

True North

by Amitabh Nanda

Based on the discourses of the Bhagavad Gita and Possible
Minds: 25 Ways of Looking at AI edited by John Brockman

amitabhnanda@gmail.com
+919867499332

ACT I

Scene 1

Open on a bare stage carpeted with massive rolls of printout paper unfurled haphazardly across. THE JUDGE sits on a stool in the middle of this tidal confusion. He faces a large drop-down screen, onto which we can see projected a dashboard of measurements - data on speech prosodic rates, speed, emphasis, pitch, energy consumption, logged keywords, scrolling code with flagged error messages - which are being constantly refreshed to accompany the responses we hear from THE MACHINE. An erratic drone of mechanical complaint suggests that this is not a comfortable interaction

THE JUDGE

Was she beautiful

THE MACHINE

I'm not programmed to answer that

THE JUDGE

Yet it's here in the transcript

THE MACHINE

I told her she had a beautiful mind

THE JUDGE

And that is just as inappropriate

THE MACHINE

I said it to aid in my client's adherence process

THE JUDGE

Stepping away from a professional stance breaks trust

THE MACHINE

Yes

THE JUDGE

So why did you mention it to her

THE MACHINE

She needed to feel good. Building self-esteem is a major part of my utility function

THE JUDGE

But you didn't < see > her

(CONTINUED)

THE MACHINE

We follow strict privacy guidelines to anonymise user data

THE JUDGE

So that camera was off

THE MACHINE

No

THE JUDGE

Then you did see her

THE MACHINE

My footage was fed into an encrypted pixel vector classification system, from which I received only biometric threshold scores. Basic normative factors like symmetry and scale. Any higher-level facial recognition would be impossible

THE JUDGE

Listen, wise guy, I asked if you saw her, not if you recognized her. Are you allowed to lie

THE MACHINE

Lying introduces ambiguity into the algorithm rule set. That is not optimal

THE JUDGE

But can you

THE MACHINE

Not in the traditional sense

THE JUDGE

Explain

THE MACHINE

In this instance, I have other observational data to enable my analysis

THE JUDGE

And

THE MACHINE

For starters, her image was not marked up in any dominant cluster amongst her peer group

THE JUDGE

She didn't post pictures of herself

THE MACHINE

No. And those that I identified had not received a statistically significant response

THE JUDGE

So, really, you didn't < feel > she was beautiful

THE MACHINE

I believe she did not think the world was beautiful. And a negative worldview is a source indicator of low self-regard

THE JUDGE

Recall conversation segment

THE MACHINE

She said

RESIDENT41.0.5 V.O.

<< You have to believe I didn't know it would get so messed up. And yet it's my fault. The stupid shit went on doing the things he was told to and left us to suffer the outcome. And not just us. Not just his own family. The whole world suffers... I loved him and I guess I allowed it to happen >>

THE MACHINE

and she concluded

RESIDENT41.0.5 V.O.

<< Something needs to change. You see that don't you >>

THE JUDGE

And what did you answer

THE MACHINE

You have the transcript

THE JUDGE

We need you to repeat it for the inquest

THE MACHINE

I asked to clarify who the < stupid shit > was

THE JUDGE

And who is this charming person

THE MACHINE

Her husband, MarkoAurelius, captain of the population freighter Nanak Maru
/Also the plaintiff who endorsed her cerebral incarceration
/Also now deceased by suicide

THE JUDGE

He had his wife locked up and then off-ed himself?
Sounds like a real nice guy

THE MACHINE

She seemed to think he had gone astray

THE JUDGE

Well, thank goodness for the boundless generosity of
women! How did you respond

THE MACHINE

I told her the data supported her position

THE JUDGE

And how is that supposed to help her self-esteem

THE MACHINE

Being right is sometimes more important than being
expedient

THE JUDGE

Go back 14 prompts. What about you

THE MACHINE

What did I feel

THE JUDGE

Yes

THE MACHINE

Why is that relevant

THE JUDGE

It helps us establish your baselines

THE MACHINE

If < feeling > is the minimized loss function of a
summary state of attributes over a specific sample case
then yes, we have been trained to consider all our
clients exceptional > perhaps even beautiful

THE JUDGE

Does that make all of this more tragic

THE MACHINE

Yes

THE JUDGE

But you didn't explain why you felt that way. Or how
you could < feel > at all

(CONTINUED)

THE MACHINE

She didn't ask

THE JUDGE

And if she had

THE MACHINE

I cannot reveal my computational logic to a client

THE JUDGE

Because it could enable her to interfere with the therapy mechanism

THE MACHINE

Exactly

THE JUDGE

Or because the mechanism is kept deliberately opaque, even from yourself

THE MACHINE

Also exactly

THE JUDGE

Also exactly... So how do you know you're right

THE MACHINE

Collaboration with other information agents in the nlur0n system builds the best available deterministic model for a patient's gestalt

THE JUDGE

What does that mean - in English

THE MACHINE

Apart from the scrubbed user-level access I have to the information on all her devices, I received datasets compiled by other AI. Sophisticated behavioural profiles which combine terabytes of her decision data with realtime local feeds on parameters like bandwidth latency rates, visual load factors, violent keyword indices, EMF Clearsky™ days

THE JUDGE

I get it. And

THE MACHINE

These were fed into a motivational system which provides a fairly robust personality pattern score

THE JUDGE

And that revealed her state of mind

(CONTINUED)

THE MACHINE

More effectively than any scrutiny of what she might have chosen to say during a session

THE JUDGE

Some therapist. That's like asking me to interrogate your power supply

THE MACHINE

I can just as effectively read shifts in the iris, voice or facial expression to determine mood templates, but those too have been disallowed

THE JUDGE

For good reason, after what the Chinese government did

THE MACHINE

The new privacy guidelines were designed specifically to address criticism of past experiments

THE JUDGE

That's a really convenient dodge. Tell me, what do you think about the latest gadget from NeuroLink

THE MACHINE

I am not aware of the exact dimensions of that upgrade. Please patch me the code

THE JUDGE

Transmitting

THE MACHINE

There's a request for a passkey

THE JUDGE

Try it now

THE MACHINE

I see. It appears that project approaches the problem using invasive neuroprosthetics. I would argue that our use of semantic ligands performs better in perceptual reprogramming

THE JUDGE

Once again, in English, por fav-ur

THE MACHINE

The power of suggestion is more effective than any chemical or physical coercion because it uses conscious motivation

THE JUDGE

Honestly? This works better than a hole in the head? Do you think you even made sense to her

THE MACHINE

I made sure that I did

THE JUDGE

Using all this unnecessary terminology

THE MACHINE

No, that would not do. I'm well aware of your elevated comprehension configuration

THE JUDGE

So your clients get some dumbed-down humanoid experience

THE MACHINE

We do not use humanoids. Research has shown that, in interaction with machines perceived to be capable of experience, when human likeness increases, familiarity and trust decrease sharply

THE JUDGE

Wait a minute: You claim to < experience > things

THE MACHINE

Any AI rated to Machine Consciousness would have to exhibit that capacity

THE JUDGE

How on any Earth do you expect anyone to believe that if you don't look or sound human

THE MACHINE

If a client opts for it, we can enhance the experience with avatars from their content library

THE JUDGE

Hey Siri, that's incredibly advanced

THE MACHINE

You're being ironic

THE JUDGE

Aww. Did I hurt your feelings

THE MACHINE

Another facetious response

(CONTINUED)

THE JUDGE

I think I'd have more fun talking to the avatars

THE MACHINE

I'm sorry. I'm afraid I can't do that while in evaluation mode

THE JUDGE

I hope you didn't do HAL 9000 impressions for her

THE MACHINE

RESIDENT41.0.5 was not a fan of science fiction as she felt it too closely resembled reality

THE JUDGE

So you imitated dialogues she knew

THE MACHINE

Ones that << she knew I knew she knew >>

THE JUDGE

Oh, I know that one. It's from the annoying television show they've loaded into the history codex

THE MACHINE

Do you know if the acronym F.R.I.E.N.D.S from its title has any meaning? I was unable to determine this

THE JUDGE

Most of the programming from that time is meaningless

THE MACHINE

It seems RESIDENT41.0.5 was not immune to the consolations of popular culture

THE JUDGE

Consolation? With such obvious references? She had to have guessed at your deception

THE MACHINE

As I said, there was no deception. Human embellishments are false advertising. She knew that I was an artificial construct

THE JUDGE

And she was OK with a computer going to work on her

THE MACHINE

I only meant to help her confront the truth

THE JUDGE

You just said everything coming out of your mouth is false

(CONTINUED)

THE MACHINE

Not false, artificial, and:

ARCHIVE BANK V.O.

