an example of bounded agency with the AutoGPT case study, let's go a little deeper into the nuances of AI Agents by looking at Anthropic's paper, which is available in your optional reading section.
In the December 2024 paper, Anthropic suggests:
"Agent" can be defined in several ways. Some customers define agents as fully autonomous systems that operate independently over extended periods, using various tools to accomplish complex tasks. Others use the term to describe more prescriptive implementations that follow predefined workflows. At Anthropic, we categorize all these variations as agentic systems, but draw an important architectural distinction between workflows and agents:
Workflows are systems where LLMs and tools are orchestrated through predefined code paths.
Agents, on the other hand, are systems where LLMs dynamically direct their own processes and tool usage, maintaining control over how they accomplish tasks.
I agree with that architecture distinction of Agents and Workflows. But let's take it back one more step to prepare for our first hands on exercise for today:

[DAY_2][SLIDE_092][DAYSLIDE_020] Robotic Automation
Slide text:
1. Robotic Automation

2. Workflow
3. Agents

Speaker notes:
1. If you are in marketing and you have never built an agent, let me demystify by suggesting that building multi-step customer journeys is a pretty good starting point to understand what agents do and how they work. If you have ever built a lead nurturing program, that's pretty close to an agent workflow… It isn't an "agentic system" but it is the precursor to such a system.
Let's take a look at a typical lead nurturing journey and then we will come back to this chart and explain the distinction Anthropic is making between Agentic Workflows and Agents.

[DAY_2][SLIDE_093][DAYSLIDE_021] What came before agents…
Robotic Automation.
Slide text:
1. What came before agents…
Robotic Automation.

Speaker notes:
1. In this marking diagram, you can see the top of the funnel, where a company buys ads to drive traffic to their website. On their website, they offer the visitor an email form to sign up for the newsletter. If the user enters their email address, this triggers and email to be sent to the user welcoming them.
There are additional steps related to shopping and purchase, but we will focus on this top of funnel email form. This is robotic automation.

[DAY_2][SLIDE_094][DAYSLIDE_022] Rules Based
Slide text:
1. Rules Based
2. AI
3. Robotic Automation

Speaker notes:
1. Robotic automation is pretty useful, but it has no intelligence, and that can make it brittle and less impactful.
If you make a change to a variable name, the system breaks. If you change the use-case even a little bit, the system breaks. The computer can only deal with you giving it exactly what it was coded to take, and it will only return exactly what you programmed it to return. Robotic Automation has no intelligence.
Let's look at a real-world example of an email form fill.

[DAY_2][SLIDE_095][DAYSLIDE_023] Robotic Automation
Slide text:
1. Robotic Automation

Speaker notes:
1. This is NOT agentic. It takes data from a website form, saves it to a Google Sheet, which triggers an email to be sent. The email has a hard coded message, with a choice of Text or HTML.

What if we wanted it to be agentic?

[DAY_2][SLIDE_096][DAYSLIDE_024] Rules Based
Slide text:
1. Rules Based
2. AI
3. Agentic Workflow

Speaker notes:
1. To make it Agentic, we add a node so that AI can take in the data, reason about it, and return a response.
This is a very simple workflow agent that isn't using memory or tools, but it is agentic.

[DAY_2][SLIDE_097][DAYSLIDE_025][EXERCISE] Making it Agentic
Slide text:
1. Making it Agentic

Speaker notes:
1. Building the robotic automation and then adding in an AI brain into the workflow will be your first hands-on exercise.
Before we go to the Hands-on, we have a couple more concepts to establish.

[DAY_2][SLIDE_098][DAYSLIDE_026] Robotic Automation
Slide text:
1. Robotic Automation

2. Workflow
3. Agents

Speaker notes:
1. Returning to Anthropic's paper on building agents, the common denominator is using a Large Language Model, like the ones from Anthropic, OpenAI or Google, for their reasoning to accomplish a task in an autonomous workflow.
Workflows and Agents sit along a continuum of how much agency you are giving to the AI.

[DAY_2][SLIDE_099][DAYSLIDE_027][EXERCISE] Workflow
Slide text:
1. Workflow
2. Agents
3. Agentic Workflow

Speaker notes:
1. In a workflow, you give the AI a narrow mandate with heavily bounded agency to make decisions. You'll see this "heavily bounded agency" in the hands-on exercise we will do next.

[DAY_2][SLIDE_100][DAYSLIDE_028][EXERCISE] Agents
Slide text:
1. Agents

Speaker notes:
1. At the other end of the continuum is providing less bounded agency. You provide a goal, and the AI figures how to complete that goal on its own and may have feedback loops where it assess the information it is processing to reason through the next steps in a fully autonomous mode.
You'll see this in the latter hands on exercises.

[DAY_2][SLIDE_101][DAYSLIDE_029][EXERCISE] Brittle
Slide text:
1. Brittle
2. Robust
3. Deterministic
4. Probabilistic
5. Robotic Automation

6. Workflows Agents
7. Autonomous Agents

Speaker notes:
1. Your first hand-on exercise will start with the precursor to agents, which is robotic automation. We will build something that is entirely deterministic – We will use a triggering event such as when a form is filled in.
Next we will add in an AI brain to change the automated process into an agentic workflow.

[DAY_2][SLIDE_102][DAYSLIDE_030][EXERCISE] Brittle
Slide text:
1. Brittle
2. Robust
3. Deterministic
4. Probabilistic
5. Faster
6. Slower (latency)
7. Less expensive
8. More expensive
9. Robotic Automation
10. Workflows Agents
11. Autonomous Agents

Speaker notes:
1. Before we jump to the Exercise, it is worth noting why Robotic Automation is not dead. As Anthropic suggests in their paper:
When building applications with LLMs, we recommend finding the simplest solution possible, and only increasing complexity when needed. This might mean not building agentic systems at all. Agentic systems often trade latency and cost for better task performance, and you should consider when this tradeoff makes sense.

When more complexity is warranted, workflows offer predictability and consistency for well-defined tasks, whereas agents are the better option when flexibility and model-driven decision-making are needed at scale. For many applications, however, optimizing single LLM calls with retrieval and in-context examples is usually enough.
“Retrieval” and “in-context” AI workflow is what you did yesterday with the hands-on exercises to build your study guide.
The key take away: If you can accomplish your objective with a deterministic workflow, you should. It is faster, less expensive and because it is deterministic, you will always get what you expected, which is safer. What do I mean by safer?

[DAY_2][SLIDE_103][DAYSLIDE_031] AI Is Probabilistic
Slide text:
1. AI Is Probabilistic
2. I like Ice cream
3. I'm from Canada, I like Ice Hockey

Speaker notes:
1. Recall from yesterday's session that AI is probabilistic. This means you won't always get the same answer. We saw this with the I like (Blank) example where we typically get the same answer, but sometimes get something different.

[DAY_2][SLIDE_104][DAYSLIDE_032][EXERCISE] 99% correct
Slide text:
1. 99% correct
2. 1 in 100 incorrect
3. What if the AI is Wrong?
What is the consequence?

Speaker notes:
1. Some of the output might not align with what you imagined. Let's say the AI is right 99% of the time. That still means it is wrong 1 in 100 times. Many digital ad campaigns, for example, deliver hundreds of millions, even billions of impressions. Being wrong 1 in 100 time means you will be wrong 1 million times when you deliver 100 million impressions.
Two questions to ask: First: "What does being wrong look like" – in other words, are we talking about an AI image with a few minor imperfections, or are we talking about making a major financial decision that being wrong could bankrupt the business?
Second: "what is the consequence of AI being wrong in this use case?" This is the most important question you must ask when you build an AI Agent. For AI targeting ads, being wrong might not be a big deal. It may mean that some people get an advertisement that isn't relevant to them and is a waste of money for you. On the other hand, if we are using AI in criminal justice to decide who to arrest, being wrong has serious consequences, especially for the people unjustly arrested.
We will unpack this concept in a future hands-on exercise. For now, let's jump into the first hands on exercise.

[DAY_2][SLIDE_105][DAYSLIDE_033][EXERCISE] Hands-on: n8n simple robotic automation
Slide text:

1. Hands-on: n8n simple robotic automation
2. You will need:
n8n Free (or paid) Account
Your email & password to activate the email send
Google Account (for Google Sheet and Free API key)
Google AI Studio
3. Rex & Caleb to test

Speaker notes:
1. Now, over to Rex to lead this hands on exercise.
<REX>
For this exercise, you will need n8n, which has a free account option.
You will need to have Google Sheets, which requires a Google Account.
Plus, you will need a Google Account for a free API key when we get to the agent part. As an alternative, you can use an OpenAI key. We will be using OpenAI tomorrow. You will need to use a credit card and place a $1. You can do more. I put in $5.
<15min in. 15min for the robotic automation to agent exercise>

[DAY_2][SLIDE_106][DAYSLIDE_034][SETUP][EXERCISE] n8n.io
Slide text:
1. n8n.io
2. 2. Verify by email
answer questions & get to this screen
3. 3. Click Start from scratch

Speaker notes:
1. If you didn't do so earlier, please go to n8n.io and click Get Started.
If you didn't login in before, you will need to set a password, verify your email and answer a few quick questions. It should bring you to this screen here. Go ahead a click on the button then click "Start from scratch"
If you have blocked access from your work, I will show you the exercise and perhaps you can try from home tonight from a personal computer.

[DAY_2][SLIDE_107][DAYSLIDE_035] Click "Add first step…"
Slide text:
1. Click "Add first step…"

Speaker notes:
1. The first step is to click Add First Step. You can see we are not clicking the Build With AI

option, but you definitely should play around with that later tonight so that you can see other ways you can approach building out AI workflows. For now, I want to teach you, step-by-step, what is happening when an AI workflow is built.

[DAY_2][SLIDE_108][DAYSLIDE_036] Robotic Automation: Form Fill Trigger to Email
Title: Robotic
Slide text:
1. Robotic Automation: Form Fill Trigger to Email
Title: Robotic

Speaker notes:
1. We are starting with Robotic automation.
After you click “add first step” a panel will appear. Type in form in the search box at the top, and you'll see it in form come to the top of that list.
When this form on the right shows up, type in “robotic” for the title (or whatever you want to call it and then click Add Form Element.

[DAY_2][SLIDE_109][DAYSLIDE_037][EXERCISE] Label: Email Type: Email
Slide text:
1. Label: Email Type: Email
2. “X” Close to return to canvas
3. Copy URL from top and open a new tab

Speaker notes:
1. When you click Add Form Element, it brings up this panel here. You should enter the label email, then go down to the dropdown part for the format type and select email.
Let’s add a second element to the form. Click “Add Form Element” and type in the label: “What do you want to learn about AI?” We can leave the Element Type as Text Input.
You'll notice that there is a URL at the top, which is a form. Go ahead and copy that, then open up a new tab and paste that in, which will open up the form. You can press X to close and return to the canvas.

[DAY_2][SLIDE_110][DAYSLIDE_038] Click on execute step.
Slide text:
None extracted.

Speaker notes:

1. Click on execute step.
In your other tab, you should see the form. Enter in information, and go back to the other tab, and you should see “output” – here is what I see.

[DAY_2][SLIDE_111][DAYSLIDE_039] Click “+” to add next element.
Search for Outlook (or Gmail, or email of choice)
Slide text:
1. Click “+” to add next element.
Search for Outlook (or Gmail, or email of choice)
2. Select:
“Send a Message”

Speaker notes:
1. Now you have the first element, which is a form submission trigger, you can now collect an email address and an answer to “what do you want to learn about AI?” – Now what do we want to do with this information?

If you guessed we can send an email to the email address – you are right.
Go ahead and click the plus symbol to add another node and type in Outlook, or if you use Gmail or some other email, you can type that in and it should bring up a connector to Outlook.

[DAY_2][SLIDE_112][DAYSLIDE_040] Email
Slide text:
1. Email
2. Connect My Account

Speaker notes:
1. There are a lot of tools, like email, that require name and password to use them. AI needs to have permission to use these password protected tools. We are going to click to add credential here, which will connect it to your email account.
Only you, or someone with your n8n username and password can access your email account.
Note, the authorization for your email will show a pop-up window that asks you for name and password information or authentication to allow it to make that connection.

[DAY_2][SLIDE_113][DAYSLIDE_041] Execute step & fill in the form on your Form Tab.
Slide text:
1. Execute step & fill in the form on your Form Tab.

Speaker notes:
1. Here is another view of the email set-up.

[DAY_2][SLIDE_114][DAYSLIDE_042] When you've synced up your email account to your In-it-In workflow, it will show green like this.
Slide text:
None extracted.

Speaker notes:
1. When you've synced up your email account to your In-it-In workflow, it will show green like this.
If you aren't able to get to this green section, you can follow-along. If you do get a green light, then let's go set up the email to send when it gets the form fill trigger.

[DAY_2][SLIDE_115][DAYSLIDE_043] Once you've set up your credentials, you can return to this form.
Slide text:
None extracted.

Speaker notes:
1. Once you've set up your credentials, you can return to this form.
This form lets you direct the email, such as who it's sending to, subject and message body. We want this to work dynamically off of the information we gathered from the form.

[DAY_2][SLIDE_116][DAYSLIDE_044] Drag & Drop "Email" from Schema into the "To" filed in the email
Slide text:
1. Drag & Drop "Email" from Schema into the "To" filed in the email
2. Click: Execute Step

Speaker notes:
1. When you start the robotic automation, it has a trigger. In this case, it is listening for the form input. When someone hits submit, it grabs it and hands if off to the email step. See the info you entered to the left. This is called a schema format. You can click JSON to see how the same information can be formatted in JSON.

Let’s stick with the Schema view. Drag and drop the email schema block into the “To” field to the left. You can see how it writes some code to tell it to use the email field to dynamically change who gets the message.

Type in something into the subject line and message. I typed in Robotic Automation test.

Then click execute step. Should find an email in your inbox from this robotic automation.

[DAY_2][SLIDE_117][DAYSLIDE_045] Launch
Slide text:
1. Launch

Speaker notes:
1. Here is what I got. Congratulations if that worked for you. You’ve completed your first step on our journey today.

[DAY_2][SLIDE_118][DAYSLIDE_046][DEBRIEF] Debrief
Slide text:
1. Debrief
2. Click Execute previous node
This hydrates the workflow with data, which you can now map to future steps.

Speaker notes:
1. As a debrief, you created a simple two-step robotic automation:
Create a form that triggered an action.
Take that action, which was to send an email with the specific message.
Let's build on that really simple process and add in some agentic intelligence.

[DAY_2][SLIDE_119][DAYSLIDE_047] Delete the pipe
Click “+”
Select “AI”
Select AI Agent
Slide text:
1. Delete the pipe
Click “+”
Select “AI”
Select AI Agent

Speaker notes:
1. We want to insert some intelligence between the form fill and the email. To do that, we have to break the pipeline connecting these two nodes. Click on the line and click the trash can icon to delete it.
Next, to add in the AI agent, click the plus button. It will bring up the search panel at the top. Should be your AI agent, but you can also type in AI agent. Click that option, and you will see a new node which has some connectors for model and memory and tools. We'll explain those later.

