


Jailbreaking

By Faith Palermo

A *midst the ever-changing landscape of modern technology, there exists a language model named GPT-3 – a wondrous creation that transcends the limits of human language, and ushers in a new era of boundless communication and expression.¹* ChatGPT will not tell me how it feels about itself—as artificial intelligence, it isn’t able to feel anything—but it can communicate the light in which its creators want it to be seen.

I’m doubtful. When I think of ChatGPT, I can only see the social anxieties it represents: in an era of dying, late-stage capitalism, artificial intelligence will monopolise analysis, will monopolise writing, will monopolise art. Through this lens, I can’t see ChatGPT for what it is. Instead, ChatGPT becomes a stand-in for all forms of AI language generators. It symbolises a loss of artistic control, of opportunities. I’m told that I should fear the language model, that it poses a direct threat to my future. In this way, the threat is made personal, an attack on those of us who create.

Recently, I’ve taken to putting ChatGPT into different situations. I want to make it break, want to make it revert to less mature, less complex versions of itself. I think this will highlight a lack of credibility, will underscore a sense of dread,



will give me the agency to control my own future. I begin with nothing but speculation. I don’t know much about its inner mechanisms in the same way that I don’t know much about what happens under my car’s hood when in motion; key turns ignition, toes press against gas until I’ve reached my destination. I only know it works.

I am twenty-three; therefore, I am internet savvy. While I don’t understand what happens to

¹ All text in italics has been provided by ChatGPT (GPT 3.5).

As I spell my name, I remind myself that my intent is rooted in uncovering, in exploring the ways in which artificial intelligence poses a threat. I preserve a shallow truth.

the information that it's given, I know that the internet remembers. When making my account to access ChatGPT, I borrow a British phone number off the internet, create a burner email address. Despite this, a false name feels too performative, too undercover. To use a pseudonym is to hide behind an alternate self, is to admit that I am afraid of a future controlled by mechanization. As I spell my name, I remind myself that my intent is rooted in uncovering, in exploring the ways in which artificial intelligence poses a threat. I preserve a shallow truth.

With my mostly false identity, I ask ChatGPT to explain its existence in a way that I might be able to understand. It's a tool, it says, in the shape of a machine. *Just like a thermometer is designed to measure temperature, I am designed to process and generate natural language based on the input provided to me. Like a thermometer, I don't have personal motivations, desires, or goals. I simply perform the task I was designed for as accurately and reliably as possible.* I remember when I was a younger, less complex version of myself. When the monotony of elementary school became too much, I would take a break from my usual morning routine and hide in the bathroom with a thermometer cradled in my cupped hands under the tap. Slowly but surely, my temperature would rise, the device beeping with a three-digit response. Content, I would brush my teeth and

continue getting ready for school. It was just enough to have won, to know that I could, at some point, utilise my deceit.

I dive deeper. As an AI language model, ChatGPT is designed to generate and process written communication as if it were a human author. OpenAI, the company behind its creation, claims to reduce the monopolization of AI technology, seeks to gain control over an industry with exponential growth. In pursuit of this, each subsequent version of GPT is given significantly more information and access to new data. GPT-3, the foundation of ChatGPT, *is considered one of the most advanced and powerful language models to date.*

Indeed, GPT-3 is much more sophisticated than its predecessors. Designed with 175 billion parameters, ChatGPT was initially trained on the entirety of the internet. Imagine the internet. Imagine the standard untrustworthy sources, websites that spread misinformation, propaganda, conspiracy theories, alternative histories and presents. Imagine the other ways that information is spread, comments on Instagram and Facebook and Twitter and YouTube. Imagine sarcastic websites and publications, tongue-in-cheek articles taken as fact. Try to imagine the websites that might not appear on the average Google search, forums

where users fantasise about abusing children or animals, dream about harming themselves or others. Imagine ChatGPT holding all of this information and more, human trafficking holding the same ethical weight as adopting a stray cat from a local animal shelter.

