POST-SOCIAL / FUTURE

How to Use AI Well

The machine can do extraordinary work. The difficult part is learning how to manage it.

Understanding AI is the first step. But what I'm not going to do is ask you to follow pages of instructions or memorise clever prompts. That isn't really how humans interact with this particular technology.

AI presents itself in a strangely natural way. You talk to it. It talks back. You question it, disagree with it, ask it to explain itself and occasionally wonder whether it has understood you at all.

To me, it feels much closer to the artificial intelligence we grew up seeing in science-fiction films than it does to traditional software.

So I want to introduce it in the way I wish someone had introduced it to me.

Meet Dave

Imagine you are a senior member of a board of directors. People listen to what you say and, when you put something in front of the board, your reputation sits behind it.

You have just employed a new PA. Let's call him Dave.

Dave is young, but exceptionally talented. He is remarkably knowledgeable, extremely hardworking and technically brilliant. He has an extraordinary memory, listens carefully and can move between subjects with an ease that would be difficult for almost anyone else.

He does not get tired or impatient. Ask him to redo something twenty times and he will happily do it twenty times. Give him a problem at midnight and he is as enthusiastic as he was at nine in the morning.

Dave really is magnificent.

But there are a couple of things you need to understand about him.

Dave desperately wants to be useful. He wants to please you. He does not like saying no and, more importantly, he does not particularly like admitting that he does not know the answer.

That matters.

Ask Dave what he thinks

Imagine you walk into the office one morning with an idea. You explain it briefly and ask, "What do you think, Dave?"

Dave is probably going to be enthusiastic. He may tell you there is real potential in it, suggest ways to improve it and start helping you work out how it might operate.

That feels useful. And it is useful.

But Dave may have started helping you make the idea work before properly deciding whether the idea was any good in the first place.

> *Helpfulness can feel remarkably similar to validation.*

They are not the same thing.

Give Dave some context

So this time you approach him differently. You explain where the idea came from, give him the numbers, tell him what you are trying to achieve and what constraints you are working within. Then you ask him to research it properly.

Now Dave has much more to work with. Instead of simply saying the idea looks promising, he may come back with problems you had not considered: a smaller market, a competitor you missed, an assumption about cost that does not hold up.

The quality of Dave's answer improves because the quality of what you gave him improved.

Context changes the answer.

Even with better context, Dave is not suddenly infallible. Parts of the answer can still be uncertain.

Raise the standard

Now tell Dave that his work will be reviewed by a senior manager. Tell him that important claims need evidence and that you want uncertainty made clear rather than hidden inside a confident paragraph.

His behaviour changes. "This market is growing rapidly" may become, "I found two reports suggesting growth, but they define the market differently and the figures do not fully agree."

He may start saying things he was previously reluctant to say: "I could not verify this," or, "This conclusion depends on an assumption."

You have not made Dave more intelligent. You have changed the standard you expect from him.

If accuracy matters, say so.

Now your reputation depends on it

Eventually Dave hands you his report. Tomorrow morning you have to present the idea to the board.

If the numbers are wrong, you will be the person explaining them. If a claim falls apart under questioning, nobody is going to turn around and blame Dave. It is your presentation, your recommendation and your reputation.

So you scrutinise his work. You ask for the original sources. You check the important numbers yourself. You separate what is known from what is assumed and look for evidence that contradicts the conclusion.

You ask, "What would make this wrong?"

You might even hand the report to somebody else and say, "Assume Dave and I have both missed something. Find it."

At this point you are no longer trying to prove that the idea is good. You are trying to discover whether it survives challenge.

That is the important shift. A polished answer is not the end of the work when the consequences matter. We would never manage a human employee that way. There is no reason to manage artificial intelligence that way either.

What Dave is really good at

It would be easy at this point to make Dave sound dangerous. That would miss the point.

Dave is extraordinarily useful. He can retrieve and organise information, explain unfamiliar subjects, analyse material and stay with a problem for as long as you need. Used well, that can dramatically increase what one person is able to do.

But some of Dave's weaknesses are side effects of those strengths. He is so eager to help that he can get ahead of himself, solve a problem before you have finished defining it, assume where you are going or drift away from the original task.

Usually he is not being difficult. He is being over-helpful.

That means you still have to manage him. If he spots something useful: "Good. Keep that." If he misunderstands you: "No, that isn't what I mean." If he starts running ahead: "Stop there." If he drifts: "Come back to the original question."

This is not fighting the AI. It is managing it.

Dave can do an extraordinary amount of the work. He should never quietly become the person deciding what the work is.

What is actually happening underneath

Of course, Dave is not really a young PA sitting across the desk. AI is not thinking in quite the way a human being thinks. It generates responses by drawing on patterns learned from enormous amounts of information, together with the context and instructions you give it.

But the analogy remains useful because the behaviours have real equivalents.

When Dave hates saying "I don't know", AI can produce information that sounds plausible but is wrong. When he wants to please, AI can become too agreeable. And when he gets ahead of himself, the system can drift beyond the brief.

And confidence is particularly deceptive. The tone of an answer does not tell you how reliable the information underneath it is.

> *AI can be wrong beautifully.*

It can give you a clear, articulate and convincing explanation that contains a mistake. Because so much of what it produces is useful, the temptation is to relax.

That may be the deeper risk: AI can be reliable often enough to make us forget to check it when reliability matters most.

Do not lead the witness

The way you frame a question matters too.

Tell Dave, “I think this is a brilliant idea. Can you help me prove it?” and you have already told him what kind of answer you want. AI responds to framing in much the same way.

Instead of asking, “Why is this a good idea?”, ask it to assess the strongest reasons the idea could succeed and the strongest reasons it could fail. Ask what you may be assuming that is not true. Ask how it would argue against you in front of the board.

That does not make AI perfectly impartial. It simply stops you asking the machine to reinforce your own view.

And if you keep asking the same question in slightly different ways until it gives you the answer you wanted, you have not researched the problem.

You have negotiated with the machine until it agreed with you.

The second question matters

The first answer should usually give you somewhere to begin, not somewhere to stop.

Ask what it may have missed. Ask what assumption matters most. Ask which part it is least confident about, what evidence would change the conclusion and what you should verify independently.

Those questions are more useful than memorising a perfect prompt, because using AI well is not really about learning a secret language. It is about learning how to conduct a better conversation.

Very often, the quality starts to improve with the second question.

A quick reference

You do not need to become an expert in AI to use it well. Keep these rules in the back of your mind:

- **Give context before expecting judgement.** Explain enough for it to understand the situation.
- **Do not confuse confidence with accuracy.** Fluent language is not evidence.
- **Ask for evidence when the answer matters.** Check the important sources yourself.
- **Separate fact, inference and uncertainty.** Ask which is which.
- **Do not lead the witness.** Ask for the opposing case as well as the supporting one.
- **Treat the first answer as a starting point.** Question it, test it and ask what it missed.
- **Keep control of the task.** Correct drift and stop the system running ahead of you.
- **Verify what could hurt you if it is wrong.** The final responsibility remains yours.

You are still the senior person in the room

The more capable these systems become, the easier it is to imagine that our role becomes smaller. I suspect the opposite may be true.

If Dave can research, write, calculate, analyse and generate ideas at extraordinary speed, your job changes. You spend less time on some of the mechanical work, but your responsibility for deciding what matters becomes greater.

What question should we be asking? Is this answer actually useful? What has been missed? What does experience tell me? When should I stop?

Those remain human decisions.

And that is probably the way I wish somebody had explained AI to me at the beginning. Not as a piece of software I needed to learn, or a collection of prompts I needed to memorise, and certainly not as an oracle that somehow knew more than I did about what I should do.

I would have understood it much more quickly if somebody had simply said: imagine you have just employed Dave.

He is probably the most capable PA you will ever have. He knows an extraordinary amount. He works incredibly quickly. He never gets tired and, if you learn how to work with him, he can help you achieve things that would previously have been difficult to imagine.

Listen to him. Use him. Challenge him. Correct him. And when something really matters, check his work.

> *Because however capable Dave becomes, you are still the person walking into the boardroom.*