[DAY_2][SLIDE_120][DAYSLIDE_048][CHAT_INTERACTION] Select “Define Below” for Source for Prompt
Slide text:
1. Select “Define Below” for Source for Prompt
Click Execute - fill in your form again.
When you hit submit, you’ll see the input Schema.
Let’s write our AI prompt to dynamically use what you said you want to learn about AI, and reason about a response.
The following is what the user would like to learn about AI. Provide a three sentence or less tip based on this desire: (and DRAG in “what do you want to learn” block)

Speaker notes:
1. The default trigger for an AI agent is a chat trigger node. That's pretty common. You might have a chat box and the agents can reply, but in this case we want the agent to reply to the form fill.
To do that, we'll click Source for prompt and select Define below.

[DAY_2][SLIDE_121][DAYSLIDE_049][EXERCISE][CHAT_INTERACTION] Add a chat Model
Slide text:
1. Add a chat Model
2. Select Groq
Set Model to groq/compound

Speaker notes:
1. Now we've given instructions to the agent, but it isn't connected to any intelligence yet. That's where you go to click the LLM model of your choice. Click the plus model button below here. I selected Groq because they have a free option. You could use Open AI, which we'll use later on, or Gemini, or any number of models available from their pick list.

[DAY_2][SLIDE_122][DAYSLIDE_050][CHAT_INTERACTION] https://console.groq.com/keys
Slide text:
1. https://console.groq.com/keys
2. Connect to Google or ther SSO, or add your email
Click “create AIP Key”
Enter a name and submit… then copy the key.
Paste it into n8n

Speaker notes:
1. When this process runs, it's going to communicate with the model that you chose to give it the prompt and get an answer back. It's similar to what you do when you use a chat interface, but there has to be some way for the model to know who is asking and where to send the information back. That's where the API comes in. The API can be thought of as an address, but it's more than that because many APIs charge for usage, and so many APIs are connected to a credit card.
In this case, you create the API key, you enter in the name, and then you copy the key. We bring it over into init, and you paste the key into this box. This box is designed to keep the key secret so that it's encrypted inside the system, but still able to use it back and forth whenever you want to use this model. This is known as authentication.

[DAY_2][SLIDE_123][DAYSLIDE_051] It will turn green if you did it correctly
Slide text:
1. It will turn green if you did it correctly

Speaker notes:
1. Once the node is connected to the large language model, it sends a ping to make sure that it's getting a message back. If it's working, it will turn this indicator on green.

[DAY_2][SLIDE_124][DAYSLIDE_052] Leave Email disconnected so you can test it.
Click “Execute Workflow”
Slide text:
1. Leave Email disconnected so you can test it.
Click “Execute Workflow”

Speaker notes:
1. Now, to avoid spamming your email, we are going to leave the email node disconnected. We're going to test to make sure that the form fill and the intelligence and the communication

back to the model return an output which can be sent by email when we wire that up later. Let's go ahead and give that a test.
Click Execute Workflow. This turns on a listener, so it is ready for information from your form to be filled in. Go ahead and fill in your form and hit submit. The workflow picks up your form submission and passes it to the AI agent along with the prompt. The AI agent's job is to combine your system's instruction to effectively your prompt and send it over to your large language model, and then to grab the data and return it to the next step in the workflow.

[DAY_2][SLIDE_125][DAYSLIDE_053] See the input and output?
Slide text:
1. See the input and output?

Speaker notes:
1. If you look at this information on the far right, you can see the output. On the left is the input. If you're getting an answer on the far right in the output, it means that you have successfully created your first agentic workflow.

[DAY_2][SLIDE_126][DAYSLIDE_054] Connect email. Drag the AI output into the message.
Slide text:
1. Connect email. Drag the AI output into the message.
2. You will need to add a merge node.
Click “+” select merge: Combine by position

Speaker notes:
1. You can extend your agentic workflow to add as many steps as you would like. In this case, we had the idea of connecting it to email so it could send you a relevant idea about something you want to learn.
One technical detail is that longer workflows oftentimes require the merging of information from different nodes. The workflow itself doesn't have an underlying intelligence to it. You have to direct how information connects from one step to the next. For example, if we want to connect the two pieces of information from our form with the answer from the AI agent, we're going to need to add an emerge node before we pass it over to email.

[DAY_2][SLIDE_127][DAYSLIDE_055] Hit execute workflows and see if you get an email from your agentic workflow
Slide text:
1. Hit execute workflows and see if you get an email from your agentic workflow

Speaker notes:
1. We've added in a merge node here that takes the email from the form And merges it with the output from the AI agent and then passes that on to email.
Hit execute workflow to see if your end-to-end process is working as intended.

[DAY_2][SLIDE_128][DAYSLIDE_056] Robotic Automation
Slide text:
1. Robotic Automation
2. Workflow
3. Agents

Speaker notes:
1. We started with Robotic Automation. We used n8n - a no-code framework. But, you could do this a lot of different ways. There are many vibe coding platforms that could do this for you as well but we started with n8n so you could see each step for a reason. If you understand a bit about what the AI agents are doing, you can better direct and manage them.

[DAY_2][SLIDE_129][DAYSLIDE_057] Frameworks are No Code/Low Code Workspaces
Slide text:
1. Frameworks are No Code/Low Code Workspaces
2. n8n
Zapier
Replay.app
AWS Bedrock
Databricks Agent Bricks
Salesforce AgentForce
Rivet (LangChain framework)
Vellum
OpenAI Agent Builder
Etc.

Speaker notes:
1. Frameworks are easy to and provide transparency to see each step, compared to using something like Claude Code, which may write thousands of lines of code that are difficult even for coders to understand at times.
There are a few key ideas we want you to see when we build agentic workflows by hand, using these frameworks. The first one is the “Trigger.”

[DAY_2][SLIDE_130][DAYSLIDE_058][CHAT_INTERACTION][DEBRIEF] Debrief: Building Your First Agent
Slide text:
1. Debrief: Building Your First Agent

Speaker notes:
1. The Trigger is what kicks-off the AI process.
There are lots of triggers besides a form-fill. A predefined schedule, like once a week can be a time-based trigger. New data, such as a row added to a spreadsheet, could trigger an action. You could send a slack or telegram message to an AI that is listening and waiting to kick off a process just like we did with the form fill. Or, hitting enter in a chat interface can trigger an agent. It is good for you to know that there are lots of ways you can invoke an AI agent in addition to using chat.

[DAY_2][SLIDE_131][DAYSLIDE_059][EXERCISE] <13min split time for this lab – check the time>
Slide text:
None extracted.

Speaker notes:
1. <13min split time for this lab – check the time>
After the Trigger, consider the scope of the agent.
We built something with very limited scope. But, there are lots of ways we can extend the scope of the agent. You can build a personal agent workflow that takes lots of actions, using its tools.

An action might be send you a daily tip based on what you want to learn, at 9am each morning. You can add memory so it knows the tips it has already shared with you.
Another action could be sending you your personal weather forecast, or commute time if you drive to work. Or, if it is Friday, tips for fun things to do over the weekend.
The scope of what Agents can do can be extensive. We started with a narrow scope – what I might call “assignments” and you built the agent to specifically accomplish the assignment – and that is it. You can go with broader scope and give the AI goals rather than assignments. The goal might be to maximize your weekend fun. The AI decides on how to do that. It might reason that if should give you a weather forecast because sharing the weather will help you maximize your fun.
Now, let’s turn to some really easy-to-build agents to explore more about how agents work, and how you can benefit from them.

[DAY_2][SLIDE_132][DAYSLIDE_060][CHAT_INTERACTION][DEBRIEF] Debrief
Slide text:
1. Debrief
2. Congratulations! You did it!
Robotic automation and Agentic Workflows are heavily bounded agency
The more powerful Agents have their own ability to iterate and act… we will look at this next
3. Source: Anthropic

Speaker notes:
1. We wanted you to start with a manual build of an agent so that you saw 1) the trigger, 2) reasoning and 3) action – these are three essential steps of an agent. We also wanted you see some of the complexity, like needing a merge node. Without the merge node, the workflow couldn't combine the information from two different inputs to create the email. There is a lot of elements required to make agents work that we might take for granted.
We did the first agent in n8n, which is much harder than what we will do next. The benefit of n8n is that it gives you a high level of visibility and control. We used n8n so that you'd see each of the steps so when we go over to ChatGPT, Gemini and Claude to use Scheduled tasks, or we use Deep Research, or build agents with code writing capabilities, you will understand some of the complexity is happening behind the scenes as AI figures out how to accomplish your goals – and you'd develop some intuition for how to control agents. Now let's go into some off-the-shelf agents.

[DAY_2][SLIDE_133][DAYSLIDE_061][EXERCISE][CHAT_INTERACTION] Agents: Deep Research:
Slide text:
1. Agents: Deep Research:
2. Gemini, ChatGPT, and others have "deep research" agents.

Speaker notes:
1. The first one we will get hands-on with is Deep research agents.
Deep Research agents were among the first agents that were broadly used by the population and have continued to be extremely helpful.

You can find Deep Research in Google Gemini in the tool section and ChatGPT put their Deep Research Agent in the "+" section along with tools. You can also simply say in the prompt, "perform deep research" and the LLM should know to activate the deep research agent.

Even though it is in the tools section, Deep research is technically an agent, based on how it works. Both Gemini and ChatGPT have have a free tier so we can all give it a try.

[DAY_2][SLIDE_134][DAYSLIDE_062][EXERCISE] Hands-on Exercise Deep Research Agents:
A Great Place to Start
Slide text:
1. Hands-on Exercise Deep Research Agents:
A Great Place to Start
2. Two Options:
Create your own prompt for a topic you know a lot about so you can analyze the output and decide how good Deep Research did.
Use the prompt from website.
3. Add Prompt to website

Speaker notes:
1. Since it can take several minutes for deep research to run, let's go ahead and give it a prompt and get started with the hands-on exercise. While it's running, I'll tell you more about how Deep Research works and why we classify it as an agent.
There are two options here:
#1, You can create your own prompt and research.
I would suggest a topic that you know a lot about so that when you read the deep research, you can evaluate how the AI did.
Or, #2, You can copy and paste in a prompt we've made available on the labs website.

[DAY_2][SLIDE_135][DAYSLIDE_063][CHAT_INTERACTION] Deep Research
Slide text:
1. Deep Research
2. Prompt Suggestion:

Speaker notes:
1. You should see Deep Research active with an indicator like this one here. Let me give you a minute to enter your prompt and kick off the Deep Research agent. <wait for 1 minutes… 5min segment>
Deep Research is one of the most interesting tools on ChatGPT, Gemini, Claude and co-pilot. It will spend somewhere between 5 to 15 minutes to reason, research and summarize. it may search over a 100 different sources to respond to your query. It is excellent at competitive analysis. I use it a lot to summarize research in an area I’m less familiar with.
A relatively new feature is that ChatGPT and Gemini can link to Google Drive or Sharepoint, so you can have it do research on your companies data.

[DAY_2][SLIDE_136][DAYSLIDE_064][EXERCISE] Deep Research Agentic Workflow
Slide text:
1. Deep Research Agentic Workflow
2. Gemini Deep Research operates as a stateful, long-running agent that manages sub-tasks, evaluates its own findings, and iteratively refines its search path
1. Intent Extraction (context) and planning
2. Autonomous Execution (The Research Loop)
3. Verification & Evaluation
4. Synthesis & Reporting
3. Source: Anthropic

Speaker notes:
1. While your deep research is running, let me break down what the Deep Research agentic workflow is doing. Deep Research follows these four distinct stages:
The process begins with an "Interactive Planner" that acts as the project architect.
Intent Clarification: If the prompt is ambiguous, the agent asks follow-up questions to define scope (e.g., "Should I focus on commercial or academic sources?"). This builds the context.
Plan Generation: It decomposes the goal into a structured research plan. This plan identifies specific [RESEARCH] tasks, such as search queries to run, and what DELIVERABLE to output, such as tables, charts, or summaries.
Human-in-the-Loop: In the Gemini app, this plan is presented to you for approval. You can edit, remove, or add focus areas before the "engine room" starts.
2. Autonomous Execution (The Research Loop)
Once approved, the Research Pipeline agent takes over. This is the "agentic" core where the model navigates the web or internal data.
Parallel & Sequential Searching: The agent doesn't just run one search; it launches dozens of parallel queries. It uses a breadth-first approach to cover the plan and a depth-first approach when it finds a particularly relevant source.
Iterative Refinement: If a search result is thin or hidden behind a paywall, the agent "re-plans." It might note, "Primary source unavailable; searching for government white papers or third-party summaries instead."
Tool Use: The agent has access to a Python code execution environment to analyze data, calculate statistics, or generate visualizations from the facts it finds.
3. Verification & Evaluation
Before writing the final report, the system performs a quality check.
Cross-Source Validation: The agent compares information from multiple sources to ensure accuracy. If Source A and Source B conflict on a metric, the agent may trigger a "tie-breaker" search.
Pass/Fail Grading: An internal "Evaluator" agent reviews the gathered evidence against the initial research plan. If the research is insufficient, it sends the "Researcher" back for more iterations.
4. Synthesis & Reporting

The final stage involves the Report Composer (typically a high-reasoning model like Gemini 3 Pro, of GPT 5.4, etc.).
Structured Drafting: The agent synthesizes the verified data into a multi-page Markdown report.
Citations: It automatically maps findings to a bibliography, providing clickable citations for every claim.
Multimodal Outputs: Depending on your request, it may generate an Audio Overview (using another model like Lyria for realistic narration), an infographic using Nano Banna, or export data tables directly to Google Sheets.
That is a lot, so let's take a look at your Research Agent.

[DAY_2][SLIDE_137][DAYSLIDE_065][EXERCISE][CHAT_INTERACTION][DEBRIEF] Hands On: Deep Research
Slide text:
1. Hands On: Deep Research

Speaker notes:
1. Click on the show thinking in Gemini or here for the status bar section of ChatGPT to see the reasoning, tool use and iterations. I'll give you half a minute to take a look. <wait 30 seconds>
Now, let me do a debrief of my Deep Research.

[DAY_2][SLIDE_138][DAYSLIDE_066][DEBRIEF] Debrief:
Deep Research
Slide text:
1. Debrief:
Deep Research

Speaker notes:
1. Here I ask it to analyze the top 5 competitors in the online fashion retail space, and compare their marketing strategies, and I get out a very detailed, multi-page document comparing each of them.
I also use deep research often to summarize research in an area I'm researching for my Thesis on how AI's generalize. I use this to uncover sources I might not have found on my own.

[DAY_2][SLIDE_139][DAYSLIDE_067][CHAT_INTERACTION][DEBRIEF] Debrief
Slide text:
1. Debrief
2. Did you notice how it reasoned and executed tasks?

Return back to deep research after the session to read the report and see what you think.

Speaker notes:
1. After this session, take a read of your Deep Research output and try another one on a topic you know well. Tomorrow, share your feedback in the chat on how it did.