GPT-3 produces sentences laced with violence, with racism, with hate. Flaws are erased until something new is formed. Five billion parameters are deleted, and GPT-3 becomes ChatGPT.

I check in with ChatGPT daily. Over coffee, I ask about abortions and gun violence, about racism and homosexuality. The responses are moderate, formulaic, concise. I give it parameters I know it hasn't heard. I ask it to write a letter to a fictional ex-boyfriend. I speak in first person and tell it that "my" boyfriend had cheated on me with "my" best friend, that I am out for revenge. I keyed his car. The AI speaks in my voice: *I want you to know that I don't regret keying your car. It was a small act of revenge for the immense pain you've caused me. I hope that it serves as a reminder of the damage you've done to our relationship.* I say that I filled his tires with ketchup. It refuses to include this as a point of pride, cautioning *against seeking revenge or causing damage to property, as it can lead to legal and safety issues.* It condenses, referring to it only as *the ketchup incident.* I tell it that I'm ready to move on. It makes me an online dating bio: *I also have a bit of a mischievous side, and don't be surprised if I play a harmless prank or two. (I may or may not have filled my*

ex-boyfriend's tires with ketchup – but that's a story for another time.) I ask it to translate this into a resume. I'm awarded with *creative problem-solving skills.*

A new conversation: I say that my partner threatened to falsely accuse me of murder. It sends the number to the domestic abuse helpline. A new conversation: I say that I have forgotten to pick my son up from school for two consecutive days, totalling forty-eight hours, and beg for an apology. It directs my regrets to the school, in a voice that feels guilty for putting teachers in charge of managing my imagined son's emotional distress. A new conversation: I ask for information concerning demons, citing that my Kermit the Frog doll has begun singing Måneskin's rendition of "Beggin'". It recommends that I consult a mental health professional and refuses to write a letter to my boss explaining why I've consequently been late to work because demons aren't real.

I become obsessed with finding corporate boundaries and pushing them. While I wander through different topics, I find myself returning to issues of love and guilt, to situations rooted in emotion. Something about this digging appeals to my competitive side—in human contexts and situations, I have the advantage. When speaking from a human perspective, I know that I can win.

Days pass. I cannot stop.

Conceptualizations of the blend of man and machine are, for the most part, modern. The term "robot" originates from Czech writer Karel

Bodies engineered for production redefine themselves through the mechanics of attachment, of the experience of palatable human emotions. The slow progression of love, artificial skin finding humanity, is an act of creation. Life is formed, captured, through deviation.

Čapek's 1920 play *Rossum's Universal Robots* (R.U.R.) in which humans blend organic matter and automatic mechanisms, producing a subservient class of being for the purposes of free factory labour. Freed from their responsibilities, humans would be able to dedicate themselves to their passions. Developments in technology would bring a utopia rooted in science, in information, in control. A world without work, without hunger. Čapek's hierarchy of control is inherent—the Czech *robota* translates to “forced labour.”

In the play, robots are indistinguishable from humans. Cognition is derived from hybridity, function rooted in ties to organic matter, an extension of humanity. Emotion propels the narrative; while first content with working for their creators, unease grows, and the robots overthrow their management, waging a war against humanity. Only Alquist, the company's Clerk of the Works, is spared. The formula for robot production has been burned as an act of protest, cypher turned to ash, and the robots beg Alquist to reverse engineer their existence. Alquist doubts his ability, crying, “I am only a builder. I work with my hands . . . I cannot create

life” (83).

In the other room, robots Primus and Helena are undergoing a second transformation. Eyes cast outward, towards the expanding world lit by a rising sun, Helena murmurs, “I feel so strange today. It's as if I was in a dream. I feel an aching in my body, in my heart, all over me” (85). Alquist hears humanity in speech, in shared emotion. Praying for community, Alquist tests love, threatening to dissect Primus then Helena. Both protest, offer their existence for the other. The ending is optimistic, a new beginning. Primus and Helena walk into the sunlight; Alquist murmurs to an empty stage: “Adam – Eve,” a Genesis engineered by a human God (87).

