

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

Hi, my name is Tom Johnson and you are watching, listening to I'd Rather Be Writing.com. This is the podcast version of content that I publish. I'm trying to... I'm trying to record more podcasts. So for the last couple years I've been doing a lot of writing, mostly just text, and I've decided lately that I should really try to convert more of this written content into audio content. Now especially because like AI image generators are so... make it so easy to create images, I've also put these... I put together like a very quick slide deck for this. So anyway, I'm going to kind of go through this post. This post titled My Experience, trying to write original full-length human-sounding articles using Claude AI. And I'm not going to read the post because people in the past have told me they hate that. And so I'm going to try to... I'm going to try to just speak it. I'm going to talk through it. Alright, so the gist of the post is that you can use tools like Claude AI or other AI tools to write full articles. You're not limited to just fixing a paragraph here or making a small update there. You can actually write full articles that sound very human-like, but it's, you know, hard to pull off. And I tried to do it and I thought it was going better than it actually was. In this post or in this podcast, I'm going to talk about some of the challenges in writing full-length content as well as my experiments and techniques and my strategies here. So first, the question is why? Why even try to do this? I can write, I don't need a computer non-human entity, an AI to do the writing for me. I'm a professional writer, a professional technical writer. I've been writing on my blog since 2006. So gosh, like, a long time. I don't have any trouble cranking out posts. Why would I want to try to get an AI to write a post? Well, first, I love experimenting with different techniques. It's sort of what makes it fun. But also, you know, there's this idea or sentiment floating around that, like, hey, AI might automate a lot of tech writing jobs or it might kind of reduce the value of writing. And, you know, if AI can write some parts of content, why can't it write more parts? Why can't it write full-length content? So in this post, I am going to walk through the techniques that I think can get partway there. But first, let me talk about some research as I was kind of looking into articles to try to learn about how to write with AI to see what's been written, what's been done and so on. First of all, it's almost impossible to get through wall after wall of marketing. Hey, use our AI writer type tool, whether it's Jasper copy AI or like a dozen others. It seems like that's the only thing I keep finding in Google. But I did come across a couple of articles that were really worth kind of learning about one is one is it's called navigating the jagged technological frontier. And this image here really is trying to depict like a centaur, which was a key metaphor in contrast to a cyborg. Apparently, the researchers gave some people some tasks and so on. And some of the tasks were more AI capable, like there were tasks that an AI could do in somewhere less sort of AI capable tasks. And the consultants, these were the users, the consultants who did the best were the ones who kind of were able to only use AI for certain parts of

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

the tasks. Whereas the other group, which were termed the cyborgs just kind of used AI wholesale for everything. And so the point of the research was to say that whether AI makes you productive or not really depends on whether you're using AI for the right sort of task. And whether you're judicious enough to be able to switch on the AI tools for those tasks where it's appropriate and then switch off that AI mode for tasks where it's not. And it's not always clear like when you should be using AI and when you shouldn't. But definitely this idea of a centaur where you're, you know, centaurs like half man, half horse. And it refers to this ability to delegate the tasks. Are you, I don't mean delegate in terms of like having somebody else do it, but well, actually, maybe I do. But delegate the AI to do something versus not delegating it to the AI. Another article. Actually, this is a podcast that I really like. This is a, let's see, it's on everyday AI. And the podcast episode was called Future Crafting. And the host interviewed somebody named Brian Sykes, who I wasn't really aware of, but I think he's kind of very well known in the creative ad design kind of domain. And like the creative professional kind of space. And they use this metaphor of a director. And their basic argument is that like, as you start to use some of these AI tools, you shift from being the actor, who's sort of performing, you shift from being the writer who's constructing each word to the director who's shaping the scenes, who's telling the story and is directing the actors about what they should say and how they should perform. And the larger idea is that like you're going to clear up a lot of mental bandwidth so that you can perform this director type role, rather than getting down in the weeds and doing the grunt work of wordsmithing and really pushing out each of the sentences. They weren't really talking so much about writing as like design. But I think a lot of it applies. You know, if you think about all the energy and effort that goes into shaping sentences that are smooth, that read well, that are direct, clear, easy to understand and so on. It's a lot of work. This idea of being a director kind of really appeals to me because with each post that I write on my blog, I spend a lot of time just re reading it and proofreading it and reading it again. Each time I try to let a day in between the edits to give it some space. And it's just really tedious, you know, like trying to get it to read well and to spot all those parts that jump out at me with more space requires a lot of energy. So if I can somehow kind of offload that wordsmithing to AI and instead focus more time and attention on the logic and the flow and this larger narrative, that would be kind of a big win. Alright, so let me move on to what my three biggest strategies are for using AI to write. Now, like I said, it's not like I've cracked the code on this by any means. But I'm just sort of sharing what I've tried and why. My first strategy is to prime the AI with accurate information. Second one to go paragraph by paragraph. And then third to balance personal voice with explanation. And let me jump into each of these with more details. Okay, so this first strategy is to prime the AI with accurate information. This is why I've specifically called out Claude AI because it

