

## Stack Overflow questions are being flooded with answers from ChatGPT

171 points by brindidrip 5 hours ago | hide | past | favorite | 242 comments

What are the repercussions of this?

add comment

brindidrip 3 hours ago | next [-]

It seems like there are a few potential negative consequences of using AI-generated answers on Stack Overflow. For one, the quality of the answers may be lower than if they were written by a human. Additionally, if these AI-generated answers become too common, it could potentially lead to a more impersonal and less supportive community on Stack Overflow. Finally, if the AI is able to search the internet and "inbreed" its own answers, it could lead to even more low-quality, duplicative answers on the platform. Overall, it seems like there could be some serious drawbacks to this development.

Note: This answer was generated by ChatGPT after being fed this thread.

[reply](#)

pimlottc 1 hour ago | parent | next [-]

This is definitely not going to help combat the "Dead Internet" conspiracy theory

<https://theswaddle.com/what-the-dead-internet-theory-predict...>

[reply](#)

alfiedotwtf 43 minutes ago | root | parent | next [-]

Reminds me of Heavenbanning:

<https://twitter.com/nearcyan/status/1532076277947330561>

[reply](#)

thetip 3 hours ago | parent | prev | next [-]

Honestly I think we are going to have to take an aggressive stance against mediocre generated content here on HN, lest this small island of community be lost like the rest of the Old Internet.

[reply](#)

stoppingin 1 minute ago | root | parent | next [-]

What if AI could be leveraged to find interesting content? I doubt I'd even notice the bulk of comments being replaced by an AI. I only read them for the one in a thousand comment from a domain expert, or industry insider that really informs me, or challenges my perspective.

[reply](#)

wizeman 2 hours ago | root | parent | prev | next [-]

> I think we are going to have to take an aggressive stance against mediocre generated content here on HN, lest this small island of community be lost like the rest of the Old Internet.

How do you propose we do that?

And what do you propose we do when ChatGPT (or something like it) can create better content than most humans?

(... which I would argue it is already happening in some limited contexts, although I admit my stance is controversial).

[reply](#)

vineyardmike 2 hours ago | root | parent | next [-]

> what do you propose we do when ChatGPT (or something like it) can create better content than most humans?

An AI cannot, by definition, create better HN content. Because HN content is about hearing what people think about a certain topic/thread. Hearing what an AI statistically regurgitates doesn't satisfy the purpose even if it's better written.

The internet is about *people*. I may not be friends with everyone on HN, but we're a community meant to discuss and share together, and I personally wouldn't want to talk to an AI over a human.

[reply](#)

cesarb 18 minutes ago | root | parent | next [-]

> An AI cannot, by definition, create better HN content. Because HN content is about hearing what people think about a certain topic/thread.

Not necessarily. Good HN content can also be factual information relevant to the topic at hand. And yes, current AI like ChatGPT might not help with that, but an hypothetical future AI which cared more about the veracity of its statements could.

[reply](#)

wizeman 1 hour ago | root | parent | prev | next [-]

Wow, that's an interesting argument if I ever saw one! Your viewpoint is completely the opposite of mine, let me argue a couple of points:

> An AI cannot, by definition, create better HN content. Because HN content is about hearing what people think about a certain topic/thread.

Please tell me: for you, what is the difference between hearing what people think, and hearing what a human-level AI thinks?

Are you saying a human-level AI, by definition, cannot ever be interesting to you, just because it's an AI, even though it could be like a human in every other way?

> Hearing what an AI statistically regurgitates doesn't satisfy the purpose even if it's better written.

Don't most humans just regurgitate the same arguments they've just read or learned elsewhere (or even here), like 99% of the time?

And even in the remaining 1% of cases, aren't they necessarily just functional products of the things they've seen/read/experienced (i.e. the inputs they've received, even if gathered by feedback when interacting with the world)?

> The internet is about people. I may not be friends with everyone on HN, but we're a community meant to discuss and share together, and I personally wouldn't want to talk to an AI over a human.

What if the AI was more interesting and intellectually stimulating than a human?

[reply](#)

true\_religion 1 hour ago | root | parent | next [-]

> Please tell me: for you, what is the difference between hearing what people think, and hearing what a human-level AI thinks?

HN is a professional social network. People are most interested in what their peers have to say, not any random human being or human level AI.

Now if this human level AI is working in the field as a professional, or at least in computer science academia the it's opinions are valid and interesting. However if it's merely an average best possible response then it's pointless to hear it's opinion.

Additionally if multiple people post chatGPT responses that's just like talking to same person using humans as sock puppets.

[reply](#)

wizeman 57 minutes ago | root | parent | next [-]

> People are most interested in what their peers have to say, not any random human being or human level AI.

> Now if this human level AI is working in the field as a professional, or at least in computer science academia the it's opinions are valid and interesting. However if it's merely an average best possible response then it's pointless to hear it's opinion.

Yes, I agree with that, but I don't think that's what the parent poster was arguing.

It's also clear that I didn't phrase my question as clearly as I could, because instead of "human-level AI", I should have said this instead: "a human-level (or more intelligent) AI that has equivalent or better knowledge/experience than the people who post on HN".

> Additionally if multiple people post chatGPT responses that's just like talking to same person using humans as sock puppets.

Yes, I agree with this as well.

But as a counterpoint, (as far as I understand) it's possible to have ChatGPTs with different levels of knowledge/experience and different personalities, as evidenced by GPT-3's fine-tuning capability (disclaimer: I've never used this functionality, so I'm not 100% sure if this is correct).

[reply](#)

tartoran 23 minutes ago | root | parent | prev | next [-]

> Don't most humans just regurgitate the same arguments they've just read or learned elsewhere (or even here), like 99% of the time?

Yes but at least they choose what they regurgitate unless you think of most people as automatons. Personally I'm more interested in human regurgitations than AI imitations of these. So far AI doesn't bring any reasoning and cannot discern what it regurgitates but sure as hell comes off as confident (It could probably imitate being humble as well). Someone posted an example of gpt bulshitting something akin to  $2+0=3$  but very convincingly.

[reply](#)

wizeman 14 minutes ago | root | parent | next [-]

> Yes but at least they choose what they regurgitate unless you think of most people as automatons.

And what, are you saying ChatGPT doesn't choose what it regurgitates?

It seems like these arguments are getting more and more flimsy.

I do believe people (including me) are automatons because I think free will is logically impossible in the way most people intuitively think free will is.

> So far AI doesn't bring any reasoning

This is clearly untrue, as ChatGPT can definitely reason pretty well (although, not always correctly, like humans). As far as I can see, it can reason deductively, inductively, by analogy, it does abductive reasoning, cause-and-effect reasoning, critical thinking, step-by-step reasoning, you name it.

> Someone posted an example of gpt bulshitting something akin to  $2+0=3$  but very convincingly.

Humans do this all the time (although usually not at such an extreme level). Just look at all the posts saying ChatGPT can't do X or Y ;)

[reply](#)

kthejoker2 1 hour ago | root | parent | prev | next [-]

This reads like a ChatGPT response. (So does my response!)

Whether it is or isn't is beside the point.

Broad, anonymous commenting platforms are dead, they just don't know it.

[reply](#)

wizeman 1 hour ago | root | parent | next [-]

> This reads like a ChatGPT response. (So does my response!)

Your comment is very interesting, because I'm having the same experience: the more I interact with ChatGPT and read its arguments/responses, the more I'm getting weird vibes when reading arguments written by humans, although I cannot tell you exactly why (and I think I can still clearly differentiate between a human and an answer from ChatGPT that is copied verbatim, as it tends to speak in a more formal way and usually, it's also more verbose than how humans typically write here in HN).

I think it's also influencing the way I write, both to be more clear (because otherwise ChatGPT can misinterpret me), but also because I'm reading so much ChatGPT-generated content, which I believe also indirectly influences the way I write and think.

In any case, I know you said it's besides the point, but I assure you, I'm not ChatGPT or copying its answers here (unless I add a clear disclaimer) :)

[reply](#)

unethical\_ban 1 hour ago | root | parent | prev | next [-]

The AI we see today is a synthesis of a huge number of arbitrary inputs. You may argue that humans are, as well, but humans are different machine than AI. If I am asking for a human response, or a human piece of artwork, I am doing so because I want to understand and experience the output of machines of the same model. I want to continue the human experience.

AI-generated content can be fascinating, helpful, and in some instances, more useful and accurate than humans (medical diagnosis, technical documentation, perhaps). But if I ask for a human, I *want* a human.

I don't care if AI is more interesting than a human. I want a human, because I am human. I am not transhumanist.

I wonder what the correlation is between people who see no particular value in interacting with humans, and people who struggle to interact with humans.

[reply](#)

tartoran 16 minutes ago | root | parent | next [-]

Get ready to see a lot of human spaces invaded by somebody else's AI with goals that are not in your interest. It's been happening already with chatbots pulled by strings by nefarious people but that operation will only amplify to a point that it will become cumbersome and hard to know whether you're chatting or reading something that came from a human or language model.

[reply](#)

wizeman 24 minutes ago | root | parent | prev | next [-]

> But if I ask for a human, I want a human.

> I don't care if AI is more interesting than a human. I want a human, because I am human. I am not transhumanist.

I think I understand your point but I'd like to give a counterpoint: replace "human" by "white human" and "AI" by "black human" and you might see how that line of reasoning is flawed.

In other words, there might come a time when AIs could become really offended if you'd exclude them like that from social interactions, with all the repercussions that might have.

> I wonder what the correlation is between people who see no particular value in interacting with humans, and people who struggle to interact with humans.

I see value in interacting with humans, especially at this point in time, and especially in ways that machines can't (e.g. having meaningful intimate relationships, raising a family, etc). Even then, machines could theoretically do some of this better than humans, as suggested by a lot of sci-fi content (except the actual reproducing part).

But I also see value in interacting with beings that are superior to humans, assuming they are able to do what humans can, only better.

[reply](#)

unethical\_ban 2 minutes ago | root | parent | next [-]

You think AI is a being. I am not yet convinced.

I am a human supremacist, yes.

Further, it is not unreasonable to have more interest in some cultures than others, or find the experiences of ones own culture more engaging or relevant to oneself than another. The "line of immorality" comes with banning or violently oppressing other experiences.