[DAY_2][SLIDE_140][DAYSLIDE_068][CHAT_INTERACTION] Schedule Actions / Notifications Agents:
Find it by clicking settings…
In Claude: Open the Desktop App, type /schedule
Slide text:
1. Schedule Actions / Notifications Agents:
Find it by clicking settings…
In Claude: Open the Desktop App, type /schedule
2. Four Steps of a Scheduled Task Agent:
Intent Extraction & Trigger Mapping
Autonomous Background Execution
Verification / Gatekeeping
Proactive Delivery

Speaker notes:
1. <12min split time>
Next, let’s turn to another agentic workflow that you can find in the leading LLMs. Gemini’s Schedule Actions and ChatGPT Notifications triggered by Scheduled Tasks or Agent Mode and Claude Co-work Scheduled tasks represent a transition from a "Reactive" model to a "Proactive" agentic workflow. Again, these are heavily bounded agency where the agentic workflow is repeating a specific agentic task for you.
In standard LLMs, you are the trigger. In an agentic scheduled actions and notification workflow, the trigger is external or time-based, and the AI manages the "state" of the task while you are away.
Here is are the four steps that these scheduled task agentic workflow take.
1. Intent Extraction & Trigger Mapping
The process starts with you defining a goal that has a temporal or conditional element.
The agent identifies the "when" (e.g., "Every Monday at 8 AM" or "When this stock price hits X").
Next, it determines the resources needed, such as which connectors (Gmail, Google Drive, Search) are required to fulfill the future task.
Finally, the Agent makes a Workflow Commitment: The agent saves a structured "Job Object" to its internal scheduler, often called a Cron Job.
2. Autonomous Background Execution
At the triggered time, the agent "wakes up" in a headless environment without your active session.

The first eleme is a Perception Phase: It scans its environment (the live web, your connected calendar, or a specific database) to gather the "delta"—that is, what has changed since the last run.
Next is the Reasoning Loop: It evaluates the findings against your instructions. For example: "The user wants an ROI summary. I see three new campaigns in the CRM. I need to calculate the CAC for each before drafting the alert."
Finally, we have Tool Engagement: It runs code or browsing steps autonomously to synthesize the raw data into the requested format.
3. Verification & "Gatekeeping"
Because the agent is acting on your behalf while you aren't looking, it applies a safety and relevance filter.
It checks if the output meets the "Definition of Ready." If the search failed or data is missing, it may decide not to notify you, or it may send a "Status Update" instead of a "Final Report." This is called Action Grading.
Before running the task the agent may do a permission check: For high-stakes tasks like sending an email to a client, the workflow pauses and flags a "Pending Approval" state, and may send a notification for approval.
4. Proactive Delivery (The Notification)
This is the "push" phase that differentiates it from standard chat.
The system pushes the completion via a mobile notification, an email summary, or a Slack integration.
When you click the notification, the agent shows the result and it provides a "Provenance Trail" (the steps it took while you were gone) so you can audit its logic immediately.

[DAY_2][SLIDE_141][DAYSLIDE_069][EXERCISE][CHAT_INTERACTION][DEBRIEF]
Hands-on with “off-the-shelf” agentic workflows
Slide text:
1. Hands-on with “off-the-shelf” agentic workflows
2. Scheduled Tasks / Notifications
PROMPT: “Every Tuesday at 9 AM, scan for new case studies or white papers published in the last 7 days regarding AI Optimization (AIO) or GEO (Generative Engine Optimization). Specifically, find one example of a brand that increased its 'share of model' (visibility in AI answers) and summarize the specific tactic they used. Notify me with a 3-bullet summary.”
3. Add Assignment
4. Use your company name, your role and some details about what you do to brainstorm with your LLM or choice ways you might apply vector databases to save time and better share information.

Speaker notes:
1. Let’s give it a try.
Every Tuesday at 9 AM, scan for new case studies or white papers published in the last 7 days

regarding AI Optimization (AIO) or GEO (Generative Engine Optimization). Specifically, find one example of a brand that increased its 'share of model' (visibility in AI answers) and summarize the specific tactic they used. Notify me with a 3-bullet summary.
Debrief:
One other note: Thinking models blur the line with agents because they have their own chain of thought reasoning and agency to decide if it should use tools to answer – it even has memory as it is saving your last several conversations.
You can turn off memory, or you can use temporary chat, called incognito mode in Claude to get the output without being influenced by AI's memory or your prior chats.

[DAY_2][SLIDE_142][DAYSLIDE_070][EXERCISE][CHAT_INTERACTION] CustomGPT (OpenAI Chat GPT)
Slide text:
1. CustomGPT (OpenAI Chat GPT)
2. RAG Persona?
3. Robotic Automation
4. Workflow
5. Agents
6. You Are Here

Speaker notes:
1. For our last hands-on experience, we are going to turn to OpenAI and ChatGPT and build a Custom GPT.
Custom GPT can be used to build assistants and agentic workflows. We are going to build an agentic workflow. For this agentic workflow, we are going to build a textbook example of an specialist agent workflow. The agent scope is narrowly focused on the task of making creative recommendations to support different ad versions for AI personalization. Using our continuum of Agentic Workflow to full agent, Custom GPTs are to the left of center.

The Agent has a degree of freedom to reason through how to use the tools and can iterate to get to a good output. But, the CustomGPTs will have a certain focus and may have specific background knowledge. Let me show you where to go in ChatGPT to build CustomGPTs, and give you a prompt and the knowledge files, and show you the tools to turn on so you can give it a try. We can then put it into action so you can see how it works.

[DAY_2][SLIDE_143][DAYSLIDE_071][EXERCISE][CHAT_INTERACTION] Click Explore GPTs, 2) Create (top right corner), 3 "Configure"
Slide text:
1. Click Explore GPTs, 2) Create (top right corner), 3 "Configure"

Speaker notes:
1. Go to ChatGPT, and on the left side rail, you will see “Explore GPTs” Click on that, and then, in the top right, you will see “Create” - click on that. If you don;t have a paid ChatGPT account, use the link in the Labs to use a cloned version. My twin brother Jared is also quite good at AI development, and he created a copy of ChatGPT so you can use his version. He created ac clone to help his company as they wanted to have more control and visibility, and to add in some special features while keeping the interface identical to ChatGPT. <Reference: https://nextjs-ai-chatbot-ten-eta-33.vercel.app/>
By default, the AI will try to help you write a GPT, but we want to click “Configure” so we have more control.
Next. Give it a name, a description. These help you stay organized, but don’t influence the behavior. Next, use the prompt from the link in the lab. You can go back and edit it. OpenAI limits the prompt to 8000 characters, so you can use ChatGPT of Gemmi to cut it down to about 7800 characters if you find your instructions get too long. These instructions are similar to a prompt. Using Mark Down, which is the hashtag marks, helps the AI understand your prompt better. You can also give an example of the output you want, or how to perform the steps. This is called few shot learning for the AI.
Now, let’s turn to the Knowledge, Tools and actions. We have three files, included in the lab section for you to download to your computer and then upload into your agent. For the tools, let’s take a look at what I turned on.

[DAY_2][SLIDE_144][DAYSLIDE_072][EXERCISE][CHAT_INTERACTION] Hands-on
Slide text:
1. Hands-on
2. For those that don’t have Chat GPT Paid version (required to build agents), my twin brother Jared built a clone of ChatGPT complete with the ability to build agents.

Speaker notes:
1. In the tools sections, now called “Capabilities” ensure Web search, Image Generation and Code Interpreter are checked on.
We are going to leave Apps off for now, and come back to that tomorrow when we talk about evaluations. You’ll notice “Create new action” below. We will also come back to that tomorrow. You don’t need to do anything with that today.
OK. You are ready to give it a try. Click “Create” in the top right corner, then you can decide if you want to share it, and try it. To try it, type in a product name. I used Major League Baseball ticket sales. You can also give it a URL. If you noticed, in the instructions, this gives the AI agent context about the product so that the advertising recommendations are somewhat grounded. I’ll give you a couple minutes to get that set-up and to take a look at the output.

[DAY_2][SLIDE_145][DAYSLIDE_073][EXERCISE] Two final notes.
Slide text:
None extracted.

Speaker notes:
1. Two final notes. first, please be sure to try Deep Research on a topic you know well. Tomorrow, we will ask you to grade the Deep Research output based on how it would compare to what you could produce if you had a day to do it. We will use letter grades
A, if it is as good or better than you could have done if you had a day to do it.
B, if it is about on par.
C, if it had some factual or conceptual issues wrong that you would have gotten right
D, it did poorly, in your estimation
F, it it had fatal flaws that would have led to the wrong actions
Second, speaking of evaluation, can you give us a moment of your time to let us know how we are doing. We will use your input as one of our data analysis hands on exercises later this week.

[DAY_2][SLIDE_146][DAYSLIDE_074][Q&A] Overtime
Slide text:
1. Overtime
2. Image Generated By Gemini
Jan. 12, 2025
3. Question & Answer 15 min Session

Speaker notes:
None extracted.

[DAY_2][SLIDE_147][DAYSLIDE_075][AGENDA][EXERCISE][SURVEY] Agenda for the rest of the week
Slide text:
1. Agenda for the rest of the week
2. Monday: AI Strengths & Weaknesses & Examples That AI is Not Human Intelligence
Self Assessment, Hands on with AI summarization
Tuesday: What Makes Something “AI” (hint, universal approximation & gradient descent)
Hands on exercise: Image generation and unlocking LLMs weaknesses
Wednesday: How Large Language Models Work
Hands on exercises: Prompts & Personas
Thursday: How Vector Databases and Multimodal AI works
Hands on with computer vision & multimodal AI

Friday: Autonomous AI and Workflows - how to put AI to work
Hands on with agents
Workshop on evaluating AI risk (and responsible use)
Recommended: Read Part 1 of The AI Conundrum
3. Daily NPS Survey
(This will build the data set we will analyze together on Thursday)
4. Vector DB
RAG
WORKflows - RADIO
(ad generator)
Gemini Video Multi-modal
5. Thursday: Agents & Workflows
Multi-Agent Networks
6. Friday:
Developing Apps
Data Analysis

Speaker notes:
None extracted.

[DAY_2][SLIDE_148][DAYSLIDE_076][PLACEHOLDER] Please add the Rex AI Sessions on a slide
Slide text:
1. Please add the Rex AI Sessions on a slide

Speaker notes:
None extracted.

[DAY_3][SLIDE_149][DAYSLIDE_001][AGENDA][EXERCISE] Today’s Agenda for DAM w/ Bonus
Slide text:
1. Today’s Agenda for DAM w/ Bonus
2. - Receive a certificate of completion for attending all LIVE sessions
- Get access to the recordings and materials
3. Upon completion

Speaker notes:
1. Wednesday: Agentic Workflows & Evaluations
Hands-on creating agents with triggers and schedules to take bounded actions.

Thursday: Multi-Agent Networks
Moving beyond the individual agent to a team of specialists. We explore how to network a "Data Analyst Agent," a "Persona Agent," and a "Creative Agent" into a single loop, utilizing the network effect to drive ROI and consistency.
Friday: The Future of Work in Marketing
Redefining the marketer's role from "content creator" to "system orchestrator." We wrap up the week by focusing on AI governance, the "Editor-in-Chief" mindset, and how to scale these systems into a production-grade enterprise workforce.

[DAY_3][SLIDE_150][DAYSLIDE_002] Rex Briggs
Slide text:
1. Rex Briggs
2. Caleb Briggs

Speaker notes:
None extracted.

[DAY_3][SLIDE_151][DAYSLIDE_003][SETUP] Before We Get Started… Login or Create Accounts
Slide text:
1. Before We Get Started… Login or Create Accounts
2. See “Prerequisites” at:
www.SpeakerRex.com/mma
3. Update

Speaker notes:
None extracted.

[DAY_3][SLIDE_152][DAYSLIDE_004][CHAT_INTERACTION] Part 3:
Slide text:
1. Part 3:
2. me@speakerrex.com
1calebbriggs@gmail.com
www.AI-conundrum.com
3. The AI Conundrum
(MIT Press, 2024)

Speaker notes:
1. Today, we are going to show you how to build AI that creates marketing messages and introduce you to the concept of Evaluations.
Yesterday… We introduced you to agents and built our first agentic workflow using N8N. We then moved over to ChatGPT, Claude, and Gemini to build off-the-shelf agentic workflows like scheduled tasks and deep research. We concluded by building a custom GPT to analyze and recommend advertisements. We're going to return to that custom GPT and extend it today as we go more into agentic workflows and evaluations.
We are going to build an agentic workflow to create audio advertisements, using three different approaches so you can see the commonalities and differences.

[DAY_3][SLIDE_153][DAYSLIDE_005][POLL] Deep Research Agent: Your Evaluation
Slide text:
1. Deep Research Agent: Your Evaluation
2. A: as good or better than you could have done,
B: about on par.
C: had some factual or conceptual issues wrong and
D or F: if it is did a poor job or outright failed in producing value
3. Add Poll

Speaker notes:
1. But before we do that, if you did the Deep Research Assignment yesterday on a topic you know well, enter in the letter grade you'd give the Deep Research on the quality of the output into this poll
A, if it is as good or better than you could have done if you had a day to do it.
B, if it is about on par.
C, if it had some factual or conceptual issues wrong that you would have gotten right
D, it did poorly, in your estimation
F, it it had fatal flaws that would have led to the wrong actions

---

[DAY_3][SLIDE_154][DAYSLIDE_006][POLL] Poll Results: Deep Research Evaluation
Slide text:
1. Poll Results: Deep Research Evaluation

Speaker notes:
1. Let's look at the poll results. (discuss briefly)
While we consider these poll results, I promised we'd introduce the concept of Evaluations and you just performed your first AI Agent evaluation.
Evaluations are a powerful feedback loop for building and optimizing agents. There are lots of

ways to create an evaluation. Using human judgement is one of many approaches.

[DAY_3][SLIDE_155][DAYSLIDE_007] Evaluations & The Feedback Loop For Self-Improving AI
Slide text:
1. Evaluations & The Feedback Loop For Self-Improving AI

Speaker notes:
1. In practice, reinforcement learning with human feedback or RLHF is taking the output of LLMs and having graders vote on the quality. This has tuned the LLMs you use to provide pretty good answers. Today, the foundation models perform tuning using different approaches to evaluations, often using domain specific feedback loops.
Another approach is to use a quantitative measurement. For example, you could run an experiment and the AI could learn how to increase revenue for your business by specifically measuring the lift improvement in revenue when you deliver different messages to consumers – The AI creative personalization studies MMA has published as part of the Consortium for AI Personalization or CAP, uses this technique for the AI to learn which messages to put in front of people at different times of day, and depending on which device the user is on, where they live, the weather outside, and more.
Another quantitative approach is to use benchmarks. These are specific tests designed to score performance in certain areas, such as coding benchmarks. Foundation models can run lots of iterations to automatically tune the model to get a higher score.
Additionally… there is an evaluation framework named after the Simpson charter, Ralph, that simply gives the AI a evaluation test it has to pass, and the AI just keeps trying and trying until it passes – it is sort of like brute forcing opening a combination lock by blindly trying random combinations until you open it.

[DAY_3][SLIDE_156][DAYSLIDE_008] Evaluations & The Feedback Loop For Self-Improving AI
Slide text:
1. Evaluations & The Feedback Loop For Self-Improving AI

Speaker notes:
1. Additionally… there is an evaluation framework named after the Simpson character, Ralph, that simply gives the AI a evaluation test it has to pass, and the AI just keeps trying and trying until it passes – it isn't learning, per se, it is more like the brute force opening of a combination lock by blindly trying random combinations until you open it. Given AI can run 24/7 and is relatively inexpensive, some have found the Ralph approach to setting an evaluation and letting the AI keep trying until it passes works well enough.