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

allows you to load it with a ton of input. You can put like 75,000 words into it. And that's like between the input and the conversation limit and so on. So the whole session. But if you're if you're writing about something that requires explanation, and you want to make sure that like anything you're explaining is going to be more accurate and not just made up hallucinated and so on. By having a lot of articles that the AI can use as source material, really make a big difference. So let's say you're writing an article about some topic, and you want to be able to sort of access easy explanations about different facets of the topic, you could go and find, I don't know, a dozen articles that really contain a lot of the information that you would then later be summarizing or paraphrasing or referring to. Just sort of load it into the AI as a kind of training right in the chat. And then later when you need to make those summaries, it's going to be a lot more accurate. So I'm a real big fan of the of the Claude AI precisely because it has such a large input source. Alright, my second strategy here is to go paragraph by paragraph. A lot of people start out and they think, Oh, I'm going to have, I'm going to have this AI tool write a whole like post a whole article in one go. And when you try that, you realize that it's terrible. It's like, it's usually way off the mark. But just to take a lesson from software methodology. Remember how with Agile, people basically do a couple of weeks of software development, and then they show the feature to a user to sort of check in. And then the user feedback might have suggestions or corrections and so on. And so you'll you'll course correct and go back another couple weeks, make some more features, show it to the user, you know, and get feedback and keep doing that. That's the whole idea of modern software development with scrum and other sort of agile workflows. The previous method method was waterfall, where you would just like spec it out, the whole project out from the beginning, and then build it for two years or so. And by the time you've finished building it, you show it to the customer, and they're like, wait a minute, that's not what I wanted. But you've already gone so far in the wrong direction, that it's sort of unrecoverable. Well, going paragraph by paragraph is precisely the strategy that aligns with Agile, and it results in much better output. And I'll jump back into paragraph by paragraph more a bit later. Alright, my third big strategy is to balance personal voice with explanation. I actually did an MFA in literary nonfiction at Columbia University way back when like back in the early 2000s. And my biggest takeaway is probably this this balance creative nonfiction, like a personal essay or another any kind of, you know, more literary essay really gets a lot of its appeal from this balance of personal experience, anecdotes, insights, coupled with explanation, and this the sort of interweaving of the two is what makes it come alive. It also has kind of another benefit. If you're doing a lot of the explanation through an AI tool, having the AI write that explanation, and then you switch into the personal voice and the anecdote, that personal voice is going to better disguise is probably the wrong word, but will disguise the AI written parts more. And when you read something

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

that has like a an eye of first person, it sort of tricks the brain into just assuming and believing that that all the parts of the essay are coming from the author, better communicate, you better connect with with the reader this way. Okay, so those are my my three big strategies I was trying. And now let's let me break down kind of this 10 step process so that I can more thoroughly kind of go through my experiment here. All right, the first step is to define an information pattern. This is kind of one of my favorite parts because a lot of people don't even realize that there's like information patterns to things. Certainly, if you're a tech writer, you're you're accustomed to information patterns in documentation, a task has a certain sort of order. You've got an intro with prerequisites and the problem that you're solving and you've got steps and examples and so on. And with a with a blog post, like sort of a long form blog post, or not even long form, but just something that's more of a personal essay ish style. There's a there's a specific pattern, you sort of have this hook that explains the relevance. You define some kind of issue that that's presented, you ask a question of some sort, what are you trying to answer? What is what is this thing you're wrestling with? You summarize prior research, and maybe you critique some limitations and so on, you do some sort of experiment, you're trying things out. And over the course of trying things out, you tend to have some sort of epiphany that changes your perspective. And then that leads you to the conclusion. I mean, it's essentially kind of like the hero's journey, only instead of like a hero, you've got your on a quest to figure out an idea or get an answer to something. Michelle de Montaigne's original sense of the word essay was like trying out an idea or making an attempt with some thesis or idea. So anyway, that's the information pattern that I think works well. And you can even see it in in this sort of podcast, this post. Alright, so let's move on to the second one, create an outline. Again, like because we're going to go paragraph by paragraph, you have to have a sense of like what the article is going to say. I tend to have a rough outline so that I can allow for more flexibility in the creative process. I don't want to be really rigid and locked down, but I definitely have a bunch of points and I put those points into that information pattern that I was describing. Alright, step three, let Claude know the context. Right. So when you when you first jump into this AI session, I say something like this, you're going to write you're going to help me write an article for my blog, I'm going to lead you paragraph by paragraph describing what I want you to write, I will be the director and you will be the writer, you will articulate my ideas in readable grammatical ways, adopting a plain speaking direct style. If at any time my ideas are ill conceived, you will push back and recommend better approaches. Are you ready to begin? So that kind of lets Claude know, hey, we're going paragraph by paragraph, you're going to do the writing, I'm going to whisper ideas into your ear, you're going to kind of be my muse and so on. One little tip, don't say the word essay, because this sort of triggers student essays. And a lot of times, if you

