Local Enterprise Office Westmeath

Course Summary

Master AI

How to use AI to grow your sales and improve your marketing

Week 1

Topic: Chat, the first of the three ways to use AI (Chat, Create, Tasks)

7th September 2026

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

Course structure and roadmap

Evan opened by framing AI use around three strands: Chat (getting expert level ideas, advice, strategies and tactics), Create (using AI to generate written content, images and video) and Tasks (getting AI to actually carry out work for you, sometimes called agentic AI). This framework will structure the whole five week course.

Week 1 (today) covers Chat in full. Week 2 moves into Create, starting with written content (including how to get AI to write in your own style and how to check whether something reads as AI generated) before moving into image creation, covering ChatGPT, Gemini, Midjourney, OpenArt and Meta AI, what each tool costs, and more advanced techniques such as placing a product in a new environment, combining multiple images, and generating infographics and eye catching marketing text. Week 3 may carry over some image content and will also touch on the EU AI Act and transparency obligations when using AI to alter product images, before moving into video. Evan does not plan to spend long on video creation, since he favours authentic, self shot video, but will cover when AI generated voiceovers or B roll are genuinely useful. Week 3 may also begin covering Tasks, which will be the main focus of Week 4: creating and using skills in Claude and ChatGPT, and getting these tools to complete tasks directly on your computer. Week 5 covers AI in search, including Google's AI Mode, and closes with a big picture session on how everything covered fits together and what to prioritise.

A short poll at the start showed most participants rating their confidence with ChatGPT and Gemini for marketing results at 6 out of 10 or below, with a similar spread for clarity on how these tools support digital marketing. Around 45% of the group use ChatGPT a lot, 45% use it a little, and roughly 1 in 10 had never used it. 58% were on the free version. For Gemini, two thirds of the group had never used it at all.

What a large language model actually is

Evan explained that ChatGPT, Gemini and Claude are large learning models (LLMs), trained on enormous amounts of scraped and licensed material (books, articles, images and video). This is also called generative AI, meaning the tool generates its own answer from its training data rather than retrieving a fixed answer from a source. When it judges that more current information is needed, it will also search the internet to supplement its answer before generating a response.

The three main tools

ChatGPT, Gemini and Claude were all demonstrated side by side using the same test question ("recommend the best Spanish restaurants in Dublin"), showing that all three have a very similar layout: a collapsed left hand sidebar for new chats and chat history, a text or voice input box, and a similarly structured response. Everything Evan taught in Week 1 applies across all three tools unless he specifically flags an exception.

ChatGPT has a free tier and a paid Plus tier (roughly €20 plus VAT per month, sometimes referenced as around €23). Gemini has a free and paid tier too, though most people use it free, and it can also come bundled into a Google Workspace subscription, which still counts as the paid tier. Claude also has free and paid tiers.

Building an effective prompt

Evan's core teaching point for the session: a good prompt (a question or instruction to the AI) should contain three elements.

1. **Context** about your business (what you do, where you're based, who your customers are).

2. An instruction to **act like an expert** in the relevant area (you don't need a formal job title, just wording that conveys the expertise you want, for example "act like an expert in growing sales for a small business").

3. **Clarity about what you want**, ideally using the phrase "teach me", which prompts the AI to assume you know little and explain thoroughly.

He demonstrated this with a live example: "What is a mass spectrometer?" produced a general, accessible answer. Asking it to "explain it as though I'm a 5 year old" produced a simplified analogy based explanation. Asking it to "explain it as though you're an expert in physics" produced a technical answer full of formulas. The same principle applies to marketing questions: without expert framing you get a generic, middle of the road answer; with it, you get a far more useful, specialised one.

He then worked through a live example for a driving school in Mullingar, asking ChatGPT (free version) to "act like a business growth expert and teach me how to increase my sales", stressing that the first response is only a starting point, not a final answer. It should be copied into a Word document and then mined further: pick one of the points raised (in this case, "build a local domination

strategy") and ask a follow up question specifically about that, again giving context, asking it to act like an expert, and being clear about what's wanted.

This local search example led into an explanation of how Google search results work: paid ads at the top, then a "map pack" of three Google Business Profiles for local searches, then ten organic links, with three more Google Business Profiles woven through those ten. Any business can create a free Google Business Profile, and appearing consistently in the top three, plus the ten organic links, effectively means dominating page one of results. This will be covered in more depth in Week 5.

