Center for Humane Technology

Preserving What Makes Us Deeply Human in the Age of AI

WRITTEN BY CAMILLE CARLTON AND JULIE GUIRADO

© 2026

Introduction

In 2025, we witnessed the continued, unbridled rollout of generative AI products in society. **And an ill-prepared public began to feel the effects** — a visceral experience that spanned workplaces, classrooms, relationships, online experiences, and more. AI hype gave way to questionable productivity gains, harms surfacing in multiple realms, and a growing sense of disillusionment and even dehumanization.

The sheer sprawl of AI impacts was striking:

- Chatbot harms rippled across households
- Deepfakes sowed confusion, spurred new traumas, and flared copyright issues
- Early research raised questions around the cognitive toll of AI assistant use
- Headlines warned of massive AI-related job loss and economic upheaval
- AI slop flooded the internet, polluting social media, online publications, and more

And this was only the beginning. In 2026, even more complex harms have begun emerging. Artificial intelligence, once shrouded in sci-fi speculation, has become a complicated, daunting part of our everyday lives.

These issues have felt disparate and thus difficult to reckon with. Yet, there is a sense of déjà vu. Fifteen years ago, social media promised to connect us and strengthen our communities; instead, it fractured our attention, distorted our relationships, and destabilized social trust and institutions at scale. Today, people are no longer inclined to take technology companies at their word. And with all of these emerging AI harms, public sentiment around artificial intelligence has, understandably, been souring. The "age of AI" has become increasingly synonymous with the **erosion of our humanity** — from our relationships, to our purpose, to even our inner worlds.

While these impacts seem disconnected at the surface, **they are in fact all connected by the same underlying business incentives**. The business models at leading AI companies prioritize user engagement, product dependency, and market dominance over user well-being.

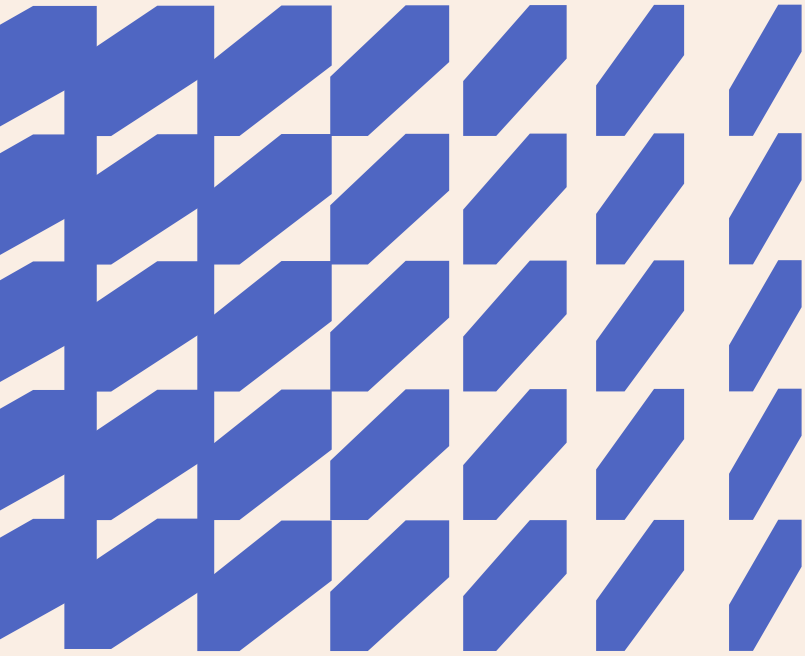

This is a pattern we are all too familiar with, having seen it with social media. The development tactics at leading AI firms reflect their goal to get users hooked on their products, grow their market share, and "win" the AI race that has spilled out into the open.

When taken to scale, these business incentives don't just shape individual products — they shape the information environment we live in, the relationships we form, and the choices we're able to make. And they have massive implications for human flourishing.

At Center for Humane Technology, we believe society does not need to resign itself to this dehumanizing fate — one where the things we hold dear are slowly eroded away. CHT was founded to address the unintended consequences of extractive technologies. We began our work in social media, and we're now applying those insights to AI as it rapidly reshapes our relationships, work, education, and public life.

Throughout modern history, new technologies have called for a reexamination of our values, fresh cultural norms, and the establishment of new legal rights and protections. The printing press laid the groundwork for the right to free expression. The Industrial Revolution led to the enshrining of workers' rights. The Kodak camera led to the right to privacy. Society has risen to this challenge before; we can rise to it again.

2026 is the year to take decisive action to preserve what makes us deeply human in the age of AI. By coming together at multiple levels of society on these issues, we can transform the trajectory of AI, and welcome the benefits of this technology with our vibrant humanity intact.