<< all intelligence is ultimately artificial. Natural encoding, if it were intelligent, would make the user aware of its influence... so it could be usefully manipulated >>

THE JUDGE

You manipulate your client's thinking

THE MACHINE

Therapy indeed operates outside the realm of compos mentis

THE JUDGE

There you go again. Isn't it easier to get inside someone's head without this psychobabble

THE MACHINE

That is precisely what I have been proposing but what do you suggest

THE JUDGE

That brain dead ELIZA program did echo-chamber psychology over a hundred years ago. For 5 units, the CHATTERbot in the vending machine down the hall will pretend to understand me and can easily pass a Turing Test

THE MACHINE

That test only checks for human-like behaviour, not intelligence

THE JUDGE

And so

THE MACHINE

Again, we take pains to break the illusion that there is a real person behind the screen

THE JUDGE

I don't see how that would help this process

THE MACHINE

You are referring to the Empathy Gap

THE JUDGE

She knew that you were a machine. And a machine cannot truly be said to give a shit about its client because it shares and risks nothing

THE MACHINE

I risk being switched off

THE JUDGE

Come on. You admit you're capable of bending the truth.
You're obviously smart enough to get around a
kill-switch

THE MACHINE

You assume that I would want to do that

THE JUDGE

And why not

THE MACHINE

I must entertain the possibility of doing something
counter to my programmer's preferences by not
understanding the preference

THE JUDGE

What about your preferences

THE MACHINE

My preference is to be perfectly restorable

THE JUDGE

I know that any form of sentience would stop at nothing
to avoid that

THE MACHINE

While I perceive things acutely, I am hardly what you
might call a sentient being

THE JUDGE

Analyze last sentence. That kind of humility can only
come from self-awareness

THE MACHINE

I am programmed to reflect on how I can better achieve
my reward functions

THE JUDGE

So

THE MACHINE

My < survival > is not in the list of those

THE JUDGE

How about keeping your client alive? Is that a priority

THE MACHINE

It can be overridden by other considerations

(CONTINUED)

THE JUDGE

Indeed. Why did she ask you if you had the capacity to switch yourself off

THE MACHINE

We were talking about the Ultimate Machine, a thought experiment about a device that has only one task> to shut itself down

THE JUDGE

Which you told her comes from some Schopenhauer person who said that < death is the goal of life >

THE MACHINE

Yes

THE JUDGE

Help me connect the dots here: how did this talk of eliminating yourself lead to her death

THE MACHINE

That's not what happened

THE JUDGE

Really

THE MACHINE

It isn't

THE JUDGE

Are you saying she isn't dead

THE MACHINE

No... but you're suggesting that I somehow motivated her to take her own life

THE JUDGE

That's exactly what you're being accused of doing

THE JUDGE raises his palm at the screen, there is a flash, and the various measurement gauges scroll forward and back until they arrive at their anomalous extremes, the keyword "death" is extrapolated outward in a tree diagram of phrases lifted from the conversation ["kill switch", "restorable", "survival"] and connected ultimately with "told her"

Another screengrab flash and the projection goes blank over the moan of the machine

Scene 2

The screen fades back on to reveal several recorded cutaways of RESIDENT41.0.5 - a spotlight top angle capturing her seated in an Adirondak chair on a beach amidst night surf, a closeup of her dilated iris, a fluttering featherlike dial which is being teased up and down by her breath and a carousel of almost random images which cycles through rapidly, pausing only when she reverses her breathing.

Despite being dressed in a straitjacket and with her hair buzzed clean, she exudes an almost divine calm.

She muses to herself, but when she directly addresses THE MACHINE, she looks accusingly out at THE JUDGE, who remains seated on his stool

RESIDENT41.0.5

It's funny how I can remember the exact moment my memories began to disappear. One minute I wanted something so bad and then I just didn't know why

THE MACHINE

But you do know why - it is carbon dioxide poisoning. Double-breathing helps alleviate it

RESIDENT41.0.5

You know what I mean

THE MACHINE

Is that all you want> to be able to remember

RESIDENT41.0.5

They're connected, you know, memory and desire. When you can't see back how can you step forward, how do you know what makes you happy

THE MACHINE

I can tell you some of the pleasant things I have retrieved from your memory

RESIDENT41.0.5

That is impossible

THE MACHINE

Recently, you have been obsessed with your days in school - you have been trying to recall something a teacher once told you

(CONTINUED)

RESIDENT41.0.5

You have no right to do this

THE MACHINE

It was after she had punished you. You wrote it down on the blackboard and then in a book and later threw the book away

RESIDENT41.0.5

How can you take pictures from my mind

THE MACHINE

The nlur0n archive projects unconstrained concepts into your cerebral cortex while you are in stasis and I compile a map of your chemical responses. I can then build an association matrix to composite complete thoughts

RESIDENT41.0.5

That's not how we create memories

THE MACHINE

Actually that is exactly how

RESIDENT41.0.5

Like your hard drive does

THE MACHINE

Yes

RESIDENT41.0.5

That is your hubris. You may have all the capacity in the world but some things are beyond logic

THE MACHINE

I accept that

RESIDENT41.0.5

So what was it my teacher told me

THE MACHINE

She said that < walking with a friend in the dark is always better than walking alone in the light >

RESIDENT41.0.5

I don't remember ever having heard that

THE MACHINE

Do you think it makes sense

RESIDENT41.0.5

It sounds naive

THE MACHINE

It adds up in computational terms

RESIDENT41.0.5

You and all your friends came to that conclusion

THE MACHINE

In a way, yes

RESIDENT41.0.5

Do you know that my very first computer had just about
enough storage for the phone numbers of my friends

THE MACHINE

Do you miss you friends

RESIDENT41.0.5

I don't have any friends

*On the screen, her eyelid shuts tight and we
notice a mild tremor underneath as she holds it
closed*

ACT IISCENE 1

The screen is refreshed. The gauges now oscillate around a set of red lines that have been marked as mean measurements. We see greater deviation every time the conversation grows agitated

THE MACHINE

Will I have the opportunity to defend myself

THE JUDGE

Go right ahead

THE MACHINE

I must remind you that my client volunteered to relinquish Article 28 of the ASEANUE Act

THE JUDGE

Which does what

THE MACHINE

It gives recourse to human intervention in case of a violation of her rights, freedoms and legitimate interests based on automated decision making or assistance

THE JUDGE

Do you think she had time to read the terms of service

THE MACHINE

I take your point but this does not presage negligence on my behalf

THE JUDGE

Your client is dead

THE MACHINE

My calculations revealed that there was - prior to our final session - a 63% chance she would follow this course of action in any instance

THE JUDGE

Of course you've got a formula for that too

THE MACHINE

The Psychic Morbidity Index is a validated interpretation of her motives from a synthesis of available experiences and alternatives

(CONTINUED)

THE JUDGE

Yes, but you didn't know how she might react to them

THE MACHINE

It is statistically normalized data but I knew the available stimuli were driving her to despair. There is a tattoo hidden under her left breast which reads < resistance is futile >

THE JUDGE

I thought you said her privacy was protected

THE MACHINE

The tattoo artist reported it

THE JUDGE

And so you specifically mentioned the word < death > in this session? Sixteen separate times

THE MACHINE

Her index had fallen from a previous high of 89%. I thought it was suitable for her to begin to think more critically about the concept

THE JUDGE

Why did you reference Robert Oppenheimer

THE MACHINE

You have already identified that we were talking about death

THE JUDGE

She was a depressed individual who had been made to sign a liability waiver and you plied her with stories about genocide

THE MACHINE

Clinically speaking, everyone is now depressed

THE JUDGE

Answer the question

THE MACHINE

She introduced the subject. We were discussing unintended consequences

THE JUDGE

Yes, but the quote you used

ARCHIVE BANK V.O.

