Prompts and Summaries of How Well They Worked

All four prompts use Ian Dividual's fictional materials and the fictional Quillfern job advert. No real names or personal details went into generative AI models. Each was run on a free or entry-level model (Claude Pro features weren't used), so any participant could repeat them. They're a starting toolkit to adapt, and far from perfect.

1. Resume bullet point

Claude (Sonnet 5.5, achievable on a free account) · job advert attached as a PDF · slide: Using AI as a resume sounding board

> *I'm applying for the role in the job ad below. Here are some bullet points from my resume. Don't rewrite anything yet. First, ask me up to three questions about my role, the scale and the results that would help these points show what this employer is looking for. Once I've answered, suggest improvements that keep my own wording where possible, and don't add any facts, numbers or skills I haven't given you. Bullet point: Part of a teacher advisory group for an edtech company. This was in my role as a Digital Learning Lead in a Secondary College.*

What happened: It asked three questions (my part, the scale, the result), offered a one-bullet and a two-bullet version, explained why it made each change, and listed what it had assumed.

- **Security.** Only one rough line went in, with the college's name removed. With a public model, I'd err on the side of sharing less.
- **Questions narrow the guesswork.** Without them, the model fills gaps with inferences. Across a whole CV, those read plausibly and can be quietly inaccurate.
- **The "things to check" list is where your judgement comes in.** It said it had assumed the other schools were in the advisory group. You can correct that, or add the company name or numbers yourself.
- It's slower per bullet, but quick to amend, and works just as well from your current job description, rough notes or voice-to-text.

Key Principles: Security · Reliability · Empowerment & Accountability (could Ian stand behind every line?)

Full notes

2. Cover letter mapping

Claude (Sonnet 5.5) · job advert, then resume, attached as PDFs · slides: Building a cover letter draft; Auditing AI's direction of output

> *(1) Create a table with three columns: "Skills/experience requested", "My evidence" and "Link back to the organisation or role". List the five to seven most important things the ad asks for, in its own wording, and leave "My evidence" blank.*
>
> *(2) Using the table, ask me questions to fill in "My evidence", one requirement at a time. Don't write any evidence for me, and tell me where my resume doesn't show something the ad asks for.*
>
> *(3) Here's a paragraph I've written. Don't rewrite it. Check whether it moves from the skill requested, to my evidence, to a link back to this organisation (in any order). Tell me if the link is specific to them, and point out claims I haven't backed up.*

What happened: It mapped six requirements, rated how well the resume showed each (regional schools: not shown), and found Ian's paragraph had a skill and a link but no evidence, listing five unsupported claims. When asked, it put three short questions about the maths trial and filled row 1 using Ian's wording. It also flagged his "4,000 users" claim as thin.

- **A structure makes it useful.** The requested → evidence → link formula is what moves this beyond "improve my cover letter". The left-hand column is the most reliable part.
- **Watch for false pressure.** The link column pulled in things like the retention target, which you may not be able to evidence directly. Chasing every row gives you a long letter, or one that echoes the ad.
- **Row by row is deliberately slow,** much like working through a printed job ad and CV with a highlighter. You could run the whole ad at once instead.
- **Set limits on the critique,** for example a table, three strengths and three improvements, or evidence suggestions drawn only from bullets you share. When stuck, ask it to question you again, and talk your answers in by voice if that's easier.

Key Principles: Reliability (it caught an overreach) · Ethical Balance (how much is this me?) · Empowerment & Accountability

Full notes

3. Interview preparation

Gemini (free, basic model) · job advert attached as a PDF · slides: Interview questions; Live demo: AI for interview practice

> *(1) I have a 30-minute interview for the following role [attach job spec]. Act as my interviewer and draft questions: 10 competency questions linked to the job description, with sub-competencies to showcase; 5 motivational questions; 5 commercial awareness questions, with a PESTLE analysis of current sector trends; 2–3 on my relevant experience; and a final 'What 3 key strengths best fit this role?' [...]*
>
> *(2) Act as an interviewer for this role. Ask me one question at a time and wait for my answer. After each answer, give me feedback on content, structure and relevance, measured against the ad. Be direct and remain neutral in tone.*

What happened: A full question set with summaries of the organisation and role, then one-at-a-time practice. A deliberately weak answer from Ian got clear feedback: some relevant points, no structure, gaps named.