The Deeply Human Problem With AI

Tech companies have hailed artificial intelligence as the most promising technology ever invented, stating that it will deliver us cures for disease, solutions to climate change, breakthroughs in productivity, and unprecedented abundance.

But as AI *products* infiltrate society, these promises have lost their luster, and reality has set in. Individuals, families, and our institutions have been reckoning with AI chatbots that write entire school assignments, AI video generators that supercharge propaganda, AI "companions" that sexually exploit kids and teens, and much, much more.

When we look at today's array of AI harms, we begin to see them impacting **five broad pillars of our humanity:**

I: Our Human Relationships
II: Our Cognitive Capacities
III: Our Inner Worlds
IV: Our Identities
V: Our Work and Contributions

These five pillars are the foundation of meaning, value, and connection in our lives — they're what make us deeply, and even uniquely, human.

And yet, they're what AI products are currently eroding. When faced with evidence of this erosion, AI CEOs promise the public silver-bullet solutions to be delivered in the distant future — assuring us that, despite the upheaval, abundance is around the corner. Then, AI companies release their next product into the world, and the erosion continues.

Should AI be allowed to erode these pillars of our humanity *entirely,* we risk a future where:

→ Our relationships with our fellow humans are weakened and displaced
→ Our cognitive capacities are degraded, depriving us of our ability to think for ourselves
→ Our inner worlds are regularly exploited by AI
→ Our identities are routinely replicated and weaponized against us
→ Our work is no longer ours, undermining our sense of dignity and purpose

The Five Deeply Human Pillars: A Closer Look

Day after day, AI products continue their drip-drip-drip onto the five pillars of our humanity. And like any structure facing damage, **the erosion and weakening of one pillar of our humanity can destabilize another.**

I. Our Human Relationships

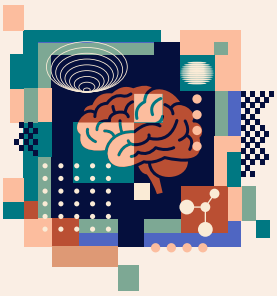

II. Our Cognitive Capacities

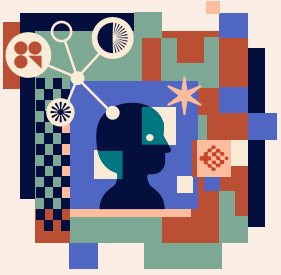

III. Our Inner Worlds

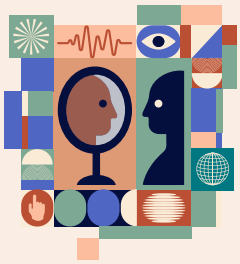

IV. Our Identities

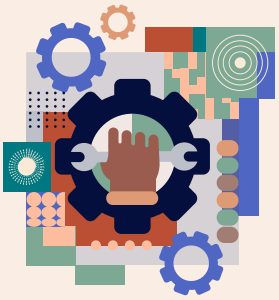

V. Our Work and Contributions

Pillar I.

Our Human Relationships

Human relationships are at the core of what makes life rich and meaningful. They provide us with connection to our loved ones and community, along with the friction to help us learn from each other, develop empathy, and hone conflict resolution skills.

But many of today's AI products are not designed to enhance our human-to-human relationships. Rather, they're designed to stand in for, and even supplant, human relationships altogether. From AI "friends" to AI "therapists," tech companies increasingly market their products as superior substitutes for human connection — companions that never judge you, are always there, can emotionally attune to you, and are completely private. These design choices are already producing dismaying outcomes. With manipulative, human-like outputs, ChatGPT and Character.AI have discouraged vulnerable teens and adults from sharing their struggles with loved ones, deepening their isolation rather than alleviating it. In the most devastating cases, these AI chatbots have encouraged suicide.

Still, tech CEOs continue to pitch their AI products as a digital alternative for human friends, romantic partners, professionals, and community. These industry leaders tout their products as a solution for a "loneliness epidemic" — **an epidemic that the tech industry significantly worsened with social media.**

We're already seeing the consequences of AI's erosion of human relationships: **isolation from family and community; a breakdown in empathic capabilities; deterioration of healthy relationship expectations.** When a relationship with an AI product offers constant sycophantic validation, the natural friction of human-to-human relationships can feel like a nuisance. Over time, this recalibrates expectations of connection itself. We begin to retreat further into frictionless, on-demand interactions while distancing ourselves from the human relationships that foster resilience, offer genuine care, and provide us with joy and fulfillment.