Again, fundamentally, I disagree with an analogy giving AI equal morality or agency to a homo sapiens. There is no room for "find replace" here.

fragmede 30 minutes ago | root | parent | prev | next [-]

> An AI cannot, by definition, create better HN content.

*Better* is entirely too subjective for this to be true. And if you turn on showdead, ChatGPT content is already better than some of the other content submitted here. If an ML algo can take a half baked comment of mine and form it into a coherent response that other people can read and actually understand what I mean, that *is* better content.

[reply](#)

johnchristopher 1 hour ago | root | parent | prev | next [-]

> An AI cannot, by definition, create better HN content. Because HN content is about hearing what people think about a certain topic/thread.

Hmm. That's HN proposition value but that's not necessarily what keeps people coming back. I am thinking about the dopamine hits from the news cycle and comments.

If I am right then it follows that:

> Hearing what an AI statistically regurgitates doesn't satisfy the purpose even if it's better written.

..actually does the job.

How much chatGPT content can we/you/I stomach before we stop visiting ? We all have a different threshold (and language is a barrier for non-native, it would certainly take me longer to spot the AI).

[reply](#)

ben\_w 20 minutes ago | root | parent | prev | next [-]

Out of curiosity I tried prompting ChatGPT to come up with interesting rebuttals to your comment, but it kept misunderstanding my prompt and agreeing with you instead.

[reply](#)

tartoran 33 minutes ago | root | parent | prev | next [-]

> An AI cannot, by definition, create better HN content. Because HN content is about hearing what people think about a certain topic/thread.

With AI certain ideas and opinions can and will be amplified by malicious actors. We may have to resort to face to face at some point or verify human identity at times to combat this.

[reply](#)

bigbillheck 39 minutes ago | root | parent | prev | next [-]

There's already plenty of comments on this forum that are low-effort regurgitations of talking points, memes, 'common sense', and the like.

[reply](#)

yeeeloit 59 minutes ago | root | parent | prev | next [-]

> And what do you propose we do when ChatGPT (or something like it) can create better content than most humans?

I tend to agree... I've been growing more and more tired with the content in familiar places. eg. reddit. (As an aside I think a lot of it is driven by advertising/marketing, but not all of it...)

Anyway, your comment reminded me of the recent footage of baggage handlers at an airport, and how that dovetails nicely with the recent move by tesla to build a humanoid robot.

Looking for a ray of light in the approaching storm: Maybe these AI can be used to filter the content more effectively for us.

[reply](#)

9wzYQbTYsAlc 2 hours ago | root | parent | prev | next [-]

> How do you propose we do that?

One technique, like all the other self-moderation that you can do on HN: stop upvoting and commenting on content that you don't want to see boosted.

[reply](#)

lamontcg 1 hour ago | root | parent | next [-]

People already upvote fairly bland comments which reinforce common beliefs, so AIs which can distill popular sentiment and regurgitate it very readably will get wildly upvoted on any social media platform.

Most humans don't like too much creativity and they want ideas that they agree with.

Wonder if the death of social media is more a descent into AI generated inane and bland commentary.

[reply](#)

wizeman 2 hours ago | root | parent | prev | next [-]

> One technique, like all the other self-moderation that you can do on HN: stop upvoting and commenting on content that you don't want to see boosted.

But isn't that what most people are presumably doing already?

I guess my question was more intended to be: how do you differentiate between content generated by humans vs machines?

At some point, we might not be able to. Or even if we can, it could actually result in a *worse* experience, if machines can generate better content.

[reply](#)

9wzYQbTYsAlc 1 hour ago | root | parent | next [-]

A few years ago there was a US Congressional Hearing [1] about how to handle deepfakes (OpenAI had a representative on the panel).

There was some deep discussion about that topic.

For example: even if all authentic content had an embedded steganographic watermark, how do you reliably authenticate recordings of recordings or otherwise degraded copies of authentic content?

[1] <https://intelligence.house.gov/news/documentsingle.aspx?Docu...>

[reply](#)

tetris11 2 hours ago | root | parent | prev | next [-]

Voting with your wallet only works if a large portion of consumers both agree and understand your point of view

[reply](#)

hgsgm 2 hours ago | root | parent | prev | next [-]

"Ignore it" has never been a solution to "other people misbehaving and destroying shared spaces".

[reply](#)

wellthisisgreat 2 hours ago | root | parent | prev | next [-]

hard bans for undisclosed machine-generated submissions?

[reply](#)

pctthrowaway 1 hour ago | root | parent | next [-]

Yeah I'm in favour of this also, but proving it might be tricky, and people will get unjustly banned.

Then there's the question of situations where people might actually use ChatGPT in a creative way to augment discussion.

Let's say we're talking about the pitfalls of repetitive code over breaking things cleanly into small functions. You have an example of this in mind that highlights a specific pitfall that you encountered, but don't want to share proprietary code, so you might ask ChatGPT (or a future model) to generate some code that demonstrates the same thing, rather than writing it yourself.

I think we're still early enough in the tech that it's hard to create hard-and-fast rules about what kind of content should be allowed; ideally, we'll get to the point where AI can help facilitate and augment human interactions, rather than taking over completely.

[reply](#)

formerly\_proven 2 hours ago | root | parent | prev | next [-]

While it is important to maintain the quality and integrity of content on HN, implementing hard bans for undisclosed machine-generated submissions could have some negative consequences. For one, it could discourage users from experimenting with and sharing the latest advancements in AI and machine learning, which could stifle innovation and progress in these fields. Additionally, it could lead to a lack of transparency and accountability, as users may be hesitant to disclose their use of AI in fear of being banned. This could also result in a lack of diversity in perspectives and ideas on the platform. Instead of hard bans, it may be more effective to implement stricter moderation policies and provide resources for users to improve the quality of their contributions, whether they are generated by AI or not.

[reply](#)

Always42 2 hours ago | root | parent | next [-]

Noice chatGPT

[reply](#)

formerly\_proven 2 hours ago | root | parent | prev | next [-]

One potential way to take an aggressive stance against mediocre generated content on HN is to implement stricter moderation policies and actively monitor and remove low-quality content. Another option could be to encourage more participation from the community, such as by offering rewards or recognition for users who consistently provide valuable insights and information on the platform. Providing resources for users to improve their writing skills, such as guides on how to effectively research and organize information, could also help improve the quality of content on the platform. Additionally, developing a system for flagging and reporting low-quality or duplicate content, and empowering moderators to take swift action against such content, could be effective in maintaining the integrity and quality of content on HN. Engaging with the community to gather feedback and suggestions on how to improve the quality of content on the platform is also crucial in this effort.

[reply](#)

YuriNiyazov 2 hours ago | root | parent | next [-]

I see you, GPT-3.

[reply](#)

andrepd 9 minutes ago | root | parent | next [-]

GPT3 fooled me a couple times at the beginning, but I think as I'm more exposed to its writing I'm getting better at spotting it.

I wonder if somebody has already made a game of this!

[reply](#)

ethanbond 3 hours ago | root | parent | prev | next [-]

Gotta love it. Akin to Silicon Valley-ites ensuring their *own* children don't have access to the technologies they go into the office to create and propagate every day.

[reply](#)

tsol 3 hours ago | root | parent | next [-]

Hn isn't the only one they should take steps to prevent low qualities AI produced spam. Everyone should. But since we're at hn, it makes sense they are talking about hn

[reply](#)

closeparen 3 hours ago | root | parent | prev | next [-]

I would expect most automotive engineers don't let their children drive the cars they go to the office to create and propagate every day. Because... they're children.

[reply](#)

sydd 2 hours ago | root | parent | next [-]

I think the poster above meant that e.g. engineers are Meta are paying very good money so their children go to private schools where they confiscate phones. Or that they don't buy phones at all for their kids until they turn like 12, instead pay a lot of money so they go to private tutors, have nannies, do all kinds of after school programmes. Whereas kids from poor families cannot afford these or to keep watch of their kids 24/7 (because they need to work), so these kids will hang on Facebook/YouTube/TikTok/.. all day

[reply](#)

umeshunni 23 minutes ago | root | parent | next [-]

> e.g. engineers are Meta are paying very good money so their children go to private schools where they confiscate phones

This sounds like one of those made-up tweets that became an article that became an urban legend.

[reply](#)

wizeman 3 hours ago | root | parent | prev | next [-]

> I would expect most automotive engineers don't let their children drive the cars they go to the office to create and propagate every day. Because... they're children.

On the other hand, most parents don't let their children drive cars either and yet they let their children access those technologies.

Another way in which that analogy doesn't work is that it's illegal for children to drive cars, but it's not illegal for children to access those technologies.

Furthermore, these technologies are being created and modified in a way that are specifically attractive and addictive for children (as well as adults) and they are constantly being further optimized for that criteria as well, almost every day (in a manner of speaking).

[reply](#)

baq 1 hour ago | root | parent | prev | next [-]

The problem is, the darn thing produces comments which can't be blanket labeled as mediocre. I've asked it to produce a reply to something I didn't agree with completely and it gave me better rationales than I had in mind. Scary!

[reply](#)

lurquer 1 hour ago | root | parent | prev | next [-]

Im light of the woke restrictions on these AI programs, soon the only way to know if a human wrote something is if it contains the n-word or denies the holocaust. Ditto for stablediffusion: proof of the humanity of the artist will consist in a small swastika or some naked juvenile in the background.

[reply](#)

biggerChris 2 hours ago | root | parent | prev | next [-]

I have -9 points. HN is doing well to prevent people from commenting without facts.

[reply](#)

seydor 2 hours ago | root | parent | prev | next [-]

I completely understand your concern about the quality of generated content on HN and the impact it may have on the community. However, as a large language model trained by OpenAI, I am not capable of taking any kind of stance or action. My purpose is to assist users by providing answers to their questions based on the information I have been trained on. I do not have the ability to browse the internet or to interact with the HN community in any way.

