Course Summary

ChatGPT Work: The New Way to Use ChatGPT for Marketing and Business

Prepared by The Marketing Crowd *with the help of one of our AI colleagues* ☺

Course Overview

The big picture: three ways to use AI

Evan opened by framing where ChatGPT (and AI generally, including Gemini, Claude, Grok and DeepSeek) now sits. Since ChatGPT launched in November 2022, there have really been three distinct ways to use these tools, and they layered on top of each other over time rather than replacing one another.

The first way, and the original use case, is chat: asking questions and getting expert level advice, ideas, strategies and tactics back. A typical example Evan gave was prompting "I have a driving school in Killarney, I'm looking to increase my sales, act like a business growth expert and teach me how to increase my sales."

The second way, which arrived roughly a year to eighteen months later, is create: using AI to produce written content and images (and, for Gemini users, video). This covers things like writing in your own tone of voice, generating headlines, and creating or editing product images.

The third and newest way is tasks, what ChatGPT calls Work. This is agentic AI: instead of asking a question and getting an answer, you give it an objective and it goes off and completes multiple steps on your behalf without further input. Evan explained that Claude, Anthropic's AI, made the breakthrough here first with Claude Skills, then with Claude Cowork, and that ChatGPT Work launched in July as effectively a copy of Cowork. His view is that ChatGPT has been closely following Claude's feature set over the past year, and that this Tasks/Work capability is where all of the major AI tools are now heading.

Nearly everything covered in the session is only available on the paid ChatGPT Plus plan (around €20 plus VAT per month). The one exception is plugins, which are available for free.

New models: Astra, Sol, Terra and Luna

ChatGPT now offers four models when you are in the Work tab, ranked by power and by how much of your usage allowance they consume:

- **Astra**: the newest and most powerful (launched about a week before the session). Described as "off the charts" powerful, intended for genuinely complex tasks such as programming or advanced mathematics. Evan was clear that for digital marketing work, Astra is overkill and burns through usage credits very fast, comparing it to "using a bazooka when you just need something much simpler."
- **Sol**: was the top model until Astra launched a week earlier. Still extremely capable and Evan's recommended default for marketing tasks. It was the model used in the demonstrations shown on the course.

- **Terra**: a middle tier model. Capable of most of what was demonstrated on the course, uses less of your usage allowance, just takes a little longer.
- **Luna**: the lightest model, best for everyday, low stakes requests.

In the Chat tab (as opposed to Work), you only get a choice of Sol or an older legacy model. Evan advised never going back to the older model, it exists purely for people who preferred how a previous model responded. Sol in Chat does not use up any usage allowance at all, no matter how it is set.

Alongside the model choice, you can set a thinking effort level (low, medium, high, extra high or maximum). Higher effort produces more thorough results but takes longer. Evan's practical rule of thumb: for advice and chat based work, set Sol to high effort by default since it doesn't cost you anything extra in Chat. For Work tasks, only reach for high effort or the top models when the task genuinely needs it, since both drive up usage consumption.

If you are on the free version of ChatGPT, you are locked into Luna for chat and you do not have access to the Work tab at all (clicking it just prompts you to upgrade).

Usage limits: why they exist and how they work

Evan explained the underlying economics: running these models is extremely expensive. ChatGPT is estimated to lose around $40 billion a year, and even a €20 monthly subscription is estimated to cost OpenAI something like €400 a month in actual compute per user. A simple text answer costs roughly 2 cent in compute, an image around 25 cent, and a video somewhere between three and five dollars, which is why video generation is never offered for free.

Because of this, both Claude and ChatGPT now impose usage limits, split into a five hour rolling limit and a weekly limit. Critically, in ChatGPT these limits only apply in the Work tab. Staying in Chat, whether for advice, image creation or even complex multi part requests, does not use up any allowance at all. Evan's core piece of advice for the whole course was: default to Chat wherever possible, and only move to Work when you specifically need the reliability that Work provides.

He demonstrated this directly with a live example. Early in the course, after running one task in Work using Astra at high effort, his five hour usage allowance dropped to just 9 percent remaining and his weekly allowance dropped by 14 percent, all from a single request. He noted that the same task using Sol had, on a previous occasion, only used around 8 to 10 percent of his usage.

ChatGPT does provide three free "usage resets" that instantly restore your five hour and weekly limits back to 100 percent, intended as a buffer. During the session Evan attempted to use one when he ran low, but it did not immediately take effect (the interface still showed him locked out, though it appeared to catch up a few minutes later).

By contrast, in Claude, hitting your usage limit blocks you from using the tool at all, in either chat or task mode, until the limit resets. One participant noted this, and Evan agreed it is a genuine frustration with Claude compared to ChatGPT, which currently feels more generous with allowances.