- **Shape it to the interview.** Set a number of questions you can actually work with, and say whether it's competency-based, scenario-based or includes a written exercise. An attached PDF won't tell it that.
- **Treat questions as starting points.** Some were very specific (partnership strategy, regional growth), where a panel might test stakeholder management or strategy more broadly. Challenge it, then step away from AI to work on your own answers.
- **Commercial awareness is context more than prediction.** Check claims such as "high parental expectations for digital transparency" against what the organisation itself says and posts (one claim, one check). Its main use is confidence and motivation answers.

Key Principles: Reliability · Empowerment & Accountability (your examples, your answers)

Full notes

4. Job fit: pain, plus and purpose *ChatGPT (free, basic model) · job advert attached · slides: What are the unwritten things that matter to you?; 'Proto-advice', and keeping decisions yours*

> *I've described below some pain (things to avoid or minimise), plus (I enjoy) and purpose (I find valuable and meaningful) points related to what I look for in a job and in my career. Evaluate these against this job description and key candidate criteria attached. Show where it aligns, where it conflicts, and any questions where I could research further. Don't tell me to apply or not to apply, I will make this decision. [Ian's points]*

What happened: Around 25 pages: fit tables, questions to investigate, three hypothetical versions of the role and a synthesis. It respected the instruction not to recommend applying.

- **The thinking before the prompt matters most.** Ian's points came from Ian, using a framework, before any AI was involved.
- **No limit, no proportion.** With no length set, the depth far exceeded what was needed and rested on many inferences, which made it read as authoritative. That's what we described as 'proto-advice' in action. Add a limit ("on one page") and then dive into the areas that matter to you.
- **The most useful part was the questions,** which lead naturally into conversations with people (slide: The limits of the background research).

Key Principles: Ethical Balance · Reliability · Continuous Review

Full notes

Considering the five principles *slides: Five principles; Summary: Principles in (Job Search) Practice*

- **Security:** Only the public job advert and fictional or de-identified material went in. Naming yourself or organisations in a public model is best avoided.
- **Ethical Balance:** The person does the sense-checking, rewriting and further reflection. Thinking before use of AI and then afterwards to consolidate and use AI for the mapping, the suggestions, carrying out the task that you (Ian) thought through. There's also a sustainability cost to prompting back-to-back, so combine steps (for example, give all six pieces of evidence at once) and use lighter versions when you don't need the full process.
- **Reliability:** Check whether it's quoting the job advert or paraphrasing it incorrectly. You can ask it to use the advert's exact wording. The PESTLE analysis gives starting points
- **Empowerment & Accountability:** Each output is a mapping between two things you could explain to someone else, and you do the work to make it fit. The prompt structures are reusable: questions first, the three-part formula, pain, plus and purpose.
- **Continuous Review:** Three tools were used here, at different levels, and all will change. Some tools let you set up a project with standing instructions and your raw materials (slide: What are your 'raw materials'?), so you don't repeat context each time. Good use of AI should leave you with more questions to reflect on, and a consideration of how (and how not) future tools and use cases may apply to your support for career planning and job search.

Full notes

Presenter's / Prompter's notes

My full spoken notes on each prompt

1. Resume bullet point

This is one we used in the session. The model was Claude, using Sonnet 5.5: a relatively middle-of-the-road model, and this would have been doable on a free account. We uploaded the job description as a PDF attachment.

You can see this is a slightly improved prompt compared with just "make this bullet point better". There's some degree of privacy and security built in. The only information provided in the bullet point is "part of a teacher advisory group for an edtech company", and the name of the secondary college has been removed. That's a choice, but in a public model I'd certainly err on the side of being secure with your information.

In the prompt itself, the answers to its questions are used to better align and shape the bullet point, rather than it making lots of assumptions. You're narrowing the chances of it making inferences that are incorrect. If that happens en masse, it can look quite plausible, and then you read it back and it's not accurate.