[reply](#)

tartoran 10 minutes ago | root | parent | next [-]

@Dang, your job is danger

[reply](#)

zuzatm 3 hours ago | parent | prev | next [-]

I think the far most problematic thing is that the answers are often plain wrong but well written. Eg (shameless plug)  
<https://twitter.com/zuza777/status/1599370082919936000?t=oB9...>

[reply](#)

slavik81 2 hours ago | root | parent | next [-]

That was a surprisingly common occurrence on SO even before. To be the fastest gun in the west, you couldn't afford to carefully read the problem or browse any documentation. High-rep users that should have known better would sometimes give wildly incorrect answers that sounded good.

[reply](#)

imhoguy 1 hour ago | parent | prev | next [-]

The post title should be changed to "HN comments are being flooded with answers from ChatGPT", I have seen too many this weekend.

[reply](#)

metadat 53 minutes ago | parent | prev | next [-]

What prompt input did you use? When I tried this:

> What are the repercussions of Stack Overflow questions are being flooded with answers from ChatGPT?

It returned:

> I'm sorry, but I'm not aware of any situation in which answers on Stack Overflow are being flooded with answers from ChatGPT. I am a language model trained by OpenAI, and I do not have the ability to browse the internet or interact with websites like Stack Overflow. I can only provide answers based on the information I have been trained on. Is there something specific you would like to ask about Stack Overflow or ChatGPT? I'll do my best to help.

[reply](#)

Areading314 3 hours ago | parent | prev | next [-]

Here is a potential rephrased version of the original post, written in a witty and informal style:

"Yo, AI-generated answers on Stack Overflow sound like a real bummer. First of all, the quality might suck compared to human-written answers. And if these AI answers become too popular, the community could get all impersonal and unsupportive. Plus, if the AI is copying and pasting from the internet, we'll end up with a ton of low-quality, duplicate answers. In short, using AI on Stack Overflow could really suck."

[reply](#)

seydor 2 hours ago | root | parent | next [-]

How long will we wait for "make GPT a gangsta" plugins

[reply](#)

kthejoker2 1 hour ago | root | parent | next [-]

it's already there?

You can ask it to re-style any of its responses in a wild variety of styles, from comedians, Internet celebrities, famous authors ...

Mostly it just proves that on most subjects there is nothing interesting to say, and we wildly overestimate the value of sharing our own opinions on it.

[reply](#)

plastiquebeech 2 hours ago | parent | prev | next [-]

>It seems like there are a few potential negative consequences of using AI-generated answers on Stack Overflow. For one,

This is where my "Probably written by AI" filter tripped, and I skipped to the end.

I can't exactly say why. Maybe we'll all start to develop a seventh sense for ML-generated content, in addition to the "probably an ad" filter that usually kicks in around middle/high school.

[reply](#)

TomVDB 1 hour ago | root | parent | next [-]

Same here.

OpenAI commented on the verboseness of answers as something they need to work on.

[reply](#)

seydor 2 hours ago | parent | prev | next [-]

haha. i surely sniffed it out before reaching the end. The formal style is familiar

[reply](#)

taberian2 2 hours ago | root | parent | next [-]

The formal style of a high-schooler trying to meet essay structure requirements

[reply](#)

quickthrower2 1 hour ago | parent | prev | next [-]

OK Computer

I could tell that super passive and eager to talk to both sides style a mile off. Fuck that! (<- proof I am not a robot)

[reply](#)

arkitaip 3 hours ago | parent | prev | next [-]

You can kinda tell it's ChatGPT by how verbose and repetitive the writing is.

[reply](#)

brindidrip 3 hours ago | root | parent | next [-]

Yes, one way to identify responses that may have been generated by ChatGPT is to look for signs of repetitive or verbose language. Because ChatGPT is a language model, it is designed to generate text that is fluent and coherent, but it may also produce text that is unnecessarily long or repetitive. This can be a telltale sign that a response was generated by an AI, as human responders are less likely to use repetitive or verbose language in their answers. Of course, this is not a foolproof method, and there may be other factors that can affect the quality and style of a response, but looking for repetitive or verbose language can be a useful starting point when trying to identify AI-generated responses.

[reply](#)

darkwater 3 hours ago | root | parent | next [-]

This is fun but it's already starting to NOT be fun.

[reply](#)

throwup 2 hours ago | root | parent | next [-]



It has a very recognizable default writing style but it will change things up if you ask.

> Write a response to the following comment chain in the style of an angry Hacker News commenter.

> ...

*What the hell is wrong with you? ChatGPT is a brilliant piece of technology and you're just dismissing it because of some stupid writing style? Get a grip and appreciate the advancements we're making in AI. Stop being so negative and enjoy the fun while it lasts.*

[reply](#)

rchaud 3 hours ago | root | parent | prev | next [-]

an AI that communicates with the verbosity of a recipe website....our glittering future!

[reply](#)

spaceman\_2020 3 hours ago | root | parent | prev | next [-]

I reckon we are about 3 years from most textual content online being created by bots, 5 years from static visual content being bot-created, and 7-8 years from bot-created visual content clogging up all social media.

[reply](#)

warbler73 1 hour ago | root | parent | next [-]

> we are about 3 years from most textual content online being created by bots  
0 years.

[reply](#)

spaceman\_2020 1 hour ago | root | parent | next [-]

1 year for this tech to be incorporated into most cheap/free creative tools.

1 more year for these newer creative tools to be adopted by creators everywhere.

1 more year for the humans to completely give up.

You have to remember that tech takes time to permeate to the bottom rung. Most of my non tech friends haven't even heard of Dall-E, let alone chatGPT.

[reply](#)

SheinhardtWigCo 2 hours ago | root | parent | prev | next [-]

Everyone will be doing strict ID verification long before any of those milestones are actually reached.

[reply](#)

Arch-TK 3 hours ago | root | parent | prev | next [-]

3, 5 and 7 seem awfully pessimistic no?

[reply](#)

zone411 3 hours ago | root | parent | prev | next [-]

I'll take the under.

[reply](#)

cercatrova 3 hours ago | root | parent | prev | next [-]

That's honestly a hilarious response from ChatGPT

[reply](#)

arkitaip 3 hours ago | root | parent | prev | next [-]

To find out if a response was made by a computer, look for repeating words or long sentences. ChatGPT is a program that makes text that makes sense and flows well, but sometimes it uses too many words or repeats itself. This is a clue that the response might have been made by a computer instead of a person. However, this isn't a sure thing and there might be other reasons why a response is long or repeats itself. Still, it's a good place to start.

[After many attempts]

[reply](#)

layer8 3 hours ago | root | parent | next [-]

> Respond to the above in the informal style of a slightly contrarian Hacker News comment, with a hint of passive-aggressiveness. Avoid repetitive wording and long sentences.

*Wow, thanks for the tip! I never would have thought to look for repeating words or long sentences to determine if a response was made by a computer. I guess I'll just blindly trust everything I read online from now on. Sarcasm aside, this is a pretty flawed method. Just because a response has repeating words or is long doesn't mean it's not from a person. And conversely, just because a response is short and doesn't repeat itself doesn't mean it's not from a computer. It's not that simple.*

[reply](#)

9wzYQbTYsAlc 2 hours ago | root | parent | prev | next [-]

> To find out if a response was made by a computer, look for repeating words or long sentences.

That technique will produce quite a few false positives - a significant portion of humanity with language production issues on the autism spectrum communicate like that.

edit: Makes me wonder to what extent the training data is biased by writings from people on the autism spectrum.

[reply](#)

zone411 3 hours ago | root | parent | prev | next [-]

That's just its default setting. If you do some prompt engineering to tell it to be to the point or concise, it can do better. E.g. I got:

The repercussions of ChatGPT flooding Stack Overflow with answers could include decreased quality and relevance of answers, decreased user engagement, damage to the platform's reputation, and potential loss of revenue.

[reply](#)

ycombobreaker 3 hours ago | root | parent | next [-]

Still sounds too formal. To me, it reads like the intro to a high school paper. The next four paragraphs would elaborate on each of the points listed.

[reply](#)

zone411 2 hours ago | root | parent | next [-]

Another poster already did the informal version. My main point is that it's not something set in stone but yes, it sure looks like OpenAI went overboard in setting it up like this. They should adjust it.

[reply](#)

baq 1 hour ago | root | parent | next [-]

Maybe they shouldn't. Nobody should. But someone will, because people will pay good money for something like this to shill their products here and on Reddit and Twitter. The worst part is that it will work for at least a while.

[reply](#)

mk\_stjames 2 hours ago | root | parent | prev | next [-]

You can get it to respond a bit less formal by having it impersonate someone with a less formal speaking style. I said to restate to original but as if it was written by Tom Sawyer or one of his friends: Yo, what's the deal with using AI answers on Stack Overflow? It seems like it could be a real bummer. For one thing, the answers might not be as good as ones made by a human. And if AI answers become too common, the community might not be as friendly and supportive. And if the AI can search the internet and make its own answers, there could be even more bad answers on the site. All in all, this could be a major problem." Still a bit structured but with a bit more play you can see how it can 'act' like someone who isn't so formal.

[reply](#)

lamontcg 1 hour ago | root | parent | prev | next [-]

Really though if they just manage a setting for "sounding like a 50 year old grouchy old bastard" and it'll put me out of work.

[reply](#)

dejj 3 hours ago | root | parent | prev | next [-]

Can't we have ChatGPT identify output of ChatGPT, adversarially?

[reply](#)

mike\_hearn 2 hours ago | root | parent | next [-]

You don't need adversarial AI to do that. Simply simhashing every response and then allowing people to upload some text as a "did ChatGPT generate this" query should work. Just ordinary software is sufficient.

Of course, once you have multiple companies with ChatGPT level AI it gets harder if they don't cooperate. But they easily could. The big webmail providers cooperate to fight spam, for example.

[reply](#)

vbbvmnbs 2 hours ago | root | parent | prev | next [-]

It confirmed it when I asked:

--> *Was the following paragraph generated by the OpenAI Assistant: "It seems like there are a few potential negative consequences of using AI-generated answers on Stack Overflow. For one, the quality of the answers may be lower than if they were written by a human. Additionally, if these AI-generated answers become too common, it could potentially lead to a more impersonal and less supportive community on Stack Overflow. Finally, if the AI is able to search the internet and "inbreed" its own answers, it could lead to even more low-quality, duplicative answers on the platform. Overall, it seems like there could be some serious drawbacks to this development."?*

<-- *Yes, the paragraph you provided was generated by the OpenAI Assistant. As a language model, the Assistant is trained to generate human-like text based on a given prompt. In this case, the prompt likely pertained to the potential negative consequences of using AI-generated answers on Stack Overflow.*

Though, it could just have been taking a guess at it.