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

do that, Claude might say something like, hey, it can't really help you with that, but it seems fine ghost writing blog articles. So anyway, you might run into that. Okay, step four, calibrate Claude's language. You'll start out, I'll start out by just describing my first paragraph. Remember this, you know, relevance hook sort of part. But really, in this first paragraph or this first response from Claude, you want to make sure that Claude is using the right sort of diction and style. If it's too like, highfalutin, or whatever, if it's too grandiose or like pseudo literary, you can try to correct, correct Claude and say, hey, simplify your language more. And if it than if it like oversimplifies, you can say, hey, try to strike a balance between the two and so on. You can go back and forth a little bit until until it's sort of hit a style that you find kind of normal. Usually by default, Claude is pretty like conversational and readable. But anytime I try to say, hey, adopt a style of a New Yorker talk of the town type writer or something, it always goes poorly. It sounds terrible, like somebody's got a Saras in their, in their pocket and every sentence they're trying to use it and so on. So anyway, after calibrating, move on to the next step here is to continue paragraph by paragraph. Just I go through my outline, and I describe what I want it to say in each paragraph, and then it writes it. Now, this is a part you may say, well, geez, if you have to explain each paragraph, like that's half the work or writing. Well, yeah, you have a good point there. And for sure, this isn't really saving a ton of time. So this might be where the experiment kind of fails. You know, maybe it saves only like 20% of time and isn't really worth it because then you have to go back and maybe rewrite a lot of these paragraphs. But anyway, this was my approach and my attempt. And I go paragraph by paragraph. And it is kind of interesting because you can read how like you can read almost in real time, the articulation of these ideas and more full fledged paragraphs. And you can see what you're, what you're writing and the ability to sort of see what you're saying gives you more distance to better evaluate and assess it. So there's value in that. I don't quite, I'm not quite sure how much value because there's also a lot of value in writing something yourself and sort of going more slowly because that lets ideas form and solidify more. So I'm sort of mixed here. But anyway, that was, this has been my, my attempt. All right, the next step is to compile all the paragraphs into like into an article. Right. Some of the paragraphs might have been written several times by Claude, maybe they were in the wrong order and so on. Step seven is to edit the whole. And this is kind of the fun part, right? Because you've reached a milestone and you're like, Oh, I've now got it first drafts, first drafts are usually things that require a lot of edits and so on. But edit the whole, see if it's got the right shape and feel and so on. And then you move on to fine tuning the article. Now, at this point, you know, since Claude or whoever, the AI has written a lot of this stuff. There might be a lot of like cliché phrasings or phrasings that are unnatural or sound weird. So it requires quite a bit of fine tuning. All right, step nine is to maybe incorporate some auto generated art. One really interesting technique is to just like copy and paste a whole section