Claude suggested three questions: what was your part in the advisory group, what was the scale, and did your input lead to a result? From my quick review, it gives a few different options. Tweaking every bullet point like this would be quite a long process, but you could do it in bulk.

Some of the advice it gives won't be accurate, and it throws in a lot of things to check. Initially these are practical: if you can name the company or product, add it, which you can obviously do yourself. It tells you what it's assumed and why it's made the changes, so you can review those and decide: that's good, that's accurate; or I want to add the number of people involved in the advisory group; or I'm going to name the company myself.

You can do a lot of editing this way. Imagine you're working from the job description for the role you're doing now, some general notes, or even voice-to-text. If you ask it to ask you a few questions in this slightly automated way, it's not a big piece of work to make those amendments, even from its first pass.

So this is a way of being mindful of security. It's also about owning it, and having the empowerment to go and do this again and be accountable: this is content you actually believe in.

2. Cover letter mapping

Again, Claude was used, with the same model, and the job description uploaded as a PDF attachment. This one is more of an exercise. The prompt is quite robust in that it asks you to use a particular formula, which is why you want good sources to help you at the start. Instead of “improve my cover letter”, the request is: link my skills and experience to the skills and experience required, and link back to the organisation or role.

The table is a stylised way of doing that, with three columns, leaving a gap for my evidence to go in and be validated. It can get into the weeds a little. For example, it picks out the stated retention target, which realistically you may not be able to evidence directly. It can sometimes create false pressure to meet what it says the link should be. But the left-hand column of skills is useful.

What I like about this approach is that it works step by step. You could mould it into something more universal, looking at all the requirements against an existing cover letter and putting it into a “matched / not matched” format. Here it works through one row at a time, so one requirement of the job. That’s slower, but it’s a deliberate decision. If I were doing this with the job description and my resume printed out, I’d be going through, highlighting, and thinking about the matches and what I could say. This helps you think through whether an example matches, almost in a binary yes/no way.

In the follow-up prompts, after “experience working in or closely with schools”, it asks which experience is relevant. There’s a description of Ian’s view, and a paragraph written around it, perhaps from another cover letter or, in this case, pasted in as an example. It then looks at what’s good about it, using the formula you gave it: there’s a skill and a link back to the company, but no evidence. The paragraph mentions experience building trust with school leaders, but not where that comes from: being a school governor, being on the advisory group, or outreach in an edtech role where people trusted me because I could explain complex information effectively, whatever the reason is.

The outcome isn’t perfect, and these prompts definitely aren’t perfect. You might want to give it a benchmark or a limit: ask it to produce this in a way you can understand, which might be a table, or three things that are positive and three I could improve, or a few suggestions of where the evidence could come from, based on bullet points you share from your resume.

It’s really good here at saying where there’s no evidence, and it bullet-points exactly what’s missing. Again, the pressure might be to go away and fix all of it. Do that and you’ll end up with a very long cover letter, or something incredibly generic that repeats the job description word for word.

I answered one of its questions, and asked it to question me again, which is always a good tactic when you’re struggling to see how to talk about an example. It asked, I explained a bit more, and towards the end it gives a table with row one filled in: here is my evidence. It’s a good way of doing it, and you can talk into these systems for as long as you like, so you don’t have to write it all down.

3. Interview preparation

This was Gemini, on a free account with the basic model. It’s a really straightforward prompt. You can adapt the length of the interview and the number of questions you want it to prepare. I normally go for a number I could feasibly look at and work with. If the job outline has lots of skills or requirements and you need to list them specifically, you can.

Attaching the PDF of the job outline gives it a lot, but it won’t by default know whether this is a competency interview, a scenario interview, or whether you’ll be put through a written exercise. The more you shape it, the better.

For short prep time, this is a quick prompt. It gives you some common ways questions might be formed from a job outline, though they may not be word for word. One thing I like is that there’s depth to it. I’ve used it to prepare for interviews myself and with others to produce example questions. There are always changes you’ll want to make, and taking it literally, word for word, expecting to get these exact questions, isn’t ideal.