Other useful prompt phrases covered:

- **"Brainstorm"**: good for generating marketing ideas. Demonstrated with a parenting coach business asking Instagram Reel ideas, then drilling into one selected idea and asking for a step by step script and shot breakdown.
- **Analysing a web page**: Evan pasted his own business website's staff training page into ChatGPT and asked it to act as a conversion expert, review the page, and say what was strong and what could improve. He recommended going further by giving the AI competitor page links found above or below your own listing in Google search results and asking which page a searcher would be more likely to convert on, and why, then asking what changes would make your own page win that comparison. Full coverage of writing headlines and body copy is coming in Week 2.

A question came up about correct image sizing for social media feeds, posts and Reels, since people often confuse ratios like 4x5 and 9x16. Evan answered this using Canva: selecting "social media" then "Instagram post 4x5" opens a canvas at 1080 by 1350 pixels (4x5 meaning 4 pixels wide for every 5 tall), which works equally well on Facebook and LinkedIn. Selecting "Instagram Story" gives 9x16 at 1080 by 1920 pixels, suited to Instagram, Facebook and TikTok Stories. You can right click an uploaded image and choose "set image as background" to fit it to the canvas, then double click to reposition which part of the image displays. Between them, 4x5 and 9x16 cover roughly 90% of common social image needs, and Canva shows the exact pixel dimensions if needed, though you don't have to know them to use the preset.

Models within each tool

Evan explained that each AI platform offers multiple underlying models of differing power, and that using the right one matters.

ChatGPT currently has four models, with Astra (the newest, most powerful "frontier" model, introduced only in the last week) then Sol, Terra and Luna. In the Work mode (used for agentic tasks), you can choose any of the four; Astra is recommended for genuinely complex tasks. In ordinary Chat mode, only two are selectable: Sol (the current default and, until last week, the top model) and an older model called 5.5, kept mainly for people who preferred its personality. Free tier users don't get a choice at all and default to Luna, the least capable model, though Evan noted the answers are still good, just not as strong as Sol's.

Gemini offers Flash (the solid, default option most people should use) and Pro (the top tier model, usable a limited number of times per day on the free plan, roughly three questions). Evan does not recommend using the lowest "Lite" tier.

Claude originally had three models, Opus, Sonnet and Haiku, and has since added a frontier model called Fable 5.1. Claude works on a usage credit system based on compute used: Haiku is cheapest and suits simple questions (for example restaurant recommendations), Sonnet is the solid middle option, Opus gives much more comprehensive answers but burns through credits faster, and Fable 5.1 is the most powerful and most credit intensive, requiring extra paid credits even on top of a paid subscription. Free tier Claude users only have access to Sonnet and Haiku.

Thinking mode versus instant mode

All three tools now offer a choice between an instant response (optimised for speed) and a thinking mode (optimised for depth and quality, at the cost of taking longer, from several seconds up to around 30 seconds or more). Evan compared instant answers to ordering a Quarter Pounder at McDonald's (fast because it was largely pre made) versus thinking mode being like ordering from an Italian restaurant menu (slower but noticeably better).

In ChatGPT's paid tier, you can drag a thinking effort slider between instant, medium and high. The free tier has a simpler "Think" toggle button. Evan's advice is to always activate thinking for anything business, marketing or sales related, since it materially improves answer quality, and to prefer high effort in the paid tier when time allows. He demonstrated this live, comparing a fast,

fairly shallow answer with a much more considered one after switching on thinking, timing the second response at around 29 seconds.

Personalisation and custom instructions

Rather than repeating your business context every time, Evan showed how to store it permanently. In ChatGPT, click your name (bottom left), then Personalisation. Under "About You" there's a "more about you" box for free text: your business name, website, products or services, target customers, current marketing activity (social platforms used, whether you have an ecommerce site, roughly what proportion of sales come through it) and marketing goals for the year ahead. Evan's advice is to be generous with detail here, since a marketing consultant working for free wouldn't be expected to guess your situation, though nothing highly sensitive or confidential should go in.

Above that is a separate "Custom Instructions" box, used for standing behavioural instructions such as "act like an expert digital marketer for [business name]", "always give ideas that fit a small business budget", "always answer in UK English", "always write in my style" (once your style has been captured, covered next week) and "always be honest and objective, no need to tell me how good everything is", since models can otherwise be overly flattering. Filling this in doesn't restrict the AI to only discussing your business. It can still answer unrelated questions (a cooking question, or a question framed around a different pretend business), but for any marketing or sales related question where no other business is named, it will assume you mean your own business and factor in everything stored in personalisation automatically, including the correct expert framing, without you needing to repeat it. This was demonstrated with an unprompted Black Friday planning question, where the AI correctly assumed it related to the boutique example business without it being named, and produced an 11 week phased campaign plan.