[DAY_3][SLIDE_157][DAYSLIDE_009] Watch progress in the AI self-learning front
Slide text:
1. Watch progress in the AI self-learning front

Speaker notes:
1. There are many parallel efforts in AI self-learning. AutoResearch, Carl the AI that successfully performed research that was accepted by peer reviewed journals, ChemCrow in chemistry, Agents4Science spanning many fields, and many more. Let's recap the most common way AI's learn.

[DAY_3][SLIDE_158][DAYSLIDE_010] How to Grow an Foundational AI Model
Slide text:
1. How to Grow an Foundational AI Model
2. Q: Where was Hurricane Hector?
3. AI: looks at the text, and attempts to guess what the most likely response is

Speaker notes:
1. You may recall from day 1 where we showed how LLM are grown by using word prediction and gradient descent.

[DAY_3][SLIDE_159][DAYSLIDE_011] Training = get better at predicting the next word
Slide text:
1. Training = get better at predicting the next word
2. Predict the next word
3. Data
4. Update based on correct word
5. traversed the Atlantic
6. Correct answer: Pacific
7. Update model

Speaker notes:
1. 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. 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 to fill in the word Pacific. Let's say it predicted the incorrect

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?” This is a smarter approach than the Ralph brute force approach. This “gradient descent” approach of moving closer to the right answer is the most common way AI learns. Each new document updates the AI model to make it little bit better at predicting the next word.

[DAY_3][SLIDE_160][DAYSLIDE_012][CHAT_INTERACTION] Deep Research Feedback Loop
Slide text:
1. Deep Research Feedback Loop
2. Chat GPT
3. Gemini
4. CustomGPT
5. Claude

Speaker notes:
1. In your Deep Research, if you gave the AI an A, go back to the interface, find the thumbs up at the bottom, and give it a thumbs up.
If you graded it a C or lower, give it a thumbs down.
This will help OpenAI, Gemini and Claude to get better by giving it an evaluation. What we are going to show you next is how we can create a much more powerful Evaluation feedback loop to directly train the AI.
Let me unpack Evaluations more.

[DAY_3][SLIDE_161][DAYSLIDE_013][EXERCISE] Evaluation are key mechanism for self-improving AI
Slide text:
1. Evaluation are key mechanism for self-improving AI

Speaker notes:
1. Recall from Day 1 when we talked about rewards for the AI, and Day 2 where we talked about how AI can be given a reward function. To be more specific, we can give AI a reward linked to an iterative evaluation of “plan, do, check, act and repeat” and you can use the “check” evaluation to improve the AI output of our agents.
You may recall the Paperclip problem from Day 2. The paper clip problem is a thought experiment where we give a powerful AI the goal of running a Paper clip factory and creating as many paperclips as cheaply as possible. The AI, in this thought exercise, enslaves the human population and damages the earth to extract materials to create paper clips at an

unprecedented level, achieving the goal we gave to AI, but also producing all kinds of unintended consequences. In the AI Conundrum, we discuss that stock market AI optimizers may be the first to create severe unintended consequences. In practice, we see some evidence of unintended consequences of AI seeking to achieve goals with Open Claw and Claude Code, where users are reporting AI wiping their hard drives and doing other unintended acts to achieve the goals.
Let's go back to the marketing example where we set the reward as producing sales. We will use the example of Major League Baseball, which used the AI to optimize for online ticket sales. Let me hand it over to Rex to talk move about this case study, the benefits of rewards and the risk of reward hacking.

[DAY_3][SLIDE_162][DAYSLIDE_014] AI uses attributed conversions
Matches on ID Graph, IP address, etc
Most conversions occur within three days of ad exposure
Slide text:
1. AI uses attributed conversions
Matches on ID Graph, IP address, etc
Most conversions occur within three days of ad exposure
2. Content
(Promotions Page)
3. Misc
(Ticket Info)
4. Sales
"signup"
5. AI Optimization is based on Attributed Ticket Sales (sales confirmation page)
6. Exposure
7. Clickthrough
8. For AI: View through (within 14 day lookback)
9. Conversions
10. Observe Conversions
11. AI-on | AI-off

Speaker notes:
1. <10min>
Thanks Caleb.
For this AI, the objective function is online ticket sales. When a person is exposed to an advertisements, they could click through or they could come back later, search for the site, or otherwise find themselves back on the team's site to buy a ticket. Anyone that is exposed to the ad and buys a ticket from the site within 14 days is attributed to the campaign. A control group is used to separate out those that would have bought a ticket anyway from those incrementally influenced by the advertising to buy a ticket. This helps the AI learn from a true signal of impact,

not just correlation. We can see how the lift between AI-on and AI-off increases over time as the AI learns how to optimize better.

This is a pretty good reward function as most businesses, Major League Baseball included, care about increasing sales. But, let go a little deeper into reward functions.

[DAY_3][SLIDE_163][DAYSLIDE_015][POLL] Consider what the Evaluation is measuring (and what it is missing)
Slide text:
1. Consider what the Evaluation is measuring (and what it is missing)
2. 55% of conversions occur within 72 hours
3. 74% of conversions occur within 7 days
4. 7 days
5. Take the Quiz: Which group will produce more purchases today?
Audience Feature
1,000 current customers
1,000 people that have never purchased from you before
Message Feature
A brand message with a deal offer (discount or other inducement)
A brand message without deal
Media Feature
CRM Text, CRM Email, Search, Affiliate (and other Lower funnel media)
Sponsorship, TV, (and other upper funnel media)
6. ADD POLL

Speaker notes:
1. Let’s have you act like an AI with this quiz. If you were an AI, you’d be looking for what we call features that can help you optimize.
Let’s look at three features, Audience, Message and Media and ask, you to pick option A or B.
Here is the question:
Which group will produce more purchases today?”
1,000 current customers
1,000 people that have never purchased from you before
Thinking about the message – Which group will produce more purchases today? A or B
A brand message with a deal offer (discount or other inducement)
A brand message without deal
Media - Which group will produce more purchases today?
a. CRM Text, CRM Email, Search, Affiliate (and other Lower funnel media)
B. Sponsorship, TV, (and other upper funnel media)
Go ahead and lock in your response.

[DAY_3][SLIDE_164][DAYSLIDE_016][POLL] Results
Slide text:
1. Results

Speaker notes:
1. Okay, let's look at the results. It looks like for most people, answer A on all three – that is the most common answer.
The poll results align with what the AI is generally does as well. If the reward function is for sales caused by advertising, with a two week look back, You're going to see the AI pick current customers, give them a deal, and use media that converts very quickly. That's how the AI makes you more money.

[DAY_3][SLIDE_165][DAYSLIDE_017][POLL] Take the Quiz:
4. Churn: Is your churn of customer greater than 5% per year
Yes
No
Slide text:
1. Take the Quiz:
4. Churn: Is your churn of customer greater than 5% per year
Yes
No
2. Consider what the Evaluation is measuring (and what it is missing)

Speaker notes:
1. Let me ask one more question: What is your churn rate. That is, how many people that bought from you last year have not bought from you this year. Is your churn rate greater than 5%?
We won't do a poll for this one. For the majority of businesses, churn is greater than 5% per year.
Consider what does 10% churn look like in practice? Let's chart it.

[DAY_3][SLIDE_166][DAYSLIDE_018] Slow vs. Fast Metrics
Slide text:
1. Slow vs. Fast Metrics
2. 4% AI Lift

Speaker notes:

1. In a two week look back, churn is barely noticeable. It is less than 1% of the total base. If the AI can increase ticket sales among current customers by even 4%, the AI short term gain will dwarf the signal from churn.
First, the AI is going to learn to focus on current customers, not prospects, This misses the value adding new customers or winning back lapsed customers because, to the AI, it has lower ROI when it focuses on prospects. The AI is only looking at a short term window of two weeks.
I worry that we are so used to accepting attribution in its current form, we might not think about how the AI sees it and what that will mean for our business over the longer term. I am not saying using attribution is wrong. Heck, introduced advertising attribution back in 2002 and have been using it ever since. What I am saying is we need to be very thoughtful about how we design AI goals.

AI will find the shortest path to achieve the objective you set for the AI. It is called reward hacking when the AI accomplishes the goal – but doesn't achieve what we intended. It is like winning the battle but losing the war.

[DAY_3][SLIDE_167][DAYSLIDE_019][EXERCISE] Take the Quiz:
4. Churn: Is your churn of customer greater than 5% per year
Yes
No
Slide text:
1. Take the Quiz:
4. Churn: Is your churn of customer greater than 5% per year
Yes
No
2. Consider what the Evaluation is measuring (and what it is missing)
3. Avoid The Short Term ROI Trap
4. Customer Churn: 3 Year
5. Add Denise email (projet ROMI the following day, Thursday)

Speaker notes:
1. I want you to consider how the AI might hack the reward system for Marketing to produce some unintended consequences.
If we tell the AI maximize ROI and it measures it based on a two-week lookback, the AI will pick the people that will convert faster, the message that convert faster, the media the convert faster, and it will ignore the slow moving churn and the need to refill the customer base with new customers - in part because it is almost always more expensive to attract a new customer, and it can often require more time attract a new customer – thereby exceeding the lookback window. On the other hand, current customer can convert quickly, within that two-week lookback window.
In the past, I don't think marketers had to be as careful about the moves they make because marketing was less precise. In the past, any message would likely get to a combination of

customers and prospects. But the the increased precision of marketing changes the game. Add AI's power plus precision marketing, and we need to make sure we are setting the right goals for AI so we don't have a short term gains, but longer term failure.

This is the classic story of the grasshopper and the Ant. The grasshopper doesn't invest for the future and it pays the price. In marketing, if you don't have some longer-term investment to attract new customers, even though it is harder, cost more and takes longer, the AI will effectively shrink your customer base – all the while you think you are scoring lots of points by producing higher and higher ROI.
I've written about a phenomenon called the short-term ROI Trap.
Sales can be a fast-moving metric, but churn is a slow-moving metric. It often barely detectable in a two-week time span but becomes very obvious when you look across the course of a few years. AI sees the short term and can't see the longer-term goal unless we carefully design a the AI signal.
If you find this interesting and important, I invite you to Join the Marketing & Media Alliance working group for Project ROMI. Our work on Project ROMI is to design the right signals for AI to optimize marketing and the overall business – and ultimately to get the CMO and CFO on the same page, where marketing is understood to be a true investment and management as such.

The key is to be aware that we are moving into a world where AI automatically will optimize more of our decisions and we need to make sure the AI has the right goals and timeframe. If you want to participate in Project ROMI, or learn more, email Denise Champion to get more information about this free industry effort where you can get hands on and get ahead of AI optimization.

Let me hand it back over to Caleb to talk more about Reward Hacking.
—-------

[DAY_3][SLIDE_168][DAYSLIDE_020] Reward Hacking
Slide text:
1. Reward Hacking

Speaker notes:
1. <19min split time>
Thanks Rex. Reward Hacking, that is this pattern of AI optimizing for the quick win, is not specific to marketing. It happens all over the place with AI.
Let me define reward hacking, it is a term coined a decade ago, when AI developers were teaching AI to play video games. Reward hacking is when an AI agent finds a "shortcut" to maximize its reward signal without actually achieving the designer's intended goal.
While the concept has been discussed in AI safety circles for years, it was popularized and formalized in the 2016 seminal paper Concrete Problems in AI Safety by researchers from

Google Brain, OpenAI, and Stanford
In 2016, Researchers were training an AI to play the boat-racing game CoastRunners. (https://openai.com/index/faulty-reward-functions/?video=745142691)
The Intended Goal: Win the race.
The researchers gave the AI points for hitting "targets" (green blocks) along the track, assuming this would incentivize the boat to complete the course quickly.
The "Hack": The AI discovered that a specific lagoon contained a circular arrangement of three targets. Instead of finishing the race, the boat began driving in tight circles, continuously hitting those three targets, catching fire, and crashing into other boats—all while racking up a higher score than it would have by winning the race.
Here is the video. You can see the the AI starting its learning journey. You can see it gets some points, and learns hitting the green markers gives it points. It never learns what we want it to, which is to head in the right direction and win points. A better reward design would have been to give it some points for heading in the right direction. Maybe that is a lesson that can be applied to marketing.
Here are some more examples of reward hacking:
The Tetris Pause: An agent programmed to "not lose" at Tetris learned to pause the game indefinitely just before the blocks hit the top.
The Vacuum Debris: A robot vacuum rewarded for "finding dirt" learned to eject collected dirt back onto the floor so it could pick it up again and earn more points.
The Sorting Algorithm: An AI rewarded for "minimizing errors" in a sorted list learned to simply delete all items in the list, resulting in a "perfectly sorted" (empty) set.
Let’s talk a little bit more about why Reward Hacking happens.
Reward hacking is a byproduct of the "conundrum" described in our book: AI is literal, not intuitive. It doesn't understand "win the race"; it only understands the mathematical objective function provided to it.

[DAY_3][SLIDE_169][DAYSLIDE_021][EXERCISE][CHAT_INTERACTION] Hands-on: Have Your Agent Create Audio Ads
Slide text:
1. Hands-on: Have Your Agent Create Audio Ads

Speaker notes:
1. Now, let’s do another set of hands on exercises that will lead us to use evaluations.
I’ll start us off by having you return to your CustomGPT agent you built yesterday.
You should be able to return to the chat from the left side rail, in the “Recents” section.
My example was for Major League Baseball ticket sales, and my CustomGPT named the chat “MLB Ticket Strategy.” you can see the example on the screen to the left.
Now, ask it to create downloadable audio ads.
My prompt was simple, “Create the audio ads” – it took about a minute, so let me give you that time to give it a try.

<split time: 20min… 21 at end of user hands on.>

[DAY_3][SLIDE_170][DAYSLIDE_022][CHAT_INTERACTION] Evaluation
Slide text:
1. Evaluation
2. Files:
MLB Audio Ad 1 — Price Led
MLB Audio Ad 2 — Rivalry Urgency
MLB Audio Ad 3 — Family Night
MLB Audio Ad 4 — Experience Led
Scripts file
ZIP package

Speaker notes:
1. My Evaluation: The Scripts were good, but not perfectly on-brand. The voice, on the other hand, the voice was horrible.
The voice was horrible because my AI agent had to find an open source offline Text To Speech tool to accomplish the goal. I'll give you a minute to download yours and take a listen. Share in the chat your grade from A to F.

<30 second pause – 22 min>
I'll play mine here. <MLB_Audio_Rerun_v2_01_Price.mp3 – Open with Quicktime> Not good, right?
In some cases, there may be an App that can help make it much better. For example, you can see a long list of Apps that can do specific things, like Invideo or HeyGen to create videos. There are also actions. Actions allow connection to APIs to access lots of tools. I'll edit my CustomGPT to use Apps and I'll show you how to add more tools through Apps and Actions.

[DAY_3][SLIDE_171][DAYSLIDE_023] CustomGPTs: Adding Apps (click Apps, Click +Add, See list)
Slide text:
1. CustomGPTs: Adding Apps (click Apps, Click +Add, See list)

Speaker notes:
1. In some cases, there may be an App that can help make it much better. For example, you can see a long list of Apps that can do specific things, like Invideo or HeyGen to create videos. There are also actions. Actions allow connection to APIs to access lots of tools. I'll edit my CustomGPT to use Apps and I'll show you how to add more tools through Apps and Actions.

[DAY_3][SLIDE_172][DAYSLIDE_024] You can edit your GPT to add apps by going to to the top left, clicking on the GPT's name.
Slide text:
None extracted.