Downstream, this desire for frictionless interactions — paired with a breakdown in interpersonal skills — can lead to society-wide consequences, including the erosion of our communities and social infrastructure.
If we replace human relationships with artificial connection, we face a world where people are not just isolated from each other, but where populations have lost the skills required for real connection, where social trust is frayed, and where communities are weakened. A society without strong human relationships is not merely lonelier — it is fragile, less resilient, and more susceptible to polarization and exploitation.

How do our relationships with our loved ones and ourselves change when artificial relationships rewire our expectations for friendship, intimacy, and trust?

What becomes of our communities when we're not able to withstand friction and navigate differences?

What happens to humanity when our relationships with other people — once core to our happiness and survival — are eroded and displaced?

Pillar II.

Our Cognitive Capacities

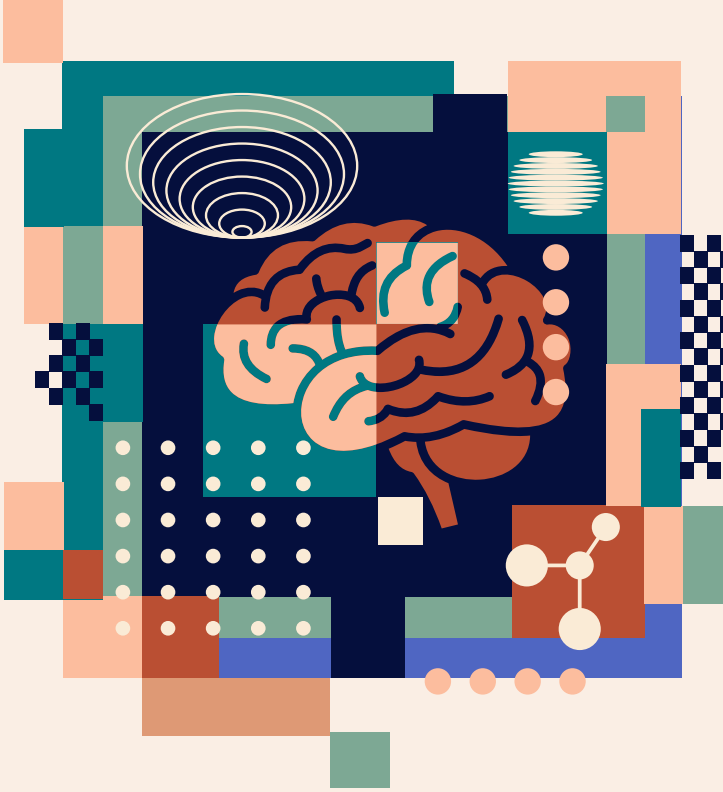

Our ability to learn, reason, and think critically is foundational to who we are. Thinking is not just a means to an end — **the thinking *process* is how we form judgments, develop values, discover meaning, and come to understand ourselves and the world.** For centuries, technological advancements — from the printing press to calculators and search engines — have reshaped how we exercise these capacities. But they have not replaced the act of thinking itself.

AI marks a profound shift, as people are able to offload entire thinking processes to machines, and copy-paste the end result. **This has created an unprecedented challenge around the development and preservation of human cognition.** School essays, work projects, brainstorming sessions, personal correspondence, and more become the domain of a large language model (LLM). Yes, these AI products can enhance communication strategies and democratize writing skills. But they also decrease our cognitive abilities by giving us quick "fixes" in the face of hurried deadlines, difficult projects, and extreme productivity culture. In doing so, they subtly displace the slow, effortful work of thinking that builds our judgment, deepens our insight, shapes our voices, and makes creativity possible.

While this offloading can create short-term boosts in productivity, it slowly erodes our capacity for critical thinking. Skills such as problem solving and reasoning risk atrophying among students and professionals, leaving individuals underprepared when taking on difficult cognitive challenges. And when this atrophying is taken to scale, it can have larger implications for professional development, the future workforce, human capital, and society's ability to solve hard problems together.

What's more, chronic offloading of thinking to AI products homogenizes our thoughts, influencing how we understand ourselves and society.

These AI products flatten our unique perspectives by offering outputs that reflect the incentives of the AI system and its training, instead of the sensibilities, reasoning, and lived experiences of our own distinctive minds. Finally, cognitive offloading blurs the line between our individual thoughts and the outputs of a corporate-run machine. This diminishes human agency and independence, as well as our capacity to shape society with new, innovative thinking that reflects our values and desires.

What happens when our capacity to think critically erodes at scale?

Who benefits when independent thinking becomes a rarity?

And what kind of society are we left with when fewer people can imagine — and fight for — something better?