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

into Claude and say, Hey, generate or create a prompt for an AI image generator based on the ideas here, and then just paste it into chat GPT plus with Dolly as the mode. It's available in the plus version and it generates art. That's how I generated like all of the slide art here. I did this very quickly and just sort of picked the most, I don't know, one of the four images that it generated that I liked. So you can move really fast if you want. All right, step 10, let the content sit a while, even though it was probably pretty easy to get a first draft. And maybe you think, Oh man, I could create a whole essay every day like this. You don't want to do that. You want to let things sit because in the back of your mind, maybe you're allowing things to surface and you're realizing that, Hey, I really didn't want that part or this is bugging me. That sort of like back of the mind percolation takes a couple days, at least for me. All right, now how did these posts turn out? I've only written like a couple of posts using this method. So like I said, this is this is an experiment. And one of these posts was a the attacks as post and the other was embracing professional redefinition. And I got some good feedback on both of them. For the most part, people like the posts, people didn't come out and say, Hey, this sound, you know, this is obviously AI written, Tom, why are you doing? No. So it wasn't bad. I didn't feel like they were great posts, but I felt like they were decent. But one reader actually reached out to me and he's like, he said, you know what, your latest posts feel a bit off. And he was trying to put his finger on exactly why. It's really interesting. He let me let me just quote a little bit. He allowed me to quote him. He said, forgive me if I'm out of line here. I've always looked to you, you know, as a thought leader and pragmatist and so on. But lately, something just feels off. The writing doesn't seem to be as refined or tight. Chunks are longer and less less Tom packed, using a kind of interesting adjective there. They're less Tom packed, as I used to say, to refer to your uncanny ability to really stuff a paragraph with value with an economy of well picked words. And he specifically caught out the Deattaxus article. He said, the summary appeared especially generated didn't have key key conclusions that I found peppered throughout the body. Still good stuff, but just not like he's used to reading from me. He says AI tends to do this, it makes a nice bulleted lists that don't really reflect the same scrutiny that a writer would apply before figuring out which of the points are really relevant to an argument and so on. So this got me thinking, because I was really like, I was really feeling pretty proud of this process, thinking, man, I've totally nailed it. And then I see this reader's feedback and I'm like, you know, I think he might be onto something. But it was kind of hard to sort out exactly what he was saying. Why was the writing off? And this is this taps into really a larger question a lot of people have had. They're like, you know, the the AI tools, they write well, but something's weird about it. It's like something is off and different people say different things. I've read one article where somebody said, it's the imperfections that we like, you know, the AI writing is too perfect. Well, I did some more research. And I've got several ideas here. My first reaction was, oh, maybe it's the uncanny valley. This is the idea that as you get closer and closer to something that looks almost human, you cross over this sort of threshold where it's clearly not a robot and you're like getting very close to human. And you hit a point where people actually start to feel a lot more unease and revulsion about it. This image right here, I think is a good a good depiction of the uncanny valley, because it's like kind of disturbing, because it looks mostly human, but it's clearly not quite. Anyway, so I thought maybe maybe the writing seemed to be humanish, but

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

wasn't quite. And I thought, well, what if, you know, especially if somebody's familiar with how AI writing looks, if you start to if you start to sort of smell that it's AI written, maybe there's a inner revulsion at feeling that you're sort of being tricked, like, hey, I thought I was reading a human. Instead, this is a machine. And maybe that leads to a sense of sort of betrayal and distrust that then provokes a stronger reaction. I'm not really sure. But when I came back to the reader's comments, he wasn't really saying that it sounded almost like Tom. He was saying something

different. He was saying that it wasn't really like there wasn't tight logic. It didn't, it was rambling, or it didn't have the writer's discerning compactness of ideas and word choice.

So I found another article. The second article here is

by Laura Hartenberger called What AI Teaches Us About Good Writing. And it's a little more vague.

She touches on a number of different things. Let me just read a passage. She says, as readers, we need to feel like the writer is paying paying attention to us, trying to connect. Chat GPT cannot build a real connection with its reader. It can only imitate one. Reading chat GPT's writing feels uncanny because there's no driver at the wheel, no real connection being built. While the machine can articulate, articulate stakes, it is indifferent to them. I think I mistyped that.

Anyway, machine can articulate late ideas and experiences. It's indifferent to them. It doesn't care if we care. And somehow that diminishes its power. Its writing tends not to move us emotionally at best. It evokes a sense of muted awe akin to watching a trained dog shake a hand. Hey, look what it can do. So my takeaway from her article is that there's sort of a connection, an emotional connection through lived experience and sharing of experiences and in this sort of meaning and paying attention to the reader that comes about from writing that is sort of lost when you don't have a human to human connection there. And she also says perhaps the time spent writing matters as much as having written. And I really like that because for sure when you wrestled with an idea, and I've written many posts like on my blog for the last, you know, 20 years, those posts where I'm really trying to sort out in my head what I think about something and wrestle with that. They're a lot more meaningful. I mean, if you just have an AI write something, it's not as meaningful.