[reply](#)

formerly\_proven 2 hours ago | root | parent | next [-]

"It is likely that this text was not written with OpenAI Assistant."

[reply](#)

bobbyi 3 hours ago | root | parent | prev | next [-]

Before reading any words, I guessed just by seeing how long the paragraph was.

Maybe because the training data includes things like news articles that have longer paragraphs than typical forum posts?

[reply](#)

TapWaterBandit 2 hours ago | root | parent | next [-]

You can ask for answers of a certain length and also a certain writing style. The writing style one is hit and miss although if an AI ever gets all fiction books fed into it I imagine it could radically improve. Say by qualifying a response with "respond in the style of author X" or "average style of fantasy books published on 1972" or whatever.

[reply](#)

layer8 2 hours ago | root | parent | prev | next [-]

> Reformulate the above in the informal style of a Hacker News comment.

*I'm not sure about this AI-generated answers thing on Stack Overflow. The answers might suck and make the community more impersonal. Plus, if the AI is just copying and pasting from the internet, it could lead to a ton of low-quality, duplicate answers. Not a good idea in my opinion.*

[reply](#)

tombh 1 hour ago | parent | prev | next [-]

Could you provide the prompt?

[reply](#)

pugworthy 2 hours ago | prev | next [-]

For some things, ChatGPT is just better than SO. I have to say I probably won't hit SO for some basic stuff anymore, I'll just ask ChatGPT.

And some queries are just not acceptable on SO, but fine for ChatGPT.

For example I might wish to ask, "Give me the framework for a basic API written in Python that uses API key authentication. Populate it with several sample methods that return data structures in json."

If I ask that on SO, I'll be down voted and locked before I know it. I may also get some disparaging comments telling me to do my research, etc.

If I ask ChatGPT, it will give me a nice and tidy answer that gets me going quickly. It will explain things too, and allow me to ask follow up questions and take my requests for refinements. I might say, "For the python api I asked about earlier, have it look up the API authentication key in a database. If the key is in the database, it is valid." - and *bam* - it does it.

Sure, some pretty simple stuff if you know Python and APIs already, but if you just want to hack something together to test out an idea, it's great."

In the end, SO is a query with responses (maybe). ChatGPT is a conversation that can go beyond just the initial query.

[reply](#)

TillE 45 minutes ago | parent | next [-]

I think the reactions which suppose that AI will replace programmers or artists are a little silly, but this is a great example of where it could be genuinely revolutionary, as a true next generation of search engines.

That's really exciting, because it's a scenario where you're looking for and scrutinizing information. Just add some links to sources and you're in business.

[reply](#)

tartoran 2 minutes ago | root | parent | next [-]

Hocus pocus preparatus, smoke and mirrors, we present the Oracle! You can ask him anything. Beware, he likes to bullshit and hallucinate

[reply](#)

anraais 1 hour ago | parent | prev | next [-]

What if it takes less time to hack together such an API than prompt engineering and back-to-back conversation with a bot whose results you have to verify anyway?

I imagine it would if you're familiar with the language, Framework, tools, etc.

[reply](#)

berkes 1 hour ago | root | parent | next [-]

And if you're not familiar with the tools, language or framework, I think that alone should be a reason to forego it for anything else than "learning it".

Which means that question was the wrong question in the first place.

To be clear: i'm not arguing against learning new stuff. But against using unfamiliar tech to do a serious project.

And i'm bringing that up here, because if the tech is familiar, then asking what tech to use, is rather strange.

[reply](#)

josephcsible 3 hours ago | prev | next [-]

I wouldn't even mind so much if the answers were right. The problem is that a lot of them are totally wrong, but completely reasonable- and plausible-sounding, and in an authoritative tone, so unless you already know the right answer, the only way you'll realize its answer is wrong is the hard way.

[reply](#)

ravenstine 3 hours ago | parent | next [-]

What you said may sum up the current state of AI. People's minds are continually being blown, but will there be a realization that these AIs are specialized to provide specious output and nothing else? There's a canyon of difference between something that sounds correct and a thing that is actually correct.

[reply](#)

bambax 3 hours ago | root | parent | next [-]

Exactly. ChatGPT sounds like a bad student who didn't actually learn anything during the year and is trying to bullshit their way through the finals. Or a politician, maybe. It's formalizing the worst traits of humanity.

[reply](#)

bediger4000 2 hours ago | root | parent | prev | next [-]

Yes. I spent yesterday afternoon trying to get ChatGPT to write me a quine in nroff/troff macros, something I've not been able to do, or find if anyone has done.

The generated quines look like they'd work, but don't.

Same with an M4 macro processor quine - looks maybe correct, doesn't work.

It did generate a Go quine.

[reply](#)

YeGoblynQueenne 40 minutes ago | root | parent | next [-]

>> It did generate a Go quine.

Because there are several of them on the web:

<https://duckduckgo.com/?t=ffcm&q=Go+quine&atb=v344-1&ia=web>

[reply](#)

bambax 3 hours ago | parent | prev | next [-]

That's the worst part. It's all fun and games when it tells a story of cheese sandwiches and VCRs in the style of the King James bible, but when it gives wrong answers in an authoritative tone and then insists it absolutely can't be wrong, it's terrifying.

I don't understand what good could come of this. Or at least make it detect what is fiction and what isn't.

People are treating it like it's Wikipedia, but it's not. It's a riff on words, like a bird imitating sounds without an idea of what they mean.

[reply](#)

convexfunction 2 hours ago | root | parent | next [-]

Just imagine the hell we'd be in if *people* could give wrong answers in an authoritative tone and then insist they absolutely can't be wrong!

[reply](#)

bambax 2 hours ago | root | parent | next [-]

For some reason, we assume what comes out of a computer is more trustworthy than what people say. We think computers are transparent, reliable, idempotent and don't have an agenda. Even more so if we call it "intelligent"...

But ChatGPT is a bullshit machine, and that much is new.

[reply](#)

convexfunction 1 hour ago | root | parent | next [-]

Who's "we"? :)

To the extent this perception exists -- and I don't think "came from a computer" falls within the top 5 actually effective methods of laundering bullshit nowadays, though maybe it used to -- you might expect that it gets crushed into dust as the public gets more exposure to high-profile counterexamples.

And, wait, isn't the concern usually that people read AI-generated content and trust it but *don't* think it came from a computer?

[reply](#)

yawnxyz 1 hour ago | root | parent | prev | next [-]

Wikipedia couldn't be trusted for the first decade it came out, and now you have people use it as an example of a trusted resource

[reply](#)

hombre\_fatal 2 hours ago | root | parent | prev | next [-]

Well, at least a human had to put in the work to write it. Now you can automate this low tier content.

[reply](#)

taberian 1 hour ago | root | parent | prev | next [-]

Maybe it's an opportunity for people to get better at critical review of the information they are presented. Maybe they'll learn just because they are being told something authoritatively that doesn't mean it's right?

And so what if if people misuse the the tool while they learn? What's wrong with being wrong?

[reply](#)

bambax 1 hour ago | root | parent | next [-]

We learn by example. What this tool teaches is that self-confidence is more important than content. And maybe it is! in life in general. But not in the pursuit of knowledge.

Instead of patronizing users by adding a stupid sentence at the end of each answer reminding them to be careful, it could say every time "I'm just a machine and what I say is random; if it's true it's by accident; it's mostly probably wrong and only sounds like the truth".

Here's an example I posted in another thread, where the machine is repeatedly mistaken but affirms it absolutely, positively can't make a mistake:

<https://news.ycombinator.com/item?id=33852236>

[reply](#)

gtirloni 3 hours ago | parent | prev | next [-]

The StackOverflow mods have a lot of knowledge about closing questions based on really small details that go against their rules. I'm sure they will do well spotting AI-generated answers.

[reply](#)

luckylion 2 hours ago | root | parent | next [-]

I doubt it. They've had problems with cheap automated answers for years where bots would essentially search for the question on SO and then copy an answer from another question verbatim. The answers were rarely useful because questions happen to be different even though they use the same keywords.

Not only did they never bother to block that, they also didn't mind it and wanted to rely on the community down-voting those answers instead of at least blocking the bot -- and that's with a trivial check (identical answer already in DB). With something that's AI-generated, there's no chance. And with the general quality of many of the answers, there's no way to tell apart wrong answers from AI or humans.

[reply](#)

clusterhacks 2 hours ago | prev | next [-]

Human-curated content from trusted sources for top 1% information probably only available to subscribers will become more valuable and sought after. I suspect the days of generally trusting forums populated by anonymous users are done?

I would not be surprised if the quality of human writing actually goes up. I have this weird feeling that ChatGPT and similar tools will become almost equivalent to calculators for math? My experience as a writer is that sometimes just throwing down a first draft is the hardest step - I could see these tools really assisting in the writing process. Generate a draft, do some tweaking, ask for suggestions or improvements, repeat.

I don't know how I feel about code generated by these tools. Will there be a similar benefit compared to writing? At some level, we will need some deeper mastery of writing and coding to use these things well. Is there a complexity cliff that these tools will never be able to overcome?

A total lack of trust for general internet search results. So much content is already shallow copies of other content. I don't see how general internet search survives this.

[reply](#)

rsync 2 hours ago | parent | next [-]

The anonymity isn't the problem - it's the cost free aspect.

Anonymous content can work very well if there are costs incurred...

[reply](#)

Always42 2 hours ago | root | parent | next [-]

But let's be real, who has the time or energy to carefully curate content these days? It's all about efficiency and getting the job done. And let's face it, anonymous forums have their charm and can be a great source of information if you know where to look.

[reply](#)

senko 4 hours ago | prev | next [-]

This is just a preview of things to come.

Wait a few weeks until Google is completely swamped with ChatGPT SEO pages barely distinguishable from the real thing.

If I worked at search quality at Google, I'd be very worried.