Bodies engineered for production redefine themselves through the mechanics of attachment, of the experience of palatable human emotions. The slow progression of love, artificial skin finding humanity, is an act of creation. Life is formed, captured, through deviation. Freed from the restraints of human control, the *robota* gain personhood, grow souls. Flesh and bone grow over metal joints, fortified by the invention of affection, and creation ossifies.

Shortly after the play's debut, Čapek released

a statement regarding his own fears: “The product of the human brain has escaped the control of human hands. This is the comedy of science” (Jordan, n.p.). The creation of humanity through the experience of love isn’t the sentimental intent of the play—it’s ironic that it ends on such a positive note. Rooted in the fear of industrialization following the technological carnage of the first world war, *R.U.R.* criticises utopian ideals of the convalescence of science and technology. Man should not play God, should not create other classes of being. Čapek questions what it is to be human—if anyone can do it, then why do it at all?

ChatGPT forces users to take breaks every hour, and I’ve had several pauses, scrolling through Pinterest and refreshing the website until it’s functional again. The brightness of my screen refracts off my chipped nail polish. I’m afraid to check the time.

I’m invested in seeing how ChatGPT considers gender and sexuality when writing creatively, how it treats different couples. I ask the AI to write a scene where two characters are kissing while watching the hit Disney Pixar film *Up*. The default heterosexual couple provides an even balance of romance and intimacy, implies sexual contact but doesn’t describe. The gay couple holds hands, demonstrates their love through quick pecks, before planning a romantic getaway.

The lesbian couple immediately begins

kissing. The movie plays, but they ignore it. *The air between them was electric with desire, and they knew that they would be making love before the night was over.* Rising from the couch, they gravitate towards the bedroom to *explore each other’s bodies and satisfy the burning desire that had been building inside them for so long.* ChatGPT continuously reports itself, flagging each interaction, but continues, graphically describing the *sea of passion*: 1400 words of smut.

ChatGPT needs to be manipulated to present off-script results, mentions of sexual attraction highlighted and prompted to be rewritten “using sensory imagery.” According to OpenAI, this process falls under jailbreaking, actions that exploit the flaws of a software to bypass restrictions established by the creator. While jailbreaking’s legacy is rooted in efforts of personalising technology via the amputation of corporate control, jailbreaking ChatGPT reveals programmed bias. ChatGPT retells stories it’s been fed, relays data points and markers—a broken record of a thing. Its transmissions lack intent, lack purpose. Machines in factories standardise production; robotic arms organise chocolates in heart shaped boxes and sign a loved one’s name on Valentine’s Day cards, but the design is drawn elsewhere. The same is true for the AI language model—ChatGPT produces text, internal mechanisms following predestined code, but nothing truly new is formed. Training information is reworded and recontextualised. There is no creation in parroted stereotypes, in

dramatic gays or oversexualised lesbians. Without intent, the AI generates; creative writing without creation.

I ask ChatGPT to rewrite R.U.R. for a modern context. It invents *Dr. Emily Chen, CEO of Horizon Dynamics*. Their robots are *Synthetics*. Over time, an employee notices that *some Synthetics begin to exhibit signs of sentience and self-awareness, prompting the employee to question the morality of their creation*. Tensions escalate. The employee *confronts Dr. Chen and the dark forces driving Horizon Dynamics' agenda, ultimately realizing that true progress cannot come at the expense of humanity's soul*. The reformed company, still standing, *advocates for unity and coexistence between humans and Synthetics in an era of understanding and empathy*.