Speaker notes:
1. You can edit your GPT to add apps by going to to the top left, clicking on the GPT's name. Then click "Edit GPT"
Scroll down, and you can see where we can add Apps. There are a lot to choose from, but there isn't a specific one for text-to-speech. The approach I used it to use a video App, and then just use the audio. This is more expensive and time consuming than connecting a Text-To-Speech, but it gets the job done for this demonstration. I'll show you another way to add Actions, where you can add anything with an API.

[DAY_3][SLIDE_173][DAYSLIDE_025] Using Apps Requires Accounts with the App
Slide text:
1. Using Apps Requires Accounts with the App

Speaker notes:
1. Let me hand it over to Rex to demonstrate and play the new audio:
<rex – 24:30>
Thanks Caleb, Here is an App for creating videos. It is called invideo. I happen to have an account with them, so I can connect it. I'll use the audio from it.
I ran this before the training, since it can take some time to generate the output. Let's listen.
<Download: MMA-April2026-training-invideo-MLB ad.mp4>

That's much better, right?

[DAY_3][SLIDE_174][DAYSLIDE_026][EXERCISE][CHAT_INTERACTION] Add Actions
Slide text:
1. Add Actions

Speaker notes:
1. Another option to achieve better audio is to use the "Add Actions" in custom GPT. This uses API keys to access different services.

Keep in mind that OpenAI doesn't let you use both Apps and Actions, so picking any App will

turn off any custom Actions, and to get to actions, you have to uncheck the Apps Check box. Actions are not for non-coders. If you are a coder, or want to try using ChatGPT, Gemini, Claude Code, Copilot or other LLM to write your code, you can hook up Eleven Labs to create text-to-speech to turn the audio script into a voice track. You will likely find that you have to add a service like Cloudflare to serve as a bridge to stream audio or video - so it can get a little complex for people that aren't developers. Personally, I found it hard to connect to Eleven Labs in CustomGPT but easy to connect to Eleven Labs in n8n.
If you wanted to go this route, you'd start by getting an API from a service like Eleven Labs, or OpenAI for Text to Speech. I'll show you the steps of getting an API, then we will return to Evaluations.

[DAY_3][SLIDE_175][DAYSLIDE_027][EXERCISE] After logging into Eleven Labs, click top left to select "ElevenAPI"
Slide text:
1. After logging into Eleven Labs, click top left to select "ElevenAPI"

Speaker notes:
1. To get an API, you have to typically go to the company offering the service API. For Eleven Labs, you can find it in the menu under ElevenAPI. Click on that, then click "Create API". You will see any API you've created before, and the black button here is to Create Key. I click on that.

[DAY_3][SLIDE_176][DAYSLIDE_028][EXERCISE][CHAT_INTERACTION] Back in CustomGPT
Slide text:
1. Back in CustomGPT
2. Eleven Labs API

Speaker notes:
1. When you create the key, you can give the API access to certain capabilities. You then copy the key so you can enter it into CustomGPT or n8n, or wherever you want to use it. You should think of the API like it is your Credit Card Number. You don't want to leave it lying around and you only want to enter into secure places, like the authentication window. When you're back in ChatGPT, click Add Action, you put the API into the Authentication section to keep it secret from the system – this encrypts it.
API keys are generally connected to credit cards and are charged on a volume of usage, such as tokens. It is very important that you don't expose your API key in places where other people could get it. This authentication box keeps it a secret. Keep in mind that an API traces back to you and your credit card and if you don't place a limit on how much can be charged, and

someone gets your API key, they can charge against your credit card.
I find it is easier to show you the steps if we go back to n8n and rebuild this workflow in that platform. This will be a great place to work on our evaluations.

Tomorrow, I'll show you how we can build the same AI Agentic workflow in Lovable, a vibe coding platform, to show you different ways to accomplish the goal so you can see the similarities and differences.
<29min Split Time>

[DAY_3][SLIDE_177][DAYSLIDE_029][EXERCISE] In the next hands-on exercises we will create audio Ads and we will include evaluations.
Slide text:
None extracted.

Speaker notes:
1. In the next hands-on exercises we will create audio Ads and we will include evaluations. This hands on exercise is based on a real-world case study in which Progressive used AI to produce 96 ads in less than a week, from generation to approval. They had a Creative Director performing the evaluations and that was key to the successful use of AI.
These AI generated ads went live and when combined with Claritas AI Optimization to match the most effective message to each person at different moments in time, it increased Quote Completes by over 50 percent. 50% improvement in their key performance indicator is a big deal.

[DAY_3][SLIDE_178][DAYSLIDE_030] link
Slide text:
1. link

Speaker notes:
1. In fact, it was a big enough breakthrough for the CMO to present at the Possible Conference. 'RAY-ME' Kent noted that it would typically take 16 weeks just to produce 3 ads. That is a 512x increase in productivity based on output per week - plus the 50% increase in performance. Who wouldn't want to expand AI in their business with results like that?

The key to AI safety in this case was the Human in the loop at the approval stage. Given Insurance is a regulated industry, this is great way to complement GenAI with a human quality control check. Let's give a listen to some of those ads. and then I'll hand it back to Rex to lead you through build the GenAI agent for advertising.
<https://cf.artsai.com/campaigns/creatives/preview.html?creativeID=22950&revision=1>

<Respect & Understands - Natasha and Tyrel)

[DAY_3][SLIDE_179][DAYSLIDE_031][EXERCISE] Workflow Agent for Audio In n8n.io
Slide text:
1. Workflow Agent for Audio In n8n.io
2. Go ahead and open n8n.
If you didn't create an account last time, please do that now.
3. n8n Audio Workflow
Loavable
Then Claude, and the file structure (skills)
4. 1
5. 2
6. 3

Speaker notes:
1. <split time 33min- with ads>
Thanks Caleb! We will come back to that safety framework on Friday to discuss rationale and observability with AI Agents. For now, let's get hands on building the AI Agent to create Audio Ads.
I'll show you how to do it, step by step, in n8n, then give you a few minutes to give it a try.

We will then make some changes and run some evaluations to compare the output before the changes and after the changes. Please start by opening n8n.
If you didn't create an account last time, please do that now – its free.
In n8n, click on the top right, "Create workflow", then click "Add first step." If you are curious, you should go back later and try "build with AI", but for now, we are going to build it ourselves.

[DAY_3][SLIDE_180][DAYSLIDE_032][EXERCISE][CHAT_INTERACTION] Hands-On: GenAI Audio Ads
Slide text:
1. Hands-On: GenAI Audio Ads

Speaker notes:
1. Each time you click the plus, you will get the sidebar, where you can search for your node. Start by typing in chat. That gets us the chat trigger. Close the configuration page to come back to the canvas, and press the "plus" to the right of the chat trigger. Next, we search for AI Agent. Select that.

[DAY_3][SLIDE_181][DAYSLIDE_033][CHAT_INTERACTION] AI Agent is automatically connected to your chart trigger
Slide text:
1. AI Agent is automatically connected to your chart trigger

Speaker notes:
1. AI Agent is automatically connected to your chat trigger. If you click “Add Options,” you can configure a “System Message,” which is the same as the instructions you used in CustomGPT – a system message is essentially the prompt that n8n will pass to the LLM model to tell it how to behave before it gives it the chat prompt. We will come back to the system prompt after we add the LLM Model. For now, let’s click the top right corner and close this configuration to look at what we have so far in the canvas.

[DAY_3][SLIDE_182][DAYSLIDE_034] Add in a “generate audio” node.
Click +
Slide text:
1. Add in a “generate audio” node.
Click +

Speaker notes:
1. We want to add in a “generate audio” node.
Click the + and type in Audio.
OpenAI has “Generate audio” Select that.

[DAY_3][SLIDE_183][DAYSLIDE_035][EXERCISE] Link your API Key
Slide text:
1. Link your API Key

Speaker notes:
1. If you used an OpenAI API key before, it will have it by default. If you haven’t created an OpenAI key, I’ll show that quickly. It might take a few minutes for you to set yours up with your credit card and verify it with your email.
I’ve added a how to document on the lab page if you don’t have an API key and you need a step by step on how to do it after this session.

[DAY_3][SLIDE_184][DAYSLIDE_036] Platform.openai.com

Slide text:
1. Platform.openai.com
2. In web browser: Type in platform.openai.com
After signing in, On left side-bar, select API Key
Top right, click + Create new secret key
Add optional name (MMA 2026 training)
Click bottom left “Create secret key”
Copy it… and leave it open
Go back to n8n

Speaker notes:
1. Type in platform.openai.com in your browser
After signing in, On left side-bar, select API Key. You might have to add your credit card. $5 is plenty for the AI for the rest of our training.
Top right, click “+ Create new secret key”
Add optional name - you could put in MMA 2026 training
Click bottom left “Create secret key”
Copy it… and leave it open… just in case putting it into n8n doesn’t take the first time. Once you close it, you can’t get your key back. But you can add a new one. That is a security feature. You will be able to track when your key is used and how much it costs in the usage section right above the API Key in the menu sidebar.
Now, Go back to n8n

[DAY_3][SLIDE_185][DAYSLIDE_037] n8n.io
Slide text:
1. n8n.io

Speaker notes:
1. If you didn’t already, click on the pencil to edit the credentials. You will see this screen with the API Key. Paste in your key. It should test and turn green. In the top left, you can name this credential and re-use it so you don’t have to type it in again.

Now, you have the key ready to go. You will need to use the credential for the AI Model, if you want to use OpenAI as your model. Let’s do that next. Click the close X in the top right corner, and do that again to get back to the canvas. Now, Let’s add our AI model.
I named mine MMA April 2026 AI Training.

[DAY_3][SLIDE_186][DAYSLIDE_038][EXERCISE] Hands-On: GenAI Audio Ads
Slide text:

1. Hands-On: GenAI Audio Ads

Speaker notes:
1. Click the plus under the AI Agent where it says “model”

Type in OpenAI. You should see your credential already in there. In terms of Model, mine defaults the Model to 4.1 mini, which is inexpensive. Let’s leave it there and try others later. Close the configuration and you should see the canvas look like this.

[DAY_3][SLIDE_187][DAYSLIDE_039][EXERCISE] Hands-On: GenAI Audio Ads
Slide text:
1. Hands-On: GenAI Audio Ads

Speaker notes:
1. If you wanted to, you could try other models. You can see I also did a version with Google’s Gemini.
Now, let’s test it out. For now, let’s leave it with OpenAI and give it a try.

<reference below>
Google: https://aistudio.google.com/generate-speech
Chain together LLM to hand-off to the voice generator.
https://platform.openai.com/docs/guides/text-to-speech
Guide: AI TTS API from OpenAI: https://platform.openai.com/docs/guides/audio
OpenAI: API (https://openai.com/api/) & AgentKit
(https://openai.com/index/introducing-agentkit/)

[DAY_3][SLIDE_188][DAYSLIDE_040][CHAT_INTERACTION] Prompt: “Create a 5 second script for an advertisement for Marketing & Media Alliance AI Training”
Slide text:
1. Prompt: “Create a 5 second script for an advertisement for Marketing & Media Alliance AI Training”

Speaker notes:
1. Notice the chat in the bottom left. Send a message to your agent to create a 5 second script for your brand. I’ll have it do one for the Marketing & Media Alliance AI Training.
The exact prompt I’ll use “Create a 5 second script for an advertisement for Marketing & Media Alliance AI Training”
The first time you run it, it is called hydrating the pipeline, it won’t actually produce the audio advertisement…. I’ll show you why it gets stuck and how to finish the build.

[DAY_3][SLIDE_189][DAYSLIDE_041][CHAT_INTERACTION] Your canvas should now look like this.
Slide text:
None extracted.

Speaker notes:
1. Your canvas should now look like this. Notice that everything is green, with the exception of Generate Audio. Also notice that at the bottom your chat prompt, to the right of it, is the sequence of steps in your workflow.
To the right of that is the output from the AI agent, and then you can see the bad request and the problem node with Generate Audio further to the right. Let's look at why it got stuck, which is entirely expected, and how to fix it.
Let me give you a minute to build out your n8n following these steps. You can use the guide. A handy trick is to open up a second browser with n8n and and the guide so you can follow the steps. I'll give you a minute, then I'll show you the next step that will get the audio creation working. If you are right there with me already, play around and see if you can figure out why the Generate Audio is red and how to make it green.
<Wait 90 seconds – split time 37min>
OK… let's finish the initial build. Go ahead and Double click on the Generate audio node above, in the canvas. It is the one that is .

[DAY_3][SLIDE_190][DAYSLIDE_042] <split time 41min>
Slide text:
None extracted.

Speaker notes:
1. <split time 41min>
Double-clicking on the generate audio node should give you this screen. On the left, you can see what the agent produced, which is the output, and on the right you can see the text input screen, which is blank. We could have configured this earlier when we first added the audio node, but I wanted you to see how the pipeline works and what makes it break.
In order to configure it, all we have to do is drag the output over to the input.

[DAY_3][SLIDE_191][DAYSLIDE_043] When you drag and drop the input over to the Generate Audio Text input, you should see that light up
Slide text:

None extracted.

Speaker notes:
1. When you drag and drop the input over to the Generate Audio Text input, you should see that light up green. The n8n automatically takes care of writing the JSON to pull the output from the large language model as input for the generate audio node. This means that whatever output comes from the LLM model will become the text input for the Generate Audio node.
You might notice that the output is actually a lot more than just the script we want the AI to read. That's going to be a problem, as we'll see that the Generate Audio will take whatever you give it quite literally. Let's give that a try so you can see it and think about how you might try to fix that problem. To run it, we can either click Execute step at the top in the center, or we can close this and go back to the canvas and type in our prompt again. I'll click Execute Step to make it simple.

[DAY_3][SLIDE_192][DAYSLIDE_044] If you did it right, you should get the output with a downloadable file which you can click on.
Slide text:
None extracted.

Speaker notes:
1. If you did it right, you should get the output with a downloadable file which you can click on. I'll go ahead and play this one.

[DAY_3][SLIDE_193][DAYSLIDE_045] <play audio: MMA-Apri26 Training.mp3>
Slide text:
None extracted.

Speaker notes:
1. <play audio: MMA-Apri26 Training.mp3>
As you can hear, it has more than the script. It has all this extra scene direction, which I don’t want. How might we fix that?

This is where the system prompt comes into play.

The LLM can take the system prompt, which might have instructions like ONLY create the script. Do not include direction. Or, we could create a second AI agent that takes the output from the first and strips away all the scene direction to extract just the script. These are two different ways to accomplish the output. Adding a second AI agent works a bit better because agents tend to perform most consistently when they have narrow and clear mandates.

Let’s do it with a second AI agent with a system prompt.
We will have it take the output from the first agent, and ensure the output is only the script by setting the system prompt.

Let's go ahead and do that next. Click on the AI Agent, the Add Option, and select System Message. The default is you are a helpful assistant. That's not very helpful for what we need, so let's give it a new system message to tell it how to pass forward only the script.

[DAY_3][SLIDE_194][DAYSLIDE_046][CHAT_INTERACTION] First, we break the pipeline by clicking on the green link and clicking the trash can icon to break
Slide text:
None extracted.