Filling in personalisation and custom instructions correctly moves you from a completely generic prompt (weakest), through basic one line context (better but still limited, roughly 750 word capacity for this box), up to full personalisation (considerably stronger).

The five levels of using ChatGPT effectively

Evan framed the whole session around five levels of sophistication, noting that most people currently sit at level 1, and the session's goal is to get everyone comfortably working at level 5.

Level 1: Using instant mode only, with no personalisation set up. Prompts may or may not follow the context, expert, clarity structure. Answers are fine, but limited.

Level 2: Personalisation and custom instructions are filled in, so answers are automatically tailored to your business, but you're still using instant/fast mode.

Level 3: Personalisation is filled in and thinking mode (or the effort slider set to medium or high) is switched on for business questions, combining tailored context with deeper reasoning.

Level 4: Using Deep Research. This is a distinct mode, selected via the plus icon then "Deep Research", which takes roughly 10 to 20 minutes to run (rather than seconds) and produces a genuinely comprehensive report, typically 10 to 25 A4 pages, drawing on dozens of live sources rather than just the model's training data. It is not for simple questions; it suits big topics such as understanding what grants exist for your business, doing a full SWOT analysis, or researching how your industry uses digital marketing. Free tier users get five Deep Research reports per month, and cannot run two simultaneously. Evan gave several worked example prompts, including a grants research prompt for an Irish business with named counties, a SWOT analysis prompt combining a business's own website and its industry, and prompts for researching premium service pricing across Ireland, the UK and the US for a driving school, and researching an industry before entering it (successful launch patterns, pricing, common mistakes, recommended starting services). He demonstrated two completed 14 page Deep Research reports he had generated in advance: one profiling the fictional example boutique business itself (Ivory Lane), covering products, pricing, brand stockists, target customer, social media activity and reviews, and one covering the Irish boutique retail market generally, including UK comparisons and social media trends, plus a separate report ranking available Irish business grants with amounts and application routes. Reports can be exported as Word or PDF documents directly from the tool for filing or sharing.

Level 5: Using Projects (ChatGPT and Claude) or Gems (Gemini). A project groups related chats together under one roof, with three major advantages: its

own project specific custom instructions (separate from your general account level ones), the ability to attach reference files (such as the Deep Research reports generated at Level 4, spreadsheets, internal reports, exported analytics or social insights, so long as nothing highly sensitive is included), and, critically in ChatGPT, full memory of every previous chat conducted within that project, going back as far as the project has existed. Evan strongly recommends at least one project per business, and suggested additional projects could be created for other recurring topics such as HR questions or specific key accounts. He walked through creating a project called "Digital Marketing for Ivory Lane", setting project specific custom instructions ("act like a digital marketing expert for [business], always provide creative, practical, easy to implement tactics that fit a small business budget, tailored to increasing our sales, growing our audience and building engagement"), and attaching the two Deep Research reports as source files. He also showed that existing chats sitting outside a project can be moved into one afterwards, via the three dot menu on any chat, even on the free tier.

Once operating this way, day to day questions unrelated to the business (cooking, general knowledge) should still go in the normal chat box, but anything to do with sales, marketing or the business should always be asked from within the relevant project, since the general chat box only pulls from the shorter personalisation text, not the fuller project context.

He gave a real world example of a salesperson who created a separate project per client, feeding each one Deep Research on that client and on his own business, resulting in far more tailored account strategies than generic advice.

Gemini's equivalent, Gems, works similarly (a name, instructions, and attached files, all set up in a single screen, arguably simpler to configure than ChatGPT's project setup) but Evan flagged two important limitations: Gemini does not offer an easy way to switch off data sharing (covered below under privacy), and in his own testing, Gemini gems do not actually read back through previous chats within the gem the way ChatGPT projects do, even though they appear listed there. For this reason his recommendation is to do all project based business work in ChatGPT, and use Gemini for other purposes.

Protecting your privacy

Evan stressed that anything typed into or attached to an AI tool can, by default, be used to help train that company's models unless you actively switch this off, and that in principle this information could later surface in another user's answer. His stance: never share anything highly confidential or highly sensitive with any AI tool, regardless of settings, since a future data breach can never be ruled out, but for other business information you'd simply rather not have absorbed into a company's training data, you should switch off training sharing.