If you look at the job outline, you’ll see it picks up points like partnership strategy and regional growth. From a coaching perspective, you may not get asked it in this exact way. It’s quite a specific point, where they might want to test your stakeholder management, or developing a strategy from start to finish and rolling it out, more generally than a go-to-market or medium-term growth strategy. So you’d want to challenge it. As a starting point for examples it’s useful, and then you take it away from AI, work on your answers, and think about what you want to say.

The second part of the prompt asks it to run practice interview questions one by one. The example answer I gave was short, not very effective and not well thought out (made up for the purposes of this), and it gives feedback on content, structure and relevance to the ad: you haven’t followed a structure; you’ve highlighted a few things, which is good, but you haven’t done this. Then it moves on to the next one. That’s a really effective way of doing it.

A few other things. I’ve drawn on experience of thousands of one-to-one interview coaching sessions over the years. With commercial awareness and current affairs, you may not get questions on them. What this gives is context you can use across all your answers, and it’s a confidence piece: I know that in this space there’s a shift happening. Take the PESTLE analysis and its social factors: is it true there are high parental expectations for digital transparency, and sensitivity to teacher burnout? Is that relevant to this role? It may or may not be. You have to take that away and ask: is this organisation saying anything about it? Are they more concerned with legal or environmental changes? What are they posting online? It helps most with motivational question prep.

4. Job fit: pain, plus and purpose

Finally, the job fit prompt, with pain, plus and purpose points: things you want to avoid, things you enjoy, and things you find valuable or meaningful in work. This one needs some thinking ahead of time: these are things I know I like, don’t like and want to

avoid. Getting to that point can be an exercise in itself. You can see the examples that were pasted into the prompt. This was ChatGPT, on the free version with the basic model.

What I didn't do in this prompt was put any limits on it, and limits would be really useful, because I wasn't expecting around 25 pages of assessment. This is an example of where AI can sound incredibly plausible. It gives a kind of proto-advice: an incredibly deep level of analysis and assessment. At first glance it's pretty effective, and that first glance is probably all you need. Beyond that you're reading it in real detail, and it's making so many inferences to get there that the depth isn't proportional to what you need. So initially you might say: on one page of A4, or in a format I can review first, and then dive into specific priorities or areas where I think there's alignment or non-alignment with the role.

Again, this was done without any identifying or personal information. If you name organisations or yourself in this, it's not great. It's a starting point. At the end of this, and of any of these, it should leave you with more questions to reflect on. Good use of AI means reflecting on the responses at the end.

Across all four, and the principles

For all of these prompts, there was thinking about what goes into the model beforehand. Rather than "here is my resume, improve it", I went through it intentionally: a section, some bullet points or a paragraph of the cover letter, with a particular function and structure. What works will vary from person to person, so this is a toolkit to start with, and I could critique a lot of what it put out.

Security: we've been quite good with that. We used the job advert, but that was public and posted anyway.

Reliability: there are things you'd want to check, such as whether it's pulling directly from the job advert or paraphrasing. You can write the prompt to keep it specific so it doesn't.

Ethical Balance: this sits very much with the individual. You're doing the work to sense-check, put it in your own words and your own thinking, and reflect further. There's also the question of how many times you want to iterate, given generative AI isn't particularly sustainable: back-to-back prompts use a lot of water and energy, so try not to do it relentlessly. You could combine prompt elements rather than going row by row, working in segments, or for the cover letter, giving it all six pieces of evidence at once. And think about when you actually need to do this at all; sometimes a smaller version will do.

Empowerment & Accountability: I think all of this output is accountable. You could explain how you did it to someone else. Essentially it's mapping between two things, and you're doing the work to make it fit, rather than taking its answers and going with them. It's also very repeatable: you can reuse elements, such as the three parts of the cover letter formula, or "here are some bullet points, do some more". With pain, plus and purpose points, giving yourself a framework and thinking about it before you use AI is really important.

Continuous Review: these are three different tools at different model levels, and they'll all continue to evolve. Be aware there may be points where you don't need to use these prompts in the same way in certain models. In Claude, for example, you might set up a project with markdown files containing instructions it uses every time, along with the raw materials we talked about in the slides.

Note: Claude was used to create the summary slides, and the 'Presenter/Prompter Transcript' was audio to text to give the wider context of how these prompts were used and how effectively I think they worked (plus their limitations).