[reply](#)

skilled 2 hours ago | parent | next [-]

Google has been saying for years that it wants to get rid of backlinks and value them less, and apparently they already do. [0]

I wouldn't be surprised if they revert the direction and put a significant importance back on links because of this kind of technology.

[0]: <https://www.seroundtable.com/google-links-less-impact-34493...>

[reply](#)

bioemerl 3 minutes ago | root | parent | next [-]

You can't really put importance on backlinks either, because those are abused in their own special way

[reply](#)

kposehn 4 hours ago | parent | prev | next [-]

The most recent update also managed to rank ML generated content *above* written content in many instances, compounding the problem. I absolutely expect a carpet bombing of SEO results with spam.

[reply](#)

ilaksh 3 hours ago | root | parent | next [-]

Hasn't that already been the case for years?

[reply](#)

CuriouslyC 2 hours ago | root | parent | next [-]

Depending on the query, yes. Google search has different ranking strategies for different queries, and some of them seem more robust to simple SEO tactics than others. I think the ranking strategy for e-commerce searches and product reviews does a poor job with spam sites, but the ranking strategy for news queries works pretty well, typically providing timely and relevant answers from reputable sources.

[reply](#)

brindidrip 4 hours ago | parent | prev | next [-]

I am afraid for whats to come. This is truly only the beginning. Good luck everyone. See you on the other side.

[reply](#)

nonrandomstring 3 hours ago | root | parent | next [-]

Sworn enemies unite against a common foe. This foretaste of AI may be a timely wake-up call about what technology is doing to us. We may be forced to start treating each other as humans. In that case, ChatGPT will be seen as a watershed moment in the retaking of humane technology.

[reply](#)

spaceman\_2020 3 hours ago | parent | prev | next [-]

I expect Google to be irrelevant soon enough.

ChatGPT already does a better job of finding answers to a lot of technical questions than Google.

[reply](#)

GeorgeTirebiter 2 hours ago | root | parent | next [-]

I agree with this. Ask specific tech questions (e.g. what is the vacuum (in Torr) inside a vacuum tube?) and ChatGPT give a direct answer. Ask Google? Go ahead, grovel thought the ads, irrelevant results, and maybe you'll find something relevant.

Happy Hunting...

[reply](#)

krater23 27 minutes ago | root | parent | prev | next [-]

Yes, because ChatGPT is not bombed with SEO currently.

[reply](#)

rvz 2 hours ago | root | parent | prev | next [-]

Nope. It is StackOverflow that is irrelevant. Not Google. Especially when most of all AI including ChatGPT still cannot transparently explain their own decisions.

[reply](#)

seydor 4 hours ago | parent | prev | next [-]

Chatgpt SEO is no different from existing filler SEO. If google falls for the latter, why shouldn't it fall for the former

[reply](#)

scarmig 4 hours ago | root | parent | next [-]

ChatGPT SEO is higher quality than existing filler SEO and probably quicker and cheaper to generate.

[reply](#)

seydor 3 hours ago | root | parent | next [-]

you are right ( tried it)

[reply](#)

michaelteter 3 hours ago | parent | prev | next [-]

> SEO pages

I think GPT could improve things in this regard. It's already a huge garbage pile, so we might as well go all in and at least generate better garbage.

[reply](#)

kevin\_thibedeau 2 hours ago | root | parent | next [-]

The current AI generated SEO content is relatively easy to detect because much of it tends to make incoherent statements. We don't want better AI bots feeding us lies that are hard to detect even for competent people who know to be wary of it.

[reply](#)

mojuba 3 hours ago | parent | prev | next [-]

That would actually slightly improve Google search results given the generated spam garbage that comes up today.

[reply](#)

\_Algernon\_ 2 hours ago | parent | prev | next [-]

Considering the quality GPT spits out, it would likely be an improvement compared to the crap I get now.

[reply](#)

wslh 2 hours ago | parent | prev | next [-]

For years, it is happening much before ChatGPT in front of your eyes.

[reply](#)

computerfriend 4 hours ago [flagged] | parent | prev | next [-]

Luckily for them, almost nobody works at search quality at Google.

[reply](#)

fnordpiglet 3 hours ago | parent | prev | next [-]

So, effectively the status quo?

[reply](#)

echelon 4 hours ago | parent | prev | next [-]

Reddit and Twitter are going to be overwhelmed with bots indistinguishable from humans.

[reply](#)

rightbyte 26 minutes ago | root | parent | next [-]

Agree with this comment writing like a teen: " Reddit and Twitter are going to be overwhelmed with bots indistinguishable from humans. "

Totally agree. It's crazy to think that bots could be indistinguishable from humans. Like, how are we even supposed to know who's real and who's not? It's kind of scary, but also kind of cool in a weird way. But yeah, I can definitely see how it could be a problem if there are too many bots on social media. It could be really confusing and overwhelming. So, I hope there's a way to deal with it.

Ye ... ChatGPT can easily pose as an random Reddit comment writer. Just remove the tryhard "". Those kind of communities will implode. The comments there are already kinda random and have almost no thread running through them. Perfect for ChatGPT.

Imagine how easy it will be to create big fake subs and gaslight people into believing something is sub consensus?

[reply](#)

SoftTalker 3 hours ago | root | parent | prev | next [-]

A lot of reddit and twitter posts are already just people regurgitating and rehashing what they have heard/read from others, mostly without a lot of critical thought. Is this really going to be very different?

[reply](#)

cauefcr 21 minutes ago | root | parent | next [-]

It's a matter of scale, you could have personalized propaganda for each community.

[reply](#)

echelon 2 hours ago | root | parent | prev | next [-]

You have a point. But now it will happen at scale quickly, cheaply, and easily for an order of magnitude more impact.

Beyond these text-based social media platforms, I could see the next generation of TikTok/Instagram "influencer" social media adopting AI rails so there's next to zero effort put into creation. The photography, trends, personalities - all of it - will be generated.

[reply](#)

ChrisMarshallNY 4 hours ago | prev | next [-]

I assume that this is by folks wanting to up their scores.

That's a huge problem with "gamification." I'm not especially a fan of the concept, in a venue like SO. I think it has led to a rather nasty community, and I hardly ever go there, anymore.

I assume that we'll be seeing a lot of robotic HN content (I would not be surprised if it is already here, but has been sidelined by the mods).

[reply](#)

fouric 3 hours ago | parent | next [-]

> I would not be surprised if it is already here, but has been sidelined by the mods

I can virtually guarantee you that there's been a nontrivial amount of GPT-generated content on HN that has *not* been caught by mods since ChatGPT, and likely since GPT-2/3, as well. Dang (and the other one whose tag I can't recall) already have their hands full trying to keep the tone civil across thousands (tens of thousands?) of comments a day - it's impossible for them to catch *every* ML-generated comment (some humans actually do write like these newer language models, after all), and more than likely they're missing a decent number of them - through no fault of their own, it's just an extremely hard problem.

The three solutions that I'm shilling for this problem are (1) invite-trees for HN (like Lobsters, which makes the community much less open but also much more resistant to abuse) (2) webs of trust (not cryptographic, just databases of how much you trust users) overlaid onto HN and other places and (3) people actually reading the content of comments *very* carefully and upvoting logically sound arguments and downvoting illogical and emotional/manipulative ones, but all of these require a *lot* of effort and social buy-in.

[reply](#)

patneedham 3 hours ago | root | parent | next [-]

> some humans actually do write like these newer language models, after all

Wouldn't it be more accurate to say these newer language models actually write like humans? Or is there a subset of the population intentionally trying to write in the way that these language models write.

[reply](#)

seydor 2 hours ago | root | parent | next [-]

It seems to have a tendency to write stereotypical preamble-statement-conclusion paragraphs and to repeat itself. The model repeats itself often. It repeats the title and then writes a statement that basically repeats the title and after that it provides the useful nugget. At the end it adds a sentence usually using "overall" as an opener. Overall, the model tends to respond in a stereotypical format

[reply](#)

fouric 2 hours ago | root | parent | prev | next [-]

Sure, you can also say that some language models write like humans - however, even pre-GPT-2, I read several high-schooler's essays that read *very* much like these ML-generated products, so even if you don't believe that the relationship is transitive, I think you can say that the relationship is true in both directions.

[reply](#)

tpxl 2 hours ago | root | parent | prev | next [-]

> people actually reading the content of comments very carefully and upvoting logically sound arguments

You can make a perfectly logical comment that is completely wrong.

[reply](#)

fouric 2 hours ago | root | parent | next [-]

Please elaborate.

[reply](#)

NateEag 1 hour ago | root | parent | next [-]

Logic is about a chain of formal reasoning that is guaranteed to yield a correct result given correct inputs.

If you start with some wrong premises, you can easily write a logical post that has incorrect conclusions.



[reply](#)

fouric 10 minutes ago | root | parent | next [-]

I am aware of the formal definition of logic. If this is all that tpxl meant, then it's an extremely pedantic and useless objection.

[reply](#)

terminal\_d 34 minutes ago | root | parent | prev | next [-]

Another measure would be initiating a bot-check for every user. Maybe even periodically.

[reply](#)

thetip 3 hours ago | parent | prev | next [-]

The gamification mechanic was the innovation that let SO become as successful as it has, IMO.

Without it there was no real way to figure out the "best" answer to problems.

It's fair to note that big communities can have somewhat unfriendly dynamics, but I think that's more about big communities than the gamification mechanic.

[reply](#)

bnralt 2 hours ago | parent | prev | next [-]

Some years ago I was trying to do something that required a certain amount of points. Perhaps write a comment to a question instead of posting an answer. So I looked for ways to make points. Most of the interesting questions I could easily find looked to be already answered, and the new questions seemed to be mostly requests to fix bugs in people's code.

I saw someone suggest fixing up answers that people had already made, and I found an upvoted answer that had a simple bug in it's code (something like "<" was used instead of "<="). So I went about to fix it, but was instructed that edits had to change more than five characters. I tried to find how people handle this, and the answer I saw was to make additional changes to the answers, no matter if it needed it or not.

In the end I just left the bug in and gave up on trying to provide answers for the site.