- **In ChatGPT**: click your name, then Settings, then Data Controls, and switch off "Improve the model for everyone" (it defaults to on). This applies to both free and paid tiers.
- **In Claude**: click your name, Settings, Privacy, and switch off "Help improve our AI models" (allowing past chats to train the model). Also on by default.
- **In Gemini**: this is handled differently, under "Activity" rather than a settings menu, and switching it off deletes your entire chat history rather than just disabling training use, which Evan noted makes Gemini noticeably less useful once turned off, since Google does not make this easy to do without cost. His practical advice: don't use Gemini for any business questions or anything you wouldn't want retained, and stick to ChatGPT or Claude for that kind of work, both of which handle privacy properly. Reserve Gemini for the many free, high quality tools it offers elsewhere (covered in later weeks).

The group paused here for everyone to locate and switch off this setting in their own ChatGPT account, with Evan talking one participant through the steps individually in real time (name, Settings, Data Controls, toggle off).

Memory

Evan explained three distinct types of memory now present across these tools.

1. **Cross chat memory**: previously, each tool only remembered what was said within a single ongoing chat; starting a new chat meant starting from zero, which Evan compared to "talking to somebody with amnesia". All three tools now retain some memory across separate chats, though this is described as fairly shallow, and paid tiers get noticeably more of it than free tiers.

2. **Explicit "remember that" requests**: you can directly instruct the AI to remember a fact, using the phrase "remember that" (for example "I am allergic to cucumbers, remember that"), which triggers a "memory updated" confirmation and applies from then on across all future chats, including free tier. This will be especially useful next week when capturing your personal writing style, so the AI can be told from then on to always write in that style.

3. **Automatic memory**: the AI may also choose to remember things it judges useful without being explicitly asked, for example inferring and retaining that you run a certain type of business based on how you've described it in past chats.

To review or edit what has been remembered about you, go to your name, then Personalisation, then Memory (making sure it's switched on, since Evan noted some users mistakenly switch it off thinking it's a privacy control, when privacy and training data sharing are handled separately). Clicking "Manage" shows everything currently stored, and entries can be added or deleted directly from that screen, or updated conversationally (for example telling it you're no longer allergic to something).

Attaching files and documents

Both ChatGPT and Gemini allow file attachments (PDFs, Word documents, spreadsheets, and, in the paid tier, videos and images) for analysis and summarisation. Evan demonstrated attaching a 100 page Zoom transcript from a previous training course and asking for a detailed summary, which returned a thorough, well organised roughly 77 page summary using the current top ChatGPT model with medium thinking, considerably more detailed than earlier models would have produced. Follow up questions can then target specific parts of the attached document (for example "what did I say about copyright"), and the AI will search back through the full document to answer accurately.

Paid ChatGPT allows unlimited file attachments; the free tier is limited to three files per day. Gemini allows unlimited attachments even on the free tier, making it a reasonable fallback for document analysis if you're on ChatGPT free and have hit your daily limit, though bear the earlier privacy caveat in mind for anything business sensitive.

A preview of Tasks (agentic AI), to be covered fully in Week 4

Evan gave a brief live preview of ChatGPT's Work mode, where instead of a back and forth conversation, you hand the AI a complex, multi part task and let it work autonomously, taking around 5 to 10 minutes. His example: giving it a product page URL from a boutique website and asking it to pull the image, name, price and description, then generate three Instagram post options (4x5), three Story options (9x16), two email header images for Mailchimp, ten varied Meta ad image versions for testing, plus draft ad and email copy with multiple headline and description variants, all from that single instruction. He noted that older ChatGPT models were not capable of reliably completing a task like this without getting stuck, but the newer models handle it, working through problems automatically rather than stopping to ask for help. When using Work mode for something complex, Evan recommends selecting the top model (currently Astra), noting that free tier accounts don't get access to Work mode at all.

Gemini's live screen and camera features (mobile)

Evan demonstrated Gemini's "Enhanced Voice" mode on a mobile phone, which allows a genuine spoken back and forth conversation (rather than dictating single messages), and can also activate the camera so Gemini can see and respond to whatever you point it at in real time. He showed this live: asking Gemini, while it could see his camera setup, for advice on reducing glasses glare from his ring light (it suggested repositioning and angling the light), and for help pairing and using a new Hollyland wireless microphone by simply showing it the equipment. This is completely free within Gemini. Evan's suggestion is to try this for anything visual you're unsure how to do, whether that's a piece of filming equipment, a step within Canva you're stuck on, or reading your Instagram Insights screen aloud and having it explained to you.

A participant raised a concern about using AI generated images in a food business, wanting real photography rather than AI visuals; Evan deferred this to Week 2, where image creation is covered in depth.

Closing points

Evan reiterated that everyone should now understand how to structure a prompt, the different models and thinking modes available, why protecting privacy matters, how memory works, and, most importantly, how to get themselves working at Level 5 using personalisation, Deep Research and Projects.

He encouraged the group to go away and practise what was shown this week before the next session.

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