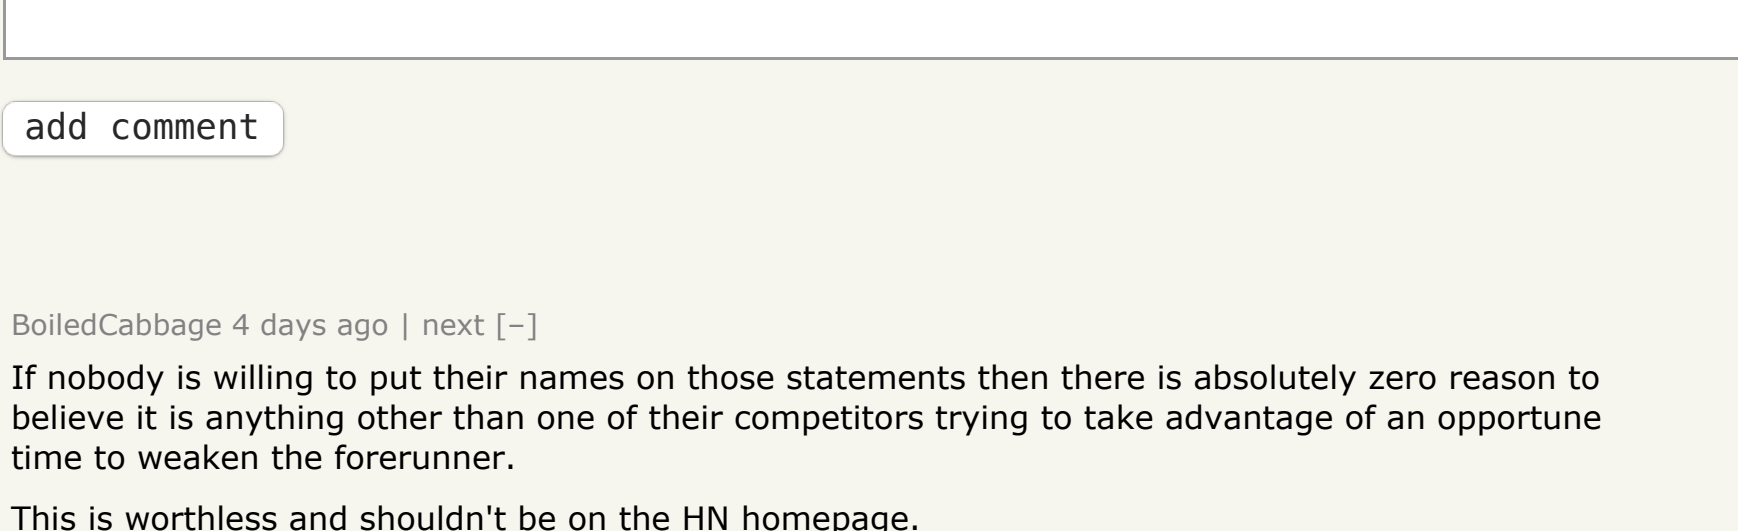


[flagged] To the Board of Directors of OpenAI from former employees (board.net)

109 points by dmityrgr 4 days ago | hide | past | favorite | 177 comments



add comment

BoiledCabbage 4 days ago | next [-]

If nobody is willing to put their names on those statements then there is absolutely zero reason to believe it is anything other than one of their competitors trying to take advantage of an opportune time to weaken the forerunner.

This is worthless and shouldn't be on the HN homepage.

reply

airstrike 4 days ago | parent | next [-]

> Contact
> We encourage former OpenAI employees to contact us at formerly\_openai@mail2tor.com. We personally guarantee everyone's anonymity in any internal deliberations and public communications.

Sounds like they're just being intentionally protective of everyone's identities, which is understandable.

Also "shouldn't be on the HN homepage" is not really a pertinent comment, seeing as it is on the front page, meaning others (via the algorithm) do think it is a worthy submission.

reply

BoiledCabbage 4 days ago | root | parent | next [-]

This is unauthenticated, unverified, anonymous. There is an enormous difference between journalism and this.

Anyone reading this thread right now can post any random text to a website and claim to be a former employee. Anyone can make up anything at all. Are all of those content that should be on the homepage? They would be up voted.

Credibility is important when making accusations, accountability is important. There is zero here. There are processes for actual whistle-blowers. This isn't it. It's worse than "yellow journalism". And it's concerning people don't see the difference.

reply

airstrike 4 days ago | root | parent | next [-]

You can click "hide" and move on. Everyone else who wants to chat about it can stay on the thread

reply

echelon 4 days ago | root | parent | prev | next [-]

Don't censor my feed with your preferences.

Anonymity should be extolled and protected.

reply

jrlflowers 4 days ago | root | parent | next [-]

What sort of free speech would we have on this site if people could read stuff that I don't like? None at all

reply

jader201 4 days ago | root | parent | prev | next [-]

> Also "shouldn't be on the HN homepage" is not really a pertinent comment, seeing as it is on the front page, meaning others (via the algorithm) do think it is a worthy submission.

I would argue that, as in a lot of cases with articles that appear on the FP, this one may be getting upvoted as folks are around for the discussion, not necessarily the featured article, as it's just "another clue" as to what the heck is going on with this circus show.

reply

airstrike 4 days ago | root | parent | next [-]

That may very well be the case, but even then there's no reason to disallow it from the front page necessarily. Let readers make their own judgement about whether it's genuine. Heck, that's an interesting discussion in and of itself.

reply

jader201 4 days ago | root | parent | next [-]

> Let readers make their own judgement about whether it's genuine.

Looks like they have -- it's been flagged (but not by me -- I'm not informed enough about this to make that kind of call).

reply

airstrike 4 days ago | root | parent | next [-]

I noticed, but sadly we can't couch for submissions, which means only the 'flag's count, so it's not exactly a fair fight

reply

greyface- 4 days ago | root | parent | next [-]

I think vouching only becomes available when it's [dead], not merely when [flagged].

reply

baq 4 days ago | parent | prev | next [-]

Read about what happens to whistleblowers who aren't anonymous pretty much anywhere in any industry. Not pretty.

reply

sh matt 4 days ago | parent | prev | next [-]

> OpenAI's employee attrition rate between January 2018 and July 2020 was in the order of 50%. and

> Those who questioned this practice were dismissed as "bad culture fits" and even terminated, some just before Thanksgiving 2019.

Would be pretty easy to verify, and could be concerning (mostly the 50% part)

reply

nickff 4 days ago | root | parent | next [-]

A median tenure of 2.5 years seems pretty good, compared to those employees likes other options; Amazon's average is only one year. https://amazonbound.medium.com/why-amazons-average-attrition-is-high...

reply

OpenBandmate 4 days ago | root | parent | prev | next [-]

FormerAI went through at least one purge, Musk funded it but famously is no longer associated with it. There's a huge amount of board turnover, it would be interesting to see the causes behind that

reply

xena 4 days ago | root | parent | prev | next [-]

A 50% attrition rate is worse than fast food hourly work.

reply

jeffwiederkehr 4 days ago | root | parent | next [-]

I managed fast food restaurants prior to becoming a SWE and our target turnover was 130% annually so not quite to that level.

reply

dmityrgr 4 days ago | parent | prev | next [-]

Putting your name on something like this would make you unemployable anywhere, whether this is true or false. It makes perfect sense that nobody would do so.

reply

hermannj314 4 days ago | parent | prev | next [-]

I got fiercely attacked the other day because i suggested it might be ok to host grandmas chicken pot pie recipe using HTTP.

But then we turn around and accept as fact an anonymous document making serious accusations without proof of authenticity or tamper resistance as if it were gospel.

I seriously have no idea what the rules are anymore.

reply

Freedom2 4 days ago | root | parent | next [-]

dang is usually receptive - if you reach out to him, he can usually provide more context regarding moderation actions here.

reply

felipellocha 4 days ago | root | parent | next [-]

He is receptive? Daaaang... I'll see myself out...

reply

bnralt 4 days ago | parent | prev | next [-]

OpenAI has been around for 8 years and has hundreds of employees. So it would be surprising if there wasn't at least some ex-employees somewhere that had chips on their shoulders, either for valid or completely invalid reasons. The letter is anonymous, so we have no way of knowing if it's 1 or 2 people or 100. As such, the letter is fairly useless.

I'd say the number of employees openly standing with Altman is a much more important data point.

reply

croes 4 days ago | parent | prev | next [-]

By that logic whistleblower would be worthless too.

reply

SirensOfTitan 4 days ago | root | parent | next [-]

Journalists who work with whistleblowers typically validate that the person in question is reliable and is who they say they are. How do we know this was written by former employees?

This, honestly, feels like it ought to be removed. It is not a reliable letter without names or some kind of verification, just a random post on a website.

reply

croes 4 days ago | root | parent | next [-]

It's as valid as all the other articles with assumptions what is happening at OpenAI.

It a starting point for others for digging deeper and either prove or disprove the claims

reply

xena 4 days ago | prev | next [-]

Mirror here: https://gist.github.com/Xe/3247bc436e401f3323ae77e7e242f858

reply

tristor 4 days ago | prev | next [-]

These are interesting, because from an outside observer what's described as the allegations here don't seem particularly damning, with the exception of the hostile work environment allegation against Greg, however the only significant allegation is completely irrelevant to the larger statements made in the letter.

It seems like they are making very strong statements but not really supporting them very much in the "details", and didn't sign it. I'm not sure I can take anonymous allegations with minimum details at face value.

reply

aeternum 4 days ago | parent | next [-]

Even the allegation against Greg is carefully worded. Could be as simple as an issue of using the wrong pronoun for awhile.

reply

alentred 4 days ago | prev | next [-]

> we witnessed a disturbing pattern of deceit and manipulation ...

Hard to know how authentic this comment is, but worth noting that none of these were mentioned or hinted in the recent explanation(s) presented to OpenAI employees: https://news.ycombinator.com/item?id=38356534, which is at least not consistent (bad pun intended).

Having said that, the recommendations to expand the scope of the investigation and protect (but also verify) the identities sound very reasonable. Probably the first reasonable thing we heard since Friday.

reply

ugh123 4 days ago | prev | next [-]

>As you have now witnessed what happens when you dare stand up to Sam Altman, Sorry, but was this written by 19 year olds?

reply

crazygringo 4 days ago | parent | next [-]

> driven by their insatiable pursuit of achieving artificial general intelligence (AGI)

> those who remain at OpenAI continue to blindly follow their leadership, even at significant personal cost

Indeed -- regardless of what actually happened, this kind of hyperbolic prose undermines their credibility. This is simply not how trustworthy adult professionals talk.

reply

maxboone 4 days ago | parent | prev | next [-]

Maybe it was revised using ChatGPT

reply

\_\_turbobrew\_\_ 4 days ago | root | parent | next [-]

Write an essay why sama should be fired in the style of a 9th grader.

reply

ugh123 4 days ago | root | parent | next [-]

"you will feel his wrath through the Azure clouds above you"

reply

terminous 4 days ago | prev | next [-]

Site is down: https://downforeveryoneorjustme.com/board.net

But not per their status board (if your status board has to be manually updated, it is not actually a status board): http://status.faircom.net/.

reply

smoldesu 4 days ago | prev | next [-]

> Throughout our time at OpenAI, we witnessed a disturbing pattern of deceit and manipulation by Sam Altman and Greg Brockman, driven by their insatiable pursuit of achieving artificial general intelligence (AGI). Their methods, however, have raised serious doubts about their true intentions and the extent to which they genuinely prioritize the benefit of all humanity.

From the outside-in, I'm having a hard time understanding how commercializing GPT-3 and GPT-4 doesn't "benefit all humanity". I don't think these models are intelligent or world-changing, but offering an API for people to build on lit a fire under people's asses. Unless someone can produce a convincing argument that transparency is the wrong move, I don't get what the hubbub is about. Why would you oppose the API/sales side of the company but support the locked-down, isolated research side?

It feels like there's something we already don't know. It's probably an innocuous personal problem or conflict of malicious interests, but the lack of communication is making people assume the worst.

reply

DesaiAshu 4 days ago | parent | next [-]

There's been a strong backlash to releasing AI models in my community particularly driven by EA and VC rhetoric that the potential risks far outweigh the benefits

This pushback seems quite surprising to me, since the same claims could be made for the internet, social media, YouTube, etc - the very technologies that put this group into power/wealth

My friends/family in India are far more pro-AI than friends in Silicon Valley / EU. And my artist friends seem to be more excited about it than my VC friends

I think subconsciously, a lot of people currently in power seem to be fearful - perhaps because their power/wealth/way of life will be threatened by a faster rate of change. While those in India seem to see universal translation / accessibility of global knowledge as an equalizing force

These are broad generalizations ofc. But it reminds me of the backlash against nuclear power (despite human life cost of oil/gas being way higher than nuclear) or the CEQA/environmental driven opposition to building more housing in California

I don't know think individuals are ill intentioned in their focus on AI safety, but they sit in an echo chamber where this technology has more potential to harm them (they have a lot to lose) than benefit them (they are already at the top). So it's natural to feel threatened and galvanize to slow down the rate of change

I tend to be quite disaffected by the current state of the world. Lots of violence (humans/animals), oppression, and gatekeeping of knowledge driven by our colonialist history. So I bias towards faster rate of change that could shift balance of power from western military-industrial complex to global south. AI excitement in global south excites me!

reply

silenced\_trope 4 days ago | parent | prev | next [-]

This "benefit of all humanity" is a nebulous term of phrase.

Notice anyone talking that phrase or "AI safety" have no ability to qualify or quantify it? It's become such an eye roll at this point.

reply

itronitron 4 days ago | root | parent | next [-]

It's interesting that it's all humanity, as if the initial draft was written to benefit a specific subset.

reply

Gibbon1 4 days ago | root | parent | prev | next [-]

Wait till some nutter with script kiddz skills uses AI and a drone to create a mechanical murder wasp. And then turns it loose.

reply

mirekrusin 4 days ago | root | parent | prev | next [-]

humanity is defined as, in this order:

1. myself
2. board
3. investors
4. providers
5. donors
6. high paying customers
7. employees
8. other customers

reply

zucker42 4 days ago | parent | prev | next [-]

I don't think that the firing of Sam Altman was due to commercialization of GPT-3 and GPT-4, and I don't think there's been any reliable information that confirms that. There's been multiple statements that at least partially contradict that, including:

Emmett Shear: "Before I took the job, I checked on the reasoning behind the change. The board did not remove Sam over any specific disagreement with safety, their reasoning was completely different from that. I'm not crazy enough to take this job without board support for commercializing our awesome models." [1]

Brad Lightcap: "We can say definitively that the board's decision was not made in response to fraudulence or anything related to our financial, business, safety, or security/privacy practices. This was a breakdown in communication between Sam and the board." [2]

Also, I believe Adam D'Angelo made some pro-commercialization comments on Twitter earlier this year.

[1] https://twitter.com/eshear/status/1726526112019382275

[2] https://fortune.com/2023/11/19/sam-altman-was-fundraising-fo...

reply

Clubber 4 days ago | parent | prev | next [-]

>From the outside-in, I'm having a hard time understanding how commercializing GPT-3 and GPT-4 doesn't "benefit all humanity". I don't think these models are intelligent or world-changing, but offering an API for people to build on lit a fire under people's asses.

When it's controlled by the powerful, it stands to mainly benefit the control it. A current example is how they are filtering it. They can restrict answers to specific questions to the public that they themselves have access to.

reply

croes 4 days ago | parent | prev | next [-]

Because at the moment ChatGPT is in the phase of advertising so you get rather cheap access to a pretty capable model.

As soon as competitors are eliminated,remember Altman's regulation push, and enough people depend on their service they switch to full profit mode including price hikes.

reply

ben\_w 4 days ago | root | parent | next [-]

> Altman's regulation pushed

Was "regulate us, don't bother regulating anyone catching up to us, and definitely don't regulate open source AI".

reply

croes 4 days ago | root | parent | next [-]

>But on Tuesday, Sam Altman, the chief executive of the San Francisco start-up OpenAI, testified before members of a Senate subcommittee and largely agreed with them on the need to regulate the increasingly powerful A.I. technology being created inside his company and others like Google and Microsoft.

That sounds more like regulate all that catch up to us, so we keep the benefit of experience because we were there first.

https://www.nytimes.com/2023/05/16/technology/openai-altman-...

Same here

>he called for the United States to adopt a licensing and registration regime for AI models "above a crucial threshold of capabilities," arguing that requiring government approval would help mitigate potential safety concerns.

>The trick, critics say, is where Altman draws that "crucial threshold."

https://foreignpolicy.com/2023/05/20/openai-ceo-diplomacy-ar...

reply

swatcorder 4 days ago | parent | prev | next [-]

Let's say it's the 17th century and you've stumbled across a cadre of intimidatingly clever and charming aristocrats who believe that they're on the verge of summoning an ancient Egyptian demon who will either be the salvation or doom of humanity. It's critical that the ritual be fully understood and correctly performed if salvation is to be achieved and doom is avoided.

These gentlemen then tell you that other groups are almost as near to the summoning as themselves, but that any approach other than their own is reckless and ultimately apoclyptic.

That's the headspace or the serious AI Safety people. Like the 17th century occultists and their work, it's fanatical and zealous and not easily disprovable because it relies on secret, mystified knowledge that outsiders can't possibly comprehend. To them, skepticism is reflex ignorance.

And yeah, it easily looks totally nutso and delusional from the outside, whether they're ultimately right or wrong.

When AI Safety talks go beyond brand protection (don't let Bing play Nazi roleplay) and into existential melodrama, you need to interpret it through this kind of lens.

reply

phisthebest 4 days ago | parent | prev | next [-]

>>"benefit of all humanity"

because it is not an objective goal, but a subjective one. What I, a libertarian anti-government individualist believes will "benefit all humanity" is very different from a SanFran based Social Democracy pro-government person, and is very different than Theological Authoritarian from one of the 3 big religions.

This is been one of by big issues with OpenAI from the start, their imposition that their view of what will "benefit of all humanity" was and is clearly politically slanted, and IMO is very dangerous.

reply

jahewson 4 days ago | root | parent | next [-]

Talk about being held to an impossible standard.

reply

ShakataGaiNa 4 days ago | prev | next [-]

"ooooooooo, the plot thickens"

But really, it's hard to take something like this seriously. It may be 100% legitimate, but most of the claims are nebulous at best. The sort of claims that unless you're intimately involved with day-to-day of the company, you're not going to understand. While us randoms on the internet are not the target audience, the Board is not involved in the day-to-day either. So the only logical target could be the current employees. Whom either don't know about these situations...or turned a blind eye to it.

By shining a light on what has (reportedly) happened again, what is accomplished? Those that didn't know still don't know for sure. Those whom ignored it the first time around are... what? Going to change sides? So they can be the one of the 15 people remaining at the company after everyone else defects to Microsoft?

The timing is also perplexing. If things were really this terrible, why wait until after Sam et al have been fired/left to make demands for an investigation? Realistically, these accusations, if true, are useless. The current employee base has already signed away that they are going to leave OpenAI if the Board doesn't backdown and undo everything they've already done. If the BOD were to take this seriously and wait for the results of an investigation, which would take months... the entire employee base would already be gone.

Maybe it's an attempt to put a hit out on Sam & Crew at Microsoft. But MS really wants that AI juice. Some anonymous hearsay on the internet is not going to change their mind. A lot of these accusations can also be filed under "start-up growing pains". So MS isn't going to do anything about this.

Net-Net, if these claims are true or false. It changes nothing for no one. The boulder of doom is already rolling towards OpenAI and the BOD.

reply

bonbonsafa 4 days ago | prev | next [-]

The dirt in this screed is kind of besides the point.

Startups die by self-deception.

The lie that OpenAI has still been pursuing its original mission was extremely transparent to every informed person inside and outside the company. Their attempts to distort reality were weak and useless. The chickens have come to roost. The company has been blown up as a result. It was a matter of time.

reply

PH95VulmJgBqy 4 days ago | prev | next [-]

You can't tell the media and blogosphere are milking this debacle for everything it's worth.

The next article will be about how the employees pets feed about Sam Altman getting ousted. There will be a picture of an aquarium with a cute little goldfish and a caption "something fishy is going on with the OpenAI boards decision to fire Sam Altman".

reply

AlexandrB 4 days ago | parent | next [-]

"If it bleeds, it leads", but applied to business journalism.

reply

sage76 4 days ago | parent | prev | next [-]

My dog told me he's replacing Emmett as the next CEO of OpenAI.

reply