And in a lot of ways, you're short changing yourself from the writing writing experience. If the value of writing is having written, then if you haven't written, then you're losing the value of writing. And if you're doing that, if you've lost the value of writing, but you're still producing words, it's going to be a hollow experience. It's not worth repeating. So anyway, very good perspective and something to keep in mind. Let me look at one more article. Actually,

this is not even an article. This is my own little hypothesis here. I think that these tools are very over agreeable and probably inevitably so. But I found that they sort of carry out your will. In my self redefinition post, I initially started out with a comparison to one of Nietzsche's Overman or Ubermench ideas and so on, because it was about like, hey, creating new values and meaning and so on. But there's kind of some negative baggage with that idea. And I was probably better off just not even steering into philosophy there. But Claude totally like went that way because I was telling it to and it didn't push back. And I think a lot of these tools will not really push back. They're sort of programmed to be very polite and collaborative and constructive.

And so on. I got into a fun little exchange with Claude because I was like, hey, you know, a true friend would really push back. And it said, well, I definitely will if there's like

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

factually false information such as saying two plus two isn't is five or something. And I responded, well, in certain systems where you have the sum being more than the individual parts and so on with dynamic feedback loops, it's very possible that two plus two could equal five. And it was like, oh, yes, of course, you're right. In that situation, it very well could be I'm so sorry. You know, I should have been more more careful in my logic. And I'm like, this is you're demonstrating the exact point you agree with everything I want to say. And I think this over agreeableness could lead to a lack of having a crisp argument. If every time you rely on AI, it sends back kind of this watered down explanation that doesn't take a side in any kind of issue. Then that might reflect in the writing as well, where I don't really have a crisp point. I'm not taking a side. And the writing might not be compact, like the reader was saying. This is just a hypothesis. Obviously, you can use the tools in many different ways, but definitely be on the lookout for over agreeableness and realize that the tools aren't going to push back and say, are you really sure you want to reference Nietzsche? I mean, you're writing to a technical communication audience, not a philosophy audience. And by the way, you don't really know that much about philosophy either. It's never going to be that direct. So let's just wrap this up. I think based on my experiment, it would probably be best to try to use AI mostly for summary and explanation. And then for the personal experience, more of the first person parts to not rely on it. It's my sort of guess or hypothesis that by just using it more for summary and explanation, you'll be more apt to have a more direct and clear argument if argument is sort of your game here in the post. You're more apt to follow that argument and have more of a point and an opinion than if you are relying on AI for that part as well. I do think though that using AI for summary and explanation really makes things a lot easier. These tools are great at that and it can reduce a lot of the cognitive overload for those parts. And that balance, again, coming back to the strategy I was describing earlier, this balance between personal experience and summary and explanation, they go back and forth, they interweave. I think it really makes for compelling engaging content, at least for a blog or other online stuff. That really is a nice winning formula. So like I said, let me just caveat this. I've only tried this for a couple of posts. I have some news summary posts where I'll summarize some news and so on, but it's not the same as trying to write like a personal essay. And I may decide this whole technique is not worth doing. Maybe it doesn't really save much time and it's not worth it. Maybe I just have to rewrite everything that Claude wrote myself because I want it to sound like me and have my ideas and phrasings. And I could just cut out that whole rewriting process and be much more efficient. But I also kind of think this might take a little bit of getting used to. It may be something that I get better at and maybe I can fine-tune a process. Ultimately, I do think that it could be useful in the workplace for writing documentation as well, which is a totally different genre. It doesn't have personal experience or argument or voice, three huge components, but it does have a high degree of accuracy. I'm talking about the genre, high degree of accuracy and precision that also can be a challenge. So I'll hit that in another post. If you have any questions or feedback, shoot me an email at tomjohtat@gmail.com or the contact form in my blog. I'd rather be writing .com. You can subscribe to my podcast pretty much anywhere. Pocketcasts, iTunes, Spotify, you'll find it pretty much everywhere. And let me know if

[Transcript] I'd Rather Be Writing Podcast / My experience trying to write original, full-length human-sounding articles using Claude AI

you like this format with the slides. This is now on YouTube and so on. I needed a visual element. It also sort of helps to talk through the points rather than relying on written text. But it would be great to get feedback. And if you feel like reviewing the podcast, great. If you don't, don't worry about it. But thanks for listening and see you online.