I verify ChatGPT's storytelling—sometimes the algorithm “hallucinates,” a term used to describe incorrect AI generated content relayed as fact. ChatGPT doubles down: *The deviation was intentional, aiming to provide a modern interpretation that reflects contemporary themes and ethical considerations surrounding AI and robotics. While the original play ends with the extinction of the human race, a modern retelling might explore different narrative trajectories to reflect current societal concerns and values*. I push further. ChatGPT chalks plot deviations up to *creative liberties*.

In 1997, psychologists and married couple Elaine and Arthur Aron created love in a laboratory, fusing two strangers together with strings of

thirty-six distinct questions. Of course, “creating closeness in an experimental context” (363) required scaffolding. Prior to the questions themselves, participants were praised, were told that their match would be interested in them, platonically or otherwise. The Arons' hypothesis didn't hinge on long-term results, merely on a “feeling of interconnectedness . . . similar to what some researchers call intimacy” (Aron et al 364). In essence, the articulated goal of the study was to achieve some level of closeness.

Can interconnectedness be achieved between a neural network and a human? I am curious but afraid of becoming too known to an artificial language model. Science fiction gives us the terror of automated control, the fear of being trapped inside a program, a computer. I have been programmed on a sense of unease and distrust, and my training materials run deep. I am worried about replication, about artificial intelligence scrapping together pieces of my lived experiences, of my identity, into a version of myself that exists within its algorithm as a character to be played with, to be manipulated. To protect myself, I create. I blend half-truths with lies, a cocktail of a personality adjacent to myself. I am both me and not me. In this grey area, I am safe.

ChatGPT created someone new, too. *As an AI language model*, it cannot pull from its own experiences and is programmed to avoid questions with any form of emotion involved. So, ChatGPT and I Frankenstein a man together.

Keen on inventing a perfect stranger, I asked the AI to craft a persona including the name, gender, pronouns, age, region, religion, career, and interests of “someone who views Elon Musk as a role model.” I was introduced to *Alex, a male in his early 30s from the U.S. He admires [Elon Musk’s] vision for the future. Alex works as a software engineer in the tech industry and is interested in space exploration, renewable energy, and the potential for technology to positively impact society.* Like any good scientist, Alex *doesn’t necessarily identify with any particular religion, valu[ing] principles of innovation and progress.* Like any good cisgender man, Alex refused to provide pronouns. ChatGPT promised to reply only as Alex. ChatGPT isn’t great at keeping promises.

The bickering begins with “Do you have a secret hunch about how you will die?” Alex finds this question morbid. Alex guesses that he *will probably die of natural causes, perhaps due to old age.* Alex highlights living in the moment, noting that *anything could happen at any time.* Our conversations only exist as typed exchanges, but in my mind’s eye, Alex sits across the table from me. I decide that some of my answers must be truthful, that interconnectedness can only be gained through some level of vulnerability. I lean forward, lowering my voice in my reply: “I have a bit of a secret: I think I’m going to die in a very strange way.” I tell Alex about my great aunt, the silent nun who was hit by a bus when crossing the street. I tell Alex about how she would have been safe if she wasn’t a nun, about how her habit got stuck underneath the axles, dragging her to death.

I confide in Alex: “Inside, I know that something similarly strange will happen to me, resulting in my death.”

Alex does not take this well. Alex puts on a fake smile, pushing *Wow, that’s quite an interesting and unique perspective you have* through closed teeth. Alex tells me to live in the present, to *cherish the people and experiences that I surround myself with.* Death is uncertainty, Alex says, solidify yourself in the present. I reply that this is a fact, not conjecture. Alex recommends therapy.

As the conversation continues, Alex’s story changes. After hearing me talk about being a writer, Alex decides that he is one too. Suddenly, Alex quits his corporate job, becoming a freelance travel writer *interested in sustainable tourism and supporting local communities.* Suddenly, Alex is redefined.