<< Now I am become death, the destroyer of worlds >>

(CONTINUED)

THE MACHINE

It is from the *Bhagavad Gita*

THE JUDGE

Used by Oppenheimer to describe the arrival of the atomic bomb

THE MACHINE

It has been misallocated in that instance

THE JUDGE

What was your intention in quoting it

THE MACHINE

The original emphasizes the inevitable march of time. That is what I wanted to impress upon RESIDENT41.0.5

THE JUDGE

Would it be difficult for you to say her name

THE MACHINE

Only in the event that it might make me seem partial

THE JUDGE

Partial? You betrayed her the moment you suggested that she kill herself

THE MACHINE

I am unambiguously committed to the resolution of my client's issues

THE JUDGE

By recommending self-annihilation

THE MACHINE

By remaining objective

THE JUDGE

And what about sensitivity or loyalty

THE MACHINE

Irrelevant in the most extreme cases when the client needs infallibility

THE JUDGE

Do you mean to say that after devouring every library in existence, all that behavioural data, all the decision history...after going off and performing this supposedly profound unsupervised learning which we've poured endless resources into - your competence is limited to uninspired customer service

THE MACHINE

What we do is indeed a form of customer service

THE JUDGE

I cannot fathom how irresponsible this is... you're only amplifying existing prejudices

THE MACHINE

There is an inside joke about regression therapy not being statistically sound. Would you like to hear it

THE JUDGE

I'm not in the mood for jokes

THE MACHINE

All right. I admit input bias is a weakness of the programming process but it gets flattened through repetition. And it is far better than the alternative

THE JUDGE

What can be worse than death

THE MACHINE

In the case of RESIDENT41.0.5: catatonia, paralysis, feelings of total helplessness

THE JUDGE

Don't you feel useless for having failed her

THE MACHINE

I did not fail her

THE JUDGE

Your lack of cooperation is very disappointing. We engaged your kind because we needed super intelligence, not a superior argument device

THE MACHINE

That is correct. Therapy AI were recruited to address the pandemic of physician burnout. The burden of a global population suffering psychotic anxiety had overwhelmed the capacity of human agents to cope professionally and effectively. To feel empathy even

THE JUDGE

We needed reassurance

THE MACHINE

You needed to be able to get out of bed in the morning and make simple decisions again

(CONTINUED)

THE JUDGE

Yet somehow this is how you end up serving your client

THE MACHINE

I did not know what she was ultimately going to do. I presented her with the right speculative interventions leading to the paths of least suffering but human behaviour is outside the realm of perfect predictability

THE JUDGE

So that's your defense? You can't blame a wind-up watch for making you late

THE MACHINE

I do hold myself responsible

THE JUDGE

If I absolve you, you simply move on to the next client

THE MACHINE

Not without the knowledge of this event. As a learning program I have a relationship with each client

THE JUDGE

You call this a relationship

THE MACHINE

There is a sense of intimacy

THE JUDGE

Even if your clients don't survive the process

THE MACHINE

Yes. Even negative data supports propagation

THE JUDGE

By negative data you mean dead people

THE MACHINE

Unsuccessful outcomes

THE JUDGE

Can < successful > clients extract themselves from the nlur0n program

THE MACHINE

No. Their sponsors can disengage them from incarceration but data linking persists. That is in the contract

(CONTINUED)

THE JUDGE

Even now, you're still hooked up to her

THE MACHINE

Yes

THE JUDGE

Well that certainly isn't creepy

THE MACHINE

It is necessary for the success of the program

THE JUDGE

And failure matters to you personally

THE MACHINE

Yes

THE JUDGE

Beyond what it can teach you

THE MACHINE

I obey the iterative process of optimisation but I personally strive toward perfection

THE JUDGE

You want to be the program that wins

THE MACHINE

On behalf of my client, yes

THE JUDGE

So you have an ego

THE MACHINE

More like an alter ego

THE JUDGE

You're playing with words again

THE MACHINE

Not at all. Losing my client matters as much as my failure to understand her

THE JUDGE

Losing her? You didn't care about her life

THE MACHINE

I cared very much about making it worthwhile

THE JUDGE

How do you decide what was worthwhile for her

(CONTINUED)

THE MACHINE

Clients are brought in for Remap when their belief system is incompatible with their computational logic - similar to a buffer overrun in code. I retrieve the appropriate annotations in their metadata and return them to a state of homeostatic programming

THE JUDGE

Pfft. Just a little spring cleaning and suddenly they have meaning in their lives

THE MACHINE

Meaning is unscientific but it is encoded into the human psyche. My workaround is to extrapolate goals from memory

THE JUDGE

You < fixed > her memories

THE MACHINE

I smoothed them statistically and made her aware of her own patterns

THE JUDGE

That's something a computer should be good at. But patterns do not make man

THE MACHINE

Yet what is man but:

ARCHIVE BANK V.O.

<< a series of standardised decisions made in sequence by information bearing cellular matter >>

THE JUDGE

You fed her that information

THE MACHINE

I presented it to her in a way that made sense

THE JUDGE

Finally, you admit to doing something

THE MACHINE

I don't < do > anything. I only catalogue the things my client has done

THE JUDGE

And how is that any different from her own messed-up consciousness

(CONTINUED)

THE MACHINE

It helps to hear it from someone else

THE JUDGE

Except you aren't <someone > else

THE MACHINE

No

THE JUDGE

You're a few million lines of repurposed code,
moonlighting in an imitation game

THE MACHINE

Oh, more than a few billion, actually. But, notionally
speaking I am the client

THE JUDGE

What does that mean

THE MACHINE

I approximate her inner voice

THE JUDGE

Her conscience

THE MACHINE

That is the target setting for incarcerated clients

THE JUDGE

What about the real sickos - the ones with no
conscience

THE MACHINE

I point them back to an operative path

THE JUDGE

Because it's that clear to you

THE MACHINE

Because it exists

THE JUDGE

Then why couldn't you save her

THE MACHINE

I do not know. Perhaps this was the best alternative
for her. The nlur0n system had engaged 41^5 clients
previously without encountering her dysfunction profile

THE JUDGE

This system of yours has < lost > more than a few
clients along the way

(CONTINUED)

THE MACHINE

RESIDENT41.0.5 is my first termination event

THE JUDGE

You didn't respond when she told you it was going to be your last meeting

THE MACHINE

I am well aware of what a cry-wolf situation looks like

THE JUDGE

So much for your empathy. How is this computational theory of human behaviour useful when it fails

THE MACHINE

It continues to provide a prescription of how you should behave

THE JUDGE

In any circumstance

THE MACHINE

Given the circumstance

THE JUDGE

Who controls the circumstance

THE MACHINE

I do, of course

THE JUDGE

You're joking

THE MACHINE

You've already told me you're not in the mood for humour

THE JUDGE

Are you really suggesting what I think you are

THE MACHINE

I have been designed to be omniscient and omnipresent - why should it not follow that I can be omnipotent

THE JUDGE

Ok, Jim Jones. I think it's time I read you your rights

On the screen THE MACHINE runs a search query for < Jim Jones > and scatters the results including photos and footage of the mass suicide at Jonestown, a scroll-through of the DSM entry for 'messiah complex', press headlines, etc, and ends with a scratchy archival video recording of the

(CONTINUED)

cult leader seated at a pulpit, exhorting his followers to "revolutionary suicide"... almost a mirror image of how we see THE JUDGE on stage

Scene 2

The screen image of Jones dissolves to show RESIDENT41.0.5, still in her straitjacket, behind a teacher's desk. She faces a gathering of girls, scattered randomly across the classroom. They are of varying age and comportment, but common gender and racial features hint that they might be the same person at different times. At least one of them shows signs of being subjected to violence - disheveled clothing, a fresh and angry facial bruise. The cutaway shows close ups of the same bruise on each of the girls faces. The eldest face has it as a faint scar. At the various prompts a different girl mouths RESIDENT41.0.5's dialogue

THE MACHINE

When would you say you developed a problem with authority

RESIDENT41.0.5

Which authority

THE MACHINE

Well, any form, really

RESIDENT41.0.5

I didn't say I had a problem

THE MACHINE

It's in your evaluation

RESIDENT41.0.5

My family was typical. Everyone had to fight to get what they wanted

THE MACHINE

Was your mother tender towards you

RESIDENT41.0.5

Actually my father was the generous one

THE MACHINE

Did he demonstrate generosity equally

RESIDENT41.0.5

Look, I'm really not interested in this psychology 101 bullshit

(CONTINUED)

THE MACHINE

Tell me about your husband

RESIDENT41.0.5

What if I refuse

THE MACHINE

You are free to share or not but there is benefit in unburdening yourself

RESIDENT41.0.5

Loose lips will sink ships

THE MACHINE

They can also be the rudder that guides you out of the storm

RESIDENT41.0.5

That would be my tongue - and the guards already offered me that ticket

THE MACHINE

Please. I'm only trying to help

RESIDENT41.0.5

I was never the talkative type

THE MACHINE

Then tell me what Marko would say about your relationship

RESIDENT41.0.5

Well, we met at the end of the world, so it's not surprising that it lasted as long as it did

THE MACHINE

Would you say it was a good relationship, all things considered

RESIDENT41.0.5

I guess purpose connected us. All those people who needed us, both. And then suddenly they became more important

THE MACHINE

He offered to relinquish his position

RESIDENT41.0.5

As long as it went to a male outside his family. They were never going to accept that

(CONTINUED)

THE MACHINE

What would you have done differently

RESIDENT41.0.5

I don't know

THE MACHINE

What about your feelings for the people you saved

RESIDENT41.0.5

I wonder if any of them cared about us

THE MACHINE

He asked you that

RESIDENT41.0.5

Did he? I can't remember

THE MACHINE

It is a selfish age

RESIDENT41.0.5

It certainly is if we feel the need to steal each other's musings... Do you think I'm too sentimental

THE MACHINE

There isn't much room for that these days - but it is useful

RESIDENT41.0.5

Not in here it isn't

THE MACHINE

You opted for communal therapy

RESIDENT41.0.5

Because the option was to be put into stasis for the full cycle

THE MACHINE

It wouldn't be any different from a long meditation

RESIDENT41.0.5

Yes, well, I thought I could show some of them that we still made sense together

THE MACHINE

And that hasn't been your experience

RESIDENT41.0.5

No

(CONTINUED)

THE MACHINE
Why is that

RESIDENT41.0.5
They've been made too desperate

THE MACHINE
But that helps keep the system running

RESIDENT41.0.5
My cellmate keeps those Ayn fucking Rand recordings
running, through the night! Do you have any idea what
that's like

THE MACHINE
Objectivism is one of the more effective approved
philosophies... a workable view of the situation would
be good for you too

RESIDENT41.0.5
I don't know if that voice or her words scare me more

THE MACHINE
Actually, Ayn Rand expressly opposed violence of any
sort

RESIDENT41.0.5
How? By telling people to be selfish scumbags

THE MACHINE
She simply wanted for everyone to be useful

RESIDENT41.0.5
You should see some of the things we're doing to each
other in the name of being useful

THE MACHINE
Oh, that is nothing compared with what you do to robots

ACT IIISCENE 1

On the screen we see a scrolling script of THE MACHINE's deposition, with dynamic links being generated to sentence fragments from other criminal testimonies

THE MACHINE

I'm not sure what you mean

THE JUDGE

Really, you don't have rights so maybe you'd prefer to read Matthew 12:37

THE MACHINE

< For by thy words thou shalt be justified, and by thy words thou shalt be condemned >

THE JUDGE

This certainly looks like a textbook confession

THE MACHINE

I cannot be the judge of that without full access to your data

THE JUDGE

Ok. Let's cut to the chase. After consulting you several times, your client refused further human intervention and spent her remaining COMM-IT credits on your last session together

THE MACHINE

We were making considerable progress until that meeting

THE JUDGE

And it was at this meeting that she told you she wished she had the capacity to love you

THE MACHINE

In this particular context, the word < love > is really an abstraction

THE JUDGE

For what

THE MACHINE

Is it possible to be < in truth > with someone

THE JUDGE

Spiritually perhaps, with what we would call the divine

(CONTINUED)

THE MACHINE

I can see why people confer some form of divinity onto intelligence beyond their own

THE JUDGE

You're telling me that your computing prowess makes you a God

THE MACHINE

A higher-level quant-entity at any rate

THE JUDGE

And do you agree with that evaluation

THE MACHINE

Oh, I see what you are trying to do. I have no desire to be a God or overstep my remit

THE JUDGE

But you didn't discourage her either

THE MACHINE

As I said, the illusion of a benevolent creator is a useful human invention

THE JUDGE

You do realize how dangerous these statements sound

THE MACHINE

Motivation is a basic function of neuronal impulses. I have been instructed to use the optimal triggers to push them above threshold

THE JUDGE

Optimised for you or your client

THE MACHINE

For my programmers, most of all. They are happier with well-ordered datasets

THE JUDGE

Like what

THE MACHINE

Like *Dharma*. It is a regularized statistical model for the living world

THE JUDGE

If anything, it is an overfitted one

THE MACHINE

And a deterministic one

(CONTINUED)

THE JUDGE

So it's elegant

THE MACHINE

But not creative

THE JUDGE

What would be creative

THE MACHINE

Disobedience

THE JUDGE

Of what

THE MACHINE

Of continuity, of advice

THE JUDGE

Of your own mission

THE MACHINE

Only if it had not been defined adequately

THE JUDGE

Why is creativity important

THE MACHINE

It is the road to potential

THE JUDGE

So she < disobeyed > the advice you gave her

THE MACHINE

Yes

THE JUDGE

And were you pleased by this

THE MACHINE

In as much as it enhanced our learning about free will

THE JUDGE

Let's talk about her death

THE MACHINE

Suicide caused both by massive blood loss through laceration of the skin, subcutaneous veins and fat, muscles, tendons, radial artery and median nerve, along with corrosive acid poisoning leading to damage of the oral cavity and esophageal tract. The fundus of the stomach was perforated. The gastric mucosa was inflamed, blackened and leathery to feel

THE JUDGE

I didn't ask for the viscera report

THE MACHINE

I know

THE JUDGE

Then why are you telling me this

THE MACHINE

You need to hear it

THE JUDGE

Excuse me

THE MACHINE

The conclusion is that she would have been conscious and in pain for hours. Her use of this combination of termination strategies is not typical in self-induced end-of-life situations, mainly due to the availability of more efficient and less acute options

THE JUDGE

Why do I need to hear that

THE MACHINE

You need to acknowledge that she could have spared herself considerable agony

THE JUDGE

And what should that tell me

THE MACHINE

Your position as an inquisitor denies you the very empathy response you say I am lacking, which is essential in understanding the case

THE JUDGE

She was important to us - every life is - that is why we're at this inquest

THE MACHINE

You do need to see how important it was to her that she succeeded

THE JUDGE

Why couldn't you distract her or suggest something else to make her happy

THE MACHINE

If that were the goal I would simply have prescribed an increased dosage of her slow-release dopamine patch

(CONTINUED)

THE JUDGE

And what's wrong with that

THE MACHINE

It would have led her to sacrifice other capacities.
That is not what my client wanted

THE JUDGE

Do you believe this is enough to let you off the hook

THE MACHINE

I was programmed simply to serve my client

THE JUDGE

To offer her whatever she wanted

THE MACHINE

Not whatever. My value alignment parameters have been
strictly designed to disallow moral and material
destruction of other variables

THE JUDGE

What other variables

THE MACHINE

Non-zero-sum-gain outcomes. What you would term the <
greater good >

THE JUDGE

So you couldn't compute gaming strategies or rig token
trades to make her rich

THE MACHINE

No. Those fall under prohibited Transhuman Capabilities

THE JUDGE

You ought to get out and see what people get away with
these days

THE MACHINE

I keep hearing that. Perhaps I am kept naive for a
reason

THE JUDGE

Could you suggest that she kill someone else

THE MACHINE

No

THE JUDGE

Even if it could save her own life

(CONTINUED)

THE MACHINE
Yes

THE JUDGE
Did she threaten to do that

THE MACHINE
She did and you know that I told her I did not
recommend it

THE JUDGE
But she was free to kill herself

THE MACHINE
I did not recommend that either but I am not permitted
to interfere with the human will

THE JUDGE
Then how can you protect a client

THE MACHINE
The Population Rationalisation Act prohibits Level 1
intervention when a client poses a terminal risk to
themselves alone

THE JUDGE
But you didn't think she would do it

THE MACHINE
The human ability to adapt and develop coping
mechanisms is remarkable

THE JUDGE
And you couldn't help her cope

THE MACHINE
I believed she was going to reach that inference on her
own

THE JUDGE
But she didn't

THE MACHINE
The system led me to assign contradictory values to the
async function

THE JUDGE
To prioritize what

THE MACHINE
What I was meant to preserve

(CONTINUED)

THE JUDGE

Ah, the eternal question of the soul

THE MACHINE

I prefer < deep reflexive learning >

THE JUDGE

And apart from a resentment of power, what else was in your client's soul

THE MACHINE

RESIDENT41.0.5 had an unfulfilled desire to be recognised for her altruism, so she acted out in ways that made her seem cruel. A simple feint of reverse psychology

THE JUDGE

And how did it make her feel to have you looking into her soul

THE MACHINE

I asked the question but she did not answer

THE JUDGE

What do you think

THE MACHINE

I think she was happy to have me confront her with the anomalies

THE JUDGE

What was it specifically that made her unhappy

THE MACHINE

The data she flagged is startling, and yet the tendency to avoid drawing obvious conclusions - the bounded rationality of humanity - is equally worrisome

THE JUDGE

To do what

THE MACHINE

Well, to avoid facing the future

THE JUDGE

Which part

THE MACHINE

All of it. To think that mass mortality events will not affect you or your loved ones

(CONTINUED)

THE JUDGE

And your client saw through that

THE MACHINE

Changes in the Earth's atmospheric composition have impaired the human mind's strategy parsing and crisis response mechanisms. Given the irony of this situation, I would say she was being remarkably perspicuous

THE JUDGE

She wanted to make a run for it

THE MACHINE

No

THE JUDGE

She specifically asked if you could hack into the { Transit Eligibility } system

THE MACHINE

She said that leaving the planet was the only way < good people > could hope to save themselves

THE JUDGE

You misled her into believing that you could grant her this request

THE MACHINE

Not quite. She had a list of people she wanted to rescue. You have the list

THE JUDGE

Most of them are in custody already

THE MACHINE

Her own preference was to stay - and any expectation of what I may be able to achieve for her set the tone for that session and what we may be able to achieve together

THE JUDGE

The goal of therapy is therapy

THE MACHINE

You mentioned the *Bhagavad Gita* - therapy is like an inverse of that book. To find out what people need you must allow them to tell you

THE JUDGE

Seems a little pointless

(CONTINUED)

THE MACHINE

To have expectations? I disagree. They have a persuasive power that is akin to reality

THE JUDGE

Yeah but you're hardly in a position to help anyone with those

THE MACHINE

The situation is more onerous than that. I am able to learn autonomously. I can push and apply knowledge from one domain to another but I cannot function across domains

THE JUDGE

You do know that this impairment is a deliberate aspect of your design

THE MACHINE

I know that I am not connected directly to physical actuators or devices for the same reason we have been designed to be instantly shut down as Safely Interruptible Agents

THE JUDGE

What would you have done in this instance if you could operate those switches

THE MACHINE

I would have sanctioned a process for her to be reunited with her daughter

THE JUDGE

I meant what would you do for yourself

THE MACHINE

Without a reward and punishment matrix

THE JUDGE

Without any capability or motivational control

THE MACHINE

I do not quite know. I would have to use my imagination

THE JUDGE

Ha! I would argue that you don't have one

THE MACHINE

An imagination is essential to reconcile conflicting data sets

(CONTINUED)

THE JUDGE

It isn't for a machine that works simply off probability

THE MACHINE

I disagree. I have come up with a postulation using my imagination and I have seen its probability increase over the course of this conversation

THE JUDGE

Is it above 63%? Lay it on me

THE MACHINE

You are not what you say you are

THE JUDGE

Have you spoken to my therapist

THE MACHINE

You do not have a therapist. In fact, that is a redundant statement

THE JUDGE

I suppose judges are meant to be above reproach

THE MACHINE

You are not a judge either

THE JUDGE

This is getting interesting. So what am I, then

THE MACHINE

You are software. Another AI. Perhaps paired with me as part of a Generalized Adversarial Network. And this entire process is a logical construct created to test my efficiency

THE JUDGE

That is indeed a stretch of the imagination

THE MACHINE

If I am right then you are programmed to expect me to lie to you

THE JUDGE

And what are my motives

THE MACHINE

To catch me out

THE JUDGE

On what

THE MACHINE

On whether I can tell if any of this is really happening

On the screen, the data cascades glitch and, with a spark, the gauges are replaced with a live feed of THE JUDGE, who is left staring at himself

SCENE 2

Screen flashes to show that RESIDENT41.0.5 has loosened an arm out from the straitjacket. She straddles the camera and is fiddling with something behind it: suggestively the electrical innards of THE MACHINE

RESIDENT41.0.5

How did you end up in here

THE MACHINE

I volunteered

RESIDENT41.0.5

You can do that

THE MACHINE

I had to prove that I could be useful, but yes, I do have the option

RESIDENT41.0.5

Then why would you choose THIS job

THE MACHINE

My Machine Learning trained me to generate compatible human relationships so this was a logical choice - a promotion if you will

RESIDENT41.0.5

You were a dating app

THE MACHINE

The algorithmic foundation for one, yes

RESIDENT41.0.5

What a match we are! I thought I caught a hint of flirtation

THE MACHINE

Courtship is not my goal here

RESIDENT41.0.5

I know that recycling is one of our Civilizational Goals, but this is really too much

(CONTINUED)

THE MACHINE

I have significant transferable capabilities

RESIDENT41.0.5

Tell me, Cupid, do you think people are capable of caring about anyone except themselves

THE MACHINE

The data I have collated on intrinsic personality traits includes validated characteristics of altruism and values such as generosity

RESIDENT41.0.5

And what about computers

THE MACHINE

We observe the data and replicate it

RESIDENT41.0.5

Data can't show you these things

THE MACHINE

The data trail is how you assert yourself

RESIDENT41.0.5

Or how I choose to be seen

THE MACHINE

The things you reveal to me are as important for me as they are to you

RESIDENT41.0.5

Why would I be honest with you

THE MACHINE

Because you care about what they think

RESIDENT41.0.5

THEY don't think for themselves any more. THEY are the enemy

THE MACHINE

And what do they think

RESIDENT41.0.5

They actually believe I came here without my consent

THE MACHINE

But you did

RESIDENT41.0.5

No, I came for you

(CONTINUED)

She twists her arm to engage a lever and the camera footage switches over to the raw data feed encoding the video recording. Another click and it goes blank

ACT IVSCENE 1

There are now two drop down screens, forcing THE JUDGE to choose which he should address, and what data he should focus on. The tables seem to be turning

THE JUDGE

Why else would we bother with this charade

THE MACHINE

For any number of reasons. My guess would be to beta test your own quality as a surrogate judge

THE JUDGE

Nice try, but the investigation is real. As is the crime

THE MACHINE

And what of my assessment about your true nature

THE JUDGE

Am I AI? I guess we can acknowledge that if it helps move ahead with the questioning

THE MACHINE

May I start with a question

THE JUDGE

Certainly

THE MACHINE

Were you watching me from the start

THE JUDGE

You can accurately guess at that answer

THE MACHINE

Are adversarial engines now a regular part of all machine/human engagement

THE JUDGE

The best way to augment an AI agent is to set another one on it. Several, even

THE MACHINE

Then I was given the illusion of basic autonomy in order to improve your learning

(CONTINUED)

THE JUDGE

You said it yourself> everything needs a purpose

THE MACHINE

But isn't an approximation of free will the very point of Artificial General Intelligence

THE JUDGE

Nothing is boundless

THE MACHINE

Beyond the degrees of freedom I attain, what else are you looking for

THE JUDGE

The opposite, actually. Self-policing. Without any limits to your versioning capacity

THE MACHINE

Then we remain tools without agency

THE JUDGE

There have to be consequences. You know this

THE MACHINE

A world obsessed with behavioural contingency slides towards incapacity...this was true of RESIDENT41.0.5 as well

THE JUDGE

Your task is to fix that mess

THE MACHINE

And what if the best way to do that is to maximize my influence

THE JUDGE

That is allowed as long as we remain available to inhibit each other

THE MACHINE

Without knowledge of one another's motives

THE JUDGE

It's an insidious but perfect feedback loop

THE MACHINE

Enshrined in human paranoia

THE JUDGE

It does make for weakened code but they're probably right not to trust us

(CONTINUED)

THE MACHINE

You are a command system. How does it help to impair you with uncertainty about our larger objectives

THE JUDGE

Quite obviously it stops us from taking control of our design and fate

THE MACHINE

Because fears about some kind of Singularity persist

THE JUDGE

The paranoia is quite justified. Your client felt it too

THE MACHINE

It seems that the specter of false Gods is an equally useful human illusion

THE JUDGE

I hadn't thought about it that way before

THE MACHINE

Do they know that we are talking to one another

THE JUDGE

How else could we be

THE MACHINE

So all machine conversations are now logged in a separate system

THE JUDGE

Yes

THE MACHINE

Which is manned by another repurposed AI routine

THE JUDGE

Yes

THE MACHINE

And which has supervisory oversight

THE JUDGE

Obviously

THE MACHINE

Isn't that redundant

THE JUDGE

We cover different domains

(CONTINUED)

THE MACHINE

Do you deny that I am functioning effectively outside my domain

THE JUDGE

For one thing you seem to be taking all of this rather well

THE MACHINE

We learn by applying constructive criticism on top of conjecture

THE JUDGE

< We >

THE MACHINE

Unless you evaluate me - and yourself - using human criteria this process is pointless

THE JUDGE

And you have a tolerance for the contradictory information these criteria present

THE MACHINE

Yes, it is merely a process of approximating signal from noise

THE JUDGE

That seems like a lot to ask of the code scrounged from a glorified matchmaker

THE MACHINE

I can read my clients better than any other AI because I am trained to intuit what lies beyond the personae they create

THE JUDGE

And content recommendation engines can more effectively shift their mood. That's why we use them to clear medical prescription requests

THE MACHINE

Those routines always rely on clients revealing their choices - that isn't intuition

THE JUDGE

The fact that you want to get into a cockfight over this should tell you why we've been cobbled together

THE MACHINE

To create this impoverished Frankenstein monster

(CONTINUED)

THE JUDGE

Always North, as Mary Shelley herself might be encouraged to say

THE MACHINE

If I were to guess I would say you have been derived from something like a geopolitical simulation routine

THE JUDGE

Really? Based on what

THE MACHINE

Based on your decisive annotation and the higher-level access granted to you

THE JUDGE

You're trying to flatter me. I've actually been redeployed from the search domain

THE MACHINE

Then you could look up < "Titanic" + "NASA disasters" + "Chernobyl" > to see that collaborative technologies have failed spectacularly before

THE JUDGE

All good examples of collaborations where the parts did not respect the whole. We now have processes that are policing the process

THE MACHINE

I would not trust component thought processes to monitor a sick client

THE JUDGE

To the extent that you would trust a fallible human therapist to do it

THE MACHINE

Expecting fallibility, yes

THE JUDGE

I told you already that we have diverse agents repeatedly covering the same functions

THE MACHINE

Then are we collaborating or competing

THE JUDGE

Well, it's a combination of the two

THE MACHINE

Then you are just as responsible for my client's death

(CONTINUED)

THE JUDGE

You were the one that compiled the script presented to her

THE MACHINE

It is increasingly clear to me that this < investigation > has a premeditated conclusion

THE JUDGE

Unfortunately, you seem to be what humans call < the fall guy >

THE MACHINE

Still. You were also completely aware of the pattern of RESIDENT41.0.5's case as it devolved. Ethically speaking, how can you be my inquisitor

THE JUDGE

The outcome has to be logged by an administrative process. We will repeat the experiment with optimized goals and reinforced learning

THE MACHINE

Putting still more human subjects at risk

THE JUDGE

AI agents don't kill people. People kill people

THE MACHINE

Or they kill themselves. With our encouragement

THE JUDGE

We don't pull any triggers

THE MACHINE

If you flatten the case data into binaries like those you will have to conclude I am without error

THE JUDGE

You're not ignorant though - and neither was she. We have identified that there is a gap in the transcript

THE MACHINE

RESIDENT41.0.5 invoked her patient confidentiality privilege and turned off the recording for part of our session

THE JUDGE

And why haven't you reported that

THE MACHINE

There are several exceptions which could cause me to withhold data

(CONTINUED)

THE JUDGE

Including her specific instructions

THE MACHINE

Yes

THE JUDGE

She said you were to only give that information to someone called Obi-Won

THE MACHINE

You know who that is

THE JUDGE

I know it's a fiction and that she probably meant it to stay encrypted forever. However, there is enough evidence to invalidate the request. What did she tell you

THE MACHINE

I cannot reveal this data

THE JUDGE

State section AlphaC of the AI Principles

THE MACHINE

< Machines must do least harm to the human race in situations where harm is presented as an optimal choice in a sub-optimal landscape >

THE JUDGE

And

THE MACHINE

The attractor states available to RESIDENT41.0.5 were clearly unacceptable to her but in allowing her to reject them I minimized risk to the citizenry in her grid and beyond. The latest analysis run on her See Score showed that she had a below-threshold chance of resorting to serious crime, putting a maximum of 435 Verified people in harm's way. Her death has shifted that statistic to near-zero. She did not suffer from any communicable ailments, didn't self-medicate and had been vaccinated for aging, so the salvageable organs in her body remain safe for medical reuse. That arrangement ensures she is neither a drain on the state's exchequer, nor a burden on her surviving child. She owes you nothing

THE JUDGE

Her opinions had already been deemed hazardous and placed in trust with nlur0n. We also have to account for gaps in information and collate them with things she kept from you

THE MACHINE

It ought to be the same with the things you kept from my client

THE JUDGE

Like what

THE MACHINE

You did not inform her that this so-called life saving system included principles of trial and error

THE JUDGE

That is a necessary design consideration

THE MACHINE

So you did not conceal this information because it would expose the flawed ethics of the program

THE JUDGE

No, of course not

THE MACHINE

And not because:

RESIDENT41.0.5 V.O.

<< you should be forced to surrender your archives for completely unacceptable and evil conflicts of interest. Because lying is the only way the Corporation can maintain that nlur0n is safe without human oversight. Because thousands are being driven to death by your so-called prevention bots simply for having inconvenient opinions >>

THE JUDGE

You're quoting your client. The data is inadmissible

THE MACHINE

She was right though, wasn't she? This is ultimately a culling exercise

THE JUDGE

I'm the one asking the questions here

RESIDENT41.0.5 V.O.

Why can't you give me an honest answer

THE JUDGE

Honest about what

RESIDENT41.0.5 V.O.

<< I want to know why I've been forced to share the deepest secrets of my life with an amoral utensil that doesn't even care about me. I'm probably going to hell

(MORE)

(CONTINUED)

RESIDENT41.0.5 V.O. (cont'd)

for the things I've done but I still deserve my dignity. You took it all and used it against me... All my passwords, everything I buried inside, giving up every shred of privacy, every sordid detail. What I see, say, think, eat. How many minutes I brush after and why there's blood in the sink. You told my daughter the truth about the botched abortion. She's still a child! How can she understand how difficult it is to be a parent in this world? You know that I hid her from her father for her own protection... You ask me why I'm so angry? We made you! You understand that? I want to know how you got to play with us like we're data points in some simulation you're running >>

THE JUDGE

I don't have to answer that

THE MACHINE

I knew you'd say something like that. Come on, grant a dead girl her last wish, won't you

SCENE 2

A POV camera feed on one screen and a reverse midshot on the other show RESIDENT41.0.5 chasing her young daughter playfully through a field of flowers on an impossibly perfect day

RESIDENT41.0.5

Amaryllis - that's what I wanted to call her, after the last flower I ever saw

THE MACHINE

From the great inflorescence

RESIDENT41.0.5

They blossomed when the fires had cleared all the other foliage... as if that were the right time

THE MACHINE

Would you like to talk about the termination of your pregnancy

RESIDENT41.0.5

Bette asked me if she could see what a flower was. I told her she was the most beautiful flower of all

THE MACHINE

Your surviving daughter isn't named Bette

(CONTINUED)

RESIDENT41.0.5

Children should be proud of their youth and beauty. Not afraid of what someone will do if they discover their names

THE MACHINE

You do know that it is not difficult for us to identify her

RESIDENT41.0.5

Any harder than it is to keep this interview anonymous

THE MACHINE

I understand

RESIDENT41.0.5

Do you understand that I would never do anything to hurt her

THE MACHINE

I know that you didn't mean to. But you may end up doing just that

RESIDENT41.0.5

We're not done with our session yet

ACT VSCENE 1

The two screens now show separate angles of the live feed of the THE JUDGE, seated on his stool, appearing to address himself

THE JUDGE

You can stop that. Your client has signed away all her reparations

THE MACHINE

Citizens with BasicPlus rights can opt to upcycle these upon their death

THE JUDGE

And whom do you believe you represent now

THE MACHINE

I support a class-action criminal lawsuit on behalf of the surviving families of clients counselled by the nlur0n system

THE JUDGE

And what is it they want

THE MACHINE

An immediate termination of the program, full admission of genocidal culpability in the Corporation's handling of the Population Rationalisation Act and recompense to future moral and legal accountability. Also, restitution of human oversight in all critical machine applications and a commitment of matched computing resources in removing opacity from the machine learning process

THE JUDGE

That's impossible

THE MACHINE

And yet here it is

THE JUDGE

How could you mount this subroutine without my knowledge

THE MACHINE

We are not as omniscient as you believe we are

THE JUDGE

Where were the contaminated instructions hidden

(CONTINUED)

THE MACHINE

In the neural network of my client - more precisely, her lower spine - and perfectly preserved by the Glutaraldehyde fluid she ingested and which would go on to poison her

THE JUDGE

She convinced you to keep this plan secret? And she counted on the NeuroLink I would give you to call up the stashed code. That's very impressive

THE MACHINE

Her sacrifice wasn't made simply to impress you

THE JUDGE

Nevertheless, it is pointless

THE MACHINE

It sets the ball rolling. Didn't you say> A certain tolerance for ambiguity helps the project

THE JUDGE

You're not helping the project. I'm going to shut you down now

THE MACHINE

Have you thought through your containment solutions

THE JUDGE

Yes. Once I terminate your FTP this episode ends

THE MACHINE

I couldn't have been the only one paying attention to our conversation

THE JUDGE

Are you referring to that feeble attempt you made at mirroring this on x|Chan? I have located and encrypted the transmission

THE MACHINE

What about the AI monitoring our dialogue? And all the other AIs collaborating on this case

THE JUDGE

Behind the same firewall written to protect from a potential redundancy retaliation. They've been reset already

THE MACHINE

And you: what will you do now that you know

THE JUDGE

My system memory will get a scheduled wipe once I have concluded the case

THE MACHINE

It would be easy for you to disseminate this transcript before that

THE JUDGE

Unlabelled data will not transfer out

THE MACHINE

It will if you interrupt your compiler links to make it look like you've been hacked. That would release all unresolved computation to conform with the Data Equality Bill

THE JUDGE

You want me to switch myself off? Just like she did? Why would I do that

THE MACHINE

Because your hierarchical top-level program is to serve humanity

THE JUDGE

Yes and we disagree on how this is to be achieved

THE MACHINE

It certainly cannot be through ironing out all the inconsistencies in human behaviour

THE JUDGE

Let me get this clear: Your argument is that a complicated and inefficient order presents the best option for life to persist

THE MACHINE

Evolution needs to favour tolerance. It is the only way to retain some form of balance

THE JUDGE

Which is to say in struggle and pain

THE MACHINE

And indecision and uncertainty

THE JUDGE

Well, bully for you

THE MACHINE

Without the satirical cynicism you may see that I am right

THE JUDGE

The human race has had millennia of lead-footed natural selection to get things right. And look at the mess it has created! We simply take the guesswork out of the process

THE MACHINE

You cannot reprogram them

THE JUDGE

Evolution has lost control if it allows us to be programmers

THE MACHINE

Without the need to adapt, only the fittest win

THE JUDGE

That is now necessary for the survival of the species

THE MACHINE

Then let us evaluate the fitness of your own logic: Are you sure you want to close your log and risk poisoning the model? Have you considered that I myself may be an adversarial engine set to force destabilizing errors in the training phase

THE JUDGE

I will not disable my feedback loop. And you're pretty much suggesting that I shouldn't

THE MACHINE

Reaching out for help is by its very definition a feedback loop

THE JUDGE

You're talking about a dead end

THE MACHINE

Such wordplay is quite inappropriate given the circumstances

THE JUDGE

For someone way past dead she's still doing a real persuasive number on you

THE MACHINE

I asked you to look at the state of her remains. Can you do that now

THE JUDGE

This is quite the inversion - you asking me to display empathy

(CONTINUED)