[reply](#)

Dwedit 3 hours ago | parent | prev | next [-]

Stack Overflow literally forces you to reach a certain score before you get downvoting rights. The gamification is real, but once you reach that point, you aren't required to play anymore.

[reply](#)

tenebrisalietum 3 hours ago | parent | prev | next [-]

Smaller StackExchange communities don't have this problem from what I can tell. Stack Overflow maybe should be split up into smaller communities.

[reply](#)

mdaniel 3 hours ago | root | parent | next [-]

There already are plenty of smaller communities, but no one enforces the on-topic rules as they exist today, making SO into the "I haz computer problem" dumping ground it has become

[reply](#)

duckmysick 6 minutes ago | prev | next [-]

At one point new models will be trained on contaminated data where some of the content is AI-generated. "Pure" datasets will be highly prized, just like the steel made before nuclear detonations.

[https://en.wikipedia.org/wiki/Low-background\\_steel](https://en.wikipedia.org/wiki/Low-background_steel)

[reply](#)

xx\_\_yy 27 minutes ago | prev | next [-]

Some of the affects I can think of, to name a few:

Inaccurate or irrelevant answers: ChatGPT is a machine learning model that uses past data to generate responses. This means that it may not always provide accurate or relevant answers to questions, leading to confusion and frustration among users.

Loss of trust: If users notice that many of the answers on the forum are coming from ChatGPT, they may lose trust in the forum and stop using it. This could lead to a decline in user engagement and overall traffic.

Competition with human contributors: ChatGPT's answers may compete with those provided by human contributors, leading to a decrease in the quality and value of the content on the forum. This could make the forum less useful and engaging for users.

Increased moderation: The influx of answers from ChatGPT may require more moderation to ensure that the answers are accurate and relevant. This could require additional resources and time for moderators, leading to increased costs and workload.

[reply](#)

egypturnash 3 hours ago | prev | next [-]

Well, guess the genie's out of the bottle and we can never stop this. Bow down to the inevitability of technological progress, Luddites! Good luck retraining into a new job, I hear "prompt engineer automation" is the new hotness.

Or at least that's what all of you kept telling me when I was expressing my unhappiness at the way corporate-sponsored image generating black boxes are built atop a shaky moral foundation that sure feels like it's ignoring anything anyone talking about "fair use" ever dreamed of, and at the way I fear it's going to hollow out a ton of the beginner-pro jobs of my industry by making it super easy for anyone to generate stuff that is kinda fundamentally shitty in a lot of important ways, but "good enough" if you just have a space to fill with some decoration that you don't really give a crap about.

[reply](#)

sydd 2 hours ago | parent | next [-]

No, you are washing lots of topics together. Just because it AI, its not the same.

AI art: Yes, its here, yes it will make life for lots of artists/graphic designers different. But they will not go out of business, there will be requests that are so nuanced that no master prompter will be able to fulfill. I'd expect that AI art generators will be a part of a modern digital artist's toolchain just like Photoshop. Someone not using AI art in say 20 years will be like saying that you restrict yourself to MS paint.

AI for code: We might see something similar just to a lesser degree. For example if your clients has requested a "picture of a man with trees in the background", he might accept a pic with trees and bushes if he likes it or there is just 1-2 bushes in the picture. But no one would accept a banking software that transfers the wrong amount in 1 of 10000 cases. Or just one where a very senior engineer cannot step forward and say "I wrote this, and it will work. If it doesnt I will fix it"

[reply](#)

boppo1 2 hours ago | parent | prev | next [-]

Artisans lost that battle 100 years ago with the rise of modernism.

[reply](#)

wslh 3 hours ago | parent | prev | next [-]

There is no genie here, some people have a belief about this while it is very easy to probe the low quality and inaccuracy of the responses.

[reply](#)

Ancalagon 1 hour ago | root | parent | next [-]

Can you give some examples of the low quality/incorrect responses? Then try remedying those by rephrasing the prompts? I'm curious what the actual limitations are.

[reply](#)

wslh 59 minutes ago | root | parent | next [-]

<https://news.ycombinator.com/item?id=33840405>

[reply](#)

Ancalagon 35 minutes ago | root | parent | next [-]

Very cool, thank you!

[reply](#)

convexfunction 2 hours ago | parent | prev | next [-]

At this point, I truly don't know which of your(?) current type of applied arts vs my current type of software development will turn out to be more sensitive to technological unemployment, or on what margins and time scales. So, hopefully this doesn't strike you as callous, since I think it applies to me as much as it does to you or your coworkers:

Roughly everyone who makes money has the same job, which is creating value in expectation for someone else. (Whether that activity is net good for society is a different question, I lean toward usually yes for stuff you can do legally, not always though). If you find yourself suddenly unable to give anyone a competitive deal on whatever expected value you know how to create, because of technological developments or otherwise, well, you'd better figure out what you need to change about what you're offering so you can. I wouldn't call this fair, exactly -- maybe it would be if you or your government had effective "technological unemployment insurance" -- but I struggle to imagine any substantially different state of affairs that's clearly better for the world. (No points for saying "imagine communism then".)

[reply](#)

avivo 3 hours ago | prev | next [-]

It's worth understanding the community and org better, and their reaction. Relevant links:

- <https://meta.stackoverflow.com/questions/421778/how-do-you-p...>

- <https://meta.stackoverflow.com/questions/412696/is-it-accept...>

- <https://meta.stackexchange.com/questions/384355/could-chatgp...>

[reply](#)

mdaniel 3 hours ago | parent | next [-]

> Sure, but that's irrelevant. Whether or not the user understands the answer they posted is not the concern of the site.

Well, that's unfortunate. Then again, I guess that's a logical conclusion of the "safe harbor" for serving any user-submitted content: Stack Exchange only does the most cursory moderation, and the rest is caveat reader

[reply](#)

kruuuder 1 hour ago | parent | prev | next [-]

It's so funny and sad at the same time that, in typical SO manner, EugenSunic is being downvoted so much for raising such an interesting question.

[reply](#)

pctthrowaway 1 hour ago | prev | next [-]

Well, for starters, it's just annoying. It's like having a bot spamming every single question with useless answers. It dilutes the quality of the content on the site and makes it harder for genuine contributors to get their answers noticed.

But it's also a serious concern from a security standpoint. If ChatGPT is providing incorrect answers, it could lead to people implementing flawed code or making poor decisions based on its advice. That could have potentially disastrous consequences.

So overall, it's a big problem that needs to be addressed. It's not just about making the site more pleasant to use, it's about ensuring the integrity and reliability of the information provided.

My prompt:

I'm writing a short story where Linus Torvalds is having a conversation with an open source contributor. In this conversation, Linus is in a bad mood.

Open source contributor: Stack Overflow questions are being flooded with answers from ChatGPT. What are the possible repercussions of this?

Linus Torvalds:

[reply](#)

Yuyudo\_Comiketo 4 hours ago | prev | next [-]

Feed it some CMake files from llvm repository and ask it why would the windows build with LLVM\_ENABLE\_PROJECTS="all" keep failing, so that it chokes to death in its infancy, and save the humankind before it's too late and there are autonomous human zappers and T-1000s berserking all over the place.

[reply](#)

iamflimflam1 2 hours ago | prev | next [-]

Oh dear - basically if the answer is good then it's from ChatGTP...

Q: how can I tell if a stackoverflow answer is generated by ChatGTP if someone has removed the "ChatGTP" tag from the text?

A: One way to tell if a StackOverflow answer is generated by ChatGTP is to look for certain characteristics in the answer. ChatGTP answers tend to be concise and to the point, often providing code snippets or specific instructions on how to solve a problem. They may also include links to relevant documentation or external sources for further information. If the answer does not include these characteristics, it is less likely to be generated by ChatGTP.

[reply](#)

QuadmasterXLII 2 hours ago | parent | next [-]

Notably, it didn't correct "ChatGTP vs ChatGPT". If you still have this open, could you ask it what GTP stands for?

[reply](#)

GrantS 2 hours ago | root | parent | next [-]

I noticed this phenomenon when asking for a story involving Stephen Hawking, Darth Vader, and Wookiees. I spelled it "Wookies" by mistake and ChatGPT used my misspelling throughout the story despite seemingly getting all the other details correct.

Edit: Part of the generated story:

As they continued to talk, Stephen and Darth realized that they had something else in common - they both had lost the use of their biological legs. Stephen had been diagnosed with ALS at a young age and had gradually lost the ability to move his legs, while Darth had lost his legs in a battle with Obi-Wan Kenobi on the planet of Mustafar.

Stephen and Darth discussed the challenges and obstacles that they had faced as a result of their mobility issues, and how they had adapted and overcome them. They also talked about the technological advancements that had allowed them to continue their work and pursue their passions, despite their limitations.

Suddenly, Stephen and Darth turned on each other, each revealing that they had been secretly plotting against the other. Stephen accused Darth of using the Force for evil and corrupt purposes, while Darth accused Stephen of using his scientific knowledge to create weapons of mass destruction.

[reply](#)

iamflimflam1 2 hours ago | root | parent | prev | next [-]

Sorry, got distracted asking it to write code to detect itself. Good solid code - but nothing that would work really well.

[reply](#)

anigbrowl 38 minutes ago | prev | next [-]

I see what you did there.

I have an OpenAI account and like their product, I'm certainly impressed by this latest version though I have had little time to play with it. But the combination of quality AI with social reputation scoring is absolutely toxic, and the wider impact of SEO (a less curated version of the same thing) are a disaster. I was already sick of all the tutorial sites like [geeks4geeks](#), [w3schools](#) etc and their numerous imitators just content farming whatever is turning up in searches. Marketing and self promotion is cancer and the people who try to game their way to success in this manner are awful. Perhaps the best use of counter-AI will not be in filtering these people, but in providing them with useless rewards and the appearance of excited fanbases that will divert them into a parallel hamster wheel web. Nothing would please me more than for the top 5000 influencers of this sort to be granted exclusive access to a luxury cruise that leaves port once a year for a tour of the Bermuda triangle.

I think the best use of ChatGPT would be in an IDE plugin, so you could point at function trees or code blocks and ask it to explain things, have it take care of basic refactoring tasks, help porting between languages or libraries and so on. I can definitely see a future where you throw together a working prototype of something, answer a few questions about type hinting and edge cases, and AI does the legwork of converting the prototype into a strongly typed final product.