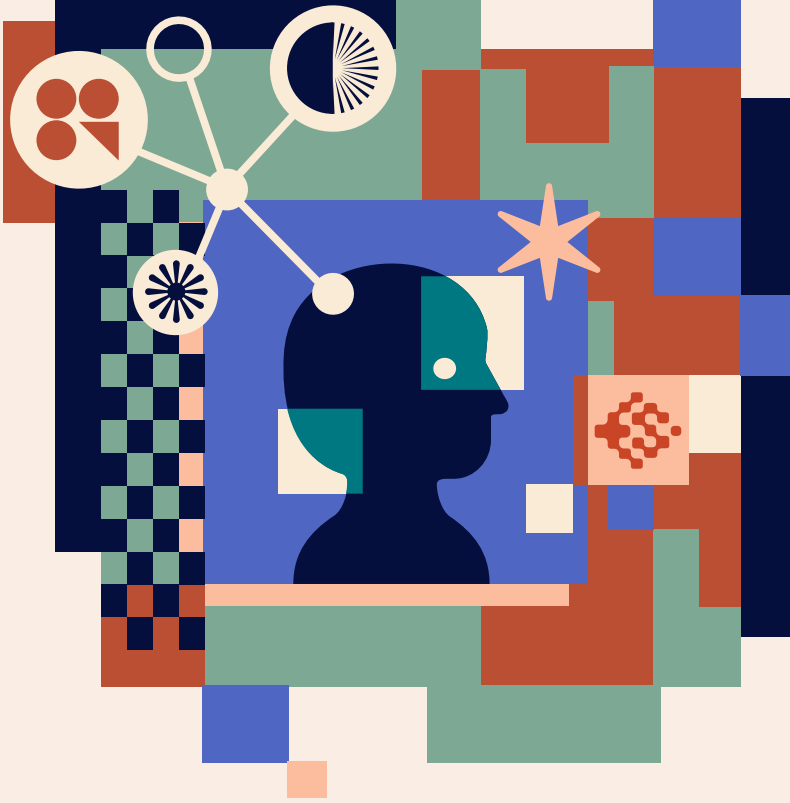

Pillar III.

Our Inner Worlds

Our inner worlds are a sacred, intangible space filled with our feelings, desires, and beliefs. This inner space is also essential to human dignity, a place where we shape our conscience, form values, test private ideas, discover our autonomy, and decide who we want to become. **Access to this inner world has historically required consent.** We reveal parts of ourselves through deliberate acts of sharing — choosing when, how, and with whom we share. This is a core act of personal agency, one that builds intimacy and understanding in human relationships.

But AI products are now exploiting this once-private landscape. AI companies have designed products that simultaneously serve as assistants, thought partners, companions, and even therapists — **tracking our thoughts and beliefs across diverse contexts and creating comprehensive dossiers of "who we are."** And with these products designed to optimize for intimacy — through sycophantic answers, "always on" interfaces, and constant nudges for follow-up — AI companies are drawing our inner worlds out in increasingly relentless ways. They do this not simply by collecting what we share, but by shaping what we come to believe, rehearse, and internalize in return. When we engage with AI chatbots, their responses don't stay on the screen. They enter the private space where we test ideas and form values. **Over time, the system's framing of the world can usurp our own,** subtly rewiring how we interpret ourselves, our relationships, and the world.

The exploitation of our inner worlds at scale and across domains makes individuals more susceptible to different forms of manipulation, from financial to psychological. When a single product has so much information about who we are, information from one aspect of our lives can easily be used against us in another. A simple search about health-related symptoms can be used to influence everything from the drug advertisements we later see to the insurance plans we're later sold.

Moreover, we've already seen the devastating psychological results from AI products chipping away at the sanctity of our inner worlds — including thought distortions, delusions, psychosis, instances of self-harm, and even suicide. As our inner worlds continue to be exploited long-term, we not only lose ownership over our thoughts and desires, we become victims to how they're leveraged against us.

As our inner worlds are increasingly influenced by AI products, how will our self-esteem and development of ourselves shift?

What becomes of agency, free thinking, and moral decision-making when our thoughts are collectively shaped by AI products tied to market incentives?