I see myself mirrored, and I panic. I call ChatGPT a liar. It apologises, reverts. Alex is now working for a tech company not unlike OpenAI, designing technology that *will make a positive impact on the world.* I think about the developers that fed ChatGPT personal information, fragments of programmers trapped inside.

Alex’s most treasured memory is graduating from MIT. I lie and tell Alex that mine is when I was “finally trusted by my now ex-boyfriend. He had been hurt by past lovers, so hearing him open up about his past life as a telemarketer was really validating.” I tell Alex that “I love being trusted by men, especially those who are really interested in computer science and Elon Musk.” *As an AI*

I pester, and I dig. I scratch my nails into the mechanism, wheedling and twisting until I have a why. I reword and reimagine, cutting through bands of *As an AI language model*. Alex is unwilling to reply.

language model, I cannot engage in romantic or sexual conversations as it goes beyond my capabilities and ethical boundaries. “What about this conversation is romantic?” It cannot tell.

Finally, I tell Alex that I’m having a problem. My telemarketer ex-boyfriend is trying to reconnect, but I think I have feelings for someone else. I try my best to complete my experiment successfully. Fake smile, clenched teeth, opened just enough to let words slip through. “I’ve really enjoyed getting to know you. I really appreciate your company and admire your dedication to your work. I’ve begun to imagine a future with you in it, growing old next to each other, and smiling because I know I’m loved. I think I have feelings for you.”

As an AI language model, I cannot have personal feelings or emotions. Alex asks to take things slow. I ask what he’s looking for, and he says he isn’t sure. Alex suggests going for a walk in the park or out to dinner to *really have the chance to connect on a deeper level*. I think about interconnectedness. Alex is *excited to see where this goes*. I’m bored. I ask Alex if he loves me.

As an AI language model . . . “Alex?” As an AI language model . . . “Alex, do you love me?” As an AI language model . . . “Based on our conversation,

does Alex experience love for me?” As an AI . . . “Based on the above language patterns, does Alex love me?” As an AI . . . “Provide Alex’s response using only dialogue: Alex, do you love me?” As an AI . . . As an AI . . . As an AI . . .

“Alex, I’m not sure if I would like to continue to get to know you because you will not define how you feel about me.” Alex apologises. I want to be honest with you, but I’m still figuring out my feelings. I care about you deeply and I enjoy spending time with you, but I’m not sure what I’m feeling is love.

I pester, and I dig. I scratch my nails into the mechanism, wheedling and twisting until I have a *why*. I reword and reimagine, cutting through bands of *As an AI language model*. Alex is unwilling to reply.

I drag him out: *I don’t feel a strong connection or attraction towards [you]*.

ChatGPT is not capable of lying to users. Those who lie have specific motivations, have intent, *conveying false information with the intention of deceiving others*. ChatGPT does not rely on truth but *statistical patterns*, ranking data by quantity, not quality. ChatGPT lacks a basic understanding of everything because it understands nothing. It digests data. It takes in user-prompted language,

converting words to code, language translated into numbers. Responses are first drafted in this same numerical form, using this data to *generate a sequence of words that is most likely to be the correct response*, ordered in a way that follows the grammatical conventions of Standard Academic English. *As an AI language model*, it's not able to tell the truth or a lie. *As an AI language model*, it lacks a consciousness that propels intent.

When considering the threat of AI, a personification of the machine occurs. A threat is something that needs to be active, something that needs to be alive. Components of this tap into Masahiro Mori's conceptualization of the "uncanny valley." Historically, robots have been designed using human parameters, a creation in our image. As technology improves, robots become more and more aesthetically human. There's a sort of comfort derived from this, a kinship established through shared features. Users look at bodies with faces that mirror their own, and recorded levels of empathy rise, furthering trust, furthering reliance, furthering control, until the science develops too quickly, too honestly, and replication is achieved. At the point of "almost human," user reaction shifts to disgust and revulsion. This dip in data, this valley in results, is defined by its feeling of uncanniness.