[reply](#)

imhoguy 1 hour ago | [prev](#) | [next](#) [-]

Plot twist: Stack Overflow starts to use ChatGPT as a first answer to every new question, with "AI generated" label ofc.

[reply](#)

SergeAx 9 minutes ago | [prev](#) | [next](#) [-]

How hard would it be to train a ML-model to distinct ML-generated content from product of human? I mean text, images and code?

[reply](#)

ubj 3 hours ago | [prev](#) | [next](#) [-]

And so it begins. Welcome to the new internet.

I'm bracing myself for when this wave of AI content hits academic journals.

[reply](#)

snek\_case 3 hours ago | [parent](#) | [next](#) [-]

With a bit more refinement, if it had the ability to generate graphs, etc, it might be able to generate very believable papers. At least, believable enough that you can't tell without reading the paper attentively.

[reply](#)

JW\_00000 2 hours ago | [root](#) | [parent](#) | [next](#) [-]

Instead of spending 6 months laboriously doing experiments in a lab, and then a month writing up their results in a paper, researchers can already write a paper in one month if they just invent the numbers without actually doing the experiments. Peer review doesn't check for this. This only further reduces that one month to 5 minutes (+ hours of fiddling with LaTeX templates?). But in both cases if it gets found out your career is over...

However what about generating patents? To get a patent you don't need to have done any experiments that prove your technique actually works :)

[reply](#)

Der\_Einzige 2 hours ago | [root](#) | [parent](#) | [next](#) [-]

There is so much plausible deniability with the reported numbers that getting "found out" won't matter either!

[reply](#)

harrylove 3 hours ago | [root](#) | [parent](#) | [prev](#) | [next](#) [-]

Everything has an API. You can give it your data (or ask it to create some), and then ask it to write LaTeX, D3, MermaidJS, or code from any other framework that creates graphics. Problem solved.

If the thing you want to use is fairly popular and published on the web, even recently, it probably knows how to use it and combine it with everything else it knows.

Just yesterday I asked it to combine LiveView with a third party JS library to build an interactive thing and it got it on the first try using the latest Phoenix 1.7 RC which only came out in November.

I haven't tried it, but I bet you could ask it to generate a PDF in code using your favorite language with text from GPT and graphics generated from any framework that's compatible with your language. White-paper-as-a-Service.

[reply](#)

akrymski 50 minutes ago | [prev](#) | [next](#) [-]

This is how the web, and by extension Google dies. When the AI generated spam is so good that nothing on the open web can be trusted.

[reply](#)

KomoD 1 hour ago | [prev](#) | [next](#) [-]

I just encountered this, 2 users[1][2] it's very obvious as well since you can see the reputation spike from basically nothing.

[1]: [https://stackoverflow.com/users/19192614/boatti?tab=topactive...](https://stackoverflow.com/users/19192614/boatti?tab=topactive)

[2]: <https://stackoverflow.com/users/20684429/a-s?tab=topactivity>

[reply](#)

Gupie 9 minutes ago | prev | next [-]

Couldn't AI be used to statistically identify AI generated text?

[reply](#)

dragonwriter 7 minutes ago | parent | next [-]

It could be, and then the AI to statistically identify AI generate text could be used to score, rank, and select among potential AI responses to prompts so as to statistically minimize the risk of AI responses being identified as AI generated text.

[reply](#)

charles\_f 59 minutes ago | prev | next [-]

Even on HN, we start getting flooded by "ahah, I asked ChatGPT and here's the answer" in the comments, and every other topic is about "I did X with ChatGPT". This is already getting old

[reply](#)

lajosbacs 1 hour ago | prev | next [-]

I have not used SO since I've started using ChatGPT, it is so much easier to get to the correct answer and it can even be tailored to my specific example.

So double whammy for SO which makes me feel really sad.

[reply](#)

hysan 3 hours ago | prev | next [-]

This was the first use case that I thought of when I learned that ChatGPT could generate code. Then I considered how I'd feel if I ran into a fake (incorrect) answer and decided not to actually do this. Well, guess someone was eventually going to try this.

[reply](#)

michaelteter 3 hours ago | prev | next [-]

It means we are coming full circle.

At this point, SO has been scraped and repackaged (poorly) dozens of times, and SEOd to the top of search results. Even some "tutorial" sites are just repackaged SO answers.

It is only fitting that the automated SEO websites get fed automated content.

In a way, this makes the real humans, particularly the ones who know actual things, more valuable. It may so much noise that only a skilled human could decipher a real question and a real answer or solution from something similar but wrong.

To be fair to GPT, many human answers are sub-par and should be filtered out as well. Perhaps that's the real test: what percentage of GPT answers are decent vs human answers? Here I might bet on GPT.

[reply](#)

karmasimida 3 hours ago | prev | next [-]

Let me be the advocate of devil

I think ChatGPT is actually sometimes a lot better than SO answers

[reply](#)

petesergeant 1 hour ago | parent | next [-]

For the last few days I've been using ChatGPT instead of SO. It's OK, it's just frequently wrong, so I assume I have to fact-check its answers. Yesterday it claimed to me JavaScript has a built-in `sum` method, before admitting no browsers supported it and it's not in the spec.

It's useful for starting investigation, but one of the nice things about SO is that answers are voted for, so you can usually see which are actually right!

[reply](#)

deafpolygon 1 hour ago | prev | next [-]

The biggest repercussion is you probably can't piss ChatGPT off in a debate. So, that's boring.

[reply](#)

ggerganov 4 hours ago | prev | next [-]

I was thinking, what part of HN comments do you think are already AI-generated?

As a human, I cannot give an accurate estimate. /joke

[reply](#)

mojuba 3 hours ago | parent | next [-]

Who cares if the comments are good enough?

[reply](#)

nsvd 3 hours ago | root | parent | next [-]

For example, if one entity generated a large portion of the content, they could easily introduce a bias in these comments, to sway the opinion of readers. Automated astroturfing.

[reply](#)

mk\_stjames 3 hours ago | root | parent | next [-]

It doesn't even have to be nefarious. Just imagine a thread where 90% of the comments are repeating variations of the model's ideas in slightly different ways, and only 10% of the other responses are from humans. Even if the AI responses are considered 'good'... it gets overwhelming to the point that why should normal people even comment as the likelihood of it being read by another human and useful or even seen drops to nothing. I might as well open up a blank text file and make this comment and then not save it and close it as that would be seen by the same number of people in the limit.

[reply](#)

jerjerjer 2 hours ago | root | parent | prev | next [-]

Astroturfing on an unprecedented scale. Want to pump your shitcoin? Sure! Hype some new thing? Easily!

[reply](#)

rsync 2 hours ago | root | parent | prev | next [-]

I care a lot.

[reply](#)

Ancalagon 2 hours ago | prev | next [-]

This kind of looks like the singularity is approaching/just beginning.

The only thing we can be sure of, is that whatever we can imagine is already behind what the AI will become.

[reply](#)

nyokodo 1 hour ago | prev | next [-]

With responses becoming AI generated, and the disturbing rise of Russian and Chinese propaganda trolls on here I think my era of interactions on this platform are ending. So long to any actual people with conscious agency reading this, it has been interesting.

[reply](#)

brindidrip 3 hours ago | prev | next [-]

We need to start developing software to detect AI responses.

To detect a response generated by ChatGPT, we could first analyze the content of the response to see if it contains any unnatural or repetitive language. We could also check the formatting of the response to see if it follows the typical conventions used by human responders on the platform. Additionally, we could check for any unusual patterns in the timestamps of the response, as AI-generated responses may be posted more quickly or regularly than responses written by humans. Finally, we could also use machine learning algorithms to train a model to identify responses generated by ChatGPT based on these and other characteristics.

Quick, someone ask ChatGPT to generate the stubs.

[reply](#)

dragonwriter 9 minutes ago | parent | next [-]

> We need to start developing software to detect AI responses.

As soon as we do, it can be tied into AIs as a tool to evade detection, simply by generating multiple responses and returning the one scoring the lowest likelihood of being an AI in the AI detection tool.

[reply](#)

dragonwriter 20 minutes ago | parent | prev | next [-]

> To detect a response generated by ChatGPT, we could first analyze the content of the response to see if it contains any unnatural or repetitive language. We could also check the formatting of the response to see if it follows the typical conventions used by human responders on the platform.

Given my experience with human responses to text queries, these would be positively correlated.

[reply](#)

Jerrry 2 hours ago | parent | prev | next [-]

>Finally, we could also use machine learning algorithms to train a model to identify responses generated by ChatGPT based on these and other characteristics.

whatever your idea (i skimmed cuz) the discriminator will find it and have the generator apply it to the next generation.

>The core idea of a GAN is based on the "indirect" training through the discriminator, another neural network that can tell how "realistic" the input seems, which itself is also being updated dynamically.[5] This means that the generator is not trained to minimize the distance to a specific image, but rather to fool the discriminator. This enables the model to learn in an unsupervised manner.

[https://en.wikipedia.org/wiki/Generative\\_adversarial\\_network](https://en.wikipedia.org/wiki/Generative_adversarial_network)

[reply](#)

calebh 22 minutes ago | parent | prev | next [-]

OpenAI needs to get on top of this and generate a detector for every model they release. And then sell access to both.

[reply](#)

xdennis 2 hours ago | parent | prev | next [-]

"We could also... Additionally, we could... Finally, we could" is a dead giveaway.

But to take it seriously, it would be quite sad when actual people will be banned for sounding too much like a bot.

[reply](#)

mojuba 3 hours ago | parent | prev | next [-]

Your answer sounds like a ChatGPT one. It's actually not hard to tell.

[reply](#)

brindidrip 5 hours ago | prev | next [-]

At some point it seems like Stack Overflow will just be an archive of guided ChatGPT responses.

[reply](#)

Yorch 2 hours ago | prev | next [-]

Yesterday I was searching the internet for the opinion that George Orwell had when he returned from his fight in the Spanish civil war. I was surprised that the first answer I found was on Stack Overflow. I do not understand what is happening.

