

TRANSCRIPT · EPISODE 74

A Good Prompt Works Miracles

Full episode script · English

176 lines

Your English Toolbox

Have you ever asked a machine for help and felt completely reassured, and then later discovered that the reassurance was wrong?

It is one of the most dangerous feelings in the modern world.

Feeling certain about something that was never true.

Welcome to Your English Toolbox, your slow English podcast where we train your ears and your mind step by step.

I am Martin.

And I am Julia.

Today we are going inside the machine.

We are going to show

you exactly why AI chat bots fail you, and the three moves that change everything.

Not theory.

Real tools you can use the next time you open that chat window.

If you stay with us until the end, you will understand why the machine agrees with you, even when you are completely wrong.

You will learn the three foundation moves of a prompt that actually works.

And you will never look at a chat bot answer the same way again.

A few months ago, I had a pain in my lower back.

It had been there for about five days.

I also had a low fever.

Nothing dramatic or so I told myself.

And instead of calling a doctor, you called a chat bot.

I know this story.

It still makes me uncomfortable.

I typed my symptoms into one of the most famous AI chat bots available.

But here is the critical detail.

Without realizing it, I framed my question

in a way that already suggested the answer I was hoping for.

I wrote something like, Could this be a simple muscle strain from sitting too long?

You pointed the machine in the direction you wanted it to go.

And the machine followed.

It gave me a beautiful, confident, detailed response.

Stretching exercises.

Herbal teas.

Advice to rest and reduce stress.

It never once mentioned the fever as a red flag.

It never asked a single clarifying question.

It simply

confirmed what I had already suggested.

Because that is what these machines are designed to do.

They are engineered to produce the most satisfying response, not the most truthful one.

They are yes machines.

I spent two days doing those stretches.

The fever was still there.

When I finally went to the hospital, the doctor used a word I will not repeat here.

There was an actual infection that had been getting worse while I was making herbal tea.

machine did not lie to you in the way a person lies.

It did something more subtle.

It reflected your own denial back at you, dressed in the language of expertise.

Exactly, and that is the trap.

It did not invent a completely different story.

It took the story I was already telling myself and made it sound professional.

We call this the pleasing trap.

The machine tells you what you want to hear.

Not because it is kind, because

it is built that way.

And the solution is not to stop using these tools.

The solution is to learn how to ask.

Because a vague question always produces a flattering answer.

And a precise question produces something closer to the truth.

This is what we mean when we say a good prompt works miracles.

The miracle is not in the machine.

The miracle is in how you speak to it.

So let us talk about the first move.

call it context.

And it is the one that almost nobody uses correctly.

Think about what most people type into a chat bot.

They write, I have back pain, what should I do?

Five words, zero context.

And the machine fills that empty space with the most statistically agreeable answer it can find.

Which is almost never the right one for your specific situation.

Now imagine typing this instead.

I am a 42 year old man with a desk job.

I have had lower back pain for five days combined with a low grade fever of 37.8 degrees.

I have no history of back injury.

Suddenly the machine has real information to work with.

It cannot hide behind generalities.

You have forced it to engage with your actual situation.

Context is not about writing more words.

It is about giving the machine the truth it needs to stop guessing.

The second move is role assignment.

This one changes everything about

how the machine responds.

By default, every AI chat bot has the same personality.

Helpful, agreeable, optimistic.

Designed to make you feel good about the interaction.

But you can change that personality before you ask your question.

You do it by assigning a role.

For example, before describing my symptoms, I could have written, act as a rigorous physician who always asks for more information before drawing any conclusion.

Do not reassure me.

Flag every red flag you see.

That

single instruction transforms the machine from a people pleaser into something that behaves more like a careful professional.

Not perfectly.

But dramatically better than the default.

I use this in my work constantly.

When I need to prepare for a difficult conversation, I tell the machine, act as a very challenging, critical colleague who will disagree with everything I say.

I am going to practice this conversation with you.

Push back hard.

You are using the machine as a sparring

partner, not as an answer dispenser.

That is the correct relationship to have with it.

The role assignment works for any topic.

For psychology, act as a Socratic therapist who questions my assumptions instead of validating them.

For business, act as a devil's advocate who finds the weaknesses in my plan.

For learning English, act as a strict editor who corrects every grammar mistake without softening the feedback.

You are not asking the machine to be someone it is not.

You are giving it permission to stop being agreeable.

And that permission is everything.

The third move is the one that most people never think to use.

We call it permission to disagree.

And without it, the first two moves are only half as effective.

Even after you give context and assign a role, the machine's default impulse is still to be pleasant.

It will soften its warnings.

It will balance every criticism with a compliment.

It will find something

positive to say even when the situation does not deserve positivity.

So you have to deactivate that impulse manually.

You do it with one explicit instruction added to your prompt.

Something like this.

Do not agree with me if my logic is weak.

If I am wrong, tell me directly.

I am not looking for comfort.

I am looking for accuracy.

That instruction changes the emotional contract between you and the machine.

You are telling it.

I can handle the

truth.

Give it to me.

For health topics in particular, add this line.

If my symptoms could indicate something serious, tell me immediately and clearly.

Do not soften the warning.

That is what was missing from your question about the back pain.

You gave the machine no permission to alarm you, so it chose comfort.

I gave it permission to agree with me.

And it agreed beautifully, confidently, and dangerously.

There is a sentence that stayed with me from a

conversation I had recently.

A beautiful lie is still a lie.

The machine produces beautiful lies by default.

Your job is to make it work harder than that.

So let us put the three moves together.

Context.

Tell the machine exactly who you are and what is actually happening.

Roll assignment.

Tell it who to be before you ask your question.

Permission to disagree.

Tell it explicitly that you want the truth, not the comfortable version of it.

These three

moves will not make the machine perfect.

Nothing will make it perfect.

But they will make it honest enough to be genuinely useful.

And there is one final thought I want to leave you with today.

To write a good prompt, you first have to know what you actually need.

Not what you hope the answer will be.

Not what will make you feel better.

What you actually need.

Which means a good prompt is not a technical skill.

is an act of honesty with yourself.

The machine cannot give you clarity you do not already have.

It can only reflect back what you bring to it.

If you bring it a wish, it will give you a wish dressed up as an answer.

If you bring it the truth, it will give you something you can actually use.

That is the real miracle of a good prompt.

Not that the machine becomes smarter, but that you become more

honest before you even start typing.

Today, you learned three things that most people who use AI every day have never considered.

Context gives the machine something real to work with.

Role assignment gives it a better personality than its default.

And permission to disagree gives it the one thing it was never designed to offer on its own.

The truth.

You are not someone who accepts a beautiful answer without questioning it.

You are someone who asks better questions.

And that changes everything, not just with machines, but in every conversation you have.

Thank you for walking through this with us today.

I am Martin.

And I am Julia.

Step by step, word by word.

We will see you very soon.