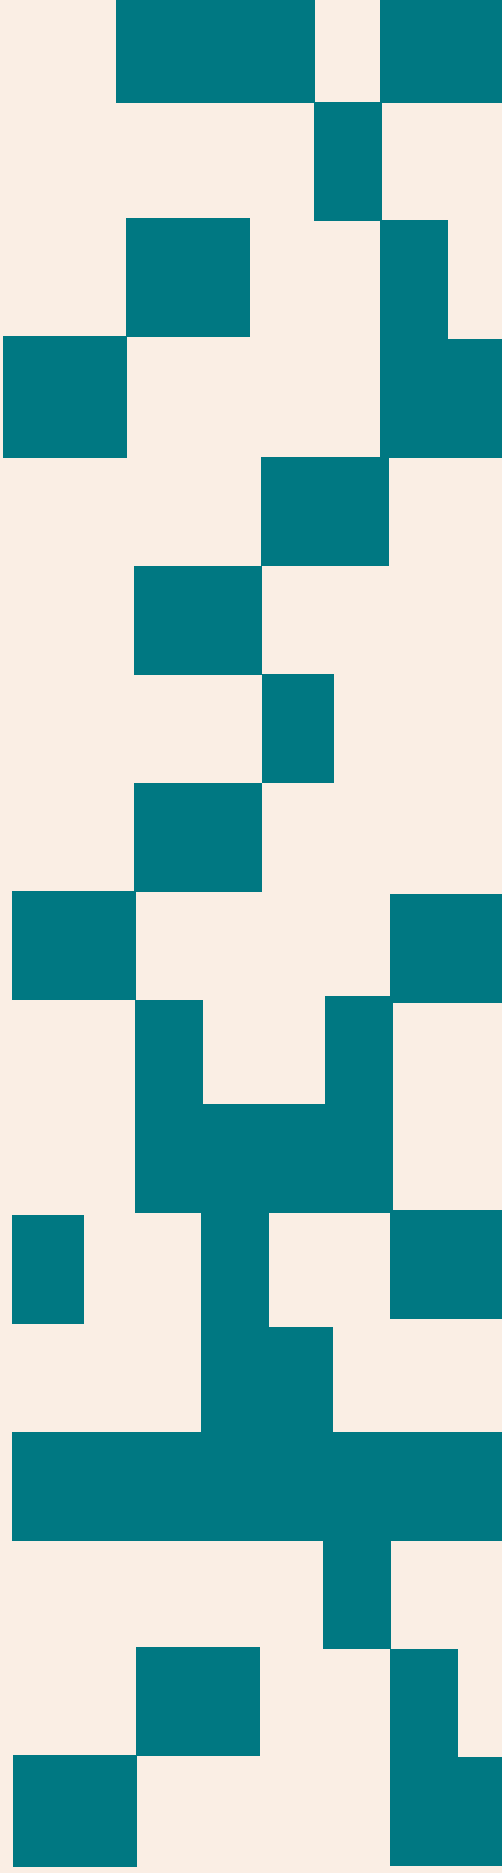

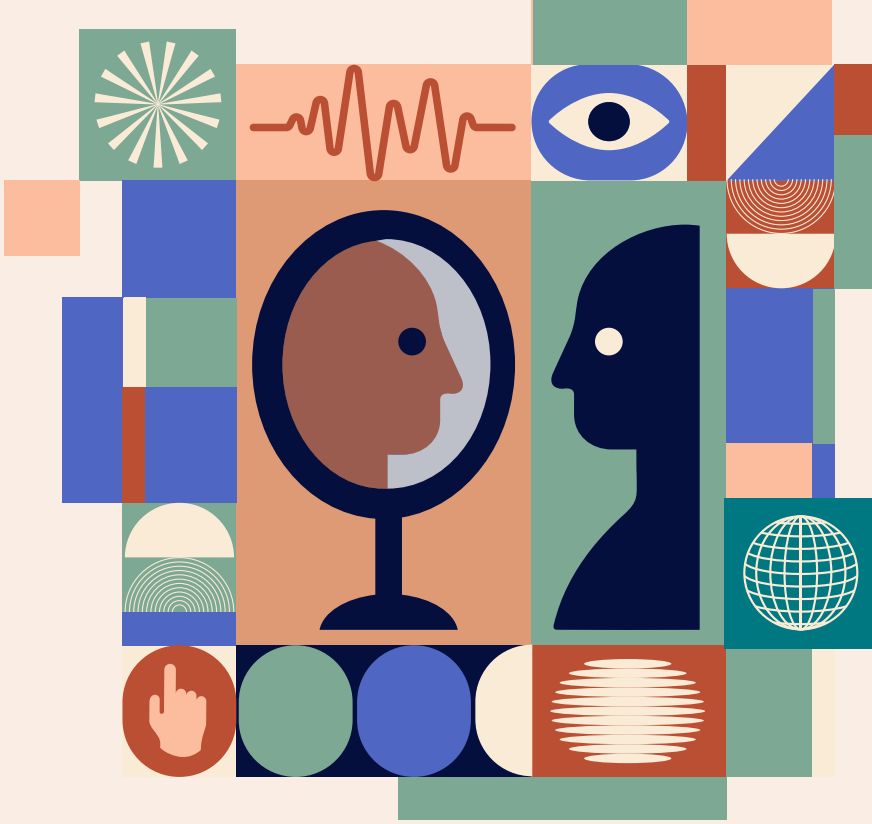

Pillar IV.