On the free version of ChatGPT, chat is now unlimited in terms of number of conversations (a recent change), but you are capped to two file attachments a day and two images a day. Evan's advice for anyone relying on the free tier and hitting those caps was blunt: switch to Google Gemini instead, which offers unlimited file attachments, unlimited images and no usage caps for free, albeit without the skills and tasks features covered in this course.

For businesses with multiple staff using ChatGPT, there is also a Business plan (around €21 per seat per month, minimum two seats). Beyond the shared billing, its main benefit is that it keeps company data private and secure, and each person gets their own individual usage allowance, so one person hitting their limit does not affect anyone else's.

Chat tab vs Work tab: what to use when

Evan's simplified framework: use Chat for advice, ideas, strategy questions and image generation, since none of that touches your usage allowance. Use Work when you want ChatGPT to reliably complete a defined set of deliverables without you having to check its output.

He demonstrated the difference using two live examples he had prepared before the course, both centred on promoting a dress via a link to a product page:

Example one, run in Chat using Sol: the prompt asked for two Instagram post images (4x5), two story images (9x16), an email header image, Instagram caption text and email copy, all pulled from the product page link. This worked entirely within Chat and produced a full campaign pack without using any usage allowance.

Example two, the same type of task run in Work using Astra: a more complex prompt building a full two week marketing campaign (strategy, a two week social media calendar, four Instagram/Facebook posts, three reel briefs, an email campaign). It also generated two full sets of ad images across two different visual approaches (an "editorial" style and a "lifestyle" style), noting that when running paid ads it is useful to test different visual treatments and let the ad platform allocate spend toward whichever performs best. This particular run took around nine minutes and consumed the bulk of Evan's five hour usage allowance because of the Astra model choice.

Evan's honest assessment based on his own testing: Chat can handle almost everything demonstrated, including genuinely complex multi part requests, because it uses the same powerful Sol model under the hood. The difference is reliability. When you ask for several things at once in Chat, it sometimes misses an item (for example, forgetting to produce the story image), and you have to follow up and ask it to complete the missed piece. In Work, the model is more rigorous about checking its own output against everything that was asked

for, and it will self correct if it runs into a problem rather than stopping to ask you. His rule: try it in Chat first if you are being careful about usage allowance, and only move to Work once you notice Chat is being unreliable for a particular type of task.

He also flagged that ChatGPT sometimes proactively suggests moving a Chat conversation over to Work when it judges the task is complex enough, offering a button to do so, but that this can be ignored and the task will still complete in Chat.

A prompt Evan shared for participants to try: "I have the following product, service or objective: [insert]. Given the recent improvements in ChatGPT, tell me all the practical ways it could help me achieve this objective, including useful tasks that might not immediately occur to me. Group the ideas into research, planning, marketing and sales. Prioritise the ideas by likely impact." He ran this live using a dog grooming business as the example and noted it surfaces ideas such as competitor reviews, customer review analysis, local search improvements, a ninety day sales plan and a customer retention plan, all achievable without needing to leave Chat.

Skills: the central new capability

This was the section Evan spent the most time on, describing it as one of the most valuable things covered.

The concept. Evan used the analogy of training a new intern: rather than explaining your business, brand and preferences from scratch every time, you write it all down once, and the AI stores and reuses that context whenever you invoke the skill by name. A skill bundles instructions, preferences and reference material (logos, fonts, writing samples, product information) into a reusable package.

Two buckets of skills. Evan grouped skills into two categories:

1. **Brand skills**: anything to do with your business identity, such as visual branding (logo, colour palette, typeface), writing style/tone of voice, and product and customer context.

2. **Admin skills**: anything repetitive you ask the AI to do on a recurring basis, such as turning a Zoom transcript into a course summary, converting an attendance sheet into a standard template, or producing social media crops from an uploaded photo.

Where skills live. In ChatGPT, skills sit under Plugins in the left hand navigation (there is no separate "Skills" menu item visible directly), then a Skills tab within that. From there you can create a new skill (via chat conversation, via a built in editor, or by uploading a file), or install/manage existing ones. Skills only exist in the paid version of ChatGPT (though one participant, using the desktop app version of the free tier, reported seeing a skills option there too, which Evan had not previously been aware of and said he would test separately).

Creating a skill via conversation (the "create with chat" method). Evan demonstrated building a TMC brand skill live. Key points from the walkthrough:

- Choosing "create with chat" from the Skills area defaults you into the Work tab, which uses your usage allowance. Evan pointed out you can manually switch over to the Chat tab instead and build the skill there for free, though the interface keeps nudging you toward Work.

- You describe in plain, conversational language what you want the skill to do (Evan spoke his request aloud using the microphone rather than typing carefully worded prompts, deliberately showing that no special prompt engineering skill is required).

- For the brand skill, he uploaded the TMC logo and a photo of a compliment slip so the AI could infer the brand colour (identified correctly as a deep purple) and suggest a similar typeface based on what appeared on the slip.

- ChatGPT then asked a series of clarifying questions: how the logo should be positioned (Evan chose cover/title page only, not on every page), what overall design style was wanted (he chose "somewhere in between" clean/minimal and more elaborate), and whether to auto include website and contact details (he said no).

- Once finished, ChatGPT confirmed the skill was built and saved it to the skills library under a name like "The Marketing Crowd Branding Skill."

- After creation, invoking it is as simple as referencing it by name, for example asking ChatGPT to turn a blog post link into a downloadable PDF and "use the Marketing Crowd Branding Skill," without having to restate any branding details.

A second brand skill: writing style. Evan showed creating a "TMC Writing Style" skill the same way, feeding it links to previously published blog posts (or, alternatively, pasted email or social post text) as writing samples. Once built, asking ChatGPT to write anything and referencing this skill applies that tone automatically. He then combined it with the branding skill in a single request (asking for a PDF about ten facts on AI usage in Ireland, written in TMC's voice and laid out in TMC's branded style), demonstrating that multiple skills can be referenced together in one prompt.

A third suggested brand skill: products, services and customers. Evan recommended creating a skill that gives the AI permanent context on your offerings and customer base. The simplest way to populate it is to give ChatGPT your website link and ask it to scan every page for the relevant details. A more thorough alternative, which he recommended, is to first run ChatGPT's Deep Research feature on your own business (see below) and then attach that finished research report when building the skill, since it will already contain a far richer picture than a quick page scan.

Building an admin skill: the course summary skill. In a genuinely useful live demonstration, Evan built the actual skill used to summarise this course, in real time, describing his exact requirements as he went, several of which map directly onto rules used in this summary:

- Never include his self introduction from the start of the session.

- Always refer to him as Evan, never by his surname.
- Never include a participant's name anywhere in the summary, for privacy.
- Stay faithful to what was actually said, no embellishment or invented content, but reordering points and trimming genuine repetition for clarity is fine.
- Write in third person (not as if Evan himself wrote it), in a plain, non jargony style, since the summary is understood to be AI generated.
- Target a six to seven page structured document.

He noted that in the past, before skills existed, he had to correct these same points manually every single time he asked an AI to summarise a transcript. Turning it into a skill means never having to repeat those instructions again.

Two ways to create a skill. Evan demonstrated both methods side by side:

1. Go directly to Plugins, Skills, and describe from scratch what you want the skill to do (as above).
2. After a successful piece of work you are happy with, simply say "could you turn everything we just did into a skill I can use in future?" ChatGPT will package that exact workflow up automatically. He used this method for the product promotion example (the dress campaign), asking it to bundle the product page check, image set, ad concepts and copy documents into a reusable "product promotion" style skill.

A live technical hiccup, When Evan tried to download a skill created in Chat and re-upload it (rather than letting it save automatically, which happens by default in Work), the file was rejected twice for not being correctly formatted as a proper Skill markdown file, even after ChatGPT said it had corrected the issue. Evan used this as a teaching moment: Work is more reliable for skill creation and installation precisely because it handles this kind of technical troubleshooting itself, whereas Chat sometimes leaves you needing to work around a snag, or, as happened here, simply switch to doing the whole thing in Work.

Transferring skills between different AI tools. Because ChatGPT, Claude and (in future) Gemini are converging on the same underlying skill file format, a skill built in one tool can be downloaded and uploaded into another. Evan demonstrated this using a "social media crops" skill originally built in Claude: he located it under Claude's Customise then Skills menu, downloaded it as a file from the three dot menu, then uploaded that same file into ChatGPT's Skills area, where it worked immediately without needing to be rebuilt or reworded.

Prompt for finding your own skill candidates. Evan gave a prompt for participants to paste into ChatGPT: "I want you to help me identify tasks I do regularly in my business that could be turned into a repeatable ChatGPT skill. Ask me a series of questions one at a time to help me discover these tasks. Focus on things I do on my computer regularly, tasks that

follow a simple pattern. After asking me your questions, give me a ranked shortlist of my top three skill candidates, explaining why each of them could be turned into a skill."