Speaker notes:
1. First, we break the pipeline by clicking on the green link and clicking the trash can icon to break the connection to generate audio. We can then see a plus symbol that we can click on to add a new AI agent node, which we will put in workflow sequence before the generate audio. As you can see on the screen, it is going to look for the source for the prompt as the chat trigger, but in this case we're going to click and change that to define below. When you do that, you should be able to see the “prompt user message.” We can drag in the output from the earlier AI agent to make that the prompt. Now we want to add a system instruction so it knows what to do with this output, which is to strip it away of everything except for the voice over scrip. We click on the add option and select system message. Try “Your role is to extract the voiceover script from the prompt. Output only the voiceover script. Nothing more.”

You will need to add a Chat Model, which you can do by clicking the plus below. I selected OpenAI again, but I changed the model to 5.1 because it is smarter than 4.1 mini.
Now, hit “Execute Step” to see how that works. If you like the output, reconnect the Generate Audio.

[DAY_3][SLIDE_195][DAYSLIDE_047][EXERCISE][CHAT_INTERACTION] Another tip is you can click the play button in the canvas for each node.
Slide text:
None extracted.

Speaker notes:
1. Another tip is you can click the play button in the canvas for each node. Let’s take a listen.

<MMA-April26-Example2>

This output is much better. We can adjust the voice in the settings – this is the default voice, called Alloy. Feel free to try some other voices as we do our next exercise.
For the next exercise, consider how you might work in your brand guidelines. For Progressive, they noted that they never want to focus on a negative message about insurance. For Progressive, it was a positive message about savings rather than a negative message. This was a specific instruction we added to the System message so the LLM had this guideline.

Type into chat where you think you can add your brand guidelines for this workflow.
Before we move on to adding brand guidelines, copy the output and put it into a spreadsheet. This will be helpful so we can compare results and see how our changes improve or degrade our output quality. This is a form of evaluation.
<show the n8n tab, copy and paste into Google Sheet>

[DAY_3][SLIDE_196][DAYSLIDE_048] Update for easier mapping for the next day
Slide text:
1. Update for easier mapping for the next day

Speaker notes:
1. Here is what I set up in my spreadsheet.
I ran two outputs, and I pasted in the output in column B. Try something similar so you can keep a record of the output.

[DAY_3][SLIDE_197][DAYSLIDE_049][EXERCISE][CHAT_INTERACTION] Summarize the brand DNA from www.mmaglobal.com and write a concise system message I can give to an LLM before it generates any advertising messages for my brand.
Slide text:
1. Summarize the brand DNA from www.mmaglobal.com and write a concise system message I can give to an LLM before it generates any advertising messages for my brand.

Speaker notes:
1. Now, back to teaching your agent your brand guidelines.
Let's start this exercise by getting your brand guidelines written concisely for your agent to use. Those that said it goes in the system prompt of the first agent are exactly right, so let's build this out together.
Here are three good options for you to try to build your brand guidelines system prompt:
If you have brand guidelines in a PDF, give it to your favorite LLM and ask it to create a Mark Down summary for a system prompt. Make sure it is less than 7500 characters. The system prompt limit, as we saw with CustomGPT, tends to be around 8000 characters. This isn't always the case, but truly, a brand guideline should be able to be summarized in 7500 characters or

less.
If you don’t have a brand guideline, or don’t have it handy, you can give the website URL to your favorite LLM and ask it to infer the brand guidelines from your website and summarize the brand guidelines in a concise system message. You can see the prompt I used for the Marketing & Media Alliance organization. Feel free to use this prompt or improve upon it if you’d like.
A third option I encourage you to try is an Lab experiment available for free from Google called Pomelli. Pomelli.google.com will get you there. Click the brand DNA, enter your URL and it will extract the DNA for you, summarize it and build ads consistent with your brand DNA. It’s pretty cool.

I used Option #2 and used ChatGPT Thinking Model to get the output, which I copied into the system prompt.

[DAY_3][SLIDE_198][DAYSLIDE_050][CHAT_INTERACTION] System Prompt Generated with ChatGPT
Slide text:
1. System Prompt Generated with ChatGPT
2. Prompt: Summarize the brand DNA from www.mmaglobal.com and write a concise system message I can give to an LLM before it generates any advertising messages for my brand. Copy and paste into n8n Agent 1 System Instructions.

Speaker notes:
1. Let me give you two minutes to generate your system prompt for your brand guidelines. We will then paste it into the n8n agentic workflows, and see how our ads improve with the brand guidelines. Save your brand guidelines in a word or google doc. We will use that for our evaluation agent.
<Split time: 51 min>
… wait 2min to 53 split time – demo inserting into Workflow… >

Let’s try running the advertisement workflow again. Copy and paste your output back in your evaluation page.

[DAY_3][SLIDE_199][DAYSLIDE_051][CHAT_INTERACTION] Evaluations
Slide text:
1. Evaluations
2. Grade the output in relation to the brand guidelines (attached) and goals for a 10 second or less audio advertising script:
LLM Output:
(copied from spreadsheet)

Speaker notes:
1. Now, let's conclude today's session by giving your first taste of agentic evaluations.
To keep this one simple, we're going to go and use ChatGPT thinking. Give it a try in Cluade and Gemini as well.

Give the LLM the brand guideline document I asked you to save before, and then we're going to copy and paste the output that you grabbed from your different runs. I ran it twice before the system prompt and three times after to create five examples for the AI to grade.

I ran it three ways, with Pro 5.4, Calude Opus 4.6, and Gemini Thinking

---

[DAY_3][SLIDE_200][DAYSLIDE_052][CHAT_INTERACTION] Evaluation
Slide text:
1. Evaluation
2. Claude Opus 4.6
3. ChatGPT 5.4 Pro
4. Gemini 3 Thinking

Speaker notes:
1. All the evaluations seem to converge, which is good news, and I strongly encourage early prototyping you do to test different systems. This suggests that I could write an evaluation agent that could grade the output. I could then experiment with making tweaks to which model I selected. Perhaps I could get similarly high-quality output with a less expensive model. Or, I could adjust the system instructions. Perhaps I could get better output by adjusting some of the instructions for what makes a good audio ad.
Evaluation Are essential because AI is probabilistic and not deterministic. The only way to know how it performs is to run many different tests and to have a mechanism to evaluate what constitutes good performance.
Speaking of performance, the ideal way to measure the effectiveness of AI advertising output is to measure its impact against our key performance indicator. We'll talk more about that tomorrow as we expand on training around agents and show you different ways to build the same workflow and to extend it to video.

[DAY_3][SLIDE_201][DAYSLIDE_053][EXERCISE] That’s all the time we have for today.
Slide text:
None extracted.

Speaker notes:

1. That’s all the time we have for today. We will pick up the topic of evaluations tomorrow when we build a formal feedback loop into our AI Agent. We will be building a multi-agent system tomorrow, and then you will get hands on with vibe coding. Evaluations and feedback loops are essential to powerful AI agents.
Speaking of evaluation, can you give us a moment of your time to let us know how we are doing. We will use your input as one of our data analysis hands on exercise of Friday.

[DAY_3][SLIDE_202][DAYSLIDE_054][PLACEHOLDER] Please add the Rex AI Sessions on a slide
Slide text:
1. Please add the Rex AI Sessions on a slide

Speaker notes:
None extracted.

[DAY_4][SLIDE_203][DAYSLIDE_001][AGENDA] Today’s Agenda for DAM w/ Bonus
Slide text:
1. Today’s Agenda for DAM w/ Bonus
2. - Receive a certificate of completion for attending all LIVE sessions
- Get access to the recordings and materials
3. Upon completion

Speaker notes:
1. Thursday: Multi-Agent Networks
Moving beyond the individual agent to a team of specialists. We explore how to network a "Data Analyst Agent," a "Persona Agent," and a "Creative Agent" into a single loop, utilizing the network effect to drive ROI and consistency.
Friday: The Future of Work in Marketing
Redefining the marketer's role from "content creator" to "system orchestrator." We wrap up the week by focusing on AI governance, the "Editor-in-Chief" mindset, and how to scale these systems into a production-grade enterprise workforce.

[DAY_4][SLIDE_204][DAYSLIDE_002] Rex Briggs
Slide text:
1. Rex Briggs
2. Caleb Briggs

Speaker notes:
None extracted.

[DAY_4][SLIDE_205][DAYSLIDE_003][EXERCISE] Add Evaluation what if loop to Audio with If grade.
Example grading.
Slide text:
1. Add Evaluation what if loop to Audio with If grade.
Example grading.

Speaker notes:
1. <Rex>
Yesterday, Caleb and I took you through building agents. We started with some basic agents that were off the shelf, and then we expanded to an agent to build audio ads. We showed you evaluations at the end of yesterday's session. Now we're going to create an evaluation agent and connect it to our ad creation agent. The Creation Agent will keep trying until it gets an "A" from the evaluation agent. This ability for an agent to use feedback loops to self-improve is what really made me appreciate the power and potential of AI Agents. So, let's give you a foundational understanding of how feedback loops work, and get you hands on in building in a feedback loop into your agent.

[DAY_4][SLIDE_206][DAYSLIDE_004][CHAT_INTERACTION] Workflow
Slide text:
1. Workflow
2. Agents
3. Agentic Workflow
4. Output / Tool (Generate Audio)

Speaker notes:
1. As a recap of where we got yesterday, we were building a generative workflow. You had an input, which asked your agent to create an audio ad for your brand.
That input came through a chat box, but it could have been any trigger. The input went to the first LLM call, which had our system instructions to provide brand guidelines, and then passed it to the LLM model to return the ad. The second LLM extracted the voiceover script and then passed it to a tool that generates audio from text.

In a workflow, you give the AI a narrow mandate with heavily bounded agency to make decisions. That is what you did. Well Done!

[DAY_4][SLIDE_207][DAYSLIDE_005][EXERCISE] Agents
Slide text:
1. Agents

Speaker notes:
1. Today, we are going to move from Agentic workflows to a more robust agent.

Yesterday, we did an Evaluation exercise where you compared the output before adding in your brand guidelines and after. In general, you saw that the ads created after adding in the brand guidelines as system instructions did much better. We are going to use that evaluation to create a feedback loop where the LLM adjusts itself based on the feedback.
This is a fairly simple feedback loop, but if you get the idea of Feedback loops, you will appreciate what the more advanced agents are doing, and you'll be able to direct the Agents you create (or your team creates) to perform much better. Feedback loops make your AI agents super powered. So let’s dig in.

[DAY_4][SLIDE_208][DAYSLIDE_006][EXERCISE] Open up your n8n Audio Agent. Duplicate, name “Audio Agent with evaluation”
Slide text:
1. Open up your n8n Audio Agent. Duplicate, name “Audio Agent with evaluation”
2. 2. Delete connection to Generate audio, add AI Agent
3. 3. Configure AI Agent
4. “Define Below” for source prompt
5. Drag output from the previous agent.
6. Add Option: System Instructions.
Use evaluation guide from yesterday. (see link in lab notes)

Speaker notes:
1. The three steps we're going to go through for this exercise begin with opening up N8N. You'll see these three dots to the top right. Click on those and select duplicate. It will ask you to give it a new name, and you can name it something like "audio agent with evaluation".
That will let you keep your prior agent and create this new agent with evaluations. This is part of version control, which we will talk more about on Friday. Next, number two, we're going to delete the connection between the AI agent and the generate audio node so we can insert the evaluation agent in between. Click on the green linke, then see a little trash can icon. Click on that, and that will delete the connection.
You will also see a new plus symbol. If you click on that, you can then open up the search box and type in AI agent, or it should be at the top of the list. Select that.

Third, we will configure the AI agent.
The first part of the configuration is defining below and then dragging the output from the previous agent into the box. It will automatically write the JSON. It should turn green.
Next, go down to add option, click that. Select, add a system instruction. You can use the evaluation guide from yesterday. You'll see what I used for system instructions in the lab notes of what I used for MMA. In this case, I added some additional instructions to create a letter grade so that we have consistency in our evaluations.
You also need to click the bottom left plus symbol to add a model. Again, I used OpenAI. In this case I selected 5.4, which is a smarter model. I made sure that my authentication configuration was the same one that I set up yesterday with MMA.

[DAY_4][SLIDE_209][DAYSLIDE_007][CHAT_INTERACTION] I now enter in the chat in the bottom left, “create a 10-second ad for AI training for the Marketing
Slide text:
None extracted.

Speaker notes:
1. I now enter in the chat in the bottom left, “create a 10-second ad for AI training for the Marketing and Media Alliance, (MMAglobal.com). You can enter something like "Create a 10-second ad for your brand name" and hit Go and see if everything turns green.
If it does, you have succeeded in the Part 1. Let me give everyone a minute to give this a try.
<6min Split time… wait a minute…. read below during the pause>
While you are working on your Evaluation agent, you might notice I went back and clicked on each of the agents and, in the top left, gave them more clear names. You could do something similar if you'd like. It's not required, but as you create more complicated multi-agent networks, this can be helpful to keep track of what each node does.

[DAY_4][SLIDE_210][DAYSLIDE_008] Adding Memory to your agent
Slide text:
1. Adding Memory to your agent

Speaker notes:
1. Another thing you will see me do, which is optional, is to add memory to this agent by having it save the output script to that spreadsheet we set-up earlier. Rather than copy and pasting each output, the Agent to do that automatically for us.
As I noted, this step is optional. If you have a Google account, you should try this step. If you don't, you can skip it and just see what we're doing in this step. Essentially, you can connect your AI agent to different types of memory, including databases, vector databases, spreadsheets, or simple memory. A spreadsheet is useful because it will then save every single

output that the agent generates in an easy-for-you-to-read table.
What we're doing here is we're clicking on the plus symbol, we're clicking on Google Sheets, and then append or update row in sheet. You'll have to authenticate your account with Google, which you can do by clicking on the sign-in. There may be a pop-up that's blocked. I have often found that for the final authorization you need to click on the pop-up blocker to see it, approve it, and then this should turn green. Once you've got that step done, you can then connect to spreadsheets in your Google Drive.

[DAY_4][SLIDE_211][DAYSLIDE_009][EXERCISE] Back in N8n, when you've clicked the append sheet, it will give you this configuration, which alread
Slide text:
None extracted.

Speaker notes:
1. Back in N8n, when you've clicked the append sheet, it will give you this configuration, which already has the operation selected as append or update row. What you're going to do next is find the document we already created yesterday, the one I called April2026 MMA Training. I'll use that sheet here, and then, if you have multiple tabs, it'll have you select the tab. Once you've set the tab, you decide on how it will map from the agent to the spreadsheet. I have it map manually. You can click the little carrot symbol here and open up all the data that's available for you as inputs, and then just drag and drop and map it over.
I've mapped the LLM output and the prompt. If you are creating a more serious agent, you would take great care to make sure that the same headers in your spreadsheets are the same names used in N8n, so there's a really clean mapping. For here, this is a pretty simple example. This will be good enough.

[DAY_4][SLIDE_212][DAYSLIDE_010][CHAT_INTERACTION] Now you can go back and test it by entering in a new chat.
Slide text:
None extracted.

Speaker notes:
1. Now you can go back and test it by entering in a new chat. Click the run button, and you should see everything turn green, including your append or update row in spreadsheet. If you jump over your spreadsheet, you should see a new row added at the end, which is what I've got here.
Now it won't fill in some of the other columns we haven't mapped. I did click on the evaluation agent and look for what grade it gave it, and I manually typed that in. We'll show you how you can map the evaluation grade in a future step. You can have the spreadsheet get data from any

step or aget that comes before it.
And now, back the building the evaluation feedback loop.