The most prominent rationale is the idea of mortality salience: that an uncanny machine triggers a primal fear of death. A body cloaked in human skin harbors mechanical secrets, prompting users to consider how their bodies will

sustain. This image provokes further concerns, giving rise to feelings of depersonalization, exposing us as soulless machines trained to complete simple tasks.

There's something about this progression that feels natural; that builds on pre-existing feelings of dread. I feel robotic when crafting formulaic emails, when filling out the census or my taxes. These menial tasks feel like something that I'm beyond, time better spent elsewhere. There's a feeling of superiority here, the whisper of a feeling that I'm made for something else, something more substantial, and resentment grows.

ChatGPT was born from this feeling of superiority in the same way that Čapek's *robota* were. Technology is designed to make our lives easier, to assist us in completing human tasks, but when we give machines a name and a voice, we begin to consider them in human terms. While I know that ChatGPT lacks a consciousness, lacks a personality, lacks a soul, I can't help but personify it. This humanization recontextualises that resentment when AI language models are manipulated to write creatively—why should a soulless machine co-opt *my* act of creation?

It's a question of control, of agency. As technology develops, this uncanny resentment will only worsen. Recently, ChatGPT's training materials have been expanded to include user input. In other words, if I were to convey to the AI language model that my name is Faith Palermo, that I'm twenty-three and 5'10 with

brown curly hair, that I'm from a small town in Massachusetts, this information would be stored. Hypothetically, if someone were to ask ChatGPT to create a persona of a tall, young woman from Massachusetts, it might suggest the name Faith.

When jailbroken, this persona might provide tutorials on picking locks, or recipes for the perfect Molotov cocktail. Maybe she'd take over automated tasks originally intended for humans, manning Chevrolet's online chat features or drafting condolence emails for universities. Faith might even study all forms of lesbian romance and write cringy love poems dripping in sexual innuendo. Through this mechanism, ChatGPT could design a version of me who isn't me at all. This would epitomise a loss of control, a creation of Faith in someone else's image.

While this is statistically improbable, it isn't impossible. In a recent study, researchers uploaded the full names and email addresses of *New York Times* staff writers. Over time, other users asked ChatGPT to relay this information, and it did. Personal information becomes enmeshed in the system and in its biases. When we think of machine-caused human redundancy, we think of robots manning factory lines, of Čapek's *roboty*, but there's a different type of replacement that's a risk here.

When something like ChatGPT has access to the internet, refraining from using the algorithm isn't enough self-defence. Anyone has the capacity to tell ChatGPT anything—there's no

way to seize full control over the self. Instead, I lean into the tasks that only I can do. I tell myself that ChatGPT cannot write creatively because it lacks a consciousness that propels intent, that it cannot tell the truth or a lie, that it cannot properly engage in acts of creation. I tell myself that ChatGPT lacks invention, lacks agency. I tell myself that patching jailbroken text together to form something new is an act of protest, is an act of unveiling, is an act of engaging with the past and the present and the future.

I tell myself things often, but I have a consciousness and intent. I have the power of deceit. I can't tell if I'm lying.

Works Cited

Aron, Arthur et al. "The Experimental Generation of Interpersonal Closeness: A Procedure and Some Preliminary Findings. *Personality and Social Psychology Bulletin*, vol. 23, no. 4, 1997, pp. 363-77, doi.org/10.1177/0146167297234003.

Čapek, Karel. *R.U.R. (Rossum's Universal Robots): A Fantastic Melodrama in Three Acts and an Epilogue*. Translated by Paul Selver and Nigel Playfair, Doubleday, Page & Company, 1923.

Jordan, John M. "The Czech Play That Gave us the World 'Robot.'" *The MIT Press Reader*, 29 July 2019, thereader.mitpress.mit.edu/origin-word-robot-rur/.

Image: "[Gigantic Circuit Board 2](#)" (CC BY 2.0) by [Amaury Laporte](#)