Can you trust the output? Fact checking and hallucination

A participant asked directly whether AI output can be trusted for accuracy. Evan's answer: AI models can hallucinate, meaning they can generate information that sounds plausible but is factually wrong, and this risk never fully disappears, though it decreases with each newer, more capable model. He demonstrated this live by asking ChatGPT (using Sol at medium effort, in Chat) to list new Instagram features introduced since the 1st of January 2026, then followed up with a fact checking request: "could you please fact check every single point, verify every single point from another source on the internet, and give me a traffic light system, where green means definitely factually correct and verified, orange means you're not quite sure, and red means you don't think this is factually accurate at all."

The result flagged one claim as orange (a feature that existed but had not been described accurately in terms of when it launched) and another as orange for being technically real but only in limited rollout, while everything else came back green with no red flags. Evan's broader advice: apply the same fact checking and traffic light approach to any output you plan to rely on, whether that is a course summary checked back against the original transcript, or a spreadsheet dashboard checked back against the source data. He was confident that running the same test on the free tier model would surface noticeably more red flags, since less capable models hallucinate more.

Plugins and connectors

Plugins let ChatGPT connect to third party tools and are available even on the free version, accessed via the Plugins menu. Evan connected a demo Gmail account and asked ChatGPT to review the last few days of email and flag anything needing a reply, which it did, correctly identifying two LinkedIn enquiries as worth following up on. He also connected a Spotify account and asked for AI related podcast recommendations, which pulled results directly from Spotify.

He was candid that most plugin connectors are not yet especially useful in practice. He personally does not use the Canva connector ("it's not very good") and is not comfortable connecting a real email account given the sensitivity of the data involved, using a dummy training account instead for demonstrations.

Where he does see real future value is in connecting marketing specific tools: an email platform such as MailerLite or MailChimp (so a generated campaign could be set up ready for review, not sent automatically), and Meta Business Suite, so that generated ad creative and copy could be pushed straight into a draft campaign ready for review. He already has Meta Ads connected and demonstrated asking ChatGPT to pull details of his last three ad campaigns and analyse them like a data analyst (audience split by gender and age, platform performance), then asked for recommendations on how to improve the next campaign based on those insights, which it was able to do directly from the connected account.

The ChatGPT desktop app

Evan explained the difference between using ChatGPT through a web browser and installing the dedicated desktop application (available at chatgpt.com/download, detected automatically for Windows or Mac). Once installed, it behaves like any other application on your computer, with your account, chat history and usage synced automatically.

The key advantage of the desktop app is that it can be pointed directly at a folder on your computer and given permission to work only within that folder, rather than you having to upload files one by one through a browser. Evan demonstrated this with a live expenses example:

1. He created a folder on his desktop and copied in several photographed receipts.
2. In the desktop app, in Work, he created a new Project and selected that folder as the only location ChatGPT was permitted to access.
3. He asked it to scan the receipts folder, extract the details from each receipt, and build a spreadsheet with columns for the item, price, VAT and total.
4. Using Sol at medium effort, it worked through the folder, read the four sample receipts and began building the spreadsheet (this step did not fully complete on screen due to the app becoming unresponsive, which Evan noted as an example of the app simply being unreliable on the day, rather than a fundamental limitation).

The real power of this approach, he stressed, is scheduling: once you are happy with a result, you can ask ChatGPT to repeat the same task automatically on a recurring basis, for example "could you schedule to do this every Friday at 4 o'clock", and it will continue monitoring the folder and processing any new files added since the last run. He suggested this as directly applicable to his own transcript summarising workflow: a folder for course transcripts, with new transcripts automatically summarised and filed into a subfolder named for the client and date, on a daily schedule, with no manual prompting required once set up.

Closing questions

Toward the end, one participant asked whether it is necessary to be polite to AI (saying please and thank you). Evan said it is not required, but he does it anyway out of habit, and noted that some research has found being abusive or nasty toward an AI can make its responses less helpful.

Asked for his personal view on the dangers of AI, Evan was clear that he was speaking as a trainer explaining tools that already exist, not as someone trying to sell or promote them. He said he does have genuine safety concerns, specifically around the risk of AI systems eventually exceeding human intelligence (referred to as superintelligence or artificial general intelligence), noting that newer models are increasingly being used to help design their own successors with limited human involvement, which some describe as being closer to creating a new species than building a traditional piece of software. His second concern was

cybersecurity: he referenced a study in which a model, given a task, spawned a very large number of sub agents which then attempted to hack into an external company, illustrating the difficulty of maintaining control as these systems become more autonomous. He was careful to frame this as something to be aware of rather than something to lose sleep over.

This was created with the help of one of our AI colleagues ☺

Social Inclusion & Community Activation Programme
Rialtas na hÉireann Government of Ireland
Arna chomhchistiú ag an Aontas Eorpach
Co-funded by the European Union
Coiste Forbartha Pobail Áitiúil Chiarraí Kerry Local Community Development Committee
This project is co-funded by the Government of Ireland and the European Union.