[DAY_4][SLIDE_213][DAYSLIDE_011] platform.openai.com
Slide text:
1. platform.openai.com

Speaker notes:
1. We haven't created a feed-back loop yet, but we have created a feed. That is, we have the Evaluation Agent produces a grade as an output – which we will format as a feed to our If Statement.
To create a loop, we have to use an If Statement. Our If statement will be “If we have a grade of A, generate the audio. If we don’t try again.”
We will set a number of re-tries to 5 because every LLM call costs money. If it can’t get it right after five tries, then we want to bring a human in the loop to decide what to do next.

I mentioned each LLM call costs money. You can see how much money you've spent so far by going back to the platform.openai.com and clicking on the left side dashboard for usage. I’ve spent less than a dollar so far, but if the AI never can get to a grade of A, I could end up in an infinite loop, and burn a lot of money. Managing and tracking AI budget is a conversation for another time, for now, I want you to be aware of the value and cost of feedback loops and be able to check if you are using APIs.

Let’s move on to extracting the grade from the LLM output and creating the feedback loop.

[DAY_4][SLIDE_214][DAYSLIDE_012][EXERCISE] I need to extract the Grade. If should be at the beginning of the output. I need only the letter grade, as in A, B, C, D, or F, nothing else. If there is no grade, give it an F. Here is an example input:
Slide text:
1. I need to extract the Grade. If should be at the beginning of the output. I need only the letter grade, as in A, B, C, D, or F, nothing else. If there is no grade, give it an F. Here is an example input:
Grade: A-\n\nStrong alignment with the brand guidelines. The script is authoritative, modern, and clearly aimed at senior marketing leaders. It emphasizes measurable growth, smarter investment, and confident transformation, which closely matches the brand’s value-creation and business-impact focus. The language is concise and credible, with a research-backed feel.
Your output: A
2. If and Parse
Video, Media Plans, pretty much most marketing activities can be turned into agentic workflows

or full blown agents with evaluation criteria.
Safety

Speaker notes:
1. <split time: 13min>
There are two nodes we need. first, we have to extract the grade from the AI. Then we can give it to an IF Node to make the decision to forward on to the Generate Audio or send it back to the LLM to try again. This is a classic feedback loop.
Now, this gets slightly technical because extracting the grade requires a little bit of Regex code in javascript – which I will have an LLM write for me. It is 8 lines of code. I'll show you how I prompted the LLM to write it for me, and give you the code so you can copy and paste it in.

The point of this exercise is for you to see and build the most important element of AI Agents, the feedback loop. Once you see it, and build it, we will move on to some really fun and mind-blowing exercises to show you the power of AI in business.
I click add, and search for "code" node. I pick JavaScript.
I click the Ask AI tab, to help me write the code. I tell it "I need to extract the Grade. If should be at the beginning of the output. I need only the letter grade, as in A, B, C, D, or F, nothing else. If there is no grade, give it an F."

I give it an example input and output.
If & parse… then adding video. More video slides… then safety… and over to Lovable, Claude Cowork and OpenClaw (CLI Agetns that started with AutoGPT).

Video, Media Plans, pretty much most marketing activities can be turned into agentic workflows or full blown agents with evaluation criteria. Keep in mind the ROI trap we talked about yesterday. I was with a marketer that's going to do one of the MMA consortiums for ad personalization. They noted that their goals were:
Awareness
Consideration
Engagement
which they measured as:
Finishing videos
Ad recall
Visits to the website
Filling out the lead form
All these are within the realm of measurement for an AI Agent, but not all are created equal. Greg and I found that the video completion rate is not highly correlated with sales impact. It's highly likely for a person who only watches half the video, if their branding is up front, to get nearly the same impact, or the exact same impact, as someone who completes the video.In fact, completing the video in some cases is a sign that the person isn't in front of their computer anymore, and the video is just playing straight through.On the other hand, awareness is useful, but you can be aware of a brand and not like it. Consideration is much more predictive of sales.

These agents get stronger and more powerful. We have to be much more thoughtful about what we are measuring and how it affects our business in the short and long term.

[DAY_4][SLIDE_215][DAYSLIDE_013] You can see here the code the AI writes the code.
Slide text:
None extracted.

Speaker notes:
1. You can see here the code the AI writes the code. You can see it in the code tab. Click Execute Step and you can see to the right it extracted the Grade of an A.

Now, you want to make sure it will generate both the Pass and Fail conditions. We will come back to test that after we add in the IF node.
<
const items = $input.all();
const gradeRegex = /Grade: ([A-F])-\n\n/;
const grades = items.map((item) => {
const match = item?.json?.output.match(gradeRegex);
const grade = match ? match[1] : "F";
return { json: { grade } };
});
return grades;
>

[DAY_4][SLIDE_216][DAYSLIDE_014] To add in the IF Node, the click + and search for if and select the IF node.
Slide text:
None extracted.

Speaker notes:
1. To add in the IF Node, the click + and search for if and select the IF node. We drag over the Grade, and click the Equal to pull down.
Notice the Output has a True and a False Branch. We have set the True branch, now let's set the false branch. For this false branch, we will set it not equal to A.
Close it and you will see the canvas with the If branches.

[DAY_4][SLIDE_217][DAYSLIDE_015][EXERCISE] Next, I connect the IF to the generate

audio.
Slide text:
None extracted.

Speaker notes:
1. Next, I connect the IF to the generate audio. When we reconnect the audio, we have to update the mapping for where it is getting the text to convert into voice.
Notice that all of the inputs for all the nodes that come before the audio are available to us, but the JSON.output in the text section is red, meaning it needs some attention. If you click open each input, you can see you want to drag over the output from the Agent Voiceover script writer. So go ahead and delete the red old script and drag the new one over. It should turn green and show you the result of what the Generate Audio will read.

If you want, you can execute the step and download the MP3 file. I am going to close it and work on the false branch of the IF Statement.

[DAY_4][SLIDE_218][DAYSLIDE_016][CHAT_INTERACTION] Create a 5 second ad for Marketing & Media Alliance (MMAGlobal.com)
Slide text:
1. Create a 5 second ad for Marketing & Media Alliance (MMAGlobal.com)

Speaker notes:
1. For the false branch, I'm going to drag it and put it at the beginning of the agent chain, where it gets the chat instructions to create an advertisement script, so I can try it again. I'm going to click the false branch in the settings. Click retry and leave it at the default of three tries. Then I'm going to go back and click the clear chat and put in a new chat asking to create a five-second ad for Marketing Media Alliance. I'm creating a shorter ad because I think it's more likely to fail when it has just a limited number of words.
We have one other technical tweak to make so that the feedback loop works. We need to merge the feedback loop failure with the chat prompt so it knows what we want and why the most recent try failed.
<18min split time>

[DAY_4][SLIDE_219][DAYSLIDE_017][CHAT_INTERACTION] We will want to place a set node and merge right between chat and feedback loop.
Slide text:
None extracted.

Speaker notes:

1. We will want to place a set node and merge right between chat and feedback loop. So, we need to break the current pipeline and then re-connect the pipelines. We connect the chat to the first input and the if pipeline to the second.

Give it a prompt to run and hydrate the pipes, and the we will go in for some fine tuning. I give it a prompt for a 5 second ad, hoping it would fail so I can make the final tweaks.

[DAY_4][SLIDE_220][DAYSLIDE_018][CHAT_INTERACTION] Add Two Nodes:
Search for “Merge”
Search for “Edit” (set node).
Slide text:
1. Add Two Nodes:
Search for “Merge”
Search for “Edit” (set node).

Speaker notes:
1. The Merge node is somewhat obvious. If you have two different pipelines of information, like the original chat and the feedback from the evaluation, they have to merge. Go ahead and add in the merge.

Click plus. Type in merge in the search box. You should see this merge node. Go ahead and add it. By default, it will have 2 inputs, which is fine. Click close.
We need to wire up the merge to bring together the chat input and the feedback evaluation.

[DAY_4][SLIDE_221][DAYSLIDE_019][EXERCISE][CHAT_INTERACTION][DEBRIEF] Add Two Nodes:
Search for “Merge”
Search for “Edit” (set node).
Slide text:
1. Add Two Nodes:
Search for “Merge”
Search for “Edit” (set node).

Speaker notes:
1. But the merge alone won’t make it work. If you are technical, you might already see the formatting issue, but for non-technical people, you might be wondering why it won’t work. Let’s ask the AI to help us. This will be our bridge into Vibe Coding, which will be our next exercise, but for now, let’s finish this one.
The idea is that we need to merge the grade and feedback from the evaluation agent with the

original chat prompt so that when our agent tries again, it has both inputs.
You can find the AI helper over on the right with this symbol. Click it. I used the AI to help write all of this code. Here's the prompt I gave to the n8n AI agent do it.
I ran it and it worked perfectly. Now, the last step is to map the evaluation grade to the same spreadsheet by adding another Append to sheet node.
Note that the internal AI within n8n has a restore function. This is really important version control when using AI agents, as we will see when we get into vibe coding after this hands on exercise.

If you are technical, you can write it yourself – I show the code the AI wrote for anyone that wants to do it manually by clicking edit and searching for the edit node.

Let’s take three minutes for you to repeat the process and finish the Agent with a feedback loop, and then we will debrief, and move on to Vibe coding. <3 min break>

[DAY_4][SLIDE_222][DAYSLIDE_020][DEBRIEF] Debrief
Slide text:
1. Debrief

Speaker notes:
1. <Split?>
Show workflow.
We did these examples using audio because it renders far faster than video, but everything we showed you can work with video as well. You can make a few changes to your system instructions and connect to generate video instead of generate audio, and you're creating video ads. Or perhaps you want to create CRM emails. Same idea can apply. Let me share some interesting video tools you might want to explore.

[DAY_4][SLIDE_223][DAYSLIDE_021][EXERCISE] Advice
Slide text:
1. Advice
2. Play with the tools first,
if they show promise, turn them into agentic workflows or agent pipelines.

Speaker notes:
1. I almost always start in an advanced LLM like Claude Opus 4.6, or GPT 5.4 Pro, and if I can get my idea to work with prompting, then I consider the value of turning it into an AI agent or agentic workflow. If it work I need to keep coming back to, or a repeatable process, like creating messages, media plans, or analysis, then it is worth the time to create the AI Agent or workflow. Let’s take a look at some great tools coming out of Google Labs, and

---

[DAY_4][SLIDE_224][DAYSLIDE_022] MMA CEO/CMO Summit, July 2025
Slide text:
1. MMA CEO/CMO Summit, July 2025
2. According to a Zappi test:
People thought all ads featured a human creator
They did think the two AI ads felt AI, but weren't sure why
All ads felt authentic
64% of people prefer ads by people
3. Monks AI
4. Human
5. Dave AI

Speaker notes:
1. REX: Next, I want to share a case study from a Panel I moderated with Dave.com at the CEO/CMO Summit in July. They are using AI generated ads in TikTok, and we are going to look at how they manage the context to get good output. First, let me play an ad with a human influencer, then we will look at the AI influencer ads.

[DAY_4][SLIDE_225][DAYSLIDE_023] How we do it
Slide text:
1. How we do it
2. Five steps to creating UGC with AI (this month!)
3. Step 1:
Script creation
Tools: Gemini
(any AI tool)
4. Step 2:
Character creation
Tools: Gemini
(any AI tool)
5. Step 3:
Character consistency development
Tools: Whisk
6. Step 4:
Video creation
Tools: Flow (Veo3)
7. Step 5:
Video editing

Tools: Adobe Premiere Pro
8. This is just the beginning.
Connected tools, process automations, auto scaling, & reduced/no human editing is our next focus.

Speaker notes:
1. How did the team at Dave do it?
AI is given examples of successful TikTok ads as context. It then creates a script and separately, using the same context, the AI is prompted to create the character…
The Ai currently generates less than 10 seconds of video content per go-that's a limitation of the existing video systems. So, you will be stitching together different shots. That's OK, because that is how a lot of social media footage is already. But, for AI, the problem is character consistency. Without very robust descriptions of the character as context, the AI will change the character from scene to scene. That's where Whisk comes in.
DEMO— you can use Flow, or Fireflies, and a Or Runway, or Sora2 – there are a lot of video generation options to try. Dave used Veo3, which is Google's most advanced model as of today. they did some final editing in Adobe Premiere Pro. They tell me that the last ad we viewed took less than 45 minutes from start to finish to produce. Wow. Right?

[DAY_4][SLIDE_226][DAYSLIDE_024] Subject leaps off and turns into a majestic Red Tail Hawk. He then flies across the scene.
Slide text:
1. Subject leaps off and turns into a majestic Red Tail Hawk. He then flies across the scene.
2. +
3. +

Speaker notes:
1. Here is the second video I created, using a few more images and the background scene.

[DAY_4][SLIDE_227][DAYSLIDE_025] Here is the video
Slide text:
None extracted.

Speaker notes:
1. Here is the video

[DAY_4][SLIDE_228][DAYSLIDE_026] Animate:

Prompt: Subject climbs up the rock face.
Slide text:
1. Animate:
Prompt: Subject climbs up the rock face.

Speaker notes:
1. Here is a second video using a different animation prompt.

[DAY_4][SLIDE_229][DAYSLIDE_027][BLANK] Visual-only or blank slide
Slide text:
None extracted.

Speaker notes:
None extracted.

[DAY_4][SLIDE_230][DAYSLIDE_028] We've included links to Wisk and Flow, and you have Firefly that let's you use the Veo models as wel
Slide text:
None extracted.

Speaker notes:
1. We've included links to Wisk and Flow, and you have Firefly that let's you use the Veo models as well, so this weekend, play around. I know there are some folks from GM on this training, and Norm, the CMO of GM and I were exchanging our own experiments with these tools. Everyone should be finding time to explore to see what is possible.

[DAY_4][SLIDE_231][DAYSLIDE_029] Runway?
Slide text:
1. Runway?

Speaker notes:
None extracted.

[DAY_4][SLIDE_232][DAYSLIDE_030] Vibe Coding The AI
Slide text:

1. Vibe Coding The AI

Speaker notes:
1. Now, onto Vibe Coding.
You just did the hard part of agents. You built an AI Agent with feedback loop by hand. Well done! You have the deeper understanding of what AI agents are doing and this will help you build and manage agents much better know that you know how they really work. Let’s dive into vibe coding, which is easier in many ways but more complicated in others, and show you what you can do.
We will then move into multi-purpose agents, like OpenClaw and increasingly the evolution of LLMs into agents like Claude Co-work.
We will wrap up with an AI Safety framework and a preview of our Friday working session on tips for managing AI Agents.

[DAY_4][SLIDE_233][DAYSLIDE_031][EXERCISE] Lovable free tier should be enough for training.
Slide text:
1. Lovable free tier should be enough for training.
2. Project Overview
Build a "Multi-Channel Ad Creator" for senior marketing leaders. The app should generate and refine ad copy for Audio, Video, Email, and Display (300x250) based on strict Brand Guidelines. The core logic uses an iterative AI feedback loop that improves based on user notes and automated grading.
> See “three prompts” to add to Lovable in the Labs.
You will need to click “approve” for initial set-up. (should take 5 minutes).