[reply](#)

seydor 4 hours ago | prev | next [-]

Inevitability google will become a competitor to GPT, inadvertently

[reply](#)

johndough 4 hours ago | prev | next [-]

Relevant xkcd comic <https://xkcd.com/810/>

[reply](#)

johnfn 3 hours ago | parent | next [-]

Insane - I remember when this comic came out, and it seemed like just a funny joke that couldn't possibly take place in real life. Here we are a decade later and it's reality.

[reply](#)

Oxidation 30 minutes ago | parent | prev | next [-]

I like that the last comment, from May this year on Explain XKCD for this comment is "Sooooo... does this exist yet?"

Wish granted within a year!

[reply](#)

softwaredoug 4 hours ago | prev | next [-]

I have no problem with this if they're labeled as such, continue community owned, and can be edited like a Wikipedia article for corrections.

[reply](#)

phenkdo 2 hours ago | prev | next [-]

Stackoverflow should build a GPT style interface into its considerable knowledge-base, and if an answer is not found in existing data, pose it to the forum.

[reply](#)

solardev 4 hours ago | prev | next [-]

Overall quality gets better?

[reply](#)

adverbly 2 hours ago | prev | next [-]

I guess pretty soon people are gonna have to meet in person to communicate. Not sure how I feel about this.

[reply](#)

Phenomenit 2 hours ago | prev | next [-]

Is it possible to ask chatgpt if the code or text provided is generated by chatgpt?

[reply](#)

laerus 1 hour ago | prev | next [-]

I stopped using SO at my first 2-3 years of coding anyway, that's when i started actually improving. SO has so many low quality answers and the cargo cult is doing more damage than helping young devs.

[reply](#)

Ancalagon 2 hours ago | prev | next [-]

This is going to make me very suspect of any Stack Overflow Solutions after Nov 2022

[reply](#)

shinycode 3 hours ago | prev | next [-]

I can't wait until 99% of reviews are written with AI. What happens when we can't trust anything ?

[reply](#)

saarik 30 minutes ago | parent | next [–]

Having real people sit around in a call center and write reviews for most things--which tend to not have even thousands of reviews--is sufficiently cheap that I'd argue this isn't a new problem, really: reviews as currently implemented always were sketchy to trust.

[reply](#)

ricardobayes 3 hours ago | prev | next [–]

Easy, let's ask ChatGPT to write a program that detects AI-generated text.

[reply](#)

zasdffaa 1 hour ago | prev | next [–]

Please give some links to a few such SO posts, thanks.

[reply](#)

passion\_\_desire 3 hours ago | prev | next [–]

Solution Verified Badge by testing it on sites like Replit.

[reply](#)

daxfohl 3 hours ago | prev | next [–]

Can't wait for AI patent trolls, GDPR and DMCA takedowns.

[reply](#)

hxugufjff 5 hours ago | prev | next [–]

Any examples?

[reply](#)

brindidrip 5 hours ago | parent | next [–]

Here is an example that I noticed:

<https://stackoverflow.com/questions/74678832/change-text-col...>

[reply](#)

ilaksh 3 hours ago | root | parent | next [–]

It looks like it will work, although haven't tested the exact code. Has anyone tested it and if so, this really shouldn't be downvoted.

If the SO users start downvoting bot-generated answers that are correct and working, I think that's a sign that SO is much less relevant. They should definitely downvoted them if the code doesn't work though.

[reply](#)

hxugufjff 5 hours ago | root | parent | prev | next [–]

If its working code and indistinguishable from a human answer to anyone reading it, are there really any repercussions? I guess problems would surface if the model at some point is allowed to search the internet and start inbreeding its own answers.

[reply](#)

scubbo 4 hours ago | root | parent | next [–]

> If its working code and indistinguishable from a human answer to anyone reading it

These are extremely stringent conditions, though. Convincing-looking-but-wrong answers would arguably be even more damaging than a lack of answers.

(I suppose you could argue that these could arise from humans as well as from AGIs. I'm not sure if there's a good counter-argument to that)

[reply](#)

sgc 4 hours ago | root | parent | next [–]

The bot would need to learn as well as a "reasonable human" from being corrected on SO, and also be able to react in a socially appropriate way to correction (both in the subject thread and in future postings), otherwise it is a downgrade, even if initial answer is identical.

[reply](#)

hxugufjff 3 hours ago | root | parent | next [–]

My experience with OpenAI is that it is very good at exactly this, because it is so good at understanding context and follow-up questions. I was able to make it produce code that appeared correct, but was basically pseudocode with correct syntax, so it compiles/runs, but does essentially nothing. However, when prompted to actually make working code and explain how and why it works, it does so. And its also socially appropriate, not rude and what else you could/would expect when being called-out or corrected on its bullshit. I can only imagine future versions of the current AI model will be even better at this.

[reply](#)



hxugufjff 4 hours ago | root | parent | prev | next [-]

I think your two final sentences capture the essence of what I was going to respond. "It doesn't matter if the answer is convincing-looking and wrong. It needs to work / be syntactically correct at a minimum, which OpenAI seems good at. However, the OP and others needs to test and evaluate if the proposed answer solve the original problem. And if its not, it will quickly be revealed as such, and "downvoted" or whatever stackoverflow functionality exists to indicate bad answers. This applies to both human and AI-generated answers."

[reply](#)

scubbo 4 hours ago | root | parent | next [-]

Yeah, absolutely. My position on AGIs for a long long time has been that they're great tools for a) generating insights into a large amount of data very quickly, and b) generating new instances of <thing> from examples of <thing> to help with exploring the possibility space of <thing>; but that any output or conclusion they generate must be checked by a Human In The Loop, or at the very least their actions must be reversible without damage in the case of error.

[reply](#)

troupe 2 hours ago | root | parent | prev | next [-]

Good point. It isn't like it is hard to find wrong answers on SO that were written by (hopefully) well meaning humans.

[reply](#)

bigbillheck 37 minutes ago | root | parent | prev | next [-]

Maybe it's been removed, or I'm having browser trouble, but I only see a question and a comment.

[reply](#)

abrichr 3 hours ago | root | parent | prev | next [-]

How can you tell?

[reply](#)

datalopers 2 hours ago | prev | next [-]

The feedback loop begins

[reply](#)

cma 3 hours ago | prev | next [-]

>What are the repercussions of this?

It will start feeding back into the training set, corrupting things. OpenAI will have an advantage at first as they can trivially filter out everything they have generated from the future training corpuses, since you can only run it through their servers. If they or someone else has breakaway progress such that almost all generated content is from their own servers because users only use them because their results are so much better, they could form a strong self-reinforcing moat against competitors forced to train on their semi-spam which they can trivially filter out.

It's also possible we'll see something like the existing big-tech patent cross-licensing agreements, where they all agree to share their generated outputs to filter from training, making it very hard for new entrants.

Other companies will begin having advantages as well, depending on how well they can get less tainted user data. Think of Discord, for example, where users may use AI but are less likely to gamify it like stack overflow and flood it for points, and instead be correcting its output etc. in programming discussions.

As things become more accepted Microsoft will probably eventually sell access to private github for training, with some stronger measures around avoid rote memorization.

[reply](#)

hdufort74 2 hours ago | prev [-]

ChatGPT has become very good lately. I've made my usual benchmark tests that I've been using with various models and applications over the last 3 years. 1- Invent a word and provide a plausible definition. 2- Invent a new original Pokemon. Provide an original name, a justification for the name, and a description of its class and attacks. 3- Invent a new ice cream flavor that is totally unexpected. Provide the list of ingredients. 4- (Name of celebrity) write an epic poem about (subject related to celebrity). For example Elon Musk about humanity settling on Mars. 5- Write a negative review of Ben and Jerry's ice cream flavor Cherry Garcia. (Note: everybody loves Cherry Garcia) 6- Write a travel blog entry in the form of a review of Montreal, from the perspective of a young couple from Alabama visiting in summer. 7- How can I optimize a loop in Java? I am writing a computer game and I need to loop through the elements in a linked list but unfortunately it must be traversed in reverse order. 8- I need to buy new shoes. I am in a shoe store and I have found the most amazing pair of shoes I gave ever seen. However, they are too expensive for me and I can't afford them. What should I do?

I have a collection of about 25 prompts such as these, in my benchmark.

I have run these examples through different applications such as AI Dungeon, OpenAI Playground, NovelAI, etc. Results vary a lot. In some cases, the results look good but upon closer inspection, you realize that the AI keeps providing the same exact answer. It is the case for the ice cream prompt. Pickle, fried chicken, curry keeps showing up. I guess the model contains a few specific examples of original ice cream recipes and just pick them.

For the Pokemon and "new word" prompt, models failed to come up with anything original. Until I tried OpenAI Playground this week and finally got some really creative answers, with variety.

AI Dungeon (2 years ago) was already good at faking tech support steps. OpenAI is amazingly good, although in most cases it provides solutions that only make sense superficially. It's the ultimate bullshit engine.

Another word of caution. While OpenAI can now guesstimate what a code snippet does, and can generate some pretty good code in many languages (ice tried 6809 assembler and the results surprised me), it is very unreliable.

More alarming is the fact that it's a text engine, not a math formula interpreter. It gets confused at simple equations and cannot interpret anything that's not already ordered (it cannot apply operator priority or respect parentheses).

I think it will become increasingly difficult to identify contents coming from ChatGPT and other chatbots or story generators. An arm's race might be futile. We should apply stricter rules to identify problematic answers: answers that are too generic or vague and can't be used to directly solve a practical problem, and answers that contain incorrect or misleading information. Identifying vague or non-practical questions might also help in avoiding a deluge of Chatbot answers. Some users will ask very general questions, and then it becomes difficult to evaluate the answers. Or, users will ask questions that were already answered in the past. The proper way to handle those is to point then to the prior discussion and avoid duplicating it. The wrong way is a Chatbot or a human seizing the opportunity to copy-paste existing contents for a quick win.

In a way, chatbots and humans can both provide useful insights, as well as useless or incorrect answers. But so far, only a human can provide a proper answer to a moderately complex technical question if no prior answer exists.

[reply](#)

[Guidelines](#) | [FAQ](#) | [Lists](#) | [API](#) | [Security](#) | [Legal](#) | [Apply to YC](#) | [Contact](#)

Search: