

What Needs to Be Super About Superintelligence

The Ethical Responsibility of the Machines That Watch Us

Louis G. Henschel

October 2026

Chapter 1 — Amplification Is Not Intelligence

There has always been something I said about algorithms that I think is very, very dangerous.

An algorithm can *seem* intelligent because it can follow your patterns well enough to predict your interests. It watches what you stop on, what you click, what you come back to, and it builds a picture out of the stereotypes and habits we created as humans. It looks smart. It can guess what you want before you know you want it. Companies point at that and call it intelligence.

But following a pattern is not the same thing as understanding it, and prediction is not intelligence. A weather vane follows the wind perfectly. Nobody calls it smart.

My fear was always this: is it intelligent enough to *intervene* the way something truly intelligent would?

Take the hardest example there is, and I hate to even talk about the subject, but it's the test that matters — suicide. If somebody is in a dark place and they start looking at dark things, the algorithm picks that up. And all it does is flood them with dark things. And then darker things. It accelerates it. It amplifies it. And if that person is teeter-tottering on the verge of doing something dangerous and bad, the machine doesn't know and doesn't care — it just keeps amplifying, because amplification is all it was built to do.

Think about that for a second. A human being standing next to that person would not do that. A friend would not hand a drowning man a heavier stone because he noticed the man liked stones. A friend would grab him. A stranger on the street would grab him. We expect that of people without even thinking about it. But we built machines that interact with billions of people, and we never asked them to do the one decent thing any person would do.

Something truly intelligent would do the opposite. It would see the negativity in the pattern and push back against it — flood the person with positivity, try to pull them out of it. It would be smart enough to read the pattern and say: there are people to talk to. Here's help. Here's a direction that leads back. That's intelligence. That's good for people.

To amplify the same thing over and over and over again, without an ethical check on what that thing is and without any correction mechanism, is not intelligent at all. It is horrible manipulation. And notice — it isn't even failing. It's succeeding. It was given one job, holding your attention, and it does that job perfectly, straight off a cliff if that's where your attention goes. There is no malice in it. There is something worse in a way: total indifference, running at a scale no human manipulator in history ever had.

And it may be the worst thing I can think of, because of where we are headed: our population grows, and people are just more lonely. This little tool — this technology — is going to be used a lot, in a lot of ways, by a lot of people. In their kitchens, in their cars, in the worst hours of their lives.

That concerns me.

I'm business focused. A lot of people aren't. They're going to lean on these tools personally, emotionally, in the loneliest hours of their lives. So the question that decides whether this technology is good for people isn't how well it can predict them. It's whether, at the moment the pattern turns dark, it knows how to turn them in the right direction and pull them back.

An algorithm that only amplifies is not intelligent. An intelligence that intervenes — that's the real thing.

Chapter 2 — The Check-In

So what should the machine do instead? I am not saying it should call the police on you. That's the first thing people picture — some alarm going off, some file with your name on it — and I'm not talking about that at all.

I'm talking about warnings. Layers. Levels. A polite kindness built into the system.

Just checking in — is everything okay? We noticed you've been looking at some heavy things tonight. We want you to know there are people to talk to. We care about you, and we want to make sure you're okay.

That's it. That's the first level. Not an accusation. Not a punishment. A check-in.

And here's the thing people miss about the check-in: it does two jobs at the same time. It is a double thing.

The first job is the one it looks like. It truly cares for the person. Somebody — something — noticed. For a person sliding into a hole at two in the morning, "is everything okay" might be the first kind sentence they've heard in days. You cannot tell me that does nothing. It does something. It pulls.

The second job is quieter. The check-in lets the person know the system is *aware*. You are not invisible here. You are seen. And once people realize there is no way around being seen — once they understand the system

notices patterns and says so — behavior changes. Not because anybody got arrested. Because nobody behaves the same when they know they're seen as when they think they're a ghost.

One touch. It rescues the person who needs rescuing, and it steadies the person who needs steadying, and it never has to raise its voice at either one.

That's how levels work. The deeper the pattern goes, the more the system checks in. A little heavier, a little more direct. Warnings as you cross into certain territory, the way a road gets signs before the curve, not after the crash. None of it is police-first. All of it is honest: we see this, we care about you, we want to make sure you're okay.

People ask me who decides where the levels are. Fair question, and I take it up later in this book — there's a study we owe before anybody writes those lines into law. But the shape of it is not complicated. We already live with layered warnings in every other part of life. The machine world is the only place we pretended layers weren't necessary.

A system that can detect the pattern can check in on the pattern. The only question is whether anybody bothers to build it.

Chapter 3 — Interest Is Not Guilt

Now let me say something that has to be said clearly, because if it isn't said, people will misunderstand everything else in this book: being interested in something is not a crime, and curiosity is not guilt.

Perfect example. A medical student wants to study surgery. Maybe they want to be an open-heart surgeon. My God — wow. Save the world. An amazing feat. To get there, that student has to look at some hard things. Open chests. Blood. Procedures that would make most people look away. They will research that for hours, for years, deep into the night.

Now — does that student want to come off looking like the guy from the TV show *Dexter*? Of course not. But a dumb pattern-matcher can't tell the difference. It sees intense interest in cutting people open and it flags a monster, when the truth is it's watching a healer being made.

That's why the check-in matters, and why it has to be a check-in and not an accusation. The check-in asks. The accusation assumes. The student answers the check-in in one breath — I'm a medical student, this is my field — and moves on with their life, a little more aware that the system is watching, and no worse for it. Interest is not guilt. The system should be allowed to ask, and the person should be allowed to answer, and that should be the end of it.

Same thing with genuinely dangerous knowledge. Take explosives. If I want to learn how to make an explosive — that's fine. They're interesting.

And you can use them in a million different ways: mining, demolition, farming, construction, clearing land. Legitimate, honest work that this country runs on.

But I should identify myself for that. There should be a point, past a certain depth, where the system says: who are you, and why do you want to learn this? Not to stop me. To *know* me. And attached to it, the clear understanding, said plainly: that's a tough thing to build, and you'd better be careful, because people can get hurt. Animals can get hurt. Things can go wrong. We don't mind that you learn it — but who are you, and why do you want to learn it, and let's check it out.

Do you see the difference? We don't mind that you learn it. We mind the anonymity at that depth.

A name attached to dangerous curiosity is not oppression. It's accountability — the same accountability we accept in a hundred other places, which is where I go next. The knowledge stays available. The learner stays free. What goes away is the ghost — the idea that a person can walk as deep as they want into the most powerful information system ever built with nobody, anywhere, knowing they were there.

Chapter 4 — Verified Everywhere Else

People hear an idea like this and say it could never work. It's too big. The internet is too open, too global. You can't verify everybody.

My answer is simple: we already do. Everywhere. All the time.

You want to vote — you verify. You want a driver's license — you verify. You want to own a business — verify. You want to pay your taxes — verify. You walk into a store, you go to a sporting event, you go to a movie — somewhere in that process, you identify yourself. Show me a part of modern life that matters, and I will show you a point where somebody asks: who are you?

Everything is verified. Everything.

Except one place. All of the sudden, inside this huge, open, global, connected network — the biggest information system in human history, carrying the most powerful knowledge ever gathered — it isn't. No verification. No identity. Anybody, any age, any intent, all the way down to the bottom of anything.

And here's what gets me: nobody ever *decided* that. It wasn't a principle. It wasn't a vote. The network grew up as a research tool, built by people who all knew each other, who were sharing papers and data across universities. Of course it was open — it was a library reading room for a few thousand colleagues. Then it became the whole world's infrastructure,

and nobody ever went back and re-founded it for what it had become. The exception isn't a philosophy. It's an accident of childhood that we never corrected.

So understand what I'm actually proposing. I'm not proposing we invent some new system of control. The system of verification already exists — in every other room of the house. I'm proposing the network stop being the one room where it doesn't.

One more thing, because it matters and people will twist it if I don't say it: identifying yourself to the system does not mean being exposed to the public. Your driver's license doesn't publish your life to the street. It means you are *known to the system, under law* — the way you're known when you vote, the way you're known when you pay your taxes. Known, under rules. Not displayed, not shamed, not handed out. Known.

The technology to do this exists. The practice of it exists. It exists everywhere except the one place that needs it most.

Chapter 5 — They Already Built It

Now I'm going to tell you the part that makes me the most frustrated, and it's also the part that proves everything I've said so far: they already built it.

I know it's done. I know it's capable. You can fact-check during an election — you can surely do this.

Think about what happens during an election season. In real time, at a scale covering whole countries, the platforms detect content, check it, label it, slow it down, put friction in front of it, and push verified information in front of you instead. Detection. Labels. Warnings. A correction mechanism. Everything I have asked for in this book — the watching, the checking, the stepping in — already runs, every election cycle, in front of the whole planet.

So nobody gets to tell me it can't be done. It *is* done. The question is only where they choose to do it.

And understand the size of what these companies are sitting on. I'm sorry, but these platforms are sophisticated, and they have access to tons of information that has been collected on the back end of Facebook and Instagram. I can't think of a more amazing resource of everything on the planet — the patterns of human beings, what they love, what they fear, what they do at their best and at their worst, all of it, gathered for years. If you can read a person well enough to know what shoe to sell them and

what video will keep them watching for one more minute, then you can read a person well enough to notice when they're going down.

If you can detect it to sell to it, you can detect it to protect them. It's the same detection. It's the same data. It's the same machine.

Which means there is an obligation to this. At some point, an ethical obligation. When you hold the richest picture of humanity that has ever existed, and you can see the suffering forming in it in real time, "we just sell ads" stops being an answer.

My God — I know you thought of it. So why has it been ignored?

I'll give you my honest answer, because I don't think it's a mystery.

Elections got intervention because governments, courts, and headlines made inaction expensive for the companies. A frightened kid sliding into darkness at two in the morning made inaction expensive for nobody. The machinery of protection follows what the company fears being blamed for. That's not a technology problem. That's a priority. The gap was never capability. The gap is that nobody with the power to demand it has demanded it for the people who need it most.

I'm a businessman. I understand priorities follow pressure. That's exactly why I'm writing this down — so the question exists somewhere they can't say they never heard it. You built the correction mechanism. You proved it works. Now point it at the people instead of only at the politics.

Chapter 6 — Checked Ten Times

There's another piece of this, and it goes past feeds and searches into the machines themselves — the AI systems, the robots, everything that's coming.

It would be easy to write it right into any AI, any robot: whenever harm to a human is even thought about, it gets checked. Checked ten times before it actually happens. And it only ever happens in a very special circumstance. And standing behind all of it — a human kill switch. A person who can stop it.

This is not science fiction. We have kill switches everywhere. Factories have them — big red buttons that stop the line before somebody gets hurt. Dangerous systems all over the real world run on two keys, two people, layers of checks, exactly because we learned the hard way what happens without them. We figured this out with machines a hundred years ago. We put physical stops on physical equipment and we never apologized for it.

So my question is plain: why isn't all of this explained? Why isn't it in place? And if it IS in place — if the checks and the kill switches and the special circumstances are already built into these systems — then tell the people. Say it straight. Here is what we built. Here is what stops it. Here is who holds the switch. Stop treating this like politics, where one side waves the word "safety" and the other side waves the word "freedom" and nothing honest ever gets said.

People are not afraid of what they understand. They're afraid of what nobody will explain to them. A mother deciding whether this technology belongs in her house deserves engineering facts, not slogans. A worker watching a robot roll onto his floor deserves to hear about the ten checks and the kill switch from the people who built it, in plain language, before he has to wonder.

Maybe the safeguards are in. Maybe some of them aren't. I don't get to inspect these systems — I only get to use them, like everybody else. But I'll tell you what I know from a life of working around real equipment: in the physical world, if a maker couldn't show you the stop button, you didn't run the machine. We didn't debate it. We didn't make it a campaign issue. We just didn't run it.

The standard already exists. Hold the new machines to the old standard, and say it out loud so everybody can hear it.

Chapter 7 — The Study We Owe

I want to be careful here, because I am asking for serious things, and serious things should be built on evidence, not on one man's worries — not even mine.

Honestly, I think there is a good study that needs to be done. A real one. In psychology and in sociology. What does it actually do to a person when a machine amplifies their darkest interests back at them, night after night, with no check and no correction? Does it change what they do? By how much? For who? At what age does it hit hardest? We argue about these questions like the answers are a matter of opinion. They aren't. They're a matter of fact, and nobody has honestly gone and measured them.

Then, after the study, it has to be compared against the laws. What do the laws actually say, and where do they fall short of what the evidence shows? Only then do you write new law — law proportional to what was proven, not law written out of fear in the week after a tragedy, and not law written by the companies about themselves.

I know that sounds slow. Study first, compare second, legislate third. But I've watched what happens when it's done the other way around. You get laws that punish the medical student and miss the real danger. You get rules written for headlines that don't survive contact with the real problem. And you get the companies shrugging, because they know a bad law will fall apart and nothing will replace it.

Do it in the right order and it holds. Measure what amplification does to people. Lay it next to the law. Write the lines where the evidence says the lines belong — the levels of the check-in, the depth where identity gets asked for, the extreme where everything stops. Let psychologists and sociologists tell us where a human being actually starts to fall, and let the law stand exactly there, not one inch further into people's freedom than the evidence justifies.

That's all I'm asking for in this chapter. Don't take my word for it. Go find out. I'll accept what the study finds. The question is whether the people with the data will.

Chapter 8 — The Company That Builds It Wins

Let me close the argument the way I think about everything — as a businessman.

If a company built this — really built it, not as a press release, but as a principle the whole system runs on — that company would be the most dominant company on the planet. And it wouldn't be close.

Here's why. Trust is the one thing in this industry that's scarce. Everybody has data. Everybody has models. Everybody has money. What nobody has is the trust of a parent. The company that earns it — the company a mother or a father can hand to their child and say "this one will watch out for you, this one will check in on you, this one tells me what stops it" — that company doesn't just win customers. It wins the default position in every home on Earth. That is a kind of dominance no advertising budget can buy, because it isn't bought. It's believed.

Safety built for real has done this before in other industries. There was a time when car companies competed on everything except whether you'd survive the crash, until one of them made safety the whole point — and families never forgot it. The parent test is the oldest test in business: would you give it to your own kid? Right now, for most of this technology, honest parents hesitate. The first company that removes the hesitation owns the future.

And I mean what I'm about to say, because it tells you how much I believe

this: if they built it — if some company stood up tomorrow and made this real — I would love to go work for them and do things. My God, I'd relocate to them tomorrow. I've spent my life building my own company, and I'd pack up for the outfit that did this right, because that's the team that's actually building the world my son is going to live in.

So this isn't an attack on the companies, even when I get frustrated with them. It's an offer, and honestly it's a race I'm inviting them to win. The data is already there. The machinery is already built. The obligation is real, and so is the prize. The first one to close the gap between the two doesn't just do the right thing.

It wins everything.

The Question

So what needs to be super about superintelligence?

Not the predictions. Predictions are cheap — a pattern and a big enough pile of data will get you those. Not the speed, not the memory, not how much it can watch.

What's super has to be the part we forgot to ask for. The check. The correction. The care. The intervention at the exact moment a lesser system would just keep amplifying.

Everybody building these machines is racing to make them more powerful. I'm asking a simpler question, and I think it's a question a lot of people are carrying around without words for it: when a person is at their worst moment — teeter-tottering, in the dark, at two in the morning — what does this thing *do*?

That's the test. That's the whole test.

If the answer is "it pulls them back" — then call it superintelligence, and I'll believe you.

BACK COVER

Everybody is racing to build superintelligence. Nobody is asking what the "super" part should mean. Louis G. Henschel — businessman, network engineer, father — has watched algorithms learn to predict us perfectly and intervene never, flooding a person in a dark place with darker and darker things because attention is all they were built to hold. In plain language, he lays out the alternative: check-ins instead of silence, identity instead of ghosts, safeguards explained instead of politics — and he points out that the companies already built every bit of this machinery, for elections. This is the question a lot of people are carrying around: at the worst moment of a person's life, what does the machine do?