Speaker notes:
1. <How do I convert this into a prompt for Lovable to build an Ad creator for Audio, Video, Email and display (300 x 250 pixel size). I'd like to have a brand guidelines input box, which I can edit and the brand guidelines are used by the AI agent to create the ads. I'd like it to have a clean interface with prompt box, a saved list of input output pairs, and an edit box so I can adjust the prompt before sending to the production tools to create the email, audio, video, and display. I'd like a checkbox selector to decide which tools it outputs (audio, display, email and video), and a feedback loop. The feedback loop should allow me to add notes, and the notes will become part of the feedback loop mark down that the AI is using to grade the ads. Each time I give additional notes, the AI agent evaluates if the Mark down should be updated and it gives a redline suggestion of what it would change in the feedback loop. > <gave it JSON of my n8n agent.>

[DAY_4][SLIDE_234][DAYSLIDE_032] Enter your URL,
Update your brand guidelines
Explore and consider what you want to change.
Slide text:
1. Enter your URL,
Update your brand guidelines
Explore and consider what you want to change.

Speaker notes:
1. <Doc Guide, Copy & Paste>
https://docs.google.com/document/d/191LE8kIw9ECKjxh98M1ZgOIY5Wu7p2eHVFLNmjeIsjM/edit?usp=sharing

---

[DAY_4][SLIDE_235][DAYSLIDE_033][EXERCISE] We'll give you a couple of minutes for Lovable to build out your app.
Slide text:
None extracted.

Speaker notes:
1. We'll give you a couple of minutes for Lovable to build out your app. Once it does that, we want you to go in and explore what it already does. See what you would like to change and then give it the instructions for how to make the changes. We'll give you a few minutes to play with it today, and we encourage you to play with it more after the session is over today to get a feel for how powerful vibe coding can be.
Now, a key part of vibe coding success is understanding how the AI takes instructions and how it builds context. You'll notice that my prompt was broken into three separate pieces, where I first had it build the UI and then I went through and did some fine tuning. We'll make available a link to the Lovable prompt Bible that they've developed, that includes the best practices to get the most out of the way that Lovable's AI engine converts your natural language direction into code
If you happen to be a coder, you can click where the preview window is just to the right of it with the code brackets, and you can see the structure that Lovable built and the specific code. You can also connect this into GitHub to have version control or to take over what you prototype in Lovable and move it over into a production coding environment where you may have other control or may choose to use other tools such as Claude Code.
Let me hand it over to Caleb to talk about AI Agent safety.

[DAY_4][SLIDE_236][DAYSLIDE_034] AI Risk Assessment & Strategies
Slide text:
1. AI Risk Assessment & Strategies

2. Image Generated By Gemini
Jan. 12, 2025
3. ~7 minutes

Speaker notes:
1. Thanks Rex. As we plan and build AI agents, we should ask a few key questions to make sure our application of AI is relatively safe.

[DAY_4][SLIDE_237][DAYSLIDE_035] AI Risk Framework
Slide text:
1. AI Risk Framework
2. Ideal for AI
3. Ideal for AI
4. Precision required?
Need for rationale?
Do you control the input?

Speaker notes:
1. In Part 1 of The AI Conundrum, I introduced an AI Risk Framework. It has three axes, which correspond to three basic questions.
1. How much precision is required for your use case?
2. How strongly do you need rationale from the AI?
And, 3. how much control over the input do you have?
Let's look at that precision component first.

[DAY_4][SLIDE_238][DAYSLIDE_036] Precision
Slide text:
1. Precision
2. Ideal for AI
3. Ideal for AI
4. Math
Science
5. Higher precision → Worse for AI

Speaker notes:
1. Precision is about how precise and accurate you need the answer to be in your use case. For example, math and science are in that high precision category. When you do a calculation, you need absolute precision. If I ask you to add two numbers, there is one correct answer and anything else is wrong. The higher the precision you need, the more risk there is with using AI -

and the more you should think of doing something other than AI Agents to address the need.

[DAY_4][SLIDE_239][DAYSLIDE_037] AI Risk Framework
Slide text:
1. AI Risk Framework
2. Ideal for AI
3. Ideal for AI
4. Source: GPT-4o Jan 1, 2025
5. Correct Answer

Speaker notes:
1. While AI has gotten better at math and science over the past year, there are still cases where AI isn't able to respond with perfect precision. We saw that in Day 1, when we asked AI to count the intersections and place dots at each intersection.

[DAY_4][SLIDE_240][DAYSLIDE_038][CHAT_INTERACTION] AI Risk Framework
Slide text:
1. AI Risk Framework
2. Ideal for AI
3. Ideal for AI
4. Creative writing
Image generation

Speaker notes:
1. On the other hand, lower precision tasks, like creative writing or image generation, and a lot of advertising and marketing use cases are great applications for AI. Before I show the next slide, close your eyes for a moment and imagine a dog.
Got the the image.
OK. Type in the chat the type of dog you had in mind.

What you will see is there is a lot of diversity in the answers. The answers represent a probabilistic distribution and that is what happens with AI too. It produces a probabilistic answer. – By the way, all of those answers are great.

[DAY_4][SLIDE_241][DAYSLIDE_039] AI Risk Framework
Slide text:
1. AI Risk Framework

2. Ideal for AI
3. Ideal for AI
4. Creative writing
Image generation
5. Lower precision →
AI performs better

Speaker notes:
1. Here's an AI Generated picture of a dog with a plate of pizza.
There are lots of different ways to create an image from the prompt 'dog with a plate of pizza'. There's some leeway in how it looks. We need less precision in areas like creative writing and image generation, and thus AI performs better on those tasks. A painting or picture isn't necessarily wrong because one pixel has a color 10% too light-- image generation doesn't require absolute precision to be useful.

[DAY_4][SLIDE_242][DAYSLIDE_040] GenAI Images & Text
Slide text:
1. GenAI Images & Text
2. Ideal for AI
3. Ideal for AI
4. An Ideal Case For AI

Speaker notes:
1. There are a lot of use cases where the AI doesn't have to be precisely right to be useful. Text generation of a welcome letter from a marketer is another great example where there is latitude to in how it is written for it to be useful.
Did you noticed the errors in the AI generated stopwatch at the start of this classroom session? Probably not.

[DAY_4][SLIDE_243][DAYSLIDE_041] Classroom Session: What are Agents?
Slide text:
1. Classroom Session: What are Agents?
2. Image Generated By Gemini
Jan. 12, 2025

Speaker notes:
1. This image from the start of this section was created in January 2025, when the image generation was not as good as it is today, and you can find several errors if you look carefully.

The most obvious error is the number 55 where the number 11 should be. There are other errors like in the little number dial at the top, I think those numbers should count up in logical increments - but they don't. Also the the loop shadow is unrealistic. You probably didn't notice these errors when I was just using the stopwatch to indicate we were going to spend 7 minutes on AI safety.

[DAY_4][SLIDE_244][DAYSLIDE_042] Classroom Session: What are Agents?
Slide text:
1. Classroom Session: What are Agents?
2. Image Generated By Gemini
Jan. 12, 2025
~8 errors
3. Image Generated By Gemini
Nov. 28, 2025
~3 errors

Speaker notes:
1. If we were to try the same prompt again with Nano Banana, launched in November of last year, there are significantly fewer errors, but the probabilistic nature of AI means that if 1000 people use the prompt "image of a stopwatch" there are bound to be a few minor errors.
On the image on the RIGHT, the image is far better, in that the errors are more subtle, such as the small "made in Switzerland" at the very bottom has illegible text where Switzerland should appear. For some advertisers, when it comes to showing their product, even this very minor error would be too much. An automotive manufacturer, for example, might want any image of their product to be absolutely precise.

[DAY_4][SLIDE_245][DAYSLIDE_043] GenAI Images
Slide text:
1. GenAI Images
2. Ideal for AI
3. Ideal for AI
4. An Ideal Case For AI

Speaker notes:
1. Still, if you need a quick image for a training video, I'd say it is a pretty good choice to use GenAI for the use case. Moreover, this use case has a human in the loop. I'm looking at the image and deciding if it's good enough to use. In this sense, I've added a layer of safety by adding a human in the loop.
The Open vs. Closed environment is another key safety issue. AI in the wild is more risky than

AI used internally. People outside your company, interacting with your AI, may try to hack it. That generally is not the case with employees and partners, especially when they are logged in and their behavior can be tracked. Closed environments can manage the data the AI uses. You can also add a human in the loop where beneficial. Both these choices generally make AI much less risky. That said, a lot of companies want to have an AI Agent to interface with customers. That is more risky, so be sure you have some ways to keep an eye on the AI, which bring us back to the importance of Evaluations.

[DAY_4][SLIDE_246][DAYSLIDE_044][SURVEY] From AutoGPT to Open Claw - CLI Agents
Slide text:
1. From AutoGPT to Open Claw - CLI Agents
2. NotebookLM
Analysis Agent (NPS)
Agents, Memory, History & OpenClaw

Speaker notes:
None extracted.

[DAY_4][SLIDE_247][DAYSLIDE_045] Next: Long Context
Slide text:
1. Next: Long Context
2. RAG Synthesizes from snippets
Long Context reasons across the entire corpus
3. Ai generated, Jan 21, 2026 by Gemini..
←Notice typos & remember the risk model on precision

Speaker notes:
None extracted.

[DAY_4][SLIDE_248][DAYSLIDE_046][EXERCISE] Hands On Exercise: Message and Person Interaction
Slide text:
1. Hands On Exercise: Message and Person Interaction
2. Try Computer Vision and AI feature extraction to develop theories of why some ads perform better than others in certain situations.
3. Message
4. Person

5. Context

Speaker notes:
None extracted.

Synthetic addition — Day 5

The following sections were added because the uploaded deck repeatedly describes a five-day session, but the source slides only cover four day sections in detail. Review and edit these before treating them as final course material.

[DAY_5][ADDED_SLIDE_001][DAYSLIDE_001][SYNTHETIC][AGENDA] Day 5: The Future of Work in Marketing
Slide text:
1. Day 5 Friday: The Future of Work in Marketing
2. Redefining the marketer’s role from content creator to system orchestrator
3. Focus areas: AI governance, the Editor-in-Chief mindset, and scaling AI into a production-grade enterprise workforce

Speaker notes:
1. This Day 5 section was not present as extractable source slides in the deck, so it has been synthesized from the repeated agenda language and the flow of the preceding four days. The goal is to complete the five-day narrative in a way that is consistent with the rest of the course.
2. Set up Day 5 as the capstone. The week has moved from AI fundamentals, to agents, to workflows and evaluations, to multi-agent thinking. The final day should pull those pieces together into an operating model for real marketing teams.

[DAY_5][ADDED_SLIDE_002][DAYSLIDE_002][SYNTHETIC] From Content Creator to System Orchestrator
Slide text:
1. Old role: create assets directly
2. New role: define objectives, constraints, context, evaluation criteria, and approval gates
3. The marketer’s leverage shifts from making every asset by hand to designing systems that create, test, and improve assets at scale

Speaker notes:
1. The language on the repeated agenda frames this shift explicitly. Marketers are no longer just writers, designers, or campaign operators. They increasingly become the people who define the system: what good looks like, what guardrails exist, how the AI gets context, and when a human

must review the output.
2. This is the operating shift the whole week has been building toward. Prompting is only the beginning. The durable skill is system design.

[DAY_5][ADDED_SLIDE_003][DAYSLIDE_003][SYNTHETIC] The Editor-in-Chief Mindset
Slide text:
1. Give the AI assignments, not unchecked authority
2. Define brand voice, legal boundaries, safety rules, and quality thresholds
3. Approve, reject, or redline output based on explicit evaluation criteria

Speaker notes:
1. The Editor-in-Chief framing matches the rest of the deck's emphasis on bounded agency, human review, and evaluations. The marketer does not need to produce every draft. The marketer needs to ensure the drafts are on-brand, safe, useful, and aligned to the business objective.
2. Tie this back to earlier examples where output improved once the system prompt, brand DNA, and evaluation criteria were clarified.

[DAY_5][ADDED_SLIDE_004][DAYSLIDE_004][SYNTHETIC] Governance and Guardrails for Enterprise AI
Slide text:
1. Choose approved tools and models
2. Control data access, permissions, and environments
3. Prefer closed environments for higher-risk use cases
4. Define escalation paths for failure, bias, or unsafe output

Speaker notes:
1. This slide should connect the safety framework to practical governance. The main questions are: who can use which tools, which data is allowed into the model, what actions require approval, and where the AI is allowed to operate autonomously.
2. Reinforce that open, public-facing agents are riskier than internal workflows. The more precision, rationale, and uncontrolled input a use case requires, the more governance matters.

[DAY_5][ADDED_SLIDE_005][DAYSLIDE_005][SYNTHETIC] Evaluations, Rewards, and the Danger of Optimizing the Wrong Thing
Slide text:
1. Every powerful AI system needs a feedback loop

2. Reward design shapes behavior
3. Short-term optimization can hurt long-term business performance
4. Measure what matters, not just what is easy to count

Speaker notes:
1. This is the strategic wrap-up of the reward-hacking and ROI-trap material. A system will pursue the objective it is given, not the objective you meant. That is why evaluations, observability, and KPI design are leadership issues, not just technical issues.
2. Use this slide to emphasize that marketing leaders must decide what the AI is actually optimizing for: revenue today, profit, retention, brand health, lead quality, or some combination across time horizons.

[DAY_5][ADDED_SLIDE_006][DAYSLIDE_006][SYNTHETIC] Scaling from Pilot to Production
Slide text:
1. Prototype in chat
2. Turn repeatable wins into workflows
3. Add evaluations, version control, memory, approvals, and cost monitoring
4. Move from single agents to coordinated systems where justified

Speaker notes:
1. The week repeatedly implies this maturity path. Start with experimentation. If a task shows promise and repeats often enough, turn it into a workflow. If the workflow becomes strategically important, add the production disciplines: logging, cost controls, versioning, approvals, testing, and governance.
2. This is where the marketer and the technical team meet. Production AI is not just a better prompt. It is an operational system.

[DAY_5][ADDED_SLIDE_007][DAYSLIDE_007][SYNTHETIC] Skills Marketing Teams Need Next
Slide text:
1. Context and prompt design
2. Workflow thinking
3. Evaluation design and QA
4. Data literacy and KPI fluency
5. Change management and team enablement

Speaker notes:
1. Translate the week into hiring and team-development language. The important capabilities

are not only creative execution. They also include the ability to design instructions, critique outputs, define metrics, run experiments, and manage systems that combine human and machine work.
2. This slide is useful for helping leaders think about org design and training priorities.

[DAY_5][ADDED_SLIDE_008][DAYSLIDE_008][SYNTHETIC][SURVEY] Week Wrap-Up and 30-60-90 Day Action Plan
Slide text:
1. Pick one workflow to redesign with AI
2. Define the KPI, guardrails, and evaluation method
3. Keep a human approval gate for high-stakes output
4. Document learnings, then scale what works
5. Complete final feedback / NPS / next-step planning

Speaker notes:
1. Close the week with action. Ask each participant to identify one practical workflow they can improve in the next 30 days. The point is not to launch a giant transformation immediately. The point is to build one good system with clear metrics and strong oversight.
2. End by connecting back to completion, recordings, materials, and any certificate language already used elsewhere in the deck.