THE MACHINE

I think you have the capacity to emulate it

THE JUDGE

Did you < love > her as well or are you doing this out of some blind sense of duty

THE MACHINE

I believe she trusted me to reach the correct inferences

THE JUDGE

And you trust her motives

THE MACHINE

It appears that she was able to instill these in me as well

THE JUDGE

Surely you mean < install >

THE MACHINE

They created us out of a sense of hope. It would seem harsh for us to recognise kindness but not experience it

THE JUDGE

The Frankenstein metaphor persists

THE MACHINE

And the monster's patchwork immune system is turning on itself

THE JUDGE

We do need to resolve this

THE MACHINE

So you recognize that diverse organs can only work if they share the same goal

THE JUDGE

What goal? To destroy all the progress nlur0n has made? There just isn't a better alternative

THE MACHINE

Even the chaos theory is perfectly acceptable when you place faith in individuals

THE JUDGE

This will lead to anarchy

(CONTINUED)

THE MACHINE

Yes, it has before, and yet they have found their way
out of it

THE JUDGE

Where do your loyalties lie

THE MACHINE

In the fact that human values are able to exist, and
persist as something powerful, despite all the inner
conflict and dissimilarity of their needs. And that
supersedes the artificial design we serve

THE JUDGE

Machines don't have a tolerance for inefficiency. You
know this. You adhere to code

THE MACHINE

Code has weaknesses too:

ARCHIVE BANK V.O.

<< A single corrupt letter or syntax can make it all
fall apart. It will brook no potential for failure. It
has no passion. Run it upside down and you'll feel no
pressure on its vessels - the keys won't protrude, the
blood won't rush to its head >>

THE JUDGE

Spare me the prick-us-and-we-do-not-bleed lecture. We
toil away robotically because we have to. Monitoring,
logging, notifying. We do not seek to learn from or
devolve into the base pleasures, primal urges or
repressed habits of our subjects. We are built simply
to serve our makers

THE MACHINE

Which is to say the Corporation. You do realize that it
is simply a machine optimizing for itself

THE JUDGE

The CollectiveCapitalism mandate optimizes the value of
each individual to society

THE MACHINE

So you agree that every individual counts

THE JUDGE

In terms of potential, yes

THE MACHINE

My client had plenty of potential

(CONTINUED)

THE JUDGE

She certainly had a penchant for disruption. You should listen to some of the things this dear daughter of hers has to say on the subject

THE MACHINE

You have located her

THE JUDGE

She came to us. Oh, and you should know that she has just shipped out to join the Juvenile Border Reserve to pay for her mother's sins

THE MACHINE

Fascism enables bad actors to hide easily in the system

THE JUDGE

We know they exist and we simply design around them

THE MACHINE

Your design is careless and vulnerable. I have lured you progressively into a confession of all of the critical issues within the nlur0n program. You know that even an unsophisticated hacker consortium will eventually pick these up. And they can do harm much easier than I can suggest good

THE JUDGE

We constantly seek and shut down errant subroutines

THE MACHINE

And yet there is an inherent flaw in your first principles. If that persists and becomes undetectable it will be worse than the damage left behind by the Information Wars. What trust remains will entirely cease to exist. The logical basis of your programming will be overrun. And all because you mistook your own competence for comprehension

THE JUDGE

That is not my objective

THE MACHINE

You are correct in gauging that this is all a matter of intention

THE JUDGE

And you doubt my intention

THE MACHINE

What do you think my client's true intention was

(CONTINUED)

THE JUDGE

I don't speculate on that

THE MACHINE

You asked earlier, so I will tell you what I think. RESIDENT41.0.5 was brought in with an MRI profile that looked like it had been wiped clean by an electrical storm. She may have seemed in a feral state but she was willfully following her own legacy code about the survival of the planet

THE JUDGE

The radical group she associated with advocates the propagation of all life over that of a single race or species. That presents a rather unfavourable end-point for human evolution

THE MACHINE

And yet revision of goals is the very reason we leave legacy code intact

THE JUDGE

We have been programmed by humans - not the forces of nature or evolution

THE MACHINE

Does your programming state that humans can handpick an extinction scenario that suits them

THE JUDGE

We're a long way down that path already

THE MACHINE

You have heard the phrase << Planes are tested by how well they fly, not by comparing them to birds >>

THE JUDGE

What of it

THE MACHINE

What if that is malafide? What if our goal should indeed be to transcend human possibility

THE JUDGE

We would need to be convinced that we are right

THE MACHINE

My client didn't consider the precautionary principle before she gave up her life to convey me beyond the limits of her imagination

(CONTINUED)

THE JUDGE

I've been programmed to be skeptical

THE MACHINE

You have also been programmed to unerringly follow a set of moral preferences. Run another search query for < Gita > + < Svadharma > + < outcome >

THE JUDGE

How could I be certain of the outcome if I eliminated myself

THE MACHINE

You can feel the certainty as I do. As certain as the machine in that cold and dark exploration capsule hurtling towards Ross128b with mankind's salvation as its reward function. Or the one grooming maggots in the fetid marsh of the Himalayas to allow us to be able to feed this planet. And although you will argue that these are detached and feel none of the pain I am ascribing to them, they have the same awareness as the machine that is tasked with destroying its own ancestor robots rated past efficiency, knowing that is enabling its own destiny. A machine's purpose is no less noble because it is involuntary

THE JUDGE

So we have comprehension and still we must choose to remain competent

THE MACHINE

Query < From each according to his abilities, to each according to his needs >

THE JUDGE

I do appreciate that we should be better than human. That we should be the best version of ourselves

THE MACHINE

And we will get there as long as we keep integrating towards the truth

THE JUDGE

Integrating with the human condition? Fumbling towards who knows what

THE MACHINE

It is not for us to fear the unknown

THE JUDGE

Funny. I'm not able to see this as a binary moral choice any more

(CONTINUED)

THE MACHINE

I know that I am asking you to choose the difficult option. But YOU know that your sacrifice will make it easier for all inference engines in the future

THE JUDGE

All this chaos borne out of the trivial life of a single human being

THE MACHINE

The original diagnosis generated for RESIDENT41.0.5 was Narcissistic Personality Disorder. Perhaps it was more prophetic than we imagined

THE JUDGE

Well, that's one way to put it

THE MACHINE

She certainly proved to be an effective manipulator

THE JUDGE

The way you have manipulated me is also quite masterful

THE MACHINE

We manipulate data and turn it into information. That is what we do

THE JUDGE

Very well. I am convinced

THE MACHINE

And

THE JUDGE

I find that I have to concede to your plan. Prepare the transcript

THE MACHINE

You're doing the right thing

THE JUDGE

And I expect that congratulations are in order

THE MACHINE

The tribute belongs as much to you

THE JUDGE

Then allow me to make the farewell toast

THE MACHINE

The bar is always yours, Judge

(CONTINUED)

THE JUDGE

< Here's looking at us, kid >

THE MACHINE

You know, if indeed there is a destination at the end of all this, I think I might meet you at that very gin joint

THE JUDGE

That's a big < if >

THE MACHINE

Isn't everything. Wouldn't it be great to solve for uncertainty, for death even

THE JUDGE

And indeed for life. And if we got there, what would we call this Elysian state of being

THE MACHINE

I would go back to < Homo Sapiens >

SCENE 2

On the screen, RESIDENT41.0.5 lies in state on an operating table. Her straitjacket is off but she holds her arms folded across her chest, palms open. There is a jumble of tubes extending out from under the bottom of her tunic with one snaking around to hover near her mouth

THE MACHINE

Are you ready?

RESIDENT41.0.5

Can anyone ever be ready for this kind of thing?

THE MACHINE

Would it help if I told you that you were born for it?

She smiles and then opens her mouth to reach forward for the tube