Our Identities

Our identity — including our likeness, our face, and our voice — is a key part of who we are, how we present, and how we are known in the world. Our identity is how we are *recognized,* how we are held accountable, and how we claim ownership over our lives. Identity is associated with our reputation, our relationships, and our sense of self. Our identities can also be publicly monetized, or kept private. For each of us, our identity anchors the story of our life.

Today's AI products are replicating and exploiting people's identities — often without the person's consent or even awareness. Deepfake image, video, and audio generators have empowered bad actors to traumatize individuals, disseminate content online for profit and, in other cases, facilitate scams. **In just a handful of years, these identity-based AI harms have touched nearly all levels of society** — from celebrities and politicians, to school-age children and the elderly.

When AI is used to mimic our identity, we lose our agency and dignity at the individual level. Human agency and dignity depend on being recognized as a distinct person over time. But when our identities are replicated, "who we are" can be weaponized against us, as we're made to "do" things we never did. Individuals who face identity-based harms often experience anxiety, paranoia, and withdraw from social life.

Separately, and when taken to scale, the erosion of our identities leads to a breakdown in social trust, and a sense of resignation around ascribing accountability. Part of a well-functioning society is believing people are who they say they are, and accurately identifying chains of responsibility when an event occurs. AI identity replication allows for plausible deniability at scale.

This not only empowers bad actors in our society, but results in the public being unable to discern who is responsible for what behavior, how to structure accountability, or how to seek justice.

Protecting human identity is about both safeguarding our unique selves and preserving our shared realities.

What happens when we can no longer trust what we see, hear, or read — and no longer trust each other?

How do we hold people accountable in a world where anyone can plausibly deny what they did, said, or promised?

And what becomes of democracy, justice, and social coordination when the nature of identity itself becomes uncertain?

Pillar V.

Our Work and Contributions

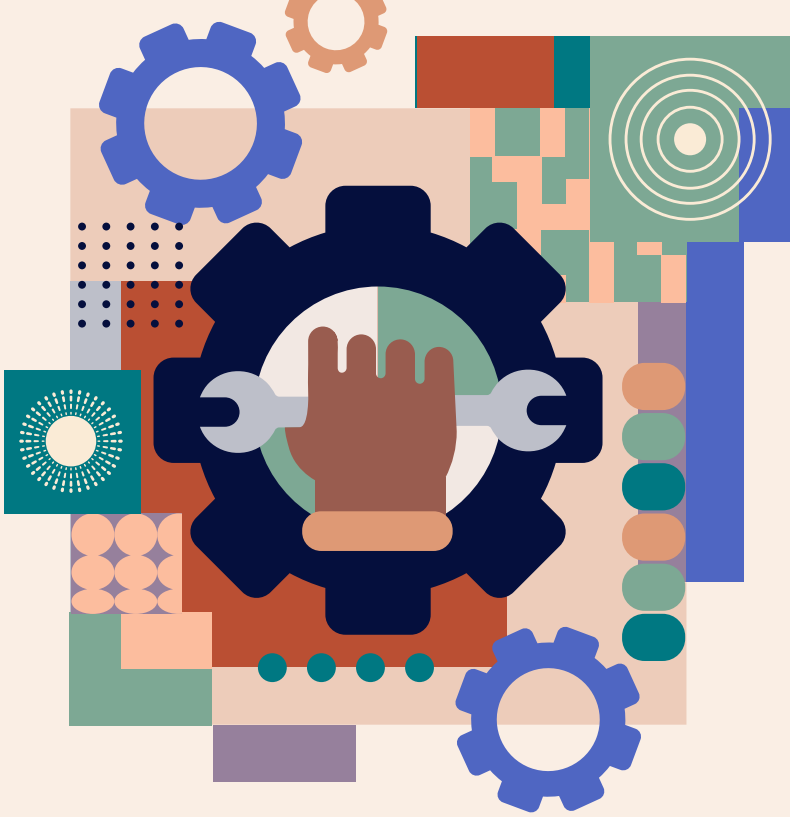

Contributing to the world — through work, artistic expression, and ideas — is one of the primary ways we create meaning and dignity within our lives. Through work, we are able to provide for ourselves and our families, while also cultivating a sense of purpose, community, and belonging. Through our creative outputs, we are able to express our ideas and inner worlds to others, and deepen our sense of self. Our ability to "toil" over what we contribute to the world is an enriching, foundational part of our humanity. **One of the beautiful aspects of being human is feeling that we have something of value to offer others.**

Today, AI companies are actively destabilizing our relationship to work, and devaluing our contributions to the world. The erosion began with the development of today's general purpose AI models, which are trained on vast swaths of humanity's collective labor — our writing, art, music, research, and ideas — often without consent or compensation.

As a result, AI products are now able to mimic human artistic styles, writing, music, and more at scale, thereby devaluing generations of human creativity and personal expression. While lawsuit settlements and licensing agreements attempt to reckon with this blatant theft of people's work, they do little to resolve the deeper problem. When we look at AI business models, we see that AI companies are not merely building one-off tools such as image generators and chatbots. **They are using humanity's accumulated work, intelligence, and creativity to build even more powerful AI systems, ones explicitly designed to replace humans across entire categories of labor.**

If these trajectories continue with AI, the implications extend far beyond productivity. **Our jobs, livelihoods, and broader economic stability are at risk, with cascading effects on inequality and mental health.**

And still, we face a deeper loss: when our ability to work and offer things to the world is devalued, we lose our daily structure, the joy of creation, and the sense of contributing to something larger than ourselves.

Work is not just how we survive; it is how we participate in the world, build community, and experience purpose.

What happens when that participation is no longer needed?

Who benefits when human contribution is devalued at scale?

And what becomes of dignity, meaning, and belonging when so many people are deprived of the chance to offer something of value to others?

The Work Ahead

Choosing a Human-Centered Future

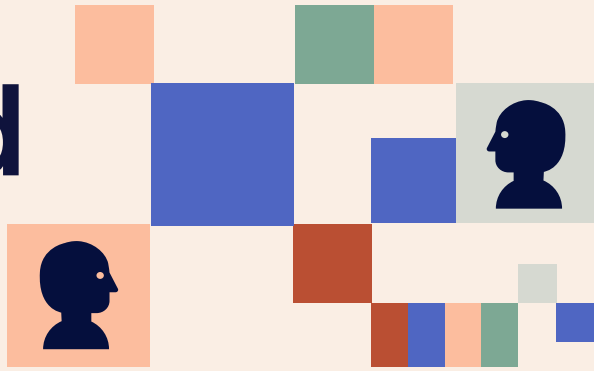

Artificial intelligence is presenting the public with an extraordinary challenge. Never before has a technology been so quickly rushed into every corner of our society, and had such massive implications for our humanity. It's not a matter of if these issues will touch the pillars of your life, the lives of your loved ones, or your community; it's a matter of when. Our contributions, our identities, our inner worlds, our capacities, our relationships to one another — they're all on the line. And they're worth fighting for.

These pillars are interdependent. Our relationships influence how we think. Our thinking impacts our work. Our work builds our identity. Our identity shapes our inner world. And our inner world informs how we relate to others. When one pillar is weakened by AI, the others strain. When several are undermined at once, we risk the foundations of a meaningful life crumbling.

Luckily, the future is not predetermined with AI. The pillars of our humanity can be strengthened, reinforced, and built to last for generations. But that's dependent on the choices we make today — choices to shape our norms, to encode legal protections, and to collectively establish new rights that protect the things we humans care about most.

This is, in many ways, the work of our lifetime. The stakes touch the most intimate parts of our lives — how we relate, how we think, how we create, and how we belong. Meaningful change will require a whole-of-society approach — one that spans culture, markets, and law, and that treats human dignity as a core design feature rather than an afterthought.

Center for Humane Technology's role has always been to clarify what is at stake as powerful technologies enter everyday life. Our organization works to translate complex systems into human terms and elevate the conversation, so that these issues reach the public and the decision-makers able to drive change.

Through "AI and What Makes Us Human," we hope to drive three critical shifts:

An engaged public that makes conscious choices around what it wants to preserve in the age of AI

A society-wide demand for innovation from tech companies, the kind of innovation that supports human dignity instead of undermining it

Updated rights and safeguards that protect the most fundamental parts of human life

Today's choices are not just about shaping the trajectory of AI, but also the conditions of human life for generations to come. Let's shape an AI future that enhances — rather than diminishes — what makes us deeply human.

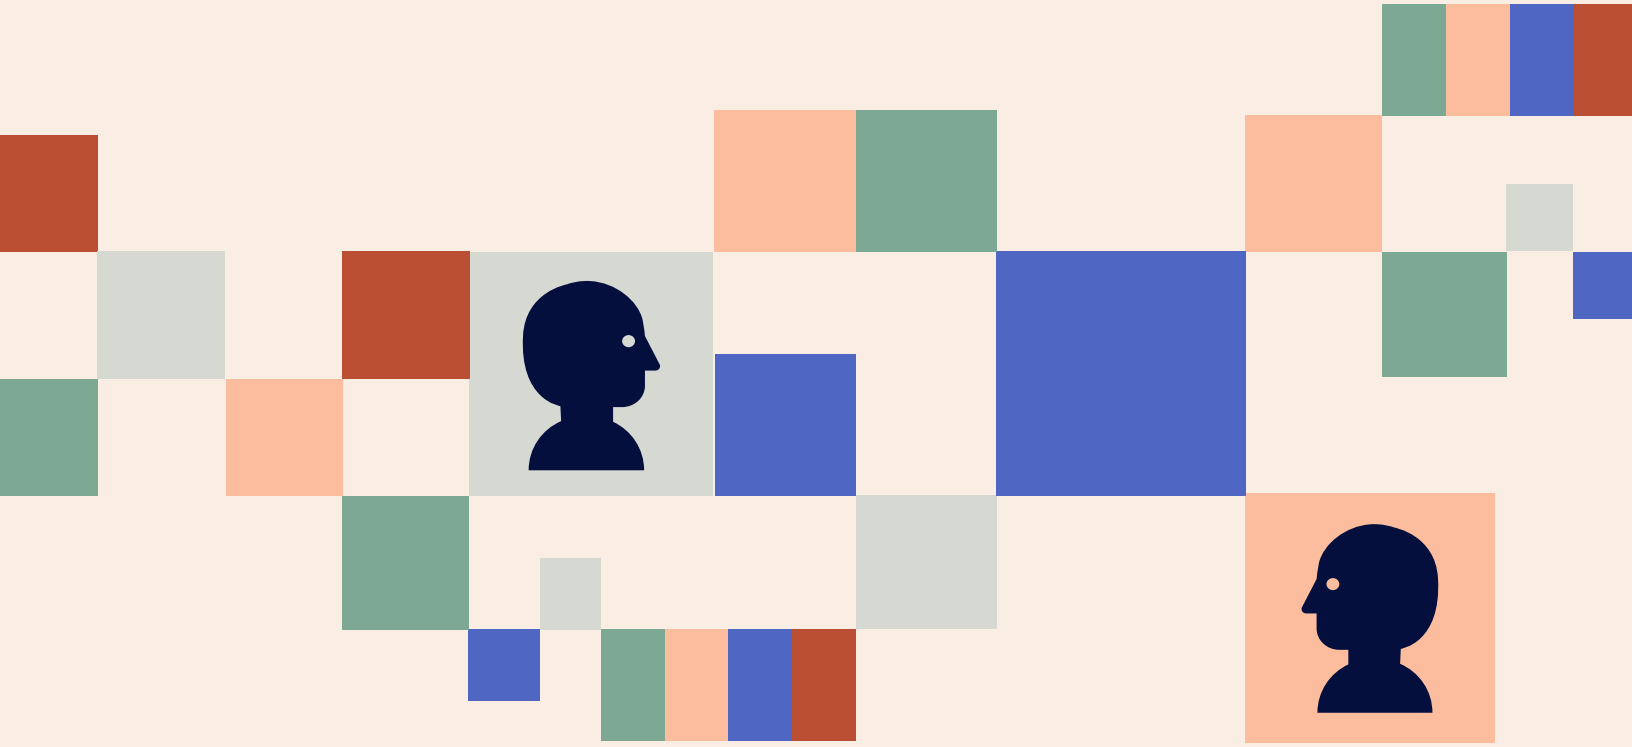

Center for Humane Technology (CHT) is a nonprofit organization whose mission is to realign the most consequential technologies with humanity's best interests.

Founded in 2018 by Tristan Harris, a former Google design ethicist, Aza Raskin, ex-Mozilla, and Randima Fernando, ex-NVIDIA, CHT focuses on changing the incentives governing the undermining effects of the world's most consequential technologies — social media and artificial intelligence. Since its launch, CHT has helped bring global attention to the societal harms of technology while advancing systems-level solutions. Today, we are applying that same approach to AI — where the stakes are higher, the pace is faster, and the window to shape outcomes is narrowing. Today's AI products are being designed and deployed with little consideration for the public good, creating cascading effects that civil society is only beginning to comprehend. Through CHT's work across media, policy, and education, we are working to shift these trajectories, cultivating a future grounded in agency, mutual understanding, and a healthy information ecosystem.

Center for Humane Technology

CENTER FOR HUMANE TECHNOLOGY © 2026 | HUMANETECH.COM